On Modal Logics of Linear Inequalities

Clemens Kupke Dirk Pattinson

Department of Computing
Imperial College London

Abstract

We consider probabilistic modal logic, graded modal logic and stochastic modal logic, where linear inequalities may be used to express numerical constraints between quantities. For each of the logics, we construct a cut-free sequent calculus and show soundness with respect to a natural class of models. The completeness of the associated sequent calculi is then established with the help of coalgebraic semantics which gives completeness over a (typically much smaller) class of models. With respect to either semantics, it follows that the satisfiability problem of each of these logics is decidable in polynomial space.

Keywords: Probabilistic modal logic, graded modal logic, linear inequalities

1 Introduction

In this paper, we consider three different, but closely related, modal logics. The first logic that we consider is probabilistic modal logic, where we think of every formula A as denoting an event $[\![A]\!]$ in a probability space. The variant of probabilistic modal logic that we consider here allows explicit comparisons of the likelihoods of individual formulas by means of linear inequalities. If $A_1, \ldots, A_m$ are formulas and $p, c_1, \ldots, c_m \in \mathbb{Q}$, then the expression $\sum_{j=1}^{m} c_j \cdot \mu([\![A_j]\!]) \geq p$ can be denoted a formula which is satisfied at a point x of a probability space if the local probability measure μ associated with point x satisfies the above inequality. An expression of this form is written as an m-ary modal operator, $L_p(c_1, \ldots, c_m)$, applied to $A_1, \ldots, A_m$.

The second logic that we consider is graded modal logic, where we may again use linear inequalities to express constraints on successors with certain properties. As before, we use m-ary modal operators of the form $L_p(c_1, \ldots, c_m)$ to express that the inequality $\sum_{j=1}^{m} c_j \cdot \sharp A_j \geq p$ holds at a particular point in a Kripke model, where $\sharp A_j$ is the number of successors of that state satisfying property A_j.

Finally, we consider a third logic, stochastic modal logic, that is a hybrid between the two. To get from probabilistic modal logic to stochastic modal logic, one needs to generalise from probability measures to arbitrary measures. To get to stochastic logic from graded modal logic, one gives up the idea of always having an integer number of successors and replaces the transition relation by a family of local real-valued measures that determine the total weight of a successor set.

For each of the logics mentioned above, we give an axiomatisation in terms of a cut-free sequent calculus and prove soundness with respect to a natural class of models: Markov models, Kripke models and what we call measurable models – the natural generalization of Markov models, where one drops the requirement of dealing with probability measures. We then establish completeness of the sequent calculus with respect to coalgebraic models. For each of the logics, we isolate a natural, coalgebraic semantics and show how the general results of coalgebraic modal logics can be used to give a rather simple completeness proof. In a third step, we relate both types of semantics, and show that the coalgebraic semantics embeds into the 'natural' semantics considered initially. Our treatment thus combines the best of both worlds for each of the logics: we establish soundness for a large class of models, whereas the logics are proved complete for a much smaller class. The complexity of each of the logics then follows by analysing the complexity of backwards proof search in the given sequent calculus.

The main contributions of this paper are the cut-free axiomatisation of three different modal logics and the completeness proof of this axiomatisation using coalgebraic methods. The sequent calculi appear to be new in each case. While the soundness proofs are certainly standard, completeness relies on coalgebraic techniques. Rather than exhibiting a fully fledged (canonical) model construction, we can make do with showing that the rules that generate the sequent systems are *one-step complete*: we interpret all logics over T-coalgebras $(X, \gamma : X \to TX)$ for suitably chosen T, where γ is the transition function. One-step completeness now stipulates that all sequents valid over the set of 'successors' TX should be derivable via modal rules whose premises are already valid over X, where X is an arbitrary set. For probabilistic and stochastic modal logic, the question of one-step completeness can be translated into a linear programming problem over the rational domain, which fails for the case of graded modal logic, where we use maximal consistent sets, but only at the level of one-step successors.

The coalgebraisation of all three logics moreover allows us to apply a number of generic (coalgebraic) results: with the help of [2] we obtain completeness and Exptime decidability of an extension of each logic with least/greatest fixpoints, [7] allows us to construct generic (tableau) algorithms for the global consequence problem, and [17,8] provides an Exptime complexity bound and optimal tableau algorithm, respectively, of hybrid extensions over arbitrary sets of global assumptions. As such, the paper does not present any new results concerning the coalgebraic interpretation of modal logics. Rather, we show how coalgebraic methods can be used to obtain results about existing modal logics.

Related Work. Probabilistic modal logic, as studied in this paper, can be seen as an extension of the probabilistic modal logic presented in [9] with linear inequalities and is a notational variant of the probabilistic logic considered in [5,4], where a complete axiomatisation in a Hilbert-style proof system and a proof of Pspace-decidability is presented. Our contribution here is a cut-free sequent system that allows for purely syntax driven implementations of satisfiability checking that are amenable to standard optimisations [10,20].

The extension of graded modal logic extended with linear inequalities considered here is a fragment of Presburger modal logic with regularity constraints [3] but subsumes Majority logic [11] and the both (standard form of) graded modal logic [6] and description logics with qualified number restrictions [1]. In absence of linear inequalities, this logic is known to be PSPACE complete [18] and PSPACE-completeness in presence of linear inequalities was shown in [3], but no complete axiomatisation appears to be known so far, which is provided here.

For stochastic modal logic, we are not aware of any results concerning completeness and complexity.

2 Preliminaries and Notation

2.1 Preliminaries on Sequent Calculi

Throughout the paper, we fix a set V of propositional variables. As we will be dealing with three different modal logics, it is convenient to isolate their syntactical differences into a *modal similarity type*, i.e. a set of modal operators with associated arities. Given a modal similarity type Λ, the set $\mathcal{F}(\Lambda)$ of Λ-formulas is given by the grammar

$$\mathcal{F}(\Lambda) \ni A, B ::= p \mid \neg A \mid A \wedge B \mid \heartsuit(A_1, \dots, A_n)$$

where $p \in \mathsf{V}$ and $\heartsuit \in \Lambda$ is n-ary. If $F \subseteq \mathcal{F}(\Lambda)$ is a set of formulas, then we write

$$\Lambda(F) = \{\heartsuit(A_1, \dots, A_n) \mid \heartsuit \in \Lambda \ n\text{-ary}, A_1, \dots, A_n \in F\}$$

for the set of formulas consisting of modalities applied to elements of F. If $\sigma : \mathsf{V} \to \mathcal{F}(\Lambda)$ is a substitution, then $A\sigma$ denotes the result of replacing all occurrences of $p \in \mathsf{V}$ in A by $\sigma(p)$.

A *sequent* is a finite multiset (so that contraction is made explicit) of formulas that we read disjunctively. We identify $A \in \mathcal{F}(\Lambda)$ with the sequent $\{A\}$ and write Γ, Δ for the (multiset) union of Γ and Δ. If $F \subseteq \mathcal{F}(\Lambda)$ is a set of formulas, we write $\mathcal{S}(F)$ for the set of those sequents that only contain elements of F, possibly negated. Substitution applies pointwise to sequents, respecting multiplicity so that $\Gamma\sigma = \{A\sigma \mid A \in \Gamma\}$. The three logics we consider in this paper can be axiomatised by *one-step rules*, that is, rules of the form

$$\frac{\Gamma_1 \quad \dots \quad \Gamma_n}{\Gamma_0}$$

where $\Gamma_1, \dots, \Gamma_n \in \mathcal{S}(\mathsf{V})$ and $\Gamma_0 \in \mathcal{S}(\Lambda(\mathsf{V}))$. If $\mathcal{R}$ is a set of one-step rules, we write $\mathcal{R} \vdash \Gamma$ if Γ is an element of the least set of sequents that is closed under the propositional rules and all substitution instances of one-step rules, that is under the rules

$$\frac{}{p, \neg p, \Delta} \qquad \frac{A, \Gamma}{\neg\neg A, \Gamma} \qquad \frac{\neg A, \neg B, \Delta}{\neg(A \wedge B), \Delta} \qquad \frac{A, \Delta \qquad B, \Delta}{A \wedge B, \Delta} \qquad \frac{\Gamma_1\sigma \quad \dots \quad \Gamma_n\sigma}{\Gamma_0\sigma, \Delta}$$

where $p \in \mathsf{V}$, $\sigma : \mathsf{V} \to \mathcal{F}(\Lambda)$ and $\Delta \in \mathcal{S}(\mathcal{F}(\Lambda))$ is a weakening context. It is easy to see that the propositional part of this calculus can be embedded into the system **GS3p** of [19] which is known to be sound and complete. The concrete syntactical presentation of the modal rules for the logics considered here is most conveniently expressed using the following notation. If $c_1, \ldots, c_m, k \in \mathbb{Q}$ are rational numbers, $a_1, \ldots, a_m \in \mathsf{V}$ are propositional variables and $\Gamma = \{a_i \mid i \in I\} \cup \{\neg a_j \mid j \notin I\}$ then

$$\Gamma \in \sum_{j=1}^{m} c_j a_j \geq k \iff \sum_{i \in I} c_i \geq k$$

so that we may use $\sum_{i=1}^{m} c_i a_i \geq k$ to denote a set of sequents that we think of the set of premises of a proof rule. We write $\mathsf{sign}(q)$ for the sign of a rational number q and, if A is a formula and $r \neq 0$, we put $\mathsf{sg}(r)A = A$ if $r > 0$ and $\mathsf{sg}(r)A = \neg A$, otherwise.

2.2 Coalgebraic Preliminaries

If $T : \mathsf{Set} \to \mathsf{Set}$ is an endofunctor, a *T-coalgebra* is a pair (X, γ) where X is a (carrier) set and $\gamma : X \to TX$ is a (transition) function. We think of T-coalgebras as playing the role of frames, and take a *T-model* to be a T-coalgebra equipped with a valuation, i.e. a triple (X, γ, π) where (X, γ) is a T-coalgebra and $\pi : \mathsf{V} \to \mathcal{P}(X)$ is a valuation of the propositional variables.

Given a similarity type Λ, we can interpret Λ-formulas over T-models provided T extends to a Λ-structure, i.e. T comes equipped with a predicate lifting (a set-indexed family of maps)

$$(\llbracket \heartsuit \rrbracket_X : \mathcal{P}(X)^n \to \mathcal{P}(TX))_{X \in \mathsf{Set}}$$

for every n-ary $\heartsuit \in \Lambda$ that satisfies the *naturality requirement*

$$(Tf)^{-1} \circ \llbracket \heartsuit \rrbracket_Y(S_1, \ldots, S_n) = \llbracket \heartsuit \rrbracket_X(f^{-1}(S_1), \ldots, f^{-1}(S_n))$$

for all $f : X \to Y$ and all $S_1, \ldots S_n \subseteq Y$. If $M = (X, \gamma, \pi)$ is a T-model, the semantics of modal formulas is now defined as expected for propositional connectives

$$\llbracket p \rrbracket_M = \pi(p) \quad \llbracket \neg A \rrbracket_M = X \setminus \llbracket A \rrbracket_M \quad \llbracket A \wedge B \rrbracket_M = \llbracket A \rrbracket_M \cap \llbracket B \rrbracket_M$$

together with the clause

$$\llbracket \heartsuit(A_1, \ldots, A_n) \rrbracket_M = \gamma^{-1} \circ \llbracket \heartsuit \rrbracket_X(\llbracket A_1 \rrbracket_M, \ldots, \llbracket A_n \rrbracket_M)$$

for the modal operators. We write $M, x \models A$ in case $x \in \llbracket A \rrbracket_M$ and $M \models A$ if $M, x \models A$ for all $x \in X$. Finally, we write $T \models \Gamma$ if $M \models \Gamma$ for all T-models M. The glue between the axiomatisation (in terms of one-step rules) and the modal semantics is provided by the following notions:

Definition 2.1 *Suppose that Λ is a modal similarity type, T a Λ-structure and $\mathcal{R}$ a set of one-step rules over Λ. We introduce the following notions in case X is a set and $\tau : \mathsf{V} \to \mathcal{P}(X)$ is a valuation:*

*(i) If $\Gamma \in \mathcal{S}(\mathsf{V})$ is a propositional sequent, we write $\llbracket \Gamma \rrbracket_{(X,\tau)} = \bigcup\{\llbracket A \rrbracket_{(X,\tau)} \mid A \in \Gamma\}$ (where $\llbracket p \rrbracket_{(X,\tau)} = \tau(p)$) for the interpretation of a sequent $\Gamma \in \mathcal{S}(\mathsf{V})$ with respect to τ and $(X, \tau) \models \Gamma$ in case Γ is τ-*valid*, i.e. $\llbracket \Gamma \rrbracket_{(X,\tau)} = X$.*

*(ii) Similarly, if $\Gamma \in \mathcal{S}(\Lambda(V))$ we write $[\![\Gamma]\!]_{(TX,\tau)} = \bigcup\{[\![A]\!] \mid A \in \Gamma\}$ (where $[\![\heartsuit(a_1, \ldots, a_n)]\!]_{(TX,\tau)} = [\![\heartsuit]\!]_X(\tau(a_1), \ldots, \tau(a_n)))$ for the interpretation of a modalised sequent with respect to τ, and $(TX, \tau) \models \Gamma$ in case Γ is τ-*valid*, i.e. $[\![\Gamma]\!]_{(TX,\tau)} = TX$.*

*(iii) Finally, a sequent $\Gamma \in \mathcal{S}(\Lambda(\mathsf{V}))$ is τ-*derivable* (with respect to $\mathcal{R}$) if there exists $\Gamma_1 \ldots \Gamma_n/\Gamma_0 \in \mathcal{R}$ and $\sigma : \mathsf{V} \to \mathsf{V}$ such that all $\Gamma_i\sigma$ are τ-valid for $1 \leq i \leq n$ and $\Gamma_0\sigma \subseteq \Gamma$.*

We can now justify the sum notation introduced earlier:

Lemma 2.2 *Suppose that $\tau : \mathsf{V} \to \mathcal{P}(X)$ is a valuation. Then*

$$\forall x \in X \left(\sum_{j=1}^{m} c_j \mathbb{1}_{\tau(a_j)}(x) \geq k \right) \iff \forall\, \Gamma \in \sum_{j=1}^{m} c_j a_j \geq k \left((X, \tau) \models \Gamma \right).$$

Moreover, we can relate one-step rules and coalgebraic semantics as follows:

Definition 2.3 *Suppose Λ is a modal similarity type and $T : \mathsf{Set} \to \mathsf{Set}$ is a Λ-structure. We say that a set $\mathcal{R}$ of one-step rules is* one-step sound *(resp.* one-step cut-free complete*) if, for all valuations $\tau : \mathsf{V} \to \mathcal{P}(X)$ and all $\Gamma \in \mathcal{S}(\Lambda(\mathsf{V}))$: Γ is τ-derivable if (resp. only if) Γ is τ-valid.*

We note that the notions of one-step soundness and one-step (cut-free) completeness do not quantify over models: both conditions can be checked locally. Importantly, these notions give rise to soundness and cut-free completeness in the standard way.

Theorem 2.4 *Suppose $\mathcal{R}$ is a set of one-step rules over a modal similarity type Λ and let T be a Λ-structure. If $\Gamma \in \mathcal{S}(\Gamma)$ and*

(i) $\mathcal{R}$ is one-step sound, then $\models \Gamma$ whenever $\mathcal{R} \vdash \Gamma$

(ii) $\mathcal{R}$ is one-step cut-free complete, then $\mathcal{R} \vdash \Gamma$ whenever $T \models \Gamma$.

The proof of the last theorem can be found in [13] but it should be remarked that this type of coherence condition between syntax and semantics is well studied: [12,15] use similar (weaker) coherence conditions to obtain soundness and completeness of a Hilbert system and [16] uses *strict completeness* to obtain what essentially amounts to a cut-free sequent system.

3 Probabilistic Modal Logic

We start our investigation into modal logics of linear inequalities by considering probabilistic modal logic where we may allow ourselves linear inequalities to specify the relationships between individual formulas. That is, we consider the modal similarity type

$$\Lambda = \{L_p(c_1, \ldots, c_m) \mid m \in \mathbb{N}, p, c_1, \ldots, c_m \in \mathbb{Q}\}$$

where the arity of $L_p(c_1, \ldots, c_m)$ is m. We interpret probabilistic modal logic over state spaces X where every point $x \in X$ induces a probability distribution μ over successor states. Informally, validity of $L_p(c_1, \ldots, c_m)(A_1, \ldots, A_m)$ at point $x \in$

X means that the linear inequality $\sum_{j=1}^{m} c_j \mu(A_j) \geq p$ holds, where $\mu(A_j)$ is the measure of the truth-set of the formula A_j, seen from point x. In particular this allows us to compare the probabilities of events:

Example 3.1 According to a recent experience of the second author with a well-known budget airline, we may consider a state space comprising all European airports, and we may think of the probability distribution associated with a particular city as giving us the probability of landing at a particular airport when boarding any flight of this carrier. In this logic, which we refrain from calling EasyLogic, we can for instance express that landing in England is 5 times as likely as landing in Scotland as $L_0(1, -5)(\mathsf{England}, \mathsf{Scotland})$ (which is reasonable to assume for carriers that are based in England). The second author however doubts that the business model of said budget airline can be axiomatised in any logic.

We axiomatise probabilistic modal logic with linear inequalities and prove soundness and completeness with respect to two different classes of models. The (complete, cut-free) axiomatisation is induced by the set $\mathcal{R}$ of one-step rules that comprises all instances of

$$(P)\frac{\sum_{i=1}^{n} r_i \left(\sum_{j=1}^{m_i} c_j^i \cdot a_j^i\right) \geq k}{\{\mathsf{sg}(r_i) L_{p_i}(c_1^i, \dots, c_{m_i}^i)(a_1^i, \dots, a_{m_i}^i) \mid i = 1, \dots, n\}}$$

where $r_1, \dots, r_n \in \mathbb{Z} \setminus \{0\}$ and $k \in \mathbb{Z}$ that satisfy the side condition

$$\sum_{i=1}^{n} r_i p_i < k \quad \text{if all } r_i < 0, \text{ and} \quad \sum_{i=1}^{n} r_i p_i \leq k \quad \text{otherwise.}$$

We first treat soundness of probabilistic modal logic, interpreted over Markov chain models before showing completeness over a class of coalgebraic models that corresponds to finitely supported Markov chains.

3.1 Markov Chain Semantics and Soundness

The first semantics of probabilistic modal logic is given with respect to Markov models.

Definition 3.2 *A* Markov model *is a triple (X, μ, π) where X is a measurable space with σ-algebra Σ_X, $\pi : \mathsf{V} \to \Sigma_X$ is a valuation and $\mu : X \times \Sigma \to [0, 1]$ is a Markov kernel, that is, $\mu(x, \cdot) : \Sigma \to [0, 1]$ is a probability measure for all $x \in X$ and $\mu(\cdot, S) : X \to [0, 1]$ is measurable for all $S \in \Sigma_X$.*

If $M = (X, \mu, \pi)$ is a Markov model, then the semantics $[\![A]\!]_M \in \Sigma_X$ is given as expected for the propositional connectives (where atomic propositions are mapped to measurable sets) and the clause for modal operators is

$$[\![L_p(c_1, \dots, c_m)(A_1, \dots, A_m)]\!]_M = \{x \in X \mid \sum_{j=1}^{m} c_j \mu(x, [\![A_j]\!]_M) \geq p]\!]$$

where we write $M, x \models A$ if $x \in [\![A]\!]_M$ and $M \models A$ if $M, x \models A$ for all $x \in X$. Finally $\mathsf{Mark} \models \Gamma$ if $M \models \bigvee \Gamma$ for all Markov models M. Note that the measurability conditions guarantee that the truth-set $[\![A]\!]_M$ of a formula is always measurable. We now show soundness of probabilistic modal logic with respect to Markov models.

Proposition 3.3 $\mathsf{Mark} \models \Gamma$ *whenever* $\mathcal{R} \vdash \Gamma$.

Proof. Suppose that $M = (X, \mu, \pi)$ is a Markov model and $\mathcal{R} \vdash \Gamma$. We show that $M \models \Gamma$ by induction on the proof of $\mathcal{R} \vdash \Gamma$ where the application of an instance of (P) is the only interesting case.

Consider the sequent Γ appearing as the conclusion of the rule

$$\frac{\sum_{i=1}^{n} r_i \sum_{j=1}^{m_i} c_i^j A_i^j \geq k}{\{\mathsf{sg}(r_i) L_{p_i}(c_i^1, \ldots, c_i^{m_i})(A_1, \ldots, A_m) \mid i = 1, \ldots, n\}}$$

the applicability of which is ensured by the side condition

$$\sum_{i=1}^{n} r_i p_i < k \quad \text{if all } r_i < 0, \text{ and} \quad \sum_{i=1}^{n} r_i p_i \leq k \quad \text{otherwise.}$$

By induction hypothesis,

$$\sum_{i=1}^{n} r_i \sum_{j=1}^{m_i} c_i^j \mathbb{1}_{[\![A_i^j]\!]_M}(x) \geq k$$

for all $x \in X$. Now suppose for a contradiction that there exists an $x \in X$ so that $M, x \not\models \Gamma$. If $\mu = \mu(x, \cdot)$ then this implies that

$$\sum_{i=1}^{n} r_i \sum_{j=1}^{m_i} c_i^j \mu([\![A_i^j]\!]_M) \geq k$$

by integrating both sides with respect to μ, and

$$\sum_{j=1}^{m_i} c_i^j \mu[\![A_i^j]\!]_M \geq p_i \quad (\text{if } r_i < 0), \text{ and} \quad \sum_{j=1}^{m_i} c_i^j \mu[\![A_i^j]\!]_M < p_i \quad (\text{if } r_i > 0)$$

for all $i = 1, \ldots, n$. In summary, this implies that

$$k \leq \sum_{i=1}^{n} r_i \sum_{j=1}^{m_i} c_i^j \mu([\![A_i^j]\!]_M) \leq \sum_{i=1}^{n} r_i p_i \leq k$$

where either the last or the penultimate inequality are strict so that $k < k$ in both cases, contradicting $M, x \not\models \Gamma$ and therefore proving the claim. $\square$

We next establish completeness over a smaller class of models, that is, Markov chains where the transition measures are finitely supported. Crucially, these fit into the framework of coalgebraic semantics:

3.2 Coalgebraic Semantics and Completeness

We write $\mathrm{supp}(f) = \{x \in X \mid f(x) \neq 0\}$ for the support of a function $f : X \to \mathbb{R}$ and consider the functor $\mathcal{D} : \mathsf{Set} \to \mathsf{Set}$ where

$$\mathcal{D}(X) = \{\mu : X \to [0,1] \mid \mathrm{supp}(\mu) \text{ finite}, \sum_{x \in X} \mu(x) = 1\}$$

that extends to a Λ-structure by stipulating that

$$[\![L_p(c_1, \dots, c_n)]\!]_X(S_1, \dots, S_n) = \{\mu \in \mathcal{D}(X) \mid \sum_{i=1}^{n} c_i \cdot \mu(S_i) \geq p\}$$

for $S_1, \dots, S_n \subseteq X$ where $\mu(S) = \sum_{x \in X} \mu(x)$. As spelled out in Section 2.2 this induces an interpretation $[\![A]\!]_M \subseteq X$ of Λ-formulas over $\mathcal{D}$-models $M = (X, \gamma, \pi)$. We now show that the set $\mathcal{R}$ of one-step rules consisting of all instances of (P) is indeed one-step complete, which is the content of the next lemma.

Lemma 3.4 *Consider a valuation $\tau : \mathsf{V} \to \mathcal{P}(X)$ and suppose that $\Gamma \in \mathcal{S}(\Lambda(\mathsf{V}))$ is τ-valid. Then Γ is τ-derivable.*

Proof. Suppose that $\Gamma = \{\mathsf{sg}(\epsilon_i) L_{p_i}(c_i^1, \dots, c_i^{m_i})(a_i^1, \dots, a_i^{m_i}) \mid i = 1, \dots, n\}$ where $\epsilon_1, \dots, \epsilon_n \in \{-1, 1\}$, the p_i and $c_i^j \in \mathbb{Q}$ and the $a_i^j \in \mathsf{V}$. Furthermore let $\tau : \mathsf{V} \to \mathcal{P}(X)$ be a valuation such that Γ is τ-valid. To see that Γ is τ-derivable, we show that there exist $k, r_1, \dots, r_n \in \mathbb{Z}$ so that

(i) $\sum_{i=1}^{n} r_i^2 > 0$ (i.e. at least one of the $r_1, \dots, r_n$ is non-zero)

(ii) $\mathsf{sign}(r_i) = \mathsf{sign}(\epsilon_i)$ for all $i = 1, \dots, n$ with $r_i \neq 0$

(iii) $\sum_{i=1}^{n} r_i \left(\sum_{j=1}^{m_i} c_j^i \cdot \mathbb{1}_{\tau(a_j^i)}(x) \right) \geq k$ for all $x \in X$

(iv) $\sum_{i=1}^{n} r_i p_i \leq k$ if at least one ϵ_i is positive, and $\sum_{i=1}^{n} r_i p_i < k$ otherwise.

We define an equivalence relation $\sim$ on X by

$$x \sim y \iff \left(x \in \tau(a_i^j) \iff y \in \tau(a_i^j) \right)$$

for all $i = 1, \dots, n$ and all $j = 1, \dots, m_i$. Assume that $x_1, \dots, x_k \in X$ are the (finitely many) representatives of the equivalence classes of X under $\sim$. Consider the matrices

$$A_0 = \begin{pmatrix} -\epsilon_1 & & 0 & 0 \\ & \ddots & & \vdots \\ 0 & & -\epsilon_n & 0 \\ -f_1(x_1) & \dots & -f_n(x_1) & 1 \\ \vdots & & \vdots & \\ -f_1(x_k) & \dots & -f_n(x_k) & 1 \end{pmatrix}$$

$$A_1 = (\quad p_1 \quad \dots \quad p_n \quad -1)$$

where $f_i = \sum_{j=1}^{m_i} c_i^j \cdot \mathbb{1}_{\tau(a_i^j)}$ and let $A = \begin{pmatrix} A_0 \\ A_1 \end{pmatrix}$. We note the following properties, where $y = (y_1, \dots, y_n, \hat{y}_1, \dots, \hat{y}_k, y_0) \in \mathbb{Q}_{\geq 0}^{n+k+1}$:

(i) if $b = (b_1, \dots, b_n, 0, \dots, 0)$ with $\sum_{i=1}^n b_i^2 > 0$, $y^T A = 0$ and $y^T b < 0$ then $y_0 > 0$.

(ii) if $y_0 = 1$ and $y^T A = 0$, then the assignment $\mu_y(x_i) = \hat{y}_i$ and $\mu_y(x) = 0$ if $x \notin \{x_1, \dots, x_k\}$ is a finitely supported probability distribution with

$$\sum_{j=1}^{m_i} c_i^j \cdot \mu_y(\tau(a_i^j)) = p_i - \epsilon_i y_i$$

for all $i = 1, \dots, n$.

For item (i) we assume (for a contradiction) that $y_0 = 0$ and consider the last column of A to obtain $\sum_{l=1}^k \hat{y}_l - y_0 = 0$ hence $\hat{y}_1 = \dots = \hat{y}_k = 0$ as $y \in \mathbb{Q}_{\geq 0}^{n+k+1}$. Now, considering the i-th column of A, we have that $0 = -\epsilon_i y_i - \sum_{l=1}^k \hat{y}_l f_l(x_i) + p_i y_0 = -\epsilon_i y_i$ whence $y_i = 0$ for all $i = 1, \dots, n$ so that, in summary, $y = 0$, contradicting $yb < 0$.

Finally, for item (ii), we first consider the last column of A and deduce from $yA = 0$ that $\sum_{l=1}^k \hat{y}_l - y_0 = 0$ so that $\sum_{l=1}^k \hat{y}_l = 1$ and μ_y is a finitely supported probability distribution as $y \in \mathbb{Q}_{\geq 0}^{n+k+1}$. Moreover, considering the i-th column of A, the equality $y \cdot A = 0$ gives

$$\begin{aligned} 0 &= -\epsilon_i y^i - \sum_{l=1}^k \hat{y}^l f_i(x_l) + y^0 p_i \\ &= -\epsilon_i y^i - \sum_{l=1}^k \hat{y}^l \cdot \sum_{j=1}^{m_i} c_i^j \cdot \mathbb{1}_{\tau(a_i^j)}(x_l) + p_i \\ &= -\epsilon_i y^i - \sum_{j=1}^{m_i} c_i^j \cdot \sum_{l=1}^{k} \hat{y}^l \cdot \mathbb{1}_{\tau(a_i^j)}(x_l) + p_i \\ &= -\epsilon_i y^i - \sum_{j=1}^{m_i} c_i^j \cdot \mu(\tau(a_i^j)) + p_i \end{aligned}$$

so that $\sum_{j=1}^{m_i} c_i^j \mu_y(\tau(a_i^j)) = p_i - \epsilon_i y_i$ as required.

According to the statement of the theorem, we distinguish the following cases.

Case 1: At least one ϵ_i is positive. The claim follows (by multiplying with a common denominator) if there exists $b = (b_1, \dots, b_n, 0, \dots, 0) \in \mathbb{Q}_{\leq 0}^n$ with $\sum_{i=1,\dots,n} b_i^2 \neq 0$ so that the system of linear inequalities

$$Ar \leq b^T \tag{1}$$

has a solution $r = (r_1, \dots, r_n, k)^T$.

Now suppose, for a contradiction, that Equation (1) does *not* have a solution for any choice of $b_1, \dots, b_n \in \mathbb{Q}_{\leq 0}$ with $\sum_{i=1}^n b_i^2 > 0$. Then, by Farkas' Lemma in the form of [14, Corollary 7.1 (e)], there exists, for every $b = (b_1, \dots, b_n, 0, \dots, 0) \in$

$\mathbb{Q}^{n+k+1}_{\leq 0}$ with $\sum_{i=1}^n b_i^2 > 0$, a vector $y_b \in \mathbb{Q}^{n+k+1}_{\geq 0}$ such that $y_b^T A = 0$ and $y_b^T \cdot b < 0$. Now consider the (non-empty) set

$$I^+ = \{i \in \{1, \dots, n\} \mid \epsilon_i > 0\}$$

and, for $i \in I^+$, the vector

$$b_i = (0, \dots, 0, -1, 0, \dots, 0)^T \in \mathbb{Q}^{n+k+1}$$

where -1 appears in the i-th coordinate. By Farkas' Lemma, this gives a vector y^i so that $(y^i)^T A = 0$ and $(y^i)^T b < 0$. In particular, $y_0^i \neq 0$ by item (i) so that we may assume that $y_0^i = 1$ by linearity for all $i \in I^+$. Moreover, $y^T b < 0$ implies that $y_i^i > 0$ for all $i \in I^+$. Now consider $y = \frac{1}{\sharp I^+} \sum_{i \in I^+} y_i$ and let $y = (y_1, \dots, y_n, \hat{y}_1, \dots, \hat{y}_k, y_0)$. By linearity, we have $y_0 = 1$, $y^T A = 0$ and $y_i = \sum_{h \in I^+} y_h^i \geq y_i^i > 0$. By item (ii) above the vector y induces a finitely supported probability measure μ so that

- $\sum_{j=1}^{m_i} c_i^j \cdot \mu(\tau(a_i^j)) = p_i - \epsilon_i y^i > p_i$ if $i \in I^+$ as $y^i > 0$ and $\epsilon_i = +1$
- $\sum_{j=1}^{m_i} c_i^j \cdot \mu(\tau(a_i^j)) = p_i - \epsilon_i y^i \leq p_i$ if $i \notin I^+$ since $\epsilon_i = -1$ and $y^i \geq 0$.

As a consequence, we have that $\mu \notin [\![\mathsf{sg}(\epsilon_i) L_{p_i}(c_i^1, \dots, c_i^{m_i})(a_i^1, \dots, a_i^{m_i})]\!]_{\mathcal{D}(X),\tau}$ for all $i = 1, \dots, n$ which contradicts our assumption that $\mathcal{D}(X), \tau \models \Gamma$. This finishes our treatment of the first case.

Case 2: $\epsilon_1 = \dots = \epsilon_n = -1$. The claim follows if we can show that there exists $b = (b_1, \dots, b_n, 0, \dots, 0) \in \mathbb{Q}^{n+k}_{\leq 0}$ with $\sum_{i=1}^n b_i^2 > 0$ so that the system

$$A_0 r \leq^T b \quad A_1 r < 0 \tag{2}$$

has a solution $r = (r_1, \dots, r_n, k)$.

Suppose for a contradiction, that (2) has no solution for all $b = (b_1, \dots, b_n, 0, \dots, 0) \in \mathbb{Q}^{n+k}_{\leq 0}$ with $\sum_{i=1}^n b_i^2 > 0$. In particular, (2) has no solution for $b = (-1, \dots, -1, 0, \dots, 0)$. By Motzkin's transposition theorem in the form of [14, Corollary 7.1 (k)], there exists

$$y = (y_1, \dots, y_n, \hat{y}_1, \dots, \hat{y}_k, y_0) \in \mathbb{Q}^n_{\geq 0}$$

so that $y^T A = y^T A_0 + y^T A_1 = 0$ and *either* $y_0 = 0$ and $yb < 0$ *or* $y_0 \neq 0$ and $yb \leq 0$.

By (i) the case $y_0 = 0$ and $y^T b < 0$ is impossible, so we may assume that $y_0 \neq 0$ and $y^T b \leq 0$, and, without loss of generality that $y_0 = 1$.

By item (ii) the vector y induces a finitely supported probability measure μ so that

$$\sum_{j=1}^{m_i} c_i^j \cdot \mu(\tau(a_i^j)) = p_i + y_i \geq p_i.$$

Hence $\mu \notin [\![\neg L_{p_i}(c_i^1, \dots, c_i^{m_i})(\tau(a_i^1), \dots, \tau(a_i^{m_i})]\!]_{\mathcal{D}(X),\tau}$ for any $i = 1, \dots, n$ which implies that $\mathcal{D}(X), \tau \not\models \Gamma$, again contradicting our assumption that $\mathcal{D}(X), \tau \models \Gamma$. Having reached a contradiction in both cases finishes the proof. □

We obtain completeness of probabilistic modal logic with respect to $\mathcal{D}$-models as a corollary of Theorem 2.4.

Corollary 3.5 *$\mathcal{R} \vdash \Gamma$ whenever $\mathcal{D} \models \Gamma$.*

We summarise our results about probabilistic modal logic with linear inequalities in the next theorem, that ties the two different semantics together.

Theorem 3.6 *Let $\Gamma \in \mathcal{S}(\Lambda)$. Then $\mathcal{D} \models \Gamma$ whenever $\mathsf{Mark} \models \Gamma$. As a consequence, the following are equivalent:*
(i) $\mathcal{R} \vdash \Gamma$ *(ii)* $\mathsf{Mark} \models \Gamma$ *(iii)* $\mathcal{D} \models \Gamma$.

Proof. We just need to show that $\mathcal{D} \models \Gamma$ whenever $\mathsf{Mark} \models \Gamma$ as the other assertions are covered in Corollary 3.5 and Proposition 3.3. So suppose that $\mathsf{Mark} \models \Gamma$ and take a $\mathcal{D}$-model $M = (X, \gamma, \pi)$ and equip X with the trivial σ-algebra $\Sigma_X = \mathcal{P}(X)$. Then $\mu(x, S) = \gamma(x)(S)$ is a Markov kernel. If $M' = (X, \mu, \pi)$ one shows by induction on the structure of formulas that $[\![A]\!]_M = [\![A]\!]'_M$ whence $M \models \Gamma$. This proves that $\mathcal{D} \models \Gamma$ as M was arbitrary. □

4 Graded Modal Logic

As for probabilistic modal logic, the graded modal logic features linear inequalities comprising the number of successors in a Kripke model. As for probabilistic modal logic, we consider modal operators $L_p(c_1, \dots, c_n)$ but $p, c_1, \dots, c_n$ are now required to be integers. In other words, we consider the modal similarity type

$$\Lambda = \{L_p(c_1, \dots, c_m) \mid m \in \mathbb{N}, p, c_1, \dots, c_m \in \mathbb{Z}\}$$

that defines the set $\mathcal{F}(\Lambda)$ of formulas of graded modal logic. For $p \in \mathbb{Z}$ we write L_p as a shorthand for the unary modality $L_p(1)$. If $A_1, \dots, A_n \in \mathcal{F}(\Lambda)$, then $L_p(c_1, \dots, c_n)(A_1, \dots, A_n)$ is valid at a point c if the linear inequality $\sum_{j=1}^{m} c_j \sharp A_j \geq p$ holds, where $\sharp A_j$ is the number of successors of c that satisfy A_j.

Example 4.1 We may use graded modal logic to reason about supporters of different football teams. Consider a Kripke model $M = (X, \gamma : X \to \mathcal{P}(X), \pi)$ where X represents individuals. We think of $x' \in \gamma(x)$ as representing that individual x "knows" x'. If $\mathsf{Arsenal}$ and $\mathsf{Chelsea}$ are propositional variables that hold for those individuals that support the respective football team, then the second author (living in North London) would satisfy the formula $L_0(1, -5)(\mathsf{Arsenal}, \mathsf{Chelsea})$ that stipulates that the individual in question knows at least 5 times as many Arsenal than Chelsea supporters – which is not valid for the first author (who resides in South London).

To obtain a sound and complete axiomatisation of graded modal logic with linear inequalities, we consider the set $\mathcal{R}$ of one-step rules that consists of all instances of

$$(G)\,\frac{\sum_{i=1}^{n} r_i \cdot \sum_{j=1}^{m_i} c_i^j a_i^j \geq 0}{\{\mathsf{sg}(r_i) L_{p_i}(c_1^i, \dots, c_{m_i}^i)(a_1^i, \dots, a_{m_i}^i) \mid i = 1, \dots, n\}}$$

where $r_1, \dots, r_n \in \mathbb{Z} \setminus \{0\}$ under the side condition

$$\sum_{r_i > 0} r_i(p_i - 1) + \sum_{r_i < 0} r_i p_i < 0$$

As before, we first discuss soundness of the ensuing sequent system with respect to Kripke frames, then provide completeness with respect to a coalgebraic semantics in terms of multigraphs, and then compare the two classes of models.

4.1 Kripke Semantics and Soundness

We define the semantics of graded modal logic with respect to image finite Kripke models, essentially following Demri and Lugiez [3] and thus generalising the definitions of Fine [6] and of Pacuit and Salame [11]. By an image finite Kripke model, we mean a triple (X, γ, π) where, as usual, X is the set of worlds, $\gamma : X \to \mathcal{P}_f(X)$ assigns a finite set of successors to every $x \in X$ and $\pi : \mathsf{V} \to \mathcal{P}(X)$ is a valuation. The semantics of $\mathcal{F}(\Lambda)$ with respect to a Kripke model $M = (X, \gamma, \pi)$ is given by the usual propositional rules, together with

$$M, x \models L_p(c_1, \dots, c_m)(A_1, \dots, A_m) \iff \sum_{j=1}^{m} c_j \cdot \sharp(\gamma(x) \cap [\![A_j]\!]_M) \geq p$$

where $[\![A]\!]_M = \{x \in X \mid M, x \models A\}$ is the truth-set of A, and $\sharp$ denotes cardinality. As usual, $M \models A$ if $M, x \models A$ for all $x \in X$. We write $\mathsf{Krip} \models \Gamma$ for $\Gamma \in \mathcal{S}(\mathcal{F}(\Lambda))$ if $M \models \bigvee \Gamma$ for all Kripke models $M = (X, \gamma, \pi)$.

Proposition 4.2 $\mathsf{Krip} \models \Gamma$ *whenever* $\mathcal{R} \vdash \Gamma$.

Proof. Consider a Kripke model $M = (X, \gamma, \pi)$ and suppose that $\mathcal{R} \vdash \Gamma$. We show that $M \models \Gamma$ by induction on the proof of $\mathcal{R} \vdash \Gamma$, where the application of an instance of (G) is the only interesting case. Consider the modal rule

$$\frac{\sum_{i=1}^{n} r_i \sum_{j=1}^{m_i} c_i^j \cdot A_i^j \geq 0}{\{\mathsf{sg}(r_i) L_{p_i}(c_i^1, \dots, c_i^{m_i})(A_i^1, \dots, A_i^{m_i}) \mid i = 1, \dots, n\}}$$

the applicability of which assumes that the side condition

$$\sum_{r_i < 0} r_i p_i + \sum_{r_i > 0} r_i (p_i - 1) > 0 \tag{3}$$

holds. By induction hypothesis, we may assume that

$$\sum_{i=1}^{n} r_i \sum_{j=1}^{m_i} c_i^j \cdot \mathbb{1}_{[\![A_i^j]\!]_M}(x) \geq 0$$

for all $x \in X$. To see that $M, x \models \Gamma$, note that the above inequality implies that

$$\sum_{i=1}^{n} r_i \sum_{j=1}^{m_i} c_i^j \cdot \sharp(\gamma(x) \cap [\![A_i^j]\!]_M) = \sum_{x' \in \gamma(x)} \sum_{i=1}^{n} r_i \sum_{j=1}^{m_i} c_i^j \cdot \mathbb{1}_{[\![A_i^j]\!]_M}(x') \geq 0 \tag{4}$$

Now suppose, for a contradiction, that $M, x \not\models \Gamma$. Then we have

$$\sum_{j=1}^{m_i} c_i^j \sharp(\gamma(x) \cap [\![A_i^j]\!]_M) < p$$

in case $r_i > 0$ and

$$\sum_{j=1}^{m_i} c_i^j \sharp(\gamma(x) \cap [\![A_i^j]\!]_M) \geq p$$

if $r_i < 0$. Combining the side condition (3) with (4) this gives

$$0 \leq \sum_{r_i<0} r_i p_i + \sum_{r_i>0} r_i(p_i - 1) < 0$$

i.e. the desired contradiction. □

This shows that graded modal logic is sound with respect to Kripke frames. We next establish completeness of graded modal logic with respect to multigraphs before we relate the different types of semantics.

4.2 Coalgebraic Semantics and Completeness

We consider the functor

$$\mathcal{B}(X) = \{f : X \to \mathbb{N} \mid \mathrm{supp}(f) \text{ is finite}\}$$

that extends to a Λ-structure by stipulating that

$$[\![L_p(c_1, \dots, c_m)]\!]_X(S_1, \dots, S_m) = \{f \in \mathcal{B}(X) \mid \sum_{j=1}^{m} c_j \cdot f(S_j) \leq p\}$$

where X is a set and $S_1, \dots, S_m \subseteq X$ and $f(S) = \sum_{x \in S} f(x)$ for $S \subseteq X$. We may think of a $\mathcal{B}$-coalgebra $(X, \gamma : X \to \mathcal{B}(X))$ as a multigraph where every edge is assigned an integer weight.

For the whole section, we fix the set $\mathcal{R}$ of one-step rules that comprises all instances of (G). Completeness of graded modal logic with linear inequalities is proved using a variant of a canonical model construction, but at the level of one-step formulas. We note two simple properties that correspond to admissibility of contraction of cut, but at the level of one-step derivations.

Lemma 4.3 *Suppose $\tau : \mathsf{V} \to \mathcal{P}(X)$ is a valuation.*

(i) if Γ, A, A is τ-derivable, then so is Γ, A.

(ii) if Γ, A and $\Gamma, \neg A$ are τ-derivable, then so is Γ.

Proof. Both are immediate from the rule format: for the first item, we obtain a new instance of (G) that witnesses derivability of Γ, A by simply adding the coefficients that induce both occurrences of A. For the second item, we are given two rule instances that witness derivability of Γ, A and $\Gamma, \neg A$ that we normalise so that the coefficients associated with A and $\neg A$ have the same magnitude and then add (the coefficients of) both rules. □

The next lemma ensures that every consistent set of formulas can be satisfied, at the one-step level, by an assignment of integer weights that is bounded.

Lemma 4.4 *Suppose that $\tau : \mathsf{V} \to \mathcal{P}(X)$ is a valuation and let $\Gamma \in \mathcal{S}(\Lambda(\mathsf{V}))$ so that Γ is not τ-derivable.*

(i) for all $a \in \mathsf{V}$, there exists $p \geq 0$ so that $\Gamma, L_p a$ is not *τ-derivable.*

(ii) If $L_p a \in \Gamma$, then $\Gamma, \neg L_{p+n} a$ is τ-derivable for all $n \geq 0$.

Proof. The first item is by contraposition: If $\Gamma, L_p A$ were derivable for all $p \geq 0$ we obtain a contradiction in terms of the inequalities in premise and side condition of the rules. For the second item, one shows that $L_p a, \neg L_{p+n} a$ is derivable for all $n \geq 0$. □

One-step completeness is now content of the following lemma.

Lemma 4.5 *Suppose that $\tau : \mathsf{V} \to \mathcal{P}(X)$ is a valuation, X is finite and $\Gamma \in \mathcal{S}(\Lambda(\mathsf{V}))$. If $\mathcal{B}(X), \tau \models \Gamma$, then Γ is τ-derivable.*

Proof. Suppose, for a contradiction, that Γ is *not* τ-derivable and pick, for all $x \in X$, pairwise distinct propositional variables b_x not occurring in Γ. Repeated application of Lemma 4.4 gives, for all (finitely many) $x \in X$, a number $k_x \in \mathbb{N}$ so that

$$\Gamma' = \Gamma \cup \{L_{k_x} b_x \mid x \in X\}$$

is not τ-derivable. By Lemma 4.3 the same holds for the sequent $\mathrm{supp}(\Gamma)$ that we may extend to a maximal subset $\mathcal{M} \subseteq \Lambda(\mathsf{V}) \cup \neg\Lambda(\mathsf{V})$ with the property that no finite subset $\Delta \subseteq \mathcal{M}$, viewed as a multiset where every element has multiplicity one, is derivable. We now define a measure $\mu : X \to \mathbb{N}$ by

$$\mu(x) = \max\{p \in \mathbb{N} \mid \neg L_p(b_x) \in \mathcal{M}\}$$

for all $x \in X$ and write $\mu(A) = \sum_{x \in A} \mu(x)$ as usual. Note that $\mu(x) \in \mathbb{N}$ by Lemma 4.4. We now claim that

$$\neg L_{\mu(\tau(a))} \in \mathcal{M} \quad \text{and} \quad L_{\mu(\tau(a))+1}(a) \in \mathcal{M}$$

for all $a \in \mathsf{V}$. For the first point, note that $\neg L_{\mu(x)} b_x \in \mathcal{M}$ by definition of μ and consider the rule

$$\frac{-\sum_{x \in \tau(a)} b_x + a \geq 0}{\{\neg L_{\mu(x)} b_x \mid x \in \tau(a)\} \cup \{L_{\mu(\tau(a))+1} a\}}$$

that witnesses $L_{\mu(\tau(a))+1} \notin \mathcal{M}$ as $\mathcal{M}$ is not derivable, and hence $\neg L_{\mu(\tau(a))+1} a \in \mathcal{M}$ by Lemma 4.3. The proof of the second point is entirely dual.

We now establish that

$$\mathsf{sg}(\epsilon) A \in \mathcal{M} \implies \mu \notin [\![\mathsf{sg}(\epsilon) A]\!]_{(\mathcal{B}(X), \tau)}$$

for all $A \in \mathcal{F}(\Lambda)$ and all $\epsilon \in \{-1, +1\}$. Let $A = L_p(c_1, \ldots, c_m)(a_1, \ldots, a_m)$. For the case $\epsilon = +1$ assume for a contradiction that $\mu \in [\![A]\!]_{(\mathcal{B}(X), \tau)}$ so that

$$\sum_{j=1}^{m} c_j \mu(a_j) \geq p.$$

Consider the rule (the side condition of which is readily established)

$$\frac{c_1 a_1 + \cdots + c_m a_m + \sum_{j=1}^m c_j a_j \geq 0}{\{\neg L_{\mu(\tau(a_j))} a_j \mid c_j < 0\} \cup \{L_{\mu(\tau(a_j))+1} a_j \mid c_j > 0\} \cup \{L_p(c_1, \ldots, c_j)(a_1, \ldots, a_j)\}}$$

that witnesses $A \notin M$ as $\mathcal{M}$ is not derivable, contradicting $A \in \mathcal{M}$. The case for $\epsilon = -1$ is entirely dual.

In summary our assumption that Γ is not τ-derivable, we obtain that $\mu \notin [\![\Gamma]\!]_{(\mathcal{B}(X),\tau)}$ contradicting that Γ is τ-valid.

$\square$

As a corollary, we obtain that the set $\mathcal{R}$ comprising all instances of (G) is one-step complete.

Proposition 4.6 $\mathcal{R} \vdash \Gamma$ *whenever* $\mathcal{B} \models \Gamma$.

Proof. By Lemma 4.5 we have that $\mathcal{R}$ is one-step cut-free complete over finite sets, which implies that $\mathcal{R}$ is one-step complete. This can either be seen as a consequence of [13, Proposition 4.5] or directly: Given that $\mathcal{B}X, \tau \models \Gamma$ let V_0 denote the propositional variables that occur in Γ and define an equivalence relation $\sim$ on X by $x \sim y \iff \forall p \in \mathsf{V}_0 (x \in \tau(p) \iff y \in \tau(p))$. Let $X_0 = X/\sim$ and $\tau_0(p) = \{[x]_\sim \mid x \in \tau(p)\}$. Then $\mathcal{B}(X_0), \tau_0 \models \Gamma$ by naturality of predicate liftings whence Γ is τ_0-derivable by Lemma 4.5 which implies τ-derivability of Γ. Completeness now follows from one-step completeness (Theorem 2.4). $\square$

Theorem 4.7 *Let* $\Gamma \in \mathcal{S}(\Lambda)$. *Then* $\mathcal{B} \models \Gamma$ *whenever* $\mathsf{Krip} \models \Gamma$. *In particular, the following are equivalent:*

(i) $\mathcal{R} \vdash \Gamma$ *(ii)* $\mathsf{Krip} \models \Gamma$ *(iii)* $\mathcal{B} \models \Gamma$

witnessing soundness and completeness of graded modal logic with linear inequalities both over Kripke frames and multigraphs.

Proof. We only need to show that $\mathcal{B} \models \Gamma$ whenever $\mathsf{Krip} \models \Gamma$, as the other claims are consequences of Proposition 4.2 and Proposition 4.6. So suppose that $\Gamma \in \mathcal{S}(\mathcal{F}(\Lambda))$ and $M = (X, \gamma, \pi)$ is a $\mathcal{B}$-model so that $M \not\models \Gamma$, i.e. there exists $x_0 \in X$ so that $M, x_0 \not\models \Gamma$. We construct a Kripke model $M' = (X', \gamma', \pi')$ by unravelling at x_0: we put

- $X' = \{x_0 \stackrel{w_1}{\to} x_1 \stackrel{w_2}{\to} \cdots \stackrel{w_n}{\to} x_n \mid n \geq 0, 0 \leq w_i < \gamma(x_i)(x_{i+1})\}$
- $\gamma'(x_0 \stackrel{w_1}{\to} \cdots \stackrel{w_n}{\to} x_n) = \{x_0 \stackrel{w_1}{\to} \cdots \stackrel{w_n}{\to} x_n \stackrel{w_{n+1}}{\to} x_{n+1} \mid 0 \leq w_{n+1} < \gamma(x_n)(x_{n+1})\}$
- $\pi'(p) = \{x_0 \stackrel{w_1}{\to} \cdots \stackrel{w_n}{\to} x_n \mid x_n \in \pi(p)\}$.

In other words, the worlds of the Kripke model (X', γ', π') are the paths from the initial point $x_0 \in X$ where we make duplicates of states according to the multiplicity of the transition. It now follows by induction on the structure of formulas that

$$M, x \models A \implies M', x' \models A$$

whenever x' is of the form $x_0 \stackrel{m_1}{\to} \cdots \stackrel{w_n}{\to} x_n$ with $x_n = x$. In particular, $M' \not\models \Gamma$ as we had to show. $\square$

5 Stochastic Logic

We may think of stochastic logic as a hybrid between probabilistic modal logic and graded modal logic. As for probabilistic modal logic, every state of a model is equipped with a measure, but we do not insist that this measure be a probability measure. If c is a state in a stochastic model and μ the ensuing measure, we may think of $\mu([\![A]\!])$ as the *total* cost of observing event A in the next transition state. As before, our formulas are linear inequalities in terms of the measures of the (truth sets of) formulas. In other words, we consider the similarity type

$$\Lambda = \{L_p(c_1, \dots, c_m) \mid m \in \mathbb{N}, p, c_1, \dots, c_m \in \mathbb{Q}\}$$

defining the formulas $\mathcal{F}(\Lambda)$ of stochastic modal logic. (Note that the syntax of stochastic modal logic is identical to that of probabilistic modal logic.) Informally speaking, the formula $L_p(c_1, \dots, c_m)(A_1, \dots, A_m)$ is valid at a point c, if the linear inequality $\sum_{j=1}^m c_j \mu(A_j) \geq p$ holds, where $\mu(A_j)$ is the measure of the truth-set of A_j as seen from point c.

Example 5.1 At the time of writing this paper, both authors frequently discussed the outcome of the (then) upcoming general election in their country of residence. To this effect, one may consider a stochastic model based on the set of inhabitants of said country. To every inhabitant c, we associate a measure that – applied to a subset S of the population – yields the overall amount of persuasion (measured as a non-negative real number) that c would have to apply in order to swing the votes of all elements of S into a particular direction. If Tory and Labour are propositional variables that denote the respective political angle, there was a heated debate whether the formula $\neg L_0(1, -1)(\mathsf{Tory}, \mathsf{Labour})$, $L_0(1, -1)(\mathsf{Labour}, \mathsf{Tory})$ or $L_0(1, -1)(\mathsf{Tory}, \mathsf{Labour}) \wedge L_0(1, -1)(\mathsf{Labour}, \mathsf{Tory})$ yields the most realistic model (both authors still hope that this did apply to $L_0(-1)(\top)$).

A sound and complete axiomatisation of stochastic modal logic will be provided by the set $\mathcal{R}$ of one-step rules comprising all instances of

$$(S)\frac{\sum_{i=1}^n r_j \sum_{j=1}^{m_i} c_i^j a_i^j \geq 0}{\{\mathsf{sg}(\epsilon_i) L_{p_i}(c_i^1, \dots, c_i^{m_i})(a_i^1, \dots, a_i^{m_i}) \mid i = 1, \dots, n\}}$$

where $r_1, \dots, r_n \in \mathbb{Z} \setminus \{0\}$ that satisfy the side condition

$$\sum_{i=1}^n r_i p_i < 0 \quad \text{if all } r_i < 0, \text{ and} \quad \sum_{i=1}^m r_i p_i \leq 0 \quad \text{otherwise.}$$

We now establish soundness of stochastic modal logic with respect to the class of finite measures, prove completeness of stochastic modal logic with respect to finitely based measures, and then align both views.

5.1 Measurable Semantics and Soundness

Definition 5.2 *A* measurable model *is a triple* (X, μ, π) *where* X *is a measurable space with* σ*-algebra* Σ_X*,* $\pi : \mathsf{V} \to \Sigma_X$ *is a valuation and* $\mu : X \times \Sigma_X \to [0, \infty)$ *is a* *measurable kernel, i.e. a function so that $\mu(\cdot, B) : X \to [0, \infty)$ is measurable for all $B \in \Sigma_X$ and $\mu(x, \cdot) : \Sigma_X \to [0, \infty)$ is a measure on X.*

Note that we require the measure that is induced my a measurable kernel is always finite. The semantics of $\mathcal{F}(\Lambda)$ with respect to measurable models is as expected for the propositional rules (note that atomic propositions are mapped to measurable sets) and the clause for a modal operator is

$$[\![L_p(c_1, \ldots, c_m)(A_1, \ldots, A_m)]\!]_M = \{x \in X \mid \sum_{j=1}^{m} c_j \cdot \mu(x, [\![A_j]\!]_M) \geq p\}$$

so that $[\![A]\!]_M$ is a measurable set for all $A \in \mathcal{F}(\Lambda)$. We write $M, x \models A$ if $x \in [\![A]\!]_M$ and $M \models A$ if $M, x \models A$ for all $x \in X$. Finally, $\mathsf{Meas} \models \Gamma$ if $M \models \bigvee \Gamma$ for all measurable models M. Soundness of stochastic modal logic over measurable models is similar to soundness of probabilistic modal logic and takes the following form.

Proposition 5.3 $\mathsf{Meas} \models \Gamma$ *whenever* $\mathcal{R} \vdash \Gamma$.

Proof. By induction on the proof of $\mathcal{R} \vdash \Gamma$ analogous to the proof of Proposition 3.3. □

This shows that stochastic modal logic is sound with respect to measurable models. We now look upon stochastic modal logic coalgebraically and establish completeness.

5.2 Coalgebraic Semantics and Completeness

To interpret stochastic modal logic over coalgebraic models, we consider the functor

$$\mathcal{M}(X) = \{\mu : X \to [0, \infty) \mid \mathrm{supp}(\mu) \text{ finite}\}$$

and we write $\mu(S) = \sum_{x \in S} \mu(x)$ whenever $\mu \in \mathcal{M}(X)$ and $S \subseteq X$. The functor $\mathcal{M}$ extends to a Λ-structure by virtue of

$$[\![L_p(c_1, \ldots, c_m)]\!]_X(S_1, \ldots, S_m) = \{\mu \in \mathcal{M}(X) \mid \sum_{j=1}^{m} c_j \cdot \mu(S_j) \geq p\}$$

where $S_1, \ldots, S_m \subseteq X$. We may think (modulo currying) of $\mathcal{M}$-coalgebras as measurable kernels with finite support. Completeness of stochastic modal logic over $\mathcal{M}$-coalgebras is established by means of the following lemma that again uses results from linear programming:

Lemma 5.4 *Consider a valuation $\tau : \mathsf{V} \to \mathcal{P}(X)$ and suppose that $\Gamma \in \mathcal{S}(\Lambda(\mathsf{V}))$ is τ-valid. Then Γ is τ-derivable.*

Proof. We proceed as in the proof of Lemma 3.4 but instead consider the matrix

$$A_0 = \begin{pmatrix} -\epsilon_1 & & 0 \\ & \ddots & \vdots \\ 0 & & -\epsilon_n \\ -f_1(x_1) & \dots & -f_n(x_1) \\ \vdots & \vdots & \\ -f_1(x_k) & \dots & -f_n(x_k) \end{pmatrix}$$
$$A_1 = \begin{pmatrix} p_1 & \dots & p_n \end{pmatrix}$$

where $f_i = \sum_{j=1}^{m_i} c_i^j \cdot \mathbb{1}_{\tau(a_i^j)}$ and let $A = \begin{pmatrix} A_0 \\ A_1 \end{pmatrix}$ and proceed as in the proof of Lemma 3.4 (where the absence of the last column means that we do not define a probability distribution). $\square$

As a corollary, we obtain that stochastic modal logic is complete over $\mathcal{M}$-coalgebras.

Corollary 5.5 *$\mathcal{R} \vdash \Gamma$ whenever $\mathcal{M} \models \Gamma$.*

In summary, we obtain the following theorem for stochastic modal logic:

Theorem 5.6 *Let $\Gamma \in \mathcal{S}(\Lambda)$. Then $\mathcal{M} \models \Gamma$ whenever $\mathsf{Meas} \models \Gamma$. In particular, the following are equivalent:*

(i) $\mathcal{R} \vdash \Gamma$ *(ii)* $\mathsf{Meas} \models \Gamma$ *(iii)* $\mathcal{M} \models \Gamma$.

Proof. We only need to show that $\mathcal{M} \models \Gamma$ whenever $\mathsf{Meas} \models \Gamma$, as the remaining assertions are the content corollary 5.5 and Proposition 5.3. So suppose that (X, γ, π) is a $\mathcal{M}$-model and $\mathsf{Meas} \models \Gamma$. We equip X with the trivial σ-algebra $\mathcal{P}(X)$ and consider the measurable kernel $\mu(x, S) = \sum_{x' \in S} \gamma(x)(x')$. Note that μ is well-defined as $\gamma(x)$ has finite support, and so defines a measurable kernel. Let $M' = (X, \mu, \pi)$. One now shows by induction on the structure of formulas that $M, x \models A \iff M', x \models A$ for all $x \in X$ which finishes the proof. $\square$

6 Complexity

Given that we have coalgebraised all three logics by equipping them with a sound and complete coalgebraic semantics, we are now in a position to use generic (coalgebraic) methods to establish complexity bounds. As we have a characterisation of universal validity in terms of a cut-free sequent calculus where the size of the formulas strictly decreases when we move from conclusion to premise, we can map the decidability problem onto backwards proof search, which we can be seen as the problem of searching a tree the length of whose branches is polynomially bounded. To see that this problem is in polynomial space, we have to agree on representations for modal operators. Here, we represent numbers in binary, that is, we put

$\mathsf{size}(n) = \lceil \log_2 n \rceil$ and $\mathsf{size}(p/q) = \mathsf{size}(p) + \mathsf{size}(q)$ which allows us to define the size of a modal operator as $\mathsf{size}(L_p(c_1, \ldots, c_m))$ as $\mathsf{size}(p) + \sum_{j=1}^{m} \mathsf{size}(c_i)$. To show decidability in PSPACE, we have to show that we can encode rules into strings of polynomial length so that all premises can be decided in NP. The formal definition is as follows.

Definition 6.1 *A set $\mathcal{R}$ of one-step rules is* PSPACE-tractable *if there exists a polynomial p such that all substitution instances of rules with conclusion $\Gamma \in \mathcal{S}(\Lambda(\mathcal{F}(\Lambda))$ can be encoded into a string of length at most $p(|\Gamma|)$ and it can be decided in* NP *whether*

- *a code represents a (substitution instance) of a rule with a given conclusion Γ*
- *a sequent belongs to the set of premises of a rule given as a code.*

We can use the methods presented in in [16] to show that the rule sets comprising of (P), (G) and (S) are indeed PSPACE-tractable.

Lemma 6.2 *If $\mathcal{R}$ comprises all instances of (P), (G) or (S), then $\mathcal{R}$ is* PSPACE-*tractable.*

Proof. It has been argued in [16, Lemma 6.16] that the coefficients r_i that occur in the rule sets (P), (G) and (S) can be polynomially bounded in the size of the linear inequalities (and hence in the size of the rule conclusions), and our argument is essentially identical to Example 6.17 of *op.cit.*. □

As a consequence, we obtain a PSPACE upper bound for all three logics considered in this paper.

Theorem 6.3 *The satisfiability problem of probabilistic modal logic, graded modal logic and stochastic modal logic (each considered with linear inequalities) is decidable in* PSPACE.

Proof. One can either invoke Theorem 6.13 of [16] or directly argue in terms of proof search where the branches of every putative proof tree are polynomially bounded in length, their nodes can be represented by strings of polynomial length, and membership in nodes can be decided in NP, all of which are consequences of tractability. This gives decidability of satisfiability in connection with the completeness (Theorem 3.6, Theorem 4.7 and Theorem 5.6). □

7 Conclusions

In this paper, we have given complete, cut-free axiomatisations of three modal logics that use linear inequalities to express constraints between probabilities of events, the number of successors in a Kripke model or, more generally, the measure of a successor set. In each case, completeness was established with the help of coalgebraic semantics, where we just had to show that a given set of (one-step) rules is one-step complete: the actual statement of completeness then follows from the general (coalgebraic) theory. As such, this paper tries to demonstrate the usefulness of the coalgebraic approach per se – we did not develop the general theory of coalgebraic logics, but just used off-the-shelf results to obtain completeness and complexity bounds. The semantics of graded and stochastic modal logic was given here in terms of image finite Kripke frames and bounded measures. It is an open problem whether the semantics can be extended to the general case in a sound fashion.

References

[1] Baader, F., D. Calvanese, D. McGuinness, D. Nardi and P. Patel-Schneider, editors, "The Description Logic Handbook," CUP, 2003.

[2] Cîrstea, C., C. Kupke and D. Pattinson, *EXPTIME tableaux for the coalgebraic μ-calculus*, in: E. Grädel and R. Kahle, editors, *Proc. CSL 2009*, number 5771 in Lect. Notes in Comp. Sci., 2009, pp. 179–193.

[3] Demri, S. and D. Lugiez, *Presburger modal logic is only PSPACE-complete*, in: U. Furbach and N. Shankar, editors, *Proc. IJCAR 2006*, Lect. Notes in Artif. Intell. **4130** (2006), pp. 541–556.

[4] Fagin, R. and J. Halpern, *Reasoning about knowledge and probability*, J. ACM **41** (1994), pp. 340–367.

[5] Fagin, R., J. Halpern and N. Megiddo, *A Logic for Reasoning about Probabilities*, Information and Computation **87** (1990), pp. 78–128.

[6] Fine, K., *In so many possible worlds*, Notre Dame J. Formal Logic **13** (1972), pp. 516–520.

[7] Goré, R., C. Kupke and D. Pattinson, *Optimal tableau algorithms for coalgebraic logics*, in: J. Esparza and R. Majumdar, editors, *Proc. TACAS 2010*, Lect. Notes in Comp. Sci., 2010.

[8] Goré, R., C. Kupke, D. Pattinson and L. Schröder, *Global caching for coalgebraic description logics*, in: J. Giesl and R. Haehnle, editors, *Proc. IJCAR 2010*, 2010, to appear.

[9] Heifetz, A. and P. Mongin, *Probabilistic logic for type spaces*, Games and Economic Behavior **35** (2001), pp. 31–53.

[10] Horrocks, I. and P. Patel-Schneider, *Optimising description logic subsumption*, J. Logic Comput. **9** (1999), pp. 267–293.

[11] Pacuit, E. and S. Salame, *Majority logic*, in: D. Dubois, C. Welty and M.-A. Williams, editors, *Proc. KR 2004* (2004), pp. 598–605.

[12] Pattinson, D., *Coalgebraic modal logic: Soundness, completeness and decidability of local consequence*, Theoret. Comput. Sci. **309** (2003), pp. 177–193.

[13] Pattinson, D. and L. Schröder, *Cut elimination in coalgebraic logics*, Information and Computation (2010), accepted for publication.

[14] Schrijver, A., "Theory of linear and integer programming," Wiley Interscience, 1986.

[15] Schröder, L., *A finite model construction for coalgebraic modal logic*, J. Log. Algebr. Program. **73** (2007), pp. 97–110.

[16] Schröder, L. and D. Pattinson, *PSPACE bounds for rank-1 modal logics*, ACM Transactions on Computational Logics **10** (2009).

[17] Schröder, L., D. Pattinson and C. Kupke, *Nominals for everyone*, in: C. Boutilier, editor, *Proc. IJCAI 2009*, 2009, pp. 917–922, online proceedings.

[18] Tobies, S., *PSPACE reasoning for graded modal logics*, J. Logic Comput. **11** (2001), pp. 85–106.

[19] Troelstra, A. and H. Schwichtenberg, "Basic Proof Theory," Number 43 in Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, 1996.

[20] Tsarkov, D., I. Horrocks and P. Patel-Schneider, *Optimizing terminological reasoning for expressive description logics*, J. Autom. Reasoning **39** (2007), pp. 277–316.