

Towards the Tree Automata Workbench MARBLES^{*}

Frank Drewes

Abstract. We sketch the conceptual ideas that are intended to become the basis for the Tree Automata Workbench MARBLES¹, an extensible system that will facilitate the experimentation with virtually any kind of algorithms on tree automata. Moreover, the system will come with a library and an application programmer’s interface that can be used by anyone wanting to apply such algorithms in research and development.

1 Introduction

Already in the 1960s, researchers in finite-automata theory realized that large parts of this theory can be generalized by replacing strings with trees, without loosing many of the positive algorithmic results, closure properties, and the like. This observation gave rise to a flourishing theory, including a large number of techniques and algorithms for analysis, modification, and synthesis of various kinds of tree recognizers, tree grammars, and tree transducers [GS84, NP92, GS97, FV98, CDG⁺07]. Throughout the rest of this paper, all devices that fall into one of these categories will be called tree automata. Today, probably more theoretical research than ever before is done in this area, motivated by a constantly growing number of applications of tree automata in fields such as verification and model checking [GK00, AJMd02, Löd02, FGVTT04], natural language processing [KG05, GKM08], XML processing [Sch07], code selection in compilers [FSW94], graph and picture generation [Eng94, Dre06], and others.

The system TREEBAG² uses tree generators to generate sets of objects over arbitrary domains. The central data type of TREEBAG is the ranked and ordered tree, with nodes labelled by symbols taken from a ranked alphabet Σ . In other words, every symbol $f \in \Sigma$ comes with a rank $k \geq 0$, such that a node labelled with f is required to have exactly k children (which are totally ordered). This means that a tree in the sense of TREEBAG is a term, i.e., a well-formed expression composed of abstract (i.e., “meaningless”) operation symbols, each having a specified rank that determines the number of subexpressions. TREEBAG deals with two types of tree automata on this type of trees, namely tree grammars and tree transducers. A tree grammar is a device that generates trees out of itself, whereas a tree transducer is one that turns input trees into output trees.


^{*} Dedicated to Hans-Jörg Kreowski on the occasion of this 60th birthday.

¹ Tree Automata Workbench = taw = a large marble, a game of marbles (Oxford New Amer. Dict.).

² Tree-Based Generator

A tree generator is a tree grammar composed with a (possibly empty) sequence of tree transducers.

In a well-known way, trees of the type described above can be assigned a semantics by choosing a domain \mathbb{A} and associating an operation on \mathbb{A} (of the appropriate arity) with each symbol in the ranked alphabet Σ considered. In other words, a Σ -algebra is specified, so that every tree evaluates to an element of \mathbb{A} . This means that a device generating trees provides the syntactic basis for a *tree-based generator* – a system consisting of a tree generator and an interpretation, thus generating elements of \mathbb{A} :



TREEBAG makes it possible to assemble tree-based generators. This is explained in slightly more detail in Section 2, because MARBLES is to a certain extent inspired by TREEBAG. However, in TREEBAG, all that can be done after having assembled a tree generator is to execute it. In contrast, the usefulness of tree automata in most application areas does not primarily lie in the fact that they can be executed. Their real advantage is that they are simple enough to be effectively analyzed and manipulated. For instance, in a model checking application, tree automata may be generated that model safety and liveness properties of a protocol to be verified. Analyzing these automata then corresponds to checking correctness criteria, thus solving a model checking problem. The idea behind MARBLES is, therefore, to make it possible to assemble *algorithms on tree automata*, thus perceiving tree automata mainly as the objects to be analyzed and manipulated, rather than as executable algorithms. The major aim is to provide researchers with a software environment and infrastructure that enables them to create, use, and experiment with algorithms on tree automata.

In addition to TREEBAG, there are several other systems that implement certain types of tree automata or algorithms on them.

AutoWrite (<http://dept-info.labri.fr/~idurand/autowrite>) is a system that allows the user to check properties of term rewrite systems by means of tree automata constructions. In particular, it allows to load, save, and combine bottom-up tree recognizers. Using the graphical user interface, one can build and manipulate bottom-up tree recognizers related to the term rewrite systems whose properties one wants to check.

Forest FIRE (<http://www.loekcleophas.com>) is a toolkit focusing on recognition, pattern matching, and parsing algorithms in connection with regular tree languages. The system has been developed on the basis of detailed taxonomies, with the major purpose of gaining a deeper conceptual understanding of how the ideas and techniques used in various tree automata constructions are related to each other.

MONA (<http://www.brics.dk/mona>) is a tool for checking the validity of formulas in the weak second-order theory of one successor (WS1S) or of two

successors (WS2S). For deciding WS2S, the decision procedures convert a given formula into a so-called guided tree automaton, a variant of a bottom-up tree recognizers, and analyse this automaton.

Tiburon (<http://www.isi.edu/licensed-sw/tiburon>) is a command-line based package of algorithms on weighted regular tree grammars, context-free string grammars, and tree transducers, including various analyzers, modifiers, and synthesizers. The devices and algorithms implemented in Tiburon are typical even for MARBLES, but Tiburon has mainly been developed for applications in Natural Language Processing, without MARBLES' emphasis on a flexible environment that can be adapted and extended to suit the needs of researchers who study tree automata from different points of view.

Timbuk (<http://www.irisa.fr/lande/genet/timbuk>) is a collection of tools for carrying out reachability proofs of term rewrite systems, among other techniques by manipulating nondeterministic bottom-up tree recognizers. It is intended to be used for the verification of programs and cryptographic protocols.

The proposed system MARBLES differs from each of these systems in several respects. Most notably, the systems above have all been developed with a particular application or problem area in mind. They are great for their particular purpose, but they are also restricted to it. In contrast, the intention behind MARBLES is to support tree automata research in general, by providing researchers with a suitable platform and infrastructure for their own extensions, making it possible to experiment with and apply tree automata algorithms of any kind.

The remainder of this paper is structured as follows. The next section presents some aspects of TREEBAG that have, in one way or the other, inspired the intended characteristics of MARBLES. In Section 3, some of the different types of trees, tree automata, and tree automata algorithms that should, in principle, be covered by MARBLES, are discussed. Section 4 presents initial ideas regarding some of the concepts needed for making this possible. Finally, Section 5 concludes the paper.

2 TREEBAG

Let us now have a slightly closer look at the concepts and design principles of TREEBAG. The following description is intentionally kept at a rather abstract level, although concrete classes of, e.g., tree grammars and algebras available in TREEBAG are sometimes mentioned as examples, mainly for readers who happen to be familiar with tree automata theory. Readers who want to inform themselves in more detail should consult the TREEBAG user manual available at <http://www.cs.umu.se/~drewes/treebag> or, for the theory behind, [Dre06].

The work on TREEBAG was started during the second half of the 1990s, when the author was a member of Hans-Jörg Kreowski's research group at the University of Bremen. Around this time, a significant part of our research was dedicated to context-free graph and collage grammars; see, e.g., [HKV91, HKL93, HKT93, DHKT95, DK96, DHK97, DK99]. Both of these can be characterized by combi-

nations of a certain type of tree grammars (namely regular tree grammars) with suitable algebras in the style of Mezei and Wright [MW67], i.e., the grammars can be viewed as tree-based generators. For graphs, this has been made explicit by Engelfriet in [Eng94], and for collages by the author in [Dre96, Dre00]. See also [DEKK03], where this characterization was used to establish certain decidability results for collage languages. In this context, one should not forget to mention Engelfriet's paper [Eng80], where he discusses symbolic computation by tree transductions, which is essentially the same idea, but now for transformation rather than generation: a tree transduction, together with algebras interpreting the input and output trees, is considered as a symbolic algorithm that performs a computation on abstract trees rather than on the concrete objects of the two domains in question.

Whereas the results mentioned above use only regular tree grammars, it is obvious that one may in fact combine arbitrary kinds of tree generators with any sort of algebra, yielding a large number of different grammatical formalisms with comparatively little effort. Being a rather straightforward implementation of this idea (in Java), `TREEBAG` allows its user to assemble tree-based generators of various kinds. There are four major abstract classes, namely *tree grammars*, *tree transducers*, *algebras*, and *displays*. The first three represent the corresponding formal concepts, whereas displays are required for actually being able to see the results of the generating process. Concrete subclasses of the four abstract classes implement particular types of tree grammars, tree transducers, algebras, and displays. For example, the classes `generators.ETOLTreeGrammar` and `generators.mtTransducer` implement ETOL tree grammars and macro tree transducers, resp. If a class such as these is available, this means that the user can define specific instances (usually in ordinary ASCII text files) and use them in assembling tree-based generators. Such instances are called components in the following.

Figure 1 shows a typical situation when working with `TREEBAG`. Window 1 is the main window of the system, the so-called worksheet. When the user loads a component, it is represented on the worksheet as a blob. These blobs represent the nodes of a directed acyclic graph whose edges determine the data flow between components. The data-flow edges are interactively established by the user, subject to a few rather obvious rules: The output of a tree grammar or tree transducer can become the input of tree transducers or algebras, and the output of an algebra can become the input of a display. The configuration in Figure 1 consists of a regular tree grammar, a free term algebra with a corresponding tree display, a top-down tree transducer, and two copies of a collage algebra, each with its corresponding collage display. With each display component, a window is associated, namely the windows numbered 3–5. Thus, these windows show the tree generated by the regular tree grammar, its interpretation by the collage grammar, and the interpretation of the transformed tree by (another instance of) the same collage algebra.

An additional window (numbered 2 in the figure) contains buttons that provide access to the user commands of the regular tree grammar. Double clicks on

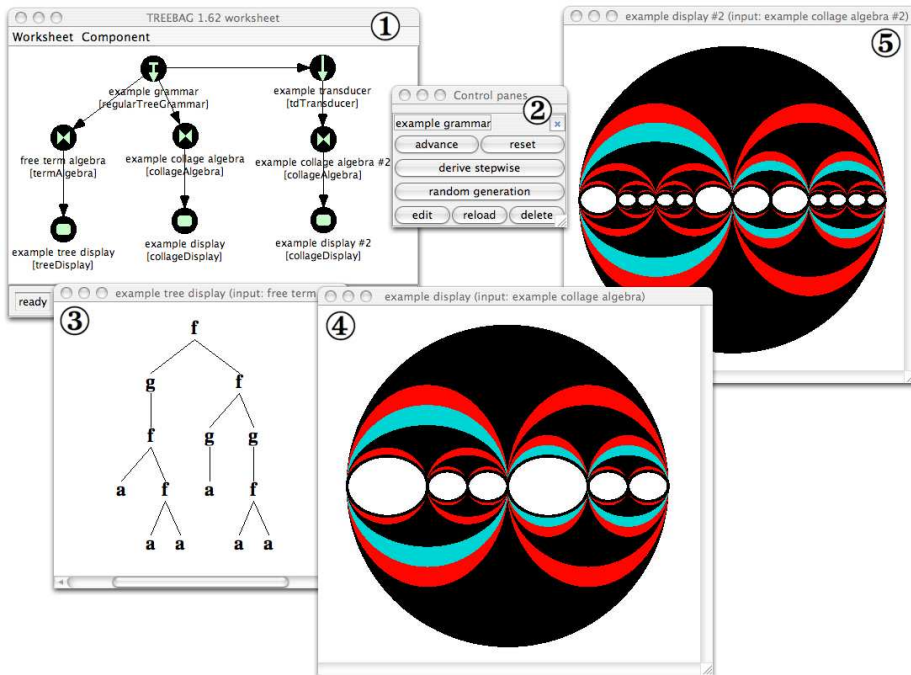


Fig. 1. A typical configuration of TREEBAG

the other components on the worksheet would open similar sections in this window for them, each one being populated by the individual commands understood by the respective component.

Let us now discuss two aspects of the design of TREEBAG which are expected to have some influence on MARBLES. In fact, these two aspects are quite closely related and can be seen as the two sides of the same coin.

From the point of view of the user, the way in which components can be interconnected depends only on their types, i.e., whether they are tree grammars, tree transducers, algebras or displays. In other words, if the user wants to connect a tree grammar and a tree transducer, this can be done regardless of whether the tree grammar at hand is a regular tree grammar, ETOL tree grammar, context-free tree grammar or whatever type of tree grammar might at some point in time be implemented in TREEBAG. Of course, users must interconnect the “right” components to achieve a particular effect desired, but the system gives users as much freedom as possible. Every concrete component class provides the user with a set of commands that can be used to interact with components of this class (recall Window 2 in Figure 1, containing buttons for the commands provided by the implementation of regular tree grammars). While the commands would be different for, e.g., ETOL tree grammars, this has no influence on the way in

which regular tree grammars or ETOL tree grammars can be connected to other components.

The person who implements new classes of tree grammars, tree transducers, algebras or displays will find out that the properties mentioned in the previous paragraph simply reflect properties of the implementation. The core of TREEBAG does not make any distinction between, e.g., different classes of tree grammars. In fact, consider the file defining the regular tree grammar used in Figure 1:

```
generators.regularTreeGrammar("example grammar"):
( { S, A },
  { f:2, g:1, a:0 },
  { S -> f[S,S],
    S -> g[A],
    A -> f[A,A],
    A -> a },
  S )
```

When the user instructs TREEBAG to load this component, it parses only the first line, to discover that the user wishes to load an instance of a component class named `generators.regularTreeGrammar`. The rest of the file uses a syntax which is unknown to (the core of) TREEBAG, as it is specific to the implementation of this class. To handle this, TREEBAG dynamically tries to load the class `generators.regularTreeGrammar` and, upon success, creates an (uninitialized) object of this class. Now, it lets this very object, which is required to contain a method called `parse`, initialize itself by parsing the remainder of the file. Each of the four abstract component types of TREEBAG requires its concrete subclasses to implement such a parsing method. To handle component-specific user commands, each concrete subclass provides two further methods. The first returns, at any point in time, the names of the user commands available at that moment (which means that the list of commands may change), while the second executes a given command.

This structure makes it possible to extend TREEBAG by new classes of tree grammars, tree transducers, algebras, and displays in an easy way, without having to change existing parts of the system. One only has to implement it as a subclass of the appropriate abstract component class and place it in the appropriate directory. Immediately afterwards (provided that everything has been done correctly), it is possible to load instances of this class onto the worksheet, interconnect them with other components, and work with them.

It may be interesting to note that the implementations of some of the classes currently available in TREEBAG make use of decomposition results from the literature. For example, a so-called branching synchronization tree grammar of nesting depth n can be decomposed into a regular tree grammar and a sequence of n top-down tree transducers (see [DE04]). During the parsing step, the implementation of this class in TREEBAG performs this decomposition and writes the $n + 1$ components onto the hard disk (in the syntax required by the respective classes). Afterwards, it uses TREEBAG's loading mechanism to load them

as internal variables hidden from the user (i.e., so that they do not appear on the worksheet). Every user command is basically forwarded to these internal components, and whichever output tree they produce is returned. In this way, the implementation of the class becomes considerably easier and less error prone than a direct one.

3 Trees, Tree Automata, and Tree Automata Algorithms

As mentioned in the introduction, the major intended purpose of MARBLES is to make it possible to apply and experiment with algorithms on tree automata. The aim is to design MARBLES in such a way that it accommodates virtually all kinds of tree automata algorithms. While this does not mean that all such algorithms should readily be implemented in the system, the design of MARBLES should enable researchers (and application programmers) interested in a particular type of tree automata algorithms to make the necessary extensions. As in the case of TREEBAG, this should be possible without requiring changes of already existing parts. However, compared to TREEBAG, the design challenge is considerably bigger for MARBLES, because its intended coverage is much wider. It seems to be reasonable to distinguish between (at least) three central categories of objects: trees, tree automata, and tree automata algorithms. Each of them may, in principle, have any number of subcategories one may wish to implement in MARBLES. In the following, some of the possible subcategories of each will be discussed to illustrate this point.

3.1 Trees

In the traditional setting (and in TREEBAG), tree automata work on trees over ranked alphabets, as explained above. This is appropriate, because trees are supposed to be evaluated by algebras by associating with every symbol of rank k a k -ary function on some domain. However, tree automata on unranked trees have received a lot of attention during recent years. Here, symbols are unranked, and a node in a tree can have any finite number of children, regardless of the symbol it is labeled with. It turns out that this variant is well suited for applications in connection with XML, because XML documents can appropriately be viewed as unranked trees. (For example, a node corresponding to a list structure in HTML may have any number of children of type *list item*.) Thus, an XML document type corresponds to a tree language of unranked trees, and a tree transducer on unranked trees corresponds to a transformation between XML document types.

While the two types of trees mentioned are the only ones that play a major role in contemporary research on tree automata, this situation may change in the future. Thus, MARBLES should allow programmers to implement other classes of trees than just these.

3.2 Tree Automata

Tree automata can be classified according to various criteria. An important observation one can make is that the resulting classifications are, to a rather large extent, orthogonal.

Perhaps the most obvious classification is the one that gave rise to the structure of TREEBAG, distinguishing between tree grammars, tree recognizers (not directly available in TREEBAG), and tree transducers. From an abstract point of view, a tree grammar is a formal device that generates output trees without requiring input. As usual, the tree recognizer is the dual concept. It takes a tree as input and computes an output value, usually in the range $\{0, 1\}$, indicating whether the tree is accepted or not. Finally, a tree transducer is a formal device transforming input trees into output trees.

The second classification distinguishes between tree automata according to the type of trees they act upon, i.e., tree automata on ranked or unranked trees. Each of the types of tree automata in the first classification can be ranked or unranked. In this sense, these two classifications are orthogonal. In fact, one may even wish to consider tree transducers that turn unranked trees into ranked ones, or vice versa.

Finally, in addition to the traditional case of tree automata, one may consider weighted ones [FV09]. Weighted tree automata deal with tree series instead of tree languages, a tree series being a mapping $\psi: T_\Sigma \rightarrow \mathbb{S}$, where T_Σ denotes the set of all trees over a given alphabet, and \mathbb{S} is a semiring. In other words, weighted tree automata generalize the traditional case, which is obtained by choosing as \mathbb{S} the Boolean semiring. Even this third classification is orthogonal to the two previous ones, provided that we define the tree automata according to the first classification in a way general enough to accommodate the weighted case.

It is interesting to note that, from an abstract point of view, but even more from the point of view of system design, weighted tree recognizers are very similar to algebras. Both take a tree as input and compute a value in some other domain.

3.3 Algorithms on Tree Automata

Many useful algorithms on tree automata have been described in the literature. For classification purposes, it is useful to distinguish between analyzers, synthesizers, and decomposition algorithms.

An analyzer for tree automata takes a tree automaton as input and analyses it with respect to certain properties. Well-known examples are algorithms that decide whether the language represented by a tree recognizer or tree grammar is empty or whether it is finite (cf., e.g., [DE98]).

A synthesizer is an algorithm that takes zero or more tree automata (and maybe some additional data) as input and yields a tree automaton as output. There are various important types of synthesizers:

- A generator is an algorithm that outputs tree automata without requiring other tree automata as input. A prominent example is given by grammatical

inference algorithms for tree automata. These are algorithms whose purpose it is to “learn” tree languages. For this, the algorithm is provided with some source of information regarding the tree language (or tree series) to be learned, such as positive and negative examples. It is then expected to construct a tree automaton representing the tree language in question. See, e.g., the references in [Dre09] for a variety of approaches.

Conceptually, a tree automaton A may be considered as a generator that outputs the constant value A .

- Conversion algorithms take a tree automaton as input and yield another tree automaton as output, usually with the same semantics as the input automaton. Well-known examples are conversions between regular tree grammars and finite-state tree recognizers and algorithms that minimize tree automata, make them deterministic, remove useless states or nonterminals, etc (see, e.g., [CDG⁺07]). There are also conversion algorithms that do not retain the semantics of the tree automaton they are applied to. For example, a macro tree transducer mtt [EV85] may be turned into a finite-state tree recognizer that accepts the pre-image of the tree transformation computed by mtt . A conversion algorithm that inverts suitable types of top-down tree transducers would be another example.
- Composition algorithms turn n tree automata ($n > 1$) into one. A wealth of such algorithms can be found in the literature. One type of example is, of course, given by composition in the strict sense. For instance, certain types of tree transductions are known to be closed under composition. Another example is the main result of [DE04], which provides an algorithm for converting a regular tree grammar g and n top-down tree transducers td_1, \dots, td_n into a branching synchronization tree grammar generating the image of $L(g)$ under $td_n \circ \dots \circ td_1$. Composition algorithms in a more general sense may not perform mathematical composition, but combine tree automata in a different way. For example, two finite-state tree recognizers can be turned into one that recognizes the intersection of the tree languages recognized by the two individual automata.

Finally, decomposition algorithms are the conceptual inverse of composition algorithms, turning one tree automaton into several others. For example, for $\{x, y\} = \{\text{top-down}, \text{bottom-up}\}$, every x tree transducer may be decomposed into two y tree transducers [Eng75]. A similar example is given by the result that every deterministic total macro tree transducer may be decomposed into a top-down tree transducer followed by a YIELD mapping [EV85].

Of course, algorithms on tree automata may additionally be classified according to the types of tree automata they work on, similarly to the fact that tree automata may be classified according to the types of trees they work on.

4 A Proposed Attribute Type System for MARBLES

As mentioned earlier, the goal behind the development of MARBLES is that it should allow its user to assemble configurations of tree automata algorithms in

a similar way as TREEBAG allows its user to assemble various sorts of tree-based generators. In particular, there should be a way to load components representing (tree automata and) tree automata algorithms, establish a data-flow relation between, and execute them. However, while TREEBAG comes with a fixed set of component types, something like this is neither possible nor desirable for MARBLES. In contrast, users should be given the possibility to define and implement their own classes of tree automata algorithms and experiment with them. The following two fictitious scenarios try to illustrate this.³

Scenario 1: Test Environment for Minimization Algorithms. Doctoral student X works in a research group using bottom-up tree recognizers for model checking purposes. A typical example is the verification of a process communication protocol P by generating a tree recognizer A_P that models P 's behavior and then analyzing A_P to establish P 's correctness. The problem is that A_P tends to be huge, and often unnecessarily huge, so that its analysis takes too much time. Unfortunately, A_P is also nondeterministic, which means that it cannot efficiently be minimized.

Therefore, in her thesis, X proposes and studies a number of efficient heuristics for reducing nondeterministic tree recognizers A in size (called minimization, for simplicity). The general technique used is to compute a suitable equivalence \equiv on the state set of A , such that the quotient automaton A/\equiv accepts the same language as A . The various heuristics studied differ only in the concrete definition (and computation) of \equiv . Besides studying the minimization algorithms theoretically to establish their correctness and worst case complexity, X wants to study empirically how they behave on real examples arising in the model checking context, in terms of size reduction and efficiency. However, X does not have the time to implement a test environment for her algorithms from scratch, in addition to her theoretical studies. Therefore, she decides to use MARBLES.

First, she notices that there is a type of tree automata algorithms called generator, a special type of synthesizer. She defines and implements a simple generator which lets the user choose the name of a protocol (from a fixed set of possible choices) and possibly some other parameters. The generator will then output nondeterministic bottom-up tree recognizers of increasing size, whenever the user presses a certain button.

Next, X discovers that there are so-called converters, and decides to implement a new type of converter as an abstract class. A concrete implementation is obtained by providing a method that, for a given bottom-up tree recognizer A , computes an equivalence relation \equiv on the states. The converter will then return A/\equiv .

Fortunately, X finds out that someone else has already implemented two useful auxiliary components. One of them is a wrapper for arbitrary converters that simply executes them, but also reports how much time the execution takes. The other one takes bottom-up tree recognizers as input and saves some statistics

³ While being fictitious, the scenarios have a real background, as they are inspired by [Kaa08] and ongoing work in our own group, resp.

about them to a file, such as the number of states and transitions. Now, X has everything needed to make the desired tests. All she has to do is to implement the different algorithms yielding the equivalence relations \equiv , load and interconnect the required components, and execute them.

Scenario 2: Simulation of Minimal Adequate Teachers Using Corpora.

The research group in which researcher Y is working has previously studied grammatical inference algorithms that, within Angluin's learning model of a *minimal adequate teacher* (MAT), construct bottom-up tree recognizers for recognizable tree languages L . Now, they want to find out whether such an algorithm can be used to learn the syntax of natural languages reasonably well, where the necessary data is taken from a corpus.⁴

The major obstacle is the MAT, an oracle capable of answering two types of queries, namely membership queries (*Is the tree t in L ?*) and equivalence queries (*Does the bottom-up tree recognizer A satisfy $L(A) = L$? If not, return a counterexample.*) Clearly, a MAT is not available in the situation sketched above. The research question is whether it can (imperfectly) be simulated on the basis of a corpus, so that the inference algorithm as a whole runs with reasonable efficiency and yields acceptable results.

Y decides to try out some approaches and to use MARBLES for that purpose. Thus, she defines two new types of algorithms, namely MATs and learners. A learner is a generator that must be connected to a MAT to create a tree automaton. During the first phase, she only wants to test different realizations of the MAT, to see whether the results are promising enough to continue. Therefore, she implements a single learner (e.g., any of those in [Dre09]). In contrast, a variety of different MATs are implemented, using different approaches for answering membership and equivalence queries based on a corpus.

To find out how good the various approaches are, Y implements a component that has access to a sufficiently large sample of positive and negative examples. It takes a tree recognizer as input, runs it on the samples, and returns statistics regarding its sensitivity and specificity. In a second phase of her research work, Y even wants to study other variants of the learner, which can be done in the same setting by replacing the one learner with another.

In scenarios such as those above, the researcher who wants to use MARBLES must implement certain extensions, new types of tree automata and algorithms that become components of MARBLES. For both the system and the user, it is necessary to know in which way instances of these components can be combined. Thus, there must be a possibility to talk about the types of components in an easy, but flexible way. A prerequisite for this is to be able to specify which basic data types exist. Here, we only focus on the perspective of the user and the GUI, which means that the only thing we need is a way to *give name* to different sorts of data.

⁴ A corpus is a manually analyzed and annotated database of sentences in a natural language.

While the tree is *the* basic data type in MARBLES, it may not be the only one. Moreover, there may be different sorts of trees. We now define basic types which make it possible to name those structures.

Definition 1 (basic type). Let $ATTR$ be a finite set of data attributes (briefly called attributes). The set $TEXP$ of all basic types is the smallest set of pairs such that, for all finite subsets T of $TEXP$ and all attributes a , we have $(a, T) \in TEXP$. If $T = \{t_1, \dots, t_n\}$, then this basic type is also denoted by $a\langle t_1, \dots, t_n \rangle$, or by a if $n = 0$.

As an example, consider trees. We may, e.g., have ranked, unranked, ordered, and unordered trees. Ranked trees may or may not be binary or monadic. Our set of data attributes could then be $ATTR = \{tree, ranked, unranked, ordered, unordered, bin, mon\}$. The basic type for ranked unordered binary trees would then be $tree\langle ranked\langle bin \rangle, unordered \rangle$. The attributes in such a basic type should be seen as assertions stating that the data in question has certain properties. In other words, the presence of an attribute restricts the data type. For example, $tree$ is a basic type meaning just any tree, and $tree\langle unordered \rangle$ means “any type of unordered trees”. Note that we, intentionally, do not associate a specific semantics with the attributes. It should, however, be *possible* to do this in MARBLES by, e.g., associating an attribute with an abstract class in the implementation. A similar remark applies to the types at the higher levels discussed next.

Next, we define what the type of an automaton looks like. We take a very general approach, where an automaton is a device that turns a finite number of input values of specified basic types and into a finite number of output values, also of specified basic types.

Definition 2 (automaton type). An automaton type is a pair (in, out) with $in \in TEXP^k$ and $out \in TEXP^l$ for some $k, l \geq 0$. Such a type will normally be written as $in \rightarrow out$. The set of all automaton types is denoted by AUT .

As an example, a tree grammar of the most general form could be described as being an automaton of type $() \rightarrow (tree)$, as it takes no input and yields any type of tree as output. Slightly more specific would be a tree grammar generating ranked trees, its type being $() \rightarrow (tree\langle ranked \rangle)$. For weighted tree automata over a semifield that work on ranked trees, the type $(tree\langle ranked \rangle) \rightarrow (semiring(semifield))$ could be an appropriate description, and for tree transducers on unranked trees one could use $(tree\langle unranked \rangle) \rightarrow (tree\langle unranked \rangle)$. Though uncommon in the literature, one may also wish to consider, e.g., tree transducers that take two trees as input and produce one output tree, the corresponding type being $(tree, tree) \rightarrow (tree)$.

Note that the concept is very general. For example, an algebra can be seen as an automaton of type $(tree\langle ranked \rangle) \rightarrow (any)$, if we let any be the most general basic type, standing for arbitrary data. Also weighted tree automata over multioperator monoids [Kui00] have this type. In fact, the concept covers even devices that do not work on trees at all.

We finally define how to distinguish between different types of algorithms on tree automata.

Definition 3 (algorithm type). *The set ALG of algorithm types is inductively defined to be the smallest set containing all triples (in, use, out) such that, for some $k, l, m \geq 0$, $in \in AUT^k$, $out \in AUT^m$, and $use \in ALG^l$. Such a triple is denoted by $in \xrightarrow{use} out$.*

The intuitive interpretation of $in \xrightarrow{use} out$ is that of an algorithm which turns inputs according to in into outputs according to out , thereby possibly making use of other algorithms given by use . A typical example is the MAT learner in Scenario 2, which could be of the type $() \xrightarrow{MAT} (TA)$, where TA is the automaton type $tree\langle ranked \rangle \rightarrow bool$.

As mentioned earlier, one of the ideas behind MARBLES is that its GUI, similar to the one of TREEBAG, should allow the user to assemble configurations of tree automata in order to experiment with them. The basic (and still somewhat tentative) plan is that every implementation of a class of tree automata or tree automata algorithms comes with a specified type according to the definitions above. When the user loads an instance of such a component, this information is used in order to determine which connections between these components are possible. For example, an algorithm of the type in Definition 3 will, from the point of view of the user, have $k + l + m$ slots representing the inputs, the used algorithms, and the outputs. For instance, if a component has an output slot s of type $tree\langle ranked \rangle \rightarrow bool$ (a recognizer for ranked trees) and another one has an input slot s' of type $tree \rightarrow bool$ (a recognizer for any sort of trees), then the data flow can be directed from s to s' .

5 Concluding Remarks

In this paper, ideas and plans regarding a successor of the system TREEBAG have been presented. While this work is still in a very preliminary phase, the overall goal is clear. MARBLES should make it possible to experiment with configurations of tree automata algorithms in a similar way as TREEBAG makes it possible to experiment with tree-based generators. Moreover, MARBLES should be extensible by researchers who are not directly involved in the development of the system itself, but want to use it for their own purposes. For this, concepts such as those presented in Section 4 seem to be a necessity, because the GUI must be able to handle extensions without explicitly being adapted.

An aspect that has not been discussed in the present paper, but which is a necessity as well, is to provide programmers with a well-documented library and a clearly structured application programmer's interface (API). Without this, it would be too difficult, error prone, and time consuming for other researchers to make their own extensions. In fact, it should also be possible to make use of the API without adopting the rest of MARBLES, and especially its GUI. This would allow programmers to give the possibility to apply tree automata algorithms in their own applications. Another aspect that has not yet been decided upon is whether and

to what extent MARBLES shall be compatible and able to interoperate with other systems dealing with tree automata, such as those mentioned in Section 1.

Acknowledgment I thank the anonymous referees for their thorough reading of the manuscript and for giving numerous useful comments.

However, most of all, I want to thank you, Hans-Jörg, for the support and inspiration during all those years, for teaching me so much about our profession, and for being a good example in all respects. Happy Birthday to you!

References

- [AJMd02] Parosh Aziz Abdulla, Bengt Jonsson, Pritha Mahata, and Julien d’Orso. Regular tree model checking. In E. Brinksma and K. Guldstrand Larsen, editors, *Proc. 14th Intl. Conf. on Computer Aided Verification (CAV’02)*, volume 2404 of *Lecture Notes in Computer Science*, pages 555–568, 2002.
- [CDG⁺07] Hubert Comon, Max Dauchet, Rémi Gilleron, Florent Jacquemard, Christof Löding, Denis Lugiez, Sophie Tison, and Marc Tommasi. *Tree Automata Techniques and Applications*. Internet publication available at <http://tata.gforge.inria.fr>, 2007. Release October 2007.
- [DE98] Frank Drewes and Joost Engelfriet. Decidability of the finiteness of ranges of tree transductions. *Information and Computation*, 145:1–50, 1998.
- [DE04] Frank Drewes and Joost Engelfriet. Branching synchronization grammars with nested tables. *Journal of Computer and System Sciences*, 68:611–656, 2004.
- [DEKK03] Frank Drewes, Sigrid Ewert, Renate Klempien-Hinrichs, and Hans-Jörg Kreowski. Computing raster images from grid picture grammars. *Journal of Automata, Languages and Combinatorics*, 8:499–519, 2003.
- [DHK97] Frank Drewes, Annegret Habel, and Hans-Jörg Kreowski. Hyperedge replacement graph grammars. In G. Rozenberg, editor, *Handbook of Graph Grammars and Computing by Graph Transformation. Vol. 1: Foundations*, chapter 2, pages 95–162. World Scientific, Singapore, 1997.
- [DHKT95] Frank Drewes, Annegret Habel, Hans-Jörg Kreowski, and Stefan Taubenberger. Generating self-affine fractals by collage grammars. *Theoretical Computer Science*, 145:159–187, 1995.
- [DK96] Frank Drewes and Hans-Jörg Kreowski. (Un-)decidability of geometric properties of pictures generated by collage grammars. *Fundamenta Informaticae*, 25:295–325, 1996.
- [DK99] Frank Drewes and Hans-Jörg Kreowski. Picture generation by collage grammars. In H. Ehrig, G. Engels, H.-J. Kreowski, and G. Rozenberg, editors, *Handbook of Graph Grammars and Computing by Graph Transformation, Vol. 2: Applications, Languages, and Tools*, chapter 11, pages 397–457. World Scientific, Singapore, 1999.
- [Dre96] Frank Drewes. Language theoretic and algorithmic properties of d -dimensional collages and patterns in a grid. *Journal of Computer and System Sciences*, 53:33–60, 1996.

- [Dre00] Frank Drewes. Tree-based picture generation. *Theoretical Computer Science*, 246:1–51, 2000.
- [Dre06] Frank Drewes. *Grammatical Picture Generation – A Tree-Based Approach*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2006.
- [Dre09] Frank Drewes. Mat learners for recognizable tree languages and tree series. *Acta Cybernetica*, 2009. To appear.
- [Eng75] Joost Engelfriet. Bottom-up and top-down tree transformations – a comparison. *Mathematical Systems Theory*, 9:198–231, 1975.
- [Eng80] Joost Engelfriet. Some open questions and recent results on tree transducers and tree languages. In R. V. Book, editor, *Formal Language Theory: Perspectives and Open Problems*, pages 241–286. Academic Press, New York, 1980.
- [Eng94] Joost Engelfriet. Graph grammars and tree transducers. In S. Tison, editor, *Proceedings of the CAAP 94*, volume 787 of *Lecture Notes in Computer Science*, pages 15–37. Springer, 1994.
- [EV85] Joost Engelfriet and Heiko Vogler. Macro tree transducers. *Journal of Computer and System Sciences*, 31:71–146, 1985.
- [FGVTT04] Guillaume Feuillade, Thomas Genet, and Valérie Viet Triem Tong. Reachability analysis over term rewriting systems. *Journal of Automated Reasoning*, 33:341–383, 2004.
- [FSW94] Christian Ferdinand, Helmut Seidl, and Reinhard Wilhelm. Tree automata for code selection. *Acta Informatica*, 31(8):741–760, 1994.
- [FV98] Zoltán Fülöp and Heiko Vogler. *Syntax-Directed Semantics: Formal Models Based on Tree Transducers*. Springer, 1998.
- [FV09] Zoltán Fülöp and Heiko Vogler. Weighted tree automata and tree transducers. In Werner Kuich, Manfred Droste, and Heiko Vogler, editors, *Handbook of Weighted Automata*. Springer, 2009.
- [GK00] Thomas Genet and Francis Klay. Rewriting for cryptographic protocol verification. In D.A. McAllester, editor, *Proc. 17th International Conference on Automated Deduction (CADE'00)*, volume 1831 of *Lecture Notes in Computer Science*, pages 271–290, 2000.
- [GKM08] Jonathan Graehl, Kevin Knight, and Jonathan May. Training tree transducers. *Computational Linguistics*, 34(3):391–427, 2008.
- [GS84] Ferenc Gécseg and Magnus Steinby. *Tree Automata*. Akadémiai Kiadó, Budapest, 1984.
- [GS97] Ferenc Gécseg and Magnus Steinby. Tree languages. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*. Vol. 3: *Beyond Words*, chapter 1, pages 1–68. Springer, 1997.
- [HKL93] Annegret Habel, Hans-Jörg Kreowski, and Clemens Lautemann. A comparison of compatible, finite, and inductive graph properties. *Theoretical Computer Science*, 110:145–168, 1993.
- [HKT93] Annegret Habel, Hans-Jörg Kreowski, and Stefan Taubenberger. Collages and patterns generated by hyperedge replacement. *Languages of Design*, 1:125–145, 1993.
- [HKV91] Annegret Habel, Hans-Jörg Kreowski, and Walter Vogler. Decidable boundedness problems for sets of graphs generated by hyperedge-replacement. *Theoretical Computer Science*, 89:33–62, 1991.
- [Kaa08] Lisa Kaati. *Reduction Techniques for Finite (Tree) Automata*. Doctoral dissertation, Uppsala University, Sweden, 2008.

- [KG05] Kevin Knight and Jonathan Graehl. An overview of probabilistic tree transducers for natural language processing. In Alexander F. Gelbukh, editor, *Proc. 6th Intl. Conf. on Computational Linguistics and Intelligent Text Processing (CICLing 2005)*, volume 3406 of *Lecture Notes in Computer Science*, pages 1–24. Springer, 2005.
- [Kui00] Werner Kuich. Linear systems of equations and automata on distributive multioperator monoids. In D. Dorninger, G. Eigenthaler, M. Goldstern, H.K. Kaiser, W. More, and W.B. Müller, editors, *Proc. 58th Workshop on General Algebra (1999)*, volume 12 of *Contributions to General Algebra*, pages 247–256, Klagenfurt, 2000. Johannes Heyn.
- [Lö02] Christof Löding. Model-checking infinite systems generated by ground tree rewriting. In M. Nielsen and U. Engberg, editors, *Proc. 5th Intl. Conf. on Foundations of Software Science and Computation Structures (FOSSACS'02)*, volume 2303 of *Lecture Notes in Computer Science*, pages 280–294, 2002.
- [MW67] Jorge Mezei and Jesse B. Wright. Algebraic automata and context-free sets. *Information and Control*, 11:3–29, 1967.
- [NP92] Maurice Nivat and Andreas Podelski, editors. *Tree Automata and Languages*. Elsevier, Amsterdam, 1992.
- [Sch07] Thomas Schwentick. Automata for XML - a survey. *Journal of Computer and System Sciences*, 73(3):289–315, 2007.



Dr. Frank Drewes

Institutionen för datavetenskap
 Umeå universitet
 S-90187 Umeå (Sweden)
 drewes@cs.umu.se
<http://www.cs.umu.se/~drewes>

Frank Drewes studied Computer Science at the University of Bremen. He was introduced to Theoretical Computer Science by Hans-Jörg Kreowski and Dr. Clemens Lautemann who, during those years, was a member of Hans-Jörg's team. After his graduation in 1990, Frank became a doctoral student supervised by Hans-Jörg. After the defense in 1996, he worked as an assistant professor in Hans-Jörg's team until he accepted an offer from the University of Umeå in 2000.