

Measuring Community Health

Vital Signs for Wikimedia Projects



Dario Taraborelli • Aaron Halfaker • Dan Andreescu

Wikimania 2014, London



気圧計

770

1030

780

1040

020

5

metrics != science



metrics as filters



Wikimetrics

Welcome to the Wikimedia Foundation's Wikimetrics homepage. This API allows you to

select a set of users, also known as a "cohort" (for example, **all users who signed up via the Thank You campaign**)

select a metric to be computed for each of these users (for example, how many **bytes they've added**) with optional parameters (for example, a **time range**)


and retrieve the response in JSON or CSV format.

You can also compute a single, aggregate value for the cohort (like the **mean** revert rate) .


[Learn More](#)[Analyze](#)

[Login](#) to see your cohorts, metrics, and reports

summer 2013



cohort-level metrics



cohort-level metrics



project-level metrics



project-level metrics



project-level metrics

Existing data sources

wikistats

Wikipedia Statistics Arabic

Monthly counts & Quarterly rankings / [Editor activity levels](#) / [Distribution namespace](#) / [Most edited articles](#) / [Zeitgeist](#)


Metrics have been collected from a **full archive dump**, which contains
See also [metrics definitions](#)

Monthly counts & Quarterly rankings: **June 2014**

Date	Wikipedians						Ar	
	total	new	edits		count			new per day
			> 5	> 100	official	> 200 ch		
Jun 2014	+1%		-5%	-4%	+2%	+2%	+7%	
May 2014	+1%		+11%	+17%	+2%	+2%	-32%	
Apr 2014	+1%		-6%	-17%	+2%	+3%	-19%	
Mar 2014	+1%		-2%	-13%	+3%	+3%	+39%	
Feb 2014	+1%		-5%	+3%	+2%	+2%	+35%	
Jan 2014	+1%		+10%	+21%	+2%	+2%	+19%	
	A	B	C	D	E	F	G	


<http://stats.wikimedia.org>

report card



<http://reportcard.wmflabs.org/>

Vital Signs



- granular measurements of new user engagement, community size, and content growth
- aggregated daily / weekly / monthly
- for every single Wikimedia project
- visualizations + raw data

https://www.mediawiki.org/wiki/Analytics/Epics/Editor_Engagement_Vital_Signs

summer 2014

4 categories of metrics

New users

- Newly registered users
- New editors
- New active editors
- Productive new editors
- Surviving new editors
- Surviving new active editors

Community

- Active editors
- Recurring old active editors
- Re-activated editors
- Unique editors
- Unique anonymous editors
- Unique editing bots
- Unique page creators
- Unique media creators

Content

- Edits
- Anonymous edits
- Bot edits
- Pages created
- Media uploaded

Curation

- Pages deleted
- Pages protected
- Pages moved
- Reverts

Principles

Relevant

Measure quantities that describe important phenomena

Replicable

Make research easily replicable and verifiable

Transparent

Provide formal specifications, remove ambiguity

Consistent

Replace proprietary, ad-hoc metric definitions; compare apples to apples

Robust

Make metrics replicable via multiple data sources at any point in time

Granular

Computable at different time scales

Anatomy of a metric

1. specification

Research:Productive new editor

Productive new editor is a [standardized user class](#) used to measure the number of first-time editors in a wiki project over time who make [productive contributions](#). It's used as a proxy for [editor productivity](#), and to a lesser extent, [editor activation](#). A "productive new editor" is a [new editor](#) who saves revisions to [content namespace](#) pages that are not [reverted](#).

Contents [hide]

1 Discussion

- 1.1 Excluding edits to deleted content
- 1.2 The n productive edits threshold
- 1.3 The t time cutoff
- 1.4 Time to revert cutoff
- 1.5 Limitations

2 Analysis

- 2.1 German Wikipedia
- 2.2 English Wikipedia
- 2.3 Spanish Wikipedia
- 2.4 French Wikipedia
- 2.5 Polish Wikipedia
- 2.6 Portuguese Wikipedia
- 2.7 Factor comparison of n and t

3 Usage

4 References


Discussion [edit]

Excluding edits to deleted content [edit]

Spammers and other non-productive new editors tend to create articles that are non-

Metrics standardization

Productive new editor



Specification

A **productive new editor**(n, t) is a [new editor](#) who completes at least n [productive edit](#)(s) within t time since [registration](#) (T).

WMF Standard

- $n = 1$ [productive edit](#)
- $t = 1$ [day](#)

Measures

[Editor productivity](#)

Aliases


[Productive newcomer](#)

Related metrics

[Newly registered user](#) • [New editor](#) • [Productive edit](#)

Anatomy of a metric


2. visualizations




Anatomy of a metric

2. visualizations


New editor



Surviving new editor



Productive new editor



Anatomy of a metric

3. discussion

Discussion

Excluding edits to deleted content

Spammers and other non-productive new editors tend to create articles that are non-productive and those articles tend to be deleted rather than the edits to the articles being reverted (and therefore excluding them from the [productive edit](#) criteria). Edits to articles that are deleted by the end of a new editor's first week since registration are *not* included in counts of productive edits.

The n productive edits threshold

Like choosing an n for any metric based on counts (e.g. [new editor](#) and [active editor](#)), choosing a threshold is somewhat arbitrary. Choosing a higher threshold will result in a smaller proportion of [newly registered users](#) being considered productive.

The t time cutoff

There are a few ways that the timespan for identifying productive edits can be drawn. The two most common ways are based on time bounds and events. A time-bounded approach is based on the use of some t cutoff to limit observations to a certain amount of time after a user registered their account. An event-based approach will use some event as the starting point to count user contributions. Another candidate time-span includes edits that a newcomer performed in their first [edit session](#). Since *productive new editor* qualifies the activity of a [new editor](#) we set $t = 1$ day, which effectively makes the class of productive new editors a proper subset of new editors. We analyze the effect of choosing a different value for t below.

Time to revert cutoff

It is theoretically possible that a revision could be reverted years after it was originally saved, observations taken at any time would truncate any future reverts (see [en:Censoring_\(statistics\)](#), specifically "right censoring"). In order to minimize this issue and compare editors' contributions fairly, a revert is only counted if it occurred within 48 hours of the original edit. Recent research suggests that, at least for the English Wikipedia, nearly all reverts take place within 48 hours^[1].

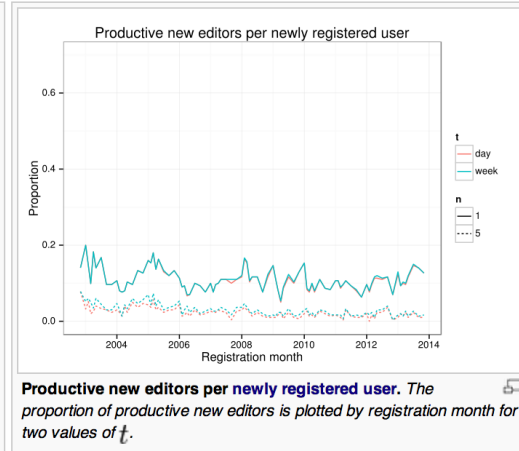

Limitations

- This metric represents productivity as a binary attribute of a user, it does not measure *how productive* a new editor is. New editors who make many productive edits and contribute substantial amounts of content will look identical (under this metric) to new editors who fix a few typos.
- The most clever vandalism/vandals may go unnoticed for more than 48 hours.


Anatomy of a metric

4. sensitivity analysis

Spanish Wikipedia [\[edit\]](#)




French Wikipedia [\[edit\]](#)



Sensitivity analysis


Does *new editor productivity* vary when we measure it over the first day or the first week?



https://meta.wikimedia.org/wiki/Research:Productive_new_editor

Sensitivity analysis

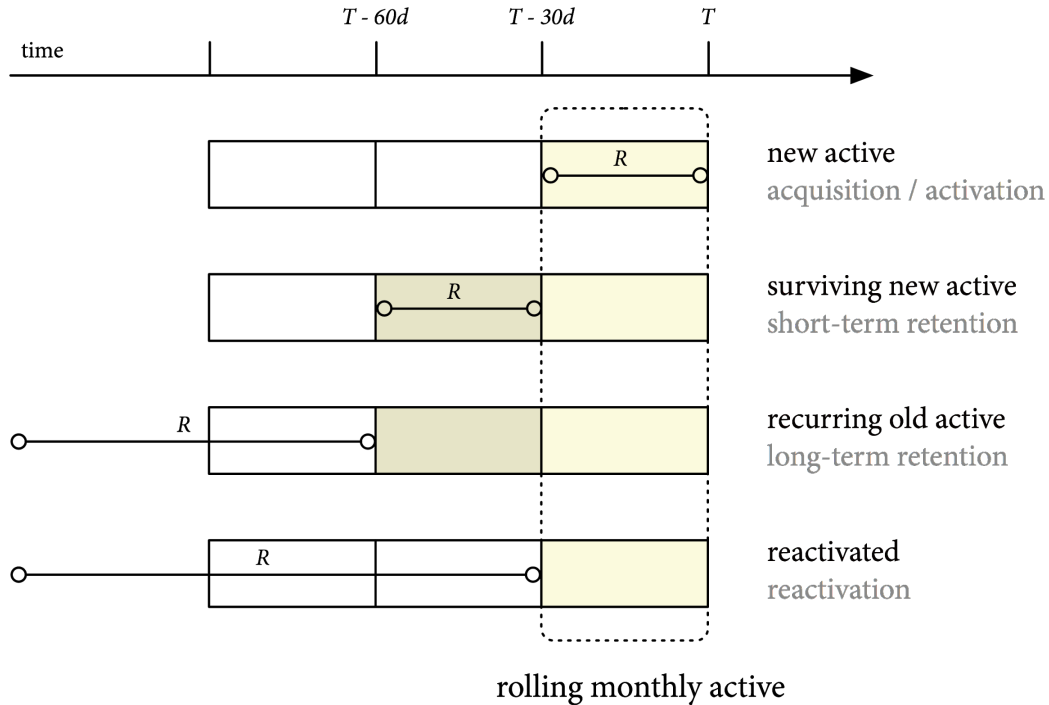
*Should we define **new editors** based on activity in the article namespace only?*



Anatomy of a metric

5. segmentation

What segments of the population of a project drive total active editor numbers



Use cases

1. Data exploration

“Newly registered users on German and Dutch Wikipedia have a higher activation rate than newbies who join the English Wikipedia”

“Spanish Wikipedia adds every day twice as many new editors than German Wikipedia, despite having only half its new user activation rate”

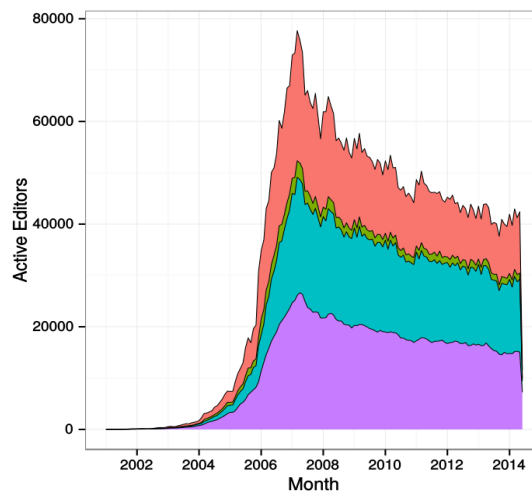

2. Natural experiments

“A change in abuse filter rules on the Italian Wikipedia significantly increased new editor survival”

Use cases

3. Projections and target setting

“To stop the active editor decline in the English Wikipedia, we should increase the retention of existing users by 87% or increase the activation of new editors by 23%”

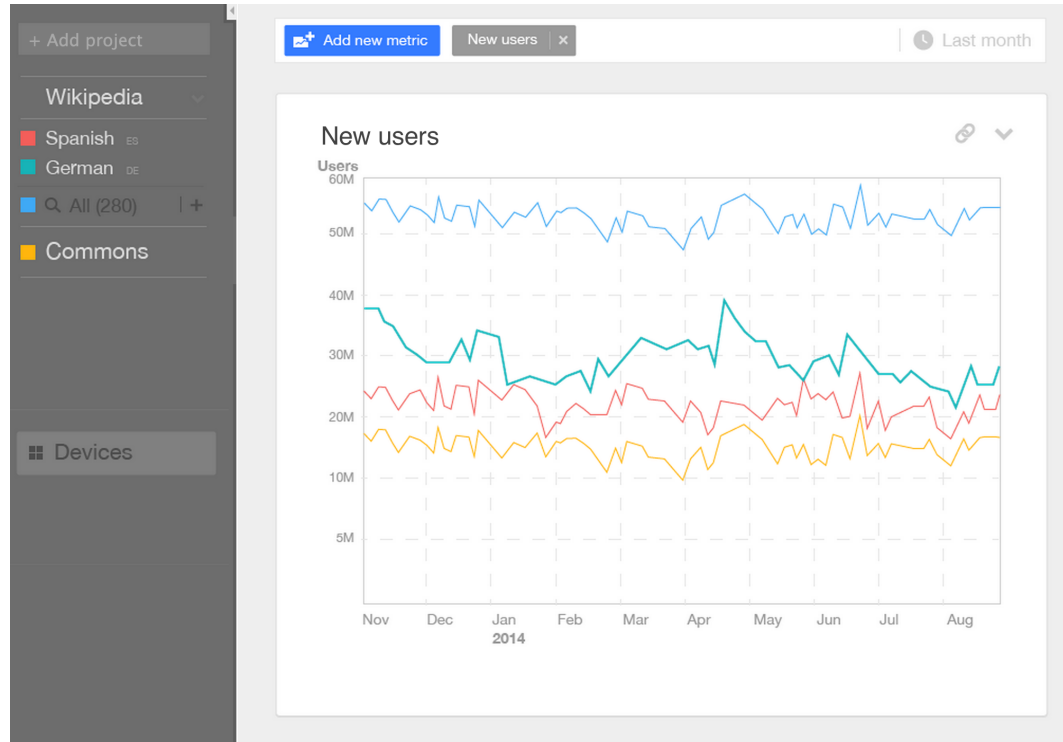


Data generation

DEMO

<https://metrics-staging.wmflabs.org/static/public/dash/>

Information architecture



<http://pauginer.github.io/prototypes/analytics-dashboard/>


Beyond basic measurements

Beyond signals based on simple edit counts: *quality, value added*

Better segmenting of the editor population (*classifying edit types*)

Readership metrics (*unique visitors, pageviews*)

Feedback and evaluation



Questions?

dario@wikimedia.org

ahalfaker@wikimedia.org

dandreescu@wikimedia.org

[[User:DarTar]]

[[User:Halfak (WMF)]]

[[User:DAndreescu]]

@readermeter

@halfak

@DanAndreescu

Read more

https://meta.wikimedia.org/wiki/Research:Metrics_standardization

Image credits

Tim Sheerman-Chase. *Fair, Barometer Detail* (CC BY) <https://www.flickr.com/photos/68932647@N00/7615682432/>

Mark Dumont. *Mad Science* (CC BY) <https://www.flickr.com/photos/23661161@N02/5826590955/>

Future scientist experiments on sister http://www.gifbay.com/gif/future_scientist_experiments_on_sister-117763/

Joaquim Alves Gaspar. *Vernier Caliper* (CC BY SA) https://commons.wikimedia.org/wiki/File:Vernier_caliper.png