

# Dependency-Aware Code Naturalness

CHEN YANG, Tianjin University, China

JUNJIE CHEN\*, Tianjin University, China

JIAJUN JIANG, Tianjin University, China

YULIANG HUANG, Tianjin University, China

Code naturalness, which captures repetitiveness and predictability in programming languages, has proven valuable for various code-related tasks in software engineering. However, precisely measuring code naturalness remains a fundamental challenge. Existing methods measure code naturalness over individual lines of code while ignoring the deep semantic relations among different lines, e.g., program dependency, which may negatively affect the precision of the measure. Despite the intuitive appeal of extending the code naturalness measure to the code dependency domain (as there are some work that have initiated the utilization of code dependency for diverse code-related tasks), this assumption remains unexplored and warrants direct investigation. In this study, we aim to perform the first empirical study to investigate whether incorporating code dependency, instead of analyzing individual lines, can enhance the precision of measuring code naturalness.

To achieve that, we first propose a new method named DAN for measuring code naturalness by incorporating the rich dependency information in the code. Specifically, DAN extracts multiple sequences of code lines by traversing the program dependency graph, where different code lines are connected by dependencies in each sequence, and then the code naturalness will be measured by taking each sequence as a whole. In this way, the dependency information can be well captured. Finally, we have conducted an extensive study to evaluate the influence of code dependency for measuring code naturalness with DAN, and compared it with the state-of-the-art methods under three emerging application scenarios of code naturalness. The results demonstrate that DAN can not only better distinguish natural and unnatural code, but also substantially boost two important downstream applications of code naturalness, i.e., distinguishing buggy and non-buggy code lines and data cleansing for training better code models, reflecting the significance of code dependency in measuring code naturalness.

CCS Concepts: • **General and reference** → **Measurement**; • **Theory of computation** → **Program analysis**; • **Software and its engineering** → **Maintaining software**.

Additional Key Words and Phrases: Naturalness, Program Dependency, Code Entropy

## ACM Reference Format:

Chen Yang, Junjie Chen, Jiajun Jiang, and Yuliang Huang. 2024. Dependency-Aware Code Naturalness. *Proc. ACM Program. Lang.* 8, OOPSLA2, Article 354 (October 2024), 23 pages. <https://doi.org/10.1145/3689794>

## 1 Introduction

Natural languages (NL) exhibit notable repetitiveness and predictability, which facilitates reliable and efficient communication among humans [Pilkington 1996; Ray et al. 2016]. This property,

\*Corresponding author

---

Authors' Contact Information: [Chen Yang](mailto:yangchenyc@tju.edu.cn), College of Intelligence and Computing, Tianjin University, Tianjin, China, yangchenyc@tju.edu.cn; [Junjie Chen](mailto:junjiechen@tju.edu.cn), College of Intelligence and Computing, Tianjin University, Tianjin, China, junjiechen@tju.edu.cn; [Jiajun Jiang](mailto:jiangjiajun@tju.edu.cn), College of Intelligence and Computing, Tianjin University, Tianjin, China, jiangjiajun@tju.edu.cn; [Yuliang Huang](mailto:huangyuliang@tju.edu.cn), College of Intelligence and Computing, Tianjin University, Tianjin, China, huangyuliang@tju.edu.cn.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2024 Copyright held by the owner/author(s).

ACM 2475-1421/2024/10-ART354

<https://doi.org/10.1145/3689794>

commonly referred to as *naturalness*, has been effectively harnessed in the area of Natural Language Processing (NLP) [Cavnar et al. 1994; Kowsari et al. 2019]. Inspired by it, programming languages, as a special form of human-oriented languages, have been also demonstrated repetitive and predictable [Hindle et al. 2012; Jiang et al. 2022]. This property of code has been successfully exploited to facilitate various code-related tasks, such as code generation [Chakraborty et al. 2022], code completion [Raychev et al. 2014], and code de-obfuscation [Raychev et al. 2015]. It also constitutes the foundation for the rapid development of statistical language modeling of source code [Wang et al. 2021a; Zheng et al. 2023].

Due to the importance of this property, how to measure the naturalness of code is one of the most fundamental tasks in this area. In the literature, some methods of measuring code naturalness have been proposed [Hindle et al. 2012; Jiang et al. 2022; Ray et al. 2016; Tu et al. 2014]. For example, Hindle et al. [Hindle et al. 2012] applied the n-gram model to measure cross entropy of source code as code naturalness. As source code tends to be more repetitive within a small module, Tu et al. [Tu et al. 2014] further incorporated a cache model to exploit localness of source code for measuring code naturalness more precisely. Then, Ray et al. [Ray et al. 2016] proposed to normalize code naturalness according to statement types in order to balance the naturalness of different syntax types. Further, to relieve the limitations of the n-gram model (e.g., trained on small datasets and thus inapplicable to cross-project scenarios), Khanfir et al. [Khanfir et al. 2022] proposed CodeBERT-NT to use the pre-trained language model, i.e., CodeBERT [Feng et al. 2020] instead of n-gram, to measure code naturalness based on the prediction confidence on masked tokens in code.

Although code naturalness can be measured more precisely with more methods being proposed, almost all of them share the same workflow, i.e., measuring naturalness of each line in a complete code snippet (e.g., a method or a file) individually. However, humans and compilers tend not to understand a single line without considering sufficient contexts, especially the lines of code that are depended on [Busjahn et al. 2015; Rahman et al. 2019]. Hence, such a way isolating each code line could negatively affect the measure of code naturalness.

In this work, we aim to perform the first empirical study to investigate whether incorporating code dependency, instead of isolating each line, can improve the precision of measuring code naturalness. Intuitively, considering all lines in the complete code snippet together in a sequential manner can be a straightforward method to include all dependency information within the code snippet. However, such a coarse-grained strategy may incur much noise because sequential lines in code may be independent with each other, which has been confirmed in our study (Section 4). Indeed, humans and compilers do not process code in a completely sequential manner [Busjahn et al. 2015; Rahman et al. 2019].

To investigate the influence of code dependency well, we first propose a more reasonable way of incorporating dependency information among lines in code for measuring code naturalness. We call it *dependency-aware code naturalness* (abbreviated as DAN). Specifically, DAN represents the complete code as a program dependency graph (PDG) by performing control- and data- flow analysis. Then, it traverses all paths in the graph in a DFS (depth first search) manner, each of which can be treated as a sequence of lines with control or data dependency. DAN measures the naturalness of each sequence with some statistical model, which is helpful to incorporate dependency information without incurring noise, and then obtains the naturalness of the complete code snippet by aggregating the naturalness of all these sequences. Note that our novelty does not lie in designing new statistical models for measuring code naturalness. In fact, the idea of DAN is general and can be applied independently of the used statistical model. To demonstrate the generalizability of our DAN idea, we create two instantiations of DAN by integrating the two state-of-the-art statistical models that have been used to measure code naturalness (i.e., the

cached n-gram model and the pre-trained CodeBERT model) for evaluation, respectively. DAN can also integrate more advanced statistical models for better performance (even though they were not investigated to measure code naturalness), which will be discussed in Section 5.2. For ease of presentation, we call them  $\text{DAN}_{n\text{gram}}$  and  $\text{DAN}_{\text{codebert}}$ , respectively.

Recent work has initiated the utilization of code dependency for diverse code-related tasks [Guo et al. 2020; Suo et al. 2024; Yang et al. 2024b], exemplified by GraphCodeBERT [Guo et al. 2020], which leverages program analysis to augment vectorization of code snippets. Despite the intuitive appeal of extending the naturalness measure to the code dependency domain, this assumption remains unexplored and warrants direct investigation. Therefore, despite conceptually simple, DAN is the first attempt to incorporate code dependency in measuring code naturalness. In this way, code semantic information can be effectively utilized to measure code naturalness to some degree.

With DAN, we conducted the extensive study to investigate the influence of code dependency on measuring code naturalness compared to the state-of-the-art n-gram-based method [Ray et al. 2016] (called Ngram-NT in our work) and CodeBERT-NT [Khanfir et al. 2022], both of which work on a per-line basis for measuring code naturalness. Specifically, we performed comparisons ( $\text{DAN}_{n\text{gram}}$  v.s. Ngram-NT and  $\text{DAN}_{\text{codebert}}$  v.s. CodeBERT-NT, each pair of methods used the same statistical model for investigating the contribution of code dependency more clearly) in three scenarios by carefully designing corresponding experiments. First of all, we investigated the ability of DAN to distinguish natural and unnatural code, which is the core task of measuring code naturalness, by intentionally transforming each piece of natural code in the widely-used HumanEval-X dataset [Zheng et al. 2023] into unnatural code based on a set of transformation rules. The results show that DAN can distinguish natural code and unnatural code better than Ngram-NT/CodeBERT-NT by achieving 41.82% and 13.41 times more difference in the measured naturalness between them on average.

We then evaluated DAN in two downstream applications of code naturalness: (1) distinguishing buggy and non-buggy code lines following the existing work [Jiang et al. 2022; Ray et al. 2016; Yan et al. 2020], and (2) cleansing training data for building better code generation models, which can generate more natural code without damaging the functionality correctness of generated code, as indicated by the existing work [Chakraborty et al. 2022; Zan et al. 2022b]. For the former, we used all the three datasets that have been used in this application, i.e., Defects4J [Just et al. 2014], GrowingBugs [gro 2023], SmartSHARK [Herbold et al. 2020]. The results show that DAN helps identify 26.67%, 14.81%, 28.48% more buggy lines within the top20% prioritized lines than Ngram-NT on the three datasets respectively, and helps identify 32.00%, 14.29%, and 40.24% more buggy lines than CodeBERT-NT accordingly. For the latter, we used two widely-studied code generation models (i.e., CodeGen-Multi [Nijkamp et al. 2022] and GPT-2 [Radford et al. 2019]) as the target models due to the model availability and fine-tuning cost, and used the APPS [Hendrycks et al. 2021] and HumanEval-X [Zheng et al. 2023] datasets since they have been widely used in the area of code generation. The results show that the training data selected by DAN builds better code generation models than the entire training set and the data selected by Ngram-NT and CodeBERT-NT respectively. Specifically, the code generation models with DAN achieve 18.08% and 16.25% higher CodeBLEU scores than those with Ngram-NT and CodeBERT-NT on average. The results demonstrate that the former can generate code that has high textual similarity with the ground-truth (human-written) code, indicating the high readability of generated code to some degree. In particular, the former even achieves higher code generation accuracy (measured through test execution on generated code) than the latter two. Overall, the results demonstrate the contribution of incorporating code dependency to improve the precision of measuring code naturalness as well as boost the two important downstream applications.

To sum up, our work makes the following major contributions:

- We are the first to empirically investigate the influence of code dependency on measuring code naturalness, compared to the state-of-the-art practice of isolating each line in code.
- We design dependency-aware code naturalness (i.e., DAN), a simple but effective method to measure code naturalness by incorporating code dependency among lines.
- We conducted the extensive study with DAN in three scenarios, including the core task of distinguishing natural and unnatural code and two downstream applications of code naturalness, demonstrating its significant contribution in all the three scenarios.

## 2 Background

### 2.1 Statistical Language Models

Language modeling fundamentally involves the construction of probability distributions over sequences of words or tokens. Typically, the conditional probability  $P(t_i|t_1, \dots, t_{i-1})$  denotes the likelihood that the token  $t_i$  will follow the sequence of preceding tokens  $h = t_1, \dots, t_{i-1}$ . However, accounting for entire sequences as context can become computationally prohibitive due to the possible vastness of sequence variation and length. Hence, the n-gram model simplifies it by considering only the most recent  $n - 1$  tokens, adhering to a Markovian assumption:

$$P_{ngram}(t_i|h) = P(t_i|t_{i-n+1}, \dots, t_{i-1}) \quad (1)$$

The n-gram model has been effectively used to capture highly repetitive regularities in both natural and programming languages, contributing significantly to various code-related tasks such as code completion [Tu et al. 2014], bug detection [Yan et al. 2020], and naturalness measure [Ray et al. 2016]. Subsequently, many pre-trained language models (e.g., CodeBERT [Feng et al. 2020]) are proposed, expanding on the capabilities of language modeling by accepting fixed-length token sequences. They leverage vast data and provide a user-friendly, out-of-the-box solution. Their robust learning mechanisms enable them to capture intricate code patterns, which are beneficial for a variety of tasks.

### 2.2 Code Naturalness

As introduced in Section 1, code naturalness is an important property of source code. Therefore, studying the accurate measure of code naturalness is critical. Hindle et al. [Hindle et al. 2012] were the first to measure code naturalness based on the n-gram model, calculating the cross-entropy for a line of code  $S = t_1 t_2 \dots t_N$  as follows:

$$H_{ngram}(S) = -\frac{1}{N} \sum_{i=1}^N \log P_{ngram}(t_i|h) \quad (2)$$

Here,  $h$  represents the token sequence before  $t_i$  in  $S$  (i.e.,  $h = t_1, \dots, t_{i-1}$ ).

Then, Tu et al. [Tu et al. 2014] improved the code naturalness measure by exploiting code localness and proposed the cached n-gram model:

$$P_{cache}(t_i|h) = \lambda \cdot P_{global}(t_i|h) + (1 - \lambda) \cdot P_{local}(t_i|h) \quad (3)$$

In Formula 3,  $P_{global}$  represents the conditional probability computed by the traditional n-gram model, which is trained with a large-scale corpus of source code, such as one project, whereas  $P_{local}$  represents the *n-gram* model that is trained over small-scale source code, such as the context of a single line in a file.  $\lambda$  is the weight reflecting both the commonalities among source code in the large-scale corpus and the particularity under certain contexts. The subsequent code naturalness calculation shares the same workflow as the one proposed by Hindle et al. [Hindle et al. 2012]. Later,

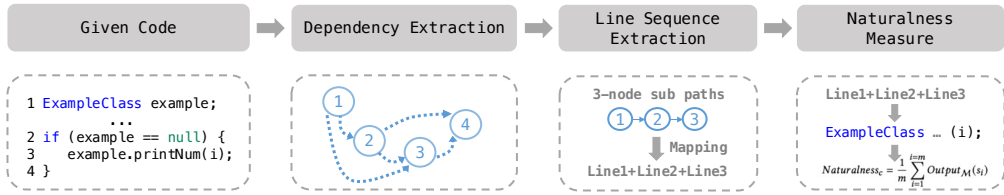


Fig. 1. Overview of DAN

Ray et al. [Ray et al. 2016] optimized its line naturalness measure by normalizing the computation for different types of code lines.

To overcome the limitation of n-gram models (e.g., trained on small datasets and thus inapplicable to cross-project scenarios), Khanfir et al. [Khanfir et al. 2022] proposed CodeBERT-NT by using the pre-trained CodeBERT model to measure code naturalness based on token predictability. CodeBERT-NT masks some tokens in each individual line and uses the minimal prediction confidence of CodeBERT on the masked tokens as the naturalness of this line.

To sum up, all these methods measure code naturalness by applying the used statistical language model to measure the naturalness of each line individually. However, they ignore essential code dependency information and thus may lead to inaccurate measure of code naturalness, since the dependency information is crucial for code comprehension as discussed in Section 1.

In what follows, we will use a simple example to illustrate this issue. Considering the code snippet shown in Figure 1 (simplified for clarity), the code is regarded as natural when each line of code is measured individually by existing methods since each code line is common and widely-used in the code snippet. However, after putting all the code lines together, it is indeed unnatural because the object `example` at line 3 will be de-referenced when it is `null`. Considering all the lines sequentially may not capture such information well since many lines between lines 1 and 2 (also after line 4) may involve too much noise and thus weaken the effect of code dependency. In such a case, incorporating accurate code dependency for measuring code naturalness may be helpful. In this work, we aim to conduct the first empirical exploration to investigate the influence of code dependency on measuring code naturalness by designing an initial (but reasonable) solution.

### 3 Methodology

In this work, we conduct the first empirical study to investigate the influence of code dependency on code naturalness. To complete this study, we propose the first method of incorporating code dependency into measuring code naturalness, i.e., dependency-aware code naturalness (DAN). In fact, there may be some other methods of utilizing code dependency, but our main contribution lies in the first exploration in this direction even though the idea of DAN is simple. If we can obtain positive conclusions from the study with DAN, designing more advanced methods of incorporating code dependency can be promising future work. In this section, we introduce DAN in detail.

In general, DAN consists of three steps as shown in Figure 1. (1) It represents the given code as a program dependency graph (PDG) by performing control- and data- flow analysis, which can help understand the semantics of code. (2) It traverses all paths in the PDG, each of which can be treated as a sequence of lines with control or data dependency. It can incorporate code dependency without much noise compared to treating all lines in code as a sequence for measuring code naturalness. (3) It applies the widely-used statistical model (i.e., n-gram or CodeBERT) to each sequence, instead of each individual line, for measuring code naturalness. With DAN, code naturalness can be measured by incorporating code semantics with the aid of code dependency information, which can intuitively




Fig. 2. An example of program dependency graph

improve the precision of measuring code naturalness. Next, we present dependency extraction in Section 3.1, line sequence extraction in Section 3.2, and naturalness measure in Section 3.3.

### 3.1 Dependency Extraction

To extract the dependency information for a given code snippet, DAN transforms the code into a PDG. The PDG serves as a representation of a program’s control and data flow dependencies among program elements. The PDG is represented as a directed graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where:

- $\mathcal{V}$  is the set of nodes in the graph, and each node corresponds to a program element, e.g., statement or function.
- $\mathcal{E}$  is the set of directed edges between nodes in  $\mathcal{V}$ , representing the dependencies between program elements.

Specifically, DAN constructs the PDG at the granularity of program statements. That is, each node  $v$  in the graph corresponds to a statement in the source code, and the edges in the graph depict the dependencies between these statements. For instance, considering the code displayed in Figure 1, the PDG constructed at the statement granularity is shown in Figure 2. For DAN, there are two types of edges in the PDG:

- $\mathcal{E}_c$ : This type of dependency represents control flow relationship between statements. For example, if a statement A depends on another statement B in a way that the execution of A depends on the outcome of B, then a control dependency edge is present from B to A in the PDG. For example, the edge between node  $v_2$  and node  $v_3$  in Figure 2 represents the control dependency between them.
- $\mathcal{E}_d$ : This type of dependency represents data flow relationship between statements. If a variable’s value is used in another statement, a data dependency edge is established from the statement using the variable to the one defining its value. For example, the edge between node  $v_1$  and node  $v_3$  represents the data dependency of the object `example`.

During the PDG construction process, DAN also analyzes the mapping between the nodes in the graph and the lines in source code, serving as the foundation for subsequent naturalness measure. For instance, node  $v_1$  in the PDG is mapped back to the object declaration at Line 1 of the source code in Figure 1. However, it is essential to note that not all nodes in the PDG have corresponding source lines. For example, node  $v_4$  in the PDG does not correspond to any lines in the source code. Also, multiple nodes in the graph sometimes correspond to the same line, such as a line of a for statement that is represented by multiple nodes in the PDG.

### 3.2 Line Sequence Extraction

Each path in the PDG represents a sequence of nodes that have dependencies. To incorporate dependency information for measuring code naturalness, DAN extracts paths from the PDG. The extraction process is performed using a Depth-First Search (DFS) strategy, which starts from a designated starting node and explores each path throughout the graph systematically. When DAN reaches a node with no successors or the node representing the termination of code execution (i.e., the exit node), this path represents a complete path. To handle cyclic paths, DAN maintains a set of visited nodes, ensuring that previously visited nodes are not revisited during the traversal process. In this way, DAN obtains a list of complete paths. For instance, DAN extracts three complete paths from the PDG of the code in Figure 1, each of which is highlighted in different colors as shown in Figure 2. Note that DAN filters out the paths that are subsumed by any other paths.

However, complete paths are typically lengthy and intricate, and thus there are usually a lot of overlapped nodes between complete paths. That results in redundant use of the dependencies between these nodes, which could incur extra computation overhead and affect the precision of code naturalness measure. Inspired by the existing work [Rahman et al. 2019] that highlighted the repetition of  $n$ -node sub-graphs of source code, DAN addresses this problem by extracting  $n$ -node sub-paths from each complete path. For example, given  $n=3$ , four distinct sub-paths can be extracted from the three complete paths in the graph shown in Figure 2, i.e.,  $\{v_1 - v_2 - v_3, v_2 - v_3 - v_4, v_1 - v_3 - v_4, v_1 - v_2 - v_4\}$ . Each sub-path represents a shorter sequence of nodes with dependencies.

To address the difficulty of using statistical models to measure the naturalness of PDG nodes, DAN further leverages the mapping obtained in the first step to convert each sub-path into a sequence of corresponding code lines. Similarly, if a sequence is a subset or duplicate of any other sequences, DAN filters it out. For instance, if there are three extracted sequences: Line1-Line2-Line3, Line2-Line3, and Line2-Line3, DAN retains only the sequence Line1-Line2-Line3. Since multiple nodes may correspond to the same code line, there may be the cases where the same line continuously appears in a sequence (e.g., Line3-Line5-Line5-Line5). Actually, this line has encompassed the dependencies of all its corresponding nodes, and thus DAN retains only one occurrence of this line in the sequence to avoid redundancy (e.g., the above example sequence is reduced as Line3-Line5).

This process results in a refined set of sequences, each of which consists of code lines with data or control dependencies. In this way, DAN can incorporate code dependency into naturalness measure by applying the widely-used statistical model to each sequence of code lines.

### 3.3 Naturalness Measure

After obtaining the set of sequences of code lines, DAN follows the similar workflow of the existing methods [Khanfir et al. 2022; Ray et al. 2016] to measure the naturalness of code. To make our method description self-contained, we also introduce the workflow briefly in this section. When measuring the naturalness of a given code snippet, DAN treats each sequence of code lines with dependencies as a unit for naturalness measure. That is, DAN applies a statistical model to each sequence of code lines instead of each code line individually, and then aggregates the naturalness of all sequences to obtain the naturalness of the given code snippet.

Specifically, given that DAN reorganizes the given code snippet into  $m$  sequences of code lines:  $c = \{s_1, s_2, \dots, s_m\}$ , where  $s_i$  represents a sequence of code lines with data or control dependencies. Then, DAN measures the naturalness of the code as follows:

$$Naturalness_c = \frac{1}{m} \sum_{i=1}^{i=m} Output_{\mathcal{M}}(s_i) \quad (4)$$

Here,  $\mathcal{M}$  represents the statistical model used to measure code naturalness and  $Output_{\mathcal{M}}(s_i)$  represents the naturalness of  $s_i$  measured by  $\mathcal{M}$ . Note that the idea of DAN is general and can be applied independently of the used statistical model. Therefore, we can use any statistical model in this step of DAN. To demonstrate the generality of our DAN idea, we create two instantiations of DAN by integrating the two widely-used statistical models in measuring code naturalness (i.e., the cached n-gram model and the pre-trained CodeBERT model) for evaluation, respectively. As described in Section 1, we call them  $DAN_{ngram}$  and  $DAN_{codebert}$ , respectively. The details of using the cached n-gram and pre-trained CodeBERT models to measure code naturalness have been introduced in Section 2. As it is infeasible to determine one syntax type for a sequence of multiple code lines,  $DAN_{ngram}$  does not perform syntax-based normalization like the state-of-the-art n-gram-based method [Ray et al. 2016].

## 4 Evaluation

With DAN, we conducted the first empirical study to investigate whether incorporating code dependency can improve the measure of code naturalness. We evaluated DAN (i.e., its two instantiations  $DAN_{ngram}$  and  $DAN_{codebert}$ ) in three scenarios for code naturalness, including its core task (i.e., distinguishing natural and unnatural code) - **RQ1**, and two applications of code naturalness: distinguishing buggy and non-buggy code lines (called buggy line prioritization) - **RQ2**, and cleansing data for training code generation models (called training data cleansing) - **RQ3**.

• **Compared Techniques.** We chose the **Ngram-NT** [Ray et al. 2016] and **CodeBERT-NT** [Khanfir et al. 2022] methods for comparison in our study, which have been introduced in Section 2. To investigate the contribution of incorporating code dependency to measure code naturalness clearly, we compared  $DAN_{ngram}/DAN_{codebert}$  with Ngram-NT/CodeBERT-NT respectively, by controlling the influence of the used statistical model. In particular, it is helpful to demonstrate the generality of our DAN idea across different statistical models.

As mentioned in Section 1, considering all code lines together in a sequential manner is an intuitive method to include all dependency information within the given code, but can incur much noise in code dependency. We regarded it as a compared technique in the study, which can help investigate the contribution of our code dependency extraction strategy in DAN. For ease of presentation, we call it **CSN** (short for complete sequence code naturalness). Due to the input length limitation in CodeBERT, we cannot use CodeBERT as the statistical model in CSN, and thus we just constructed the instantiation of CSN by using the cached n-gram model as the statistical model (called **CSN<sub>ngram</sub>**).

Besides, DAN extracts n-node sub-paths from each complete path to relieve the lengthiness and intricacy of complete paths. CSN may also suffer from this issue. Hence, we improve CSN by splitting the complete sequence with all code lines used in it into a set of n-line subsequences, then measuring naturalness of each n-line subsequence, and finally aggregating the naturalness of all subsequences, for more sufficient comparison. We call this variant of **CSN-sub**. According to the used statistical model, we created two instantiations, i.e., **CSN-sub<sub>ngram</sub>** and **CSN-sub<sub>codebert</sub>**. We set n to 3 by default following the existing study [Rahman et al. 2019], which demonstrated the repetitiveness of 3-node sub-graphs of code, and extracted all sub-paths with length 3 from complete paths. However, we also investigated the influence of n by setting n to 2 in DAN.

Overall, we had two sets of comparisons:  $DAN_{ngram}$  vs.  $DAN_{ngram}(n=2)$  vs. Ngram-NT v.s.  $CSN_{ngram}$  v.s.  $CSN-sub_{ngram}$  and  $DAN_{codebert}$  vs.  $DAN_{codebert}(n=2)$  vs. CodeBERT-NT v.s.  $CSN-sub_{codebert}$ , so as to evaluate the importance of incorporating code dependency in measuring code naturalness regardless of the used statistical models. As fine-tuning code models is resource-intensive and



time-consuming, we only compared  $DAN_{ngram}$  vs. Ngram-NT and  $DAN_{codebert}$  vs. CodeBERT-NT for the third scenario (i.e., training data cleansing).

• **Implementations.** We implemented DAN in Java and Python. We implemented PDG construction using Soot [Vallée-Rai et al. 1999] (a widely-used program analysis tool), and constructed the mapping between nodes and code lines based on the debugging information provided by Soot (i.e., line number tables and local variable tables). We implemented the cached n-gram model based on nltk 3.8.1 [Loper and Bird 2002], javalang 0.13.0 [jav 2023], and scikit-learn 1.2.2 [skl 2023]. For Ngram-NT, we replicated the experimental results in the original paper [Jiang et al. 2022] based on our re-implementation, which validates the correctness of our re-implementation to a large extent. For CodeBERT-NT, we directly used their released tool [Khanfir et al. 2022].

Although the idea of DAN is general, we implemented and evaluated it on Java projects currently, which may not represent the subjects under other programming languages. This is a potential threat in our study, and we discussed it in detail in Section 5.4. Our decision to first focus on Java projects was influenced by the high-impact study conducted by Ray et al. [Ray et al. 2016] and Khanfir et al. [Khanfir et al. 2022], which were also carried out on Java projects. This choice allows for more direct comparison between our study and the existing work.

We conducted all the experiments on a server with Intel(R) Xeon(R) Silver 4214 @ 2.20GHz CPU, 256GB memory, and Ubuntu 18.04 operating system.

#### 4.1 RQ1: Effectiveness Comparison on Distinguishing Natural and Unnatural Code

**4.1.1 Dataset.** To answer RQ1, it is essential to create a collection of both natural code and unnatural code. Here, we treated the code in the HumanEval-X dataset [Zheng et al. 2023], which extends the widely-used HumanEval dataset [Chen et al. 2021b] to multiple languages (including the Java version we focused on), as natural code. This dataset consists of 164 hand-written programming problems and the corresponding answers, which has been widely used as the reference to evaluate code quality [Athiwaratkun et al. 2022; Nijkamp et al. 2022; Touvron et al. 2023; Zheng et al. 2023]. In HumanEval-X, 34.99% of statements are dependent on other statements and each dependency sequence involves 7 statements on average. The code in HumanEval-X is meticulously crafted by experienced developers and has been recognized as high-quality code [Athiwaratkun et al. 2022; Zheng et al. 2023]. Hence, this dataset is suitable as natural code in this experiment.

We then constructed the corresponding unnatural code from each piece of natural code in HumanEval-X. Specifically, we implemented a set of semantic-preserving transformation rules for de-naturalizing code following the existing work [Chakraborty et al. 2022; Patra and Pradel 2021]. By applying them to a piece of natural code, we can obtain a set of unnatural but equivalent code. The reasons for obtaining unnatural code in this way are two-fold: (1) There is no existing open-source dataset of unnatural code to our best knowledge; (2) Such a way can control the difference between each pair of natural code and unnatural code, which can help evaluate the ability of distinguishing natural code and unnatural code more purely by keeping the remaining part of code the same (except the de-naturalized part of code).

In total, we implemented three equivalent transformation rules to de-naturalize code. The general idea of these rules is to transform frequently-encountered natural code that developers can typically write, into the synthetic unnatural form. The existing studies [Casalnuovo et al. 2020a,b] have demonstrated that the de-naturalized code with these transformation rules is more challenging to read and understand. We used an example shown in Figure 3 to facilitate illustrating the three rules in the following:

**Dead Code Insertion:** It inserts a block of dead code that will never be executed at random positions in the original code. To ensure the property of dead code, the inserted block is a loop or

```

1 public int getMax(int[] arr)
2 {
3     int len = arr.length;
4     int max = arr[0];
5     int i = 0;
6     while (i < len) {
7         if (arr[i] > max) {
8             max = arr[i];
9         }
10        i++;
11    }
12    return max;
13 }
14
15
16

```

(a) Original code

```

public int getMax(int[] arr)
{
    int len = arr.length;
    int max = arr[0];
    int i = 0;
    while (i < len) {
        while (i != i) {
            max = arr[i];
        }
        if (arr[i] > max) {
            max = arr[i];
        }
        i++;
    }
    return max;
}

```

(b) Dead Code Insertion

```

1 public int getMax(int[] arr)
2 {
3     int len = arr.length;
4     int max = arr[0];
5     int i = 0;
6     while (i < len) {
7         if (arr[i] > max) {
8             max = arr[i];
9         }
10        for (int j=0; j<len;) {
11            j += 1;
12        }
13        i++;
14    }
15    return max;
16 }

```

(c) Confusing Code Insertion

```

public int getMax(int[] VAL_1)
{
    int len = VAL_1.length;
    int max = VAL_1[0];
    int i = 0;
    while (i < len) {
        if (VAL_1[i] > max) {
            max = VAL_1[i];
        }
        i++;
    }
    return max;
}

```

(d) Variable Renaming

Fig. 3. Examples of de-naturalization

branch block, which is guarded by an unsatisfiable condition (e.g., a variable is not equal to itself). The content in the branch or loop body is the statements transplanted from the original code. As shown in Figure 3b, a while loop with an identically-false condition is inserted at Lines 7-9.

**Confusing Code Insertion:** It inserts a block of code, which can be executed without altering the original functionality, at random positions in the natural code. Specifically, the inserted block is a loop that can be executed several times and the content in the loop body is a statement with increment or decrement operation on a variable. For example, it inserts a for loop that can be executed without altering the original functionality at Lines 10-12 as shown in Figure 3c.

**Variable Renaming:** It renames variables to VAR\_i (an unnatural name). When renaming a variable, we analyze the data-flow of this variable and rename all the occurrences of this variable in the original code. For example, we rename the variable arr in Figure 3a to VAL\_1 in Figure 3d.

For each piece of natural code in HumanEval-X, we applied a transformation rule to create four distinct unnatural versions, each exhibiting different degrees of de-naturalization. They are produced based on  $k$ -order ( $k=\{1,2,3,4\}$ ) transformation with a specific transformation rule (applying the transformation rule  $k$  times to the piece of natural code), respectively. Consequently, we crafted four sets of unnatural code per transformation rule from HumanEval-X, accumulating 12 sets across

Table 1. Effectiveness on distinguishing natural and unnatural code

Method	Dead Code Insertion				Confusing Code Insertion				Variable Renaming			
	1	2	3	4	1	2	3	4	1	2	3	4
DAN <sub>ngram</sub>	<b>12.17%</b>	<b>19.08%</b>	<b>23.53%</b>	<b>26.26%</b>	<b>24.86%</b>	<b>38.76%</b>	<b>46.88%</b>	53.24%	<b>17.70%</b>	30.12%	<b>42.23%</b>	<b>48.45%</b>
DAN <sub>ngram</sub> (n=2)	11.30%	18.01%	22.50%	25.29%	24.68%	38.11%	46.31%	<b>53.72%</b>	17.65%	<b>30.21%</b>	42.16%	48.05%
Ngram-NT	7.96%	14.05%	19.95%	22.82%	17.96%	29.46%	37.60%	42.97%	9.13%	19.09%	25.12%	31.72%
CSN <sub>ngram</sub>	8.58%	10.66%	15.39%	15.80%	19.88%	31.05%	37.85%	43.40%	14.18%	20.29%	32.85%	35.96%
CSN-sub <sub>ngram</sub>	10.06%	16.66%	20.39%	25.33%	17.44%	25.87%	34.16%	38.78%	6.77%	14.30%	18.78%	26.13%
DAN <sub>codebert</sub>	<b>9.46%</b>	10.11%	<b>13.20%</b>	13.75%	<b>19.37%</b>	<b>30.41%</b>	<b>32.43%</b>	<b>32.68%</b>	<b>14.83%</b>	<b>26.28%</b>	<b>35.65%</b>	<b>37.87%</b>
DAN <sub>codebert</sub> (n=2)	8.90%	9.62%	11.15%	12.87%	8.48%	15.51%	18.28%	16.05%	10.72%	17.56%	32.75%	36.85%
CodeBERT-NT	7.90%	5.25%	7.61%	7.06%	2.02%	-0.25%	0.35%	0.74%	10.59%	15.70%	29.49%	35.91%
CSN-sub <sub>codebert</sub>	9.35%	<b>13.50%</b>	10.79%	<b>14.66%</b>	10.33%	8.26%	6.59%	5.91%	10.28%	18.43%	28.37%	28.96%

three distinct rules. This comprehensive collection of data is instrumental in assessing how well different methods can distinguish natural code and its unnatural counterparts with different degrees of de-naturalization. Note that we did not use loop transformation and block/operation-swap transformation in the existing work [Chakraborty et al. 2022], since they cannot guarantee the transformed code is really unnatural (e.g., changing  $a==b$  to  $b==a$ ). Regarding these two rules, code naturalness and readability can be retained to a large extent since they retain structural information to a large extent and all textual semantics without changing identifiers.

**4.1.2 Metrics.** To evaluate the ability of distinguishing natural code and unnatural code for each method of measuring code naturalness, we first measured the naturalness of each piece of natural code and the corresponding unnatural code transformed from it in each set of unnatural code. For ease of presentation, we denote the naturalness of the original natural code as  $Naturalness_n$  and denote that of the corresponding unnatural one as  $Naturalness_u$ . Then, we measured the normalized difference in naturalness between each pair of natural and unnatural code:  $diff = \frac{Naturalness_u - Naturalness_n}{Naturalness_n}$ . The difference provides insights into how distinctly a method can capture changes in code naturalness, reflecting its sensitivity to a large extent. Also, larger differences provide higher confidence to trust measure results to some degree. If a piece of code becomes unnatural after (slight) code changes during evolution, the measure with higher sensitivity is more effective to identify it.

**4.1.3 Results and Analysis.** Table 1 shows the comparison effectiveness of DAN<sub>ngram</sub> v.s. Ngram-NT and DAN<sub>codebert</sub> v.s. CodeBERT-NT in distinguishing natural code and unnatural code produced by different orders of transformation with different rules. For each set of unnatural code, we presented the average normalized difference values and put all the detailed results at our project homepage. From Table 1, we found that DAN<sub>ngram</sub> outperforms Ngram-NT, and DAN<sub>codebert</sub> outperforms CodeBERT-NT across different degrees of unnatural code produced by different transformation rules. Specifically, DAN<sub>ngram</sub> improves upon Ngram-NT by an average of 41.82%, while DAN<sub>codebert</sub> improves upon CodeBERT-NT by an average of 13.41 times, in terms of average normalized difference values across all the 12 sets of unnatural code constructed from the set of natural code. In particular, CodeBERT-NT even incorrectly identifies unnatural code as being more natural in one case (i.e., the average normalized difference value is negative), whereas DAN<sub>codebert</sub> avoids such misjudgments. The results demonstrate the superiority of DAN in distinguishing natural and unnatural code regardless of the used statistical models. In particular, we performed a paired sample Wilcoxon signed-rank test [Woolson 2007] at the significant level of 0.05 to further confirm the superiority of DAN over the state-of-the-art counterparts in statistics by obtaining the p-values

all smaller than 0.05. The effect size between  $DAN_{ngram}/DAN_{codebert}$  and  $Ngram-NT/CodeBERT-NT$  is 0.36 and 0.65.

Moreover, the results also demonstrate that  $DAN_{ngram}$  outperforms  $CSN-sub_{ngram}$  and  $DAN_{codebert}$  outperforms  $CSN-sub_{codebert}$  in the vast majority of cases. On average,  $DAN_{ngram}$  shows an improvement of 58.65% over  $CSN-sub_{ngram}$ , and  $DAN_{codebert}$  shows an average improvement of 136.75% over  $CSN-sub_{codebert}$ , in terms of average normalized difference values. In addition,  $DAN_{ngram}$  always performs better than  $CSN_{ngram}$ . For different  $n$  values,  $DAN$  with  $n=3$  outperforms that with  $n=2$  in the vast majority of cases, but the latter still significantly outperforms the baselines. The results not only confirm the contribution of relatively broad dependency information compared to the local context information incorporated by  $CSN-sub$ , but also confirm the poor performance of directly treating the code as a whole sequence due to incurring too much noise (even though including all dependency information), demonstrating the importance of incorporating code dependency in a reasonable manner for measuring code naturalness (like  $DAN$ ).

Here, we noticed that CodeBERT-based methods ( $DAN_{codebert}$  and CodeBERT-NT) perform worse than  $n$ -gram-based methods ( $DAN_{ngram}$  and  $Ngram-NT$ ) in many cases. The reason may be that CodeBERT was pre-trained on extensive open-source projects, which may not work well on our datasets constructed with minimal de-naturalization, while the cached  $n$ -gram model considers localness well.

We also constructed four sets of unnatural code by mixing the three rules for  $k$ -order ( $k=\{1,2,3,4\}$ ) transformation on each piece of natural code for more sufficient evaluation. Specifically, we randomly selected a rule from the three for each transformation. All the conclusions are consistent with the above ones. Due to the space limit, we put the detailed results at our project homepage.

**Finding 1:** Incorporating code dependency for measuring code naturalness with  $DAN$  achieves a stronger ability of distinguishing natural and unnatural code than the state-of-the-art  $Ngram-NT$  and CodeBERT-NT methods that isolate each code line for measuring code naturalness. The code dependency extraction strategy contribute to the overall effectiveness of  $DAN$  compared to the intuitive methods of treating code as a whole sequence.

## 4.2 RQ2: Effectiveness Comparison on Buggy Line Prioritization

As demonstrated by the existing studies [Jiang et al. 2022; Ray et al. 2016], buggy code lines often exhibit worse naturalness than non-buggy lines. Therefore, code naturalness has been used to identify buggy lines for early inspection in order to improve the efficiency of software development and maintenance [Khanfir et al. 2022; Yan et al. 2020]. In this RQ, we investigated whether incorporating code dependency into measuring code naturalness with  $DAN$  can help improve the effectiveness of buggy line prioritization.

Note that we do not aim to compete with the state-of-the-art static bug finders and fault localization techniques, but use this downstream application to demonstrate the superiority of  $DAN_{ngram}$  over  $Ngram-NT$  and  $DAN_{codebert}$  over CodeBERT-NT.

**4.2.1 Datasets.** To answer RQ2, we collected all the datasets used for buggy line prioritization from the existing studies [Jiang et al. 2022; Khanfir et al. 2022], i.e., Defects4J [Just et al. 2014], GrowingBugs [gro 2023], and SmartSHARK [Herbold et al. 2020]. Note that we are the first to use all the three datasets for evaluating code naturalness measures in the scenario of buggy line prioritization. Each dataset contains a collection of real-world bugs from different Java projects. In total, we used 1,346 real-world bugs in this experiment, including 614 bugs from Defects4J, 324 bugs from GrowingBugs, and 368 bugs from SmartSHARK. Across all three datasets, 27.22% of

Table 2. Effectiveness on buggy line prioritization

Method	Defects4j						GrowingBugs						SmartSHARK					
	20%	40%	60%	80%	100%	MRR	20%	40%	60%	80%	100%	MRR	20%	40%	60%	80%	100%	MRR
DAN <sub>ngram</sub>	0.133	0.240	0.333	0.406	0.474	0.118	0.155	0.259	0.344	0.425	0.499	0.178	0.194	0.356	0.491	0.592	0.667	0.129
DAN <sub>ngram</sub> (n=2)	0.132	0.239	0.333	0.406	0.474	0.117	0.152	0.256	0.343	0.424	0.498	0.177	0.181	0.343	0.483	0.585	0.662	0.126
Ngram-NT	0.105	0.192	0.277	0.356	0.433	0.111	0.135	0.228	0.324	0.408	0.485	0.174	0.151	0.309	0.450	0.565	0.649	0.125
CSN-sub <sub>ngram</sub>	0.095	0.166	0.237	0.311	0.392	0.081	0.118	0.193	0.270	0.347	0.433	0.118	0.097	0.181	0.274	0.381	0.489	0.075
DAN <sub>codebert</sub>	0.066	0.126	0.185	0.240	0.309	0.099	0.064	0.131	0.190	0.255	0.331	0.130	0.115	0.206	0.290	0.362	0.440	0.128
DAN <sub>codebert</sub> (n=2)	0.065	0.121	0.180	0.239	0.306	0.096	0.059	0.117	0.181	0.251	0.329	0.126	0.098	0.182	0.269	0.347	0.429	0.112
CodeBERT-NT	0.050	0.089	0.136	0.188	0.263	0.081	0.056	0.102	0.156	0.215	0.294	0.114	0.082	0.147	0.208	0.285	0.373	0.107
CSN-sub <sub>codebert</sub>	0.061	0.116	0.177	0.236	0.307	0.086	0.058	0.112	0.178	0.245	0.323	0.089	0.098	0.184	0.269	0.347	0.429	0.105

statements are dependent on others and each dependency sequence involves 25 statements per file. We discarded the bugs with compilation errors since the underlying tool (i.e., Soot) extracts code dependency at the Jimple code. The main reason lies in that many bugs from SmartSHARK are very old and thus it is difficult for them to restore the environment and dependency required for compilation. Nevertheless, the number of bugs used in our study is the largest among all the existing studies on buggy lines prioritization [Jiang et al. 2022; Ray et al. 2016]. Each bug is equipped with a buggy version and a bug-fixed version. Following the existing studies [Jiang et al. 2022; Ray et al. 2016], we identified the code lines in the buggy version that have either been deleted or modified in the bug-fixed version as **buggy lines**. The remaining lines in the same files as buggy lines are identified as **non-buggy lines**.

**4.2.2 Process.** Following the existing work [Jiang et al. 2022; Yan et al. 2020], we applied Ngram-NT and CodeBERT-NT to each line in the buggy files for each bug, respectively. It measures the naturalness of each line, and then prioritizes all the lines in the buggy files as the descending order of their measured naturalness scores. The lines with higher priorities are more likely to be buggy. For DAN<sub>ngram</sub> and DAN<sub>codebert</sub>, we applied them to measure the naturalness of each sequence that consists of multiple lines with dependencies and then prioritized all the sequences in the buggy files in the descending order, respectively. To break the tie for the lines within each sequence, we prioritized them by measuring the naturalness of each of these lines same as Ngram-NT/CodeBERT-NT.

**4.2.3 Metrics.** In this experiment, we adopted the widely-used metrics, i.e., **RAUC-k** (Raw Area Under the Curve) and **MRR** (Mean Reciprocal Rank), following the existing studies on prioritization tasks [Chen et al. 2023; Wang et al. 2021b; Yang et al. 2023]. RAUC-k measures the effectiveness of a prioritization technique for the first  $k$  results in a prioritized list [Wang et al. 2021b]. Specifically, it transforms the prioritization result produced by a technique to a plot, where the x-axis represents the number of prioritized lines and the y-axis represents the number of identified buggy lines. Then, it calculates the ratio of the area under the curve for a technique to that for the perfect prioritization. To perform comprehensive comparison, we considered  $k$  to be 20%, 40%, 60%, 80%, 100% of the number of prioritized lines, respectively. MRR measures the effectiveness of a prioritization technique based on its ability to rapidly identify the first correct result [Worster and Haines 2004]. Specifically, MRR treats each buggy file as a query and computes the reciprocal of the rank of the first correctly identified buggy line for each query, and then averages these reciprocal ranks across all queries to measure the prioritization effectiveness. Larger RAUC-k and MRR values mean better prioritization results, indicating better effectiveness of measuring code naturalness in the application of buggy line prioritization.

**4.2.4 Results and Analysis.** Table 2 shows the comparison results in terms of average RAUC-k and MRR across all bugs in each dataset. Due to the space limit, we put the detailed results on each bug

by each method of measuring code naturalness at our project homepage. From Table 2,  $DAN_{ngram}$  outperforms Ngram-NT and  $DAN_{codebert}$  outperforms CodeBERT-NT on all three datasets in terms of all the metrics. For example, on average across all three datasets,  $DAN_{ngram}/DAN_{codebert}$  achieves an improvement of 14.81%~28.48%/14.29%~40.24% over Ngram-NT/CodeBERT-NT in terms of RAUC-20% and 2.30%~6.31%/14.04%~22.22% in terms of MRR. In particular, the improvements of  $DAN_{ngram}$  over Ngram-NT and  $DAN_{codebert}$  over CodeBERT-NT are larger when  $k$  is smaller, indicating that we can identify buggy lines more efficiently especially when the given inspection time is limited. The results also show that  $DAN_{ngram}/DAN_{codebert}$  outperforms CSN-sub $_{ngram}/CSN-sub_{codebert}$  respectively in buggy line prioritization across all three datasets. For example,  $DAN_{ngram}$  and  $DAN_{codebert}$  improve upon their CSN-sub counterparts by 31.36%~100.00% and 8.20%~17.35% respectively, across all three datasets in terms of RAUC-20%. For different  $n$  values, the results demonstrate that DAN with  $n=3$  consistently outperforms that with  $n=2$ , but the latter still significantly outperforms the baselines. Overall, the results show the stable and obvious superiority of DAN over the baselines in identifying buggy lines among all code lines.

Furthermore, we performed a paired sample Wilcoxon signed-rank test [Woolson 2007] at the significance level of 0.05 to check whether each metric value achieved by  $DAN_{ngram}/DAN_{codebert}$  is significantly better than that of Ngram-NT/CodeBERT-NT in statistics. We found that all the  $p$ -values are smaller than 0.05 and the effect size between  $DAN_{ngram}/DAN_{codebert}$  and Ngram-NT/CodeBERT-NT is 0.33/0.55 in terms of MRR across the three datasets, further confirming the superiority of DAN.

**Finding 2:** Incorporating code dependency for measuring code naturalness with DAN can more effectively identify buggy lines from all the code lines than the state-of-the-art Ngram-NT and CodeBERT-NT, demonstrating that DAN provides more accurate measure for code naturalness in the application of buggy line prioritization.

### 4.3 RQ3: Effectiveness Comparison on Training Data Cleansing

Code generation holds great promise in automating various aspects of software development and the naturalness (such as readability and maintainability) of generated code is an important property (besides the widely-watched performance metric – functionality correctness of generated code) [Austin et al. 2021]. Cleansing training data with code naturalness measure (such as DAN) is a potential way of building the code generation models ensuring this property. Moreover, the existing work has indicated that small-scale high-quality training data can achieve performance levels in code generation that are comparable to models trained on larger datasets [Gunasekar et al. 2023]. Indeed, sampling training data with high code naturalness can also ensure the high quality of data to a large extent due to guaranteeing syntactic correctness as well, and thus such a data cleansing strategy will not damage code generation performance too much. In this RQ, we investigated whether cleansing training data by emphasizing code naturalness with DAN (compared to the state-of-the-art Ngram-NT and CodeBERT-NT) can lead to better models that generate code with a more natural and readable structure without sacrificing overall code generation performance.

**4.3.1 Datasets.** To answer RQ3, our subjects include the models and datasets for code generation. Here, we cleansed training data in each dataset, and then trained or fined-tuned each model based on the cleansed data, in order to build a better model. In the study, we used two pre-trained code generation models, i.e., CodeGen-Multi [Nijkamp et al. 2022] and GPT-2 [Radford et al. 2019], as the target models due to the availability, popularity, and reasonable training cost. Moreover, we used APPS [Hendrycks et al. 2021] and HumanEval-X [Zheng et al. 2023] as the datasets, which has

been widely used in the task of code generation [Le et al. 2022; Nijkamp et al. 2022; Zan et al. 2022a; Zhang et al. 2023; Zheng et al. 2023]. HumanEval-X has been introduced in Section 4.1, for which we randomly sampled 90% of data as training data and regarded the remaining data as test data. APPS consists of 10,000 programming problems gathered from various open-access programming websites. It has been officially divided into training and test data, each of which contains 5,000 programming problems. These problems are categorized into three levels according to their degrees of difficulty. Following the existing work [Olausson et al. 2023], we randomly sampled 300 problems from the APPS test set according to the original distribution of the three levels, in order to balance evaluation cost and conclusion generality. As the current implementation of DAN is for Java code, we used the Java versions of HumanEval-X and APPS. The former is provided by the existing work [Zheng et al. 2023], while the latter is constructed by us with the aid of ChatGPT [ope 2023]. In the Java version of APPS, 57.95% of statements are dependent on other statements and each dependency sequence involves 3 statements on average. We also released the Java version of APPS at our project homepage for future research.

**4.3.2 Process.** For each dataset, we first measured the naturalness of each training code snippet by each method, and ranked all the training code snippets in the ascending order of the measured naturalness scores. Then, we selected top-50% training data as the cleansed data, which are used to fine-tune each pre-trained model. In this way, a fine-tuned model is built based on cleansed training data by each method of measuring code naturalness.

For sufficient comparison, we also used all training data in each dataset without cleansing to fine-tune each pre-trained model. Besides, we included a baseline, which randomly selected 50% training data from each dataset to fine-tune each pre-trained model. In total, for a pair of dataset and pre-trained model, we obtained six fine-tuned models, among which two models are based on the cleansed training data by Ngram-NT and  $DAN_{ngram}$ , two models are based on the cleansed training data by CodeBERT-NT and  $DAN_{codebert}$ , one model is based on randomly sampled training data, and one model is based on all training data. For the former five, the size of cleansed training data for model fine-tuning is the same (i.e., 50% of the whole training data) for fair comparison. The fine-tuning process follows the common practice in the existing work [Chen et al. 2021a; Yan et al. 2023]. Finally, we compared the performance of these models on the test data in the corresponding dataset. Note that we did not cleanse test data to avoid the overfitting issue.

**4.3.3 Metrics.** We measured the performance of fine-tuned models using two metrics: CodeBLEU [Ren et al. 2020] and AvgPassRatio (Average Pass Ratio) [Tian and Chen 2023]. CodeBLEU measures the textual similarity between generated code and ground-truth code by considering both syntax match (the proportion of matched subtrees between their ASTs) and dataflow match (the proportion of matched def-use edges in their dataflow graphs). Since the ground-truth code is meticulously crafted by developers, CodeBLEU can help measure the readability of generated code. That is, higher CodeBLEU scores mean that the generated code looks more similar to the developer-written ground truth, indicating higher readability of generated code to a large extent. It is also possible that the generated code is significantly different from the ground-truth code but is also natural and readable. For these cases, CodeBLEU could produce false negatives. However, we did not find such cases by manually analyzing a small set of data via random sampling. Indeed, CodeBLEU's effectiveness in measuring readability has been demonstrated by numerous studies [Ren et al. 2020; Yan et al. 2023; Yang et al. 2024c]. Therefore, the threat from CodeBLEU may be not serious. Currently, it is still an open challenge to design an automated metric that can entirely eliminate false negatives in terms of readability [Zan et al. 2022b], which can be regarded as our future work. Even though CodeBLEU has been also used to measure code generation performance,

Table 3. CodeBLEU effectiveness on training data cleansing

Method	HumanEval-X		APPS	
	CodeGen	GPT-2	CodeGen	GPT-2
All	0.43	0.14	0.34	0.33
Random	0.40	0.07	0.33	0.32
Ngram-NT	0.37	0.17	0.34	0.32
DAN <sub>ngram</sub>	0.46	0.21	0.37	0.37
CodeBERT-NT	0.47	0.14	0.40	0.31
DAN <sub>codebert</sub>	0.49	0.20	0.42	0.35

some recent work has pointed out its limitation in this aspect [Roy et al. 2021], and thus we just used it to measure code readability.

Besides, we used AvgPassRatio to measure code generation performance (the functionality correctness of generated code). AvgPassRatio first calculates the ratio of passing test cases to all executed test cases on each piece of generated code, and then calculates the average ratio across the entire test set. Larger AvgPassRatio values mean better code generation performance.

**4.3.4 Results and Analysis.** Table 3 shows the CodeBLEU score of each fine-tuned model based on different sets of (cleansed) training data for each pair of pre-trained model and dataset, respectively. Row “All” presents the result of each fine-tuned model based on all training data in each dataset. The last five rows present the results of each fine-tuned model based on the cleansed training data in each dataset by five different methods, respectively.

From Table 3, the fine-tuned models with the aid of DAN can generate code with higher readability (measured by CodeBLEU) than those based on all training data, demonstrating the importance of cleansing training data for model building. This aligns with the existing work [Chakraborty et al. 2022; Gunasekar et al. 2023], reinforcing the conclusion that training with natural code can enhance the readability of generated code. However, random sampling cannot achieve better model performance than using all training data on all the subjects, highlighting the necessity of designing effective data cleansing methods. Notably, DAN<sub>ngram</sub> achieves an average improvement of 18.08% over Ngram-NT in facilitating generating natural code, and DAN<sub>codebert</sub> achieves an average improvement of 16.25% over CodeBERT-NT, across all the subjects.

Furthermore, both DAN<sub>ngram</sub> and DAN<sub>codebert</sub> do not make their fine-tuned models sacrifice overall code generation performance. In terms of AvgPassRatio, their fine-tuned models even slightly improve the code generation performance by 0.90%~8.66% over the fine-tuned models with all training data and 0.12%~6.85% over the fine-tuned models with randomly sampled training data. However, both Ngram-NT and CodeBERT-NT make at least one fine-tuned models perform worse than fine-tuning with all training data or randomly sampled training data in terms of AvgPassRatio. The results demonstrate the superiority of DAN over the state-of-the-art methods of measuring code naturalness in cleansing training data for building better models, which can generate more natural code without damaging the overall functionality correctness of the generated code.

**Finding 3:** Incorporating code dependency for measuring code naturalness with DAN can more effectively cleanse training data for building better models (that generate more natural code without sacrificing overall code generation performance), which outperforms the data cleansing methods with the state-of-the-art Ngram-NT and CodeBERT-NT, random sampling, and the way of using all training data.



Table 4. Effectiveness comparison between DAN and its variant w/o sub path (SP) in distinguishing unnatural code

Transformation Rules	DAN <sub>ngram</sub> w/o SP	DAN <sub>ngram</sub>
Dead Code Insertion	12.50%	<b>20.26%</b>
Confusing Code Insertion	24.99%	<b>40.94%</b>
Variable Renaming	33.54%	<b>34.63%</b>

Table 5. Effectiveness of DAN with more advanced statistical models

Transformation Rules	CodeLlama	DAN <sub>CodeLlama</sub>
Dead Code Insertion	2.64%	<b>5.19%</b>
Confusing Code Insertion	5.32%	<b>7.09%</b>
Variable Renaming	2.85%	<b>10.83%</b>

## 5 Discussion

### 5.1 Contribution of Sub-path Extraction

As presented in Section 3, one important component in DAN is to extract n-node sub-paths from each complete path for facilitating the measure of code naturalness. Here, we conducted an experiment to empirically evaluate the contribution of this component by comparing DAN<sub>ngram</sub> with its variant without this component. Specifically, this variant transforms a complete path (instead of a sub-path) into a sequence of code lines for naturalness measure. We repeated the experiment of RQ1 based on DAN<sub>ngram</sub> and this variant. Table 4 presents the comparison results between them in terms of the average normalized difference values. The results demonstrate that DAN<sub>ngram</sub> significantly outperforms the variant of DAN<sub>ngram</sub> without the sub-path extraction component in distinguishing natural code and unnatural code. The average improvement of DAN<sub>ngram</sub> over this variant is 56.59% in terms of average normalized difference values across all the 12 cases. The results confirm the contribution of the sub-path extraction component in DAN<sub>ngram</sub>.

### 5.2 Effectiveness of DAN with More Advanced Statistical Models

As discussed before, our DAN idea can be generally applied independently of the used statistical model. In our study, we chose the cached n-gram model and the pre-trained CodeBERT model as the underlying statistical model in DAN for evaluation, respectively. This is because there are existing studies that proposed the corresponding methods of measuring code naturalness using the two models. Here, we further investigated whether integrating a more advanced large language model in DAN can also outperform the corresponding method of directly using the large language model to measure code naturalness of each line, which can be helpful to further confirm the generalizability of DAN and the orthogonality between our DAN idea and the underlying statistical models. Specifically, we used the CodeLlama-13b-hf model [Roziere et al. 2023] as the representative large language model for investigation. Accordingly, we constructed the instantiation DAN<sub>CodeLlama</sub> and the baseline CodeLlama-NT (which employs CodeLlama-13b-hf to measure code naturalness of each line individually). We repeated the experiment of RQ1 based on DAN<sub>CodeLlama</sub> and CodeLlama-NT. Table 5 presents the comparison results between them in terms of the average normalized difference values. The results demonstrate that DAN<sub>CodeLlama</sub> outperforms CodeLlama-NT in distinguishing natural and unnatural code, achieving an average improvement of 1.24 times across all 12 cases in terms of average normalized difference values. These results further confirm the

generalizability of DAN, demonstrating its consistent enhancement on code naturalness measure regardless of the underlying statistical models.

### 5.3 Future Work

Although DAN does not necessarily achieve state-of-the-art performance on the two downstream applications as our primary goal is to demonstrate the superiority of  $DAN_{ngram}$  over Ngram-NT and  $DAN_{codebert}$  over CodeBERT-NT, future work could explore how DAN can be integrated with state-of-the-art methods to enhance their effectiveness. For example, in the scenario of buggy line prioritization, spectrum-based fault localization (SBFL) is a state-of-the-art method for buggy line prioritization based on test coverage. However, SBFL suffers from the tie problem [Li et al. 2019], where lines within the same block are given the same priority. DAN could address this issue by distinguishing between lines that belong to different sub-paths, thereby breaking ties in SBFL and potentially improving its accuracy. In the scenario of training data cleansing, code evaluation involves multiple dimensions, such as syntactic validity, semantic correctness, and naturalness. While state-of-the-art cleansing methods focus primarily on validity and diversity [Gunasekar et al. 2023; Guo et al. 2024], DAN offers a complementary approach by enhancing code naturalness, which could lead to more comprehensive improvements in code-generation performance.

Moreover, the model-agnostic nature of DAN presents opportunities for further exploration across a broader range of models and tasks, such as code recommendation or code retrieval. DAN is also suitable to distinguish LLM-generated and human-written code as a more precise measure in code naturalness, since the recent work [Xu and Sheng 2024] indicates that slight alterations to LLM-generated code typically cause higher perplexity to the model, making the code appear more unnatural. This further exploration could reveal whether DAN can serve as a static analysis way to address nonlinearities in tasks involving code, in contrast to approaches that rely on learning nonlinear context.

Currently, DAN depends on PDGs, which have scalability limitations. Future research could focus on alleviating scalability challenges. Additionally, DAN has been built on the assumption that naturalness has to be line-level, and implemented on Java, with some experiments conducted on the relatively small HumanEval dataset. In the future, we plan to evaluate DAN on a more diverse set of subjects, including different granularities, programming languages, and datasets.

### 5.4 Threats to Validity

The *internal* threat mainly arises from the implementation of DAN and the compared techniques. To reduce this threat, we conducted thorough code review and replicated the original results of Ngram-NT in its original study based on our re-implementation [Ray et al. 2016].

The threat to *external* validity mainly lies in the subjects used. As presented in Section 4, we implemented and evaluated DAN on Java code, which may not represent the subjects under other programming languages. In fact, the idea of DAN is general and can be implemented to the code written in other programming languages, as long as there is a program analysis tool supporting the extraction of code dependency as DAN. In the future, we will extend and evaluate DAN on a more diverse set of subjects, encompassing different programming languages, to reduce this threat.

The threats to *construct* validity mainly lie in the randomness involved in our study, i.e., the construction of unnatural code in RQ1 and the baseline of randomly sampling training data in RQ3, the setting of  $n$  in the sub-path extraction component, and the metrics. For the first kind of randomness, we used each transformation rule to construct four pieces of unnatural code with different degrees of de-naturalization from each piece of natural code. Such an extensive dataset can help reduce this threat. For the second kind of randomness, we repeated the experiment five times and calculated the average results. For the second threat, we set  $n$  to 3 following the existing

study [Rahman et al. 2019]. In the future, we will investigate the influence of this parameter to further reduce this threat. For the last threat, we have employed the widely-used metrics by following previous studies [Chen et al. 2023; Ren et al. 2020; Tian and Chen 2023; Wang et al. 2021b; Yang et al. 2023]. In fact, we also analyzed the overhead of our technique by taking Defects4j (real-world Java projects) as the representative benchmark. On average, the time spent on extracting dependencies and sub-paths by DAN is just 3s per file. The time spent on measuring naturalness per file by DAN<sub>ngram</sub>/DAN<sub>codebert</sub>/Ngram-NT/CodeBERT-NT is 0.0099s/0.0283s/0.0049s/0.0209s. Overall, the overhead of DAN is acceptable especially when considering its significant effectiveness.

## 6 Related Work

In the literature, there are some existing methods of measuring code naturalness [Hindle et al. 2012; Ray et al. 2016; Tu et al. 2014], which are very relevant to our work. Besides the state-of-the-art Ngram-NT [Ray et al. 2016] and CodeBERT-NT [Khanfir et al. 2022] that have been presented in Section 2 and compared in our study, there are some other methods, upon which Ngram-NT are designed. For example, Hindle et al. [Hindle et al. 2012] applied the n-gram model and cross entropy to measure code naturalness. Tu et al. [Tu et al. 2014] incorporated a cache model to exploit localness of source code for measuring code naturalness more precisely. Our work goes beyond the per-line measure for code naturalness and incorporates code dependency to improve the precision of measuring code naturalness.

Code naturalness has been utilized for several software engineering applications [Ray et al. 2016; Tu et al. 2014; Yan et al. 2020], such as buggy line prioritization and training data cleansing as evaluated in Section 4. Besides, Raychev et al. [Raychev et al. 2014] used the naturalness of code APIs to build the language model for code completion. Allamanis et al. [Allamanis et al. 2014] learned the style of a codebase and utilized software naturalness to suggest revisions for improving stylistic consistency. Dantas et al. [Dantas et al. 2019] proposed to assist search space exploration in program repair using code naturalness. YAN et al. [Yan et al. 2020] proposed a just-in-time defect prediction framework based on code naturalness. Different from them, our work investigated the contribution of incorporating code dependency for code naturalness measure by designing DAN and improved the effectiveness of several naturalness applications. In the future, we will improve the effectiveness of more naturalness applications based on DAN.

There are also some empirical studies on code naturalness [Jiang et al. 2022; Rahman et al. 2019]. Rahman et al. [Rahman et al. 2019] conducted an investigation into the impact of syntax tokens on naturalness measure and explored how different code representations exhibit different levels of repetition, which prompts a suggestion for future research to concentrate on new code representations for code naturalness. However, they did not involve developing new methods to improve the measure of code naturalness, which contrasts with our method of utilizing code dependency to achieve it. Jiang et al. [Jiang et al. 2022] conducted an extensive study to investigate the naturalness of buggy, non-buggy, and bug-fixing code, and confirmed that buggy code are less natural. Different from this work, we investigated whether incorporating code dependency can improve the measure of code naturalness and thus facilitate the downstream applications.

Some studies have looked at leveraging code dependencies for the tasks related to code embedding, such as the pre-trained GraphCodeBERT [Guo et al. 2020] and GraphCode2Vec models [Ma et al. 2022]. Different from them, our work aims to investigate the influence of code dependency on code naturalness measure, which can be regarded as a basis for these tasks.

Code naturalness is also related to code smell, since there is often a correlation between lower code naturalness and the presence of code smell. Code that is less natural tends to be harder to understand and maintain, which can lead to the introduction of code smells. For example, complex or convoluted code structures, poor naming conventions, and inconsistent formatting can all

contribute to decreased naturalness and increase the likelihood of code smells. Conversely, code that exhibits higher naturalness, with clear and concise structure, meaningful variable names, and consistent formatting, is generally less prone to containing code smells. However, it is important to note that while there is a correlation, it is not a strict rule, code may lack naturalness due to factors such as complex logic, tight constraints, or specialized requirements, rather than due to inherent issues like code smells. Conversely, code that appears natural and readable may still contain code smells if certain design or implementation choices introduce inefficiencies or make maintenance difficult. In essence, while naturalness can be an indicator of code quality, it does not guarantee the absence or presence of code smell, and vice versa. Both aspects should be considered independently.

## 7 Conclusion

In this work, we have conducted the first empirical study to investigate whether incorporating code dependency, instead of isolating each code line like existing methods, can improve the precision of measuring code naturalness. To achieve that, we first proposed a reasonable yet general method named DAN for incorporating the rich code dependency information in measuring code naturalness. The evaluation results of DAN in three emerging applications of code naturalness demonstrate that the code dependency information is significant in measuring code naturalness, encouraging more future research on it.

## Data-Availability Statement

We released our artifact and experimental data for replication and practical use [Yang et al. 2024a].

## Acknowledgments

We thank the anonymous reviewers for their valuable suggestions. The work has been supported by the National Natural Science Foundation of China Grant Nos. 62322208, 12411530122, 62232001, 62202324, CCF Young Elite Scientists Sponsorship Program (by CAST).

## References

- Accessed: 2023. GrowingBugs. <https://github.com/jiangyanjie/GrowingBugs/>.
- Accessed: 2023. javalang. <https://github.com/c2nes/javalang/>.
- Accessed: 2023. OpenAI ChatGPT. <https://chat.openai.com/>.
- Accessed: 2023. scikit-learn. <https://scikit-learn.org/stable/>.
- Miltiadis Allamanis, Earl T Barr, Christian Bird, and Charles Sutton. 2014. Learning Natural Coding Conventions. *ACM* (2014).
- Ben Athiwaratkun, Sanjay Krishna Gouda, Zijian Wang, Xiaopeng Li, Yuchen Tian, Ming Tan, Wasi Uddin Ahmad, Shiqi Wang, Qing Sun, Mingyue Shang, et al. 2022. Multi-lingual Evaluation of Code Generation Models. In *The Eleventh International Conference on Learning Representations*.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732* (2021).
- Teresa Busjahn, Roman Bednarik, Andrew Begel, Martha Crosby, James H Paterson, Carsten Schulte, Bonita Sharif, and Sascha Tamm. 2015. Eye movements in code reading: Relaxing the linear order. In *2015 IEEE 23rd International Conference on Program Comprehension*. IEEE, 255–265.
- Casey Casalnuovo, Kevin Lee, Hulin Wang, Prem Devanbu, and Emily Morgan. 2020a. Do programmers prefer predictable expressions in code? *Cognitive science* 44, 12 (2020), e12921.
- Casey Casalnuovo, E Morgan, and P Devanbu. 2020b. Does surprisal predict code comprehension difficulty. In *Proceedings of the 42nd Annual Meeting of the Cognitive Science Society*. Cognitive Science Society Toronto, Canada.
- William B Cavnar, John M Trenkle, et al. 1994. N-gram-based text categorization. In *Proceedings of SDAIR-94, 3rd annual symposium on document analysis and information retrieval*, Vol. 161175. Las Vegas, NV, 14.
- Saikat Chakraborty, Toufique Ahmed, Yangruibo Ding, Premkumar T Devanbu, and Baishakhi Ray. 2022. NatGen: generative pre-training by “naturalizing” source code. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 18–30.

- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021b. Evaluating large language models trained on code.(2021). *arXiv preprint arXiv:2107.03374* (2021).
- Songqiang Chen, Shuo Jin, and Xiaoyuan Xie. 2021a. Testing your question answering software via asking recursively. In *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 104–116.
- Zhichao Chen, Junjie Chen, Weijing Wang, Jianyi Zhou, Meng Wang, Xiang Chen, Shan Zhou, and Jianmin Wang. 2023. Exploring better black-Box test case prioritization via log analysis. *ACM Transactions on Software Engineering and Methodology* 32, 3 (2023), 1–32.
- Altino Dantas, Eduardo F de Souza, Jerffeson Souza, and Celso G Camilo-Junior. 2019. Code naturalness to assist search space exploration in search-based program repair methods. In *Search-Based Software Engineering: 11th International Symposium, SSBSE 2019, Tallinn, Estonia, August 31–September 1, 2019, Proceedings 11*. Springer, 164–170.
- Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiao Cheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, et al. 2020. CodeBERT: A Pre-Trained Model for Programming and Natural Languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. 1536–1547.
- Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, et al. 2023. Textbooks Are All You Need. *arXiv preprint arXiv:2306.11644* (2023).
- Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, LIU Shujie, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, et al. 2020. GraphCodeBERT: Pre-training Code Representations with Data Flow. In *International Conference on Learning Representations*.
- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y Wu, YK Li, et al. 2024. DeepSeek-Coder: When the Large Language Model Meets Programming—The Rise of Code Intelligence. *arXiv preprint arXiv:2401.14196* (2024).
- Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, Collin Burns, Samir Puranik, Horace He, Dawn Song, et al. 2021. Measuring coding challenge competence with apps. *arXiv preprint arXiv:2105.09938* (2021).
- Steffen Herbold, Alexander Trautsch, Benjamin Ledel, Alireza Aghamohammadi, Taher Ahmed Ghaleb, Kuljit Kaur Chahal, Tim Bossenmaier, Bhavet Nagaria, Philip Makedonski, Matin Nili Ahmadabadi, Kristof Szabados, Helge Spieker, Matej Madeja, Nathaniel Hoy, Valentina Lenarduzzi, Shangwen Wang, Gema Rodríguez-Pérez, Ricardo Colomo-Palacios, Roberto Verdecchia, Paramvir Singh, Yihao Qin, Debasish Chakroborti, Willard Davis, Vijay Walunj, Hongjun Wu, Diego Marcilio, Omar Alam, Abdullah Aldaej, Idan Amit, Burak Turhan, Simon Eismann, Anna-Katharina Wickert, Ivano Malavolta, Matus Sulir, Fatemeh Fard, Austin Z. Henley, Stratos Kourtzanidis, Eray Tuzun, Christoph Treude, Simin Maleki Shamasbi, Ivan Pashchenko, Marvin Wyrich, James Davis, Alexander Serebrenik, Ella Albrecht, Ethem Utku Aktas, Daniel Strüber, and Johannes Erbel. 2020. Large-Scale Manual Validation of Bug Fixing Commits: A Fine-grained Analysis of Tangling. *arXiv:2011.06244 [cs.SE]*
- Abram Hindle, Earl T Barr, Zhendong Su, Mark Gabel, and Premkumar Devanbu. 2012. On the naturalness of software. In *2012 34th International Conference on Software Engineering (ICSE)*. IEEE.
- Yanjie Jiang, Hui Liu, Yuxia Zhang, Weixing Ji, Hao Zhong, and Lu Zhang. 2022. Do bugs lead to unnaturalness of source code?. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1085–1096.
- René Just, Darioush Jalali, and Michael D Ernst. 2014. Defects4J: A database of existing faults to enable controlled testing studies for Java programs. In *Proceedings of the 2014 international symposium on software testing and analysis*. 437–440.
- Ahmed Khanfir, Matthieu Jimenez, Mike Papadakis, and Yves Le Traon. 2022. Codebert-nt: code naturalness via codebert. In *2022 IEEE 22nd International Conference on Software Quality, Reliability and Security (QRS)*. IEEE, 936–947.
- Kamran Kowsari, Kiana Jafari Meimandi, Mojtaba Heidarysafa, Sanjana Mendu, Laura Barnes, and Donald Brown. 2019. Text classification algorithms: A survey. *Information* 10, 4 (2019), 150.
- Hung Le, Yue Wang, Akhilesh Deepak Gopmare, Silvio Savarese, and Steven Chu Hong Hoi. 2022. Coderl: Mastering code generation through pretrained models and deep reinforcement learning. *Advances in Neural Information Processing Systems* 35 (2022), 21314–21328.
- Xia Li, Wei Li, Yuqun Zhang, and Lingming Zhang. 2019. Deepfl: Integrating multiple fault diagnosis dimensions for deep fault localization. In *Proceedings of the 28th ACM SIGSOFT international symposium on software testing and analysis*. 169–180.
- Edward Loper and Steven Bird. 2002. NLTK: the Natural Language Toolkit. In *Proceedings of the ACL-02 Workshop on Effective tools and methodologies for teaching natural language processing and computational linguistics-Volume 1*. 63–70.
- Wei Ma, Mengjie Zhao, Ezekiel Soremekun, Qiang Hu, Jie M. Zhang, Mike Papadakis, Maxime Cordy, Xiaofei Xie, and Yves Le Traon. 2022. GraphCode2Vec: Generic Code Embedding via Lexical and Program Dependence Analyses. In *Proceedings of the 19th International Conference on Mining Software Repositories (Pittsburgh, Pennsylvania) (MSR '22)*.

- Association for Computing Machinery, New York, NY, USA, 524–536. <https://doi.org/10.1145/3524842.3528456>
- Erik Nijkamp, Bo Pang, Hiroaki Hayashi, Lifu Tu, Huan Wang, Yingbo Zhou, Silvio Savarese, and Caiming Xiong. 2022. Codegen: An open large language model for code with multi-turn program synthesis. *arXiv preprint arXiv:2203.13474* (2022).
- Theo X Olausson, Jeevana Priya Inala, Chenglong Wang, Jianfeng Gao, and Armando Solar-Lezama. 2023. Demystifying GPT Self-Repair for Code Generation. *arXiv preprint arXiv:2306.09896* (2023).
- Jibesh Patra and Michael Pradel. 2021. Semantic bug seeding: a learning-based approach for creating realistic bugs. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 906–918.
- Adrian Pilkington. 1996. *The Language Instinct: The New Science of Language and Mind*: by Steven Pinker, 1994, The Penguin Press, London pp. 494, ISBN 0 713 99099 6 (hbk.). *Language and Literature* 5, 1 (1996), 71–74.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.
- Musfiqur Rahman, Dharani Palani, and Peter C Rigby. 2019. Natural software revisited. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 37–48.
- Baishakhi Ray, Vincent Hellendoorn, Saheel Godhane, Zhaopeng Tu, Alberto Bacchelli, and Premkumar Devanbu. 2016. On the "naturalness" of buggy code. In *Proceedings of the 38th International Conference on Software Engineering*. 428–439.
- Veselin Raychev, Martin Vechev, and Andreas Krause. 2015. Predicting program properties from "big code". *ACM SIGPLAN Notices* 50, 1 (2015), 111–124.
- Veselin Raychev, Martin Vechev, and Eran Yahav. 2014. Code completion with statistical language models. *Acm Sigplan Notices* 49, 6 (2014).
- Shuo Ren, Daya Guo, Shuai Lu, Long Zhou, Shujie Liu, Duyu Tang, Neel Sundaresan, Ming Zhou, Ambrosio Blanco, and Shuai Ma. 2020. CodeBLEU: a Method for Automatic Evaluation of Code Synthesis. *CoRR* abs/2009.10297 (2020).
- Devjeet Roy, Sarah Fakhoury, and Venera Arnaoudova. 2021. Reassessing automatic evaluation metrics for code summarization tasks. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1105–1116.
- Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950* (2023).
- Chenyao Suo, Junjie Chen, Shuang Liu, Jiajun Jiang, Yingquan Zhao, and Jianrong Wang. 2024. Fuzzing MLIR Compiler Infrastructure via Operation Dependency Analysis. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*.
- Zhao Tian and Junjie Chen. 2023. Test-Case-Driven Programming Understanding in Large Language Models for Better Code Generation. *arXiv:2309.16120* [cs.SE]
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).
- Zhaopeng Tu, Zhendong Su, and Premkumar Devanbu. 2014. On the localness of software. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*. 269–280.
- Raja Vallée-Rai, Phong Co, Etienne Gagnon, Laurie Hendren, Patrick Lam, and Vijay Sundaresan. 1999. Soot-a Java bytecode optimization framework. In *Proceedings of the 1999 conference of the Centre for Advanced Studies on Collaborative research*. 13.
- Yue Wang, Weishi Wang, Shafiq Joty, and Steven CH Hoi. 2021a. CodeT5: Identifier-aware Unified Pre-trained Encoder-Decoder Models for Code Understanding and Generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 8696–8708.
- Zan Wang, Hanmo You, Junjie Chen, Yingyi Zhang, Xuyuan Dong, and Wenbin Zhang. 2021b. Prioritizing test inputs for deep neural networks via mutation analysis. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. IEEE, 397–409.
- Robert F Woolson. 2007. Wilcoxon signed-rank test. *Wiley encyclopedia of clinical trials* (2007), 1–3.
- Andrew Worster and Ted Haines. 2004. Advanced statistics: understanding medical record review (MRR) studies. *Academic emergency medicine* 11, 2 (2004), 187–192.
- Zhenyu Xu and Victor S Sheng. 2024. Detecting AI-Generated Code Assignments Using Perplexity of Large Language Models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 23155–23162.
- Ming Yan, Junjie Chen, Jie M Zhang, Xuejie Cao, Chen Yang, and Mark Harman. 2023. Coco: Testing code generation systems via concretized instructions. *arXiv preprint arXiv:2308.13319* (2023).
- Meng Yan, Xin Xia, Yuanrui Fan, Ahmed E Hassan, David Lo, and Shanping Li. 2020. Just-in-time defect identification and localization: A two-phase framework. *IEEE Transactions on Software Engineering* 48, 1 (2020), 82–101.

- Chen Yang, Junjie Chen, Xingyu Fan, Jiajun Jiang, and Jun Sun. 2023. Silent Compiler Bug De-duplication via Three-Dimensional Analysis. In *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis*. 677–689.
- Chen Yang, Junjie Chen, Jiajun Jiang, and Yuliang Huang. 2024a. *Dependency-aware code naturalness*. <https://doi.org/10.5281/zenodo.12783666>
- Chen Yang, Junjie Chen, Bin Lin, Jianyi Zhou, and Ziqi Wang. 2024b. Enhancing LLM-based Test Generation for Hard-to-Cover Branches via Program Analysis. *arXiv preprint arXiv:2404.04966* (2024).
- Guang Yang, Yu Zhou, Wenhua Yang, Tao Yue, Xiang Chen, and Taolue Chen. 2024c. How important are good method names in neural code generation? a model robustness perspective. *ACM Transactions on Software Engineering and Methodology* 33, 3 (2024), 1–35.
- Daoguang Zan, Bei Chen, Dejian Yang, Zeqi Lin, Minsu Kim, Bei Guan, Yongji Wang, Weizhu Chen, and Jian-Guang Lou. 2022a. CERT: Continual Pre-training on Sketches for Library-oriented Code Generation. In *The 2022 International Joint Conference on Artificial Intelligence*.
- Daoguang Zan, Bei Chen, Fengji Zhang, Dianjie Lu, Bingchao Wu, Bei Guan, Yongji Wang, and Jian-Guang Lou. 2022b. Large language models meet nl2code: A survey. *arXiv preprint arXiv:2212.09420* (2022).
- Shun Zhang, Zhenfang Chen, Yikang Shen, Mingyu Ding, Joshua B Tenenbaum, and Chuang Gan. 2023. Planning with large language models for code generation. *arXiv preprint arXiv:2303.05510* (2023).
- Qinkai Zheng, Xiao Xia, Xu Zou, Yuxiao Dong, Shan Wang, Yufei Xue, Zihan Wang, Lei Shen, Andi Wang, Yang Li, et al. 2023. Codegeex: A pre-trained model for code generation with multilingual evaluations on humaneval-x. *arXiv preprint arXiv:2303.17568* (2023).

Received 2024-04-06; accepted 2024-08-18