Dynamics of Team Library Adoptions: An Exploration of GitHub Commit Logs

Pamela Bilo Thomas
University of Notre Dame
pthomas4@nd.edu

Rachel Krohn
University of Notre Dame
rkrohn@nd.edu

Tim Weninger
University of Notre Dame
tweninge@nd.edu

Abstract—When a group of people strives to understand new information, struggle ensues as various ideas compete for attention. Steep learning curves are surmounted as teams learn together. To understand how these team dynamics play out in software development, we explore Git logs, which provide a complete change history of software repositories. In these repositories, we observe code additions, which represent successfully implemented ideas, and code deletions, which represent ideas that have failed or been superseded. By examining the patterns between these commit types, we can begin to understand how teams adopt new information. We specifically study what happens after a software library is adopted by a project, *i.e.*, when a library is used for the first time in the project. We find that a variety of factors, including team size, library popularity, and prevalence on Stack Overflow are associated with how quickly teams learn and successfully adopt new software libraries.

Index Terms—information adoption, GitHub, software libraries, innovation

I. INTRODUCTION

Online collaboration systems like GitHub provide a powerful setting in which to study the process of learning new information and adopting new technology in a group setting. By investigating how software developers adopt and use software libraries in this specific context, we may better understand how humans learn new technical information and incorporate concepts [1].

This work aims to answer the following research questions:

- RQ1 What are the events that happen when a team adopts a library for the first time?
- RQ2 Are commits containing new libraries more likely to have deletions than other types of commits?
- RQ3 Do the answers to these questions vary by library type, team size, or the amount of information available on Stack Overflow?
- RQ4 Do team members fight over the adoption and usage of a new library?

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASONAM '19, August 27-30, 2019, Vancouver, Canada

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6868-1/19/08/\$15.00

<http://dx.doi.org/10.1145/3341161.3342928>

II. DATA ACQUISITION AND CLEANING

First, we issued a query to the GitHub search API for projects written primarily in Python. GitHub returned repository IDs of the 1,000 most popular Python projects on the site. We then found all GitHub users who made at least one commit to a repository in this set and retrieved all their Python projects. We did this breadth-first style crawling two additional times, culminating in 259,923 projects with 89,311 contributing GitHub users.

Of these, we were able to clone 259,613 Python repositories to disk. These repositories constitute about 13% of all Python repositories on GitHub as of September 2018. The full dataset of cloned projects occupies about 8 TB of disk space. Source code and repository IDs are available at <https://github.com/tweninger/growingpains>.

Additionally, we downloaded all question posts from Stack Overflow from its inception until September 2018. Appropriate tagging tends to increase viewership of questions [2], so we filtered out any posts that were not tagged as Python posts, then extracted all libraries mentioned in each question.

An important complication that arises in Git commit histories is that stashes, reverts, branches, and other Git commands can result in a non-monotonic chronology of edits. Fortunately, each commit keeps a pointer to its parent(s). Because the order of actions is more important than exact times, we enforce a monotonic chronology according to the commit graph.

We combed through each Git commit log to find which libraries had been imported by using Regex pattern matching to identify import statements, such as ‘from ℓ import f ’ and ‘import ℓ as f ’, where ℓ is a library and f is a function. We searched the log to find lines which referenced the functions contained in a library import using pattern matching to search for the libraries ℓ and functions f , along with indicator characters such as $.$ and $($. We gathered the author name, library used, and commit type.

Formally, for a project p and a time t , we define a *library adoption* to be an event (p, ℓ, t) representing the first time t that ℓ is found in p . The distribution of commit-activity per project, illustrated in Fig. 1a, resembles a shifted power law distribution. Because of this dynamic, 50% of projects were found to have 10 or fewer commits (*i.e.*, median of 10) and 90% of projects have 100 or fewer commits.

The distribution of the number of libraries adopted per project, illustrated in Fig. 1b, also resembles a shifted power

Fig. 1: (a) Number of commits per project follows a shifted power law distribution. (b) Number of libraries used per project also follows a shifted power law distribution. (c) When libraries are adopted into repositories, it tends to occur early in the repository history.

TABLE I: Statistics surrounding newly adopted library ℓ

Avg LOC that reference ℓ	31.34
Median LOC that reference ℓ	4
Avg inserted LOC after first adoption that reference ℓ	2.09
Avg deleted LOC after first adoption that reference ℓ	1.62

law distribution, albeit with a larger offset than the commit-activity distribution of Fig. 1a. However, the number of adoptions is less evenly distributed: 54% of projects adopted 10 or fewer distinct libraries and 98% of projects adopted 100 or fewer libraries.

Across all commits of all projects, we find that library adoptions occur more frequently within the first few commits. Figure 1c shows that a project's first commit adopts 6.4 libraries on average (with a median of 2 not illustrated). A project's second through fifth commits adopt 3.3, 1.1, 0.8, and 0.65 libraries on average (with median values of 0 not illustrated). In general, the average number of adoptions per commit appears to follow a Zipfian decay.

III. ACTIVITY CHANGE AFTER ADOPTION

A simple (albeit poor) indicator of productivity in software projects is the number of lines of code (LOC) that are added and/or removed in a commit. Table I shows there is a wide gap between the average and median LOC that reference a library ℓ . Additionally, average LOC drops quickly after the first commit, as shown in Fig. 2.

After the initial commit, we find that most of the following commits have only a small positive net productivity, and that the volume of activity of lines of code referencing ℓ in Fig. 1c tends towards zero rather quickly after the adoption.

Stack Overflow. To find the impact of Stack Overflow on the speed of library adoptions, we plot the number of users of ℓ by the mean average number of Stack Overflow posts (across all adoption times) in Fig. 3 that existed when ℓ was referenced. We observe a strong positive correlation between the number of library users and the number of Stack Overflow mentions for standard libraries ($R^2=0.625$, $p < 0.001$) and PyPi libraries ($R^2=0.410$, $p < 0.001$). There is a small positive correlation between usage of unknown libraries and Stack Overflow posts ($R^2=0.08$, $p < 0.001$).

We also investigate the impact of Stack Overflow popularity on library growth. Formally, we compute the growth of a library ℓ within a project as follows. If $x=0$, then let $y_x=1$; otherwise $y_x = y_{x-1} \left(\frac{\sum_{i=0}^{x-1} (n_i) + n_x}{\sum_{i=0}^{x-1} (n_i)} \right)$, where n_i is the number of changed lines of code in commit i that contain ℓ .

We plot the median growth (in LOC referencing ℓ) as a function of the number of commits after the adoption in Fig. 4. The primary distinction is in the growth rates for libraries with more than 1000 Stack Overflow posts. 100 commits after the adoption, the adoption of a highly mentioned library on Stack Overflow will have approximately 350% growth (on average) compared to only 250% growth for less mentioned libraries.

When a library does not appear on Stack Overflow, the growth rate is similar to libraries that have over 1000 posts. Libraries that do not appear at all in Stack Overflow mostly consist of libraries that were written by developers who are also the authors committing the library to the repository. This may explain why growth is large in unknown libraries – the adopters know how to use the library because they wrote it.

Project Team Size. Next, we investigate differences in library adoptions as a function of team size. Git and GitHub directly store how team members collaborate and the types of activities that they perform [3]. For example, researchers have found that diversity and team makeup have a significant impact on the productivity of GitHub teams [4], and larger teams tend to process more pull requests on average [5].

We observe in Fig. 5 that larger teams exhibit slower adoption growth and appear to be less agile than larger teams. A possible explanation could be due to perspective differences between two or more committers to a project. Users might feel more comfortable making more commits or experimenting with newly adopted libraries in smaller teams, or if they are working alone. Also, we can theorize that larger teams might do more work while using fewer commits, which results in there being less opportunity for growth later.

Additionally, Fig 5 shows us that the distribution of team sizes has a power law-like heavy tail wherein 59% of projects have only a single committer; 24% and 7% of projects have two and three distinct committers respectively.

Fig. 2: Library additions, deletions and net-usage (in LOC) after the adoption event.

Fig. 3: Number of users of ℓ in GitHub dataset as a function of the number of Stack Overflow posts about ℓ

IV. CODE FIGHTS

In the context of library adoptions in collaborative projects, we informally define a code fight as a series of commits that include back-and-forth additions and deletions of the same code containing a newly adopted ℓ . A *round* is a series of commits by one user that is uninterrupted by the other user. Formally, let $n^{(r)}$ represent the net change in lines of code referencing ℓ in round r ; $r = 0$ indicates the round of the adoption event. Also let $n^{\leq r}$ be the sum of all lines of code referencing ℓ up to and including r , i.e., the running total.

A *code fight* occurs if there exists any r such that $(1 - \epsilon)n^{\leq(r-1)} \leq n^{\leq r}$, where we set ϵ to represent the percent reduction that must occur, with $\epsilon \in \{.10, .20, .30, .40, .50\}$.

We observe that fights are relatively rare, occurring between 1 and 3 times for every 100,000 commits on average. Also, choice for ϵ has a limited effect on the probability of a fight. The probability of a fight increases with team size, but with diminishing returns that resemble a Zipfian Distribution.

Next, we analyze what happens during a two-person fight. Technically speaking, the first round of a fight is the adoption event and the second round of the fight is the removal of at

least $100(1 - \epsilon)$ percent lines of the adopter's code. After this point, the two fighters may continue with more rounds of back and forth commits.

Despite the dropoff in number of fights, the adopter tends to fight back with more lines of code. We observe that odd-numbered rounds, corresponding to the adopter, have more net LOC referencing ℓ per round, than the deleter's round that comes afterwards. This is shown in Fig 6. We see that the larger the original deletion of the code was, the less likely the adopter is to fight back with lines of code.

We define a fight's winner as the user who was the last committer referencing ℓ . By our definition of rounds, the deleter wins approximately 90% of the fights because the adopter only fights back 10% of the time.

What role does experience play in winning a fight? The current work maintains the standard set by prior studies [6] and therefore defines experience as the time since the user's first commit (in any project). The more experienced committer wins the fights between 70% and 80% of the time. Results from alternative ϵ values were nearly identical to $\epsilon = 0.1$. Interestingly, the more experienced users have about a 75% win probability even when the experience differences are less than a week or even a day.

We observe that common debugging libraries `pdb`, `pprint`, and `syslog` comprise three of the top four most common causes of fights. It is not surprising to see the `distutils` library counted among the top fight starters. This particular library is used to generate source code distributions i.e., code releases, but it is strongly encouraged that users use the `setuptools` library instead. So most cases importing `distutils` is likely an error.

V. DISCUSSION

RQ1: What does it look like when a team adopts a library for the first time?

In Fig. 1c we observe that library adoptions tend to happen early in a project's history. We can expect that it is difficult to adopt a new library once a project has matured. Perhaps this is because new libraries may introduce instability into a repository, or because the primary innovation within a project occurs early on in its lifespan.

Fig. 4: Growth of library usage after adoption grouped by Stack Overflow usage. Q1, Median, and Q3 growth in lines of code referencing ℓ are represented from bottom to top respectively.

Fig. 5: (a) Median percentage change in lines of code referencing ℓ after adoption. (b) Like the commit and adoption distributions illustrated in Figs. 1a and 1b, the team size distribution follows a power law.

Fig. 6: When a two-person fight occurs, the adopter u (indicated by odd-numbered x) tends to commit more code than v on average.

RQ2: Are commits containing new libraries more likely to have deletions than other types of commits? Once an adoption has occurred, we track how long it takes the library to become stable within the project by examining how many additions and deletions occur in the commits after a library is adopted. In Fig. 2, we observe that activity involving a newly adopted library is relatively high after a commit occurs. Over time, the number of lines of code referencing the adopted library stabilizes. We can safely conclude that users tend to write most lines of code that involve a newly adopted library relatively soon (within 10-15 commits) after library adoption.

RQ3: Do the answers to these questions vary by library type, team size, or the amount of information available on Stack Overflow? When team sizes are larger, the lines of library code do not grow as quickly relative to the first adoption commit as they do when team sizes are

smaller. This may be because larger team projects require more communication and planning and are therefore less agile than small teams or individual projects. Additionally, we showed that the number of times a library appears in Stack Overflow is highly correlated with the number of adoptions.

RQ4: What does it look like when team members fight over new library usage? When working on a team, there is bound to be conflict. Different team members have various opinions about which library is best to use in a repository. The probability of these fights occurring increases with team size. The winner of these fights tends to be more experienced.

Implications. We see that the number of commits and adoptions per project, along with team size, follow a power law distribution. We found positive correlations between the number of times libraries appear in Stack Overflow and GitHub. We discovered that popular libraries on Stack Overflow have faster rates of adoption for projects in Git. Additionally, smaller teams are more agile and they can adopt to using new libraries more quickly than larger teams, when productivity is measured as a function of median percentage growth. We also find that code fights are rare, but when they occur, they tend to be won by more experienced coders, and involve libraries which are used for debugging purposes.

ACKNOWLEDGEMENTS

This work is sponsored by grant W911NF-17-1-0448 from the US Army Research Office and grant W911NF-17-C-0094 from DARPA.

REFERENCES

- [1] R. Krohn and T. Weninger, "Library adoption in public software repositories," *Journal of Big Data*, vol. 6, no. 1, p. 36, 2019.
- [2] A. K. Saha, R. K. Saha, and K. A. Schneider, "A discriminative model approach for suggesting tags automatically for stack overflow questions," in *Mining Software Repositories*. IEEE Press, 2013, pp. 73–76.
- [3] J. Middleton, E. Murphy-Hill, D. Green, A. Meade, R. Mayer, D. White, and S. McDonald, "Which contributions predict whether developers are accepted into github teams," in *Mining Software Repositories*. IEEE Press, 2018.
- [4] B. Vasilescu, A. Serebrenik, and V. Filkov, "A data set for social diversity studies of github teams," in *Mining Software Repositories*. IEEE Press, 2015, pp. 514–517.
- [5] B. Vasilescu, Y. Yu, H. Wang, P. Devanbu, and V. Filkov, "Quality and productivity outcomes relating to continuous integration in github," in *Foundations of Software Engineering*. ACM, 2015, pp. 805–816.
- [6] R. Rotabi, C. Danescu-Niculescu-Mizil, and J. Kleinberg, "Competition and selection among conventions," in *WWW*. 2017, pp. 1361–1370.