

Recruitment Criteria

Recruitment criteria for the "Repository" recruitment channel of the publication "Committed to Trust: A Qualitative Study on Security & Trust in Open Source Software Projects".

Our general recruitment approach in the repository channel was a stratified sampling in quartiles of GitHub repositories ranked by both a popularity and an activity score. We based this score on repository-level metrics provided by the GitHub API such as the number of commits and committers as well as the number of stars and forks.

Our initial repository dataset was downloaded in July 2021 from GH Archive (<https://www.gharchive.org/>), a service providing historical GitHub repository data, publicly available for further analysis. We limited our dataset to code repositories that received at least 40 commits from at least 20 distinct committers in the previous six months, which sets a minimum threshold for any given selected project's activity. This was done with the intent of excluding inactive and personal projects, in which our inquiry would either not reach active contributors or where interpersonal trust processes are irrelevant.

The resulting 15,256 repositories were enriched with up-to-date data from GitHub's API, such as programming language usage, topic tags, as well as star and fork counts. Users usually give out stars as a means of bookmarking a project, or to explicitly value a project's merit. A project is "forked" into a user's namespace for them to be able to make changes to its code base and consequently create a pull request for their changes to be accepted into the main source tree. The combination of the number of a project's stars and forks can thus serve as a proxy for its popularity. To ensure that the selected projects recently went through the onboarding of new contributors, we only proceeded with those that gained new committers in July 2021, and which had not contributed to the project before. After excluding duplicate repositories as well as repositories exclusively containing markup languages, we arrived at a set of 4,456 projects for final consideration.

We joined the popularity and activity indicators to a combined ranking and divide the set of projects into quartiles. This ensured high diversity across the indicators, while minimizing the amount of strata. We then iteratively selected and contacted projects from each stratum (e.g., first project from 1st quartile, first project from 2nd quartile, and so on) until we reached interview saturation.

Figure 1 shows this approach in a diagram.

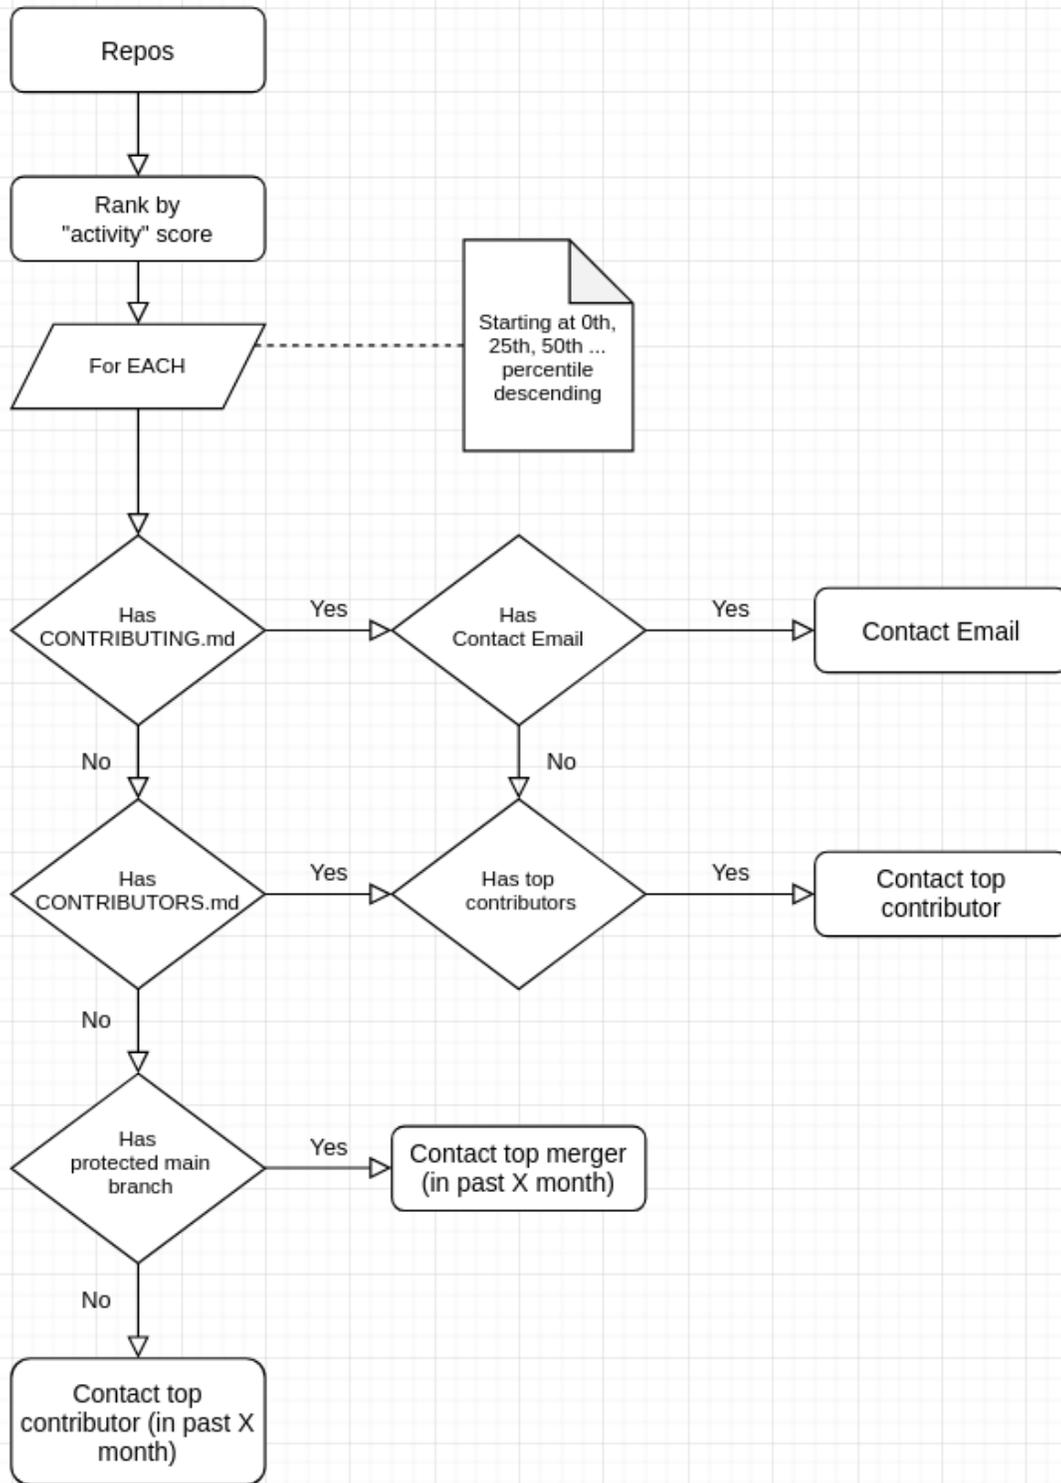


Figure 1: Diagram of our selection methodology.