



# A Syntax-Guided Edit Decoder for Neural Program Repair

Qihao Zhu

Key Laboratory of HCST, MoE  
DCST, Peking University  
Beijing, China  
Zhuqh@pku.edu.cn

Wenjie Zhang

Key Laboratory of HCST, MoE  
DCST, Peking University  
Beijing, China  
zhang\_wen\_jie@pku.edu.cn

Zeyu Sun

Key Laboratory of HCST, MoE  
DCST, Peking University  
Beijing, China  
szy\_@pku.edu.cn

Kang Yuan

Stony Brook University  
New York, US  
kang.yuan@stonybrook.edu

Lu Zhang

Key Laboratory of HCST, MoE  
DCST, Peking University  
Beijing, China  
zhanglucs@pku.edu.cn

Yuan-an Xiao

Key Laboratory of HCST, MoE  
DCST, Peking University  
Beijing, China  
xiaoyuanan@pku.edu.cn

Yingfei Xiong\*

Key Laboratory of HCST, MoE  
DCST, Peking University  
Beijing, China  
xiongyf@pku.edu.cn

## ABSTRACT

Automated Program Repair (APR) helps improve the efficiency of software development and maintenance. Recent APR techniques use deep learning, particularly the encoder-decoder architecture, to generate patches. Though existing DL-based APR approaches have proposed different encoder architectures, the decoder remains to be the standard one, which generates a sequence of tokens one by one to replace the faulty statement. This decoder has multiple limitations: 1) allowing to generate syntactically incorrect programs, 2) inefficiently representing small edits, and 3) not being able to generate project-specific identifiers.

In this paper, we propose Recoder, a syntax-guided edit decoder with placeholder generation. Recoder is novel in multiple aspects: 1) Recoder generates edits rather than modified code, allowing efficient representation of small edits; 2) Recoder is syntax-guided, with the novel provider/decider architecture to ensure the syntactic correctness of the patched program and accurate generation; 3) Recoder generates placeholders that could be instantiated as project-specific identifiers later.

We conduct experiments to evaluate Recoder on 395 bugs from *Defects4J v1.2*, 420 additional bugs from *Defects4J v2.0*, 297 bugs from *IntroClassJava* and 40 bugs from *QuixBugs*. Our results show that Recoder repairs 53 bugs on *Defects4J v1.2*, which achieves

26.2% (11 bugs) improvement over the previous state-of-the-art approach for single-hunk bugs (TBar). Importantly, to our knowledge, *Recoder is the first DL-based APR approach that has outperformed the traditional APR approaches on this benchmark*. Furthermore, Recoder repairs 19 bugs on the additional bugs from *Defects4J v2.0*, which is 137.5% (11 bugs) more than TBar and 850% (17 bugs) more than SimFix. Recoder also achieves 775% (31 bugs) and 30.8% (4 bugs) improvement on *IntroClassJava* and *QuixBugs* over the baselines respectively. These results suggest that Recoder has better generalizability than existing APR approaches.

## CCS CONCEPTS

• **Software and its engineering**; • **Computing methodologies**  
→ *Software testing and debugging*; *Neural networks*;

## KEYWORDS

Automated program repair, Neural networks

### ACM Reference Format:

Qihao Zhu, Zeyu Sun, Yuan-an Xiao, Wenjie Zhang, Kang Yuan, Yingfei Xiong, and Lu Zhang. 2021. A Syntax-Guided Edit Decoder for Neural Program Repair. In *Proceedings of the 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '21)*, August 23–28, 2021, Athens, Greece. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3468264.3468544>

## 1 INTRODUCTION

Automated program repair (APR) aims to reduce bug-fixing effort by generating patches to aid the developers. Due to the well-known problem of weak test suites [55], even if a patch passes all the tests, the patch still has a high probability of being incorrect. To overcome this problem, existing approaches have used different means to guide the patch generation. A typical way is to learn from existing software repositories, such as learning patterns from existing patches [2, 23, 24, 28, 40, 42, 58], and using program code to guide the patch generation [24, 43, 60, 71, 72].

\*Corresponding author.  
HCST: High Confidence Software Technologies.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
*ESEC/FSE '21*, August 23–28, 2021, Athens, Greece  
© 2021 Association for Computing Machinery.  
ACM ISBN 978-1-4503-8562-6/21/08...\$15.00  
<https://doi.org/10.1145/3468264.3468544>

Deep learning is known as a powerful machine learning approach. Recently, a series of research efforts have attempted to use deep learning (DL) techniques to learn from existing patches for program repair [8, 19, 36, 66]. A typical DL-based approach generates a new statement to replace the faulty statement located by a fault localization approach. Existing DL-based approaches are based on the encoder-decoder architecture [3]: the encoder encodes the faulty statement as well as any necessary code context into a fixed-length internal representation, and the decoder generates a new statement from it. For example, Hata et al. [19] and Tufano et al. [66] adopt an existing neural machine translation architecture, NMT, to generate the bug fix; SequenceR [8] uses a sequence-to-sequence neural model with a copy mechanism; DLFix [36] further treats the faulty statement as an AST rather than a sequence of tokens, and encodes the context of the statement.

However, despite multiple existing efforts, DL-based APR approaches have not yet outperformed traditional APR approaches. Since deep learning has outperformed traditional approaches in many domains, in this paper we aim to further improve the performance of DL-based APR to understand whether we could outperform traditional APR using a DL-based approach. We observe that, though existing DL-based APR approaches have proposed different encoder architectures for APR, the decoder architecture remains to be the standard one, generating a sequence of tokens one by one to replace the original faulty program fragment. The use of this standard decoder significantly limits the performance of DL-based APR. Here we highlight three main limitations.

**Limitation 1: Including syntactically incorrect programs in the patch space.** The goal of the decoder is to locate a patch from a patch space. The smaller the patch space is, the easier the task is. However, viewing a patch as a sequence of tokens unnecessarily enlarges the patch space, making the decoding task difficult. In particular, this space representation does not consider the syntax of the target programming language and includes many syntactically incorrect statements, which can never form a correct patch.

**Limitation 2: Inefficient representation of small edits.** Many patches only modify a small portion of a statement, and re-generating the whole statement leads to an unnecessarily large patch space. For example, let us consider the patch of defect Closure-14 in the Defects4J benchmark [26], as shown in Figure 1. This patch only changes one token in the statement, but under existing representation, it is encoded as a sequence of length 13. The program space containing this patch would roughly contain  $n^{13}$  elements, where  $n$  is the total number of tokens. On the other hand, let us consider a patch space including only one-token change edits. To generate that patch, only selecting a token in the faulty statement and a new token for replacement is needed. This patch space contains only  $mn$  elements, where  $m$  is the number of tokens in the faulty statement. Therefore, the size of the patch space is significantly reduced.

**Limitation 3: Not being able to generate project-specific identifiers.** Source code of programs often contains project-specific identifiers like variable names. Since it is impractical to include all possible identifiers in the patch space, existing DL-based APR approaches only generate identifiers that have frequently appeared in the training set. However, different projects have different sets of project-specific identifiers, and therefore only considering identifiers in the training set may exclude possible patches from the

```
- cfa.createEdge(fromNode, Branch.UNCOND, finallyNode);
+ cfa.createEdge(fromNode, Branch.ON_EX, finallyNode);
```

Figure 1: The Patch for Closure-14 in Defects4J

```
- return cAvailableLocaleSet.contains(locale);
+ return availableLocaleSet().contains(locale);
```

Figure 2: The Patch for Lang-57 in Defects4J

patch space. For example, Figure 2 shows the patch for defect Lang-57 in Defects4J. To generate this patch, we need to generate the identifier “availableLocaleSet”, which is a method name of the faulty class, and is unlikely to be included in the training set. As a result, existing DL-based approaches cannot generate patches like this.

In this paper, we propose a novel DL-based APR approach, Recoder, standing for **repair decoder**. Similar to existing approaches, Recoder is based on the encoder-decoder architecture. To address the limitations above, the decoder of Recoder has following two novel techniques.

**Novelty 1: Syntax-Guided Edit Decoding with Provider/Decoder Architecture** (concerning limitation 1 & 2). To address limitation 2, the decoder component of Recoder produces a sequence of edits rather than a new statement. Our edit decoder is based on the idea of the syntax-guided decoder in existing neural program generation approaches [56, 63, 64, 75]. For an unexpanded non-terminal node in a partial AST, the decoder estimates the probability of each grammar rule to be used to expand the node. Based on this, the decoder selects the most probable sequence of rules to expand the start symbol into a full program using a search algorithm such as beam search. We observe that edits could also be described by a grammar. For example, the previous patch for defect Closure-14 could be described by the following grammar:

$$\begin{aligned} \text{Edit} &\rightarrow \text{Insert} \mid \text{Modify} \mid \dots \\ \text{Modify} &\rightarrow \text{modify}(\text{NodeID}, \text{NTS}^1) \end{aligned}$$

Here `modify` represents replacing an AST subtree denoted by its root node ID (*NodeID*) in the faulty statement with a newly generated subtree (*NTS*<sup>1</sup>).

However, directly applying the existing syntax-guided decoder to the grammar above would not form an effective program repair approach, because the choice of expanding different non-terminal nodes may need to be deduced along with different types of dependencies. First, the expansion of some non-terminals depends on the local context, e.g., the choice of *NodeID* depends on the faulty statement, and the neural network needs to be aware of the local context to make a suitable choice. Second, to guarantee syntax correctness (limitation 1), dependency exists among the choices for expanding different non-terminal nodes, e.g., when *NodeID* expands to an ID pointing to a node with non-terminal *JavaExpr*, *NTS* should also expand to *JavaExpr* to ensure syntactic correctness. These choices cannot be effectively pre-defined, and thus the existing syntax-guided decoders, which only select among a set of pre-defined grammar rules, do not work here.

<sup>1</sup>“NTS” stands for “non-terminal symbol in an AST”.

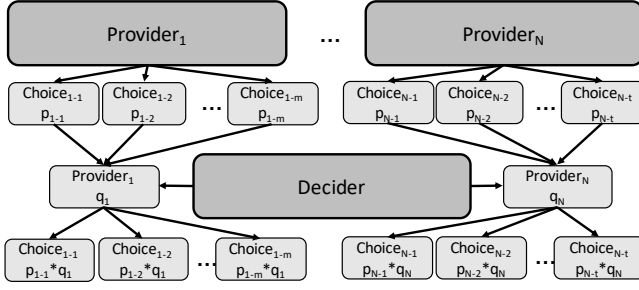


Figure 3: Provider/Decoder Architecture

To overcome these problems, Recoder introduces a provider/decoder architecture, as shown in Figure 3. A provider is a neural component that provides a set of choices for expanding a non-terminal and estimates the probability  $p_i$  of each choice. A basic provider is the rule predictor, which, similar to existing syntax-guided decoders, estimates the probability of each grammar rule to expand the node. Fixing the Closure-14 example needs another provider, namely the subtree locator, which estimates the probability of each subtree in the faulty statement to be replaced. On the other hand, the decoder is a component that estimates the probability of  $q_j$  using each provider. In this example, when expanding *Edit*, the probability of using the rule predictor is 1, and the probability of using the subtree locator is 0; when expanding *Modify*, the probability of using the rule predictor is 0 and the probability of using the subtree locator is 1 (the located subtree decides both the content of *NodeID* and the root symbol of *NTS*). Finally, the choices provided by all providers form the final list of choices, while the probability of each choice is the product of the probability predicted by its provider and the probability of the provider itself, i.e.,  $p_i * q_j$ .

In this example, for each non-terminal, we use the choices from only one provider, and thus the probabilities of providers are either 0 or 1. Later we will see that expanding some non-terminals requires comparing the choices of multiple providers, and the probabilities of providers could be a real number between 0 and 1.

**Novelty 2: Placeholder Generation** (concerning limitation 3). To generate project-specific identifiers, a direct idea is to add another provider that selects an identifier from the local context. However, to implement such a provider, the neural component needs to access all of the name declarations within the current project. This is a difficult job, as the neural component could hardly encode all source code from the whole project.

Instead of relying on the neural network to generate project-specific identifiers, in Recoder the neural network generates placeholders for such identifiers, and these placeholders are instantiated with all feasible identifiers when applying the edits. A feasible identifier is an identifier compatible with constraints in the programming language, such as the type system. As for defect Lang-57 shown in Figure 2, Recoder first generates a placeholder for “availableLocaleSet”, and it will be replaced with all methods accessible in the local context that takes no arguments and returns an object with a member method “contains”. Each replacement forms a new patch. The key insight is that, when considering constraints in the programming language, the number of choices for replacing a placeholder with an identifier is small, and thus instantiating the placeholders with all possible choices is feasible.

To train the neural network to generate placeholders, we replace infrequent user-defined identifiers in the training set with placeholders. In this way, the neural network learns to generate placeholders for these identifiers.

Our experiment is conducted on four benchmarks: (1) 395 bugs from Defects4J v1.2 for comparison with existing approaches. (2) 420 additional bugs from Defects4J v2.0, (3) 297 bugs from IntroClassJava, and (4) 40 bugs from QuixBugs to evaluate the generalizability of Recoder. The results show that Recoder correctly repairs 53 bugs on the first benchmark, which are 26.2% (11 bugs) more than TBar [40] and 55.9% (19 bugs) more than SimFix [24], two best-performing single-hunk APR approaches on Defects4J v1.2; Recoder also correctly repairs 19 bugs on the second benchmark, which are 137.5% (11 bugs) more than TBar and 850.0% (17 bugs) more than SimFix. On IntroClassJava and QuixBugs, Recoder repairs 35 bugs and 17 bugs respectively, which also achieves better performance than the existing APR tools that were evaluated on the two benchmarks. The results suggest that Recoder has better performance and better generalizability than existing approaches. To our knowledge, *this is the first DL-based APR approach that has outperformed traditional APR approaches*. To summarize, this paper makes the following contributions:

- We propose a syntax-guided edit decoder for APR with a provider/decoder architecture to accurately predict the edits and ensure that the edited program is syntactically correct and uses placeholders to generate patches with project-specific identifiers.
- We design Recoder, a neural APR approach based on the decoder architecture described above.
- We evaluate Recoder on 395 bugs from Defects4J v1.2 and 420 additional bugs from Defects4J v2.0. The results show that Recoder significantly outperforms state-of-the-art approaches for single-hunk bugs in terms of both repair performance and generalizability.

## 2 EDITS

We introduce the syntax and semantics of edits and their relations to providers in this section. The neural architecture to generate edits and implement providers will be discussed in the next section.

### 2.1 Syntax and Semantics of Edits

Figure 4 shows the syntax of edits. Note that our approach is not specific to a particular programming language and can be applied to any programming language (called the *host language*) that has a concept similar to the statement. In particular, it is required that when a statement is present in a program, a sequence of statements can also be present at the same location. In other words, inserting a statement before any existing statement would still result in a syntactically correct program. To ensure syntactic correctness of the edited program, the syntax of edits depends on the syntax of the host language. In Figure 4, “HL” refers to the host programming language our approach applies to. In the following we explain each rule in Figure 4 in order.

As defined by Rule 1 and Rule 2, an *Edits* is a sequence of *Edit* ended by a special symbol end. An *Edit* can be one of two edit operations, insert and modify.

1. *Edits* → *Edit; Edits* | end
2. *Edit* → *Insert* | *Modify*
3. *Insert* → *insert*(*<HLStatement>*)
4. *Modify* → *modify*(  
     *<ID of an AST Node with a NTS>*,  
     *<the same NTS as the above NTS>*)
5. *<Any NTS in HL>* →  
     *copy*(*<ID of an AST Node with the same NTS>*)  
     | *<The original production rules in HL>*)
6. *<HLIdentifier>* → *placeholder*  
     | *<Identifiers in the training set>*)

“HL” stands for “host language”. “NTS” stands for “non-terminal symbol”. “*<HLStatement>*” is the non-terminal in the grammar of the host language representing a statement. “*<HLIdentifier>*” is the non-terminal in the grammar of the host language representing an identifier.

Figure 4: The Syntax of Edits

Rule 3 defines the syntax of insert operation. The insert operation inserts a newly generated statement before the faulty statement. As shown in Rule 3, the insert operation has one parameter, which is the statement to insert. Here *<HLStatement>* refers to the non-terminal in the grammar of the host language that represents a statement. This non-terminal could be expanded into a full statement, or a copy operation that copies a statement from the original program, or a mixture of both. This behavior will be explained later in Rule 5.

Rule 4 defines the syntax of modify operation. The modify operation replaces an AST subtree in the faulty statement with a new AST subtree. The modify operation has two parameters. The first parameter is the ID of the root node in the AST subtree to be replaced. The ID of a node is defined as the order of a node in the pre-order traversal sequence, e.g., the 6th visited node has the ID of 6. The second parameter is an AST subtree whose root node has the same symbol, i.e., the root node cannot be changed. In this way, the replacement ensures syntactic correctness. To ensure that there is an actual change, the subtree to be replaced should have more than one node, i.e., the root node should have a non-terminal symbol.

For both insert and modify, we need to generate a new AST subtree. It is noticeable that in many patches, the AST subtree being inserted or modified is not completely original; some of its subtrees may be copied from other parts of the program. Taking advantage of this property, copy operation is introduced to further reduce the patch space. Rule 5 defines the syntax of this operation. It is a meta-rule applied to any non-terminal symbol of the host language. For any non-terminal symbol in the host language, we add a production rule that expands it into a copy operation. The original production rules for this non-terminal are also kept, so that when generating the edits, the neural network could choose to directly generate a new subtree or to copy one.

The copy operation has one parameter, which identifies the root node of the AST subtree to be copied. The AST subtree can be selected from the faulty statement or its context. In our current implementation, we allow copying from the method surrounding the faulty statement. Also, to ensure syntactic correctness, the root

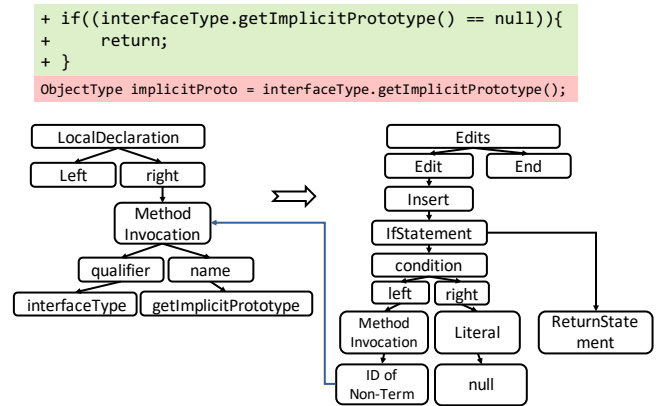


Figure 5: Example of Insert Operation (Closure-2)

node of the subtree to be copied should have the same non-terminal symbol as the symbol being extended.

Finally, Rule 6 introduces placeholder into the grammar. Normally, the grammar of a programming language uses a terminal symbol to represent an identifier. To enable the neural network to generate concrete identifiers as well as the placeholder, we change identifier nodes into non-terminals, which expand to either placeholder or one of the frequent identifiers in the training set. In our current implementation, an identifier is considered frequent if it appears more than 100 times in the training set.

When applying the edits, the placeholder tokens are replaced with feasible identifiers within the context. We first collect all identifiers in the current projects by performing a lexical analysis and collect the tokens whose lexical type is *<HLIdentifier>*, the symbol representing an identifier in the host language. Then we filter identifiers based on the following criteria: (1) the identifier is accessible from the local context, and (2) replacing the placeholder with the identifier would not lead to type errors. The remaining identifiers are feasible identifiers.

Figure 5 and Figure 6 show two example patches represented by edits. The patch in Figure 5 inserts an if statement, and the conditional expression contains a method invocation that is copied from the faulty statement. The patch in Figure 6 replaces the qualifier of a method invocation with another invocation, where the name of the method is a placeholder to be instantiated later.

**THEOREM 2.1.** *The edited programs are syntactically correct.*

**PROOF.** It is easy to see that the theorem holds by structural induction on the grammar of the edits. First, the requirement on the host programming language ensures that inserting a statement before another statement is syntactical correct. Second, when replacing a subtree with modify, the root symbol of the subtree remains unchanged. Third, the new subtree in insert and modify is generated by either using the grammar rules of the host language, or copying a subtree with the same root symbol. Finally, instantiating a placeholder ensures syntactic correctness because we only replace a placeholder with a token whose lexical type is *<HLIdentifier>*. □

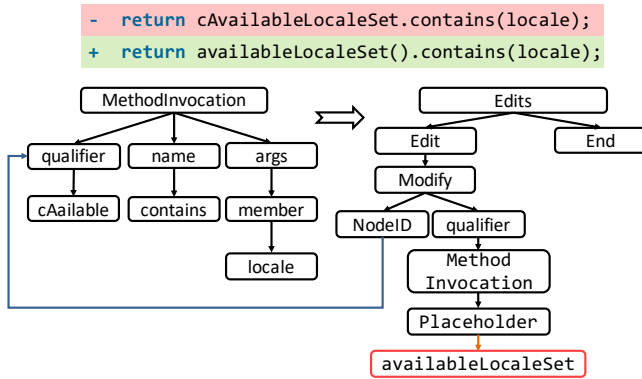


Figure 6: Example of Modify Operation (Lang-57)

Table 1: Providers for non-terminals

Component	Associated Non-terminals
Rule Predictor	$\langle Edits, Edit, Insert, \langle HLIIdentifier \rangle, \langle Any\ NTS\ in\ HL \rangle \rangle$
Subtree Locator	$\langle Modify \rangle$
Tree Copier	$\langle Any\ NTS\ in\ HL \rangle$

## 2.2 Generation of Edits

Since the choice of expanding a non-terminal may depend on the local context or a previous choice, we use providers to provide choices and estimate their probabilities. Our current implementation has three types of providers. Table 1 shows these providers and their associated non-terminals.

For non-terminals  $\langle Edits, Edit, Insert \rangle$  and  $\langle HLIIdentifier \rangle$ , the rule predictor is responsible for providing choices and estimates the probability of each production rule. The rule predictor consists of a neural component and a logic component. After the neural component assigns the probability for each production rule, the logic component resets the probability of rules whose left-hand side is not the corresponding non-terminal to zero and normalizes the remaining probabilities.

For  $\langle Modify \rangle$ , the subtree locator is responsible for providing the choices. The subtree locator estimates the probability of each AST subtree with a size larger than 1 in the faulty statement. The choice of a subtree  $t$  means that we should expand  $\langle Modify \rangle$  into  $\text{modi fy}(ID, NTS)$  where  $ID$  is the root ID of  $t$  and  $NTS$  is the root symbol of  $t$ .

For any non-terminal in the grammar of the host language (note that  $\langle HLIIdentifier \rangle$  is a terminal symbol in the host language), both the rule predictor and the tree copier are responsible to provide the choices. The tree copier estimates the probabilities of each AST subtree with a size larger than 1 in the method surrounding the faulty statement. The choice of a subtree  $t$  means that we should expand the non-terminal into  $\text{copy}(ID)$ , where  $ID$  is the root ID of  $t$ . Similar to the rule predictor, the tree copier employs a logic component after the neural component to reset the probabilities of subtrees whose root symbols are different from the non-terminal symbol being expanded.

Finally, the decoder assigns a probability to each provider. The decoder also includes a similar logic component, which resets the probability of a provider to zero if that provider is not responsible

for the current non-terminal symbol. For example, if the symbol being expanded is  $\langle Modify \rangle$ , the decoder resets the probability of rule predictor and tree copier to zero.

## 3 MODEL ARCHITECTURE

The design of our model is based on the state-of-the-art syntax-guided code generation model, TreeGen [64]. It is a tree-based Transformer [67] that takes a natural language description as input and produces a program as output. Since our approach takes a faulty statement and its context as input and produces edits as output, we replace the components in TreeGen for encoding natural language description and decoding the program.

Figure 7 shows an overview of our model. The model performs one step in the edit generation process, which is to predict probabilities of choices for expanding a non-terminal node. Beam search is used to find the best combination of choices for generating the complete edits. The model consists of four main components:

- The **code reader** that encodes the faulty statement and its context.
- The **AST reader** that encodes the partial AST of the edits that have been generated.
- The **tree path reader** that encodes a path from the root node to a non-terminal node which should be expanded.
- The **edit decoder** that takes the encoded information from the previous three components and produces a probability of each choice for expanding the non-terminal node.

Among them, the AST reader and the tree path reader are derived from TreeGen, where the code reader and the edit decoder are newly introduced in this paper. In this section, we focus on describing the latter two components in detail.

### 3.1 Code Reader

The code reader component encodes the faulty statement and the method surrounding the faulty statement as its context, where the faulty statement is localized by a fault localization technique. It uses the following three inputs. (1) **AST traversal sequence**. This is a sequence of tokens following the pre-order traversal of the AST,  $c_1, c_2, \dots, c_L$ , where  $c_i$  is the token encoding vector of the  $i$ th node embedded via word embedding [49]. (2) **Tag embedding**. This is a sequence of tags following the same pre-order traversal of the AST, where each tag denotes which of the following cases the corresponding node belongs to: 1. in the faulty statement, 2. in the statement before the faulty statement, 3. in the statement after the faulty statement, or 4. in other statements. Each tag is embedded via an embedding-lookup table. We denote the tag embedding as  $t_1, t_2, \dots, t_L$ . (3) **AST-based Graph**. Considering that the former two inputs do not capture the neighbor relations between AST nodes, in order to capture such information, we treat an AST as a directional graph where the nodes are AST nodes and the edges link a node to each of its children and its left sibling, as shown in Figure 8(b). This graph is embedded as an adjacent matrix.

The code reader uses three sub-layers to encode the three inputs above, as discussed in the following sections.

**3.1.1 Self-Attention.** The self-attention sub-layer encodes the AST traversal sequence, following the Transformer [67] architecture to capture the long dependency information in the AST.

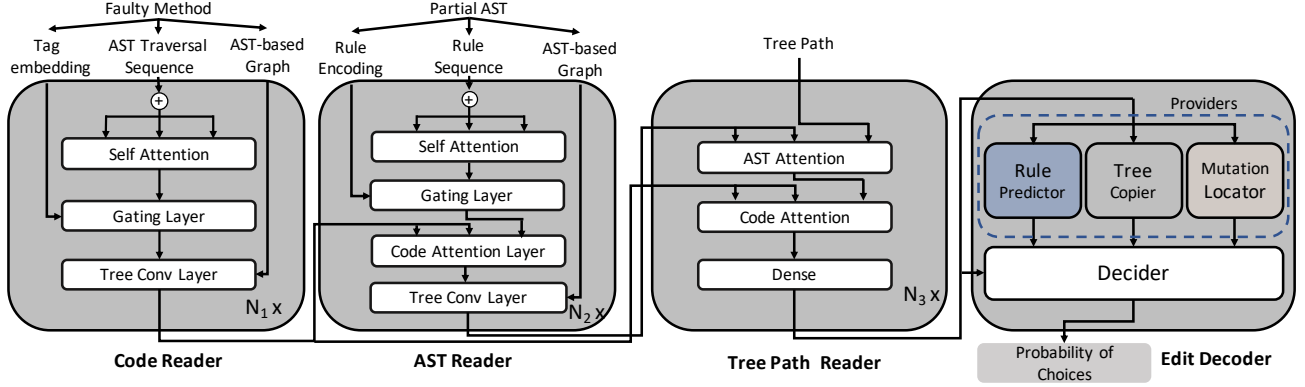


Figure 7: Overview of Recoder

Given the embedding of the input AST traversal sequence, we use position embedding to represent positional information of the AST token. The input vectors are denoted as  $c_1, c_2, \dots, c_L$ , and the position embedding of  $i$ th token is computed as

$$p_{(i,2j)} = \sin(pos/(10000^{2j/d})) \quad (1)$$

$$p_{(i,2j+1)} = \cos(pos/(10000^{2j/d})) \quad (2)$$

where  $pos = i + step$ ,  $j$  denotes the element of the input vector and  $step$  denotes the embedding size. After we get the vector of each position, it is directly added to the corresponding input vector, where  $e_i = c_i + p_i$ .

Then, we adopt multi-head attention layer to capture non-linear features. Following the definition of Vaswani et al. [67], we divide the attention mechanism into  $H$  heads. Each head represents an individual attention layer to extract unique information. The single attention layer maps the query  $Q$ , the key  $K$ , and the value  $V$  into a weighted-sum output. The computation of the  $j$ th head layer can be represented as

$$head_j = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (3)$$

where  $d_k = d/H$  denotes the length of each extracted feature vector, and  $Q$ ,  $K$  and  $V$  are computed by a fully-connected layer from  $Q$ ,  $K$ ,  $V$ . In the encoder, vectors  $Q$ ,  $K$  and  $V$  are all the outputs of the position embedding layer  $e_1, e_2, \dots, e_L$ . The outputs of these heads are further joint together with a fully-connected layer, which is computed by

$$Out = [head_1; \dots; head_H] \cdot W_h \quad (4)$$

where  $W_h$  denotes the weight of the fully-connected layer and  $Out$  denotes the outputs  $a_1, a_2, \dots, a_L$  of the self-attention sub-layer.

**3.1.2 Gating Layer.** This sub-layer takes the outputs of the previous layer and the tag embedding as input. Gating mechanism, as defined in TreeGen [64], is used in this layer. It takes three vectors named  $q, c_1, c_2$  as input and aims to incorporate  $c_1$  with  $c_2$  based on  $q$ . The computation of gating mechanism can be represented as

$$\alpha_i^{c_1} = \exp(q_i^T k_i^{c_1}) / \sqrt{d_k} \quad (5)$$

$$\alpha_i^{c_2} = \exp(q_i^T k_i^{c_2}) / \sqrt{d_k} \quad (6)$$

$$h_i = (\alpha_i^{c_1} v_i^{c_1} + \alpha_i^{c_2} v_i^{c_2}) / (\alpha_i^{c_1} + \alpha_i^{c_2}) \quad (7)$$




Figure 8: Example of AST-based Graph

where  $d_k = d/H$  is a normalization factor,  $H$  denotes the number of heads, and  $d$  denotes the hidden size;  $q_i$  is computed by a fully-connected layer over the control vector  $q_i$ ;  $k_i^{c_1}, v_i^{c_1}$  is computed by another fully-connected layer over vector  $c_1$ ;  $k_i^{c_2}$  and  $v_i^{c_2}$  are also computed by the same layer with different parameters over the vector  $c_2$ .

In our model, we treat the outputs of the self-attention sub-layer  $a_1, a_2, \dots, a_L$  as  $q$  and  $c_1$ , and the tag embedding  $t_1, t_2, \dots, t_L$  as  $c_2$ . Thus, embedding of the  $i$ th AST node of the gating-layer can be represented as  $u_i = \text{Gating}(a_i, a_i, t_i)$ .

**3.1.3 Tree Conv Layer.** This sub-layer takes the output  $u_i$  of the previous layer and the AST-based graph  $G$  (represented as an adjacency matrix) as input. We adopt a GNN [62, 76] layer to process the inputs, and the encoding of the neighbors  $r_i$  is computed as

$$g_i = W_g \sum_{r^j \in G} A_{r_i, r^j}^n u^j \quad (8)$$

where  $W_g$  is the weight of a fully-connected layer and  $\hat{A}$  is a normalized adjacency matrix of  $G$ . The computation of the normal operation proposed by Kipf and Welling [30] is represented as  $\hat{A} = S_1^{-1/2} A S_2^{-1/2}$ , where  $A$  is the adjacency matrix of  $G$ , and  $S_1, S_2$  are the diagonal matrices with a summation of  $A$  in columns and rows. Then, the encoding of the neighbors is directly added to the input vector.

In summary, the code reader has  $N_1$  blocks of these three sub-layers, and yields the features of the input AST,  $t_1, t_2, \dots, t_L$ , which would be used for the AST reader and the tree path reader.

## 3.2 AST Reader

The AST reader encodes the partial generated AST of the edit, which has the same structure as the one in TreeGen [64]. This component takes three inputs derived from the partial generated

AST as code reader. The rule sequence is represented as real-value vectors and then is fed into a self-attention layer. We then integrate the output of the self-attention layer with the rule encoding via a gating layer as Equation 7. We also adopt a multi-head attention layer over the outputs of the code reader and the gating layer like the decoder-encoder attention in Transformer. Finally, we use a tree convolutional layer like the code reader to extract the structural information. More details of this component can be found in the publication of TreeGen [64].

### 3.3 Tree Path Reader

The tree path reader encodes the information of the non-terminal node to be expanded and is the same as the one in TreeGen [64]. This component represents the non-terminal node as a path from the root to the node to be expanded and transforms the nodes in this path into real-value vectors. As shown in Figure 7, these vectors are fed into two attention layers like Equation 4. Finally, a set of two fully-connected layers, where the first layer has a *GELU* [20] activation function, are followed to extract features for edit decoder. More details of this component can be found in the publication of TreeGen [64]. We denote the output of the tree path reader as  $\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_T$ .

### 3.4 Edit Decoder

The edit decoder takes the output of the tree path reader,  $\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_T$  with length  $T$ , as input. These vectors are produced by the tree path reader and contain the encoded information from all the inputs: the faulty statement with its surrounding method, the partial AST generated so far, and the tree path denoting the node to be expanded.

**3.4.1 Provider.** As mentioned before, there are currently three types of providers: rule predictor, tree copier, and subtree locator. These providers take the vector  $\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_T$  as input and output the probability of choices for different non-terminals.

**Rule Predictor.** The rule predictor estimates the probability of each production rule in the grammar of edits. The neural component of this decoder consists of a fully-connected layer. The output of the fully-connected layer is denoted as  $\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_T$ . Then, these vectors are normalized via softmax, which computes the normalized vectors  $\mathbf{p}_1^r, \mathbf{p}_2^r, \dots, \mathbf{p}_T^r$  by

$$\mathbf{p}_k^r(m) = \frac{\exp\{\mathbf{s}_k^m\}}{\sum_{j=1}^{N_r} \exp\{\mathbf{s}_k^j\}} \quad (9)$$

where  $N_r$  denotes the number of production rules in the grammar of edits, and  $m$  denotes the  $m$ th dimension of the vector  $\mathbf{p}_k^r$  (i.e., the production rule with ID  $m$ ). In particular, invalid rules whose left-hand side is not the corresponding non-terminal are not allowed in our approach. For these rules, the logic component resets the output of the fully-connected layer to  $-\infty$ . Thus, the probability of invalid rules will be zero after softmax normalization.

**Tree Copier.** This provider is designed for any non-terminal symbol in the grammar of edits to choose a subtree in the local context. The neural component is based on a pointer network [68]. The computation can be represented as

$$\boldsymbol{\theta}_i = \mathbf{v}^T \tanh(W_1 \mathbf{d}_i + W_2 \mathbf{t}) \quad (10)$$

where  $\mathbf{t}$  denotes the output of the code reader, and  $\mathbf{v}, W_1, W_2$  denote the trainable parameters. The logic component also resets  $\boldsymbol{\theta}$  to  $-\infty$  if the root symbol of the corresponding subtree is different from the symbol being expanded. These vectors are then normalized via softmax as Equation 10. We denote the normalized vector as  $\mathbf{p}_1^t, \mathbf{p}_2^t, \dots, \mathbf{p}_T^t$ .

**Subtree Locator.** This component outputs an ID of the subtree in the faulty statement for not-terminal symbol, *Modify*, in the grammar of edits. The computation of this component is the same as the tree copier. We denote the output vector of this provider as  $\mathbf{p}_1^s, \mathbf{p}_2^s, \dots, \mathbf{p}_T^s$ .

**3.4.2 Decider.** For these three providers, the decider estimates the probability of using each provider. The neural component also takes the output of the tree path reader,  $\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_T$ , as input, and produces the probability of using each provider as output. The computation can be represented as  $\boldsymbol{\lambda}_i = W \mathbf{d}_i + \mathbf{b}$ , where  $W$  and  $\mathbf{b}$  denote the parameters of a fully-connected layer. The logic component resets  $\boldsymbol{\lambda}$  to  $-\infty$  if the corresponding provider is not responsible for the symbol being expanded following Table 1. Then, the vectors are normalized via softmax as Equation 9. We denote the normalized vectors as  $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \dots, \boldsymbol{\lambda}_T$ . The final probability of each choice can be computed as

$$\mathbf{o}_i = [\boldsymbol{\lambda}_i^r \mathbf{p}_i^r; \boldsymbol{\lambda}_i^t \mathbf{p}_i^t; \boldsymbol{\lambda}_i^s \mathbf{p}_i^s] \quad (11)$$

where  $\mathbf{o}_i$  will be the probability vector of the next production rule at  $i$ th step during patch generation.

## 3.5 Training and Inference

During training, the model is optimized by maximizing the negative log-likelihood of the oracle edit sequence and we do not use the logic component in the providers and decider. Here we would like Recoder to learn the distribution of the rules handled by the logic component. If the logic component is present at training, Recoder would not be trained for a large portion of rules. During inference, these unseen rules would distort the distribution of output, making Recoder fail to distinguish the part of rules that it is supposed to distinguish.

When generating edits, inference starts with the rule *start* : *start*  $\rightarrow$  *Edits*, expanding a special symbol *start* to *Edits*. The recursive prediction terminates if every leaf node in the predicted AST is a terminal. We use beam search with a size of 100 to generate multiple edits.

Generated edits may contain placeholders. Though the number of choices for a single placeholder is small, the combination of multiple placeholders may be large. Therefore, we discard patches containing more than one placeholder symbol during beam search.

## 3.6 Patch Generation and Validation

Patches are generated according to the result of the fault localization technique. In our approach, the model described above is invoked for each suspicious faulty statement according to the result of fault localization. For each statement, we generate 100 valid patch candidates via beam search: when beam search generates a valid patch, we remove it from the search set and continue to search for the next patch until 100 candidates are generated in total

for that statement. After patches are generated, the final step is to validate them via the test suite written by developers. The validation step filters out patches that do not compile or fail a test case. All generated patches are validated until a plausible patch (a patch that passes all test cases) is found.

## 4 EXPERIMENT SETUP

We have implemented Recoder for the Java programming language. In this and the next sections we report our experiments on repairing Java bugs.

### 4.1 Research Questions

Our evaluation aims to answer the following research questions:

**RQ1: What is the performance of Recoder?**

To answer this question, we evaluated our approach on the widely used APR benchmark, *Defects4J v1.2*, and compared it with traditional and DL-based APR tools.

**RQ2: What is the contribution of each component in Recoder?**

To answer this question, we started from the full model of Recoder, and removed each component in turn to understand its contribution to performance.

**RQ3: What is the generalizability of Recoder?**

To answer this question, we first conducted an experiment on 420 additional bugs from *Defects4J v2.0*. To our best knowledge, this is the first APR approach that has been applied to this benchmark. We compared Recoder with the previous two best-performing APR approaches for single-hunk bugs on *Defects4J v1.2*, namely TBar [40] and SimFix [24]. In addition, we also applied Recoder to other two benchmarks, *QuixBugs* and *IntroClassJava*, via *RepairThemAll* [10] framework, which allows the execution of automatic program repair tools on benchmarks of bugs.

### 4.2 Dataset

The neural network model in our approach needs to be trained with a large number of history patches. To create this training set, we crawled Java projects created on GitHub [16] between March 2011 and March 2018, and downloaded 1,083,185 commits where the commit message contains at least one word from the following two groups, respectively: (1) *fix, solve*; (2) *bug, issue, problem, error*. Commits were filtered to include only patches that modify one single statement or insert one new statement, corresponding to two types of edits that our approach currently supports. To avoid data leak, we further discarded patches where (1) the project is a clone to *Defects4J* project or a program repair project using *Defects4J*, or (2) the method modified by the patch is the same as the method modified by any patch in *Defects4J v1.2* or *v2.0*, based on AST comparison. There are 103,585 valid patches left after filtering, which are further split into two parts: 80% for training and 20% for validation.

We used four benchmarks to measure the performance of Recoder. The first one contains 395 bugs from *Defects4J v1.2* [26], which is a commonly used benchmark for automatic program repair research. The second one contains 420 additional bugs from *Defects4J v2.0* [26]. *Defects4J v2.0* introduces 438 new bugs compared

with *Defects4J v1.2*. However, GZoltar [57], the fault localization approach used by our implementation as well as two baselines (TBar and SimFix), failed to finish on the project *Gson*, so we excluded 18 bugs in *Gson* from our benchmark. The third one contains 40 bugs from *QuixBugs* [37], which is a benchmark with 40 buggy algorithmic programs specified by test cases. The last one, *IntroClassJava* [11], consists of 297 buggy Java programs generated from the *IntroClass* [34] benchmark for C.

### 4.3 Fault Localization

In our experiment, two settings for fault localization are used. In the first setting, the faulty location of a bug is unknown to APR tools, and they rely on existing fault localization approaches to localize the bug. Recoder uses Ochiai [1] (implemented in GZoltar [57]), which is widely used in existing APR tools [24, 40]. In the second setting, the actual faulty location is given to APR tools. This is to measure the capability of patch generation without the influence of a specific fault localization tool, as suggested and adopted in previous studies [6, 44, 66].

### 4.4 Baselines

We selected existing APR approaches as the baselines for comparison. Since Recoder generates only single-hunk patches (patches that only change a consecutive code fragment), we chose 10 traditional single-hunk APR approaches that are often used as baselines in existing studies: jGenProg [35], HDRRepair [32], Nopol [73], CapGen [69], SketchFix [22], TBar [40], FixMiner [31], SimFix [24], PraPR [15], AVATAR [39]. In particular, TBar correctly repairs the highest number of bugs on *Defects4J v1.2* as far as we know. We also selected DL-based APR approaches that adopt the encoder-decoder architecture to generate patches and have been evaluated on *Defects4J* as baselines. Four approaches have been chosen based on this criteria, namely, SequenceR [66], CODIT [6], DLFix [36], and CoCoNuT [44].

For *Defects4J v1.2*, the performance data of the baselines are collected from existing papers [40, 41]. For additional bugs from *Defects4J v2.0*, two best-performing single-hunk APR approaches on *Defects4J v1.2*, TBar and SimFix, are adapted and executed for comparison. For *QuixBugs* and *IntroClassJava*, we directly chose the APR tools used in *RepairThemAll* [10] and DL-based APR tools which have experimented on these two benchmarks as baselines: jGenProg [35], RSRepair [55], Nopol [73], and CoCoNuT [44]. We also directly used the result reported in the original papers [10, 44].

### 4.5 Correctness of Patches

To check the correctness of the patches, we manually examined every patch if it is the same with or semantically equivalent to the patch provided by *Defects4J*, as in previous works [15, 24, 36, 40, 44]. To reduce possible errors made in this process, every patch is examined by two of the authors individually and is considered correct only if both authors consider it correct. The kappa score of the experiment is 0.98. Furthermore, we also publish all the patches generated by Recoder for public judgment<sup>2</sup>.

<sup>2</sup>The source code of Recoder, generated patches, and an online demo are available at <https://github.com/pkuzqh/Recoder>



```
- this(time, RegularTimePeriod.DEFAULT_TIME_ZONE, Locale.getDefault());
+ this(time, zone, Locale.getDefault());
```

Figure 9: Chart-8 - A bug fixed by Recoder with Modify operation

```
- if (result != null) {
+ if(((result != null) && !result.isNoType())){
```

Figure 10: Closure-104 - A bug fixed by Recoder with placeholder generation

### 4.6 Implementation Details

Our approach is implemented based on PyTorch [53], with parameters set to  $N_1 = 5, N_2 = 9, N_3 = 2$ , i.e., the code reader contains a stack of 5 blocks, the AST reader contains a stack of 9 blocks, and the decoder contains a stack of 2 blocks, respectively. Embedding sizes for all embedding vectors are set to 256, and all hidden sizes are set following the configuration of TreeGen [64]. During training, dropout [21] is used to prevent overfitting, with the drop rate of 0.1. The model is optimized by Adam [29] with learning rate 0.0001. These hyper-parameters and parameters for our model are chosen based on the performance on validation set.

We set a 5-hour running-time limit for Recoder, following existing studies [24, 36, 44, 61].

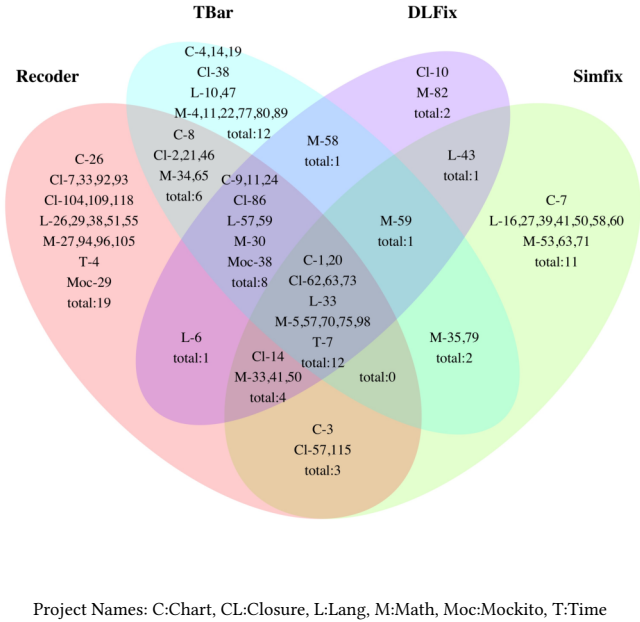
## 5 EXPERIMENTAL RESULTS

### 5.1 Performance of Recoder (RQ1)

5.1.1 *Results without Perfect Fault Localization.* We first compare Recoder with the baselines in the setting where no faulty location is given. Results as Table 2 shown only include baselines that have been evaluated under this setting. As shown, Recoder correctly repairs 53 bugs and outperforms all of the previous single-hunk APR techniques on Defects4J v1.2. In particular, Recoder repairs 26.2% (11 bugs) more bugs than the previous state-of-the-art APR tool for single-hunk bugs, TBar. Within our knowledge, Recoder is the first DL-based APR approach that has outperformed the traditional APR approaches.

We show a few example patches that are possibly generated with the help of the novel techniques in Recoder. As shown in Figure 9, Chart-8 is a bug that DLFix fails to fix. The correct patch only changes a parameter of the method invocation while DLFix needs to generate the whole expression. By contrast, Recoder generates a *modify* operation that changes only one parameter. Figure 10 shows a bug only repaired by Recoder. This patch relies on a project-specific method, “isNoType”, and thus cannot be generated by many of the existing approaches. However, Recoder fixes it correctly by generating a placeholder and then instantiating it with “isNoType”.

5.1.2 *Results with Perfect Fault Localization.* Table 3 shows the result where the actual faulty location is provided. As before, only baselines that have been evaluated under this setting are listed. Recoder still outperforms all of the existing APR approaches, including traditional ones. Also, compared with Recoder using Ochiai for fault localization, this model achieves a 36.5% improvement. The result implies that Recoder can achieve better performance with better fault localization techniques.



Project Names: C:Chart, CL:Closure, L:Lang, M:Math, Moc:Mockito, T:Time

Figure 11: Degree of Complementarity.

5.1.3 *Degree of Complementarity.* We further investigate to what extent Recoder complements the three best-performing existing approaches for fixing single-hunk bugs, TBar, SimFix, and DLFix. Figure 11 reveals the overlaps of the bugs fixed by different approaches. As shown, Recoder fixes 19 unique bugs when compared with three baselines. Moreover, Recoder fixes 34, 28, 27 unique bugs compared with SimFix, TBar, and DLFix, respectively. This result shows that Recoder is complementary to these best-performing existing approaches for single-hunk bugs.

### 5.2 Contribution of Each Component (RQ2)

To answer RQ2, we conducted an ablation test on Defects4J v1.2 to figure out the contribution of each component. Since the ablation test requires much time, we only conducted the experiment based on Ochiai Fault Localization scenario.

Table 4 shows the results of the ablation test. We respectively removed three edit operations, modify, copy, and insert, as well as the generation of placeholders. As shown in the table, removing any of the components leads to a significant drop in performance. This result suggests that the two novel techniques proposed in Recoder are the key to its performance.

### 5.3 Generalizability of Recoder (RQ3)

The results on Defects4J v2.0, QuixBugs and IntroClassJava are shown in Table 5 and Table 6. As shown, on Defects4J v2.0, all three approaches repair a smaller proportion of bugs, suggesting that the additional bugs on Defects4J v2.0 are probably more difficult to repair. Nevertheless, Recoder still repairs most bugs compared with baselines, 19 in total, achieving 137.5% (11 bugs) improvement over TBar and 850.0% (17 bugs) improvement over SimFix. We believe that the considerable performance drops of TBar and SimFix are caused by their design: TBar is based on validated patterns on

**Table 2: Comparison without Perfect Fault Localization**

Project	jGenProg	HDRRepair	Nopol	CapGen	SketchFix	FixMiner	SimFix	TBar	DLFix	PraPR	AVATAR	Recoder
Chart	0/7	0/2	1/6	4/4	6/8	5/8	4/8	<b>9/14</b>	5/12	4/14	5/12	8/14
Closure	0/0	0/7	0/0	0/0	3/5	5/5	6/8	8/12	6/10	12/62	8/12	<b>17/31</b>
Lang	0/0	2/6	3/7	5/5	3/4	2/3	<b>9/13</b>	5/14	5/12	3/19	5/11	<b>9/15</b>
Math	5/18	4/7	1/21	12/16	7/8	12/14	14/26	<b>18/36</b>	12/28	6/40	6/13	15/30
Time	0/2	0/1	0/1	0/0	0/1	1/1	1/1	1/3	1/2	0/7	1/3	<b>2/2</b>
Mockito	0/0	0/0	0/0	0/0	0/0	0/0	0/0	1/2	1/1	1/6	<b>2/2</b>	<b>2/2</b>
Total	5/27	6/23	5/35	21/25	19/26	25/31	34/56	42/81	30/65	26/148	27/53	<b>53/94</b>
P(%)	18.5	26.1	14.3	<b>84.0</b>	73.1	80.6	60.7	51.9	46.2	17.6	50.9	56.4

In the cells,  $x/y$  denotes the number of correct patches, and  $y$  denotes the number of patches that can pass all the test cases.

**Table 3: Comparison with Perfect Fault Localization**

Project	SequenceR	CODIT	DLFix	CoCoNuT	TBar	Recoder
Chart	3	4	5	7	<b>11</b>	<b>11</b>
Closure	3	3	11	9	17	<b>26</b>
Lang	3	3	8	7	<b>13</b>	10
Math	4	6	13	16	<b>22</b>	19
Time	0	0	2	1	2	<b>3</b>
Mockito	0	0	1	<b>4</b>	3	2
Total	13	16	40	44	68	<b>71</b>

**Table 4: Ablation Test for Recoder on Defects4J v1.2**

Project	-modify	-subtreecopy	-insert	-placeholder	Recoder
Chart	4	6	7	8	8
Closure	6	12	12	11	16
Lang	3	6	5	5	10
Math	7	8	9	9	15
Time	1	1	1	1	2
Mockito	2	1	1	1	2
Total	23	34	35	35	53

**Table 5: Comparison on the 420 additional bugs**

Project	# Used Bugs	Bug IDs	TBar	SimFix	Recoder
Cli	39	1-5,7-40	1/7	0/4	<b>3/3</b>
Clousre	43	134 - 176	0/5	<b>1/5</b>	0/7
JacksonDataBind	112	1-112	0/0	0/0	0/0
Codec	18	1-18	<b>2/6</b>	0/2	<b>2/2</b>
Collections	4	25-28	0/1	0/1	0/0
Compress	47	1-47	1/13	0/6	<b>3/9</b>
Csv	16	1-16	1/5	0/2	<b>4/4</b>
JacksonCore	26	1-26	0/6	0/0	0/4
Jsoup	94	1-94	3/7	1/5	<b>7/13</b>
JXPath	22	1-22	0/0	0/0	0/4
Total	420	-	8/50	2/25	<b>19/46</b>

Defects4J v1.2, which may not generalize beyond the projects in Defects4J v1.2; SimFix relies on similar code snippets in the same project, but new projects in Defects4J v2.0 are much smaller, and thus the chance to find similar code snippets become smaller. On the other hand, Recoder is trained from a large set of patches collected from different projects and is thus more likely to generalize to new projects. On QuixBugs and IntroClassJava, Recoder also repaired 775% (31 bugs) and 30.8% (4 bugs) more bugs on IntroClassJava and

**Table 6: Comparison on IntroClassJava and QuixBugs**

Project	# Used Bugs	jGenProg	RSRepair	Nopol	CoCoNuT	Recoder
IntroClassJava	297	1/4	4/22	3/32	-	<b>35/56</b>
QuixBugs	40	0/3	2/4	1/4	13/20	<b>17/17</b>
Total	337	1/7	6/26	4/36	13/20	<b>52/73</b>

**Performance of Recoder on Different Size of Dataset****Figure 12: Adequacy of the Dataset.**

QuixBugs over the baselines respectively, further confirming the effectiveness and generalizability of Recoder.

## 6 DISCUSSION

### 6.1 Adequacy of Dataset

To understand the adequacy of our training data, we trained Recoder on differently sized subsets of the original training dataset and calculate the loss over the single-hunk bugs in the Defects4J v1.2 dataset. For each subset, we train 5 models with different random seeds and report the average performance of these models. Figure 12 shows the sensitivity analysis of dataset size for Recoder. As shown, the loss and the diversity both decrease with the increase of training subset size. Therefore, the performance of Recoder may further increase if more training data are provided.

### 6.2 Limitations of Recoder

Recoder shares limitations common to most APR approaches: 1) Bugs to be fixed should be reproducible by failing test cases. 2)

Effective fault localization is needed to identify the faulty statement. Recoder also shares the limitations common to DL-based approaches: The performance would degrade if the training set and the testing set have different distributions.

## 7 RELATED WORK

**DL-based APR Approaches.** APR has been approached using different techniques, such as heuristics or random search [33, 35, 45, 54], semantic analysis [7, 12, 22, 27, 39, 46–48], manually defined or automatically mined repair patterns [2, 14, 23, 24, 31, 42, 52, 59], and learning from source code [24, 43, 69–72]. For a thorough discussion of APR, existing surveys [13, 17, 50, 51] are recommended to readers.

A closely related series of work is APR based on deep learning. The mainstream approaches treat APR as a statistical machine translation that generates the fixed code with faulty code. DeepFix [18] learns the syntax rules via a sequence-to-sequence model to fix syntax errors. Nguyen et al. [52] and Tufano et al. [66] also adopt a sequence-to-sequence translation model to generate the patch. They use sequence-to-sequence NMT with a copy mechanism. Chakraborty et al. [6] propose CODIT, which learns code edits by encoding code structures in an NMT model to generate patches. Li et al. [36] propose a Tree-LSTM to encode the abstract syntax tree of the faulty method to generate the patches. CoCoNuT, as proposed by Lutellier et al. [44], adopts CNN to encode the faulty method and generates the patch token by token. Compared to them, our paper is the first work that aims to improve the decoder and employs a syntax-guided manner to generate edits, with the provider/decider architecture and placeholder generation.

**Multi-Hunk APR.** Most of the above approaches generate single-hunk patches—patches that change one place of the program. Recently, Saha et al. [61] propose Hercules to repair multi-hunk bugs by discovering similar code snippets and applying similar changes. Since lifting single-hunk repair to multiple-hunk repair is a generic idea and can also be applied to Recoder, we did not directly compare Recoder with the multi-hunk repair tools in our evaluation. Nevertheless, though Recoder only repairs single-hunk bugs, we notice that it still outperforms Hercules by repairing 5 more bugs on the Defects4J v1.0 (including all projects from v1.2 except for Mockito), the dataset Hercules has been evaluated on.

**DL-based Code Generation.** Code generation aims to generate code from a natural language specification and has been intensively studied during recent years. With the development of deep learning, Ling et al. [38] propose a neural machine translation model to generate the program token by token. Being aware that code has the constraints of grammar rules and is different from natural language, Yin and Neubig [75] and Rabinovich et al. [56] propose to generate the AST of the program via expanding from a start node. To integrate the semantic of identifiers, OCoR [77] proposes to encode the identifiers at character level. To alleviate the long dependency problem, a CNN decoder [63] and TreeGen (a tree-based Transformer) [64] are proposed to generate the program. In this paper, we significantly extend TreeGen to generate the edit sequence for program repair.

**DL-based Code Edit Generation.** Several existing DL-based approaches also use the idea of generating edits on programs [4, 5, 9,

65, 74]. Tarlow et al. [65] view a program as a sequence of tokens and generate a sequence of token-editing commands. Brody et al. [4] view a program as a tree and generate node-editing or subtree-editing commands. Dinella et al. [9] view a program as a graph and generate node-editing commands. Compared with our approach, there are three major differences. First, the existing approaches are not syntax-guided and treat an edit script as a sequence of tokens. As a result, they may generate syntactically incorrect edit scripts and do not ensure syntactic correctness of the edited program. On the other hand, our approach introduces the provider/decider architecture and successfully realizes the syntax-guided generation for edits. Second, none of the existing approaches support placeholder generation, and thus are ineffective in generating edits with project-specific identifiers. Third, the editing commands they use are atomic and are inefficient in representing large changes. For example, to insert a variable declaration, in our approach there is one insert operation: `insert(int var = 0;)`. However, the existing approaches have to represent this change as a sequence of 5 insertions, where each insertion inserts one token.

## 8 THREATS TO VALIDITY

**Threats to external validity.** mainly lie in the evaluation dataset we used. First, though our approach applies to different programming languages, so far, we have only implemented and evaluated it on Java, so future work is needed to understand its performance on other programming languages. Second, though we have evaluated on Defects4J v2.0, QuixBugs, and IntroClassJava, it is yet unknown how our approach generalizes to different datasets [25]. This is a future work to be explored.

**Threats to internal validity.** mainly lie in our manual assessment of patch correctness. To reduce this threat, two authors have independently checked the correctness of the patches, and a patch is considered correct only if both authors consider it correct. The generated patches also have been released for public assessment.

## 9 CONCLUSION

In this paper, we propose Recoder, a syntax-guided edit decoder with placeholder generation for automated program repair. Recoder uses a novel provider/decider architecture to ensure accurate generation and syntactic correctness of the edited program and generates placeholders for project-specific identifiers. In the experiment, Recoder achieved 26.2% improvement (11 bugs) over the existing state-of-the-art APR approach for single-hunk bugs on Defects4J v1.2. Importantly, Recoder is the first DL-based APR approach that has outperformed traditional APR techniques on this benchmark. Further evaluation on three other benchmarks shows that Recoder has better generalizability than some state-of-the-art APR approaches.

## ACKNOWLEDGMENTS

This work is sponsored by the National Key Research and Development Program of China under Grant No. 2017YFB1001803, National Natural Science Foundation of China under Grant Nos. 61922003, and a grant from ZTE-PKU Joint Laboratory for Foundation Software.

## REFERENCES

- [1] Rui Abreu, Peter Zoetewij, and Arjan J. C. Van Gemund. 2007. On the Accuracy of Spectrum-based Fault Localization. In *Testing: Academic & Industrial Conference Practice & Research Techniques-mutation*.
- [2] Johannes Bader, Andrew Scott, Michael Pradel, and Satish Chandra. 2019. Getafix: Learning to fix bugs automatically. *Proceedings of the ACM on Programming Languages* 3, OOPSLA (2019), 1–27.
- [3] Dmztry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. *CoRR abs/1409.0473* (2015).
- [4] Shaked Brody, Uri Alon, and Eran Yahav. 2020. A Structural Model for Contextual Code Changes. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 215 (Nov. 2020), 28 pages. <https://doi.org/10.1145/3428283>
- [5] Yingkui Cao, Zeyu Sun, Yanzhen Zou, and Xie Bing. 2020. Structurally-enhanced approach for automatic code change transformation. *Ruan Jian Xue Bao/Journal of Software*, 2021,32(4):1006 (Jul 2020). <https://doi.org/10.13328/j.cnki.jos.006227>
- [6] Saikat Chakraborty, Miltiadis Allamanis, and Baishakhi Ray. 2018. CODIT: Code Editing with Tree-Based Neural Machine Translation. *arXiv preprint arXiv:1810.00314* (2018).
- [7] Liushan Chen, Yu Pei, and Carlo A Furiá. 2017. Contract-based program repair without the contracts. In *2017 32nd IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 637–647.
- [8] Z. Chen, S. J. Komrusch, M. Tufano, L. Pouchet, D. Poshyvanik, and M. Monperrus. 2019. SEQUENCER: Sequence-to-Sequence Learning for End-to-End Program Repair. *IEEE Transactions on Software Engineering* (2019), 1–1. <https://doi.org/10.1109/TSE.2019.2940179>
- [9] Elizabeth Dinella, Hanjun Dai, Ziyang Li, Mayur Naik, Le Song, and Ke Wang. 2020. Hoppity: Learning Graph Transformation to Detect and Fix Bugs in Programs. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=SJeqs6EFvB>
- [10] Thomas Durieux, Fernanda Madeiral, Matias Martinez, and Rui Abreu. 2019. Empirical Review of Java Program Repair Tools: A Large-Scale Experiment on 2,141 Bugs and 23,551 Repair Attempts. In *Proceedings of the 27th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '19)*. <https://arxiv.org/abs/1905.11973>
- [11] Thomas Durieux and Martin Monperrus. 2016. *IntroClassJava: A Benchmark of 297 Small and Buggy Java Programs*. Technical Report. Universite Lille 1. <https://hal.archives-ouvertes.fr/hal-01272126/document>
- [12] Xiang Gao, Bo Wang, Gregory J Duck, Ruyi Ji, Yingfei Xiong, and Abhik Roychoudhury. 2021. Beyond Tests: Program Vulnerability Repair via Crash Constraint Extraction. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 30, 2 (2021), 1–27.
- [13] L. Gazzola, D. Micucci, and L. Mariani. 2019. Automatic Software Repair: A Survey. *IEEE Transactions on Software Engineering* 45, 1 (2019), 34–67. <https://doi.org/10.1109/TSE.2017.2755013>
- [14] Ali Ghanbari, Samuel Benton, and Lingming Zhang. 2019. Practical program repair via bytecode mutation. In *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 19–30.
- [15] A. Ghanbari and L. Zhang. 2019. PraPR: Practical Program Repair via Bytecode Mutation. In *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 1118–1121. <https://doi.org/10.1109/ASE.2019.00116>
- [16] GitHub. 2020. <https://github.com/>. GitHub.
- [17] Claire Le Goues, Michael Pradel, and Abhik Roychoudhury. 2019. Automated program repair. *Commun. ACM* 62, 12 (2019), 56–65.
- [18] Rahul Gupta, Soham Pal, Aditya Kanade, and Shirish Shevade. 2017. DeepFix: Fixing Common C Language Errors by Deep Learning. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence* (San Francisco, California, USA) (AAAI'17). AAAI Press, 1345–1351.
- [19] Hideaki Hata, Emad Shihab, and Graham Neubig. 2018. Learning to Generate Corrective Patches using Neural Machine Translation. *CoRR abs/1812.07170* (2018). [arXiv:1812.07170](http://arxiv.org/abs/1812.07170) <http://arxiv.org/abs/1812.07170>
- [20] D. Hendrycks and K. Gimpel. 2016. Bridging Nonlinearities and Stochastic Regularizers with Gaussian Error Linear Units. (2016).
- [21] Geoffrey E. Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2012. Improving neural networks by preventing co-adaptation of feature detectors. *CoRR abs/1207.0580* (2012). [arXiv:1207.0580](http://arxiv.org/abs/1207.0580) <http://arxiv.org/abs/1207.0580>
- [22] Jinru Hua, Mengshi Zhang, Kaiyuan Wang, and Sarfraz Khurshid. 2018. Sketchfix: A tool for automated program repair approach using lazy candidate generation. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 888–891.
- [23] Jiajun Jiang, Luyao Ren, Yingfei Xiong, and Lingming Zhang. 2019. Inferring program transformations from singular examples via big code. In *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 255–266.
- [24] Jiajun Jiang, Yingfei Xiong, Hongyu Zhang, Qing Gao, and Xiangqun Chen. 2018. Shaping Program Repair Space with Existing Patches and Similar Code. In *International Symposium on Software Testing & Analysis*. 298–309.
- [25] Yanjie Jiang, Hui Liu, Nan Niu, Lu Zhang, and Yamin Hu. 2021. Extracting Concise Bug-Fixing Patches from Human-Written Patches in Version Control Systems. In *43rd IEEE/ACM International Conference on Software Engineering, ICSE 2021, Madrid, Spain, 22-30 May 2021*. IEEE, 686–698. <https://doi.org/10.1109/ICSE43902.2021.00069>
- [26] René Just, Dariouh Jalali, and Michael D. Ernst. 2014. Defects4J: A database of existing faults to enable controlled testing studies for Java programs. In *Proceedings of the International Symposium on Software Testing and Analysis (ISSTA)*. San Jose, CA, USA, 437–440.
- [27] Shalini Kaleeswaran, Varun Tulsian, Aditya Kanade, and Alessandro Orso. 2014. Minthint: Automated synthesis of repair hints. In *Proceedings of the 36th International Conference on Software Engineering*. 266–276.
- [28] D. Kim, J. Nam, J. Song, and S. Kim. 2013. Automatic patch generation learned from human-written patches. In *2013 35th International Conference on Software Engineering (ICSE)*. 802–811. <https://doi.org/10.1109/ICSE.2013.6606626>
- [29] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. *CoRR abs/1412.6980* (2015).
- [30] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [31] Anil Koyuncu, Kui Liu, Tegawendé F Bissyandé, Dongsun Kim, Jacques Klein, Martin Monperrus, and Yves Le Traon. 2020. Fixminer: Mining relevant fix patterns for automated program repair. *Empirical Software Engineering* (2020), 1–45.
- [32] Xuan Bach D Le, David Lo, and Claire Le Goues. 2016. History driven program repair. In *2016 IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, Vol. 1. IEEE, 213–224.
- [33] C. Le Goues, M. Dewey-Vogt, S. Forrest, and W. Weimer. 2012. A systematic study of automated program repair: Fixing 55 out of 105 bugs for \$8 each. In *2012 34th International Conference on Software Engineering (ICSE)*. 3–13. <https://doi.org/10.1109/ICSE.2012.6227211>
- [34] Claire Le Goues, Neal Holtschulte, Edward K. Smith, Yuriy Brun, Premkumar Devanbu, Stephanie Forrest, and Westley Weimer. 2015. The ManyBugs and IntroClass Benchmarks for Automated Repair of C Programs. *IEEE Transactions on Software Engineering (TSE)* 41, 12 (December 2015), 1236–1256. <https://doi.org/10.1109/TSE.2015.2454513> DOI: 10.1109/TSE.2015.2454513
- [35] C. Le Goues, T. Nguyen, S. Forrest, and W. Weimer. 2012. GenProg: A Generic Method for Automatic Software Repair. *IEEE Transactions on Software Engineering* 38, 1 (2012), 54–72. <https://doi.org/10.1109/TSE.2011.104>
- [36] Yi Li, Shaohua Wang, and Tien N. Nguyen. 2020. DLFix: Context-Based Code Transformation Learning for Automated Program Repair. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering (Seoul, South Korea) (ICSE '20)*. Association for Computing Machinery, New York, NY, USA, 602–614. <https://doi.org/10.1145/3377811.3380345>
- [37] Derrick Lin, James Koppel, Angela Chen, and Armando Solar-Lezama. 2017. QuixBugs: a multi-lingual program repair benchmark set based on the quixey challenge. 55–56. <https://doi.org/10.1145/3135932.3135941>
- [38] Wang Ling, Phil Blunsom, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, Fumin Wang, and Andrew Senior. 2016. Latent Predictor Networks for Code Generation. In *ACL*. 599–609.
- [39] Kui Liu, Anil Koyuncu, Dongsun Kim, and Tegawendé F Bissyandé. 2019. Avatar: Fixing semantic bugs with fix patterns of static analysis violations. In *2019 IEEE 26th International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 1–12.
- [40] Kui Liu, Anil Koyuncu, Dongsun Kim, and Tegawendé F Bissyandé. 2019. Tbar: Revisiting template-based automated program repair. In *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 31–42.
- [41] Kui Liu, Shangwen Wang, Anil Koyuncu, Kisub Kim, Tegawendé F Bissyandé, Dongsun Kim, Peng Wu, Jacques Klein, Xiaoguang Mao, and Yves Le Traon. 2020. On the efficiency of test suite based program repair: A systematic assessment of 16 automated repair systems for java programs. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*. 615–627.
- [42] Fan Long, Peter Amidon, and Martin Rinard. 2017. Automatic inference of code transforms for patch generation. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*. 727–739.
- [43] Fan Long and Martin Rinard. 2016. Automatic patch generation by learning correct code. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. 298–312.
- [44] Thibaud Lutellier, Hung Viet Pham, Lawrence Pang, Yitong Li, and Lin Tan. 2020. CoCoNuT: combining context-aware neural translation models using ensemble for program repair. In *ISSTA '20: 29th ACM SIGSOFT International Symposium on Software Testing and Analysis*.
- [45] Matias Martinez and Martin Monperrus. 2016. ASTOR: A Program Repair Library for Java (Demo). In *Proceedings of the 25th International Symposium on Software Testing and Analysis (Saarbrücken, Germany) (ISSTA 2016)*. Association for Computing Machinery, New York, NY, USA, 441–444. <https://doi.org/10.1145/2931037.2948705>
- [46] Sergey Mechtaev, Alberto Griggio, Alessandro Cimatti, and Abhik Roychoudhury. 2018. Symbolic execution with existential second-order constraints. In *Proceedings*

- of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering. 389–399.
- [47] Sergey Mechtaev, Jooyong Yi, and Abhik Roychoudhury. 2015. Directfix: Looking for simple program repairs. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, Vol. 1. IEEE, 448–458.
- [48] Sergey Mechtaev, Jooyong Yi, and Abhik Roychoudhury. 2016. Angelix: Scalable multiline program patch synthesis via symbolic analysis. In *Proceedings of the 38th international conference on software engineering*. 691–701.
- [49] Tomas Mikolov, Kai Chen, G. S. Corrado, and J. Dean. 2013. Efficient Estimation of Word Representations in Vector Space. In *ICLR*.
- [50] Martin Monperrus. 2017. Automatic Software Repair: a Bibliography. *ACM Computing Surveys* 51 (2017), 1–24. <https://doi.org/10.1145/3105906>
- [51] Martin Monperrus. 2020. *The Living Review on Automated Program Repair*. Technical Report hal-01956501. HAL. <https://www.monperrus.net/martin/repair-living-review.pdf>
- [52] H. D. T. Nguyen, D. Qi, A. Roychoudhury, and S. Chandra. 2013. SemFix: Program repair via semantic analysis. In *2013 35th International Conference on Software Engineering (ICSE)*. 772–781. <https://doi.org/10.1109/ICSE.2013.6606623>
- [53] pytorch. 2020. <https://pytorch.org/>. pytorch.
- [54] Yuhua Qi, Xiaoguang Mao, Yan Lei, Ziyang Dai, and Chengsong Wang. 2014. The Strength of Random Search on Automated Program Repair. In *Proceedings of the 36th International Conference on Software Engineering (Hyderabad, India) (ICSE 2014)*. Association for Computing Machinery, New York, NY, USA, 254–265. <https://doi.org/10.1145/2568225.2568254>
- [55] Zichao Qi, Fan Long, Sara Achour, and Martin Rinard. 2015. An Analysis of Patch Plausibility and Correctness for Generate-and-validate Patch Generation Systems (*ISSTA*). 24–36.
- [56] Maxim Rabinovich, Mitchell Stern, and Dan Klein. 2017. Abstract Syntax Networks for Code Generation and Semantic Parsing. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, Regina Barzilay and Min-Yen Kan (Eds.). Association for Computational Linguistics, 1139–1149. <https://doi.org/10.18653/v1/P17-1105>
- [57] André Ribeiro and Rui Abreu. 2010. The GZoltra Project: A Graphical Debugger Interface. 215–218. [https://doi.org/10.1007/978-3-642-15585-7\\_25](https://doi.org/10.1007/978-3-642-15585-7_25)
- [58] Reudismam Rolim, Gustavo Soares, Loris D’Antoni, Oleksandr Polozov, Sumit Gulwani, Rohit Gheyi, Ryo Suzuki, and Björn Hartmann. 2017. Learning syntactic program transformations from examples. In *Proceedings of the 39th International Conference on Software Engineering, ICSE 2017, Buenos Aires, Argentina, May 20-28, 2017*, Sebastián Uchitel, Alessandro Orso, and Martin P. Robillard (Eds.). IEEE / ACM, 404–415. <https://doi.org/10.1109/ICSE.2017.44>
- [59] Reudismam Rolim, Gustavo Soares, Loris D’Antoni, Oleksandr Polozov, Sumit Gulwani, Rohit Gheyi, Ryo Suzuki, and Björn Hartmann. 2017. Learning syntactic program transformations from examples. In *2017 IEEE/ACM 39th International Conference on Software Engineering (ICSE)*. IEEE, 404–415.
- [60] Ripon K. Saha, Yingjun Lyu, Hiroaki Yoshida, and Mukul R. Prasad. 2017. ELIXIR: Effective Object Oriented Program Repair. In *ASE (Urbana-Champaign, IL, USA)*. IEEE Press. <http://dl.acm.org/citation.cfm?id=3155562.3155643>
- [61] Seemanta Saha et al. 2019. Harnessing evolution for multi-hunk program repair. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 13–24.
- [62] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. 2009. The Graph Neural Network Model. *IEEE Transactions on Neural Networks* 20, 1 (2009), 61–80.
- [63] Zeyu Sun, Qihao Zhu, Lili Mou, Yingfei Xiong, Ge Li, and Lu Zhang. 2019. A Grammar-Based Structural CNN Decoder for Code Generation. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*. AAAI Press, 7055–7062. <https://doi.org/10.1609/aaai.v33i01.33017055>
- [64] Zeyu Sun, Qihao Zhu, Yingfei Xiong, Yican Sun, Lili Mou, and Lu Zhang. 2020. TreeGen: A Tree-Based Transformer Architecture for Code Generation. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. AAAI Press, 8984–8991. <https://aaai.org/ojs/index.php/AAAI/article/view/6430>
- [65] Daniel Tarlow, Subhdeep Moitra, Andrew Rice, Zimin Chen, Pierre-Antoine Manzagol, Charles Sutton, and Edward Aftandilian. 2019. Learning to Fix Build Errors with Graph2Diff Neural Networks. *CoRR* abs/1911.01205 (2019). arXiv:1911.01205 <http://arxiv.org/abs/1911.01205>
- [66] M. Tufano, C. Watson, G. Bavota, M. di Penta, M. White, and D. Poshyvanyk. 2018. An Empirical Investigation into Learning Bug-Fixing Patches in the Wild via Neural Machine Translation. In *2018 33rd IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 832–837. <https://doi.org/10.1145/3238147.3240732>
- [67] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, unfiledudukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (Long Beach, California, USA) (NIPS’17)*. Curran Associates Inc., Red Hook, NY, USA, 6000–6010.
- [68] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer Networks. *ArXiv* abs/1506.03134 (2015).
- [69] M. Wen, J. Chen, R. Wu, D. Hao, and S. Cheung. 2018. Context-Aware Patch Generation for Better Automated Program Repair. In *2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE)*. 1–11. <https://doi.org/10.1145/3180155.3180233>
- [70] Qi Xin and Steven P. Reiss. 2017. Leveraging syntax-related code for automated program repair. In *Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 660–670.
- [71] Yingfei Xiong, Bo Wang, Guirong Fu, and Linfei Zang. 2018. Learning to synthesize. In *Proceedings of the 4th International Workshop on Genetic Improvement Workshop*. 37–44.
- [72] Yingfei Xiong, Jie Wang, Runfa Yan, Jiachen Zhang, Shi Han, Gang Huang, and Lu Zhang. 2017. Precise condition synthesis for program repair. In *2017 IEEE/ACM 39th International Conference on Software Engineering (ICSE)*. IEEE, 416–426.
- [73] Jifeng Xuan, Matias Martinez, Favio Demarco, Maxime Clement, Sebastian Lame-las Marcote, Thomas Durieux, Daniel Le Berre, and Martin Monperrus. 2016. Nopol: Automatic repair of conditional statement bugs in java programs. *IEEE Transactions on Software Engineering* 43, 1 (2016), 34–55.
- [74] Michihiro Yasunaga and Percy Liang. 2020. Graph-based, Self-Supervised Program Repair from Diagnostic Feedback. arXiv:2005.10636 [cs.SE]
- [75] Pengcheng Yin and Graham Neubig. 2017. A Syntactic Neural Model for General-Purpose Code Generation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, Regina Barzilay and Min-Yen Kan (Eds.). Association for Computational Linguistics, 440–450. <https://doi.org/10.18653/v1/P17-1041>
- [76] Wenjie Zhang, Zeyu Sun, Qihao Zhu, Ge Li, Shaowei Cai, Yingfei Xiong, and Lu Zhang. 2020. NLocalSAT: Boosting Local Search with Solution Prediction. *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (Jul 2020)*. <https://doi.org/10.24963/ijcai.2020/164>
- [77] Qihao Zhu, Zeyu Sun, Xiran Liang, Yingfei Xiong, and Lu Zhang. 2020. OCoR: An Overlapping-Aware Code Retriever. In *35th IEEE/ACM International Conference on Automated Software Engineering, ASE 2020, Melbourne, Australia, September 21-25, 2020*. IEEE, 883–894. <https://doi.org/10.1145/3324884.3416530>