



# JEMMA: An extensible Java dataset for ML4Code applications

Anjan Karmakar<sup>1</sup> · Miltiadis Allamanis<sup>2</sup> · Romain Robbes<sup>3</sup>

Accepted: 13 December 2022 / Published online: 10 March 2023  
© The Author(s) 2023

## Abstract

Machine Learning for Source Code (ML4Code) is an active research field in which extensive experimentation is needed to discover how to best use source code's richly structured information. With this in mind, we introduce JEMMA: An Extensible Java Dataset for ML4Code Applications, which is a large-scale, diverse, and high-quality dataset targeted at ML4Code. Our goal with JEMMA is to lower the barrier to entry in ML4Code by providing the building blocks to experiment with source code models and tasks. JEMMA comes with a considerable amount of pre-processed information such as metadata, representations (e.g., code tokens, ASTs, graphs), and several properties (e.g., metrics, static analysis results) for 50,000 Java projects from the 50K-C dataset, with over 1.2 million classes and over 8 million methods. JEMMA is also extensible allowing users to add new properties and representations to the dataset, and evaluate tasks on them. Thus, JEMMA becomes a workbench that researchers can use to experiment with novel representations and tasks operating on source code. To demonstrate the utility of the dataset, we also report results from two empirical studies on our data, ultimately showing that significant work lies ahead in the design of context-aware source code models that can reason over a broader network of source code entities in a software project—the very task that JEMMA is designed to help with.

**Keywords** Software engineering · Machine learning · Empirical datasets

---

Communicated by: Massimiliano Di Penta

✉ Anjan Karmakar  
akarmakar@unibz.it

Miltiadis Allamanis  
miltiadis.allamanis@microsoft.com

Romain Robbes  
rrobbes@unibz.it

<sup>1</sup> Free University of Bozen-Bolzano, Bolzano, Italy

<sup>2</sup> Microsoft Research, Cambridge, UK

<sup>3</sup> Université de Bordeaux CNRS, Bordeaux, France

## 1 Introduction

Software systems are complex networks of interacting entities. This makes them extremely challenging to develop, understand, and modify—despite the constant need to do so. In this context, appropriate tool support for source code can make developers faster and more productive. A variety of such tools have been proposed over the years, ranging from Integrated Development Environments (IDEs), testing tools, static analyzers, version control systems, and issue tracking systems, to name a few.

**Machine Learning for Source Code** In recent years, significant research effort has been undertaken towards developing tools based on machine learning models of source code (Allamanis et al. 2018) to handle several *tasks*.

A *task*, in the machine learning paradigm, is a type of action that a machine learning model is trained to perform. Code completion is a good example of a task—that a machine learning model can be trained to perform. There can be several other types of tasks, such as code summarization, defect prediction, classification and translation tasks, and many more.

This line of work was developed from the observation that simple statistical models of source code, such as n-gram models, were surprisingly effective for tasks such as code completion (Hindle et al. 2016). Since then such probabilistic models of source code have come a long way. In the present day, large-scale machine learning models of source code, based on the Transformer architecture, e.g. CuBERT (Kanade et al. 2020), PLBART (Ahmad et al. 2021), CodeBERT (Feng et al. 2020), and GraphCodeBERT (Guo et al. 2020) have achieved state-of-the-art performance on a number of Software Engineering (SE) tasks such as code generation, code search, code summarization, clone detection, code translation, and code refinement.

Largely by increasing the capacity of models and training datasets, deep learning based code completion has transitioned from the token level (Karampatsis et al. 2020) to completing entire snippets of code (Chen et al. 2021), the latter being now available on IDEs as an extension named *GitHub Copilot*<sup>1</sup>.

In parallel, other works on modeling of source code have observed that source code has a well-known structure compared to natural language. Source code can be unambiguously parsed into structured representations, such as Abstract Syntax Trees (ASTs); functions and methods can have control flows and data flows; functions and methods can interact with each other via calls, parameters and return values. Therefore, even though modeling source code as a series of tokens—analogueous to words in a sentence or a paragraph—has proven to be effective, another view shows that accounting for the structure of source code to be more effective.

A fair amount of research has addressed this issue in source code modeling, by proposing the incorporation of the inherent structural information of source code. Several works model source code as Abstract Syntax Trees (Mou et al. 2016; Alon et al. 2018; LeClair et al. 2019). Allamanis et al. were among the first to model source code snippets as graphs, including a wide variety of structural information, from data flow information, control flow information, lexical usage information, to call information (Allamanis et al. 2017).

The space of possibilities to model source code is vast, from text to tokens to advanced graphs—although each comes with its own issues and challenges. Thus, while being mindful of how we represent source code with as much information as possible, we also need

---

<sup>1</sup><https://copilot.github.com>

to make sure that the models trained on such representations are scalable and reliable for a number of source code tasks and corresponding applications.

**From Snippets to Projects** An important limitation of the current breed of deep learning models for source code is that the vast majority of the work has so far focused much more on single code snippets, methods, or functions, rather than on the intricate inter-relationships among source code elements, particularly when these relationships cross file boundaries.

Since source code is interconnected and interdependent, we argue that reasoning over a single method or function is fundamentally *inadequate* for several kinds of tasks. For instance, defect prediction tasks, e.g. predicting null pointer exceptions, resource leaks, may benefit from reasoning over associated code entities across the project. In fact, (Li et al. 2019) in their study construct a *global* context by connecting associated method entities based on Program Dependence Graph (PDG) and Data Flow Graph (DFG) to achieve state-of-the-art performance on bug prediction.

Even for tasks where the need for additional context may not be apparent, we note that most methods which have multiple calls to other callee methods are in fact dependent on supporting context—since the callee methods logically contribute to the overall functionality of the parent method.

Our views are supported by recent studies which show that encoding additional context while training machine learning models of code significantly improves model performance on a number of tasks. For instance, (Tian and Treude 2022) find that adding further context from the call hierarchy (i.e., caller and callee context) improves performance on the clone detection task by 8%. Li et al. (2021) include additional context from caller-callee methods and sibling methods in the same enclosing class, to train a model on the task of method-naming, and improve upon the state-of-the-art F-score by 11.9%. Liu et al. (2022) by encoding a broader context at the project-level, including comments, documentation, and nested scopes, improve on the method-naming task further. Lu et al. (2022) make use of additional code with lexical-similarity as external context to establish state-of-the-art performance on the CodeXGLUE benchmark (Lu et al. 2021) for code completion task.

In this paper, Section 5 provides further evidence that adding contextual information along with the input representations significantly improves the model performance on a method-call completion task, across four state-of-the-art transformer models: BERT, CodeBERTa, CodeBERT, and GraphCodeBERT.

The benefits of including a larger context while modeling source code is demonstrated in the studies mentioned above as well as our own. Therefore, from this current stage, we must gradually move towards building context-aware models that can reason over larger neighborhoods of interacting entities.

This is not only a paradigm shift but also a clear indication of the potential need for large-scale code datasets from which additional contexts can be constructed and used in training robust and context-aware source code models.

The major reasons for the lack of such work is that the necessary data is not yet collected, organized, and is missing at scale, or they support just a single task. The following section highlights the absence of such datasets for code that have the right mix of source code granularity, size, scale, and detail of information to allow researchers to research on models that go beyond single code snippets.

Datasets that are large focus either on individual code snippets at the method-level, or at best, source files; while other datasets are either too small, or lack significant preprocessing. Choosing good quality data in sufficient quantity, downloading and storing the data, extracting valuable information from the data or simply running tools to preprocess the data

and gather additional information, and then building an experimental infrastructure in place, requires a large amount of time and effort—even before a single experiment is run. This is all the more true when this has to be done for source code models, where some of the pre-processing and analysis tools can be extremely time-consuming and resource-intensive at scale. Therefore, in this paper, we contribute such a dataset: JEMMA.

**JEMMA as a Dataset** JEMMA has multiple levels of granularity: from methods, to classes, to packages, and entire projects. It consists of over 8 million Java method snippets along with substantial metadata; pre-processed source code representations—including graph representations that comes with control- and data-flow information; call-graph information for all methods at the project-level; and a variety of additional properties and metrics.

We make sure that all the processed data are clean, consistent, as well as comprehensive—using common data validation, filtering, deduplication, and data curation techniques. Corrupted, incomplete, and blank/null values were corrected where possible; valid results were accurately and consistently mapped to source code entities based on data curation principles; some outputs were filtered at source based on an expected range, on expected datatypes, and/or on formatting conventions; while deduplication eliminated redundant or corrupted entries. Furthermore, the availability of supplementary data down to AST-node level, resulting from our extensive processing, ensured comprehensiveness at scale, for millions of source code entities defined within JEMMA. All of which contribute to the overall quality of the data presented.

JEMMA is built upon the 50K-C dataset of compilable Java projects (Martins et al. 2018), and complements it with significant processing, measured in years of compute time. Section 3 presents all the components of the JEMMA Dataset.

**JEMMA as a Workbench** JEMMA is not a static dataset: we purposefully designed it to be extensible in a variety of ways. Concretely, JEMMA comes with a set of tools to: add metrics or labels to source code snippets (e.g., by utilizing static analysis tools); define prediction tasks based on metrics, properties, or the representation themselves; process the code snippets and existing representations to generate new representations of source code; and run supported models on a task. We describe how to extend the dataset, along with several examples in Section 4. This extensibility is critical, because it transforms JEMMA into a workbench with which users can experiment with the design of ML models of code and tasks, while saving a lot of time in pre-processing the data.

Traditionally, a database workbench is described as a tool that can be used to view, create, and edit tables, indexes, stored procedures, and other database metadata objects. Thereby, if we now extend the same concept to working with our collection of data, the *JEMMA Workbench* is a set of tools and helper functions that helps in several operations such as viewing, creating, retrieving from, and appending to datasets (independent of how they are stored), among many other tasks that do not involve working directly with the datasets.

**Empirical Studies with JEMMA** In Sections 5 and 6, we show how JEMMA can be used to gain insights via empirical studies. The first is a study on the *non-localness* of software, and how it impacts the performance of models on a variant of the code completion task. This study shows how the data from JEMMA can be used to gain insights into how the models perform on code samples, highlighting what performance issues exist and what we can do to address such issues by adding project-wide context (Section 5).

The second is the study of the size of entities that constitute software projects, and how it relates to the context size of popular machine-learning (ML) models. The second study

confirms that significant work lies ahead in designing models that efficiently encode large contexts (Section 6).

While these examples are related to empirical analyses in the sub-field of Machine Learning for Software Engineering, we can envision further uses for JEMMA in empirical studies. For example, empirical studies on fault-prone or misleading method names, or the impact of complexity on other code properties, or the challenges of coupling in large projects, and others.

Finally, in Section 7 we document the limitations of JEMMA, and then conclude with a summary of our work in Section 8.

## 2 Related Work

With the gradual evolution of machine learning techniques suitable for processing code—where data plays a central role—a multitude of efforts have been made for collecting and organizing quality data. Such datasets have not only contributed to the development of competent models of source code, but also opened the avenues for empirical analysis of these models. In this section, we outline some of the datasets from both genres.

### 2.1 Datasets for Machine Learning on Code

Since machine learning requires considerable amounts of data, multiple datasets have been produced, usually as a means of validating a specific machine learning method, rather than as a principled standalone effort. This has resulted in datasets that contain input data either not far from raw text, or that contain a lossy view of the underlying analyzed software systems.

**Code Datasets** Allamanis and Sutton (2013) collected a set of over 1 billion Java code tokens and provide the code text per file for training n-gram models. Later, Karampatsis et al. (2020) extended this with additional datasets for C and Python; and a different extension of the dataset was provided by Alon et al. (2018). Raychev et al. (2016) released Py150 and JS150, two datasets of 150,000 Python and Javascript functions parsed into ASTs. Unfortunately, these datasets are limited to small programs or code snippets only at the method level. In comparison, JEMMA provides code entities in multiple granularities across several representation types—creating a wide range of modelling opportunities.

Several datasets focus on specific tasks, such as the BigCloneBench Svajlenko and Roy (2015) dataset for large-scale clone detection in Java. ManyTypes4Py Mir et al. (2021) is a Python dataset aimed at evaluating type inference in Python, and Devign (Zhou et al. 2019), provides labeled code with coarse-grained source code vulnerability detection in mind. JEMMA, on the other hand, is not task-specific and supports multiple tasks out of the box.

Datasets with specific representations of code have been common. CoCoGum (Wang et al. 2020) use class context represented as abstracted UML diagrams, for code summarization, at the file-level. Allamanis et al. (2017) and Allamanis et al. (2020) extract control, data flow graphs, along with syntax within a single file. Representation-specific datasets are useful but they limit cross-representational and cross-architectural analyses for tasks. JEMMA supports several representations including raw source code, tokens, ASTs, and graphs, at both method-level and class-level, building up to even coarser granularities.

Datasets of code from student assignments, programming competitions, and other smaller programs, have also been created. Among them, Google Code Jam<sup>2</sup> and POJ-104 (Mou et al. 2016) are clone detection tasks (clone detection in this case is formulated as a program classification task). COSET (Wang and Christodorescu 2019), and CodeNet (Puri et al. 2021) also feature smaller programs, but complement them with additional metrics and labels. Although these datasets have many desirable properties, they do not represent source code used in real-life software systems and thus it is unclear if learning on these datasets can generalize to general-purpose software. Semi-synthetic datasets, such as NAPS (Zavershynskiy et al. 2018), also fall into the same category. JEMMA balances the preceding concerns by building upon organic projects coming from a diverse set of domains (e.g., games, websites, standalone applications, etc) and of development standards (ranging from student projects to industry-grade open-source projects) which add a healthy factor of generalization for source code modeling.

**Code Datasets with Natural Language** Natural language presents an interesting, yet separate, modality from source code and is central to the NLP task of semantic parsing (i.e., text-to-code). A few datasets have focused on this: CodeNN (Iyer et al. 2016), CoNaLa (Yin et al. 2018), and StaQC (Yao et al. 2018). Datasets such as NL2Bash (Lin et al. 2018) provide data for semantic parsing from natural language to Bash commands, while Spider (Yu et al. 2018) is a dataset for the text-to-SQL semantic parsing task. Finally, (Barone and Sennrich 2017), CodeSearchNet (Husain et al. 2019), and (LeClair and McMillan 2019) pair natural language documentation comments with code, targeting code search and code summarization applications.

All these datasets provide dumps of source code snippets per file, and while it is possible to parse the code text and perform some intra-procedural analyses for the few file-level datasets, information about external dependencies is commonly lost rendering it impossible to extract accurate semantic data. JEMMA successfully mitigates such issues by providing a dataset of inter-related code entities across granularities, along with comprehensive intra- and inter-procedural relationship information coming from data-flow, control-flow, call graphs, etc.

**Code Datasets with Higher-Level Representations** While the above datasets focus on code snippets or files, some work has extracted datasets aiming for representations that capture information beyond a single file. However, commonly these datasets opt for an application-specific representation that loses information that could be useful for other research. For example, (DeFreez et al. 2018) extract path embeddings over functions in Linux. LambdaNet (Wei et al. 2020) extracts type dependency graphs in TypeScript code but removes most code text information; their dataset is also limited to 300 projects, which range from 500 to 10,000 lines of code. The dataset by Bansal et al. was also refined and used in a source code summarization approach that defined a *project-level* encoder, that considers functions in up to 10 source code files in a project (Bansal et al. 2021). The scale and breadth of information present in JEMMA keeps necessary code information intact, be it across procedures or within procedures, and across files at the project-level.

**Code Datasets with Changes** Given the importance of software maintenance in the development lifecycle, a few datasets have focused on edit operations in code. The goal of these

---

<sup>2</sup><https://code.google.com/codejam/contests.html>

datasets is to foster research in Neural Program Repair. ManySSubs4J (Karampatsis and Sutton 2020) and Bugs2Fix (Tufano et al. 2019) both fall in this category: they are corpora of *small* bug fixes extracted from GitHub commits. These datasets often focus on local changes (e.g., diffs) and ignore the broader context.

Although our dataset does not come with changes, it provides increased modeling opportunities for users as it comes with inter-procedural relationship information among code entities for all 50K projects. Neural Program Repair workflows could benefit at the dataset creation stage by leveraging bug-related properties in our dataset.

## 2.2 Datasets for Empirical Studies

Several corpora of complete software systems have been built with the primary goal to conduct traditional empirical studies, without direct considerations necessary for machine learning research.

**The Qualitas Corpus and its Descendants** The Qualitas corpus (Tempero E et al. 2010), is an influential corpus of 111 large-scale Java systems that was used for a large number of empirical studies of Java and the characteristics of the systems implemented in it. While this dataset was source code only, it was post-processed in various ways, producing several derived datasets. The Qualitas.class corpus, (Terra et al. 2013) is a version of the Qualitas corpus that is also compiled in `.class` files. The QUAATLAS corpus (De Roover et al. 2013), is a post-processed version of the Qualitas corpus that allows better support for API usage analysis. XCorpus (Dietrich et al. 2017), is a subset of the Qualitas corpus (70 programs) complemented by 6 additional programs, that can all be automatically executed via test cases (natural, or generated).

**Java Datasets** (Lämmel et al. 2011) gathered a dataset of Java software from Sourceforge, that had 1,000 projects that were parsed into ASTs. The BOA dataset and infrastructure, by (Dyer et al. 2013), provides an API for pre-processing software repositories, such as providing and analyzing ASTs, for 32,000 Java projects. The 50K-C dataset of (Martins et al. 2018) contains 50,000 Java projects that were selected because they could be automatically compiled. A follow-up effort is the Normalized Java Resource (Palsberg and Lopes 2018) (NJR). A first release, NJR-1, provides 293 programs on which 12 different static analyzers can run (Utture et al. 2020), but has a stated goal of gathering 100,000 *runnable* Java projects, but is still a work in progress.

**Other Datasets** (Spinellis 2017) released a dataset that contains the entire history of Unix as a single Git repository. (Geiger et al. 2018) present a graph-based dataset of commit history of real-world android apps. The entire Maven software ecosystem was released as a dataset with higher-level metrics, such as changes and dependencies (Raemaekers et al. 2013). The Maven Dependency Graph by (Benelallam et al. 2019) provides a snapshot of the Maven Central as a graph, modeling all its dependencies. Fine-GRAPe is a dataset of fine-grained API usage across the Maven software ecosystem (Sawant and Bacchelli 2017). Finally, both Software Heritage (Pietri et al. 2019) and World of Code (Ma et al. 2021) are very large-scale efforts that aim to gather the entirety of open-source software as complete and up-to-date datasets. The main goal of World of Code is to enable analytics, while the main goal of Software Heritage is preservation (although it also supports analytics). The Perceval tool (Gonzalez-Barahona et al. 2018) also promises automatic and incremental data gathering from almost any tool related to contributing to open source development, among

other sources. We find that although there are similarities between JEMMA and Perceval, the differences highlight that both tools can complement each other well. Perceval can be used to fetch raw project data from a wide variety of data sources. On the other hand JEMMA focuses on source code, and can be used to take care of the analysis of data, pre-processing of data, task definition, and training of models out of the box.

### 2.3 The 50K-C Dataset

Having surveyed the landscape of existing datasets, we conclude that most machine learning datasets focus on small-scale entities such as functions, methods, or single classes. The ones that offer higher-level representations are specific and too small in scale. The corpora of systems used for empirical studies provide a better starting point, as they can be pre-processed to extract additional information. Of the existing datasets, the most suitable option that is large enough and that allows the most pre-processing is the 50K-C dataset of 50,000 compilable projects.

Since JEMMA builds upon 50K-C, we provide detailed background information on it in this section. The 50K-C dataset is a collection of 50,000 compilable Java projects, with a total of almost 1.2m Java class files, its compiled bytecode, dependent jar files, and build scripts. It is divided into three subsets:

- `projects`: It contains the 50,000 java projects, as zipped files. The projects are organized into 115 subfolders each with about 435 projects.
- `jars`: It contains the 5,362 external jar dependencies, which are required for successful project builds. This is important as missing dependencies is the common cause of failing to compile code at scale.
- `build_results`: It contains the build outputs for the 50,000 projects, including compiled bytecode, build metadata, and original build scripts. In addition to the above data, a mapping between each project and its GitHub URL is also provided. The bytecode is readily available for a variety of tasks, such as running static analysis tools, or, if the projects can also be executed, as input for testing, and dynamic analysis tools.

Beyond the size of the dataset, the fact that the projects are compilable is the main reason we chose to build upon 50K-C. The extensive pre-processing that we perform on top of 50K-C requires the use of static analysis tools, to do things such as call graph extraction, and to extract valuable metrics about the systems. Since the vast majority of static analysis tools operate on bytecode, 50K-C was the most suitable option that combines both scale and the ability to automate the analysis at such scale.

**Selection Criteria** The dataset authors downloaded close to 500k Java projects, attempted to compile all of them, and selected 50k projects among the ones that could be compiled. Two filters were applied: projects that were Android applications were excluded, and projects that were clones were also excluded—using the DéjàVu clone repository (Lopes et al. 2017), and the Sourcerer CC tool (Sajjani et al. 2016). We find that the projects have a diverse set of domains (e.g., games, websites, standalone apps, etc), and development levels (ranging from student projects to industry-grade open-source projects).

The dataset consists of both large-scale projects with as many as 5k classes, and smaller projects with as low as 5 classes. While the larger projects are good representatives of real-world projects, the smaller projects are valuable too, since machine learning models of code still need to make significant headway in code understanding which necessitates reasoning on projects across all sizes.



### 3 The JEMMA Dataset

Our goal with the JEMMA project is to provide the research community with a large, comprehensive, and extensible dataset for Java that can be used to advance the field of source code modeling. The JEMMA datasets consist of a large collection of code samples in varying granularities, with wide-ranging and diverse metadata, a range of supported source code representations, and several properties. In addition, it also includes source code information related to code structure, data-flow, control-flow, caller-callee relationships etc.

For every project in the JEMMA Dataset, we gather data at the project-level, and provide information on all the packages and classes. Furthermore, for every class, we parse and provide data on all the methods—including respective metadata, several representations and properties. The detail of data provided for every method entity is comprehensive, with data at the level of AST with data-flow, control-flow, lexical-usage, and call-graph edges among others. In addition to necessary information, such as line numbers and position numbers of code tokens, supplementary information such as token types, node types, etc, are also provided. More details are presented in the following sections.

JEMMA also comes equipped with *Workbench* tools that allow users to perform a variety of tasks out of the box, such as: transforming code samples into intermediate source code representations, making tailored selections of entities to define tasks and forming custom datasets, or to run supported models, among others (Section 4 provides more details).

**Statistics** The original 50K-C dataset contains a total of 50,000 projects. It has 85 projects with over 500 classes (with a maximum of 5549 classes in a project), 1264 projects with 101–500 classes, 2751 projects with 51–100 classes, 10693 projects with 21–50 classes, 14322 projects with 11–20 classes, and 20885 projects with 10 or fewer classes (with a minimum of 5 classes per project). We have collected metadata for all of these projects. Overall, the data consists of 1.2 million Java classes, which define over 8 million unique method entities.

**Granularity** JEMMA supports multiple granularities. We have processed and catalogued data starting from the project-level descending to smaller entities, which means a spectrum of granularities of code can be accessed.

However, since a method is the most basic unit of behavior in an object-oriented program, and it is a self-contained program segment that carries out some specific, well-defined logical task, we collect all the properties at the (*finer*) method level. For instance, method entities can be sampled from the datasets and used independently to run code-analysis tools. In this sense, methods can be considered as our primary entity-of-interest.

Starting from the method-level, larger contextual entities at the class-, package-, and project-level can then be created by building upon the smaller entities. And since most prevalent models of source code accept input samples at the method-level, we do provide all source code properties and representations at this (*finer*) method-level by default; properties and representations of larger entities-of-interest can always be built upon smaller entities within it.

**Compilability** Successful compilation ensures that the source code snippets from the projects have been type-checked and parsed successfully, and are valid Java code. Having full-scale compilable projects gives us the assurance that the source code is complete and self-contained; and thus, all the inter-relationships among code entities can be captured and

studied. Additionally, static and dynamic analysis tools can be run to generate information for new code tasks.

Some of the tools that we use to post-process the data require the ability to compile the code, rather than just analyzing compiled code. These tools insert themselves in the compilation process (for instance, Infer (Calcagno et al. 2015)). Therefore, we also have to be able to compile the code on demand. Practically, we found that recompilation was not 100% replicable. Of the 50,000 projects, we were able to compile about 92% of the projects; a failed compilation is usually linked to a missing or inaccessible dependency. Nevertheless, 100% the project entities were successfully processed and catalogued, along with their corresponding properties and representations.

**Runtime Considerations** The analyses and post-processing that we apply to the projects is very computationally intensive for some of the tools. For instance, the analyses run just by a single tool—the Infer static analyzer (Calcagno et al. 2015)—can take on the order of half an hour for a single medium-sized project. Analyzing 50,000 projects with a number of tools and then post-processing the outputs is thus both time-consuming and resource-intensive.

Since projects can vary significantly in size, depending on whether it's a small project or a large one the processing times can be very different. For smaller projects, with 1-20 classes, on average it takes 20-30 minutes of processing time overall. For medium-size projects, with 21-50 classes, it takes over an hour. For large projects, with 51-100 classes, it takes a few hours. For the rest of the projects with more than 100 classes, it can take just 4-5 hours or as much as a couple of days depending on the size of the project.

We have gone ahead and done most of the necessary pre-/post- processing and only a small portion (about 6-9% of the data at the time of writing) is still processing and it will be made available shortly.

**Storage and Sharing Considerations** The amount of data produced is considerable. To maximize accessibility, we provide it as a set of Comma Separated Values (CSV) files, so that users can choose and download the data that they need. Note that only the metadata and the original source data are absolutely necessary; other data can be downloaded on a per-need basis. The JEMMA *Workbench* allows the recomputation of the other properties, if, for some properties, it is more efficient to recompute them than to download them. The data is uploaded on Zenodo; due to its size, it is provided as multiple artifacts. We present the components of the dataset, along with their DOIs (links to the download page), and sizes later on.

**Interacting with the Data** Most of the data from the JEMMA are organized in Comma-Separated-Values (CSV) files, consequently basic analyses can be run with tools such as *csvstat*. Furthermore, our *Workbench* APIs can be used to gather extensive statistics of the projects, classes, methods, bytecode, and data and control-flow information.

The JEMMA datasets are grouped into three major parts: data at the metadata level (Section 3.1), data at the property level (Section 3.2), and at the representation level (Section 3.3). In addition, we also provide project-wide callgraph information for the 50,000 projects, uniquely identifying and associating source and destination nodes in the callgraph with the help of the metadata defined by JEMMA (Section 3.4). This allows for accessing project-wide data on the whole, for different granularities of code entities.

Figure 1 gives a glimpse of the extent and detail of the data contribution made by JEMMA. The top-left corner represents the raw data from 50K-C, which we catalog by adding UUIDs (symbolized by colored squares). The rest of the figures depicts the additional pre-/post-processing we performed: the colored gears represent external tools that we

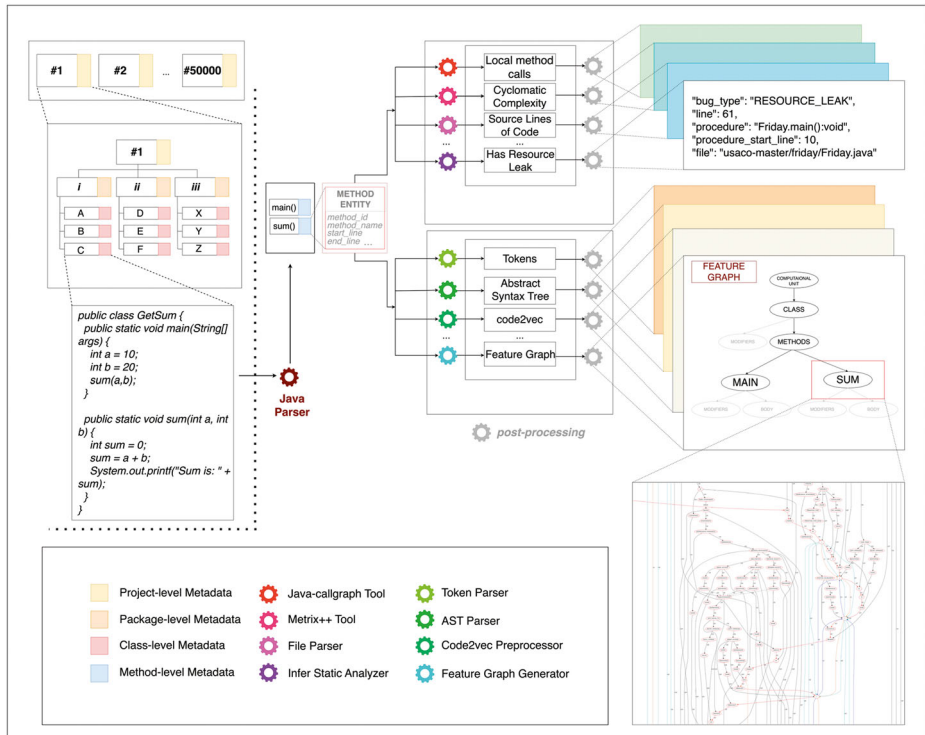


Fig. 1 Overview of data-level contributions

run to collect additional data (properties and representations), while the grey gears represent further post-processing that we perform on the tool outputs to integrate it in our dataset.

### 3.1 JEMMA: Metadata

In this section we present the metadata for the JEMMA datasets. The metadata is made available in CSV (comma-separated values) files. This allows for easy processing, even with simple command-line tools. The metadata is organized in four parts, from the largest units of interest to the smallest: projects, packages, classes, and methods. The units of interest can then be inter-related systematically. The metadata serves two major purposes:

1. Uniquely identify a source code entity with an UUID.
2. Gather basic and often-used information on each source code entity.

Taken together, these two purposes allow users to extend the dataset with additional properties. The UUID allows us to uniquely identify an entity in the dataset, and the supplementary metadata helps disambiguate entities (file paths, parent relationships, location information in the file, etc). In Section 4 we show how this metadata can be used to add an additional property to source code entities.

Since the metadata formalizes the organization of the data, and establishes the relationship between projects, packages, classes, and methods, JEMMA users can leverage it

to construct custom data queries and make selections from the large collection of data at different granularities.

**Projects** For the project-level metadata, we provide a single CSV file that lists all the projects in the 50K-C dataset along with their corresponding metadata—*project\_id*, *project\_path*, *project\_name*. The UUID is referenced by the entities contained in the project. The project path is relative to the root directory of the 50K-C dataset<sup>3</sup>, and can be used to access the raw source code of the project.

**Packages** For the package-level metadata, a single CSV file lists all the packages present in the projects. The metadata comprises of the UUID of the parent project as *project\_id*, the UUID assigned to the package as *package\_id*, the relative path of the package as *package\_path*, and the name of the package directory as *package\_name*.

**Classes** For the class-level metadata, we provide a single CSV file that lists all the classes in the 50K-C dataset along with their corresponding metadata: *project\_id*, *package\_id*, *class\_id*, *class\_path*, *class\_name*. Similarly to the projects, the class path is a relative path starting from the 50K-C dataset's root directory, that allows to access the raw source code of the class.

**Methods** At the method-level, the metadata is more extensive. Just having the name of a method might not be enough to disambiguate methods. Thus, the metadata is a CSV file lists all the methods in the 50K-C dataset along with their corresponding metadata: *project\_id*, *package\_id*, *class\_id*, *method\_id*, *method\_path*, *method\_name*, *start\_line*, *end\_line*, *method\_signature*.

### 3.2 JEMMA: Properties

JEMMA leverages the UUIDs assigned to projects, classes, and methods as a way to attach additional properties to these entities. Thus, a property can be an arbitrary value that is associated to an entity, such as a metric. Even though we have gathered several properties associated with code entities, it should be noted that a particular property may not be available or may not apply for a given code entity. Users can add new properties associated with code entities as contributions to the dataset, where the property should be given a unique name and be stored in the correct location for it to be visible to JEMMA *Workbench* APIs. (Section 4 provides more details).

Next, we list the tools used to obtain the properties:

- The *Infer* static analyser (Calcagno et al. 2015) is a tool that provides advanced abstract interpretation-based analyses for several languages, including Java. Examples of the analyses that Infer can run include an interprocedural analysis to detect possible null pointer dereferences. Infer can also perform additional analyses such as taint analysis, resource leak detection, and estimate the run-time cost of methods. We chose Infer mainly because it can perform inter-procedural analysis that reasons across procedure boundaries, while being able to scale to large codebases.

---

<sup>3</sup>[http://mondego.ics.uci.edu/projects/jbf/downloads/50K-C\\_projects.tgz](http://mondego.ics.uci.edu/projects/jbf/downloads/50K-C_projects.tgz)

**Table 1** List of tools used and properties obtained

Property	Tool used
[TLOC] Total Lines of Code	Metrix++
[SLOC] Source Lines of Code	Metrix++
[CMFX] McCabe or Cyclomatic Complexity	Metrix++
[MXIN] Maximum indent depth of nesting	Metrix++
[NPTH] Npath Complexity	PMD
[NMTK] Number of Code Tokens	Parser
[NMFR] Number of parameters	Parser
[NUID] Number of unique identifiers	Parser
[NMOP] Number of operators	Parser
[NMLT] Number of literals	Parser
[NMRT] Number of return statements	Parser
[NAME] Name of source code entity	Parser
[NUPC] Number of unique parent callers	java-callgraph
[NUCC] Number of unique child callees	java-callgraph
[NMNC] Number of non-local calls	java-callgraph
[NMLC] Number of local calls	java-callgraph
[NLDF] Presence of Null Dereference	Infer
[RSLK] Presence of Resource Leak	Infer

- *Metrix++* is a tool that can compute a variety of basic metrics on source code entities, such as lines of code, code complexity, and several others<sup>4</sup>. We chose *Metrix++* since it is suitable for processing large codebases, processing thousands of files per minute; it recognizes various types of entities including classes, interfaces, namespaces, functions, comments; and supports multiple metrics.
- *PMD* is a static code analysis tool<sup>5</sup> that can compute a variety of static analysis warnings and metrics, such as the *npath complexity* metric, among many, many others. We used *PMD* because it is inexpensive while reviewing large codebases; and it is trusted by industry practitioners and researchers. *PMD* can also be used to identify defects and problems in code entities which can be useful for future works.
- The *java-callgraph*<sup>6</sup> extractor is a tool for generating call graphs in Java. We used this tool to extract project-wide call graphs, from which callers and callees were identified and linked to their respective *UUIDs* at the post-processing stage. The *java-callgraph* generator tool was used since it was capable of generating both static and dynamic call-graphs suitable for our dataset of compilable code entities.

Table 1 provides the list of properties that are currently defined at the finer method-level granularity in *JEMMA*. It also maps the tools used to obtain the properties. Later, we provide a table links to the datasets for all of the data.

Other tools that could be run to extend the dataset include static analysis tools, such as *FindBugs* (Hovemeyer and Pugh 2004), *SpotBugs*, or similar tools such as *Error Prone* and

<sup>4</sup><https://metrixplusplus.github.io/metrixplusplus/>

<sup>5</sup><https://pmd.github.io>

<sup>6</sup><https://github.com/gousiosg/java-callgraph>

NullAway. The warnings and outputs from these tools can serve as metrics for code entities. Bespoke static analyses from Soot or other static analysis research frameworks, or clone detection (Cordy JR and Roy CK 2011) tools could be run as well. These properties could be useful to conduct studies similar to the ones from (Habib and Pradel 2018).

### 3.3 JEMMA: Representations

Machine learning models are trained on a collection of feature vectors derived from the input data. For source code machine learning models the input data can be the raw text of a source code entity. For example, for the Java method shown in Fig. 2a, a corresponding source code representation could be its raw tokens as shown in Fig. 2b.

```
public static void init(RobotController rci,int bigBoxSize){
    rc = rci;
    width = rc.getMapWidth()/bigBoxSize;
    height = rc.getMapHeight()/bigBoxSize;
    MapAssessment.assessMap(bigBoxSize, rci);
    updateInternalMap(rc);
}
```

(a) Java method as original text representation

TOKENS (TKNA)

```
public static void init ( RobotController rci ,
int bigBoxSize ) { rc = rci ; width = rc . getMapWidth
( ) / bigBoxSize ; height = rc . getMapHeight ( ) /
bigBoxSize ; MapAssessment . assessMap ( bigBoxSize ,
rci ) ; updateInternalMap ( rc ) ; }
```

TOKENS (TKNB)

```
PUBLIC,STATIC,VOID,init,LPAREN,RobotController,rci,COMMA,
INT,bigBoxSize,RPAREN,LBRACE,rc,EQ,rci,SEMI,width,EQ,rc,
DOT,getMapWidth,LPAREN,RPAREN,SLASH,bigBoxSize,SEMI,height,
EQ,rc,DOT,getMapHeight,LPAREN,RPAREN,SLASH,bigBoxSize,SEMI,
MapAssessment,DOT,assessMap,LPAREN,bigBoxSize,COMMA,rci,
RPAREN,SEMI,updateInternalMap,LPAREN,rc,RPAREN,SEMI,RBRACE
```

(b) Java method represented as tokens

**Fig. 2** An example of a Java method and two of its possible token representations

Since source code is highly structured, the design space for representations is vast and diverse. This has been explored to some extent, with approaches that model source code as sequences of tokens or subtokens via RNNs (Pradel and Sen 2018), LSTMs (Karampatsis et al. 2020), or Transformers (Feng et al. 2020). Other approaches have leveraged the structure of code either via ASTs (Mou et al. 2016) or linearized ASTs (Alon et al. 2019). Yet other approaches use more expressive structures incorporating, for instance, data flow, and use Graph Neural Networks (GNNs) to represent code (Allamanis et al. 2017).

Our goal with JEMMA is to provide the building blocks to experiment with the design space of representations. Since extracting the relevant information is costly in terms of computational resources, a significant effort went into adding several basic representations at the method level, ranging from the raw source code to the information behind a very complete graph representation. At the representation level, we provide several ready-to-use source code representations (compatible with different models) for over 8 million method snippets. The method level representations that we provide are described in the following subsections.

### 3.3.1 Raw text (TEXT)

First and foremost, the original code for each method is provided by default, with no preprocessing. This allows approaches that need to operate on the raw text to do so directly (e.g., a model that implements its own tokenization). The default method text includes its annotations and also comments within the code snippet, if any. The whitespace for each method text is also preserved intentionally (it can be easily stripped off at any point). The raw text can also be used to re-generate the other representations as needed.

### 3.3.2 Tokens (TKNA, TKNB)

Each Java method snippet is tokenized with an ANTLR4 grammar (Parr 2013) and made available to the user preprocessed. The tokenized code includes method annotations, if any, but does not include natural language comments. However, with the entire raw text of method snippets made available by default, users are free to include comments in their custom tokenizations.

For every method snippet in our dataset, we provide the corresponding string of tokens. In fact, we provide two types of tokenized strings. First, a simple space-separated string of tokens. This representation is meant to be directly passed to an existing ML model that has its own tokenizer, without any further processing. The downside is that some literals that include spaces may be treated as more than one token, or symbols and special characters may be ignored while using certain tokenizers (e.g., natural-language tokenizers). Should this be an issue, users may use the second representation.

In the second type of tokenized representation the tokens are made available as a comma-separated string with the symbols and operators replaced with suitable string tokens (commas in literals are replaced suitably with <LITCOMMA> tokens). This representation is recommended for users who would tokenize the code themselves, or would want to avoid literals being split into several tokens, or avoid ambiguities with symbols and special characters when using natural-language tokenizers.

### 3.3.3 Abstract Syntax Tree (ASTS)

An Abstract Syntax Tree, or AST, is a tree representation of the source code of a computer program that conveys the structure of the source code.

In an AST, nodes can either be terminal nodes (leaf nodes), which are the tokens of the grammar, or non-terminal nodes (internal nodes), representing the higher-level units such as method calls, code blocks, etc. This information is represented for each method as a set of nodes, followed by a set of node pairs representing child edges.

### 3.3.4 code2vec (C2VC) and code2seq (C2SQ)

The code2vec (Alon et al. 2019) and code2seq (Alon et al. 2018) representations are derivatives of the AST representation. The goal of these approaches is to linearize ASTs by sampling a fixed number of AST paths (i.e., selecting 2 AST nodes at random and utilizing the connected path between them).

The difference between the approaches is that code2vec represents each identifier and each path as unique symbols leading to large vocabularies, and consequently Out-Of-Vocabulary (OOV) issues, while code2seq models identifiers and paths as sequences of symbols from smaller vocabularies, which alleviates the same issues. However, the downside is that the code2seq representation is significantly larger. Both kinds of inputs are fed to models that use the attention mechanism to select a set of AST paths that are relevant to the model's training objective (by default, method naming).

We have generated the code2vec and code2seq representations of every method in the dataset by running the pre-processing scripts, which can serve as a ready-to-use input to the code2vec and code2seq path-attention models. Furthermore, if mutations to the code snippets are necessary, our *Workbench* tools enable users to easily transform raw code snippets into the corresponding representations using the original code2vec and code2seq preprocessors.

### 3.3.5 Feature Graph (FTGR)

A Feature Graph is a feature-rich graph representation of a source code entity. It is built on top of the abstract syntax tree, but containing multiple edge types to model syntactic, semantic, lexical, data-flow, and control-flow relationships between nodes. (Allamanis et al. 2017)

The Feature Graph representation is comprised of a set of nodes, and then node pairs representing different edge types; the nodes are also presented in a sequence to capture the order of tokens. Specific edge types can be filtered as needed (such as to produce the AST representations, or to reduce the size of the graph (Hellendoorn et al. 2019b)). The full list of included edges are:

- Child edges encoding the AST.
- NextToken edges, encoding the sequential information of code tokens.
- LastRead, LastWrite, and ComputedFrom edges that link variables together, and provide data flow information.
- LastLexicalUse edges link lexical usage of variables (independent of data flow).
- GuardedBy and GuardedByNegation edges connecting a variable used in a block to conditions of a control flow.
- ReturnTo edges link return tokens to the method declaration.
- FormalArgName edges connect arguments in method calls to the formal parameters.

This representation is significantly feature-rich, as it includes, for instance, all the source code tokens and their precise locations in the original source code, the signatures of all the



methods called in the class, the source code comments, if any, including a variety of data-flow, control-flow, lexical-usage, and hierarchical edges. Derivative representations such as AST with dataflow information, a subset of the feature graph representation, which corresponds to the data used by models such as GraphCodeBERT can also be produced from the feature graph representations. The feature graph representation is obtained from Andrew Rice's feature graph extraction tool<sup>7</sup>.

### 3.4 JEMMA: Callgraphs

Since many relationships among source code entities are not simply hierarchical containment relationships, we also provide a very useful additional data: the project's call graph (CG), in which methods calling each other are explicitly linked. Thanks to our meta-data, these method call information can then be used to combine representations to create interesting global contexts for large-scale source code models.

Previous techniques are useful to design representations at the level of methods. However, designing models that reason over larger entities requires more data. Hierarchical relationships can be already inferred from the meta-data. In addition, since software systems are composed of modules that interact with each other, caller-callee relationships are crucial to model systems accurately. For this, we use a Java callgraph extractor tool, to extract project-wide call graphs, from which callers and callees are identified and linked to their respective UUIDs through post-processing (links to external calls are still recorded but we do not assign UUIDs to them).

Method signatures are used to disambiguate methods with similar names. Note that for polymorphic calls, the call graph provides links to the statically identified type, not to all possible types. Additional post-processing would be possible to add these links to the call graph. In previous work, we have seen that the use of polymorphism in Java is significant (Milojkovic et al. 2015), so this would be a useful addition.

## 4 Extending and Using JEMMA

Table 2 presents the links to the actual datasets with JEMMA: meta-data, properties, representations, and callgraphs. These are standalone CSV files that can be used on their own, but to make it easy for users to access and use them in common usage scenarios we have added a *Workbench* component to JEMMA.

When building JEMMA, we intended it to be large-scale, yet extensible, flexible, and most importantly, easy to use. We have implemented several tools to help with this, and as a result, researchers can readily use JEMMA as a *Workbench* to experiment with variants of datasets, models, and tasks while minimizing the processing that is involved.

The JEMMA *Workbench* tools and implementations, written in Python, are accessible through a set of APIs, which helps developers interface with it when writing machine-learning code and take advantage of several pre-implemented functionalities from viewing, creating, retrieving from, and appending to datasets, defining task labels, generating custom/variant code representations, to training and evaluating supported models on such datasets.

<sup>7</sup><https://github.com/acr31/features-javac>

**Table 2** JEMMA dataset artifacts, locations, and sizes

Artifact	DOI	Size
<i>Metadata: Projects</i>	<a href="https://doi.org/10.5281/zenodo.5807578">https://doi.org/10.5281/zenodo.5807578</a>	4.7 MB
<i>Metadata: Packages</i>	<a href="https://doi.org/10.5281/zenodo.5807586">https://doi.org/10.5281/zenodo.5807586</a>	42.2 MB
<i>Metadata: Classes</i>	<a href="https://doi.org/10.5281/zenodo.5808902">https://doi.org/10.5281/zenodo.5808902</a>	269.7 MB
<i>Metadata: Methods</i>	<a href="https://doi.org/10.5281/zenodo.5813089">https://doi.org/10.5281/zenodo.5813089</a>	2.8 GB
<i>Properties: [TLOC]</i>	<a href="https://doi.org/10.5281/zenodo.5813102">https://doi.org/10.5281/zenodo.5813102</a>	335.5 MB
<i>Properties: [SLOC]</i>	<a href="https://doi.org/10.5281/zenodo.5813094">https://doi.org/10.5281/zenodo.5813094</a>	335.0 MB
<i>Properties: [NUID]</i>	<a href="https://doi.org/10.5281/zenodo.5813028">https://doi.org/10.5281/zenodo.5813028</a>	335.6 MB
<i>Properties: [NTID]</i>	<a href="https://doi.org/10.5281/zenodo.5813029">https://doi.org/10.5281/zenodo.5813029</a>	336.7 MB
<i>Properties: [NMTK]</i>	<a href="https://doi.org/10.5281/zenodo.5813032">https://doi.org/10.5281/zenodo.5813032</a>	342.5 MB
<i>Properties: [NMRT]</i>	<a href="https://doi.org/10.5281/zenodo.5813034">https://doi.org/10.5281/zenodo.5813034</a>	333.3 MB
<i>Properties: [NMPR]</i>	<a href="https://doi.org/10.5281/zenodo.5813053">https://doi.org/10.5281/zenodo.5813053</a>	333.3 MB
<i>Properties: [NMOP]</i>	<a href="https://doi.org/10.5281/zenodo.5813055">https://doi.org/10.5281/zenodo.5813055</a>	334.5 MB
<i>Properties: [NMLT]</i>	<a href="https://doi.org/10.5281/zenodo.5813054">https://doi.org/10.5281/zenodo.5813054</a>	333.4 MB
<i>Properties: [NAME]</i>	<a href="https://doi.org/10.5281/zenodo.5813308">https://doi.org/10.5281/zenodo.5813308</a>	432.0 MB
<i>Properties: [MXIN]</i>	<a href="https://doi.org/10.5281/zenodo.5813081">https://doi.org/10.5281/zenodo.5813081</a>	267.0 MB
<i>Properties: [CMPX]</i>	<a href="https://doi.org/10.5281/zenodo.5813084">https://doi.org/10.5281/zenodo.5813084</a>	267.1 MB
<i>Properties: [NUPC]</i>	<a href="https://doi.org/10.5281/zenodo.7019128">https://doi.org/10.5281/zenodo.7019128</a>	333.3 MB
<i>Properties: [NUCC]</i>	<a href="https://doi.org/10.5281/zenodo.7019176">https://doi.org/10.5281/zenodo.7019176</a>	333.6 MB
<i>Properties: [NMNC]</i>	<a href="https://doi.org/10.5281/zenodo.7019960">https://doi.org/10.5281/zenodo.7019960</a>	334.0 MB
<i>Properties: [NMLC]</i>	<a href="https://doi.org/10.5281/zenodo.7020084">https://doi.org/10.5281/zenodo.7020084</a>	333.2 MB
<i>Properties: [NLDF]</i>	<a href="https://doi.org/10.5281/zenodo.1096080">https://doi.org/10.5281/zenodo.1096080</a>	333.6 MB
<i>Properties: [RSLK]</i>	<a href="https://doi.org/10.5281/zenodo.1096082">https://doi.org/10.5281/zenodo.1096082</a>	334.0 MB
<i>Represent.: (TEXT)</i>	<a href="https://doi.org/10.5281/zenodo.5813705">https://doi.org/10.5281/zenodo.5813705</a>	3.8 GB
<i>Represent.: (TKNA)</i>	<a href="https://doi.org/10.5281/zenodo.5813717">https://doi.org/10.5281/zenodo.5813717</a>	3.3 GB
<i>Represent.: (TKNB)</i>	<a href="https://doi.org/10.5281/zenodo.5813730">https://doi.org/10.5281/zenodo.5813730</a>	4.6 GB
<i>Represent.: (ASTS) *</i>	<a href="https://doi.org/10.5281/zenodo.5813880">https://doi.org/10.5281/zenodo.5813880</a>	4.1 GB
<i>Represent.: (FTGR) *</i>	<a href="https://doi.org/10.5281/zenodo.5813933">https://doi.org/10.5281/zenodo.5813933</a>	5.2 GB
<i>Represent.: (C2VC) *</i>	<a href="https://doi.org/10.5281/zenodo.5813993">https://doi.org/10.5281/zenodo.5813993</a>	6.1 GB
<i>Represent.: (C2SQ) *</i>	<a href="https://doi.org/10.5281/zenodo.5814059">https://doi.org/10.5281/zenodo.5814059</a>	10.9 GB
<i>Callgraphs: Projects</i>	<a href="https://doi.org/10.5281/zenodo.6758937">https://doi.org/10.5281/zenodo.6758937</a>	7.2 GB

On a high level, the *Workbench* supports the following types of operations. The next sections describe the usage of the *Workbench* tools to perform some of these operations in more detail.

- GET meta-data, properties, representations, callers/ees, n-hop context
- ADD meta-data, properties, representations, callers/ees
- GEN (create/adapt) representations
- RUN (train/evaluate) supported models on a task

These operations are supported by a set of APIs. The JEMMA *Workbench* implementations, along with an exhaustive list of APIs, are made available online.<sup>8</sup> We demonstrate the usage of some of the APIs in this section.

**Potential Uses for the Data** In the following paragraphs, we describe some of the potential ways in which data can be utilized directly with the help of the *Workbench* tools.

*Create a multitude of new datasets:* Since JEMMA is built on top of the 50K-C Dataset, we have catalogued all of the 50,000 projects and their child code entities, made them uniquely-identifiable, and provided a range of properties associated with them, along information on inter-relationships. Using this information a multitude of datasets can be prepared, not just specifically for ML4Code but also for other purposes, e.g. creating a dataset of projects based on the project size, or creating a dataset of method snippets with complexities based on a criterion, and so on, for a diverse range of use-cases.

*Define ML task datasets:* Users can decide on a modelling objective and retrieve the training data from JEMMA. This may sound straightforward but preparing a sound dataset for model training is one of the most important steps, and it is often time-consuming given the amount of data cleaning and transformations involved before model training. The JEMMA *Workbench* tools help users choose from a range of pre-processed source code representations across 8M samples, and filter them based on a range of properties, and even use the properties as prediction labels. The representations and properties can be used, either singularly, or in combination, to generate thousands of combinations of clean and balanced ML task datasets ready to be trained. Section 4.2.1 demonstrates a similar example.

*Retrieve code information:* Within JEMMA we catalogue code at the project, package, class, and method-level. Furthermore, we process source code into feature graphs yielding feature-rich code information at the AST node-level. This enables users to access a diverse range of granularities from coarse file-level to finer node-level information. Not only that, information such as the data-flow, control-flow, etc. between nodes are also available at the node level—providing a remarkable level of detail for code entities. In addition, call-graph links provide information on the inter-procedural relationships across entities within projects. This affords users to access code entities at scale, in different granularities, with detailed and intricate information based on various intra- and inter-procedural relationships.

**Potential Operations on the Retrieved Data** Here we describe potential operations that can be performed on the data once they have been queried.

*Working with representations.* Different model architectures process input in different formats. Graph models work with code represented as graphs while some others may process code represented as ASTs. Using the *Workbench* tools users can easily undertake operations such as extensions and abstractions. For example, users can abstract from an existing representation to create sparse representations (e.g. from feature graphs to just data-flow graphs) or add new information to existing representations to make it more feature-rich. Our tools help users create abstracted or extended alternatives of existing representations. In addition, the *Workbench* tools make it extremely easy for users to re-generate representations from scratch when necessary (see Section 4.2.2).

*Conduct analyses.* Having accessed the data at scale, one of the most obvious things that users can do is to conduct statistical analyses. In addition, users can conduct a multitude of empirical studies to test a range of hypotheses utilizing the aggregated array of information on millions of code entities.

<sup>8</sup><https://github.com/giganticode/jemma/>

*Training and evaluating models.* Since JEMMA was prepared with ML4Code in mind, users can easily train/evaluate a number of models, conduct inference, and establish benchmarks for tasks. The diversity of representations facilitates training on several different types of model architectures, from graph-based models, to models that take ASTs as input, and other architectures such as code2seq which reason over a bag of AST-paths. This enables users to model source code in various formats and combinations and extract valuable insights.

In the next sections, we concretely demonstrate how JEMMA can be extended and used, emphasizing on some essential use-cases.

## 4.1 Extending JEMMA

In Section 4.1.1 we describe how JEMMA can be extended with a new property; in Section 4.1.2 we describe how it can be extended with a new representation; and in Section 4.1.3 we show new projects can be added to JEMMA.

### 4.1.1 Adding a New Property

The simplest way to extend JEMMA is to add a new property. This could be any property of interest that can be computed for a source code entity. Examples include defining a new source code metric, or the result of a static analysis tool indicating the presence (or absence) of a specific source code characteristic.

To extend JEMMA with a new property, the workflow has three main steps: a) accessing a set of source code entities, b) generating associated property values, and c) merging the associated property values to the dataset. JEMMA facilitates accessing the correct code input by providing the location and metadata for code entities, and several initial representations (raw text, ASTs, etc.). An associated property could then be obtained either directly (e.g. method name) or by means of a code analysis tool (e.g. cyclomatic complexity).

Figure 3 shows how output metrics from the *Metrix++* tool can be associated with the methods in JEMMA and added to the dataset as properties. The yellow highlights mark the *Workbench* API calls in the code snippets.

### 4.1.2 Adding a New Representation

Different machine learning models of code require different source code representations as input. Some models reason over tokenized source code, while other models reason over more complex structures such as ASTs. Each representation comes with its own set of advantages and drawbacks, while one is extremely feature-rich the other is simple, scalable, and practical. Therefore, the work on representations is still an active area of research—as researchers are continuing to develop new source code representations, or improving the present ones, e.g., by augmenting them with further information.

JEMMA makes it quite simple to do both: create new representations, and modify existing ones. There are three main steps to extend JEMMA with a representation: a) accessing a set source code entities, b) generating associated representations, and c) merging the representations to JEMMA.

The raw source text, or even representations such as the AST, could be accessed directly to produce new representations for associated code entities. And with an array of source code representations readily made available for over 8 million code entities, simplifying or augmenting such representations to create others would save a lot of pre-processing

```

import os, csv, argparse
import pandas as pd
import jemma.jemma_utils as ju
import jemma.jemma_contrib as jc
from subprocess import Popen, PIPE

def run_cmd(cmd):
    process = Popen(cmd, shell=True, stdout=PIPE, stderr=PIPE)
    res, err = process.communicate()
    return res

parser = argparse.ArgumentParser()
parser.add_argument("--metrixplusplus_path", required=True)
parser.add_argument("--output_path", required=True)
parser.add_argument("--project_path", required=True)
args = parser.parse_args()

# generating metrics for project
project_name = ju.get_project_name(project_path=args.project_path)
project_id = ju.get_project_id(project_name=project_name)
output_dbfile = os.path.join(args.output_path, project_name + ".db")
output_cvfile = os.path.join(args.output_path, project_name + ".csv")
output_mtfile = os.path.join(args.output_path, project_name + "_cmpx.csv")

collect_cmd = ("python2 " + args.metrixplusplus_path +
               " collect --db-file=" + output_dbfile +
               " --std.code.lines.code --std.code.complexity.maxindent" +
               " --std.code.complexity.cyclomatic -- " + args.project_path)

export_cmd = ("python2 " + args.metrixplusplus_path +
              " export --db-file=" + output_dbfile +
              " > " + output_cvfile + " -- " + args.project_path)

collect_status = run_cmd(collect_cmd)
export_status = run_cmd(export_cmd)

# linking method_ids
target_data = []
with open(output_cvfile, "r") as rd:
    csv_reader = csv.reader(rd)
    csv_header = next(csv_reader)

    for i, row in enumerate(csv_reader):
        (filepath, regionname, regiontype, md, st, en, cyclomatic, maxindent, nl) = row
        if regiontype == "function":
            class_name = filepath.split("/")[-1][-5]
            class_id = ju.get_class_id(project_id=project_id, class_name=class_name)
            method_id = ju.get_method_id(class_id=class_id, method_name=regionname)
            target_data.append([method_id, cyclomatic])

# exporting and adding metrics to the jemma datasets
target_df = pd.DataFrame(target_data, columns=["method_id", "CMPX"])
jc.add_metrics(metric="CMPX",
               metrics=target_df["CMPX"].tolist(),
               methods=target_df["method_id"].tolist())

```

**Fig. 3** Defining and adding new property to JEMMA: This snippet shows how to add metrics from the *Metrix++* code analysis tool as a new property to the JEMMA dataset

time for the users. In addition, newer representations could also be derived from existing representations based on specific model architectures and needs.

The feature graph representation which we include in our dataset (see Section 3.3.5) is built upon the abstract syntax tree (AST) of the code, and is extended with a number of additional edges, depicting various inter-relationships between the AST nodes (e.g., data-flow, control-flow, lexical-usage, call-edges among others). In addition to other necessary information such as line numbers and position numbers of every source code token, supplementary information such as token types, node types, are also provided. Thus, with this representation, the detail of data provided for every code entity is comprehensive.

Several derivative representations can be created directly from this one representation by choosing the necessary edge types from the feature graph. For example, for models that require the AST representation as input, choosing just the *Child* edges of the feature graph representation would result in the AST representation. Yet for models that reason over the dataflow information, choosing the *LastRead* and *LastWrite* edges of the feature graph would result in a new dataflow-based representation. Experimenting with these variants is important, to obtain a better understanding of the trade-offs between the kinds of information available, what they can contribute to model performance, and the difficulty of obtaining the information.

Beyond deriving descendant representations, adding further edge types to the feature graph is always possible—making it even more feature-rich, and JEMMA facilitates such extensions by providing the base representations for several million code entities. In a similar manner, the other representations included with JEMMA could also be simplified, modified, augmented to create new representations.

Once new representations are created, they are associated with corresponding source code entities by means of UUIDs. The representations can then be added to JEMMA using the *Workbench* APIs—quite similar to that of adding new properties as demonstrated in Fig. 3.

### 4.1.3 Adding a New Project

With the evolution of source code over time, and the inclusion of new libraries and modern technologies — source code datasets (especially the ones built for modeling code) must also keep at pace if they are to remain relevant. We built JEMMA on top of the 50K-C Dataset of 50,000 compilable Java projects, however, we want it to be extensible. So, we have provided mechanisms to include additional projects into the fold of JEMMA.

Adding a new project to JEMMA involves three main steps: 1) forking the jemma repository, 2) generating the meta-data, representations, properties, call-graphs — by running the relevant scripts, and 3) making a pull request to add the new-generated data. We provide a simple bash script that helps users generate all the relevant data in one go—generating meta-data and cataloging code entities within the project, generating representations, generating properties, and generating project-level call-graphs. Once the data for the new project is ready, users can then make a pull request to append the data to JEMMA *Datasets*. A detailed tutorial is provided in our documentation. Figure 4 lists the command-line procedure to add a new project to JEMMA.

## 4.2 Using JEMMA

In this section, we describe various scenarios in which JEMMA can be put to use. In Section 4.2.1 we describe how a property can be used to define a prediction task, while

```

# STEP 1: fork the repo at: https://github.com/giganticodedemma.git
# STEP 2: clone the forked repo and make your changes
$ git clone https://github.com/<your-github-username>/jemma.git
$ cd jemma/jemma/
$ git checkout -b <your-branch-name>

# INFO: run the add_project script -4 passing the path to local project
$ sh add_project.sh <full/path/to/local/compiled/project>

# INFO: data is generated in the contrib/ directory
# INFO: add, commit, push to your forked repo
$ git add contrib/
$ git commit -m <your-commit-message>
$ git push -u origin <your-branch-name>

# STEP 3: create a pull request on github from your forked repo

```

**Fig. 4** Procedure to add a new project to JEMMA. `add_project` runs all the sub-scripts necessary to generate all data for the new project

discussing ways in which JEMMA can help avoid common pitfalls and biases. In Section 4.2.2 we explain how source code representations can be used for tasks such as mutation detection and masked prediction.

In Section 4.2.3 we describe how models can be trained and evaluated on prediction tasks using the *Workbench* tools, and finally, in Section 4.2.4 we describe how new and extended representations can be formulated with a greater context.

#### 4.2.1 Defining Tasks Based on Properties

Once a property is defined in JEMMA, it can be used in a variety of ways. One such way is to use them as prediction labels for a prediction task. A good example of such a prediction task is complexity prediction, i.e., given a snippet of code as input, a source code model must predict its cyclomatic complexity (property) as output. While this may appear trivial (taking a random sample of entities that have that property defined, and splitting it into training, validation, and test sets), in practice it is often more complex. This is because care must be taken that the data does not contain biases that provide an inaccurate estimate of model performance. In this context, there are several groups of issues that JEMMA helps contend with while defining the task datasets.

**Rare Data** The first is that some property values may be very rare, making them hard to learn at all. Examples of this would be uncommon bugs and errors such as, e.g., resource leaks. Since JEMMA is large to start with (over 8 million method entities), the scale of data makes it much more likely that there is enough data to learn in the first place, compared to other alternatives.

**Defining Task Labels** Once a property is defined, the *Workbench* tools provide flexibility in the use of property values as task prediction labels. For instance, for classification tasks, the *Workbench* API endpoints allow users to query and obtain a balanced set of prediction labels, ready for training.


Furthermore, the *Workbench* tools allow the selection of property values as prediction labels that satisfy some criteria based on other properties in the dataset. A subset of the data can also be selected, if one wants to define a task for which data is more scarce, in order to incentivize sample-efficient models. Similar to a Database Workbench, the *JEMMA Workbench* allows several such operations in the context of defining prediction labels for a machine learning task, and in managing and retrieving large amounts of specific information.

Creating datasets for training machine learning models can often be time-consuming and frustrating. Since we start with a lot of data, it might be necessary to filter it down. With the tools included as a part of our *Workbench*, users can query *JEMMA* and obtain clean, complete, and balanced datasets.

**Investigating and Mitigating Biases** When defining a task, care must be taken that the models learn from the right signal, and not from the correlated signal that may be easier to learn from, but is not actually a predictive factor. Such issues have been observed in related fields, such as in Computer Vision (Beery et al. 2018) and NLP (McCoy et al. 2019; Gururangan et al. 2018).

In source code, other issues might be present, e.g., a random sample of methods may contain a lot of small methods (including many easy to learn getters and setters), which may inflate performance. For instance, the performance of method naming models is much higher on very short methods (3 lines), than it is for longer methods (Alon et al. 2018). To mitigate this, *JEMMA Workbench* tools can be used to filter and leverage the already existing properties to empirically investigate the performance of models on the tasks and get insights.

In the case of the code complexity example, Fig. 5 shows the relationship between the size of methods (SLOC) and their complexity as a hexbin plot. We observe that there is an overall tendency for shorter methods to be less complex, and longer methods to be more complex. On the other hand, there also methods that are very long, but have very low



**Fig. 5** Hexbin plot of cyclomatic complexity (y-axis) vs. source lines of code (x-axis)



complexity (along the bottom axis). This information can be used to properly balance the data, for instance, by making sure that examples that are short and complex, and examples that are long and simple, are also included in the training and evaluation datasets.

**Avoiding Data Leakage** Multiple studies have shown that code duplication is prevalent in large source code corpora (Schwarz et al. 2012), (Lopes et al. 2017), and that it impacts machine learning models (Allamanis 2019). Since JEMMA is built on top of 50K-C, we benefit from its selection of projects, which intentionally limited duplication. 50K-C's filtering significantly reduces the risk of leakage across projects.

However, since source code can also be repetitive within projects, it could also be a potential source of data leakage. Models that are trained and tested with files from the same project can see their performance affected (LeClair and McMillan 2019). Since JEMMA keeps the metadata of which project a method belongs to, it is easy to define training, validation and test splits that all contain code from different projects, if necessary.

#### 4.2.2 Defining Tasks Based on Representations

JEMMA can also be used to define tasks that operate on the source code representations themselves, rather than predicting a source code property. These tasks are usually of two forms: a) masked code prediction tasks, and b) mutation detection tasks.

- (a) **Masked code prediction tasks.** In a masking task, one or more parts of the representation are masked with a special token (e.g., “<MASK>”), and the model is tasked with predicting the masked parts of the representations. Examples of this would include the method naming task, where the name of the method is masked, or a method call completion task, where a method call is masked in the method's body. A simpler variant of this would be to use a multiple-choice format, where the model has to recognize which of several possibilities is the token that was masked.
- (b) **Mutation detection tasks.** In a mutation detection task, the representation is altered with a fixed probability, presumably in a way that would cause a bug (for instance, two arguments in a method call can be swapped (Pradel and Sen 2018)). The task is to detect that the representation has been altered. This can either be formulated as a binary classification task (altered vs not altered), or, as a “pointing” task, where the model should learn to highlight which specific portion of the given input was altered (Vasic M et al. 2019).

For both of these tasks, the input representation needs to be modified in some way. JEMMA can help with this. For simple modifications (e.g., masking the first occurrence of an operator), it is enough to directly change the default textual representation, and then use the *Workbench* APIs to re-generate the other representations. Figure 6 shows an example of how to generate new representations for a masking task—method call completion.

When doing these kinds of changes, particular care has to be given to data leakage issues. For instance, for a method naming task, the name of the method should be masked in the method's body if it occurs there. Other bias issues can affect these tasks as well, such as a method naming task that over-emphasizes performance on getters and setters. JEMMA can be used to analyse the performance of the models on the task and extract insights that may affect the design of the task.

The snippet above shows how representations can be re-generated for masked code snippets. The `gen_representation` call handles running all the necessary tools in the background, such that given any source code snippet, representations can be generated on the fly.

```

gen_comp_representations.py

import pandas as pd
import argparse
import jemma.jemma_utils as ju
import jemma.jemma_task_utils as jt

parser = argparse.ArgumentParser()
parser.add_argument("--methods_csv_path", required=True)
args = parser.parse_args()

methods_df = pd.read_csv(args.methods_path, header=0, low_memory=False)
method_ids = methods_df['method_id'].tolist()
metrics = ju.get_metrics(metric="COMP", methods=method_ids)
representations = ju.get_representations(representation="TEXT", methods=method_ids)

for (method_id, comp_metric, repr) in zip(method_ids, metrics, representations):
    pattern = '.' + comp_metric + '('
    mask = '<MASK>'

    repr = []
    for line in repr.split('\n'):
        if pattern in line:
            line = line.replace(comp_metric, mask)
            repr.append(line)
            break # we mask just the first call encountered
        repr.append(line)

    comp_repr = '\n'.join(repr)
    tokn_repr = jt.gen_representation(representation="TKNB",
                                    method_id=method_id,
                                    custom_method_text=comp_repr) # tokens

    c2vc_repr = jt.gen_representation(representation="C2VC",
                                    method_id=method_id,
                                    custom_method_text=comp_repr) # code2vec

    c2sq_repr = jt.gen_representation(representation="C2SQ",
                                    method_id=method_id,
                                    custom_method_text=comp_repr) # code2seq

    ftgr_repr = jt.gen_representation(representation="FTGR",
                                    method_id=method_id,
                                    custom_method_text=comp_repr) # ft_graph

```

**Fig. 6** Generating new representations for a masking task: This example shows how to generate new representations for a masking task

### 4.2.3 Running Models

Once a task is defined, the JEMMA *Workbench* APIs make it easy to run supported models on the task. Several basic baselines are pre-implemented, and models hosted on the huggingface<sup>9</sup> platform are supported out of the box.

The JEMMA *Workbench* API also facilitates the interaction with other libraries, in particular to run models using the code2vec and code2seq architectures, as well as Graph Neural Network (GNN) models implemented with the ptgnn<sup>10</sup> library.

Finally, since models based on the transformer architecture, currently, have been the state of the art for a variety of tasks, JEMMA allows to easily interface with HuggingFace's

<sup>9</sup><https://huggingface.co/models>

<sup>10</sup><https://github.com/microsoft/ptgnn>

```

run_models.py

import pandas as pd
import argparse
import jemma.jemma_utils as ju
import jemma.jemma_task_utils as jt

parser = argparse.ArgumentParser()
parser.add_argument("--methods_csv_path", required=True)
args = parser.parse_args()

# Selecting data for complexity task
# criteria for choosing data depends on the user
# for now, we'll consider selecting the dataset from the 10 largest projects

target_projects = [
    '45410069-e58f-4214-a22d-5457968501b7', 'a18aa329-1ff8-4458-9dd0-4ed054844fd',
    'b8133cb2-0e41-468b-b85d-54f11a527a0b', '906abd61-b7a4-4ab6-9f72-38ceff674620',
    '6e440c8f-99e1-4f20-b00a-9bebe0e18fec', 'a4d054f6-0b62-48d6-a9e0-af94b0f75aa5',
    'cfe3c3c7-79cf-49ae-a14a-7e2740e322eb', '6de59384-daed-4838-ab53-0a4f9db5c90b',
    '5fbcfde7-0018-44fc-a5d0-f4ded035454b', '265bee2a-557c-4114-9a5d-9f97c86f86a2']

methods_df = pd.read_csv(args.methods_path, header=0, low_memory=False)
methods_df = methods_df[methods_df['project_id'].isin(target_projects)]
method_ids = methods_df['method_id'].tolist()

cmpx_metrics = ju.get_metrics(metric="CMPX", methods=method_ids)
cmpx_representations = ju.get_representations(representation="TKNB", methods=method_ids)

target_data = list(zip(method_ids, cmpx_metrics, cmpx_representations))
target_data_df = pd.DataFrame(target_data, columns=["method_id", "CMPX", "TKNB"])

# Balancing the classes in the dataset - basic balancing based on minimum count of a label
df_g = target_data_df.groupby("CMPX")
df_g = pd.DataFrame(df_g.apply(lambda x: x.sample(df_g.size().min()).reset_index(drop=True)))
df_g.to_csv(args.output_csv_path, index=False)

# splitting train/test sets
total_method_ids = df_g['method_id'].tolist()
train_method_ids = total_method_ids[(int(df_g.shape[0] * 0.8))]
test_method_ids = total_method_ids[(int(df_g.shape[0] * 0.2)):]

# running a transformer model
res = jt.run_models(metric="CMPX",
                    representation="TKNB",
                    train_methods=train_method_ids,
                    test_methods=test_method_ids,
                    models=["microsoft/codebert-base"])

print(res)

```

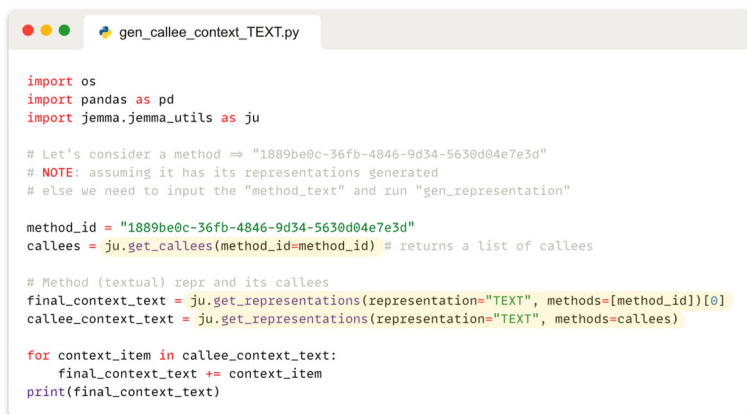
**Fig. 7** Running a Transformer model: Evaluating a transformer model on a prediction task, specifically the cyclomatic complexity prediction task

Transformer library (Wolf et al. 2019b). This allows a variety of pre-trained models to be fine-tuned on the tasks defined with JEMMA (such as CodeBERT (Feng et al. 2020), GraphCodeBERT (Guo et al. 2020) etc.). Figure 7 shows how to run a Transformer model on the method complexity task using the *Workbench*.

#### 4.2.4 Defining Representations with Larger Contexts

One of our goals with JEMMA is to allow experimentation with novel source code representations. In particular, we want users to be able to define representations that can take into account a larger context than a single method, or a single file, as is done with the vast majority of work today.

The key to building such extended representations is to have access to the necessary contextual information. The extensive pre-processing we did to create JEMMA gives us all the



```

import os
import pandas as pd
import jemma.jemma_utils as ju

# Let's consider a method => "1889be0c-36fb-4846-9d34-5630d04e7e3d"
# NOTE: assuming it has its representations generated
# else we need to input the "method_text" and run "gen_representation"

method_id = "1889be0c-36fb-4846-9d34-5630d04e7e3d"
callees = ju.get_callees(method_id=method_id) # returns a list of callees

# Method (textual) repr and its callees
final_context_text = ju.get_representations(representation="TEXT", methods=[method_id])[0]
callee_context_text = ju.get_representations(representation="TEXT", methods=callees)

for context_item in callee_context_text:
    final_context_text += context_item
print(final_context_text)

```

**Fig. 8** Building a Context: This example shows how to combine a textual representation of a method with additional context from its direct callees

relevant tools to gather that information. The metadata of JEMMA documents the containment hierarchies (e.g., which files belong to which project, and which classes belong to which package etc.) and provides the ability to uniquely and unambiguously identify source code entities at different granularities. In addition, the call graph data documents which are the immediate callers and callees of each individual method. Since the call graphs link to each method identified by their UUID, all the properties of the methods, including their representations, can be accessed easily and systematically. Thus, from navigating the call graph and the containment hierarchy, various types of global contexts can be defined at the class-, package-, or even project-level. We present two simple examples in the [Appendix](#).

Figures 8 and 9 show how to combine representations of a given method with the representations of its direct callees to include greater context. We encourage users to experiment with more complex representations adding context information that go beyond a single



```

import os
import pandas as pd
import jemma.jemma_utils as ju

# Let's consider a method => "1889be0c-36fb-4846-9d34-5630d04e7e3d"
# NOTE: assuming it has its representations generated
# else we need to input the "method_text" and run "gen_representation"

method_id = "1889be0c-36fb-4846-9d34-5630d04e7e3d"
callees = ju.get_callees(method_id=method_id) # returns a list of callees

# Method (bag of paths) repr and its callees
final_context_c2vc = ju.get_representations(representation="C2VC", methods=[method_id])[0]
callee_context_c2vc = ju.get_representations(representation="C2VC", methods=callees)

for context_item in callee_context_c2vc:
    final_context_c2vc += context_item
print(final_context_c2vc)

```

**Fig. 9** Building a Context: This example shows how to combine a code2vec representation of a method with additional context from its direct callees

method. The extensive pre-processing of data, at the scale of tens of thousands of projects, combined with the *Workbench* makes it possible to do so easily.

## 5 Empirical Study I: On the Extent of Non-Localness of Software

Since software is made of a lot of interacting entities across files, packages, and projects—modeling source code by learning on smaller entities of code (e.g., methods) can at best provide a localized understanding of source code for a given task. We hypothesize that in order to improve our source code models comprehensively beyond localized understanding, we must include non-local contextual features while modeling source code. As a result, we must study the extent of non-localness of software and whether non-local context is useful.

In this section, we demonstrate the utility of *JEMMA Dataset* and *Workbench* through an empirical investigation. We study the extent to which software is made up of interacting functions and methods in a sample of projects contained in *JEMMA* by analysing their call graphs. We observe how often method calls are local to a file, cross file boundaries, or are calls to external APIs. Then, we analyze the performance on the method call code completion task through the lens of call types when non-local context is added. We pose the following research questions for this part of our study.

- RQ. 1. To what extent are method calls non-local?
- RQ. 2. What is the effect of adding non-local context for the method call completion task?

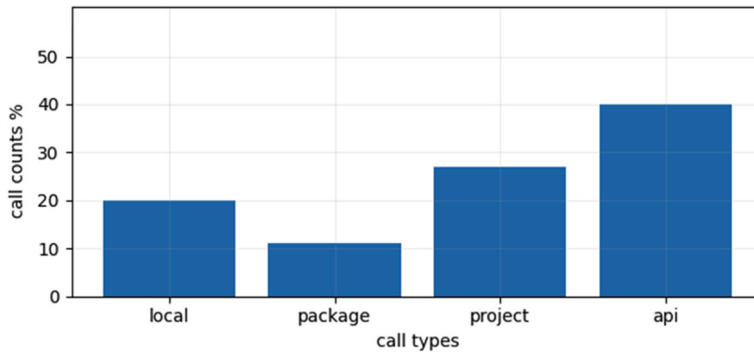
### 5.1 Extent of Non-Localness of Code

Since methods generally do not exist in isolation, a large number of associations can be found among source code entities in project-wide contexts. Thus, we pose the first research question to determine the extent of non-local association of methods with other source code entities defined the project and beyond. In other words, we attempt to determine the extent of non-localness dependence of software.

To determine the extent of non-localness of software, we first track the interacting method entities within projects. For each method defined in a project, we count the number of unique callers and number of unique callees in the call graph and find that over 70% of the methods have at least one unique caller, and over 72% of the methods have at least one unique callee. This confirms that software is highly interconnected.

But to what extent is the interconnected-ness strictly non-local? To study this we measure the frequency of the various types of method calls in these projects. We classify all calls into four categories as listed below:

- *Local calls*. The entity is defined in the same file; thus, a machine learning model that has a file context would be likely to see it.
- *Package calls*. The entity is defined in the same Java package (i.e., the classes as in the same file directory).
- *Project calls*. The called entity is defined in the project, but in a different package than the caller.
- *API calls*. The called entity is not defined in the project, but is a call to an imported library.



**Fig. 10** Distribution of calls by type

Figure 10 shows the distribution of the calls. We can see that only 20% of calls are *local calls*; these are the calls whose callees are visible to the models that learn from the entire file context, such as CoCoGum (Wang et al. 2020); the remaining 80% of calls are *non-local* and are not visible to models that learn from the file context only. Of these, 12% are *package calls*; thus a model that builds a context of the classes in the same directory to absorb a larger context than the file would have visibility into these callees. On the other hand, 28% of calls are *project calls*, thus models would need either a larger context, or the ability to select from this larger context in order to have visibility in the callees. Finally, *API calls* constitute 40% of all calls (inflated by the vast majority of standard library calls). While these are out of reach for most models, a silver lining is that, in practice, it is often possible to learn common *API* call usages as modern large-scale source code models do.

## 5.2 Impact of Non-Localness on Code Completion

Having some insight into the extent of non-localness of software, we turn to the issue of whether adding non-local context can have an impact on the performance of models.

The study of (Hellendoorn et al. 2019a) investigated the differences between code completions from synthetic benchmarks and real-world completions. It found that synthetic benchmarks underweighted the frequencies of method completions relative to real-world completions, and that those were the most difficult. They also observed that among method completions, the hardest ones were the ones from project internal identifiers. Hellendoorn's study offers valuable insights but has limitations: the data for the real-world completions was relatively small (15,000 completions); the study evaluated RNNs with a closed vocabulary, which were unable to learn new identifiers. Since then, open-vocabulary models (Karampatsis et al. 2020; Ciniselli et al. 2021) have considerably improved the state of the art.

The key point from (Hellendoorn et al. 2019a) that motivates our study is the observation: *code completion models struggle the most with project internal identifiers, e.g. method-call completions*. This is because models such as RNNs have a very limited context size, so they are unable to know which identifiers are defined in the project. And since this information is spread over the entire project, it motivates our choice to design a code completion task with much more data-points that focuses particularly on method-call completions considering a project-wide context.

**Table 3** Method call completion (by call types) without vs. with context, % improvement. Scores for: A - BERT, B - CodeBERTa, C - CodeBERT, D - GraphCodeBERT

	<i>Local</i>			<i>Package</i>			<i>Project</i>			<i>API</i>		
	n/c	c	±	n/c	c	±	n/c	c	±	n/c	c	±
A	0.102	0.159	56%	0.112	0.154	37%	0.144	0.181	26%	0.296	0.336	14%
B	0.171	0.284	66%	0.255	0.362	42%	0.278	0.370	33%	0.524	0.606	16%
C	0.137	0.187	36%	0.144	0.176	22%	0.182	0.209	15%	0.376	0.397	6%
D	0.142	0.188	32%	0.147	0.176	20%	0.184	0.209	14%	0.379	0.398	5%

We use the JEMMA *Workbench* to analyse the performance of three state-of-the-art Transformer code models, with the natural-language BERT model as the baseline, on a derivative of the code completion task: method-call completion.

We first train our models without any additional context; and report the exact-match<sup>11</sup> accuracies for method-call predictions. We then train our models including context from caller-callee method entities defined across the project, comparing the results. A brief explanation on how larger contexts may be constructed is presented in the Appendix A. For our experiment, we have considered the context information from a method's 1-hop neighborhood, considering all possible callee names. Furthermore, informed by the observations from the previous section, we separately analyse the performance of models on different strata of the test set, according to the categories defined above: *local calls*, *package calls*, *project calls* and *API calls*.

**Task definition** We define the method call completion task as a masking task: for each method snippet, we mask one single method call in the code snippet at random. These methods can be present in the same class (18% of the dataset), in another class in the same package (10%), in another package in the system (26%), or imported from a dependency (46%). The goal of the task is for a source code model to predict the *exact* method name that was masked. We sample 100K methods from the JEMMA *Datasets*, splitting 80K samples as training data, 5K as validation data, and 15K as test data for training and evaluation.

**Model Performance** We analyze the performance of three large-scale Transformer models of code: CodeBERTa, CodeBERT, and GraphCodeBERT. We use the BERT model as the baseline model for this task. All of these models accept sequences of tokens as input, so we use the token representation for training.

Table 3 shows the accuracy across the call types when the models were trained without context and with additional context. We observe an improvement across all models, and across all call-types, when additional context was included. This shows that tasks like method-call completion—an integral component for the success of code completion, rely on the information beyond the local context and could benefit from additional project-wide context.

Furthermore, we can clearly see that accuracies are much higher for the *API calls* than the other categories, with the second highest being the *local calls*; the *project calls* and the *package calls* having the lowest performance. While we can expect that different models

<sup>11</sup>We intentionally chose against string similarity-based metrics, since the completions need to be *exactly* the same as the *APIs* or method-names actually defined in the project.

would perform differently, the margin between *API calls* and the other types of calls is wide enough to demonstrate that the models perform much better at predicting *API calls* than calls defined within the project.

### 5.3 Implications

From the observations in the previous sections we see that: a) a large number of method calls are non-local, i.e., collectively 80% of the method calls are not local to the same parent class, and b) source code models struggle to predict call completions of methods defined in the same project, but improve when additional non-local context is added.

This nudges us to explore the notion of designing and training source code models in a way that it can reason over a larger context of information, at least at the project-level. It becomes necessary to determine ways in which models could be made aware of the inter-relationships that exist among code entities by providing a feature-rich representation with as much context information that we can possibly fit. With the depth and extent of data that we have gathered, and with the help of our *Workbench*, users can easily construct extended contexts beyond the method-level for use in training context-aware source code models in the future.

## 6 Empirical Study II: OOV is the Next OOV

The studies in the previous section show that software entities have inter-relationships which when considered can affect the performance of models. This section provides data to inform the design of possible model architectures that can absorb a larger context. In particular, we focus on the *size* of this context, as deep learning models can be strongly affected by the input size.

Machine learning models of code once struggled with *Out-Of-Vocabulary* (OOV) issues (Hellendoorn and Devanbu 2017), until more recent models introduced and adopted an open vocabulary (Karampatsis et al. 2020).

We argue that the next problem to address is the *Out-Of-Window* (OOW) issue: all modern state-of-the-art models tend to have a fixed input size, which may not be enough to fit the additional context needed. How to best use this limited resource is thus, an open problem. To that effect, we pose the following research questions in this section:

- RQ. 1. Given the need for fitting additional context, are English-based model tokenizers comparable to language-specific tokenizers?
- RQ. 2. From the perspective of context size, what types of code entities fit modern transformer models at different input size limits?

### 6.1 Transformers, Window Sizes, and Tokenizers

For many machine learning tasks, Transformer-based models (Vaswani et al. 2017) are now the state of the art. Some transformer models that have achieved state-of-the-art performance on source code tasks include CodeBERT (Feng et al. 2020), CodeBERTa (Wolf et al. 2019a), PLBART (Ahmad et al. 2021), CodeT5 (Wang et al. 2021), CodeGen (Nijkamp et al. 2022), GraphCodeBERT (Guo et al. 2020). Codex (Chen et al. 2021) is yet another of these large pre-trained Transformer models, that has demonstrated compelling competence on a variety of tasks without necessarily needing fine-tuning, ranging from program



synthesis, program summarization (Chen et al. 2021), to even program repair (Prenner and Robbes 2021).

However, all Transformers that follow the classic architecture have fixed window sizes: for CodeBERT, it is 512 tokens, while for the largest Codex model (codex-davinci), it is 4,096 tokens. If an input is longer than the window, it is generally truncated. Transformers rely on self-attention, where the attention heads *attend* to each pair of tokens: the complexity is hence quadratic, which renders very large windows prohibitive in terms of training time and inference time. This raises the question: for a given window size, how much code can we expect to fit?

Since Transformers are open-vocabulary models, the tokens that they take as input are actually subtokens, common subsequences of characters learned from a corpus, rather than entire tokens. A word that would be unrecognized by a closed-vocabulary model will, instead, be split up in several more common subtokens. This means that the number of lexical tokens in a method does not match the length of the method in terms of subtokens, and depends on the corpus that was used to train the subword tokenizer. It is important to note that both CodeBERT and Codex are not models trained from scratch on source code: given the amount of time needed to train such a model from scratch, previous models trained on English (RoBERTa for CodeBERT, a version of GPT-3 for Codex) were fine-tuned on source code instead. This means that both CodeBERT and Codex use a subword tokenizer that was not learned for source code, but for English, which might lead to sub-optimal tokenization.

To estimate the number of tokens that a method will take in the model's input window, we first selected a sample of 200,000 Java methods from JEMMA, and used several subword tokenizers to estimate the ratio of subtokens that each subword tokenizer will produce. We first noticed that the choice of subword tokenizer has a significant impact on the produced tokenization, and consequently the amount of code that can fit in a model's input window. We used the following tokenizers for our analyses:

- RoBERTa tokenizer. A byte-level BPE tokenizer, trained on a large English corpus, with a vocabulary of slightly more than 50,000 tokens. A similar tokenizer is used by CodeBERT and Codex.
- CodeBERTa tokenizer. The tokenizer used by CodeBERTa. This tokenizer was trained on source code from the CodeSearchNet corpus, which comprises of 2 million methods in 6 programming languages, including Java.
- Java BPE tokenizer. A tokenizer similar to CodeBERTa tokenizer, trained on 200,000 Java methods from Maven, instead of several languages.
- Java Parser. A standard tokenizer from a Java Parser, that does not perform sub-tokenizations. We use this as a baseline for our analyses.

We tokenized Java source code using the tokenizers above, keeping the Java Parser (standard tokenizer) as the baseline, and then calculated the average percentage- increase or decrease in the number of generated tokens. The CodeBERTa tokenizer learned on multiple programming languages, on average, generates 98 tokens per 100 tokens of the baseline Java Parser tokenizer. This is expected since some common token sequences can be merged in a single token (e.g., `()`; can be counted as one token instead of three tokens). The learned Java BPE tokenizer is even more efficient, using on average 85% of the tokens (i.e. it generates 85 tokens per 100 tokens of the standard tokenizer). This is possible since, for instance, specific class names will be common enough that they can be represented by a single token (e.g., `ArrayIndexOutOfBoundsException`). On the other hand,

the RoBERTa tokenizer is considerably less efficient, needing 126% of the lexical tokens compared to the baseline.

With an equal vocabulary size, the most efficient language-specific encoding can fit close to 25% more effective tokens in the same window size. For a window size of 512, a Java-specific tokenizer will, on average, be able to effectively fit 602 actual tokens, while the English-specific tokenizer—used by both CodeBERT and Codex—will be able to fit only 409 actual tokens. For example, for tokens such as *ArrayIndexOutOfBoundsException*, efficient language-specific code tokenizers will tokenize it as a single token, rather than six separate tokens.

This establishes that language-specific code tokenizers are more efficient in tokenizing source code compared to their English-language counterparts. And since almost all model architectures have a maximum input size limit, the tasks that specifically rely on additional context information can benefit from efficient tokenizers, whereby input source code snippets can be represented in less number of tokens leaving space for additional context information. From this perspective, efficient tokenizers can be helpful because having the possibility of including additional context can ultimately improve model performance.

## 6.2 Fitting Code Entities

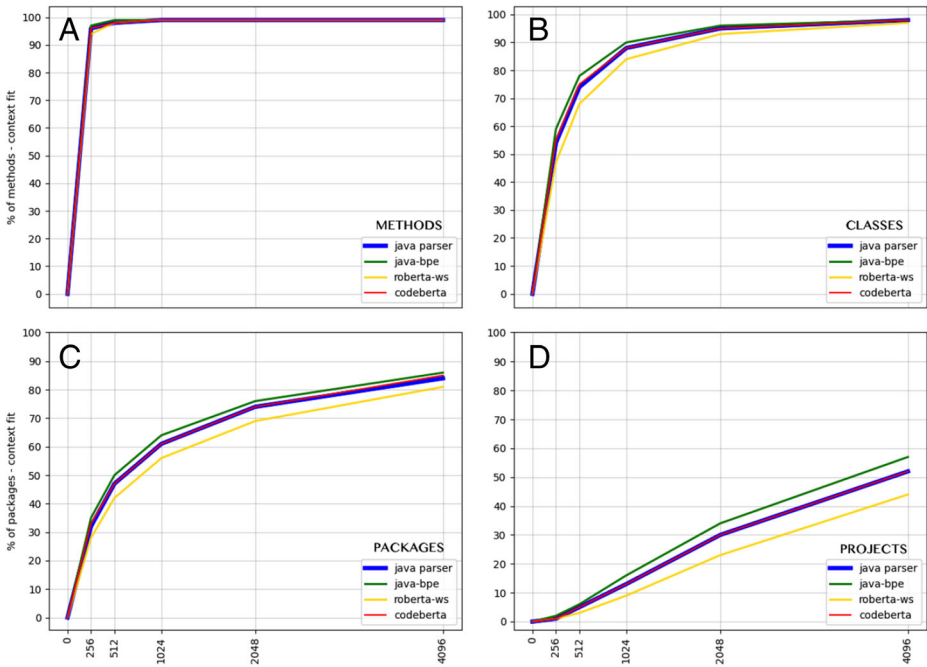
Taking the same 400 projects as in the code completion study in Section 5, we tokenize the methods and the classes in these projects with the four tokenizers above. We then estimate the size of higher-level entities (packages and projects) by summing the token sizes of the classes in them. We compare these sizes against a range of Transformer window size thresholds:

- ◇ *Small*. A window size of 256 tokens, representing a small transformer model
- ◇ *Base*. A window size of 512 tokens, representing a model with the same size as CodeBERT (Feng et al. 2020).
- ◇ *Large*. A window size of 1,024 tokens, which is the context size used by the largest GPT-2 model (Radford et al. 2019).
- ◇ *XL*. A window size of 2,048 tokens, which is the context size used by the largest GPT-3 model (Brown et al. 2020).
- ◇ *XXL*. A window size of 4,096, which is the context size used by the largest Codex model (Chen et al. 2021).

It is important to note that these models are very expensive to train. In practice, training a model with a *Base* window size of 512 tokens, from scratch, is a significant endeavour inaccessible for most academic groups, leaving fine-tuning as the only practical option. Only industry research groups or large consortiums of academics may have the resources necessary to train such large models. Even conducting inference on the largest of models becomes impractical due to their size.

### 6.2.1 Methods

Figure 11 (top-left) shows the percentage of methods that fit within different window size thresholds. We can see that even the *Small* model (with a maximum input size of 256 tokens) is able to comfortably fit the vast majority of methods (over 94%). The choice of tokenization still matters, as more efficient tokenization can make up to 97% of methods fit in the *Small* model. Overall, a *Base* model with a window size of 512 tokens can fit 99% of the



**Fig. 11** Percentage context-fit for: (a) methods; (b) classes; (c) packages; (d) projects

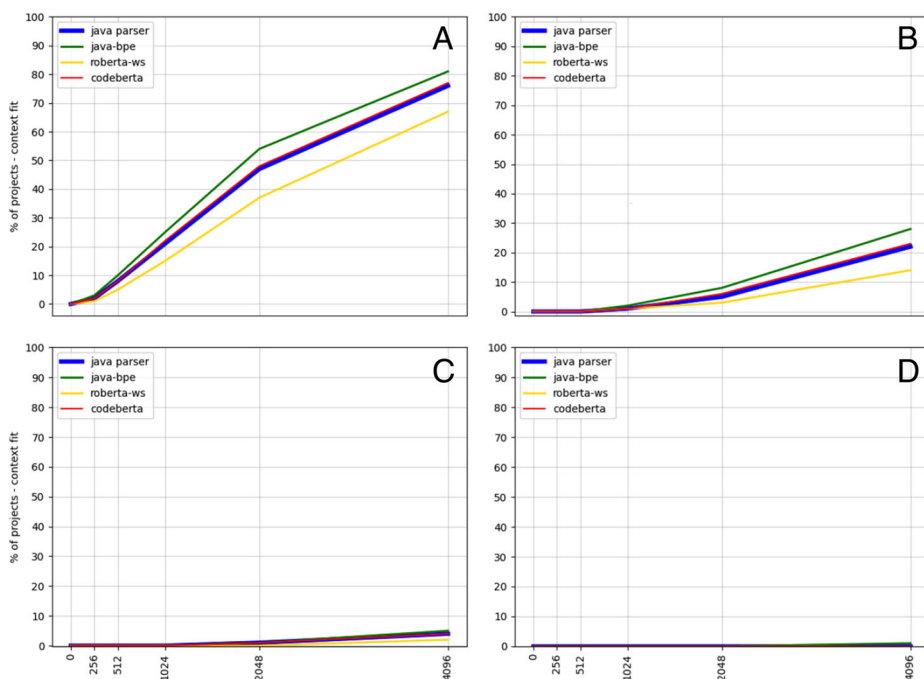
methods in our sample, while only extreme outliers do not fit even in the *XXL* models with a limit of 4096 tokens.

### 6.2.2 Classes

We tokenize the entire source file to compute the context size needed for classes. Figure 11 (top-right) shows the percentage of classes that fit within different window size thresholds. We can see that models with smaller window sizes are beginning to struggle. A *Small* model with a token limit of only 256 tokens will be able to process between 47-59% of the classes. A *Base* model would instead be able to process between 68 and 78% of the classes, while a *Large* model would fit up to 90% of the classes. *XL* models can fit almost more than 95% of the classes on average, but some outliers (2-3%) will remain even for a Codex-sized model.

### 6.2.3 Packages

Figure 11 (bottom-left) shows the percentage of packages that fit within different window size thresholds. Models with smaller window sizes struggle significantly, with a *Small* model able to fit only a 30 to 35% of the packages, and a *Base* model 42 to 50%, depending on the tokenization. A *Large* model succeeds in 55 to 65% of the cases. We can clearly see that even the models with the largest token limits start to struggle while fitting packages into context: 69-76% fit in a window size of 2048 tokens, and 81-86% fit in a window-size of 4096 tokens.



**Fig. 12** Percentage of context-fit for full projects by project sizes. A: up to 20 classes; B: 21-50 classes; C: 51-100 classes, D: more than 100 classes

## 6.2.4 Projects

On average, only half the projects can fit in the window sizes, as seen in Fig. 11 (bottom-right). But since we expect that larger projects would behave differently, we present a context-fit graph for projects based on size (Fig. 12). We observe that while models with large window sizes are able to fit 66-81% of small-sized projects that have 20 or fewer classes, the rate drops drastically as project size increases—falling to 14-28% for medium-sized projects. Beyond this, very few (less than 6%) of the larger projects can fit any window-size. Of note, the largest projects that do not fit the model window sizes, being the most complex, are likely the ones for which the source code models might be the most useful.

## 6.3 Implications

In addressing our research questions, we find that: a) language-specific code tokenizers outperform English tokenizers, and b) code entities at the method- and class-level can comfortably fit the models with the largest of input windows, but fitting larger contexts beyond class-level may still not be practical.

We find that a model that efficiently encodes its code input using a code-specific tokenizer, is able to encode the same data in less space. This leads to a greater amount of context-fitting. Therefore, we need to encourage researchers and model architects to adopt such changes, instead of relying on sub-tokenizations from tokenizers trained on English text.

It's worth noting that classical Transformer models exhibit a quadratic complexity in terms of the input size due to the attention layers. This contributes to their issues in scaling beyond a threshold limit. Thus, reasoning at the scale of packages or projects would require a rethink of the architecture, such as using a Transformer variant that better handles longer sequences such as a Reformer (Kitaev et al. 2020), or another *efficient* Transformer (Tay et al. 2020b) which exhibits lower complexities as input size increases. Whether this is sufficient is uncertain: *efficient* transformers can struggle with very long sequences, as exhibited in specialized benchmarks (Tay et al. 2020a).

While we focused specifically on Transformers as they have a fixed context size window, other models will also be challenged by large input sizes. The ASTs and graph representations of classes, packages, and projects will also have scaling issues as the number of nodes to consider will grow very quickly. Furthermore, Graph Neural Networks can also struggle with long-distance relationships in the graph (Alon and Yahav 2020). Clearly, significant work is needed to find architectures that can fit contexts at the project-level, especially if the model size is to be kept small enough to be manageable.

On the other hand, we see promise in an approach that is able to select the input relevant to the task. Of note, recent work has started to go in this direction for code summarization, both at the file level (Clement CB et al. 2021) and multiple files (Bansal et al. 2021). Significant work lies ahead in devising techniques that truly take into account a larger global context, thus addressing the “*Out-Of-Window*” (OOW) problem; at a minimum, JEMMA provides the data at scale, and tools to investigate this.

## 7 Limitations

JEMMA is the only effort we are aware of in gathering enough data that is preprocessed sufficiently to enable empirical research of machine learning models that can reason on a more global context than the file or method level. Nevertheless, it has several limitations. Some of these issues are inherited from our use of 50K-C, while others are due to limitations in our pre-processing; while the former will be hard to overcome (barring extensive additional data collection), the latter could be mitigated by further processing from our side.

### 7.1 Limitations Stemming from the Use of 50K-C

**Monolingual** JEMMA is comprised of projects in the Java programming language only. This poses issues as to whether models that work well for Java would also work well for other languages. The reason for this limitation is twofold: 1) adding other languages at a similar scale would drastically increase the already extremely significant time we invested in pre-processing data, and 2) restricting to one language frees us from tooling issues: we don't need to settle on a “common denominator” in tool support (e.g., Infer supports few programming languages, and many of its analyses are limited to a single programming language).

**Monoversion** JEMMA is comprised of snapshots of projects, rather than multiple project versions. This prevents us from using it for tasks that would rely on multiple versions, or commit data, such as some program repair tasks. On the other hand, this frees us from issues related to the evolution of software systems, such as performing origin analysis (Godfrey and Zou 2005), which is essential as refactorings are very common in software evolution, and can lead to discontinuities in the history of entities, particularly for the most changed

ones (Hora et al. 2018). Omitting versions also considerably reduces the size of the dataset, which is already rather large as it is.

**Static Data Only** While the projects included in 50K-C were selected because they could be compiled, 50K-C provide no guarantees that they can be run. Indeed, it is hard to know if a project can run, even if it can be compiled. In case it can run, the project likely expects some input of some sort. This leaves running test cases as the only option to reliably gather runtime data. In our previous work in Smalltalk, where we performed an empirical study of 1,000 Smalltalk projects, we could run tests for only 16% of them (Callaú et al. 2014). Thus, JEMMA makes no attempt at gathering properties that comes from dynamic analysis tools at this time. In the future, JEMMA's property mechanism could be used to document whether a project has runnable test cases, as a first step towards gathering runtime information. We could also expand the dataset with the 76 projects coming from XCorpus, which were selected because they are runnable (Dietrich et al. 2017).

## 7.2 Limitations Stemming from our Pre-Processing

**Incomplete Compilation** While the projects in 50K-C were selected because they were successfully compiled, we were not able to successfully recompile all of them. Roughly 18% of the largest projects could not be compiled; this number trends down for smaller projects. We are not always sure of the reasons for this, although we suspect that issues related to dependencies might come into play. This could add a bias to our data, in case the projects that we are unable to compile are markedly different from the ones that we could compile. Nevertheless, all of the meta-data, call-graphs, and almost all of the properties and representations could be generated even for uncompiled projects.

**Imprecisions in Call Graphs** The call graph extraction tool that we use has some limitations that we inherit. In particular, handling methods called via reflection is a known problem for static analysis (Bodden et al. 2011); the call graph extraction tool does not handle these cases. A second issue is related to polymorphism, where it is impossible to know, in the absence of runtime information, which of the implementations can be called. In this case, our call graph has an edge to the most generic method declaration.

**Inner Classes** Our handling of inner classes is limited. Since inner classes are contained in methods, the models can have access to their definitions. However, we do not assign UUIDs to them or to the methods defined in them, as this would significantly increase the complexity of our model (in terms of levels of nesting in the hierarchy), while these cases are overall rare. Additional pre-processing could handle these cases, but we do not expect this to become necessary.

**Class-Level Data** Since most machine learning models of code take method-level samples as input, we work with this representation in our experiments, although we include larger contexts. As a consequence, our modeling of classes and packages is limited in this paper. While information about, for instance, the class attributes is not explicitly modeled in our work, it is easily accessible in the file-level feature graph representations, so that models that wish to use this information can access it.

**Incomplete Preprocessing** At the time of writing, not all the representation data is present for all the projects, due to the very computationally expensive processing that is needed. We

started with the largest projects, and worked our way down to the smaller ones. All of the metadata is present for all of the projects. However, some of the smaller projects (the ones with less than 20 classes) will have their representations computed and added to JEMMA in the coming weeks. A second category of incomplete processing is that some tools will occasionally fail on some very specific input (e.g., the parser used by an analysis tool may handle some edge cases differently than the official parser).

## 8 Conclusion

In this article, we presented JEMMA, a dataset and workbench to support research in the design and evaluation of source code machine learning models. Seen as a dataset, JEMMA is built upon the 50K-C dataset of 50,000 compilable Java projects, which we extend in several ways. We add multiple source code representations at the method level, to allow researchers to experiment on the effectiveness of these, and their variations. We add a project-level call graph, so the researchers can experiment with models that consider multiple methods, rather than a single method or a single file. Finally, we add multiple source code properties, obtained by running source code static analyzers—ranging from basic metrics to advanced analyses characteristics based on abstract interpretation.

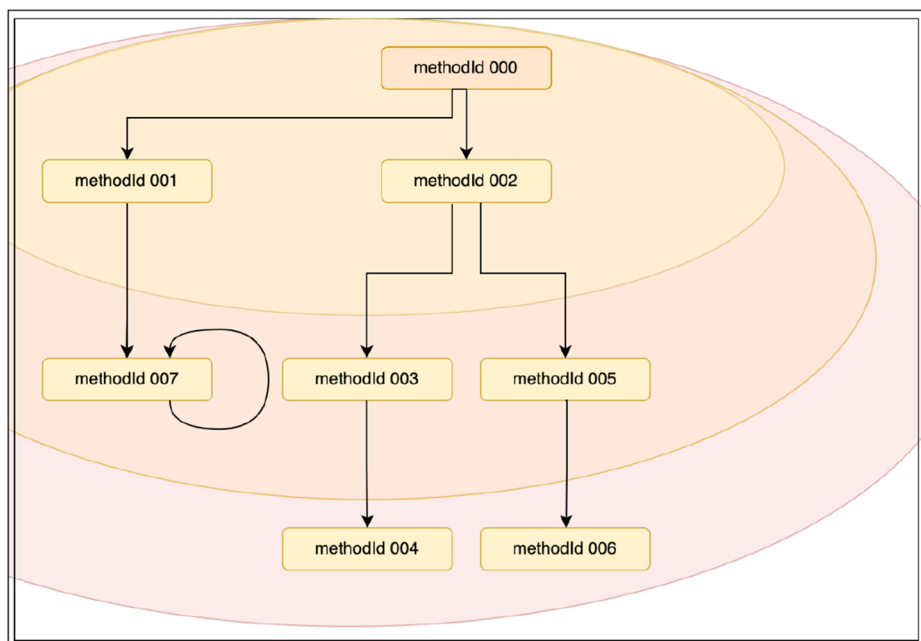
JEMMA *Workbench*, its toolchain and corresponding APIs, help achieve a variety of objectives. JEMMA can extend itself with new properties and representations. It can be used to define machine learning tasks, using the properties and the representations themselves as basis for prediction tasks. The properties defined in JEMMA can be used to get insight into the performance of tasks and pinpoint possible sources of bias. Finally, JEMMA provides all the tools to experiment with new representations that combine the existing ones, allowing the definition of models that can learn from larger contexts than a single method snippet.

Alongside, we have provided examples of usage of JEMMA. We have shown how JEMMA can be used to define a metric prediction and a method call completion task. We have also shown how JEMMA can be used for empirical studies. In particular, we investigated how the performance of our code completion task was impacted by the type of identifier to predict, showing that models performed much better on API method calls than on method calls defined in the project, indicating the need for models that take into account the project's context. Finally, we have shown that taking into account this global context will be challenging, by studying its size. While state-of-the-art transformer models such as CodeBERT can fit most methods in the dataset, fitting package-level or higher context is much more challenging, even for the largest models such as OpenAI's Codex model. This indicates that significant effort lies ahead in defining models able to process this amount of data, a task that we hope JEMMA will support the community in achieving.

## Appendix A

### A.1 Building Larger Contexts

Figure 13 is an illustration of how larger contexts can be gradually created from callers/callees over 1-hop, 2-hop, and greater neighborhood of method calls. For a given method (000), its direct callee methods (001 and 002) provide additional context information from its immediate neighborhood (1-hop or level-1 neighborhood). Further callees



**Fig. 13** Step-wise inclusion of greater context information, based on caller-callee relationships, from 1-hop, 2-hop, and 3-hop neighborhoods

(003, 005, 007) of the direct callee methods provide indirect context information from a further neighborhood (2-hop or level-2 neighborhood) and so on.


Based on the call graph, a given method can have multiple sources of context information from different neighborhoods (level- $n$ ) of callers and callees. Then, depending on the task, specific parts of the multiple sources of context information may be added to the input to establish the context information (e.g., method names, call statements, identifiers, entire method text). For our experiment in Section 5 we have considered the context information from a method's 1-hop neighborhood, considering all possible callee method names.

This example is representative of one of the ways in which how larger contexts may be constructed. Other techniques might consider including context from sibling methods in the same class, or from methods based on data-flow, control-flow, or dependency graphs. The

**Table 4** Method call completion (by size) without vs. with context, and % improvement. Scores for: A - BERT, B - CodeBERTa, C - CodeBERT, D - GraphCodeBERT

	01-20			21-50			51-100			> 100		
	n/c	c	±	n/c	c	±	n/c	c	±	n/c	c	±
A	0.139	0.175	26%	0.155	0.203	31%	0.170	0.220	29%	0.194	0.237	22%
B	0.290	0.360	24%	0.315	0.419	33%	0.340	0.449	32%	0.331	0.426	29%
C	0.177	0.200	13%	0.211	0.247	17%	0.219	0.252	15%	0.245	0.277	13%
D	0.175	0.201	15%	0.212	0.247	17%	0.225	0.253	12%	0.249	0.278	12%





**Fig. 14** Call-type comparison for each project size

use of context information while modeling source code is still nascent, therefore, further research on how to effectively build and use project-wide contexts is an open problem.

## A.2 Model Accuracies by System Sizes

Table 4 shows the accuracies for the method-call completion task, with and without additional context, split across project sizes. Once again, we observe that including additional non-local context improves the accuracy for all of the models. The CodeBERTa model posts the highest accuracies across all project sizes. Interestingly, it is also the model with the most efficient source code tokenizer compared to the others. This matches the observations in Section 5 where the results are split across call-types.

To observe the call-type accuracy comparison (with and without context-information) for each project size, we have included Fig. 14. The rows represent the models. The columns represent the project sizes. The red bars record the accuracies without context, while the green bars record the accuracies with added context — for each call type: class, package, project, api. We can clearly see that for each model, the accuracies are similar across project sizes for the different call types, especially for package, project, api calls. Most importantly, all models appear to improve their scores with added context across project-sizes and call-types.

**Funding** Open access funding provided by Libera Università di Bolzano within the CRUI-CARE Agreement. This work was partially funded by the IDEALS and ADVERB projects of the Free University of Bozen-Bolzano. The authors have no other competing interests to declare that are relevant to the content of this article.

**Data Availability Statements** The particular datasets generated during and/or analysed during the current study (specifically, in Section 5) are made available in the following repository:

<https://github.com/giganticode/jemma/tree/master/jemma/paper/>.

The official JEMMA repository provides documentation, all links to datasets, and any other relevant information: <https://github.com/giganticode/jemma/>

## Declarations

**Conflict of Interests** The authors have no relevant financial or non-financial interests to disclose. The authors have no financial or proprietary interests in any material discussed in this article. Neither are there any usage of third-party artifacts. The authors declare that they have no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Ahmad W, Chakraborty S, Ray B, Chang KW (2021) Unified pre-training for program understanding and generation. In: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.211>
- Allamanis M (2019) The adverse effects of code duplication in machine learning models of code. In: Proceedings of the 2019 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software, pp 143–153
- Allamanis M, Sutton C (2013) Mining source code repositories at massive scale using language modeling. In: 2013 10Th working conference on mining software repositories, MSR, IEEE, pp 207–216
- Allamanis M, Brockschmidt M, Khademi M (2017) Learning to represent programs with graphs. arXiv:171100740
- Allamanis M, Barr ET, Devanbu P, Sutton C (2018) A survey of machine learning for big code and naturalness. ACM Computing Surveys (CSUR) 51(4):1–37
- Allamanis M, Barr ET, Ducousso S, Gao Z (2020) Typilus: Neural type hints. In: Proceedings of the 41st acm sigplan conference on programming language design and implementation
- Alon U, Yahav E (2020) On the bottleneck of graph neural networks and its practical implications. arXiv:200605205
- Alon U, Brody S, Levy O, Yahav E (2018) code2seq: Generating sequences from structured representations of code. arXiv:180801400
- Alon U, Zilberstein M, Levy O, Yahav E (2019) code2vec: Learning distributed representations of code. Proceedings of the ACM on Programming Languages 3(POPL):1–29
- Bansal A, Haque S, McMillan C (2021) Project-level encoding for neural source code summarization of subroutines. arXiv:210311599
- Barone AVM, Sennrich R (2017) A parallel corpus of Python functions and documentation strings for automated code documentation and code generation. arXiv:170702275
- Beery S, Van Horn G, Perona P (2018) Recognition in terra incognita. In: Proceedings of the European conference on computer vision (ECCV), pp 456–473
- Benelallah A, Harrand N, Soto-Valero C, Baudry B, Barais O (2019) The maven dependency graph: A temporal graph-based representation of maven central. In: Proceedings of the 16th International Conference on Mining Software Repositories, IEEE Press, MSR '19, p 344–348. <https://doi.org/10.1109/MSR.2019.00060>
- Bodden E, Sewe A, Sinschek J, Oueslati H, Mezini M (2011) Taming Reflection: Aiding static analysis in the presence of reflection and custom class loaders. In: 2011 33rd international conference on software engineering, ICSE, IEEE, pp 241–250
- Brown TB, Mann B, Ryder N, Subbiah M, Kaplan J, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, et al. (2020) Language models are few-shot learners. arXiv:200514165
- Calcagno C, Distefano D, Dubreil J, Gabi D, Hooimeijer P, Luca M, O'Hearn P, Papakonstantinou I, Purbrick J, Rodriguez D (2015) Moving fast with software verification. In: NASA Formal Methods Symposium, Springer, pp 3–11

- Callaú O, Robbes R, Tanter E, Röthlisberger D, Bergel A (2014) On the use of type predicates in object-oriented software: The case of smalltalk. In: Proceedings of the 10th ACM Symposium on Dynamic languages, pp 135–146
- Chen M, Tworek J, Jun H, Yuan Q, Pinto HPdO, Kaplan J, Edwards H, Burda Y, Joseph N, Brockman G, et al. (2021) Evaluating large language models trained on code. arXiv:210703374
- Ciniselli M, Cooper N, Pascarella L, Poshyvanyk D, Di Penta M, Bavota G (2021) An empirical study on the usage of bert models for code completion. arXiv:210307115
- Clement CB, Lu S, Liu X, Tufano M, Drain D, Duan N, Sundaresan N, Svyatkovskiy A (2021) Long-range modeling of source code files with ewash: Extended window access by syntax hierarchy. arXiv:210908780
- Cordy JR, Roy CK (2011) The nicad clone detector
- De Roover C, Lämmel R, Pek E (2013) Multi-dimensional exploration of api usage. In: 2013 21st international conference on program comprehension, ICPC, IEEE, pp 152–161
- DeFreez D, Thakur AV, Rubio-González C (2018) Path-based function embedding and its application to error-handling specification mining
- Dietrich J, Schole J, Sui L, Tempero E (2017) Xcorpus—an executable corpus of java programs. Journal of Object Technology
- Dyer R, Nguyen HA, Rajan H, Nguyen TN (2013) Boa: A Language and infrastructure for analyzing ultra-large-scale software repositories. In: 2013 35th international conference on software engineering, ICSE, IEEE, pp 422–431
- Feng Z, Guo D, Tang D, Duan N, Feng X, Gong M, Shou L, Qin B, Liu T, Jiang D, et al. (2020) Codebert: A pre-trained model for programming and natural languages. arXiv:200208155
- Geiger FX, Malavolta I, Pascarella L, Palomba F, Di Nucci D, Bacchelli A (2018) A graph-based dataset of commit history of real-world android apps. In: Proceedings of the 15th International Conference on Mining Software Repositories, Association for Computing Machinery, New York, NY, USA, MSR '18, p 30–33. <https://doi.org/10.1145/3196398.3196460>
- Godfrey MW, Zou L (2005) Using origin analysis to detect merging and splitting of source code entities. IEEE Trans Softw Eng 31(2):166–181
- Gonzalez-Barahona J, Dueñas S, Cosentino V, Robles G (2018) Perceval: Software project data at your will. <https://doi.org/10.1145/3183440.3183475>
- Guo D, Ren S, Lu S, Feng Z, Tang D, Liu S, Zhou L, Duan N, Svyatkovskiy A, Fu S, et al. (2020) Graphcodebert: Pre-training code representations with data flow. arXiv:200908366
- Gururangan S, Swayamdipta S, Levy O, Schwartz R, Bowman SR, Smith NA (2018) Annotation artifacts in natural language inference data. arXiv:180302324
- Habib A, Pradel M (2018) How many of all bugs do we find? a study of static bug detectors
- Hellendoorn VJ, Devanbu P (2017) Are deep neural networks the best choice for modeling source code?
- Hellendoorn VJ, Proksch S, Gall HC, Bacchelli A (2019a) When code completion fails: A case study on real-world completions
- Hellendoorn VJ, Sutton C, Singh R, Maniatis P, Bieber D (2019b) Global relational models of source code. In: International conference on learning representations
- Hindle A, Barr ET, Gabel M, Su Z, Devanbu P (2016) On the naturalness of software. Commun ACM 59(5):122–131
- Hora A, Silva D, Valente MT, Robbes R (2018) Assessing the threat of untracked changes in software evolution. In: Proceedings of the 40th International Conference on Software Engineering, pp 1102–1113
- Hovemeyer D, Pugh W (2004) Finding bugs is easy. Acm Sigplan Notices 39(12):92–106
- Husain H, Wu HH, Gazit T, Allamanis M, Brockschmidt M (2019) Codesearchnet challenge: Evaluating the state of semantic code search. arXiv:190909436
- Iyer S, Konstas I, Cheung A, Zettlemoyer L (2016) Summarizing source code using a neural attention model
- Kanade A, Maniatis P, Balakrishnan G, Shi K (2020) Learning and evaluating contextual embedding of source code. In: International Conference on Machine Learning, PMLR, pp 5110–5121
- Karampatsis RM, Sutton C (2020) How often do single-statement bugs occur? the manysstubs4j dataset
- Karampatsis RM, Babii H, Robbes R, Sutton C, Janes A (2020) Big Code!= big vocabulary: Open-vocabulary models for source code
- Kitaev N, Kaiser Ł, Levskaya A (2020) Reformer: The efficient transformer. arXiv:200104451
- LeClair A, McMillan C (2019) Recommendations for datasets for source code summarization. arXiv:190402660
- LeClair A, Jiang S, Mcmillan C (2019) A neural model for generating natural language summaries of program subroutines. In: 2019 IEEE/ACM 41st international conference on software engineering, ICSE, IEEE, pp 795–806

- Li Y, Wang S, Nguyen TN, Van Nguyen S (2019) Improving bug detection via context-based code representation learning and attention-based neural networks. *Proc ACM Program Lang* 3(OOPSLA). <https://doi.org/10.1145/3360588>
- Li Y, Wang S, Nguyen T (2021) A context-based automated approach for method name consistency checking and suggestion. In: 2021 IEEE/ACM 43rd international conference on software engineering, ICSE, IEEE, pp. 574–586
- Lin XV, Wang C, Zettlemoyer L, Ernst MD (2018) Nl2bash: A corpus and semantic parser for natural language interface to the linux operating system. arXiv:180208979
- Liu F, Li G, Fu Z, Lu S, Hao Y, Jin Z (2022) Learning to recommend method names with global context. arXiv:220110705
- Lämmel R, Pek E, Starek J (2011) Large-scale, ast-based api-usage analysis of open-source java projects. pp 1317–1324. <https://doi.org/10.1145/1982185.1982471>
- Lopes CV, Maj P, Martins P, Saini V, Yang D, Zitny J, Sajnani H, Vitek J (2017) Déjàvu: a map of code duplicates on github. *Proceedings of the ACM on Programming Languages*, 1 (OOPSLA): 1–28
- Lu S, Guo D, Ren S, Huang J, Svyatkovskiy A, Blanco A, Clement C, Drain D, Jiang D, Tang D, et al. (2021) Codexglue: A machine learning benchmark dataset for code understanding and generation. arXiv:210204664
- Lu S, Duan N, Han H, Guo D, Hwang Sw, Svyatkovskiy A (2022) Reacc: A retrieval-augmented code completion framework. arXiv:220307722
- Ma Y, Dey T, Bogart C, Amreen S, Valiev M, Tutko A, Kennard D, Zaretzki R, Mockus A (2021) World of code: Enabling a research workflow for mining and analyzing the universe of open source vcs data. *Empir Softw Eng* 26(2):1–42
- Martins P, Achar R, Lopes CV (2018) 50K-c: a dataset of compilable, and compiled, java projects. In: 2018 IEEE/ACM 15th international conference on mining software repositories, MSR, IEEE, pp 1–5
- McCoy RT, Pavlick E, Linzen T (2019) Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. arXiv:190201007
- Milojkovic N, Caracciolo A, Lungu MF, Nierstrasz O, Röthlisberger D, Robbes R (2015) Polymorphism in the spotlight: Studying its prevalence in java and smalltalk. In: 2015 IEEE 23rd International Conference on Program Comprehension, IEEE, pp 186–195
- Mir AM, Latoskinas E, Gousios G (2021) Manytypes4py: A benchmark python dataset for machine learning-based type inference. arXiv:210404706
- Mou L, Li G, Zhang L, Wang T, Jin Z (2016) Convolutional neural networks over tree structures for programming language processing. In: Thirtieth AAAI Conference on Artificial Intelligence
- Nijkamp E, Pang B, Hayashi H, Tu L, Wang H, Zhou Y, Savarese S, Xiong C (2022) A conversational paradigm for program synthesis. arXiv:220313474
- Palsberg J, Lopes CV (2018) Njr: A normalized java resource. In: Companion Proceedings for the ISSTA/ECOOP 2018 Workshops, pp 100–106
- Parr T (2013) The definitive ANTLR 4 reference. Pragmatic Bookshelf
- Pietri A, Spinellis D, Zacchiroli S (2019) The software heritage graph dataset: public software development under one roof. In: 2019 IEEE/ACM 16th international conference on mining software repositories, MSR, IEEE, pp 138–142
- Pradel M, Sen K (2018) Deepbugs: a learning approach to name-based bug detection. *Proceedings of the ACM on Programming Languages* 2(OOPSLA):1–25
- Prenner JA, Robbes R (2021) Automatic program repair with openai’s codex: Evaluating quixbugs. arXiv:2111.03922
- Puri R, Kung DS, Janssen G, Zhang W, Domeniconi G, Zolotov V, Dolby J, Chen J, Choudhury M, Decker L, et al. (2021) Project codenet: A large-scale ai for code dataset for learning a diversity of coding tasks. arXiv:210512655
- Radford A, Wu J, Child R, Luan D, Amodei D, Sutskever I et al (2019) Language models are unsupervised multitask learners. *OpenAI blog* 1(8):9
- Raemaekers S, Van Deursen A, Visser J (2013) The maven repository dataset of metrics, changes, and dependencies. In: 2013 10th working conference on mining software repositories, MSR, IEEE, pp 221–224
- Raychev V, Bielik P, Vechev M (2016) Probabilistic model for code with decision trees. *ACM SIGPLAN Not* 51(10):731–747
- Sajnani H, Saini V, Svajlenko J, Roy CK, Lopes CV (2016) Sourcererc: Scaling code clone detection to big-code. In: Proceedings of the 38th International Conference on Software Engineering, pp 1157–1168
- Sawant AA, Bacchelli A (2017) Fine-grape: fine-grained api usage extractor—an approach and dataset to investigate api usage. *Empir Softw Eng* 22(3):1348–1371
- Schwarz N, Lungu M, Robbes R (2012) On how often code is cloned across repositories

- Spinellis D (2017) A repository of unix history and evolution. *Empir Softw Eng* 22(3):1372–1404
- Svajlenko J, Roy CK (2015) Evaluating clone detection tools with bigclonebench. In: 2015 IEEE International conference on software maintenance and evolution, ICSME. IEEE, pp 131–140
- Tay Y, Dehghani M, Abnar S, Shen Y, Bahri D, Pham P, Rao J, Yang L, Ruder S, Metzler D (2020a) Long range arena: A benchmark for efficient transformers. arXiv:201104006
- Tay Y, Dehghani M, Bahri D, Metzler D (2020b) Efficient transformers: A survey. arXiv:200906732
- Tempero E, Anslow C, Dietrich J, Han T, Li J, Lumpe M, Melton H, Noble J (2010) The qualitas corpus: A curated collection of java code for empirical studies. In: 2010 Asia Pacific Software Engineering Conference, IEEE, pp 336–345
- Terra R, Miranda LF, Valente MT, Bigonha RS (2013) Qualitas. class corpus: a compiled version of the qualitas corpus. *ACM SIGSOFT Softw Eng Notes* 38(5):1–4
- Tian F, Treude C (2022) Adding context to source code representations for deep learning. arXiv:220800203
- Tufano M, Watson C, Bavota G, Penta MD, White M, Poshyvanyk D (2019) An empirical study on learning bug-fixing patches in the wild via neural machine translation. *ACM Trans Softw Eng Methodol (TOSEM)* 28(4):1–29
- Utture A, Kalhauge CG, Liu S, Palsberg J (2020) Njr-1 dataset. <https://doi.org/10.5281/zenodo4839913>
- Vasic M, Kanade A, Maniatis P, Bieber D, Singh R (2019) Neural program repair by jointly learning to localize and repair. arXiv:190401720
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017) Attention is all you need
- Wang K, Christodorescu M (2019) Coset: A benchmark for evaluating neural program embeddings. arXiv:190511445
- Wang Y, Du L, Shi E, Hu Y, Han S, Zhang D (2020) Cocogum: Contextual code summarization with multi-relational gnn on umls. Tech. rep., Microsoft, Tech. Rep. MSR-TR-2020-16, May 2020.[Online]. Available: <https://...>
- Wang Y, Wang W, Joty S, Hoi SC (2021) Codet5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation. arXiv:210900859
- Wei J, Goyal M, Durrett G, Dillig I (2020) Lambdanet: Probabilistic type inference using graph neural networks. arXiv:200502161
- Wolf T, Debut L, Sanh V, Chaumond J, Delangue C, Moi A, Cistac P, Rault T, Louf R, Funtowicz M, Brew J (2019a) Huggingface’s transformers: State-of-the-art natural language processing. CoRR abs/1910.03771
- Wolf T, Debut L, Sanh V, Chaumond J, Delangue C, Moi A, Cistac P, Rault T, Louf R, Funtowicz M, et al. (2019b) Huggingface’s transformers: State-of-the-art natural language processing. arXiv:191003771
- Yao Z, Weld DS, Chen WP, Sun H (2018) StaQC: A systematically mined question-code dataset from Stack Overflow. In: Proceedings of the 2018 World Wide Web Conference, pp 1693–1703
- Yin P, Deng B, Chen E, Vasilescu B, Neubig G (2018) Learning to mine aligned code and natural language pairs from stack overflow
- Yu T, Zhang R, Yang K, Yasunaga M, Wang D, Li Z, Ma J, Li I, Yao Q, Roman S, et al. (2018) Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task. arXiv:180908887
- Zavershynskiy M, Skidanov A, Polosukhin I (2018) Naps: Natural program synthesis dataset. arXiv:180703168
- Zhou Y, Liu S, Siow J, Du X, Liu Y (2019) Devign: Effective vulnerability identification by learning comprehensive program semantics via graph neural networks. arXiv:190903496