Contents lists available at ScienceDirect

The Journal of Systems and Software

journal homepage: www.elsevier.com/locate/jss

GAP: Forecasting commit activity in git projects



^a Software Engineering Lab, University of Mons Avenue Maistriau 15, B-7000 Mons, Belgium
^b Software Engineering and Technology group, Eindhoven University of Technology Groene Loper 5, AE 5612, Eindhoven, the Netherlands

^c Lab On Reengineering (LORE - Ansymo), University of Antwerp Middelheimlaan 1, Antwerpen B-2020, Belgium

ARTICLE INFO

Article history: Received 12 November 2019 Revised 22 February 2020 Accepted 9 March 2020 Available online 10 March 2020

Keywords: Git Commit activity Developer abandonment Distributed software development Prediction model

ABSTRACT

Abandonment of active developers poses a significant risk for many open source software projects. This risk can be reduced by forecasting the future activity of contributors involved in such projects. Focusing on the commit activity of individuals involved in git repositories, this paper proposes a practicable probabilistic forecasting model based on the statistical technique of survival analysis. The model is empirically validated on a wide variety of projects accounting for 7528 git repositories and 5947 active contributors. We found that a model based on the last 20 observed days of commit activity per contributor provides the best concordance. We also found that the predictions provided by the model are generally close to actual observations, with slight underestimations for low probability predictions and slight overestimations for higher probability predictions. This model is implemented as part of an open source tool, called GAP, that predicts future commit activity.

© 2020 Elsevier Inc. All rights reserved.

1. Introduction

Software development is inherently a social activity, especially in open source software (OSS) development, as OSS projects are developed and maintained by communities of interacting contributors (Crowston and Howison, 2005). The evolution of OSS projects depends on the presence of volunteers or paid contributors, i.e., company employees supporting the project, that will take on the maintenance effort (Riehle et al., 2014). However, company support in OSS is rather scarce (Dias et al., 2018), requiring the help of volunteers to maintain OSS projects. The risk of contributors becoming inactive can be detrimental for the project as less than 50% of such abandoned projects find new contributors to take on their maintenance (Avelino et al., 2019). It is therefore important to retain existing contributors to minimize such risks, as well as to prevent the knowledge loss that occurs when important contributors abandon the projects they are involved in (Rigby et al., 2016).

The topic of contributor retention has been extensively studied by the research community. Most studies performed an *a posteriori* analysis and investigated the effects of contributor abandonment on the project sustainability over time, revealing that contributor

* Corresponding author.

retention is one of the many factors that plays a role in a project's sustainability over time (Gamalielsson and Lundell, 2014; Steinmacher et al., 2016; Constantinou and Mens, 2017a; Fronchetti et al., 2019; Miller et al., 2019). However, OSS contributors and community managers can benefit from an *a priori* approach that will identify contributors at risk of abandoning the project early. Such an approach would consist of monitoring the activity of contributors to detect probable future abandonments, and to mitigate the associated risks as early as possible, e.g.,by engaging with the contributor, or by identifying other qualified contributors to take over the work.

Existing project monitoring solutions provide useful information about a project's contributor community, e.g., Bitergia's analytics service¹ analyses project community activity, performance, diversity, size and demography. Although these analytics are highly relevant to get insight in the past and current project activity, they cannot be utilized to assess the risk of losing existing contributors in the near future. Currently, there is no methodology in the literature to predict when the next contributor activity will take place, neither automated tool support to perform such an analysis. This paper fills that gap by providing a method to forecast the future development (commit) activity of OSS contributors





E-mail addresses: alexandre.decan@umons.ac.be, alexandre.decan@lexpage.net (A. Decan), e.constantinou@tue.nl (E. Constantinou), tom.mens@umons.ac.be (T. Mens), henrique.rocha@uantwerp.be (H. Rocha).

¹ bitergia.com/bitergia-analytics/.

on git repositories. By doing so, our work enables community managers to properly assess abandonment risks and tackle such issues before they become detrimental to the project continuation.

We propose a simple probabilistic forecasting model that relies on the previous and recent contributions of individual contributors to estimate the elapsed amount of time until their next contribution. We validate the forecasting model on a wide range of diverse projects. To this end, we use a large collection of git repositories involving 5947 distinct software developers contributing to 7528 software repositories, covering a wide range of projects in terms of longevity (short/long-lived), technical size (number of commits) and community size (small/large groups of contributors). The model is operationalised through the Git Activity Predictor (GAP), an open source tool that supports different output formats and visualisations of the predictions.

The remainder of this paper is structured as follows. Section 2 presents the related work. Section 3 introduces the probabilistic activity forecasting model. Section 4 validates the model through an empirical case study on Cargo. Section 5 presents the GAP tool implementing this model. Section 6 discusses the benefits that GAP can bring to OSS project communities, possible limitations of the underlying model, and presents potential directions of future work. Section 7 discusses the threats to validity. Finally, Section 8 concludes.

2. Related work

The evolution of OSS projects depends on the presence of volunteers or paid contributors, i.e., company employees supporting the project, that will take on the maintenance effort (Riehle et al., 2014). Voluntary work is an important asset playing a big role in OSS project long term survival and continuity. It is essential both to attract newcomers and retain developers to maintain a critical mass of core developers in a project (Qureshi and Fang, 2011). Coelho and Valente reported that 41% of failed open source projects cited a reason related to a lack of time or interest of the developer team (Coelho and Valente, 2017).

Many studies have focused on developer attraction and retention. Steinmacher et al. conducted a literature review identifying barriers encountered by newcomers to OSS projects and proposed a classification of these barriers based on socio-technical attributes, e.g.,technical knowledge, experience, and social interactions (Steinmacher et al., 2015). In a follow-up work, they presented FLOSScoach, a dedicated portal created to support newcomers to OSS projects (Steinmacher et al., 2016). Fronchetti et al. found through a study of 450 OSS projects that the most influencing factors to attract contributors are the number of stars, the time to merge pull requests, and the number of programming languages used (Fronchetti et al., 2019).

Lin et al. analysed the evolution of industrial OSS projects to identify how the duration and type of contributions affect developer turnover (Lin et al., 2017). They found that developers that remain active for longer periods of time mainly focus on contributing to the source code and on modifying existing files instead of creating new ones. Iaffaldano et al. conducted interviews and analysed the reasons for contributors taking temporary breaks (Iaffaldano et al., 2019). Their analysis suggests that longer temporary breaks can lead a contributor to a permanent abandonment state. Constantinou and Mens investigated socio-technical factors in relation to developer retention in the RubyGems and npm ecosystems (Constantinou and Mens, 2017a). They found that both social and technical activities influence developer abandonment, and that socially active developers are less likely to abandon.

Developer retention has been shown to influence project sustainability. The tight relationship between an author and its contributed source code makes software development susceptible to knowledge loss when authors leave the project and abandon their code (Rigby et al., 2016). Gamalielsson and Lundell showed that the long-term involvement of developers is favourable for a project's sustainability (Gamalielsson and Lundell, 2014). Constantinou and Mens studied the Ruby ecosystem on GitHub and found a high impact of contributor abandonment on the sustainability of Ruby projects (Constantinou and Mens, 2017b). Avelino et al. studied the abandonment of open source projects by their core developers and the motivation and difficulties faced by new core developers taking over these projects (Avelino et al., 2019). They report that only 41% of the studied projects recovered after all their core developers abandoned them. Miller et al. looked at the reasons why established open source contributors disengage. They suggest that data-driven systems could be developed to help identify groups at risk (Miller et al., 2019).

These studies reveal the need for project community monitoring tools to predict which individuals are more likely to abandon soon, and to mitigate such risks early. Our work aims to address this need, by proposing an activity prediction model that can be used as a basis to assess the likelihood of contributor abandonment.

Although there are studies on the analysis of developer contributions in open-source projects, these studies focus mainly on measuring or characterising developer contributions, rather than on when these contributions will occur. For example, Amor et al. proposed to characterize the complete development activity by combining data and metrics from different data sources (versioning system, mailing list, issue tracker, etc.) to measure the total effort invested in a project (Amor et al., 2006). This idea was further extended by Gousios et al. who proposed a model delivering accurate developer contribution measurements based on a combination of traditional contribution metrics and data mined from software repositories (Gousios et al., 2008). Mockus et al. studied how much effort remains to be spent on a specific software project and how that effort will be distributed over time. They proposed a model to predict software project effort, schedule and defects so as to plan resource allocation (Mockus et al., 2003). Weicheng et al. observed that developers often do not commit for a long time. They explored developer commit habits and their relationship with file version evolution on eight GitHub projects using four metrics (Weicheng et al., 2013), including the time interval between two commits. They showed that the average time between commits ranges from 2 to 5 days, depending on the number of changed files. They concluded that knowing the habits of developers is helpful to make a better work plan.

Our work is inspired by the above studies, but focuses on forecasting the future activity of developers in terms of commit activity. The model we propose is based on the technique of survival analysis (Aalen et al., 2008). This technique analyses "time to event" data with the aim to estimate the survival rate of a given population. It has been used in other empirical software engineering studies, e.g.,to estimate the survival of open source projects over time (Samoladas et al., 2010), to analyse the use and removal of functions in PHP code (Kyriakakis and Chatzigeorgiou, 2014), to analyse the survival of database access libraries in Java code (Decan et al., 2017), to compare the release cycle of reusable components in software ecosystems (Decan et al., 2019), to analyse retention of developers in software ecosystems (Constantinou and Mens, 2017a; Lin et al., 2017), or to understand how contributors join the core team of a FLOSS project (Bird et al., 2007).

3. Probabilistic activity forecasting model

The goal of this paper is to provide a practicable forecasting model that estimates the elapsed amount of time until the next contribution of a git author. Such a model relies on recent information about previous commits of individual contributors to estimate the expected amount of time until their next activity will take place. In Section 3.1 we evaluate the statistical modelling technique that we deem most appropriate for this purpose. Section 3.2 provides more details on the selected technique.

3.1. Selected technique

We propose a model based on the technique of survival analysis, also known as event history analysis (Aalen et al., 2008). This technique analyses "time to event" data with the aim to estimate the survival rate of a given population, i.e.,the expected time duration until a specific "event of interest" happens. Survival analysis models take into account the fact that some observed subjects may be "censored" when the event of interest was not (yet) observed for them. Survival analysis is a natural choice when it comes to predict the likelihood of an event such as death of an biological organism, failure of mechanical component, or the occurrence of a commit activity in a project.

While other techniques could have been considered to predict future activity of git contributors, they were excluded for practical reasons. *Auto-Regressive Integrated Moving Average (ARIMA)* is a well-known statistical technique for predicting future points in a time series (Shumway and Stoffer, 2011), but it requires the time series to be stationary. In our data, we found that the number of nonseasonal differences (i.e.,the degree of differencing) needed for stationarity varies in time and depends on the considered contributor and project. As such, ARIMA cannot be applied to project contributors without prior individual validation. Doing so would require a lot more historical data that is not necessarily available for smaller projects or for less active contributors.

More advanced models, such as semi-parametric proportional hazards models or fully parametric accelerated failure time models, have the advantage that they can take into account the effect of covariates on events (Collett, 2003). Such models often lead to more accurate predictions, but come at the extra cost of more demanding assumptions: they require that all individuals have the same hazard function, up to a scaling factor. Moreover, these models need to be trained with appropriate covariates. Identifying such covariates is not an easy task and collecting the data for them is challenging in practice, even for simple covariates (e.g., activity of the contributor in other projects, other kinds of activities carried out by the contributor, etc.). While these data can be collected once to feed and train the model, it is still required to collect new data each time the model is used for predictions on a new contributor, a new project or a new date. This makes such models difficult to use in practice. In comparison, a survival analysis model only requires to collect, for each individual, a limited number of previous contributions within the same project.

3.2. Survival analysis model

The prediction model we propose relies on the Kaplan-Meier estimator (Kaplan and Meier, 1958), a common non-parametric statistic to estimate survival functions. A survival function defines the probability of surviving past time t or, equivalently, the probability that the event has not occurred yet at time t.

In our case, it provides probabilities for individual git contributors to remain active in function of the time elapsed since their previous recorded day of activity. We decided to focus on daily commit activity for pragmatic reasons: it is unlikely that a project community manager would need to predict the activity of a contributor up to the nearest second or even hour. While an even coarser granularity of time (e.g.,weeks instead of days) might lead to more accurate models, it would be too coarse to provide useful insights.



Fig. 1. Survival probability for event "the contributor commits to the git repository" of project serde-rs/serde on GitHub, computed at four different time points.

Concretely, given an individual contributor and a project, for each day when a git activity is recorded, we compute the duration until the next day of commit activity of this contributor. The last n computed durations are then fed to the model and used to estimate the survival rate of the considered event.

As an example, Fig. 1 shows Kaplan-Meier survival functions at four different time points in 2018 and in 2019, obtained by extracting and analysing the commit activity of an anonymized contributor of project serde-rs/serde on GitHub. For each considered time point, the model is built based on the last n = 20 time deltas (the elapsed time between the previous commit day and the next one) for the commit activity of the contributor.²

The figure shows the survival probability for the event "the contributor commits to the git repository" as a function of the time elapsed since the previous day of activity of this contributor. Each coloured curve shows the probability (on the y-axis) that the next commit *will not* take place in a specific number of days (value on x-axis). Based on this probability, the complement probability that a commit *will* take place in a specific number of days can be easily computed. For example, consider the orange curve corresponding to October 1st 2018. There is a 45% probability (= 1.0 - 0.55) that the next commit will take place within 2 days, as illustrated by the dashed grey lines in Fig. 1.

To have a first indication of the expected accuracy of the proposed forecasting model, we applied it on all days of activity of an anonymized git contributor on project serde-rs/serde on GitHub. Concretely, for each date of recorded activity, a prediction model was built based on the last n = 20 time deltas. These models were used to predict the duration until next contribution with probabilities 0.5, 0.6, 0.7, 0.8 and 0.9. The predicted durations are compared with the actual duration between the currently considered day and the day of the next recorded activity.

Fig. 2 shows the *actual* days of commit activity of the contributor, and the *predicted* probabilities as areas with different shades of blue. Dots • correspond to *correctly predicted* days of activity (they belong to a blue shared area), while crosses \times correspond to those days of activity for which the model *failed to correctly predict* the time until the next commit. The figure provides initial insight (for a single contributor on a single project only) in the accuracy of the model, in that there are considerably more correct predictions than there are incorrect ones.

4. Validation of the probabilistic model

This section carries out a full-fledged empirical study to evaluate the accuracy of the model on a large and varied dataset containing thousands of projects and contributors. A replication package of the model and its evaluation is provided at doi.org/10.5281/zenodo.3666048.

² The motivation of the choice of n = 20 will be discussed in detail in Section 4.2.



Fig. 2. Rolling application of the probabilistic model for an anonymized contributor on git project serde-rs/serde on GitHub. Dots • correspond to *correctly predicted* days of activity while crosses \times correspond to *incorrectly predicted* days of activity. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)



Fig. 3. Distributions of several characteristics for the considered projects and contributors.

4.1. Selected dataset

To validate the probabilistic forecasting model, we need a large dataset containing thousands of git repositories involving thousands of contributors. The dataset should include a wide range of software projects serving different purposes, and exhibiting a wide variation in terms of longevity and size. The dataset should exclude git repositories that have been created merely for experimental or personal reasons, or that only show sporadic traces of commit activity (Kalliamvakou et al., 2014).

Registries of reusable software packages (e.g.,npm for JavaScript, Cargo for Rust, Maven for Java, or PyPI for Python) are good candidates to find such large datasets, as they typically host thousands of language-specific software packages at different levels of maturity and popularity and serving different purposes. However, not all packages belonging to such registries necessarily have an associated git repository.

To choose an appropriate dataset, we considered all 36 package registries available in the libraries.io 1.2.0 dataset of package managers (Nesbitt and Nickolls, 2018). We retained the Cargo package registry as an adequate case study since it includes over ten thousand distinct packages covering 47 different application domains³, and a very high proportion of these packages (84.5%) have an associated git repository. For comparison, only 68.4% of all npm packages report an associated git repository, and even less for PyPI (61.0%) and Maven (47.5%).

Over 98% of all git repositories associated to Cargo projects were hosted on GitHub. We cloned and locally analysed all those GitHub repositories that were available at the time of our analysis. This corresponded to 75.2% of the initial dataset, accounting for 9208 projects. For each project, the commits of all branches were extracted. Commits found in multiple branches (i.e.,commits that were merged in other branches) were deduplicated and considered only once. This left us with over 1.2M commits.

To safeguard the validity of our results, we addressed the issue of multiple identities for distinct contributors within and across all git repositories. Following the approach of Avelino et al. (Avelino et al., 2019), the GitHub usernames were retrieved using the GitHub API and mapped to contributors' name and email

pairs. An alternative way to resolve multiple identities would have been to resort to an identity merging algorithm (Robles and González-Barahona, 2005; Bird et al., 2006; Goeminne and Mens, 2013; Kouters et al., 2012). However, such algorithms require manual inspection and validation of the results, resulting in a very labour-intensive and error-prone process.

We also removed 6975 contributor/project pairs with a single day of activity as it is obviously not possible to make predictions based on a single day of activity.

Our final dataset includes 7528 projects, 5947 developers, and accounts for 374,622 distinct days of commit activity when considering all 17,340 contributor/project pairs. Fig. 3 shows the distributions of several characteristics of the considered projects and contributors by means of boxen plots (Hofmann et al., 2011). We observe that the dataset covers a wide range of projects and contributors with a diversity in terms of longevity, technical size (number of commits), social size (number of contributors per project and number of projects per contributor) and popularity (number of stars and forks).

It is important to note that not all contributor/project pairs can be used to train the forecasting model. For a given n, the model requires at least n durations (corresponding to at least n + 1 days of git commit activity) to be trained, and an extra one for its validation. Some contributor/project pairs will not have a commit-day history length of at least n + 2 samples and therefore have to be excluded for validating the model.

Fig. 4 shows the number of contributor/project pairs (left y-axis) and the number of samples in the dataset (right y-axis) with respect to the value of *n*. We observe a rapid decrease in the



Fig. 4. Number of contributor/project pairs (left y-axis) and number of samples in the dataset (right y-axis) in function of *n*.

³ https://crates.io/categories.

number of pairs for low values of n, a consequence of the presence of many short-lived contributions. For instance, going from n = 1 to n = 2 results in 3231 fewer pairs (-18.6%) accounting for 17K samples (-4.9%). However, the set of samples on which the model is evaluated remains large even for higher values of n. For instance, the dataset still contains 214,993 samples for n = 20, 141,679 samples for n = 50, and 88,644 samples for n = 100.

4.2. Evaluation of the probabilistic model

Using the selected dataset, for each recorded day of activity of each contributor in each project, n = 20 durations corresponding to the time intervals between the 21 previously recorded days of activity were fed individually into the model, resulting in one model per day of activity, per contributor and per project. Next, the values predicted by the model were compared with the actual duration between the considered day of activity and the next one, following a rolling-origin/walk-forward validation method (Stein, 2007).

Because of the probabilistic nature of the activity prediction model, it cannot be evaluated in a straightforward way (Wilks, 2011). This is because the model does not return a prediction in the form of "Alice will contribute again in 5 days". Instead, it provides a probability distribution for the next contribution of Alice, e.g.,there is a 70% chance that Alice will contribute in the next 10 days and a 90% chance that she will contribute in the next 20 days. Therefore, we first check whether the returned probability distributions correspond to the distribution of actual observations. In other words, we verify if the values predicted with a probability p actually correspond to 100*p percent of the actual observations.

To this end, our model is evaluated using a reliability plot (Wilks, 2011), shown in Fig. 5. It is a variant of a Q - Q (quantile-quantile) plot showing for each probability the proportion of predictions that are actually observed (i.e.,that are less than or equal to the prediction). A point (p, y) in the figure indicates that 100**y* percent of the values predicted for probability *p* are correct upper bounds when they are compared with the actual observations. A perfectly concordant model is obtained when y = p, i.e.,when 100**p* percent of the observations are less than or equal to the predicted values associated with probability *p*. As can be observed in Fig. 5, high probability values have a very small distance between the blue line and the ideal situation represented by the red dotted line. This means that the higher the probability is, the better the concordance of the model.

High probability forecasts ($p \ge 0.8$) are very slightly underconfident (they are predicted less often than they are observed), and low probability forecasts (especially below 0.5) are strongly overconfident (they are predicted far more often than they are actually observed). The observed overconfidence for low probability predictions can be explained by the low dispersion of values that are predicted with lower probabilities *p*: from 94.2% (for *p* = .1) to



Fig. 5. Reliability plot relating the estimated probability distribution to the distribution of actual observations.



Fig. 6. Deviation from perfect concordance w.r.t. the number of past observations used to train the model.

36.9% (for p = .5) of the predictions are of 0 days, a duration that was actually observed in around 43.4% of the cases considered. This is a consequence of many contributors having many periods of time with a (mostly) daily activity, and explains why at least 44.1% of the predictions are actually observed even when the probability is low (p = .1).

While the evaluation results presented above rely on a model based on the last n = 20 past observed durations, we experimentally evaluated the model on a range of increasing values for n, respectively 10, 20, 30, 50, 75 and 100. Larger values of n imply that we require and consider a longer commit history, incidentally excluding spontaneous contributors with short-lived contributions. Smaller values of *n* indicate that we only rely on the most recent days of activity to forecast when future contributions will take place. Fig. 6 shows the deviation of the model with respect to the perfect concordance for different values of n. Smaller deviation indicates a closer match to the perfect concordance and consequently a better model. We computed the sum of squared deviations from perfect concordance for each considered value of n to compare the resulting models. We found that considering the last 20 observed durations produces the lowest score (0.24) compared to the other values that have a score comprised between 0.25 (+4%, for n = 30) and 0.34 (+41%, for n = 100).

4.3. Accuracy of the forecasting model

The accuracy of the forecasting model is evaluated by analysing the error between the predicted and actual (observed) date of the next commit. Since such errors can be either overestimations or underestimations, we computed the number of observations for which the predicted value is above (overestimation), below (underestimation), or equal to the observed value. Fig. 7 shows the proportion of overestimations, underestimations and correct estimations for probabilities ranging from 0.1 to 0.9.

We observe that predictions made with lower probabilities are mostly underestimations while predictions made with higher probabilities are mostly overestimations. This should not come as a surprise since the estimations provided by the prediction model



Fig. 7. Proportion of overestimations, underestimations and exact estimations made by the forecasting model, for probability values ranging from 0.1 to 0.9.



Fig. 8. Distribution of the differences between predicted and observed durations of inactivity.

increase with the probabilities. On average, the model predicts lower values for lower probabilities, and higher values for higher probabilities.

Fig. 8 presents the distributions of the differences between predicted and observed durations of inactivity for probabilities ranging from 0.1 to 0.9. Values closer to zero correspond to more accurate predictions. The dashed gray line corresponds to a "perfect" match between the predicted and observed time until the next commit. Values above the dashed line indicate overestimations, while values below indicate underestimations of the time until the next commit.

We observe that most differences are quite close to 0. For probabilities p < .5, half of the differences between the estimated value and the observed value are comprised between 0 (25th percentile) and -5 days (75th percentile). For probabilities p > .5, positive differences gradually replace negative differences as the probability increases, confirming what was observed for Fig. 7. For instance, for p = .8 half of the differences range between 0 and 7 days with an average of 1.4 days, close to the median value. For p = .9, the distribution of the positive differences is much more skewed, with a median value of 6 days, an average of 12.9 days and a 75th percentile of 18 days.

5. The GAP tool

Based on the probabilistic forecasting model of Section 3, the Git Activity Predictor (GAP) was implemented in Python 3.5+. It is packaged and installable through pip, the official package manager for Python.⁴. GAP makes it possible to forecast developer activity in the form of git commits. Version 0.11.0 of GAP was used for this paper.

The tool has also been integrated by the Bitergia company⁵ in their software development analytics and project monitoring dashboard GrimoireLab. This will facilitate its accessibility and take up by open source project communities.

Fig. 9 presents the general architecture and process followed by GAP. Given a list of git repositories and an optional identity mapping file, GAP analyses the repositories to generate per contributor probabilistic forecasting models for the next commit activity of the contributor at a specific point in time. The forecasts can be reported in four different formats: (i) simple text, (ii) commaseparated values (csv), (iii) JSON, and (iv) bar activity charts.

GAP comes with a command-line interface, summarised in Fig. 10. In its simplest form, it only requires a path to the git repository of the software project that needs to be analysed. One or more such repositories can be provided as input, depending on the scope of the analysis. For example, a company may desire to analyse the commit activity of all its paid contributors in the open source projects the company is involved in; package maintainers may desire to analyse the global activity of all the packages they are maintaining; a package developer may wish to analyse the activity in a project and all its dependents; and so on. By default, GAP analyses all branches of each repository but the list of branches to analyse can be specified.

A file mapping multiple authors into specific identities can be provided to group contributors into teams, to merge multiple identities of the same contributor into one, to anonymise author identities, or to exclude specific authors from the analysis. Parameters can be provided to set the date of the analysis (current day by default), the number of observations used by the model (20 by default), the list of probability values (0.5, 0.6, 0.7, 0.8 and 0.9 by default), and the minimal recency of the last activity (30 days by default). Predictions can be expressed either as dates or as relative time differences (by default).

GAP analyzes the commit history per contributor and predicts the future activity according to the forecasting model of Section 3, producing activity forecasts for each contributor. Fig. 11 shows the output of GAP for a randomly chosen Rust project on git, analysed on 2019-05-16. The authors listed in the output are autogenerated names that we have provided to the identity mapping file anon.csv to guarantee author anonymity.

The output shows, for each recently active contributor (first column), the time difference in days since the last recorded commit activity (second column), and the expected number of days until the next predicted day of activity according to a certain probability threshold ranging between 0.5 and 0.9 (last five columns). Predicted values greater than zero indicate that the next day of activity is expected in the future, while negative values indicate that the predicted day of activity should have taken place before the specified day of analysis.

The output of GAP can also be used as a basis for a projectlevel dashboard that visualises the project's past and estimated future commit activities, as in the example provided in Fig. 12, generated by GAP by appending --plot to the command-line argument. In Fig. 12, the orange squares represent the dates of recorded commit activity of each contributor. The blue rectangles indicate the predicted duration until the next activity for different probabilities. For instance, the figure shows that *M. Fry* (the third contributor in Fig. 12) has been inactive for 8 days (counting backwards from the date of the analysis) but there is a 70% probability that this contributor will contribute within the day, and 80% within the next 8 days. On the other hand, the inactivity of *J. Lopez* (the second contributor in Fig. 12) already exceeds the duration predicted with a probability of 90%, indicating an unexpected irregularity in the activity frequency.

The line with "bots (grouped)" in the figure shows the activity of bots. Two bots were manually identified in the dataset for this particular project repository and we used the identity mapping file to merge them into a single identity. As expected, bots reveal a much more frequent activity pattern. With few exceptions, they commit on a daily basis. Hence, there is 80% probability that either bot will commit again the next day.

6. Discussion and future work

Keeping track of contributors' activity or lack of activity can be very important for the project's sustainability (Gamalielsson and Lundell, 2014; Constantinou and Mens, 2017a). Longer periods of inactivity of a contributor may lead to a permanent disengagement and can have a negative effect on a project. For projects involving a large number of contributors, it is not feasible to contact them directly and ask when their next activity is going to take place. Since each contributor has different activity rhythms, it is difficult to assess manually when someone is falling behind their regular schedule. Consequently, a forecasting model

⁴ pip install git+https://github.com/alexandredecan/gap.

⁵ https://bitergia.com.



Fig. 9. Presentation of the tool architecture.

\$ gap -h
usage: gap [-h] [--date DATE] [--obs OBS]
[--probs [PROB [PROB]]]

```
[--probs [PROB [PROB ...]]]
[--limit LIMIT] [--mapping [MAPPING]]
[--branches [BRANCH [BRANCH ...]]]
[--as-dates]
[--text | --csv | --json | --plot [PLOT]]
[PATH [PATH ...]]
```

Fig. 10. List of GAP command-line arguments.

| \$ gap ./repo | date | 2019- | 05-16 | 5r | nappi | ng and | on.csv |
|----------------|------|-------|-------|-----|-------|--------|--------|
| author | last | 0.5 | 0.6 | 0.7 | 0.8 | 0.9 | |
| M. Johnston | 0 | 3 | 3 | 4 | 6 | 7 | |
| D. Young | 0 | 1 | 1 | 2 | 3 | 4 | |
| B. Rodriguez | 0 | 1 | 1 | 3 | 3 | 6 | |
| D. Johnson | 0 | 2 | 3 | 4 | 4 | 8 | |
| Z. Andrews | 0 | 4 | 5 | 8 | 10 | 23 | |
| J. Berry | 0 | 3 | 6 | 7 | 29 | 131 | |
| bots (grouped) | 0 | 1 | 1 | 1 | 1 | 2 | |
| L. Owen | -1 | 3 | 7 | 7 | 17 | 22 | |
| S. Allen | -2 | 3 | 7 | 12 | 25 | 38 | |
| J. Smith | -4 | 0 | 3 | 6 | 7 | 21 | |
| J. Schultz | -4 | 5 | 10 | 13 | 23 | 38 | |
| M. Fry | -9 | -6 | -3 | 0 | 8 | 21 | |
| J. Lopez | -17 | -11 | -11 | -9 | -6 | -1 | |
| S. Lewis | -20 | -12 | -10 | -3 | 50 | 96 | |

Fig. 11. Example of running GAP on a Rust project's git repository. (Author names are anonymised and replaced by autogenerated ones.) Reported values represent the predicted number of days until the next commit activity will take place. Negative values indicate that no activity has been recorded within the predicted time interval. For example, J. Lopez was predicted with 80% probability to have committed 6 days before 2019-05-16.

\$ gap ./repo --date 2019-05-16 --mapping anon.csv --plot



Fig. 12. Proposed activity forecasting visualisation for a selected git project. GAP was applied on 2019-05-16 for this example. (Autogenerated names are used for the project contributors.).

needs to take into account different activity behaviours for each person (laffaldano et al., 2019). The value of GAP relies on such occasions. GAP identifies the likelihood of a new contribution based on the past activity profile of each individual contributor. Therefore, the forecasts produced by GAP are specific to each individual. The tool can therefore be used to verify which contributors are close to becoming inactive and plan accordingly, possibly mitigating abandonment issues.

As discussed in Section 4.3, estimations provided by GAP are in general close to the observed durations. Forecasts with a higher probability are mostly overestimations, i.e.,the predicted time for the next activity is usually higher than the actual time a contributor usually takes. For lower probabilities, most of the observed differences between predictions and actual observations are slight underestimations. When we consider the scenario of identifying people getting close to becoming inactive, overestimation or underestimation is not a real issue. Indeed, since underestimations mostly occur for lower probabilities, they mainly affect short-term predictions (e.g.,predicting two days of inactivity instead of three). Overestimations are a priori of no consequence since the activity will be observed before the predicted date.

Project managers can decide on the most adequate abandonment risk policy for dealing with inactive contributors. For instance, some managers might prefer overestimating in order to have some leeway before contacting (and possibly bothering) an inactive contributor. Others might prefer to underestimate the predictions to mitigate the loss of certain contributors as early as possible. Either way, GAP can be configured to give a better estimation based on each project manager's specific policy.

Although a project manager could use the forecasting model on git commit activity for the aforementioned reasons, git commits constitute only a subset of all activity that could be made by project contributors. The forecasting model can be adapted to take into account activities at another level of granularity (e.g.,pull requests instead of commits), other technical activities (e.g.,code reviews, issue triaging), or even social activities (e.g.,end-user support, discussions on mailing list). The model can be easily adapted to consider individually each of these types of activities. However, it would fail to deal with multiple simultaneous activities. The reason is that the used statistical Kaplan-Meier estimator is not able to handle competing events and competing risks: "this method can handle only one single event at a time: all other events are treated as censored observations and the complement of the Kaplan-Meier estimate (1-KM) is interpreted as the probability of the event of interest in a hypothetical world in which the competing event does not exist" (Noordzij et al., 2013). Consequently, a promising future work is to explore different techniques (such as the Cumulative Incidence Function (Austin et al., 2016)) that are able to deal with competing events and, by extension, deal with simultaneous activities of different types.

Not only can the model be used for real time monitoring of a project or an ecosystem community, but it could also lead to interesting insights in the context of an *a posteriori* analysis. Indeed, even if the proposed model is a predictive model, the underlying statistical technique of survival analysis is essentially descriptive. The study of the evolution of the survival function of contributors through time constitutes an interesting research track, which can be useful to understand the evolution of contributors' behaviour (e.g.,contribution phases or cycles), or the evolution of contribution profiles (e.g.,an occasional contributor becoming regular, or vice-versa). At ecosystem scale, such a study could provide useful insights, e.g.,on inter-project migrations or abandonment, as "a dead or dying project can also be indicated [...] through lost energy from formerly involved contributors, who may have moved on to other projects that better capture their interest".⁶

This paper mainly focused on "when" the next commit activity will take place, and not on the reasons for such an activity. While from a practical point of view, more advanced models based on Accelerated Failure Time and Proportion Hazard Models have been discarded for practical reasons (because of the amount of data needed to train and use them), such models may be particularly interesting in the context of an *a posteriori* analysis. Since they can take into account the effect of covariates on the probability of occurrence of an event, they could lead to interesting insights to identify and understand the factors that could influence the commit activity of contributors.

7. Threats to validity

Following the structure recommended by Wohlin et al. (2000) we discuss the threats that might affect the validity of our findings, and how we have tried to mitigate them.

Construct validity concerns the relation between the theory behind the experiment and the observed findings. In our work, the main sources of such threats are related to distorted git commit histories due to the presence of *micro-commits* and *commit squashing* (Kalliamvakou et al., 2014).

The practice of micro-commits tends to result in sequences of many small commits that should be considered together as single atomic changes. Micro-commits can lead to overestimating the activity and productivity of a contributor. Since micro-commit sequences tend to take place in short time spans (e.g., the commits are spread over not more than a couple of hours), our findings should not be affected by their presence, as our model relies on the duration between *days* of commit activity. As such, even a huge number of commits contributed during the same day will still be counted as a single event (one day of commit activity by the contributor).

The practice of commit squashing corresponds to the *a posteriori* grouping of several commits. Depending on a project's policy, such commit squashes are mainly observed in pull requests. The rationale is to deliberately group all commits related to a contribution into a single commit in order to ease the reading of the project's git history. Commit squashing inevitably leads to an underestimation of the actual contributions, as it implies the removal of potentially many commits. Unfortunately, there is no reliable technique to identify and/or split these commits to retrieve the original ones. However, it should have little influence on the accuracy of the predictions made by our model as long as squashing is consistently performed on the git history of a given contributor within a given project.

Another threat to construct validity stems from the presence of contributors with multiple identities (different git "author" field values), since the commit activity of a given individual could be spread over these different identities. To mitigate this threat, we followed an approach similar to Avelino et al. (2019) by merging the multiple identities using GitHub usernames in our dataset, as explained in Section 4. It cannot be excluded that some GitHub accounts are shared by multiple contributors (false positives). In that case, GAP reports on the combined activity of this account instead of the individual activity of each contributor using it. It cannot be excluded either that some contributors use multiple GitHub accounts at the same time (false negatives). In that case, GAP reports on the activity of each account individually, and not on the combined activity.

Internal validity concerns choices and factors internal to the study that could influence the observations we made. The main threat comes from the range of observations selected to feed to the model. We decided to rely on a variable time period containing exactly n observations to ensure that the model is trained with a fixed amount of data regardless of a contributor's activity frequency. Choosing a fixed time period (e.g.,n days or weeks) might lead to either a large or a small number of observations, depending on the contributor's activity frequency. The choice of n = 20 is discussed and supported in Section 4.2, where we showed the impact of this parameter on the concordance of the model.

While the technique of survival analysis makes no assumption on the unit of time required to measure time to event, the choice of that unit has a direct consequence of how commits are aggregated and fed to the model. We selected *days* as unit of time because it is a good compromise between practicability and accuracy. A larger granularity (e.g.,*weeks*) would lead to more accurate predictions but provide less interesting insights for project managers. A smaller granularity (e.g.,*hours*) would lead to less accurate results, without making them substantially more informative from a practical point of view.

Conclusion validity concerns the degree to which the conclusions we derived from our data analysis are reasonable. Our conclusions are mostly based on the results of the validation of the model. Since we carefully validated the model following a rolling-origin/walk-forward validation method (Stein, 2007) based on observed historical data, the conclusions we have drawn are unlikely to be affected by such threats.

External validity concerns whether the results and conclusions can be generalized outside the scope of this study. The main threat to external validity relates to the dataset we used to validate the model. To alleviate a potential selection bias related to the profile of contributors and projects, we selected a large dataset containing a wide variety of projects and contributors. The fact that the selected projects are all distributed through a package registry constitutes a possible bias. Indeed, one could expect these projects to be more mature, exhibiting more stable and less impulsive activity patterns than projects being in an early development phase. We are confident that the selected dataset is sufficiently heterogeneous in order not to affect the conclusions obtained, but we do not claim that it is representative of all possible git projects and contributors.

Although the technique used in this paper can be easily transposed to any type of activity and any type of version control system, we cannot guarantee that it will result in accurate predictions. For example, git repositories associated with noncode-based, personal or experimental projects may exhibit totally different activity patterns that cannot be captured by the model. The same argument also holds for other types of activity, such as pull requests, code reviews, social communication, etc. Other version control systems may give rise to a different notion of contributor activity. Due to its decentralized nature, git encourages the creation of commits regardless of the degree of completion of the work performed in the local branch. In contrast, centralised

⁶ github.com/todogroup/todogroup.github.io/blob/master/content/en/guides/ shutting-down.md.

systems such as Subversion encourage contributors to commit changes to a repository only when these changes reflect a single purpose (fixing of a specific bug, addition of a new feature, etc.). This difference in how work is committed to the repository may influence the observed frequency of contributions, and can therefore affect the generalization of our findings.

8. Conclusion

Keeping track of the presence or lack of activity of project contributors can be very important for the project's sustainability. Existing methods and solutions follow an *a posteriori* approach, since they provide information and insights on the past project activity. In practice, community managers also need indicators of potential risks stemming from contributors abandoning their projects. Focusing on the specific activity of contributing to git software repositories, there are currently no techniques or automated tools for predicting the future commit activity of project contributors.

This paper proposes a technique and associated tool to address these issues. In doing so, we hope to enable project managers and practitioners to monitor the future development activity in their git repositories and deal with contributor abandonment risks early, before their departure becomes potentially detrimental to the project. We presented a probabilistic forecasting model to estimate when the next commit activity of a given git contributor will occur. This model relies on survival analysis, a statistical technique that we used to estimate a contributor's probability of committing again, based on the past durations between days of observed commit activity for that contributor. The model relies on the last 20 observed durations between days of activity per contributor, which presents several benefits: (i) it does not require to collect or process a large amount of data; (ii) little prior knowledge is required, making it applicable to short-lived as well as long-lived projects; and (iii) since the analysis requires not more than a few seconds per project repository, it can be used for real-time monitoring, even for large communities of project contributors.

We evaluated the model on more than 7.5K git repositories corresponding to a wide variety of projects. We showed that estimations provided by our approach are generally close to the observed durations. By analysing the accuracy of the prediction model, we found that forecasts with a higher probability are mostly overestimations, while most of the observed differences between predictions for lower probabilities and actual observations are slight underestimations. We found that relying on the last n = 20 observed durations between days of activity provides the results with the lowest deviation.

To showcase the practicality of the model, we operationalised it through the Git Activity Predictor (GAP), an open source command-line utility. This tool enables practitioners and researchers to benefit from the proposed forecasting model and to apply it on sets of git repositories. GAP has built-in support for identity mapping to anonymise contributors and to merge identities (e.g.,to deal with multiple identities, or to combine bot activity), thus both preserving privacy and providing accurate results. It supports different output formats, including a bar activity chart visualisation of the recent and predicted commit activities of contributors. GAP has also been integrated as part of the open source GrimoireLab software development analytics platform.⁷

We plan to evaluate and extend the probabilistic model and tool for other types of project forecasting activities (e.g., issue report interaction, code review or mailing list activity), and to combine different types of activity to better understand and predict the activity dynamics of individual project contributors. In addition, the model could be extended to the team level, to understand and forecast activity for groups of collaborating contributors within the same project.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Alexandre Decan: Conceptualization, Methodology, Software, Validation, Investigation, Data curation, Writing - original draft, Writing - review & editing, Visualization. **Eleni Constantinou:** Conceptualization, Investigation, Data curation, Writing - original draft, Writing - review & editing. **Tom Mens:** Conceptualization, Writing - original draft, Writing - review & editing, Supervision, Project administration, Funding acquisition. **Henrique Rocha:** Writing - review & editing.

Acknowledgments

This research was supported by the Fonds de la Recherche Scientifique – FNRS under Grants number T.0017.18, R.60.04.18.F (FRQ-FNRS collaborative research project "SECOHealth") and O.0157.18F-RG43 (Excellence of Science project "SECO-Assist").

Supplementary material

Supplementary material associated with this article can be found, in the online version, at 10.1016/j.jss.2020.110573

References

- Aalen, O., Borgan, O., Gjessing, H., 2008. Survival and Event History Analysis: A Process Point of View. Springer doi:10.1007/978-0-387-68560-1.
- Amor, J.J., Robles, G., Gonzalez-Barahona, J.M., 2006. Effort estimation by characterizing developer activity. In: Int'l Workshop on Economics driven software engineering research. ACM, pp. 3–6. doi:10.1145/1139113.1139116.
- Austin, P.C., Lee, D.S., Fine, J.P., 2016. Introduction to the analysis of survival data in the presence of competing risks. Circulation 133 (6), 601–609. doi:10.1161/ CIRCULATIONAHA.115.017719.
- Avelino, G., Constantinou, E., Valente, M.T., Serebrenik, A., 2019. On the abandonment and survival of open source projects: An empirical investigation. Int'l Symp. Empirical Software Engineering and Measurement (ESEM). IEEE doi:10. 1109/ESEM.2019.8870181.
- Bird, C., Gourley, A., Devanbu, P., Gertz, M., Swaminathan, A., 2006. Mining email social networks. In: Int'l Conf. Mining Software Repositories. ACM Press, pp. 137– 143. doi:10.1145/1137983.1138016.
- Bird, C., Gourley, A., Devanbu, P., Swaminathan, A., Hsu, G., 2007. Open borders? Immigration in open source projects. Int'l Workshop on Mining Software Repositories (MSR) doi:10.1109/MSR.2007.23.
- Coelho, J., Valente, M.T., 2017. Why modern open source projects fail. In: Joint Meeting on Foundations of Software Engineering (FSE). ACM, pp. 186–196. doi:10.1145/3106237.3106246.
- Collett, D., 2003. Modelling Survival Data in Medical Research, Second Edition. Chapman & Hall/CRC Texts in Statistical Science, Taylor & Francis.
- Constantinou, E., Mens, T., 2017. An empirical comparison of developer retention in the RubyGems and NPM software ecosystems. Innov. Syst. Softw. Eng. 13 (2), 101–115. doi:10.1007/s11334-017-0303-4.
- Constantinou, E., Mens, T., 2017. Socio-technical evolution of the Ruby ecosystem in GitHub. In: Int'l Conf. Software Analysis, Evolution and Reengineering (SANER), pp. 34–44. doi:10.1109/SANER.2017.7884607.
- Crowston, K., Howison, J., 2005. The social structure of free and open source software development. First Monday 10 (2). doi:10.5210/fm.v10i2.1207.
- Decan, A., Goeminne, M., Mens, T., 2017. On the interaction of relational database access technologies in open source Java projects. In: Bagge, A., Mens, T., Osman, H. (Eds.), Post-proceedings of the 8th Seminar on Advanced Techniques and Tools for Software Evolution. CEUR-WS.org, pp. 26–35.
- Decan, A., Mens, T., Grosjean, P., 2019. An empirical comparison of dependency network evolution in seven software packaging ecosystems. Empir. Softw. Eng. 24, 381-416. doi:10.1007/s10664-017-9589-y.
- Dias, L.F., Steinmacher, I., Pinto, G., 2018. Who drives company-owned OSS projects: internal or external members? J. Braz. Comput. Soc. 24 (1). doi:10.1186/ s13173-018-0079-x.

⁷ chaoss.github.io/grimoirelab/

Fronchetti, F., Wiese, I., Pinto, G., Steinmacher, I., 2019. What attracts newcomers to onboard on OSS projects? tl;dr: Popularity. In: Open Source Systems. Springer, pp. 91–103.

- Gamalielsson, J., Lundell, B., 2014. Sustainability of open source software communities beyond a fork: how and why has the LibreOffice project evolved? J. Syst. Softw. 89, 128–145. doi:10.1016/j.jss.2013.11.1077.
- Goeminne, M., Mens, T., 2013. A comparison of identity merge algorithms for software repositories. Sci. Comput. Program. 78 (8), 971–986. doi:10.1016/j.scico. 2011.11.004.
- Gousios, G., Kalliamvakou, E., Spinellis, D., 2008. Measuring developer contribution from software repository data. In: Int'l Working Conf. Mining Software Repositories (MSR). ACM, pp. 129–132. doi:10.1145/1370750.1370781.
- Hofmann, H., Kafadar, K., Wickham, H., 2011. Letter-value plots: Boxplots for large data. Technical Report. had.co.nz.
- Iaffaldano, G., Steinmacher, I., Calefato, F., Gerosa, M., Lanubile, F., 2019. Why do developers take breaks from contributing to OSS projects? A preliminary analysis. In: Int'l Workshop on Software Health (SoHeal). IEEE, pp. 9–16. doi:10.1109/ SoHeal.2019.00009.
- Kalliamvakou, E., Gousios, G., Blincoe, K., Singer, L., German, D.M., Damian, D., 2014. The promises and perils of mining GitHub. In: Int'l Conf. Mining Software Repositories. ACM, pp. 92–101. doi:10.1145/2597073.2597074.
- Kaplan, E.L., Meier, P., 1958. Nonparametric estimation from incomplete observations. J. Am. Stat. Assoc. 53 (282), 457–481. doi:10.1080/01621459.1958. 10501452.
- Kouters, E., Vasilescu, B., Serebrenik, A., van den Brand, M.G.J., 2012. Who's who in Gnome: using LSA to merge software repository identities. In: Int'l Conf. Software Maintenance. IEEE, pp. 592–595. doi:10.1109/ICSM.2012.6405329.
- Kyriakakis, P., Chatzigeorgiou, A., 2014. Maintenance patterns of large-scale PHP web applications. In: Int'l Conf. Software Maintenance and Evolution, pp. 381–390. doi:10.1109/ICSME.2014.60.
- Lin, B., Robles, G., Serebrenik, A., 2017. Developer turnover in global, industrial open source projects: Insights from applying survival analysis. In: Int'l Conf. Global Software Engineering (ICGSE) doi:10.1109/ICGSE.2017.11.
- Miller, C., Widder, D.G., Kästner, C., Vasilescu, B., 2019. Why do people give up FLOSSing? A study of contributor disengagement in open source. In: Open Source Systems. Springer, pp. 116–129.
- Mockus, A., Weiss, D.M., Zhang, P., 2003. Understanding and predicting effort in software projects. In: Int'l Conf. Software Engineering, pp. 274–284. doi:10.1109/ ICSE.2003.1201207.
- Nesbitt, A., Nickolls, B., 2018. Libraries.io open source repository and dependency metadata. 10.5281/zenodo.1196312
- Noordzij, M., Leffondré, K., van Stralen, K.J., Zoccali, C., Dekker, F.W., Jager, K.J., 2013. When do we need competing risks methods for survival analysis in nephrology? Nephrol. Dial. Transpl. 28 (11), 2670–2677. doi:10.1093/ndt/gft355.
- Qureshi, I., Fang, Y., 2011. Socialization in open source software projects: a growth mixture modeling approach. Organ. Res. Methods 14 (1), 208–238. doi:10.1177/ 1094428110375002.
- Riehle, D., Riemer, P., Kolassa, C., Schmidt, M., 2014. Paid vs. volunteer work in open source. In: 2014 47th Hawaii International Conference on System Sciences, pp. 3286–3295. doi:10.1109/HICSS.2014.407.
- Rigby, P.C., Zhu, Y.C., Donadelli, S.M., Mockus, A., 2016. Quantifying and mitigating turnover-induced knowledge loss: case studies of Chrome and a project at Avaya. In: Int'l Conf. Software Engineering. ACM, pp. 1006–1016. doi:10.1145/ 2884781.2884851.
- Robles, G., González-Barahona, J.M., 2005. Developer identification methods for integrated data from various sources. In: Int'l Conf. Mining Software Repositories. ACM doi:10.1145/1082983.1083162.
- Samoladas, I., Angelis, L., Stamelos, I., 2010. Survival analysis on the duration of open source projects. Inf. Softw. Technol. 52 (9), 902–922. doi:10.1016/j.infsof. 2010.05.001.
- Shumway, R., Stoffer, D., 2011. Time Series Analysis and Its Applications With R Examples, third Springer doi:10.1007/978-1-4419-7865-3.

- Stein, R.M., 2007. Benchmarking default prediction models: pitfalls and remedies in model validation. J. Risk Model Validation 1 (1), 77–113. doi:10.21314/JRMV. 2007.002.
- Steinmacher, I., Conte, T.U., Treude, C., Gerosa, M.A., 2016. Overcoming open source project entry barriers with a portal for newcomers. In: International Conference on Software Engineering. ACM, pp. 273–284. doi:10.1145/2884781.2884806.
- Steinmacher, I., Graciotto Silva, M.A., Gerosa, M.A., Redmiles, D.F., 2015. A systematic literature review on the barriers faced by newcomers to open source software projects. Inf. Softw. Technol. 59 (C), 67–85. doi:10.1016/j.infsof.2014.11.001.
- Weicheng, Y., Beijun, S., Ben, X., 2013. Mining Github: Why commit stops exploring the relationship between developer's commit pattern and file version evolution. In: Asia-Pacific Software Engineering Conf. (APSEC), 2, pp. 165–169. doi:10.1109/APSEC.2013.133.
- Wilks, D., 2011. Forecast verification. In: Statistical Methods in the Atmospheric Sciences. In: International Geophysics, 100. Academic Press, pp. 301–394. doi:10. 1016/B978-0-12-385022-5.00008-7.
- Wohlin, C., Runeson, P., Host, M., Ohlsson, M.C., Regnell, B., Wesslen, A., 2000. Experimentation in Software Engineering - An Introduction. Kluwer.

Alexandre Decan (orcid.org/0000-0002-5824-5823) conducted his doctoral studies at the University of Mons on the subject of data quality in relational databases. He obtained a PhD in Science in 2013 for the thesis entitled "Certain Query Answering in First-Order Languages". He is post-doctoral researcher at the Software Engineering Lab of the University of Mons (Belgium), where he has authored many wellcited publications related to the maintenance and evolution of software ecosystems. He has been actively involved in several research projects such as the UMONS ARC "ECOS", the Walloon ERDF project portfolio "IDEES", the FNRS-FRQ collaborative research project "SecOHealth" and the joint Belgian FNRS-FWO Excellence of Science project "SECO-ASSIST".

Eleni Constantinou (orcid.org/0000-0002-4242-2581) obtained her PhD in 2015 from the Aristotle University of Thessaloniki and was subsequently a postdoctoral researcher at the University of Mons. Since October 2019, she is an assistant professor at the Eindhoven University of Technology (TU/e). Her research interests include mining software repositories, software evolution, empirical software engineering, software ecosystems and software health. She has published her research in several peer-reviewed conferences and journals. She was the scientific coordinator of the FNRS-FRQ collaborative research project "SecoHealth", and is an external collaborator of the "SECO-ASSIST" project.

Prof. Dr. Tom Mens (orcid.org/0000-0003-3636-5020) obtained a PhD in Science in 1999 at the Vrije Universiteit Brussel. He is full professor at the University of Mons in October 2003, where he is directing the Software Engineering Lab. His current research interests include software evolution, empirical software engineering, open source software ecosystems, and software health. He published numerous scientific articles in peer-reviewed international software engineering conferences and journals. He was project leader of the FNRS-FRQ collaborative research project "Seco-Health" and is project leader of the Belgian FNRS-FWO Excellence of Science project "SECO-ASSIST".

Henrique Rocha (orcid.org/0000-0002-9154-0277) finished his PhD in 2016 at UFMG (Brazil) working in the research group of Applied Software Engineering (ASERG). He was a post-doctoral researcher at the RMOD team of INRIA Lille (France) from May 2017 to December 2018), focusing on Blockchain Oriented Software Engineering. Since January 2019, he is a post-doctoral researcher at the University of Antwerp (Belgium) focusing on software testing and software ecosystems in the context of the "SECO-ASSIST" project. His research interests include software testing, software ecosystems, applied software engineering, software maintenance and evolution, and Blockchain technology.