

Inference of Node Replacement Recursive Graph Grammars

Jacek P. Kukluk, Lawrence B. Holder, and Diane J. Cook

{kukluk, holder, cook} @cse.uta.edu
Department of Computer Science and Engineering
University of Texas at Arlington
Box 19015, Arlington, TX 76019

Abstract

In this paper we describe an approach to learning node replacement graph grammars. This approach is based on previous research in frequent isomorphic subgraphs discovery. We extend the search for frequent subgraphs by checking for overlap among the instances of the subgraphs in the input graph. If subgraphs overlap by one node we propose a node replacement grammar production. We also can infer a hierarchy of productions by compressing portions of a graph described by a production and then infer new productions on the compressed graph. We validate this approach in experiments where we generate graphs from known grammars and measure how well our system infers the original grammar from the generated graph. We briefly discuss other grammar inference systems indicating that our study extends classes of learnable graph grammars.

Keywords: Grammar Induction, Graph Grammars, Graph Mining.

1. Introduction

String grammars are fundamental to linguistics and computer science. Graph grammars can represent relations in data which strings cannot. Graph grammars can represent hierarchical structures in data and generalize knowledge in graph domains. They have been applied as analytical tools in physics, biology, and engineering [Gernert97, Milo02]. In this paper we study the problem of grammar inference. We introduce an algorithm which builds on previous work in discovering frequent subgraphs in a graph [Cook94]. We check if subgraphs overlap and if they overlap by one node, we use this node and subgraph structure to propose a node replacement graph grammar. A vast amount of research has been done in string grammar inference [Sakakibara97]. We found only a few studies in graph grammar inference, which we now describe.

Jeltsch and Kreowski [Jeltsch90] did a theoretical study of inferring hyperedge replacement graph grammars from simple undirected, unlabeled graphs. Their paper leads through an example where from four complete bipartite graphs $K_{3,1}$, $K_{3,2}$, $K_{3,3}$, $K_{3,4}$, the authors describe the inference of a grammar that can generate a more general class of bipartite graphs $K_{3,n}$, where $n \geq 1$. The authors define four operations that lead to a final hyperedge replacement grammar. The operations are: INIT, DECOMPOSE, RENAME, and REDUCE. The INIT operation will start the process from a grammar which has all sample graphs in its productions and therefore it generates only the sample graphs. Then, the DECOMPOSE operation transforms the initial productions into productions that are more general but can still produce every graph from the sample graphs. RENAME allows for changing names of non-terminal labels. REDUCE removes redundant productions. They start the process from a grammar which has all the sample graphs in its productions. Then they transform the initial productions into productions that are more general but can still produce every graph from the sample graphs. Their approach guarantees that the final grammar will generate graphs that contain all sample graphs.

Oates et al. [Oates03] discuss the problem of inferring probabilities of every grammar rule for stochastic hyperedge replacement context free graph grammars. They call their program Parameter Estimation for Graph Grammars (PEGG). They assume that the grammar is given. Given a structure of a grammar S and a finite set of graphs E generated by grammar S , they ask what are the probabilities θ associated with every rule of the grammar. Their strategy is to look for a set of parameters θ that maximizes the probability $p(E | S, \theta)$.

In terms of similarity to string grammar inference we consider the Sequitur system developed by Nevill-Manning and Witten [Nevill97]. Sequitur infers hierarchical structure by replacing substrings based on grammar rules. The new, compressed string is searched for substrings which can be described by

grammar rules, and they are then compressed with the grammar and the process continues iteratively. Similarly, in our approach we replace the part of a graph described by the inferred graph grammar with a single node and we look for grammar rules on the compressed graph and repeat this process iteratively until the graph is fully compressed.

The most relevant work to this research is Jonyer et al.'s approach to node-replacement graph grammar inference [Jonyer02, Jonyer04]. Their system starts by finding frequently occurring subgraphs in the input graphs. Frequent subgraphs are those that when replaced by single nodes minimize the description length of the graph. They check if isomorphic instances of the subgraphs that minimize the measure are connected by one edge. If they are, a production $S \rightarrow PS$ is proposed, where P is the frequent subgraph. P and S are connected by one edge. Our approach is similar to Jonyer's in that we also start by finding frequently occurring subgraphs, but we test if the instances of the subgraphs overlap by one node. Jonyer's method of testing if subgraphs are adjacent by one edge limits his grammars to description of "chains" of isomorphic subgraphs connected by one edge. Since an edge of a frequent subgraph connecting it to the other isomorphic subgraph can be included to the subgraph structure, testing subgraphs for overlap allows us to propose a class of grammars that have more expressive power than the graph structures covered by Jonyer's grammars. For example, testing for overlap allows us to propose grammars which can describe tree structures, while Jonyer's approach does not allow for tree grammars.

In our approach we use the frequent subgraph discovery system Subdue developed by Cook and Holder [Cook94]. We would like to mention other systems developed to discover frequent subgraphs and therefore having potential to be modified into a system which can infer a graph grammar. Kuramochi and Karypis [Kuramochi01] implemented the FSG system for finding all frequent subgraphs in large graph databases. FSG starts by all frequent one and two edge subgraphs. Then, in each iteration, it generates candidate subgraphs by expanding the subgraphs found in the previous iteration by one edge. In every iteration, the algorithm checks how many times the candidate subgraph occurs within an entire graph. The candidates, whose frequency is below a user-defined level, are pruned. The algorithm returns all subgraphs occurring more frequently than the given level. In the candidate generation phase, computation costs of testing graphs for isomorphism are reduced by building a unique code for the graph (canonical labeling).

Yan and Han introduced gSpan [Yan02], which does not require candidate generation to discover frequent substructures. The authors combine depth first search and lexicographic order in their algorithm. Their algorithm starts from all frequent one-edge graphs. The labels on these edges together with labels on incident nodes define a code for every such graph. Expansion of these one-edge graphs maps them to longer codes. The codes are stored in a tree structure such that if $\alpha = (a_0, a_1, \dots, a_m)$ and $\beta = (a_0, a_1, \dots, a_m, b)$, then the β code is a child of the α code. Since every graph can map to many codes, the codes in the tree structure are not unique. If there are two codes in the code tree that map to the same graph and one is smaller than the other, the branch with smaller code is pruned during depth first search traversal of the code tree. Only the minimum code uniquely defines the graph. Code ordering and pruning reduces the cost of matching frequent subgraphs in gSpan.

The challenge of using frequent subgraph mining systems like gSpan or FSG to infer graph grammars would be the modification to allow subgraph instances to overlap. Overlapping substructures are available as an option in the Subdue system [Cook94]. Also, Subdue allows for identification of one substructure with the best compression score, which we can modify to identify one grammar production with the best score, while FSG and gSpan return all candidate subgraphs above a user-defined frequency level leaving interpretation and final selection for the user.

In the remainder of the paper we introduce the definition of the discussed graph grammars. Next we introduce an algorithm which we describe informally using an example. Afterwards, we give a more formal description. Then, we show experiments to reveal the advantages and limitations of our method. We close with some conclusions and future directions.

2. Node replacement recursive graph grammar

We give the definition of a graph and graph grammars which is relevant to our implementation. The defined graph has labels on vertices and edges. Every edge of the graph can be directed or undirected. The definition of a graph grammar describes the class of grammars that can be inferred by our approach. We emphasize the role of recursive productions in the name of the grammar, because the type of inferred productions are such that the non-

terminal label on the left side of the production appears one or more times in the node labels of a graph on the right side. It is the main characteristic of our grammar productions. Our approach can also infer non-recursive productions. The embedding mechanism of the grammar consists of connection instructions. Every connection instruction is a pair of vertices that indicate where the production graph can connect to itself in a recursive fashion. Our graph generator can generate a larger class of graph grammars than defined below. We will describe the grammars used in generation later in the paper.

A labeled graph G is a 6-tuple,

$$G = (V, E, \mu, \nu, \eta, L), \text{ where}$$

V - is the set of nodes, $E \subseteq V \times V$ - is the set of edges, $\mu : V \rightarrow L$ - is a function assigning labels to the nodes, $\nu : E \rightarrow L$ - is a function assigning labels to the edges, $\eta : E \rightarrow \{0, 1\}$ - is a function assigning direction property to edges (0 if undirected, 1 if directed). L - is a set of labels on nodes and edges.

A node replacement recursive graph grammar is a tuple $Gr = (\Sigma, \Delta, \Gamma, P)$, where

Σ - is an alphabet of node labels,

Δ - is an alphabet of terminal node labels, $\Delta \subseteq \Sigma$,

Γ - is an alphabet of edge labels, which are all terminals,

P - is a finite set of productions of the form (d, G, C) , where $d \in \Sigma - \Delta$, G is a graph, and there are two types of productions:

(1) *recursive productions* of the form (d, G, C) , with $d \in \Sigma - \Delta$, G is a graph, where there is at least one node in G labeled d . C is an embedding mechanism with a set of connection instructions, $C \subseteq V \times V$, where V is the set of nodes of G . A connection instruction $(v_i, v_j) \in C$ implies that derivation can take place by replacing v_i in one instance of G with v_j in another instance of G . All the edges incident to v_i are incident to v_j . All the edges incident to v_j remain unchanged

(1) *non-recursive production*, there is no node in G labeled d (our inference system does not infer an embedding mechanism for these productions).

3. The algorithm

We will describe an algorithm informally allowing for an intuitive understanding of the idea. An

example in Figure 1 shows a graph composed of three overlapping substructures. The algorithm generates candidate substructures and evaluates them using the following measure of compression,

$$\frac{\text{size}(G)}{\text{size}(S) + \text{size}(G|S)}$$

where G is the input graph, S is a substructure and $G|S$ is the graph derived from G by compressing each instance of S into a single node. $\text{size}(g)$ can be computed simply by summing the number of nodes and edges: $\text{size}(g) = \text{vertices}(g) + \text{edges}(g)$. Another successful measure of $\text{size}(g)$ is the Minimum Description Length (MDL) discussed in detail in [Cook94]. Either of these measures can be used to guide the search and determine the best graph grammar. In our experiments we used only the size measure.

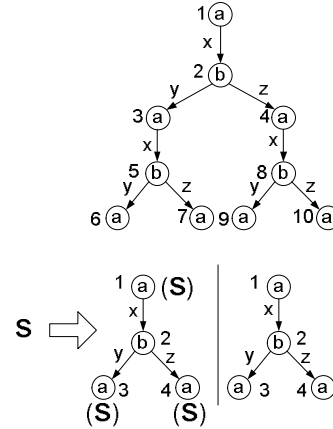


Figure 1: A graph with overlapping substructures and a graph grammar representation of it.

In Figure 1, the subgraphs overlap at nodes 3 and 4. The algorithm starts by finding nodes with the same label. There are seven nodes labeled "a" and three nodes labeled "b". The single node labeled "a" becomes a candidate substructure with seven instances $I_1=\{1\}$, $I_2=\{3\}$, $I_3=\{4\}$, $I_4=\{6\}$, $I_5=\{7\}$, $I_6=\{9\}$, $I_7=\{10\}$. The numbers in parentheses refer to the nodes in Figure 1. This initial substructure will be expanded by a node and an edge in each iteration of the algorithm's main discovery loop. Similarly, the initial substructure of a node labeled "b" and its instances are determined. Both of these substructures are expanded simultaneously. Let us follow the expansion of only one substructure, which starts from all nodes labeled "b." Table 1 gives the instance expansion at every step and a substructure value We expand the instances I by edge labeled y and a vertex labeled a, which gives us the set of instances I^P . Instances I can also be expanded by edge z or x.

Similarly, we expand \mathbf{I}' by edge z and a vertex a , which gives us \mathbf{I}'' . \mathbf{I}' can also be expanded by edge x . We omit in Table 1 alternative expansions of \mathbf{I} by z , x and \mathbf{I}' by x . These additional expansions are part of our algorithm. They lead to the same solution. When the set of instances \mathbf{I}'' is expanded by the edge with label x , we detect an overlap, i.e., two or more instances share the same node. The overlapping instances of the substructure allow us to propose the recursive graph grammar shown on the bottom of Figure 1. This grammar can compress the entire graph to one node and has a better substructure value than any other substructure discovered so far.

The grammar from Figure 1 consists of a graph isomorphic to three overlapping substructures and connection instructions. We find connection instructions when we check for overlaps. In this example there are two connection instructions 1-3 and 1-4. Hence, in generation of a graph from the grammar, in every derivation step an isomorphic copy of the subgraph definition will be connected to the existing graph by connecting node 1 of the subgraph to either a node 3 or a node 4 in the existing graph. The grammar shown on the bottom in Figure 1 cannot only regenerate the graph shown on the top, but also generate generalizations of this graph. Generalization in this example means that the grammar describes graphs composed from one or more star looking substructures of four nodes labeled “a”, “b”. All these substructures overlap on a node with label “a”.

Our graph grammar inference method is based on Cook et al.’s [Cook94] substructure discovery system called Subdue. Subdue is looking for repetitive, highly-compressing subgraphs. The algorithm starts by finding all nodes with the same label. It maintains a list of the best subgraphs found so far. In each iteration new candidates for the best subgraphs are created by expanding all the subgraphs in the list by one edge or edge and a node. Then, candidates for the best subgraphs are evaluated. In the evaluation process, every occurrence of a candidate subgraph within the entire graph is temporarily replaced by a new node. The compression achieved with this replacement is measured by calculating minimum

description length or size (number of nodes + number of edges) of an original and compressed graph. Only subgraphs with the highest compression ratio remain in the list of the best subgraphs.

The input to our algorithm is a graph G which can be one connected graph or set of disconnected graphs. G can have directed edges or undirected edges. The algorithm assumes labels on nodes and edges. The algorithm processes the list of substructures Q . In Figure 2 we see an example of a substructure definition. A substructure consists of a graph definition and a set of instances from the input graph that are isomorphic to the graph definition. The example in Figure 2 is a continuation of the example in Figure 1. The numbers in parentheses refer to nodes of the graph in Figure 1.

The algorithm starts with a list of substructures where every substructure is a single node and its instances are all nodes in the graph with this node label. The best substructure is initially the first substructure in the Q list. We extend each substructure in Q in all possible ways by a single edge and a node or only by single edge if both nodes are already in the graph definition of the substructure. We allow instances to grow and overlap, but any two instances can overlap by only one node. We keep all extended substructures in $newQ$. We evaluate substructures in $newQ$. The recursive substructure is evaluated along with non-recursive substructures and is competing with non-recursive substructures. The total number of substructures considered is determined by the input parameter *Limit*. We compress G with best substructure. Compression replaces every instance of best substructure with a single node. This node is labeled with a non-terminal label. The compressed graph is further processed until it cannot be compressed any more. In consecutive iterations best substructure can have one or more non-terminal labels. It allows us to create a hierarchy of grammar productions. The input parameter *Beam* specifies the width of a beam search, i.e., the length of Q . For more details about the algorithm see [Cook94, Jonyer02, Jonyer04].

Table 1. Expansion of instances which start from nodes labeled “b” in Figure 1.

Expansion	Instances	$\frac{size(G)}{size(S)+size(G S)}$
initial instances	$\mathbf{I} =\{ I_1=\{2\}, I_2=\{5\}, I_3=\{8\} \}$	$19/(1+19)=0.95$
\mathbf{I} expanded by y	$\mathbf{I}' =\{ I_1=\{2, 3\}, I_2=\{5, 6\}, I_3=\{8, 9\} \}$	$19/(3+13)=1.19$
\mathbf{I}' expanded by z	$\mathbf{I}'' =\{ I_1=\{2, 3, 4\}, I_2=\{5, 6, 7\}, I_3=\{8, 9, 10\} \}$	$19/(5+7)=1.58$
\mathbf{I}'' Expanded by x	$\mathbf{I}''' =\{ I_1=\{2, 3, 4, 1\}, I_2=\{5, 6, 7, 3\}, I_3=\{8, 9, 10, 4\} \}$ (overlap !)	$19/(7+1)=2.38$

Figure 2 assists us in explaining conversion of substructure S into recursive substructure. Every instance graph has two positive integers assigned to it. One integer, in parentheses in Figure 2, is the number of a node in the processed graph G . The second integer is a node number of an instance graph. The instances are isomorphic to the substructure graph definition and instance node numbers are assigned to them according to this isomorphism. Given pair of instances (I_1, I_2) we examine if there is a node $v \in G$, which also belongs to I_1 and I_2 . We find two overlapping nodes, [3] and [4], examining node numbers in parentheses in the example in Figure 2. Having the number of node $v \in G$ we find corresponding to v two node numbers of instance graphs $v_I \in I_1$ and $v_I' \in I_2$. The pair of integers (v_I, v_I') is a connection instruction. There are two connection instructions in Figure 2: 1-3 and 1-4. If (v_I, v_I') is not already in list of connections instructions for recursive substructure, we include it.

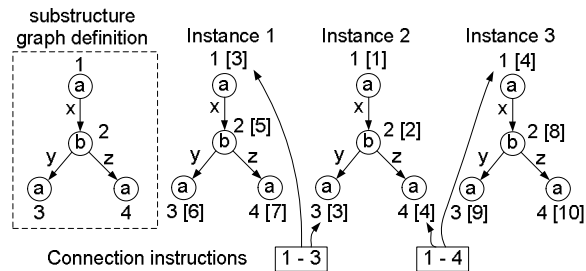


Figure 2: Substructure and its instances while determining connection instructions (continuation of the example from Figure 1)

A recursive instance is a connected subgraph of G which can be described by the discovered grammar production. It means that for every subset of instances $\{I_m, I_{m+1}, \dots, I_l\}$ from the instance list of substructure S , in which union $I_m \cup I_{m+1} \cup \dots \cup I_l$ is a connected graph, we create one recursive instance $IR_k = I_m \cup I_{m+1} \cup \dots \cup I_l$. The recursive instances are no longer isomorphic as instances of S and they vary in size. Every recursive instance is compressed to a single node in the evaluation process.

Subdue uses a heuristic search whose complexity is polynomial in the size of the input graph [Cook00]. Our modification does not change the complexity of this algorithm. The overlap test is the main computationally expensive addition of our grammar discovery algorithm. Analyzing informally, the number of nodes of an instance graph is not larger

than V , where V is the number of nodes in the input graph. Checking two instances for overlap will not take more than $O(V^2)$ time. The number of pairs of instances is no more than V^2 , so the entire overlap test will not take more than $O(V^4)$ time.

4. Hierarchy of productions

In our first example from Figure 1, we described a grammar with only one production. Now we would like to introduce a complex example to illustrate the inference of a grammar which describes a more general tree structure. In Figure 3 we have a tree with all nodes having the same label.

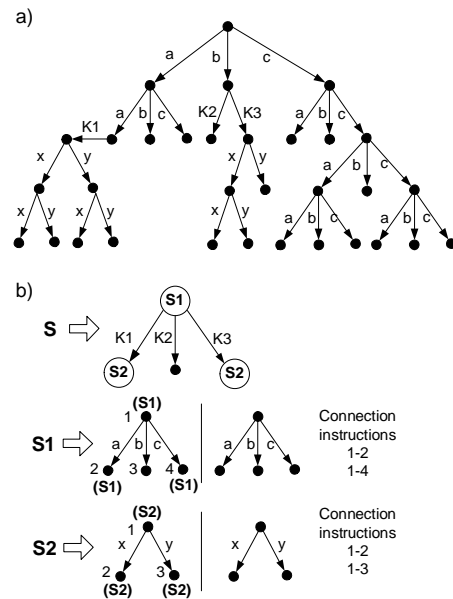


Figure 3: The tree (a) and inferred tree grammar (b).

There are two repetitive subgraphs in the tree. One has three edges labeled “a,” “b,” and “c.” The other has two edges with labels “x” and “y.” There are also three edges $K1, K2,$ and $K3$ which are not part of any repetitive subgraph. In the first iteration we find grammar production $S1$, because overlapping subgraphs with edges “a,” “b,” and “c” score the highest in compressing the graph. Examining production $S1$, we notice that node 3 is not involved in connection instructions. It is consistent with the input graph where there are no two subgraphs overlapping on this node. The compressed graph, at this point, contains the node $S1$, edges $K1, K2, K3$ and subgraphs with edges “x” and “y.” In the second iteration our program finds all overlapping substructures with edges “x” and “y” and proposes

production S2. Compressing the tree with production S2 results in a graph which we use as an initial production S, because the graph can be compressed no further. In Figure 3 productions for S1 and S2 have graphs as terminals. We will omit drawing terminal graphs in subsequent figures. The tree used in this example was used in our experiments, and the grammar on the right in Figure 3 is the actual inferred grammar.

5. Experiments

5.1. Methodology

Having our system implemented, we faced the challenge of evaluating its performance. There are an infinite number of grammars as well as graphs generated from these grammars. In our experiments we restricted grammars to node replacement grammars with two productions. The second production replaces a non-terminal node with a single terminal node. In Figure 4 we give an example of such a grammar. The grammar on the left is of the form used in generation. The grammar on the right is the inferred grammar in our experiment. The inferred grammar production is assumed to have a terminating alternative with the same structure as the recursive alternative, but with no non-terminals. We omit terminating production in Figure 4.

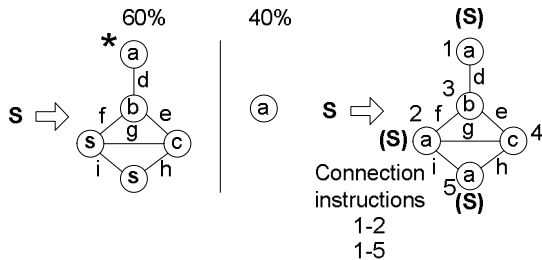


Figure 4: Example of graph grammar used in the experiments.

We developed a graph generator to generate graphs from a known grammar. We can generate directed or undirected graphs with labels on nodes and edges. Our generator produces a graph by replacing a non-terminal node of a graph until all nodes and edges are terminal. The generation process expands the graph as long as there are any non-terminal edges or nodes. Since selection of a production is random according to the probability distribution specified in the input file, the number of nodes of a generated graph is also random. We place limits on the size of the generated graph with two parameters: *minNodes* and *maxNodes*. We generate graphs from the grammar until the number of nodes is between *minNodes* and *maxNodes*. We distinguish two different distorting

operations to the graph generated from grammar: corruption and added noise. Corruption involves the redirection of randomly selected edges. The number of edges of a graph multiplied by *noise* gives the number of redirected edges, where *noise* is a value from 0 to 1. We redirect an edge $e = (v_1, v_2)$ by replacing nodes v_1 and v_2 with two new, randomly selected graph nodes v'_1 and v'_2 . When we add noise, we do not destroy generated graph structure. We add new nodes and new edges with labels assigned randomly from labels used in already generated graph structure. We compute the number of added nodes from the formula $(noise / (1 - noise)) * number_of_nodes$. The number of added edges we find from $(noise / (1 - noise)) * number_of_edges$. A new edge connects two nodes selected randomly from existing nodes of the generated structure and newly added nodes.

We associate probabilities with productions used in generation. These probabilities define how often a particular production is used in derivations. Assigning probabilities to productions helps us to control the size of the generated graph. Our inference system does not infer probabilities. Oates et al. [Oates03] addresses the problem of inferring probabilities assuming that the productions of a grammar are given. We are considering inferring probabilities along with productions as a future work.

We examined grammars with one, two, and three non-terminals. The first productions of the grammars have an undirected, connected graph with labels on nodes and edges on the right side. We use all possible connected simple graphs with three, four, and five nodes as the structures of graphs used in the productions. There are twenty nine different simple connected undirected unlabeled graphs [Read98]. We show them in Figure 7. Our graph generator generates graphs from the known grammar that is based on one of the twenty nine graph structures. Then we use our inference system to infer the grammar from the generated graph. We measure an error between the original and inferred grammar. We use MDL as a measure of the complexity of a grammar. Our results describe the dependency of the grammar inference error on complexity, noise, number of labels, and size of generated graphs.

5.2. MDL as a measure of complexity of a grammar

We seek to understand the relationship between graph grammar inference and grammar complexity, and so need a measure of grammar complexity. One

such measure is the Minimum Description Length (MDL) of a graph, which is the minimum number of bits necessary to completely describe the graph. The MDL measure, which while not provably minimal, is designed to be a near-minimal encoding of a graph. See [Cook94] for a more detailed discussion.

Since all the grammars in our experiments have two productions and the second production replaces a non-terminal with a single node, the complexity of the grammar will vary depending only on the graph on the right side of the first production. We would like our results for one, two and three non-terminal grammars to be comparable; therefore we do not want our measure of complexity of a grammar to be dependent on the number of non-terminals. In every graph used in the productions we reserve three nodes. We give the same label to these nodes. When we generate a graph, we replace one, two, or three labels of these nodes with the non-terminal S when we need a grammar with one, two or three non-terminals. However, when we measure MDL of a graph we leave the original three labels unchanged. In our experiments we always use that same non-terminal label. In the general case a production can contain different non-terminals. Every non-terminal would need to be counted as a different label of a graph and MDL would increase with increasing number of non-terminals.

5.3. Error

We introduce a measure to compare the original grammar to the inferred grammar. Our definition of an error has two aspects. First, there is the structural difference between the inferred and the original graph used in the productions. Second, there is the difference between the number of non-terminals and the number of connection instructions. If there is no error, the number of non-terminals in the original grammar is the same as the number of connection instructions in the inferred grammar. We compute the structural difference between graphs with an algorithm for inexact graph match initially proposed by Bunke and Allermann [Bunke83]. For more details see also [Cook94]. We would like our error to be a value between 0 and 1; therefore, we normalize the error by having in the denominator the sum of the size of the graph used in the original grammar and the number of non-terminals. We do not allow an error to be larger than 1; therefore, we take the minimum of 1 and our measure as a final value.

$$Error = \min\left(1, \frac{\text{matchCost}(g_1, g_2) + |\#CI - \#NT|}{\text{size}(g_1) + \#NT}\right),$$

where

$\text{matchCost}(g_1, g_2)$ is the minimal number of operations required to transform g_1 to a graph isomorphic to g_2 , or g_2 to a graph isomorphic to g_1 . The operations are: insertion of an edge or node, deletion of a node or an edge, or substitution of a node or edge label. $\#CI$ is the number of inferred connection instructions. $\#NT$ is the number of non-terminals in the original grammar. $\text{size}(g_1)$ is the sum of the number of nodes and edges in the graph used in the grammar production

5.4. Experiment 1: Error as a function of noise and complexity of a grammar

We used twenty nine graphs from Figure 7 in grammar productions. We assigned different labels to nodes and edges of these graphs except three nodes used for non-terminals. We generated graphs with noise from 0 to 0.9 in 0.1 increments. For every value of noise and MDL we generated thirty graphs from the known grammar and inferred the grammar from the generated graph. We computed the inference error and averaged it over thirty examples. We generated 8700 graphs to plot each of the three graphs in Figure 5. The first plot shows results for grammars with one non-terminal. The second and the third plot show results for grammars with two and three non-terminals. We did not corrupt the generated graph structure in experiments in Figure 5. As noise we added nodes and edges to the generated graph structure. Figure 6 has the same results as Figure 5 with the difference that we corrupted the graph structure generated from the grammar and then we added nodes and edges to the graph.

We see trends in the plots in Figure 5 and Figure 6. Error decreases as MDL increases. A low value of MDL is associated with small graphs, with three or four nodes and a few edges. These graphs, when used on the right hand side of a grammar production, generate graphs with fewer labels than grammars with high MDL. Smaller numbers of labels in the graph increase the inference error, because everything in the graph looks similar, and the approach is more likely to propose another grammar which is very different than the original. As expected, the error increases as the noise increases in experiments with corrupted graph structure. However, there is little dependency of an error from the noise if the graph generated from the grammar is not corrupted (Figure 5).

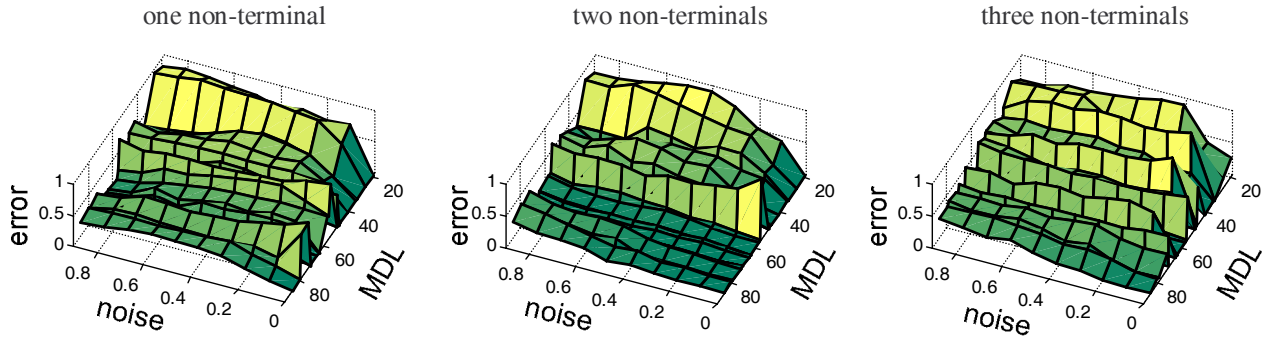


Figure 5: Error as a function of noise and MDL where graph structure was not corrupted.

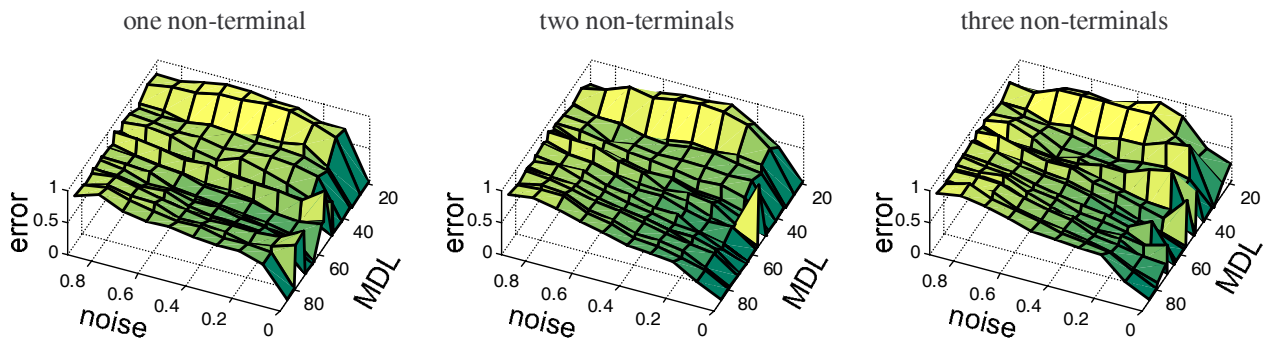


Figure 6: Error as a function of noise and MDL where graph structure was corrupted.

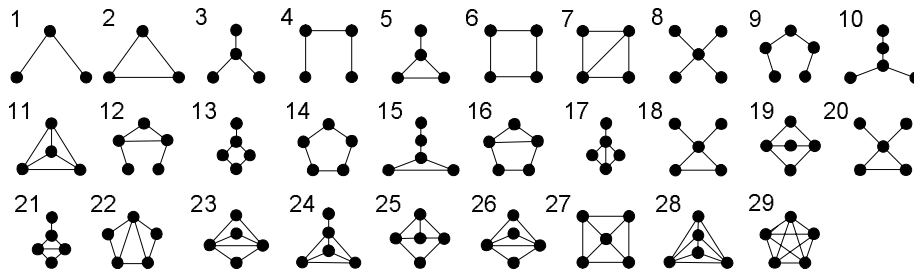


Figure 7: Twenty nine simple connected graphs ordered according to non-decreasing MDL value.

Table 2: Twenty nine simple graphs ordered according to increasing average inference error of six experiments in Figure 5 and Figure 6. The numbers in the table refer to structures in Figure 7.

Number of non-terminals	1	Not Corrupted	21 17 22 15 8 10 23 28 20 27 29 19 26 12 16 3 18 4 24 25 9 5 7 14 6 13 11 1 2
	2		21 23 22 15 18 16 17 20 19 9 28 12 10 14 26 13 27 25 8 24 29 4 5 7 3 6 11 2 1
	3		21 15 23 16 17 19 18 14 9 13 28 12 27 26 25 24 5 10 4 29 22 6 20 7 11 2 8 1 3
Number of non-terminals	1	Corrupted	8 10 12 21 17 15 20 23 16 19 18 22 13 14 9 27 4 28 25 3 7 29 24 6 26 5 11 1 2
	2		9 17 19 16 21 13 18 8 15 14 10 12 25 27 23 22 24 20 26 28 4 3 6 5 29 7 11 1 2
	3		9 19 14 12 18 16 13 15 21 17 4 23 10 25 27 26 5 6 24 20 28 22 29 8 7 3 11 1 2

We average the value of an error over ten values of noise which gives us the value we can associate with the graph structure. It allowed us to order graph

structures used in the grammar productions based on average inference error. In Figure 7 we show all twenty nine connected simple graphs with three, four

and five nodes used in productions ordered in non-decreasing MDL value of a graph structure. In Table 2 we give an order of graph structures for six experiments with corrupted and non-corrupted structures and one, two, and three non-terminals. The numbers in the table refer to structure numbers in Figure 7. We see in Table 2 that graph number 21 is close to the beginning of the list in all six experiments. Graphs number 1, 2, and 11 are close to the end of all six lists. We conclude that when graph number 21 is used in the grammar production, it is the easiest for our inference algorithm to find the correct grammar. When graph number 1, 2, or 11 is used in the grammar production and generated graphs have noise present, we infer grammars with some error. We also observe a tendency of decreasing error with increasing MDL in the graph orders in Table 2. Graph 29 has the highest MDL, because it has the most nodes and edges. In five experiments graph 29 is closer to the end of the list

Quantitative definition of an error allows us to automate the process and perform tests on thousands of graphs. The error is caused by a wrongly inferred graph structure used in the production or number of connection instructions which is too large or too small. However, there are cases where the inferred grammar represents the graph well, but the graph in the production has a different structure. For example, we observed that the grammar with MDL=55.58 and graph number 11 causes an error even if we infer the grammar from graphs with no corruption and zero noise. The inferred graph structure contains two overlapping copies of the graph used in the original grammar production. We illustrate it in Figure 8. The structure has significant error, yet does subjectively capture the recursive structure of the original grammar.

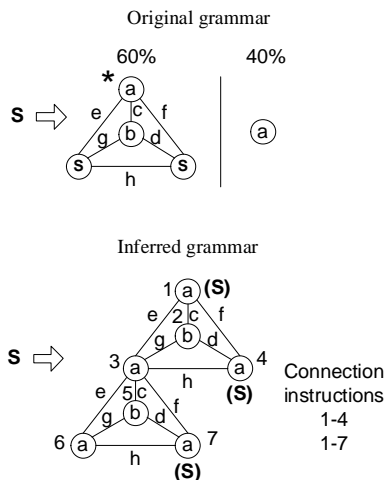


Figure 8: An inference error where larger graph structure is proposed.

5.5. Experiment 2: Error as a function of number of labels and complexity of a grammar

We would like to evaluate how error depends on the number of different labels used in a grammar. We restricted graph structures used in productions to graphs with five nodes. Every graph structure we labeled with 1, 2, 3, 4, 5 or 6 different labels. For every value of MDL and number of labels we generated 30 different graphs from the grammar and computed average error between them and the learned grammars. The generated graphs were without corruption and without noise. We show the results for one, two, and three non-terminals in Figure 10. Below the three dimensional plots, for clarity, we give two dimensional plots with triangles representing the errors. The larger and lighter the triangle the larger is the error. We see that the error increases as the number of different labels decreases. We see on the two dimensional plots the shift in error towards graphs with higher MDL when the number of non-terminals increases.

The average error for graphs with only one label is 1 or very close to 1. The most frequent inference error results from the tendency of our algorithm to propose one-edge grammars when inferred from a graph with only one label. We illustrate this in Figure 9 where we see a production with a pentagon using only one label “a”. The inferred grammar has one edge with two connection instructions 1-1 and 1-2. Since all the edges in the generated graph have the same label and all the nodes have the same label, this grammar compresses the graph well and is evaluated highly by our compression-based measure. However, this one-edge grammar cannot generate even a single pentagon. An evaluation measure which penalizes grammars for an inability to generate an input graph would bias the algorithm away from single-edge grammars and could correct the one-edge grammar problem. However, this approach would require graph-grammar parsing, which is computationally complex.

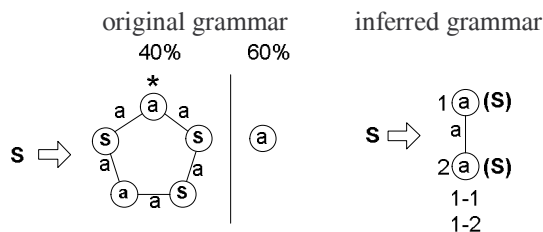


Figure 9: Error where inferred grammar is reduced to production with single edge.

5.6. Experiment 3: Error as a function of size of a graph and complexity of a grammar

We generated graphs from grammars with two non-terminals and noise=0.2. The number of nodes of the generated graphs was from the interval $[x, x+20]$,

where we change x from 20 to 420. For each value of x and MDL we generated thirty graphs and compute average inference error over them. We show in Figure 11 the results for corrupted and not corrupted graph structure. We concluded that there is no dependency between the size of a sample graph and inference error.

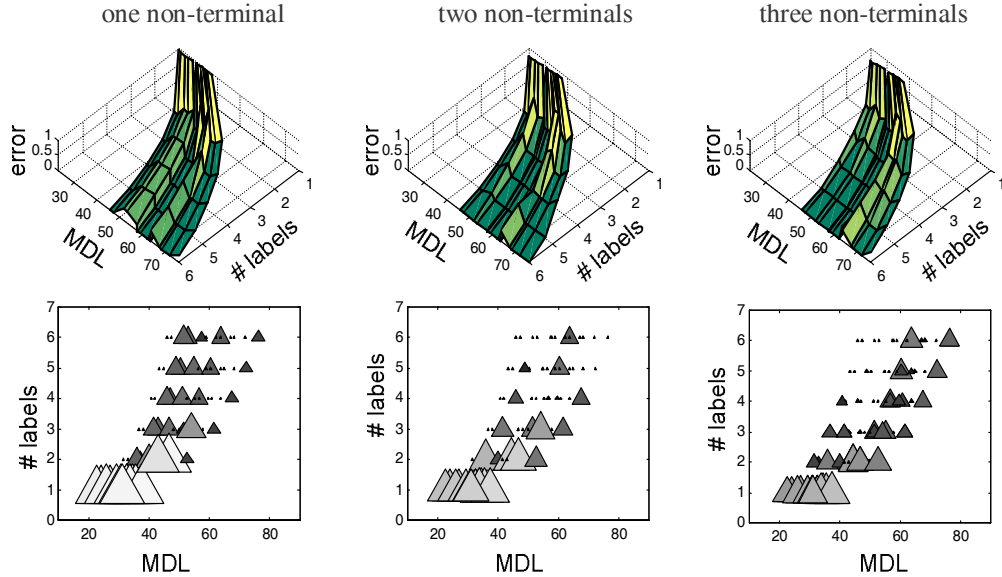


Figure 10 : Error as a function of MDL and number of different labels used in a grammar definition.

5.7. Experiment 4: Limitations

In Figure 12 we show an example illustrating the limits of our approach. In Figure 12 (a) we have a graph consisting of overlapping squares. All labels on nodes are the same, and we omit them. The squares do not overlap by one node but by an edge. Our algorithm assumes that only one node overlaps in the instances of the substructure and therefore infers the grammar shown in Figure 12(b). The inferred grammar can generate chains, an example of which is shown in Figure 4 (c). The original input graph is not in the set of graphs generated by the inferred grammar. An extension of our method to overlapping edges would allow us to infer the correct grammar in this example.

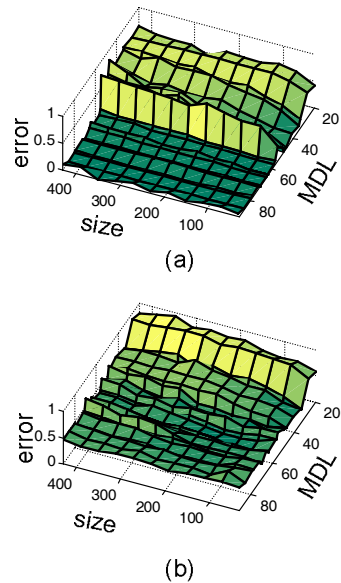


Figure 11: Error as a function of MDL and size of generated graphs (noise=0.2, two non-terminals): (a) uncorrupted graph structure, (b) corrupted graph structure

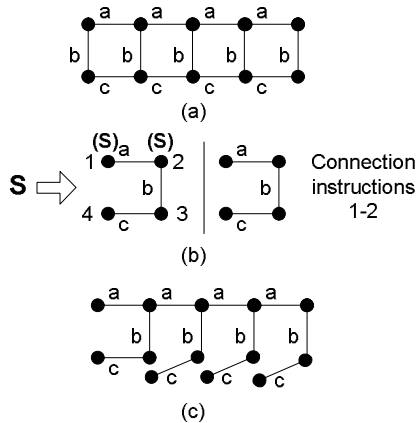


Figure 12: Graph with overlapping squares (a), inferred grammar (b), and graph generated from inferred grammar (c)

Figure 13 shows another example illustrating the limits of our algorithm. The first graph in the first production on the left is a square with two non-terminals labeled S1, and the graph of the second production is a triangle with one non-terminal labeled S. Our algorithm is not designed to find alternating productions of this type. We generated a graph from the grammar on the left, and the grammar we inferred is on the right in Figure 13. The inferred grammar has one production in which the graph combines both the triangle and square. The set of graphs generated by alternating squares and triangles according to the grammar from the left does not match exactly the set of graphs of the inferred grammar. Nevertheless, we consider it an accurate inference, because the inferred grammar will describe the majority of every graph generated by the original grammar.

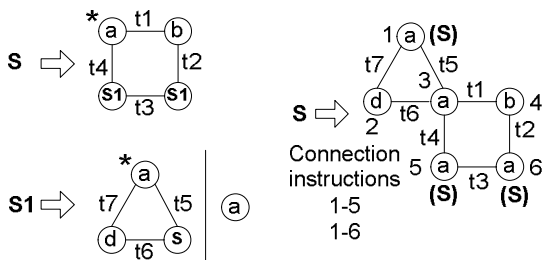


Figure 13: The grammar with alternating productions (left) and inferred grammar (right).

5.8. Experiment 5: Chemical structure

As an example from the real-world domain of chemistry, we use the structure of cellulose with hydrogen bonding as the input graph in our next experiment. Figure 14 shows the structure of the molecule and the grammar production we found in

this structure. The grammar production we found captures the underlying motif of the chemical structure. It shows the repetitive connected component, the basic building block of the structure. We can search for such underlying building motifs in different domains, hoping that they will improve our understanding of chemical, biological, computer, and social networks.

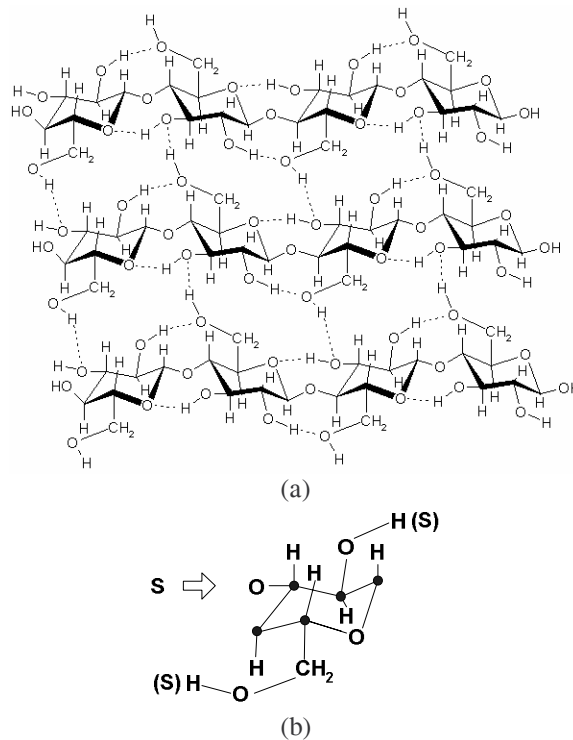


Figure 14: The structure of cellulose with hydrogen bonding (a) and the inferred grammar production (b).

6. Conclusion and future work

We described an algorithm for inferring certain types of graph grammars we call recursive node replacement graph grammars. The algorithm is based on previous work in frequent substructure discovery. It checks if frequent subgraphs overlap by a node and proposes a graph grammar if they do. The algorithm we described has its limitations: the left side of the production is limited to one single node; only connecting two single nodes is allowed in derivations. The algorithm finds recursive productions if repetitive patterns occur within an input graph and they overlap. If such patterns do not exist, the algorithm finds non-recursive productions and builds hierarchical structure of the input data. Grammar productions with graphs of higher complexity measured by MDL are inferred with

smaller error. There is little dependency of error on noise if the generated graphs are not corrupted. The error of grammar inference increases as the number of different labels used in the grammar decreases. There is no dependency between the size of a sample graph and inference error. If all labels on nodes are the same and all labels on edges are the same, the algorithm produces a grammar which has only one edge in the graph definition. One-edge grammars over-generalize if the input graph is a tree, and they are inaccurate in many other graphs. This tendency to find one-edge grammars from large, connected graphs is due to the fact that one-edge grammars score high because they can compress the graph well.

Grammars inferred by the approach developed by Jonyer et al. [Jonyer04] were limited to chains of isomorphic subgraphs which must be connected by a single edge. Since the connecting edge can be included in the production's subgraph, and isomorphic subgraphs will overlap by one vertex, our approach can infer Jonyer et al.'s class of grammars. As we noticed in our experiment shown in Figure 12 when the subgraphs overlap by more than one node, our algorithm still looks for overlap on only one node and infers a grammar which cannot generate the input graph. Therefore one extension to the algorithm would be a modification which would allow for overlap larger than a single node.

The evaluation method can be modified to avoid one-edge grammar productions in graphs with one label. We are exploring other domains where data can be represented as a graph composed from smaller structures to further test our inference system and examine it as a data mining tool for these domains. We are continuing our research in graph grammar inference to develop methods allowing for discovery of more powerful classes of graph grammars than discussed in this paper.

References

- [Bunke83] H. Bunke and G. Allermann, *Inexact graph matching for structural pattern recognition*. Pattern Recognition Letters, 1(4) 245-253. 1983
- [Cook00] D. Cook and L. Holder, *Graph-Based Data Mining*. IEEE Intelligent Systems, 15(2), pages 32-41, 2000.
- [Cook94] D. Cook and L. Holder, *Substructure Discovery Using Minimum Description Length and Background Knowledge*. Journal of Artificial Intelligence Research, Vol 1, (1994), 231-255
- [Doshi02] S. Doshi, F. Huang, and T. Oates, *Inferring the Structure of Graph Grammar from Data*. Proceedings of the International Conference on Knowledge Based Computer Systems (KBCS'02)
- [Gernert97] D. Gernert, *Graph grammars as an analytical tool in physics and biology*. Biosystems 1997, vol. 43, no. 3, pp. 179-187(9)
- [Jeltsch90] E. Jeltsch, H. Kreowski, *Grammatical Inference Based on Hyperedge Replacement*. Graph-Grammars. Lecture Notes in Computer Science 532, 1990: 461-474
- [Jonyer02] I. Jonyer, L. Holder, and D. Cook, *Concept Formation Using Graph Grammars*, Proceedings of the KDD Workshop on Multi-Relational Data Mining, 2002.
- [Jonyer04] I. Jonyer and L. Holder, and D. Cook, *MDL-Based Context-Free Graph Grammar Induction and Applications*. International Journal of Artificial Intelligence Tools, Volume 13, No. 1 pages 65-79, 2004.
- [Kuramochi01] M. Kuramochi and G. Karypis, *Frequent subgraph discovery*. In Proceedings of IEEE 2001 International Conference on Data Mining (ICDM '01), 313-320, 2001.
- [Milo02] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon, *Network Motifs: Simple Building Blocks of Complex Networks*, Science. Vol 298, Issue 5594, 824-827, 2002
- [Nevill97] G. Nevill-Manning and H. Witten, *Identifying Hierarchical Structure in Sequences: A linear-time algorithm*. Journal of Artificial Intelligence Research, Vol 7, (1997), 67-82
- [Oates03] T. Oates, S. Doshi, and F. Huang, *Estimating Maximum Likelihood Parameters for Stochastic Context-Free Graph Grammars*. In T. Horváth and A. Yamamoto, editors, Proceedings of the 13th International Conference on Inductive Logic Programming, volume 2835 of Lecture Notes in Artificial Intelligence, pages 281--298. Springer-Verlag, 2003.
- [Read98] R. Read and R. Wilson, *An Atlas of Graphs*. Oxford University Press, 1998
- [Sakakibara97] Y. Sakakibara, *Recent advances of grammatical inference*. Theoretical Computer Science, 185:15-45, 1997.
- [Yan02] X. Yan and J. Han, *gSpan: Graph-based substructure pattern mining*. In IEEE International Conference on Data Mining, Maebashi City, Japan, December 2002.