

Canonical Nondeterministic Automata

Robert S. R. Myers¹, Jiří Adámek¹, Stefan Milius² and Henning Urbat¹

¹ Institut für Theoretische Informatik
Technische Universität Braunschweig, Germany

² Lehrstuhl für Theoretische Informatik
Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany

Abstract. For each regular language L we describe a family of canonical nondeterministic acceptors (nfas). Their construction follows a uniform recipe: build the minimal dfa for L in a locally finite variety \mathcal{V} , and apply an equivalence between the finite \mathcal{V} -algebras and a category of finite structured sets and relations. By instantiating this to different varieties we recover three well-studied canonical nfes (the *átomaton*, the *jiromaton* and the minimal xor automaton) and obtain a new canonical nfa called the *distromaton*. We prove that each of these nfes is minimal relative to a suitable measure, and give conditions for state-minimality. Our approach is coalgebraic, exhibiting additional structure and universal properties.

1 Introduction

One of the core topics in classical automata theory is the construction of state-minimal acceptors for a given regular language. It is well-known that the difficulty of this task depends on whether one has deterministic or nondeterministic acceptors in mind. First, every regular language $L \subseteq \Sigma^*$ is accepted by a unique minimal *deterministic* finite automaton (dfa): its states Q_L are the derivatives of L , i.e.,

$$Q_L = \{w^{-1}L : w \in \Sigma^*\} \quad \text{where } w^{-1}L = \{v \in \Sigma^* : wv \in L\},$$

the transitions are $K \xrightarrow{a} a^{-1}K$ for $K \in Q_L$ and $a \in \Sigma$, the initial state is L , and a state is final iff it contains the empty word. This construction is due to Brzozowski [10], and is the basis for efficient dfa minimization algorithms. For *nondeterministic* finite automata (nfes) the situation is significantly more complex: a regular language may have many non-isomorphic state-minimal nfes, and generally there is no way to identify a “canonical” one among them. However, several authors have recently proposed nondeterministic acceptors that *are* in some sense canonical (though not necessarily state-minimal), e.g. the *átomaton* of Brzozowski and Tamm [9], the *jiromaton*¹ of Denis, Lemay and Terlutte [11], and the *minimal xor automaton* of Vuillemin and Gama [19]. In each case, the respective nfa is formed by closing the set Q_L of derivatives under certain algebraic operations and taking a minimal set of generators as states. Specifically,

¹ In [11] the authors called their acceptor “canonical residual finite state automaton”. We propose the shorter “jiromaton” because this is analogous to the *átomaton* terminology.

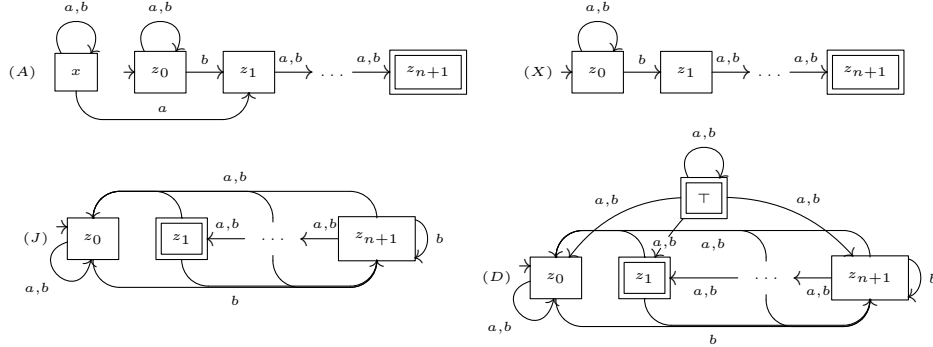
1. the states of the átomaton are the atoms of the boolean algebra generated by Q_L , obtained by closing Q_L under finite union, finite intersection and complement;
2. the states of the jiomaton are the join-irreducibles of the join-semilattice generated by Q_L , obtained by closing Q_L under finite union;
3. the states of the minimal xor automaton form a basis for the \mathbb{Z}_2 -vector space generated by Q_L , obtained by closing Q_L under symmetric difference.

In this paper we demonstrate that all these canonical nfas arise from a coalgebraic construction. For this purpose we first consider *deterministic* automata interpreted in a locally finite variety \mathcal{V} , where *locally finite* means that finitely generated algebras are finite. A *deterministic \mathcal{V} -automaton* is a coalgebra for the functor $T_{\mathcal{Z}} = \mathbb{2} \times \text{Id}^{\mathcal{Z}}$ on \mathcal{V} , for a fixed two-element algebra $\mathbb{2}$. In Section 2 we describe a Brzozowski-like construction that yields, for every regular language, the minimal deterministic finite \mathcal{V} -automaton accepting it. Next, for certain varieties \mathcal{V} of interest, we derive an equivalence between the full subcategory \mathcal{V}_f of finite algebras and a suitable category $\overline{\mathcal{V}}$ of finite structured sets, whose morphisms are relations preserving the structure. In each case, the objects of $\overline{\mathcal{V}}$ are “small” representations of their counterparts in \mathcal{V}_f , based on specific generators of algebras in \mathcal{V}_f . The equivalence $\mathcal{V}_f \cong \overline{\mathcal{V}}$ then induces an equivalence between deterministic finite \mathcal{V} -automata and coalgebras in $\overline{\mathcal{V}}$ which are *nondeterministic* automata.

This suggests a two-step procedure for constructing a canonical nfa for a given regular language L : (i) form L 's minimal deterministic \mathcal{V} -automaton, and (ii) use the equivalence of \mathcal{V}_f and $\overline{\mathcal{V}}$ to obtain an equivalent nfa. Applying this to different varieties \mathcal{V} yields the three canonical nfas mentioned above. For the átomaton one takes $\mathcal{V} = \text{BA}$ (boolean algebras). Then the minimal deterministic BA-automaton for L arises from the minimal dfa by closing its states Q_L under boolean operations. The category $\overline{\mathcal{V}} = \overline{\text{BA}}$ is based on Stone duality: $\overline{\text{BA}}$ is the dual of the category of finite sets, so it has a objects all finite sets, as morphisms all converse-functional relations, and the equivalence functor $\text{BA}_f \rightarrow \overline{\text{BA}}$ maps each finite boolean algebra to the set of its atoms. This equivalence applied to the minimal deterministic BA-automaton for L gives precisely L 's átomaton. Similarly, by taking $\mathcal{V} = \text{join-semilattices}$ and $\mathcal{V} = \text{vector spaces over } \mathbb{Z}_2$ and describing a suitable equivalence $\mathcal{V}_f \cong \overline{\mathcal{V}}$, we recover the jiomaton and the minimal xor automaton, respectively. Finally, for $\mathcal{V} = \text{distributive lattices}$ we get a new canonical nfa called the *distromaton*, which bears a close resemblance to the universal automaton [15].

Example 1.1. Consider the language $L = (a + b)^*b(a + b)^n$ where $n \in \omega$. Its minimal dfa has $\geq 2^n$ states and its (A) átomaton, (X) minimal xor automaton, (J) jiomaton and

(D) distromaton are the nfas with $\leq n + 3$ states depicted below (see Section 3.3).



The minimal xor automaton accepts L by \mathbb{Z}_2 -weighted acceptance, which is the usual acceptance in this case. It is a state-minimal nfa, as is the jromaton. The state-minimality of the latter follows from a general result (Theorem 4.4).

Generally, the sizes of the four canonical nfas and the minimal dfa are related as follows:

- (a) all the four canonical nfas can have exponentially fewer states than the minimal dfa;
- (b) the minimal xor automaton and jromaton have no more states than the minimal dfa;
- (c) the átomaton and distromaton have the same number of states, although their structure can be very different.

In Section 4 we characterize the átomaton, jromaton, minimal xor automaton and distromaton by a minimality property. This provides an explanation of the canonicity of these acceptors that is missing in the original papers. We then use this additional structure to identify conditions on regular languages that guarantee the *state-minimality* of the canonical nfas. That is, there exists a natural class of languages where *canonical* state-minimal nfas exist and can be computed relatively easily.

Related work. Our paper unifies the constructions of canonical nfas given in [9, 11, 19] from a coalgebraic perspective. Previously, several authors have studied coalgebraic methods for constructing minimal and canonical representatives of machines, including Adámek, Bonchi, Hülsbusch, König, Milius and Silva [1], Adámek, Milius, Moss and Sousa [2] and Bezhanishvili, Kupke and Panangaden [5]. Only the first of these three papers, however, treats the case of nondeterministic automata explicitly – in particular, there the átomaton is recovered as an instance of projecting coalgebras in a Kleisli category into a reflective subcategory. This approach is methodologically rather different from the present paper where a categorical equivalence (rather than a reflection) is the basis for the construction of nfas.

In [9] the authors propose a surprisingly simple algorithm for constructing the átomaton of a language L : take the minimal dfa for L 's reversed language, and reverse this dfa. These steps form a fragment of a classical dfa minimization algorithm due to Brzozowski. Recently Bonchi, Bonsangue, Rutten and Silva [7] gave a (co-)algebraic explanation of this procedure, based on the classical duality between observability and reachability of dfas. We provide another explanation in Section 3.3.

A coalgebraic treatment of linear weighted automata (of which xor automata considered here are a special case) appears in [6]; this paper also provides a procedures for computing the minimal linear weighted automaton.

Finally, our work is somewhat related to work on coalgebraic trace semantics [12]. However, while that work considers coalgebras whose carrier is a the free algebra of a variety we consider coalgebras whose carriers are arbitrary algebras from the given variety; this means we consider coalgebras over an Eilenberg-Moore category (cf. [8, 13]).

2 Deterministic Automata

We start with recalling the concept of a finite automaton. Throughout this paper let us fix a finite input alphabet Σ .

- Definition 2.1.** (a) A nondeterministic finite automaton (nfa) is a triple $N = (Z, R_a, F)$ consisting of a finite set Z of states, transition relations $R_a \subseteq Z \times Z$ for each $a \in \Sigma$ and final states $F \subseteq Z$. Morphisms of nfes are the usual bisimulations, i.e., relations that preserve and reflect transitions and final states. If N is equipped with initial states $I \subseteq Z$ we write $N = (Z, R_a, F, I)$. In this case, N accepts a language $\mathcal{L}_N(I) \subseteq \Sigma^*$ in the usual way.
- (b) A deterministic finite automaton (dfa) is an nfa with a single initial state whose transition relations are functions.

Although the goal of our paper is constructing canonical nondeterministic automata, we first consider deterministic ones from a coalgebraic perspective. Given an endofunctor $T : \mathcal{V} \rightarrow \mathcal{V}$ of a category \mathcal{V} , a T -coalgebra (Q, γ) consists of a \mathcal{V} -object Q and a \mathcal{V} -morphism $\gamma : Q \rightarrow TQ$. A coalgebra homomorphism into another coalgebra $\gamma' : Q' \rightarrow TQ'$ is a \mathcal{V} -morphism $h : Q \rightarrow Q'$ such that $Th \circ \gamma = \gamma' \circ h$. This defines a category $\text{Coalg}(T)$. If it exists, its terminal object νT is called the *final T -coalgebra*.

Assumption 2.2. From now on \mathcal{V} is a locally finite variety with a specified two-element algebra $\mathfrak{2} = \{0, 1\}$. That is, \mathcal{V} is the category of algebras for some finitary signature and equations, its morphisms being the usual algebra homomorphisms. That \mathcal{V} is *locally finite* means its finitely generated algebras are finite, equivalently its finitely generated *free* algebras are finite.

- Example 2.3.* (a) The category Set_* of pointed sets is a locally finite variety, given by the signature with a constant 0 and no equations. Let $\mathfrak{2} \in \text{Set}_*$ have point 0.
- (b) The category BA of boolean algebras is a locally finite variety: a boolean algebra on n generators has at most 2^{2^n} elements. $\mathfrak{2}$ is the 2-chain $0 < 1$.
- (c) The category $\text{Vect}(\mathbb{Z}_2)$ of vector spaces over the binary field \mathbb{Z}_2 is a locally finite variety. Here $\mathfrak{2} = \mathbb{Z}_2$ as a one-dimensional vector space.
- (d) The category JSL of (join-)semilattices with a least element 0 is locally finite: the finite powerset $\mathcal{P}_f X$ is the free semilattice on X , so a semilattice on n generators has at most 2^n elements. $\mathfrak{2}$ is the 2-chain $0 < 1$.
- (e) The category DL of distributive lattices with a least and largest element \perp and \top is locally finite. Again, $\mathfrak{2}$ is the 2-chain $0 < 1$.

Definition 2.4. If Q is a join-semilattice then $q \in Q$ is join-irreducible if (i) $q \neq 0$ and (ii) $q = r \vee r'$ implies $q = r$ or $q = r'$. The set of join-irreducibles is written $J(Q) \subseteq Q$.

Definition 2.5. A T -coalgebra (Q', γ') is a subcoalgebra of (Q, γ) if there exists an injective coalgebra homomorphism $m : (Q', \gamma') \hookrightarrow (Q, \gamma)$, and a quotient coalgebra of (Q, γ) if there exists a surjective coalgebra homomorphism $e : (Q, \gamma) \twoheadrightarrow (Q', \gamma')$.

Definition 2.6. A deterministic \mathcal{V} -automaton is a coalgebra for the functor

$$T_\Sigma : \mathcal{V} \rightarrow \mathcal{V}, \quad T_\Sigma = \mathbb{2} \times \text{Id}^\Sigma = \mathbb{2} \times \text{Id} \times \cdots \times \text{Id}.$$

Remark 2.7. Hence, by the universal property of the product, a deterministic \mathcal{V} -automaton $Q \rightarrow \mathbb{2} \times Q^\Sigma$ is given by an algebra Q of states, a \mathcal{V} -morphism $\gamma_\varepsilon : Q \rightarrow \mathbb{2}$ defining final states via $\gamma_\varepsilon^{-1}(\{1\})$ and, for each $a \in \Sigma$, a \mathcal{V} -morphism $\gamma_a : Q \rightarrow Q$ representing the a -transitions. In particular, deterministic Set-automata are precisely the classical (possibly infinite) deterministic automata without initial states, shortly *da*'s.

- Example 2.8.* (a) A deterministic Set_* -automaton is a *da* whose carrier is a pointed set and whose point is a non-final sink state; these are the partial automata of [18].
- (b) A deterministic BA-automaton is a *da* with a boolean algebra structure on the states Q such that (i) the final states form an ultrafilter, (ii) $q \xrightarrow{a} q'$ and $r \xrightarrow{a} r'$ implies $q \vee r \xrightarrow{a} q' \vee r'$ and $\neg q \xrightarrow{a} \neg q'$, and (iii) \perp is a non-final sink state.
- (c) A deterministic $\text{Vect}(\mathbb{Z}_2)$ -automaton is a *da* with a \mathbb{Z}_2 -vector space structure on the states Q such that (i) the final states $F \subseteq Q$ satisfy $0 \notin F$ and also $q + r \in F$ iff either $q \in F$ or $r \in F$ but not both, (ii) $q \xrightarrow{a} q'$ and $r \xrightarrow{a} r'$ implies $q + r \xrightarrow{a} r + r'$, and (iii) 0 is a non-final sink state.
- (d) A deterministic JSL-automaton is a *da* with a join-semilattice structure on the states Q such that (i) the final states form a prime filter, (ii) $q \xrightarrow{a} q'$ and $r \xrightarrow{a} r'$ implies $q + r \xrightarrow{a} q' + r'$, and (iii) 0 is a non-final sink state. Recall that a *prime filter* is an upwards closed $F \subseteq Q$ where $0 \notin F$ and $q + q' \in F$ iff $q \in F$ or $q' \in F$.
- (e) A deterministic DL-automaton is a *da* with a distributive lattice structure on the states Q such that (i) the final states form a prime filter, (ii) $q \xrightarrow{a} q'$ and $r \xrightarrow{a} r'$ implies $q \vee r \xrightarrow{a} q' \vee r'$ and $q \wedge r \xrightarrow{a} q' \wedge r'$, and (iii) \perp is a non-final sink state and \top is a final one.

Remark 2.9. For finitary endofunctors T , Milius [16] introduced the concept of a locally finitely presentable coalgebra: it is a filtered colimit of coalgebras carried by finitely presentable objects. In the present context the finitely presentable objects are precisely the finite algebras in \mathcal{V} , so we speak about *locally finite coalgebras*. A T_Σ -coalgebra is locally finite iff from each state only finitely many states are reachable by transitions.

Remark 2.10. 1. The final T_Σ -coalgebra in Set is $\nu T_\Sigma = \mathcal{P}\Sigma^*$, the set of formal languages over Σ , with transitions $L \xrightarrow{a} a^{-1}L$ for $a \in \Sigma$ and final states precisely those languages containing ε . Importantly, νT_Σ arises as the ω^{op} -limit of T_Σ 's terminal sequence $(T_\Sigma^n 1)_{n < \omega}$, see [4]. Since for any variety \mathcal{V} the forgetful functor from \mathcal{V} to Set creates limits, the final T_Σ -coalgebra νT_Σ in \mathcal{V} exists and lifts the one in Set, so νT_Σ has underlying set $\mathcal{P}\Sigma^*$ and the transitions and final states are as above.

2. The *final locally finite* T_Σ -coalgebra is denoted by ρT_Σ . In $\mathcal{V} = \text{Set}$ this is the subalgebra of $\nu T_\Sigma = \mathcal{P}\Sigma^*$ given by the set of all regular languages over Σ . This generalizes to any locally finite variety \mathcal{V} : ρT_Σ is a subcoalgebra of νT_Σ and its underlying set is the set of regular languages.

Example 2.11. (a) In Set_* the carrier of the final coalgebra νT_Σ has the constant \emptyset , which ρT_Σ inherits.

- (b) In BA, νT_Σ has the usual set-theoretic boolean algebra structure. The principal filter $\uparrow\varepsilon$ is an ultrafilter and the transition maps $L \mapsto a^{-1}L$ are boolean morphisms.
- (c) In $\text{Vect}(\mathbb{Z}_2)$ the vector space structure on νT_Σ and ρT_Σ is given by symmetric difference and \emptyset is the zero vector.
- (d) In JSL the join-semilattice structure on νT_Σ is union and \emptyset . The final states form a one-generated upset $\uparrow\varepsilon$ which is a prime filter because the language $\{\varepsilon\}$ is join-irreducible in νT_Σ . The transition maps are join-semilattice morphisms.
- (e) In DL we have the usual set-theoretic lattice structure on νT_Σ . The final states form a prime filter and the transition maps are lattice morphisms.

Notation 2.12. Let (Q, γ) be a locally finite T_Σ -coalgebra. The unique coalgebra homomorphism into ρT_Σ is written:

$$\mathcal{L}_\gamma : Q \rightarrow \rho T_\Sigma.$$

The function \mathcal{L}_γ sends $q \in Q$ to the regular language $\mathcal{L}_\gamma(q) \subseteq \Sigma^*$ the state q accepts.

Definition 2.13. Let $V \in \mathcal{V}$ denote the free algebra on one generator g . Then a pointed T_Σ -coalgebra (Q, γ, q_0) is a T_Σ -coalgebra (Q, γ) with a morphism $q_0 : V \rightarrow Q$. The latter may be viewed as the initial state $q_0(g) \in Q$. The language accepted by (Q, γ, q_0) is $\mathcal{L}_\gamma(q_0)$. We say that (Q, γ, q_0) is

1. *reachable* if it is generated by q_0 , i.e., no proper subcoalgebra contains q_0 ;
2. *simple* if it has no proper quotients, i.e., for every quotient coalgebra $e : (Q, \gamma) \rightarrow (Q', \gamma')$ the map e is bijective;
3. *minimal* if it is reachable and simple.

Lemma 2.14. (Q, γ, q_0) is reachable iff the algebra Q is generated by those $q \in Q$ reachable from q_0 by transitions. It is simple iff \mathcal{L}_γ is injective.

Brozowski's construction of the minimal dfa for a regular language (see Introduction) generalizes to deterministic \mathcal{V} -automata as follows:

Construction 2.15. For any regular language $L \subseteq \Sigma^*$ let $A_\mathcal{V}^L$ be the pointed T_Σ -coalgebra (Q_L, γ, L) where:

1. Q_L is the subalgebra of $\nu T_\Sigma = \mathcal{P}\Sigma^*$ generated by all derivatives $w^{-1}L$ ($w \in \Sigma^*$).
2. The transitions are $K \xrightarrow{a} a^{-1}K$ for $a \in \Sigma$ and $K \in Q_L$.
3. $K \in Q_L$ is final iff $\varepsilon \in K$.

Lemma 2.16. For every regular language $L \subseteq \Sigma^*$, $A_\mathcal{V}^L$ is a well-defined finite pointed T_Σ -coalgebra.

Proof. L is regular so it has only finitely many distinct derivatives $w^{-1}L$. Hence Q_L is a finite algebra because \mathcal{V} is a locally finite variety. It remains to show that $\gamma_a : Q_L \rightarrow Q_L$ and $\gamma_\varepsilon : Q_L \rightarrow \mathbb{2}$ as specified in points 2. and 3. are well-defined \mathcal{V} -morphisms. Recall the final locally finite T_Σ -coalgebra $(\rho T_\Sigma, \gamma_\rho)$. Then

$$\gamma_\varepsilon = Q_L \hookrightarrow \rho T_\Sigma \xrightarrow{(\gamma_\rho)_\varepsilon} \mathbb{2}$$

is a \mathcal{V} -morphism since ρT_Σ is a lifting of the da of regular languages, see Remark 2.10. Furthermore $(\gamma_\rho)_a : \rho T_\Sigma \rightarrow \rho T_\Sigma$ is defined $(\gamma_\rho)_a(K) = a^{-1}K$ i.e. the derivative $a^{-1}(-)$ preserves the algebraic operations. Thus Q_L is closed under derivatives, so γ_a is a well-defined algebra morphism. \square

- Example 2.17.* (a) In Set_* , we have $Q_L = \{\emptyset\} \cup \{w^{-1}L : w \in \Sigma^*\}$.
 (b) In BA, Q_L is the closure of $\{\emptyset\} \cup \{w^{-1}L : w \in \Sigma^*\}$ under union and complement.
 (c) In $\text{Vect}(\mathbb{Z}_2)$, Q_L is the closure of $\{w^{-1}L : w \in \Sigma^*\}$ under symmetric difference.
 (d) In JSL, Q_L is the closure of $\{\emptyset\} \cup \{w^{-1}L : w \in \Sigma^*\}$ under union.
 (e) In DL, Q_L is the closure of $\{\emptyset, \Sigma^*\} \cup \{w^{-1}L : w \in \Sigma^*\}$ under union and intersection.

Remark 2.18. The category $\text{Coalg}(T_\Sigma)$ of T_Σ -coalgebras has a factorization system (surjective homomorphism, injective homomorphism) lifting the usual factorization system (surjective, injective) = (regular epi, mono) in \mathcal{V} .

Construction 2.19 (see [2]). These factorizations give a two-step minimization of any finite pointed T_Σ -coalgebra (Q, γ, q_0) :

1. Construct the reachable subcoalgebra $(R, \delta) \hookrightarrow (Q, \gamma)$ generated by q_0 .
2. Factorize the unique T_Σ -coalgebra homomorphism $\mathcal{L}_\delta : (R, \delta) \rightarrow (\rho T_\Sigma, \gamma_\rho)$ as:

$$(R, \delta) \xrightarrow{s} (R', \delta') \xrightarrow{m} (\rho T_\Sigma, \gamma_\rho)$$

Then $(R', \delta', s(q_0))$ is minimal.

Theorem 2.20. *Let $L \subseteq \Sigma^*$ be a regular language. Then $A_{\mathcal{V}}^L$ is (up to isomorphism) the unique minimal pointed \mathcal{V} -automaton accepting L . It arises from any pointed finite \mathcal{V} -automaton (Q, γ, q_0) accepting L by Construction 2.19.*

Proof. Viewed as a da, $A_{\mathcal{V}}^L$ is a subautomaton of the da ρT_Σ of regular languages. Then the state L accepts L . It is reachable because every state is a \mathcal{V} -algebraic combination of those states reachable from L by transitions i.e. L 's derivatives. It is simple because different states accept different languages, so it is minimal.

Now let (Q, γ, q_0) be any pointed T_Σ -coalgebra accepting L and (R, δ, q_0) its reachable subautomaton, so every $q' \in R$ arises as a \mathcal{V} -algebraic combination of those states reachable from q_0 by transitions. Now $\mathcal{L}_\delta : R \rightarrow \rho T_\Sigma$ is an automata morphism, so the languages of states reachable from q_0 are precisely the derivatives of L . Since \mathcal{L}_δ is an algebra morphism its image is Q_L . \square

3 From Deterministic to Nondeterministic Automata

We now know that each regular language L has *many* canonical deterministic acceptors: one for each locally finite variety \mathcal{V} containing a two-element algebra $\mathbb{2}$. However this canonical acceptor $A_{\mathcal{V}}^L$ is generally larger than the minimal dfa in Set because one has to close under the \mathcal{V} -algebraic operations on the regular languages. In this section we will show how these larger *deterministic* machines induce smaller *nondeterministic* ones. Let us outline our approach:

1. We restrict attention to finite da's in \mathcal{V} , i.e., T_{Σ} -coalgebras with finite carrier.
2. For each of our varieties \mathcal{V} of interest, we describe an equivalence G of categories between the finite algebras \mathcal{V}_f and another category $\overline{\mathcal{V}}$ where (i) $\overline{\mathcal{V}}$'s objects are "small" representations of their counterparts in \mathcal{V}_f , and (ii) $\overline{\mathcal{V}}$'s morphisms are relations, not functions (see Lemmas 3.4, 3.8 and 3.10).
3. From G we derive equivalences \mathbb{G} and \mathbb{G}_* between (pointed) deterministic finite \mathcal{V} -automata and (pointed) coalgebras in $\overline{\mathcal{V}}$ which are *nondeterministic* finite automata, see Lemma 3.17.
4. Applying this equivalence to the minimal deterministic \mathcal{V} -automaton $A_{\mathcal{V}}^L$ gives a canonical nondeterministic acceptor for L . This is illustrated in Section 3.3.

3.1 The Equivalence between \mathcal{V}_f and $\overline{\mathcal{V}}$

For each of our varieties \mathcal{V} of interest there is a well-known description of the dual category of \mathcal{V}_f : we have Stone duality ($\text{BA}_f \cong \text{Set}_f^{op}$), Priestley duality ($\text{DL}_f \cong \text{Poset}_f^{op}$), where Poset_f is the category of finite posets and monotone functions, and the self-dualities $\text{JSL}_f \cong \text{JSL}_f^{op}$ and $\text{Vect}_f(\mathbb{Z}_2) \cong \text{Vect}_f(\mathbb{Z}_2)^{op}$. We now describe each of these dually equivalent categories as a category $\overline{\mathcal{V}}$ of finite structured sets and relations. The idea is to represent the finite algebras in \mathcal{V} in terms of a minimal set of generators.

- Example 3.1.* (a) For any $Q \in \text{Set}_*$ the subset $Q \setminus \{0\}$ generates Q ; that means that we can always drop one element.
- (b) Any finite boolean algebra $Q \in \text{BA}_f$ is generated by its atoms $\text{At}(Q)$, these being the join-irreducible elements.
- (c) Any finite join-semilattice $Q \in \text{JSL}_f$ is generated by its join-irreducibles $J(Q)$.
- (d) A finite dimensional vector space $Q \in \text{Vect}_f(\mathbb{Z}_2)$ is generated by any basis $B \subseteq Q$, although there is no canonical choice of a basis.
- (e) Any finite distributive lattice $Q \in \text{DL}_f$ is generated by its join-irreducibles $J(Q)$.

In the case of Set_{*f} , BA_f and $\text{Vect}_f(\mathbb{Z}_2)$ we can replace each algebra by a set of generators and each algebra morphism by a relation between these generators.

Definition 3.2. Let $\overline{\text{Set}}_*$ be the category Par_f of finite sets and partial functions. $\overline{\text{BA}}$ is obtained from the category Rel_f of finite sets and relations by restricting to relations whose converse is a function. Finally $\overline{\text{Vect}}(\mathbb{Z}_2)$ has the same objects and morphisms as Rel_f although now the composition of $R_1 \subseteq X \times Y$ and $R_2 \subseteq Y \times Z$ is defined by

$$R_2 \bullet R_1 := \{(x, z) : |\{y : (x, y) \in R_1, (y, z) \in R_2\}| \text{ is odd}\}.$$

Notation 3.3. Given a basis GQ of a vector space Q , for each basis vector $z \in GQ$ denote by $\pi_z : Q \rightarrow \{0, 1\}$ the projection onto the z -coordinate.

Lemma 3.4. *The following functors G are equivalences of categories where $f : Q \rightarrow Q'$ is any \mathcal{V}_f -morphism:*

1. $G : \text{Set}_{*f} \rightarrow \text{Par}_f$ defined by

$$GQ = Q \setminus \{0\} \quad Gf(z) = \begin{cases} f(z) & \text{if } f(z) \neq 0, \\ \text{undefined} & \text{otherwise.} \end{cases}$$

2. $G : \text{BA}_f \rightarrow \overline{\text{BA}}$ where $GQ = \text{At}(Q)$ is the set of atoms and $Gf = \{(z, z') \in \text{At}(Q) \times \text{At}(Q') : z' \leq_{Q'} f(z)\}$.

3. $G : \text{Vect}_f(\mathbb{Z}_2) \rightarrow \overline{\text{Vect}(\mathbb{Z}_2)}$ where GQ chooses a basis and $Gf = \{(z, z') \in GQ \times GQ' : \pi_{z'} \circ f(z) = 1\}$.

Finite join-semilattices are represented using closure spaces:

Definition 3.5. *For any set X a closure operator (shortly, a closure) on X is a function $\text{cl}_X : \mathcal{P}X \rightarrow \mathcal{P}X$ such that for all $S, S' \subseteq X$:*

$$\frac{S \subseteq S'}{\text{cl}_X(S) \subseteq \text{cl}_X(S')}, \quad \text{cl}_X(S) \supseteq S, \quad \text{cl}_X \circ \text{cl}_X = \text{cl}_X.$$

A closure space $X = (X, \text{cl}_X)$ is a set with a closure defined on it. It is finite if X is finite, strict if $\text{cl}_X(\emptyset) = \emptyset$, separable if $x \neq x'$ implies $\text{cl}_X(x) \neq \text{cl}_X(x')$, and topological if $\text{cl}_X(A \cup B) = \text{cl}_X(A) \cup \text{cl}_X(B)$ for all $A, B \subseteq X$. A subset $S \subseteq X$ is closed if $\text{cl}_X(S) = S$ and open if its complement is closed.

Finite posets are well-known to be equivalent to finite T_0 topological spaces, which amount to finite separable topological closures. For finite join-semilattices we instead use *finite strict closures* i.e. we do not require separability or preservation of unions.

Example 3.6. Each finite join-semilattice Q has an associated finite strict closure space $GQ = (J(Q), \text{cl}_{J(Q)})$ where $J(Q) \subseteq Q$ is the set of join-irreducibles and

$$\text{cl}_{J(Q)}(S) = \{j \in J(Q) : j \leq \sum_{s \in S} s\} \quad \text{for any } S \subseteq J(Q).$$

For example the closure space associated to the free join-semilattice $\mathcal{P}n$ is $(n, \text{id}_{\mathcal{P}n})$, identifying $J(\mathcal{P}n)$ with n .

Definition 3.7. *The category $\overline{\text{JSL}}$ has as objects all finite strict closure spaces as morphisms all continuous relations. Here a relation $R \subseteq X \times Y$ between two finite strict closure spaces X and Y is called continuous if, for all $x \in X$ and $S \subseteq X$,*

1. $R[x] \subseteq Y$ is closed, and
2. if $x \in \text{cl}_X(S)$ then $R[x] \subseteq \text{cl}_Y(R[S])$.

The composition of $R_1 \subseteq X \times Y$ and $R_2 \subseteq Y \times Z$ is defined by

$$R_2 \bullet R_1 := \{(x, z) \in X \times Z : z \in \mathbf{cl}_Z(R_2 \circ R_1[x])\},$$

and the identity morphism on X is $\text{id}_X = \{(x, x') \in X \times X : x' \in \mathbf{cl}_X(\{x\})\}$.

The following equivalence was derived from a similar one due to Moshier [14].

Lemma 3.8. *The functor $G : \text{JSL}_f \rightarrow \overline{\text{JSL}}$, defined on objects Q as in Example 3.6 and for morphisms $f : Q \rightarrow Q'$ by*

$$Gf = \{(j, j') \in J(Q) \times J(Q') : j' \leq_{Q'} f(j)\},$$

is an equivalence of categories.

Proof. (Sketch) We describe the opposite equivalence $H : \overline{\text{JSL}} \rightarrow \text{JSL}_f$ and also the unit and counit. Given $X = (X, \mathbf{cl}_X)$ then $HX = \{S \subseteq X : \mathbf{cl}_X(S) = S\} \subseteq \mathcal{P}X$ is the join-semilattice of closed subsets where $0_{HX} = \emptyset$ and $S +_{HX} S' = \mathbf{cl}_X(S \cup S')$. Given a continuous relation $R \subseteq X \times Y$ then $HR = \lambda S. \mathbf{cl}_Y(R[S]) : HX \rightarrow HY$ is the corresponding algebra morphism. The unit $\eta : \text{Id} \Rightarrow HG$ is defined

$$\eta_Q = \lambda q \in Q. \{j \in J(Q) : j \leq_Q q\},$$

and for $X = (X, \mathbf{cl}_X)$ the counit $\varepsilon : GH \Rightarrow \text{Id}$ is defined:

$$\varepsilon_X = \{(K, x) \in J(HX) \times X : K \in J(HX), x \in K\}.$$

It is well-typed because $J(HX) \subseteq HX \subseteq \mathcal{P}X$. □

Definition 3.9. $\overline{\text{DL}}$ has finite posets as objects and as morphisms those relations $R \subseteq P \times Q$ such that:

1. Each $R[p] \subseteq Q$ is downclosed,
2. If $p \leq_P p'$ then $R[p] \subseteq R[p']$,
3. R preserves all intersections of downclosed subsets.

id_P is the relation $\{(p, p') \in P \times P : p' \leq_P p\}$ and composition is relational composition.

Lemma 3.10. *The functor $G : \text{DL}_f \rightarrow \overline{\text{DL}}$ where $GQ = J(Q)$ (considered as a subposet of Q) and for morphisms $f : Q \rightarrow Q'$*

$$Gf = \{(z, z') \in J(Q) \times J(Q') : z' \leq_{Q'} f(z)\}$$

is an equivalence of categories.

Proof. G is restriction of the equivalence $\text{JSL}_f \cong \overline{\text{JSL}}$ described above. The closure spaces associated to distributive lattices are precisely the separable topological ones, so we can replace them by finite posets. This gives the first two conditions on morphisms, where closed means downwards closed. However semilattice morphisms between distributive lattices need not preserve meets. This is captured by the third condition. □

3.2 From Determinism to Nondeterminism

We first restrict the endofunctor T_Σ of Definition 2.6 to finite algebras:

$$T_\Sigma = \mathbb{2} \times \text{Id}^\Sigma : \mathcal{V}_f \rightarrow \mathcal{V}_f$$

Then for each of our five equivalences $G : \mathcal{V}_f \rightarrow \bar{\mathcal{V}}$ described in the previous section we have a corresponding functor

$$\bar{T}_\Sigma = \mathbb{1} \times \text{Id}^\Sigma : \bar{\mathcal{V}} \rightarrow \bar{\mathcal{V}}$$

where $\mathbb{1} = G\mathbb{2} \in \bar{\mathcal{V}}$. In each case $\mathbb{1}$ has carrier $\{1\}$.

\mathcal{V}	$\mathbb{1}$
Set_*	$\mathbb{2} \setminus \{0\} = \{1\}$
BA	$\text{At}(\mathbb{2}) = \{1\}$ the unique atom
$\text{Vect}(\mathbb{Z}_2)$	$\{1\}$ unique basis of $\mathbb{2} = \mathbb{Z}_2$
JSL	$(J(\mathbb{2}), \text{cl})$ where $J(\mathbb{2}) = \{1\}$, $\text{cl} = \text{id}_{\mathcal{P}\{1\}}$
DL	$J(\mathbb{2}) = \{1\}$ a discrete poset

Lemma 3.11. *There is an equivalence $\mathbb{G} : \text{Coalg}(T_\Sigma) \rightarrow \text{Coalg}(\bar{T}_\Sigma)$ defined by*

$$\mathbb{G}(Q, \gamma) = (GQ, \gamma') \text{ on objects} \quad \text{and} \quad \mathbb{G}f = Gf \text{ on morphisms,}$$

where $\gamma' : GQ \rightarrow \mathbb{1} \times (GQ)^\Sigma$ is the $\bar{\mathcal{V}}$ -morphism uniquely determined by the morphisms $G\gamma_\varepsilon : GQ \rightarrow \mathbb{1}$ and $G\gamma_a : GQ \rightarrow GQ$ for each $a \in \Sigma$.

Given a \bar{T}_Σ -coalgebra $\delta : Z \rightarrow \bar{T}_\Sigma Z = \mathbb{1} \times Z^\Sigma$ we write its component maps as $\delta_\varepsilon : Z \rightarrow \mathbb{1}$ and $\delta_a : Z \rightarrow Z$ for $a \in \Sigma$. Notice that these are *relations* rather than functions, so \bar{T}_Σ -coalgebras are *nondeterministic* automata.

Example 3.12. (a) When $\mathcal{V} = \text{Set}_*$ a \bar{T}_Σ -coalgebra $\delta : X \rightarrow \bar{T}_\Sigma X$ consists of:

1. A finite set X .
2. A partial function $\delta_\varepsilon : X \rightarrow \{1\}$ whose domain defines the final states.
3. A partial function $\delta_a : X \rightarrow X$ for each $a \in \Sigma$, defining the transitions.

Hence \bar{T}_Σ -coalgebras are *partial dfas*. The equivalence \mathbb{G} assigns to each deterministic Set_* -automaton (Q, γ) the partial dfa $(Q \setminus \{0_Q\}, \delta)$ whose final states are the given ones and $q \xrightarrow{a} q'$ iff $\gamma_a(q) = q' \neq 0_Q$.

(b) When $\mathcal{V} = \text{BA}$ a \bar{T}_Σ -coalgebra $\delta : X \rightarrow \bar{T}_\Sigma X$ consists of:

1. A finite set X .
2. A converse-functional relation $\delta_\varepsilon \subseteq X \times \{1\}$ whose domain defines a single final state.
3. Converse-functional relations $\delta_a \subseteq X \times X$ for $a \in \Sigma$.

Hence \bar{T}_Σ -coalgebras are *reverse-deterministic nfas*, i.e., reversing all transitions yields a dfa. The equivalence \mathbb{G} assigns to each deterministic BA-automaton (Q, γ) an nfa $(\text{At}(Q), \delta)$ whose states are Q 's atoms. Moreover, its single final state is the unique atom generating the ultrafilter $\gamma_\varepsilon^{-1}(\{1\})$ and $z \xrightarrow{a} z'$ iff $z' \leq_Q \gamma_a(z)$.

(c) If $\mathcal{V} = \text{Vect}(\mathbb{Z}_2)$ then a \bar{T}_Σ -coalgebra $\delta : X \rightarrow \bar{T}_\Sigma X$ consists of:

1. A finite set X .
 2. An arbitrary relation $\delta_\varepsilon \subseteq X \times \{1\}$, amounting to an arbitrary set of final states by taking the domain.
 3. Arbitrary relations $\delta_a \subseteq X \times X$ for each $a \in \Sigma$.
- Hence \overline{T}_Σ -coalgebras are classical nfes. The equivalence \mathbb{G} assigns to a deterministic $\text{Vect}(\mathbb{Z}_2)$ -automaton (Q, γ) the nfa (Z, δ) for some chosen basis $Z \subseteq Q$. The final states are $Z \cap \gamma_\varepsilon^{-1}(\{1\})$ and $z \xrightarrow{a} z'$ iff $\pi_{z'} \circ \gamma_a(z) = 1$, cf. Notation 3.3.
- (d) If $\mathcal{V} = \text{JSL}$ then a \overline{T}_Σ -coalgebra $\delta : Z \rightarrow \overline{T}_\Sigma Z$ consists of:
1. A finite strict closure space $Z = (Z, \text{cl}_Z)$.
 2. A continuous relation $\delta_\varepsilon \subseteq Z \times \{1\}$, equivalently δ_ε 's domain $F \subseteq Z$ is an open set of final states.
 3. Continuous relations $\delta_a \subseteq Z \times Z$.

We call \overline{T}_Σ -coalgebras *nondeterministic closure automata*. The equivalence \mathbb{G} assigns to each deterministic JSL-automaton (Q, γ) the nondeterministic closure automaton $((J(Q), \text{cl}_Q), \delta)$ whose states are Q 's join-irreducibles. The open set of final states is $J(Q) \cap \gamma_\varepsilon^{-1}(\{1\})$ and $z \xrightarrow{a} z'$ iff $z' \leq_Q \gamma_a(z)$.

Note that every nfa can be turned into a nondeterministic closure automaton by endowing the states with the identity closure, so classical nfes form a proper subclass.

- (e) If $\mathcal{V} = \text{DL}$ then a \overline{T}_Σ -coalgebra $\delta : P \rightarrow \overline{T}_\Sigma P$ consists of:
1. A finite poset P .
 2. A non-empty relation $\delta_\varepsilon \subseteq P \times \{1\}$ whose domain is a filter (i.e., a down-directed upset), these being the final states.
 3. Transition relations $\delta_a \subseteq P \times P$ such that:
 - (i) $\delta_a[p]$ is downclosed for each $p \in P$.
 - (ii) $p \leq_P q$ implies $\delta_a[p] \subseteq \delta_a[q]$.
 - (iii) $\delta_a[\bigcap_I A_i] = \bigcap_I \delta_a[A_i]$ for downclosed A_i .

Note that reverse-deterministic nfes are the special case where P is discrete. An important non-discrete example is the *universal automaton* [15], we recall it after Corollary 3.21.

The equivalence \mathbb{G} assigns to each deterministic DL-automaton (Q, γ) the \overline{T}_Σ -coalgebra $(J(Q), \delta)$ where $J(Q)$ is a subposet of Q . The final states form the upwards closed set $J(Q) \cap \gamma_\varepsilon^{-1}(1)$ and $z \xrightarrow{a} z'$ iff $z' \leq_Q \gamma_a(z)$.

Remark 3.13. A morphism $f : (Z, \delta) \rightarrow (Z', \delta')$ of \overline{T}_Σ -coalgebras is, by definition, a $\overline{\mathcal{V}}$ -morphism $f : Z \rightarrow Z'$ satisfying $\overline{T}_\Sigma f \circ \delta = \delta' \circ f$, or equivalently:

$$\delta_\varepsilon = \delta'_\varepsilon \circ f, \quad \delta'_a \circ f = f \circ \delta_a \quad (a \in \Sigma).$$

For $\mathcal{V} = \text{Set}_*$, BA and DL, these morphisms are those relations (from $\overline{\mathcal{V}}$) which (i) reflect and preserve transitions and (ii) have $z \in Z$ final iff some $z' \in f[z]$ is final. The cases $\mathcal{V} = \text{JSL}, \text{Vect}(\mathbb{Z}_2)$ are different because composition in $\overline{\mathcal{V}}$ is not relational.

3.3 Canonical Nondeterministic Automata

So far we have seen equivalences between deterministic and nondeterministic automata without initial states. Next, for each of our five running examples $\mathcal{V} = \text{Set}_*, \text{BA}, \text{Vect}(\mathbb{Z}_2), \text{JSL}, \text{DL}$ we will extend $\mathbb{G} : \text{Coalg}(T_\Sigma) \rightarrow \text{Coalg}(\overline{T}_\Sigma)$ to an equivalence of pointed coalgebras.

Definition 3.14. $\text{Coalg}_*(T_\Sigma)$ is the category whose objects are the pointed T_Σ -coalgebras and whose morphisms $f : (Q, \gamma, q_0) \rightarrow (Q', \gamma', q'_0)$ are those T_Σ -coalgebra homomorphisms $f : (Q, \gamma) \rightarrow (Q', \gamma')$ preserving initial states, i.e., $f \circ q_0 = q'_0$.

Using the equivalence $G : \mathcal{V}_f \rightarrow \bar{\mathcal{V}}$, a pointed \bar{T}_Σ -coalgebra is a \bar{T}_Σ -coalgebra (Z, δ) equipped with a $\bar{\mathcal{V}}$ -morphism $i : GV \rightarrow Z$. And pointed \bar{T}_Σ -coalgebra homomorphisms are those \bar{T}_Σ -coalgebra homomorphisms f from (Z, δ) to (Z', δ') such that $f \circ i = i'$. Just as a morphism $q_0 : V \rightarrow Q$ corresponds to an initial state $q_0(g)$, it turns out that a morphism $i : GV \rightarrow Z$ corresponds to a set of initial states $I = i[g] \subseteq Z$, as one would expect for nfas.

Example 3.15. For each \mathcal{V} we describe the possible sets of initial states $I \subseteq Z$ for a \bar{T}_Σ -coalgebra (Z, δ) .

- (a) If $\mathcal{V} = \text{Set}_*$ then $V = \{0, g\}$ and $GV = \{g\}$. Partial functions $i : \{g\} \rightarrow Z$ are determined by their codomain $I = i[g]$. Then I is either empty or any singleton subset.
- (b) If $\mathcal{V} = \text{BA}$ then $V = \{\perp, g, \neg g, \top\}$ and $GV = \{g, \neg g\}$. Given $i \subseteq \{g, \neg g\} \times Z$ then $i[g], i[\neg g]$ partition Z so i is determined by $I = i[g]$. Then I is any subset of Z .
- (c) If $\mathcal{V} = \text{Vect}(\mathbb{Z}_2)$ then $V = \{0, g\}$ and $GV = \{g\}$, so the arbitrary relation $i \subseteq \{g\} \times Z$ is determined by its codomain $I = i[g]$. Then I is any subset of Z .
- (d) If $\mathcal{V} = \text{JSL}$ then $V = \{0, g\}$ and $GV = \{g\}$ with closure $\text{id}_{\mathcal{P}\{g\}}$. The relation $i \subseteq \{g\} \times Z$ is determined by $I = i[g]$. By continuity $I \subseteq Z$ is any closed subset.
- (e) If $\mathcal{V} = \text{DL}$ then $V = \{\perp, g, \top\}$ is a 3-chain and $GV = \{g, \top\}$ a 2-chain. Given $i \subseteq \{g, \top\} \times Z$ then $i[g] \subseteq i[\top]$ and $i[\{g, \top\}] = Z$ implies $i[\top] = Z$, so i is determined by $I = i[g]$. Then I is any downclosed subset of Z .

By reinterpreting point preservation relative to I we can finally define the category of pointed \bar{T}_Σ -coalgebras.

Definition 3.16. For each of our five running examples, $\text{Coalg}_*(\bar{T}_\Sigma)$'s objects are triples (Z, δ, I) where (Z, δ) is a \bar{T}_Σ -coalgebra and $I \subseteq Z$ is restricted as in Example 3.15. The pointed \bar{T}_Σ -coalgebra homomorphisms $f : (Z, \delta, I) \rightarrow (Z', \delta', I')$ are \bar{T}_Σ -coalgebra homomorphisms $f : (Z, \delta) \rightarrow (Z', \delta')$ such that:

1. If $\mathcal{V} = \text{Set}_*, \text{BA}$ or DL then $I' = f[I]$.
2. If $\mathcal{V} = \text{JSL}$ then I' is the closure of $f[I]$.
3. If $\mathcal{V} = \text{Vect}(\mathbb{Z}_2)$ then $I' = \{z' \in Z' : |I \cap \check{f}[z']| \text{ is odd}\}$.

where $\check{f} \subseteq Z' \times Z$ is the converse relation.

Lemma 3.17. There is an equivalence of pointed coalgebras $\mathbb{G}_* : \text{Coalg}_*(T_\Sigma) \rightarrow \text{Coalg}_*(\bar{T}_\Sigma)$ defined by

$$\mathbb{G}_*(Q, \gamma, q_0) = (\mathbb{G}(Q, \gamma), I) \quad \mathbb{G}_*f = Gf$$

where $I = Gq_0[g] \subseteq GQ$.

Let us spell out the equivalence \mathbb{G}_* for each of our varieties \mathcal{V} . For the rest of this section fix a T_Σ -coalgebra $A = (Q, \gamma, q_0)$ and a regular language $L \subseteq \Sigma^*$. We give an explicit description of the nfa G_*A and, in particular, of the canonical nfa for L obtained by applying \mathbb{G}_* to $A_{\mathcal{V}}^L$ from Construction 2.15.

(a) **The Minimal Partial Dfa.** If $\mathcal{V} = \text{Set}_*$ then \mathbb{G}_*A is the partial dfa $(Q \setminus \{0_Q\}, \delta, I)$ that arises from A by deleting the state 0_Q along with all in- and outgoing transitions. Hence the initial states are $I = \{q_0\}$ if $q_0 \neq 0_Q$ and $I = \emptyset$ if $q_0 = 0_Q$. Clearly \mathbb{G}_*A (viewed as an nfa) accepts A 's language.

In particular, $\mathbb{G}_*(A_{\text{Set}_*}^L)$ is the *minimal partial dfa* of L . It has states

$$\mathcal{Q}_L = \{w^{-1}L : w \in \Sigma^*\} \setminus \{\emptyset\},$$

transitions $K \xrightarrow{a} a^{-1}K$ whenever $a^{-1}K \neq \emptyset$, and a state is final iff it contains ε . The initial states are $\{L\}$ if $L \neq \emptyset$ and \emptyset otherwise. Hence the minimal partial dfa is the trim part of L 's minimal dfa (obtained by deleting its sink state, if it exists).

(b) **The Átomaton.** If $\mathcal{V} = \text{BA}$ then \mathbb{G}_*A is the nfa $(\text{At}(Q), \delta, I)$ with initial states $I = \{q \in \text{At}(Q) : q \leq_Q q_0\}$. It accepts A 's language. In particular, $\mathbb{G}_*(A_{\text{BA}}^L)$ is called the *átomaton* of L , see [9]. Its states

$$\mathcal{Q}_L = \text{At}(\langle \{w^{-1}L : w \in \Sigma^*\} \rangle_{\nu T_\Sigma})$$

are the atoms of the finite boolean subalgebra of $\mathcal{P}\Sigma^*$ generated by L 's derivatives. An atom K is an initial state if $K \subseteq L$, the final states are the atoms containing ε , and one has transitions $K \xrightarrow{a} K'$ whenever $K' \subseteq a^{-1}K$. Explicitly constructing \mathcal{Q}_L can be difficult. Fortunately, a simpler method is known [9]:

1. Construct the minimal dfa for L 's reversed language.
2. Construct its reversed nfa i.e. flip initial/final states and reverse all transitions.

The átomaton is isomorphic to the resulting nfa as we now explain coalgebraically. Let $T'_\Sigma = 2 \times \text{Id}^\Sigma : \text{Set}_f \rightarrow \text{Set}_f$. Then the usual reversal of finite pointed deterministic automata defines a dual equivalence:

$$\begin{aligned} H &: (\text{Coalg}_*(T'_\Sigma))^{op} \rightarrow \text{Coalg}_*(\overline{T}_\Sigma) \\ Hf^{op} &= \{(z', z) : z \in f^{-1}(\{z'\})\} \subseteq Z' \times Z, \end{aligned}$$

Since reachability (no proper subobjects) and simplicity (no proper quotients) are *dual* concepts (see Definition 2.13), a T'_Σ -coalgebra is minimal iff its image under H is minimal, implying the above description.

- Example 3.18.*
1. The átomaton for $L = (a + b)^*b(a + b)^n$ in Example 1.1 arises by constructing the minimal dfa for the reversed language $\text{rev}(L)$ and taking the reverse nfa. Its atoms are $\{(a + b)^*a(a + b)^n, L\} \cup \{(a + b)^j : 0 \leq j \leq n\}$.
 2. The átomaton can have exponentially many more states than the minimal dfa, e.g. for $L = (a + b)^nb(a + b)^*$ it has $\geq 2^n$ states.

(c) **The Minimal Xor Automaton.** If $\mathcal{V} = \text{Vect}(\mathbb{Z}_2)$ then \mathbb{G}_*A is the nfa (Z, δ, I) where $Z \subseteq Q$ is a basis and $I = \{z \in Z : \pi_z(q_0) = 1\}$, see Notation 3.3. It accepts A 's language by \mathbb{Z}_2 -weighted nondeterministic acceptance: a word $w \in \Sigma^*$ is accepted iff its number of accepting paths is odd (this is different than the usual acceptance condition of standard nondeterministic automata).

The nfa $\mathbb{G}_*(A_{\text{Vect}(\mathbb{Z}_2)}^L)$ is called the *minimal xor automaton* of L , see [19]. Note that its construction depends on the choice of a basis, so the minimal xor automaton is only determined up to isomorphism in the category of pointed \overline{T}_Σ -coalgebras. We provide a new way to construct it:

1. Construct L 's átomaton (Z, R_a, F, I) and determine the collection $C \subseteq \mathcal{P}Z$ of all subsets of Z which are reachable from I .
2. Find any minimal $\mathcal{Q} \subseteq \mathcal{P}Z$ whose closure under set-theoretic symmetric difference equals C 's closure.
3. Build the nfa $(\mathcal{Q}, R'_a, \mathcal{Q} \cap F, I)$ where $R'_a(y, y')$ iff $\pi_{y'}(R_a[y]) = 1$ and $I = \{y \in \mathcal{Q} : \pi_y(I) = 1\}$.

Briefly, closure under boolean operations implies closure under symmetric difference. Then $A_{\text{Vect}(\mathbb{Z}_2)}^L \subseteq A_{\text{BA}}^L$ as da's, leading to the above algorithm. Since the basis \mathcal{Q} has $|\mathcal{Q}| \leq |C| = |\{w^{-1}L : w \in \Sigma^*\}|$ it follows that *the minimal xor automaton is never larger than the minimal dfa of L* , see [19].

Example 3.19. Take the átomaton of Example 1.1, with states $Z = \{x\} \cup \{z_i : 0 \leq i \leq n+1\}$ and reachable subsets $C = \{S \subseteq Z : x \notin S, z_0 \in S\}$. One can verify that (i) the closure of $\mathcal{Q} = \{\{z_i\} : 0 \leq i \leq n+1\}$ under symmetric difference is the closure of C and (ii) \mathcal{Q} is minimal. The induced nfa is the minimal xor automaton of Example 1.1. Alternatively $\mathcal{Q} = \{\{z_0, z_i\} : 0 \leq i \leq n+1\} \subseteq C$ yields a different nfa.

(d) **The Jiromaton.** If $\mathcal{V} = \text{JSL}$ then \mathbb{G}_*A is the nondeterministic closure automaton $(J(Q), \delta, I)$ with initial states $I = \{z \in J(Q) : z \leq_Q q_0\}$ where $J(Q)$ is the closure space of Example 3.6. The underlying nfa (forgetting the closure) accepts A 's language. In particular, $\mathbb{G}_*(A_{\text{JSL}}^L)$'s underlying nfa is called the *jiromaton* of L , see [11]. Its states

$$\mathcal{Q}_L = J(\langle \{w^{-1}L : w \in \Sigma^*\} \rangle_{\nu T_\Sigma})$$

are the join-irreducibles of the finite join-subsemilattice of $\mathcal{P}\Sigma^*$ generated by L 's derivatives. Since the latter form the minimal generating set, \mathcal{Q}_L consists of those L -derivatives not arising as unions of other derivatives – the *prime* derivatives. Therefore, the jiromaton has no more states than the minimal dfa. Its structure is analogous to the átomaton: $K \in \mathcal{Q}_L$ is initial iff $K \subseteq L$, final iff $\varepsilon \in K$ and $K \xrightarrow{a} K'$ iff $K' \subseteq a^{-1}K$.

An algorithm to construct the jiromaton from any nfa accepting L is given in [11].

Example 3.20. In the jiromaton of Example 1.1, the state z_0 accepts L and state z_i accepts $L + (a+b)^{i-1}$ for each $i > 0$. These are the prime derivatives of L . The closure is defined $\text{cl}_Z(\emptyset) = \emptyset$, $\text{cl}_Z(S) = \{z_0\} \cup \{S\}$ for $S \neq \emptyset$. It is topological: the closed sets are the downsets of the poset where $z_0 \leq_Z z_i$ for all $0 \leq i \leq n+1$.

(e) **The Distromaton.** If $\mathcal{V} = \text{DL}$ then $\mathbb{G}_* A = (J(Q), \delta, I)$ with initial states $I = \{z \in J(Q) : z \leq_Q q_0\}$. Forgetting $J(Q)$'s poset structure, the underlying nfa accepts A 's language. We call $\mathbb{G}_*(A_{\text{DL}}^L)$ the *distromaton* of L . Its states

$$\mathcal{Q}_L = J(\langle \{w^{-1}L : w \in \Sigma^*\} \rangle_{\nu T_\Sigma})$$

are the join-irreducibles of the sublattice of $\mathcal{P}\Sigma^*$ generated by L 's derivatives. One can close under intersections and then unions (which cannot add or remove join-irreducibles) so \mathcal{Q}_L consists of finite intersections $\bigcap_i w_i^{-1}L$ not arising as finite unions of other such intersections. The structure is again analogous to the átomaton and the jromaton: $K \in \mathcal{Q}_L$ is initial iff $K \subseteq L$, final iff $\varepsilon \in K$ and $K \xrightarrow{a} K'$ iff $K' \subseteq a^{-1}K$. There is another way to construct the distromaton, analogous to the construction of the átomaton:

1. Take the minimal pointed dfa $(Z, \xrightarrow{a}, z_0, F)$ for the reversed language $\text{rev}(L)$ where Z is ordered by language-inclusion.
2. Build the pointed \overline{T}_Σ -coalgebra (Z^{op}, δ, F) with final states $\downarrow_Z z_0$ and $z' \in \delta_a[z]$ iff $z' \xrightarrow{a} y \geq_Z z$.

The initial states F are downclosed in Z^{op} and the final states are upclosed in Z^{op} , as required. The proof that this is isomorphic to the distromaton is analogous to our earlier argument regarding the átomaton. Briefly, let $T'_\Sigma = 2 \times \text{Id}^\Sigma : \text{Poset}_f \rightarrow \text{Poset}_f$ where 2 is the two-chain. Then there is a dual equivalence

$$H : (\text{Coalg}_*(T'_\Sigma))^{op} \rightarrow \text{Coalg}_*(\overline{T}_\Sigma),$$

which ‘reverses’ finite pointed deterministic automata equipped with a compatible ordering. The minimal T'_Σ -coalgebra for L is the usual minimal dfa, now equipped with the language-inclusion ordering. Its image under H is again minimal, yielding the above description of the distromaton.

Corollary 3.21. *L 's átomaton and distromaton have the same number of states, namely, the number of states of the minimal dfa for the reversed language $\text{rev}(L)$.*

Example 3.22. The distromaton in Example 1.1 has order $z_0 \leq_Z z_i$ and $z_i \leq_Z \top$ for all $0 \leq i \leq n + 1$. We have the state \top because Σ^* is not the union of non-empty intersections of L 's derivatives, see Example 3.20. It arises from the jromaton by adding a final sink state, see Corollary 4.6.

We finally mention the well-studied *universal automaton* for L [15]. It is the nfa with states

$$\mathcal{Q} = \left\{ \bigcap_{w \in I} w^{-1}L : I \subseteq_\omega \Sigma^* \right\}$$

ordered by inclusion, where K is final iff $\varepsilon \in K$ and $K \xrightarrow{a} K'$ iff $K' \subseteq a^{-1}K$. The distromaton is never larger and often much smaller because one restricts to the join-irreducible intersections. However the universal automaton has its own advantages: in a sense every state-minimal nfa lies inside it.

4 State Minimality and Universal Properties

This final section is split into three parts.

1. We prove L 's jiomaton is minimal amongst all nondeterministic acceptors of L relative to a suitable measure (Section 4.1).
2. We give a sufficient condition on L such that the jiomaton is state-minimal and the distromaton and átomaton have at most one more state (Section 4.2).
3. We characterize each of our canonical nfas amongst subclasses of nondeterministic acceptors (Section 4.3).

4.1 The Jiomaton is Minimal

There is a measure on finite nondeterministic automata such that L 's jiomaton is smaller than any other nfa accepting L . For any nfa $N = (Q, R_a, F)$ and $I \subseteq Q$ let $\mathcal{L}_N(I) \subseteq \Sigma^*$ be the accepted language. Define the following measures:

$$|N| = |Q|, \quad \text{acc}(N) = |\{\mathcal{L}_N(I) : I \subseteq Q\}|, \quad \text{tr}(N) = \sum_{a \in \Sigma} |R_a|.$$

These are the number of states, the number of distinct languages accepted and the number of transitions. Let J_L be L 's jiomaton without initial states. Recall that isomorphisms of nfas are bijective bisimulations (see Definition 2.1).

Theorem 4.1. *The jiomaton J_L is (up to isomorphism) the unique nfa accepting L such that for every nfa N accepting L :*

- (1) $\text{acc}(J_L) \leq \text{acc}(N)$,
- (2) *If additionally $\text{acc}(J_L) = \text{acc}(N)$ then either:*
 - (a) $|J_L| < |N|$ or
 - (b) $|J_L| = |N|$ and $\text{tr}(N) \leq \text{tr}(J_L)$.

Proof. Since J_L 's individual states accept derivatives of L , it follows that J_L accepts precisely the unions of derivatives of L . Any nfa N accepting L accepts these languages, so $\text{acc}(J_L) \leq \text{acc}(N)$. Suppose $\text{acc}(J_L) = \text{acc}(N)$, so N accepts precisely the unions of L 's derivatives. Then each prime derivative has a distinct state in N accepting it, as it cannot arise as the union of other derivatives, so $|J_L| \leq |N|$. Lastly if $\text{acc}(J_L) = \text{acc}(N)$ and $|J_L| = |N|$ then there is language preserving bijection between N 's states and the set of prime derivatives P_L , so assume N 's carrier is P_L . Given $K \xrightarrow{a} K'$ in N we must have $K' \subseteq a^{-1}K$, so there is a corresponding transition in J_L . Hence $\text{tr}(N) \leq \text{tr}(J_L)$ and (2) holds. Moreover, in case $\text{tr}(N) = \text{tr}(J_L)$ the previous argument shows that N and J_L are isomorphic. Thus the conditions (1) and (2) determine J_L up to isomorphism. \square

4.2 Conditions for Canonical State-minimality

In the following let d_L and n_L be the minimal number of states of a dfa (respectively nfa) accepting the regular language L . For any state-minimal nfa $N = (n_L, R_a, F)$ accepting L via $I \subseteq n_L$, one can construct a simple pointed T_Σ -coalgebra (Q, γ', L) whose equivalent nondeterministic closure automaton is another state-minimal acceptor of L . First view N as the T_Σ -coalgebra $(\mathcal{P}n_L, \gamma)$ via the subset construction. Factorizing the unique homomorphism \mathcal{L}_γ we obtain (Q, γ') where Q is the semilattice of languages accepted by N . Then (Q, γ') is equivalent to a nondeterministic closure automaton accepting L . Since $\mathcal{P}n_L \twoheadrightarrow Q$ implies $n_L = |J(\mathcal{P}n_L)| \geq |J(Q)|$, by forgetting the closure we obtain a state-minimal nfa accepting L .

Hence instead of working with state-minimal nfes we may work with simple T_Σ -coalgebras which are *supercoalgebras* of A_{JSL}^L . This follows because A_{JSL}^L 's carrier is the semilattice S_L of unions of L 's derivatives, which Q necessarily contains. We now provide a condition ensuring that $|J(S_L)|$ is the minimal size of an nfa accepting L and hence L 's jiromaton is *state-minimal*.

Definition 4.2. *A regular language L is intersection-closed if every binary intersection of L 's derivatives is a union of L 's derivatives.*

- Example 4.3.*
1. $L = (a + b)^*b(a + b)^n$ where $n \in \omega$ is intersection-closed.
 2. \emptyset, Σ^* and $\{w\}$ for $w \in \Sigma^*$ are intersection-closed.
 3. Fix $n \in \omega, t \in \mathbb{R}$ and $k_i \in \mathbb{R}$ ($1 \leq i \leq n$). Then the language $L = \{w \in 2^n : \sum_i k_i w_i \geq t\}$ (modeling the behaviour of an artificial neuron) is intersection-closed.
 4. Every linear subspace $L \subseteq \mathbb{Z}_2^n$ (viewed as a language over the alphabet $\{0, 1\}$) is intersection-closed.

Theorem 4.4. *If L is intersection-closed then its jiromaton is state-minimal.*

Proof. By assumption the carrier S_L of A_{JSL}^L is closed under both unions and non-empty intersections, so $D = S_L \cup \{\Sigma^*\}$ is a distributive lattice of languages. Let N be any state-minimal nfa accepting L via initial states I , and $S \subseteq \mathcal{P}\Sigma^*$ be the semilattice of languages accepted by N (by varying I). The nfa N must at least accept L 's derivatives. Since S is closed under unions we have $S_L \subseteq S$. By the surjective morphism $\mathcal{P}n_L \twoheadrightarrow S$ it follows that $|N| \geq |J(S)|$, so it suffices to prove that $|J(S)| \geq |J(S_L)|$. Let $S_* = S \cup \{\Sigma^*\}$ be the semilattice obtained by adding a top element if necessary. We have a JSL_f -morphism $\iota : D \hookrightarrow S_*$. The meets in D are also meets in S_* so the same function defines a JSL_f -morphism $\iota : D^{op} \hookrightarrow S_*^{op}$. By the self-duality of JSL_f we obtain a surjective morphism $\iota' : S_* \twoheadrightarrow D$, hence $|J(S_*)| \geq |J(D)|$. If $D = S_L$ then $S_* = S$, so $|J(S)| \geq |J(S_L)|$ and we are done. Otherwise $\Sigma^* \notin S_L$ and we now prove $\Sigma^* \notin S$. By state minimality N is reachable, so each state q accepts a subset of some L -derivative. Then if $\Sigma^* \in S$ we deduce Σ^* is the union of L 's derivatives, so $\Sigma^* \in S_L$ – a contradiction. Consequently $|J(D)| = 1 + |J(S_L)|$ and $|J(S_*)| = 1 + |J(S)|$ hence $|J(S)| \geq |J(S_L)|$ again. \square

Remark 4.5. The converse of this theorem is generally false: the language $L = \overline{\{aa\}}$ is not intersection-closed, but its jiromaton is state-minimal.

Corollary 4.6. *If L is intersection-closed then its átomaton and distromaton have at most one more state than the jiromaton.*

Proof. By the above proof the distromaton may only have an additional final sink state – otherwise it has the same transition structure. By Corollary 3.21 the átomaton has the same number of states. \square

By Corollary 3.21 we further deduce:

Corollary 4.7. *If $L \subseteq \Sigma^*$ is intersection-closed then any state-minimal nfa accepting L has (i) $d_{\text{rev}(L)}$ states if Σ^* is a union of L 's derivatives and (ii) $d_{\text{rev}(L)} - 1$ otherwise.*

Theorem 4.8. *If $d_L = 2^{n_L}$ then the jiromaton of L is state-minimal.*

Proof. Let $N = (n_L, R_a, F)$ be a state-minimal nfa accepting L via $I \subseteq n_L$. View it as a pointed T_Σ -coalgebra $A = (\mathcal{P}n_L, \gamma, I)$ via the subset-construction. By assumption $d_L = |\mathcal{P}n_L|$, so this is a state-minimal dfa accepting L ; in particular, it is a reachable pointed T_Σ -coalgebra. Then the surjective morphism $A \twoheadrightarrow A_{\text{JSL}}^L$ implies that A_{JSL}^L has no more than n_L join-irreducibles, so the jiromaton is state-minimal. \square

4.3 Characterizing the Canonical Nfas

Although the canonical nfas are generally not state-minimal, they are state-minimal amongst certain subclasses of nfas.

Theorem 4.9. *The átomaton of a regular language L is state-minimal amongst all nfas accepting L whose accepted languages are closed under complement.*

Proof. Assume the weaker condition that an nfa N accepts every language in the boolean algebra $\mathcal{B} \subseteq_\omega \mathcal{P}\Sigma^*$ generated by L 's derivatives. By an earlier argument, N induces a simple T_Σ -coalgebra (\mathcal{Q}, γ) whose states are the languages N accepts and $|\mathcal{N}| \geq |\mathcal{J}(\mathcal{Q})|$. By assumption $\mathcal{Q} \supseteq \mathcal{B}$ (a distributive lattice), so $|\mathcal{J}(\mathcal{Q})| \geq |\mathcal{J}(\mathcal{B})|$ by the proof of Theorem 4.4. The join-irreducibles of a finite boolean algebra are its atoms, so N has no less states than the átomaton. \square

The next result is from [19]. It follows because quotients and subspaces of finite-dimensional vector spaces cannot have larger dimension.

Theorem 4.10 ([19]). *Any canonical xor nfa for L is state-minimal amongst nfas accepting L via \mathbb{Z}_2 -weighted acceptance.*

We give a mild generalization of a result in [11]. Recall that nfas accepting L also accept all unions of its derivatives. Then we can conclude from Theorem 4.1:

Corollary 4.11. *The jiromaton of a regular language L is state-minimal amongst nfas accepting precisely the unions of L 's derivatives.*

Example 4.12. Let N be an nfa accepting L via initial states I . If every singleton set of states is reachable from I then N accepts precisely the unions of L 's derivatives. Thus, it is no smaller than L 's jiromaton.

Theorem 4.13. *The distromaton of a regular language L is state-minimal amongst all nfas accepting L whose accepted languages are closed under intersection.*

Proof. Reuse the proof of Theorem 4.9. Again we actually have a stronger result: the distromaton is state-minimal amongst all nfas which can accept every intersection of L 's derivatives. \square

5 Conclusions and Future Work

It is often claimed in the literature that canonical nondeterministic automata do not exist, usually as a counterpoint to the minimal dfa. On the contrary we have shown that they *do* exist and moreover arise from the minimal dfa interpreted in a locally finite variety. In so doing we have unified previous work from three sources [9, 11, 19] and introduced a new canonical nondeterministic acceptor, the distromaton. We also identified a class of languages where canonical *state-minimal* nfas exist. These results depend heavily on a coalgebraic approach to automata theory, providing not only new structural insights and construction methods but also a new perspective on what a state-minimal acceptor actually is.

In this paper we introduced nondeterministic closure automata, viz. \overline{T}_Σ -coalgebras in the category of closure spaces, mainly as a tool for constructing the jromaton. However, nondeterministic closure automata bear interesting structural properties themselves, which we did not discuss here in depth. We expect that a proper investigation of these machines will lead to further insights about nondeterminism, in particular additional and more general criteria for the (state-)minimality of nfas.

Another point we aim to investigate in more detail are the algorithmic aspects of the state-minimization problem for nfas. Although this problem is known to be PSPACE-complete in general, the canonicity of our nfas suggests that – at least for certain natural subclasses of nfas – efficient state-minimization procedures may be in reach. We leave the study of such complexity-related issues for future work.

References

- [1] Adámek, J., Bonchi, F., Hülsbusch, M., König, B., Milius, S., Silva, A.: A coalgebraic perspective on minimization and determinization. In: FoSSaCS. pp. 58–73 (2012)
- [2] Adámek, J., Milius, S., Moss, L.S., Sousa, L.: Well-pointed coalgebras (extended abstract). In: FoSSaCS. pp. 89–103 (2012)
- [3] Adámek, J., Milius, S., Velebil, J.: Iterative algebras at work. *Math. Structures Comput. Sci.* 16(6), 1085–1131 (2006)
- [4] Barr, M.: Terminal coalgebras in well-founded set theory. *Theor. Comput. Sci.* 114(2), 299–315 (1993)
- [5] Bezhanishvili, N., Kupke, C., Panangaden, P.: Minimization via duality. In: WoLLIC. pp. 191–205 (2012)
- [6] Bonchi, F., Bonsangue, M.M., Boreale, M., Rutten, J.J.M.M., Silva, A.: A coalgebraic perspective on linear weighted automata. *Inform. and Comput.* 211, 77–105 (2012)
- [7] Bonchi, F., Bonsangue, M.M., Rutten, J.J.M.M., Silva, A.: Brzozowski’s algorithm (co)algebraically. In: *Logic and Program Semantics*. pp. 12–23 (2012)
- [8] Bonsangue, M.M., Milius, S., Silva, A.: Sound and complete axiomatizations of coalgebraic language equivalence. *ACM Trans. Comput. Log.* 14(1:7) (2013)
- [9] Brzozowski, J., Tamm, H.: Theory of Átomata. In: *Proc. 15th International Conference on Developments in Language Theory (DLT’11)*. Lecture Notes Comput. Sci., vol. 6795, pp. 105–116. Springer (2011)
- [10] Brzozowski, J.A.: Canonical regular expressions and minimal state graphs for definite events. In: *Mathematical theory of Automata*, pp. 529–561. Volume 12 of MRI Symposia Series, Polytechnic Press, Polytechnic Institute of Brooklyn, N.Y. (1962)

- [11] Denis, F., Lemay, A., Terlutte, A.: Residual finite state automata. *Fund. Inform.* XX, 1–30 (2002)
- [12] Hasuo, I., Jacobs, B., Sokolova, A.: Generic trace semantics via coinduction. *Log. Methods Comput. Sci.* 3(4:11) (2007)
- [13] Jacobs, B., Silva, A., Sokolova, A.: Trace semantics via determinization. In: Pattinson, D., Schröder, L. (eds.) *Proc. CMCS 2012. Lecture Notes Comput. Sci.*, vol. 7399, pp. 109–129. Springer (2012)
- [14] Jipsen, P.: Categories of algebraic contexts equivalent to idempotent semirings and domain semirings. In: Kahl, W., Griffin, T.G. (eds.) *RAMICS. Lecture Notes in Computer Science*, vol. 7560, pp. 195–206. Springer (2012)
- [15] Lombardy, S., Sakarovitch, J.: The universal automaton. In: Flum, J., Grädel, E., Wilke, T. (eds.) *Logic and Automata. Texts in Logic and Games*, vol. 2, pp. 457–504. Amsterdam University Press (2008)
- [16] Milius, S.: A sound and complete calculus for finite stream circuits. In: *Proc. 25th Annual Symposium on Logic in Computer Science (LICS'10)*, pp. 449–458. IEEE Computer Society (2010)
- [17] Myers, R.: *Rational Coalgebraic Machines in Varieties: Languages, Completeness and Automatic Proofs*. Ph.D. thesis, Imperial College, London (2011)
- [18] Silva, A., Bonchi, F., Bonsangue, M.M., Rutten, J.J.M.M.: Generalizing determinization from automata to coalgebras. *Log. Methods Comput. Sci.* 9(1:9) (2013)
- [19] Vuillemin, J., Gama, N.: *Efficient Equivalence and Minimization for Non Deterministic Xor Automata*. Research report, LIENS (May 2010), <http://hal.inria.fr/inria-00487031>

A Appendix

In this appendix, which we include for the convenience of the referees, we provide technical details and proofs we omitted due to space constraints.

A.1 Deterministic \mathcal{V} -Automata

Let \mathcal{V} be a locally finite variety with a specified two-element algebra $\mathbb{2}$.

Proof (Details of Remark 2.9). A T -coalgebra is called *locally finitely presentable* (lfp) if it is a filtered colimit of finitely presentable T -coalgebras i.e. ones whose carrier is a finitely presentable object [16]. In locally finite varieties, T_Σ -coalgebras are lfp iff they are directed unions of finite T_Σ -coalgebras, so we call them *locally finite* T_Σ -coalgebras.

Colimits of categories of coalgebras are always constructed in their underlying category and directed colimits in \mathcal{V} are constructed in Set i.e. we simply take the union of these finite coalgebras. Consequently each state reaches only finitely many other states. Conversely any such T_Σ -coalgebra can be viewed as a directed union of its finite subcoalgebras – a particular filtered colimit. \square

Proof (Details of Remark 2.10.2). For any locally finitely presentable category \mathcal{V} and finitary functor $T : \mathcal{V} \rightarrow \mathcal{V}$ one can construct the *rational fixpoint* ρT i.e. the filtered colimit of all finitely presentable T -coalgebras [3]. At this level of generality it is known that the rational fixpoint is the final object in the category of lfp coalgebras.

The functor $T_\Sigma : \mathcal{V} \rightarrow \mathcal{V}$ lifts $T'_\Sigma = 2 \times \text{Id}^\Sigma : \text{Set} \rightarrow \text{Set}$. Since (i) colimits of coalgebras are constructed in \mathcal{V} , (ii) filtered colimits in varieties are constructed in Set and (iii) an algebra in \mathcal{V} is finitely presentable iff it is finite, it follows that ρT_Σ lifts $\rho T'_\Sigma$. The latter T'_Σ -coalgebra is known to be the da of regular languages i.e. they are those languages arising as the filtered colimit of all finite deterministic automata without initial states.

Furthermore the unique final map from any finite T'_Σ -coalgebra (Q, γ) to $\rho T'_\Sigma$ assigns to each state its language. Again this generalizes to T_Σ -coalgebras by the lifting and we have denoted it by \mathcal{L}_γ , see Notation 2.12. \square

Proof (Lemma 2.14). Let $A = (Q, \gamma, q_0)$ be a finite T_Σ -coalgebra. Since T_Σ preserves intersections they also exist in $\text{Coalg}(T_\Sigma)$ and are constructed on the level of \mathcal{V} . Consequently we can construct generated subcoalgebras. In particular the subcoalgebra $Q_{q_0} \subseteq Q$ generated by $q_0 \in Q$ is precisely the subalgebra of Q generated by those states $q' \in Q$ reachable from q by transitions, see [17]. This clearly implies the first statement of the lemma.

Now suppose A is simple: whenever $f : (Q, \gamma, q_0) \rightarrow (Q', \gamma', q')$ is surjective then it is bijective. Take the unique final map $\mathcal{L}_\gamma : (Q, \gamma) \rightarrow \rho T_\Sigma$ and factorize it into a surjective morphism f followed by an injective one, see Remark 2.18. Then f is also a pointed coalgebra morphism (mapping q_0 to $f(q_0)$) and hence bijective by simplicity. It follows that \mathcal{L}_γ is injective. Conversely suppose \mathcal{L}_γ is injective and $f : (Q, \gamma, q_0) \rightarrow (Q', \gamma', q')$ is surjective. We have $\mathcal{L}_\gamma = \mathcal{L}_{\gamma'} \circ f$ by finality, so f is injective and hence bijective. \square

Proof (Details of Remark 2.18). Let \mathcal{V} be a category with an (E, M) -factorization system and $T : \mathcal{V} \rightarrow \mathcal{V}$ be a functor such that $f \in M$ implies $Tf \in M$. Then it is well known that this factorization system lifts to $\text{Coalg}(T)$. In particular any variety \mathcal{V} has the factorization system (surjective hom, injective hom) and $T_\Sigma = \mathbb{2} \times \text{Id}^\Sigma$ preserves injections. \square

A.2 From Determinism to Nondeterminism

Proof (Lemma 3.4). The equivalence between pointed sets and partial functions is well known. For the second equivalence observe (i) BA_f is dually equivalent to Set_f , (ii) Set_f is a non-full subcategory of Rel_f , (iii) the latter category is self-dual by taking converse relations. Then Gf arises by following these three steps. The third statement follows by the well known equivalence of the category of finite dimensional vector spaces over some fixed field \mathbb{F} , and the category whose objects are the natural numbers and whose morphisms are matrices over \mathbb{F} . That is, we use relations rather than matrices and an altered form of relational composition rather than matrix multiplication over \mathbb{Z}_2 . \square

Proof (Lemma 3.8). See Section A.3 below. \square

Proof (Lemma 3.11). By universality of the product in \mathcal{V}_f , a finite T_Σ -coalgebra (Q, γ) amounts to \mathcal{V}_f -morphisms $\gamma_\varepsilon : Q \rightarrow \mathbb{2}$ and $\gamma_a : Q \rightarrow Q$ for each $a \in \Sigma$. Likewise a T_Σ -coalgebra homomorphism $f : (Q, \gamma) \rightarrow (Q', \gamma')$ amounts to a \mathcal{V}_f -morphism $f : Q \rightarrow Q'$ such that $\gamma_\varepsilon = \gamma'_\varepsilon \circ f$ and $f \circ \gamma_a = \gamma'_a \circ f$ for each $a \in \Sigma$. Then everything is defined in terms of the category \mathcal{V}_f , so the equivalence $G : \mathcal{V}_f \rightarrow \overline{\mathcal{V}}$ allows us to define an equivalent category by simply applying G to every morphism $\gamma_\varepsilon, \gamma_a$ and f . By universality of the product in $\overline{\mathcal{V}}$ it follows that $\mathbb{G} : \text{Coalg}(T_\Sigma) \rightarrow \text{Coalg}(\overline{T}_\Sigma)$ defines an equivalence. \square

Proof (Lemma 3.17). By Lemma 3.11 we know \mathbb{G} is an equivalence of coalgebras. The pointed T_Σ -coalgebras are defined by restricting T_Σ -coalgebras to have structure purely contained in \mathcal{V}_f i.e. \mathcal{V}_f -morphisms $q_0 : V \rightarrow Q$ which must be preserved. By the equivalence $G : \mathcal{V}_f \rightarrow \overline{\mathcal{V}}$ we can require the same conditions but now for morphisms $i : GV \rightarrow Z$. Reinterpreting i as a subset $I \subseteq Z$ (see Example 3.15) we again obtain an equivalent category i.e. $\text{Coalg}_*(\overline{T}_\Sigma)$. \square

A.3 Finite Join-Semilattices and Closure Spaces

We prove the equivalence of JSL_f and $\overline{\text{JSL}}$, the category of finite strict closure spaces. Let us first recall a few basic properties of finite join-semilattices.

Lemma A.1. *For every $Q \in \text{JSL}_f$, the join-irreducibles $J(Q) \subseteq Q$ form the unique minimal generating set.*

Proof. Suppose Q is generated by $G \subseteq Q$ i.e. every $q \in Q$ arises as a possibly empty join of elements in G . Then $J(Q) \subseteq G$ because join-irreducibles are not the join of distinct elements. Finally $J(Q)$ generates Q : we can continually split apart any $q \in Q \setminus \{0_Q\}$ to obtain a set of join-irreducibles whose join is q . This process terminates because Q is finite.

Lemma A.2. *The category of finite join-semilattices is self-dual.*

Proof. This duality restricts the self-duality of complete join-semilattices with arbitrary-join preserving morphisms. Specifically, we have the dual equivalence $H : (\text{JSL}_f)^{op} \rightarrow \text{JSL}_f$ defined by

$$HS = S^{op} \quad \frac{f : A \rightarrow B}{Hf = \lambda b. \sum_A f^{-1}(\downarrow_B b) : B^{op} \rightarrow A^{op}}$$

Here $x +_{A^{op}} y := x \wedge_A y$ and $0_{A^{op}} := \top_A$, recalling that every finite join-semilattice is a complete lattice and hence all meets exist. Equivalently A^{op} takes the meet-semilattice structure on A as its join-semilattice structure.

One can prove this by using the adjoint functor theorem for posets. That is, finite join-semilattices are finite posets with all joins (colimits) and join-semilattice morphisms are monotone maps that preserve all joins. Consequently each $f : A \rightarrow B$ has a right adjoint $f' : B^{op} \rightarrow A^{op}$ where the order is reversed because right adjoints preserve all meets. The uniqueness of adjoints implies that this is an equivalence. Its explication yields the above action on the morphisms.

Recall that a finite strict closure space is a pair (X, cl_X) where X is a finite set and $\text{cl}_X : \mathcal{P}X \rightarrow \mathcal{P}X$ is a strict closure operator, which means that

$$\frac{A \subseteq B}{\text{cl}_X(A) \subseteq \text{cl}_X(B)} \quad A \subseteq \text{cl}_X(A) \quad \text{cl}_X \circ \text{cl}_X = \text{cl}_X \quad \text{cl}_X(\emptyset) = \emptyset$$

We will use notation associated to the power set monad $\mathcal{P} : \text{Set} \rightarrow \text{Set}$, namely its unit $\eta_X : X \rightarrow \mathcal{P}X$ defined $\eta_X(x) = \{x\}$ and the multiplication $\mu_X : \mathcal{P}\mathcal{P}X \rightarrow \mathcal{P}X$ defined $\mu_X(S) = \bigcup S$, for any set X . Any function $f : X \rightarrow \mathcal{P}Y$ induces a relation \mathbf{f} and a union preserving function f^\uparrow as follows:

$$\mathbf{f} \subseteq X \times Y, \quad \mathbf{f}(x, y) \text{ iff } y \in f[x]$$

$$f^\uparrow : \mathcal{P}X \rightarrow \mathcal{P}Y, \quad f^\uparrow = \mu_Y \circ \mathcal{P}f = \lambda S. f[S]$$

where each f , \mathbf{f} and f^\uparrow determines the other two. In fact for any function $\theta : \mathcal{P}X \rightarrow \mathcal{P}Y$ define $\theta^\downarrow =_{def} \theta \circ \eta_X : X \rightarrow \mathcal{P}Y$. Then

$$(f^\uparrow)^\downarrow = \mu_Y \circ \mathcal{P}f \circ \eta_X = f.$$

Note that cl_X need not preserve unions, so in general $(\text{cl}_X^\downarrow)^\uparrow \neq \text{cl}_X$.

Recall from Definition 3.7 that the morphisms of closure spaces are continuous relations. For the upcoming proofs it is handier to replace these relations $\mathcal{R} \subseteq X \times Y$ by their corresponding functions $f : X \rightarrow \mathcal{P}Y$.

Definition A.3. *A morphism $f : (X, \text{cl}_X) \rightarrow (Y, \text{cl}_Y)$ of closure spaces is a function $f : X \rightarrow \mathcal{P}Y$ such that:*

1. $\text{cl}_Y \circ f = f$
[$f(x)$ is closed in Y .]

$$2. \text{cl}_Y \circ f^\uparrow \circ \text{cl}_X = \text{cl}_Y \circ f^\uparrow$$

[If x lies in the closure of $A \subseteq X$ then $f(x)$ lies inside the closure of $f[A] \subseteq Y$.]

where \circ denotes functional composition. The identity morphism $\text{id}_X : (X, \text{cl}_X) \rightarrow (X, \text{cl}_X)$ is defined $\text{id}_X =_{\text{def}} \text{cl}_X^\downarrow = \text{cl}_X \circ \eta_X$. Given any two morphisms $f : (X, \text{cl}_X) \rightarrow (Y, \text{cl}_Y)$ and $g : (Y, \text{cl}_Y) \rightarrow (Z, \text{cl}_Z)$, we define their composite as $g \bullet f =_{\text{def}} \text{cl}_Z \circ g^\uparrow \circ f$.

The category of finite strict closure spaces and their morphisms is denoted Cls . It is clearly isomorphic to $\overline{\text{JSL}}$. In the following we will work with Cls rather than $\overline{\text{JSL}}$.

- Example A.4.* (a) For any finite set X we have the discrete closure space $(X, \text{id}_{\mathcal{P}X})$. The Cls -morphisms $f : (X, \text{id}_{\mathcal{P}X}) \rightarrow (Y, \text{cl}_Y)$ are precisely the functions $f : X \rightarrow \mathcal{P}Y$ where $f(x)$ is closed in Y for all $x \in X$. This follows because the second condition is superfluous i.e. every subset of X is closed.
- (b) The Cls -morphisms between two discrete spaces $(X, \text{id}_{\mathcal{P}X})$ and $(Y, \text{id}_{\mathcal{P}Y})$ are precisely the functions $f : X \rightarrow \mathcal{P}Y$ because every subset of Y is closed, so both conditions are superfluous.
- (c) The Cls -morphisms $f : (X, \text{cl}_X) \rightarrow (Y, \text{id}_{\mathcal{P}Y})$ are precisely the functions $f : X \rightarrow \mathcal{P}Y$ such that $f^\uparrow \circ \text{cl}_X = f^\uparrow$. Here the first condition is superfluous.
- (d) Every closure space has a unique morphism *to* and *from* the empty space $(\emptyset, \text{cl}_\emptyset)$, so the latter is the initial and final object of Cls .

Lemma A.5. 1. $(g \circ f)^\uparrow = g \circ f^\uparrow$ for any function $f : X \rightarrow \mathcal{P}Y$ and union preserving function $g : \mathcal{P}Y \rightarrow \mathcal{P}Z$.

2. $f^\uparrow \circ \text{cl}_X^\downarrow = f$ for any Cls -morphism $f : X \rightarrow \mathcal{P}Y$.
3. $f^\uparrow \circ (\text{cl}_X^\downarrow)^\uparrow = f^\uparrow$ for any Cls -morphism $f : X \rightarrow \mathcal{P}Y$.
4. $\text{cl}_X \circ (\text{cl}_X^\downarrow)^\uparrow = \text{cl}_X$ for any closure cl_X .
5. $\text{cl}_Y \circ (\text{cl}_Y \circ f)^\uparrow = \text{cl}_Y \circ f^\uparrow$ for any function $f : X \rightarrow \mathcal{P}Y$ and closure cl_Y .
6. $\text{cl}_X(\bigcup_i A_i) = \text{cl}_X(\bigcup_i \text{cl}_X(A_i))$ for any closure cl_X and index set I .

Proof. 1. For any subset $A \subseteq X$ we compute

$$(g \circ f)^\uparrow(A) = \bigcup_{a \in A} g \circ f(a) = g\left(\bigcup_{a \in A} f(a)\right) = g \circ f^\uparrow(A)$$

2. The inclusion $f(x) \subseteq f^\uparrow \circ \text{cl}_X^\downarrow(x)$ is immediate. The converse direction is an instance of the second condition i.e. if x' lies in the closure of $\{x\}$ then $f(x')$ lies inside the closure of $f(x)$, in fact $f(x') \subseteq f(x)$ because $f(x)$ is closed in Y by the first condition.
3. follows by applying the first statement to the second statement.
4. First observe that $(\text{cl}_X^\downarrow)^\uparrow(A) = \bigcup_{x \in A} \text{cl}_X(\{x\}) \subseteq \text{cl}_X(A)$ by monotonicity, hence $\text{cl}_X \circ (\text{cl}_X^\downarrow)^\uparrow(A) \subseteq \text{cl}_X \circ \text{cl}_X(A) = \text{cl}_X(A)$ by monotonicity and also idempotency. This covers one inclusion. Conversely $A \subseteq (\text{cl}_X^\downarrow)^\uparrow(A)$ as every $x \in A$ lies in $\text{cl}_X(\{x\}) \subseteq (\text{cl}_X^\downarrow)^\uparrow(A)$, so the reverse inclusion follows by monotonicity.

5. We compute:

$$\begin{aligned}\mathbf{cl}_Y \circ (\mathbf{cl}_Y \circ f)^\uparrow(A) &= \mathbf{cl}_Y(\bigcup\{\mathbf{cl}_Y(f(x)) : x \in A\}) \\ &= \mathbf{cl}_Y(\bigcup\{f(x) : x \in A\}) \\ &= \mathbf{cl}_Y \circ f^\uparrow(A)\end{aligned}$$

using 6. in the second step.

6. We have the inclusion \subseteq by monotonicity and $A_i \subseteq \mathbf{cl}_X(A_i)$. The reverse inclusion follows because each $\mathbf{cl}_X(A_i) \subseteq \mathbf{cl}_X(\bigcup_i A_i)$.

Lemma A.6. *Cls is a well-defined category.*

Proof. Firstly $f \bullet \text{id}_X = f \bullet \mathbf{cl}_X^\downarrow = \mathbf{cl}_Y \circ f^\uparrow \circ \mathbf{cl}_X^\downarrow = \mathbf{cl}_Y \circ f = f$ using Lemma A.5.2 and the first condition. Next:

$$\text{id}_Y \bullet f = \mathbf{cl}_Y^\downarrow \bullet f = \mathbf{cl}_Y \circ (\mathbf{cl}_Y^\downarrow)^\uparrow \circ f = \mathbf{cl}_Y \circ f = f$$

where we have used Lemma A.5.4. It remains to verify the conditions hold for composite morphisms. Firstly:

$$\mathbf{cl}_Z \circ (g \bullet f) = \mathbf{cl}_Z \circ \mathbf{cl}_Z \circ g^\uparrow \circ f = \mathbf{cl}_Z \circ g^\uparrow \circ f = g \bullet f$$

and secondly:

$$\begin{aligned}\mathbf{cl}_Z \circ (g \bullet f)^\uparrow \circ \mathbf{cl}_X &= \mathbf{cl}_Z \circ (\mathbf{cl}_Z \circ g^\uparrow \circ f)^\uparrow \circ \mathbf{cl}_X \\ &= \mathbf{cl}_Z \circ (g^\uparrow \circ f)^\uparrow \circ \mathbf{cl}_X && \text{by Lemma A.5.5} \\ &= \mathbf{cl}_Z \circ g^\uparrow \circ f^\uparrow \circ \mathbf{cl}_X && \text{by Lemma A.5.1} \\ &= \mathbf{cl}_Z \circ g^\uparrow \circ \mathbf{cl}_Y \circ f^\uparrow \circ \mathbf{cl}_X && \text{by 2nd condition} \\ &= \mathbf{cl}_Z \circ g^\uparrow \circ \mathbf{cl}_Y \circ f^\uparrow && \text{by 2nd condition} \\ &= \mathbf{cl}_Z \circ g^\uparrow \circ f^\uparrow && \text{by 2nd condition} \\ &= \mathbf{cl}_Z \circ (\mathbf{cl}_Z \circ g^\uparrow \circ f)^\uparrow && \text{by Lemma A.5.5} \\ &= \mathbf{cl}_Z \circ (g \bullet f)^\uparrow\end{aligned}$$

An analogous computation proves associativity.

We are going to prove that Cls is equivalent to the category JSL_f of finite join-semilattices with all finite join preserving morphisms. But first we need to recall some constructions one can associate to closures.

Remark A.7. Closure operators $\mathbf{cl}_X : \mathcal{P}X \rightarrow \mathcal{P}X$ are precisely the monads of the poset $\mathcal{P}X$ (ordered by inclusion), viewed as a category. The Eilenberg-Moore category of \mathbf{cl}_X is the poset of closed subsets of X :

$$\text{EM}(\mathbf{cl}_X) =_{\text{def}} \{S \in \mathcal{P}X : \mathbf{cl}_X(S) = S\} = \{S \in \mathcal{P}X : \mathbf{cl}_X(S) \subseteq S\}$$

Moreover, the induced adjunction is:

$$\begin{array}{ccc} & & I \\ \text{EM}(\mathbf{cl}_X) & \xleftrightarrow{\quad \top \quad} & \mathcal{P}X \\ & & K \end{array}$$

where $I : \text{EM}(\text{cl}_X) \hookrightarrow \mathcal{P}X$ is the inclusion and K is defined as $\text{cl}_X : \mathcal{P}X \rightarrow \mathcal{P}X$ with restricted codomain $\text{EM}(\text{cl}_X)$. We have the adjoint relationship $\text{cl}_X(A) \subseteq B$ iff $A \subseteq \text{cl}_X(B)$ for any set $A \in \mathcal{P}X$ and closed set $B \in \text{EM}(\text{cl}_X)$.

Lemma A.8. *For any strict closure space (X, cl_X) the associated subposet $\text{EM}(\text{cl}_X) \subseteq \mathcal{P}X$ is a finite join-semilattice with bottom. Its bottom element is \emptyset and its binary join is defined $A_1 \oplus A_2 =_{\text{def}} \text{cl}_X(A_1 \cup A_2)$.*

Proof. Since arbitrary intersections of closed subsets of X are closed, the poset $Z = \text{EM}(\text{cl}_X)$ forms a complete sub-meet-semilattice of $\mathcal{P}X$ with top element $\bigcap \emptyset = X$. Therefore Z is also a (complete) join-semilattice with bottom element \emptyset and the following binary joins:

$$A_1 \oplus_Z A_2 = \bigcap \{B \in Z : A_1, A_2 \subseteq B\} = \bigcap \{B \in Z : A_1 \cup A_2 \subseteq B\} = \text{cl}_X(A_1 \cup A_2).$$

Definition A.9. (a) The functor $C : \text{Cls} \rightarrow \text{JSL}_f$ is defined:

$$C(X, \text{cl}_X) = \text{EM}(\text{cl}_X) \quad \frac{f : (X, \text{cl}_X) \rightarrow (Y, \text{cl}_Y)}{Cf = \text{cl}_Y \circ f^\uparrow : \text{EM}(\text{cl}_X) \rightarrow \text{EM}(\text{cl}_Y)}$$

where we have restricted the domain of $f^\uparrow : \mathcal{P}X \rightarrow \mathcal{P}Y$ to the sub join-semilattice $\text{EM}(\text{cl}_X) \subseteq \mathcal{P}X$. We have also restricted the codomain of $\text{cl}_Y : \mathcal{P}Y \rightarrow \mathcal{P}Y$ to $\text{EM}(\text{cl}_Y)$, which is certainly allowed because cl_Y constructs closed subsets of Y .

(b) The functor $G : \text{JSL}_f \rightarrow \text{Cls}$ is defined on objects S by

$$GS = (J(S), \text{cl}_{GS}), \quad \text{cl}_{GS}(A) = \{j \in J(S) : j \leq_S \bigoplus_S A\},$$

where $J(S) \subseteq S$ is the set of join-irreducibles, and on morphisms $\phi : S \rightarrow S'$ by

$$G\phi(j) = (\downarrow \phi(j)) \cap J(S') : J(S) \rightarrow \mathcal{P}(J(S')) \quad \text{for all } j \in J(S).$$

Hence $G\phi(j)$ is the set of join-irreducibles in S' that are less than or equal to $\phi(j)$.

Lemma A.10. *The functor $C : \text{Cls} \rightarrow \text{JSL}_f$ is well-defined.*

Proof. It is well-defined on objects by Lemma A.8 i.e. $CX = \text{EM}(\text{cl}_X)$ is a finite join-semilattice with bottom. We now check that the function $Cf = \text{cl}_Y \circ f^\uparrow : CX \rightarrow CY$ preserves all finite joins. Firstly:

$$Cf(\perp_{CX}) = \text{cl}_Y \circ f^\uparrow(\emptyset) = \text{cl}_Y(\emptyset) = \emptyset = \perp_{CY}$$

Secondly:

$$\begin{aligned} Cf(A_1 \oplus_{CX} A_2) &= \text{cl}_Y \circ f^\uparrow \circ \text{cl}_X(A_1 \cup A_2) && \text{def. of } Cf \text{ and } \oplus_{CX} \\ &= \text{cl}_Y \circ f^\uparrow(A_1 \cup A_2) && \text{second condition} \\ &= \text{cl}_Y(f^\uparrow(A_1) \cup f^\uparrow(A_2)) && f^\uparrow \text{ preserves unions} \\ &= \text{cl}_Y(\text{cl}_Y(f^\uparrow(A_1)) \cup \text{cl}_Y(f^\uparrow(A_2))) && \text{Lemma A.5.6} \\ &= Cf(A_1) \oplus_{CY} Cf(A_2) \end{aligned}$$

Moreover we have $C \text{id}_X = C \text{cl}_X^\perp = \text{cl}_X \circ (\text{cl}_X^\perp)^\uparrow = \text{cl}_X = \text{id}_{CX}$ by Lemma A.5.4 and the fact that cl_X is the identity morphism on $CX = \text{EM}(\text{cl}_X)$. Finally we show C preserves composition of morphisms.

$$\begin{aligned}
C(g \bullet f) &= \text{cl}_Z \circ (g \bullet f)^\uparrow \\
&= \text{cl}_Z \circ (\text{cl}_Z \circ g^\uparrow \circ f)^\uparrow \\
&= \text{cl}_Z \circ (g^\uparrow \circ f)^\uparrow && \text{by Lemma A.5.5} \\
&= \text{cl}_Z \circ g^\uparrow \circ f^\uparrow && \text{by Lemma A.5.1} \\
&= \text{cl}_Z \circ g^\uparrow \circ \text{cl}_Y \circ f^\uparrow && \text{by second condition} \\
&= Cg \circ Cf
\end{aligned}$$

Lemma A.11. *The functor $G : \text{JSL}_f \rightarrow \text{Cls}$ is well-defined.*

Proof. Given $S \in \text{JSL}_f$, we first verify that $\text{cl}_{GS} : \mathcal{P}(J(S)) \rightarrow \mathcal{P}(J(S))$ is a strict closure operator. If $A \subseteq B \subseteq J(S)$ then if $j \in \text{cl}_{GS}(A)$ we have $j \leq_S \bigoplus_S A \leq_S \bigoplus_S B$ and hence $j \in \text{cl}_{GS}(B)$, so the closure is monotonic. Next $A \subseteq \text{cl}_{GS}(A)$ because if $j \in A$ then we certainly have $j \leq_S \bigoplus_S A$. Next, cl_{GS} is idempotent because the sum of all join-irreducibles less than or equal to $\bigoplus_S A$ is itself less than or equal to $\bigoplus_S A$. Finally it is also strict i.e. $\text{cl}_{GS}(\emptyset) = \emptyset$, since by definition \perp_S is not join-irreducible.

Hence G is well-defined on objects. To verify that it is well-defined on morphisms, let $\phi : S \rightarrow S'$ be a join-semilattice morphism. By definition $G\phi : J(S) \rightarrow \mathcal{P}(J(S'))$ sends $j \in J(S)$ to the set of join-irreducibles in S' less than or equal to $\phi(j)$. We show that $G\phi$ is a Cls-morphism. Indeed, for the first condition we compute:

$$\begin{aligned}
\text{cl}_{GS'} \circ G\phi(j) &= \text{cl}_{GS'}(\{j' \in J(S') : j' \leq_{S'} \phi(j)\}) \text{ def. } G\phi \\
&= \{j' \in J(S') : j' \leq_{S'} \phi(j)\} && \text{def. } \text{cl}_{GS'}, \phi(j) \text{ join of jirrs} \\
&= G\phi(j) && \text{def. } G\phi
\end{aligned}$$

In the second step we use that $\phi(j) = \bigoplus_{S'} \{j' \in J(S') : j' \leq_{S'} \phi(j)\}$ and hence $G\phi(j)$ is already closed in GS' . To verify the second condition we consider each side of the equation: for $A \subseteq J(S)$ we have

$$\begin{aligned}
&\text{cl}_{GS'} \circ (G\phi)^\uparrow \circ \text{cl}_{GS}(A) \\
&= \text{cl}_{GS'} \circ (G\phi)^\uparrow(\{j \in J(S) : j \leq_S \bigoplus_S A\}) \\
&= \text{cl}_{GS'}(\bigcup\{(\downarrow_{S'} \phi(j)) \cap J(S') : j \in J(S), j \leq_S \bigoplus_S A\}) \\
&= \{j' \in J(S') : j' \leq_{S'} \bigoplus_{S'} \bigcup\{(\downarrow_{S'} \phi(j)) \cap J(S') : j \in J(S), j \leq_S \bigoplus_S A\}\}
\end{aligned}$$

$$\begin{aligned}
&\text{cl}_{GS'} \circ (G\phi)^\uparrow(A) \\
&= \text{cl}_{GS'}(\bigcup\{(\downarrow_{S'} \phi(j)) \cap J(S') : j \in A\}) \\
&= \{j' \in J(S') : j' \leq_{S'} \bigoplus_{S'} \bigcup\{(\downarrow_{S'} \phi(j)) \cap J(S') : j \in A\}\}
\end{aligned}$$

Given any join-irreducible j' less than or equal to the upper sum $\bigoplus_{S'} \bigcup\{(\downarrow_{S'} \phi(j)) \cap J(S') : j \in J(S), j \leq_S \bigoplus_S A\}$, then it is necessarily less than or equal to $\phi(\bigoplus_S A)$ because every $\phi(j)$ in the sum is too. Then $j' \leq_{S'} \bigoplus_{a \in A} \phi(a)$ because ϕ preserves joins. It follows that j' is less than or equal to the lower sum $\bigoplus_{S'} \bigcup\{(\downarrow_{S'} \phi(j)) \cap J(S') : j \in A\}$ since the latter's summands contain every join-irreducible

less than or equal to any $\phi(a)$. This proves the inclusion $\mathbf{cl}_{GS'} \circ (G\phi)^\uparrow \circ \mathbf{cl}_{GS}(A) \subseteq \mathbf{cl}_{GS'} \circ (G\phi)^\uparrow(A)$ and the reverse inclusion follows by monotonicity.

Finally, we show that G preserves identities and composition. First, we have

$$G\text{id}_S(j) = \{j' \in J(S) : j' \leq_S j\} = \mathbf{cl}_{GS}^\downarrow(j) = \text{id}_{GS}.$$

Furthermore, given $\phi : S \rightarrow S'$ and $\psi : S \rightarrow S''$ in JSL_f , we prove $G(\psi \circ \phi) = G\psi \bullet G\phi$. Indeed, for all $j \in J(S)$ and $j'' \in J(S'')$ we have

$$\begin{aligned} & j'' \in (G\psi \bullet G\phi)(j) \\ \Leftrightarrow & j'' \in \mathbf{cl}_{GS''} \circ (G\psi)^\uparrow \circ G\phi(j) && \text{def. } \bullet \\ \Leftrightarrow & j'' \leq \oplus_{S''} ((G\psi)^\uparrow \circ G\phi(j)) && \text{def. } \mathbf{cl}_{GS''} \\ & = \oplus_{S''} \{\tilde{j} \in J(S'') : \exists j' \in J(S') . (j' \leq_{S'} \phi(j) \wedge \tilde{j} \leq \psi(j'))\} && \text{def. } G, (-)^\uparrow \\ & = \oplus_{S''} \{\psi(j') : j' \in J(S') \text{ and } j' \leq_{S'} \phi(j)\} && (*) \\ & = \psi \circ \phi(j) && (**) \\ \Leftrightarrow & j'' \in G(\psi \circ \phi)(j) && \text{def. } G \end{aligned}$$

For the step (*) we used that $\psi(j') = \oplus_{S''} \{\tilde{j} \in J(S'') : \tilde{j} \leq \psi(j')\}$. And for (**) we used that $\phi(j) = \oplus_{S'} \{j' \in J(S') : j' \leq \phi(j)\}$ and ψ preserves this finite join.

Theorem A.12. *The functors $C : \text{Cls} \rightarrow \text{JSL}_f$ and $G : \text{JSL}_f \rightarrow \text{Cls}$ define an equivalence of categories. The unit and counit are defined:*

$$\begin{aligned} \eta : \text{Id} \Rightarrow CG \quad \eta_S : S \rightarrow \text{EM}(\mathbf{cl}_{GS}) \quad \eta_S(s) &= \{j \in J(S) : j \leq_S s\} \\ \varepsilon : GC \Rightarrow \text{Id} \quad \varepsilon_X : J(\text{EM}(\mathbf{cl}_X)) \rightarrow \mathcal{P}X \quad \varepsilon_X(A) &= A \end{aligned}$$

Proof. We first consider η . Each component η_S is a well-defined function because $\eta_S(s) = \mathbf{cl}_{GS}(J(S) \cap \downarrow_S s)$ is closed in GS . Furthermore $\eta_S(\perp_S) = \emptyset = \perp_{CGS}$ and

$$\begin{aligned} \eta_S(s_1 \oplus_S s_2) &= \{j \in J(S) : j \leq_S s_1 \oplus_S s_2\} \\ &= \mathbf{cl}_{GS}(\{j \in J(S) : j \leq_S s_1\} \cup \{j \in J(S) : j \leq_S s_2\}) \\ &= \eta_S(s_1) \oplus_{CGS} \eta_S(s_2) \end{aligned}$$

so it is a well-defined join-semilattice morphism. Moreover, η_S is an isomorphism. Indeed, if $\eta_S(s) = \eta_S(s')$ then s and s' have the same generators beneath them and hence $s = s'$, which proves that η_S is injective. And given any closed set $\mathbf{cl}_{GS}(A) \in CGS$ where $A \subseteq J(S)$ then $\eta_S(\bigoplus_S A) = \mathbf{cl}_{GS}(A)$, so η_S is also surjective. Next we prove naturality of η , i.e., $CG\phi \circ \eta_S = \eta_{S'} \circ \phi$ for every join-semilattice morphism $\phi : S \rightarrow S'$. We compute:

$$\begin{aligned} CG\phi \circ \eta_S(s) &= CG\phi(J(S) \cap \downarrow_S s) \\ &= \mathbf{cl}_{GS} \circ (G\phi)^\uparrow (J(S) \cap \downarrow_S s) \\ &= \mathbf{cl}_{GS}(\{j' \in J(S') : \exists j \in J(S) . (j \leq_S s \wedge j' \leq_{S'} \phi(j))\}) \\ &= \{j' \in J(S') : j' \leq_{S'} \phi(s)\} \\ &= \eta_{S'} \circ \phi(s) \end{aligned}$$

We explain the penultimate step. Given any $j' \in J(S')$ and $j \in J(S)$ such that $j \leq_S s$ and $j' \leq_{S'} \phi(j)$ then by monotonicity $j' \leq_{S'} \phi(j) \leq_{S'} \phi(s)$. Conversely given any $j' \in J(S')$ with $j' \leq_{S'} \phi(s)$ then we have the inequality $j' \leq_{S'} \bigoplus_{S'} \bigcup_{j \leq_S s} \{k' \in J(S') : k' \leq_{S'} \phi(j)\}$ because $\phi(s) = \bigoplus_{j \leq_S s} \phi(j)$ i.e. the lower set is a subset of the upper set. Then we have proved that η is natural isomorphism.

Next we show each ε_X is a well-defined Cls-morphism. Firstly each $\varepsilon_X(A) = A$ is closed in X because $J(\text{EM}(\text{cl}_X)) \subseteq \text{EM}(\text{cl}_X)$ is a subcollection of the closed sets in X . Secondly suppose we have $A \in J(\text{EM}(\text{cl}_X))$ and $\mathcal{A} \subseteq J(\text{EM}(\text{cl}_X))$ such that A lies in the closure of \mathcal{A} . Then we must show that $\varepsilon_X(A) = A$ lies in the closure of $\varepsilon_X^\uparrow(\mathcal{A}) = \mathcal{A}$ which is immediate. To show ε_X is a Cls-isomorphism we explicitly define its inverse:

$$\varepsilon_X^{-1} : X \rightarrow \mathcal{P}(J(\text{EM}(\text{cl}_X))) \quad \varepsilon_X^{-1}(x) = \{A \in J(\text{EM}(\text{cl}_X)) : A \subseteq \text{cl}_X^\downarrow(x)\}$$

Let us check it is a well-defined Cls-morphism. Firstly $\text{cl}_X^\downarrow(x)$ is closed in X and therefore arises as the union (join) of some join-irreducibles $A_i \in J(\text{EM}(\text{cl}_X)) \subseteq \text{EM}(\text{cl}_X)$. Thus $\varepsilon_X^{-1}(x) = \text{cl}_{GCX}(\{A_i : i \in I\})$ is closed in GCX . To verify the second condition we first calculate:

$$\begin{aligned} \text{cl}_{GCX} \circ (\varepsilon_X^{-1})^\uparrow(B) &= \text{cl}_{GCX} \circ \bigcup \{\varepsilon_X^{-1}(b) : b \in B\} \\ &= \text{cl}_{GCX}(\{A \in J(CX) : \exists b \in B. A \subseteq \text{cl}_X^\downarrow(b)\}) \\ &= \{A \in J(CX) : A \leq_{CX} \bigoplus_{CX} \{A' \in J(CX) : b \in B, A' \subseteq \text{cl}_X^\downarrow(b)\}\} \\ &= \{A \in J(CX) : A \leq_{CX} \bigoplus_{CX} \{\text{cl}_X^\downarrow(b) : b \in B\}\} \\ &= \{A \in J(CX) : A \subseteq \text{cl}_X(\bigcup \{\text{cl}_X^\downarrow(b) : b \in B\})\} \\ &= \{A \in J(CX) : A \subseteq \text{cl}_X(B)\} \end{aligned}$$

Then it follows that:

$$\begin{aligned} \text{cl}_{GCX} \circ (\varepsilon_X^{-1})^\uparrow \circ \text{cl}_X(B) &= \{A \in J(CX) : A \subseteq \text{cl}_X \circ \text{cl}_X(B)\} \\ &= \{A \in J(CX) : A \subseteq \text{cl}_X(B)\} \\ &= \text{cl}_{GCX} \circ (\varepsilon_X^{-1})^\uparrow(B) \end{aligned}$$

so we know ε_X^{-1} is a well-defined Cls-morphism. Next we check ε_X^{-1} is the inverse of ε_X :

$$\begin{aligned} \varepsilon_X \bullet \varepsilon_X^{-1}(x) &= \text{cl}_X \circ \varepsilon_X^\uparrow \circ \varepsilon_X^{-1}(x) \\ &= \text{cl}_X \circ \varepsilon_X^\uparrow(\{A \in J(CX) : A \subseteq \text{cl}_X^\downarrow(x)\}) \\ &= \text{cl}_X(\bigcup \{A \in J(CX) : A \subseteq \text{cl}_X^\downarrow(x)\}) \\ &= \text{cl}_X^\downarrow(x) \\ &= \text{id}_X(x) \end{aligned}$$

$$\begin{aligned} \varepsilon_X^{-1} \bullet \varepsilon_X(B) &= \text{cl}_{GCX} \circ (\varepsilon_X^{-1})^\uparrow \circ \varepsilon_X(B) \\ &= \text{cl}_{GCX} \circ (\varepsilon_X^{-1})^\uparrow(B) \\ &= \{A \in J(CX) : A \subseteq \text{cl}_X(B)\} && \text{by earlier calculation} \\ &= \{A \in J(CX) : A \leq_{CX} \bigoplus_{CX} \{B\}\} \\ &= \text{cl}_{GCX}^\downarrow(B) \\ &= \text{id}_{GCX}(B) \end{aligned}$$

Thus each component ε_X is a well-defined Cls-isomorphism. It only remains to check that ε defines a natural transformation: for any Cls-morphism $f : X \rightarrow \mathcal{P}Y$ we must show that $f \bullet \varepsilon_X = \varepsilon_Y \bullet GCf$. Indeed,

$$\begin{aligned}
\varepsilon_Y \bullet GCf(A) &= \mathbf{cl}_Y \circ \varepsilon_Y^\uparrow \circ GCf(A) \\
&= \mathbf{cl}_Y \circ \varepsilon_Y^\uparrow (J(CY) \cap \downarrow_{CY} Cf(A)) \\
&= \mathbf{cl}_Y (\cup (J(CY) \cap \downarrow_{CY} Cf(A))) \\
&= \mathbf{cl}_Y (\cup (J(CY) \cap \downarrow_{CY} (\mathbf{cl}_Y \circ f^\uparrow)(A))) \\
&= \mathbf{cl}_Y (\cup \{A' \in J(CY) : A' \leq_{CY} \mathbf{cl}_Y \circ f^\uparrow(A)\}) \\
&= \bigoplus_{CY} \{A' \in J(CY) : A' \leq_{CY} \mathbf{cl}_Y \circ f^\uparrow(A)\} \\
&= \mathbf{cl}_Y \circ f^\uparrow(A) \\
&= \mathbf{cl}_Y \circ f^\uparrow \circ \varepsilon_X(A) \\
&= f \bullet \varepsilon_X(A)
\end{aligned}$$

So both η and ε are natural isomorphisms and the equivalence is proved.