
Structured Generative Models of Natural Source Code

Chris J. Maddison[†]

University of Toronto

CMADDIS@CS.TORONTO.EDU

Daniel Tarlow

Microsoft Research

DTARLOW@MICROSOFT.COM

Abstract

We study the problem of building generative models of *natural* source code (NSC); that is, source code written by humans and meant to be understood by humans. Our primary contribution is to describe new generative models that are tailored to NSC. The models are based on probabilistic context free grammars (PCFGs) and neuro-probabilistic language models (Mnih & Teh, 2012), which are extended to incorporate additional source code-specific structure. These models can be efficiently trained on a corpus of source code and outperform a variety of less structured baselines in terms of predictive log likelihoods on held-out data.

1. Introduction

Source code is ubiquitous, and a great deal of human effort goes into developing it. An important goal is to develop tools that make the development of source code easier, faster, and less error-prone, and to develop tools that are able to better understand pre-existing source code. To date this problem has largely been studied outside of machine learning. Many problems in this area do not appear to be well-suited to current machine learning technologies. Yet, source code is some of the most widely available data with many public online repositories. Additionally, massive open online courses (MOOCs) have begun to collect source code homework assignments from tens of thousands of students (Huang et al., 2013). At the same time, the software engineering community has recently observed that it is useful to think of source code as natural—written by humans and meant to be understood by other humans (Hindle et al., 2012). This *natural source code* (NSC) has a great deal of statistical regularity that is ripe for study in machine learning.

The combination of these two observations—the availability of data, and the presence of amenable statistical structure—has opened up the possibility that machine learning tools could become useful in various tasks related to source code. At a high level, there are several potential areas of contributions for machine learning. First, code editing tasks could be made easier and faster. Current autocomplete suggestions rely primarily on heuristics developed by an Integrated Development Environment (IDE) designer. With machine learning methods, we might be able to offer much improved completion suggestions by leveraging the massive amount of source code available in public repositories. Indeed, Hindle et al. (2012) have shown that even simple n -gram models are useful for improving code completion tools, and Nguyen et al. (2013) have extended these ideas. Other related applications include finding bugs (Kremenek et al., 2007), mining and suggesting API usage patterns (Bruch et al., 2009; Nguyen et al., 2012; Wang et al., 2013), as a basis for code complexity metrics (Allamanis & Sutton, 2013), and to help with enforcing coding conventions (Allamanis et al., 2014). Second, machine learning might open up whole new applications such as automatic translation between programming languages, automatic code summarization, and learning representations of source code for the purposes of visualization or discriminative learning. Finally, we might hope to leverage the large amounts of existing source code to learn improved priors over programs for use in programming by example (Halbert, 1984; Gulwani, 2011) or other program induction tasks.

One approach to developing machine learning tools for NSC is to improve specific one-off tasks related to source code. Indeed, much of the work cited above follows in this direction. An alternative, which we pursue here, is to develop a generative model of source code with the aim that many of the above tasks then become different forms of query on the same learned model (e.g., code completion is conditional sampling; bug fixing is model-based de-

Proceedings of the 31st International Conference on Machine Learning, Beijing, China, 2014. JMLR: W&CP volume 32. Copyright 2014 by the author(s).

[†]Work done primarily while author was an intern at Microsoft Research.

```
for (m[canplace] = -0; i[0] <= temp; m++) {
    co.nokori = Dictionary;
    continue;
}
```

(a) Log-bilinear PCFG

```
for (int i = words; i < 4; ++i) {
    for (int j = 0; j < i; ++j) {
        if (words.Length % 10 == 0) {
            Math.Max(j+j, i*2 + Math.Abs(i+j));
        }
    }
}
```

(b) Our Model

Figure 1. Samples of for loops generated by learned models. Our model captures hierarchical structure and other patterns of variable usage and local scoping rules. Whitespace edited to improve readability.

noising; representations may be derived from latent variables (Hinton & Salakhutdinov, 2006) or from Fisher vectors (Jaakkola & Haussler, 1998)). We believe that building a generative model focuses attention on the challenges that source code presents. It forces us to model all aspects of the code, from the high level structure of classes and method declarations, to constraints imposed by the programming language, to the low level details of how variables and methods are named. We believe building good models of NSC to be a worthwhile modelling challenge for machine learning research to embrace.

In Section 2 we introduce notation and motivate the requirements of our models—they must capture the sequential and hierarchical nature of NSC, naturalness, and code-specific structural constraints. In Sections 3 and 4 we introduce Log-bilinear Tree-Traversal models (LTTs), which combine natural language processing models of trees with log-bilinear parameterizations, and additionally incorporate compiler-like reasoning. In Section 5 we discuss how efficiently to learn these models, and in Section 7 we show empirically that they far outperform the standard NLP models that have previously been applied to source code. As an introductory result, Fig. 1 shows samples generated by a Probabilistic Context Free Grammar (PCFG)-based model (Fig. 1 (a)) versus samples generated by the full version of our model (Fig. 1 (b)). Although these models apply to any common imperative programming language like C/C++/Java, we focus specifically on C#. This decision is based on (a) the fact that large quantities of data are readily available online, and (b) the recently released Roslyn C# compiler (MSDN, 2011) exposes APIs that allow easy access to a rich set of internal compiler data structures and processing results.

2. Modelling Source Code

In this section we discuss the challenges in building a generative model of code. In the process we motivate our

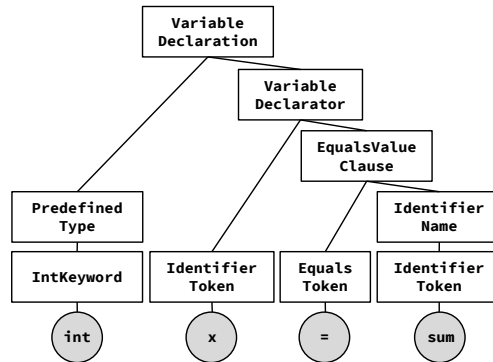


Figure 2. Example Roslyn AST. Rectangles are internal nodes and shaded circles are tokens.

choice of representation and model and introduce terminology that will be used throughout.

Hierarchical Representation. The first step in compilation is to lex code into a sequence of *tokens*, $(\alpha_t)_{t=1}^T = \alpha$. Tokens are strings such as “sum”, “.”, or “int” that serve as the atomic syntactic elements of a programming language. However, representing code as a flat sequence leads to very inefficient descriptions. For example, in a C# for loop, there must be a sequence containing the tokens for, (, an initializer statement, a condition expression, an increment statement, the token), then a body block. A representation that is fundamentally flat cannot compactly represent this structure, because for loops can be nested. Instead, it is more efficient to use the hierarchical structure that is native to the programming language. Indeed, most source code processing is done on tree structures that encode this structure. These trees are called abstract syntax trees (ASTs) and are constructed either explicitly or implicitly by compilers after lexing valid sequences of code. The leaf nodes of the AST are the tokens produced by the lexer. The *internal nodes* $\{n_i\}_{i=1}^N$ are specific to a compiler and correspond to expressions, statements or other high level syntactic elements such as Block or ForStatement. The *children tuple* C_i of an internal node n_i is a tuple of nodes or tokens. An example AST is shown in Fig. 2. Note how the EqualsValueClause node has a subtree corresponding to the code = sum. Because many crucial properties of the source code can be derived from an AST, they are a primary data structure used to reason about source code. For example, the tree structure is enough to determine which variables are in scope at any point in the program. For this reason we choose the AST as the representation for source code and consider generative models that define distributions over ASTs.

Modelling Context Dependence. A PCFG seems like a natural choice for modelling ASTs. PCFGs generate ASTs from the root to the leaves by repeatedly sampling children

Algorithm 1 Sampling from LTTs.

```

1: initialize empty stack  $S$ 
2: sample  $(n, h_0) \sim p(n, h_0)$ 
3: push  $n$  onto  $S$ 
4:  $(i, t) \leftarrow (1, 1)$ 
5: while  $S$  is not empty do
6:   pop the top node  $n$  from  $S$ 
7:   if  $n$  is an internal node then
8:      $n_i \leftarrow n$ 
9:     sample  $h_i \sim p(h_i | h_{i-1})$ 
10:    sample  $C_i \sim p(C_i | n_i, h_i)$ 
11:    push  $n$  for  $n \in \text{REVERSED}(C_i)$  onto  $S$ 
12:     $i \leftarrow i + 1$ 
13:   else
14:      $\alpha_t \leftarrow n$ 
15:      $t \leftarrow t + 1$ 
16:   end if
17: end while
    
```

tuples given a parent node. The procedure recurses until all leaves are tokens producing nodes n_i in a depth-first traversal order and sampling children tuples independently of the rest of the tree. Unfortunately this independence assumption produces a weak model; Fig. 1 (a) shows samples from such a model. While basic constraints like matching of parentheses and braces are satisfied, most important contextual dependencies are lost. For example, identifier names (variable and method names) are drawn independently of each other given the internal nodes of the AST, leading to nonsense code constructions.

The first source of context dependence in NSC comes from the naturalness of software. People have many stylistic habits that significantly limit the space of programs that we might expect to see. For example, when writing nested for loops, it is common to name the outer loop variable i and the inner loop variable j . The second source of context dependence comes from additional constraints inherent in the semantics of code. Even if the syntax is context free, the fact that a program conforms to the grammar does not ensure that a program compiles. For example, variables must be declared before they are used.

Our approach to dealing with dependencies beyond what a PCFG can represent will be to introduce *traversal variables* $\{h_i\}_{i=0}^N$ that evolve sequentially as the nodes n_i are being generated. Traversal variables modulate the distribution over children tuples by maintaining a representation of context that depends on the state of the AST generated so far.

3. Log-bilinear Tree-Traversal Models

LTTs are a family of probabilistic models that generate ASTs in a depth-first order (Algorithm 1). First, the stack is initialized and the root is pushed (lines 1-4). Elements are popped from the stack until it is empty. If an internal node

n_i is popped (line 6), then it is expanded into a children tuple and its children are pushed onto the stack (lines 10-11). If a token α_t is popped, we label it and continue (line 14). This procedure has the effect of generating nodes in a depth-first order. In addition to the tree that is generated in a recursive fashion, traversal variables h_i are updated whenever an internal node is popped (line 9). Thus, they traverse the tree, evolving *sequentially*, with each h_i corresponding to some partial tree of the final AST. This sequential view will allow us to exploit context, such as variable scoping, at intermediate stages of the process (see Section 4).

Algorithm 1 produces a sequence of internal nodes $(n_i)_{i=1}^N$, traversal variables $(h_i)_{i=0}^N$, and the desired tokens $\{\alpha_t\}_{t=1}^T$. It is defined by three distributions: (a) the prior over the root node and traversal variables, $p(n, h)$; (b) the distribution over children tuples conditioned on the parent node and h , denoted $p(C | n, h)$; and (c) the transition distribution for the h s, denoted $p(h_i | h_{i-1})$. The joint distribution over the elements produced by Algorithm 1 is

$$p(n_1, h_0) \prod_{i=1}^N p(C_i | n_i, h_i) p(h_i | h_{i-1}) \quad (1)$$

Thus, LTTs can be viewed as a Markov model equipped with a stack—a special case of a Probabilistic Pushdown Automata (PPDA) (Abney et al., 1999). Because depth-first order produces tokens in the same order that they are observed in the code, it is particularly well-suited. We note that other traversal orders produce valid distributions over trees such as right-left or breadth-first. Because we compare to sequential models, we consider only depth-first.

Parameterizations. Most of the uncertainty in generation comes from generating children tuples. In order to avoid an explosion in the number of parameters for the children tuple distribution $p(C | n, h)$, we use a log-bilinear form. For all distributions other than $p(C | n, h)$ we use a simple tabular parameterization.

The log-bilinear form consists of a real-valued vector representation of (n_i, h_i) pairs, $\mathbf{R}_{con}(n_i, h_i)$, a real-valued vector representation for the children tuple, $\mathbf{R}_{ch}(C_i)$, and a bias term for the children, $\mathbf{b}_{ch}(C_i)$. These are combined via an inner product, which gives the negative energy of the children tuple

$$-E(C_i; n_i, h_i) = \mathbf{R}_{ch}(C_i)^T \mathbf{R}_{con}(n_i, h_i) + \mathbf{b}_{ch}(C_i)$$

As is standard, this is then exponentiated and normalized to give the probability of sampling the children: $p(C_i | n_i, h_i) \propto \exp\{-E(C_i; n_i, h_i)\}$. We take the support over which to normalize this distribution to be the set of children tuples observed as children of nodes of type n_i in the training set.

The representation functions rely on the notion of an R matrix that can be indexed into with objects to look up D dimensional real-valued vectors. R_x denotes the row of the R

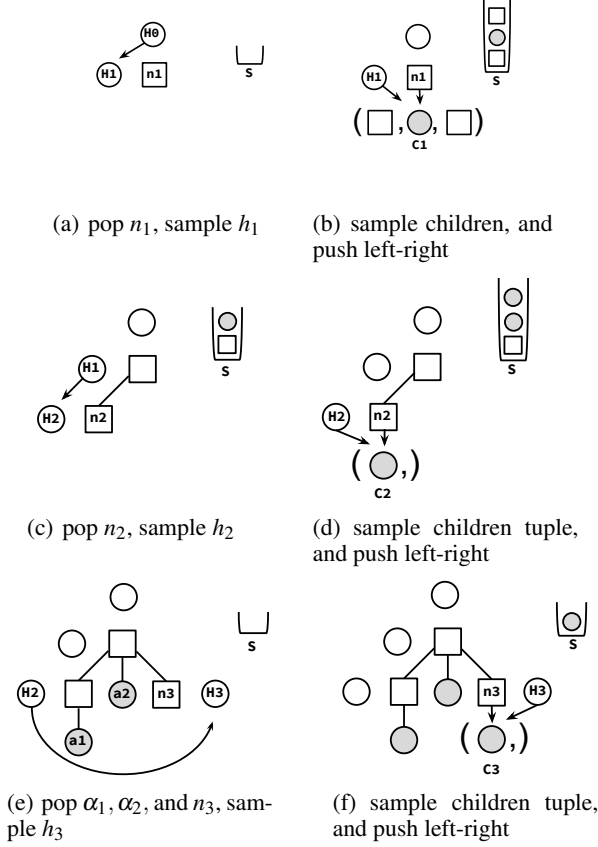


Figure 3. Example sampling run from an LTT. Rectangles are internal nodes, shaded circles are tokens, circles are traversal variables, and stack S is shown in the state after the computations described in subcaption. Parentheses indicate tuples of nodes and arrows indicate conditioning. Popping of tokens omitted for brevity, but note that they are labelled in the order encountered.

matrix corresponding to any variable equal to x . For example, if $n_i = \text{int}$ and $n_j = \text{int}$, then $R_{n_i} = R_{n_j}$. These objects may be tuples and in particular $(\text{type}, \text{int}) \neq \text{int}$. Similarly, b_x looks up a real number. In the simple variant, each unique C sequence receives the representation $R_{ch}(C_i) = R_{C_i}$ and $b_{ch}(C_i) = b_{C_i}$. The representations for (n, h) pairs are defined as sums of representations of their components. If h_i is a collection of variables (h_{ij} representing the j th variable at the i th step) then

$$R_{con}(n_i, h_i) = W_0^{con} R_{n_i} + \sum_{j=1}^H W_j^{con} R_{h_{ij}} \quad (2)$$

The W^{con} s are matrices (diagonal for computational efficiency) that modulate the contribution of a variable in a position-dependent way. In other variants the children tuple representations will also be defined as sums of their component representations. The log-bilinear parameterization has the desirable property that the number of parameters

grows linearly in the dimension of h , so we can afford to have high dimensional traversal variables without worrying about exponentially bad data fragmentation.

4. Extending LTTs

The extensions of LTTs in this section allow (a) certain traversal variables to depend arbitrarily on previously generated elements of the AST; (b) annotating nodes with richer types; and (c) letting R_{ch} be compositionally defined, which becomes powerful when combined with deterministic reasoning about variable scoping.

We distinguish between *deterministic* and *latent* traversal variables. The former can be computed deterministically from the current partial tree (the tree nodes and tokens that have been instantiated at step i) that has been generated while the latter cannot. To refer to a collection of both deterministic and latent traversal variables we continue to use the unqualified “traversal variables” term.

Deterministic Traversal Variables. In the basic generative procedure, traversal variables h_i satisfy the first-order Markov property, but it is possible to condition on any part of the tree that has already been produced. That is, we can replace $p(h_i | h_{i-1})$ by $p(h_i | h_{0:i-1}, n_{1:i}, \alpha_{1:t})$ in Eq. 1. Inference becomes complicated unless these variables are deterministic traversal variables (inference is explained in Section 5) and the unique value that has support can be computed efficiently. Examples of these variables include the set of node types that are ancestors of a given node, and the last n tokens or internal nodes that have been generated. Variable scoping, a more elaborate deterministic relationship, is considered later.

Annotating Nodes. Other useful features may not be deterministically computable from the current partial tree. Consider knowing that a `BinaryExpression` will evaluate to an object of type `int`. This information can be encoded by letting nodes take values in the cross-product space of the node type space and the annotation space. For example, when adding type annotations we might have nodes take value `(BinaryExpression, int)` where before they were just `BinaryExpression`. This can be problematic, because the cardinality of the parent node space increases exponentially as we add annotations. Because the annotations are uncertain, this means there are more choices of node values at each step of the generative procedure, and this incurs a cost in log probabilities when evaluating a model. Experimentally we found that simply annotating expression nodes with type information led to worse log probabilities of generating held out data: the cost of generating tokens decreased because the model had access to type information, the increased cost of generating type annotations along with nodetypes outweighed the improvement.

Identifier Token Scoping. The source of greatest uncertainty when generating a program are children of IdentifierToken nodes. IdentifierToken nodes are very common and are parents of all tokens (e.g. variable and method names) that are not built-in language keywords (e.g., IntKeyword or EqualsToken) or constants (e.g., StringLiterals). Knowing which variables have previously been declared and are currently in scope is one of the most powerful signals when predicting which IdentifierToken will be used at any point in a program. Other useful cues include how recently the variable was declared and what the type the variable is. In this section we a scope model for LTTs.

Scope can be represented as a set of variable feature vectors corresponding to each to a variable that is in scope.¹ Thus, each feature vector contains a string identifier corresponding to the variable along with other features as (key, value) tuples, for example (type, int). A variable is “in scope” if there is a feature vector in the scope set that has a string identifier that is the same as the variable’s identifier.

When sampling an identifier token, there is a two step procedure. First, decide whether this identifier token will be sampled from the current scope. This is accomplished by annotating each IdentifierToken internal node with a binary variable that has the states global or local. If local, proceed to use the local scope model defined next. If global, sample from a global identifier token model that gives support to all identifier tokens. Note, we consider the global option a necessary smoothing device, although ideally we would have a scope model complete enough to have all possible identifier tokens.

The scope set can be updated deterministically as we traverse the AST by recognizing patterns that correspond to when variables should be added or removed from the scope. We implemented this logic for three cases: parameters of a method, locally declared variables, and class fields that have been defined prior in the class definition. We do not include class fields defined after the current point in the code, and variables and methods available in included namespaces. This incompleteness necessitates the global option described above, but these three cases are very common and cover many interesting cases.

Given the scope set which contains variable feature vectors $\{v_\alpha\}$ and parent node (IdentifierToken, local), the probability of selecting token child α is proportional to $p(\alpha | n_i, h_i) \propto \exp\{-E(\alpha; n_i, h_i)\}$, where we normalize only over the variables currently in scope. Specifically, we

¹Technically, we view the scope as a deterministic traversal variable, but it does not contribute to R_{con} .

let $R_{ch}(\alpha)$ and $b_{ch}(\alpha)$ be defined as follows:

$$R_{ch}(\alpha) = \sum_{u=1}^V W_u^{ch} R_{v_{\alpha u}} \quad b_{ch}(\alpha) = \sum_{u=1}^V b_{v_{\alpha u}}. \quad (3)$$

For example, if a variable in scope has feature vector (numNodes, (type, int), (recently-declared, 0)), then its corresponding R_{ch} would be a context matrix-modulated sum of representations $R_{numNodes}$, $R_{(type,int)}$, and $R_{(recently-declared,0)}$. This representation will then be combined with the context representation as in the basic model. The string identifier feature numNodes is the same object as token nodes of the same string, thus they share their representation.

5. Inference and Learning in LTTs

In this section we briefly consider how to compute gradients and probabilities in LTTs.

Only Deterministic Traversal Variables. If all traversal variables h_i can be computed deterministically from the current partial tree, we use the compiler to compute the full AST corresponding to program α_m . From the AST we compute the only valid setting of the traversal variables. Because both the AST and the traversal variables can be deterministically computed from the token sequence, all variables in the model can be treated as observed. Since LTTs are directed models, this means that the total log probability is a sum of log probabilities at each production, and learning decomposes into independent problems at each production. Thus, we can simply stack all productions into a single training set and follow standard gradient-based procedures for training log-bilinear models. More details will be described in Section 7, but generally we follow the Noise-Contrastive Estimation (NCE) technique employed in Mnih & Teh (2012).

Latent Traversal Variables. In the second case, we allow latent traversal variables that are not deterministically computable from the AST. In this case, the traversal variables couple the learning across different productions from the same tree. For simplicity and to allow efficient exact inference, we restrict these latent traversal variables to just be a single discrete variable at each step (although this restriction could easily be lifted if one was willing to use approximate inference). Because the AST is still a deterministic function of the tokens, computing log probabilities corresponds to running the forward-backward algorithm over the latent states in the depth-first traversal of the AST. We can formulate an EM algorithm adapted to the NCE-based learning of log-bilinear models for learning parameters. The details of this can be found in the Supplementary Material.

6. Related Work

The LTTs described here are closely related to several existing models. Firstly, a Hidden Markov Model (HMM) can be recovered by having all children tuples contain a token and a Next node, or just a token (which will terminate the sequence), and having a single discrete latent traversal variable. If the traversal variable has only one state and the children distributions all have finite support, then an LTT becomes equivalent to a Probabilistic Context Free Grammar (PCFG). PCFGs and their variants are components of state-of-the-art parsers of English (McClosky et al., 2006), and many variants have been explored: internal node annotation Charniak (1997) and latent annotations Matsuzaki et al. (2005). Aside from the question of the order of the traversals, the traversal variables make LTTs special cases of Probabilistic Pushdown Automata (PPDA) (for definition and weak equivalence to PCFGs, see Abney et al. (1999)). Log-bilinear parameterizations have been applied widely in language modeling, for n -gram models (Saul & Pereira, 1997; Mnih & Hinton, 2007; Mnih & Teh, 2012) and PCFG models (Charniak, 2000; Klein & Manning, 2002; Titov & Henderson, 2007; Henderson & Titov, 2010). To be clear, our claim is not that general Tree Traversal models or the log-bilinear parameterizations are novel; however, we believe that the full LTT construction, including the tree traversal structure, log-bilinear parameterization, and incorporation of deterministic logic to be novel and of general interest.

The problem of modeling source code is relatively understudied in machine learning. We previously mentioned Hindle et al. (2012) and Allamanis & Sutton (2013), which tackle the same task as us but with simple NLP models. Very recently, Allamanis & Sutton (2014) explores more sophisticated nonparametric Bayesian grammar models of source code for the purpose of learning code idioms. Liang et al. (2010) use a sophisticated non-parametric model to encode the prior that programs should factorize repeated computation, but there is no learning from existing source code, and the prior is only applicable to a functional programming language with quite simple syntax rules. Our approach builds a sophisticated and learned model and supports the full language specification of a widely used imperative programming language.

7. Experimental Analysis

In all experiments, we used a dataset that we collected from TopCoder.com. There are 2261 C# programs which make up 140k lines of code and 2.4M parent nodes in the collective abstract syntax trees. These programs are solutions to programming competitions, and there is some overlap in programmers and in problems across the programs. We created training splits based on the user identity, so the set of users in the test set are disjoint from those in the training

or validation sets (but the training and validation sets share users). The overall split proportions are 20% test, 10% validation, and 70% train. The evaluation measure that we use throughout is the log probability under the model of generating the full program. All logs are base 2. To make this number more easily interpretable, we divide by the number of tokens in each program, and report the average log probability per token.

Experimental Protocol. All experiments use a validation set to choose hyperparameter values. These include the strength of a smoothing parameter and the epoch at which to stop training (if applicable). We did a coarse grid search in each of these parameters and the numbers we report (for train, validation, and test) are all for the settings that optimized validation performance. For the gradient-based optimization, we used AdaGrad (Duchi et al., 2011) with stochastic minibatches. Unless otherwise specified, the dimension of the latent representation vectors was set to 50. Occasionally the test set will have tokens or children tuples unobserved in the training set. In order to avoid assigning zero probability to the test set, we locally smoothed every children distribution with a default model that gives support to all children tuples. The numbers we report are a lower bound on the log probability of data under for a mixture of our models with this default model. Details of this smoothed model, along with additional experimental details, appear in the Supplementary Materials. There is an additional question of how to represent novel identifiers in the scope model. We set the representations of all features in the variable feature vectors that were unobserved in the training set to the all zeros vector.

Baselines and Basic Log-bilinear Models. The natural choices for baseline models are n -gram models and PCFG-like models. In the n -gram models we use additive smoothing, with the strength of the smoothing hyperparameter chosen to optimize validation set performance. Similarly, there is a smoothing parameter in the PCFG-like models that is chosen to optimize validation set performance. We explored the effect of the log-bilinear parameterization in two ways. First, we trained a PCFG model that was identical to the first PCFG model but with the parameterization defined using the standard log-bilinear parameterization. This is the most basic LTT model, with no traversal variables (LTT- \emptyset). The result was nearly identical to the standard PCFG. Next, we trained a 10-gram model with a standard log-bilinear parameterization, which is equivalent to the models discussed in (Mnih & Teh, 2012). This approach dominates the basic n -gram models, allowing longer contexts and generalizing better. Results appear in Fig. 4.

Deterministic Traversal Variables. Next, we augmented LTT- \emptyset model with deterministic traversal variables that in-

Method	Train	Valid	Test
2-gram	-4.28	-4.98	-5.13
3-gram	-2.94	-5.05	-5.25
4-gram	-2.70	-6.05	-6.31
5-gram	-2.68	-7.22	-7.45
PCFG	-3.67	-4.04	-4.23
LTT- \emptyset	-3.67	-4.04	-4.24
LBL 10-gram	-3.19	-4.62	-4.87

Figure 4. Baseline model log probabilities per token.

Method	Train	Valid	Test
LTT- \emptyset	-3.67	-4.04	-4.24
LTT-Seq	-2.54	-3.25	-3.46
LTT-Hi	-2.28	-3.30	-3.53
LTT-HiSeq	-2.10	-3.06	-3.28

Figure 5. LTT models augmented with deterministically determinable latent variables.

clude hierarchical and sequential information. The hierarchical information is the depth of a node, the kind of a node’s parent, and a sequence of 10 ancestor history variables, which store for the last 10 ancestors, the kind of the node and the index of the child that would need to be recursed upon to reach the current point in the tree. The sequential information is the last 10 tokens that were generated.

In Fig. 5 we report results for three variants: hierarchy only (LTT-Hi), sequence only (LTT-Seq), and both (LTT-HiSeq). The hierarchy features alone perform better than the sequence features alone, but that their contributions are independent enough that the combination of the two provides a substantial gain over either of the individuals.

Latent Traversal Variables. Next, we considered latent traversal variable LTT models trained with EM learning. In all cases, we used 32 discrete latent states. Here, results were more mixed. While the latent-augmented LTT (LTT-latent) outperforms the LTT- \emptyset model, the gains are smaller than achieved by adding the deterministic features. As a baseline, we also trained a log-bilinear-parameterized standard HMM, and found its performance to be far worse than other models. We also tried a variant where we added latent traversal variables to the LTT-HiSeq model from the previous section, but the training was too slow to be practical due to the cost of computing normalizing constants in the E step. See Fig. 6.

Scope Model. The final set of models that we trained incorporate the scope model from Section 4 (LTT-HiSeq-Scope). The features of variables that we use are the identifier string, the type, where the variable appears in a list sorted by when the variable was declared (also known as a de Bruijn index), and where the variable appears in a list sorted by when the variable was last assigned a value.

Method	Train	Valid	Test
LTT- \emptyset	-3.67	-4.04	-4.24
LTT-latent	-3.23	-3.70	-3.91
LBL HMM	-9.61	-9.97	-10.10

Figure 6. LTT-latent models augmented with latent variables and learned with EM.

Method	Train	Valid	Test
LTT-HiSeq (50)	-2.10	-3.06	-3.28
LTT-HiSeq-Scope (2)	-2.28	-2.65	-2.78
LTT-HiSeq-Scope (10)	-1.83	-2.29	-2.44
LTT-HiSeq-Scope (50)	-1.54	-2.18	-2.33
LTT-HiSeq-Scope (200)	-1.48	-2.16	-2.31
LTT-HiSeq-Scope (500)	-1.51	-2.16	-2.32

Figure 7. LTT models with deterministic traversal variables and scope model. Number in parenthesis is the dimension of the representation vectors.

Here, the additional structure provides a large additional improvement over the previous best model (LTT-HiSeq). See Fig. 7.

Analysis. To understand better where the improvements in the different models come from, and to understand where there is still room left for improvement in the models, we break down the log probabilities from the previous experiments based on the value of the parent node. The results appear in Fig. 8. In the first column is the total log probability number reported previously. In the next columns, the contribution is split into the cost incurred by generating tokens and trees respectively. We see, for example, that the full scope model pays a slightly higher cost to generate the tree structure than the Hi&Seq model, which is due to it having to properly choose whether IdentifierTokens are drawn from local or global scopes, but that it makes up for this by paying a much smaller cost when it comes to generating the actual tokens.

In the Supplementary Materials, we go further into the breakdowns for the best performing model, reporting the percentage of total cost that comes from the top parent kinds. IdentifierTokens from the global scope are the largest cost (30.1%), with IdentifierTokens covered by our local scope model (10.9%) and Blocks (10.6%) next. This suggests that there would be value in extending our scope model to include more IdentifierTokens and an improved model of Block sequences.

Samples. Finally, we qualitatively evaluate the different methods by drawing samples from the models. Samples of for loops appear in Fig. 1. To generate these samples, we ask (b) the PCFG and (c) the LTT-HiSeq-Scope model to generate a ForStatement. For (a) the LBL n -gram model, we simply insert a for token as the most recent token. We also initialize the traversal variables to reasonable values:

Method	Total	Token	Tree
LTT- \emptyset	-4.23	-2.78	-1.45
LTT-Seq	-3.53	-2.28	-1.25
LTT-Hi	-3.46	-2.18	-1.28
LTT-HiSeq	-3.28	-2.08	-1.20
LTT-HiSeq-Scope	-2.33	-1.10	-1.23

Figure 8. Breakdowns of test log probabilities by whether the cost came from generating the tree structure or tokens. For all models $D = 50$.

e.g., for the LTT-HiSeq-Scope model, we initialize the local scope to include `string[]` words. We also provide samples of full source code files (`CompilationUnit`) from the LTT-HiSeq-Scope model in the Supplementary Material, and additional for loops. Notice the structure that the model is able to capture, particularly related to high level organization, and variable use and re-use. It also learns quite subtle things, like `int` variables often appear inside square brackets.

8. Discussion

Natural source code is a highly structured source of data that has been largely unexplored by the machine learning community. We have built probabilistic models that capture some of the structure that appears in NSC. A key to our approach is to leverage the great deal of work that has gone into building compilers. The result is models that not only yield large improvements in quantitative measures over baselines, but also qualitatively produce far more realistic samples.

There are many remaining modeling challenges. One question is how to encode the notion that the point of source code is to *do something*. Relatedly, how do we represent and discover high level structure related to trying to accomplish such tasks? There are also a number of specific sub-problems that are ripe for further study. Our model of Block statements is naive, and we see that it is a significant contributor to log probabilities. It would be interesting to apply more sophisticated sequence models to children tuples of Blocks. Also, applying the compositional representation used in our scope model to other children tuples would be interesting. Similarly, it would be interesting to extend our scope model to handle method calls. Another high level piece of structure that we only briefly experimented with is type information. We believe there to be great potential in properly handling typing rules, but we found that the simple approach of annotating nodes to actually hurt our models.

More generally, this work’s focus was on generative modeling. An observation that has become popular in machine learning lately is that learning good generative models can be valuable when the goal is to extract features from the data. It would be interesting to explore how this might be

applied in the case of NSC.

In sum, we argue that probabilistic modeling of source code provides a rich source of problems with the potential to drive forward new machine learning research, and we hope that this work helps illustrate how that research might proceed forward.

Acknowledgments

We are grateful to John Winn, Andy Gordon, Tom Minka, and Thore Graepel for helpful discussions and suggestions. We thank Miltos Allamanis and Charles Sutton for pointers to related work.

References

- Abney, Steven, McAllester, David, and Pereira, Fernando. Relating probabilistic grammars and automata. In *ACL*, pp. 542–549. ACL, 1999.
- Allamanis, Miltiadis and Sutton, Charles. Mining source code repositories at massive scale using language modeling. In *MSR*, pp. 207–216. IEEE Press, 2013.
- Allamanis, Miltiadis and Sutton, Charles A. Mining idioms from source code. *CoRR*, abs/1404.0417, 2014.
- Allamanis, Miltiadis, Barr, Earl T, and Sutton, Charles. Learning natural coding conventions. *arXiv preprint arXiv:1402.4182*, 2014.
- Bruch, Marcel, Monperrus, Martin, and Mezini, Mira. Learning from examples to improve code completion systems. In *ESEC/FSE*, pp. 213–222. ACM, 2009.
- Charniak, Eugene. Statistical parsing with a context-free grammar and word statistics. *AAAI/IAAI*, 2005:598–603, 1997.
- Charniak, Eugene. A maximum-entropy-inspired parser. In *ACL*, pp. 132–139, 2000.
- Duchi, John, Hazan, Elad, and Singer, Yoram. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, pp. 2121–2159, 2011.
- Gulwani, Sumit. Automating string processing in spreadsheets using input-output examples. In *ACM SIGPLAN Notices*, volume 46, pp. 317–330. ACM, 2011.
- Halbert, Daniel Conrad. *Programming by example*. PhD thesis, University of California, Berkeley, 1984.
- Henderson, James and Titov, Ivan. Incremental sigmoid belief networks for grammar learning. *JMLR*, 11:3541–3570, 2010.
- Hindle, Abram, Barr, Earl T, Su, Zhendong, Gabel, Mark, and Devanbu, Premkumar. On the naturalness of software. In *ICSE*, pp. 837–847. IEEE, 2012.
- Hinton, Geoffrey E and Salakhutdinov, Ruslan R. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
- Huang, Jonathan, Piech, Chris, Nguyen, Andy, and Guibas, Leonidas. Syntactic and functional variability of a million code submissions in a machine learning MOOC. In *AIED*, 2013.

- Jaakkola, Tommi and Haussler, David. Exploiting generative models in discriminative classifiers. In Kearns, Michael J., Solla, Sara A., and Cohn, David A. (eds.), *NIPS*, pp. 487–493. The MIT Press, 1998. ISBN 0-262-11245-0.
- Klein, Dan and Manning, Christopher D. Fast exact inference with a factored model for natural language parsing. In *NIPS*, pp. 3–10, 2002.
- Kremenek, Ted, Ng, Andrew Y, and Engler, Dawson R. A factor graph model for software bug finding. In *IJCAI*, 2007.
- Liang, P., Jordan, M. I., and Klein, D. Learning programs: A hierarchical Bayesian approach. In *ICML*, pp. 639–646, 2010.
- Matsuzaki, Takuya, Miyao, Yusuke, and Tsujii, Jun’ichi. Probabilistic cfg with latent annotations. In *ACL*, pp. 75–82. ACL, 2005.
- McClosky, David, Charniak, Eugene, and Johnson, Mark. Effective self-training for parsing. In *ACL*, pp. 152–159. ACL, 2006.
- Mnih, Andriy and Hinton, Geoffrey. Three new graphical models for statistical language modelling. In *ICML*, pp. 641–648, 2007.
- Mnih, Andriy and Teh, Yee Whye. A fast and simple algorithm for training neural probabilistic language models. In *ICML*, pp. 1751–1758, 2012.
- MSDN. Microsoft Roslyn CTP, 2011. URL <http://msdn.microsoft.com/en-gb/roslyn>.
- Nguyen, Anh Tuan, Nguyen, Tung Thanh, Nguyen, Hoan Anh, Tamrawi, Ahmed, Nguyen, Hung Viet, Al-Kofahi, Jafar, and Nguyen, Tien N. Graph-based pattern-oriented, context-sensitive source code completion. In *ICSE*, ICSE 2012, pp. 69–79. IEEE Press, 2012. ISBN 978-1-4673-1067-3.
- Nguyen, Tung Thanh, Nguyen, Anh Tuan, Nguyen, Hoan Anh, and Nguyen, Tien N. A statistical semantic language model for source code. In *ESEC/FSE*, pp. 532–542. ACM, 2013.
- Saul, Lawrence and Pereira, Fernando. Aggregate and mixed-order Markov models for statistical language processing. In *EMNLP*, 1997.
- Titov, Ivan and Henderson, James. Incremental Bayesian networks for structure prediction. In *ICML*, 2007.
- Wang, Jue, Dang, Yingnong, Zhang, Hongyu, Chen, Kai, Xie, Tao, and Zhang, Dongmei. Mining succinct and high-coverage api usage patterns from source code. In *MSR*, pp. 319–328. IEEE Press, 2013.