

The OCEAN mailing list data set: Network analysis spanning mailing lists and code repositories

Melanie Warrick
warrick.melanie@gmail.com
Google, Inc.
Mountain View, California, USA

Samuel F. Rosenblatt
samuel.f.rosenblatt@uvm.edu
University of Vermont
Burlington, Vermont, USA

Jean-Gabriel Young
jean-gabriel.young@uvm.edu
University of Vermont
Burlington, Vermont, USA

Amanda Casari
amcasari@google.com
Google, Inc.
Mountain View, California, USA

Laurent Hébert-Dufresne
laurent.hebert-dufresne@uvm.edu
University of Vermont
Burlington, Vermont, USA

James Bagrow
james.bagrow@uvm.edu
University of Vermont
Burlington, Vermont, USA

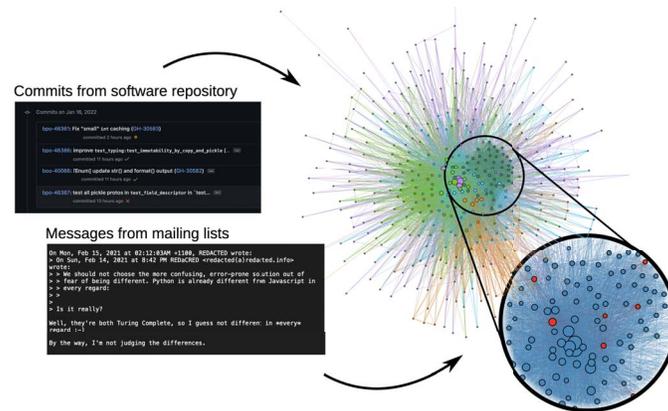


Figure 1: Layering social communication data onto the technical network structure of an open source project. The main network is produced by projecting the account-commit-file data of the CPython repository on a “graph of who collaborates with whom” which unveils the modular structure of the project and community (nodes are colored to highlight this using Louvain modularity maximization [4]). We then zoom on a particular module to highlight the role of certain nodes in the social layer of the community using our mailing list data and find that 6 of the top 10 nodes ranked by counts of toxic words are found in this single module.

ABSTRACT

Communication surrounding the development of an open source project largely occurs outside the software repository itself. Historically, large communities often used a collection of mailing lists to discuss the different aspects of their projects. Multimodal tool use, with software development and communication happening on different channels, complicates the study of open source projects as a sociotechnical system. Here, we combine and standardize mailing lists of the Python community, resulting in 954,287 messages from 1995 to the present. We share all scraping and cleaning code

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

MSR '22, May 23–24, 2022, Pittsburgh, PA, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9303-4/22/05... \$15.00

<https://doi.org/10.1145/3524842.3528479>

to facilitate reproduction of this work, as well as smaller datasets for the Golang (122,721 messages), Angular (20,041 messages) and Node.js (12,514 messages) communities. To showcase the usefulness of these data, we focus on the CPython repository and merge the technical layer (which GitHub account works on what file and with whom) with the social layer (messages from unique email addresses) by identifying 33% of GitHub contributors in the mailing list data. We then explore correlations between the valence of social messaging and the structure of the collaboration network. We discuss how these data provide a laboratory to test theories from standard organizational science in large open source projects.

CCS CONCEPTS

• **Human-centered computing** → **Social network analysis**; • **Software and its engineering** → *Open source model*; • **Applied computing** → Document management and text processing.

KEYWORDS

datasets, sociotechnical systems, network analysis, text tagging

ACM Reference Format:

Melanie Warrick, Samuel F. Rosenblatt, Jean-Gabriel Young, Amanda Casari, Laurent Hébert-Dufresne, and James Bagrow. 2022. The OCEAN mailing list data set: Network analysis spanning mailing lists and code repositories. In *19th International Conference on Mining Software Repositories (MSR '22), May 23–24, 2022, Pittsburgh, PA, USA*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3524842.3528479>

1 INTRODUCTION

The practice of mining and analysing software repositories provides a unique window into software development and the social practice of problem solving as a whole. Unfortunately, these analyses often come with the caveat that much of the collaborative work other than coding appears in public but away from the software repository itself [11]. The community relies on data streams that, unlike git history for example, do not come in standard formats [15]: bug reports, mailing lists, online forums, etc.

For large open source communities (OSC), social interactions are rarely limited to a single platform [32], from direct messages, to targeted forums or massively distributed mailing lists. Importantly, discussions there concern a much wider range of contributions to the community than coding itself, with efforts around the development of norm for the community or organisation of social events [24, 34?]. These interactions have been shown to be critical to the growth and health of the community, with constructive and timely replies being positively correlated with future participation [16].

Data describing social interactions within OSCs are notoriously hard to analyse given their unstructured format [15, 26]. Past work has therefore focused on manual readings of samples of messages, for example from newcomers [16] or major developers and community members [26], or on automated analysis of specific windows chosen for their fixed format [8]. Altogether, several challenges exist to mining mailing lists [3, 14]. Here, we present data with scraping software to address two specific problems: Large communities tend to use multiple mailing lists for different topics (e.g. development, ideas, announcements, etc.) and these different lists do not operate with the same software, encoding, or data structure.

2 THE OCEAN MAILING LIST DATA

We present the data collected as part of the Open-source Complex Ecosystem And Networks (OCEAN) partnership between Google Open Source and the University of Vermont. Our code base for data collection is hosted at <https://github.com/google/project-OCEAN> and contains a suite of utilities to compile mailing lists based on their hosting service. This is currently limited to mailing lists hosted through Mailman, Google Groups, or Piplermail.

We include 14 mailing lists as part of this initial release but the data set could be straightforwardly extended to include other OSCs which use mailing lists as a main form of communication. We here show all lists currently available, in chronological order:

- (1) pipemail-python-dev, started on 1995-03,
- (2) pipemail-python-list, started on 1999-02,
- (3) pipemail-python-announce-list, started on 1999-04,
- (4) mailman-python-announce-list, started on 1999-04,
- (5) mailman-python-dev, started on 1999-04,
- (6) mailman-python-ideas, started on 2006-12,

- (7) pipemail-python-ideas, started on 2006-12,
- (8) gg-nodejs, started on 2009-06,
- (9) gg-angular, started on 2009-09,
- (10) gg-golang-checkins, started on 2009-11,
- (11) gg-golang-dev, started on 2009-11,
- (12) gg-golang-nuts, started on 2009-11,
- (13) gg-golang-announce, started on 2011-05,
- (14) gg-golang-codereviews, started on 2013-12.

This database was aggregated by consulting with open source community members to identify both active mailing lists used by the selected projects, as well as mailing lists which are known to contain historical decision making information.

Our codebase then standardizes date format, text information and origins of messages under one large database per community. The most important fields of our data sets are as follows:

Field	Type	Example
from_name	string	Example Person
from_email	string	example.person@gmail.com
to_name	string	python-dev
to_email	string	python-dev@python.org
message_id	string	<cc76ef2 ... @googlegroups.com>
subject	string	parameters not working
date	datetime	2016-05-02T11:45:35
body_text	string	I'm trying to...
original_url	string	https://groups.google.com/...

In the earlier years of the archives, the date format used by the mail clients was quite variable. While our tool supports a series of queries to parse dates, “to,” “from,” and references fields, it does not always succeed, so we also preserved the raw data under additional fields to allow easy format checks.

Key descriptive statistics of the mailing lists data are shown in Fig. 2. In terms of volume of both messages and newcomers, all mailings lists most recently peaked between 2011 and 2015. Across all communities, the ratio of newcomers to established addresses tend to be relatively fixed but varied with some mailing lists driven by newcomers (Angular) or existing addresses (GoLang). In all cases, we find skewed distributions of messages per unique email address, with outliers most often representing bots. These bots can be easily identified within the data, but are preserved as they often forward actual social messages from other platforms.

3 CURRENT USE OF THE DATABASE

The OCEAN mailing list database was recently shared with researchers at Galois, inc. as part of a DARPA funded project on the impact of toxicity on collaborations in large and important open source projects. As part of this project, we here focus on a subset of the Python community: the CPython repository containing the reference implementation of the language and where most Python development occurs. The database was leveraged as one element in a new tool, the LAGOON platform for Leveraging AI to Guard Online Open-source Networks, which does identity merging and comes packaged with a UI for inspecting and analyzing the ingested open source community [9]. Through this tool, our collaboration highlights that 33.2% of all CPython contributors can be identified

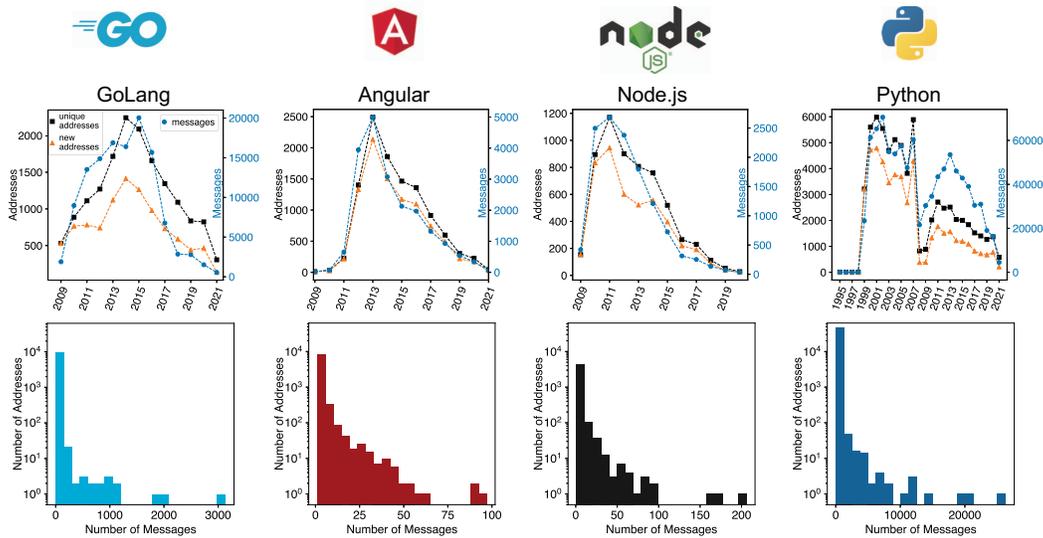


Figure 2: (top row) Volumes of addresses and messages of different mailing list datasets through time. Messages and addresses are on different scales. (bottom row) Histograms of messages per address. Three addresses with outlying volumes of automated messages (7,592, 19,492, and 33,660) are not shown in the GoLang histogram (lower left).

in the mailing list data. This provides us with a large sampling of the content and valence of discussions within the community.

As a simple case study, we can tag CPython contributors by the amount of toxic language they send on the mailing list. Previous research in management science has explored the positive association between power and negative ties [17], and it is unclear if these results generalize to OSS communities which lack clear hierarchies. For simplicity, we here study this issue with a naive definition of toxic language, based on a previous study which developed an annotated corpus and lexicon for harassment research [25]. This corpus was manually edited to account for words which tend to have a different meaning in technical discussions than in colloquial language (e.g. “primitive”). We then summarize the social messages of a given users simply by counting how many toxic words they have used while messaging the community.

In Fig. 1, we show where toxic nodes are found within the collaboration network of CPython, a projection of the account-commit-files network on the set of edges of accounts that have touched files in common). To do so, we first identify modules (e.g. teams) in the community using the Louvain modularity maximization algorithm [4]. We then find that toxic language is not randomly distributed within the collaboration network of CPython, with 6 of the top 10 nodes ranked by counts of toxic words found within a single central module.

In Fig. 3, we quantify the correlations between toxic word usage and network structure in the collaboration network. In the right panel we first find, unsurprisingly, that (1) accounts that send more messages are more likely to be flagged as toxic (more opportunity to use toxic words) and that (2) regardless of toxicity, accounts that

send more messages also tend to have more collaborators in the technical layer of the network. More surprisingly, in the middle panel, we find that (3) toxic accounts tend to gain less collaborators for every commit to the repository. This result is further explored in the right panel where we see that accounts that are toxic in a given year tend to have a smaller average number of collaborators per file.

Further research is needed to fully understand these results and, in particular, to delineate between toxic accounts that as contributors are central technical dependencies and toxic accounts that operate on the periphery of a project and offer minimal contributions. Many different mechanisms can be postulated to explain the observed correlations: perhaps having few collaborators lead to toxic languages, or conversely, toxic languages lead to disengagement of collaborators. Altogether, these simple correlations showcase the value of merging social communication data like mailing lists with technical data from the repository itself, but warrant further research to disentangle the mechanisms at play.

4 ADDITIONAL RELATED WORK

There are several data sets aiming to capture activities around software development outside of its code base. Besides the analyses of mailing lists already mentioned [3, 10, 14, 30], relevant studies includes discussions on unusual platforms such as blogs [30], Slack messages [6, 7] or IRC messages [28, 33], though the data for these studies is not always made available in easy to parse format, if at all. Going further afield, datasets of Stack Overflow messages and events [1, 2, 33] are also related to our work in that they pertain to a community and are disconnected from the code base. However,

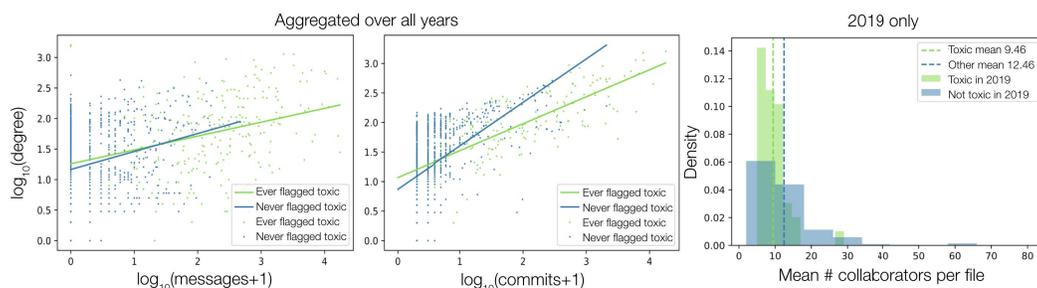


Figure 3: Correlations between activity in the social (messages and toxicity) and technical (commits, collaborators or “degree”, and files) layers of the CPython repository. Degree refers to a user’s number of co-editors in the projected version of the user-commit-file network. Accounts that use toxic words tend to have less collaborators per commit and per file.

discussions on Stack Overflow focus on using the software rather than developing it [18], and thus capture a different aspect of a community’s behavior.

Our example application—natural language processing and sentiment analysis of data pertaining to open source communities—is also a very active field [5]. Work in this area uses diverse data source such as commit message [29], issue reports [19, 21], email lists [26, 27], or Q&A platform, to name a few [5]. The main challenges here is the unusual nature of technical communication, which differs significantly from the type of text usually analyzed with NLP [20].

5 CHALLENGES, IMPROVEMENTS, AND FUTURE WORK

From the projects included in our data repository, we find the practice of using mailing lists as the main communication channel for open source projects has been steadily decreasing in popularity over the last decade or so. This presents an obvious limitation of relying solely on mailing lists as a proxy for social interactions within a community, but also an provides an archival opportunity as our database creates a largely static window into the past social interactions around current major projects in open source software development. For continued study of active projects within these communities, our data will need to be supplemented by forum discussions, bug reports, and newer platform conversation channels such as messaging platforms, social media, and blog rolls [12, 30].

In addition to the raw text data associated with messages, we expect to process the mailing lists to automatically identify text structures such as email signatures, code and quotes in replies [23]. Quotes, in particular, are potentially powerful as they provide a window into moderation norms and practices. Indeed, targeted toxic emails are often moderated and do not appear as a sent message in the mailing list, but targeted recipients can still receive this message directly rather than through the list. For example, by sending the email to both `targetemail.com` and to the mailing list, the target receives the email regardless of moderation. The target can then reply to the message, which will be quoted in the reply therefore bypassing the original moderation. Identifying and flagging these quotes will open interesting avenues of research as moderated texts are virtually never included in available data.

A key point of discussion among the project team was how to address personal identifiable information that is already accessible from the data sources, available under the platform terms and conditions, and shared following the community guidelines. Best practices for aggregating, sharing, and working with open source data gleaned from community projects is still mixed [13]; while transparency is valued in open source, communities also value their personal privacy while working in open spaces. Since we are sharing this database as both a dataset, the list of sources this data was aggregated from, and the source code used to produce it, the team chose to follow the CHAOSS Community Data Policy, which states that “our community data is part of our public history,” disclosing all data as assembled from the original sources to “preserve the authenticity of [our] community data” [22].

Our repository for data compiled using the OCEAN tool will grow and allow the public history of an OSC to be studied in order to evaluate and improve the effectiveness of community norms. Future work will integrate the raw mailing list data with natural language processing tools to classify messages and evaluate how communication, miscommunication, and toxicity can affect collaborations. We then envision studying how the social structure found in the mailing list data reflects the collaborative structure found in the software repository, as well as the architecture of the software produced; from toxic messages and drive-by contributions, to social support and long-term innovations.

ACKNOWLEDGMENTS

This material is based upon work supported by Google Open Source under the Open Source Complex Ecosystems And Networks (OCEAN) project and by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR00112190092. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of Google Open Source or DARPA.

AUTHOR CONTRIBUTIONS

M.W. built the data set. S.F.R. ran the analyses. J.G.Y., A.C., L.H.D. and J.P.B. wrote the original draft, which all authors edited. Data are hosted at: <https://doi.org/10.6084/m9.figshare.19082540.v2> [31].

REFERENCES

- [1] Alberto Bacchelli. 2013. Mining Challenge 2013: Stack Overflow. In *The 10th Working Conference on Mining Software Repositories*. to appear.
- [2] Sebastian Baltes, Christoph Treude, and Stephan Diehl. 2019. SOTorrent: Studying the Origin, Evolution, and Usage of Stack Overflow Code Snippets. In *Proceedings of the 16th International Conference on Mining Software Repositories (MSR 2019)*.
- [3] 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.
- [4] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008, 10 (2008), P10008.
- [5] Fabio Calefato, Filippo Lanubile, Federico Maiorano, and Nicole Novielli. 2018. Sentiment polarity detection for software development. *Empirical Software Engineering* 23, 3 (2018), 1352–1382.
- [6] Preetha Chatterjee, Kostadin Damevski, Nicholas A Kraft, and Lori Pollock. 2020. Software-related slack chats with disentangled conversations. In *Proceedings of the 17th International Conference on Mining Software Repositories*. 588–592.
- [7] Preetha Chatterjee, Kostadin Damevski, Lori Pollock, Vinay Augustine, and Nicholas A Kraft. 2019. Exploratory study of slack Q&A chats as a mining source for software engineering tools. In *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. IEEE, 490–501.
- [8] Hoa Khanh Dam, Bastin Tony Roy Savarimuthu, Daniel Avery, and Aditya Ghose. 2015. Mining software repositories for social norms. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, Vol. 2. IEEE, 627–630.
- [9] Sourya Dey and Walt Woods. 2022. LAGOON: An Analysis Tool for Open Source Communities. In *2022 IEEE/ACM 19th International Conference on Mining Software Repositories (MSR)*.
- [10] Nicolas Ducheneaut. 2005. Socialization in an open source software community: A socio-technical analysis. *Computer Supported Cooperative Work (CSCW)* 14, 4 (2005), 323–368.
- [11] Nadia Eghbal. 2020. *Working in Public: The Making and Maintenance of Open Source Software*. Stripe Press, San Francisco, CA.
- [12] Hongbo Fang, Daniel Klug, Hemank Lamba, James Herbsleb, and Bogdan Vasilescu. 2020. Need for Tweet: How Open Source Developers Talk About Their GitHub Work on Twitter. In *2020 IEEE/ACM 17th International Conference on Mining Software Repositories (MSR)*. 322–326.
- [13] Nicolas E Gold and Jens Krinke. 2022. Ethics in the mining of software repositories. *Empirical Software Engineering* 27, 1 (2022), 1–49.
- [14] Anja Guzzi, Alberto Bacchelli, Michele Lanza, Martin Pinzger, and Arie Van Deursen. 2013. Communication in open source software development mailing lists. In *2013 10th Working Conference on Mining Software Repositories (MSR)*. IEEE, 277–286.
- [15] Ahmed E Hassan. 2008. The road ahead for mining software repositories. In *2008 Frontiers of Software Maintenance*. IEEE, 48–57.
- [16] Carlos Jensen, Scott King, and Victor Kuechler. 2011. Joining free/open source software communities: An analysis of newbies' first interactions on project mailing lists. In *2011 44th Hawaii international conference on system sciences*. IEEE, 1–10.
- [17] Giuseppe Labianca and Daniel J. Brass. 2006. Exploring the Social Ledger: Negative Relationships and Negative Asymmetry in Social Networks in Organizations. *Academy of Management Review* 31, 3 (2006), 596–614. <https://doi.org/10.5465/amr.2006.21318920> arXiv:<https://doi.org/10.5465/amr.2006.21318920>
- [18] Saraj Singh Manes and Olga Baysal. 2019. How often and what StackOverflow posts do developers reference in their GitHub projects?. In *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. IEEE, 235–239.
- [19] Mika Mäntylä, Bram Adams, Giuseppe Destefanis, Daniel Graziotin, and Marco Ortu. 2016. Mining valence, arousal, and dominance: possibilities for detecting burnout and productivity?. In *Proceedings of the 13th international conference on mining software repositories*. 247–258.
- [20] Mika V Mäntylä, Fabio Calefato, and Maelick Claes. 2018. Natural language or not (nlon) a package for software engineering text analysis pipeline. In *Proceedings of the 15th International Conference on Mining Software Repositories*. 387–391.
- [21] Mika V Mäntylä, Nicole Novielli, Filippo Lanubile, Maelick Claes, and Miikka Kuuttila. 2017. Bootstrapping a lexicon for emotional arousal in software engineering. In *2017 IEEE/ACM 14th International Conference on Mining Software Repositories (MSR)*. IEEE, 198–202.
- [22] CHAOSS Project. 2022. *CHAOSS Project's Data Policies*. Retrieved January 25, 2022 from <https://chaoss.community/about/data-policies/>
- [23] Musfiqur Rahman, Peter Rigby, Dharani Palani, and Tien Nguyen. 2019. Cleaning stackoverflow for machine translation. In *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. IEEE, 79–83.
- [24] Eric Raymond. 1999. The cathedral and the bazaar. *Knowledge, Technology & Policy* 12, 3 (1999), 23–49.
- [25] Mohammadreza Rezvan, Saeedeh Shekarpour, Lakshika Balasuriya, Krishnaprasad Thirunarayan, Valerie L Shalin, and Amit Sheth. 2018. A quality type-aware annotated corpus and lexicon for harassment research. In *Proceedings of the 10th ACM Conference on Web Science*. 33–36.
- [26] 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)*. 23–23. <https://doi.org/10.1109/MSR.2007.35>
- [27] Athanasios-Ilias Rousinopoulos, Gregorio Robles, and Jesús M González-Barahona. 2014. Sentiment analysis of free/open source developers: preliminary findings from a case study. *Revista Eletrônica de Sistemas de Informação* 13, 2 (2014).
- [28] Emad Shihab, Zhen Ming Jiang, and Ahmed E Hassan. 2009. Studying the use of developer IRC meetings in open source projects. In *2009 IEEE International Conference on Software Maintenance*. IEEE, 147–156.
- [29] Vinayak Sinha, Alina Lazar, and Bonita Sharif. 2016. Analyzing developer sentiment in commit logs. In *Proceedings of the 13th international conference on mining software repositories*. 520–523.
- [30] Patrick Wagstrom, Jim Herbsleb, and Kathleen Carley. 2005. A social network approach to free/open source software simulation. In *Proceedings First International Conference on Open Source Systems*. 16–23.
- [31] Melanie Warrick, Samuel F. Rosenblatt, Jean-Gabriel Young, Amanda Casari, Laurent Hébert-Dufresne, and James Bagrow. 2022. OCEAN mailing list data from open source communities. (3 2022). <https://doi.org/10.6084/m9.figshare.19082540.v2>
- [32] Andrea Wiggins, James Howison, and Kevin Crowston. 2008. Social dynamics of FLOSS team communication across channels. In *IFIP International Conference on Open Source Systems*. Springer, 131–142.
- [33] Annie T. T. Ying. 2015. Mining Challenge 2015: Comparing and combining different information sources on the Stack Overflow data set. In *The 12th Working Conference on Mining Software Repositories*. to appear.
- [34] Jean-Gabriel Young, Amanda Casari, Katie McLaughlin, Milo Z. Trujillo, Laurent Hébert-Dufresne, and James P. Bagrow. 2021. Which contributions count? Analysis of attribution in open source. In *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. 242–253. <https://doi.org/10.1109/MSR52588.2021.00036>
- [35] Ligu Yu, Srini Ramaswamy, Alok Mishra, and Deepti Mishra. 2011. Communications in global software development: an empirical study using GTK+ OSS repository. In *OTM Confederated International Conferences "On the Move to Meaningful Internet Systems"*. Springer, 218–227.