

A Circus of Circuits: Connections Between Decision Diagrams, Circuits, and Automata

Antoine Amarilli

Marcelo Arenas

YooJung Choi

Mikaël Monet

Guy Van den Broeck

Benjie Wang

Abstract

This document is an introduction to two related formalisms to define Boolean functions: binary decision diagrams, and Boolean circuits. It presents these formalisms and several of their variants studied in the setting of knowledge compilation. Last, it explains how these formalisms can be connected to the notions of automata over words and trees.

Contents

1	Introduction	2
2	Preliminaries	3
3	Binary Decision Diagrams	3
3.1	Basic classes	3
3.2	Binary decision diagrams with or-nodes	6
3.3	Completion and zero-suppressed semantics	8
3.4	Decision trees	10
4	Boolean Circuits	10
4.1	Definitions	11
4.2	Sentential decision diagrams	13
4.3	Formulas	14
4.4	Smoothness	15
5	From Binary Decision Diagrams to Boolean Circuits	15
6	Automata	17
6.1	Automata on words	18
6.2	Automata on trees	18
6.3	Computing Boolean circuits and binary decision diagrams from automata	19
7	Conclusion and Extensions	22

1 Introduction

This document is about *Boolean functions* and formalisms to represent them. Given their naturalness, many communities in theoretical and applied computer science use Boolean functions and study their representations. The *circuit complexity* community [Vollmer, 1999], for instance, studies classes of *Boolean circuits* defined by conditions such as the *depth*, the *size*, and the *kinds of gates* that are allowed. Other communities study *Boolean functions* [Wegener, 1987], or ways to represent them concisely, in particular as *decision diagrams* [Wegener, 2004]. Meanwhile, the *knowledge compilation* community, motivated by the practical use case of *SAT solvers*, investigates formalisms to represent Boolean functions as *binary decision diagrams*, or *Boolean circuits*, while ensuring the tractability of certain tasks. A central work in this area, but dating back to 2002, is the *knowledge compilation map* [Darwiche and Marquis, 2002]; the area also uses tools from neighboring fields such as *communication complexity* [Kushilevitz, 1997]. Last, the *database theory* community has also investigated Boolean functions, in particular in the setting of *provenance* [Green et al., 2007], including representations such as *provenance circuits* [Deutch et al., 2014]: in this area, restricted Boolean circuit classes were in particular studied for query evaluation on *probabilistic databases* [Jha and Suciu, 2013]. All told, this illustrates that the notions of *decision diagrams* (e.g., OBDDs) and *restricted circuit classes* (e.g., d-DNNFs) are studied to a large extent by separate communities.

The goal of this document is twofold. First, we propose a unified introduction to *binary decision diagrams* and *Boolean circuits*, using consistent terminology across the two paradigms. We hope that this can serve as an introduction to researchers familiar with one of the two formalisms, as a way to understand the connections with the other formalism. Second, we present a lesser-known correspondence that relates *automata on words and trees* to these binary decision diagrams and Boolean circuit classes, specifically, to the setting of *ordered* binary decision diagrams and *structured* Boolean circuit classes. The connection goes through the notion of *provenance circuits*, which relates automata to a Boolean function informally describing their behavior on words of a specific length: intuitively, provenance circuits are obtained by unraveling the automaton up to that length. The point of this correspondence is that well-known conditions on finite automata (e.g., determinism, unambiguity) relate to conditions on the resulting ordered binary decision diagrams (for word automata) or structured Boolean circuits (for tree automata).

Thus, our hope is that this document can serve as a “Rosetta stone” to highlight connections between binary decision diagrams, Boolean circuits, and automata, and encourage further interaction between the communities studying these formalisms.

Document structure. The document is structured in the following way. First, in Section 2, we give some common preliminaries. We then define the two main formalisms we use to represent Boolean functions: we define *binary decision diagrams* in Section 3, and then define *Boolean circuits* in Section 4. We then explain in Section 5 in which sense binary decision diagrams can be seen as a special case of Boolean circuits. Then, we introduce in Section 6 the notions of automata on words and on trees, and we explain how we can define *provenance circuits* for automata: this allows us to draw a correspondence between conditions on automata and conditions on provenance circuits.

Acknowledgements. This work was initiated while the authors were visiting the Simons Institute for the Theory of Computing, and was done in part during the corresponding program at the institute.

2 Preliminaries

We write sets of variables with uppercase boldface letters, such as $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$, and single variables with uppercase non-bold letters, such as $X \in \mathbf{X}, Y \in \mathbf{Y}$, etc.

Assignments and Boolean functions. Let \mathbf{X} be a finite set of variables. An *assignment* of \mathbf{X} is a function $\mathbf{a} : \mathbf{X} \rightarrow \{0, 1\}$. We denote by $\text{Assign}(\mathbf{X})$ the set of all assignments of \mathbf{X} . A *Boolean function over \mathbf{X}* is a function $f : \text{Assign}(\mathbf{X}) \rightarrow \{0, 1\}$. An assignment \mathbf{a} of \mathbf{X} is *satisfying* if $f(\mathbf{a}) = 1$ (also denoted $\mathbf{a} \models f$). We denote by $\text{Models}(f) \subseteq \text{Assign}(\mathbf{X})$ the set of all satisfying assignments of f , and $\#\text{Models}(f)$ the size of this set.

CNFs and DNFs. Let \mathbf{X} a set of variables. A *literal* is an expression of the form X or $\neg X$ for $X \in \mathbf{X}$. A *clause* is a disjunction of literals, for instance $X_1 \vee X_2 \vee \neg X_3$. A *formula in conjunctive normal form*, or *CNF* for short, is a formula that is a conjunction of clauses, i.e., an expression of the form $\bigwedge_{i=1}^n C_i$ where each C_i is a clause. We call a conjunction of literals a *term*. A *formula in disjunctive normal form*, or *DNF*, is a disjunction of terms i.e., an expression of the form $\bigvee_{i=1}^n t_i$ where each t_i is a term.

Directed graphs, DAGs, labeled DAGs. A *directed graph* is a tuple $G = (N, E)$ where N is the set of nodes and $E \subseteq N \times N$ is the set of edges. We say that G is *acyclic* if it contains no cycles, and call it a *directed acyclic graph* or *DAG*.

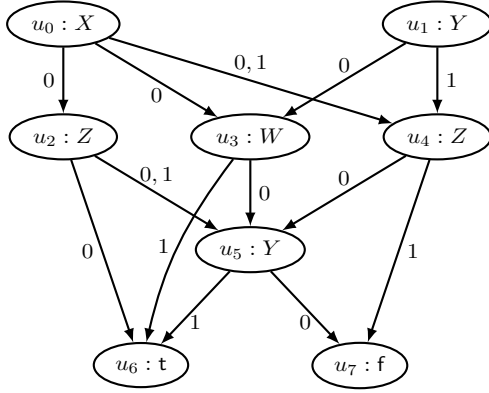
For a finite set Σ of labels, a Σ -*node-labeled directed graph*, or just node-labeled directed graph, is a tuple $G = (N, E, \lambda)$ where (N, E) is a directed graph and $\lambda : N \rightarrow \Sigma$ is the node labeling function. A Σ -*labeled directed graph*, or just labeled directed graph, is a node-labeled directed graph whose edges additionally carry labels. Formally, a Σ -labeled directed graph is a tuple $G = (N, E, \lambda)$ such that $E \subseteq N \times \Sigma \times N$ and $(N, \{(x, y) \mid \exists \ell : (x, \ell, y) \in E\}, \lambda)$ is a Σ -node-labeled directed graph. A node u of G is called a *source* if it does not have any incoming edges, a *sink* if it does not have any outgoing edges, and an *internal node* if it has some outgoing edges. If a labeled directed graph is acyclic, we call G a *labeled DAG*. Notice that the previous definition allows a labeled directed graph G to have multiple edges between the same pair of nodes, but each one with a different label. We define the *size* of G , written $|G|$, to be $|N| + |E| + |\Sigma|$.

3 Binary Decision Diagrams

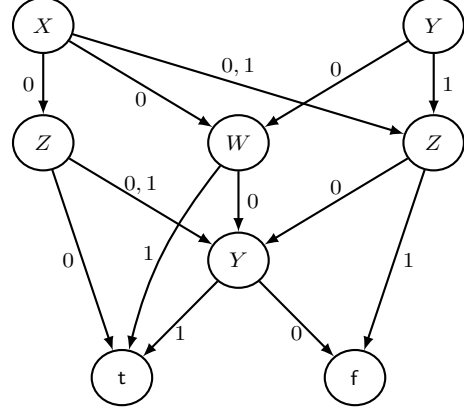
We start by giving formal definitions of (nondeterministic) binary decision diagrams and their semantics, and present two conditions on diagrams: variable structuredness (corresponding to free or ordered binary decision diagrams), and ambiguity levels. We then comment on an alternative way to express nondeterminism. After that, we present the case of binary decision diagrams without sharing (aka decision trees), and discuss the notion of completeness for binary decision diagrams, along with an alternative semantics called the zero-suppressed semantics.

3.1 Basic classes

A *nondeterministic binary decision diagram* (nBDD [Bollig and Wegener, 1997, Amarilli et al., 2020]) over a set of variables \mathbf{X} is a labeled DAG \mathcal{D} such that: (i) each edge is labeled with either the symbol 0 (a 0-*edge*) or the symbol 1 (a 1-*edge*); (ii) each sink is labeled with t (for true) or f (for false); and (iii) each internal node is labeled with a variable $X \in \mathbf{X}$ and has at least one



(a) A graphical representation of an nBDD including node identifiers and their labels.



(b) The usual graphical representation of an nBDD, where node identifiers are not included.

Figure 1: An nBDD over the set of variables $\mathbf{X} = \{X, Y, Z, W\}$.

outgoing 0-edge and at least one outgoing 1-edge. Given an assignment $\mathbf{a} \in \text{Assign}(\mathbf{X})$, a *run* π of \mathcal{D} following \mathbf{a} is a sequence u_1, u_2, \dots, u_k of nodes of \mathcal{D} such that u_1 is a source of \mathcal{D} , u_k is a sink of \mathcal{D} , and for each $i \in \{1, \dots, k-1\}$, letting X_i be the variable that labels u_i , there is an edge in \mathcal{D} from u_i to u_{i+1} with label $\mathbf{a}(X_i)$. Note that there is always at least one run following \mathbf{a} , which we can obtain by starting at an arbitrary source and following edges until we reach a sink: indeed, condition (iii) ensures that every node has an outgoing edge with the correct label. If π follows \mathbf{a} , we also say that \mathbf{a} is *consistent* with π : note that π can be consistent with different assignments if some variables do not occur as the label of any node of π . The run π is *accepting* if the label of the sink at the end of π is t . An assignment \mathbf{a} is *accepted* by \mathcal{D} , denoted by $\mathcal{D}(\mathbf{a}) = 1$ or $\mathbf{a} \models \mathcal{D}$, if there exists an accepting run of \mathcal{D} following \mathbf{a} ; otherwise \mathbf{a} is *rejected* by \mathcal{D} , denoted by $\mathcal{D}(\mathbf{a}) = 0$ or $\mathbf{a} \not\models \mathcal{D}$. In this paper, we sometimes refer to a run without explicitly mentioning an assignment that it follows. Finally, observe that \mathcal{D} represents a Boolean function over the set of variables \mathbf{X} , and we will often identify \mathcal{D} with the Boolean function that it represents.

Example 3.1. An nBDD \mathcal{D} over the set of variables $\mathbf{X} = \{X, Y, Z, W\}$ is shown in Figure 1a. For each node we include its identifier and label; for instance, $u_0 : X$ indicates that node u_0 has label X . Moreover, we depict edges with their labels; for instance, $(u_0, 0, u_2)$ is the only edge from u_0 to u_2 , while $(u_0, 0, u_4)$, $(u_0, 1, u_4)$ are the two edges from u_0 to u_4 (the symbol $0, 1$ next to the edge from u_0 to u_4 is used to denote two edges, one with label 0 and the other one with label 1). The two sources of \mathcal{D} are u_0 and u_1 , while the two sinks of \mathcal{D} are u_6 and u_7 .

Consider the assignment $\mathbf{a}_1 \in \text{Assign}(\mathbf{X})$ such that $\mathbf{a}_1(X) = \mathbf{a}_1(Y) = \mathbf{a}_1(Z) = \mathbf{a}_1(W) = 0$. Then we have that \mathbf{a}_1 is consistent with the run $\pi_1 = u_0, u_2, u_6$. Notice that an assignment can be consistent with many different runs of an nBDD; for instance, \mathbf{a}_1 is consistent with π_1 as well as with the run $\pi_2 = u_1, u_3, u_5, u_7$. An assignment is accepted by an nBDD if there exists at least one accepting run of the nBDD that is consistent with it; for instance $\mathcal{D}(\mathbf{a}_1) = 1$ since \mathbf{a}_1 is consistent with the accepting run π_1 ($\mathcal{D}(\mathbf{a}_1) = 1$ despite the fact that \mathbf{a}_1 is consistent with the run $\pi_2 = u_1, u_3, u_5, u_7$ and the label of sink u_7 is f).

We depict in Figure 1b the same nBDD as in Figure 1a, but without including node identifiers. The usual graphical representation of an nBDD is the one given in Figure 1b.

Table 1: Classification of nondeterministic binary decision diagrams based on ambiguity level (nondeterministic, unambiguous, or deterministic) and variable structuredness (unrestricted, free, or ordered).

	Unrestricted	Free	Ordered
Nondeterministic	nBDD	nFBDD	nOBDD
Unambiguous	uBDD	uFBDD	uOBDD
Deterministic	BDD	FBDD	OBDD

Nondeterministic binary decision diagrams are classified according to two dimensions.

- **Variable structuredness (free, ordered).** Let \mathcal{D} be an nBDD over a set of variables \mathbf{X} . Then \mathcal{D} is *free* (nFBDD) if, for every run π of \mathcal{D} , no two distinct nodes in π have the same label. In addition, \mathcal{D} is *ordered* (nOBDD) if there exists a linear order $<$ on the set \mathbf{X} such that, if a node u_1 appears before a node u_2 in some run of \mathcal{D} , then, letting $X_1 \in \mathbf{X}$ be the label of u_1 and $X_2 \in \mathbf{X}$ the label of u_2 , we have $X_1 < X_2$. Notice that an nOBDD is in particular an nFBDD.

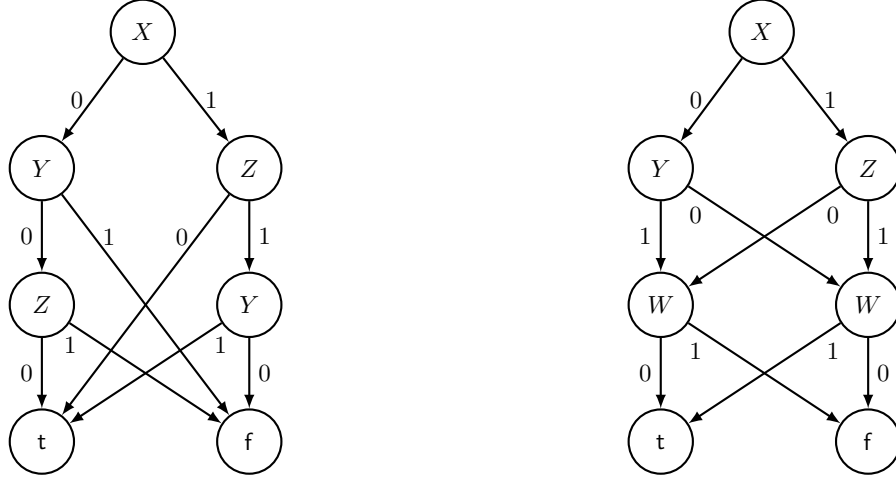
The notions of nFBDD and nOBDD are defined in terms of the runs of an nBDD. Thanks to condition (iii) of the definition of an nBDD, such notions can be equivalently defined in terms of the paths of an nBDD. In particular, it would be equivalent to say that an nBDD \mathcal{D} is free if for every (directed) path π in \mathcal{D} , no two distinct nodes in π have the same label, and likewise for the condition that the nBDD is ordered.

- **Ambiguity level.** Let \mathcal{D} be an nBDD over a set of variables \mathbf{X} . Then \mathcal{D} is *unambiguous* (uBDD) if, for every assignment $\mathbf{a} \in \text{Assign}(\mathbf{X})$, there exists at most one accepting run of \mathcal{D} that is consistent with \mathbf{a} . Moreover, \mathcal{D} is *deterministic*, which is referred to as BDD [Lee, 1959, Wegener, 2004]), if, for every assignment $\mathbf{a} \in \text{Assign}(\mathbf{X})$, there exists exactly one run of \mathcal{D} that is consistent with \mathbf{a} .

As mentioned before, thanks to condition (iii) of the definition of an nBDD, for every nBDD \mathcal{D} over a set of variables \mathbf{X} and every $\mathbf{a} \in \text{Assign}(\mathbf{X})$, there exists at least one run of \mathcal{D} that is consistent with \mathbf{a} . In the case where \mathcal{D} is deterministic, such a run must be unique. This leads to the following equivalent definition of a BDD, which is the most commonly used in the literature: \mathcal{D} is a BDD if and only if \mathcal{D} has a single source, which is called the *root* of \mathcal{D} , and every internal node of \mathcal{D} has exactly one outgoing 0-edge and exactly one outgoing 1-edge.

The combination of the previous two dimensions gives rises to 9 different classes of nondeterministic binary decision diagrams, which are shown in Table 1. The most widely used models among these classes are the deterministic variants, i.e., binary decision diagrams (BDD) [Lee, 1959], free binary decision diagrams (FBDD) [Fortune et al., 1978, Blum et al., 1980], and ordered binary decision diagrams (OBDD) [Bryant, 1986]. An FBDD is also referred to as a *read-once branching program* in the literature, where the term *nondeterministic read-once branching program* [Razgon, 2014] is used to denote¹ nFBDDs.

¹Note that [Razgon, 2014] also defines a notion of nFBDD, or *normalized FBDD*, which is different from the nFBDDs that we consider.



(a) A free binary decision diagram.

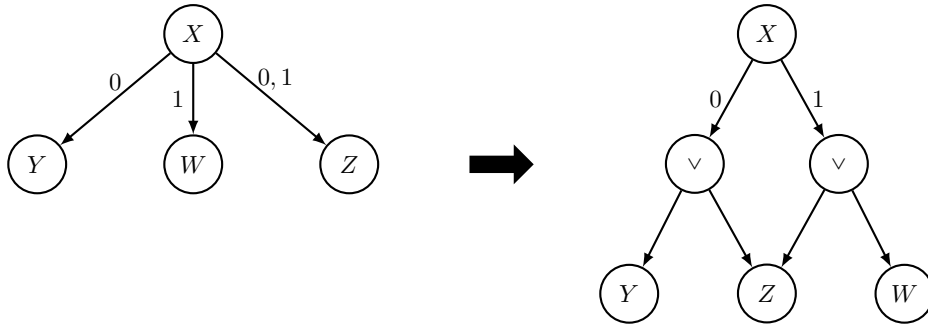
(b) An ordered binary decision diagram.

Figure 2: Two deterministic binary decision diagrams over the set of variables $\mathbf{X} = \{X, Y, Z, W\}$.

Example 3.2. The *nBDD* in Figure 1 is not free: indeed, for the run u_1, u_3, u_5, u_7 , the nodes u_1 and u_5 have the same label Y . On the other hand, the *nBDD* shown in Figure 2a is free and deterministic, so it is an *FBDD*. Moreover, the *nBDD* shown in Figure 2b is ordered and deterministic, because every run in it follows the linear order $X < Y < Z < W$, so it is an *OBDD*. The *nBDD* of Figure 2a, on the other hand, is not ordered.

3.2 Binary decision diagrams with or-nodes

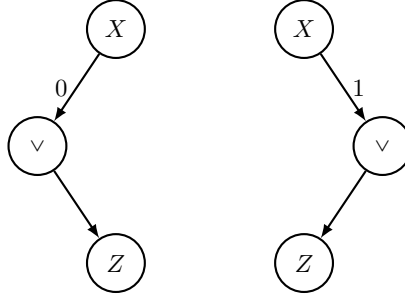
Nondeterministic binary decision diagrams have also been defined by extending binary decision diagrams with *or-nodes* [Amarilli et al., 2020, Bollig et al., 2010] (they are also called *guessing nodes* [Razborov, 1991]). An or-node u in a BDD is an internal node such that u is labeled with the disjunction symbol \vee instead of a variable, and the outgoing edges of u are not labeled. Acceptance of a BDD with or-nodes is defined as for the case of BDDs; in particular, if a run goes through an or-node, then it chooses an outgoing edge of this or-node, which does not impose any restriction on the values assigned to variables. In this sense, or-nodes can be used to encode nondeterministic choices as shown in the following example:



Part of an *nBDD* is shown in the left-hand side of this figure, while its representation as a BDD with or-nodes is shown in the right-hand side. In particular, if the variable X is assigned value 0 in

the nBDD, then a run can move either to the node with label Y or to the node with label Z . Such a choice is represented in the BDD by connecting the outgoing 0-edge of variable X to an or-node, which in turn is connected (by means of edges without labels) to the nodes with labels Y and Z ; in this way, we indicate that if X is assigned value 0 in a run of the BDD, then this run must move to an or-node, from which it must choose whether to move either to variable Y or to variable Z . It is straightforward to see that this idea can be used to translate in polynomial time an nBDD into an equivalent BDD with or-nodes.

In the other direction, a BDD with or-nodes can be translated in polynomial time into an equivalent nBDD by applying the transformation shown in Figure 3, which we explain next. Part of a BDD with or-nodes is shown in the left-hand side of this figure, while its representation as an nBDD is shown in the right-hand side. More precisely, for a pair of nodes u, v that are labeled by variables in a BDD with or-nodes \mathcal{D} , an *or-path* from u to v is a path π from u to v in \mathcal{D} such that every node in π except for u and v is an or-node. For example, the following are or-paths from the node X to the node Z in Figure 3:



Notice that the first edge in an or-path must be labeled 0 or 1; in the first case, the or-path is said to be a *0-or-path*, while in the second is said to be a *1-or-path*. Moreover, given a node u that is labeled by a variable, the *or-closure* of u is defined as the set of nodes v such that v is labeled by a variable and there exists an or-path from u to v . In the BDD with or-nodes in Figure 3, the or-closure of X consists of the nodes with labels Y, Z, W and V . Then in the translation of a BDD with or-nodes into an nBDD, for every node v in the or-closure of a node u , a 0-edge from u to v is included in the nBDD if there exists a 0-or-path from u to v , and a 1-edge from u to v is included in the nBDD if there exists a 1-or-path from u to v . An example of such a transformation is shown in Figure 3. It is straightforward to see that this idea can be used to translate in polynomial time any BDD with or-nodes \mathcal{D} into an equivalent nBDD \mathcal{D}' .² Observe that the complexity of that operation is essentially that of computing the transitive closure of the binary decision diagram, which can be done in polynomial time. Furthermore, there is no need to compute this transitive closure if no or-node is connected to another or-node: if we impose this condition on the input, then the translation becomes linear-time.

The different conditions on variable structuredness and ambiguity level for nBDDs can be directly extended to BDDs with or-nodes. Thus, for example, we can talk about FBDDs with or-nodes and OBDDs with or nodes. The procedures to transform nBDDs into BDDs with or-nodes and vice-versa preserve all such conditions. In this sense, up to the translation that we presented, the models of nBDDs and BDDs with or-nodes are completely interchangeable: we use nBDDs in the sequel. We last note that some restricted cases of BDDs with or-nodes have been considered in the literature, e.g., PBDDs, which are (nondeterministic) disjunctions of OBDDs with different orders [Bollig and Wegener, 1997]. Again, these models can be represented within the framework based on nBDDs that is used in this paper.

²Note that, if the root of \mathcal{D} is an or-node, then \mathcal{D}' will have multiple sources.

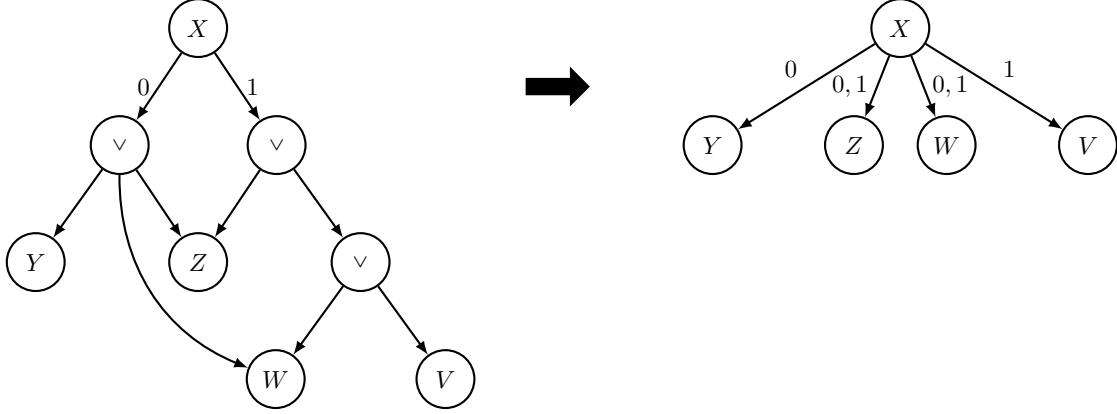


Figure 3: Transformation of a BDD with or-nodes into an nBDD.

3.3 Completion and zero-suppressed semantics

We present in this section the *completion* transformation on BDDs, and the alternative semantics of BDDs called the *zero-suppressed semantics*.

Completion. The notion of *complete* nBDDs has been studied, e.g., in [Bolig, 2016]. An nBDD is *complete* if all variables are tested on all runs. More precisely, for each run $\pi = u_1, u_2, \dots, u_k$, the set of variables that occur as labels of $\{u_1, u_2, \dots, u_k\}$ is equal to the set \mathbf{X} of all variables. In particular, if the binary decision diagram is free, then every such path must consist of exactly $|\mathbf{X}|$ internal nodes followed by a sink. Further, if the binary decision diagram is ordered, the sequence of variables tested by every run is exactly the linear order on the variables that the binary decision diagram follows.

Completeness can be useful for certain tasks. Consider for instance the counting problem for FBDDs: the input is a FBDD \mathcal{D} over a set of variables \mathbf{X} , and the output is $\#\text{Models}(\mathcal{D})$, i.e., the number of assignments of $\text{Assign}(\mathbf{X})$ that satisfy \mathcal{D} . If \mathcal{D} is complete, then the following linear-time bottom-up algorithm is correct. Annotate all t-sinks with 1 and all 0-sinks with 0. Then for an internal node n , annotate it by the sum of the annotations of the nodes that can be reached from n by following only one edge, then output the sum of the annotations of all source nodes. If \mathcal{D} was not complete, this would not return the correct result, as this might have missed some satisfying assignments.

We can always complete an nBDD in polynomial time by adding extra internal nodes before sinks. More precisely, assuming $\mathbf{X} = \{X_1, \dots, X_k\}$, we replace each sink u by a sequence of nodes u_1, \dots, u_k, u_{k+1} where the label of u_i is X_i for every $i \in [1, k]$, the label of u_{k+1} is the same as the label of u , and there is a 0-edge and a 1-edge from u_i to u_{i+1} for every $i \in [1, k]$. Notice that this translation does not affect unambiguity or determinism; however, the resulting nBDD is in general not free or ordered. The complexity of the procedure is $O(|\mathcal{D}| \cdot |\mathbf{X}|)$, where $|\mathcal{D}|$ is the size of the input nBDD \mathcal{D} and \mathbf{X} is the set of variables.

More interestingly, we can always complete an nFBDD in polynomial time while ensuring that the result is still an nFBDD, as shown in [Wegener, 2000] (see also [Amarilli et al., 2020]). To do this, we rewrite the nFBDD from the sources to the sinks in polynomial time while ensuring that, for every node u , all paths from a source to u test the same set S_u of variables, and every variable in S_u is tested exactly once in each such path. The base case is that of a source u labeled

with variable X which tests the set $S_u = \{X\}$. Inductively, considering an internal node u labeled with X with incoming edges from nodes u_1, \dots, u_ℓ testing sets $S_{u_1}, \dots, S_{u_\ell}$, and letting $S = \bigcup_i S_{u_i}$, we replace each edge from u_i to u by a sequence of nodes like in the previous paragraph which test the variables of the (possibly empty) set $S \setminus S_i$. By inductive assumption, all paths reaching the end of such a sequence are testing precisely S . As the binary decision diagram is an nFBDD, we know that the label X of node u does not belong to S_i for every $i \in [1, \ell]$. Hence, $x \notin S$ and all paths reaching node u after the rewriting are indeed testing all variables of $S_u = S \cup \{X\}$ precisely once. Last, for sinks, we proceed as in the previous paragraph to make sure that any remaining variables are tested. Note that this translation again does not affect unambiguity or determinism, and the result is complete and is still free, and again the complexity is $O(|\mathcal{D}| \cdot |\mathbf{X}|)$ for \mathcal{D} the input nFBDD and \mathbf{X} its set of variables.

We last observe that this process on nFBDDs, when applied to nOBDDs, can be performed in a way that gives an nOBDD as a result, if the sequences of variables that we insert always follow the order $<$ of the nOBDD. For this we can observe that the sets S_n of the previous proofs are always a prefix of the order $<$. Again the transformation does not affect unambiguity or determinism, and the complexity is the same. In this specific case, a simpler description of this algorithm is then the following. Assume that the nOBDD conforms to the linear order $X_1 < X_2 < \dots < X_k$. In order to complete it, we can replace each edge $X_i \xrightarrow{a} X_j$ by the sequence of edges $X_i \xrightarrow{a} X_{i+1} \xrightarrow{0,1} X_{i+2} \xrightarrow{0,1} \dots \xrightarrow{0,1} X_j$, where the intermediate steps represent fresh nodes that test the indicated variables.

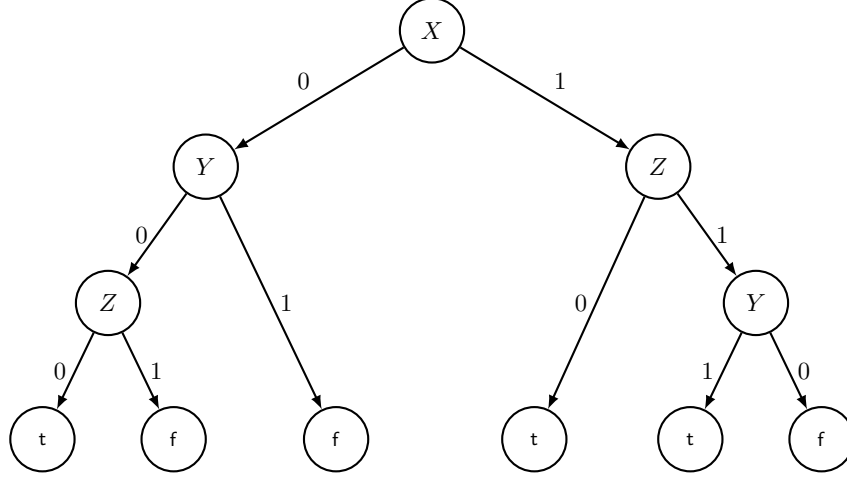
Zero-suppressed semantics. An alternative way to define the semantics of an nBDD is via the so-called *zero-suppressed semantics*, see [Wegener, 2000, Section 8.1] or [Minato, 1993]. Intuitively, in a zero-suppressed nBDD, it is assumed that every variable not mentioned along a run takes the value 0, unlike the standard semantics presented so far where the value of such variables is unconstrained. More formally, given a set of variables \mathbf{X} , a zero-suppressed nBDD \mathcal{D} over \mathbf{X} is defined like an nBDD, but with the following modification in the definition of the acceptance condition for \mathcal{D} . Let $\mathbf{a} \in \text{Assign}(\mathbf{X})$ and $\pi = u_1, u_2, \dots, u_k$ be a path from a source to a sink of \mathcal{D} , and let X_i be the label of u_i for each $i \in [1, k-1]$. Then π is said to be *a run consistent with \mathbf{a}* if there exists an edge from u_i to u_{i+1} with label $\mathbf{a}(X_i)$ for each $i \in [1, k-1]$, and $\mathbf{a}(X) = 0$ for each variable $X \in \mathbf{X} \setminus \{X_1, \dots, X_{k-1}\}$.

The purpose of zero-suppressed semantics is that, in practice, it can yield smaller binary decision diagrams. For example, a Boolean function over a set of variables \mathbf{X} with the only satisfying assignment being the one that assigns 0 to all variables can be represented in the zero-suppressed semantics simply as a single source node that is also a t-sink. In contrast, using the usual semantics requires $|\mathbf{X}|$ nodes. However, as we explain next, there are simple polynomial-time conversions between the two semantics.

For an nBDD that is complete, given that all runs must test every variable, there is no difference between the standard semantics and the zero-suppressed semantics. For this reason, every nBDD can be converted in polynomial time to an equivalent zero-suppressed nBDD, which preserves the variable structuredness and ambiguity level of the nBDD being translated. Conversely, given a zero-suppressed nBDD \mathcal{D} , it can be transformed into an equivalent nBDD by considering a variant of the completion procedure presented before: in that procedure, we only add nodes that check that the missing variables are mapped to 0. This procedure makes \mathcal{D} complete without changing the function that it represents in the zero-suppressed semantics. In addition, it preserves the variable structuredness and ambiguity level of the zero-suppressed nBDD being translated.

3.4 Decision trees

If we disallow sharing (i.e., we disallow internal nodes having multiple incoming edges), we get decision trees. More precisely, an nBDD \mathcal{D} is a *nondeterministic decision forest* (nDF) if, ignoring the sinks, the underlying graph of \mathcal{D} is a tree. If, in addition, it has only one source node, then it is a *nondeterministic decision tree* (nDT). This notion of *nondeterministic decision trees* has already been studied, e.g., in [Wegener, 2000, Theorem 10.1.4]. For example, the nBDD in Figure 2a is an nDT, which is usually depicted in the following way, implicitly indicating that leaves are not considered in the underlying structure that defines an nDT:



On the other hand, the nBDD in Figure 2b is not an nDT.

As a special case of nBDDs, nDFs and nDTs inherit property definitions such as variable structuredness and ambiguity levels. As with binary decision diagrams, the most widely used subclass are the nDTs that are deterministic, which are usually referred to simply as *decision trees* (DT) [Safavian and Landgrebe, 1991, Wu et al., 2008]. For example, the nBDD in Figure 2a is a DT, which can be easily verified considering the previous figure.

It is important to mention that nDFs and DTs are usually assumed to be free because, unlike general binary decision diagrams, an nDF can always be transformed in polynomial time into an equivalent free nDF (and likewise for DTs). More precisely, for every node u of an nDF \mathcal{D} , there exists a unique path from the corresponding source of \mathcal{D} to u . The labels of the vertices and edges traversed on this path specify a partial assignment for a set of variables. If the node u tests a variable X that was already tested on this path, then the value a of X is already specified in the partial assignment. Thus, we can remove u , re-connect the outgoing edges from u labeled by a to the predecessor of u in the path, and remove the other outgoing edges from u (which are labeled $1 - a$). Repeating this process in a traversal of \mathcal{D} makes this nDF free. Observe that this process does not work with general binary decision diagrams as there may be different paths reaching a node u with conflicting values for the variable tested by u .

4 Boolean Circuits

We now move from binary decision diagrams to the more general formalism of Boolean circuits. The section is structured similarly to the previous section: we give the basic definitions of a Boolean circuit and its semantics, and define notions of structuredness and ambiguity levels. We then mention the definition of Sentential Decision Diagrams (SDDs) as a special case of Boolean circuits.

We then mention the case of Boolean circuits without sharing (aka formulas), and the notion of smoothness for Boolean circuits (corresponding to completeness for binary decision diagrams).

4.1 Definitions

A *circuit* C over variables \mathbf{X} is a rooted $\{\wedge, \vee, \neg, \mathbf{t}, \mathbf{f}\} \cup \mathbf{X}$ -node labeled DAG $C = (N, W, \lambda)$ together with a designated sink $g_0 \in N$ called the *output gate* of C . The vertices N are called *gates*, the edges W are called *wires*. We often abuse notation and write $g \in C$ to mean $g \in N$. An *input* of a gate $g \in C$ is a gate $g' \in C$ that has a wire to g , i.e., we have $(g', g) \in W$. The gates of C can be of several kinds:

- If $\lambda(g) = \mathbf{t}$, then g is a *constant true-gate* (also sometimes called *constant 1-gate*), and it must then have no input;
- If $\lambda(g) = \mathbf{f}$, then g is a *constant false-gate* (also sometimes called *constant 0-gate*), and it must then have no input;
- If $\lambda(g) \in \mathbf{X}$, then g is a *variable gate*, and it must then have no input;
- If $\lambda(g) = \neg$, then g is a *negation gate*, and it must then have exactly one input;
- If $\lambda(g) = \vee$ (resp., $\lambda(g) = \wedge$), then g is a \vee -gate (resp., \wedge -gate), and it must then have at least one input.

The Boolean circuit C represents a Boolean function over \mathbf{X} in the following way. Given an assignment $\mathbf{a}: \mathbf{X} \rightarrow \{0, 1\}$, we extend it to give a Boolean value to all gates of the Boolean circuit by bottom-up induction:

- The value of a constant true-gate g is $\mathbf{a}(g) = 1$;
- The value of a constant false-gate g is $\mathbf{a}(g) = 0$;
- The value of a variable gate g annotated with variable X is $\mathbf{a}(g) := \mathbf{a}(X)$;
- The value of a negation gate g is $\mathbf{a}(g) := 1 - \mathbf{a}(g')$, where g' is the input of g ;
- The value of an \wedge -gate g (resp., \vee -gate g) with input gates g_1, \dots, g_n is $\mathbf{a}(g) := \bigwedge_i \mathbf{a}(g_i)$ (resp., $\mathbf{a}(g) := \bigvee_i \mathbf{a}(g_i)$).

The Boolean function defined by C is then the one that maps assignments \mathbf{a} to the value $\mathbf{a}(g_0)$ of the output gate of the Boolean circuit. We sometimes abuse terminology and call the variable gates the *inputs* of the Boolean circuit.

The size $|C|$ of a Boolean circuit is its number of wires. For a gate g of C , we denote by $\text{Vars}(g)$ the set of variables that have a directed path to g in C , and we denote by C_g the Boolean circuit over \mathbf{X} whose output gate is g .

In the rest of this document, we will always consider Boolean circuits in *negation normal form* (NNF), where negation gates are always applied to variables; formally, for any negation gate g , then its one input must be a variable gate. We can equivalently see NNF Boolean circuits as positive Boolean circuits (circuits without negation) defined directly on the literals (i.e., the variables and their negations).

Just like binary decision diagrams, we classify Boolean circuits according to two dimensions.

Table 2: Classification of NNF Boolean circuits based on ambiguity level (nondeterministic, unambiguous, or deterministic) and variable structuredness (unrestricted, free, or ordered).

	Unrestricted	Decomposable	Structured
Arbitrary	NNF	DNNF	SDNNF
Deterministic (i.e., unambiguous)	d-NNF	d-DNNF	d-SDNNF
Decision	dec-NNF	dec-DNNF	dec-SDNNF

- **Variable structuredness.** A Boolean circuit is called *decomposable* if, intuitively, \wedge -gates partition the variables into disjoint sets. Formally, an \wedge -gate g of C is *decomposable* if it has exactly two input gates $g_1 \neq g_2$ such that we have $\text{Vars}(g_1) \cap \text{Vars}(g_2) = \emptyset$. A Boolean circuit C is *decomposable* if every \wedge -gate of C is. Note that this is a syntactic condition that can easily be checked in time $O(|C| \cdot |\mathbf{X}|)$. A decomposable NNF Boolean circuit is called a *DNNF*. We point out that DNNFs are sometimes defined without the restriction that \wedge -nodes always have two inputs; we impose this for convenience, and this is without much loss of generality as this can be enforced in linear time (using constant gates).

Further, a DNNF is *structured* if the partitions that are defined by the \wedge -gates are compatible, in the following sense. A *v-tree* over the set of variables \mathbf{X} is a rooted full binary tree T whose leaves are in bijection with \mathbf{X} . We always identify each leaf with the associated element of \mathbf{X} . For a node $n \in T$, we abuse notation and denote by $\text{Vars}(n)$ the set of variables in the subtree rooted at n . A DNNF D is *structured by the v-tree T* if there exists a mapping ρ labeling each \wedge -gate of D with a node of T that satisfies the following: for every \wedge -gate g of D with two inputs g_1, g_2 , the node $\rho(g)$ *structures* g , i.e., $\rho(g)$ is not a leaf and, letting n_1 and n_2 be its two children in some order, we have $\text{Vars}(g_i) \subseteq \text{Vars}(T_{n_i})$ for $i = 1, 2$. A DNNF is *structured*, written *SDNNF*, if there exists a v-tree that structures it.

- **Ambiguity level.** A Boolean circuit is *unambiguous*, which is (unfortunately) called for historical reasons *deterministic*, if, intuitively, the inputs to \vee -gates are mutually exclusive. Formally, an \vee -gate g of C is *deterministic* if for every pair $g_1 \neq g_2$ of input gates of g , the Boolean functions over \mathbf{X} captured by C_{g_1} and C_{g_2} are disjoint; that is, we have $\text{Models}(C_{g_1}) \cap \text{Models}(C_{g_2}) = \emptyset$. We call C *deterministic* if each \vee -gate is. Note that being deterministic is a semantic condition, not a syntactic one.

A stronger notion of determinism is that of determinism in the sense of BDDs, which we call *decision* to distinguish it from determinism. An \vee -gate is *decision* if it has exactly two inputs g_0 and g_1 and there is a variable X such that g_1 is an \wedge -gate with X as input, and g_0 is an \wedge -gate with a negation gate of X as input. The Boolean circuit C is *decision* if all \vee -gates are decision. Note that being decision is a syntactic condition that can be checked in linear time, and a decision circuit is always deterministic.

The combination of the previous two dimensions again gives rise to 9 different classes of Boolean circuits, which are shown in Table 2. A Boolean circuit which is both decomposable and deterministic is called a *d-DNNF*, and if it is in addition structured then we have a *d-SDNNF*. A Boolean circuit which is both decomposable and decision is called a *dec-DNNF*, and *dec-SDNNF* if it is structured. Note that dec-DNNFs are also called *decision-DNNFs* [Lagniez and Marquis, 2017]; they can also be seen as *AND-FBDDs* with decomposable ANDs [Beame et al., 2013].

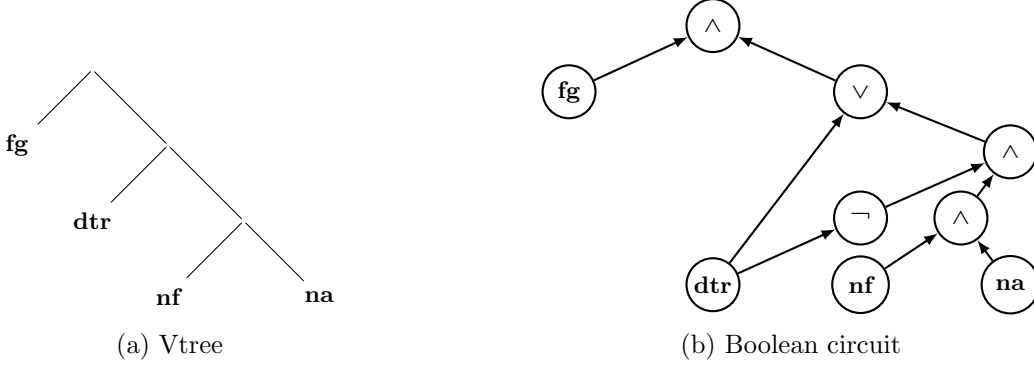


Figure 4: A deterministic and decomposable Boolean circuit as a classifier and its vtree.

Last, we point out that ambiguity levels are typically only considered in combination with decomposability, i.e., there is no standard notation for a deterministic NNF such as *d-NNF* (or *dec-NNF*). This is because, to the best of our knowledge, there is no interesting task that can be tractably solved specifically on such Boolean circuits.

Example 4.1. *The following example is taken from [Arenas et al., 2023]. We want to classify papers submitted to a conference as rejected (Boolean value 0) or accepted (Boolean value 1). Papers are described by Boolean variables **fg**, **dtr**, **nf** and **na**, which stand for “follows guidelines”, “deep theoretical result”, “new framework” and “nice applications”, respectively. The Boolean classifier for the papers is given by the Boolean circuit in Figure 4. The input of this Boolean circuit are the variables **fg**, **dtr**, **nf** and **na**, each of which can take value either 0 or 1, depending on whether the variable is present (1) or absent (0). The nodes with labels \neg , \vee or \wedge are logic gates, and the associated Boolean value of each one of them depends on the logical connective represented by its label and the Boolean values of its inputs. The output value of the Boolean circuit is given by the top node in the figure.*

*The Boolean circuit in Figure 4 is decomposable, because each \wedge -gate has two inputs, and the sets of variables of its inputs are pairwise disjoint. For instance, in the case of the top node in Figure 4, the left-hand side input has $\{\mathbf{fg}\}$ as its set of variables, while its right-hand side input has $\{\mathbf{dtr}, \mathbf{nf}, \mathbf{na}\}$ as its set of variables, which are disjoint. Also, this Boolean circuit is deterministic as for every \vee -gate two of its inputs cannot be given value 1 by the same Boolean assignment for the variables. For instance, in the case of the only \vee -gate in Figure 4, if a Boolean assignment for the variables gives value 1 to its left-hand side input, then variable **dtr** has to be given value 1 and, thus, such an assignment gives value 0 to the right-hand side input of the \vee -gate. In the same way, it can be seen that if a Boolean assignment for the variables gives value 1 to the right-hand side input of this \vee -gate, then it gives value 0 to its left-hand side input.*

*Last note that the Boolean circuit is structured by the vtree shown at the left. Further, the circuit is not decision; however, it could be made decision by adding a \wedge -gate having as inputs the variable gate **dtr** and a constant 1-gate, and making that gate an input to the \vee -gate instead of **dtr**.*

4.2 Sentential decision diagrams

Sentential decision diagrams (SDDs) are a restricted class of structured Boolean circuits, satisfying a new ambiguity level called *strong determinism*. An \vee -gate g is said to be strongly deterministic if it satisfies the following criteria. Firstly, the inputs of g are all \wedge -gates, written g_1, \dots, g_m , and they are all structured by some v-tree node n , i.e., we have $\rho(g_i) = n$ for all $1 \leq i \leq m$. Further,

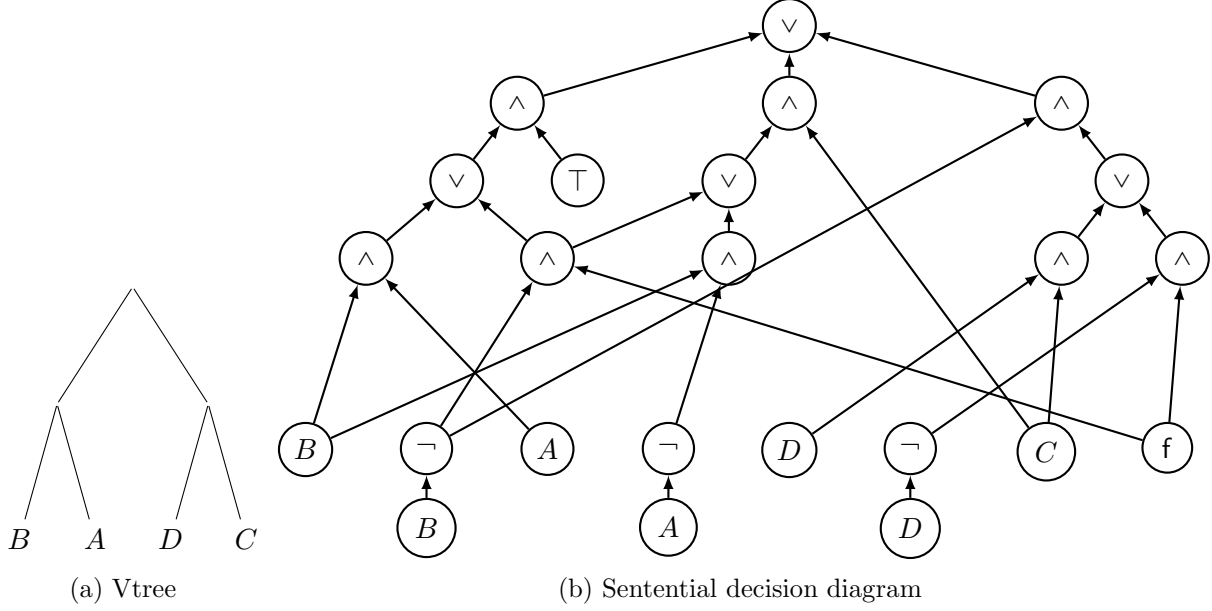


Figure 5: An example of a sentential decision diagram (SDD) and its vtree.

each g_i for $1 \leq i \leq m$ has precisely two inputs p_i and s_i , and, letting n_1, n_2 be the children of n in the order given by the v-tree, we have $\text{Vars}(p_i) \subseteq \text{Vars}(n_1)$ and $\text{Vars}(s_i) \subseteq \text{Vars}(n_2)$. Then we call p_1, \dots, p_m the *primes* and s_1, \dots, s_m the *subs* of g , and we require that the primes p_i, \dots, p_m are mutually exclusive, namely, we require that $p_i \wedge p_j$ is unsatisfiable for any $i \neq j$.

A Boolean circuit C is then *strongly deterministic* if every \vee -gate of C is. This gives rise to the class of *strongly deterministic SDNNF* [Pipatsrisawat and Darwiche, 2010], which, as the name suggests, satisfies a stronger notion of determinism than a d-SDNNF while being more general than dec-SDNNF. A *sentential decision diagram (SDD)* [Darwiche, 2011] is a strongly deterministic DNNF which further ensures that, for every \vee -gate g , letting p_1, \dots, p_m be its primes, then they are *exhaustive*, formally, $\bigvee_i p_i = \top$.

Example 4.2. In Figure 5b, we show an example of a sentential decision diagram from [Darwiche, 2011], in circuit notation. The inputs of this Boolean circuit are the variables A, B, C, D . It can be seen that structuredness holds with respect to the shown vtree; for example, the \wedge -node immediately below the root \vee -node has one child with variables $\{B, A\} \subseteq \{B, A\}$ and one child with variables $\{C\} \subseteq \{D, C\}$. Strong determinism also holds; for example, for the root \vee -node, its children are all \wedge -nodes structured by the root vtree node, and the primes correspond to the logical formulae $B \wedge A$, $B \wedge \neg A$ and $\neg B$, which are clearly mutually exclusive (and also exhaustive).

4.3 Formulas

As for binary decision diagrams and their restriction to decision trees, one can also be interested in circuit classes in which we disallow sharing.

When the underlying graph of a Boolean circuit is a tree (noting that we may have multiple variable gates labeled by the same variable), then we call the result a *Boolean formula* (not to be confused with a Boolean function). If the Boolean circuit is in NNF, then the Boolean formula is in NNF. Examples of Boolean formulas in NNF include Boolean formulas in conjunctive normal form (CNF), and Boolean formulas in disjunctive normal form (DNF). Note that Boolean formulas

may still contain several occurrences of the same variable or literal; when we further impose that the Boolean circuit has a tree structure and that there is only one gate for each variable, we obtain what is called a *read-once formula* [Angluin et al., 1993].

4.4 Smoothness

Recall that, for binary decision diagrams, we introduced a notion of *completeness* which intuitively requires binary decision diagrams to test all variables. We now present the analogous requirement on circuits, which is called *smoothness*.

The notion of *smooth* Boolean circuits has been introduced in [Darwiche, 2001]; see also [Shih et al., 2019]. An \vee -gate g of a Boolean circuit C is *smooth* if for every input g' of g we have $\text{Vars}(g) = \text{Vars}(g')$, and C is called *smooth* if all its \vee -gates are smooth. Smoothness is the Boolean circuit analogue of *completeness* for binary decision diagrams. *Smoothing* is the process of taking a Boolean circuit C as input, and constructing a Boolean circuit C' that is equivalent to C and that is smooth. This can be done naïvely in quadratic time as follows. Letting \mathbf{X} be the variables, compute in time $O(|C| \cdot |\mathbf{X}|)$ the set $\text{Vars}(g)$ for every $g \in \mathbf{X}$. Then, for every \vee -gate g and input gate g' such that $\text{Vars}(g') \subsetneq \text{Vars}(g)$, replace g' by an \wedge -gate of g' and of $\bigwedge_{X \in \text{Vars}(g) \setminus \text{Vars}(g')} (X \vee \neg X)$. It is clear that the resulting Boolean circuit is equivalent and smooth, and that this can be done in $O(|C| \cdot |\mathbf{X}|)$. Furthermore, if C is in NNF, then so is C' , and if C is decomposable, then C' also is. We note that smoothing can be done more efficiently for structured Boolean circuit classes [Shih et al., 2019].

5 From Binary Decision Diagrams to Boolean Circuits

In this section, we connect the notions of binary decision diagrams (Section 3) and Boolean circuits (Section 4). We do so by explaining why binary decision diagrams are in fact specific classes of Boolean circuits, with a direct translation from binary decision diagrams to Boolean circuits. Furthermore, we show how the conditions that we have defined on diagrams can be preserved by this transformation.

Linear v-trees. The translation from binary decision diagrams to Boolean circuits that we will present, when applied to ordered binary decision diagrams (nOBDDs), will produce structured Boolean circuits whose v-trees have a specific shape. We define these v-trees, called *right-linear v-trees*, and explain how to obtain them from the order respected by the nOBDD.

Formally, given an order $<$ on variables \mathbf{X} , we define the *right-linear vtree* $T_{<}$ obtained from $<$ in the following way:

- if \mathbf{X} consists of a single variable x then $T_{<}$ is the singleton tree whose root is a leaf node labeled by x ;
- otherwise, letting x be the smallest element of \mathbf{X} according to $<$, we let $T_{<}$ be the tree whose root is an internal node having x as left child and having as right child the root of the tree $T_{<'}$ obtained from the order $<'$ which is the restriction of $<$ to $\mathbf{X} \setminus \{x\}$.

Note that the choice of right-linear v-trees is arbitrary; similar constructions could be made with a left-linear v-tree.

Converting binary decision diagrams to Boolean circuits. We are now ready to state the immediate translation from binary decision diagrams to Boolean circuits, and to explain how conditions on the Boolean circuit can be rephrased to conditions on the binary decision diagram. This is a folklore observation which has previously appeared, e.g., as Proposition 3.4 of [Amarilli et al., 2020]:

Proposition 5.1. *Given a $nBDD$ \mathcal{D} , we can convert it in linear time to a Boolean circuit C that describes the same Boolean function. The translation preserves the following:*

- *Variable structuredness:*
 - *If the input $nBDD$ is free, then the resulting Boolean circuit is decomposable (DNMF).*
 - *If the input $nBDD$ is ordered with an order $<$, then the resulting Boolean circuit is structured with the right-linear vtree $T_{<}$.*
- *Ambiguity level:*
 - *If the input $nBDD$ is unambiguous then the resulting Boolean circuit is deterministic.*
 - *If the input $nBDD$ is deterministic then the resulting Boolean circuit is decision.*
- *Other conditions:*
 - *If the input is a decision tree, then the output is a Boolean formula.*
 - *If the input is complete, then the result is smooth.*

Proof. We use the following general linear-time translation from binary decision diagrams to Boolean circuits, building a Boolean circuit C from the input binary decision diagram \mathcal{D} :

- We replace true sinks and false sinks respectively by constant true and false gates.
- We replace internal nodes n on a variable X with a gate defined like in the definition of decision gates above; formally, let $g_0^1, \dots, g_0^{k_0}$ (resp., $g_1^1, \dots, g_1^{k_1}$) be the translations of the nodes to which n had edges labeled with 0 (resp., with 1). Construct a gate g_0 (resp., g_1) to be an \vee -gate of the gates $g_0^1, \dots, g_0^{k_0}$ (resp., of $g_1^1, \dots, g_1^{k_1}$). We then translate n to a \vee -gate whose inputs are an \wedge -gate conjoining g_1 and X , and an \wedge -gate conjoining g_0 and $\neg X$. (If g_0 has only one input gate, then we replace it by that input, and likewise for g_1 .)
- Last, the output (root) gate of the Boolean circuit is an \vee -gate taking the disjunction of the translations of all the sources of the $nBDD$; or, if the $nBDD$ has only one source, then it is the gate that translates this source.

The translation process is illustrated in Figure 6. One can check that the translation runs in linear time and produces a Boolean circuit C with the same semantics as the $nBDD$ \mathcal{D} , i.e., C that represents the same Boolean function as \mathcal{D} . Further, we can check that C has the stated properties:

- If the input $nBDD$ \mathcal{D} is free, for every \wedge -gate g in C created when translating a node n of the $nBDD$, then g will conjoin a literal for the variable x tested by n with a gate g' , and the gates g'' having a directed path to g' are gates g'' that are translations of nodes n' of \mathcal{D} to which n has a path (along with intermediary gates introduced in the translation), so these gates g'' cannot test x because \mathcal{D} is free. Hence, g is decomposable.

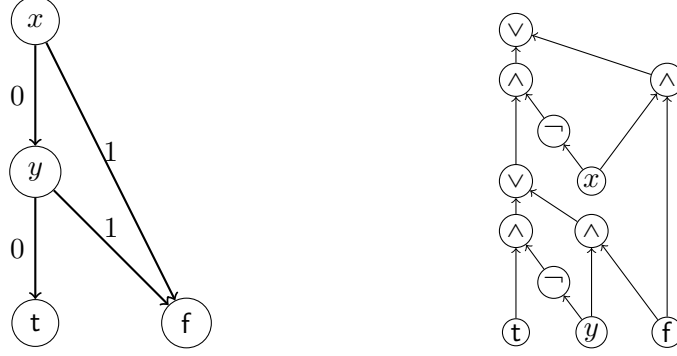


Figure 6: Left: OBDD for the Boolean function $\neg x \wedge \neg y$. Right: equivalent dec-SDNNF.

- If the input nBDD \mathcal{D} is ordered by an order $<$, then any \wedge -gate g in C created when translating a node n of \mathcal{D} will conjoin a literal for the variable x tested by n with a gate that depends on variables tested by nodes to which n has a directed path, therefore, C is structured by the right-linear \vee -tree $T_{<}$.
- If the input nBDD \mathcal{D} is unambiguous, then, for every node n of \mathcal{D} , for every assignment \mathbf{a} of the variables there is at most one accepting run which is compatible with \mathbf{a} and that goes through n . This ensures that, considering the \vee -gates g_0 and g_1 created in C when translating n , there cannot be two inputs of such a gate that are made true by \mathbf{a} . Further, for every assignment, we know that at most one source of \mathcal{D} is the beginning of a successful run, so the \vee -gate which is the output gate of C also does not have two mutually satisfiable inputs. (Note that the \vee -gates created as the translation of n above are decision gates, so they are always deterministic.)
- If the input nBDD \mathcal{D} is deterministic, then the gates of the form g_0 and g_1 in the translation had only one input, so they were merged with that input; and likewise \mathcal{D} has only one source so we did not create an \vee -gate as the output gate of C . Hence, in this case, all \vee -gates in C are those that translate a node of \mathcal{D} , and they are decision gates.
- If the input \mathcal{D} is a decision tree, then there is no sharing in the process described above (creating different copies of variable gates labeled by the same variable), so it creates a Boolean formula.
- If the input \mathcal{D} is complete, then an easy induction shows that the set of variables having a directed path to a gate g of the Boolean circuit created to translate a node n of \mathcal{D} is precisely the set of variables tested by the nodes n' of \mathcal{D} to which n have a directed path, so as \mathcal{D} is complete we conclude that C is smooth. \square

6 Automata

This section presents the notions of *word automata* and *tree automata*. It then explains how to relate these notions to binary decision diagrams and Boolean circuits, respectively, via the notion of *provenance circuits* for automata. The use of Boolean circuits for provenance representations is originally from [Deutsch et al., 2014], and its use for automata is in [Amarilli et al., 2015]; but the relation between word automata and binary decision diagrams has been studied in other contexts, e.g., [Bollig, 2016].

The section first defines word automata, then tree automata, and then explains how to compute provenance circuits for these automata.

6.1 Automata on words

We define standard notions from formal language theory, before defining word automata.

Alphabets, words, languages. An *alphabet* is a finite set Σ of *letters*. A *word* on Σ is a finite (possibly empty) sequence $\mathbf{w} = w_1, \dots, w_n$ of letters from Σ ; its *length* $|\mathbf{w}|$ is n . The set of all words over Σ is denoted by Σ^* . A *language* over Σ is a subset of Σ^* .

Word automata. A *non-deterministic finite automaton* (NFA) $A = (Q, I, F, \delta)$ over Σ consists of a finite set Q of *states*, a set $I \subseteq Q$ of *initial states*, a set $F \subseteq Q$ of *final states*, and a *transition relation* $\delta \subseteq Q \times \Sigma \times Q$. We define $|A|$, the *size* of A , to be $|A| := |\Sigma| + |Q| + |\delta|$. A *partial run* of $A on a word $\mathbf{w} \in \Sigma^*$ is a sequence of states $\rho = q_0, q_1, \dots, q_{|\mathbf{w}|}$ such that $(q_i, w_{i+1}, q_{i+1}) \in \delta$ for every $i \in \{0, \dots, |\mathbf{w}| - 1\}$. We say that ρ *starts at* q_0 and *ends at* $q_{|\mathbf{w}|}$. A *run* of $A on \mathbf{w} is a partial run of A on \mathbf{w} which starts at an initial state. An *accepting run* of $A on \mathbf{w} is a run of A on \mathbf{w} which ends in a final state. For a word \mathbf{w} , we say that \mathbf{w} is *accepted by* A if there is an accepting run of A on \mathbf{w} . The *language accepted by* A , denoted $L(A)$, is the set of words over Σ that are accepted by A .$$$

The automaton A is called *unambiguous* (UFA) if for every word $\mathbf{w} \in \Sigma^*$ there exists at most one accepting run of A on \mathbf{w} . Note that this is a semantic condition, which is not straightforward to verify syntactically. By contrast, we say that the automaton A is *deterministic* (DFA) if it satisfies the following syntactic condition: there is precisely one initial state in I , and, for every state $q \in Q$ and letter $a \in \Sigma$, there is at most one state $q' \in Q$ such that the transition (q, a, q') is in δ . In this case, we can equivalently see δ as a partial function from $Q \times \Sigma$ to Q . Note that, for any deterministic automaton A and word \mathbf{w} , then there is at most one run of A on \mathbf{w} (accepting or not). Hence, a deterministic automaton is necessarily unambiguous, but the converse is not true: there are some unambiguous automata that are not deterministic.

We say that a word automaton A is *trimmed* if, for every state $q \in Q$, there is a word \mathbf{w} in the language of A such that q occurs in an accepting run ρ of A on \mathbf{w} . Notice that this is equivalent to saying that, for every state $q \in Q$, there exists a path in the underlying directed graph of the automaton from some initial state to q (we say that q is *accessible*), and a path from q to some final state (we say that q is *co-accessible*). Given a word automaton A , we can easily convert it in linear time to a trimmed automaton which is equivalent (i.e., that recognizes the same language), and this conversion preserves unambiguity and determinism.

Note that an unambiguous word automaton A which is trimmed satisfies in particular a stronger requirement (*): for any word \mathbf{w} , there cannot be two different runs of A on \mathbf{w} that lead to the same state q . Indeed, if this were to happen, then as q is co-accessible we could complete \mathbf{w} to a word on which A has two accepting runs, contradicting unambiguity. This property (*) is generally not satisfied if the automaton A is not trimmed.

6.2 Automata on trees

We now define the notions of tree languages and automata over trees.

Trees. Let Σ be an alphabet. A Σ -*tree* (T, λ) is a finite, rooted, ordered binary tree T (all internal nodes have exactly two children, and these are ordered) such that every node n of T is labeled by a

letter $\lambda(n) \in \Sigma$. We say that T is the *skeleton* of (T, λ) . We denote by $\mathcal{T}(\Sigma)$ the set of all Σ -trees. A *tree language* over Σ is a (potentially infinite) set of Σ -trees. The *complement* of a tree language L is the set of Σ -trees that are not in L .

Tree automata. We only consider *bottom-up* finite tree automata in this document. A *non-deterministic (bottom-up) finite tree automaton* (NFTA) $A = (Q, F, \Delta, \iota)$ over Σ consists of a finite set of *states* Q , a set $F \subseteq Q$ of final states, an *initialization relation* $\iota \subseteq \Sigma \times Q$, and a *transition relation* $\Delta \subseteq Q \times Q \times \Sigma \times Q$. We define $|A|$, the *size* of A , to be $|A| := |\Sigma| + |Q| + |\Delta| + |\iota|$. Let (T, λ) be a Σ -tree. A *run of A on (T, λ)* is a Q -tree (T, λ') having the same skeleton T and satisfying the following conditions:

- every leaf is labeled by a state given by applying the initialization relation to the label of that leaf: formally, for every leaf n of T we have that $(\lambda(n), \lambda'(n)) \in \iota$;
- every internal node is labeled by a state given by applying the transition relation to the two states labeling the two children: formally, for every internal node n of T , letting n_1, n_2 be the (ordered) children of n , we have that $(\lambda'(n_1), \lambda'(n_2), \lambda(n), \lambda'(n)) \in \Delta$.

We say that run (T, λ') *ends at state q* when $q = \lambda'(r)$ for r the root of T . The run (T, λ') is *accepting* if, letting n be the root of T , we have $\lambda'(n) \in F$. If there exists an accepting run of A on (T, λ) then we say that (T, λ) is accepted by A . The *language* $L(A)$ of A is the set of Σ -trees accepted by A .

The automaton A is called *unambiguous* (UFTA) if for any Σ -tree (T, λ) , there exists at most one accepting run of A on (T, λ) : again, this is a semantic criterion. We say that A is a *deterministic bottom-up finite tree automaton* (DFTA) if it satisfies the following syntactic criterion: (1) for every letter $a \in \Sigma$, there is at most one state q such that $(a, q) \in \iota$; and (2) for each pair of states $(q_1, q_2) \in Q$ and letter $a \in \Sigma$, there is at most one state q such that $(q_1, q_2, a, q) \in \delta$. For a deterministic automaton, we can equivalently see ι as a partial function from Σ to Q , and see δ as a partial function from $Q \times Q \times \Sigma$ to Q . Again, if A is deterministic, then for any Σ -tree (T, λ) there is at most one run of A on (T, λ) (accepting or not). Hence, again, a deterministic tree automaton is necessarily unambiguous, but the converse does not necessarily hold.

We say that A is *trimmed* if, for every state $q \in Q$, there is a Σ -tree (T, λ) in the language accepted by A and a run ρ of A on (T, λ) in which q appears. Given a tree automaton, we can convert it in linear time to an automaton which is equivalent (recognizes the same language) and is trimmed, and this conversion preserves unambiguity and determinism.

Note that a UFTA A which is trimmed satisfies again a stronger requirement (*): for every Σ -tree (T, λ) and state $q \in Q$ there is at most one run of A on (T, λ) that ends at q .

6.3 Computing Boolean circuits and binary decision diagrams from automata

We now explain how to compute Boolean circuits and binary decision diagrams from automata. We first give the construction for the most general formalisms possible (NFAs and NFTA), giving nOBDD and structured DNNF circuits respectively. Then, we observe how these constructions can apply to restricted automata (unambiguous and deterministic), and which conditions these ensure on the resulting binary decision diagrams and Boolean circuits.

We assume for now that the alphabet used by automata is $\Sigma = \{0, 1\}$; we will later explain why other alphabets can also be handled.

Definition 6.1. Let A be an NFA over alphabet $\Sigma = \{0, 1\}$, and let $n \in \mathbb{N}$. The provenance of A on n is a Boolean function $\phi_{A,n}$ over variables $\mathbf{X} = \{X_1, \dots, X_n\}$ defined in the following way:

for any assignment \mathbf{a} of \mathbf{X} , letting $w_{\mathbf{a}} = \mathbf{a}(X_1) \dots \mathbf{a}(X_n)$ be the corresponding word of length n over Σ , we have that \mathbf{a} satisfies $\phi_{A,n}$ iff A accepts $w_{\mathbf{a}}$. A provenance circuit (resp, provenance decision diagram) of A on n is then a Boolean circuit (resp, binary decision diagram) representing $\phi_{A,n}$.

Likewise, if A is a NFTA over alphabet $\Sigma = \{0, 1\}$ and T is an unlabeled full binary tree, the provenance of A on T is a Boolean function $\phi_{A,T}$ whose variables \mathbf{X} are the nodes of T and which is defined in the following way: for any assignment \mathbf{a} of \mathbf{X} , letting $T_{\mathbf{a}}$ be the Σ -tree obtained from T by labeling the nodes according to \mathbf{a} , we have that \mathbf{a} satisfies $\phi_{A,T}$ iff A accepts $T_{\mathbf{a}}$. A provenance circuit (resp., provenance decision diagram) of A on T is a Boolean circuit (resp., binary decision diagram) representing $\phi_{A,T}$.

We show in this section the following two results on word automata and tree automata:

Proposition 6.2. *Let A be an NFA over alphabet $\Sigma = \{0, 1\}$ and let $n \in \mathbb{N}$. We can build in time $O(|A| \times n)$ a provenance diagram C of A on n which is a complete n OBDD with variable order X_1, \dots, X_n . Further, if A is unambiguous then C is an uOBDD, and if A is deterministic then C is an OBDD.*

Proposition 6.3. *Let A be a NFTA over alphabet $\Sigma = \{0, 1\}$ and let T be an unlabeled full binary tree. We can build in time $O(|A| \times |T|)$ a provenance circuit C of A on T which is a smooth structured DNNF (smooth SDNNF), together with its vtree. Further, if A is unambiguous then C is a d -SDNNF.*

Note that, in Proposition 6.3, there is no discussion of the case where the input tree automaton is deterministic. We are not aware of a standard Boolean circuit class corresponding to deterministic automata, though a notion of *upwards-deterministic Boolean circuits* is introduced for that purpose in [Amarilli et al., 2017].

We briefly comment on the relationship of these results to [Amarilli et al., 2015]. In the latter work, the provenance for automata is defined on an alphabet consisting of a fixed part together with a Boolean annotation. For instance, for word automata, the alphabet is $\Sigma \times \{0, 1\}$, and the provenance is defined on an input word of Σ^* , to describe which of the Boolean annotations of the word are accepted. The results above, with alphabet $\{0, 1\}$, allow us to recapture this setting. Indeed, we can modify automata working on a larger alphabet to first read a binary representation of the letter in Σ followed by the Boolean annotation. Applying the results above, and fixing the inputs corresponding to letters to the intended values, gives us the provenance circuit in the sense of [Amarilli et al., 2015]. This uses the fact that the Boolean circuit and binary decision diagram classes that we consider are closed under the *conditioning* operation where we force a variable to be equal to a specific value.

We first prove Proposition 6.2 as a warm-up:

Proof of Proposition 6.2. We assume that the input automaton A is *complete* in the sense that, for every state q and every letter $b \in \Sigma$, there is at least one transition for letter b on state q . We can clearly make A complete in linear time up to adding a sink state.

We build a binary decision diagram with decision nodes $g_{i,q}$ for each $1 \leq i \leq n$ and for each state q : the node $g_{i,q}$ is labeled with the variable X_i . The initial nodes are g_{1,q_0} for each initial state $q_0 \in I$. We also have sinks $g_{n+1,q}$ for each state q : the sink $g_{n+1,q}$ is a 0-sink if $q \notin F$, and it is a 1-sink if $q \in F$.

Now, for every $1 \leq i \leq n$, we add the following edges: for each transition from a state q to state q' when reading symbol $b \in \Sigma$, we add an edge to $g_{i,q}$ labeled b to $g_{i+1,q'}$.

The construction satisfies the time bound. We modify the result in linear time to trim it, removing all nodes that do not have a path from a starting node. The result is an nOBDD, because each decision node has at least one outgoing 0-edge and 1-edge, and the variables are ordered in the right way. Note that it is also complete in the sense that every run tests all variables.

We show correctness by finite induction: for every $0 \leq i \leq n$, given a word w of length i , the reachable nodes $g_{i+1,q}$ by the partial valuation corresponding to w are those for the states q that we can reach when reading w . The base case is immediate: when reading the empty word, we can get precisely to the initial state. Now, for the induction step, if the states that we can reach when reading a word of length i are correct, then when reading an extra letter $b \in \Sigma$ we can go precisely to the states having a transition labeled b from the previously reachable states, so the result is correct. The finite induction applied to $i = n$ confirms that we can reach a final state (i.e., the word is accepted) iff we can reach a 1-sink.

Now, observe that if the automaton is unambiguous, then for any valuation there is at most one way to reach the 1-sink, because there is at most one accepting run of the automaton on the corresponding word. Further, if the automaton is deterministic, then for any valuation there is at most one consistent path, because each decision node has at most one outgoing 0-edge and 1-edge. \square

We next prove Proposition 6.3. To do this, we first need to do a small adjustment: v-trees are defined in Section 4 so that the variables are at the leaves, but tree automata read labels on all tree nodes, including internal nodes.

Definition 6.4. Let T be a finite rooted ordered binary tree with nodes N . Its leaf-push is the finite rooted ordered binary tree T' obtained by applying bottom-up the following transformation: the transformation of a leaf node is this leaf node, and we replace each internal node n with children n_1 and n_2 by an internal node n' having as children one leaf n and an internal node n'' having as children the transformations of n_1 and n_2 .

Note that the leaves of T' are precisely in bijection with the nodes of T .

Proof of Proposition 6.3. We do not assume this time that the input tree automaton A is complete, but must instead assume that it is *trimmed*, so as to satisfy property (*) above.

We build a Boolean circuit bottom-up, with \vee -gates $g_{n,q}$ for each tree node n of the input tree T . The output gate will be an \vee -gate g doing the disjunction of all the $g_{r,q}$ for q final and for r the root of T .

Now, for every leaf n of the input tree T , we add a variable gate for n as an input of the gate $g_{n,q}$ for each $q \in \iota(1)$, and we add a gate for the negation of n as an input of the gate $g_{n,q}$ for each $q \in \iota(0)$.

For every internal node n of the input tree T with children n_1 and n_2 , for every pair of states q_1 and q_2 , for every $b \in \{0, 1\}$, for every state q having a transition from q_1, q_2, b , we add as input to $g_{n,q}$ a \wedge -gate having as first input either n or $\neg n$ depending on whether $b = 1$ or $b = 0$, and as second input a \wedge -gate having as first input g_{n_1,q_1} and as second input g_{n_2,q_2} .

We then remove from the obtained Boolean circuit all gates that have no directed path to the output gate.

The construction satisfies the time bound. Note that the circuit is clearly in NNF. We claim that it is decomposable, and that it is structured by a v-tree which is the leaf-push of the tree T . Indeed, we can easily show by bottom-up induction that the domain of each gate $g_{n,q}$ is a subset of the nodes of the subtree of T rooted at n (including n), so that the \wedge -gates are indeed structured according to the leaf-push. We additionally point out that the circuit is smooth: this relies on the fact that we removed gates having no directed path to the output gate.

We show correctness again by finite induction: for every tree node n of T , given a labeling of the subtree T_n of T rooted at n , the gates $g_{n,q}$ satisfied by the corresponding partial assignment are precisely those for the states q that we can reach when reading the labeling of T_n in A . The base case corresponds to leaves, where indeed the gates $g_{n,q}$ that are satisfied follow the initial function ι by construction.

Now, for the induction step, let us consider an internal node n of T with children n_1 and n_2 . By induction hypothesis, we assume that the states that we can reach when reading a labeling λ_1 of the tree T_{n_1} and when reading a labeling λ_2 of the tree T_{n_2} are the ones given by the invariant. Let us show that the invariant is satisfied when reading a labeling λ of T_n . Note that λ is defined by a labeling λ_1 on T_{n_1} , a labeling λ_2 on T_{n_2} , and a label $b \in \Sigma$ on n . The gates $g_{n,q}$ that are satisfied are precisely those for which there is a transition from q_1 and q_2 and b to q and for which g_{n_1,q_1} and g_{n_2,q_2} are satisfied: this allows us to conclude from the induction hypothesis. Hence, using the induction result on the root r of T (where $T_r = T$), we conclude that the output gate of the Boolean circuit is true on an assignment \mathbf{a} iff there is a final state of A that we can reach on the labeling of T according to \mathbf{a} .

Now, observe that if the automaton is unambiguous, then each \vee -gate is deterministic. Indeed, if the output gate had two inputs that are mutually satisfiable, then it witnesses the existence of a labeling of the tree T on which A reaches two different final states, contradicting the unambiguity of A . Likewise, if a gate $g_{n,q}$ has two inputs that are mutually satisfiable, as they are satisfied by the same valuation the literal on n must be the same so it must be the case that we can simultaneously satisfy g_{n_1,q_1} and g_{n_2,q_2} , and g_{n_1,q'_1} and g_{n_2,q'_2} , for two 2-tuples $(q_1, q_2) \neq (q'_1, q'_2)$. This witnesses that, on this labeling of T_n , the automaton has two distinct runs leading to state q at the root. This is a contradiction of property (*). \square

7 Conclusion and Extensions

We have introduced in this document the notions of binary decision diagrams, Boolean circuits, and automata. We have explained in which sense binary decision diagrams can be seen as a special case of Boolean circuits. We have also explained how automata can be translated to structured Boolean circuits (for tree automata) or ordered binary decision diagrams (for word automata).

We close the document by reviewing topics which are not presently covered by the document, but could be covered in further versions of the document or by follow-up works:

- The connections between *SDDs and automata*, between *strongly deterministic Boolean circuits and automata*, or the question of which Boolean circuits can be associated to *deterministic automata* (a related notion is *upwards-determinism* in [Amarilli et al., 2017]).
- The notion of *k-unambiguity* for automata (i.e., having *at most k* accepting runs), and the classes of Boolean circuits and binary decision diagrams to which this corresponds.
- The notion of *width* for structured Boolean circuits [Capelli and Mengel, 2019] and for binary decision diagrams, and its connection to the number of states of automata.
- The class of *circuits having a compatible order* [Amarilli et al., 2017], which is intermediate between structuredness and decomposability (i.e., it restricts the possible conjunctions but in a weaker way than requiring a fixed v-tree).
- The question of the *complexity of various problems* on the various classes of Boolean circuits and binary decision diagrams, in the spirit of the knowledge compilation map [Darwiche and

Marquis, 2002]; or their *closure under operations* (e.g., given two Boolean circuits in some formalism, can we tractably compute their disjunction, their conjunction, etc., in the same formalism?); or the *complexity of translating from one formalism to another* (some results of this kind are surveyed in [Amarilli et al., 2020])

- The complexity, given an input Boolean circuit or binary decision diagram, of *testing which conditions it satisfies*: this is generally tractable for syntactic criteria (e.g., decomposability, decision), but can be intractable for semantic criteria (e.g., determinism).
- The extension from Boolean circuits to different circuit types. These include in particular *multivalued circuits*, which are defined over larger sets than the Boolean values; such circuits can be useful in a database context because they relate to the notion of *factorized databases* [Olteanu and Závodný, 2015]. We can also interpret circuits more generally over semirings (of which the Boolean semiring is a special case), which has been used in the setting of *provenance circuits for semiring provenance* [Deutch et al., 2014], and which has been exploited for tractable model counting [Kimmig et al., 2017]. Last, we can consider *arithmetic circuits* [Shpilka and Yehudayoff, 2010], which are circuits computing a polynomial over a given field \mathbb{F} , with internal nodes corresponding to multiplication or addition. The analogue of decomposability is then to require that the polynomial computed is *syntactically multilinear*, i.e., that in the expanded form of the polynomial, in every monomial, each variable has exponent either zero or one. Further, the analogue of determinism is then that every monomial in the expanded form of the polynomial has coefficient either zero or one.
- The connections to *probabilistic circuits* [Choi et al., 2020], which are circuit classes that define probability distributions rather than Boolean functions. The field of probabilistic circuits also defines circuit properties analogous to those discussed in this document, and studies the complexity of various *probabilistic queries* on different classes of probabilistic circuits. We could also study the connections to various tractable probabilistic models such as *sum-product networks* [Poon and Domingos, 2011], *arithmetic circuits* [Darwiche, 2003], *cutset network* [Rahman et al., 2014], *and-or graphs* [Dechter and Mateescu, 2007], *probabilistic sentential decision diagrams* [Kisa et al., 2014], and more, which can be understood via translations to various classes of probabilistic circuits.
- The connections to other formalisms to represent languages over words, e.g., *context-free grammars*, for which Boolean circuit representations are implicit in [Amarilli et al., 2022], and which have recently been related to factorized databases in [Kimelfeld et al., 2023].
- The use of *depth reduction* techniques, to rewrite circuits and binary decision diagrams to circuits of lower depth [Valiant et al., 1983]
- As a converse to provenance circuits, the *conversion of circuits into automata*, although this is not obvious to define because, unlike automata, circuits generally do not have a “uniform” behavior.

References

Antoine Amarilli, Pierre Bourhis, and Pierre Senellart. Provenance circuits for trees and treelike instances. In *ICALP*, 2015. doi: 10.1007/978-3-662-47666-6_5.

- Antoine Amarilli, Pierre Bourhis, Louis Jachiet, and Stefan Mengel. A circuit-based approach to efficient enumeration. In *ICALP*, 2017. doi: 10.4230/LIPIcs.ICALP.2017.111.
- Antoine Amarilli, Florent Capelli, Mikaël Monet, and Pierre Senellart. Connecting knowledge compilation classes and width parameters. *Theory Comput. Syst.*, 64(5), 2020. doi: 10.1007/s00224-019-09930-2.
- Antoine Amarilli, Louis Jachiet, Martín Muñoz, and Cristian Riveros. Efficient enumeration algorithms for annotated grammars. In *PODS*, 2022. doi: 10.1145/3517804.3526232.
- Dana Angluin, Lisa Hellerstein, and Marek Karpinski. Learning read-once formulas with queries. *JACM*, 40(1), 1993. doi: 10.1145/138027.138061.
- Marcelo Arenas, Pablo Barceló, Leopoldo Bertossi, and Mikaël Monet. On the complexity of SHAP-score-based explanations: Tractability via knowledge compilation and non-approximability results. *Journal of Machine Learning Research*, 24(63), 2023.
- Paul Beame, Jerry Li, Sudeepa Roy, and Dan Suciu. Lower bounds for exact model counting and applications in probabilistic databases. In *UAI*, 2013.
- Manuel Blum, Ashok K. Chandra, and Mark N. Wegman. Equivalence of free boolean graphs can be decided probabilistically in polynomial time. *Inf. Process. Lett.*, 10(2), 1980. doi: 10.1016/S0020-0190(80)90078-2.
- Beate Bollig. On the minimization of (complete) ordered binary decision diagrams. *Theory of Computing Systems*, 59, 2016. doi: 10.1007/s00224-015-9657-x.
- Beate Bollig and Ingo Wegener. Complexity theoretical results on partitioned (nondeterministic) binary decision diagrams. In *MFCS*, 1997. doi: 10.1007/s002240000128.
- Beate Bollig, Martin Sauerhoff, Detlef Sieling, Ingo Wegener, Yves Crama, and Peter L. Hammer. Binary decision diagrams. In Yves Crama and Peter L. Hammer, editors, *Boolean Models and Methods in Mathematics, Computer Science, and Engineering*. Cambridge University Press, 2010. ISBN 978-0-521-84752-0.
- Randal E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Trans. Computers*, 35(8), 1986. doi: 10.1109/TC.1986.1676819.
- Florent Capelli and Stefan Mengel. Tractable QBF by knowledge compilation. In *STACS*, 2019. doi: 10.4230/LIPIcs.STACS.2019.18.
- YooJung Choi, Antonio Vergari, and Guy Van den Broeck. Probabilistic circuits: A unifying framework for tractable probabilistic models. 2020. URL <http://starai.cs.ucla.edu/papers/ProbCirc20.pdf>.
- Adnan Darwiche. On the tractable counting of theory models and its application to truth maintenance and belief revision. *Journal of Applied Non-Classical Logics*, 11(1-2), 2001. doi: 10.3166/jancl.11.11-34.
- Adnan Darwiche. A differential approach to inference in bayesian networks. *Journal of the ACM (JACM)*, 50(3), 2003. doi: 10.1145/765568.765570.
- Adnan Darwiche. SDD: A new canonical representation of propositional knowledge bases. In *IJCAI*, 2011.

- Adnan Darwiche and Pierre Marquis. [A knowledge compilation map](#). *Journal of Artificial Intelligence Research*, 17, 2002. doi: 10.1613/jair.989.
- Rina Dechter and Robert Mateescu. [AND/OR search spaces for graphical models](#). *Artificial intelligence*, 171(2-3), 2007. doi: 10.1016/j.artint.2006.11.003.
- Daniel Deutch, Tova Milo, Sudeepa Roy, and Val Tannen. [Circuits for Datalog Provenance](#). In *ICDT*, 2014. doi: 10.5441/002/ICDT.2014.22.
- Steven Fortune, John E. Hopcroft, and Erik Meineche Schmidt. The complexity of equivalence and containment for free single variable program schemes. In *ICALP*, 1978. doi: 10.1007/3-540-08860-1_17.
- Todd J Green, Grigoris Karvounarakis, and Val Tannen. [Provenance semirings](#). In *PODS*, 2007. doi: 10.1145/1265530.1265535.
- Abhay Jha and Dan Suciu. [Knowledge compilation meets database theory: Compiling queries to decision diagrams](#). *Theory of Computing Systems*, 52, 2013. doi: 10.1145/1938551.1938574.
- Benny Kimelfeld, Wim Martens, and Matthias Niewerth. [A unifying perspective on succinct data representations](#). *arXiv preprint arXiv:2309.11663*, 2023. doi: 10.48550/arXiv.2309.11663.
- Angelika Kimmig, Guy Van den Broeck, and Luc De Raedt. [Algebraic model counting](#). *Journal of Applied Logic*, 22, 2017. doi: 10.1016/j.jal.2016.11.031.
- Doga Kisa, Guy Van den Broeck, Arthur Choi, and Adnan Darwiche. [Probabilistic sentential decision diagrams](#). In *KR*, 2014.
- Eyal Kushilevitz. [Communication complexity](#). In *Advances in Computers*, volume 44. 1997. doi: 10.1016/S0065-2458(08)60342-3.
- Jean-Marie Lagniez and Pierre Marquis. [An improved decision-DNNF compiler](#). In *IJCAI*, volume 17, 2017.
- Chang-Yeong Lee. [Representation of switching circuits by binary-decision programs](#). *The Bell System Technical Journal*, 38(4), 1959. doi: 10.1002/j.1538-7305.1959.tb01585.x.
- Shin-ichi Minato. [Zero-suppressed BDDs for set manipulation in combinatorial problems](#). In *IDAC*, 1993. doi: 10.1145/157485.164890.
- Dan Olteanu and Jakub Závodný. [Size bounds for factorised representations of query results](#). *ACM Transactions on Database Systems (TODS)*, 40(1), 2015. doi: 10.1145/2656335.
- Thammanit Pipatsrisawat and Adnan Darwiche. [A lower bound on the size of decomposable negation normal form](#). In *AAAI*, volume 24, 2010. doi: 10.1609/AAAI.V24I1.7600.
- Hoifung Poon and Pedro Domingos. [Sum-product networks: A new deep architecture](#). In *ICCV workshops*, 2011.
- Tahrima Rahman, Prasanna Kothalkar, and Vibhav Gogate. [Cutset networks: A simple, tractable, and scalable approach for improving the accuracy of Chow-Liu trees](#). In *ECML/PKDD*, 2014. doi: 10.1007/978-3-662-44851-9_40.

- Alexander A Razborov. Lower bounds for deterministic and nondeterministic branching programs. In *FCT*, 1991. doi: 10.1007/3-540-54458-5_49.
- Igor Razgon. No small nondeterministic read-once branching programs for CNFs of bounded treewidth. In *IPEC*, 2014. doi: 10.1007/978-3-319-13524-3_27.
- S. Rasoul Safavian and David A. Landgrebe. A survey of decision tree classifier methodology. *IEEE Trans. Syst. Man Cybern.*, 21(3), 1991.
- Andy Shih, Guy Van den Broeck, Paul Beame, and Antoine Amarilli. Smoothing structured decomposable circuits. In *NeurIPS*, 2019.
- Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends in Theoretical Computer Science*, 5(3–4), 2010. doi: 10.1561/04000000039.
- LG Valiant, S Skyum, S Berkowitz, and C Rackoff. Fast parallel computation of polynomials using few processors. *SIAM Journal on Computing*, 12(4), 1983. doi: 10.1007/3-540-10856-4_79.
- Heribert Vollmer. *Introduction to circuit complexity: a uniform approach*. Springer Science & Business Media, 1999. ISBN 978-3-540-64310-4.
- Ingo Wegener. *The complexity of Boolean functions*. John Wiley & Sons, Inc., 1987. ISBN 0-471-91555-6.
- Ingo Wegener. *Branching programs and binary decision diagrams: Theory and applications*. SIAM, 2000. ISBN 978-0-89871-458-6.
- Ingo Wegener. BDDs—design, analysis, complexity, and applications. *Discret. Appl. Math.*, 138(1-2), 2004. doi: 10.1016/S0166-218X(03)00297-X.
- Xindong Wu, Vipin Kumar, J. Ross Quinlan, Joydeep Ghosh, Qiang Yang, Hiroshi Motoda, Geoffrey J. McLachlan, Angus F. M. Ng, Bing Liu, Philip S. Yu, Zhi-Hua Zhou, Michael S. Steinbach, David J. Hand, and Dan Steinberg. Top 10 algorithms in data mining. *Knowl. Inf. Syst.*, 14(1), 2008. doi: 10.1007/s10115-007-0114-2.