

Canonical Automata via Distributive Law Homomorphisms

Stefan Zetsche*

University College London
s.zetsche@cs.ucl.ac.uk

Gerco van Heerdt†

Royal Holloway, University of London
University College London

Matteo Sammartino‡

Alexandra Silva§

Cornell University
University College London

The classical powerset construction is a standard method converting a non-deterministic automaton into a deterministic one recognising the same language. Recently, the powerset construction has been lifted to a more general framework that converts an automaton with side-effects, given by a monad, into a deterministic automaton accepting the same language. The resulting automaton has additional algebraic properties, both in the state space and transition structure, inherited from the monad. In this paper, we study the reverse construction and present a framework in which a deterministic automaton with additional algebraic structure over a given monad can be converted into an equivalent succinct automaton with side-effects. Apart from recovering examples from the literature, such as the canonical residual finite-state automaton and the átomaton, we discover a new canonical automaton for a regular language by relating the free vector space monad over the two element field to the neighbourhood monad. Finally, we show that every regular language satisfying a suitable property parametric in two monads admits a size-minimal succinct acceptor.

1 Introduction

The existence of a unique minimal *deterministic* acceptor is an important property of regular languages. Establishing a similar result for *non-deterministic* acceptors is significantly more difficult, but nonetheless of great practical importance, as non-deterministic automata (NFA) can be exponentially more succinct than deterministic ones (DFA). The main issue is that a regular language can be accepted by several size-minimal NFAs that are not isomorphic. A number of sub-classes of non-deterministic automata have been identified in the literature to tackle this issue, which all admit canonical representatives: the *átomaton* [12], the *canonical residual finite-state automaton* (short *canonical RFSA* and also known as *jiromaton*) [13], the *minimal xor automaton* [42], and the *distromaton* [33].

In this paper we provide a general categorical framework that unifies constructions of canonical non-deterministic automata and unveils new ones. Our framework adopts the well-known representation of side-effects via *monads* [32] to generalise non-determinism in automata. For instance, an NFA (without initial states) can be represented as a pair $\langle X, k \rangle$, where X is the set of states and $k: X \rightarrow 2 \times \mathcal{P}(X)^A$ combines the function classifying each state as accepting or rejecting with the function giving the set of next states for each input. The powerset forms a monad $\langle \mathcal{P}, \{-\}, \mu \rangle$, where $\{-\}$ creates singleton

*The author has been supported by GCHQ via the VeTSS grant “Automated black-box verification of networking systems” (4207703/RFA 15845).

†The author has been supported by the EPSRC via the Standard Grant CLeVer (EP/S028641/1).

‡The author has been supported by the EPSRC via the Standard Grant CLeVer (EP/S028641/1).

§The author has been supported by the ERC via the Consolidator Grant AutoProbe 101002697 and by a Royal Society Wolfson Fellowship.




Figure 1: Generalised determinisation of automata with side-effects in a monad.

sets and μ takes the union of a set of sets. This allows describing the classical powerset construction, converting an NFA into a DFA, in categorical terms [39] as depicted on the left of Figure 1, where $k^\sharp: \mathcal{P}(X) \rightarrow 2 \times \mathcal{P}(X)^A$ represents an equivalent DFA, obtained by taking the subsets of X as states, and $\langle \epsilon, \delta \rangle: 2^{A^*} \rightarrow 2 \times (2^{A^*})^A$ is the automaton of languages. There then exists a unique automaton homomorphism obs , assigning a language semantics to each set of states.

As seen on the right of Figure 1 this perspective further enables a *generalised determinisation* construction [39], where $2 \times (-)^A$ is replaced by any (suitable) functor F describing the automaton structure, and \mathcal{P} by a monad T describing the automaton side-effects. $\Omega \xrightarrow{\omega} F\Omega$ is the so-called *final coalgebra*, providing a semantic universe that generalises the automaton of languages.

Our work starts from the observation that the deterministic automata resulting from this generalised determinisation constructions have *additional algebraic structure*: the state space $\mathcal{P}(X)$ of the determinised automaton defines a free complete join-semilattice (CSL) over X , and k^\sharp and obs are CSL homomorphisms. More generally, TX defines a (free) algebra for the monad T , and k^\sharp and obs are T -algebra homomorphisms.

With this observation in mind, our question is: can we exploit the additional algebraic structure to “reverse” these constructions? In other words, can we convert a deterministic automaton with additional algebraic structure over a given monad to an equivalent succinct automaton with side-effects, possibly over another monad? To answer this question, the paper makes the following contributions:

- We present a general categorical framework based on bialgebras and distributive law homomorphisms that allows deriving canonical representatives for a wide class of succinct automata with side-effects in a monad.
- We strictly improve the expressivity of previous work [4, 19]: our framework instantiates not only to well-known examples such as the canonical RFSA (Example 16) and the minimal xor automaton (Example 18), but also includes the átomaton (Section 5.3) and the distromaton (Section 5.4), which were not covered in [4, 19]. While other frameworks restrict themselves to the category of sets [19], we are able to include canonical acceptors in other categories, such as the *canonical nominal RFSA* (Example 17).
- We relate vector spaces over the unique two element field with complete atomic Boolean algebras and consequently discover a previously unknown canonical mod-2 weighted acceptor for regular languages—the *minimal xor-CABA automaton* (Section 5.5)—that in some sense is to the minimal xor automaton what the átomaton is to the canonical RFSA (Figure 9).
- We introduce an abstract notion of *closedness* for succinct automata that is parametric in two monads (Definition 30), and prove that every regular language satisfying a suitable property admits a canonical size-minimal representative among closed acceptors (Theorem 32). By instantiating the latter we subsume known minimality results for canonical automata, prove the xor-CABA automaton minimal, and establish a size comparison between different acceptors (Section 6.1).

An extended version of this paper is available at [44].




Figure 2: The minimal DFA for $\mathcal{L} = (a + b)^*a$.

2 Overview of the approach

In this section, we give an overview of the ideas of the paper through an example. We show how our methodology allows recovering the construction of the átomaton for the regular language $\mathcal{L} = (a + b)^*a$, which consists of all words over $A = \{a, b\}$ that end in a . For each step, we hint at how it is generalised in our framework.

The classical construction of the átomaton for \mathcal{L} consists in closing the *residuals*¹ of \mathcal{L} under all Boolean operations, and then forming a non-deterministic automaton whose states are the *atoms*² of the ensuing complete atomic Boolean algebra (CABA)—that is, non-empty intersections of complemented or uncomplemented residuals. In our categorical setting, this construction is obtained in several steps, which we now describe.

2.1 Computing residuals

We first construct the minimal DFA accepting \mathcal{L} as a coalgebra of type $M_{\mathcal{L}} \rightarrow 2 \times (M_{\mathcal{L}})^A$. By the well-known Myhill-Nerode theorem [34], $M_{\mathcal{L}}$ is the set of residuals for \mathcal{L} . The automaton is depicted in Figure 2.

In our framework, we consider coalgebras over an arbitrary endofunctor $F: \mathcal{C} \rightarrow \mathcal{C}$ ($F = 2 \times (-)^A$ and $\mathcal{C} = \text{Set}$ in this case). Minimal realisations, generalising minimal DFAs, exist for a wide class of functors F and categories \mathcal{C} , including all the examples in this paper.

2.2 Taking the Boolean closure

We close the minimal DFA under all Boolean operations, generating an equivalent deterministic automaton that has additional algebraic structure: its state space is a CABA. This is achieved via a double powerset construction—where sets of sets are interpreted as full disjunctive normal form—and the resulting coalgebra is of type $\mathcal{P}^2(M_{\mathcal{L}}) \rightarrow 2 \times (\mathcal{P}^2(M_{\mathcal{L}}))^A$. Our construction relies on the so-called *neighbourhood monad* \mathcal{H} , whose algebras are precisely CABAs, and yields a (free) *bialgebra* capturing both the coalgebraic and the algebraic structure; the interplay of these two structures is captured via a *distributive law*. We then minimise this DFA to identify Boolean expressions evaluating to the same language. As desired, the resulting state space is precisely the Boolean closure of the residuals of \mathcal{L} . Formally, we obtain the minimal bialgebra for \mathcal{L} depicted in Figure 3.

This step in our framework is generalised as closure of an F -coalgebra w.r.t (the algebraic structure induced by) any monad S for which a suitable distributive law λ with the coalgebra endofunctor F exists. The first step of the closure yields a free λ -bialgebra, comprised of both an F -coalgebra and an S -algebra over the same state space. In a second step, minimisation is carried out in the category of λ -bialgebras, which guarantees simultaneous preservation of the algebraic structure and of the language semantics.

¹A language is a *residual* or *left quotient* of $\mathcal{L} \subseteq A^*$, if it is of the form $v^{-1}\mathcal{L} = \{u \in A^* \mid vu \in \mathcal{L}\}$ for some $v \in A^*$.

²A non-zero element $a \in B$ is called *atom*, if for all $x \in B$ such that $x \leq a$ one finds $x = 0$ or $x = a$.




Figure 3: The minimal CABA-structured DFA for $\mathcal{L} = (a + b)^*a$, where $1 \equiv [\{\{x\}, \{x, y\}\}]$, $2 \equiv [\emptyset]$, $3 \equiv [\{\emptyset\}]$, $4 \equiv [\{\{x, y\}, \emptyset\}]$, $5 \equiv [\{\{x\}, \{y\}, \{x, y\}\}]$, $6 \equiv [\{\{y\}\}]$, $7 \equiv [\{\{y\}, \emptyset\}]$, $8 \equiv [\{\{x\}, \{y\}, \{x, y\}, \emptyset\}]$.

2.3 Constructing the átomaton

This step is the key technical result of our paper. Atoms have the property that their Boolean closure generates the entire CABA. In our framework, this property is generalised via the notion of *generators* for algebras over a monad, which allows one to represent a bialgebra as an equivalent *free* bialgebra over its generators, and hence to obtain succinct canonical representations (Proposition 14). In Section 4 we apply this result to obtain the canonical RFSA, the canonical nominal RFSA, and the minimal xor automaton for a given regular language.

However, to recover the átomaton from the minimal CABA-structured DFA of the previous step, in addition a subtle change of perspective is required. In fact, we are still working with the “wrong” side-effect: the non-determinism of bialgebras so far is determined by \mathcal{H} , whereas we are interested in an NFA, whose non-determinism is captured by the powerset monad \mathcal{P} . As is well-known, every element of a CABA can be obtained as the join of the atoms below it. In other words, those atoms are also generators of the underlying CSL, which is an algebra for \mathcal{P} . We formally capture this idea as a map between monads $\mathcal{H} \rightarrow \mathcal{P}$. Crucially, we show that this map lifts to a *distributive law homomorphism* and allows translating a bialgebra over \mathcal{H} to a bialgebra over \mathcal{P} , which can be represented as a free bialgebra over atoms—the átomaton for \mathcal{L} , which is shown in Figure 4.

In Section 5 we generalise this idea to the situation of two monads S and T involved in distributive laws with the coalgebra endofunctor F . In particular, Corollary 21 is our free representation result, spelling out a condition under which a bialgebra over S can be represented as a free bialgebra over T , and hence admits an equivalent succinct representation as an automaton with side-effects in T . Besides the átomaton and the examples in Section 4, this construction allows us to capture the distromaton and a newly discovered canonical acceptor that relates CABAs with vector spaces over the two element field.

3 Preliminaries

We assume basic knowledge of category theory (including functors, natural transformations, and adjunctions) [5]. In this section we recall the relevant notions for our technical development: coalgebras, monads, algebras over a monad, distributive laws, and bialgebras.

Unpointed deterministic automata are basic examples of coalgebras in the category of sets and functions: they are of the type $k: X \rightarrow FX$, where $FX = 2 \times X^A$ and k pairs the final state function and the transition function assigning a next state to each letter $a \in A$. Coalgebra has emerged as a unifying framework to study infinite data types and state-based systems [37].

Definition 1. (Coalgebra) A *coalgebra* for an endofunctor F in a category \mathcal{C} is a tuple $\langle X, k \rangle$ consisting of an object X in \mathcal{C} and a morphism $k: X \rightarrow FX$.

Crucial in the theory of coalgebras is the notion of homomorphism, which allows to relate states of coalgebras of the same type. A homomorphism $f: \langle X, k_X \rangle \rightarrow \langle Y, k_Y \rangle$ between F -coalgebras is a morphism $f: X \rightarrow Y$ satisfying $k_Y \circ f = Ff \circ k_X$. The category of F -coalgebras and homomorphisms is denoted by $\text{Coalg}(F)$. If it exists, the final object of this category is of particular importance.

Definition 2. (Final coalgebra) An F -coalgebra $\langle \Omega, k_\Omega \rangle$ is *final* if every F -coalgebra $\langle X, k \rangle$ admits a unique homomorphism $\text{obs}_{\langle X, k \rangle}: \langle X, k \rangle \rightarrow \langle \Omega, k_\Omega \rangle$.

The unique final coalgebra homomorphism can be understood as the observable behaviour of a system. For example, for the functor $FX = 2 \times X^A$, the final F -coalgebra is the set of all languages $\mathcal{P}(A^*)$ and the final coalgebra homomorphism assigns to a state x of an unpointed deterministic automaton the language in $\mathcal{P}(A^*)$ it accepts³ when given the initial state x .

In the context of computer science, monads have been introduced by Moggi as a general perspective on exceptions, side-effects, non-determinism, and continuations [30–32].

Definition 3. (Monad) A *monad* on a category \mathcal{C} is a tuple $\langle T, \eta, \mu \rangle$ consisting of an endofunctor $T: \mathcal{C} \rightarrow \mathcal{C}$ and natural transformations $\eta: \text{id}_{\mathcal{C}} \Rightarrow T$ and $\mu: T^2 \Rightarrow T$ satisfying $\mu \circ T\mu = \mu \circ \mu_T$ and $\mu \circ \eta_T = \text{id}_T = \mu \circ T\eta$.

By a slight abuse of notation we will refer to a monad simply by its underlying endofunctor.

Non-determinism is typically modelled by the *powerset monad* \mathcal{P} , whose underlying endofunctor \mathcal{P} assigns to a set X the set of subsets $\mathcal{P}X$; whose unit maps an element x to the singleton $\eta_X(x) = \{x\}$; and whose multiplication flattens subsets by taking their union $\mu_X(\Phi) = \bigcup_{U \in \Phi} U$. Other monads that play a role for us are the *nominal powerset monad* \mathcal{P}_n [29], the *neighbourhood monad* \mathcal{H} [23], the *monotone neighbourhood monad* \mathcal{A} [23], and the *free vector space monad* \mathcal{R} over the unique two element field [22]. The formal definitions are given in [44, Definition 40].

The concept of a monad can also be seen as an alternative to Lawvere theory as a category theoretic formulation of universal algebra [14, 28].

Definition 4. (Algebra over a monad) An *algebra* over a monad T on \mathcal{C} is a tuple $\langle X, h \rangle$ consisting of an object X in \mathcal{C} and a morphism $h: TX \rightarrow X$ satisfying $h \circ \mu_X = h \circ Th$ and $h \circ \eta_X = \text{id}_X$.


Every object admits a *free algebra* $\langle TX, \mu_X \rangle$. A homomorphism $f: \langle X, h_X \rangle \rightarrow \langle Y, h_Y \rangle$ between T -algebras is a morphism $f: X \rightarrow Y$ satisfying $h_Y \circ Tf = f \circ h_X$. The category of T -algebras and homomorphisms is denoted by $\text{Alg}(T)$.

Example 5. • The category $\text{Alg}(\mathcal{P})$ is isomorphic to the category of complete join-semi lattices (CSL) and functions that preserve all joins [22].

- The category $\text{Alg}(\mathcal{H})$ is isomorphic to the category of complete atomic Boolean algebras (CABA) and Boolean algebra homomorphisms that preserve all meets and all joins [23].
- The category $\text{Alg}(\mathcal{A})$ is isomorphic to the category of completely distributive lattices (CDL) and functions that preserve all meets and all joins [23].
- The category $\text{Alg}(\mathcal{R})$ is isomorphic to the category of vector spaces over the unique two element field (\mathbb{Z}_2 -Vect) and linear maps [22].

Distributive laws have originally occurred as a way to compose monads [6], but now also exist in a wide range of other forms [40]. For our particular case it is sufficient to consider distributive laws between a monad and an endofunctor, sometimes referred to as *Eilenberg-Moore laws* [25].

³For a deterministic automaton given by $\varepsilon: X \rightarrow 2$ and $\delta: X \rightarrow X^A$, acceptance is coinductively defined as a function $\text{obs}: X \rightarrow 2^{A^*}$ by $\text{obs}(x)(\varepsilon) = \varepsilon(x)$ and $\text{obs}(x)(av) = \text{obs}(\delta(x)(a))(v)$.

Figure 4: The átomaton for $\mathcal{L} = (a + b)^*a$.

Definition 6. (Distributive law) A distributive law between a monad T and an endofunctor F on \mathcal{C} is a natural transformation $\lambda : TF \Rightarrow FT$ satisfying $F\eta_X = \lambda_X \circ \eta_{FX}$ and $\lambda_X \circ \mu_{FX} = F\mu_X \circ \lambda_{TX} \circ T\lambda_X$.

For example, every algebra $h : TB \rightarrow B$ for a set monad T induces a distributive law λ^h between T and F with $FX = B \times X^A$ defined by

$$\lambda_X^h := (h \times \text{st}) \circ \langle T\pi_1, T\pi_2 \rangle, \quad (1)$$

where st denotes the usual strength function⁴ [21]. We are particularly interested in canonical algebra structures for the output set $B = 2$. For instance, the algebra structures defined by $h^{\mathcal{P}}(\varphi) = h^{\mathcal{P}}(\varphi) = \varphi(1)$ and $h^{\mathcal{C}}(\Phi) = h^{\mathcal{C}}(\Phi) = \Phi(\text{id}_2)$, where we identify subsets with their characteristic functions. In these cases we will abuse notation and write λ^T instead of λ^{h^T} .

Example 7. (Generalized determinisation [36]) Given a distributive law, one can model the determinisation of a system with dynamics in F and side-effects in T (sometimes referred to as *succinct* automaton) by lifting a FT -coalgebra $\langle X, k \rangle$ to the F -coalgebra $\langle TX, k^\sharp \rangle$, where $k^\sharp := (F\mu_X \circ \lambda_{TX}) \circ Tk$. As one verifies, the latter is in fact a T -algebra homomorphism of type $k^\sharp : \langle TX, \mu_X \rangle \rightarrow \langle FTX, F\mu_X \circ \lambda_{TX} \rangle$. For instance, if the distributive law λ is induced by the disjunctive \mathcal{P} -algebra $h^{\mathcal{P}} : \mathcal{P}2 \rightarrow 2$ with $h^{\mathcal{P}}(\varphi) = \bigvee_{u \in \varphi} u = \varphi(1)$, the lifting k^\sharp is the DFA in CSL obtained from an NFA k via the classical powerset construction.

The example above illustrates the concept of a bialgebra: the algebraic part $\langle TX, \mu_X \rangle$ and the coalgebraic part $\langle TX, k^\sharp \rangle$ of a lifted automaton are compatible along the distributive law λ .

Definition 8. (Bialgebra) A λ -bialgebra is a tuple $\langle X, h, k \rangle$ consisting of a T -algebra $\langle X, h \rangle$ and an F -coalgebra $\langle X, k \rangle$ satisfying $Fh \circ \lambda_X \circ Tk = k \circ h$.

A homomorphism between λ -bialgebras is a morphism between the underlying objects that is simultaneously a T -algebra homomorphism and an F -coalgebra homomorphism. The category of λ -bialgebras and homomorphisms is denoted by $\text{Bialg}(\lambda)$. The existence of a final F -coalgebra is equivalent to the existence of a final λ -bialgebra, as the next result shows.

Lemma 9. [24] Let $\langle \Omega, k_\Omega \rangle$ be the final F -coalgebra, then $\langle \Omega, h_\Omega, k_\Omega \rangle$ with $h_\Omega := \text{obs}_{\langle T\Omega, \lambda_\Omega \circ Tk_\Omega \rangle}$ is the final λ -bialgebra satisfying $\text{obs}_{\langle X, h, k \rangle} = \text{obs}_{\langle X, k \rangle}$. Conversely, if $\langle \Omega, h_\Omega, k_\Omega \rangle$ is the final λ -bialgebra, then $\langle \Omega, k_\Omega \rangle$ is the final F -coalgebra.

For instance, for the distributive law in Example 7, the final bialgebra is carried by the final coalgebra $\mathcal{P}(A^*)$ and also has a free \mathcal{P} -algebra structure that takes the union of languages.

The generalized determinisation procedure in Example 7 can now be rephrased in terms of a functor between the category of coalgebras with dynamics in F and side-effects in T on the one side, and the category of bialgebras on the other side.

Lemma 10. [24] Defining $\text{exp}_T(\langle X, k \rangle) := \langle TX, \mu_X, (F\mu_X \circ \lambda_{TX}) \circ Tk \rangle$ and $\text{exp}_T(f) := Tf$ yields a functor $\text{exp}_T : \text{Coalg}(FT) \rightarrow \text{Bialg}(\lambda)$.

We will sometimes refer to the functor which arises from the one above by precomposition with the canonical embedding of F -coalgebras into FT -coalgebras.

Corollary 11. Defining $\text{free}_T(\langle X, k \rangle) := \langle TX, \mu_X, \lambda_X \circ Tk \rangle$ and $\text{free}_T(f) := Tf$ yields a functor $\text{free}_T : \text{Coalg}(F) \rightarrow \text{Bialg}(\lambda)$ satisfying $\text{free}_T(\langle X, k \rangle) = \text{exp}_T(\langle X, F\eta_X \circ k \rangle)$.

⁴For any two sets X, A the strength function $\text{st} : T(X^A) \rightarrow (TX)^A$ is defined by $\text{st}(U)(a) = T(\text{ev}_a)(U)$, where $\text{ev}_a(f) = f(a)$.




Figure 5: (a) The minimal CSL-structured DFA for $\mathcal{L} = (a + b)^*a$; (b) The canonical RFSA for $\mathcal{L} = (a + b)^*a$.

4 Succinct automata from bialgebras

In this section we introduce the foundations of our theoretical contributions. We begin with the notion of a *generator* [4] for an algebra over a monad and demonstrate how it can be used to translate a bialgebra into an equivalent free bialgebra. While the treatment is very general, we are particularly interested in the case in which the bialgebra is given by a deterministic automaton that has additional algebraic structure over a given monad, and the translation results in an automaton with side-effects in that monad. We will demonstrate that the theory in this section instantiates to the canonical RFSA [13], the canonical *nominal* RFSA [29], and the minimal xor automaton [42].

Definition 12. (Generator and basis) A *generator* for a T -algebra $\langle X, h \rangle$ is a tuple $\langle Y, i, d \rangle$ consisting of an object Y , a morphism $i: Y \rightarrow X$, and a morphism $d: X \rightarrow TY$ such that $(h \circ Ti) \circ d = \text{id}_X$. A generator is called a *basis* if it additionally satisfies $d \circ (h \circ Ti) = \text{id}_{TY}$.

A generator for an algebra is called a *scoop* by Arbib and Manes [4]. Here, we additionally introduce the notion of a basis. Intuitively, one calls a set Y that is embedded into an algebraic structure X a generator for the latter if every element x in X admits a decomposition $d(x) \in TY$ into a formal combination of elements of Y that evaluates to x . If the decomposition is moreover *unique*, that is, $h \circ Ti$ is not only a *surjection* with right-inverse d , but a *bijection* with two-sided inverse d , then a generator is called a basis. Every algebra is generated by itself using the generator $\langle X, \text{id}_X, \eta_X \rangle$, but not every algebra admits a basis. We are particularly interested in classes of set-based algebras for which *every* algebra admits a *size-minimal* generator, that is, no generator has a carrier of smaller size. In such a situation we will also speak of *canonical* generators.

Example 13. • A tuple $\langle Y, i, d \rangle$ is a generator for a \mathcal{P} -algebra $L = \langle X, h \rangle \simeq \langle X, \vee^h \rangle$ iff $x = \bigvee_{y \in d(x)} i(y)$ for all $x \in X$. Note that if $Y \subseteq X$ is a subset, then $i(y) = y$ for all $y \in Y$. If L satisfies the descending chain condition, which is in particular the case if X is finite, then defining $i(y) = y$ and $d(x) = \{y \in J(L) \mid y \leq x\}$ turns the set of join-irreducibles⁵ $J(L)$ into a size-minimal generator $\langle J(L), i, d \rangle$ for L , cf. [44, Lemma 55].

- A tuple $\langle Y, i, d \rangle$ is a generator for a \mathcal{R} -algebra $V = \langle X, h \rangle \simeq \langle X, +^h, \cdot^h \rangle$ iff $x = \sum_{y \in Y} d(x)(y) \cdot^h i(y)$ for all $x \in X$. As it is well-known that every vector space can be equipped with a basis, every \mathcal{R} -algebra V admits a basis. One can show that a basis is size-minimal, cf. [44, Lemma 52].

It is enough to find generators for the underlying algebra of a bialgebra to derive an equivalent free bialgebra. This is because the algebraic and coalgebraic components are tightly intertwined via a distributive law.

Proposition 14. Let $\langle X, h, k \rangle$ be a λ -bialgebra and let $\langle Y, i, d \rangle$ be a generator for the T -algebra $\langle X, h \rangle$. Then $h \circ Ti: \exp_T(\langle Y, Fd \circ k \circ i \rangle) \rightarrow \langle X, h, k \rangle$ is a λ -bialgebra homomorphism.

⁵A non-zero element $x \in L$ is called *join-irreducible* if for all $y, z \in L$ such that $x = y \vee z$ one finds $x = y$ or $x = z$.




Figure 6: The orbit-finite representation of the canonical nominal RFSA for $\mathcal{L} = \{vawau \mid v, w, u \in A^*, a \in A\}$.

Intuitively, the bialgebra $\langle X, h, k \rangle$ is a deterministic automaton with additional algebraic structure in the monad T and say initial state $x \in X$, while the equivalent free bialgebra is the determinisation of the succinct automaton $Fd \circ k \circ i: Y \rightarrow FTY$ with side-effects in T and initial state $d(x) \in TY$. The following result further observes that if one considers a basis for the underlying algebraic structure of a bialgebra, rather than just a generator, then the equivalent free bialgebra is in fact isomorphic to the original bialgebra.

Proposition 15. *Let $\langle X, h, k \rangle$ be a λ -bialgebra and let $\langle Y, i, d \rangle$ be a basis for the T -algebra $\langle X, h \rangle$. Then $h \circ Ti: \exp_T(\langle Y, Fd \circ k \circ i \rangle) \rightarrow \langle X, h, k \rangle$ is a λ -bialgebra isomorphism.*

We conclude this section by illustrating how Proposition 14 can be used to construct the canonical RFSA [13], the canonical nominal RFSA [29], and the minimal xor automaton [42] for a regular language \mathcal{L} over some alphabet A . All examples follow three analogous steps:

1. We construct the minimal⁶ pointed coalgebra $M_{\mathcal{L}}$ for the (nominal) set endofunctor $F = 2 \times (-)^A$ accepting \mathcal{L} . For the case $A = \{a, b\}$ and $\mathcal{L} = (a+b)^*a$, the coalgebra $M_{\mathcal{L}}$ is depicted in Figure 2.
2. We equip the former with additional algebraic structure in a monad T (which is related to F via a canonically induced distributive law λ) by generating the λ -bialgebra $\text{free}_T(M_{\mathcal{L}})$. By identifying semantically equivalent states we consequently derive the minimal⁷ (pointed) λ -bialgebra $\langle X, h, k \rangle$ for \mathcal{L} .
3. We identify canonical generators $\langle Y, i, d \rangle$ for $\langle X, h \rangle$ and use Proposition 14 to derive an equivalent succinct automaton $\langle Y, Fd \circ k \circ i \rangle$ with side-effects in T .

Example 16. (The canonical RFSA) Using the \mathcal{P} -algebra structure $h^{\mathcal{P}}: \mathcal{P}2 \rightarrow 2$ with $h^{\mathcal{P}}(\varphi) = \varphi(1)$, we derive a canonical distributive law $\lambda^{\mathcal{P}}$ between F and the powerset monad \mathcal{P} . The minimal pointed $\lambda^{\mathcal{P}}$ -bialgebra for $\mathcal{L} = (a+b)^*a$ with its underlying CSL structure is depicted in Figure 5a; the construction can be verified with the help of [44, Lemma 47]. The partially ordered state space $L = \{\{\emptyset\} \leq \{\{x\}\} \leq \{\{y\}\}\}$ is necessarily finite, thus satisfies the descending chain condition, which turns the set of join-irreducibles into a size-minimal generator $\langle J(L), i, d \rangle$ with $i(y) = y$ and $d(x) = \{y \in J(L) \mid y \leq x\}$, cf. [44, Lemma 55]. In this case, the join-irreducibles are given by all non-zero states. The \mathcal{P} -succinct automaton consequently induced by Proposition 14 is depicted in Figure 5b; it can be recognised as the canonical RFSA, cf. e.g. [33].

Example 17. (The canonical nominal RFSA) It is not hard to see that F extends to a functor on the category of nominal sets; the usual strength function is equivariant [44, Lemma 46]; and $h^{\mathcal{P}_n}: \mathcal{P}_n 2 \rightarrow 2$ with $h^{\mathcal{P}_n}(\varphi) = \varphi(1)$ defines a \mathcal{P}_n -algebra, which induces a canonical distributive law $\lambda^{\mathcal{P}_n}$ between F and the nominal powerset monad \mathcal{P}_n . As in [29], let $\mathcal{L} = \{vawau \mid v, w, u \in A^*, a \in A\}$, then $a^{-n}\mathcal{L} =$

⁶Minimal in the sense that every state is reachable by an element of A^* and no two different states observe the same language.

⁷Minimal in the sense that every state is reachable by an element of $T(A^*)$ and no two different states observe the same language.

$a^{-2}\mathcal{L} = A^*$ for $n \geq 2$, and $v^{-1}\mathcal{L} = \cup_{a \in A} a^{-|v|_a} \mathcal{L}$, where $|v|_a$ denotes the number of a 's that occur in v . In consequence, the nominal CSL underlying the minimal pointed $\lambda^{\mathcal{P}_n}$ -bialgebra is generated by the orbit-finite nominal set of join-irreducibles $\{\mathcal{L}\} \cup \{a^{-1}\mathcal{L} \mid a \in A\} \cup \{A^*\}$, which is equipped with the obvious $\text{Perm}(A)$ -action and satisfies the inclusion $\mathcal{L} \subseteq a^{-1}\mathcal{L} \subseteq A^*$. The orbit-finite representation of the \mathcal{P}_n -succinct automaton consequently induced by Proposition 14 is depicted in Figure 6.

Example 18. (The minimal xor automaton) The \mathcal{R} -algebra structure $h^{\mathcal{R}} : \mathcal{R}2 \rightarrow 2$ with $h^{\mathcal{R}}(\varphi) = \varphi(1)$ induces a canonical distributive law $\lambda^{\mathcal{R}}$ between F and the free vector space monad \mathcal{R} over the two element field. The minimal pointed $\lambda^{\mathcal{R}}$ -bialgebra accepting $\mathcal{L} = (a+b)^*a$ is depicted in Figure 7a and coincides with the bialgebra freely generated by the F -coalgebra in Figure 2. The construction can be verified using [44, Lemma 50]. The underlying vector space structure necessarily has a basis; we choose the size-minimal generator $\langle Y, i, d \rangle$ with $Y = \{\{x\}, \{x, y\}\}$, $i(y) = y$, and $d(\emptyset) = \emptyset$, $d(\{x\}) = \{\{x\}\}$, $d(\{y\}) = \{\{x\}, \{x, y\}\}$, $d(\{x, y\}) = \{\{x, y\}\}$, which is sufficient by [44, Lemma 52]. The \mathcal{R} -succinct automaton induced by Proposition 14 is depicted in Figure 7b; it can be recognised as the minimal xor automaton, cf. e.g. [33].

5 Changing the type of succinct automata

This section contains a generalisation of the approach in Section 4. The extension is based on the observation that in the last section we implicitly considered *two* types of monads: (i) a monad S that describes the additional algebraic structure of a given deterministic automaton; and (ii) a monad T that captures the side-effects of the succinct automaton that is obtained by the generator-based translation. In Proposition 14, the main result of the last section, the monads coincided, but to recover for instance the $\hat{\text{a}}$ tomaton [12] we will have to extend Proposition 14 to a situation where S and T can differ.

5.1 Relating distributive laws

We now introduce the main technical ingredient of our extension: *distributive law homomorphisms*. As before, we present the theory on the level of arbitrary bialgebras, even though we will later focus on the case where the coalgebraic dynamics are those of deterministic automata. Distributive law homomorphisms will allow us to shift a bialgebra over a monad S to an equivalent bialgebra over a monad T , for which we can then find, analogous to Section 4, an equivalent succinct representation. The notion we use is an instance of a much more general definition that allows to relate distributive laws on two different categories. We restrict to the case where both distributive laws are given over the same behavioural endofunctor F .

Definition 19. (Distributive law homomorphism [35, 43]) Let $\lambda^S : SF \rightarrow FS$ and $\lambda^T : TF \rightarrow FT$ be distributive laws between monads S and T and an endofunctor F , respectively. A distributive law homomorphism $\alpha : \lambda^S \rightarrow \lambda^T$ consists of a natural transformation $\alpha : T \Rightarrow S$ satisfying $\mu^S \circ \alpha_S \circ T\alpha = \alpha \circ \mu^T$, $\alpha \circ \eta^T = \eta^S$ and $\lambda^S \circ \alpha_F = F\alpha \circ \lambda^T$.

The above definition is such that α induces a functor between the categories of λ^S - and λ^T -bialgebras.

Lemma 20. [10, 27] *Let $\alpha : \lambda^S \rightarrow \lambda^T$ be a distributive law homomorphism. Then $\alpha\langle X, h, k \rangle := \langle X, h \circ \alpha_X, k \rangle$ and $\alpha(f) := f$ defines a functor $\alpha : \text{Bialg}(\lambda^S) \rightarrow \text{Bialg}(\lambda^T)$.*

The next result is a straightforward consequence of Proposition 14, and may be strengthened to an isomorphism in case one is given a basis instead of a generator, analogous to Proposition 15. It can be seen as a road map to the approach we propose in this section.




Figure 7: (a) The minimal \mathbb{Z}_2 -Vect structured DFA for $\mathcal{L} = (a + b)^*a$ (freely-generated by the DFA in Figure 2); (b) Up to the choice of a basis, the minimal xor automaton for $\mathcal{L} = (a + b)^*a$.

Corollary 21. *Let $\alpha : \lambda^S \rightarrow \lambda^T$ be a homomorphism between distributive laws and $\langle X, h, k \rangle$ a λ^S -bialgebra. If $\langle Y, i, d \rangle$ is a generator for the T -algebra $\langle X, h \circ \alpha_X \rangle$, then $(h \circ \alpha_X) \circ Ti : \exp_T(\langle Y, Fd \circ k \circ i \rangle) \rightarrow \langle X, h \circ \alpha_X, k \rangle$ is a λ^T -bialgebra homomorphism.*

5.2 Deriving distributive law relations

We now turn to the procedure of deriving a distributive law homomorphism. In practice, coming up with a natural transformation and proving that it lifts to a distributive law homomorphism can be quite cumbersome.

Fortunately, for certain cases, there is a way to simplify things significantly. For instance, as the next result shows, if, as in (1), the involved distributive laws are induced by algebra structures h^S and h^T for an output set B , respectively, then one of the conditions is implied by a less convoluted constraint.

Lemma 22. *Let $\alpha : T \Rightarrow S$ be a natural transformation satisfying $h^S \circ \alpha_B = h^T$, then $\lambda^S \circ \alpha_F = F \alpha \circ \lambda^T$.*

The next result shows that for the neighbourhood monad there exists a family of *canonical* choices of distributive law homomorphisms parametrised by Eilenberg-Moore algebra structures on the output set $B = 2$. While it is well-known that such algebras induce a monad morphism, for instance in the coalgebraic modal logic community [16, 26, 38], its commutativity with canonical distributive laws has not been observed before. Moreover, we provide a new formalisation in terms of the strength function, which allows the result to be lifted to strong monads and arbitrary output objects on other categories than the one of sets and functions.

Proposition 23. *Any algebra $h : T2 \rightarrow 2$ over a set monad T induces a homomorphism $\alpha^h : \lambda^{\mathcal{H}} \rightarrow \lambda^h$ between distributive laws by $\alpha_X^h := h^{2^X} \circ \text{st} \circ T(\eta_X^{\mathcal{H}})$.*

The rest of the section is concerned with using Proposition 23 and Corollary 21 to derive canonical acceptors based on induced distributive law homomorphisms.

5.3 Example: The átomaton

We will now justify the previous informal construction of the átomaton. As hinted before, the átomaton can be recovered by relating the neighbourhood monad \mathcal{H} —whose algebras are complete *atomic* Boolean algebras (CABAs)—and the powerset monad \mathcal{P} . Formally, as a consequence of Proposition 23 we obtain the following.

Corollary 24. *Let $\alpha_X : \mathcal{P}X \rightarrow \mathcal{H}X$ satisfy $\alpha_X(\varphi)(\psi) = \bigvee_{x \in X} \varphi(x) \wedge \psi(x)$, then α constitutes a distributive law homomorphism $\alpha : \lambda^{\mathcal{H}} \rightarrow \lambda^{\mathcal{P}}$.*

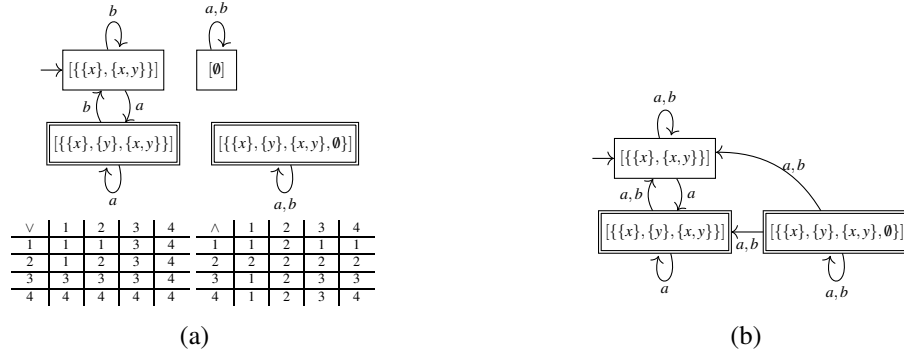


Figure 8: (a) The minimal CDL-structured DFA for $\mathcal{L} = (a+b)^*a$, where $1 \equiv [\{\{x\}, \{x,y\}\}]$, $2 \equiv [\emptyset]$, $3 \equiv [\{\{x\}, \{y\}, \{x,y\}\}]$, $4 \equiv [\{\{x\}, \{y\}, \{x,y\}, \emptyset\}]$; (b) The distromaton for $\mathcal{L} = (a+b)^*a$.

The next statement follows from a well-known Stone-type duality [41] representation theorem for CABAs.

Lemma 25. *Let $\alpha_X : \mathcal{P}X \rightarrow \mathcal{H}X$ satisfy $\alpha_X(\varphi)(\psi) = \bigvee_{x \in X} \varphi(x) \wedge \psi(x)$. If $B = \langle X, h \rangle$ is a \mathcal{H} -algebra, then $\langle \text{At}(B), i, d \rangle$ with $i(a) = a$ and $d(x) = \{a \in \text{At}(B) \mid a \leq x\}$ is a basis for the \mathcal{P} -algebra $\langle X, h \circ \alpha_X \rangle$.*

The átomaton for the regular language $\mathcal{L} = (a+b)^*a$, for example, can now be obtained as follows. First, we construct the minimal pointed $\lambda^{\mathcal{H}}$ -bialgebra accepting \mathcal{L} , which is depicted in Figure 3 together with its underlying CABA structure B . The construction can be verified with the help of [44, Lemma 48]. Using the distributive law homomorphism α of Corollary 24, it can be translated into an equivalent pointed $\lambda^{\mathcal{P}}$ -bialgebra with underlying CSL-structure $\alpha(B)$. By Lemma 25 the atoms $\text{At}(B)$ of B form a basis for $\alpha(B)$. In this case the atoms are given by $[\{\{x\}, \{x,y\}\}]$, $[\{\{y\}\}]$ and $[\{\emptyset\}]$. The \mathcal{P} -succinct automaton consequently induced by Corollary 21 is depicted in Figure 4; it can be recognised as the átomaton, cf. e.g. [33].

5.4 Example: The distromaton

We shall now use our framework to recover another canonical non-deterministic acceptor: the *distromaton* [33]. As the name suggests, it can be constructed by relating the monotone neighbourhood monad \mathcal{A} —whose algebras are completely *distributive* lattices—and the powerset monad \mathcal{P} . Formally, the relationship can be established by the same natural transformation we used for the átomaton.

Corollary 26. *Let $\alpha_X : \mathcal{P}X \rightarrow \mathcal{A}X$ satisfy $\alpha_X(\varphi)(\psi) = \bigvee_{x \in X} \varphi(x) \wedge \psi(x)$, then α constitutes a distributive law homomorphism $\alpha : \lambda^{\mathcal{A}} \rightarrow \lambda^{\mathcal{P}}$.*

The distromaton for the regular language $\mathcal{L} = (a+b)^*a$, for example, can now be obtained as follows. First, we construct the minimal pointed $\lambda^{\mathcal{A}}$ -bialgebra for \mathcal{L} , depicted in Figure 8a with its underlying CDL structure h . The construction can be verified with the help of [44, Lemma 49]. Using the distributive law homomorphism α in Corollary 26, it can be translated into an equivalent pointed $\lambda^{\mathcal{P}}$ -bialgebra with underlying CSL structure $L = h \circ \alpha_X$. Its partially ordered state space $[\emptyset] \leq [\{\{x\}, \{x,y\}\}] \leq [\{\{x\}, \{y\}, \{x,y\}\}] \leq [\{\{x\}, \{y\}, \{x,y\}, \emptyset\}]$ is necessarily finite, which turns the set of join-irreducibles into a size-minimal generator $\langle J(L), i, d \rangle$ for L , where $i(y) = y$ and $d(x) = \{y \in J(L) \mid y \leq x\}$. In this case, the join-irreducibles are given by all non-zero states. The \mathcal{P} -succinct automaton consequently induced by Corollary 21 is depicted in Figure 8b and can be recognised as the distromaton, cf. [33].




Figure 9: The minimal xor-CABA automaton is to the minimal xor automaton what the átomaton is to the canonical RFSA.

5.5 Example: The minimal xor-CABA automaton

We conclude this section by relating the neighbourhood monad \mathcal{H} with the free vector space monad \mathcal{R} over the unique two element field \mathbb{Z}_2 . In particular, we derive a new canonical succinct acceptor for regular languages, which we call the *minimal xor-CABA automaton*.

Intuitively, the next result says that every CABA can be equipped with a symmetric difference like operation that turns it into a vector space over the two element field.

Corollary 27. *Let $\alpha_X : \mathcal{R}X \rightarrow \mathcal{H}X$ satisfy $\alpha_X(\varphi)(\psi) = \bigoplus_{x \in X} \varphi(x) \cdot \psi(x)$, then α constitutes a distributive law homomorphism $\alpha : \lambda^{\mathcal{H}} \rightarrow \lambda^{\mathcal{R}}$.*

Since every vector space admits a basis, above result leads to the definition of a new acceptor of regular languages. Let α denote the homomorphism in Corollary 27 and F the endofunctor given by $FX = 2 \times X^A$.

Definition 28. (Minimal xor-CABA automaton) Let $\langle X, h, k \rangle$ be the minimal x -pointed $\lambda^{\mathcal{H}}$ -bialgebra accepting a regular language $\mathcal{L} \subseteq A^*$, and $B = \langle Y, i, d \rangle$ a basis for the \mathcal{R} -algebra $\langle X, h \circ \alpha_X \rangle$. The *minimal xor-CABA automaton* for \mathcal{L} with respect to B is the $d(x)$ -pointed \mathbb{Z}_2 -weighted automaton $Fd \circ k \circ i$.

In Figure 9 it is indicated how the canonical acceptors of this paper, including the minimal xor-CABA automaton, are based on relations between pairs of monads.

For the regular language $\mathcal{L} = (a + b)^*a$ above definition instantiates as follows. First, as for the átomaton, we construct the minimal pointed $\lambda^{\mathcal{H}}$ -bialgebra $\langle X, h, k \rangle$ for \mathcal{L} ; it is depicted in Figure 3. As one easily verifies, the \mathbb{Z}_2 -vector space $\langle X, h \circ \alpha_X \rangle$ is induced by the symmetric difference operation \oplus on subsets. Using the notation in Figure 3, we choose the basis $\langle Y, i, d \rangle$ with $Y = \{4, 6, 7, 8\}$; $i(y) = y$; and $d(1) = 7 \oplus 8$, $d(2) = \emptyset$, $d(3) = 6 \oplus 7$, $d(4) = 4$, $d(5) = 6 \oplus 7 \oplus 8$, $d(6) = 6$, $d(7) = 7$, $d(8) = 8$. The induced $d(1) = 7 \oplus 8$ -pointed \mathcal{R} -succinct automaton accepting \mathcal{L} , i.e. the minimal xor-CABA automaton, is depicted in Figure 10.

6 Minimality

In this section we restrict ourselves to the category of (nominal) sets. We show that every language satisfying a suitable property parametric in monads S and T admits a size-minimal succinct automaton of type T accepting it. As a main result we obtain Theorem 32, which is a generalisation of parts of [33, Theorem 4.8]. In Section 6.1 we instantiate the former to subsume known minimality results for canonical automata, to prove the xor-CABA automaton minimal, and to establish a size-comparison between different acceptors.

Given a distributive law homomorphism $\alpha : \lambda^S \rightarrow \lambda^T$, let $\text{ext} : \text{Coalg}(FT) \rightarrow \text{Coalg}(FS)$ be the functor given by $\text{ext}(\langle X, k \rangle) = \langle X, F\alpha_X \circ k \rangle$ and $\text{ext}(f) = f$. Moreover, let $\text{exp}_U : \text{Coalg}(FU) \rightarrow \text{Bialg}(\lambda^U)$ for $U \in \{S, T\}$ denote the functor introduced in Lemma 10.

Proposition 29. *Let $\alpha : \lambda^S \rightarrow \lambda^T$ be a distributive law homomorphism. Then $\alpha_X : TX \rightarrow SX$ underlies a natural transformation $\alpha : \text{exp}_T \Rightarrow \alpha \circ \text{exp}_S \circ \text{ext}$ between functors of type $\text{Coalg}(FT) \rightarrow \text{Bialg}(\lambda^T)$.*




Figure 10: The minimal xor-CABA automaton for $\mathcal{L} = (a + b)^*a$.

In the above situation a T -succinct automaton admits *two* semantics, induced by lifting the former either to a bialgebra over λ^S or λ^T . The next definition introduces a notion of *closedness* that captures those cases in which the image of both semantics coincides.

Definition 30. (α -closed succinct automaton) Let $\alpha : \lambda^S \rightarrow \lambda^T$ be a distributive law homomorphism. We say that a T -succinct automaton \mathcal{X} is α -closed if the unique diagonal below is an isomorphism:

$$\begin{array}{ccc}
 \text{exp}_T(\mathcal{X}) & \xrightarrow{\text{obs}} & \text{im}(\text{obs}_{\text{exp}_T(\mathcal{X})}) \\
 \text{obs} \circ \alpha_X \downarrow & \swarrow & \downarrow \\
 \text{im}(\text{obs}_{\alpha(\text{exp}_S(\text{ext}(\mathcal{X})))}) & \xrightarrow{\quad} & \Omega
 \end{array}$$

Next we show that succinct automata obtained from certain generators are α -closed.

Lemma 31. *Let $\alpha : \lambda^S \rightarrow \lambda^T$ be a distributive law homomorphism and $\langle X, h, k \rangle$ a λ^S -bialgebra. If $\langle Y, i, d \rangle$ is a generator for $\langle X, h \circ \alpha_X \rangle$, then $\langle Y, Fd \circ k \circ i \rangle$ is α -closed.*

We are now able to state our main result, which is a generalisation of parts of [33, Theorem 4.8].

Theorem 32 (Minimal succinct automata). *Given a language $\mathcal{L} \in \Omega$ such that there exists a minimal pointed λ^S -bialgebra \mathbb{M} accepting \mathcal{L} and the underlying algebra of $\alpha(\mathbb{M})$ admits a size-minimal generator, there exists a pointed α -closed T -succinct automaton \mathcal{X} accepting \mathcal{L} such that:*

- for any pointed α -closed T -succinct automaton \mathcal{Y} accepting \mathcal{L} we have that $\text{im}(\text{obs}_{\text{exp}_T(\mathcal{X})}) \subseteq \text{im}(\text{obs}_{\text{exp}_T(\mathcal{Y})})$;
- if $\text{im}(\text{obs}_{\text{exp}_T(\mathcal{X})}) = \text{im}(\text{obs}_{\text{exp}_T(\mathcal{Y})})$, then $|X| \leq |Y|$, where X and Y are the carriers of \mathcal{X} and \mathcal{Y} , respectively.

For a T -succinct automaton \mathcal{X} let us write $\text{obs}_{\mathcal{X}}^{\dagger} := \text{obs}_{\text{exp}_T(\mathcal{X})} \circ \eta_X^T : X \rightarrow \Omega$ for a generalisation of the semantics of non-deterministic automata. The next result provides an equivalent characterisation of α -closedness in terms of $\text{obs}_{\mathcal{X}}^{\dagger}$ that will be particularly useful in Section 6.1.

Lemma 33. *Let $\alpha : \lambda^S \rightarrow \lambda^T$ be a distributive law homomorphism. For any T -succinct automaton \mathcal{X} it holds that $\text{im}(\text{obs}_{\text{exp}_T(\mathcal{X})}) = \text{im}(h \circ \alpha_{\Omega} \circ T(\text{obs}_{\mathcal{X}}^{\dagger}))$ and $\text{im}(\text{obs}_{\alpha(\text{exp}_S(\text{ext}(\mathcal{X})))}) = \text{im}(h \circ S(\text{obs}_{\mathcal{X}}^{\dagger}))$, where $\langle \Omega, h, k \rangle$ is the final λ^S -bialgebra.*

6.1 Applications to canonical automata

In this section we instantiate Theorem 32 to characterise a variety of canonical acceptors from the literature as size-minimal representatives among subclasses of α -closed succinct automata, i.e. those automata whose images of the two semantics induced by α coincide. We begin with the canonical RFSA and the minimal xor automaton, for which α is the identity and α -closedness therefore is trivial.

In [13] the canonical RFSA for \mathcal{L} has been characterised as size-minimal among those NFAs accepting \mathcal{L} for which states accept a residual of \mathcal{L} . More recently, it was shown that the class in fact

can be extended to those NFAs accepting \mathcal{L} for which states accept a *union* of residuals of \mathcal{L} [33]. The next result recovers the latter as a consequence of the second point in Theorem 32. We write \overline{Y} for the algebraic closure⁸ of a subset $Y \subseteq X$ of some T -algebra X .

Corollary 34. *The canonical RFSA for \mathcal{L} is size-minimal among non-deterministic automata \mathcal{Y} accepting \mathcal{L} with $\overline{\text{im}(\text{obs}_{\mathcal{Y}}^{\dagger})}^{\text{CSL}} \subseteq \overline{\text{Der}(\mathcal{L})}^{\text{CSL}}$.*

The second condition in Theorem 32 is always satisfied for a *reachable* succinct automaton \mathcal{Y} . Since for \mathbb{Z}_2 -weighted automata it is possible to find an equivalent reachable \mathbb{Z}_2 -weighted automaton with less or equally many states (which for NFA is not necessarily the case), the minimal xor automaton is minimal among *all* \mathbb{Z}_2 -weighted automata, as was already known from for instance [42].

Corollary 35. *The minimal xor automaton for \mathcal{L} is size-minimal among \mathbb{Z}_2 -weighted automata accepting \mathcal{L} .*

For the átomaton, the distromaton, and the minimal xor-CABA automaton the distributive law homomorphism α in play is non-trivial; α -closedness translates to the below equalities between closures. In all three cases it is possible to waive the inclusion induced by the second point in Theorem 32.

Corollary 36. *The átomaton for \mathcal{L} is size-minimal among non-deterministic automata \mathcal{Y} accepting \mathcal{L} with $\overline{\text{im}(\text{obs}_{\mathcal{Y}}^{\dagger})}^{\text{CSL}} = \overline{\text{im}(\text{obs}_{\mathcal{Y}}^{\dagger})}^{\text{CABA}}$.*

The above result can be shown to be similar to [33, Theorem 4.9], which characterises the átomaton as size-minimal among non-deterministic automata whose accepted languages are *closed under complement*. The result below is very similar to a characterisation of the distromaton as size-minimal among non-deterministic automata whose accepted languages are *closed under intersection* [33, Theorem 4.13].

Corollary 37. *The distromaton for \mathcal{L} is size-minimal among non-deterministic automata \mathcal{Y} accepting \mathcal{L} with $\overline{\text{im}(\text{obs}_{\mathcal{Y}}^{\dagger})}^{\text{CSL}} = \overline{\text{im}(\text{obs}_{\mathcal{Y}}^{\dagger})}^{\text{CDL}}$.*

The size-minimality result for the newly discovered minimal xor-CABA automaton is analogous to the ones for the átomaton and the distromaton.

Corollary 38. *The minimal xor-CABA automaton for \mathcal{L} is size-minimal among \mathbb{Z}_2 -weighted automata \mathcal{Y} accepting \mathcal{L} with $\overline{\text{im}(\text{obs}_{\mathcal{Y}}^{\dagger})}^{\mathbb{Z}_2\text{-Vect}} = \overline{\text{im}(\text{obs}_{\mathcal{Y}}^{\dagger})}^{\text{CABA}}$.*

We conclude with a size-comparison between acceptors that is parametric in the closure of derivatives.

Corollary 39. • *If $\overline{\text{Der}(\mathcal{L})}^{\mathbb{Z}_2\text{-Vect}} = \overline{\text{Der}(\mathcal{L})}^{\text{CABA}}$, then the minimal xor automaton and the minimal xor-CABA automaton for \mathcal{L} are of the same size.*

- *If $\overline{\text{Der}(\mathcal{L})}^{\text{CSL}} = \overline{\text{Der}(\mathcal{L})}^{\text{CDL}}$, then the canonical RFSA and the distromaton for \mathcal{L} are of the same size.*
- *If $\overline{\text{Der}(\mathcal{L})}^{\text{CSL}} = \overline{\text{Der}(\mathcal{L})}^{\text{CABA}}$, then the canonical RFSA and the átomaton for \mathcal{L} are of the same size.*

⁸If $Y = \text{im}(f)$ for some morphism f with codomain $\langle X, h \rangle$, the closure is given by the induced T -algebra structure on $\text{im}(h \circ T f)$.

7 Related work

One of the main motivations for the present paper is provided by active learning algorithms for the derivation of succinct state-based models [2]. A major challenge in learning non-deterministic models is the lack of a canonical target acceptor for a given language [13]. The problem has been independently approached for different variants of non-determinism, often with the idea of finding a subclass admitting a unique representative [7, 15] such as e.g. the canonical RFSA, the minimal xor automaton, or the átomaton.

A more general and unifying perspective on learning automata that may not have a canonical target was given by Van Heerdt [17, 18, 20]. One of the central notions in this work is the concept of a scoop, originally introduced by Arbib and Manes [4] and here referred to as a generator. The main contribution in [20] is a general procedure to find irreducible sets of generators, which thus restricts the work to the category of sets. In the present paper we generally work over arbitrary categories, although we assume the existence of a minimal set-based generator in Theorem 32. Furthermore, the work of Van Heerdt has no size-minimality results.

Closely related to the present paper is the work of Myers et al. [33], who present a coalgebraic construction for canonical non-deterministic automata. They cover the canonical RFSA, the minimal xor automaton, the átomaton, and the distromaton. The underlying idea in [33] for finding succinct representations is similar to ours: first they build the minimal DFA for a regular language in a locally finite variety, then they apply an equivalence between the category of finite algebras and a suitable category of finite structured sets and relations. On the one hand, the category of finite algebras in a locally finite variety can be translated into our setting by considering a category of algebras over a monad preserving finite sets. In fact, modulo this translation, many of the categories considered here already appear in [33], e.g. vector spaces, Boolean algebras, complete join-semi lattices, and distributive lattices. On the other hand, their construction seems to be restricted to the category of sets and non-deterministic automata, while we work over arbitrary monads on arbitrary categories. Their work does not provide a general algorithm to construct a succinct automaton, i.e., the specifics vary with the equivalences considered, while we give a general definition and a soundness argument in Corollary 21. While Myers et al. give minimality results for a wide range of acceptors, each proof follows case-specific arguments. In Theorem 32 we provide a unifying minimality result for succinct automata that generalises parts of [33, Theorem 4.8] and subsumes most of their results [33, Theorem 4.9, Theorem 4.10, Corollary 4.11, Theorem 4.13].

8 Discussion and future work

We have presented a general categorical framework based on bialgebras and distributive law homomorphisms for the derivation of canonical automata. The framework instantiates to a wide range of well-known examples from the literature and allowed us to discover a previously unknown canonical acceptor for regular languages. Finally, we presented a theorem that subsumes previously independently proven minimality results for canonical acceptors, implied new characterisations, and allowed us to make size-comparisons between canonical automata.

In the future, we would like to cover other examples, such as the canonical probabilistic RFSA [15] and the canonical alternating RFSA [3, 7]. Probabilistic automata of the type in [15] are typically modelled as TF -coalgebras instead of FT -coalgebras [24], and thus will need a shift in perspective. For alternating RFSA we expect a canonical form can be constructed in the spirit of this paper, from generators for algebras over the neighbourhood monad, by interpreting the join-dense atoms of a CABA

as a full meet of ground elements.

Generally, it would be valuable to have a more systematic treatment of the range of available monads and distributive law homomorphisms [45], making use of the fact that distributive law homomorphisms compose.

Further generalisation in another direction could be achieved by distributive laws between monads and endofunctors on different categories. For instance, we expect that operations on automata as the product can be captured by homomorphisms between distributive laws of such more general type.

Finally, we would like to lift existing double-reversal characterisations of the minimal DFA [11], the átomaton [12], the distromaton [33], and the minimal xor automaton [42] to general canonical automata. The work in [8, 9] gives a coalgebraic generalisation of Brzozowski's algorithm based on dualities between categories, but does not cover the cases we are interested in. The framework in [1] recovers the átomaton as the result of a minimisation procedure, but does not consider other canonical acceptors.

References

- [1] Jiří Adámek, Filippo Bonchi, Mathias Hülsbusch, Barbara König, Stefan Milius & Alexandra Silva (2012): *A coalgebraic perspective on minimization and determinization*. In: *International Conference on Foundations of Software Science and Computational Structures*, Springer, pp. 58–73, doi:10.1007/978-3-642-28729-9_4.
- [2] Dana Angluin (1987): *Learning regular sets from queries and counterexamples*. *Information and computation* 75(2), pp. 87–106, doi:10.1016/0890-5401(87)90052-6.
- [3] Dana Angluin, Sarah Eisenstat & Dana Fisman (2015): *Learning regular languages via alternating automata*. In: *IJCAI*, pp. 3308–3314.
- [4] Michael A Arbib & Ernest G Manes (1975): *Fuzzy machines in a category*. *Bulletin of the Australian Mathematical Society* 13(2), pp. 169–210, doi:10.1017/S0004972700024412.
- [5] Steve Awodey (2010): *Category theory*. Oxford university press.
- [6] Jon Beck (1969): *Distributive laws*. In: *Seminar on triples and categorical homology theory*, Springer, pp. 119–140, doi:10.1007/BFb0083084.
- [7] Sebastian Berndt, Maciej Liškiewicz, Matthias Lutter & Rüdiger Reischuk (2017): *Learning residual alternating automata*. In: *Thirty-First AAAI Conference on Artificial Intelligence*.
- [8] Filippo Bonchi, Marcello M Bonsangue, Helle H Hansen, Prakash Panangaden, Jan JMM Rutten & Alexandra Silva (2014): *Algebra-coalgebra duality in Brzozowski's minimization algorithm*. *ACM Transactions on Computational Logic (TOCL)* 15(1), pp. 1–29, doi:10.1145/2490818.
- [9] Filippo Bonchi, Marcello M Bonsangue, Jan JMM Rutten & Alexandra Silva (2012): *Brzozowski's algorithm (co) algebraically*. In: *Logic and Program Semantics*, Springer, pp. 12–23, doi:10.1007/978-3-642-29485-3_2.
- [10] Marcello M Bonsangue, Helle Hvid Hansen, Alexander Kurz & Jurriaan Rot (2013): *Presenting distributive laws*. In: *International Conference on Algebra and Coalgebra in Computer Science*, Springer, pp. 95–109, doi:10.1007/978-3-642-40206-7_9.
- [11] Janusz A Brzozowski (1962): *Canonical regular expressions and minimal state graphs for definite events*. In: *Proc. Symposium of Mathematical Theory of Automata*, pp. 529–561.
- [12] Janusz A. Brzozowski & Hellis Tamm (2014): *Theory of átomata*. *Theor. Comput. Sci.* 539, pp. 13–27, doi:10.1016/j.tcs.2014.04.016.
- [13] François Denis, Aurélien Lemay & Alain Terlutte (2001): *Residual finite state automata*. In: *STACS 2001*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 144–157, doi:10.1007/3-540-44693-1_13.

- [14] Samuel Eilenberg, John C Moore et al. (1965): *Adjoint functors and triples*. *Illinois Journal of Mathematics* 9(3), pp. 381–398, doi:10.1215/ijm/1256068141.
- [15] Yann Esposito, Aurélien Lemay, François Denis & Pierre Dupont (2002): *Learning probabilistic residual finite state automata*. In: *International Colloquium on Grammatical Inference*, Springer, pp. 77–91, doi:10.1007/3-540-45790-9_7.
- [16] Helle Hvid Hansen, Clemens Kupke & Raul Andres Leal (2014): *Strong completeness for iteration-free coalgebraic dynamic logics*. In: *IFIP International Conference on Theoretical Computer Science*, Springer, pp. 281–295, doi:10.1007/978-3-662-44602-7_22.
- [17] Gerco van Heerdt (2016): *An abstract automata learning framework*. Master’s thesis, Radboud University Nijmegen.
- [18] Gerco van Heerdt (2020): *CALF: Categorical automata learning framework*. Ph.D. thesis, University College London.
- [19] Gerco van Heerdt, Joshua Moerman, Matteo Sammartino & Alexandra Silva (2019): *A (co)algebraic theory of succinct automata*. *J. Log. Algebraic Methods Program.* 105, pp. 112–125, doi:10.1016/j.jlamp.2019.02.008.
- [20] Gerco van Heerdt, Matteo Sammartino & Alexandra Silva (2020): *Learning automata with side-effects*. In: *Coalgebraic Methods in Computer Science - 15th IFIP WG 1.3 International Workshop, CMCS 2020, Colocated with ETAPS 2020, Dublin, Ireland, April 25-26, 2020, Proceedings, Lecture Notes in Computer Science 12094*, Springer, pp. 68–89, doi:10.1007/978-3-030-57201-3_5.
- [21] Bart Jacobs (2006): *A bialgebraic review of deterministic automata, regular expressions and languages*, pp. 375–404. Springer Berlin Heidelberg, Berlin, Heidelberg, doi:10.1007/11780274_20.
- [22] Bart Jacobs (2011): *Bases as coalgebras*. In: *International conference on algebra and coalgebra in computer science*, Springer, pp. 237–252, doi:10.2168/LMCS-9(3:23)2013.
- [23] Bart Jacobs (2015): *A recipe for state-and-effect triangles*. In: *6th Conference on Algebra and Coalgebra in Computer Science (CALCO 2015)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, doi:10.4230/LIPIcs.CALCO.2015.116.
- [24] Bart Jacobs, Alexandra Silva & Ana Sokolova (2012): *Trace semantics via determinization*. In: *International Workshop on Coalgebraic Methods in Computer Science*, Springer, pp. 109–129, doi:10.1007/978-3-642-32784-1_7.
- [25] Bart Jacobs, Alexandra Silva & Ana Sokolova (2015): *Trace semantics via determinization*. *Journal of Computer and System Sciences* 81(5), pp. 859–879, doi:10.1016/j.jcss.2014.12.005.
- [26] Bartek Klin (2004): *A coalgebraic approach to process equivalence and a coinduction principle for traces*. *Electronic Notes in Theoretical Computer Science* 106, pp. 201–218, doi:10.1016/j.entcs.2004.02.029.
- [27] Bartek Klin & Beata Nachyla (2015): *Presenting morphisms of distributive laws*. In: *6th Conference on Algebra and Coalgebra in Computer Science (CALCO 2015)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [28] Fred EJ Linton (1966): *Some aspects of equational categories*. In: *Proceedings of the Conference on Categorical Algebra*, Springer, pp. 84–94, doi:10.1007/978-3-642-99902-4_3.
- [29] Joshua Moerman & Matteo Sammartino (2020): *Residual nominal automata*. In: *CONCUR*, 171, pp. 44:1–44:21, doi:10.4230/LIPIcs.CONCUR.2020.44.
- [30] Eugenio Moggi (1988): *Computational lambda-calculus and monads*. University of Edinburgh, Department of Computer Science, Laboratory for Foundations of Computer Science.
- [31] Eugenio Moggi (1990): *An abstract view of programming languages*. University of Edinburgh, Department of Computer Science, Laboratory for Foundations of Computer Science.
- [32] Eugenio Moggi (1991): *Notions of computation and monads*. *Information and computation* 93(1), pp. 55–92, doi:10.1016/0890-5401(91)90052-4.

- [33] Robert S. R. Myers, Jiri Adamek, Stefan Milius & Henning Urbat (2015): *Coalgebraic constructions of canonical nondeterministic automata*. *Theor. Comput. Sci.* 604, pp. 81–101, doi:10.1016/j.tcs.2015.03.035.
- [34] Anil Nerode (1958): *Linear automaton transformations*. *Proceedings of the American Mathematical Society* 9(4), pp. 541–544, doi:10.1090/S0002-9939-1958-0135681-9.
- [35] John Power & Hiroshi Watanabe (2002): *Combining a monad and a comonad*. *Theoretical Computer Science* 280(1-2), pp. 137–162, doi:10.1016/S0304-3975(01)00024-X.
- [36] Jan Rutten, Marcello Bonsangue, Filippo Bonchi & Alexandra Silva (2013): *Generalizing determinization from automata to coalgebras*. *Logical Methods in Computer Science* 9, doi:10.2168/LMCS-9(1:9)2013.
- [37] Jan JMM Rutten (2000): *Universal coalgebra: a theory of systems*. *Theoretical computer science* 249(1), pp. 3–80, doi:10.1016/S0304-3975(00)00056-6.
- [38] Lutz Schröder (2008): *Expressivity of coalgebraic modal logic: The limits and beyond*. *Theoretical Computer Science* 390(2-3), pp. 230–247, doi:10.1016/j.tcs.2007.09.023.
- [39] Alexandra Silva, Filippo Bonchi, Marcello M Bonsangue & Jan JMM Rutten (2010): *Generalizing the powerset construction, coalgebraically*. In: *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2010)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [40] Ross Street (2009): *Weak distributive laws*. *Theory and Applications of Categories [electronic only]* 22, pp. 313–320.
- [41] Paul Taylor (2002): *Subspaces in abstract stone duality*. *Theory and Applications of Categories* 10(13), pp. 301–368.
- [42] Jean Vuillemin & Nicolas Gama (2010): *Efficient equivalence and minimization for non deterministic xor automata*.
- [43] Hiroshi Watanabe (2002): *Well-behaved translations between structural operational semantics*. *Electronic Notes in Theoretical Computer Science* 65(1), pp. 337–357, doi:10.1016/S1571-0661(04)80372-4.
- [44] Stefan Zetsche, Gerco van Heerdt, Matteo Sammartino & Alexandra Silva (2021): *Canonical automata via distributive law homomorphisms*. *arXiv preprint arXiv:2104.13421*.
- [45] Maaïke Zwart & Dan Marsden (2019): *No-go theorems for distributive laws*. In: *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, IEEE, pp. 1–13, doi:10.1109/LICS.2019.8785707.