

Rank-1 Modal Logics are Coalgebraic

Lutz Schröder¹ and Dirk Pattinson²

¹ Department of Computer Science, University of Bremen, and DFKI Lab Bremen

² Department of Computing, Imperial College London

Abstract. Coalgebras provide a unifying semantic framework for a wide variety of modal logics. It has previously been shown that the class of coalgebras for an endofunctor can always be axiomatised in rank 1. Here we establish the converse, i.e. every rank 1 modal logic has a sound and *strongly* complete coalgebraic semantics. As a consequence, recent results on coalgebraic modal logic, in particular generic decision procedures and upper complexity bounds, become applicable to arbitrary rank 1 modal logics, without regard to their semantic status; we thus obtain purely syntactic versions of these results. As an extended example, we apply our framework to recently defined deontic logics.

Introduction

In recent years, coalgebras have received a steadily growing amount of attention as general models of state-based systems [18], encompassing such diverse systems as labelled transition systems, probabilistic systems, game frames, and neighborhood frames [21]. On the logical side, modal logic has emerged as the adequate specification language for coalgebraically modelled systems. A variety of different frameworks have been proposed; here, we work with *coalgebraic modal logic* [15], which allows for a high level of generality while retaining a close relationship to the established syntactic and semantic tradition of modal logic.

In fact, one can reverse the viewpoint that coalgebraic modal logic is a specification language for coalgebras and regard coalgebra as a generic semantics for modal logics of essentially arbitrary nature, including non-normal and non-monotone ones. Under this perspective, coalgebraic modal logic is a generic notion of modal logic that subsumes e.g. Hennessy-Milner logic, graded modal logic [4], majority logic [13], probabilistic modal logic [12, 7], and coalition logic [16], but also modal operators of higher arity as e.g. in conditional logic [3].

It has been shown in [20] that every coalgebraic modal logic can be axiomatized by formulas of rank 1, i.e. with nesting depth of modal operators uniformly equal to 1 (logics of arbitrary rank are obtained by restricting the relevant class of coalgebras, which play the role of generic frames); such axioms may be regarded as concerning precisely the single next transition step. Here, we establish the converse: given a modal logic \mathcal{L} of rank 1, we construct a functor $M_{\mathcal{L}}$ that provides a sound and strongly complete semantics for \mathcal{L} ; i.e. *coalgebraic modal logic subsumes all rank-1 modal logics*. The functor $M_{\mathcal{L}}$, which can be viewed

as a generalization of the neighbourhood frame functor, is moreover a canonical semantics for \mathcal{L} in a precise categorical sense; a finitary modification of $M_{\mathcal{L}}$ provides a canonical finitely branching semantics.

Besides rounding off the picture in a pleasant way, these results make the rapidly expanding meta-theory of coalgebraic modal logic applicable to arbitrary rank-1 modal logics, even when the latter are given purely syntactically or equipped with a semantics that fails to be, or has not yet been recognized as, coalgebraic. This includes results on the Hennessy-Milner property [19] and bisimulation-somewhere-else [10], and most notably generic decidability and complexity results [20, 21], of which we now obtain purely syntactic versions. As an extended example, we discuss applications of these results to recently defined variants of deontic logic [5].

1 Coalgebraic Modal Logic

We briefly recapitulate the basics of the coalgebraic semantics of modal logic. Coalgebraic modal logic in the form considered here has been introduced in [15], generalising previous results [9, 17, 11, 14]. For the sake of readability, we restrict the exposition to unary modalities. However, we emphasize that all our results extend in a straightforward way to polyadic operators, found e.g. in conditional and default logics [19].

A *modal signature* is just a set Λ of (unary) modal operators. The set $\mathcal{F}(\Lambda)$ of Λ -formulas ϕ is defined by the grammar

$$\phi ::= \perp \mid \phi \wedge \psi \mid \neg\phi \mid L\phi$$

where L ranges over all modalities in Λ . Other boolean operations are defined as usual; propositional atoms can be expressed as constant modalities.

Generally, we denote the set of propositional formulas over a set V by $\text{Prop}(V)$, generated by the basic connectives \neg and \wedge , and the set of propositional tautologies by $\text{Taut}(V)$. We use variables ϵ etc. to denote either nothing or \neg . Thus, a *literal* over V is a formula of the form ϵa , with $a \in V$. A *clause* is a finite, possibly empty, disjunction of literals. The set of all clauses over V is denoted by $\text{Cl}(V)$. We denote by $\text{Up}_{\Lambda}(V)$ the set $\{La \mid L \in \Lambda, a \in V\}$. If $V \subseteq \mathcal{F}(\Lambda)$, we also regard propositional formulas over V as Λ -formulas. We sometimes explicitly designate V as consisting of *propositional variables*; these retain their status across further applications of Up_{Λ} and Prop (e.g. V is also the set of propositional variables for $\text{Up}_{\Lambda}(\text{Prop}(V))$). An *L -substitution* is a substitution σ of the propositional variables by elements of a set L ; for a formula ϕ over V , we call $\phi\sigma$ an *L -instance* of ϕ . If $L \subset \mathcal{P}(X)$ for some X , then we also refer to σ as an *L -valuation*.

Definition 1. A *rank-1 clause* (in Λ) over a set V of propositional variables is an element of $\text{Cl}(\text{Up}_{\Lambda}(\text{Prop}(V)))$. A *rank-1 (modal) logic* is a pair $\mathcal{L} = (\Lambda, A)$, where A is a set of rank-1 clauses in Λ .

Note that the definition of rank-1 clause rules out axioms involving purely propositional components, such as $\Box a \rightarrow a$ (results covering also such more general axioms are under way); the archetypal rank-1 logic is K .

Given a rank-1 logic $\mathcal{L} = (A, A)$, we inductively define \mathcal{L} -derivability $\vdash_{\mathcal{L}}$ from a set $\Phi \subseteq \mathcal{F}(A)$ as follows:

$$\frac{\phi \in \mathbf{Taut}(\mathcal{F}(A))}{\Phi \vdash_{\mathcal{L}} \phi} \quad \frac{\phi \in \Phi}{\Phi \vdash_{\mathcal{L}} \phi} \quad \frac{\Phi \vdash_{\mathcal{L}} \phi, \phi \rightarrow \psi}{\Phi \vdash_{\mathcal{L}} \psi} \quad \frac{\psi \in A}{\Phi \vdash_{\mathcal{L}} \psi\sigma} \quad \frac{\Phi \vdash_{\mathcal{L}} \phi \leftrightarrow \psi}{\Phi \vdash_{\mathcal{L}} L\phi \leftrightarrow L\psi}$$

where σ is an $\mathcal{F}(A)$ -substitution. The last rule above is referred to as the *congruence rule*. We write $\vdash_{\mathcal{L}} \psi$ instead of $\emptyset \vdash_{\mathcal{L}} \psi$.

It has been shown in [20] that rank-1 clauses may be equivalently replaced by *one-step rules* ϕ/ψ , where $\phi \in \mathbf{Prop}(V)$ and $\psi \in \mathbf{Cl}(\mathbf{Up}_A(V))$. We shall present rank-1 logics as pairs $\mathcal{L} = (A, R)$, with R a set of one-step rules, when convenient; in this case, the penultimate clause above is replaced by

$$\frac{\Phi \vdash_{\mathcal{L}} \phi\sigma \quad \phi/\psi \in R}{\Phi \vdash_{\mathcal{L}} \psi\sigma}.$$

The extension of R by the congruence rule is denoted R_C . Coalgebraic modal logic interprets modal formulas over coalgebras, which abstract from concrete notions of reactive system; here, the interpretation of modalities is given by a choice of predicate liftings. We recall the formal definitions:

Definition 2. Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be a functor, referred to as the *signature functor*, where \mathbf{Set} is the category of sets. A T -coalgebra is a pair $C = (X, \xi)$ where X is a set (of *states*) and ξ is a function $X \rightarrow TX$ called the *transition function*. A *morphism* $(X_1, \xi_1) \rightarrow (X_2, \xi_2)$ of T -coalgebras is a map $f : X_1 \rightarrow X_2$ such that $\xi_2 \circ f = Tf \circ \xi_1$. States x, y in coalgebras C, D are *behaviourally equivalent* if there exist coalgebra morphisms $f : C \rightarrow E$ and $g : D \rightarrow E$ such that $f(x) = g(y)$. A *predicate lifting* for T is a natural transformation $\lambda : Q \rightarrow Q \circ T^{\text{op}}$, where Q denotes the contravariant powerset functor $\mathbf{Set}^{\text{op}} \rightarrow \mathbf{Set}$.

We view coalgebras as generalised transition systems: the transition function maps states to a structured set of successors

Assumption 3. We can assume w.l.o.g. that T preserves injective maps ([2], proof of Theorem 3.2). For convenience, we will in fact sometimes assume that $TX \subseteq TY$ if $X \subseteq Y$. Moreover, we assume that T is non-trivial, i.e. $TX = \emptyset \implies X = \emptyset$ (otherwise, $TX = \emptyset$ for all X).

Recall that a functor is ω -accessible if it preserves directed colimits.

Lemma 4. ([1], Proposition 5.2) *For a set functor T , the following are equivalent:*

1. T is ω -accessible
2. T preserves directed unions
3. For every set X , $TX = \bigcup_{Y \subseteq X \text{ finite}} TY$ (recall Assumption 3).

The coalgebraic semantics of modal logics is defined as follows. Given a modal signature Λ , a Λ -structure consists of a signature functor T and an assignment of a predicate lifting $\llbracket L \rrbracket$ for T to every modal operator $L \in \Lambda$. The satisfaction relation \models_C between states x of a T -coalgebra $C = (X, \xi)$ and Λ -formulas is defined inductively, with the usual clauses for the boolean operations. The clause for a modal operator L is

$$x \models_C L\phi \iff \xi(x) \in \llbracket L \rrbracket \llbracket \phi \rrbracket_C,$$

where $\llbracket \phi \rrbracket_C = \{x \in X \mid x \models_C \phi\}$. We drop the subscripts C when C is clear from the context. When we speak of a coalgebraic modal logic informally, we mean a Λ -structure; if the interpretation of modalities is clear from the context, this structure is simply referred to as T .

Satisfaction of Λ -formulas is invariant under behavioural equivalence [15]. Conversely, Λ has the *Hennessey-Milner property* for T , i.e. states that satisfy the same Λ -formulas are behaviourally equivalent, if T is ω -accessible and Λ is *separating* in the sense that $t \in TX$ is determined by the set $\{(L, A) \in \Lambda \times \mathcal{P}(X) \mid t \in \llbracket L \rrbracket(A)\}$ [19].

Definition 5. Given a Λ -structure T , we write $\Phi \models_T \psi$ for a Λ -formula ψ and a set $\Phi \subseteq \mathcal{F}(\mathcal{L})$ if, for every state x in every T -coalgebra, $x \models \psi$ whenever $x \models \Phi$ (i.e. $x \models \phi$ for all $\phi \in \Phi$). The logic \mathcal{L} is *sound* over T if $\Phi \models_T \psi$ whenever $\Phi \vdash_{\mathcal{L}} \psi$, and *strongly (weakly) complete* if $\Phi \vdash_{\mathcal{L}} \psi$ ($\vdash_{\mathcal{L}} \psi$) whenever $\Phi \models_T \psi$ ($\emptyset \models_T \psi$).

The requirement that axioms are of rank 1 means that every axiom makes assertions precisely about the next transition step. This allows us to capture soundness as a property exhibited in a single transition step as follows. Given a set X and a $\mathcal{P}(X)$ -valuation τ , we define interpretations $\llbracket \phi \rrbracket_{\tau} \subseteq X$ and $\llbracket \psi \rrbracket_{\tau} \subseteq TX$ for $\phi \in \mathbf{Prop}(V)$ and $\psi \in \mathbf{Prop}(\mathbf{Up}_{\Lambda}(\mathbf{Prop}(V)))$ by the usual clauses for boolean operators and by $\llbracket L\phi \rrbracket_{\tau} = \llbracket L \rrbracket \llbracket \phi \rrbracket_{\tau}$. We write $X, \tau \models \phi$ if $\llbracket \phi \rrbracket_{\tau} = X$, correspondingly for TX .

Definition 6. A rank-1 clause ψ (one-step rule ϕ/ψ) is *one-step sound* for a Λ -structure T if $TX, \tau \models \psi\tau$ for each set X and each $\mathcal{P}(X)$ -valuation τ (such that $X, \tau \models \phi$). An \mathcal{L} -structure for a rank-1 logic \mathcal{L} with signature Λ is a Λ -structure for which all axioms (or rules) of \mathcal{L} are one-step sound.

It is easy to see that one-step soundness implies soundness, so \mathcal{L} is *sound for all \mathcal{L} -structures*. Additional conditions guarantee weak completeness [20]. In general, this is all one can hope for, as many coalgebraic modal logics fail to be compact [20]. However, it will turn out that \mathcal{L} is indeed *strongly* complete for the canonical \mathcal{L} -structure constructed below.

Example 7. We give a brief description of some coalgebraic modal logics, illustrating in particular the fact that many interesting modal logics are axiomatised in rank 1. We mostly omit the definition of predicate liftings and the axiomatisations; for these and further examples, cf. [20, 21].

1. The Kripke semantics of the modal logic K , defined in terms of a single operator \Box and the axioms $\Box\top$ and $\Box(a \rightarrow b) \rightarrow \Box a \rightarrow \Box b$, is obtained as the K -structure given by the covariant powerset functor \mathcal{P} and $\llbracket \Box \rrbracket(A) = \mathcal{P}(A) \subset \mathcal{P}(X)$ for $A \subseteq X$.

2. *Graded modal logic* (GML) [4] has operators \Diamond_k for $k \in \mathbb{N}$ of the nature ‘in more than k successor states, it is the case that’. A GML-structure is given by the finite multiset functor, which takes a set X to the set of maps $A : X \rightarrow \mathbb{N}$ with finite support, where $A(x) = n$ is read ‘multiset A contains x with multiplicity n ’. Over this functor, one can also interpret the additional operator W of *majority logic* [13], read ‘in at least half of the successor states, it is the case that’.

3. *Probabilistic modal logic* (PML) [12, 7] has modal operators L_p for $p \in [0, 1] \cap \mathbb{Q}$, read ‘in the next step, it is with probability at least p the case that’. A PML-structure is given by the finite distribution functor, which takes a set X to the set of finitely supported probability distributions over X .

2 From Rank-1 Logics to Coalgebraic Models

In this section we construct for a given rank-1 modal logic \mathcal{L} a *canonical \mathcal{L} -structure* $M_{\mathcal{L}}$ for which \mathcal{L} is (sound and) *strongly complete*. Moreover, we consider a finitely branching substructure $M_{\mathcal{L}}^{fin}$ of $M_{\mathcal{L}}$ which is canonical among the finitely branching \mathcal{L} -structures. For $M_{\mathcal{L}}^{fin}$, \mathcal{L} is (sound and) weakly complete and has the Hennessy-Milner property, i.e. states satisfying the same formulas are behaviourally equivalent. This tradeoff is typical: the Hennessy-Milner property holds only over finitely branching systems, while strong completeness will fail over such systems due to the breakdown of compactness.

The construction of the canonical structure resembles the construction of canonical models using maximally consistent sets, but works, like many concepts explained in the previous section, at the single step level:

Definition 8. Let $\mathcal{L} = (A, A)$ be a rank-1 logic, and let X be a set. *One-step derivability* $\Phi \vdash_{\mathcal{L}}^X \psi$ of $\psi \in \text{Prop}(\text{Up}_A(\mathcal{P}(X)))$ from $\Phi \subseteq \text{Prop}(\text{Up}_A(\mathcal{P}(X)))$ is defined inductively by

$$\frac{\phi \in \Phi}{\Phi \vdash_{\mathcal{L}}^X \phi} \quad \frac{\phi \in \text{Taut}(\text{Up}_A(\mathcal{P}(X)))}{\Phi \vdash_{\mathcal{L}}^X \phi} \quad \frac{\Phi \vdash_{\mathcal{L}}^X \phi \rightarrow \psi \quad \Phi \vdash_{\mathcal{L}}^X \phi}{\Phi \vdash_{\mathcal{L}}^X \psi} \quad \frac{\psi \in A}{\Phi \vdash_{\mathcal{L}}^X \psi\tau}$$

where τ is a $\mathcal{P}(X)$ -valuation. (In the last clause, elements of $\text{Prop}(\mathcal{P}(X))$ are implicitly interpreted as elements of $\mathcal{P}(X)$ in the obvious way. If \mathcal{L} is presented by rules ϕ/ψ , the last clause is modified accordingly, with additional premise $X, \tau \models \phi$.) The set Φ is *one-step consistent* if $\Phi \not\vdash_{\mathcal{L}}^X \perp$, and *maximally one-step consistent* if Φ is maximal w.r.t. \subseteq among the one-step consistent subsets of $\text{Prop}(\text{Up}_A(\mathcal{P}(X)))$.

The canonical \mathcal{L} -structure $M_{\mathcal{L}}$ for \mathcal{L} is now given by the functor $M_{\mathcal{L}}$ that takes a set X to the set of maximally one-step consistent subsets of $\text{Prop}(\text{Up}_A(\mathcal{P}(X)))$. For a map $f : X \rightarrow Y$, $M_{\mathcal{L}}(f)$ is defined by

$$M_{\mathcal{L}}(f)(\Phi) = \{\phi \in \text{Prop}(\text{Up}_A(\mathcal{P}(Y))) \mid \phi \circ f \in \Phi\},$$

where σ_f is the substitution $A \mapsto f^{-1}[A]$. This definition is justified by

Lemma 9. *For $\Phi \in M_{\mathcal{L}}(X)$, the set $M_{\mathcal{L}}(f)(\Phi)$ is maximally one-step consistent.*

Remark 10. From the point of view of Stone duality, a rank-1 logic defines a functor $L : \mathbf{BA} \rightarrow \mathbf{BA}$ on the category \mathbf{BA} of boolean algebras. In this framework, the functor $M_{\mathcal{L}}$ arises as the composition $M_{\mathcal{L}} = USL\bar{Q}$ where $\bar{Q} : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{BA}$ is the contravariant powerset functor, $S : \mathbf{BA}^{\text{op}} \rightarrow \mathbf{Stone}$ is part of the duality between Stone spaces and boolean algebras, and $U : \mathbf{Stone} \rightarrow \mathbf{Set}$ is the forgetful functor; see [10] for details.

The interpretation of modal operators by predicate liftings for $M_{\mathcal{L}}$ is now obvious:

Theorem and Definition 11. *The assignment*

$$\llbracket L \rrbracket A = \{\Phi \in M_{\mathcal{L}}(X) \mid LA \in \Phi\}.$$

defines an \mathcal{L} -structure $M_{\mathcal{L}}$, the canonical \mathcal{L} -structure.

Note that this immediately implies soundness of \mathcal{L} over $M_{\mathcal{L}}$. We now turn to *strong completeness*, which is established by a canonical model construction that generalizes the standard notion of canonical Kripke structure. The carrier of the canonical model is the set C of maximally consistent sets of \mathcal{L} -formulas. The key to the construction is the *existence proof* (rather than the explicit construction) of a suitable $M_{\mathcal{L}}$ -coalgebra structure on C , a technique first employed in [20]:

Lemma and Definition 12 (Existence Lemma). *There exists a canonical model, i.e. an $M_{\mathcal{L}}$ -coalgebra structure $\zeta : C \rightarrow M_{\mathcal{L}}C$ such that*

$$\zeta(A) \in \llbracket L \rrbracket \hat{\phi} \quad \text{iff} \quad L\phi \in A$$

for all $L \in \Lambda$, $\phi \in \mathcal{L}$, $A \in C$, where $\hat{\phi} = \{B \in C \mid \phi \in B\}$.

Lemma 13 (Truth Lemma). *For canonical models (C, ζ) , $A \models_{(C, \zeta)} \phi$ iff $\phi \in A$.*

Theorem 14 (Strong completeness). *The logic \mathcal{L} is strongly complete for $M_{\mathcal{L}}$.*

Finally, we consider the Hennessy-Milner property (cf. Section 1). The functor $M_{\mathcal{L}}$ fails to be ω -accessible for obvious cardinality reasons. Intuitively, $M_{\mathcal{L}}$ -models have unbounded branching, while the Hennessy-Milner property can only be expected for finitely branching systems (as is the case already for standard Kripke models). We thus consider a subfunctor $M_{\mathcal{L}}^{\text{fin}}$ of $M_{\mathcal{L}}$ that captures precisely the finitely branching models.

In order to construct $M_{\mathcal{L}}^{\text{fin}}$, we can rely on the following general mechanism. We define the ω -accessible part T^{fin} of a set functor T by

$$T^{\text{fin}} X = \bigcup_{Y \subseteq X \text{ finite}} TY \quad \subseteq TX$$

(recall Assumption 3). It is easy to see that T^{fin} is a subfunctor of T . By Lemma 4, T^{fin} is ω -accessible. Moreover, T^{fin} agrees with T on finite sets. A predicate lifting λ for T restricts to a predicate lifting λ^{fin} for T^{fin} given by $\lambda_X^{fin} A = \lambda_X A \cap T^{fin} X$.

We define the *canonical finitely branching \mathcal{L} -structure* $M_{\mathcal{L}}^{fin}$ as the ω -accessible part of $M_{\mathcal{L}}$, with modal operators interpreted by restricted predicate liftings as described above. We then obtain

Theorem 15. *\mathcal{L} is weakly complete and has the Hennessy-Milner property for $M_{\mathcal{L}}^{fin}$.*

Example 16. We give explicit descriptions (up to natural isomorphism) of $M_{\mathcal{L}}$ and $M_{\mathcal{L}}^{fin}$ in some concrete cases.

1. For $\mathcal{L} = (\{\Box\}, \emptyset)$, $M_{\mathcal{L}}$ is the neighbourhood frame functor $Q \circ Q^{op}$.
2. For the standard modal logic K (Example 7.1), M_K^{fin} is the finite powerset functor, while M_K is the filter functor [6].
3. For graded modal logic GML (Example 7.2), M_{GML}^{fin} is a modification of the finite multiset functor where elements of multisets may have infinite multiplicity.
4. For probabilistic modal logic PML (Example 7.3), M_{PML}^{fin} is a modification of the finite distribution functor where events A are assigned ‘probabilities’ PA which are downclosed subsets of the rational interval $[0, 1] \cap \mathbb{Q}$. Thus, the space of ‘probabilities’ essentially consists of the interval $[0, 1]$ and an additional copy of $[0, 1] \cap \mathbb{Q}$, where the second copy of $q \in [0, 1] \cap \mathbb{Q}$ is infinitesimally greater than the first. The distributions $P \in M_{\mathcal{L}}^{fin}(X)$ are required to obey the axiomatization of PML [21] w.r.t. the canonical semantics; it is presently unclear whether this requirement can be replaced by a simpler condition.

3 An Adjunction between Syntax and Semantics

We now set up an adjoint correspondence between rank-1 logics and set functors as their semantic counterparts. This establishes the canonical structure of a rank-1 logic as indeed canonical in a precise sense, i.e. as a universal model capturing all other ones. This situation is analogous (although not in any obvious sense technically related) to similar correspondences in equational logics and type theory: e.g. to a single-sorted equational theory, interpreted over cartesian categories (i.e. categories with finite products) with a distinguished object, one associates a Lawvere theory, which is again a cartesian category with a distinguished object and may simultaneously be regarded as an initial model and as a semantic representation of the given theory. The situation is dual for modal logics: the canonical structure serves as a *final* model of the given rank-1 logic, into which all other models may be mapped.

We make the categorical setting precise by collecting all rank 1 modal logics in a category **ModL** with morphisms $(A_1, A_1) \rightarrow (A_2, A_2)$ all maps $h : A_1 \rightarrow A_2$ such that the induced translation of formulas maps axioms in A_1 to derivable

formulas in (A_2, A_2) . The category of semantic structures is the category $\mathbf{Fn} = [\mathbf{Set}, \mathbf{Set}]$ of set functors and natural transformations. We have a functor $\mathbf{Th} : \mathbf{Fn}^{\text{op}} \rightarrow \mathbf{ModL}$ which takes a functor T to the logic (A_T, A_T) , where A_T is the set of all predicate liftings for T , and A_T is the set of all rank-1 clauses over A_T which are one-step sound for T . Given a natural transformation $\mu : T \rightarrow S$, $\mathbf{Th}(\mu) : \mathbf{Th}(S) \rightarrow \mathbf{Th}(T)$ is the morphism taking a predicate lifting $\lambda : Q \rightarrow Q \circ S^{\text{op}}$ for S to the predicate lifting $Q\mu \circ \lambda$ for T . Note that, in this terminology, an \mathcal{L} -structure is just a morphism of the form $h : \mathcal{L} \rightarrow \mathbf{Th}(T)$. In particular, the canonical \mathcal{L} -structure can be cast as a morphism $\eta_{\mathcal{L}} : \mathcal{L} \rightarrow \mathbf{Th}(M_{\mathcal{L}})$. The arrows $\eta_{\mathcal{L}}$ are part of the announced adjunction:

Theorem 17. *The canonical \mathcal{L} -structure $\eta_{\mathcal{L}}$ is universal; i.e. for each \mathcal{L} -structure $h : \mathcal{L} \rightarrow \mathbf{Th}(T)$, there exists a unique natural transformation $h^{\#} : T \rightarrow M_{\mathcal{L}}$ such that $\mathbf{Th}(h^{\#})\eta_{\mathcal{L}} = h$.*

In other words, the canonical structure is the final \mathcal{L} -structure, where a morphism of \mathcal{L} -structures is a natural transformation between the associated functors which is compatible with the interpretation of modal operators. A similar result holds for the canonical finitely branching \mathcal{L} -structure $M_{\mathcal{L}}^{\text{fin}}$, which now becomes a morphism $\eta_{\mathcal{L}}^{\text{fin}} : \mathcal{L} \rightarrow \mathbf{Th}(M_{\mathcal{L}}^{\text{fin}})$.

Theorem 18. *The \mathcal{L} -structure $\eta_{\mathcal{L}}^{\text{fin}}$ is universal among the finitely branching \mathcal{L} -structures; i.e. for each \mathcal{L} -structure $h : \mathcal{L} \rightarrow \mathbf{Th}(T)$ with T ω -accessible, there exists a unique natural transformation $h^{\#} : T \rightarrow M_{\mathcal{L}}^{\text{fin}}$ such that $\mathbf{Th}(h^{\#})\eta_{\mathcal{L}}^{\text{fin}} = h$.*

Theorems 17 and 18 allow us to replace rank-1 logics by functors in the definition of the coalgebraic semantics: an \mathcal{L} -structure may equivalently be regarded as a natural transformation $T \rightarrow M_{\mathcal{L}}$; analogously, an \mathcal{L} -structure over an ω -accessible functor T may be regarded as a natural transformation $T \rightarrow M_{\mathcal{L}}^{\text{fin}}$. We have

Proposition 19. *An \mathcal{L} -structure T is separating iff the associated natural transformation $T \rightarrow M_{\mathcal{L}}$ is injective.*

Thus, we have the following classification result.

Theorem 20. *Up to natural isomorphism, the ω -accessible \mathcal{L} -structures for which \mathcal{L} has the Hennessy-Milner property are precisely the subfunctors of the canonical finitely branching \mathcal{L} -structure $M_{\mathcal{L}}^{\text{fin}}$.*

4 Applications

A benefit of the coalgebraic semantics constructed above is that we can now apply results on coalgebraic modal logic to arbitrary rank-1 modal logics, even when the latter lack a formal semantics. This includes in particular the generic decidability and complexity results of [20, 21], of which we now obtain purely syntactic versions.

In [20], a generic finite model construction was given which yields criteria for decidability and upper complexity bounds for coalgebraic modal logics. The generic complexity bounds generally do not match known bounds in particular examples, typically *PSPACE*. This is remedied in [21], where a generic *PSPACE* decision procedure for coalgebraic modal logics based on a shallow model construction is given, at the price of stronger assumptions on the logic.

A crucial role in the algorithmic methods of [20] is played by the following localised version of the satisfiability problem:

Definition 21. The *one-step satisfiability* problem for a Λ -structure T is to decide, given a finite set X and a conjunctive clause ψ over $\text{Up}_\Lambda(\mathcal{P}(X))$, whether ψ is *one-step satisfiable*, i.e. $\llbracket \psi \rrbracket \subset TX$ is non-empty.

The satisfiability problem of a coalgebraic modal logic is

- decidable if its one-step satisfiability problem is decidable
- in *NEXPTIME* if one-step satisfiability is in *NP*
- in *EXPTIME* if one-step satisfiability is in *P*

(cf. [20]). This instantiates to the canonical structure as follows.

Lemma 22. *One-step satisfiability in $M_{\mathcal{L}}$ is one-step consistency in \mathcal{L} .*

Corollary 23. *The consistency problem of a rank-1 logic \mathcal{L} (i.e. deciding whether an \mathcal{L} -formula ϕ is consistent) is*

- decidable if one-step consistency (over finite sets) is decidable
- in *NEXPTIME* if one-step consistency is in *NP*
- in *EXPTIME* if one-step consistency is in *P*.

Corollary 24. *The consistency problem of $\mathcal{L} = (\Lambda, \mathbf{R})$ is decidable if Λ is finite and \mathbf{R} is recursive (i.e. it is decidable whether a one-step rule ϕ/ψ is contained in \mathbf{R} up to propositional equivalence of premises).*

The generic *PSPACE*-algorithm of [21] relies on a notion of *strictly one-step complete* rule set. Rather than repeating the definition here, we recall that strict one-step completeness follows from one-step completeness (i.e. $TX, \tau \models \psi$ implies $\vdash_{\mathcal{L}}^X \psi$ for all $\psi \in \text{Prop}(\text{Up}_\Lambda(V))$) in combination with *resolution closedness*. The latter refers to a notion of rule resolution where propositional resolvents of the conclusions of two rules are formed and the premises are combined by conjunction, with possible subsequent elimination of propositional variables; cf. [21] for a formal definition. As an example, consider the rules

$$(N) \frac{a}{\Box a} \quad (RR) \frac{a \wedge b \rightarrow c}{\Box a \wedge \Box b \rightarrow \Box c} \quad (RK_n) \frac{\bigwedge_{i=1}^n a_i \rightarrow b}{\bigwedge_{i=1}^n \Box a_i \rightarrow \Box b} \quad (n \geq 0).$$

The rule set $\{(N), (RR)\}$ presents the modal logic *K*, and its resolution closure consists of the rules (RK_n) . Cf. [21] for further examples.

In [21], a shallow model property is proved based on strictly one-step complete rule sets. The canonical semantics allows us to turn this into a shallow proof property:

Definition 25. A set Σ of formulas is called *closed* if it is closed under subformulas and negation, where $\neg\neg\phi$ is identified with ϕ . The smallest closed set containing a given formula ϕ is denoted $\Sigma(\phi)$. A subset H of Σ is called a Σ -*Hintikka set* if $\perp \notin H$ and, for $\phi \wedge \psi \in \Sigma$, $\phi \wedge \psi \in H$ iff $\phi, \psi \in H$, and, for $\neg\phi \in \Sigma$, $\neg\phi \in H$ iff $\phi \notin H$.

Theorem 26. Let $\mathcal{L} = (\Lambda, \mathbf{R})$, where \mathbf{R} is resolution closed. Then $\vdash_{\mathcal{L}} \phi$ iff for every $\Sigma(\phi)$ -Hintikka set H containing $\neg\phi$, there exist a clause $\bigvee_{i=1}^n \epsilon_i L_i \rho_i$ over $\Sigma(\phi)$ and a rule $\psi / \bigvee_{i=1}^n (\epsilon_i L_i a_i)$ in \mathbf{R}_C such that $\vdash_{\mathcal{L}} \psi[\rho_i/a_i]_{i=1, \dots, n}$ and $\epsilon_i L_i \rho_i \notin H$ for all i .

Theorem 26 implies that for \mathcal{L} as in the statement, every provable formula ϕ has a tableau proof of linear depth which mentions only propositional combinations of subformulas of ϕ , in particular mentions only the modal operators contained in ϕ .

Remark 27. We hope to generalize Theorem 26 to more general classes of logics (i.e. beyond rank 1), possibly using purely proof-theoretic methods. This would also imply wider applicability of the generic *PSPACE* algorithm discussed below.

Corollary 28. Let $\mathcal{L} = (\Lambda, \mathbf{R})$ be a rank-1 logic with \mathbf{R} resolution closed, and let $\Lambda_0 \subset \Lambda$. Let $\mathcal{L}_0 = (\Lambda_0, \mathbf{R}_0)$ where \mathbf{R}_0 consists of all \mathbf{R} -rules that only mention Λ_0 -operators. Then \mathcal{L} conservatively extends \mathcal{L}_0 , i.e. $\Phi \vdash_{\mathcal{L}} \phi$ implies that $\Phi \vdash_{\mathcal{L}_0} \phi$ for all $\phi \in \mathcal{F}(\mathcal{L}_0)$ and all $\Phi \subseteq \mathcal{F}(\mathcal{L}_0)$.

Applied to majority logic (Example 7), this immediately leads to a complete axiomatisation of the majority operator alone.

Example 29. In the presentation of [21], a resolution closed set of rules for majority logic (Example 7.2) was given, consisting of the rules

$$(M_m) \frac{\sum_{i=1}^n a_i + \sum_{r=1}^v c_r + m \leq \sum_{j=1}^k b_j + \sum_{s=1}^w d_s}{\bigwedge_{i=1}^n \diamond_{k_i} a_i \wedge \bigwedge_{r=1}^v W c_r \rightarrow \bigvee_{j=1}^k \diamond_{l_j} b_j \vee \bigvee_{s=1}^w W d_s} \quad (m \in \mathbb{Z})$$

with side conditions $\sum_{i=1}^n (k_i + 1) - \sum_{j=1}^k l_j + w - 1 - \max(m, 0) \geq 0$ and $v - w + 2m \geq 0$ (the sums in the premise refer to the — propositionally expressible — arithmetic of characteristic functions, cf. [21]). By Corollary 28, the rules

$$(W_m) \frac{\sum_{r=1}^v c_r + m \leq \sum_{s=1}^w d_s}{\bigwedge_{r=1}^v W c_r \rightarrow \bigvee_{s=1}^w W d_s} \quad (m \in \mathbb{Z})$$

with side conditions $w - 1 - \max(m, 0) \geq 0$ and $v - w + 2m \geq 0$ form a complete axiomatisation of the majority operator W .

Theorem 26 suggests an obvious recursive algorithm for checking provability (or, dually, consistency). In order to ensure that this algorithm is feasible, we need to make sure that we never need to prove ‘small’ clauses by instantiating propositional variables with identical formulas in ‘large’ rules. We thus further

require *reduction closedness* of the rule set, in the sense that every rule instance where a literal is duplicated in the conclusion can be replaced by an instance of another rule where all literals are distinct; cf. [21] for a formal definition.

The main result of [21] states that the satisfiability problem of a \mathcal{L} -structure T is in *PSPACE* if T has a strictly one-step complete reduction closed rule set which is *PSPACE-tractable*, which essentially means that the rules applied according to Theorem 26 have representations whose size is polynomial in the matched clause and from which the clauses of the premise are easily extracted (again, cf. [21] for a formal definition). Applying this result to the canonical \mathcal{L} -structure, we obtain a purely syntactic criterion for a rank-1 logic to be in *PSPACE*:

Theorem 30. *The consistency (provability) problem of $\mathcal{L} = (\mathcal{A}, \mathcal{R})$ is in PSPACE if \mathcal{R} has a resolution closure which has a PSPACE-tractable reduction closure.*

5 Example: Deontic Logic

A typical application area for the above results are modal logics that come from a philosophical background, such as epistemic and deontic logics, which are often defined either without any reference to semantics at all or with a neighbourhood semantics essentially equivalent to the canonical semantics described above. Deontic logics [8], which have received much recent interest in computer science as logics for obligations of agents, are moreover often axiomatised in rank 1.

Standard deontic logic [3] is just the modal logic *KD*. This has been criticized on the grounds that it entails the *deontic explosion*: if O is the modal obligation operator ‘it ought to be the case that’, the *K*-axiom $(Oa \wedge Ob) \leftrightarrow O(a \wedge b)$ implies that in the presence of a single deontic dilemma, everything is obligatory, i.e. $Oa \wedge O\neg a \rightarrow Ob$. Some approaches to this problem are summarized in [5], where the novel solution is advocated to restrict at least one direction of *K* to the case that $a \wedge b$ is *permitted*, i.e. to $P(a \wedge b)$, where P is the dual $\neg O \neg$ of O . This leads to the axioms

$$\begin{aligned} \text{(PM)} \quad & O(a \wedge b) \wedge P(a \wedge b) \rightarrow Oa \\ \text{(PAND)} \quad & Oa \wedge Ob \wedge P(a \wedge b) \rightarrow O(a \wedge b) \end{aligned}$$

(in [5], (PM) is formulated as a rule (RPM)). Two systems are proposed (both including the congruence rule): given the further axioms (N) $O\top$, (P) $\neg O\perp$, and

$$\text{(ADD)} \quad (Oa \wedge Ob) \rightarrow O(a \wedge b),$$

DPM.1 is determined by (PM), (N), and (ADD), while DPM.2 is given by (PM), (PAND), (N), and (P). A further system PA, consisting of (PAND), (P), (N), and the standard monotonicity axiom is rejected, as it still leads to a form of deontic explosion where everything permitted is obligatory in the presence of a dilemma.

It is shown in [5] that DPM.1 and DPM.2 are sound and *weakly* complete w.r.t. the obvious classes of neighbourhood frames, and that both logics are decidable; the proofs are rather involved. In our framework, the situation presents itself as follows. The neighbourhood semantics of [5] is easily seen to be precisely the canonical semantics; the new insight here is that the semantics is coalgebraic. The rest is for free: by Theorem 14, both DPM.1 and DPM.2 are even *strongly* complete (the reason that the strong completeness proof fails in [5] is that an explicit construction of a canonical model is attempted). Decidability is immediate by Corollary 24; the finite model property (proved in [5] using filtrations) follows from the results of [20]. Moreover, the resolution closures of DPM.1 and DPM.2 enjoy the pleasant proof theoretic properties listed in Theorem 26. (The same holds for PA, and in fact for rather arbitrary variations of the axiom system.) A challenge that remains is to establish that DPM.1 and DPM.2 are in *PSPACE* by the method described at the end of the previous section, the main problem being to harness closure under reduction.

6 Conclusion

We have established that every modal logic \mathcal{L} of rank 1 has a canonical coalgebraic semantics for which \mathcal{L} is sound and strongly complete. Moreover, \mathcal{L} has a canonical finitely branching coalgebraic semantics for which \mathcal{L} is sound and weakly complete and has the Hennessy-Milner property, and from which all finitely branching semantics for which \mathcal{L} has the Hennessy-Milner property are obtained as substructures. This is a converse to the previous insight that every coalgebraic modal logic can be axiomatized in rank 1 [20]. It allows us to formulate purely syntactic versions of semantics-based generic decidability and complexity criteria for coalgebraic modal logic [20, 21], including e.g. the result that *every recursively axiomatised rank-1 logic with finitely many modal operators is decidable*. We have applied this framework to recently defined versions of deontic logic which accommodate deontic dilemmas [5]. In particular, we have obtained decidability and strong completeness for these logics as immediate consequences of our generic results, while the original work has rather involved proofs and moreover establishes only decidability and weak completeness. Application of the generic *PSPACE* upper bound [21] to these logics remains an open problem.

We emphasise that the restriction to rank 1 is not an inherent limitation of the coalgebraic approach — the fact that coalgebraic modal logics are of rank 1 is due to the interpretation of these logics over the whole class of coalgebras for the relevant functor (in analogy to the standard modal logic K), and logics outside rank 1 may be modelled by passing to suitable subclasses of coalgebras. Ongoing work is aimed at pushing the generic results beyond strict rank 1; preliminary results have been obtained for axioms that combine rank 1 with rank 0, i.e. a coalgebraic counterpart of KT . A further point of interest is to obtain completeness and decidability results for coalgebraic modal logics with iteration, i.e. the coalgebraic counterpart of CTL.

References

- [1] J. Adámek and H.-E. Porst. On tree coalgebras and coalgebra presentations. *Theoret. Comput. Sci.*, 311:257–283, 2004.
- [2] M. Barr. Terminal coalgebras in well-founded set theory. *Theoret. Comput. Sci.*, 114:299–315, 1993.
- [3] B. Chellas. *Modal Logic*. Cambridge, 1980.
- [4] K. Fine. In so many possible worlds. *Notre Dame J. Formal Logic*, 13:516–520, 1972.
- [5] L. Goble. A proposal for dealing with deontic dilemmas. In *Deontic Logic in Computer Science, DEON 2004*, volume 3065 of *LNAI*, pages 74–113. Springer, 2004.
- [6] H.-P. Gumm. Functors for coalgebras. *Algebra Universalis*, 45:135–147, 2001.
- [7] A. Heifetz and P. Mongin. Probabilistic logic for type spaces. *Games and Economic Behavior*, 35:31–53, 2001.
- [8] R. Hilpinen. Deontic logic. In L. Goble, editor, *The Blackwell Guide to Philosophical Logic*. Blackwell, 2001.
- [9] B. Jacobs. Towards a duality result in the modal logic of coalgebras. In *Coalgebraic Methods in Computer Science*, volume 33 of *ENTCS*. Elsevier, 2000.
- [10] C. Kupke, A. Kurz, and D. Pattinson. Ultrafilter extensions for coalgebras. In *Algebra and Coalgebra in Computer Science, CALCO 2005*, volume 3629 of *LNCS*, pages 263–277. Springer, 2005.
- [11] A. Kurz. Specifying coalgebras with modal logic. *Theoret. Comput. Sci.*, 260:119–138, 2001.
- [12] K. Larsen and A. Skou. Bisimulation through probabilistic testing. *Inform. Comput.*, 94:1–28, 1991.
- [13] E. Pacuit and S. Salame. Majority logic. In *Principles of Knowledge Representation and Reasoning, KR 04*, pages 598–605. AAAI Press, 2004.
- [14] D. Pattinson. Semantical principles in the modal logic of coalgebras. In *Theoretical Aspects of Computer Science, STACS 01*, volume 2010 of *LNCS*, pages 514–526. Springer, 2001.
- [15] D. Pattinson. Expressive logics for coalgebras via terminal sequence induction. *Notre Dame J. Formal Logic*, 45:19–33, 2004.
- [16] M. Pauly. A modal logic for coalitional power in games. *J. Logic Comput.*, 12:149–166, 2002.
- [17] M. Rößiger. Coalgebras and modal logic. In *Coalgebraic Methods in Computer Science*, volume 33 of *ENTCS*. Elsevier, 2000.
- [18] J. Rutten. Universal coalgebra: A theory of systems. *Theoret. Comput. Sci.*, 249:3–80, 2000.
- [19] L. Schröder. Expressivity of coalgebraic modal logic: the limits and beyond. In *Foundations of Software Science And Computation Structures*, volume 3441 of *LNCS*, pages 440–454. Springer, 2005. Extended version to appear in *Theoret. Comput. Sci.*
- [20] L. Schröder. A finite model construction for coalgebraic modal logic. In *Foundations Of Software Science And Computation Structures*, volume 3921 of *LNCS*, pages 157–171. Springer, 2006. Extended version to appear in *J. Logic Algebraic Programming*.
- [21] L. Schröder and D. Pattinson. PSPACE reasoning for rank-1 modal logics. In *Logic in Computer Science*, pages 231–240. IEEE, 2006. Presentation slides available under www.informatik.uni-bremen.de/~lschrode/slides/rank1pspace.pdf.