A Structural Model for Contextual Code Changes

SHAKED BRODY, Technion, Israel

URI ALON, Technion, Israel

ERAN YAHAV, Technion, Israel

We address the problem of predicting *edit completions* based on a learned model that was trained on past edits. Given a code snippet that is partially edited, our goal is to predict a *completion of the edit for the rest of the snippet*. We refer to this task as the EDITCOMPLETION task and present a novel approach for tackling it. The main idea is to directly represent structural edits. This allows us to model the likelihood of the edit itself, rather than learning the likelihood of the edited code. We represent an edit operation as a path in the program's Abstract Syntax Tree (AST), originating from the source of the edit to the target of the edit. Using this representation, we present a powerful and lightweight neural model for the EDITCOMPLETION task.

We conduct a thorough evaluation, comparing our approach to a variety of representation and modeling approaches that are driven by multiple strong models such as LSTMs, Transformers, and neural CRFs. Our experiments show that our model achieves a 28% relative gain over state-of-the-art sequential models and 2× higher accuracy than syntactic models that learn to generate the edited *code*, as opposed to modeling the *edits* directly.

Our code, dataset, and trained models are publicly available at <https://github.com/tech-srl/c3po/>.

CCS Concepts: • **Software and its engineering** → **General programming languages**; • **Computing methodologies** → **Neural networks**.

Additional Key Words and Phrases: Edit Completions, Neural Models of Code, Machine Learning

ACM Reference Format:

Shaked Brody, Uri Alon, and Eran Yahav. 2020. A Structural Model for Contextual Code Changes. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 215 (November 2020), 28 pages. <https://doi.org/10.1145/3428283>

1 INTRODUCTION

Software development is an evolutionary process. Programs are being maintained, refactored, fixed, and updated on a continuous basis. Program *edits* are therefore at the very core of software development. Poor edits can lead to bugs, security vulnerability, unreadable code, unexpected behavior, and more. The ability to suggest a good edit in code is therefore crucial.

We introduce the EDITCOMPLETION task: predict *edit completions* based on a learned model that was trained on past edits. Given a code snippet that is partially edited, our goal is to predict an *edit completion* that *completes the edit for the rest of the snippet*. The edit completion is represented technically as a sequence of edit operations that we refer to as an *edit script*.

Authors' addresses: Shaked Brody, Technion, Israel, shakedbr@cs.technion.ac.il; Uri Alon, Technion, Israel, urialon@cs.technion.ac.il; Eran Yahav, Technion, Israel, yahave@cs.technion.ac.il.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2020 Copyright held by the owner/author(s).

2475-1421/2020/11-ART215

<https://doi.org/10.1145/3428283>

(a) The predicate of the if statement in C was edited to include a null check for `attack`. Thus, in P , the checking of `attack != null` and the ternary operator can be removed.

(b) The signature of `GetFileCharacteristics` in C was edited to return a `FileCharacteristic` object instead of modifying an output parameter. Thus, in P , the method should return a `FileCharacteristic` object instead of returning `true`.

Fig. 1. Examples of **EDITCOMPLETION**. The input consists of a program fragment $\mathcal{P}$ and edits that occurred in the context that transformed C into C' . The output is $\Delta_{\mathcal{P}}$ – an edit script that describes the likely edit. Applying $\Delta_{\mathcal{P}}$ to $\mathcal{P}$ results in $\mathcal{P}'$ – the code after the edit.

Problem Definition. Let $\mathcal{P}$ be a given program fragment and C be the surrounding context of $\mathcal{P}$ before any edits were applied. Let Δ_C denote the edits that were applied to C , and $C' = \Delta_C(C)$ the resulting edited context. The goal in our **EDITCOMPLETION** task is to predict an edit function $\Delta_{\mathcal{P}}$, such that applying $\Delta_{\mathcal{P}}$ to $\mathcal{P}$ results in the program fragment after the edit: $\Delta_{\mathcal{P}}(\mathcal{P}) = \mathcal{P}'$. Our underlying assumption is that the distribution of edits in $\mathcal{P}$ can be inferred from the edits Δ_C that occurred in its context. We thus model the probability: $Pr(\Delta_{\mathcal{P}} \mid \Delta_C)$. We present a new approach for representing and predicting $\Delta_{\mathcal{P}}$ in the **EDITCOMPLETION** task, named C^3 : Contextual Code Changes.

Motivating Examples. Consider the **EDITCOMPLETION** examples in Figure 1a and Figure 1b. These illustrate the significance of edits in the context C and how they can help in suggesting a likely edit for $\mathcal{P}$. In Figure 1a, the edit in the context consists of changing the if statement predicate,

Fig. 2. An example of two edits. These examples are different and the edits operate on different values. However, observing the structure of these edits reveals the similarity between them and allows a learning model to generalize better. This similarity is expressed as almost identical AST paths. For simplicity, only the program fragment that should be edited $\mathcal{P}$ is shown, without the context $\mathcal{C}$.

resulting in a null check for the variable `attack`. After the edit in the context, the value of `attack` in $\mathcal{P}$ cannot be null. Therefore, the ternary statement that checks `attack` for nullness in $\mathcal{P}$ can be removed. Our model successfully predicted the needed edit $\Delta_{\mathcal{P}}$, which is applied to $\mathcal{P}$ to yield $\mathcal{P}'$.

Figure 1b shows another example, in which the edit in the context is a modification of a function signature. In $\mathcal{C}'$, the return type was changed to `FileCharacteristics`, and the output parameter `fileCharacteristics` for the function was removed. $\mathcal{P}$ consists of an assignment to the parameter `fileCharacteristics`, and a return statement of `true` value. The edit in the context implies a necessary edit in $\mathcal{P}$, in which the assignment statement has to be removed (since `fileCharacteristics` is no longer defined) and the return statement must include a variable of type `FileCharacteristics`. Our model successfully predicted the correct edit for $\mathcal{P}$. $\mathcal{P}'$ consists of returning an object of type `FileCharacteristics`.

Edit Completion vs. Code Completion. It is important to note that `EDITCOMPLETION` and code completion are completely different tasks. The goal of code completion is to predict missing fragments of a program, given a partial program as context. In contrast, the goal of `EDITCOMPLETION` is to predict additional edits in a partial sequence of edit operations. That is, while code completion operates on code, `EDITCOMPLETION` operates on *code edits*.

Representing Code Edits. The main design decision in learning code edits is *how to represent the edit*, i.e., how to represent the difference between the code in its original form and its desired, altered, form. Naïvely, differencing programs can be performed by treating the code as text and using *text-diff* algorithms for line differencing [Hunt and McIlroy 1975] or inline differencing

[Birney et al. 1996]. In contrast, we model the *difference between the abstract syntax trees* (ASTs) of the original and the edited code. This allows us to naturally use paths in the AST (AST paths) to model edits.

Our Approach. We present a novel approach for EDITCOMPLETION: predicting contextual code changes – C^3 . Code changes can be described as a sequence of edit operations, such as “move a node, along with its underlying subtree, to be a child of another node” or “update the value of a node to be identical to the value of another node”. Such edit operations can be naturally represented as paths between the source node and the target node, along with the relationship between them and the edit command, i.e., “move” or “update”. AST paths provide a natural way to express binary relationships between nodes (and thus subtrees) in the AST. We use AST paths to represent Δ_C – edits that occurred in the context and transformed C into C' , such that $\Delta_C(C) = C'$. We also use AST paths to represent Δ_P – the edits that should be applied to P . We thus model the probability $Pr(\Delta_P | \Delta_C)$, where both the input Δ_C and the output Δ_P are represented as AST paths.

Representing edits as paths allows a learning model to generalize well across different examples. Consider the two examples in Figure 2. In Figure 2a, the edit modifies a series of LINQ calls, converting `Where(<predicate>).FirstOrDefault()` into `FirstOrDefault(<predicate>)`. The edit in Figure 2b modifies `Where(<predicate>).First()` into `First(<predicate>)`. We elaborate on the representation of edits as paths in Section 2 and Section 4. For now, it suffices to note that there is a sequence of three edit operations in each of the figures (numbered ①, ②, ③). Although the predicates are different and these edits operate on different values, the structure of the edits in Figure 2a and Figure 2b is identical. This similarity is expressed in the AST paths that represent these edits. For example, consider the identical structure of the path ① in the two figures, where it operates on a different value in each figure (`FirstOrDefault` and `First`).

Our use of AST paths allows the model to generalize these edits, even though these edits are not identical and their predicates are different.

We apply a Pointer Network [Vinyals et al. 2015] to point to paths in the AST of P and create an edit operation sequence, i.e., an edit script. While prior work used AST paths to *read* programs and predict a label [Alon et al. 2019a,c], we generate an edit script by *predicting* AST paths, i.e., making AST paths *the output* of our model.

Previous Approaches. In related tasks, such as bug fixing and program repair, previous approaches have mostly represented code as a flat token stream [Chen et al. 2019; Tufano et al. 2018; Vasic et al. 2019]; although this allows the use of NLP models out-of-the-box, such models do not leverage the rich syntax of programming languages.

Yin et al. [2019] suggested a system that learns to represent an edit and uses its representation to apply the edit to another code snippet. Although it sounds similar, the task that Yin et al. [2019] addressed and our task are dramatically different. Yin et al. [2019] addressed the (easier) variant and assume that the edit that needs to be applied is *given* as part of the input. This is done in the form of “before” and “after” versions of another code with *the same edit applied*; their task is only to apply the given edit on a given code. Thus, in the task of Yin et al. [2019], the assumption is that $\Delta_C = \Delta_P$. In contrast, we do not assume that the edit Δ_P is given; we condition on edits that occurred in the context (Δ_C), but these edits are different than the edits that need to be applied to P . Our model needs to predict the edit to P itself, i.e., *predict what needs to be edited and how*. Other work did use syntax but did not represent the structure of the edit itself. Dinella et al. [2020] proposed a model for detecting and fixing bugs using graph transformations, without considering context changes (i.e., $\Delta_C = \emptyset$). Their method can predict *unary* edit operations on the AST. In

contrast, we predict *binary* edit operations. Thus, our representation is much more expressive. For example, consider the edit of moving a subtree. This edit can be represented as a *single binary* operation; alternatively, this edit can require *multiple unary* operations.

Modeling Code Likelihood vs. Modeling Edit Likelihood. In general, there are two main approaches for learning to edit a given code snippet. Assume that we wish to model the probability of a code snippet $\mathcal{Y}$ given another code snippet $\mathcal{X}$. Much prior work [Chen et al. 2019; Mesbah et al. 2019] followed the approach of generating $\mathcal{Y}$ directly. Attempting to model $\mathcal{Y}$ given $\mathcal{X}$ modeled the probability $Pr(\mathcal{Y} | \mathcal{X})$. This approach is straightforward, but requires modeling the likelihood of $\mathcal{Y}$, which is a problem that is more difficult than necessary. In contrast, it can be much more effective to model *the likelihood of the edit*, which transforms $\mathcal{X}$ into $\mathcal{Y}$, without modeling the likelihood of $\mathcal{Y}$ itself; hence, $Pr(\Delta_{\mathcal{X} \rightarrow \mathcal{Y}} | \mathcal{X})$. Our modeling of the *edit* follows the latter approach: $Pr(\Delta_{\mathcal{P}} | \Delta_{\mathcal{C}})$. In this work, we learn to predict the *edit* ($\Delta_{\mathcal{P}}$) that transforms $\mathcal{P}$ into $\mathcal{P}'$, instead of predicting the entire program ($\mathcal{P}'$). By applying $\Delta_{\mathcal{P}}$ to $\mathcal{P}$, generating $\mathcal{P}'$ is straightforward: $\Delta_{\mathcal{P}}(\mathcal{P}) = \mathcal{P}'$. Learning to predict *the edit* instead of learning to predict *the edited code* makes our learning task much easier and provides much higher accuracy, as we show in Section 6.

We show the effectiveness of C^3 on EDITCOMPLETION on a new dataset, scraped from over 300,000 commits in GitHub.

Our approach significantly outperforms textual and syntactic approaches that either model the code or model only the edit, and are driven by strong neural models.

Contributions. The main contributions of this paper are:

- We introduce the EDITCOMPLETION task: given a program $\mathcal{P}$ and edits that occurred in its context, predict the likely edits that should be applied to $\mathcal{P}$.
- C^3 – a novel approach for representing and predicting contextual edits in code. This is the first approach that represents structural edits directly.
- Our technique directly captures the relationships between subtrees that are changed in an edit using *paths* in the AST. The output of our technique is an edit script that is executed to edit the program fragment $\mathcal{P}$.
- A prototype implementation of our approach, called C^3PO , for Contextual Code Changes via Path Operations. C^3PO is implemented using a strong neural model that predicts the likely edit by pointing to an AST path that reflects that edit.
- A new EDITCOMPLETION dataset of source code edits and their surrounding context edits, scraped from over 300,000 commits in GitHub.
- An extensive empirical evaluation that compares our approach to a variety of representation and modeling approaches, driven by strong models such as LSTMs, Transformers, and neural CRFs. Our evaluation shows that our model achieves over 28% relative gain over state-of-the-art strong sequential models, and over 2× higher accuracy than syntactic models that do not model edits directly.
- A thorough ablation study that examines the contribution of syntactic and textual representations in different components of our model.

2 OVERVIEW

In this section, we demonstrate our approach using a simple EDITCOMPLETION example. The main idea is to represent all valid edit operations in $\mathcal{P}$ as AST paths, and predict a sequence of these paths. Since every path is associated with an edit operation, by pointing to a sequence of paths, we, in fact, predict an edit script.

```
+ public virtual Navigation AddNavigation(string name, ForeignKey foreignKey, bool pointsToPrincipal)
```

(a)

(b)

(c)

(d)

 $\mathcal{P} \rightarrow \mathcal{P}'$

```
- productType.AddNavigation(  
  new Navigation(  
    featuredProductFk,  
    "FeaturedProductCategory",  
    pointsToPrincipal: false);  
+ productType.AddNavigation(  
  "FeaturedProductCategory",  
  featuredProductFk,  
  pointsToPrincipal: false);
```

(e)

Fig. 3. An EDITCOMPLETION example from our test set. Figure 3a shows the edit that transforms $\mathcal{C}$ into $\mathcal{C}'$ – overloading the function AddNavigation. Figure 3e shows $\mathcal{P}$ and $\mathcal{P}'$ as code in red and green, respectively. Figure 3b depicts the partial AST and the first three edit operations of the edit. Figure 3c shows the AST after applying the first three operations, and shows the next three operations as AST paths. Figure 3d illustrates the AST after performing all operations, resulting in an AST that corresponds to $\mathcal{P}'$. Every edit operation is represented by an AST path having the same color and number as the edit command. Dotted contours represent subtrees that will be affected by applying these operations. Already-affected subtrees are surrounded by dashed contours.

2.1 Motivating Example

High-level Overview. Consider the edit that occurred in the context of Figure 3a – insertion of a new definition of the method AddNavigation, which overloads previous definitions. After applying this edit, it is possible to use this new signature when calling AddNavigation. Consider the original code snippet $\mathcal{P}$ at the top of Figure 3e. The edit in the context allows us to simplify the call to AddNavigation using the new signature, as shown in the “edited” code snippet $\mathcal{P}'$ at the bottom of Figure 3e. Consider the partial AST of $\mathcal{P}$ in Figure 3b. The desired edit can be described as an

edit script consisting of six edit operations to the AST of $\mathcal{P}$. Consider the first operation: ① **MOV**. The meaning of this operation is to move the node Expr with its subtree to be the leftmost child of the node Unit. This edit operation can be represented by the **red** ① path: Expr $\rightarrow$ Arg $\rightarrow$ ArgList $\rightarrow$ Call $\rightarrow$ Expr $\rightarrow$ Unit. Note how this path directly captures the syntactic relationship between the node Expr and the node Unit, allowing our model to predict a **MOV** operation as part of the edit script.

In Figure 3c we can see the result of applying the following first three operations: ① **MOV**, ② **MOV**, ③ **MOV**, moving subtrees to new locations in the tree. The last three commands are **DEL** operations, expressing deletion of a node and its underlying subtree. These operations can be represented using paths as well. For instance, ④ **DEL** is represented by the **green** ④ path: Navigation $\rightarrow$ Call $\rightarrow$ Expr $\rightarrow$ Unit $\rightarrow$ DEL, where DEL is an artificial node that we add as a child of the AST's root. In Figure 3d we can see the AST after applying all six operations. After executing all six operations, our model produces $\mathcal{P}'$, shown in Figure 3e.

Path Extraction. To inform the model about the available edits it can use for prediction, we parse the AST of $\mathcal{P}$ to extract all AST paths that represent valid edits. Every path can represent different edit “commands” that use the same path. For example, consider the **blue** ② path in Figure 3b: Name $\rightarrow$ Call $\rightarrow$ ArgList $\rightarrow$ Arg $\rightarrow$ Expr $\rightarrow$ Call. This path can represent a move operation – **MOV**, i.e., moving the node Name with its subtree, to be the leftmost child of Call; alternatively, this path can represent an insertion operation – **INS**, i.e., copy Name with its subtree, and insert it as the leftmost child of Call. To distinguish between different edit operations that are represented using the same AST path, each path is encoded as a vector once, and projected into three vectors using different learned functions. Each resulting vector corresponds to a different kind of edit operation. For example, the **orange** ③ path in Figure 3b can represent either “move” (**MOV**), “update” (**UPD**) or “insert” (**INS**) operations. In this case, this path was projected using the learned function that represents “move”.

Edit Script Prediction. We predict one edit operation at each step by *pointing* at a path and its associated operation from among the valid edit operations. This results in an *edit script*. For example, in Figure 3, our model finds that the **red** ① path with **MOV** is most likely to be the first operation. Then, given this edit, our model finds that the **blue** ② path with **MOV** is most likely to be the next operation, and so on, until we predict a special “end of sequence” (EOS) symbol.

Modeling Code Likelihood vs. Modeling Edit Likelihood. Modeling edits using AST paths provides an effective way to model *only the difference* between $\mathcal{P}$ and $\mathcal{P}'$. For example, consider the **red** ① path that moves the subtree rooted at Expr from its original place to be the first child of Unit. To predict this edit, our model only needs to select the **red** ① path out of the other available operations. In contrast, a model that attempts to generate $\mathcal{P}'$ entirely [Chen et al. 2019], would need to generate the entire subtree from scratch in the new location.

Pairwise Edit Operations. Most edit operations, such as “move” and “update”, can be described as *pairwise* operations, having the “source” and the “target” locations as their two arguments. AST paths provide a natural way to represent pairwise relations, originating from the “source” location, and reaching the “target” location through the shortest path between them in the tree. In contrast, prior work that used only *unary* edit operations such as HOPPITY [Dinella et al. 2020] are limited to inserting each node individually, and thus use *multiple* edit commands to express the ① **MOV** operation. Our model represents this edit operation as a single AST path – the **red** ① path.

Key aspects. The example in Figure 3 demonstrates several key aspects of our method:

- Edits applied to the context of $\mathcal{P}$ can provide useful information for the required edit to $\mathcal{P}$.

- Pairwise edit operations can be naturally represented as AST paths.
- A neural model, trained on these paths, can generalize well to other programs, thanks to the direct modeling of code edits as paths.
- By *pointing* at the available edit operations, the task that the model addresses becomes *choosing* the most likely edit, rather than generating $\mathcal{P}'$ from scratch, and thus significantly eases the learning task.

3 BACKGROUND

In this section, we provide the necessary background. First we define abstract syntax trees (ASTs) and AST paths in Section 3.1. In Section 3.2, we use these definitions to describe how to represent code edits using AST paths and perform AST differencing. Finally, in Section 3.3 and Section 3.4, we describe the concept of *attention* and *pointer networks*, which are crucial components in our neural architecture (described in Section 5).

3.1 AST Paths

Given a programming language $\mathcal{L}$ and its grammar, we use V to denote the set of *nonterminals*, and T to denote the set of *terminals* in the grammar. The Abstract Syntax Tree (AST) of a program can be constructed in the standard manner, defined as follows:

Definition 3.1. (Abstract Syntax Tree) Given a program $\mathcal{P}$ written in a programming language $\mathcal{L}$, its Abstract Syntax Tree $\mathcal{A}$ is the tuple $(A, B, r, X, \delta, \varphi)$, where A is the set of non-leaf nodes, such that each $n \in A$ is of type that belongs to V ; B is the set of leaves such that each $n \in B$ is of type that belongs to T ; $r \in A$ is the root of the tree; X is a set of values taken from $\mathcal{P}$; δ is a function $\delta : A \rightarrow (A \cup B)^*$ that maps *nonterminals* nodes to their children; φ is a mapping $\varphi : B \rightarrow X$ that maps a *terminal* node to a value.

An AST path is simply a sequence of nodes in the AST, formally:

Definition 3.2. (AST Path) Given an AST $\mathcal{A} = (A, B, r, X, \delta, \varphi)$, an AST path is a sequence of nodes $p = n_1, n_2, \dots, n_k$, where $n_i \in A \cup B$, such that for every consecutive pair of nodes n_i and n_{i+1} , either $n_i \in \delta(n_{i+1})$ or $n_{i+1} \in \delta(n_i)$. We follow Alon et al. [2018] and associate each node's *child index* with its type.

For example, consider the blue ② path in Figure 3b. The path starts in the node `Name`, goes up to its parent node `Call`, then goes down to its right-most child `ArgList`, and so on.

AST paths are a natural way to describe relationships between nodes in the AST, and can serve as a general representation of relationships between elements in programs. For example, Alon et al. [2018, 2019c] used paths between leaves in the AST as a way to create an aggregated representation of the AST.

In this work, we use AST paths to model relationships between arbitrary nodes in the tree (both terminals and nonterminals) to model the effect of edit operations.

3.2 AST Differencing

An edit in a program can be represented as a sequence of operations on its AST. To compute the difference between two programs, we compute the difference between the two ASTs using algorithms such as GumTree [Falleri et al. 2014]. Given two programs $\mathcal{P}$ and $\mathcal{P}'$, along with their ASTs $\mathcal{A}$ and $\mathcal{A}'$, GumTree outputs an *edit script*, consisting of instructions to change $\mathcal{A}$ so it becomes $\mathcal{A}'$. Each operation in the script is either **MOV**, **DEL**, **UPD**, or **INS** and operates on one or two nodes. The command **MOV** n_s, n_t stands for moving a subtree inside the AST. This operation takes the source node n_s to be moved and the target n_t node, which will be the left sibling of n_s after the

Fig. 4. Example of AST edit operations. Figure 4a depict the AST before the change. Figure 4b shows the result of **MOV** operation – moving C to be the right sibling of D. Figure 4c shows the result of **DEL** – removing C. Figure 4d shows the result of **UPD** – updating C to Z. Figure 4e shows the result of **INS** – inserting E to be the right sibling of D.

move. The command **DEL** n_s stands for removing the node n_s from the tree. We use the command **UPD** v, n_t , to update the value of the node n_t to become v . Lastly, to represent insertion, we use **INS** n_s, n_t , where n_s is the root of a subtree to be inserted and n_t is the target node that will be the left sibling of n_s after the insertion.

Figure 4 demonstrates all operations: Figure 4a illustrates the AST before the edits; Figure 4b shows the result of **MOV** C, D; Figure 4c depicts the command **DEL** C; Figure 4d shows the update of C to the value Z, i.e., **UPD** Z, C; Figure 4e illustrates the command **INS** E, D – the insertion of node E as a right sibling of D.

In general, AST differencing algorithms consist of two steps. The first step maps nodes from $\mathcal{A}$ to $\mathcal{A}'$, where each node belongs to a single mapping at most and mapped nodes share the same type. The second step uses the mapping and aims to produce a short edit script. The GumTree algorithm focuses on the first step of mapping, since there are known quadratic optimal algorithms [Chawathe et al. 1996] for the second step.

GumTree [Falleri et al. 2014] breaks the mapping stage into three steps. The first step is a top-down algorithm that finds isomorphic subtrees across $\mathcal{A}$ and $\mathcal{A}'$. The roots of these subtrees are called *anchors mapping*. The second step is a bottom-up algorithm that looks for *containers mapping*; these are node pairs among $\mathcal{A}$ and $\mathcal{A}'$, such that their descendants share common *anchors*. Finally, the last step looks for additional mappings between the descendants of the *containers mapping* pairs.

Applying the GumTree algorithm for the mapping stage and using known techniques to produce the edit script results in an end-to-end efficient algorithm. The complexity of this algorithm is $O(n^2)$ in the worst case, where n is the number of nodes in the larger among $\mathcal{A}$ and $\mathcal{A}'$, i.e., $n = \max(|\mathcal{A}|, |\mathcal{A}'|)$.

3.3 Attention

An attention mechanism computes a learned weighted average of some input vectors, given another input *query* vector. Usually, attention is used by a neural model to align elements from different modalities. For example, in neural machine translation (NMT) [Bahdanau et al. 2014], attention allows the model to “focus” on different words from the source language while predicting every word in the target language, by computing a different weighted average at every step. This ability has shown a significant improvement across various tasks such as translation [Bahdanau et al. 2014; Luong et al. 2015; Vaswani et al. 2017], speech recognition [Chan et al. 2016], and code summarization and captioning [Alon et al. 2019a].

Formally, given a set of k vectors $Z = z_1, z_2, \dots, z_k \in \mathcal{R}^d$ (usually, an encoding of the input of the model) and a query vector $q \in \mathcal{R}^d$ (usually, the hidden state of a decoder at a certain time step t), attention performs the following computation. The first step computes a “score” for each input vector z_i . For example, Luong et al. [2015] use a learned matrix $W_a \in \mathcal{R}^{d \times d}$ to compute the score s_i of the vector z_i :

$$s_i = z_i \cdot W_a \cdot q^\top \quad (1)$$

Next, all scores are normalized into a pseudo-probability using the softmax function:

$$\alpha_i = \frac{e^{s_i}}{\sum_{j=1}^k e^{s_j}} \quad (2)$$

where every normalized score is between zero and one $\alpha_i \in [0, 1]$, and their sum is one: $\sum \alpha_i = 1$. Then, a *context vector* is computed as a weighted average of the inputs $z_1, z_2, \dots, z_k$, such that the weights are the computed weights α :

$$c = \sum_i^k \alpha_i \cdot z_i$$

This dynamic weighted average can be computed iteratively at different prediction time steps t , producing different attention scores α_t and thus a different context vector c_t . This offers a decoder the ability to focus on different elements in the encoded inputs at each prediction step.

3.4 Pointer Networks

A *pointer network* [Vinyals et al. 2015] is a variant of the seq2seq paradigm [Sutskever et al. 2014], where the output sequence is a series of *pointers* to the encoded inputs, rather than a sequence from a separate vocabulary of symbols. This mechanism is especially useful when the output sequence is composed *only* of elements from the input, possibly permuted and repeated. For example, the problem of *sorting* a sequence of numbers can be naturally addressed using pointer networks: the input for the model can be the unsorted sequence of numbers, and the output is the sorted sequence, where every output prediction is a *pointer* to an element in the input sequence.

Pointing can be performed in a manner similar to attention: at each decoding step, Equation (1) and Equation (2) compute input scores, similar to attention. Then, the resulting normalized scores α_i can be used for classification over the encoded inputs, as the output probability of the model.

Pointer networks and attention share almost the same implementation, but they are different in principle. Attention computes a dynamic average c_t at each decoding iteration. Then, c_t is used in the prediction of this time step, among a different closed set of possible classes. For example, the possible classes can be the words in the target language. In pointer networks, on the other hand, the possible classes at each decoding step are the elements *in the input sequence itself*.

Another difference is that in pointer networks there is a label associated with each “pointing” step. Each “pointing” distribution α is directly supervised by computing a cross-entropy loss with a reference label. In other words, each pointing can be measured for its correctness, and the mostly-pointed input is either correct or incorrect. In contrast, attention is not directly supervised; the model’s attention distribution α is internal to the model. The attention distribution α is usually neither “correct” nor “incorrect”, because the attention is used for a follow-up prediction.

4 REPRESENTING EDITS WITH AST PATHS

In the EDITCOMPLETION task that we consider in this work, the input contains multiple edit operations that occurred in the context, and the output is a series of edit operations that should be performed. The main challenge is *how to represent edits in a learning model*? We look for a representation that is *expressive* and *generalizable*. The representation should be *expressive*, such that different edits

are reflected differently; this would allow a model to consider the difference between examples. However, just representing every edit uniquely is not enough, because the representation should also be *generalizable*, such that similar edits would be reflected similarly. This would allow a model to generalize better, even if the edit that should be predicted at test time does not look exactly like an edit that was observed at training time.

Representing edits using AST paths provides an expressive and generalizable solution. An edit operation, such as “move”, can be represented as the path in the AST from the subtree that should be moved, up to its new destination. This path includes the *syntactic relation* between the source and the target of the move. Different move operations would result in different paths and thus, this representation is expressive. Similar moves will result in similar paths and thus, this representation is generalizable. In this section, we explain how AST paths can naturally represent such edit operations.

We represent edit operations as follows:

- (1) The **MOV** (move) operation has two arguments: the first is the source node – the root of the subtree to be moved, and the second is the target node. The meaning of “**MOV** n_s, n_t ” is that node n_s moves to be the right sibling of node n_t . To support moving a node to be the leftmost child, we augment the AST with Placeholder nodes, that are always present as the leftmost child nodes of all *nonterminal* nodes.
- (2) The **UPD** (update) operation has two arguments: the first argument is a node with a source value, and the second argument is a node whose value needs to be updated. For instance, if the value of node n_t needs to be updated to x , and the value of node n_s is x , we denote this by: “**UPD** n_s, n_t ”.
- (3) The **INS** (insert) operation has two arguments: the first argument is the subtree to be copied, and the second is the target node. The operation “**INS** n_s, n_t ” means that the subtree rooted at n_s should be copied and inserted as the right sibling of n_t . If n_s should be inserted as a leftmost child, the target node will be the appropriate Placeholder node.
- (4) The **DEL** (delete) operation has one argument, which is a subtree to be deleted. We represent **DEL** as a path that originates from the root of the subtree to be deleted, into a special DEL target node that we artificially add as a child of the AST’s root. So in practice, we represent “**DEL** n_s ” as “**MOV** n_s, n_{DEL} ” where n_{DEL} is the DEL node.

Since all four operations can be represented using two nodes n_s and n_t from the AST of $\mathcal{P}$, the AST path from n_s to n_t is a natural way to represent an edit operation. Figure 5 demonstrates a **MOV** operation and its associated path representation. Figure 5a depicts the path $\text{Arg}_1 \rightarrow \text{ArgList} \rightarrow \text{Arg}_3$, which can be associated with **MOV** and represent the operation **MOV** $\text{Arg}_1, \text{Arg}_3$, i.e., moving the first argument to be the last. Figure 5b shows the AST after the movement.

To represent insertions (**INS**) and updates (**UPD**) in the *context* that transformed C into C' , we augment the AST with additional UPD and INS nodes. To represent all update operations **UPD** n_s, n_t , we add the necessary n_s nodes as children of UPD. For example, in Figure 6, there are two update operations that involve two source nodes: y and Bar . Thus, we add these nodes as children of UPD and represent the operations with paths that originate from these nodes. The orange ① path, for instance, represents the update of Foo to become Bar . In the case of insertion of a new subtree, we represent this operation with a path that originates from INS and ends in the root of the subtree. Consider the purple ③ path in Figure 6. This path represents that the subtree with the root Name was inserted as the leftmost child of Type . We augment the AST with additional UPD and INS nodes as additional children of the AST’s root, along with the special DEL node.

Fig. 7. A high-level overview of our architecture. On the left, the partial AST that represents the context C . The red paths represent the transformation from C to C' . On the right, we can see the partial AST of P and its paths that represent possible valid predictions. The model attends to the paths that transform C to C' to point to a path of P that corresponds to an edit operation.

5.1 High-Level View

At a high-level, our model reads the edits that occurred in the context and predicts edits that should be performed in the program. Since there may be multiple edits in the context, our model uses *attention* to compute their dynamic weighted average. To predict the edit in the program, our model enumerates all possible edits, expresses them as AST paths, and *points* to the most likely edit. Thus, the input of the model is a sequence of AST paths from the augmented C , and the output is a sequence of AST paths from P . Our model is illustrated in Figure 7.

Our model follows the encoder-decoder paradigm. The encoder encodes all valid paths of the input code (P) and the paths of the input context (transforming C to C') into continuous vectors. The decoder generates an edit sequence by pointing to the set of paths in P while attending to the paths of C and C' . First, to consider edits that occurred in the context, our model encodes the sequence of context paths that transformed C into C' , as a set of vectors. Then, the model performs a series of predictions, where each such prediction is an edit that should be applied to P . At each prediction step, the model *attends* (as explained in Section 3.3) to the context paths. Using the resulting attention vector, the model *points* (as explained in Section 3.4) to a single path in P . The path in P that the model points to, is translated to the edit that should be applied in this step. In the next step, the chosen edit from the previous step is used to compute the attention query of the next step.

An edit operation that occurred in the context can be represented as an AST path (Section 4). We denote the sequence of paths that represent the edits in the context as $\Delta_C = \text{Paths}(C, C')$. The edit function that should be predicted is also represented as AST paths, where each path is associated with an edit operation. We denote the sequence of AST paths that represent the edits that should be predicted as $\Delta_P = \text{Paths}(P, P')$; we use these vectors as our *classes* from which we make predictions.

Using the above notations, we model the conditional probability: $Pr(\Delta_P \mid \Delta_C)$.

5.2 Encoder

Given a sequence of paths $\text{Paths}(C, C')$, we encode all paths using a Path Encoder (Section 5.2.1). Then, since it is a sequence that has a meaningful order, the context paths go through an LSTM [Hochreiter and Schmidhuber 1997], resulting in the sequence of vectors Z_C .

We enumerate all valid edits that can be applied to $\mathcal{P}$ and denote this set as $Paths(\mathcal{P}, \overline{\mathcal{P}})$. We then encode these paths using the Path Encoder (Section 5.2.1), which results in the set of vectors $Z_{\overline{\mathcal{P}}}$. Every path in $Paths(\mathcal{P}, \overline{\mathcal{P}})$ can represent different edit operations (Section 5.2.2), i.e., both “update” and “move”. Thus, every path vector $z \in Z_{\overline{\mathcal{P}}}$ is projected to represent different edit operations, resulting in the set of vectors Z_{Op} , which represent the set of classes from which the model can predict.

5.2.1 Path Encoder. Given a set of AST paths, our goal is to create a vector representation z_i for each path $v_1 \dots v_l$. The vocabulary of nodes of the AST is limited to a fixed-size vocabulary from the grammar of the language. In contrast, the *values* of AST leaves correspond to the tokens in the textual representation of the program. Therefore, the vocabulary of these tokens is unbounded. To address the issue of the unbounded vocabulary of terminal values, we follow previous work [Allamanis et al. 2015, 2016; Alon et al. 2019a], and split these values into *subtokens*. For example, the value `toString` will be split into `to` and `string`. We represent each path as a sequence of node types using an LSTM, and use subtoken embeddings to represent terminal values (the tokens).

Node Representation. Each AST path is composed of nodes $v_1, \dots, v_l$. Each node is taken from a limited vocabulary of 88 symbols of the programming language. Terminal nodes also have a user-defined token value. Every node has an associated child index, i.e., its index among its sibling nodes [Alon et al. 2018]. We represent each node using a learned embedding matrix E^{nodes} and a learned embedding matrix for its child indices E^{index} . We sum the vector of the node type w with the vector of its child index i to represent the node:

$$encode_node(w) = E_i^{index} + E_w^{nodes}$$

The first and the last node of an AST path may be terminals whose values are tokens in the code.

¹ We use a learned embedding matrix $E^{subtokens}$ to represent each subtoken:

$$encode_value(w) = \sum_{s \in split(w)} E_s^{subtokens} \quad (3)$$

where w is a value associated with a terminal node.

Path Representation. We encode the path $v_1, \dots, v_l$ by applying an LSTM:

$$h_1, \dots, h_l = LSTM_{path}(encode_node(v_1), \dots, encode_node(v_l))$$

We concatenate the last state vector with an encoding of the values associated with the first and the last nodes in the path, and then pass them through a learned fully connected layer W_{path} and a nonlinearity:

$$encode_path(v_1 \dots v_l) = \tanh(W_{path} \cdot [h_l; encode_value(\varphi(v_1)); encode_value(\varphi(v_l))])$$

where φ is the function that retrieves a terminal node’s associated value (Section 3.1). If v_1 or v_l are nonterminals, and thus do not have an associated value, we encode the first and the last nodes instead of their values; i.e., $encode_node(v)$ instead of $encode_value(\varphi(v))$.

To express the order of context paths $Paths(C, C')$, we pass these through another LSTM:

$$Z_C = LSTM_C(PathEncoder(Paths(C, C'))) \quad (4)$$

Applying the path encoder on $Paths(\mathcal{P}, \overline{\mathcal{P}})$ results in $Z_{\mathcal{P}}$:

$$Z_{\mathcal{P}} = PathEncoder(Paths(\mathcal{P}, \overline{\mathcal{P}}))$$

¹As opposed to code2vec [Alon et al. 2019c], our paths can originate from and end in nonterminal nodes.

5.2.2 Operation Encoder. To represent different operations (i.e., **MOV**, **UPD**, **INS**) that share the same path $z \in Z_{\mathcal{P}}$, we project z using different learned matrices W_{MOV} , W_{UPD} , W_{INS} :

$$z_{MOV} = z \cdot W_{MOV} \quad z_{UPD} = z \cdot W_{UPD} \quad z_{INS} = z \cdot W_{INS}$$

such that z_{MOV} , z_{UPD} , and z_{INS} are used for pointing to **MOV**, **UPD**, and **INS** edits, which are all described by the same encoded path z . This creates our set of possible classes to point to:

$$Z_{Op} = \bigcup_{z \in Z_{\mathcal{P}}} \{z_{MOV}, z_{UPD}, z_{INS}\} \quad (5)$$

We use Z_{Op} as the representations of the *classes* over which our model outputs a distribution.

5.3 Decoder

The decoder generates an edit script given the outputs of the encoder. At each decoding time step, the decoder predicts a single edit operation, by pointing to a single vector from Z_{Op} , while attending to the sequence of vectors Z_C . The decoder consists of three main components: an LSTM [Hochreiter and Schmidhuber 1997], attention [Bahdanau et al. 2014], and a pointer [Vinyals et al. 2015].

The decoder LSTM operates by receiving an input vector at each time step; then, it uses this input vector to update the LSTM's internal state, and uses the updated state as the query for attention. Given the current state, we compute an attention vector of the vectors in Z_C , and use the resulting vector to point to a (prediction) vector in Z_{Op} . In the next time step, the input vector for the LSTM is the last pointed to vector from the previous step. The initial hidden state of the LSTM is an elementwise average of paths in $Z_{\mathcal{P}}$ and in Z_C .

Attention. We employ attention as described in Section 3.3, where the query is h_t – the hidden state of the decoder LSTM at time step t . At each time step, we compute a scalar score for every vector $z_i \in Z_C$. This score is computed by performing a dot product between each context vector $z_i \in Z_C$ and a learned matrix W_a and h_t . We then normalize all scores with a softmax function to get the normalized weights α^t :

$$\alpha^t = \text{softmax}(Z_C \cdot W_a \cdot h_t^\top)$$

We then compute a weighted average of Z_C to get the attention vector c_t

$$c_t = \sum_{z_i \in Z_C} \alpha_i \cdot z_i$$

Pointing. Given the vector c_t , we compute a pointing score for each valid edit that is represented as $z \in Z_{Op}$. The resulting scores are normalized using softmax; these normalized scores constitute the model's output distribution.

We perform a dot product of every $z \in Z_{Op}$ with another learned weight matrix W_p and c_t . This results in a scalar score for every valid prediction in Z_{Op} . We then apply a softmax, resulting in a distribution over the vectors in Z_{Op} :

$$\hat{y}_t = \text{softmax}(Z_{Op} \cdot W_p \cdot c_t) \quad (6)$$

We use this distribution $\hat{y}_t$ as the model's prediction at time step t . At training time, we train all learnable weights to maximize the log-likelihood [Rubinstein 1999] of $\hat{y}_t$ according to the true label. At test time, we compute the *argmax* of $\hat{y}_t$ to get the prediction: our model predicts the edit operation that is correlated with the element having the highest pointing score. The output of the decoder across time steps can be (unambiguously) translated to an edit script.

Table 1. Statistics over our dataset.

	Train	Validation	Test
# projects	38	8	7
# examples	39504	4468	5934
Avg. number of paths	474	514	405
Avg. number of edit operations	2.6	2.5	2.7
Avg. number of MOV	36.4%	38.3%	41.4%
Avg. number of DEL	48.6%	50%	50.8%
Avg. number of INS	5%	4.5%	2.8%
Avg. number of UPD	10.1%	7.3%	5%
Avg. size of moved subtrees (MOV)	3.48	2.95	2.85
Avg. size of deleted subtrees (DEL)	4.49	4.97	4.39
Avg. size of inserted subtrees (INS)	1.27	2.09	1.26

6 EXPERIMENTS

We implemented our approach for `EDITCOMPLETION` in a neural model called **C³PO**, short for Contextual Code Changes via Path Operations. The main contributions of our approach are (a) the syntactic representation of code edits; and (b) modeling of the likelihood of code edits, rather than modeling the likelihood of the edited code. Thus, these are the main ideas that we wish to evaluate. We compare our model with baselines that represent each of the different paradigms (Table 2) on a new dataset. Our model shows significant performance improvement over the baselines.

6.1 Dataset

We introduce a new `EDITCOMPLETION` dataset of code edits in C#. We scraped the 53 most popular C# repositories from GitHub and extracted all commits since the beginning of the project's history. From each commit, we extracted edits in C# files along with the edits in their surrounding context. Note that a given edit can be part of the edits in the surrounding context (Δ_C) of one example and can be the edit to be predicted (Δ_P) of another example. In other words, the same edit can have different roles in different examples. We verified that both examples reside in the same split (i.e., either both examples are in the training set, or both examples are in the test set), without leakage between the sets.

For each edit, we considered a context radius of 10 lines, above and 10 lines below the edit. Representing long sequences is computationally difficult for baselines that use Transformers [Vaswani et al. 2017] because Transformers have a quadratic time and space complexity. We thus limited the context to 10 lines before and after the edit for these baselines. To make a fair comparison, we limited this in our model as well. We filtered out examples having more than 50 nodes in the AST of $\mathcal{P}$. Choosing 50 nodes at most captured the vast majority of examples (81%). While the technique works for any number of nodes, we picked a limit of 50 to keep the time and cost of experiments reasonable.

To make the task even more challenging, we filtered out examples for which: (a) the edit in $\mathcal{P}$ consists of only **DEL** operations; and (b) edits that both $\mathcal{P}$ and its context contain only **UPD** operations such that all updates in $\mathcal{P}$ are included in the updates of C , since these usually reflect simple renaming that is easily predicted by modern IDEs. Following recent work on the adverse effects of code duplication [Allamanis 2019; Lopes et al. 2017], we split the dataset into training-validation-test *by project*. This resulted in a dataset containing 39.5K/4.4K/5.9K train/validation/test set examples,

Table 2. A high-level taxonomy of our model and the baselines.

	Textual	Syntactic
Code Likelihood	SequenceR [Chen et al. 2019]	Path2Tree [Aharoni and Goldberg 2017]
Edit Likelihood	LaserTagger+CRF [Malmi et al. 2019]	C ³ PO (this work)

respectively. We trained all models and baselines on the training set, performed tuning and early-stopping using the validation set, and report final results on the test set. Table 1 shows a summary of the statistics of our dataset. A list of the repositories we used to create our dataset are shown in Appendix A. We make our new dataset publicly available at <https://github.com/tech-srl/c3po/>.

6.2 Baselines

The two main contributions of our approach that we wish to examine are: (a) the syntactic representation of code edits; and (b) modeling edit likelihood, rather than modeling code likelihood. Since we define the new task of `EDITCOMPLETION`, we picked strong neural baselines and adapted them to this task, to examine the importance of these two main contributions.

Table 2 shows a high-level comparison of our model and the baselines. Each model can be classified across two properties: whether it uses a syntactic or textual representation of the edit, and whether it models the *likelihood of the code* or models the *likelihood of the edit*. We put significant effort into performing a fair comparison to all baselines, including subtoken splitting as in our model, lowercasing the subtokens, and replacing generated UNK tokens with the tokens that were given the highest attention score.

LaserTagger [Malmi et al. 2019] - is a textual model that models the *edit likelihood*. LaserTagger learns to apply textual edits to a given text. The model follows the framework of *sequence tagging*, i.e., classifying each token in the input sequence. Each input token is classified into one of: *KEEP_φ*, *DELETE_φ* and *SWAP*, where ϕ belongs to a vocabulary of all common phrases obtained from the training set. While LaserTagger leverages edit operations, it does not take advantage of the syntactic structure of the input. Since the original implementation of LaserTagger uses a pre-trained BERT NLP model, which cannot be used for code, we carefully re-implemented a model in their spirit, without BERT. We used the same preprocessing scripts and sequence tags as Malmi et al. [2019], and encoded the input using either a bidirectional LSTM or a Transformer [Vaswani et al. 2017] (LaserTagger_{LSTM} and LaserTagger_{Transformer}, respectively). We further strengthened these models with neural Conditional Random Fields (CRFs) [Ma and Hovy 2016]. To represent context edits, we employed a sequence alignment algorithm [Birney et al. 1996] and extracted the textual edits. We encoded these context edits using a bidirectional LSTM and concatenated the resulting vector to the model's encoded input.

SequenceR is a re-implementation of Chen et al. [2019]. SequenceR follows the sequence-to-sequence paradigm from Neural Machine Translation (NMT) with attention [Luong et al. 2015] and a copy mechanism [Gu et al. 2016]. The input is the subtokenized code snippet, along with the textual edits in the context. The output is the edited code. Hence, this method does not take advantage of syntax or edit operations. We carefully re-implemented this approach because SequenceR *abstracts away* identifier names, and replaces identifier names with generic names. For example `int x = 0` becomes `int varInt = 0`. Since our model uses identifier names and we found that identifier names help our model, to perform a fair comparison – we kept identifier

names in SequenceR as well. While the original SequenceR uses LSTMs with copy and attention (SequenceR_{LSTM}), our re-implementation allowed us to strengthen this baseline by replacing the LSTM with a Transformer [Vaswani et al. 2017] and a copy mechanism (SequenceR_{Transformer}). We evaluated both SequenceR_{LSTM}, which follows the original model of Chen et al. [2019], and the strengthened SequenceR_{Transformer} baseline.

Path2Tree follows Aharoni and Goldberg [2017]. This baseline leverages the syntax and models the code likelihood. In this baseline, we performed a pre-order traversal of the AST and represented the AST as a serialized sequence of nodes. Using this sequential serialization of the AST, we could employ strong neural seq2seq models. The input consists of the paths that represent edits in the context (as in our model), along with a serialized sequence that represents the AST of $\mathcal{P}$. The output of the model is the sequence that represents the AST of $\mathcal{P}'$. As the neural underlying seq2seq model, we used both a Transformer (Path2Tree_{Transformer}) with a copy mechanism and a BiLSTM with attention and copy mechanisms (Path2Tree_{LSTM}).

6.3 Setup

From each sample in our dataset, we (a) extracted all paths of $Paths(\mathcal{P}, \overline{\mathcal{P}})$ that describe possible valid edit operations; and (b) extracted the paths that represent the transformation of C to C' , i.e., $Paths(C, C')$. We did not filter, discard any of these paths, or limited the paths lengths.

We used input embedding dimensions of 64, LSTM cells with a single layer, and 128 units. This resulted in a very lightweight model of only 750K learnable parameters. We trained our model on a Tesla V100 GPU using the Adam optimizer [Kingma and Ba 2014] with a learning rate of 0.001 to minimize the cross-entropy loss. We applied a dropout [Hinton et al. 2012] of 0.25.

In the baselines, we used BiLSTMs with 2 layers having an embedding and hidden state of size 512; this resulted in 10M learned parameters in SequenceR_{LSTM} and in Path2Tree_{LSTM} resulting in 10M learned parameters. We used the original hyperparameters of the Transformer [Vaswani et al. 2017] to train Transformers in SequenceR_{Transformer} and Path2Tree_{Transformer}, resulting in 45M learned parameters. LaserTagger_{LSTM} uses BiLSTMs with 2 layers having a hidden state size of 128 and an embedding size of 64. This model contained 1M learned parameters. For LaserTagger_{Transformer}, we used 4 layers of Transformer encoders, with 4 layers and 8 attention heads, an embedding size of 64, and a hidden state size of 512. For both, the context encoders use BiLSTMs with 2 layers having a hidden state size of 128 and an embedding size of 64. We experimented with LaserTaggers that contain a context encoder that uses Transformer and setups that contained larger dimensions, but they achieved slightly lower results. In the other baselines, larger dimensions contributed positively to the performance.

Evaluation Metric. To perform a fair comparison across all examined models, we had to use a metric that would be meaningful and measurable in all models and baselines. We thus measured exact-match accuracy across all models and baselines. The accuracy of each model is the percentage of examples in the test set for which the entire target sequence was predicted correctly.

6.4 Results

Performance. Table 3 depicts the main results of our evaluation: C³PO gains more than 11% absolute accuracy over LaserTagger_{Transformer}, which performed the best of all baselines. These results emphasize the need for structural representation of both edits and context. C³PO achieves accuracy that is twice that of the syntactic baseline Path2Tree. Although this baseline uses AST paths to represent the changes in the context of $\mathcal{P}$ and to represent $\mathcal{P}$ with its underlying AST, its performance is inferior compared to our C³PO. This is because Path2Tree does not model the edit operations directly and thus needs to generate the entire AST of $\mathcal{P}'$.

Table 3. Our model achieves significantly higher accuracy than the baselines.

Model	Acc	Parameters	Training Time
SequenceR _{LSTM} [Chen et al. 2019] +copy	30.7	10M	14h
SequenceR _{Transformer} [Chen et al. 2019] +copy	32.6	45M	20h
LaserTagger _{LSTM} [Malmi et al. 2019] +CRF	40.9	1M	10h
LaserTagger _{Transformer} [Malmi et al. 2019] +CRF	41.4	1.6M	20h
Path2Tree _{LSTM} [Aharoni and Goldberg 2017] +copy	22.5	10M	24h
Path2Tree _{Transformer} [Aharoni and Goldberg 2017] +copy	25.6	45M	24h
C ³ PO (this work)	53.2	750K	9h

These results show the significance of our model’s two main contributions. *Modeling the edit* has the most significant contribution as expressed in the advantage of our model over both versions of Path2Tree, and in the advantage of both versions of LaserTagger over both versions of SequenceR. Syntactic representation over textual representation also has a significant contribution, which is expressed in the superiority of our model over both versions of LaserTagger. Using these two key contributions, our model performs significantly better than all models, while being much more lightweight in terms of learnable parameters. The same results are visualized in Figure 8.

Fig. 8. Visualization of the accuracy score of our model compared to the baselines. The values are the same as in Table 3. Our model achieves significantly higher accuracy than the baselines.

Path2Tree Lower Performance Compared to SequenceR. Although Path2Tree represents the edits syntactically (which we believe to be a better representation, in general) and SequenceR represents edits textually, the results of Path2Tree are lower than those of SequenceR.

We believe that the main limitation of Path2Tree is that it cannot easily generalize between the program fragment $\mathcal{P}$ and the given context $\mathcal{C}$. This occurs because $\mathcal{C}$ is represented as paths, while $\mathcal{P}'$ is generated as a tree. In contrast, SequenceR represents all inputs and outputs the same, i.e., as sequences of tokens. We also performed initial experiments with a Tree2Tree baseline that encodes $\mathcal{C}$, $\mathcal{C}'$, and $\mathcal{P}$ as trees and generates $\mathcal{P}'$ as a tree, thus potentially having a better generalization ability than Path2Tree. However, Tree2Tree achieved much lower results than Path2Tree, because the tree encoding created very large inputs, especially in $\mathcal{C}$. These, prevented the model from properly capturing the edits that occurred in the context ($\Delta_{\mathcal{C}}$), while the encoding of $\Delta_{\mathcal{C}}$ as paths is much more succinct and focused (and performed better than Tree2Tree).

Fig. 9. The upper figure depicts the accuracy as a function of the context radius size. The lower figure shows the number of examples per radius size in the test set.

6.5 Scalability Analysis

We conducted an analysis of our model that shows the performance of C^3PO as a function of the context radius size and the number of nodes.

Figure 9 shows the accuracy of C^3PO compared to the context radius size, i.e., the number of lines between the beginning of C and $\mathcal{P}$. As shown, the accuracy of C^3PO remains stable when the context radius increases. This hints that the context radius can be further increased without sacrificing accuracy. In our experiments, we put this limitation only to limit the size of the dataset.

Figure 10 shows the accuracy of C^3PO compared to the number of nodes in the AST of $\mathcal{P}$. As the size of $\mathcal{P}$ increases, our model shows a natural descent, and the accuracy stabilizes for sizes of 31 nodes and above. As shown in the lower part of Figure 10, the number of examples also decreases with the size of the edit: the most common edits have 11 to 15 nodes, in which our model achieves an accuracy of 80%.

6.6 Qualitative Analysis

We manually examined the predicted examples and discuss two representative cases.

Figure 11 shows an example in which the modification of a method signature in the context affects $\mathcal{P}$, which lies in the method body. The context of $\mathcal{P}$, shown in Figure 11a, includes a change in the signature of the method `GetFileCharacteristic`. The name of the method was changed to `GetAppender` and its return type was updated from `FileCharacteristic` to `BaseFileAppender`.

Consider $\mathcal{P}$ in Figure 11b. $\mathcal{P}$ is a return statement, located in the body of the changed method `GetFileCharacteristic`. Since the return type of the method was updated to `BaseFileAppender`, the return statements inside the method must be changed as well. The renaming of the method to

Fig. 10. The upper figure depicts the accuracy as a function of the number of nodes in $\mathcal{P}$. The lower figure shows the number of examples compared to the number of nodes in $\mathcal{P}$.

GetAppender may have also hinted to our model that the appender object itself should be returned. Our model successfully predicted the desirable edit, altering the return statement from `return appender.GetFileCharacteristic` to `return appender;`. This example shows how context edits are important in predicting edits in a program, by providing information about (a) return type changes and (b) method renaming.

Figure 12 illustrates a case where the edit in the context is conceptually similar to the edit in $\mathcal{P}$, but is not identical. Figure 12a shows a variable declaration statement, where part is cast to the type `MethodCallExpression` and assigned to the newly-declared variable `methodExpression`. In the edited context, the keyword `var` was updated to an explicit type `MethodCallExpression`. Figure 12b shows an edit that is similar in spirit: $\mathcal{P}$ consists of an initialization statement, where the variable `nameParts` is assigned a new `Stack<string>`. Using the edit in the context, our model predicted the edit of `var` to `Stack<string>` in $\mathcal{P}$. This edit consists of an insertion of a new subtree, since `Stack<string>` is represented as a subtree of five nodes. In contrast, the edit in the context is represented as an **UPD** edit, because it only needs to update the value of a single node. This example demonstrates a class of examples where the edit in the context hints edits that are similar in spirit in $\mathcal{P}$, but are not identical and should be performed differently.

7 ABLATION STUDY

We conducted an extensive ablation study to examine the importance of different components in our model. We focused on two axes: the representation of $\Delta_{\mathcal{P}}$ and the representation of $\Delta_{\mathcal{C}}$. This allowed us to examine the origin for the advantage of our model over the strongest baselines, to understand whether it comes from the syntactic representation of the *context* or the syntactic representation of $\mathcal{P}$.

Fig. 11. An example where the edit of a method signature affects the edit of $\mathcal{P}$ which lies in the method body. Figure 11a illustrates the edit in the context and the paths that describe the transformation from $\mathcal{C}$ to $\mathcal{C}'$. Figure 11b shows the predicted edit operations along with their associated paths in $\mathcal{P}$.

Table 4. Variations on our model. † marks results that are copied from Table 3.

	No Context	Textual Context	Path-Based Context
Textual $\mathcal{P}$	35.5	41.4 [†]	39.5
Path-Based $\mathcal{P}$	46.5	48.5	53.2 [†]
			C ³ (this work)

In our model, $\mathcal{P}$ is represented using its syntactic structure, i.e., a path-based representation. Alternatively, $\mathcal{P}$ can be represented using its textual representation. The representation of $\mathcal{P}$ determines the representation of $\mathcal{P}'$. They must be represented similarly, otherwise the model would need to “translate” $\mathcal{P}$ into a different representation to predict $\mathcal{P}'$. However, the representation of the context $\mathcal{C}$ can theoretically be different than that of $\mathcal{P}$.

We thus took our model and examined different representations of the context: path-based context (as in our original model), textual context, and “no context”. For each type of context representation, we also experimented with different types of representations for $\mathcal{P}$: syntactic representation, as in our original model, and textual representation of $\mathcal{P}$. For textual representation of $\mathcal{P}$ we used LaserTagger_{Transformer} [Malmi et al. 2019], which we found to be the strongest textual baseline in Section 6. All the hybrid models were re-trained, and their performance is shown in Table 4.

Fig. 12. An example in which the edit in the context is conceptually similar to the edit of $\mathcal{P}$. Figure 12a illustrates the edit that occurred in the context and the paths that describe the transformation from $\mathcal{C}$ to $\mathcal{C}'$. Figure 12b shows the predicted edit operations along with their associated paths in $\mathcal{P}$.

Contribution of Context. According to our observations, the contribution of the changes in the context is considerable, for both the textual and path-based representations of $\mathcal{P}$. Ignoring changes in the context (the left “No Context” column of Table 4) results in lower accuracy. This motivates our task of predicting edits *given* the context. Program edits are correlated with edits that occurred in the context and predicting edits should consider the context edits.

$\mathcal{P}$ Representation. We observed that across all different settings of context representation, a syntactic representation of $\mathcal{P}$ performs better than a textual representation of $\mathcal{P}$. That is, even if the context is textual (the right column of Table 4), a model benefits from a syntactic representation of $\mathcal{P}$. This advantage is even clearer in the case of “No context”, where the path-based representation of $\mathcal{P}$ achieves more than 10% absolute accuracy over the textual representation of $\mathcal{P}$. A path-based representation of $\mathcal{P}$ allows us to model the edit in $\mathcal{P}$ directly, which makes the learning task much easier and more generalizable.

Context Representation. As Table 4 shows, the representation of the context *should be compatible* with the representation of $\mathcal{P}$. If $\mathcal{P}$ is textual, a textual context performs better; if $\mathcal{P}$ is syntactic, a syntactic context performs better. We hypothesize that matching the context representation to the program representation allows the model to better utilize the context and makes it easier to model the correlation between edits occurring in the context to edits that should be applied to $\mathcal{P}$.

8 RELATED WORK

Representing Programs in Learning Models. The representation of programs in learning models is a question that is even more imperative than the learning algorithm we employed or neural architecture. In the last few years, several approaches have been proposed. Early work used the straightforward textual representation, essentially learning from the flat token stream [Allamanis et al. 2016; Chen et al. 2019; Iyer et al. 2016; Tufano et al. 2018; Vasic et al. 2019]. Although this leverages NLP learning approaches, such models do not leverage the rich syntax of programming languages, and eventually perform worse than other representations, despite their use of strong NLP models. Another line of work represent programs as graphs. These usually augment the AST of a program with additional semantic edges and use a graph neural network to learn from the resulting graph [Allamanis et al. 2018; Brockschmidt et al. 2019; Fernandes et al. 2019; Hellendoorn et al. 2020; Yin et al. 2019]. Graphs provide a natural way to represent programs and allow us to easily augment programs with domain knowledge such as semantic analysis. However, it is unclear how well can these models perform in the absence of full semantic information – given partial code, given code that cannot be compiled, or languages that are difficult to analyze semantically. As in [Alon et al. 2018], we leverage AST paths to represent programs. AST paths were shown to be an effective representation for predicting variable names, method names [Alon et al. 2019c], natural language descriptions [Alon et al. 2019a], and code completion [Alon et al. 2019b]. In our task, AST paths allow us to model edits directly, along with the syntactic relationship between the source and the target node of the edits.

Representing Edits. Much work has been proposed on representing edits. Yin et al. [2019] proposed a model that learns to apply a given code edit on another given code snippet. Although this sounds similar to the task we address; with EDITCOMPLETION there is no specific edit in our input that needs to be applied. In contrast to Yin et al. [2019], the model must predict what should be edited and how, instead of applying a given edit. In our work, there is no guarantee that the edit that needs to be predicted is included in the context. Furthermore, there could be several edits in the context. Thus, our model needs to choose and predict the most likely edit itself, while the edits that occurred in the context may only be related to the edit that needs to be predicted.

SequenceR [Chen et al. 2019] used state-of-the-art NMT models to predict bug fixes on single-line buggy programs. Our work is different from their approach when it comes to the representation of the input and the output. Chen et al. [2019] represent the code as a token stream, while our approach represents edits as AST paths. Further, their approach attempts to generate the entire edited program, whereas our model models only the edit. We demonstrated the advantage of our approach over SequenceR empirically in Section 6.

One problem connected to ours is the task of fixing compilation errors. Tarlow et al. [2019] follows the encoder-decoder paradigm, using an encoder that consists of a graph neural network (GNN) that encodes a multi-graph built from the AST and the compilation error log messages. The decoder is a Transformer [Vaswani et al. 2017] that outputs a sequence representing the predicted edit. DeepDelta [Mesbah et al. 2019] used an NMT model in which the input consists of compilation errors and an AST path from the problematic symbol in the code to the root of the tree. The output of their model is a sequence that represents the edit script. In our work, pairwise AST paths allow us to model the desired edit directly, instead of predicting an edit using multiple predictions.

Recently, Dinella et al. [2020] proposed a model called HOPPITY to detect and fix bugs in source code using graph transformations. The main difference between our approach and theirs is that HOPPITY does not model edit operations directly, as our model does. Rather, it models a graph that represents the input, and uses the resulting node representations to predict actions. This modeling makes their model predict *unary* edit operations, while our model predicts *binary* edits: HOPPITY

can only predict *single-node edits* in each step, such as deleting a subtree root, inserting a *single node*, and changing a *single node* value. Thus, edits like moving large subtrees require multiple insertion operations of a single node at a time. In our approach, moving and inserting a subtree can be performed by a *single* edit operation. Dinella et al. [2020] evaluated their model on examples that contain *three* single-node operations at most. However, as shown in Appendix A, the average size of moved subtrees in our train set is 3.48. Such edits would have required HOPPITY to generate the entire subtree in the new position (three operations) and delete the subtree in its original place (one operation), resulting in four operations in total. Hence, our average case is larger than the cases examined by HOPPITY.

CC2Vec [Hoang et al. 2020] represent edits in version-control-systems (e.g., GitHub). However, their approach represents edits *only textually*. CC2Vec was demonstrated on the tasks of predicting commit messages, predicting bug fixes, and defect prediction; however, their model could not predict the edit itself the way we do in this paper.

Chakraborty et al. [2018] proposed a two-step model that aims to apply edits in code. The first step of their model encodes the sequence that represents a pre-order traversal of the AST of the original code and generates the sequence that represents the AST of the edited code. In the second step, they assign the values of terminal nodes to concrete values. Their approach predicts the edit by synthesizing the entire AST. In Section 6 we showed the advantage of modeling the likelihood of edits over modeling the likelihood of the code. Additionally, our model is trained end-to-end, while Chakraborty et al. [2018] trains different components of their model separately.

9 CONCLUSION

We presented a novel approach for representing and predicting edits in code. Our focus is on learning the likelihood of the edit itself, rather than learning the likelihood of the new program. We use paths from the Abstract Syntax Tree to represent code edits that occurred in a given context, and use these paths to *point* to edits that should be predicted.

We demonstrate the effectiveness of our approach by using the EDITCOMPLETION task to predict edits in a section of code, given edits in its surrounding context. We speculate that our direct modeling of the likelihood of edits, and use of the rich structure of code, are the main components that contribute to the strength of our model. We affirm this conjecture in a thorough evaluation and ablation study. Our method performs significantly better than strong neural baselines that leverage syntax but do not model edits directly, or those that model edits but do not leverage syntax.

We believe our approach can serve as a basis for a variety of models and tools that require the modeling and prediction of code edits. Examples include bug fixing, an EDITCOMPLETION assistant in the programmer’s IDE, and automatically adapting client code to changes in public external APIs. Further, we believe our work can serve as the basis for a future “neural code reviewer”, to save human effort and time. To these ends, we make all our code, dataset, and trained models publicly available at <https://github.com/tech-srl/c3po/>.

A DATASET

Table 5 lists the GitHub repositories we used to create our dataset.

Table 5. Our dataset repositories.

Repository	User	Split
corefx	dotnet	Train
shadowsocks-windows	shadowsocks	Train
CodeHub	CodeHubApp	Train
coreclr	dotnet	Train
roslyn	dotnet	Train
PowerShell	PowerShell	Train
WaveFunctionCollapse	mxgmn	Train
SignalR	SignalR	Train
ShareX	ShareX	Train
Nancy	NancyFx	Train
dapper-dot-net	StackExchange	Train
mono	mono	Train
Wox	Wox-launcher	Train
AutoMapper	AutoMapper	Train
RestSharp	restsharp	Train
BotBuilder	Microsoft	Train
SparkleShare	hbons	Train
Newtonsoft.Json	JamesNK	Train
MonoGame	MonoGame	Train
MaterialDesignInXamlToolkit	MaterialDesignInXAML	Train
ReactiveUI	reactiveui	Train
msbuild	Microsoft	Train
aspnetboilerplate	aspnetboilerplate	Train
orleans	dotnet	Train
Hangfire	HangfireIO	Train
Sonarr	Sonarr	Train
dnSpy	0xd4d	Train
Psychson	brandonlw	Train
acat	intel	Train
SpaceEngineers	KeenSoftwareHouse	Train
PushSharp	Redth	Train
cli	dotnet	Train
StackExchange.Redis	StackExchange	Train
akka.net	akkadotnet	Train
framework	accord-net	Train
monodevelop	mono	Train
Opserver	opserver	Train
ravendb	ravendb	Train
OpenLiveWriter	OpenLiveWriter	Validation
Mvc	aspnet	Validation
GVFS	Microsoft	Validation
OpenRA	OpenRA	Validation
Rx.NET	dotnet	Validation
MahApps.Metro	MahApps	Validation
FluentValidation	JeremySkinner	Validation
ILSpy	icsharpcode	Validation
ServiceStack	ServiceStack	Test
choco	chocolatey	Test
duplicati	duplicati	Test
CefSharp	cefsharp	Test
NLog	NLog	Test
JavaScriptServices	aspnet	Test
EntityFrameworkCore	aspnet	Test

REFERENCES

- Roei Aharoni and Yoav Goldberg. 2017. Towards String-To-Tree Neural Machine Translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Vancouver, Canada, 132–140. <https://doi.org/10.18653/v1/P17-2021>
- Miltiadis Allamanis. 2019. The Adverse Effects of Code Duplication in Machine Learning Models of Code. In *Proceedings of the 2019 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software* (Athens, Greece) (*Onward! 2019*). Association for Computing Machinery, New York, NY, USA, 143–153. <https://doi.org/10.1145/3359591.3359735>

- Miltiadis Allamanis, Earl T. Barr, Christian Bird, and Charles Sutton. 2015. Suggesting Accurate Method and Class Names. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering* (Bergamo, Italy) (ESEC/FSE 2015). ACM, New York, NY, USA, 38–49. <https://doi.org/10.1145/2786805.2786849>
- Miltiadis Allamanis, Marc Brockschmidt, and Mahmoud Khademi. 2018. Learning to Represent Programs with Graphs. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=BJOFETxR->
- Miltiadis Allamanis, Hao Peng, and Charles Sutton. 2016. A Convolutional Attention Network for Extreme Summarization of Source Code. In *Proceedings of The 33rd International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 48)*, Maria Florina Balcan and Kilian Q. Weinberger (Eds.). PMLR, New York, New York, USA, 2091–2100. <http://proceedings.mlr.press/v48/allamanis16.html>
- Uri Alon, Shaked Brody, Omer Levy, and Eran Yahav. 2019a. code2seq: Generating Sequences from Structured Representations of Code. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=H1gKYo09tX>
- Uri Alon, Roy Sadaka, Omer Levy, and Eran Yahav. 2019b. Structural Language Models of Code. arXiv:1910.00577 [cs.LG]
- Uri Alon, Meital Zilberstein, Omer Levy, and Eran Yahav. 2018. A general path-based representation for predicting program properties. *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation - PLDI 2018* (2018). <https://doi.org/10.1145/3192366.3192412>
- Uri Alon, Meital Zilberstein, Omer Levy, and Eran Yahav. 2019c. Code2vec: Learning Distributed Representations of Code. *Proc. ACM Program. Lang.* 3, POPL, Article 40 (Jan. 2019), 29 pages. <https://doi.org/10.1145/3290353>
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* (2014).
- Ewan Birney, Julie D. Thompson, and Toby J. Gibson. 1996. PairWise and SearchWise: Finding the Optimal Alignment in a Simultaneous Comparison of a Protein Profile against All DNA Translation Frames. *Nucleic Acids Research* 24, 14 (07 1996), 2730–2739. <https://doi.org/10.1093/nar/24.14.2730> arXiv:<https://academic.oup.com/nar/article-pdf/24/14/2730/7064078/24-14-2730.pdf>
- Marc Brockschmidt, Miltiadis Allamanis, Alexander L. Gaunt, and Oleksandr Polozov. 2019. Generative Code Modeling with Graphs. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Bke4KsA5FX>
- Saikat Chakraborty, Miltiadis Allamanis, and Baishakhi Ray. 2018. Tree2Tree Neural Translation Model for Learning Source Code Changes. *ArXiv abs/1810.00314* (2018).
- William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals. 2016. Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 4960–4964.
- Sudarshan S. Chawathe, Anand Rajaraman, Hector Garcia-Molina, and Jennifer Widom. 1996. Change Detection in Hierarchically Structured Information. *SIGMOD Rec.* 25, 2 (June 1996), 493–504. <https://doi.org/10.1145/235968.233366>
- Zimin Chen, Steve Kommrusch, Michele Tufano, Louis-Noël Pouchet, Denys Poshyvanyk, and Martin Monperrus. 2019. SequenceR: Sequence-to-Sequence Learning for End-to-End Program Repair. *CoRR abs/1901.01808* (2019). arXiv:1901.01808 <http://arxiv.org/abs/1901.01808>
- Elizabeth Dinella, Hanjun Dai, Ziyang Li, Mayur Naik, Le Song, and Ke Wang. 2020. HOPPITY: LEARNING GRAPH TRANSFORMATIONS TO DETECT AND FIX BUGS IN PROGRAMS. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=SJeqs6EFvB>
- Jean-Rémy Falleri, Floréal Morandat, Xavier Blanc, Matias Martinez, and Martin Monperrus. 2014. Fine-grained and accurate source code differencing. In *ACM/IEEE International Conference on Automated Software Engineering, ASE '14, Vasteras, Sweden - September 15 - 19, 2014*. 313–324. <https://doi.org/10.1145/2642937.2642982>
- Patrick Fernandes, Miltiadis Allamanis, and Marc Brockschmidt. 2019. Structured Neural Summarization. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=H1ersoRqtm>
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O. K. Li. 2016. Incorporating Copying Mechanism in Sequence-to-Sequence Learning. *CoRR abs/1603.06393* (2016). arXiv:1603.06393 <http://arxiv.org/abs/1603.06393>
- Vincent J. Hellendoorn, Charles Sutton, Rishabh Singh, Petros Maniatis, and David Bieber. 2020. Global Relational Models of Source Code. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=B1lnbRNtWr>
- 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>
- Thong Hoang, Hong Jin Kang, Julia Lawall, and David Lo. 2020. CC2Vec: Distributed Representations of Code Changes. arXiv:2003.05620 [cs.SE]
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Comput.* 9, 8 (Nov. 1997), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- James W. Hunt and M. Douglas McIlroy. 1975. An algorithm for differential file comparison.

- Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, and Luke Zettlemoyer. 2016. Summarizing Source Code using a Neural Attention Model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, 2073–2083. <https://doi.org/10.18653/v1/P16-1195>
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. <http://arxiv.org/abs/1412.6980> cite arxiv:1412.6980Comment: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015.
- Cristina V Lopes, Petr Maj, Pedro Martins, Vaibhav Saini, Di Yang, Jakub Zitny, Hitesh Sajjani, and Jan Vitek. 2017. DéjàVu: a map of code duplicates on GitHub. *Proceedings of the ACM on Programming Languages* 1, OOPSLA (2017), 1–28.
- Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective Approaches to Attention-based Neural Machine Translation. *CoRR* abs/1508.04025 (2015). arXiv:1508.04025 <http://arxiv.org/abs/1508.04025>
- Xuezhe Ma and Eduard Hovy. 2016. End-to-end Sequence Labeling via Bi-directional LSTM-CNNs-CRF. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 1064–1074.
- Eric Malmi, Sebastian Krause, Sascha Rothe, Daniil Mirylenka, and Aliaksei Severyn. 2019. Encode, Tag, Realize: High-Precision Text Editing. In *EMNLP-IJCNLP*.
- Ali Mesbah, Andrew Rice, Emily Johnston, Nick Glorioso, and Edward Aftandilian. 2019. DeepDelta: learning to repair compilation errors. In *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 925–936.
- Reuven Rubinfeld. 1999. The cross-entropy method for combinatorial and continuous optimization. *Methodology and Computing in Applied Probability* 1, 2 (1999), 127–190.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to Sequence Learning with Neural Networks. *CoRR* abs/1409.3215 (2014). arXiv:1409.3215 <http://arxiv.org/abs/1409.3215>
- 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. arXiv:1911.01205 [cs.LG]
- Michele Tufano, Cody Watson, Gabriele Bavota, Massimiliano Di Penta, Martin White, and Denys Poshyvanyk. 2018. An Empirical Study on Learning Bug-Fixing Patches in the Wild via Neural Machine Translation. *CoRR* abs/1812.08693 (2018). arXiv:1812.08693 <http://arxiv.org/abs/1812.08693>
- Marko Vasic, Aditya Kanade, Petros Maniatis, David Bieber, and Rishabh Singh. 2019. Neural Program Repair by Jointly Learning to Localize and Repair. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=ByloJ20qtm>
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 5998–6008. <http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer Networks. arXiv:1506.03134 [stat.ML]
- Pengcheng Yin, Graham Neubig, Miltiadis Allamanis, Marc Brockschmidt, and Alexander L. Gaunt. 2019. Learning to Represent Edits. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=BJl6AjC5F7>