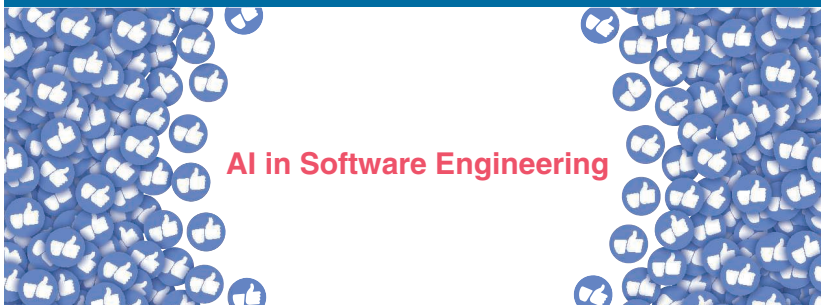


AI in Software Engineering at Facebook

Johannes Bader, Sonia Seohyun Kim, Frank Sifei Luan, Satish Chandra, and Erik Meijer, Facebook, Inc.

// How can artificial intelligence help software engineers better do their jobs and advance the state of the practice? We describe three productivity tools that learn patterns from software artifacts: code search using natural language, code recommendation, and automatic bug fixing. //



AI in Software Engineering

©SHUTTERSTOCK/ALEX GONTAR

ARTIFICIAL INTELLIGENCE (AI) and, more specifically, the machine learning (ML) subarea of AI, has had a transformative impact on almost every major industry today, ranging from retail, to pharmaceuticals, to finance. Not surprisingly, it is beginning to transform the software development industry as well, though significant potential remains untapped.

The underlying basis for the transformative impact of ML is the vast amount of data that are available to be analyzed and mined, from which clever ML algorithms can extract patterns and insights. In software engineering, one of the most easily accessible data are source code itself. For example, GitHub hosts millions of projects, which, together, add up to billions of lines of code; most companies have large proprietary code repositories as well. Other

examples of sources of data include the following:

- incremental changes between repository versions of code
- a large number of tests and their outcomes during continuous integration
- online forums, such as Stack Overflow, in which developers interact with each other.

What are some of the useful insights to be extracted from these data? How can we use ML to extract those insights? Since software engineering is a lot about developer productivity, in the rest of this introduction, we give several examples of scenarios in which we have used ML to help developers work more efficiently; in later sections, we give technical details of how these tools work. Toward the end of the article, we present a broader picture of additional ways in which ML-based insights can help in software engineering.

Code Search Using Natural Language

Consider the life of a developer who has to implement a function, for example, for hiding the Android soft keyboard programmatically. One way to tackle this problem is to study Android application programming interfaces (APIs) and then implement the function, but APIs may take a long time to comprehend. It would be much more efficient to derive inspiration from existing code that serves a related purpose. One way to find a relevant code snippet is with a quick search on Stack Overflow. However, if the question is not already answered on Stack Overflow, posting a new question and waiting for a response has a long latency.

On the other hand, copious amounts of relevant Android code are available on GitHub. The problem is that it is hard to find such relevant snippets directly from a collection of repositories. We have created a technique that can help retrieve a pertinent code snippet directly from source code, starting with just rough keywords. While the search does not come with the explanation that a Stack Overflow post has, it retrieves potentially useful information in real time.

Code Recommendation

Even when one does have a start on which APIs to use for a certain task at hand, the task is not done. When writing code, developers are curious about how other programmers have written similar code, to get reassured or discover considerations they might have missed. If they directly search on a large code corpus for an API name, they might get tens of thousands of results. What they instead want is a small set of sample usages from the repository that gives them some additional information.

Consider an example usage of an Android API method `decodeStream`:

```
Bitmap bitmap = BitmapFactory.  
decodeStream(input);
```

However, if one were to look at related code elsewhere in the repository, one variation is to make sure the app does not crash on an exception:

```
try {  
    Bitmap bitmap = BitmapFactory.decode  
        Stream(input); ...  
} catch (IOException e) {...}
```

This is a different search scenario that we call *code recommendation*. The input is a code snippet, and the

output is a small list of related code fragments that show only a few representative variations of information that occur commonly enough. We will discuss our approach to building such a code recommendation engine in the “Code Recommendation” section.

though there might exist alternate, semantically equivalent ways. A tool that recommends fixes must suggest the one that the developer finds natural in a given context. We will talk about a tool that discovers and learns bug-fixing patterns from data.

When writing code, developers are curious about how other programmers have written similar code, to get reassured or discover considerations they might have missed.

Automatically Fixing Routine Bugs

Code evolves constantly. At Facebook, the Android app repository alone sees thousands of commits per week. Since many of these commits are fixes to various issues, we can use ML to figure out the patterns to these fixes and automatically suggest an appropriate fix.

More specifically, we have found that fixes to static analysis warnings often come from a large palette of code patterns. The following shows an example fix (inserted code in green) of Infer’s warning on potential `NullPointerException` (NPE) (null dereferences) in Java:

```
if (this.lazyProvider == null || shouldSkip) {  
    return false;  
}  
Provider p = this.lazyProvider.get();
```

The notable point is that developers have a strong preference for a certain way to fix a warning, even

Takeaways

These are just some of the many initiatives we have started and incorporated into practice at Facebook. Additional work includes predictive regression test selection,¹ triaging for crashes,² and code autocompletion. Our F8 presentation³ demonstrates how these tools are integrated into the Facebook development environment.

Our thesis is that even simple ML methods can help remove a lot of inefficiencies in the day-to-day life of a developer. No longer should they spend a lot of time looking for information over a repository, finding relevant information from hundreds of code fragments, or fixing simple, predictable bugs manually. In the next section, we describe technical details for the three topics we have introduced.

Code Search

Background

The ability to search over large code corpora can be a powerful

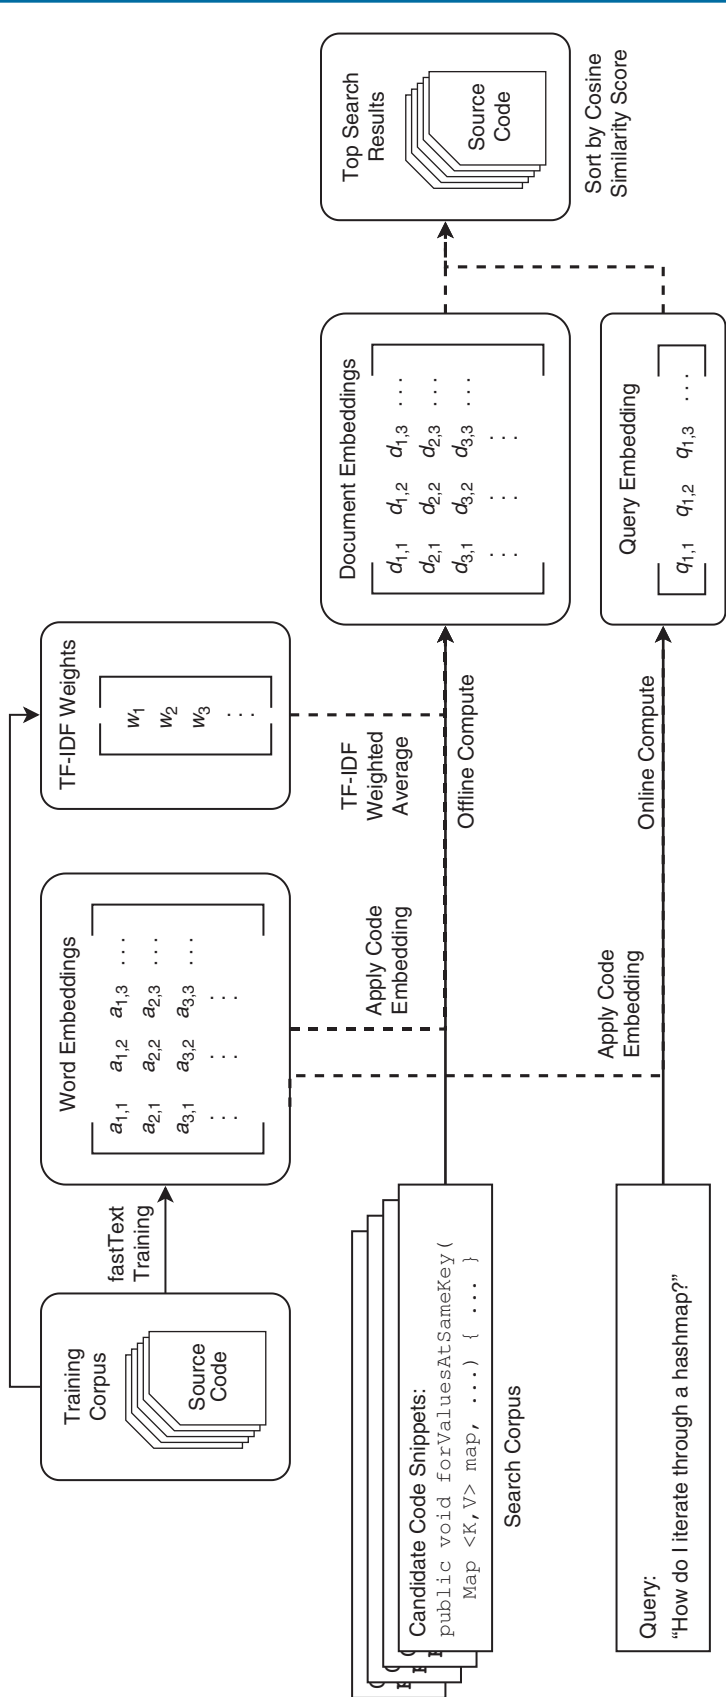


FIGURE 1. The NCS model training and search retrieval. NCS extracts information from the source code, builds word embeddings, and uses TF-IDF weighting to get a document embedding for each code snippet. The query is mapped to the shared vector space, and the most relevant code snippets are ranked with cosine similarity.

productivity booster. Therefore, we have explored ways to search directly over the provided code corpora using basic natural language processing and information-retrieval techniques.

There have been previous works in code search, such as CoCaBu⁴ (a code-search tool that augments natural language queries by adding correlated code vocabulary from Internet forums) and Sourcerer⁵ (a code-search framework that searches over open source projects available on the Internet). However, these tools are not applicable for internal use since most of our developers work with proprietary APIs and frameworks, which are rarely discussed on the Internet.

Thus, we came up with an approach to directly search over the given corpus. Our tool, called *neural code search (NCS)*,⁶ aims to find relevant code snippet examples given a query in natural language.

How Does It Work?

NCS is built using the idea of embeddings, which are vector representations of code that aim to capture the intent of a piece of code in a form suitable for ML. Our hypothesis is that the tokens in source code are generally meaningful, and embeddings derived from these tokens can capture the intent of the code snippet well enough for a code search. NCS creates embeddings at the granularity of a method body.

As shown in Figure 1, NCS works in the following steps.

Extract Information. NCS first extracts relevant tokens from source code to create a “natural language” document. The information NCS extracts includes method

names, comments, class names, and string literals.

Build Word Embeddings. NCS then builds word embeddings using FastText,⁷ which gives vector representations for each word in the corpus. Similar to Word2Vec,⁸ FastText performs unsupervised training such that words appearing in similar contexts have similar vector representations. For example, the embedding of *button* is the closest with the embeddings of *click*, *popup*, and *dismissible* when trained on an Android code corpus.

Build Document Embeddings. Finally, to create a document embedding for each method body in the corpus, NCS computes a weighted average from its tokenized words and its respective word embeddings, as shown in (1), where d is a set of words in a document, C is the corpus containing all documents, and u is a normalizing function. This document embedding serves to capture the overall semantic meaning of the method body. NCS weights the words using term frequency-inverse document frequency (TF-IDF) (2), a well-known weighting technique in information retrieval. The top portion of Figure 1 shows the NCS model training part.

$$v_d = u\left(\sum_{w \in d} u(v_w) \cdot \text{tfidf}(w, d, C)\right) \quad (1)$$

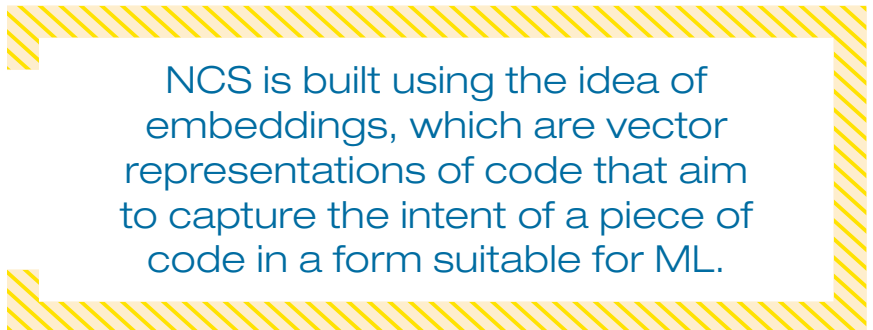
$$\text{tfidf}(w, d, C) = \frac{1 + \log \text{tf}(w, d)}{\log |C| \cdot \text{df}(w, C)} \quad (2)$$

Search Retrieval. Upon receiving a search query, NCS tokenizes the query and uses the same trained word embeddings to represent it as a vector. It is important to note that the tokenization will turn the natural language query to a series of main keywords that captures the essence of the query. For example, the query

“How to get the ActionBar height?” will be tokenized to “get action bar height.” NCS then compares this vector to the document embeddings, as discussed previously. NCS ranks the document embeddings by cosine similarity using Facebook AI Similarity Search,⁹ a standard similarity search algorithm that operates on high-dimensional data, and returns the top

- “How to get the ActionBar height?”
- “How to find MAC address of an Android device programmatically?”

Sachdev et al.⁶ include more details on the training and evaluation of NCS. We further investigated whether deep learning models lead to better code-search results.¹¹



results. The bottom portion of Figure 1 shows the search retrieval part.

Evaluation

We evaluated the effectiveness of NCS on a set of Stack Overflow questions, with the post title as the query and a code snippet from the accepted answer as the desired code answer. Given a query, we measured whether NCS was able to retrieve a correct answer from a large search corpus (GitHub repositories). Out of 287 questions, NCS correctly answered 98 questions in the top 10 results. This evaluation data set, along with the search corpus, is publicly available from Li et al.¹⁰

Some examples of Stack Overflow questions that NCS answers well are as follows:

- “How to delete a whole folder and content?”
- “How to convert an image into Base64 string?”

Developer Feedback

The usage of NCS at Facebook was somewhat different from the way we had envisioned it. Developers did not often write Stack Overflow-style questions; instead, they mostly searched with keyword queries, such as “contract number amount.” Although the raw query types were different, with the tokenization step where we break down both code snippets and the queries into keywords, we were able to deploy NCS with no adaptations to the model at Facebook.

At Facebook, NCS is integrated into the main code-search tools (e.g., the website and IDE) as a complement to the existing exact-match code-search capabilities. Initially, the NCS results and the exact-match (grep-like) results were shown together. Sometimes, though, developers were looking only for exact matches and got confused by the interleaving of results. Consequently, exact-match

provide usage examples to help with integration. API documentation tools provide helpful usage templates, but these are limited to API queries rather than arbitrary code snippets. Code-to-code search engines return exhaustive code matches, whereas our goal is to provide concise recommendations by clustering together similar results. Aroma is able to overcome all of these shortcomings.

How Does Aroma Work?

Aroma indexes the code corpus by creating sparse vector representations of each method body. To do so, it first parses the source code to get a simplified parse tree. Aroma uses this representation because it allows the rest of the algorithm to be language agnostic.

Aroma then extracts features (presented in Figure 2) from the parse tree to capture the code structure and semantics. Aroma creates the feature set of a code snippet by aggregating the features of all tokens in that code snippet. After obtaining the vocabulary of all features, Aroma assigns a unique index to each feature, then converts the feature set to a sparse vector. Given a query code snippet, Aroma runs the following phases to create recommendations.

Feature-Based Search. Aroma takes the query code snippet and creates a vector representation using the same steps in indexing. It then computes a list of top (e.g., 1,000) candidate methods that have the most overlap with the query. This computation is very efficient by utilizing parallel sparse matrix multiplication.

Clustering. Aroma then clusters together similar-looking method bodies. Instead of showing similar or duplicate code, we want to create a single, idiomatic code recommendation from

them. Aroma performs a fine-grained analysis on the candidate methods and finds clusters based on similarities among the method bodies.

Intersecting. The final step is to create one code recommendation for each cluster of method bodies. The intersecting algorithm works by taking the first code snippet as the “base” code and then iteratively pruning it with respect to every other method in the cluster. Its goal is to return only the common coding idiom among the cluster, by removing extraneous lines that may be just situational in a spe-

```
options.inSampleSize = 2;
options.inJustDecodeBounds = ...
bitmap = BitmapFactory.decodeStream(in,
    null, options);
}
```

Both snippets contain a few lines of similar code but also different lines specific to themselves. Aroma’s intersection algorithm compares the base snippet with the second snippet, keeping only the lines that are common in both. It then compares these lines with the next method body. The remaining lines are returned as a code recommendation:

Code-to-code search engines return exhaustive code matches, whereas our goal is to provide concise recommendations by clustering together similar results.

cific method. Refer to our paper¹² for full algorithm details.

As a concrete example, suppose the following two code snippets are in one cluster and that the first one is the “base” code snippet:

```
//Base snippet
InputStream is = ...;
final BitmapFactory.Options options = new
    BitmapFactory.Options();
options.inSampleSize = 2;
Bitmap bmp = BitmapFactory.decodeStream
    (is, null, options);
ImageView imageView = ...;
```

```
//2nd snippet
BitmapFactory.Options options = new Bitmap
    Factory.Options();
while (...){
```

```
//A code recommendation
final BitmapFactory.Options options = new
    BitmapFactory.Options();
options.inSampleSize = 2;
Bitmap bmp = BitmapFactory.decodeStream
    (is, null, options);
```

Other code recommendations are created from other clusters in the same way. Aroma’s algorithm ensures that these recommendations are substantially different from one another, so developers can learn a diverse range of coding patterns.

Results

We instantiated Aroma on a large code corpus of Android GitHub repositories and performed Aroma searches with code snippets chosen

from the 500 most popular Stack Overflow questions with the Android tag. We observed that Aroma provided useful recommendations for a majority of these snippets. Moreover, when we used half of the snippet as the query, Aroma exactly recommended the second half of the code snippet in 37 out of 50 cases.

Developer Feedback

At Facebook, Aroma is integrated into the Visual Studio Code IDE. The developer selects a portion of code to be used as a query, and, in response, Aroma presents a set of code recommendations. From Aroma's feedback workgroups, this integration received mixed feedback: developers were unsure about the use case. Is it a "teacher" to show better code? Is it a warning about potential code duplication? In the end, developers were most interested in seeing examples of API usage. We have since developed a new tool for generating code examples¹³ to address this need.

Bug Fixing

Background

Large code repositories also come with a long history of commits (i.e., code changes), recording how the code base evolved into its current state. If we can find repetitive patterns using ML among these changes, then we can automate the routine work that engineers repetitively do. At Facebook, we have found that one common class of repetitive changes encompasses bug fixes. Therefore, we built a tool called *Getafix*, which learns bug-fixing patterns and automatically offers fix suggestions.

Getafix has goals similar to those of existing automated program repair techniques, but it fills a previously unoccupied spot in the design space: single/

few shot prediction of natural-looking fixes, but for specific kinds of bugs. In contrast to generate-and-validate approaches,¹⁴ we focus on learning patterns from past fixes for specific bug types and leverage information known about bug instances (e.g., blamed variable). Getafix does not attempt to find generic solutions from any sort of ingredient space or by generically mutating the code. It tends to produce actual, human-like fixes by construction, as it takes nothing but past human fixes as inspiration.

How Does It Work?

For clarity, we focus on a specific type of bug that can crash Android apps: Java NPE. The following code snippet shows an example of an NPE and a possible fix:

```
public int getWidth() {
    @Nullable View v = this.getView();
    return v.getWidth(); //Bug: NPE if v is null
    return v != null ? v.getWidth() : 0;
}
```

At Facebook, we use the Infer¹⁵ static analyzer to detect and warn about potential NPEs (the line highlighted in red). From the Infer records, we identify commits that fix the potential NPE (the line highlighted in green). We scrape hundreds of such bug-fixing commits from the version history and use them as training data for Getafix.

Edit Extraction. To find repetitive patterns of bug fixes ("fix patterns") from these training data, Getafix splits commits into fine-grained abstract syntax tree (AST) edits. Getafix first parses each file touched by a commit into a pair of ASTs: one for the source code before the changes made, and another for after the change. Getafix then applies a tree differencing algorithm

similar to GumTree¹⁶ to each pair of ASTs to predict the edits (insertions, deletions, moves, or updates) that likely represent the difference. For the example fix described, Getafix will extract the following edit: `v.getWidth() → v!=null? v.getWidth(): 0`.

Clustering. Getafix takes a data-driven approach, called *antiunification*, by clustering the set of AST edits yielded by the previous step by similarity: it merges the most similar pair of edits in the set into a new edit pattern, abstracting away details only where necessary. An example follows:

```
Edit A:    v.getWidth() → v!=null
? v.getWidth(): 0
Edit B:    lst.size() → lst!=null
? lst.size(): 0
Antiunification: α.β() → α!=null
? α.β(): 0
```

Antiunification has the desirable property of merging edit patterns in the most information-preserving way possible. Getafix repeats this step as often as possible, putting the resulting edit pattern back into the set in place of its constituents, hence reducing the size of the set and allowing edit patterns to be merged and abstracted even further. This process results in a hierarchy of edit patterns, with the original edits as leaf nodes and increasingly abstract edit patterns closer to the roots.

Fix Prediction. With such a hierarchy of fix patterns for NPE warnings, Getafix can automatically fix future warnings: when Infer produces a new, previously unseen NPE warning, Getafix retrieves all patterns that are applicable from our hierarchy of fix patterns. It then applies those candidate patterns to the code, generating candidate fixes, which are ranked statistically using

a metric comparable to TF-IDF. To limit computational cost, one or, at most, a few of the top-ranked candidate patterns are then validated (e.g., by running Infer and making sure the warning disappeared).

The best passing candidate fix is offered to the engineer as a suggestion he or she can accept or reject at the click of a button. Getafix suggests only one fix to limit the cognitive load and provide a straightforward user experience. We do require a final human confirmation since Getafix uses statistical learning and ranking techniques, so there is no formal guarantee of correctness despite certain forms of validation. For more details about Getafix, refer to Bader et al.¹⁷

Results

Of the Infer NPE warnings fixed by Facebook engineers since the Getafix service was rolled out, 42% were fixed by accepting our fix suggestion, and, in 9% of the cases, engineers wrote a semantically identical fix (which goes to show that developers are very particular about the fix suggestions they accept). Note that our pattern-learning phase takes any set of changes as input, so a different scenario we have successfully started automating is the discovery and application of “lint” rules. Changes made in response to code review are often fixes to common antipatterns that were pointed out by a reviewer, and finding and fixing these antipatterns can be baked into a lint rule.

Developer Feedback

We show fix suggestions for warnings during code review and in the IDE wherever possible. We found that warnings that came with a fix suggestion were more actionable and addressed (whether via accepting the

suggested fix or hand-writing one) more often than plain warnings. Individual reactions ranged from ignoring our suggestions to expressing excitement about their level of sophistication in internal feedback groups.

We found that semantic equivalence is insufficient to our engineers and that syntactic differences do matter to them: for instance, we sometimes predict using a ternary conditional and, in several cases, observed developers adopt this fix but negating the condition and swapping the “then” and “else” expressions. At this point, the “accept with one click” experience we provide is ineffective, so we strive to suggest natural-looking fixes exactly as our engineers expect, so syntax and even details like idiomatic white space must be human-like. Our ML-based approach learns patterns that look natural by construction (learned from real fixes) and also learns how to rank among them in a principled way, which would be strenuous to replicate manually.

The Big Picture

We now take a step back to discuss how these ML-based techniques fit

in the broader picture of the software development process. In fact, these techniques have the potential to influence not just writing or fixing code but almost all stages of the software lifecycle.

Figure 3 shows a way to think about modern software development, as organized in three stages, recurring in a cycle (not depicted). The workflow begins with an individual developer’s work, which involves editing the code to implement new functionality, or in response to an issue and making sure the code compiles and passes at least some lightweight quality control (e.g., linters or unit tests). Next is the team stage: once the developer is satisfied with a code change he or she is making, it is sent in for code review, and, perhaps simultaneously, more extensive testing and verification are kicked off. Either of these can require the jumping back into the individual workflow.

Once code gets released and enters production, new issues can arise that were not caught by previous stages. The process must account for how such issues are tracked. The production stage would typically also include

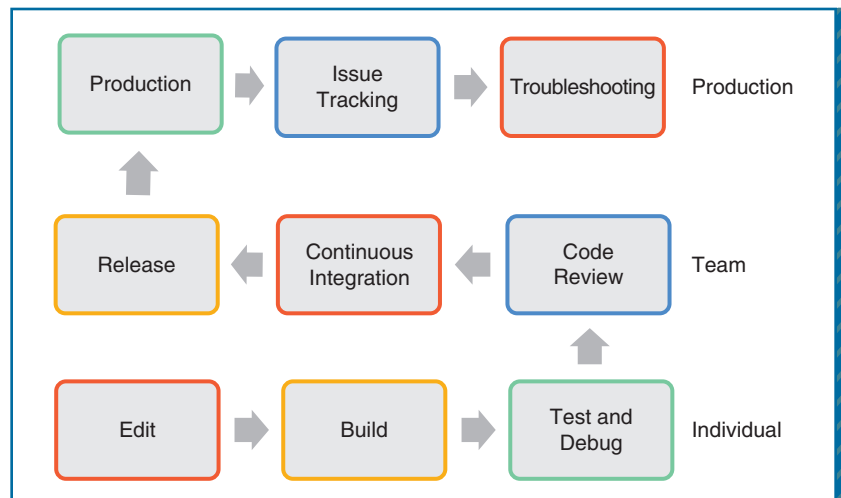



FIGURE 3. The common workflows in software development.



JOHANNES BADER is a software engineer at Facebook, Inc., Menlo Park, California, 94025, USA. His research interests include automatic program repair and programming languages and verification. Bader received his M.S. in computer science from Karlsruhe Institute of Technology in 2016. Further information about him can be found at johannes-bader.com. Contact him at mail@johannes-bader.com.



SONIA SEOHYUN KIM is a software engineer at Facebook, Inc., Menlo Park, California, 94025, USA. Her research interests include applying machine learning to generate code. Kim received her B.S. in applied computational math and minor in computer science from Caltech. Contact her at skim131@fb.com.



FRANK SIFEI LUAN is a software engineer at Facebook, Inc., Menlo Park, California, 94025, USA. His research interests include machine learning and programming languages. He received his bachelor's degrees in computer science and statistics from the University of Chicago in 2017. Further information about him can be found at franklsf.org. Contact him at lsf@berkeley.edu.



SATISH CHANDRA is a software engineer at Facebook, Inc., Menlo Park, California, 94025, USA. His research interests include programming languages and software engineering, including program analysis, type systems, software synthesis, bug finding and repair, software testing and test automation, and, most recently, applications of machine learning to developer tools. Chandra received his Ph.D. in computer science at the University of Wisconsin-Madison. He is an Association for Computing Machinery Distinguished Scientist. Contact him at schandra@acm.org.



ERIK MEIJER is an engineering director at Facebook, Inc., Menlo Park, California, 94025, USA. His research interests include the intersection of programming languages, software engineering, systems, and machine learning. He is an honorary professor of programming language design at the school of Computer Science of the University of Nottingham, member of the *ACM Queue* editorial board, and member of the standing council for Engineering Research Visioning Alliance. Contact him at erikm@fb.com.

some telemetry that helps with bug isolation. Feedback from production kicks off a new cycle, starting again from the individual stage.

The previous sections talked about concepts that are primarily applicable in the code-editing phase of the individual workflow. In addition to those ideas, the most visible developer-facing use of ML in the code-editing phase is autocomplete, which has been widely studied and deployed. More ambitiously, ML techniques can also help developers complete code via program synthesis.

Significant opportunities exist in the other states—for example, we had previously mentioned our own work on predictive regression test selection¹ and triaging crashes²—but a detailed discussion of these is outside the scope of this brief article. Here, in our view, are some of the most promising but relatively untapped opportunities for using ML pertinent to aspects of the team and production states.

- *Code review*: Code review, while widely regarded as essential for maintaining software quality, is also a significant time commitment for software engineers. ML techniques can help automate routine code reviews (such as formatting and best coding practices). More ambitiously, perhaps ML can also automatically resolve a routine code-review comment.
- *Assessing the risk of a code change*: In principle, any code change increases the riskiness of an application. Arguably, the entire testing and verification pipeline exists essentially to reduce this risk. Can we design ML-based techniques that provide a quantitative assessment of the risk of a code change, complementing the usual testing and verification

pipeline? Advances here will impact both testing (by prioritizing tests related to riskier changes) and release management (by carrying out additional quality control for riskier code releases). By comparison, techniques for assessing the impact of a change (e.g., Reb et al.¹⁸) take a binary view of affectedness and, due to the limitations of static analysis, often would be overly pessimistic in their assessment.

- **Troubleshooting:** For widely deployed applications, customers send their feedback implicitly (telemetry or crashes) and, sometimes, explicitly by sending comments. The volume of this feedback can be huge. This is another area where ML can help in multiple ways: not only in triaging these reports, but clustering them to identify common issues, finding important clues from telemetry logs and code changes that could be connected to the issue at hand.

With renewed interest in ML and an emerging uniformity of software development processes (common repositories as well as continuous integration and release), industry is ripe for absorbing these ideas into the mainstream. At Facebook, we certainly are transforming our development process to be as data driven as possible. 🍷

References

1. M. Machalica, A. Samykin, M. Porth, and S. Chandra, “Predictive test selection,” in *Proc. 41st Int. Conf. Softw. Eng. Softw. Eng. Pract. (ICSE-SEIP 19)*, 2019, pp. 91–100. doi: 10.1109/ICSE-SEIP.2019.00018.

2. R. Qian, Y. Yu, W. Park, V. Murali, S. Fink, and S. Chandra, “Debugging crashes using continuous contrast set mining,” in *ICSE-SEIP '20: Proc. ACM/IEEE 42nd Int. Conf. Software Engineering: Software Engineering in Practice*, June 2020, pp. 61–70. <https://doi.org/10.1145/3377813.3381369>.
3. J. Bader, S. Chandra, S. S. Kim, and F. S. Luan, “F8: Using machine learning for developer productivity.” Facebook, 2019. <https://developers.facebook.com/videos/2019/using-machine-learning-for-developer-productivity/>
4. R. Sirres et al., “Augmenting and structuring user queries to support efficient free-form code search,” in *Proc. 40th Int. Conf. Softw. Eng. (ICSE'18)*, 2018, p. 945. doi: 10.1145/3180155.3182513.
5. H. Sajjani, V. Saini, J. Svajlenko, C. K. Roy, and C. V. Lopes, “SourceCC: Scaling code clone detection to bigcode,” in *Proc. 38th Int. Conf. Softw. Eng. (ICSE'16)*, 2016, pp. 1157–1168. doi: 10.1145/2884781.2884877.
6. S. Sachdev, H. Li, S. Luan, S. Kim, K. Seohyun, and S. Chandra, “Retrieval on source code: A neural code search,” in *Proc. 2nd ACM SIGPLAN Int. Workshop on Mach. Learn. Program. Languages*, 2018, pp. 31–41.
7. P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, “Enriching word vectors with subword information,” 2016, arXiv: 1607.04606.
8. T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, “Distributed representations of words and phrases and their compositionality,” 2013, arXiv: 1310.4546.
9. J. Johnson, M. Douze, and H. Jégou, “Billion-scale similarity search with GPUs,” 2017, arXiv:1702.08734.
10. H. Li, S. Kim, and S. Chandra, “Neural code search evaluation dataset,” 2019, arXiv: 1908.09804 [cs.SE].
11. J. Cambronero, H. Li, S. Kim, K. Sen, and S. Chandra, “When deep learning met code search,” in *Proc. 27th ACM Joint Meeting European Softw. Eng. Conf. Symp. Found. Softw. Eng. (ESEC/FSE 2019)*, 2019, pp. 964–974. doi: 10.1145/3338906.3340458.
12. S. Luan, D. Yang, C. Barnaby, K. Sen, and S. Chandra, “Aroma: Code recommendation via structural code search,” in *Proc. ACM Program. Languages*, Oct. 2019, vol. 3, no. OOPSLA, pp. 152:1–152:28. doi: 10.1145/3360578.
13. C. Barnaby, K. Sen, T. Zhang, E. Glassman, and S. Chandra, *Exempla Gratis (E.G.): Code Examples for Free*. New York: Association for Computing Machinery, 2020, pp. 1353–1364.
14. X. B. D. Le, D. Lo, and C. Le Goues. “History Driven Program Repair,” in *Proc. IEEE 23rd Int. Conf. Softw. Anal., Evolution, Reeng. (SANER)*, 2016, vol. 1, pp. 213–224.
15. C. Calcagno et al., “Moving fast with software verification,” in *Proc. NASA Formal Method Symp.*, 2015.
16. J.-R. Falleri, F. Morandat, X. Blanc, M. Martinez, and M. Monperrus, “Fine-grained and accurate source code differencing,” in *Proc. ACM/IEEE Int. Conf. Automat. Softw. Eng. (ASE'14)*, 2014, pp. 313–324. doi: 10.1145/2642937.2642982.
17. J. Bader, A. Scott, M. Pradel, and S. Chandra, “Getafix: Learning to fix bugs automatically,” in *Proc. ACM Program. Languages*, Oct. 2019, vol. 3, no. OOPSLA. doi: 10.1145/3360585.
18. X. Ren, F. Shah, F. Tip, B. G. Ryder, and O. C. Chesley, “Chianti: A tool for change impact analysis of Java programs,” in *Proc. 19th Annu. ACM SIGPLAN Conf. Object-Oriented Programming, Syst., Languages, Appl.*, Oct. 2004, pp. 432–448. doi: 10.1145/1035292.1029012.