




Contents lists available at ScienceDirect

Information and Software Technology

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

Mining the Technical Roles of GitHub Users

João Eduardo Montandon ^{*a}, Marco Tulio Valente ^b, Luciana L. Silva ^c^a Technical College (COLTEC), Federal University of Minas Gerais, Belo Horizonte, Brazil^b Department of Computer Science, Federal University of Minas Gerais, Belo Horizonte, Brazil^c Department of Computer Science, Federal Institute of Minas Gerais, Ouro Branco, Brazil

ARTICLE INFO

Keywords:

Technical roles
 Technical expertise
 Developers profiles
 Machine learning
 GitHub

ABSTRACT

Context: Modern software development demands high levels of technical specialization. These conditions make IT companies focus on creating cross-functional teams, such as frontend, backend, and mobile developers. In this context, the success of software projects is highly influenced by the expertise of these teams in each field.

Objective: In this paper, we investigate machine-learning based approaches to automatically identify the technical roles of open source developers.

Method: For this, we first build a ground truth with 2284 developers labeled in six different roles: backend, frontend, full-stack, mobile, devops, and data science. Then, we build three different machine-learning models used to identify these roles.

Results: These models presented competitive results for precision (0.88) and AUC (0.89) when identifying all six roles. Moreover, our results show that programming-languages are the most relevant features to predict the investigated roles.

Conclusion: The approach proposed in this paper can assist companies during their hiring process, such as by recommending developers with the expertise required by job positions.

1. Introduction

Software development is a human-intensive activity, which makes the success of software projects highly influenced by the quality and expertise of their development workforce [1]. Besides, software systems are complex engineering artifacts, which demand high levels of technical specialization in different areas [2]. These conditions make IT companies focus on creating cross-functional teams with experts in several positions, such as databases, security, frontend design, backend design, mobile apps, etc.

Recently, these companies are increasingly using social coding platforms—such as GitHub—to search for suitable candidates to their opened positions [3–5]. Previous works have proposed solutions to use the data of such platforms to discover some skills of these candidates. They mainly focus on identifying which technologies do software developers most work on, such as libraries, frameworks, and languages [6–14]. A few others have dedicated to discerning developers' soft skills, including communication, teamwork, responsibility, etc [5,6,15]. Lastly, some studies aimed on identifying the roles of developers in open source projects, like bug fixer, bug triager, core developer, and

tester [16–23].

Usually, IT companies organize their development teams accordingly to the technology the developers master, e.g., *frontend*, *backend*, *mobile*, and others [24]. For instance, *frontend* developers are specialized on the application's interface, as opposed to *backend* who are responsible for core features. In this context, we currently lack approaches for inferring developers' technical roles. Therefore, our key goal in this paper is to *identify the technical roles played by developers using information available in open source platforms*.

Does the identification of technical roles matter?

We define technical roles as the ones derived from expertise on particular programming technologies (e.g., programming languages) and/or architectural components (e.g., frontend frameworks). In this sense, distinguishing candidates' technical roles are important as they represent a good proxy for the technologies and techniques each one master [24]. For instance, data scientists might have a deeper understanding of data analyzing techniques. By contrast, frontend developers should master Web APIs concepts, such as RESTful APIs, AJAX, etc. In

* Corresponding author.

E-mail addresses: joao.montandon@dcc.ufmg.br (J.E. Montandon), mtov@dcc.ufmg.br (M.T. Valente), luciana.lourdes.silva@ifmg.edu.br (L.L. Silva).<https://doi.org/10.1016/j.infsof.2020.106485>

Received 5 March 2020; Received in revised form 27 October 2020; Accepted 4 November 2020

Available online 12 November 2020

0950-5849/© 2020 Elsevier B.V. All rights reserved.

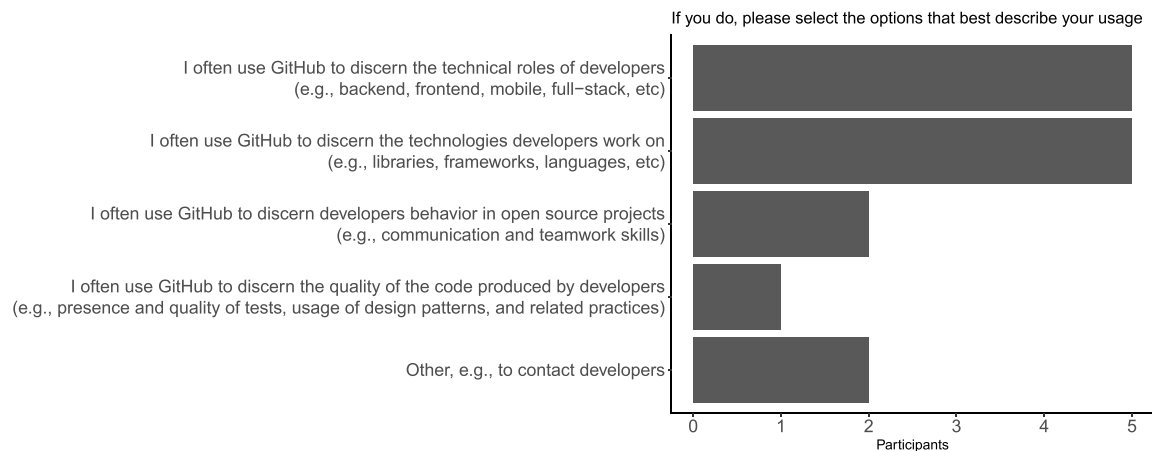


Fig. 1. Survey responses from the recruiters that use GitHub in their hiring process.

fact, source code hosting platforms—e.g., GitHub—are currently looking for alternatives to make their users profile richer by better expressing their skills, accomplishments, and interests.¹

Indeed, technical roles are one of the first aspects considered by recruiters when hiring developers for a specific position. To provide evidence of this claim, we inspected the jobs listed by Stack Overflow Jobs²—the part of the Q&A forum where companies post job offers—on October 25, 2018. Our key finding can be summarized as follows:

3234 out of 5027 Stack Overflow job posts (64%) include in their description at least one of the technical roles investigated in this paper.

Additionally, we conducted a preliminary survey to better understand how GitHub is used when hiring software developers. For this, we contacted all Stack Overflow users self-described as technical recruiters who were active in the platform since 2018, and that publicly provided their e-mail for contact. We emailed these users with a brief survey where we asked:

Do you often use GitHub in your hiring process? If you do, please select the options that best describe your usage (you can mark multiple options).

We provided four distinct options for the second question, which were formulated after consulting the literature about software expertise [3,5,14,25]. To avoid possible bias introduced by the order of the options, we configured our survey system to present them in random order for each contact. In total, we contacted 30 recruiters and received 7 answers (23% response rate). Six developers confirmed they use GitHub for hiring purposes. The survey results are summarized in Fig. 1.

Considering our intention in this paper, the main finding of this survey can be summarized as follows:

Five out of six technical recruiters use GitHub to discern the *technical roles* of the candidates, which is exactly our central goal in this paper.

Although not directly linked with our goals here, another frequent usage of GitHub includes identifying which technologies do developers work on (also with 5 answers) and to discern developers' behavior in open source projects (2 answers).

Proposed study

In this study, we investigate an approach centered on supervised machine-learning techniques to identify developers' technical roles by

considering their contributions to public GitHub projects. More specifically, we infer developers' technical roles by using features extracted from their public GitHub profiles (e.g., programming languages, short bio, etc) and their GitHub projects. We first build a ground truth with the technical roles of 2284 GitHub developers. Then, we executed the proposed technique to identify developers in six popular technical roles:³ *backend*, *frontend*, *full-stack*, *mobile*, *devops*, and *data science*. Lastly, we evaluated our approach by answering four research questions.

(RQ.1) How accurate are machine learning classifiers on identifying developers' technical roles? Our best model scored meaningful results for precision (0.75) and AUC (0.70), while it reported lower ones for recall and F1 (0.49 and 0.59, respectively). However, in the context of this work, precision-based metrics gain more importance, as technical recruiters are more interested in correctly identifying candidates that fit their needs (i.e., precision) instead of identifying them (i.e., recall). When considered independently, our models provide the best results for *data science* and *frontend* roles (0.86 and 0.77; precision). By contrast, we observed lower results for *backend* (0.62; precision).

(RQ.2) What are the most relevant features to distinguish technical roles? Features associated with programming languages are relevant for all roles. Individually, *data science*'s top-10 most relevant features have the largest representativeness rate: 33.2%. By contrast, *backend* top-10 features scores only 6.8% of the total.

(RQ.3) Do technical roles influence each other during classification? In this RQ, we investigate the gains achieved with the use of Classifier Chains in our problem. Suppose that a developer is predicted as having a role R_1 . After making this prediction, R_1 is used as input to the classifiers of R_2 , R_3 , R_4 , and R_5 roles (assuming we analyze five roles, as in our first three RQs). Specifically, we check the improvements achieved with Classifier Chains by setting up and training 120 different classification models (i.e., covering all possible permutations of five roles). Overall, such models present minor improvements in recall (+ 0.05) at the cost of precision scores (− 0.05) when compared to the ones in RQ.1.

(RQ.4) How effectively can we identify full-stack developers? Differently from the technical roles analyzed in the first three RQs, the *full-stack* role is derived from the combined expertise in *backend* and *frontend* technologies. Therefore, we conducted a separated study to identify *full-stack* developers and verify how they impact the classification of *backend* and *frontend* developers. The proposed classifier performed very well when identifying FULLSTACK, achieving 0.99 for precision and 0.71 for recall. Moreover, the addition of FULLSTACK developers significantly improved the identification of both BACKEND and frontend—0.87 and 0.86 for precision, respectively—after specific adjustments.

¹ <https://twitter.com/natfriedman/status/1133700043695886336>

² <https://stackoverflow.com/jobs>

³ Accordingly to <https://insights.stackoverflow.com/survey/2018>

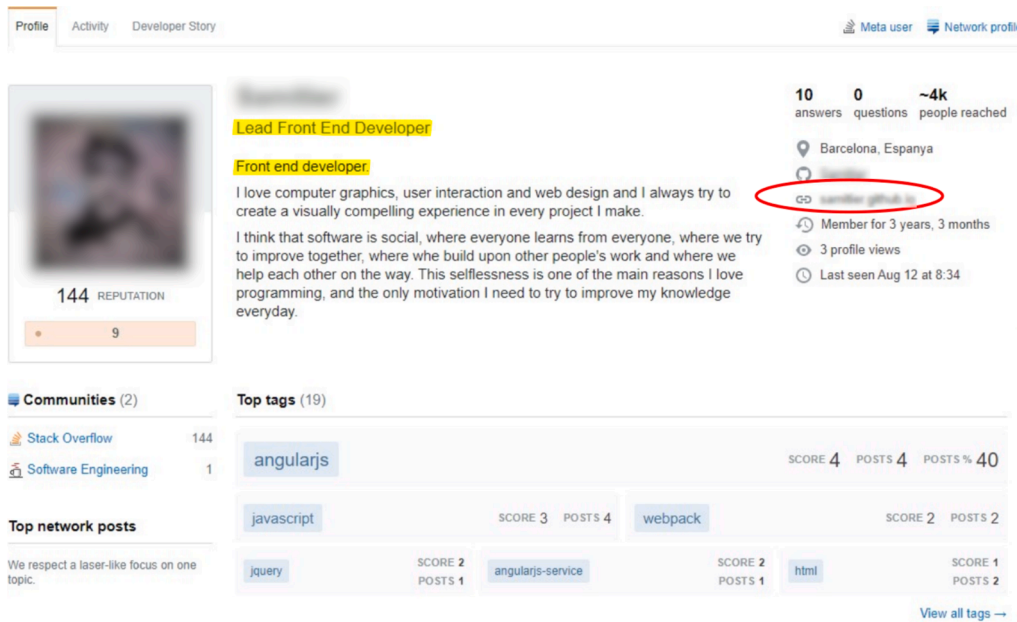


Fig. 2. Example of Stack Overflow developer profile. The data used in the ground truth creation is highlighted.

Our contributions

We show that—based on high-level features retrieved from developers’ public profiles and projects in GitHub—it is possible to infer major technical roles commonly used in industry to recognize software developers. Secondly, we make publicly available our ground truth with the roles of 2284 GitHub developers. This ground truth can motivate and support further research in the area.

Paper structure

The rest of this paper has nine sections. First, Section 2 documents the data and methods used in the study, including the ground truth creation, the features extracted from GitHub, and the machine learning setup. Section 3 presents the results for the first three research questions. Section 4 describes the exploratory study conducted on full-stack developers. In Section 5.3, we manually analyze some developers classified by our approach to qualitatively interpret the classification results. In Section 5 we discuss the implications to both practitioners and software engineering community. Section 6 discusses threats to validity and Section 7 presents Related Work. Finally, Section 8 concludes this work.

2. Study design

2.1. Technical roles

The focus of this paper is to propose an approach to unveil the technical roles of GitHub developers. We used the Annual Developer Survey conducted by Stack Overflow in 2018 to select the roles considered in our study. More than 100,000 developers from 183 countries answered this 30-minute survey where, among other questions, they answered which technical roles they associate with the tasks they normally perform. In this paper, we study the top-4 most popular roles, according to this survey: BACKEND (58%), FULLSTACK (48%), FRONTEND (38%), and MOBILE (20%). Moreover, we included DEVOPS (10.4%) and DATASCIENCE (7.7%) roles in our analysis as both are also frequently featured in the top trending IT positions. Specifically, we first focus on five roles: BACKEND, FRONTEND, DEVOPS, DATASCIENCE and MOBILE (RQ.1, RQ.2, and RQ.3). Due to its particularity, we investigate the FULLSTACK role separately (RQ.4).

2.2. Ground truth

As the proposed machine learning model relies on features extracted from GitHub to classify developers’ technical roles, we fully avoided using GitHub’s data to generate our ground truth. Instead, we relied exclusively on Stack Overflow’s data to build this ground truth, since this data would not be used further in our study. Other works use data from Stack Overflow as a reliable source to assess developers’ expertise [8,15,26]. On the other hand, after building the ground truth, our goal is to identify technical roles of GitHub users, i.e., by only considering GitHub data.

We followed a two-step process to build the ground truth, as explained next.

Stack Overflow data gathering: We used Stack Exchange Data Explorer (SEDE)⁴ to collect Stack Overflow users with GitHub profiles. SEDE is a publicly available tool that allows querying data available in Stack Exchange platform, using a web interface. On June 2nd, 2020, we queried SEDE for Stack Overflow users who have a link to GitHub. This query returned 27,051 developers. We then extracted—through regular expressions—the GitHub username from the provided links, and used GitHub GraphQL API⁵ to retrieve GitHub data (as described in Section 2.3) for developers with valid usernames. During this procedure, we discarded developers with URLs pointing to invalid GitHub usernames. We ended up with a dataset composed of 24,889 developers.

Labeling: In this step, we analyzed the profile information provided by Stack Overflow to label developers’ roles. Fig. 2 depicts an example of Stack Overflow’s user profile. Among the information available in his profile, this developer described himself as a “Lead Front End Developer”. Therefore, we relied on developers’ description text in Stack Overflow to determine their roles. More specifically, we elaborated five distinct regular expressions to identify each role, as described in Table 1. These regular expressions were configured to consider spaces, e.g., *Backend*, *Front end* and to ignore case (e.g., *Mobile* and *mobile*). Considering the profile in Fig. 2, our labeling process classify him as FRONTEND (“Lead Front End Developer” and “Front end developer.”).

Selected developers: By following the aforementioned steps, 1802

⁴ <https://data.stackexchange.com/>
⁵ <https://developer.github.com/v4/>

Table 1
Regular expressions used to identify technical roles.

Role	Regular expression	Examples
Backend	<code>/\bback.{0,1}end\b/i</code>	<i>Back end, backend</i>
Frontend	<code>/\bfront.{0,1}end\b/i</code>	<i>Frontend, front-end</i>
DevOps	<code>/\bdev.{0,1}ops\b/i</code>	<i>DevOps, dev-ops</i>
DataScience	<code>/\bdata.{0,1}scientist\b/i</code>	<i>Data Scientist, data scientist</i>
Mobile	<code>/\bmobile\b/i</code>	<i>mobile, Mobile</i>

developers were labeled with at least one technical role. We discarded 140 (7.7%) developers with less than five GitHub public repositories. The reason is that our models depend on developers being active on GitHub to infer their roles. Lastly, 1662 developers were included in our ground truth, containing 2022 roles assignments.

FRONTEND is the role with more developers (820), followed by MOBILE (453) and BACKEND (450). Table 2 shows how these labels are distributed. As we can observe, 1340 developers (80%) are associated with just one role. Most of them are FRONTEND (545; 33%), or MOBILE (352; 21%). Considering the 287 developers who have two or more roles, the majority are both BACKEND and FRONTEND (198; 69%). By contrast, DATA-SCIENCE is the one with the smallest intersection with other roles, i.e., one developer in all situations. 32 developers were labeled in three roles, most of them in BACKEND, FRONTEND, and MOBILE (23; 72%). Only three developers were assigned in four roles. Finally, we found no developers who self-labeled themselves in all five roles.

2.3. Data collection

After using Stack Overflow’s data to leverage our ground truth, we collected GitHub’s data for each developer in this golden set. This data will be later transformed into a list of features to feed our prediction models (see Section 2.4). To provide a code agnostic solution (and therefore analyze repositories in multiple programming languages), we collected mostly textual data about the developers’ profiles and the projects they own.⁶ The data we collected fits into six distinct categories:

- **Programming Language:** Previous works report that programming languages are a reliable proxy to assess developers’ expertise area [3, 5,25,27,28]. For each developer, we first retrieved her list of projects, and then extracted both *commits* and the *main programming language* of each one. Next, we derived the commits information into three dimensions: (a) the total number of commits; (b) the number of commits performed by the developer (i.e., she is the author); and (c) the rate between both of them. Lastly, we grouped the projects accordingly to the programming language collected earlier and averaged the resulting values. We ended up with the following features for each developer, in each language: *Language (total)*, *Language (author)*, and *Language (rate)*.
- **Short Bio:** Short description of the developers’ main activities, manually provided by each one. We decided to include this category in our analysis as some developers use this information to describe their field of activity, e.g., “I develop iOS/Android apps with Swift/Kotlin language”.
- **Projects’ Names:** These names may contain keywords referencing specific technologies used in the projects, e.g., *android-data-binding-1*, *docker-example*, etc. As observed in previous works, these technologies are also indicators of expertise [9,14,26,29]. Therefore, for each developer, we collected the name of her GitHub projects and merged them into a single data point.

⁶ It is important to note that we have discarded forked projects in our analysis.

Table 2
Distribution of developers among the analyzed roles.

# Roles	Role					# Devs.
	Backend	Frontend	Mobile	DevOps	DataScience	
One	✓					178
		✓				545
			✓			352
				✓		119
					✓	146
Two	✓	✓				198
	✓		✓			29
	✓			✓		9
	✓				✓	1
		✓	✓			39
Three		✓	✓			4
			✓	✓		5
				✓	✓	1
					✓	1
	✓	✓	✓			23
Four	✓	✓		✓		7
	✓	✓			✓	1
	✓	✓	✓	✓		1
Five	✓	✓	✓	✓		3

- **Projects’ Topics:** For the same reason as above, topics are generally used by GitHub developers to declare technologies and tools used by their projects, e.g., *ionic*, *gulp*, and *webpack*. We also collected the GitHub topics of the developers’ projects and merged them into a single data point as well.
- **Projects’ Descriptions:** This information can also provide clues about the technologies used by GitHub projects [30]. For instance, the description “a React.js contact manager” indicates that *ReactJS* is used in the project. For this reason, we collected the short description of each developer’s project and merged them into a single data point.
- **Projects’ Dependencies:** We decided to include the list of dependencies since 3rd-party libraries are largely used in modern software systems [7,10,14,30]. For each project, we extracted its list of dependencies using a GitHub GraphQL endpoint.⁷ Next, we assembled these features by summing up the number of times each dependency was referenced in each developer’s projects. We initially collected data on 17,556 dependencies, but we restricted ourselves to the top-1000 most used ones as they are present in 81% of all analyzed projects.

2.4. Selected features

To transform the data described in Section 2.3 into a list of features, which can be used as input for a machine learning algorithm, we applied the following transformation steps:

Correlation analysis: Initially, we ended up with 477 and 1000 features for *Programming Languages* and *Projects’ Dependencies* categories, respectively. Due to this high number, we followed a correlation analysis procedure to remove the ones with high correlations in each category. For this, we used *corr* function from *pandas*⁸ Python library to generate a Spearman correlation matrix. Then, we identified pairs of features with a correlation greater than 0.7, as previously adopted in the literature [31]; in such cases, we discarded the feature that has the highest correlation rate when compared to the other ones. In total, we discarded

⁷ <https://developer.github.com/v4/object/dependencygraphmanifest/>

⁸ <https://pandas.pydata.org/>

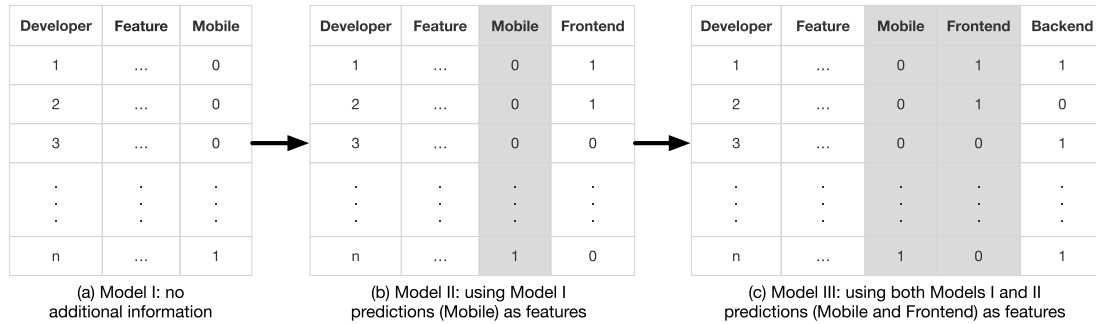


Fig. 3. Multi-label classification using classifier chains.

260 features for *Programming Languages*. Generally these correlations happened with features belonged to the same language, but covering different domains, e.g., *C (author)* and *C (rate)*, *CSS (author)* and *CSS (rate)*, etc. In most cases, we maintain the *rate*-based ones since they usually have the lowest correlation rate with the others. As for *Projects' Dependencies*, we discarded 202. Differently, the discarded dependencies are more diverse, including specific (*ionic-native-statusbar*, a specific plugin for ionic framework) and general (*unicorn*, an http webserver) purpose libraries.

Bag-of-words: Machine learning algorithms do not deal with text directly. For this reason, we used a bag-of-words to transform each textual category into a list of features. Bag-of-words is a simple and efficient technique to extract features from text [32]. In summary, each word is mapped to a feature that describes its frequency in a document. As in other studies [33–35], we performed the following steps to apply this technique on *Short Bio*, *Projects' Names*, *Projects' Topics*, and *Projects' Descriptions*. First, we manually stripped out HTML tags, punctuation, and numbers through regular expressions. Then, we removed stop words (e.g., of, him, the, and, etc.) using the default English list provided by *sklearn.text* module. Finally, we used *TfidfVectorizer*⁹ class from *sklearn* library to apply bag-of-words using the TF-IDF outcome as feature values. Following common guidelines for text processing [36], we considered only words in a certain document frequency range: [0.04, 0.15] for *Projects' Descriptions*, [0.03, 0.25] for *Projects' Names*, [0.01, 0.25] for *Projects' Topics*, and [0.01, 0.20] for *Short Bio*. These limits were defined after executing Random Forest classifier with 100 different bag-of-words configurations by randomly selecting different document frequency values in each configuration. We selected the configuration which presents improvements in most evaluated metrics. The result of

Table 3
Dataset with five target labels.

Developer	BACKEND	FRONTEND	MOBILE	DEVOPS	DATA SCIENCE
D_1	...	1	1	0	0
D_2	...	1	0	0	1
D_3	...	0	0	0	0
⋮	⋮	⋮	⋮	⋮	⋮
D_n	...	0	0	1	0

⁹ http://scikit-learn.org/stable/modules/generated/sklearn.feature_extraction.text.TfidfVectorizer.html#sklearn.feature_extraction.text.TfidfVectorizer

each configuration is available in our replication package for further reference.

After these steps (i.e., *Correlation Analysis* and *Bag-of-Words*), we ended up with 1471 features, including 798 from *Projects' Dependencies* category, 217 from *Programming Languages*, 169 from *Projects' Descriptions*, 155 from *Projects' Names*, 69 for *Short Bio*, and 63 from *Projects' Topics*.

2.5. Machine learning setup

Multi-label problem: Usually, classification problems rely on a single-label, i.e., each instance is associated with a single label. On the other hand, a multi-label classification problem can associate more than one label to each instance [35,37]. Particularly, our dataset is a five-label classification problem since there are five distinct roles each developer can be an expert on. Table 3 illustrates this dataset. In this example, developer D_1 is associated with BACKEND and FRONTEND roles. Likewise, developer D_2 is tied with BACKEND and DEVOPS. On the other hand, D_3 is associated just with DATA SCIENCE.

Problem transformation: We can approach a multi-label classification problem in different ways. For instance, we can transform a multi-label dataset into a single-label one [38,39]. Alternatively, some algorithms work directly with multi-label datasets [37]. In this paper, we use two common transformation techniques for dealing with multi-label datasets:

1. **Binary Relevance (BR):** This technique splits the original dataset into multiple binary classification problems, i.e., one binary classification problem for each label [38,40]. Then, such models are fitted independently. The results can be shown either independently (one result for each model) or aggregated (weighted average of all models). In this study, we calculated the aggregated results using a *micro-average* strategy, as it is recommended for multilabel problems.¹⁰
2. **Classifier Chains (CC):** This technique links the binary classifiers along a chain in a way that each classifier uses results from earlier ones in its predictions [39]. The goal is to take advantage of eventual dependencies that might exist among target labels. Fig. 3 illustrates a classification scenario using CC considering three technical roles: MOBILE, FRONTEND, and BACKEND. Each model includes previous predictions as features and propagates its results along the chain. Fig. 3a presents the first model in the chain (Model I). As we can see, this model makes its predictions using only the available features. Model II (Fig. 3b) predicts its values using all features plus Model I predicted

¹⁰ https://scikit-learn.org/stable/modules/model_evaluation.html#from-binary-to-multiclass-and-multilabel

labels. Finally, the process is propagated to Model III, using all features plus both Models I and II predicted labels. In this way, *CC* tackles the label independence problem faced by *BR* [41,42].

Machine learning classifier: We initially decided to use Random Forest [43] and Naive Bayes [44] classifiers to train and test our models for identifying the technical roles of GitHub users. We selected Random Forest due to its robustness to noise and outliers [45]. Moreover, Random Forest was successfully used in many application areas, including software engineering problems [31,35,46–49]. As we are dealing mostly with textual information, we also decided to include Naive Bayes in our study as this classifier generally presents good results in textual scenarios, such as spam filtering [50] and news classification [51]. We used the implementation provided by *scikit-learn* [52] for both classifiers along with 10-fold cross-validation to select the best model. Basically, cross-validation splits the data into k folds (in our case, $k = 10$), where $k - 1$ folds are used to fit the model and the remaining one is used to test the predictions.

Evaluation metrics: We use six metrics to evaluate the quality of the model's predictions: precision, recall, F1-score, AUC (Area Under the Curve), Jaccard Coefficient, and Hamming Loss. Precision measures the correctness, while recall measures the completeness of the model predictions. F1-score is the harmonic mean of precision and recall. AUC represents the area under the Precision-Recall (PR) curve. Unlike ROC, PR curves are more recommended for assessing unbalanced datasets [36]. Jaccard Coefficient—also known as multi-label accuracy [53]—is the division of the number of correctly predicted labels by all true labels. Hamming Loss is the ratio of wrong labels by the number of labels [54]. As it is a loss measure, the lower the value, the better the model. Finally, we compared all metrics against a Stratified baseline. This baseline considers the dataset distribution to perform its predictions. For instance, in a dataset where 10% of the samples are positive labels, this baseline limits its positive predictions to 10% as well.

3. Results

(RQ.1) How accurate are machine learning classifiers in identifying technical roles?

Table 4 (a) presents the general results for the first three roles using Binary Relevance (*BR*). The Random Forest classifier presented the best results overall, scoring 0.77 for precision and 0.71 for AUC. Moreover, apart from recall, Random Forest outperformed Naive Bayes for all evaluated metrics. Finally, both classifiers performed significantly better than the Stratified Baseline. The baseline results are, for example, only 0.33 for precision.

Table 4 (b) presents the classification results for each role, separately. Overall, Random Forest presented better results for precision

Table 4a
Binary relevance overall results.

Metric	Stratified Baseline	Random Forest	Naive Bayes
Precision	0.33	0.77	0.51
Recall	0.32	0.49	0.62
F1	0.33	0.59	0.56
AUC	0.41	0.71	0.59
Jaccard Coeff.	0.20	0.42	0.39
Hamm. Loss	0.32	0.16	0.24

(a) Overall results

Table 4b
Binary relevance results for each role.

Role	Precision	Recall	F1
Stratified Baseline			
BACKEND	0.28	0.28	0.28
FRONTEND	0.48	0.46	0.47
MOBILE	0.28	0.28	0.28
DEVOPS	0.08	0.06	0.07
DATA SCIENCE	0.09	0.09	0.09
Random Forest			
BACKEND	0.62	0.12	0.18
FRONTEND	0.77	0.78	0.77
MOBILE	0.78	0.38	0.51
DEVOPS	0.70	0.13	0.20
DATA SCIENCE	0.86	0.66	0.74
Naive Bayes			
BACKEND	0.33	0.40	0.36
FRONTEND	0.69	0.82	0.75
MOBILE	0.51	0.46	0.47
DEVOPS	0.24	0.47	0.32
DATA SCIENCE	0.45	0.85	0.58

(b) Results for each role

when compared to Naive Bayes, and both outperformed the Stratified Baseline. Regarding Random Forest, the classifier presented high precision rates—i.e., above 0.7—for 4 out of 5 roles: DATA SCIENCE (0.86), MOBILE (0.78), FRONTEND (0.77) and DEVOPS (0.70). On the other hand, only FRONTEND achieved results as good as for recall (0.78) and F1 (0.77). Further, the classifier scored poor results especially for BACKEND role (precision, recall, and F1 all equal to 0.28).

When it comes to Naive Bayes, we did not identify any score above 0.7 for precision; FRONTEND has the highest one (0.69), followed by MOBILE (0.51) and DATA SCIENCE (0.45). In fact, we observed such good results in three scenarios only: FRONTEND and DATA SCIENCE for recall (0.82 and 0.85, respectively), and FRONTEND for F1 (0.75). Even though, Naive Bayes was superior to the baseline in all scenarios.

Summary: Random Forest presented the best results overall (e.g., 0.77, precision; 0.71, AUC), outperforming both Naive Bayes and the baseline. When analyzed individually, DATA SCIENCE and FRONTEND scored the best results for most metrics (e.g., 0.86 and 0.77; precision), while BACKEND showed the worse ones (e.g., 0.62; precision).

(RQ.2) What are the most relevant features to distinguish technical roles?

Fig. 4 shows the top-10 most relevant features by technical role. To identify these features, we executed the Random Forest classifier for each role independently, using the same parameters as in RQ.1. We then selected the top-10 features with the highest value from the feature relevance ranking provided by each model. In Fig. 4, the colors and shapes identify the category of each feature (programming languages, projects' names, projects' descriptions, projects' topics, projects' dependencies, or short bio). Furthermore, we annotate the category associated with each feature in its description. Finally, we normalized the ranking with respect to the feature with the highest value.

Features associated with *programming languages* are largely predominant for all five roles, representing 38 out of 50 features: 7 for DATA SCIENCE, DEVOPS, and MOBILE; 8 for FRONTEND, and 9 for BACKEND. From these, 19 (50%) are rate-based, 13 (34%) consider the total number of commits, and 6 (16%) include only the commits performed by the

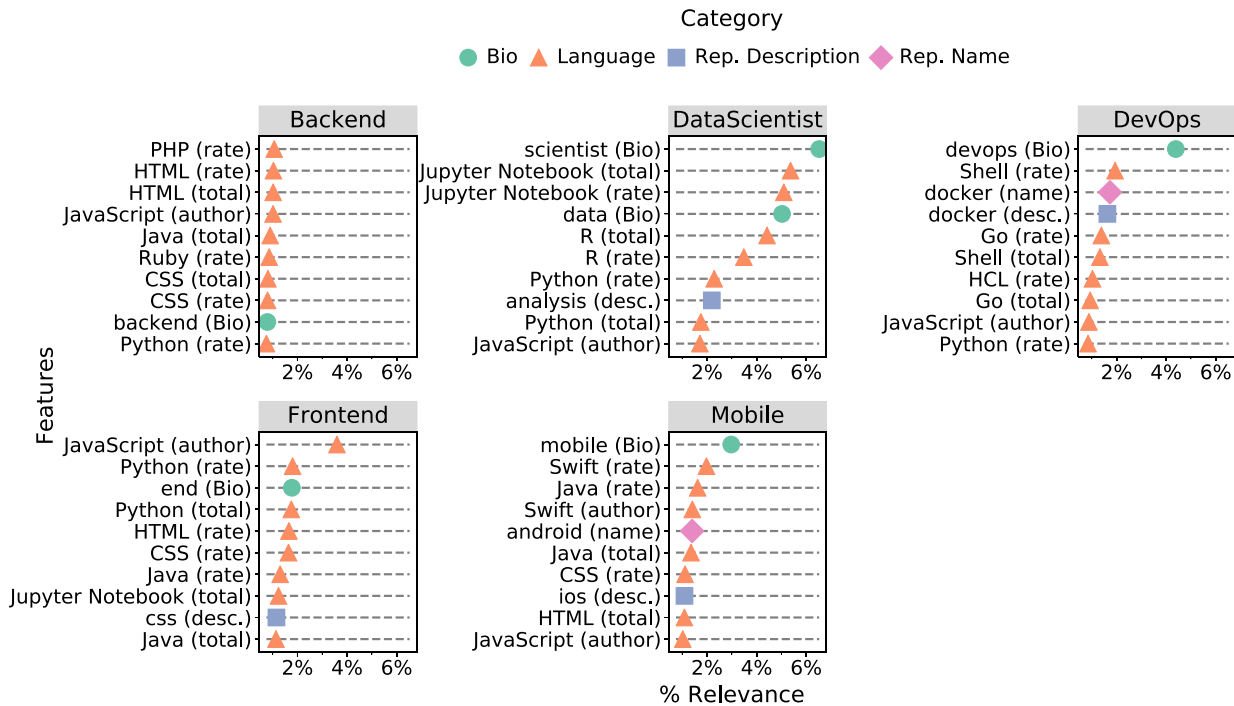


Fig. 4. Most relevant features for each technical role. Colors and shapes represent each feature category: developers bio ((Bio), green circle), projects’ descriptions ((desc.), blue square), projects’ names ((name), pink diamond), and programming languages ((rate, author, and total), orange triangle). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

author. In other words, 66% of the relevant programming languages features do take into account actual developers’ contributions.

Next, short bio appears as the second most frequent with 6 occurrences, followed by projects’ descriptions and projects’ names (3 and 2 occurrences, respectively). Nevertheless, short bio shows up as the most important feature in three roles. Most of these features directly describes the role that is analyzed: scientist for DATA SCIENCE, devops for DEV OPS, mobile for MOBILE, and backend for BACKEND. Surprisingly, projects’ topics and projects’ dependencies are not present in any ranking position.

When we analyze the results for each role, we see that DATA SCIENCE has the highest relevant features. Six features stand out with more than 3% of relevance rate: scientist (Bio) (6.5%), Jupyter Notebook (total) (5.4%), Jupyter Notebook (rate) (5.1%), data (Bio) (5.0%), R (total) (4.4%), and R (rate) (3.5%). Although these values may sound low, our model includes 1662 features. Also, these features are rather specific to the DATA SCIENCE field; from them only Jupyter Notebook (total) is also present in other roles (FRONTEND).

For MOBILE, mobile (Bio) stands out with 3.0% of relevance rate, followed by Swift (rate) (2.0%) and Java (rate) (1.6%). More specifically, three out of 10 features are directly associated to the iOS development platform (Swift (rate), Swift (author), and ios (desc.)). Likewise, other three features are linked to the Android platform: Java (rate), android (name), and Java (total). The features related to iOS and Android development represent 4.5% and 4.3% of the top-10 listed features, respectively.

Regarding DEV OPS, we observe that devops (Bio) plays a prominent position in the ranking, with 4.4%. The next ones are Shell (rate) (1.9%), docker (name) (1.7%), and docker (desc.) (1.6%). FRONTEND ranking presents a similar distribution, where JavaScript (author) is the prominent feature with 3.6%, whereas the next five are valued between 1.8% (Python (rate)) and 1.6% (CSS (rate)).

Finally, we observe that no feature stands out in BACKEND ranking. The most relevant feature is PHP (rate), with 1.1%, while Python (rate) is the 10th feature in the ranking, with 0.8%. Even though, 5 out of the top-10

features are linked with backend development: PHP (rate), Java (total), Ruby (rate), backend (Bio), and Python (rate).

Summary: Features related to programming languages are predominant for all five roles. In DATA SCIENCE role, six features stand out: scientist (Bio), Jupyter Notebook (total), Jupyter Notebook (rate), data (Bio), R (total), and R (rate). On the other hand, fewer features are presented such importance in other roles: devops (Bio) for DEV OPS, mobile (Bio) for MOBILE, JavaScript (author) for FRONTEND. Lastly, no feature stands out from the others for BACKEND.

(RQ.3) Do technical roles influence each other during classification?

As explained in Section 2.5, Classifier Chains (CC) is a technique to deal with multi-label classification problems. When using this technique, Binary Classifiers are ordered in such a way that each model uses previous predictions as input. The rationale is that, given two labels A and B, if A influences B then we might want to use A to improve B’s predictions.

Particularly, we executed the CC technique with Random Forest for all possible role combinations. In total, the classifier was executed for 5! combinations, i.e., 120 different combinations. Fig. 5 presents the results in six different graphics, one for each metric considered in this study. For comparison reasons, the results are reported in relation to RQ.1, which is used as a baseline in this RQ and it is represented by a dashed line in the figure. The horizontal axis represents each CC configuration and the vertical axis the number of points each metric has above or below the baseline. This means that the CC technique performed better than the baseline when its results are above the dashed line, or worse otherwise.

Overall, all models show a minor increase of recall up to almost + 0.05 (configurations 72, 89, and 96), at a cost of lower precision (− 0.05; configurations 47, 80, 93, and 113). We also observe minor improvements for both F1 and Jaccard Coefficient (+0.02, configurations

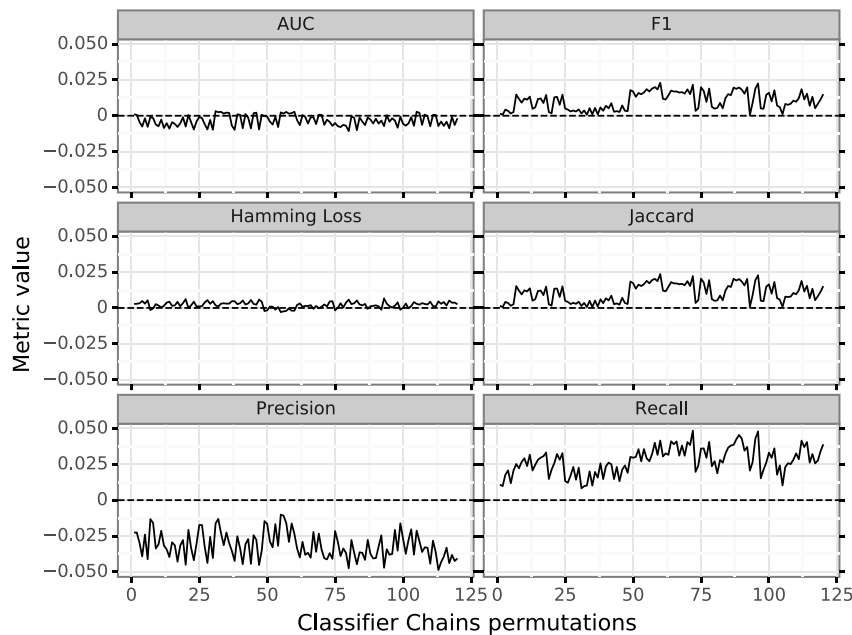


Fig. 5. Overall results using Classifier Chain, reported in relation to the RQ.1 baseline (i.e., dashed line).

60, 72, and 96; both), while AUC presented a minimal decrease (−0.01, configurations 80 and 83). By contrast, we did not find any relevant change to the Hamming Loss score. Interestingly, the highest improvements for recall, F1, and Jaccard—and also some of the minor setbacks for precision—happened between configurations 49 and 60, when MOBILE is classified first, followed by BACKEND or FRONTEND.

Summary: The results reveal that including technical roles predictions to the classifying process does not bring significant improvements. Overall, we obtained minor improvements for recall, Jaccard, and F1 on the cost of precision. These results are more consistent when MOBILE is classified first, followed by BACKEND or FRONTEND.

4. Understanding the FULLSTACK role

Differently from the previously analyzed roles, FULLSTACK developers are defined by their ability to work through the whole application stack. In fact, the industry sees a FULLSTACK developer as someone who can perform both FRONTEND and BACKEND tasks. The following quote, retrieved from a popular Medium post,¹¹ illustrates this point:

Being a Full-Stack developer [...] means that you are able to work on both sides and understand what is going on when building an application.

Due to this particularity, we decided to investigate the FULLSTACK role in a specific section. First, we extended the ground truth described in Section 2.2 to include developers self-identified as FULLSTACK on Stack Overflow.¹² The new ground truth is composed by 2284 developers; 853 of them are FULLSTACK. From these, 783 were not labeled either as BACKEND or FRONTEND, which contradicts the FULLSTACK’s definition. To fix this inconsistency, we adapted the labeling process to consider every FULLSTACK developer as an expert in both BACKEND and FRONTEND roles.

After applying the same process described in Sections 2.3 and 2.4,

¹¹ <https://medium.com/coderbyte/a-guide-to-becoming-a-full-stack-developer-in-2017-5c3c08a1600c>, Accessed at 2019-04-04.

¹² We used the following regex to label developers as FULLSTACK: `/\bfull\.(0,1)stack\b/i`.

Table 5

Binary relevance results after including FULLSTACK, random forest only.

Role	Precision	Recall	F1
FULLSTACK	0.99	0.71	0.83
BACKEND	0.87 (+ 0.25)	0.63 (+ 0.51)	0.73 (+ 0.55)
FRONTEND	0.86 (+ 0.09)	0.89 (+ 0.11)	0.87 (+ 0.10)
MOBILE	0.80 (+ 0.03)	0.34 (− 0.04)	0.47 (− 0.04)
DEVOPS	0.75 (+ 0.05)	0.06 (− 0.07)	0.11 (− 0.09)
DATA SCIENCE	0.86 (+ 0.00)	0.62 (− 0.04)	0.71 (− 0.03)
Overall¹³	0.88 (+ 0.11)	0.69 (+ 0.20)	0.77 (+ 0.18)

this extended dataset ends up with 1567 features: 819 from *Projects’ Dependencies* category, 219 from *Programming Languages*, 212 from *Projects’ Descriptions*, 146 from *Projects’ Topics*, 101 from *Projects’ Names*, and 70 from *Short Bio*.

(RQ.4) How effectively can we identify full-stack developers?

Table 5 presents the classification results for Random Forest using the Binary Relevance approach (RQ.1) and including FULLSTACK. Besides, we included the results for each of the five previous roles, along with the difference for the one described in RQ.1.

Overall, all metrics reported significantly better results after including FULLSTACK developers in the ground truth. Precision, recall and F1 increase to 0.88 (+ 0.11), 0.69 (+ 0.20), and 0.77 (+ 0.18), respectively. Likewise the new AUC, Jaccard Coefficient, and Hamming Loss reach 0.89 (+ 0.18), 0.69 (+ 0.10), and 0.13 (− 0.03), respectively.

Considering each role, we observe that FULLSTACK presented very high results: its precision is 0.99, recall is 0.71, and F1 is 0.83. When it comes to the other roles, we noticed an interesting improvement for both FRONTEND and BACKEND. FRONTEND’s value for F1 increased from 0.77 to 0.87 (+ 0.10), mainly due to recall improvement (from 0.78 to 0.89; + 0.11). Most important, BACKEND gained 55 points in F1 when compared to the one in RQ.1 (from 0.28 to 0.73). Recall is also largely responsible for this difference as it scored 0.63 against 0.28 before; an increase of 51 points. By contrast, we observe a minor negative impact in F1 for MOBILE

(− 0.04), DEVOPS (− 0.09), and DATASCIENCE (− 0.03). Even so, the precision score for these roles has either increased (MOBILE, + 0.03; DEVOPS, + 0.05) or stayed the same (DATASCIENCE).

We accredit this side effect as a consequence of expanding the dataset with the new FULLSTACK developers.¹³ Particularly, in our second dataset, we found several FULLSTACK developers with missing BACKEND or FRONTEND labels (more precisely, 783 developers, as we mentioned before). In other words, once they are FULLSTACK, for these developers it is implicit that they should be also viewed as BACKEND or FRONTEND. Consequently, after labelling FULLSTACK developers as BACKEND and FRONTEND we provided new information to the classifiers, which allowed them to improve their predictions.

Summary: The proposed classifier performed very well when identifying FULLSTACK developers (*precision* = 0.99). Moreover, FULLSTACK developers have an interesting collateral effect, since they contributed to improving results of both BACKEND (+25%) and FRONTEND (+9%).

5. Discussion

In this section, we start discussing the implications of our work to both academia and practitioners (Section 5.1). Then, we argue about how the results for precision and recall should be interpreted in the context of this study (Section 5.2). Lastly, we analyze the performance of our approach by inspecting developers profiles in three different scenarios (Section 5.3).

5.1. Implications

We shed light on the importance of unveiling technical roles; an expertise topic not yet extensively investigated by the software engineering community. More specifically, we performed two preliminary studies to demonstrate the importance of this topic. First, we inspected posts from Stack Overflow Jobs platform and find out that 64% of 5027 offers are for one of the six technical roles studied in this paper. Second, we conducted a survey with technical recruiters to understand which characteristics do they look for in GitHub profiles. Five out of six recruiters indicated they search for clues to discern the *technical roles* of the candidates.

Besides, the proposed machine learning approach has most of its usage in hiring processes. In this context, when hiring software developers, technical recruiters can benefit from the technical roles inferred by the proposed models. This usage might occur in two main ways. First, by *proactively* identifying developers with the skills expected by existing job positions in the company. Secondly, by *reactively* evaluating the profile of candidates who have already applied to existing job positions, to assure they have the expected skills. In both cases, the inferred technical roles should be used as a piece of additional information during the hiring process, which certainly includes other selection instruments, such as technical interviews, reference letters, etc.

5.2. A note on precision and recall

In general, our approach showed its effectiveness in revealing the technical roles of software developers given their GitHub profiles. For instance, our model scored 0.88 for precision and 0.89 for AUC when considering all six roles analyzed in this study (see Table 5); by contrast we achieved a lower result for recall (0.69). When analyzed individually, we observe that the precision for each role ranges from 0.75 (DEVOPS) to 0.99 (FULLSTACK). In fact, except for DEVOPS all other roles

achieved at least 80% of precision. On the other hand, only FRONTEND reached similar value for recall (0.89); the other ones remained between 0.06 (DEVOPS) and 0.71 (FULLSTACK).

Essentially, precision answers the following question “*What proportion of positive identifications was actually correct?*”, whereas recall seeks to answer “*What proportion of actual positives was identified correctly?*”.¹⁴ In the context of technical recruiters, we advocate that precision is more important than recall as they normally need to identify a small set of candidates for the existing positions; i.e., they do not need to locate all developers matching their job positions in the GitHub universe. Furthermore, companies avoid at most a false positive (a bad hiring), as it is more expensive [55,56]. On the other hand, we understand that such limitation might represent an issue for software developers, as qualified developers can eventually be discarded for not scoring meaningful results in the features considered in this study.

5.3. Manual analysis

We manually inspected developers profiles to analyze their predictions so we can better understand the proposed classification results. Specifically, we analyze developers in three distinct scenarios: correct classifications (True Positive); wrongly classified in a given role (False Positive); and not classified in a given role, despite being so (False Negative).

5.3.1. True positive scenario

Developers in this group have had their roles correctly predicted by our approach. In this scenario, we inspected developer D_{844} , classified as MOBILE Developer. D_{844} owns 22 public GitHub projects, most of them related to mobile development; four projects have *Swift* as their main programming language, and other five have the word *Android* in their description. Moreover, another four projects are implemented in Java and contain experiments about chart animations in Android. In total, 19 out of 22 projects are directly related to mobile development. In his personal website, D_{844} describes himself as “Android & iOS Engineer | App Maker”. Through D_{844} ’s GitHub page, we can also reach his LinkedIn profile and find out he has been working with mobile development since 2014.

5.3.2. False positive scenario

This group contains developers wrongly predicted as experts in any of the analyzed roles. We inspected the profile of D_{1341} , which is not a DEVOPS developer, but was predicted as such. In his LinkedIn profile, D_{1341} identifies himself as a “Data Engineer”, also mentioning he works in this position since 2016. D_{1341} received most endorsements in *Python* (19), *Data Science*, and *Machine Learning* (18, both). By contrast, D_{1341} has few endorsements in DEVOPS-specific technologies: one for *AWS* and two for *Linux*. He also obtained certifications in “Machine Learning” and “Scalable Microservices with Kubernetes”. D_{1341} is very active on GitHub, performing 711 contributions in the last year. This developer owns 30 projects in 13 different languages, most of them in *Python* (five projects). Further, *Docker* appears in the description of two other projects described as an assistant tool for data analysis. Therefore, due to his limited experience with DEVOPS-based tools and languages, D_{1341} should not be classified as a DEVOPS.

5.3.3. False negative scenario

Developers are considered False Negative when they are not classified as experts in a given role, despite being so. For this scenario, we examined D_{68} , who is a FRONTEND developer, but was not identified as such. In his short bio, D_{68} describes himself as a “Frontend Engineer”. On

¹³ Overall value is calculated using the same strategy as described in Section 2.5.

¹⁴ Questions retrieved from <https://developers.google.com/machine-learning/crash-course/classification/precision-and-recall>

GitHub, this developer has performed 1018 contributions over the last year, where 51% of them were commits. Even though, D_{68} is the owner of only 12 projects; five of them are directly related to frontend technologies (e.g., *JavaScript*, *vue.js*, etc). The remaining projects are implemented in other languages, such as *Java*, *PHP*, and *Python*. Moreover, his projects' descriptions do not mention the use of `FRONTEND` technologies. In fact, five projects do not contain any description at all. In other words, the *programming languages* and *short bio* used by D_{68} are the only features directly related to `FRONTEND`, but their presence was not enough to label him correctly.

False negative predictions gather our attention as they directly impact the recall's performance. As our classifiers scored lower values for this metric in `MOBILE` and `DEVOPS` roles, we decided to extend this analysis by inspecting 10 new randomly selected developers, five in each role. Overall, we observe these developers do not hold specific information to classify them in such roles. For instance, the five selected `DEVOPS` developers maintain 70 projects in total. However, only one is mainly implemented using *Shell Script*, which is the second most relevant feature for this role (see Fig. 4). One developer mentioned the *devops* keyword in his bio, but this was not reinforced by the other relevant features; only one of his 18 projects is implemented using a relevant language (*Python*). Similar behavior is noted when analyzing `MOBILE` developers. The majority of the projects analyzed are not written in `MOBILE`-based languages, such as *Swift* and *Java*. Interestingly, most of them are implemented with `FRONTEND`-based ones, such as *JavaScript* and *TypeScript*. We analyzed this characteristic further and found out that some of these projects rely on cross-platform `MOBILE` frameworks—e.g., *cordova*, *ionic*, etc—which are not in the top-10 most relevant features. In fact, two out of the five developers explicitly stated they are `FRONTEND` developers with experience in building mobile web apps. Lastly, relevant keywords like *android*, *ios*, and *mobile*, are indeed mentioned in the repositories of one developer but in a format that prevents their extraction by our bag-of-words tokenization process, e.g., *iOSUIAutomation*, *MyFirstAndroid*, etc.

6. Threats to validity

The following issues are possible threats to our results:

Target roles:We analyzed six technical roles in this paper. Although they represent a restricted number of roles, we selected the ones among the most popular, based on a recent large-scale survey conducted by Stack Overflow. Moreover, they cover 64% of the jobs recently listed at Stack Overflow jobs, as mentioned in the Introduction.

Ground truth:Our dataset is restricted to developers who have profiles on both Stack Overflow and GitHub. Besides, the ground truth is based on developers' public activities on GitHub. Evidently, they represent a subset of all activities of several developers. This limitation may increase false negatives in particular scenarios. For example, a developer is *mobile* but the publicly available data characterizes him/her as *backend* developer. A potential direction for tackling this problem could be the usage of developers' activities from additional data sources. Furthermore, we followed an automated process to label the 2284 developers in the ground truth. This process is subjected to bias since it relies on the description provided by the developer on Stack Overflow. Further work should consider a semi-automatic labeling strategy [57] or semi-supervised learning techniques [58] to evaluate the accuracy of our labeling process.

Multi-label classification:Other classification techniques can be applied in multi-label problems, e.g., Label Powerset [59]. However, we rely on the two most used techniques to handle multi-label classification. Moreover, we used two well-known classification algorithms—Random Forest and Naive Bayes. Despite that, further work should consider other classification algorithms, such as XGBoost [60].

Thresholds:As usual in machine learning studies, we acknowledged that our results depend on different thresholds, which are used for example to discard correlated features (Pearson ≥ 0.7), to limit the

number of features used in projects' descriptions ([0.04, 0.15]), projects' names ([0.01, 0.25]), projects' topics ([0.01, 0.25]), and short bio ([0.01, 0.20]), as reported in Section 2.4. As a general guideline, we always use conservative thresholds. Furthermore, we experimented with other threshold values in all cases and selected the ones that presented the best results.

7. Related work

Related work approaches the expertise of software developers in two major standpoints. The first one, known as *domain expertise*, address expertise in the context of specific projects, ignoring developers' technical background. For instance, these works involve the identification of source code file owners [19,21,61], core developers [20,22,62,63], or organizational and contribution roles [18,23,64]. The second category, known as *technical expertise*, focuses on revealing developers' technical abilities. Most studies in this second group focus on identifying expertise in 3rd-party technologies [7,8,13,14,29,30], source code quality [3,27], and soft skills [25].

7.1. Domain expertise

Avelino et al. [63] relied on developers commit activity to estimate truck-factors (i.e., a source code authorship metric). Furthermore, they tested their approach against 119 open source systems to identify their core developers. Honsel et al. [20] use multi-dimensional Hidden Markov Models to predict the contribution behavior of developers. For each developer in a project, they collect data (commits, bug comments, and mailing lists) and classify them into *core*, *major*, and *minor* classes. Previously, Alonso et al. [65] build a rule-based model to derive developers' expertise based on the frequency of their commits. Their model links a list of categories—extracted from the directory structure of each project—to each developer based on the frequency of changes they perform in each directory. Da Silva et al. [19] introduced a technique to leverage expertise in system's artifacts at source code method-level. The authors rely on a GPU-based approach to leverage such expertise information. Kagdi et al. [66] mine software repositories to recommend developers who master a particular component of a project. Other works investigated expertise in the context of organizational and contribution roles [18,23]. For example, Bhattacharya et al. [18] analyzed the use of graph models to classify developers on seven organizational roles: *patch tester*, *assist*, *triager*, *bug analyst*, *core developer*, *bug fixer*, and *patch-quality improver*. Lastly, Yu and Ramaswamy [61] also mined software repositories to predict developers' roles, but they focus on the interactions among them (e.g., frequency of co-editing). Note that the works in this subsection aim at recommending developers to maintain specific projects. This scenario fits well in the OSS development field. On the other hand, the IT industry is also interested in identifying developers' technical abilities regardless their previous knowledge in specific projects [6, 67,68].

7.2. Technical expertise

In fact, most related work addresses the technical expertise problem towards *3rd-party technologies*. In recent work, we [14] proposed two methods to assess the expertise level of 575 developers in three popular JavaScript libraries. Basically, we used features collected from GitHub activity—e.g., commits frequency, code churn, etc.—to classify developers into three expertise levels: *novice*, *intermediate*, and *expert*. Previously, Teyton et al. [7,69] proposed a search engine to locate—based on syntactical changes—developers with experience on some technical skills, e.g., *Java*, *Linux*, *Apache*, etc. Similarly, Wan et al. [13] developed a probabilistic model to assess developers' coding abilities on software technologies. Lastly, Greene and Fischer [29] also proposed a tool that extracts and presents developers' skills data from GitHub, including skills on programming languages, libraries, and frameworks.

Other authors proposed matching the technologies that developers use with the ones required in specific tasks. For example, Hauff and Gousios [30] extracted user activity from GitHub and job advertisements to recommend matching open jobs to developers. They leverage a vector space model—using data extracted from GitHub profiles and job advertisements of a large UK job portal—and rely on cosine similarity to match developers and jobs. Venkataramani et al. [8] proposed a model that, by using developers' activity information on GitHub, recommends Stack Overflow questions for these developers to answer. By contrast, our work focus on detecting developers' expertise into a higher granularity level, i.e., technical roles.

Other works approached technical expertise differently. For example, Agrawal et al. [21] used clustering techniques to analyze which technical factors are more relevant for *backend* developers and *testers*. Differently from our work, the authors consider only commits, timestamps, and file names, whereas we collect a wider set of features, i.e., programming languages, developers' bio, and projects' metadata. Other researchers investigated the importance of *source code quality* when analyzing candidates for opened positions. Marlow et al. [3,27] interviewed both developers and recruiters to understand which information is more relevant when evaluating online profiles in developer communities. In general, interviewees reported some cues used by them when assessing these profiles, including, which programming languages and libraries do developers use, as well as the readability of their code. Lastly, there are studies mentioning the importance of *soft skills* in the software development environment. Baltes and Diehl [25] presented a conceptual theory structuring the relevant factors in software development. Among the several abilities listed in their work, the authors highlighted the importance that proper personal skills—such as communication, critical thinking, and communication—has on software development.

8. Conclusion

The increasing complexity and relevance of modern software systems are fostering the specialization of software developers in particular components and technologies. As a result, when hiring developers, companies usually do not look for developers with a broad range of skills, but for ones who can work with specific technologies and in specific tasks. Motivated by this context, in this paper we described an approach centered on supervised machine learning to predict six widely popular technical roles of developers nowadays: `BACKEND`, `FRONTEND`, `MOBILE`, `FULLSTACK`, `DATASCIENCE`, and `DEVOPS`. Using features extracted from the public profiles of GitHub users, we obtained great results for identifying all six roles in terms of precision (0.88) and AUC (0.89). By contrast, we observed lower results for `DEVOPS` and `MOBILE` regarding recall, e.g., 0.06 and 0.34, respectively. Even so, we believe that this study can offer good assistance to technical recruiters as, in their context, identifying correct candidates (precision) is more relevant. As future work, we plan to extend our research in the following lines: (a) improve the classification process by adding new features (e.g., frameworks) and by testing others classifiers (e.g., XGBoost [60]); (b) validate our results (and the practical usage of our approach) through a large-scale survey with real developers, and (c) implement a browser plug-in to inform the technical roles of GitHub existing profiles.

Replication package: Our data and our scripts—in a Jupyter Notebook format—are publicly available at the following URL: <https://doi.org/10.5281/zenodo.3986172>.

CRedit authorship contribution statement

João Eduardo Montandon: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing - original draft, Writing - review & editing. **Marco Tulio Valente:** Conceptualization, Methodology, Validation, Resources, Writing - original draft, Writing - review & editing. **Luciana L. Silva:**

Conceptualization, Methodology, Validation, Writing - original draft, Writing - review & editing.

Declaration of Competing Interest

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.

References

- [1] T. DeMarco, T. Lister, *Peopleware: Productive Projects and Teams*, second ed., Dorset House, 1999.
- [2] F.P. Brooks Jr., *The Mythical Man-Month*, Addison-Wesley, 1995.
- [3] J. Marlow, L. Dabbish, Activity traces and signals in software developer recruitment and hiring. Conference on Computer Supported Cooperative Work (CSCW), ACM Press, 2013, pp. 1–11.
- [4] A. Capiluppi, A. Serebrenik, L. Singer, Assessing technical candidates on the social web, *IEEE Softw.* 30 (1) (2013) 45–51.
- [5] A. Sarma, X. Chen, S. Kuttal, L. Dabbish, Z. Wang, Hiring in the global stage: profiles of online contributions. International Conference on Global Software Engineering (ICGSE), 2016, pp. 1–10.
- [6] L. Singer, F. Figueira Filho, B. Cleary, C. Treude, M.-A. Storey, K. Schneider, Mutual assessment in the social programmer ecosystem. ACM Computer Supported Cooperative Work (CSCW), 2013, pp. 103–116.
- [7] T. Teyton, J.-R. Falleri, F. Morandat, X. Blanc, Find your library experts. Working Conference on Reverse Engineering (WCRE), 2013, p. 202—211.
- [8] R. Venkataramani, A. Gupta, A. Asadullah, B. Muddu, V. Bhat, Discovery of technical expertise from open source code repositories. International Conference on World Wide Web (WWW), 2013, pp. 97–98.
- [9] C. Teyton, M. Palyart, J.-R. Falleri, F. Morandat, X. Blanc, Automatic extraction of developer expertise categories and subject descriptors. International Conference on Evaluation and Assessment in Software Engineering (EASE), 2014, pp. 1–10.
- [10] W. Huang, W. Mo, B. Shen, Y. Yang, N. Li, CPDScore: modeling and evaluating developer programming ability across software communities. International Conference on Software Engineering and Knowledge Engineering (SEKE), 2016, pp. 87–92.
- [11] E. Constantinou, G.M. Kapitsaki, Identifying developers' expertise in social coding platforms. Euromicro Conference on Software Engineering and Advanced Applications (SEAA), 2016, pp. 63–67.
- [12] Z. Wang, H. Sun, Y. Fu, L. Ye, Recommending crowdsourced software developers in consideration of skill improvement. IEEE/ACM International Conference on Automated Software Engineering (ASE), 2017, pp. 717–722.
- [13] Y. Wan, L. Chen, G. Xu, Z. Zhao, J. Tang, J. Wu, SCSSMiner: mining social coding sites for software developer recommendation with relevance propagation, *World Wide Web* 21 (2018) 1–21.
- [14] J.E. Montandon, L.L. Silva, M.T. Valente, Identifying experts in software libraries and frameworks among GitHub users. International Conference on Mining Software Repositories (MSR), 2019, pp. 1–12.
- [15] I. Ahmed, What makes a good developer? An empirical study of developers' technical and social competencies. IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC), 2018, pp. 319–321.
- [16] Y. Ye, K. Kishida, Toward an understanding of the motivation open source software developers. International Conference on Software Engineering (ICSE), 2003, pp. 419–429.
- [17] G. Robles, J.M. Gonzalez-Barahona, I. Herraiz, Evolution of the core team of developers in libre software projects. Working Conference on Mining Software Repositories (MSR), 2009, pp. 167–170.
- [18] P. Bhattacharya, I. Neamtiu, M. Faloutsos, Determining developers' expertise and role: a graph hierarchy-based approach. International Conference on Software Maintenance and Evolution (ICSME), 2014, pp. 11–20.
- [19] J.R. Da Silva, E. Clua, L. Murta, A. Sarma, Niche vs. breadth: calculating expertise over time through a fine-grained analysis. International Conference on Software Analysis, Evolution, and Reengineering (SANER), 2015, pp. 409–418.
- [20] V. Honsel, S. Herbold, J. Grabowski, Hidden Markov models for the prediction of developer involvement dynamics and workload. International Conference on Predictive Models and Data Analytics in Software Engineering (PROMISE), 2016, pp. 1–10.
- [21] K. Agrawal, M. Aschauer, T. Thonhofer, S. Bala, A. Rogge-Solti, N. Tomsich, Resource classification from version control system logs. IEEE International Enterprise Distributed Object Computing Workshop (EDOCW), 2016, pp. 249–258.
- [22] M. Joblin, S. Apel, C. Hunsen, W. Maurer, Classifying developers into core and peripheral: an empirical study on count and network metrics. International Conference on Software Engineering (ICSE), 2017, pp. 164–174.
- [23] E. Constantinou, G.M. Kapitsaki, Developers expertise and roles on software technologies. Asia-Pacific Software Engineering Conference (APSEC), 2017, pp. 365–368.
- [24] J.E. Montandon, C. Politowski, L.L. Silva, M.T. Valente, F. Petrillo, Y.-G. Guéhenec, What skills do IT companies look for in new developers? A study with stack overflow jobs, *Inf. Softw. Technol.* 129 (2021) 1–6.
- [25] S. Baltes, S. Diehl, Towards a theory of software development expertise. Symposium on the Foundations of Software Engineering (FSE), 2018, pp. 1–14.

- [26] R. Saxena, N. Pedanekar, I know what you coded last summer. Conference on Computer Supported Cooperative Work and Social Computing (CSCW), 2017, pp. 299–302.
- [27] J. Marlow, L. Dabbish, J. Herbsleb, Impression formation in online peer production : activity traces and personal profiles in GitHub. Conference on Computer Supported Cooperative Work (CSCW), 2013, pp. 117–128.
- [28] D. Ford, Z. Barik, L. Rand-Pickett, C. Parnin, The tech-talk balance: what technical interviewers expect from technical candidates. International Workshop on Cooperative and Human Aspects of Software Engineering (CHASE), 2017, pp. 43–48.
- [29] G.J. Greene, B. Fischer, CVExplorer: identifying candidate developers by mining and exploring their open source contributions. International Conference on Automated Software Engineering (ASE), 2016, pp. 804–809.
- [30] C. Hauff, G. Gousios, Matching GitHub developer profiles to job advertisements. Working Conference on Mining Software Repositories (MSR), 2015, pp. 362–366.
- [31] L. Bao, Z. Xing, X. Xia, D. Lo, S. Li, Who will leave the company?: a large-scale industry study of developer turnover by mining monthly work report. Working Conference on Mining Software Repositories (MSR), 2017, pp. 170–181.
- [32] Y. Goldberg, Neural Network Methods in Natural Language Processing (Synthesis Lectures on Human Language Technologies), eleventh ed., Morgan & Claypool Publishers, 2017.
- [33] S. Kim, J. E. James Whitehead, Y. Zhang, Classifying software changes: clean or buggy? IEEE Trans. Softw. Eng. 34 (2) (2008) 181–196.
- [34] J.S. Shirabad, T.C. Lethbridge, S. Matwin, Mining the maintenance history of a legacy software system. International Conference on Software Maintenance (ICSM), 2003, pp. 95–104.
- [35] S. Beyer, C. Macho, M. Pinzger, M. Di Penta, Automatically classifying posts into question categories on stack overflow. International Conference on Program Comprehension (ICPC), 2018, pp. 211–221.
- [36] F. Provost, T. Fawcett, Data Science for Business: What You Need to Know About Data Mining and Data-analytic Thinking, first ed., O'Reilly Media, Inc., 2013.
- [37] G. Tsoumakas, I. Katakis, I. Vlahavas, Random k-labelsets for multilabel classification, IEEE Trans. Knowl. Data Eng. 23 (7) (2011) 1079–1089.
- [38] O. Luaces, J. Díez, J. Barranquero, J.J. del Coz, A. Bahamonde, Binary relevance efficacy for multilabel classification, Prog. Artif. Intell. 1 (4) (2012) 303–313.
- [39] J. Read, B. Pfahringer, G. Holmes, E. Frank, Classifier chains for multi-label classification, Mach. Learn. 85 (3) (2011) 333–359.
- [40] G. Tsoumakas, I. Katakis, I. Vlahavas, Mining Multi-Label Data, Springer, pp. 667–685.
- [41] G. Madjarov, D. Kocev, D. Gjorgjevikj, S. Deroski, An extensive experimental comparison of methods for multi-label learning, Pattern Recognit. 45 (9) (2012) 3084–3104.
- [42] M.-L. Zhang, Z.-H. Zhou, A review on multi-label learning algorithms, IEEE Trans. Knowl. Data Eng. 26 (8) (2014) 1819–1837.
- [43] L. Breiman, Random forests, Mach. Learn. 45 (1) (2001) 5–32.
- [44] M.E. Maron, Automatic indexing: an experimental inquiry, J. ACM (JACM) 8 (3) (1961) 404–417.
- [45] Y. Tian, M. Nagappan, D. Lo, A.E. Hassan, What are the characteristics of high-rated apps? A case study on free android applications. International Conference on Software Maintenance and Evolution (ICSME), 2015, pp. 301–310.
- [46] T. Menzies, A. Butcher, D. Cok, A. Marcus, L. Layman, F. Shull, B. Turhan, T. Zimmermann, Local versus global lessons for defect prediction and effort estimation, IEEE Trans. Softw. Eng. 39 (6) (2013) 822–834.
- [47] F. Peters, T. Menzies, A. Marcus, Better cross company defect prediction. Working Conference on Mining Software Repositories (MSR), 2013, pp. 409–418.
- [48] F. Provost, T. Fawcett, Robust classification for imprecise environments, Mach. Learn. 42 (3) (2001) 203–231.
- [49] J. Coelho, M.T. Valente, L.L. Silva, E. Shihab, Identifying unmaintained projects in GitHub. International Symposium on Empirical Software Engineering and Measurement (ESEM), 2018, pp. 14–21.
- [50] V. Metsis, I. Androutsopoulos, G. Paliouras, Spam filtering with naive Bayes—which naive Bayes?. Conference on Email and Anti-Spam (CEAS), 2006, pp. 1–9.
- [51] A. McCallum, K. Nigam, A comparison of event models for naive Bayes text classification. International Conference on Machine Learning, 1998, pp. 41–48.
- [52] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: machine learning in Python, J. Mach. Learn. Res. 12 (2011) 2825–2830.
- [53] M.R. Boutell, J. Luo, X. Shen, C.M. Brown, Learning multi-label scene classification, Pattern Recognit. 37 (9) (2004) 1757–1771.
- [54] K. Dembczyński, W. Waegeman, W. Cheng, E. Hüllermeier, On label dependence and loss minimization in multi-label classification, Mach. Learn. 88 (1) (2012) 5–45.
- [55] G. McDowell, Cracking the Coding Interview: 189 Programming Questions and Solutions, CareerCup, LLC, 2015.
- [56] M. Behroozi, C. Parnin, T. Barik, Hiring is broken: what do developers say about technical interviews?. Visual Languages and Human Centric Computing (VLHCC), 2019, pp. 1–9.
- [57] S. Vajda, Y. Rangoni, H. Cecotti, Semi-automatic ground truth generation using unsupervised clustering and limited manual labeling: application to handwritten character recognition, Pattern Recognit. Lett. 58 (2015) 23–28.
- [58] O. Chapelle, B. Schölkopf, A. Zien, Semi-Supervised Learning, Cambridge, MIT Press, 2006.
- [59] F. Herrera, F. Charte, A.J. Rivera, M.J. del Jesus, Multilabel Classification: Problem Analysis, Metrics and Techniques, Springer, 2016.
- [60] T. Chen, C. Guestrin, XGBoost: a scalable tree boosting system. International Conference on Knowledge Discovery and Data Mining (KDD), 2016, pp. 785–794.
- [61] L. Yu, S. Ramaswamy, Mining CVS repositories to understand open-source project developer roles. International Workshop on Mining Software Repositories (MSR), 2007, pp. 1–10.
- [62] G. Avelino, L. Passos, A. Hora, M.T. Valente, Assessing code authorship: the case of the Linux kernel. 13th International Conference on Open Source Systems (OSS), 2017, pp. 151–163.
- [63] G. Avelino, L. Passos, A. Hora, M.T. Valente, Measuring and analyzing code authorship in 1+118 open source projects, Sci. Comput. Program. 1 (1) (2019) 1–34.
- [64] J. Anvik, L. Hiew, G.C. Murphy, Who should fix this bug?. International Conference on Software Engineering (ICSE), 2006, pp. 361–370.
- [65] O. Alonso, P.T. Devanbu, M. Gertz, Expertise identification and visualization from CVS. International Working Conference on Mining Software Repositories (MSR), 2008, pp. 125–128.
- [66] H. Kagdi, M. Hammad, J.I. Maletic, Who can help me with this source code change?. International Conference on Software Maintenance (ICSM), 2008, pp. 157–166.
- [67] B. Prabhakar, C.R. Litecky, K. Arnett, IT skills in a tough job market, Commun. ACM 48 (10) (2005) 91–94.
- [68] C. Litecky, A. Aken, A. Ahmad, H.J. Nelson, Mining for computing jobs, IEEE Softw. 27 (2009) 78–85.
- [69] C. Teyton, M. Palyart, J.-R. Falleri, F. Morandat, X. Blanc, Automatic extraction of developer expertise. International Conference on Evaluation and Assessment in Software Engineering (EASE), 2014, pp. 1–10.