Interactive Learning of Node Selecting Tree Transducers

Julien Carme, Rémi Gilleron, Aurélien Lemay, Joachim Niehren

▶ **To cite this version:**

HAL Id: inria-00087226

https://hal.inria.fr/inria-00087226v5

Submitted on 9 Apr 2010

Interactive Learning of Node Selecting Tree Transducers[*]

Julien Carme, Rémi Gilleron, Aurélien Lemay, Joachim Niehren

INRIA Futurs and Lille University, LIFL, Mostrare project

Version of April 9, 2010

Key words Web information extraction, wrapper induction, grammatical inference, tree automata, monadic queries.

Abstract We develop new algorithms for learning monadic node selection queries in unranked trees from annotated examples, and apply them to visually interactive Web information extraction.

We propose to represent monadic queries by bottom-up deterministic Node Selecting Tree Transducers (Nstts), a particular class of tree automata that we introduce. We prove that deterministic Nstts capture the class of queries definable in monadic second order logic (Mso) in trees, which Gottlob and Koch (2002) argue to have the right expressiveness for Web information extraction, and prove that monadic queries defined by Nstts can be answered efficiently. We present a new polynomial time algorithm in Rpni-style that learns monadic queries defined by deterministic Nstts from completely annotated examples, where all selected nodes are distinguished.

In practice, users prefer to provide partial annotations. We propose to account for partial annotations by intelligent tree pruning heuristics. We introduce *pruning* Nstt*s* - a formalism that shares many advantages of Nstts. This leads us to an interactive learning algorithm for monadic queries defined by pruning Nstts, which satisfies a new formal active learning model in the style of Angluin (1987).

We have implemented our interactive learning algorithm and integrated it into a visually interactive Web information extraction system – called Squirrel– by plugging it into the Mozilla Web browser. Experiments on

---

realistic Web documents confirm excellent quality with very few user inter-
actions during wrapper induction.

Contents

1 Node Selection Tasks in Web Information Extraction

The Web's information overload calls for intelligent services that provide
high-quality collected and value-added information. Information from di-
verse sources sites relevant to the service needs to be integrated and made
accessible in a uniform manner. Web documents in HTML are ill-suited for
automated processing, which nevertheless is an inevitable prerequisite. This
problem is addressed by much work on so-called *Web wrappers.* These are
programs that extract the relevant information from HTML documents and
translate it into a more machine-friendly format.

   Manual wrapper programming is often not manageable. Even for ex-
perts, this task remains time consuming and error prone. Furthermore,
wrappers must frequently be adapted to changes of the document sources.
This is why new tools have been proposed that assist humans in wrapper
creation. These are of two kinds:

– *Machine learning techniques* are used for inducing *string based wrappers.*
  Most systems (Kushmerick, 1997; Hsu & Dung, 1998; Freitag & Mc-
  Callum, 1999; Freitag & Kushmerick, 2000; Kushmerick, 2000; Muslea,
  Minton, & Knoblock, 2001; Chidlovskii, 2001) rely on finite-state ap-
  proaches (patterns, automata and transducers, Hidden Markov Models).
  See in Kushmerick (2002) for a comparison.
– *Visual specification* of *node selection queries in trees* (Baumgartner,
  Flesca, & Gottlob, 2001; Gottlob & Koch, 2002). Suitable query specifi-
  cation languages behind the graphical interface are monadic Datalog or
  similar tree logics.

Some of the most recent approaches propose machine learning techniques
for inducing *tree-based wrappers.* One of them is multi-view inductive logic
programming (Cohen, Hurst, & Jensen, 2003) which enriches string-based

---

Search Results for: *Bruss*

Found *26* Total matches in *1.32067* seconds

Score: *1*
Name: Ingo Bruss   **First Entered:** 03/07/96
Email: bruss@mach.uni/karsruhe.de

Score: *1*
Name: Trevor Bruss   **First Entered:** 20/05/97
Email: bruss@pz621.inland.com

Score: *1*
Name: Trevor Bruss   **First Entered:** 04/08/98
Email: bruss@virgo.valpo.edu

Figure 1   A rendered HTML document returned by some Web search machine.

machine learning techniques by tree views. Another is based on finite-state techniques (Kosala, Bruynooghe, Van den Bussche, & Blockeel, 2003; Raeymaekers, Bruynooghe, & Van den Bussche, 2005). This approach differs from ours in that it represents monadic queries by a local tree languages. Induction starts from learning algorithms from positive examples. For a more detailed comparison see Section 6.

In the present article – extending on Carme, Lemay, and Niehren (2004a) and Carme, Gilleron, Lemay, and Niehren (2005) – we combine features of both, machine learning techniques and visual specification. We define an interactive setting for inducing tree based wrappers. We are interested in machine learning techniques that can infer node selection queries in trees. We consider an HTML document as a tree structure obtained by parsing the document. An example for a rendered HTML document is given in Figure 1, and its parse tree in Figure 2. The HTML document contains the email addresses of three persons, called Ingo Bruss, Trevor Bruss, and Trevor Bruss again, each of which is described in its own HTML table. The names of these persons as well as their email addresses are contained in hyperlinks, which correspond to nodes of the parse tree that are labeled by tag `A`.

Suppose that we want to extract all email addresses in this HTML document. With respect to the parse tree, this wish can be expressed by the monadic node selection query that selects all hyperlinks in the last row of a table. In XPath like notation (but with additional axes), this query can be defined as follows:

```
/descendant-or-self::TABLE/lastchild::*/descendant::A/child::*
```

Start at the root, go downwards to a node labeled by TABLE, move to its last child, then move to some A-labeled offspring and finally select one of its children. The same query can be expressed by a conjunctive formula of

Figure 2 The parse tree of the HTML document in Figure 1.

first-order logic with one free node variable x:

$$\{x \mid \exists x_1 \exists x_2 \exists x_3.\ label_{\texttt{TABLE}}(x_1) \wedge lastchild(x_1, x_2)$$
$$\wedge descendant(x_2, x_3) \wedge label_{\texttt{A}}(x_3) \wedge child(x_3, x)\}$$

One might opt for a more precise query that further constrains the path from the TABLE to the A labeled node. For tree structures generated by the same search machine, however, such a refinements seem unnecessary to correctly select the informative nodes.

Another task of interest might be to extract all pairs of names and email addresses. This binary query can be expressed by a conjunctive first-order formula with two free variables. In what follows, however, we will restrict ourselves to monadic queries. This is because most existing systems rely on monadic queries for treating more complex binary or n-ary cases (Baumgartner et al., 2001; Muslea et al., 2001; Kosala et al., 2003).

We thus consider the problem of extracting sets of informative nodes in trees, that are generated homogeneously by some back-end database or search machine. We fix a formal query representation language and split the selection task into two subproblems:

1. find a representation of an appropriate monadic query for informative nodes in trees; here we want to use machine learning techniques, rather than manual query specification or programming.
2. answer the query in a tree: given a tree and a query definition, compute the set of nodes that the query selects in the tree.

Suitable query representation languages thus need to satisfy a number of constraints.

Expressiveness: All useful monadic queries should be definable in the chosen query representation language. Monadic second-order logic (MSO) is the yardstick for measuring the expressiveness of query languages for trees and thus wrapper language in information extraction from tree structured documents (Gottlob & Koch, 2002). It allows to define all regular

node selection queries in trees, which is adequate for many information extraction tasks.

Efficiency: We need efficient algorithms for answering monadic queries given a tree a representation of the query in our language.

Learnability: The objective is to induce monadic queries from trees in some of the nodes are annotated that are to be extracted. In the simpler case, we consider completely annotated Web pages in which all nodes are annotated that are subject to extraction. Because complete annotation may be unfeasible, we also consider partial annotations. Therefore we need to consider the problem of learning monadic queries from completely or partially annotated trees.

In this article, we propose two new query representation formalisms for monadic queries that are based on tree automata, *Node Selecting Tree Transducers (*NsTTs*) and *Pruning Node Selecting Tree Transducers (pNsTTs)*. We introduce NsTTs and pNsTTs in Section 2, with the objective to illustrate the basic ideas and concepts related to tree automata for unranked trees and monadic queries by examples.

In Section 3, we show that these query representation languages satisfy some of the algorithmic learnability requirements. What we need is an efficient test to determine whether a tree automaton is functional and thus an NsTT, or cut-functional and thus a pNsTT. The algorithm for cut-functionality is elaborated in the appendix (Section A).

Our results on efficiency and expressiveness of NsTTs and pNsTTs are postponed to the appendix as well (Sections B and C). There, we show that monadic queries defined by deterministic NsTTs or pNsTTs can be answered efficiently with linear time combined complexity. We then show that this holds even though deterministic NsTTs and pNsTTs capture the class of monadic queries definable in MSO. Both query languages NsTTs and pNsTTs are thus unrestricted from the view point of expressiveness[1].

The subsequent sections are devoted to learning algorithms for node selection queries, that are represented by NsTTs or pNsTTs. In Section 4, we present a new polynomial time algorithm in RPNI-style that we call tRPNI which learns monadic queries defined by deterministic NsTTs from completely annotated examples (for the target query). It relies on a subprocedure for testing membership to the class of NsTTs, i.e. for testing functionality.

In Section 5, we model the interactive process of learning monadic queries as active learning in the style of Angluin (1987), in order to measure the user's effort during the interaction. Only partially annotated trees are available in interactive Web information extraction. We solve this problem by using intelligent tree pruning heuristics and we present a practically feasible interactive learning algorithm for monadic queries defined by pNsTTs.

---

[1] Completeness with respect to expressiveness and learning are independent questions. Our learning algorithm tRPNI for NsTTs is a complete, while our algorithm $\mathrm{RPNI_{prune}}$ inducing pNsTTs is not (for realistic pruning functions).

It relies on a sub-procedure for testing membership to the class of pNsTTs, i.e. cut-functionality.

Finally, Section 6 discusses the integration of our learning algorithms into visually interactive Web information extraction. We have implemented our core learning algorithm in OCaml and have turned it into a visually interactive Web information extraction system – called SQUIRREL– by integration into the Mozilla Web browser. Experiments on realistic Web documents confirm excellent quality with very few user interactions – annotations and corrections – during wrapper induction.

2 Node Selecting Tree Transducers: NsTTs and pNsTTs

Parse trees of HTML or XML documents are sibling ordered unranked trees whose nodes may have an unbounded list of children. In this section, we develop query representation languages for monadic node selection queries in unranked trees, that we will use for query induction by methods of grammatical inference. We rely on tree automata for unranked trees. We propose two kinds of tree automata, that we call *node selecting tree transducers (*NsTT*s) and *pruning node selecting tree transducers (p*NsTT*s)*.

We start with a gentle introduction to NsTTs and pNsTTs, whose objective is to illustrate the basic ideas and concepts related to tree automata for unranked trees and monadic queries by examples. Algorithmic learnability aspects will be discussed in Section 3. Questions on efficiency and expressiveness for NsTTs and pNsTTs are answered in the appendix (Sections B and C).

2.1 Monadic queries in unranked trees

Let $\mathbb{N} = \{1, 2, \ldots\}$ be the set of natural numbers without 0 and $\mathbb{B} = \{0, 1\}$ be the set of Booleans. Let Σ be a finite set, the set of *node labels*, also called the *signature*.

An *unranked tree* over Σ is a finite tree with Σ-labeled nodes whose edges are totally ordered from the left to the right. An unranked tree t over Σ is a pair consisting of a label $a \in \Sigma$ and a possibly empty list of unranked trees $t_1, \ldots, t_n$: the following abstract syntax:

$$t ::= a(t_1, \ldots, t_n) \qquad \text{where } n \in \mathbb{N} \cup \{0\} \text{ and } a \in \Sigma$$

As a running example, we consider the following unranked tree with labels in $\Sigma = \{L, F, A, H, W\}$:

$$\textit{films} = L(F(A, H), F(A, W), F(A, H))$$

This tree represents a list (L) of three films (F), two of which are directed by Hitchcock (H) and one by Wenders (W). The letter (A) represents the

Figure 3 The unranked tree $films = L(F(A, H), F(A, W), F(A, H))$. We want to query for the actor lists (A) in the films by Hitchcock (H).

actor lists of these films (which would rather expand to a larger subtree in realistic examples). We draw trees as graphs, as for instance in Figure 3.

As in this example, we will systematically ignore textual contents of HTML or XML documents in this article. This design decision clearly simplifies our tasks, which nevertheless remains useful for practical applications to Web information extraction, as we will see.

We write $nodes(t)$ for the set of nodes of a tree t. We identify a node with its relative address, so that $nodes(t) \subseteq \mathbb{N}^*$. The node 21, for instance, is the first child of the second child of the root. In our example tree $films$, this node is labeled by A. We write $t(v)$ for the the label of some $v \in nodes(t)$, for instance:

$$films(21) = A$$

We say that two trees have *the same shape* if they have the same set of nodes. This definition works, since we identify nodes with their relative address from the root.

Definition 1 *A* monadic query *in unranked Σ-trees is a function q that maps unranked trees t over Σ to subsets of their nodes:*

$$q(t) \subseteq nodes(t)$$

The function $nodes$ itself, for instance, is a monadic query that maps a tree t to the set of all its nodes. A less trivial example is the query $actors\_hitchcock$ which asks for the set of all actor lists in films by Hitchcock. Applied to the example tree $films$, this query yields the following set of nodes:

$$actors\_hitchcock(films) = \{11, 31\}$$

Each actor list of a Hitchcock film corresponds to some node labeled by A, that has a parent labeled by F, and a next sibling to the right that is labeled by H. Furthermore, the F-parent should belong to a list of films, so that its own parent is labeled by L and all its siblings by F.

2.2 Tree automata for unranked trees

Tree automata are machines that recognize sets of trees. Tree automata are traditionally developed for binary or ranked trees (Comon, Dauchet,

Gilleron, Jacquemard, Lugiez, Tison, & Tommasi, 1997), and indeed, the situation gets different for unranked trees.

A large number of different automata notions for unranked trees and querying have been proposed recently. Besides so called *unranked tree automata (UTAs)* by Brüggemann-Klein, Wood, and Murata (2001) that are most popular in the database community (Libkin, 2005) and identical with the much earlier *pseudo-automata* by Thatcher (1967), there are so called *query automata* (Neven & Schwentick, 2002), forest grammars of Neumann and Seidl (1998), automata on binary encodings of unranked trees (Frick, Grohe, & Koch, 2003) and the algebraic recognizably notion of *stepwise tree automata* by Carme, Niehren, and Tommasi (2004b).

With respect to query induction as proposed in this article, we need an automata notion for unranked trees that comes with a good notion of bottom-up determinism, so that the Myhill-Nerode theorem on minimization holds. Martens and Niehren (2007) discuss the available alternatives. They show that unique minimization fails for deterministic UTAs, and argue that stepwise tree automata yield smaller minimal automata than the inverted standard binary encoding (up to a linear factor depending on the size of the alphabet) and *parallel UTAs* (Raeymaekers & Bruynooghe, 2004; Cristau, Löding, & Thomas, 2005). We thus choose to work with stepwise tree automata.

Stepwise tree automata can be understood as standard tree automata that operate on binary encodings of unranked trees (not the standard encoding however). This fact will come in handy for our more formal results on NSTTs and pNSTTs in Section 3. In this section, however, we will use finite word automata syntax to present stepwise tree automata. Semantically, however, they behave such as tree automata. The relationship to standard tree automata will be clarified on need (in Section 3.2).

Definition 2 A stepwise tree automaton *for unranked trees over Σ is a finite word automaton A with alphabet$(A) = $ states$(A) \uplus \Sigma$ such that*

- *there is a unique initial state: initial$(A) = \{q_I\}$.*
- *if $q_0 \xrightarrow{a} q'$ in rules(A) for some $a \in \Sigma$ then $q_0 = q_I$, and*
- *no transition $q_1 \xrightarrow{q_2} q$ in rules(A) satisfies $q = q_I$ nor $q_2 = q_I$.*

A run r of a stepwise tree automaton A on an unranked tree t labels the nodes of t by states of A. It is a function of type $r : nodes(t) \to states(A)$ that can be understood as a tree over $states(A)$ having the same shape as t. Value $r(v)$ is the state that run r assigns to node $v \in nodes(t)$. If v has n children $v1, \ldots, vn$ then $r(v)$ is a state obtained by running A on the following word:

$$t(v) \; r(v1) \; \ldots \; r(vn)$$

This word starts with the label of v in t and continues by the sequence of states that r assigns to the children of v. A run is *successful* if it maps the root of t to a final state of A. A stepwise tree automaton A accepts an unranked tree t if there exists a successful run of A on t. The language

Figure 4 A stepwise tree automaton that accepts film lists, and the *films* tree annotated by a successful run of this automaton. State 6 of the root is obtained by evaluating the word $L454$ by the automaton.

$L(A)$ of a stepwise tree automaton A is the set of all unranked trees that A accepts. An example for a stepwise tree automaton is given in Figure 4. This automaton accepts all lists containing films by Hitchcock and Wenders. The *films* tree annotated by a run of this automaton is drawn in Figure 4.

Definition 3 *A stepwise tree automaton A is (bottom-up) deterministic, if it is a left-to-right deterministic finite automaton, i.e., for ever every $q \in states(A)$ and $a \in alphabet(A)$ there exists at most one q' in $states(A)$ such that $q \xrightarrow{a} q'$ belongs to $rules(A)$.*

What might be less clear at the time being is that every stepwise tree automaton can be determinized without altering its tree language, and that all regular languages of unranked trees are recognized by a unique minimal deterministic stepwise tree automaton (up to isomorphism). This will become obvious from the tree automata perspective adopted in Section 3.

2.3 Nstts for unranked trees

The next goal is to use stepwise tree automata in order to define monadic queries in unranked trees. There are many possibilities to do so, Nstts provide one of them, pNstts another. Two further methods will be discussed and used in Section B.

An *annotated tree* over Σ is a tree over $\Sigma \times \mathbb{B}$. Every annotated tree can be decomposed into a product $t \times \beta$ between a Σ-tree t and a $\mathbb{B}$-tree β of the same shape. For convenience, we will often write a_b for pairs (a, b). For instance:

$$L(F(A, H)) \times 0(0(1, 0)) = L_0(F_0(A_1, H_0))$$

A *completely annotated tree for a query q in Σ-trees* is an annotated tree of the form:

$$t \times q(t)$$

**Figure 5** Three film lists completely annotated by the query *actors_hitchcock*

Figure 6 A deterministic NSTT defining the query *actors_hitchcock* and one of its successful runs.

for some Σ-tree t. Here, we freely consider the set $q(t)$ as a $\mathbb{B}$-tree of the shape of t, which relabels selected nodes of t to 1 and all others to 0. Examples for three completely annotated trees for the query *actors_hitchcock* are given in Figure 5.

The idea of NSTT*s* is to represent a monadic query q by a tree automaton that recognizes the set of completely annotated trees for q except for some $t \times q(t)$ where $q(t) = \emptyset$. Such languages L of annotated trees are *functional* in that for every Σ-tree t there exists at most one Boolean tree β such that $t \times \beta \in L(A)$ (namely $\beta = q(t)$).

Definition 4 *An* NSTT *over* Σ *is a stepwise tree automaton* A *over* $\Sigma \times \mathbb{B}$ *that recognizes a functional language of annotated trees. Every* NSTT *A defines a monadic query* q_A *such that for all unranked* Σ*-trees* t:

$$q_A(t) = \{v \mid \exists \mathbb{B}\text{-tree } \beta \text{ s.t. } t \times \beta \in L(A), \beta(v) = 1\}$$

As example, we present in Figure 6 a deterministic NSTT that defines the query *actors_hitchcock*. Let us recall that we freely write a_b instead of (a, b).

A query q is *represented by an* NSTT A if A defines the query q, i.e. $q_A = q$. As we do not care about trees t such that $q(t) = \emptyset$, a query q may be represented by several NSTTs. Two NSTTs A and A' are said to be *equivalent* if $q_A = q_{A'}$.

2.4 Pruning NSTT*s for unranked trees*

One of the disadvantages of NSTTs is that they need to model all subtrees, even those that are irrelevant for node selection. The NSTT in Figure 4, for

Figure 7 Two trees obtained by pruning the tree *films* annotated by query *actors_hitchcock*. The Boolean annotations are compatible with this query.

instance, has rules to recognize films that are not directed by Hitchcock (but by Wenders). In Web information extraction, such information is typically not available or difficult to learn from the very few given examples.

We propose pNsTTs to solve this problem. These are machines that operate like NsTTs except they first prune trees non-deterministically and then select nodes from the pruned tree. This way, they are able to ignore all pruned subtrees.

Let T be a special constant, that we use to mark leaves in pruned trees where complete subtrees have been cut off. A *pruned tree* over a signature Σ is an unranked tree over the signature

$$\Sigma_{\mathrm{T}} = \Sigma \cup \{\mathrm{T}\}$$

whose inner nodes are labeled in Σ. Only leaves of pruned trees may be labeled by T. For every Σ-tree t, let $cuts(t)$ be the set of Σ_{T}-trees obtained by pruning some subtrees of t, i.e., by substituting some complete subtrees of t by T. Two pruned trees t_1 and t_2 are said to be *compatible* if they can be obtained by pruning from some common third tree, i.e., if there is a tree t such that $t_1, t_2 \in cuts(t)$.

See Figure 7 for some examples of pruned annotated trees, which have the signature $(\Sigma \times \mathbb{B})_{\mathrm{T}}$. It will be convenient to identify pruned annotated trees over $(\Sigma \times \mathbb{B})_{\mathrm{T}}$ with trees over $(\Sigma \times \mathbb{B}) \uplus \{\mathrm{T}_{\mathrm{T}}\}$, i.e., we implicitly identify T with the pair (T, T) which is $\mathrm{T} = \mathrm{T}_{\mathrm{T}}$. This permits us to decompose all trees over $(\Sigma \times \mathbb{B})_{\mathrm{T}}$ into products $t \times \beta$ of trees t over Σ_{T} and trees β over $\mathbb{B}_{\mathrm{T}}$. Intuitively, the annotation of T by T means that we don't know whether nodes in pruned subtrees will be selected or not.

We call a language L of $(\Sigma \times \mathbb{B})_{\mathrm{T}}$-trees *cut-functional* if for all $t_1 \times \beta_1, t_2 \times \beta_2 \in L$ such that t_1 and t_2 are compatible, it holds that β_1 and β_2 are compatible. This means that different prunings t_1 and t_2 of the same third tree t will never carry contradictory information. A Σ_{T}-tree t is called *unpruned* if T does not occur in the set of node labels used by t. Note that cut-functional languages of unpruned trees (and thus with signature $\Sigma \uplus \mathbb{B}$) are functional.

Consider for instance the two annotated prunings of *films* in Figure 7. Their Boolean annotations are compatible in that both of them require the root to be unselected. For all other nodes, there exist no overlapping annotations. The fact that the left-most (resp. right-most) A-node is to be selected is available on the left-most (resp. right-most) pruning only.

**Figure 8** A deterministic pNSTT defining the query *actors_hitchcock*.

Figure 9 Two runs of different prunings of *films* by the pNSTT in Figure 8. Each of them selects one of the A-nodes in Hitchcock's films.

Definition 5 *A* pNSTT *over Σ is a tree automaton over $(\Sigma \times \mathbb{B})_T$ whose language is cut-functional. Every pNSTT A defines the query q_A such that for all Σ-trees t:*

$$q_A(t) = \{v \in nodes(t) \mid \exists\ t' \in cuts(t), \exists \beta' :\ t' \times \beta' \in L(A), \beta'(v) = 1\}$$

The nodes of $q_A(t)$ can thus be computed non-deterministically, by guessing a pruning of t' of t and inspecting the Boolean annotation β' such that $t' \times \beta' \in L(A)$, which is unique by cut-functionality if it exists.

A deterministic pNSTT that recognizes the actors lists of Hitchcock's films is given in Figure 8. Indeed, it leaves the structure of all films by other directors than Hitchcock undefined, since these can be pruned to T.

A query q is *represented by a* pNSTT A if A defines the query q, i.e. $q_A = q$. A query q may be represented by several pNSTTs.

3 Testing Functionality in Polynomial Time

Queries defined by NSTTs and pNSTTs have three important algorithmic or descriptive properties that render them suitable for interactive query induction.

Membership: whether a deterministic tree automaton over $\Sigma \times \mathbb{B}$ is functional, and hence an NSTT can be decided in polynomial time. Whether a tree automaton over $(\Sigma \times \mathbb{B})_T$ is cut-functional and thus a pNSTT can be decided in polynomial time.

Efficiency: monadic queries defined by NSTTs or pNSTTs can be answered in polynomial time

Expressiveness: deterministic NSTTs can express all monadic second order definable monadic queries, as well as deterministic pNSTTs.

In this section, we prove the first property, by presenting efficient algorithms for membership testing, on which our learning algorithms will rely. Efficiency will be treated in Section B. It is fundamental because correctness of the hypothesis NSTT or pNSTT must be checked frequently in the interactive setting. Also, an information extraction program must be efficient. Expressiveness will be treated in Section C. We show that deterministic NSTTs and pNSTTs can express all MSO-definable queries. The learning algorithm for NSTTs will be able to infer all MSO-definable queries, when given enough examples (which is unrealistic in practice). The choice of a practical pruning function will restrict the class of learnable queries by the learning algorithm of pNSTTs.

3.1 Tree automata for binary trees

Properties depending on determinism and functionality are easier to establish for standard tree automata on binary trees. As we will show, they can be lifted to stepwise tree automata by exploiting a particular binary encoding of unranked trees, called Currying in (Carme et al., 2004b). Note that the notion of determinism for tree automata on unranked trees works out smoothly for stepwise tree automata in contrast to more standard notions of tree automata for unranked trees (Martens & Niehren, 2007).

Standard tree automata (Comon et al., 1997) are traditionally defined for binary trees, i.e., unranked trees where every node has either no or exactly 2 children. For sake of simplicity, we will develop our algorithms for binary trees, and then lift them to unranked trees via binary encoding.

A *tree automaton* A for binary trees over some signature Σ consists of a finite set $states(A)$, a subset $final(A) \subseteq states(A)$, and a set of rules of two kinds where $a \in \Sigma$ and $p, p_1, p_2 \in states(A)$.

Standard tree automata can be seen as particular stepwise tree automata, except that they decompose binary rules into three steps.

As a consequence, the notions of runs, successful runs, and accepted tree languages carry over immediately. We write $runs_A(t)$ for the set of runs of automaton A on tree t and $succ\_runs_A(t)$ for the subset of successful runs. The notion of bottom-up determinism carries over from stepwise to standard tree automata, too. A tree automaton is bottom-up deterministic if none of its rules have the same left hand side. It is well known that every standard tree automaton can be determinized without affecting its tree language. Further notions that carry over to binary trees without any changes are functionality, cut-functionality, NSTTs, and pNSTTs.

Figure 10 Curried binary encoding: curry(L (F (A,H),F (A,W),F (A,H))) = L @(F @A @H)@(F @A @W)@(F @A @H).

3.2 Binary encoding of unranked trees

We next relate deterministic stepwise tree automata for unranked trees to deterministic standard tree automata over binary trees. This shows that subsequent learnability results for deterministic standard tree automata can be lifted to deterministic stepwise tree automata. The base idea is to encode unranked trees into binary trees by using Currying. The better known first-child next-sibling encoding fails to do this job; it does not preserve determinism.

Some more care is needed when determinizing stepwise tree automata. The usual determinization algorithm for finite automata fails since it alters the tree language of a stepwise tree automaton. Instead, a stepwise tree automaton needs to be understood as a standard tree automata operating on Curryied encodings. Determinisation for standard tree automata is then doing the job.

For applying Currying we consider unranked trees as lambda terms. The term $L(F, F, F)$ for instance, is a lambda term that describes the application of function L to a triple of arguments. Its Curried encoding is the binary tree $(((L@F)@F)@F)$ which describes the application of function L to its arguments stepwise one by one. The Curried binary encoding of the unranked tree *films* is displayed in Figure 10.

Stepwise tree automata A over Σ can be understood as syntax for standard tree automata A^b for binary trees over $\Sigma \uplus \{@\}$, which operates on Curried binary encodings of unranked trees. The translation is as follows, where $a \in \Sigma$ and $p_0 \in initial(A)$, and $p_1, p_2, p \in states(A) - initial(A)$:

$$p_0 \xrightarrow{a} p \in rules(A) \quad \text{iff} \quad a \to p \in rules(A^b)$$
$$p_1 \xrightarrow{p_2} p \in rules(A) \quad \text{iff} \quad @(p_1, p_2) \to p \in rules(A^b)$$

The translation of the stepwise automaton in Figure 4 is given in Figure 11. Generally, it holds that:

$$A \text{ is deterministic} \quad \Leftrightarrow \quad A^b \text{ is deterministic}.$$

This follows immediately from the definitions. As an immediate consequence, it follows that every stepwise tree automaton can be determinized.

Labels: $\{@, A, H, W, F, L\}$
States: $\{1, 2, 3, 4, 5, 6, 8, 9\}$
Rules: $A \to 1$ $@(8,1) \to 9$
        $H \to 2$ $@(9,2) \to 4$
        $F \to 8$ $@(6,4) \to 6$
        $L \to 6$ $@(6,5) \to 5$
        $W \to 3$ $@(9,3) \to 5$
 $@(6,5) \to 6$
Final: $\{6\}$

Figure 11 A standard tree automaton for binary encodings of film lists, corresponding to the stepwise tree automaton in Figure 4.

Furthermore, the language of both automata correspond modulo Currying, i.e., for all unranked trees t over $\Sigma \uplus \{@\}$:

$$t \in L(A) \;\Leftrightarrow\; curry(t) \in L(A^b)$$

This is because both automata produce runs that correspond precisely except for intermediate results stored in @-labeled nodes (which could be annotated to the edges of unranked trees). Since $states(A^b) = states(A) - initial(A)$, it also follows that there exists unique minimal deterministic stepwise tree automata for all regular languages of unranked trees.

The correspondence lifts to monadic queries defined by NSTTs or pNSTTs even though this needs a little care. An NSTT A over Σ is a tree automaton over $\Sigma \times \mathbb{B}$, so that A^b is an automaton over $(\Sigma \times \mathbb{B}) \uplus \{@\}$. Unfortunately, A^b is not an NSTT over $\Sigma \times \{@\}$, but this problem can be resolved by replacing @ by $@_0$, which means that we do not want to extract any auxiliary nodes introduced by the binary encoding, which is of course consistent with our intuition. Modulo the equation $@ = @_0$ it holds for every tree automaton A over $\Sigma \times \mathbb{B}$, and Σ-tree t:

$$A \text{ is functional } \Leftrightarrow A^b \text{ is functional}$$
$$q_A(t) = q_{A^b}(curry(t))$$

Hence, A is an NSTT for unranked trees over Σ if and only if A^b is an NSTT for binary trees over $\Sigma \uplus \{@\}$, and that both automata define the same query modulo the natural correspondence between nodes in unranked trees to leafs in binary encodings. As a consequence, we can reduce functionality testing and NSTT query answering on unranked trees to the respective problems on binary trees.

The same arguments carry over to pNSTTs without any further changes. Modulo the equation $@ = @^0$ it holds for every tree automaton A over $(\Sigma \times \mathbb{B})_T$, and Σ-tree t:

$$A \text{ is cut-functional } \Leftrightarrow A^b \text{ is cut-functional}$$
$$q_A(t) = q_{A^b}(curry(t))$$

3.3 Testing functionality and cut-functionality

We next show how to test whether a tree automaton for binary trees over $\Sigma \times \mathbb{B}$ is functional. As argued above, this is sufficient to solve the same problem for unranked trees

We reduce functionality testing to unambiguity testing, a problem for tree automata that can be solved in polynomial time (Seidl, 1989).

A tree automaton is *unambiguous* if it permits at most one successful run per tree. Note that every deterministic tree automaton is unambiguous, since it permits at most one run per tree, may it be successful or not.

The *projected automaton* $\pi(A)$ over Σ has $states(\pi(A)) = states(A) \times \mathbb{B}$, $final(\pi(A)) = final(A) \times \mathbb{B}$, and the rules are given by the following inference scheme where $a \in \Sigma$ and $b, b_1, b_2 \in \mathbb{B}$:

$$\frac{(a,b)(p_1,p_2) \to p \in rules(A)}{a((p_1,b_1),(p_2,b_2)) \to (p,b) \in rules(\pi(A))} \qquad \frac{(a,b) \to p \in rules(A)}{a \to (p,b) \in rules(\pi(A))}$$

Lemma 1 $r \in runs_A(t \times \beta)$ *iff* $r \times \beta \in runs_{\pi(A)}(t)$.

Lemma 2 *A deterministic tree automaton A over $\Sigma \times \mathbb{B}$ is functional and thus an* NSTT *if and only if its projection $\pi(A)$ to Σ is unambiguous.*

Proof For the one direction, let A be a deterministic NSTT. Suppose $r_1 \times \beta_1, r_2 \times \beta_2 \in succ\_runs_{\pi(A)}(t)$ for some $t \in tree_\Sigma$. Lemma 1 yields $r_1 \in succ\_runs_A(t \times \beta_1)$ and $r_2 \in succ\_runs_A(t \times \beta_2)$. The functionality of A implies that $t \times \beta_1 = t \times \beta_2$ and thus $\beta_1 = \beta_2$. The unambiguity of A yields $r_1 = r_2$ so that $r_1 \times \beta_1 = r_2 \times \beta_2$. This proves that $\pi(A)$ is unambiguous.

For the converse, assume that $\pi(A)$ is unambiguous and suppose $t \times \beta_1, t \times \beta_2 \in L(A)$. Let $r_1 \in succ\_runs_A(t \times \beta_1)$ and $r_2 \in succ\_runs_A(t \times \beta_2)$. Lemma 1 yields $r_1 \times \beta_1, r_2 \times \beta_2 \in succ\_runs_{\pi(A)}(t)$. The unambiguity of $\pi(A)$ implies $\beta_1 = \beta_2$, i.e., $L(A)$ is functional.

Proposition 1 *Whether a deterministic tree automaton over $\Sigma \times \mathbb{B}$ is functional and hence an* NSTT *can be tested in cubic time.*

It is sufficient to prove this proposition for standard tree automata over binary trees, so that it carries over to stepwise tree automata for unranked trees.

Proof Let A be a deterministic tree automaton over $\Sigma \times \mathbb{B}$. By Lemma 2 it is sufficient to compute the Σ-projection $\pi(A)$ and to test whether this automaton is unambiguous. This can be done in cubic time.

It remains to show that we can equally decide membership to the class of deterministic pNSTT in polynomial time. This is stated by the following proposition that is proved in Section A.

Proposition 2 *Cut-functionality of deterministic* pNSTT*s can be tested in cubic time.*

4 Learning NsTTs from Completely Annotated Examples

We have shown that monadic node selection queries can be represented by deterministic NsTTs which are tree automata. It is known that regular tree languages represented by deterministic bottom-up tree automata are learnable from polynomial time and data (Oncina & García, 1993). But this result holds only when positive and negative examples are available, i.e. trees belonging to the target language or not. In our query learning framework, however, completely annotated trees for the target query will be available instead. These express a number of positive and implicit negative examples. We will prove in this section that monadic node selection queries represented by NsTTs are learnable from polynomial time and data, provided that the data contains completely annotated examples for the target query. The difficulty is to show, that completely annotated examples are indeed equivalent to arbitrary samples of positive and negative examples.

4.1 Identification from polynomial time and data

We recall the learning model from polynomial time and data (Gold, 1967, 1978; de la Higuera, 1997), and adapt it to the problem of tree language inference using deterministic tree automata. An example is a couple (t, b) where t is a tree and b is its Boolean label. Positive examples are labeled by 1 and negative examples by 0. A (tree) sample is a finite set of examples. A tree automaton A is compatible with an example (t, b) if $(t \in L(A)) \Leftrightarrow b$, it is compatible with a sample if it is compatible with all its examples. We also say in this case that the sample is consistent with A. The *size* of a sample is the total number of nodes of all trees in the sample.

Definition 6 *Tree languages represented by a class of tree automata $\mathcal{C}$ are said to be* identifiable from polynomial time and data *if there exist two polynomials p_1 and p_2 and an algorithm learner such that:*

- *with as input a sample S of size m, learner returns a tree automaton $A \in \mathcal{C}$ compatible with S in $O(p_1(m))$ time;*
- *for each tree automaton $A \in \mathcal{C}$ of size n, there exists a sample – called the characteristic sample for A – char(A) of cardinality less than $p_2(n)$ such that, with input a sample S that contains char(A), learner returns a tree automaton $A' \in \mathcal{C}$ equivalent with A.*

Regular tree languages represented by deterministic tree automata are identifiable from polynomial time and data (Oncina & García, 1993). In the sequel of the paper we mainly use the existence of a polynomial time learning algorithm *learner*. It is set to be the RPNI algorithm for deterministic tree automata which is defined in (Oncina & García, 1993) as an extension of the RPNI algorithm for deterministic string automata (Oncina & Garcia, 1992).

The reader should note that the model for tree languages differs from the model for string languages on one important point: it is required that the *cardinality* of the characteristic sample is polynomial in the size of the target automaton whereas for string languages it is required that the *size* of the characteristic sample is polynomial in the size of the target automaton. This is of course weaker and due to the fact that the size of trees – measured as the number of nodes – in the characteristic sample can be exponential. Indeed, consider the tree language containing a single balanced binary tree of height n on the alphabet $\Sigma = \{a, f\}$ where a is a constant and f a binary symbol. The characteristic sample must contain this tree which has $2^{n+1} - 1$ nodes while the corresponding tree automaton has only $n + 1$ states. This blow-up is indeed due to the intrinsic nature of trees. This problem could be solved with a new representation schema for trees in the characteristic sample: represent a tree by a deterministic automaton recognizing the language reduced to this unique tree. The proof is out of scope of the paper because we will only consider the *learner* defined with the RPNI algorithm presented in (Oncina & García, 1993).

The learning model for tree languages is not well suited for node selection queries in trees. We can represent a node selection query by a regular tree language over $\Sigma \times \mathbb{B}$. But, for learning node selection queries, only completely annotated trees for the target query will be available. Thus, if q is the target query, only positive examples $(t \times q(t), 1)$ will be available. We will refer to them as completely annotated examples. Also, negative examples are useless as they can be deduced from completely annotated ones: for a completely annotated example $t \times \beta$, every other annotated tree $t \times \beta'$ ($\beta' \neq \beta$) is negative. Two query representations are equivalent if they represent the same query. This leads to the following definition:

Definition 7 *Node selection queries represented in a class of automata C are said to be* identifiable from polynomial time and data from completely annotated examples *if there exist two polynomials p_1 and p_2 and an algorithm learner such that:*

- *for all sets S of annotated trees of size m, learner(S) returns a query representation $A \in C$ in time $O(p_1(m))$ such that q_A is compatible with S;*
- *for all automata $A \in C$ of size n, there exists a set of completely annotated examples for q_A – called the characteristic sample for A – char(A) of cardinality less than $p_2(n)$ such that, for every sample S that contains char(A), learner(S) returns an automaton $A' \in C$ such that $q_A = q_{A'}$.*

We will simply talk of identifiability when we do not impose polynomial constraints.

4.2 tRPNI: an RPNI extension for deterministic NSTTs

We now want to prove that node selection queries represented by deterministic NSTTs are identifiable from polynomial time and data. We first show

that one can easily deduce identifiability of node selection queries represented by deterministic NSTTs in this framework from the previous result of identifiability of tree languages represented by deterministic automata – polynomial constraints set aside.

Note that the proof presented here leads to an exponential time algorithm – as we will have to enumerate explicitly all negative examples induce by a functional set of completely annotated examples. Polynomiality can nevertheless be reached by using functionality directly. This leads us to a new variation of the learning algorithm RPNI which can identify node selection queries represented by NSTTs.

Lemma 3 *Node selection queries represented by deterministic NSTTs are identifiable from completely annotated examples.*

Proof First note that annotated tree languages represented by deterministic NSTTs are identifiable as they are deterministic tree automata over $\Sigma \times \mathbb{B}$. Let *learner* be an algorithm according to Definition 6 and let *char* be a function that computes the characteristic sample for any deterministic tree automaton over $\Sigma \times \mathbb{B}$ according to Definition 6.

To prove the lemma, we have to show that we can infer an NSTT defining the target query from completely annotated examples. A completely annotated example $t \times \beta$ for a query q defines a positive example $(t \times \beta, 1)$ and negative ones $(t \times \beta', 0)$ for all $\beta' \neq \beta$. Thus, we define a function pn that takes as an input a set S of annotated trees and outputs a sample over $\Sigma \times \mathbb{B}$.

$$pn(S) = \{(t \times \beta', b) \mid t \times \beta \in S, b \Leftrightarrow (\beta = \beta')\}$$
$$\cup \ \{(t \times \emptyset, 0) \mid t \in T_\Sigma)$$

We next define *learner'* and *char'* in the following manner:

$$learner'(S) = learner(pn(S))$$
$$L_q = \{t \times q(t) \mid q(t) \neq \emptyset\}$$
$$char'(q) = \{t \times \beta \in L_q \mid (t \times \beta', b) \in char(L_q)\}$$

First note that $char'(q)$ is a set of completely annotated examples for q. Second, we show that $char(L_q) \subseteq pn(char'(q))$.

- For positive examples $(t \times \beta', 1) \in char(L_q)$, we have $t \times \beta' \in L_q$ and thus $t \times \beta' \in char'(q)$, so that $(t \times \beta', 1) \in pn(char'(q))$.
- For negative examples, lets assume that $(t \times \beta', 0) \in char(L_q)$. If $q(t) \neq \emptyset$ then $t \times q(t) \in L_q$ and thus $t \times q(t) \in char'(q)$. Furthermore, $\beta' \neq q(t)$ so that $(t \times \beta', 0) \in pn(char'(q))$.

Let q be a node selection query and $S \supseteq char'(q)$ a sample consistent with L_q. Then $pn(S) \supseteq char(q)$ and $pn(S)$ consistent with L_q. As *learner* identifies regular tree languages represented by deterministic tree automata over $\Sigma \times \mathbb{B}$, $learner(pn(S)) = learner'(S)$ outputs a deterministic tree automaton A such that $L(A) = L_q$, as required for identifying queries represented by NSTTs.

The functions $learner'$ and $char'$ witness that monadic queries representable by deterministic NSTTs can be identifed for completely annotated examples (but not yet from polynomial time and data).   □

Our next objective is to adapt the functions $learner'$ and $char'$ in the previous proofs, in order to establish identifiability of NSTT-queries from polynomial time and data. There are two problems. First, algorithm $learner'$ computes the set of implicit examples $pn(S)$ induced by functionality, which may be of exponential in cardinality compared to S. This problem be solved by an alternative learning algorithm that always test the current automaton for functionality, rather than testing for consistency w.r.t. negative examples in $pn(S)$. And fortunately, functionality can be tested in P-time for deterministic automata over $\Sigma \times \mathbb{B}$. Second, we need to show for all NSTTs A that the sample $char'(q_A)$ is always of polynomial cardinality in the size of A. This is the case, since we can compute in polynomial time another NSTT recognizing L_{q_A}, so that L_{q_A} is definable by some deterministic NSTT of size polyomial in that of A.

Theorem 1 *Monadic node selection queries represented by deterministic NSTTs are identifiable from polynomial time and data from completely annotated examples.*

Proof First, let us recall the RPNI-algorithm (Oncina & García, 1993; Oncina & Garcia, 1992; Lang, 1992). RPNI inputs a sample S of positive and negative examples. It first computes a deterministic automaton which recognizes the set of positive examples in the sample called the initial automaton and denoted NSTT(S) – it is the largest deterministic tree automaton that recognizes exactly S. It then merges states exhaustively in some fixed order. A merging operation applies to the recent deterministic automaton A and two states $q_1, q_2 \in states(q)$ and returns a deterministic automaton det-merge(A, q_1, q_2). A deterministic merge is a merge followed by recursive merges needed to preserve determinism. For example, merging q_1 and q_2 in an automaton with rules $f(q_1) \to q_3$ and $f(q_2) \to q_4$ requires merging q_3 with q_4. A merging operation is licensed only if det-merge(A, q_1, q_2) is consistent with all negative examples in the sample. The algorithm tRPNI (described in Figure 12) behaves as RPNI except that it checks differently whether deterministic merging operation are licensed. It tests whether the language L of det-merge(A, q_i, q_j) is functional and whether $L \cap \{t \times \emptyset \mid t\} = \emptyset$. It thereby avoids to enumerate implicit negative examples. It checks functionality in polynomial time in the automaton size (Proposition 1), at this at most quadratically many times. Also, as the NSTT should not recognize trees that has no selectionned elements, we remove those trees (kept in a set N), and the algorithm has to check that the language $\pi_\Sigma(L) \cap N = \emptyset$. Therefore:

– The polynomial time condition for $learner$ equal to tRPNI is satisfied. Consistency with the input sample is also given: by construction, the inferred automaton always recognizes trees of S.

---

tRPNI(S)

$N = \{t \in S \mid t \times \beta \in S, \beta = \emptyset\}$
$A \leftarrow \text{NSTT}(\{t \times \beta \mid t \times \beta \in S, \beta \neq \emptyset\})$
For $i = 1$ **to** $|states(A)|$ **do**
    For $j = 0$ **to** $i - 1$ **do**
        $A' \leftarrow \text{det-merge}(A, q_i, q_j)$
        If A' is functional
        and $L(A) \cap \{s \mid s$ is tree over $\Sigma \times \{0\}\} \neq \emptyset$
        and $\pi_\Sigma(L(A)) \cap N = \emptyset$
        then $A \leftarrow A'$
            exit inner loop
        else skip
Output: A

Figure 12 Learning from completely annotated examples.

– The characteristic set of completely annotated examples is given in the proof of Lemma 3. Given an NSTT A such that $q = q_A$ we first compute another automaton $\tilde{A}$ so that $L(\tilde{A}) = L_q$. This can be done in polynomial time in the size of A, by intersecting A with an automaton recognizing $\{t \times \beta \mid t \in T_\Sigma, \beta \neq \emptyset\}$. We then compute $char(L_q)$ in polynomial time from $\tilde{A}$, and $char'(q)$ in polynomial time from $char(L_q)$. It follows that the cardinality of $char'(q)$ is polynomial in the size of A. Furthermore, given any sample S containing $char'(q)$, $t\text{RPNI}(S)$ behaves as the algorithm $learner'(S)$ described in this lemma (except that it does not explicitly enumerate negative examples) and therefore induces the target NSTT. The polynomial cardinality of the characteristic set of annotated trees is guaranteed by the polynomial cardinality of the characteristic sample of the target NSTT.

5 Interactive Learning of pNSTTs

The above framework has two main drawbacks:

– completely annotated examples are required, which may be a problem in practical cases where sometimes hundreds of nodes would have to be selected by the user. Therefore, one should deal with examples that are only partially annotated.
– Furthermore, it does not take into account the interactions between the user and the algorithm. The user does not simply provide a set of annotated examples. He should for example be able to give some annotations to the algorithm in order to have a candidate query that he could then correct if he is not happy with its behaviour.

To take into account those two points, a new learning model has to be defined. We present here the active learning model. The SQUIRREL system that we develop here is a learning algorithm for monadic node selection queries in this framework that allows the user - presented here as an oracle - to interact with a core learning algorithm, which will be a variant of tRPNI able to deal with partial annotations.

5.1 Active Learning Model

Deterministic finite automata are learnable in the MAT model (Angluin, 1987), i.e., by an active learner who may ask membership and equivalence queries to his teacher. This result extends to deterministic bottom-up tree automata (Sakakibara, 1990; Drewes & Hogberg, 2003). Such queries, however, are not well suited for active learning of node selection queries in trees. Indeed, a node selection query may be represented by a regular tree language over $\Sigma \times \mathbb{B}$. Thus, a membership query would need as input an annotated tree which is not manageable in practice. Moreover, a membership query answers yes or no whether the annotated tree is correctly annotated, but, if the answer is negative, a membership query gives no information on the annotation error. Also, recall that a monadic query can be represented by several languages of annotated trees. Therefore equivalence of tree languages is not adapted for equivalence of monadic node selection queries.

We introduce two new types of queries that the learner may ask to the teacher, *correct labeling queries (*CLQ*s)* and *query equivalence queries (*QEQ*s)*. This can be done with respect to an arbitrary formalism for representing queries – for instance by NSTTs. We denote query representations by A and the query they represent by q_A. A CLQ about target q is defined as follows:

```
┌──────────── Correct Labeling Query (CLQ) ────────────┐
│                                                      │
│  Input:     an annotated $\Sigma \times \mathbb{B}$-tree $t \times \beta$        │
│  Output:    yes if $q(t) = \beta$, else some node $v \in nodes(t)$   │
│             such that $\beta(v) \neq q(t)(v)$                      │
│                                                      │
└──────────────────────────────────────────────────────┘
```

A CLQ asks whether the current annotation hypothesis $t \times \beta$ is correct and requires a failure witness v otherwise. A QEQ about target q has the form:

```
┌──────────── Query Equivalence Query (QEQ) ────────────┐
│                                                       │
│  Input:     a query representation $A$                  │
│  Output:    yes if $q = q_A$, else some $\Sigma$-tree $t$ s. t. $q_A(t) \neq q(t)$. │
│                                                       │
└───────────────────────────────────────────────────────┘
```

A QEQ asks whether the current query hypothesis A is correct, and requires a failure witness t otherwise.

Before presenting the main result, let us discuss about queries in practice. The user (or query designer) plays the role of the oracle. This is not a problem for CLQ: the user can check on the current web page if it is well annotated and correct it otherwise. This may however be a problem for QEQ. We do not want to show the inferred NSTT to the user, as we want a "black-box" system usable even for non-expert users. The only thing that the user can check is whether the inferred query annotates correctly web pages that he proposes. The user simply answers "yes" to the QEQ (by actually stopping the learning process) when he is satisfied with the behavior of the proposed query on web pages processed so far. The possible consequence is that this query may not behave correctly on some unseen web pages, especially web pages that have a general design which is different to the one of processed web pages. Solutions to this problem are discussed in Section 6 and in the conclusion. Nevertheless QEQs and CLQs allow to define a rigorous framework for interactive learning of monadic queries. We have the following result:

Theorem 2 *Monadic node selection queries represented by deterministic NSTTs are learnable in polynomial time by active learning processes with correct labeling queries and query equivalence queries.*

Proof Let q be the target query and L be a target functional language over $\Sigma \times \mathbb{B}$ defining q. Let $\mathcal{A}$ be a learning algorithm with membership and equivalence queries for deterministic tree automata. At each call of $\mathcal{A}$ to a membership query $t \times \beta \in L$, we ask the correct labeling oracle with input $t \times \beta$ and answer yes if the CLQ answers yes and no if the CLQ outputs a node. At each call of $\mathcal{A}$ to an equivalence query, we ask the query equivalence query oracle with the same hypothesis tree automaton and answers yes if the QEQ outputs yes. If the QEQ outputs a Σ-tree t, we have to construct the annotated tree $t \times \beta$ which is a counterexample for the equivalence query. To do so, ask the correct labeling oracle until we get the annotated tree $t \times \beta$. The reader should note that the number of calls to the correct labeling oracle is bounded by the number of nodes to be selected in t. The modified algorithm is a learning algorithm with CLQs and QEQs. It outputs in polynomial time a minimal tree automaton A of the functional language L, i.e. a deterministic NSTT defining the target query q.

5.2 The SQUIRREL Algorithm

The outer loop of the SQUIRREL algorithm in Figure 13 halts when the teacher considers the current query to be equivalent to the target query. Otherwise a new tree is investigated. The inner loop halts when the teacher confirms that the current tree is correctly annotated by the current hypothesis query. Otherwise, it asks for further partial annotations for the current tree and learns from them.

A *partially annotated tree* over Σ is a triple $\langle t, p_+, p_- \rangle$ consisting of a Σ-tree t, a set $p_+ \subseteq nodes(t)$ of *positively annotated* nodes, a set $p_- \subseteq nodes(t)$

SQUIRREL

// fix a representation formalism for queries //
// we write q_A for the query represented by A //
let A represent the empty query
let S be the empty set of completely annotated trees
Loop until QEQ (A) returns yes
    let t be the output of QEQ
    let $p_+ = p_- = \emptyset$ be partial annotations for t
    Loop until CLQ $(t \times q_A(t))$ returns yes
        let v be output of CLQ
        // add the correct annotation of v to p_+ or p_- //
        If $v \in q_A(t)$ **then** add v to p_- **else** add v to p_+
        // learn a new query representation//
        $A \leftarrow \text{RPNI}_{\text{prune}}(S, \langle t, p_+, p_- \rangle)$
    Add $t \times q_A(t)$ to the set S
Output : A

Figure 13 SQUIRREL active learning algorithm.

of *negatively annotated* nodes, and $p_+ \cap p_- = \emptyset$. It is *consistent* with a query q if $p_+ \subseteq q(t)$ and $p_- \cap q(t) = \emptyset$. SQUIRREL is parametrized by an arbitrary algorithm RPNI$_{\text{prune}}$, which we will fix later on to be a variant of the tRPNI algorithm for NSTTS. However, we could use every other algorithm that computes query representations from a set S of (completely) annotated trees (the set of correctly annotated trees seen so far) and a partially annotated tree $\langle t, p_+, p_- \rangle$ (the tree for which the teacher is answering CLQ queries).

5.3 Pruning

In order to be able to benefit from partial annotations, we propose to prune subtrees that seem irrelevant for distinguishing positively annotated nodes. All remaining nodes are assumed to be unselected if not annotated positively (which seems reasonable for appropriate pruning heuristics). Thereby, we obtain complete annotations on pruned trees. Note that negatively annotated nodes may be pruned. This complicates the situation a little, as wrapper induction needs to ensure consistency with all negative annotations, pruned or not.

At the same time, pruning helps resolving the second drawback because the size of the target automaton is reduced. The idea is similar to that of windowing in wrapper induction from texts (see e.g.(Freitag & Kushmerick, 2000; Chidlovskii, 2001)).

*5.4 t*Rpni *with Pruning*

The learning algorithm Rpni$_{\text{prune}}$ in Figure 14 is parametrized by a *pruning function prune*, which is an arbitrary function mapping $\Sigma \times \mathbb{B}$-trees to $(\Sigma \times \mathbb{B})_{\text{T}}$-trees such that for every possible argument t:

$$prune(t) \in cuts(t)$$

First, we prune all input trees in S and the partially annotated input tree $\langle t, p_+, p_- \rangle$. This is done by computing a completely annotated tree $t \times p_+$ in which all nodes of p_+ are labeled by 1 and all others by 0, to which *prune* is then applied. Let S' be the resulting set of completely annotated pruned trees.

Second, Rpni$_{\text{prune}}$ infers a pNstt by a variant of the tRpni algorithm which tests for cut-functionality. It starts with an initial pNstt for S', denoted by pNstt(S'), as usual with Rpni. This is the deterministic pNstt whose language consists of all pruned trees in S', such that for all states $q \in states(p$Nstt$(S'))$ there exists exactly one tree t over $(\Sigma \times \mathbb{B})_{\text{T}}$ with a successful run by A. Then, we use a state merging for generalization. The function det-merge takes as input an automaton A and two states q_i and q_j of $states(A)$ and outputs an automaton where q_i and q_j are merged such that det-merge(A, q_i, q_j) is made deterministic.

The iteration continues with the merged automaton if it is cut-functional and consistent with all negative annotations in S and $\langle t, p_+, p_- \rangle$, or else refuses the merging step and continues with the previous automaton. *Consistency* here means that there does not exist any tree $t' \times \beta \in S$ and node $v \in q_A(t')$ such that $\beta(v) = 0$, and there does not exist any node $v \in p_-$ such that $v \in q_A(t)$.

Improvements.   The most classical one is to tweak the merging order, as illustrated by evidence driven state merging (Lang, Pearlmutter, & Price, 1998). Here, we choose a merging order which favors merges leading to pNstts with "horizontal" recursion, since this type of recursion is more frequent in semi-structured documents. The second improvement is to use background knowledge in form of typing automata (Kermorvant & de la Higuera, 2002). For web documents, we use a very general typing automaton that recognizes all well-formed web documents. Merges between states of different types are forbidden. This restricts the class of pNstts that can be inferred.

To conclude, the reader should note that Rpni$_{\text{prune}}$ responds in polynomial time (as tRpni and Rpni). As the cut-functionality of a tree automaton A can be tested in $O(|A|^3)$ (see Proposition 2), which is the most complex task of the inner loop, and as this test has to be performed a quadratic amount of time, the overall complexity of tRpni is $O(|S|^5)$ (at worst, the working automaton is always the initial automaton, which has a size at worst equal to the size of the input sample). Note that we consider here the size of an automaton as the sum of its number of node and its number

---

RPNI$_{\text{prune}}$ $(S, \langle t, p_+, p_- \rangle)$

// prune all input trees //
$S' = \{ prune(t' \times \beta) \mid t' \times \beta \in S, \ \beta \neq \emptyset \} \cup prune(t \times p_+)$
// compute a deterministic pNsTT by RPNI//
$A \leftarrow p\text{NsTT}(S')$
For $i = 1$ **to** $|states(A)|$ **do**
    For $j = 0$ **to** $i - 1$ **do**
        $A' \leftarrow \text{det-merge}(A, q_i, q_j)$
        If A' is cut-functional
        and A' consistent with S and $\langle t, p_+, p_- \rangle$
        and $L(A') \cap \{s \mid s$ is tree over $(\Sigma \times \{0\})_\text{T}\} = \emptyset$
        then $A \leftarrow A'$
            exist inner loop
        else skip
Output: A

**Figure 14** Learning from partially annotated trees.

Figure 15 The pruned trees applying the strategy *paths_only* (left) and *paths_extended* (right) to the tree *films*

of rules, and the size of a sample as the total number of nodes of trees in the sample. For large trees and a strongly pruning function there would be a speed-up of the learning process. However there is a trade-off in that we loose completeness in learnability. We discuss this last point in the next subsection.

5.5 Pruning Heuristics

The choice of the pruning function determines the quality of RPNI$_{\text{prune}}$. To be able to learn a target query, some conditions should be imposed on the pruning function. Namely that it should be possible to find a characteristic pruned set for the target language, given the pruning function. And that it should be possible to find a cut-functional pruned language satisfying the target, for a given pruning function. Therefore the class of learnable functions depends on the pruning function. The time complexity of the learning algorithm is measured w.r.t. the size of pruned trees. Thus, the more aggressive the pruning function is, the more efficient the learning algorithm

will be. Generally, we will only consider pruning functions that keep all
positively annotated nodes. We have considered three pruning strategies:

- The function *paths_only* cuts all nodes except those that are ancestors
  of or equal to positively labeled nodes. An example is given in Figure 15.
  Note that more than a single path is kept so that induction is not sim-
  ply reduced to learning path expressions. The problem of defining the
  class of Nstts that can actually be identified by $\text{RPNI}_{\text{prune}}$ with the
  pruning function *paths_only* is still open. The experiments show that
  pruning by *paths_only* yields significantly better learning efficiency than
  less aggressive pruning functions. Also experimental results show that
  expressiveness is sufficient for monadic queries based on the structure.
- The function *paths_extended* additionally keeps all the nodes whose
  parents are on such paths. An example is given in Figure 15. The prob-
  lem of defining the class of Nstt that can actually be identified by
  $\text{RPNI}_{\text{prune}}$ with the pruning function *paths_extended* is still open. The
  class is larger than for the pruning function *paths_only* because some
  node selection queries using textual contents are learnable. For instance,
  a query for actor lists where the next sibling contains the textual value
  "Hitchcock". As soon as textual values are not involved, experimental re-
  sults have shown lower learning efficiency and equivalent expressiveness
  than for the pruning function *paths_only*.
- The function *identity*. $\text{RPNI}_{\text{prune}}$ restricted to completely annotated trees
  behaves as the tRpni algorithm for deterministic Nstts. All Mso-definable
  node selection queries could be the target of $\text{RPNI}_{\text{prune}}$. However exper-
  imental results show that a large number of annotated examples is re-
  quired to achieve learning. Also, in the interactive setting, learning from
  a small number of annotated positions could introduce wrong general-
  izations.

6 Squirrel Web Information Extraction System

We have developed a software tool, Squirrel (Figure 16). We first describe
the system. Then we give experimental results in the off-line setting in order
to compare our system with others. Finally we give experimental results in
order to evaluate the number of interactions in the interactive setting.

System description. Squirrel preprocesses pages via the parser included
in Mozilla. This yields a Dom Xhtml tree. We ignored textual values. At-
tributes were ignored too except for `ID` and `CLASS`, which were merged with
the name of the tag, i.e. `<DIV ID=title>` gives a node labeled `DIVtitle`.

The user is the query designer. When a Clq is made by Squirrel, the
query designer annotates an incorrect node (unwanted or missing). In a real
setting, the query designer may annotate more than one position. When a
Qeq is made by Squirrel, the query designer decides to halt the induction
process or chooses a new Web page.

Figure 16 Screenshot of Mozilla Firefox with SQUIRREL: a toolbar has been added to help the user to annotate data in the web page and to dialog with the learning algorithm.

The SQUIRREL algorithm calls the $\text{RPNI}_{\text{prune}}$ algorithm. For the experimental results we have chosen the pruning function *paths_only*.

Off-line Experiments. The aim of these experiments is to compare $\text{RPNI}_{\text{prune}}$ with state-of-the-art algorithms. Thus we do not consider the interactive framework and $\text{RPNI}_{\text{prune}}$ is given as an input completely annotated examples for the target query.

We considered benchmarks from RISE[2], the only public repository of annotated documents designed to evaluate information extraction algorithms on Web pages. We had to discard some of the datasets, because the extraction tasks they defined were not HTML element extraction tasks, and could

---

[2] http://www.isi.edu/~muslea/RISE/index.html

therefore not be addressed by our algorithm. Typically, extracting a given word in a sentence is a problem outside the scope of our work.

The remaining datasets we used are these ones:

- 18 of Kushmerick's WIEN benchmark (Freitag & Kushmerick, 2000) used by Muslea in (Muslea et al., 2001) to establish a comparison between Wien and Stalker. They are denoted by S-*XX* in table of results. Each of them is constituted with 8 pages, and is associated with one to four extraction tasks.
- Bigbook, a set 235 pages extracted from an address book. Two tasks are associated with this set: name extraction and address extraction.
- Okra, a set of 250 pages which corresponds to a people search engine output. Three tasks are associated with this set: name, mail address and score extraction.

For each dataset, we evaluate our algorithm using a 8-fold cross-validation protocol. The dataset is divided in 8 parts. A wrapper dedicated to one extraction task is learned from 7 parts, and tested on the last one. True positives, false negatives and false positives are counted and a F1-measure score is computed. Each part of the dataset is used as test set for each extraction task in a distinct experience, and an average F-measure is computed.

It is rather difficult to compare our own results to the ones of other systems, considering that we cannot test them ourselves and that we have to use published results whose experimental protocol are never exactly the same. Following Muslea in (Muslea et al., 2001), we establish a comparison using a qualitative criteria: we consider three possible behaviours for each system on each benchmark: perfect (average F-measure of 100%), imperfect (average F-measure between 90% and 100%) and failure (F-measure below 90%). All results are presented in Fig. 17. Results concerning WIEN (Kushmerick, 2000) and Stalker (Muslea et al., 2001), two well known information extraction systems, are given for the sake of comparison. They have been directly taken from (Muslea et al., 2001).

These results show that SQUIRREL behaves satisfyingly in a large majority of the benchmarks, and achieves performances on these benchmarks which are equals or greater than WIEN and STALKER. A closer analysis of the failures of Squirrel (S-6 and S-11) shows that they have expected causes: in both cases, the structure is not informative enough because textual contents of leaves are necessary to define the query. Therefore SQUIRREL with the pruning function *paths_only* can not infer a relevant wrapper. One way to overcome this problem is to preprocess the data sets in order to distinguish text nodes which should be useful for the induction process, consider trees with these textual values, and use the pruning function *paths_extended*.

To end the comparisons with existing systems, we now compare our system with the system presented in (Raeymaekers et al., 2005). They represent a monadic node selection query q by a tree language L_q where each tree contains a single marked node. To decide if a node of a tree t

	Stalker	Wien	**Squirrel**
S-4	√	√	√
S-5	√	√	√
S-6	≃	−	−
S-7	√	−	√
S-8	√	√	√
S-9	≃	−	√
S-10	√	√	√
S-11	≃	−	−
S-12	√	√	√
S-13	√	√	√
S-14	√	√	√
S-19	√	√	√
S-20	√	√	√
S-22	√	√	√
S-23	√	√	√
S-24	≃	−	√
S-25	√	√	√
S-28	√	√	√
S-30	√	√	√
Okra	√	√	√
BigBook	√	√	√

Figure 17 Comparison between Stalker, Wien and Squirrel on the Rise corpus. The symbol − means *failure*, that is average F-measure lower than 90%, ≃ means *imperfect*, that is average F-measure between 90% and 100%, and √ means *perfect*, that is average F-measure equal to 100%.

is to be selected, the node is marked (replaced by a special symbol), if the tree is in L_q, then the original node is to be selected. Whereas we use the Rpni algorithm from positive and negative examples, their base learning algorithm is from positive examples only. They consider a subclass of the class of regular languages which is chosen to be the class of (k, l)-contextual tree languages where k and l are parameters. With parameter tuning, they achieve perfect results on benchmarks presented Fig. 17. They succeed in S-6 and S-11 because they use a preprocessing procedure to distinguish important textual values. They also present a method to learn the parameters from positive and negative examples.

Benchmarking interactive learning.   For the experiments, we used the benchmarks Okra, Bigbook and real datasets[3] from different Web sites:

– the `Yahoo` dataset contains 80 pages of the Yahoo Web directory. The task is to extract a set of URLs from a list of URLs. There are different lists on a page, the pages have different structures and the position of the list is not fixed;

---

[3] http://www.grappa.univ-lille3.fr/~carme/corpus.html

	#Qᴇǫ	#Cʟǫ
Okra–names	1.6	3.48
Bigbook-addresses	1	3.02
Yahoo	6.18	11.36
E-bay	1.06	2.62
NYTimes	1.44	1.44
Google	1.86	4.78

Figure 18 #Qᴇǫ denotes the number of equivalence queries and #Cʟǫ denotes the number of correct labeling queries. Results are averaged over 100 experiments.

- the `eBay` dataset contains 40 pages of products. The task is to extract product names from a table. The product name can be mixed with other information in the cell;
- the `Google` dataset contains 40 results of the search engine. The task is to extract the links. The structure of pages is not fixed and there are numerous auxiliary links;
- the `New York Times` dataset contains 40 article pages with different structures. The task is to extract texts of the articles.

In order to evaluate the efficiency of the algorithm in the interactive mode, we simulated the behavior of a query designer. Squirrel is interfaced with a user simulator S: at each call of the correct labeling oracle, S answers the first wrong annotation of the page in reading order of the source file (which is the first positive annotation of the page for the first call) ; at each call of the equivalence oracle, S returns a Web page on which the pruned Nstt is not correct.

In an interactive setting, time complexity is important. For the different data sets, on a standard computer, the average time for learning the current hypothesis and annotating the current Web page is between 1 and 2 seconds and less than 5 seconds on the more intricate cases.

We also measured the number of Cʟǫs and Qᴇǫs which are needed to achieve perfect results on the set of Web pages. The results are given in Figure 18. The number of Cʟǫs is the number of corrections made by the query designer. The results show that this number is quite small. We should note that this number is over-valued because in a real setting, the query designer may correct more than one wrong annotation at a time, and we can hope that his choice of correction is more relevant than the one used in our protocol. The number of Qᴇǫs is the number of Web pages necessary to achieve perfect results. The number of necessary Web pages for the `Yahoo` dataset is greater than for other datasets because the dataset is heterogeneous. We should note that the number of necessary Web pages is lower than the number of Web pages viewed by the query designer. Indeed, in our protocol the user simulator selects only Web pages on which the current hypothesis is incorrect while, in the real setting, the query designer may choose a Web page on which the current hypothesis is correct. Therefore

	1	2	5	10
	random			
Yahoo	71.8	80.9	82.1	93.7
NYTimes	91.2	92.4	95.3	100
Google	98.7	99.1	99.5	99.8
	active			
Yahoo	71.8	86.1	98.4	99.4
NYTimes	91.2	100	100	100
Google	98.7	100	100	100

Figure 19 F-measure of inferred query on test set, using 1,2 5 or 10 documents chosen randomly (top) or chosen accordingly to our strategy (bottom)

we now discuss the choice of the next Web page to be processed in the interactive setting.

Towards active learning We should render our model active in the sense that the learner intelligently proposes the next Web page among a set of Web pages – the Web pages of the target Web site.

We considered a method estimating the value of a tree according to the number of extracted nodes by the current hypothesis wrapper. We supposed that there is a linear dependence between the number of extracted nodes and the total number of nodes. This lead to the following strategy: let L be the set of Web pages not seen so far by SQUIRREL; for every Web page p in L, we compute the number x of nodes of p and y the number of selected nodes by the current hypothesis wrapper; we used the test of Grubbs for outlier selection to select the web page (the farthest point (x, y) from the regression line). The selected Web page is processed by the main loop of SQUIRREL.

We only considered the three most difficult datasets. We computed performance with a learning set of annotated documents of respectively 1, 2, 5, 10 Web pages. These Web pages were either randomly chosen or chosen according to the strategy given above. The experimental results are given in Figure 19. They show that our naive strategy for active learning leads to significant improvements. It remains to find an intelligent active learning strategy as a future work.

7 Conclusion

We have presented a wrapper induction system with the following characteristics:

- it generates tree-based wrappers which select nodes in trees;
- it uses new machine learning techniques based on grammatical inference and active learning;
- it is integrated in a visual environment;

– it has simple interactions with the query designer.

From the theoretical point of view, it is proved that NSTTs and pNSTTs can express all MSO-definable node selection queries. However any pruning function, different from identity, will restrict the class of learnable node selection queries. In future work, We have to specify the subclass of node selection queries which is learnable depending on the chosen pruning strategy.

From the practical point of view, we have proved that our system achieves state-of-the-art performance on standard benchmarks and that the number of interactions with the query designer is quite small. Nevertheless there are still a number of possible improvements. First, we should render our algorithm active with intelligent strategies. Second, we do not consider textual values and omit some attribute values. Therefore, we should adapt our approach to take textual information into account. Finally, our current system is limited to one-slot information extraction since we only consider monadic node selection queries over trees. Work is under progress to extend our approach to multi-slot information extraction, i.e., the problem of learning n-ary regular queries in trees.

References

Angluin, D. (1987). Learning regular sets from queries and counterexamples. *Information and Computation*, *75*(2), 87–106.

Baumgartner, R., Flesca, S., & Gottlob, G. (2001). Visual web information extraction with lixto. In *28th International Conference on Very Large Data Bases*, pp. 119–128.

Brüggemann-Klein, A., Wood, D., & Murata, M. (2001). Regular tree and regular hedge languages over unranked alphabets: Version 1..

Carme, J., Gilleron, R., Lemay, A., & Niehren, J. (2005). Interactive learning of node selecting tree transducer. In *IJCAI Workshop on Grammatical Inference*.

Carme, J., Lemay, A., & Niehren, J. (2004a). Learning node selecting tree transducer from completely annotated examples. In *7th International*

Colloquium on Grammatical Inference, Vol. 3264 of *Lecture Notes in Artificial Intelligence*, pp. 91–102. Springer Verlag.

Carme, J., Niehren, J., & Tommasi, M. (2004b). Querying unranked trees with stepwise tree automata. In *19th International Conference on Rewriting Techniques and Applications*, Vol. 3091 of *Lecture Notes in Computer Science*, pp. 105 – 118. Springer Verlag.

Chidlovskii, B. (2001). Wrapping web information providers by transducer induction. In *Proc. European Conference on Machine Learning*, Vol. 2167 of *Lecture Notes in Artificial Intelligence*, pp. 61 – 73.

Cohen, W., Hurst, M., & Jensen, L. (2003). *Web Document Analysis: Challenges and Opportunities*, chap. A Flexible Learning System for Wrapping Tables and Lists in HTML Documents. World Scientific.

Comon, H., Dauchet, M., Gilleron, R., Jacquemard, F., Lugiez, D., Tison, S., & Tommasi, M. (1997). Tree automata techniques and applications. Available on: http://www.grappa.univ-lille3.fr/tata.

Cristau, J., Löding, C., & Thomas, W. (2005). Deterministic automata on unranked trees. In *15th International Symposium on Fundamentals of Computation Theory*, Vol. 3623 of *Lecture Notes in Computer Science*, pp. 68–79. Springer Verlag.

de la Higuera, C. (1997). Characteristic sets for polynomial grammatical inference. *Machine Learning*, *27*, 125–137.

Drewes, F., & Hogberg, J. (2003). Learning a regular tree language from a teacher. In *D.L.T. 2003*, Vol. 2710 of *Lecture Notes in Computer Science*, pp. 279–291.

Freitag, D., & Kushmerick, N. (2000). Boosted wrapper induction. In *AAAI/IAAI*, pp. 577–583.

Freitag, D., & McCallum, A. K. (1999). Information extraction with hmms and shrinkage. In *Proceedings of the AAAI-99 Workshop on Machine Learning for Information Extraction*.

Frick, M., Grohe, M., & Koch, C. (2003). Query evaluation on compressed trees. In *18th IEEE Symposium on Logic in Computer Science*, pp. 188–197.

Gold, E. (1967). Language identification in the limit. *Inform. Control*, *10*, 447–474.

Gold, E. (1978). Complexity of automaton identification from given data. *Inform. Control*, *37*, 302–320.

Gottlob, G., & Koch, C. (2002). Monadic queries over tree-structured data. In *17th Annual IEEE Symposium on Logic in Computer Science*, pp. 189–202 Copenhagen.

Hsu, C.-N., & Dung, M.-T. (1998). Generating finite-state transducers for semi-structured data extraction from the web. *Information Systems*, *23*(8), 521 – 538.

Kermorvant, C., & de la Higuera, C. (2002). Learning language with help. In *6th International Colloquium on Grammatical Inference*, Vol. 2484 of *Lecture Notes in Artificial Intelligence*, pp. 161–173. Springer Verlag.

Kosala, R., Bruynooghe, M., Van den Bussche, J., & Blockeel, H. (2003). Information extraction from web documents based on local unranked tree automaton inference. In *18th International Joint Conference on Artificial Intelligence*, pp. 403–408. Morgan Kaufmann.

Kushmerick, N. (1997). *Wrapper Induction for Information Extraction*. Ph.D. thesis, University of Washington.

Kushmerick, N. (2000). Wrapper induction: Efficiency and expressiveness. *Artificial Intelligence*, *118*(1-2), 15–68.

Kushmerick, N. (2002). Finite-state approaches to web information extraction. In *Proc. 3rd Summer Convention on Information Extraction*.

Lang, K. J., Pearlmutter, B. A., & Price, R. A. (1998). Results of the abbadingo one DFA learning competition and a new evidence-driven state merging algorithm. *Lecture Notes in Computer Science*, *1433*, 1–12.

Lang, K. (1992). Random DFA's can be approximately learned from sparse uniform examples. In *Proc. 5th Annu. Workshop on Comput. Learning Theory*, pp. 45–52. ACM Press, New York, NY.

Libkin, L. (2005). Logics over unranked trees: an overview. In *Automata, Languages and Programming: 32nd International Colloquium*, No. 3580 in Lecture Notes in Computer Science, pp. 35–50. Springer Verlag.

Martens, W., & Niehren, J. (2007). On the minimization of XML schemas and tree automata for unranked trees. *Journal of Computer and System Science*, *73*(4), 550–583.

Muslea, I., Minton, S., & Knoblock, C. A. (2001). Hierarchical wrapper induction for semistructured information sources. *Autonomous Agents and Multi-Agent Systems*, *4*(1/2), 93–114.

Neumann, A., & Seidl, H. (1998). Locating matches of tree patterns in forests. In *Foundations of Software Technology and Theoretical Computer Science*, pp. 134–145.

Neven, F., & Schwentick, T. (2002). Query automata over finite trees. *Theoretical Computer Science*, *275*(1-2), 633–674.

Neven, F., & Van Den Bussche, J. (2002). Expressiveness of structured document query languages based on attribute grammars. *Journal of the ACM*, *49*(1), 56–100.

Niehren, J., Planque, L., Talbot, J.-M., & Tison, S. (2005). N-ary queries by tree automata. In *10th International Symposium on Database Programming Languages*, Vol. 3774 of *Lecture Notes in Computer Science*, pp. 217–231. Springer Verlag.

Oncina, J., & Garcia, P. (1992). Inferring regular languages in polynomial update time. In *Pattern Recognition and Image Analysis*, pp. 49–61.

Oncina, J., & García, P. (1993). Inference of recognizable tree sets. Tech. rep., Departamento de Sistemas Informáticos y Computación, Universidad de Alicante. DSIC-II/47/93.

Raeymaekers, S., & Bruynooghe, M. (2004). Minimization of finite unranked tree automata. Manuscript.

Raeymaekers, S., Bruynooghe, M., & Van den Bussche, J. (2005). Learning (k,l)-contextual tree languages for information extraction. In *Proceedings of ECML'2005*, Vol. 3720 of *Lecture Notes in Artificial Intelligence*, pp. 305–316.

Sakakibara, Y. (1990). Learning context-free grammars from structural data in polynomial time. *Theoretical Computer Science*, *76*, 223 – 242.

Seidl, H. (1989). On the finite degree of ambiguity of finite tree automata. *Acta Informatica*, *26*(6), 527–542.

Thatcher, J. W. (1967). Characterizing derivation trees of context-free grammars through a generalization of automata theory. *Journal of Computer and System Science*, *1*, 317–322.

Thatcher, J. W., & Wright, J. B. (1968). Generalized finite automata with an application to a decision problem of second-order logic. *Mathematical System Theory*, *2*, 57–82.

A Testing Cut-Functionality in Polynomial Time

We present a direct proof for Proposition 2 in the case of binary trees (and thus for unranked trees), stating that cut-functionality of deterministic tree automata can be decided in polynomial time.

We characterize cut-functionality of a tree automaton A with signature $(\Sigma \times \mathbb{B})_T$ by a binary relation $inc_A \subseteq states(A)^2$, testing for two states whether they can be reached by two runs of A on some incompatibly annotated compatible trees. For all two states $p, p' \in states(A)$:

$$inc_A(p, p') \text{ iff } \begin{cases} \exists t, t' \text{ compatible trees over } \Sigma_T \\ \exists \beta, \beta' \text{ incompatible trees over } \mathbb{B}_T \\ \exists r \in runs_A(t \times \beta) : \ r(\varepsilon) = p \\ \exists r' \in runs_A(t' \times \beta') : \ r'(\varepsilon) = p' \end{cases}$$

Lemma 4 *A tree automaton A over $(\Sigma \times \mathbb{B})_T$ is cut-functional if and only if $inc_A(p, p')$ doesn't hold for any two states $p, p' \in final(A)$.*

Proof Suppose $inc_A(p, p')$ holds for some $p, p' \in final(A)$. Then there exist runs $r \in runs_A(t \times \beta)$ and $r' \in runs_A(t' \times \beta')$ for incompatible trees β and β' and compatible trees t and t'. These runs are successful since $p, p' \in final(A)$. Thus, A is not cut-functional. The argument can be reversed.

It remains to show that the relation inc_A can be computed in polynomial time. We prove that it can be computed by a saturation procedure over states. We first need to define another binary relation sim_A on states. sim_A holds if two states can be reached by compatible annotated trees:

Definition 8 *Let A be a deterministic tree automaton over $(\Sigma \times \mathbb{B})_T$, p and p' two states of A, then $sim_A(p, p')$ holds if there exist two compatible $(\Sigma \times \mathbb{B})_T$-trees $t \times \beta, t' \times \beta'$, a run r of A over $t \times \beta$ such that $r(\varepsilon) = p$, and a run r' of A over $t' \times \beta'$ such that $r'(\varepsilon) = p'$.*

We now give an inductive characterization of the relation inc_A which is based on the relation sim_A and on the rule set of the tree automaton. We call an automaton *trimmed* if it has no useless state, i.e., if for each state p there is a run r of A on some tree such that $r(\varepsilon) = p$.

Lemma 5 *Let A be a trimmed deterministic tree automaton over $(\Sigma \times \mathbb{B})_T$, $inc_A(p, p')$ holds if and only if there exist rules satisfying one of the cases in Figure 20*

Proof Let us suppose that $inc_A(p, p')$. Then there exist two $(\Sigma \times \mathbb{B})_T$-trees $t \times \beta, t' \times \beta'$ such that t and t' are compatible but $t \times \beta$ and $t' \times \beta'$ are not compatible, a run r of A over $t \times \beta$ such that $r(\varepsilon) = p$, and a run r' of A over $t' \times \beta'$ such that $r'(\varepsilon) = p'$.

$$(i) \quad (a, b) \rightarrow p \in rules(A)$$
$$\quad \wedge (a, \neg b) \rightarrow p' \in rules(A)$$
$$(ii) \quad (a, b)(p_1, p_2) \rightarrow p \in rules(A)$$
$$\quad \wedge (a, \neg b)(p_1', p_2') \rightarrow p' \in rules(A)$$
$$\quad \wedge sim_A(p_1, p_1') \wedge sim_A(p_2, p_2')$$
$$(iii) \quad (a, b)(p_1, p_2) \rightarrow p \in rules(A)$$
$$\quad \wedge (a, b')(p_1', p_2') \rightarrow p' \in rules(A)$$
$$\quad \wedge inc_A(p_1, p_1') \wedge sim_A(p_2, p_2')$$

$$(iv) \quad (a, b)(p_1, p_2) \rightarrow p \in rules(A)$$
$$\quad \wedge (a, b')(p_1', p_2') \rightarrow p' \in rules(A)$$
$$\quad \wedge sim_A(p_1, p_1') \wedge inc_A(p_2, p_2')$$
$$(v) \quad (a, b)(p_1, p_2) \rightarrow p \in rules(A)$$
$$\quad \wedge (a, b')(p_1', p_2') \rightarrow p' \in rules(A)$$
$$\quad \wedge inc_A(p_1, p_1') \wedge inc_A(p_2, p_2')$$

Figure 20   Cases of Lemma 5.

- First, note that neither t nor t' is equal to T (which would correspond to a tree entirely pruned), as otherwise $\beta = T$ (resp. $\beta' = T$), and as T is compatible with every tree, we would have β and β' compatible.
- If t is a single leaf tree with $t = a$ then $t' = a$ also as t and t' are compatible. In this case $\beta = b$ and $\beta' = \neg b$ as they are not compatible. As $r(\varepsilon) = p$ and $r'(\varepsilon) = p'$, this implies that there are two rules $(a, b) \rightarrow p$ and $(a, \neg b) \rightarrow p'$ and (i) holds.
- If t is a tree of the form $t = a(t_1, t_2)$ then, as t and t' are compatible, there are trees t_1' and t_2' such that $t' = a(t_1', t_2')$ such that t_1 is compatible with t_1' and t_2 is compatible with t_2'. Let $t \times \beta = (a, b)(t_1 \times \beta_1, t_2 \times \beta_2)$ and $t' \times \beta' = (a, b')(t_1' \times \beta_1', t_2' \times \beta_2')$. Let us consider the rules $(a, b)(p_1, p_2) \rightarrow p$ and $(a, b')(p_1', p_2') \rightarrow p'$ which are applied at the root of the trees $t \times \beta$ and $t' \times \beta'$ in the runs r and r'. As t_1 and t_1' are compatible, and t_2 and t_2' are compatible, the states satisfy $sim_A(p_1, p_1')$ or $inc_A(p_1, p_1')$ and $sim_A(p_2, p_2')$ or $inc_A(p_2, p_2')$. First, consider the case where the Boolean values at the root of the trees $t \times \beta$ and $t' \times \beta'$ are different. In this case, one of the four conditions (ii), (iii), (iv) or (v) holds. If the Boolean values at the root of the trees $t \times \beta$ and $t' \times \beta'$ are equal, then $inc_A(p_1, p_1')$ or $inc_A(p_2, p_2')$ must hold and one of the conditions (iii), (iv) or (v) holds.

We now prove the other direction of the equivalence.

- (i) If $((a, b) \rightarrow p \in rules(A) \wedge (a, \neg b) \rightarrow p' \in rules(A))$ then $inc_A(p, p')$.
- (ii) Let us consider two rules $(a, b)(p_1, p_2) \rightarrow p$ and $(a, \neg b)(p_1', p_2') \rightarrow p'$ in $rules(A)$ such that $sim_A(p_1, p_1')$ and $sim_A(p_2, p_2')$). Then $sim_A(p_1, p_1')$ implies that there are two compatible trees $t_1 \times \beta_1$ and $t_1' \times \beta_1'$, and runs $r_1 \in runs_A(t_1 \times \beta_1)$ and $r_1' \in runs_A(t_1' \times \beta_1')$ such that $p_1 = r_1(\varepsilon)$ and $p_1' = r_1'(\varepsilon)$. We have the same property for the right sub-term (the corresponding states, runs and terms are denoted by $p_2, p_2', r_2, r_2', t_2 \times \beta_2$ and $t_2' \times \beta_2'$). Therefore we can consider trees $t \times \beta = (a, b)(t_1 \times \beta_1, t_2 \times \beta_2)$ and $t' \times \beta' = (a, \neg b)(t_1' \times \beta_1', t_2' \times \beta_2')$ associated to p and p' through runs $r = p(r_1, r_2)$ and $r' = p'(r_1', r_2')$. $t \times \beta$ and $t' \times \beta'$ are not compatible whereas t and t' are, therefore $inc_A(p, p')$ holds.
- (iii) We consider two rules satisfying (iii). We consider the two trees that are associated with the property $inc_A(p_1, p_1')$, and then we construct two

trees proving that $inc_A(p, p')$ holds. The construction is similar to the one used in the precedent point. In this case, incompatibility between Boolean values comes from the left subtrees $t_1 \times \beta_1$ and $t_1' \times \beta_1'$ because of $inc_A(p_1, p_1')$.

– the cases (iv) and (v) are similar.

Lemma 6 *Let A be a trimmed deterministic tree automaton over $(\Sigma \times \mathbb{B})_T$, $sim_A(p, p')$ stands if and only if there exist rules satisfying one of the following cases:*

$$
\begin{aligned}
&(i)\quad &&(p = p')\\
&(ii)\quad &&(T \to p) \in rules(A)\\
&(iii)\quad &&(T \to p') \in rules(A)\\
&(iv)\quad &&\exists((a, b)(p_1, p_2) \to p \in rules(A)\\
&&&\land\, (a, b)(p_1', p_2') \to p' \in rules(A)\\
&&&\land\, (sim_A(p_1, p_1')) \land (sim_A(p_2, p_2'))
\end{aligned}
$$

Proof If $sim_A(p, p')$, then there exist two compatible $(\Sigma \times \mathbb{B})_T$-trees $t \times \beta$ and $t' \times \beta'$, a run r of A over $t \times \beta$ such that $r(\varepsilon) = p$, and a run r' of A over $t' \times \beta'$ such that $r'(\varepsilon) = p'$.

– If $t = T$, then there is a rule $T \to p$ in A and (ii) holds.
– If $t \times \beta = (a, b)$, then either $t' \times \beta' = (a, b)$ or $t' \times \beta' = (T, T)$ because $t \times \beta$ and $t' \times \beta'$ are compatible. In the first case, because of determinism, $p = p'$ and (i) holds. In the second case, there is a rule $T \to p'$ in A and (iii) holds.
– If there are some trees $t_1 \times \beta_1$ and $t_2 \times \beta_2$ such that $t \times \beta = (a, b)(t_1 \times \beta_1, t_2 \times \beta_2)$, then either t' is equal to T (and (iii) applies) or the root of t' is also (a, b) as $t \times \beta$ and $t' \times \beta'$ are compatible. Let $t' \times \beta' = (a, b)(t_1' \times \beta_1', t_2' \times \beta_2')$. Let us consider the rules $(a, b)(p_1, p_2) \to p$ and $(a, b)(p_1', p_2') \to p'$ which are applied at the root of the trees $t \times \beta$ and $t' \times \beta'$ in the runs r and r'. We should note that $t_1 \times \beta_1$ and $t_1' \times \beta_1'$ are compatible, and $t_2 \times \beta_2$ and $t_2' \times \beta_2'$ are compatible, thus the states satisfy $sim_A(p_1, p_1')$ and $sim_A(p_2, p_2')$, and (iv) holds.

We now prove the other direction of the equivalence.

– (i) If $p = p'$, let $t \times \beta$ be any tree that has a run in A whose root is p, $t \times \beta$ is compatible with itself and $sim_A(p, p')$.
– (ii) and (iii) If $T \to p \in rules(A)$, then we can consider any tree $t' \times \beta'$ that has a run in A whose root is p' then $t' \times \beta'$ is compatible with T and $sim_A(p, p')$. Case (iii) is similar.
– (iv) If $(a, b)(p_1, p_2) \to p$ and $(a, b)(p_1', p_2') \to p'$ in $rules(A)$ with $sim_A(p_1, p_1')$ and $sim_A(p_2, p_2')$. The two relations $sim_A(p_1, p_1')$ and $sim_A(p_2, p_2')$ imply that there exist trees $t_1 \times \beta_1$, $t_1' \times \beta_1'$, $t_2 \times \beta_2$ and $t_2' \times \beta_2'$ that are associated to states p_1, p_1', p_2 and p_2' through runs r_1, r_1', r_2 and r_2' on A. Let $t \times \beta = (a, b)(t_1 \times \beta_1, t_2 \times \beta_2)$ and $t' \times \beta' = (a, b)(t_1' \times \beta_1', t_2' \times \beta_2')$, then $t \times \beta$ and $t' \times \beta'$ are compatible, $t \times \beta$ has a run $p(r_1, r_2)$ in A that uses the rule $(a, b)(p_1, p_2) \to p$ and $t' \times \beta'$ has a run $p'(r_1', r_2')$ in A that uses the rule $(a, b)(p_1', p_2') \to p'$. Therefore $sim_A(p, p')$ holds.

Proposition 2 *Cut-functionality of a deterministic tree automaton A over $(\Sigma \times \mathbb{B})_T$ can be tested in cubic time.*

Proof First, the relation sim_A is computed by a saturation process: Initialize with the first three rules in Lemma 6: consider the set of $\{(p, p) \mid p \in states(A)\}$ union $\{(p, p') \mid \exists(T \to p \in rules(A))\}$ union $\{(p, p') \mid \exists(T \to p' \in rules(A))\}$; and then use the fourth rule (iv) inductively. There is at most a quadratic amount of linear tests to perform, therefore the construction is in cubic time. We use a similar construction for the relation inc_A: First use the rule (i) in Lemma 5, second the rule (ii), and then apply inductively the rules (iii), (iv) and (v). Again, the construction is in cubic time.

B Efficient Query Answering for Nstts and pNstts

We next present algorithms for answering queries defined by Nstts or pNstts. Query answering will be needed in our interactive learning environment for testing the current query hypothesis on given trees.

B.1 Alternative definitions of monadic queries

We present two further ways for expressing monadic queries by tree automata. They will be useful for query answering for Nstts and pNstts.

Every tree automaton A with signature $\Sigma \times \mathbb{B}$ defines a MSO-definable monadic query q_A in Σ-trees such that for all such trees t :

$$q_A(t) = \{v \in nodes(t) \mid \beta(v) = 1 \text{ and } t \times \beta \in L(A)\}$$

We can compute $q_A(t)$ for a given binary tree t and automaton A by enumerating all annotated trees $t \times \beta$, filtering those that belong to $L(A)$, and collecting the nodes that such β's annotates by 1. This way of answering queries defined by tree automata is highly non-deterministic, and thus inefficient. But we can do better, as we will see.

Run-based automata queries offer another way to define monadic queries, for which efficient algorithms are known. They are defined by a tree automaton A over Σ and a set of selection states $S \subseteq states(A)$, and satisfy for all binary trees t:

$$q_{A,S}(t) = \{v \in nodes(t) \mid \exists r \in succ\_runs_A(t), \ r(v) \in S\}$$

This query selects all those nodes of t are labeled by some selection state in S in some successful run of B.

Proposition 3 (Classical Result) *Given a tree automaton A over Σ, a set $S \subseteq states(A)$, and a binary Σ-tree t one can compute the set $q_{A,S}(t)$ in time $O(|A| * |t|)$.*

The algorithm proceeds in two phases. In a bottom-up phase, every node v of t is annotated by all states of B into which the subtree rooted by v can be evaluated (determinization on the fly). In a second top-down phase, one underlines all those states that are used by successful runs. Finally, one selects all those nodes that are annotated by underlined states from the selection set S.

B.2 NSTTs and pNSTTs

We first consider monadic queries defined by tree automata A over $\Sigma \times \mathbb{B}$. We show that we can express them by run-based queries via automata projection (which will be useful later on too).

Lemma 7 *Queries by an automata A over $\Sigma \times \mathbb{B}$ can be defined equivalently by runs of the Σ projected automaton $\pi(A)$: $q_A = q_{\pi(A), states(A) \times \{1\}}$.*

As a side product, Lemmas 7 show that every NSTT-defined query can be defined as run-based query by some unambiguous tree automaton.

Proposition 4 *Given a tree automaton A over $\Sigma \times \mathbb{B}$ and a binary Σ-tree t one can compute the set $q_A(t)$ in time $O(|A| * |t|)$.*

By Lemma 7 we can convert the query q_A into a run-based query by automata projection, which requires linear time. This later query can be answered in linear time by Proposition 3.

Theorem 3 *Queries $q_A(t)$ by NSTTs or pNSTTs over binary trees can be computed in linear time $O(|A| * |t|)$.*

For NSTTs this follows immediately from Proposition 4. Queries by pNSTTs can be answered the same way after conversion by the following lemma.

Lemma 8 *Every pNSTTs can be transformed in linear time into an automata over $\Sigma \times \mathbb{B}$ that defines the same query.*

Note that the resulting automaton will not necessarily be an NSTT.

Proof Let A be a pNSTT with signature $(\Sigma \times \mathbb{B})^T$. The idea of the transformation is to compute an automaton $\mathsf{e}(A)$ over $\Sigma \times \mathbb{B}$ that recognizes all expansions of trees in $L(A)$, obtained by replacing T-leaves by arbitrary trees over $\Sigma \times \{0\}$.

$$states(\mathsf{e}(A)) = states(A) \uplus \{T\}$$
$$final(\mathsf{e}(A)) = final(A)$$

The rules of $\mathsf{e}(A)$ are induced by the schemata in 21, which are valid for all symbols $a \in \Sigma$, $p \in states(A)$, and left hand sides l of rules in A: all rules of A are preserved by $\mathsf{e}(A)$ except those with the symbol T, which is

$$\frac{l \rightarrow p \in rules(A) \qquad l \neq T}{l \rightarrow p \in rules(\mathsf{e}(A))} \qquad \frac{a \in \Sigma}{(a,0) \rightarrow T \in rules(\mathsf{e}(A))}$$

$$\frac{T \rightarrow p \in rules(A)}{T \xrightarrow{\varepsilon} p \in rules(\mathsf{e}(A))} \qquad \frac{a \in \Sigma}{(a,0)(T,T) \rightarrow T \in rules(\mathsf{e}(A))}$$

Figure 21   Rules of automata $\mathsf{e}(A)$ that recognizes expansions of trees in $L(A)$.

moved from the signature into the states. Automaton $\mathsf{e}(A)$ may label all leaves by state T, as well as inner nodes whose descendants it has labeled by T before. At an arbitrary time point, $\mathsf{e}(A)$ can decide to move from state T to some state p of A for which $T \rightarrow p$ is a rule of A. Subsequently, $\mathsf{e}(A)$ behaves equally to A.

Note that we freely permit ε-rules, which can easily be eliminated while preserving the language. For these particular automata, ϵ-rules can be eliminated in linear time, since no two ϵ rules can be composed.

It now remains to show that $q_A = q_{\mathsf{e}(A)}$ for every pNSTT A. We omit these details.

C Expressiveness of NSTTs and pNSTTs

We show that deterministic NSTTs as well as deterministic pNSTTs capture the class of monadic MSO-definable queries.

C.1 MSO-definable queries

In a logical perspective, a tree t is seen as a logical structure with domain $nodes(t)$, unary relations $label_a$ for every $a \in \Sigma$, and two binary relations $firstchild$ and $nextsibling$. We assume an infinite set of first-order variables x, y, z, and an infinite set of monadic second-order variables S. Formulas of MSO in unranked (or binary) trees have the following syntax:

$$\phi ::= label_a(x) \mid firstchild(x,y) \mid nextsibling(x,y)$$
$$\mid \quad \phi \rightarrow \phi' \mid S(x) \mid \forall x.\phi \mid \forall S.\phi$$

Standard logical connectives such as conjunction $\phi \wedge \phi'$, negation $\neg\phi$, existential quantification $\exists x.\phi$ and $\exists S.\phi$ can be expressed as usual, as well set inclusion $S \subseteq S'$. Further relations such as $child$ and $lastchild$ are definable in MSO as well.

Models of MSO formulas are defined in the usual Tarskian style. They consist of a tree structure t and a variable assignment α mapping first-order variables to $nodes(t)$ and second-order variables to $2^{nodes(t)}$. We write $t, \alpha \models \phi$ if t, α is a model of ϕ, and consider models of ϕ as equal if they coincide on the free variables of ϕ.

Every MSO formula $\phi(x)$ with a single free variable x defines a monadic query $q_{\phi(x)}$ in unranked (resp. binary) trees, so that for all trees t:

$$q_{\phi(x)}(t) = \{\alpha(x) \in nodes(t) \mid t, \alpha \models \phi(x)\}$$

C.2 Nstt *and p*Nstt*-defined queries*

The idea behind completely annotated examples (as recognized by Nstts) is to annotate all nodes that are selected in some tree in a single annotation. In Mso, we can express the same idea by collecting all nodes x selected by some query $\phi(x)$ in a single set S.

$$\phi'(S) \;\equiv\; \forall x. \; (S(x) \leftrightarrow \phi(x))$$

More generally, we can use Mso formulas $\phi(S)$ with a single free set variable S to define monadic queries, such that for all trees t:

$$q_{\phi(S)}(t) = \{v \in nodes(t) \mid t, \alpha \models \phi(S), \; v \in \alpha(S)\}$$

A monadic query is definable by a Mso formula with a single free first-order variable if and only if it is definable by an Mso formula with a single free second-order variable.

Annotated trees can be understood as the models of Mso formulas with a single set valued variable. A model t, α of Mso formulas with a single free variable S becomes the annotated tree $t \times \alpha(S)$, which annotates all nodes in t belonging to $\alpha(S)$ by 1 and all others by 0.

Theorem 4 (Thatcher and Wright (1968)) *The set of annotated trees corresponding to models of an* Mso *formula $\phi(S)$ over signature Σ is a regular tree language over $\Sigma \times \mathbb{B}$, and vice versa.*

The theorem states that a monadic query in binary trees over Σ definable by an Mso formula with one free set variable if and only it is definable by some tree automaton over the signature $\Sigma \times \mathbb{B}$. The original theorem is even more general in that it applies to n-ary queries which are definable by formulas with n free variables (Niehren, Planque, Talbot, & Tison, 2005).

Theorem 5 *A query is definable in* Mso *iff it is definable by a deterministic* Nstt *iff it is definable by a deterministic p*Nstt*.*

Proof Consider a query that is defined by some Mso-formula $\phi(x)$. The same query is defined by the formula $\phi'(S) \equiv \forall x (x \in S \leftrightarrow \phi(x))$. Formula $\phi'(S)$ is functional in that the value of S is uniquely determined by the tree structure of models of $\phi'(S)$. Let A be the tree automaton over $\Sigma \times \mathbb{B}$ that recognizes the set of annotated trees corresponding to the models of $\phi'(S)$, according to Thatcher and Wright's theorem 4. The language of this automaton is functional, given that the value of S in models $t, \alpha \models \phi'(S)$ is uniquely defined. This proves that A is an Nstt. By construction, it defines the same query as $\phi'(S)$ and thus $\phi(x)$. The same automaton is also a pNstt.

We have shown in Section B.2 that monadic queries represented by deterministic Nstts can be identified with run-based queries defined by unambiguous tree automata. Theorem 5 shows that runs of unambiguous tree

automata can express all Mso definable monadic queries. This result has
been shown before by Neven and Van Den Bussche (2002). The simple
Mso-based proof presented here stems from Niehren et al. (2005) where
the expressiveness of n-ary queries by unambiguous tree automata has been
studied too.