

Open Source Metrics (Sunday 11 am)

Measure all the things!

(Editor disclaimer: these notes are still pretty raw. I tried to add in some links and context in places but may be off here and there)

Speakers: Josianne Marsan from CHAOSS, Felipe Hoffa Developer Advocate from Google, Josh Simmons from GSOC staff

Governing board from CHAOSS - Linux stuff on metrics. <https://chaoss.community>

How do we take all these tools and data and turn them into something useful, like dashboards

Two committees working with communities on data. Looking at project health and metrics design.

Google has staff working with TODOGroup also working on metrics. Varies what a company or a community wants to see.

Nice chart on top contributors on GitHub comparing repos on one axis and committer count on another (and stars to indicate circle size). But put together by one guy - how can we make more interesting charts like that and get projects to come up with consistent metrics? Links appreciated! Of course other closely related stats are available such as PRs, issue discussion, and such, can we improve the chart? GitHub archive download lacks org info?

So much data available these days. But plenty of people are just pulling out semi-random non producible stuff from it that looks cool but doesn't really lead to over-time consistent usefulness. CHAOSS may help here

What's the hypothesis? What are the factors? What do you expect to learn? Need to figure some of that out before even going looking for data. What are you trying to prove or learn?

Josianne from CHAOSS has been talking to a lot of people and encourages further discussion.

What metrics do we care about as GSOC orgs? What defines our project health?
Contributor count? PR count?

No one-size-fits-all for metrics

Metrics that exist get used more than metrics that do not exist. Communicate the caveats.

In the very beginning of CHAOSS very fundamental questions were asked: Could they come up with consistent metrics and explain them well? Highlight not trying to compare apples to oranges, but self-analyze when useful

One typical example of useful metrics is diversity - but even if you get something out of self-analyzing your own project's diversity could letting others read those results be met with a backlash from lack of context?

Blog entry by Felipe:

<https://towardsdatascience.com/these-are-the-real-stack-overflow-trends-use-the-pageviews-c439903cd1a>

Stack Overflow has a "Trends" section - but how to use those metrics? Tensorflow has more recent questions than Git but probably Git is still used more (more page views)

Some projects handle development very differently: Git focuses on very tiny changes, which then might result in some higher metrics

As a project owner looking for analytics: how are we managing our community? PR review time? Forum signups vs forum engagement?

"So you want to ask Stack Overflow questions? (Gabi mysql backend)" - how long does it take a question to be answered? (Editor: this was up on the projector and I thought I could find it online easily and provide a link. Now not so sure but maybe

<https://medium.freecodecamp.org/always-end-your-questions-with-a-stack-overflow-bigquery-and-other-stories-2470ebcda7f> is close?)

Use analytics to help identify gaps

Some users prefer anonymity. You can test out various responses if you show up with a female, male, or generic avatar/name.

How do you analyze “softer” metrics like student retention? (Editor: this is one of the extensions I’m aiming for in my Trello adventure in better tracking contributors)

Google uses very subjective approaches to try to judge OSS orgs for GSOC. One is whether it looks like the project is active in general (not compared to other orgs) via GitHub stats (for a project hosted there, of course)

Again various so heavily on the project and even the lifecycle of a given project that it is hard to make comparisons or attempt consistent analysis. A brand new project will likely have a higher ratio of abstract questions than an old one.

Has CHAOSS done research on sub-projects within a single community to be able to draw more accurate conclusions within such a context? As opposed to entire projects compared to others.

Running sentiment analysis on comments might be an interesting exercise.

A CHAOSS student this year in fact did a sentiment analysis as part of GSOC via GitHub comments.

One interesting observation (drawn from presumed gender analysis on Stack Overflow comments) showed that females tended to thank users more than males, despite a policy discouraging thanks in general - which was then changed

Even personal context matters hugely - whether somebody got asked a question at work or online. Or if a given question had already been asked in the past and kept being asked

Very popular first words on issues: “It would be nice ...” - with the starting phrase having an apparent impact on closure rate (starting with “Would it be possible...” had a higher closure rate)

TODOGroup repo on GitHub to check out:

<https://github.com/todogroup/awesome-oss-mgmt>

Find Josianne in Quebec at Universite Laval (

<https://www4.fsa.ulaval.ca/enseignants/josianne-marsan/>) or on the CHAOSS website

(for instance read

<https://www.linuxfoundation.org/blog/2017/09/chaoss-project-creates-tools-to-analyze-software-development-and-measure-open-source-community-health/>) -

alternatively reach out to the CHAOSS community at large, especially the two related committees on metrics (<https://chaoss.community/metrics>) and software (

<https://chaoss.community/software>)

Find Felipe Hoffa on Twitter: <https://twitter.com/felipehoffa>