Sustainability Forecasting for Apache Incubator Projects

Likang Yin

lkyin@ucdavis.edu

University of California, Davis

Davis, California, USA

Qi Xuan

xuanqi@zjut.edu.cn

Zhejiang University of Technology

Hangzhou, Zhejiang, China

Zhuangzhi Chen

zzch@zjut.edu.cn

Zhejiang University of Technology

Hangzhou, Zhejiang, China

Vladimir Filkov

vfilkov@ucdavis.edu

University of California, Davis

Davis, California, USA

ABSTRACT

Although OSS development is very popular, ultimately more than 80% of OSS projects fail. Identifying the factors associated with OSS success can help in devising interventions when a project takes a downturn. OSS success has been studied from a variety of angles, more recently in empirical studies of large numbers of diverse projects, using proxies for sustainability, e.g., internal metrics related to productivity and external ones, related to community popularity. The internal socio-technical structure of projects has also been shown important, especially their dynamics. This points to another angle on evaluating software success, from the perspective of self-sustaining and self-governing communities.

To uncover the dynamics of how a project at a nascent development stage gradually evolves into a sustainable one, here we apply a socio-technical network modeling perspective to a dataset of Apache Software Foundation Incubator (ASFI), sustainability-labeled projects. To identify and validate the determinants of sustainability, we undertake a mix of quantitative and qualitative studies of ASFI projects' socio-technical network trajectories. We develop interpretable models which can forecast a project becoming sustainable with 93+% accuracy, within 8 months of incubation start. Based on the interpretable models we describe a strategy for real-time monitoring and suggesting actions, which can be used by projects to correct their sustainability trajectories.

CCS CONCEPTS

• Software and its engineering → Open source model.

KEYWORDS

OSS Sustainability; Apache Incubator; Sociotechnical System

ACM Reference Format:

Likang Yin, Zhuangzhi Chen, Qi Xuan, and Vladimir Filkov. 2021. Sustainability Forecasting for Apache Incubator Projects. In *Proceedings of*

the 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '21), August 23–28, 2021, Athens, Greece. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3468264.3468563>

1 INTRODUCTION

Open source has democratized software development. Developers flock to OSS projects hoping to add certain functionality, contribute to a worthy goal, and sharpen their skills. However, more than 80% of OSS projects become abandoned over time, especially the smaller and younger projects [44]. Certainly not all OSS projects are meant to be widely used or even to persist beyond a college semester. However, even the large, popular OSS software, widely used in our daily lives and by fortune 500 companies, started out as small projects. Thus, from a societal perspective, it is important to ask: Why do some nascent OSS projects succeed and become self-sustaining while others do not [40]? And can the latter be helped?

To aid developer communities to build and maintain sustainable OSS projects, nonprofit organizations like the Apache Software Foundation (ASF) have been established. ASF runs the ASF Incubator (ASFI) [15], where nascent projects aiming to be a part of the ASF community are provided with stewardship and mentor-like guidance to help them eventually become self-sustaining and even top-level projects in ASF. Projects in ASFI, called *podlings*, are required to adhere to ASF rules and regulations, including keeping all commits and emails public. When certain conditions are satisfied, project developers and ASF committees decide if a podling should be *graduated*, referred to as a *successful* sustainability outcome. Otherwise they get *retired*. Per ASFI, 'A major criterion for graduation is to have developed an open and diverse meritocratic community. Graduation tests whether a project has learned enough and is responsible enough to sustain itself as such a community'¹.

Podlings in ASFI receive a mentor, file monthly reports, and get feedback. In spite of this support, many podlings fail. Most ASFI committers do not lack coding expertise, but graduating from the incubator requires more: it asks for effective teamwork and sustainable, community development. These requirements are most challenging to meet. From comments in ASFI, we see that developers are confused by the expectation of 'The Apache Way', especially initially. E.g., from project *Flex* on 05 Jan 2012, '...I think there is a need to keep things as simple as possible for people who [are]

This work is licensed under a Creative Commons Attribution International 4.0 License.

ESEC/FSE '21, August 23–28, 2021, Athens, Greece

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8562-6/21/08.

<https://doi.org/10.1145/3468264.3468563>

¹<https://incubator.apache.org/guides/graduation.html>

already confused with what’s going on with the move to Apache...’ The frustrations sometimes persist beyond the initial period. A comment from project *Flex* on 06 Jan 2012, states ‘...many people are confused and lost as to what to do. Who is providing that direction for them?’ In large part, understanding how to achieve the graduation requirements seems to be the culprit, likely due to the abstract nature of those concepts. Another reason is that comparison to others is difficult. From project *Rave* on 28 June 2011: ‘[sporadic adherence to requirements] makes it very confusing and difficult to compare against what other projects as some are doing too little while others are doing too much...’

Thus, there is a need to connect the proverbial dots on how to get from the point of entry into ASFI to checking all the graduation requirement boxes. That brings us to the motivation of our paper. The extrinsically labeled ASFI dataset offers heretofore unavailable, fine-grained records of historical trajectories of projects under policies and regulations of ASFI. We posit that:

The process that ASFI projects follow toward becoming sustainable can be modeled as a function of a small set of project features, so that the outcome (graduation/retirement) can be predicted early on, from the successes and failures of others, allowing for trajectory adjustment if needed.

To deliver on that, in addition to the historical, outcome-labeled ASFI data, we also need a theoretical framework that can capture the complexity of OSS development. Over the past 30 years research in organizational behavior and management has documented the evolution in project management practices [8], as they have moved toward more successful models [2]. The *socio-technical* view of an organization [30] has emerged as one of the more successful hybrid models recognizing the benefits that integrated treatment of the technical aspect (code, machines, device, etc.) and the social aspect (people, communication, well-being, etc.) has on an organization [36]. Likewise, OSS projects have been effectively studied from the *socio-technical system* (STS) perspective [1], with the social side capturing humans and their communication channels, and the technical capturing the content and structure of the software [43].

Here, inspired by the socio-technical systems modeling perspective, and the availability of extrinsically labeled historical data of project sustainability from ASFI, our goals are: (1) to identify socio-technical features distinguishing projects that graduate from the ASF Incubator from those that do not (i.e., find the determinants of OSS project sustainability), and (2) to build temporal forecasting models that can predict sustainability outcomes at any time point in the project development, and thus (3) to offer practical and timely advice on intervening to correct a project’s course, especially early on. To approach these goals, we conduct a mix of quantitative and qualitative empirical studies. We start by gathering project technical traces (commits) and social traces (emails) from the ASF Incubator website². From those, we construct the temporal social and technical networks for each project, and perform exploratory data analyses, deep-dive case studies, build an accurate forecasting model, and finally implement the interpretable model presenting timely advice. We illustrate our workflow in Figure 1.

²<http://incubator.apache.org/projects/>

Figure 1: The workflow of our mixed-methods study.

Our contributions in this paper are:

- We provide a novel longitudinal dataset of hundreds of OSS projects’ development traces under ASF regulation, with extrinsically labeled project sustainability status.
- We propose the first OSS project sustainability forecast measure modeled from tens of socio-technical network and project features. Our model shows excellent predictive performance ($\geq 93\%$ accuracy as early as 8 months into incubation).
- We find that ASF incubator projects with fewer but more focused committers and more but distributed (participating in asynchronous discussions) communicators are more likely to gain momentum to self-sustainability.
- We describe a strategy for real-time monitoring of the sustainability forecast for any project, derived from an interpretable version of our DNN model.

This paper is a first step toward showing that end-results of OSS projects’ sustainability trajectories can be effectively forecast and possibly corrected upward, if needed. Our motivation goes beyond ASFI as many more nascent projects fail outside of ASFI, so self-monitoring and self-adjustment may be more pertinent to them.

2 BACKGROUND AND THEORIES

We present background on the Apache Software Foundation (ASF) and OSS success, and then we introduce related theories through which we generate our Research Questions.

2.1 Apache Software Foundation Incubator

Community over code is the tenet of the Apache Software Foundation (ASF) community [22]. Their belief is that if they take good care of the community, good software will emerge from it.

However, conflicts are ubiquitous in OSS projects [27], and not even ASF can escape them. To minimize conflicts, ASF requires projects to make all communication publicly available on the mailing lists, summarized popularly as *if it did not happen in the mailing lists, it did not happen* [42]. The communication records benefit developers as they reflect on previous decisions and trace their precursors, therefore improving efficiency and productivity.

The ASF community adopts a democratic way in many of their affairs. For example, contributors are invited to vote +1 (yes), 0

(okay), or -1 (no) to project-wide changes. However, ASF committers can live in different time zones, and their response to a project decision can delay largely. Regarding that, ASF community adopts the *Lazy Consensus* [31]. Moreover, the ASF community also believes in *Earned Rights*, that newcomers should be treated the same way if they have proven their technical skills.

The goal of the ASF Incubator (ASFI) is to help projects become self-sustained and eventually join ASF. Like many OSS developers, ASFI committers work at will and there are no formal obligations on them. Thus, ASFI projects are required to show they are able to recruit new committers, and fill existing technical debt [38]. However, attracting new committers is difficult as they can be affected by both social and technical barriers [46]. To address this issue, the ASFI community has established a set of specific rules that emphasize providing mentorship to newcomers [47].

During incubation, ASFI projects need to adopt ASF procedures to develop and cultivate the projects' community, and standardize their working style. To graduate from the incubator and finally become a part of the ASF, projects are required to demonstrate they can self-govern and be self-sustained [12]. The specific requirement of sustainability can vary from one project to another [13, 18]. A project's graduation is ultimately approved for its self-sustainability by ASF's Project Management Committee via several rounds of public voting³.

2.2 OSS Projects Success and Sustainability

Recently, substantial work has focused on modeling the success of OSS projects [34, 37, 48]. Even though there is no universally agreed definition of OSS success [19], there are two main perspectives. The first one is from the development process viewpoint, which is often measured by technical metrics of software [21], e.g., code defect density [39], response time [35], and error resolution rate [28]. The second one is more from the social angle, including contributor growth [59], community participation [32], and communication patterns [56]. In effect, K. Crowston et al. [11] studied the operationalizing success measures under the context of FLOSS projects. N. Cerpa et al. [7] provide survey evidence that factors in the logistic regression models are not as predictive as expected. D. Surian et al. [49] identify discriminative rich patterns from socio-technical networks to predict project success in the context of *SourceForge* projects. J. Coelho et al. [9] conduct mixed-method analysis on GitHub projects, and they find that most of modern OSS projects fail due to project characteristics and team issues. M. Valiev et al. [53] conducted studies on sustained activity within the PyPI ecosystem, and find that relative position in the dependency network has a significant impact on project sustainability. None of these consider forecasting over time and thus are not useful for real-time monitoring, which is the major contribution of this work.

Although OSS success and sustainability measure similar aspects of projects [58], they are, in fact, not the same thing. There are two main differences between the two. First, OSS success is measured statically while sustainability is measured dynamically. Second, sustainability is a measure related more to the human and social aspect (e.g., the ability to take responsible collective action, and

an open and inclusive atmosphere, etc.) than the technical aspect (defect density, technical advantage, etc.).

Therefore, the sustainability of an OSS project becomes even more important when it is a part of a larger ecosystem [20, 53]. Such OSS projects are inter-dependent to each other, the sustainability and stability of one project can introduce tremendous network effect to its ecosystem.⁴ Therefore, the sustainability of OSS projects becomes even more significant as they can influence many other OSS projects in the ecosystem that rely on them.

2.3 Socio-technical Systems Theory

Socio-technical structure plays an important role in achieving collective success in OSS projects [4, 14, 56]. A Socio-Technical System (STS) comprises two entities [51]: the social system where members continuously create and share knowledge via various types of individual interactions, and the technical system where the members utilize the technical hardware to accomplish certain collective tasks. The theory of STS is often referenced when studying how the technical system is able to provide efficient and reliable individual interactions [24], and how the social subsystem becomes contingent in the interactions and further affects the performance of the technical subsystem [16].

OSS projects have been studied from the network view [14]. González-Barahona et al. [23] proposed using technical networks, where nodes are the modules in the CVS repository and edges indicate two modules share common committers, to study the organization of ASF projects.

Moreover, in socio-technical systems, governance can be applied through long-term or short-term interventions. Smith et al. [45] proposed two conceptual approaches: 'Governance on the outside' objectifies the socio-technical and is managerial in approach. 'Governance on the inside' is more reflexive about the role of governance in co-constituting the socio-technical. From that perspective, the ASFI community is a unique system that has both outside influence (regulations from ASF committee) and inside governance (motivated by the project managers).

2.4 Contingency Theory

Contingency theory is the notion that there is no one best way to govern an organization. Instead, each decision in an organization must depend on its internal structure, contingent upon the external context (e.g., stakeholder [52], risk [10], schedule [54], etc.). Joslin et al. [26] find that project success associates with the methodologies (e.g., process, tools, methods, etc.) adopted by the project. And not a single organizational structure is equally effective in all cases. As the organizational context changes over time, to maintain consistency, the project must adapt to its context accordingly. Otherwise, conflicts and inefficiency occur [3], i.e., *one size does not fit all*.

To address the conflicts caused by incompatible fitting to the project's context, previous work suggests thinking holistically. Lehtonen et al. [29] consider the project environment as all measurable spatio-temporal factors when a project is initiated, processed, adjusted, and finally terminated, suggesting that the same factor can have an opposite influence on the projects under a different context.

³<https://incubator.apache.org/guides/graduation.html>

⁴A developer abruptly deleting the widely used, 11-line of *left-pad* code, led to cascading disruption of other OSS projects in *npm* ecosystem.

In the domain of software engineering, Joslin et al. [26] considers project governance to be part of the project context, concluding that project governance can impact the use and effectiveness of project methodologies.

In the context of OSS projects seen as socio-technical systems, contingency theory implies that observing and tracking multiple facets/features of the projects may lead to more effective models of system evolution.

3 HYPOTHESES AND RESEARCH QUESTIONS

Our goal is to build effective models for forecasting ASFI project sustainability. Here, we generate our hypotheses and formulate research questions based on prior work and the pertinent theories.

3.1 Hypotheses

STS theory suggests that publicly observable participation and decentralized contributions to software projects foster sustainable collaborations. The ASFI ecosystem is in a form of a typical STS where the technical activities build a shared code artifact and the social ones mediate knowledge and organizational details to individuals. Since all activities are logged, various socio-technical metrics can be calculated.

Our main hypothesis is that the STS formalism and the full availability of the projects' longitudinal digital traces, will make it possible to build an accurate model of sustainability. According to contingency theory, no single organizational structure is equally effective under all circumstances. Thus, across ASFI projects, we expect to see that the same socio-technical factors may have different contributions to sustainability. Finally, we posit, per contingency theory, that the roles of some social-technical factors may vary over time. Moreover, we expect to see that similar ASFI projects can end with divergent outcomes (graduation or retirement) over time based on actions they have undertaken.

3.2 Research Questions

We formalize the above into our RQs, as follows. The first is a validation of contingency theory hypotheses, that there exist measurable differences between graduated projects and retired projects, along with multiple features. Namely,

RQ₁ Are there significant differences among STS measures, between graduated projects and retired projects?

Next, STS theory holds that project sustainability is associated with social and technical network features. Contingency theory implies there will be multiple such features in play. Thus, we expect that a quantitative temporal model may be fitted well to the available ASFI data, so long as a sufficient number of features and projects are available.

RQ₂ How well can we predict the sustainability based on temporal traces of ASFI projects? And, can we identify the determinants along with their weights and directions?

To make the model useful in practice it needs more than just accurate predictions of outcomes, it also needs to generate timely advice on whether the project should stay the course or implement specific corrective action to improve the graduation trajectory. We formulate this as:

RQ₃ Can we monitor project sustainability status in a continuous manner? When and how should projects react to the monitoring advice?

4 DATA AND METHODS

We collected historical trace data of commits, emails, incubation length, sponsor information, and incubation outcome for 263 ASFI projects, which have available archives of both commits and emails from 03/29/2003 to 10/31/2019 [57]. Among them, 176 projects have already graduated, 46 have retired, and 41 are still in incubation. The latter, projects still in incubation, were not studied in this paper.

We collected the ASFI data from two sources: ASF mailing lists and SVN commits. The mailing list archives are open access and can be accessed through the archive web page, http://mail-archives.apache.org/mod_mbox/. They contain all emails and commits from the project's ASF entry date, and are current. We constructed URLs for individual project files in the ASF incubator as *Project URL*. The project URLs use the pattern: `project-name/(YYYYMM).mbox`. For example, for project *hama*, the full URL is http://mail-archives.apache.org/mod_mbox/hama-dev/201904.mbox. Each such file contains a month of mailing list messages from the project, for the date specified in the URL. Here *dev* stands for 'emails among developers'.

However, we find that many projects, especially these over ten years old, which used SVN, used a bot in the *dev* mailing list to record all commits, thus a message from *dev* is not always an email. Similar emails were sent to the 'commits' mailing list, which, thus, contains some emails. Therefore, we collected both *dev* and *commits* mailing lists files for the 222 graduated or retired ASF Incubator projects through the archive web page⁵.

4.1 Data Pre-processing

ASF manages and records the communications among people by globally assigning an exclusive email name to each developer at the project-level. However, some developers still prefer to use their personal email/name instead of the assigned one, which in turn complicates the identification of distinct developers [61]. We performed de-aliasing for those developers with multiple aliases and/or email addresses, as follows. We first remove titles (e.g., jr.) and common words in the name (e.g., admin, lists, group) from usernames, then we match with both the original order and switched first/last name order whenever names contain exactly one comma to eliminate ambiguous styles. Then we match each developer with her/his multiple email addresses (if any).

Many projects, especially those over ten years old that used SVN, utilized a bot for extensive mailings (empirical evidence shows 26% of popular GitHub projects use bots) [55], thus forming outliers in the dataset. Some broadcast emails are automatically generated by the issue tracking tool (e.g., JIRA), and no developer would reply to them. We eliminated the broadcast messages that no one replied to, and we find many of them are generated by automated tools. We find some developers contributed many commits by directly changing/uploading massive non-source code files (e.g., data, configuration, and image files). Since those can form outliers in the dataset, commits to files with extension data: .json, .xml, .yaml, .yml,

⁵Data and scripts are available: <https://doi.org/10.5281/zenodo.4564072>

.jar; text/configurations: .config, .info, .ini, .txt, .md, and image: .jpg, .gif, .pdf, .png are eliminated.

As result, we identify 21,328 unique contributors (who either committed code or posted emails). Among them, 1,469 only committed code, and 18,205 only posted/replied to discussions without committing code. The remaining 1,654 contributors engaged in both activities. We identify 2,764,309 commits, modifying a total of 404,455 source code files. We collect 879,812 emails, from them we identify 19,859 developers who participated in discussions (by sending or receiving emails). Among them, 19,573 proactively engaged in discussion activities (i.e., sending emails), the remaining 286 developers collaborated in a passive way (only received emails).

4.2 STS and Socio-technical Networks

We use socio-technical networks to anchor our study of OSS STS. Network science approaches have been prominent in studying complex dynamics of OSS projects [5, 49], although the specific definition may vary with domain context [33].

In this paper, we define the projects' socio-technical structure using social (email-based) and technical (code-based) networks, induced from their emails to the mailing lists and commits to source files, as follows. Similar to the approach by Bird et al. [4], we form a social (email) network for each project, at each month, from the communications between developers: a directed edge from developer *A* to *B* exists if *B* has replied to *A*'s post in a thread or if *A* has emailed *B* directly (which is contained in the "in-reply-to" field). The technical (code) collaboration networks are formed for each project, at each month, by including an undirected edge between developer *A* and developer *B* if both developer *A* and *B* has committed to the same coding source file(s) *F* that month (excluding the SVN branch names).

4.3 Features/Metrics of Interest

The socio-technical and project features/variables that we chose for this study have been identified based on our discussion and consideration of the underlying theories. All our data is longitudinal. All metrics are aggregated over monthly intervals, for each project, from the start to the end of its incubation [60]. We started with 29 variables, given their statistics as Supplementary Material.

Variable Selection We used Lasso regression [50] (L1 regularization) to identify a smaller set of 18 linearly independent variables, plus the outcome, described in the following. We used *R*'s library *glmnet* [17] for the Lasso regression, with $\lambda = 0.001$.

Outcome: Graduation Status. Graduation *grad_status* is a binary variable (0='Retired' or 1='Graduated') indicating the projects graduation status in the incubator, as discussed above.

Longitudinal Project Metrics: The number of Active Developers *num_act_devs* is the count of contributors who have been active by either making commits or participating in discussions. Number of commits *num_commits* is the count of source code commits made by all committers in the project. The process of excluding the commits that do not contain source code is described in the Data Section. The number of Emails *num_emails* is the number of emails (including both thread starter emails and reply-to emails). *num_files* is the total number of unique source code files created during the incubation. To measure the continuity of activities, we

define *c_interruption* and *e* as the sum of the time intervals of the top 3 longest interruptions between successive commits and successive emails, respectively. The commit percentage *top_c_fract* and email percentage *top_e_fract* are the percentages of respective activities performed by the top 10% contributors.

Longitudinal Socio-Technical Project Metrics: For each project network, for each month, we constructed the technical and social networks, and from them calculate the number of active nodes, *c_nodes*, and edges *c_edges* in the technical network; *e_nodes* and *e_edges* in the social network. The prefix *c_* in a variable's name indicates it is of the technical (code) network, while the prefix *e_* in a variable's name indicates it is of the social (email) network. Additionally, we calculated the mean degrees *c_mean_degree* and *e_mean_degree* (sum of all nodes' degree divided by the number of nodes) in the technical network and social network, respectively. We calculate the clustering coefficients *c_c_coef*, *e_c_coef* as the number of connected triplets divided by the number of all triplets in the corresponding monthly network. The long-tail-ness *c_long_tail*, *e_long_tail* is calculated as the degree of the 75th percentile of nodes in the network, for the monthly networks, in the technical and social network, respectively. To get a sense of the range and variability in these variables, we show them aggregated over all months and projects in Table 1.

4.4 Models

We needed a modeling approach able to learn and forecast from longitudinal data, have excellent performance, and be interpretable.

4.4.1 LSTM-Based Learning Model. Long short-term memory [25] is a variant of Recurrent Neural Networks (RNNs), designed to learn and model sequential data and is less sensitive to the problem of gradient disappearance and gradient explosion when training on long

Table 1: Statistics of the 176 graduated and 46 retired projects in the ASFI dataset. *c_* and *e_* correspond to technical networks and social networks, respectively.

Statistic	Mean	St. Dev.	5%	95%
<i>grad_status</i>	0.79	0.41	0	1
<i>num_files</i>	1,821.87	3,346.23	122.45	5436.6
<i>num_emails</i>	3,963.12	4,930.54	262.65	12463.6
<i>num_commits</i>	12,451.84	27,373.41	453.8	36359.7
<i>num_act_devs</i>	121.23	119.85	25	415.05
<i>c_interruption</i>	0.20	0.20	0.03	0.60
<i>e_interruption</i>	0.11	0.14	0.01	0.32
<i>top_c_fract</i>	0.65	0.19	0.38	0.94
<i>top_e_fract</i>	0.71	0.11	0.49	0.85
<i>c_nodes</i>	15.44	17.09	2	49.9
<i>c_edges</i>	120.15	276.19	1	531.5
<i>c_c_coef</i>	0.78	0.25	0	1
<i>c_long_tail</i>	10.65	11.93	0	33.85
<i>c_mean_degree</i>	8.14	7.45	1	23.75
<i>e_nodes</i>	113.38	115.07	22	408.15
<i>e_edges</i>	399.23	562.38	47.1	1315.9
<i>e_c_coef</i>	0.43	0.10	0.28	0.58
<i>e_long_tail</i>	10.33	7.15	3	24.95
<i>e_mean_degree</i>	6.07	1.91	3.88	9.62

Figure 2: The descriptive variables between graduated projects (in green, left) and retired projects (in red, right). The corresponding p-value of the Student's t-test is in the brackets, suggesting significant statistical differences exist between them.

sequences. To obtain sequence data for each project, we aggregated historical ASFI records into monthly data, for each month from the incubation start date to the project graduation/retirement from the incubator. We interpret the monthly LSTM output probability as the *graduation forecast*, i.e., the probability of the project eventually graduating.

We prepared the data as follows. We randomly divided the projects into training and test sets in an 8-to-2 ratio. Because we have variables that are of very different magnitudes, and many of those are not normally distributed, thus we choose to use the *MinMaxScaler* function to standardize all prediction variables.

We implemented a 3-layer LSTM model: a 64 neurons LSTM layer, followed by a 0.3 rate drop-out layer, and a dense layer with the *softmax* function to yield the predicted outcome of the classification task (graduate/retire). If the probability of graduating is higher than 0.5 then we predict the project will graduate, otherwise, we predict it will retire. We validated the cutoff choice by examining the distribution of sustainability forecasts. During training, we used a binary cross-entropy as the loss function and *Adam* as the optimizer. Since the length of the temporal data of each project varies, instead of using zero-padding, which could possibly introduce variance to the model, we chose to use a slower but more reliable way by only feeding one training sample at a time. We use the accuracy, precision, recall, and F_1 -measure to evaluate the performance of the LSTM model using the *classification_report* function from the *sklearn* package.

To get the *graduation forecast* for a project at month $m+1$, we cap the project history at month m , i.e., we only use the first m months in the model. We interpret the outcome yield of the LSTM model as the *graduation forecast*. Projects are not being calculated and used in the prediction when the current time exceeds their incubation lengths. We generated graduation forecasts for each month, and thus obtained the *graduation forecast trajectories*. Repeating the above process 10 times, selecting different training/test split each time, produced our error bounds.

4.4.2 LIME-Based Interpretable Model. Black-box deep learning models, like LSTMs, are less ideal for decision making than interpretable approximations of deep learning models. One such approach is the Local Interpretable Model-agnostic Explanations (LIME) [41]. Given a pre-trained model and an input instance, LIME reasons how the black-box model yielded the output, by probing the model along each of the features. LIME yields a magnitude and

a sign (positive or negative) that characterize the contribution of each feature toward explaining the outcome.

Assumption LIME assumes that any complex model is linear (i.e., interpretable) at a local scale. Therefore, given an input instance, LIME first artificially generates large enough samples that are presumably very close to the given sample (by some distance measure). Then, by training on the predictions of those newly generated samples given by the complex model, LIME can locally approximate the complex model using linear models, thereby presenting the coefficients of variables of the input instance.

Procedure We first constructed a LIME explainer using the *RecurrentTabularExplainer* function from the Python LIME package (version 0.2.0). That package was designed for explaining RNN-type models with tabular data. For each project, we use *explain_instance* function, setting the parameter *num_features* to the product of the incubation length and the number of features. In this way, we can obtain the coefficients of all features over all time. The parameter *num_samples* is set to 5,000 (by default), which is empirically sufficient for convergent results. Next, LIME probes the pre-trained LSTM model 5,000 times by feeding it the newly generated samples. LIME uses a similarity/distance function to measure the importance of each new sample on the locality of the instance to be explained. Lastly, LIME fits a weighted linear model dataset, and the explanations all come from the final linear model. Since the LIME framework requires all samples to have the same shape, we divided our projects into several buckets where the projects have at least n months of temporal data in the n -th bucket.

Project-specific vs Overall Modeling We used the LIME results in two modeling ways: *project-specific* level and *overall* level. In the former, we used LIME to obtain the monthly coefficient of each variable and then aggregated them over all months to obtain a project-specific coefficient. In the latter, we aggregated project-specific coefficients over all ASFI projects to obtain the coefficients for each variable over all projects.

5 RESULTS AND DISCUSSIONS

5.1 RQ₁: Graduated vs. Retired Projects

To perform exploratory data analysis, we first contrast ASFI graduated and retired projects along the technical (code-based) and social (email-based) dimensions in our data.

Box-plot comparisons of the incubation months, number of commits, number of emails, number of nodes in the social networks, and number of nodes in the technical networks are shown in Figure 2.

We observe notable differences as follows. The median incubation months of retired projects is significantly higher than of the graduated projects, suggesting that retirement is not an easy decision, and that perhaps necessary time is given to projects to change their trajectories and achieve graduation.

Graduated projects also tend to have more code commits and more email communications, implying that, in terms of criteria for graduation, the ASF community values both technical contribution (as commits) and social communication (as emails), and both of them may be of importance in building a sustainable community. Such results also motivate us to expand our research goals from descriptive data to inferential data with more complex network features.

Across both the social and technical networks, the network size varies for the graduated projects, indicating that projects of any size can be sustainable, and it also suggests that project size can be used as a control in modeling. The notable difference between graduated and retired projects, and the lack of variance in network sizes in the retired projects suggests recruitment difficulties in the latter, exposing them to significant risks as people leave.

Answer to RQ₁: We observe significant differences across multiple key measures between graduated projects and retired projects. Notably, retired projects tend to stay longer in the incubator than graduated ones, and the productivity and diversity of graduated projects are higher compared to retired projects.

5.2 RQ₂: Interpretable Forecasting

Here we present the results of training an LSTM model on our ASFI data, and an LSTM-derived LIME model on the same data, using the methods as described above. We use those models for forecasting by each month into the project the eventual graduation outcome. First, we show the performance curves of the LSTM model, over time, in Figure 3. The overall accuracy, and F1-value, and their standard errors (grey area), for the full model and the model without the socio-technical variables, suggest an excellent and stable predictive performance of the trained LSTM model, with significant contribution from the socio-technical networks. As early as month 8 the accuracy is 93%, and staying above after that.

The total incubation time varies significantly across ASFI projects. While for most projects the incubation time is between 8 months and 25 months, some spend more than 30 months in incubation, while others only 6 months, as shown in the inset plot in Figure 3. Thus, the model’s performance decreases for projects with below 8 and above 25 incubation months, due to insufficient data above and below those values.

Next, we apply LIME to understand and interpret the LSTM model and derive regression-like coefficients for the features in the socio-technical networks. To illustrate how to interpret the results at a project-specific level, we show an example of LIME’s output for a graduated project, ‘Empire-DB’, in Figure 4. Note that the coefficients are aggregated over all incubation months. Looking at the median over all months provides model coefficient stability and avoids emphasizing very small effects which nevertheless dominate in some months. The magnitudes of the coefficients tend to be small

Figure 3: Performance metrics of the full LSTM model across incubation months (top 2 curves), showing the significant contribution of the socio-technical metrics. Curves are plotted using loess. Grey area shows the standard errors. The red vertical line shows 93% accuracy at 8 months of incubation. The inset shows project density over total incubation time.

Figure 4: The coefficients of all variables from a graduated project (‘Empire-DB’), aggregated over all incubation months, showing that LIME delivers stable estimation at the project-level.

because the model predicts probabilities within $[0, 1]$, and there are tens of them.

There, we see that the technical network clustering coefficient c_c_coef is negatively associated with successful graduation. A high clustering coefficient in the technical network of people and source files indicates a high overlap of developers’ activities on the same files. One possible reason for the negative effect introduced by c_c_coef is that, work may not be well distributed among the team members. Another reason might be that the artifact is not well-conceived. Yet a third reason can be that the number of developers on the project is small and they must all “tend to fires” wherever they might be. Interestingly, the fraction of commits (top_c_fract) and emails (top_e_fract) by the top 10% developers who tend to be the influencers in projects are positively associated with graduation, implying that the efforts of the top 10% help sustain the project. Figure 4 also shows the feature coefficients of a selected project in both size and direction across all incubation months, which provides further insight, and confidence in the methods utility.

Next, by only counting the signs of the project-level coefficients across all projects, we can identify whether there is an overall

Figure 5: The overall-level coefficients of all variables of interest (blue is positive, while red is negative to graduation). It shows that some variables have same effect on almost all projects, while others do not.

consistent direction in which that feature is contributing to the prediction. E.g., if a feature is consistently negative to the outcome across all projects, i.e., that feature’s coefficient is negative in most project models, then it has the same, overall negative effect.

In Figure 5, we show the count of aggregate signs of feature coefficients across all projects. There, *blue* indicates a positive effect and *red* a negative one. Visually, if most of the bar is a single color, then that variable has a consistent effect direction, i.e., same sign coefficient, among all projects. Overall, perhaps surprisingly, we find that for almost all projects, the number of nodes in the technical networks *c_nodes* has a negative effect on graduation for minor projects while the number of nodes in the social networks *e_nodes* is positively associated with graduation. This is consistent with prior research findings that communication is more determinant of success and onboarding than coding activities [6]. Other variable appear to have inconsistent effect across projects, e.g., *num_files* and *c_mean_degree*.

Since LIME fits model coefficients for all months, we can also examine the dynamics of feature coefficients. Figure 6 shows that when broken down into 4 intervals, the effect of *num_act_devs* becomes less positive, and less important, over time, perhaps due to the project becoming more stable. Contrariwise, the negative effect of the mean degree in technical networks (*c_mean_degree*), diminishes in latter development, arguably, again, because of increased project stability over time.

Answer to RQ₂: Effective models of project sustainability can be built from tens of socio-technical and project features. Stable and interpretable models can be derived yielding feature coefficients at both project-specific and overall levels. Notably, overall, projects with fewer but more centralized committers and those with more but distributed communicators are more likely to become self-sustainable in the ASF incubator.

5.3 Case Study: Change of Fate

To understand in depth why trajectories may change, with the help of our interpretable model we randomly selected three qualitatively different example projects (2 graduated and 1 retired), showing up-

Figure 6: The overall-level coefficient of two selected variables: mean degree in technical network (in red) and the number of active developers (in green) in different incubating quarters.

Figure 7: The graduation forecast of the marginal projects. *Commonsrdf* (ID: 82, in green) and *Etc* (ID: 103, in blue) are graduated projects that almost failed while retired project *Ariatosca* (ID: 256, in red) almost succeeded.

or down-turn points in their sustainability trajectories, as illustrated in Figure 7.

§1 In project *Commonsrdf* (ID: 82), the graduation forecast starts high but experiences a downturn in the first half of incubation, and then it rises again. Our model identifies the lack of both email and commit activities are associated with the downturn. There is also an overall decreasing trend in the number of active developers (from 27 to 11, then to 4 eventually). To investigate why their promising trajectory changed, and then changed again, we looked through their discussions on the mailing list.

In the month with the lowest graduation forecast, the project released a major version of their software, following which the committers showed less activity. We also noticed that the ASF incubator Project Management Committee (PMC) routinely sent an email to the project asking for the periodically report⁶, but such reports were not requested for more than 2 months after the release. Lastly, we saw email arguments on the technical direction for the project’s future development. Unsurprisingly the project was going to fail if no one observes such situation and takes action to intervene. However, an email from one of the major contributors, *Dev1*, appears to have initiated the turning point. Below we quote it and the ensuing discussion.

⁶<https://cwiki.apache.org/confluence/display/INCUBATOR/Reports>

Dev₁ “Folks, I’ve seen very little traffic for the last few months. . . I am concerned that there is perhaps no longer a viable community around this podling. . .”, and seriously asked “Do people still think this project has/can build the momentum to move forwards towards graduation?”

Dev₂ “. . . we lost one of the main pillars of this projects . . . So our mentors are right. We’re in a situation where the project has no momentum at all, and honestly I have no idea what’s best to do. . .”

Dev₃ “. . . there was a small group of 5 core committers to begin with. As of right now, the number is 22. We’ve actually done pretty well. . .”

Then, *Dev₄* responded with an email titled ‘Values and Terms’ and suggested a technical directions, with detailed reasoning: “. . . if you actually tried to use this (algorithm) would (i) hurt speed, and (ii) hurt the perception of speed. . .” and clearly states that “I’d be inclined to go another step further and add a generic parameter. . .”

After this discussion, the community became more engaged and increased some activities, and the community felt more confident about the upcoming routine report. Eventually, the project graduated in the end (our forecast went up to 80%).

§2 In project *Etch* (ID: 103) there was a major depletion of senior developers after a milestone was reached in the middle of incubation. Then commits stopped for almost one year. Our graduation forecast reflects that: it dropped from 79% to only 19%.

After a long time being inactive, one developer sent a broadcast email titled ‘Future of Etch’. Many developers participated in the discussion thread, with seeming agreement that their project is either to be retired or changes are needed. The project mentor brought up the lack of diversity as a possible cause for stagnation, since all developers came from the same company. Some developers concurred, and they feel “...continual pressure to wrangle new committers...”, and consider that the ASF values “...extroverted tendencies of the committers rather than the merits of technology...”.

Eventually, the contributors reached an agreement that the project technology is and will be valuable in the future. Among them, one developer stated that they “...do not want to see it retired...” The developers then made a list of future objectives, and worked to make the community thriving again.

§3 Project *Ariatosca* almost succeeded, but eventually failed (the graduation forecast dropped from 96% to 47%). We find that the major reason is that all senior contributors left the project due to their busy day work. At the very end of the project, there were new(er) developers who wanted to contribute to the project. However, since newcomers could only contribute by creating Pull Requests (PRs), and PRs required a senior committer to accept and merge the code changes, they could not submit their code. Attempts at setting meetings with the original developers failed due to busy schedules, and the project was eventually retired.

5.4 RQ₃: Actionable Recommendation

We start with a caveat. Precise actionable models require interventions and randomized experiments. Our models are not based on such experiments, and any actionability we derive from them must, therefore, be less powerful than those. At best, our experiments can be considered *natural experiments*, a subclass of quasi-experiments where the class assignment is not controlled by the

experimenter. Thus, any interventions we suggest here must be validated experimentally in order to avoid large uncertainties in outcomes. With that caveat in mind, we sought to answer, to the best of our experimental methods, the following intervention question: “What action should a project take and when?”, in order to increase its graduation forecast in our model.

Here we propose a pragmatic, laissez-faire-unless-needed prospective strategy: to continuously monitor the graduation forecast for significant downturns, and if detected, suggest interventions that may improve the forecast. We deconstruct the intervention question above into two parts: 1) What is a significant downturn? and 2) how to interpret the variables and coefficients in our fitted model into actions. For the following, recall that our interpretable sustainability model gives a graduation forecast from the historical project trace data and the socio-technical project structure, available until that time. It also yields the coefficient of each significant model feature along with its direction for every month.

§1 Identifying significant downturns. We want to identify downturns that dominate any naturally occurring noise or jitter in the forecast. We looked across all projects for how long it takes for a forecast to bounce-up from any downturn. We found that 149 out of the 222 projects (about 67.1%) experienced an upturn or downturn event of a 0.3 magnitude in their forecast probability curve. Of those projects that have not experienced downturns, the average incubation time was 16 months, versus 28 months for projects that have had downturns, implying an association between downturns and length of stay in the incubator. We note that the upturn events tend to be associated with ASF mentors’ intervention. This is consistent with one of the responsibilities of ASF mentors: to watch over the projects they are mentoring and keep them on track. However, downturn events do not correlate similarly with extrinsic events; we think they tend to be more related to projects’ internal events, e.g., ebbing social interest, core committers leaving, etc. Figure 8 shows that the median bounce-up time is about 2.5, and, respectively, 3.5 months for drops of 0%–5%, and, respectively, > 5% in the graduation forecast. We also noted that graduated projects seem to bounce-up faster than retired projects. So, we use the median of the latter as the baseline, and define a drop in the forecast of greater than 5% over a period of one or two months to indicate a significant downturn event. This ad hoc approach, while an approximation, is a natural signal processing way to account for inherent uncertainty in the signal.

§2 Actions that improve the forecast. Our model yields real-numbered coefficients for each significant feature in each month, which we aggregate for stability over all months, and obtain the medians as in Fig. 4. Increasing the values of features with positive coefficients and/or decreasing the values of the negative coefficient features results in an increase of the graduation forecast. Thus, once a month or two with a significant downturn is detected, the project developers can look at the most positive and most negative medianed significant features from the fitted model and consider increasing, respectively decreasing, them. This is an approximation and is valid to the extent that the median is representative of the values over all months, which is more the case in the earlier months than the latter ones, see Figure 6. The earlier months are the ones

Table 2: Positive Actions for Each Feature.

	Positive Action (+)
num_act_dev	Contribute frequently; Advertise, Recruit.
num_emails	Reach out; Ask questions; Encourage communication.
num_commits	Commit frequently; commit smaller; use CI.
num_files	Split files; refactor code; encourage modularity.
c_interruption	Go on vacation often; contribute in bursts.
e_interruption	Email seldom; discourage discussion.
top_e_fract	Encourage core emailers to respond more.
top_c_fract	Core contributors commit exclusively.
c_nodes	Establish technical mentorship; encourage commits.
c_edges	Commit to same files as others; document code well.
c_c_coef	Encourage collaborations, pair programming.
c_mean_degree	Encourage commits by minor contributors.
c_long_tail	Mentor collaborations with newcomers.
e_nodes	Mentor low communicators.
e_edges	Reply to questions; ask questions.
e_c_coef	Encourage non-hierarchical communications.
e_mean_degree	Communicate with minor emailers.
e_long_tail	Foster communication-heavy culture.

we care more about in practice, as we care about being most helpful to nascent projects.

Some of the socio-technical and project features may be difficult to interpret in practice. To aid with this step, we suggest a project should compile an *action table* to summarize possible actions that may positively (or negatively in the reverse way) change the value of model features. (Multiple actions may affect the same feature.) In Table 2 we provide one such action table: a mapping between all of our model features and a non-exhaustive set of actions that we identified as likely to move each variable in the positive direction. (The negative actions are not shown, but are complements of those.) E.g., the `e_c_coef`, which counts the number of triangles in the social network, can be increased by increasing emails to everyone and not just the prominent developers or thread starters; conversely, communicating hierarchically in a tree-like fashion would decrease `e_c_coef` as it will eliminate triangles.

§3 Ecosystem and project-specific fine-tuning. Our strategy can be further fine-tuned in several ways. First, there are common patterns in the graduation forecasts over all projects. Figure 9 shows that for graduated projects (in green) there is an apparent upward trend in the forecast in the first 6 months, suggesting that the early-stage development deserves more attention from project managers. From month 6 to month 12 we do not observe a significant change, and the decreasing variance also tends to support such an argument. However, we find increasing variance in months 12 through 18. One possible reason is that many graduated projects achieve their milestone in that period, and slow further commits and discussions, thus lowering the graduation forecast. These considerations can be taken into account during forecast monitoring, with more frequent monitoring chosen or increased attention paid at times around milestones and releases. The ASF committees can also ask for more frequent project reports, during the first 6 months and 12 months of incubation, as project reports were seen to be an incentive to productivity in our case study.

Figure 8: Bounce-up after downturn in graduation forecasts for graduated (green) and retired projects (red).

We recognize that some socio-technical elements are more difficult to change compared to others. This is, project-specific, in that some projects can easier modify some features than can other projects. E.g., if the interpretable model suggests increasing `c_edges` and decreasing `c_interruption` this may be easier done in smaller projects than larger ones due to the difficulty of influencing many people at once. Thus these action tables should ideally be project-specific, designed and updated as the project evolves.

Mechanistically, here are some possible reasons for why features, e.g., number of commits/files, may have different effects across projects. First, having much more than the usual amount of code changes and file touches in a given period of time may be indicative of refactoring or even change of direction for a project. This can be the result of a project pivoting which may lead to longer stay in the incubator and possible retirement. Second, increases in the number of files may result in insufficient human resources to handle them (e.g., bugs, documentation, mentoring, and training), resulting in code quality issues and technical debt, especially for smaller projects. This can lead to loss of interest in the small community and eventual disengagement. Third, for overleveraged projects, increase in commits and/or file touches may mean not having time for proper communication, which can slow progress and result in community collapse. In effect, the coefficients of the same factor can vary under different project contexts.

Lastly, a word of caution. In terms of expectations that including more features to intervene on, may lead to faster bounce-up, we note that empirical evidence shows socio-technical features have ways of getting correlated pairwise over time in the same system. Thus, even if not correlated, the effect of increasing multiple features simultaneously may not be additive.

Answer to RQ₃: Our strategy can be used as a monitoring tool that feeds suggestions into developers' decision process. Project-specific features can be selected from the suggestions for intervention when experiencing downturns. The algorithm can be made bespoke by introducing more frequent monitoring around releases/milestones.

5.4.1 Actionable Strategy Example. Here we apply our strategy on a project from Figure 7: *Commonsrdf*, and show specific recommendations following a detected downturn.

Figure 9: Graduation forecasts for all graduated (green) and retired projects (red) over the first 24 incubation months.

While monitoring project *Commonsrdf*⁷ we would have observed a significant downturn at months 4 and 5 (> 5% drop), see Figure 7. For that project, our interpretable model yields as the 3 features with the highest positive medians overall: *top_c_fract*, *top_e_fract*, and *e_nodes*; the three with the most negative median are: *c_c_coef*, *c_interruption*, and *c_nodes*. Consulting Table 2, it calls for the project to increase the commit and email contributions by core developers and encourage more committers to communicate. It also calls for the project to decrease collaborations, decrease commit interruptions, and decrease the number of developers that commit. A continuous integration may also be recommended to this project, to keep people on track with smaller, more frequent commits.

What actually happened in the project is that that initial period of downturn was missed; as we saw in our case study for this project, the project manager sent an email titled 'Anybody there?' in month 8, when productivity was already significantly reduced. Using our strategy, the downturn could have been identified and possibly avoided 3 months sooner.

6 THREATS TO VALIDITY AND CONCLUSION

Threats. First, our commit and email data is from only hundreds of projects ASF incubator projects. Thus, generalizing the implications beyond ASF, or even beyond the ASF Incubator projects carries potential risks. Expanding the dataset beyond ASF, e.g., with additional projects from other open-sourced incubator projects can lower this risk. Second, we do not consider communications other than through the ASF mailing lists. However, ASF's policies and regulations insist on the use of mailing lists, which lowers this risk. Lastly, interpreting deep learning models is still an art, and LIME models are approximations. They may be particularly sensitive to correlated features. We lower such risk by eliminating correlated variables before training. Taking the actionable suggestions given in this paper may result in changes of more than one variable, e.g., increase the active developers may also increase the number of commits.

Conclusion. Understanding why many nascent projects have failed may help others improve their individual practice, organizational management, and institutional structure. Here we showed

that quantitative network science approaches combined with state-of-the-art AI methods can effectively model ASF incubator project sustainability, from a novel longitudinal dataset of socio-technical contributions in ASF projects, more narrow in scope than general OSS projects but with extrinsic graduation/sustainability labels. We also demonstrated the combined power of mixed methods: both quantitative and qualitative studies. Through case studies, we identified specific reasons for success and failure of projects, complementing our models. Finally, we developed a strategy for translational use of the models in practice. Our methods make it straight forward to track a project's trajectory as it progresses toward sustainability, and even offer advice for correcting trajectories upwards. Future work is needed to offer validation of this or similar strategies experimentally.

ACKNOWLEDGEMENTS

We are grateful to the National Science Foundation for funding this project, under Grant #2020751. This work was also partially supported by the National Natural Science Foundation of China under Grant #61973273, and by the Zhejiang Provincial Natural Science Foundation of China under Grant #LR19F030001. We greatly thank the FSE 2021 reviewers for their constructive comments.

REFERENCES

- [1] Chintan Amrit and Jos Van Hillegersberg. 2010. Exploring the impact of socio-technical core-periphery structures in open source software development. *Journal of information technology* 25, 2 (2010), 216–229.
- [2] Erling S Andersen, Anders Dysvik, and Anne Live Vaagaasar. 2009. Organizational rationality and project management. *International Journal of Managing Projects in Business* (2009).
- [3] Donald W Barclay. 1991. Interdepartmental conflict in organizational buying: The impact of the organizational context. *Journal of Marketing Research* 28, 2 (1991), 145–159.
- [4] Christian Bird, Alex Gourley, Prem Devanbu, Michael Gertz, and Anand Swaminathan. 2006. Mining email social networks. In *Proceedings of the 2006 international workshop on Mining software repositories*. 137–143.
- [5] Christian Bird, Nachiappan Nagappan, Harald Gall, Brendan Murphy, and Premkumar Devanbu. 2009. Putting it all together: Using socio-technical networks to predict failures. In *2009 20th International Symposium on Software Reliability Engineering*. IEEE, 109–119.
- [6] Casey Casalnuovo, Bogdan Vasilescu, Premkumar Devanbu, and Vladimir Filkov. 2015. Developer onboarding in GitHub: the role of prior social links and language experience. In *Proceedings of the 2015 10th joint meeting on foundations of software engineering*. 817–828.
- [7] Narciso Cerpa, Matthew Bardeen, Barbara Kitchenham, and June Verner. 2010. Evaluating logistic regression models to estimate software project outcomes. *Information and Software Technology* 52, 9 (2010), 934–944.
- [8] Theodore Chaikalis and Alexander Chatzigeorgiou. 2014. Forecasting java software evolution trends employing network models. *IEEE Transactions on Software Engineering* 41, 6 (2014), 582–602.
- [9] Jailton Coelho and Marco Tulio Valente. 2017. Why modern open source projects fail. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*. 186–196.
- [10] Terry Cooke-Davies. 2002. The "real" success factors on projects. *International journal of project management* 20, 3 (2002), 185–190.
- [11] Kevin Crowston, James Howison, and Hala Annabi. 2006. Information systems success in free and open source software development: Theory and measures. *Software Process: Improvement and Practice* 11, 2 (2006), 123–148.
- [12] Kevin Crowston and Ivan Shamshurin. 2017. Core-periphery communication and the success of free/libre open source software projects. *Journal of Internet Services and Applications* 8, 1 (2017), 10.
- [13] Leticia Duboc, Stefanie Betz, Birgit Penzenstadler, Sedef Akinli Kocak, Ruzanna Chitchyan, Ola Leifler, Jari Porras, Norbert Seyff, and Colin C Venters. 2019. Do we really know what we are building? Raising awareness of potential Sustainability Effects of Software Systems in Requirements Engineering. In *2019 IEEE 27th International Requirements Engineering Conference (RE)*. IEEE, 6–16.
- [14] Nicolas Ducheneaut. 2005. Socialization in an open source software community: A socio-technical analysis. *Computer Supported Cooperative Work (CSCW)* 14, 4

⁷http://mail-archives.apache.org/mod_mbox/commonsrdf-dev/

- (2005), 323–368.
- [15] Juan C Dueñas, Félix Cuadrado, Manuel Santillán, José L Ruiz, et al. 2007. Apache and Eclipse: Comparing open source project incubators. *IEEE software* 24, 6 (2007), 90–98.
 - [16] Gerhard Fischer and Thomas Herrmann. 2011. Socio-technical systems: a meta-design perspective. *International Journal of Sociotechnology and Knowledge Development (IJSKD)* 3, 1 (2011), 1–33.
 - [17] Jerome Friedman, Trevor Hastie, and Rob Tibshirani. 2009. glmnet: Lasso and elastic-net regularized generalized linear models. *R package version 1*, 4 (2009).
 - [18] Jonas Gamalielsson and Björn Lundell. 2014. Sustainability of Open Source software communities beyond a fork: How and why has the LibreOffice project evolved? *Journal of Systems and Software* 89 (2014), 128–145.
 - [19] Bahar Gezici, Nurseda Özdemir, Nebi Yılmaz, Evren Coşkun, Ayça Tarhan, and Oumout Chouseinoglou. 2019. Quality and Success in Open Source Software: A Systematic Mapping. In *2019 45th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*. IEEE, 363–370.
 - [20] Amir Hossein Ghapanchi. 2015. Predicting software future sustainability: A longitudinal perspective. *Information Systems* 49 (2015), 40–51.
 - [21] Amir Hossein Ghapanchi, Aybuke Aurum, and Graham Low. 2011. A taxonomy for measuring the success of open source software projects. *First Monday* 16, 8 (2011).
 - [22] Mohammad Gharehyazie, Daryl Posnett, Bogdan Vasilescu, and Vladimir Filkov. 2015. Developer initiation and social interactions in OSS: A case study of the Apache Software Foundation. *Empirical Software Engineering* 20, 5 (2015), 1318–1353.
 - [23] Jesús M González-Barahona, Luiz Lopez, and Gregorio Robles. 2004. Community structure of modules in the Apache project. In *Proceedings of the 4th International Workshop on Open Source Software Engineering*. IET, 44–48.
 - [24] Thomas Herrmann, Marcel Hoffmann, Gabriele Kunau, and Kai-Uwe Loser. 2004. A modelling method for the development of groupware applications as socio-technical systems. *Behaviour & Information Technology* 23, 2 (2004), 119–135.
 - [25] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
 - [26] Robert Joslin and Ralf Müller. 2016. The impact of project methodologies on project success in different project environments. *International Journal of Managing Projects in Business* (2016).
 - [27] Bakhtiar Khan Kasi. 2014. Minimizing software conflicts through proactive detection of conflicts and task scheduling. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*. 807–810.
 - [28] Jennifer W Kuan. 2001. Open source software as consumer integration into production. Available at SSRN 259648 (2001).
 - [29] Päivi Lehtonen and Miia Martinsuo. 2006. Three ways to fail in project management and the role of project management methodology. *Project Perspectives* 28, 1 (2006), 6–11.
 - [30] Hsiu-Fen Lin and Gwo-Guang Lee. 2006. Effects of socio-technical factors on organizational intention to encourage knowledge sharing. *Management decision* (2006).
 - [31] Suresh Marru, Lahiru Gunathilake, Chathura Herath, Patanachai Tangchaisin, Marlon Pierce, Chris Mattmann, Raminder Singh, Thilina Gunaratne, Eran Chinthaka, Ross Gardler, et al. 2011. Apache airavata: a framework for distributed applications and computational workflows. In *Proceedings of the 2011 ACM workshop on Gateway computing environments*. 21–28.
 - [32] Nora McDonald and Sean Goggins. 2013. Performance and participation in open source software on github. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. 139–144.
 - [33] Andrew Meneely and Laurie Williams. 2011. Socio-technical developer networks: Should we trust our measurements?. In *Proceedings of the 33rd International Conference on Software Engineering*. 281–290.
 - [34] Vishal Midha and Prashant Palvia. 2012. Factors affecting the success of Open Source Software. *Journal of Systems and Software* 85, 4 (2012), 895–905.
 - [35] Audris Mockus, Roy T Fielding, and James D Herbsleb. 2002. Two case studies of open source software development: Apache and Mozilla. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 11, 3 (2002), 309–346.
 - [36] Marc Palyart, Gail C Murphy, and Vaden Masrani. 2017. A study of social interactions in open source component use. *IEEE Transactions on Software Engineering* 44, 12 (2017), 1132–1145.
 - [37] JH Piggott. 2013. Open source software attributes as success indicators. *Univ. of Twente* (2013).
 - [38] Aniket Potdar and Emad Shihab. 2014. An exploratory study on self-admitted technical debt. In *2014 IEEE International Conference on Software Maintenance and Evolution*. IEEE, 91–100.
 - [39] Cobra Rahmani and Deepak Khazanchi. 2010. A study on defect density of open source software. In *2010 IEEE/ACIS 9th International Conference on Computer and Information Science*. IEEE, 679–683.
 - [40] Uzma Raja and Marietta J Tretter. 2012. Defining and evaluating a measure of open source project survivability. *IEEE Transactions on Software Engineering* 38, 1 (2012), 163–174.
 - [41] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why should I trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1135–1144.
 - [42] Peter C Rigby and Ahmed E Hassan. 2007. What can oss mailing lists tell us? a preliminary psychometric text analysis of the apache developer mailing list. In *Fourth International Workshop on Mining Software Repositories (MSR'07: ICSE Workshops 2007)*. IEEE, 23–23.
 - [43] Warren Sack, Françoise Détiennne, Nicolas Ducheneaut, Jean-Marie Burkhardt, Dilan Mahendran, and Flore Barcellini. 2006. A methodological framework for socio-cognitive analyses of collaborative design of open source software. *Computer Supported Cooperative Work (CSCW)* 15, 2-3 (2006), 229–250.
 - [44] Charles M Schweik and Robert C English. 2012. *Internet success: a study of open-source software commons*. MIT Press.
 - [45] Adrian Smith and Andy Stirling. 2007. Moving outside or inside? Objectification and reflexivity in the governance of socio-technical systems. *Journal of Environmental Policy & Planning* 9, 3-4 (2007), 351–373.
 - [46] Igor Steinmacher, Tayana Conte, Marco Aurélio Gerosa, and David Redmiles. 2015. Social barriers faced by newcomers placing their first contribution in open source software projects. In *Proceedings of the 18th ACM conference on Computer supported cooperative work & social computing*. 1379–1392.
 - [47] Igor Steinmacher, Marco Aurelio Graciotto Silva, Marco Aurelio Gerosa, and David F Redmiles. 2015. A systematic literature review on the barriers faced by newcomers to open source software projects. *Information and Software Technology* 59 (2015), 67–85.
 - [48] Chandrasekar Subramaniam, Ravi Sen, and Matthew L Nelson. 2009. Determinants of open source software project success: A longitudinal study. *Decision Support Systems* 46, 2 (2009), 576–585.
 - [49] Didi Surian, Yuan Tian, David Lo, Hong Cheng, and Ee-Peng Lim. 2013. Predicting project outcome leveraging socio-technical network patterns. In *2013 17th European Conference on Software Maintenance and Reengineering*. IEEE, 47–56.
 - [50] Robert Tibshirani. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* 58, 1 (1996), 267–288.
 - [51] Eric Trist. 1981. *The evolution of socio-technical systems: A conceptual framework and an action research program*. Ontario Ministry of Labour.
 - [52] J Rodney Turner and Ralf Müller. 2004. Communication and co-operation on projects between the project owner as principal and the project manager as agent. *European management journal* 22, 3 (2004), 327–336.
 - [53] Marat Valiev, Bogdan Vasilescu, and James Herbsleb. 2018. Ecosystem-level determinants of sustained activity in open-source projects: A case study of the PyPI ecosystem. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 644–655.
 - [54] Stephen Wearne and AAR Stanbury. 1989. A study of the reality of project management: WG Morris and GH Hough, John Wiley, UK (1987) £ 29.95, ISBN 0471 915513 pp 295. *International Journal of Project Management* 7, 1 (1989), 58.
 - [55] Mairieli Wessel, Bruno Mendes De Souza, Igor Steinmacher, Igor S Wiese, Ivanilton Polato, Ana Paula Chaves, and Marco A Gerosa. 2018. The power of bots: Characterizing and understanding bots in oss projects. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–19.
 - [56] Jing Wu, Khim-Yong Goh, and Qian Tang. 2007. Investigating success of open source software projects: A social network perspective. *ICIS 2007 Proceedings* (2007), 105.
 - [57] L. Yin, Z. Zhang, Q. Xuan, and V. Filkov. 2021. Apache Software Foundation Incubator Project Sustainability Dataset. In *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. IEEE Computer Society, Los Alamitos, CA, USA, 595–599. <https://doi.org/10.1109/MSR52588.2021.00081>
 - [58] Marcelo Serrano Zanetti. 2012. The co-evolution of socio-technical structures in sustainable software development: Lessons from the open source software communities. In *2012 34th International Conference on Software Engineering (ICSE)*. IEEE, 1587–1590.
 - [59] Marcelo Serrano Zanetti, Ingo Scholtes, Claudio Juan Tessone, and Frank Schweitzer. 2013. The rise and fall of a central contributor: Dynamics of social organization and performance in the gentoo community. In *2013 6th International Workshop on Cooperative and Human Aspects of Software Engineering (CHASE)*. IEEE, 49–56.
 - [60] Feng Zhang, Ahmed E Hassan, Shane McIntosh, and Ying Zou. 2016. The use of summation to aggregate software metrics hinders the performance of defect prediction models. *IEEE Transactions on Software Engineering* 43, 5 (2016), 476–491.
 - [61] Jiaxin Zhu and Jun Wei. 2019. An empirical study of multiple names and email addresses in oss version control repositories. In *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. IEEE, 409–420.