A van Benthem Theorem for Fuzzy Modal Logic

Paul Wild
Lutz Schröder
Friedrich-Alexander-Universität Erlangen-Nürnberg

Dirk Pattinson
Australian National University, Canberra

Barbara König
Universität Duisburg-Essen

Abstract

We present a fuzzy (or quantitative) version of the van Benthem theorem, which characterizes propositional modal logic as the bisimulation-invariant fragment of first-order logic. Specifically, we consider a first-order fuzzy predicate logic along with its modal fragment, and show that the fuzzy first-order formulas that are non-expansive w.r.t. the natural notion of bisimulation distance are exactly those that can be approximated by fuzzy modal formulas.

CCS Concepts • **Theory of computation → Modal and temporal logics**; *Description logics*;

Keywords Fuzzy modal logic, behavioural metrics, correspondence theory, modal characterization theorems, description logics

ACM Reference Format:
Paul Wild, Lutz Schröder, Dirk Pattinson, and Barbara König. 2018. A van Benthem Theorem for Fuzzy Modal Logic.

1 Introduction

Fuzzy logic is a form of multi-valued logic originally studied by Łukasiewicz and Tarski [27] and later popularized as a logic of *vagueness* by Zadeh [52]. It is based on replacing the standard set of Boolean truth values with a different lattice, most often, like in the present paper, the unit interval. Saying that a formula ϕ has truth value $r \in [0, 1]$ then means that ϕ holds with *degree* r, which would apply to typical vague qualifications such as a given person being tall (in contrast to assigning a *probability* $p \in [0, 1]$ to ϕ, which would be read as saying that ϕ is either completely true with probability p or completely false with probability $1 - p$, as in 'the die under the cup shows a 3 with probability p').

Beyond the original propositional setup, fuzzy truth values appear in variants of more expressive logics, notably in *fuzzy first-order logics* [10, 20, 31] and in various *fuzzy modal logics*. The latter go back to many-valued modal logics based on making valuations in Kripke models [29, 30, 32, 35, 38, 41] or additionally also the accessibility relation [19] many-valued, and are nowadays maybe most popular in their incarnation as fuzzy description logics (e.g. [21, 37, 42, 45, 51]; see [28] for an overview). Many-valued modal fixpoint logics are also used in software model checking (e.g. [7, 26]).

Like in the classical case, fuzzy modal logics typically embed into their first-order counterparts. In the classical setting, the core result on this embedding is *van Benthem's theorem*, which states that a first-order formula ϕ is equivalent to a modal formula if and only if ϕ is invariant under bisimulation [46]. This is a form of expressive completeness: Modal logic expresses only bisimulation-invariant properties, but for such properties it is as expressive as first-order logic. Briefly, the aim of the current paper is to provide a counterpart of this theorem for a fuzzy modal logic.

There is a wide variety of possible semantics for the fuzzy propositional connectives (see [28] for an overview), employing, e.g., the additive structure (*Łukasiewicz logic*), the multiplicative structure (*product logic*) or the Heyting algebra structure (*Gödel logic*) of the unit interval. For technical reasons, we work with the simplest possible semantics where conjunction is interpreted as minimum and all other connectives are derived using the classical encodings, effectively a fragment of Łukasiewicz logic often called *Zadeh logic*. That is, we consider *Zadeh fuzzy modal logic*, more precisely *Zadeh fuzzy K* or in description logic terminology *Zadeh fuzzy* $\mathcal{ALC}$ [42], with *Zadeh fuzzy first-order logic* as the first-order correspondence language, essentially the Zadeh fragment of Novak's Łukasiewicz fuzzy first order logic [31].

It has long been recognized that for quantitative systems, notions of *behavioural distance* are more natural than two-valued bisimilarity [48]. In such a metric setting, bisimulation invariance becomes non-expansivity w.r.t. behavioural distance (e.g. if one views classical bisimilarity as a $\{0, 1\}$-valued pseudometric, then non-expansivity means that distance 0 is preserved, which is precisely bisimulation invariance). The first step in our program is therefore to establish a notion of behavioural distance for fuzzy relational systems. We consider three different ways to define such a behavioural metric: via the modal logic, via a bisimulation game (similarly as in work on probabilistic systems [13]), or via a fixpoint characterization based on the Kantorovich lifting (similarly as in [4]). We show that they all coincide; in particular we obtain a Hennessy-Milner type theorem (behavioural distance equals logical distance). This gives us a stable notion of behavioural metric for fuzzy relational systems.

Our main result then says that *the fuzzy modal formulas lie dense in the bisimulation-invariant first-order formulas*, where by bisimulation-invariant we now mean non-expansive w.r.t. behavioural distance. In other words, every bisimulation-invariant fuzzy first-order formula can be modally approximated. The proof follows a strategy introduced for the classical case by Otto [33], going via locality w.r.t. an adapted notion of Gaifman distance to show that every bisimulation-invariant fuzzy first-order formula is already non-expansive w.r.t. *depth-k* behavioural distance for some k (this distance arises, e.g., by limiting the bisimulation game to k rounds). The key part of our technical development is, then, to establish a fuzzy counterpart of what in the classical case is a triviality: The classical proof ends in remarking that *every* state property (without any assumption of first-order definability) of relational transition systems that is invariant under depth-k bisimilarity is expressible by a modal formula of modal rank k. In the fuzzy setting, this becomes a non-trivial result of independent interest: *The fuzzy modal formulas of modal rank k lie dense in the fuzzy state properties that are non-expansive w.r.t. depth-k behavioural distance.*

Proofs are mostly omitted or only sketched; full proofs are in the appendix.

Related Work Van Benthem's theorem was later shown by Rosen [34] to hold also over finite structures. Modal characterization theorems have since been proved in various settings, e.g. logics with frame conditions [11], coalgebraic modal logics [40], fragments of XPath [1, 18, 44], neighbourhood logic [22], modal logic with team semantics [24], modal μ-calculi (within monadic second order logics) [15, 23], PDL (within weak chain logic) [9], modal first-order logics [43, 47], and two-dimensional modal logics with an $S5$-modality [50]. All these results concern two-valued logics; we are not aware of any previous work of this type for fuzzy modal logics.

There is, however, work on behavioural distances and fuzzy bisimulation in connection with fuzzy modal logic. We discuss only fuzzy notions of bisimulation, omitting work on classical behavioural equivalence for fuzzy transition systems and fuzzy automata. Balle et al. [5] consider bisimulation metrics for weighted automata in order to characterize approximate minimization. Cao et al. [8] study a notion of behavioural distance for fuzzy transition systems, where the lifting of the metric is derived from a transportation problem (the dual of the Kantorovich metric), but without considering modal logics. Fan [16] proves a Hennessy-Milner type theorem for a fuzzy modal logic with Gödel semantics and a notion of fuzzy bisimilarity. In [17] she considers an application to social network analysis and also observes that Łukasiewicz logic is problematic in this context (since the operators do not preserve non-expansivity). Eleftheriou et al. [14] show a Hennessy-Milner theorem for Heyting-valued modal logics as introduced by Fitting [19].

While we work in a fuzzy setting, we were inspired by related work on probabilistic systems: Desharnais et al. studied behavioural distances on logics [12] as well as a game characterization of probabilistic bisimulation [13]. A Hennessy-Milner theorem for the probabilistic case is presented in [48], based on a coalgebraic semantics.

2 Fuzzy Modal Logic

We proceed to recall the syntax and semantics of *Zadeh fuzzy K* or equivalently *Zadeh fuzzy $\mathcal{ALC}$* [42], along with its first-order correspondence language. For simplicity we restrict the exposition to the unimodal case; the development extends straightforwardly to the multimodal case by just adding more indices. Formulas ϕ, ψ of *fuzzy modal logic* are given by grammar

$$\phi, \psi ::= c \mid p \mid \phi \ominus c \mid \neg\phi \mid \phi \wedge \psi \mid \Diamond\phi$$

where p ranges over a fixed set At of *propositional atoms* and $c \in \mathbb{Q} \cap [0, 1]$ over rational truth constants. The syntax is thus mostly the same as for standard modal logic; the only additional ingredients are the truth constants and *modified subtraction* $\ominus$ as used in real-valued modal logics for probabilistic systems [48]. Further logical connectives are defined by the classical encodings, e.g. $\phi \vee \psi$ abbreviates $\neg(\neg\phi \wedge \neg\psi)$, and $\phi \rightarrow \psi$ abbreviates $\neg\phi \vee \psi$; also, we introduce a dual modality $\Box$ as $\Box\phi := \neg\Diamond\neg\phi$. The *rank* $\mathsf{rk}(\phi)$ of a formula ϕ is the maximal nesting depth of the modality $\Diamond$ *and propositional atoms* in ϕ. Formally, $\mathsf{rk}(\phi)$ is thus defined recursively by $\mathsf{rk}(c) = 0$, $\mathsf{rk}(p) = 1$, $\mathsf{rk}(\Diamond\phi) = 1 + \mathsf{rk}(\phi)$, and obvious clauses for the remaining constructs. We write $\mathcal{L}_k$ for the set of modal formulas of rank at most k.

The *semantics* of the logic is defined over *fuzzy relational models* (or just *models*)

$$\mathcal{A} = (A, (p^{\mathcal{A}})_{p \in \mathsf{At}}, R^{\mathcal{A}})$$

consisting of a set A of *states*, a map $p^{\mathcal{A}} : A \rightarrow [0, 1]$ for each $p \in \mathsf{At}$, and a map $R^{\mathcal{A}} : A \times A \rightarrow [0, 1]$; we will drop superscripts $\mathcal{A}$ when clear from the context. That is, propositional atoms are interpreted as fuzzy predicates on the state set, and states are connected by a binary fuzzy transition relation, where *fuzzy* is short for $[0, 1]$-valued (as usual, we use *crisp* as an informal opposite of fuzzy, i.e. crisp means two-valued). Fuzzy relational models are a natural fuzzification of Kripke models, and in fact the instantiation of latticed Kripke models over DeMorgan lattices [7, 26] to the lattice $[0, 1]$; they arise from *fuzzy transition systems* (e.g. [8]; *fuzzy automata* go back as far as [49]) by adding propositional atoms. Unless stated otherwise, we adhere to the convention that models are denoted by calligraphic letters and their state sets by the corresponding italic.

We use $\wedge$, $\vee$ to denote meets and joins in $[0, 1]$. A modal formula ϕ is then assigned a fuzzy truth value $\phi_{\mathcal{A}}(a)$, or just

$\phi(a)$, at every state $a \in A$, defined inductively by

$$c(a) = c \qquad p(a) = p^{\mathcal{A}}(a)$$
$$(\phi \ominus c)(a) = \max(\phi(a) - c, 0)$$
$$(\neg\phi)(a) = 1 - \phi(a)$$
$$(\phi \wedge \psi)(a) = \phi(a) \wedge \psi(a)$$
$$(\Diamond\phi)(a) = \textstyle\bigvee_{a' \in A}(R^{\mathcal{A}}(a, a') \wedge \phi(a')).$$

For brevity, we often conflate formulas and their evaluation functions in both notation and vernacular, e.g. in statements claiming that certain modal formulas form a dense subset of some set of state properties.

Remark 2.1. As indicated above, we thus equip the propositional connectives with Zadeh semantics. This corresponds to widespread usage but is not without disadvantages in comparison to Łukasiewicz semantics, which defines the conjunction of $a, b \in [0, 1]$ as $\max(a + b - 1, 0)$; e.g. implication is the residual of conjunction in Łukasiewicz semantics but not in Zadeh semantics (see [25] for a more detailed discussion). We will later point out where this choice becomes most relevant; roughly speaking, Łukasiewicz semantics is not easily reconciled with behavioural distance.

The modal syntax as given above is essentially identical to the one used by van Breughel and Worrell to characterize behavioural distance in probabilistic transition systems [48]. Semantically, fuzzy models differ from probabilistic ones in that they do not require truth values of successor edges to sum up to 1, and moreover in the probabilistic setting the modality $\Diamond$ is interpreted by expected truth values instead of suprema. The semantics of the propositional connectives, on the other hand, is in fact the same in both cases.

Example 2.2. We can see fuzzy K as a logic of fuzzy transition systems (e.g. [8]). E.g. the formula $\Diamond\Box 0$ then describes, roughly speaking, the degree to which a deadlocked state can be reached in one step. Formally, $(\Box 0)(y)$ is the degree to which a state y in a model $\mathcal{A}$ is deadlocked, i.e. the infimum over $1 - r$ where r ranges over the degrees $R^{\mathcal{A}}(y, z)$ to which any state z is a successor of y. Then, $(\Diamond\Box 0)(x)$ is the supremum of $\min(R^{\mathcal{A}}(x, y), (\Box 0)(y))$ over all y.

In the reading of fuzzy K as the description logic fuzzy $\mathcal{ALC}$ [42] (with only one role for simplicity), the underlying fuzzy relation would be seen as a vague connection between individuals, such as a 'likes' relation between persons. In this reading, the formula

$$\Box(\text{soft-spoken} \wedge \Diamond\text{reasonable})$$

describes people who only like people who are soft-spoken and like some reasonable person, with all these terms understood in a vague sense.

As indicated previously, the first-order correspondence language for fuzzy modal logic in this sense is *Zadeh fuzzy first-order logic* over a single binary predicate R and a unary predicate p for every propositional atom p. Formulas ϕ, ψ of what we briefly term *fuzzy first-order logic* or *fuzzy FOL* are thus given by the grammar

$$\phi, \psi ::= c \mid p(x) \mid R(x, y) \mid x = y \mid \phi \ominus c \mid \neg\phi \mid \phi \wedge \psi \mid \exists x.\phi$$

where $c \in [0, 1] \cap \mathbb{Q}$, $p \in \mathsf{At}$, and x, y range over a fixed countably infinite set of variables. We have the usual notions of free and bound variables. The *quantifier rank* $\mathrm{qr}(\phi)$ of a formula ϕ is defined, as usual, as the maximal nesting depth of quantifiers in ϕ (unlike for the modal rank, we do not let atomic formulas count towards the quantifier rank). The semantics is determined as the evident extension of the modal semantics, with the existential quantifier interpreted as supremum and '=' as crisp equality. Formally, a formula $\phi(x_1, \dots, x_n)$ with free variables among $x_1, \dots, x_n$ is interpreted, given a fuzzy relational model $\mathcal{A}$ and a vector $\bar{a} = (a_1, \dots, a_n)$ of values for the free variables, as a truth value $\phi(\bar{a}) \in [0, 1]$, given by

$$p(x_i)(\bar{a}) = p^{\mathcal{A}}(a_i) \qquad R(x_i, x_j)(\bar{a}) = R^{\mathcal{A}}(a_i, a_j)$$
$$(x_i = x_j)(\bar{a}) = 1 \text{ if } a_i = a_j, \text{ and } 0 \text{ otherwise}$$
$$(\exists x_0.\phi(x_0, \dots, x_n))(\bar{a}) = \textstyle\bigvee_{a_0 \in A} \phi(a_0, \bar{a})$$

and essentially the same clauses as in the modal case for the other connectives.

We thus have a variant of the classical *standard translation*, that is, a truth-value preserving embedding ST_x of fuzzy K into fuzzy FOL, indexed over a variable x naming the current state and defined inductively by $\mathsf{ST}_x(p) = p(x)$,

$$\mathsf{ST}_x(\Diamond\phi) = \exists y.\,(R(x, y) \wedge \mathsf{ST}_y(\phi)),$$

and commutation with all other constructs. Fuzzy K thus becomes a fragment of fuzzy FOL, and the object of the present paper is to characterize their relationship.

Coalgebraic view Recall that an F-*coalgebra* (A, α) for a set functor $F : \mathsf{Set} \to \mathsf{Set}$ consists of a set A of *states* and a map $\alpha : A \to FA$. The set FA is thought of as containing structured collections over A, so that α assigns to each state a a structured collection $\alpha(a)$ of successors. Coalgebras thus provide a general framework for state-based systems [36]. We will partly use coalgebraic techniques in our proofs, in particular final chain arguments. We therefore note that fuzzy relational models are coalgebras for the set functor G given by

$$\mathsf{G} = [0, 1]^{\mathsf{At}} \times \mathsf{F}$$

where $\mathsf{F}X = [0, 1]^X$ is the fuzzy version of the covariant powerset functor. That is, F acts on maps $f : X \to Y$ by taking fuzzy direct images,

$$\mathsf{F}f(g)(y) = \textstyle\bigvee_{f(x)=y} g(x).$$

Explicitly, a fuzzy relational model $\mathcal{A}$ corresponds to the G-coalgebra (A, α) given by $\alpha(a) = (h, g)$ where $h(p) = p^{\mathcal{A}}(a)$ for $p \in \mathsf{At}$ and $g(a') = R^{\mathcal{A}}(a, a')$ for $a' \in A$.

A G-*coalgebra morphism* $f : (A, \alpha) \to (B, \beta)$ between G-coalgebras $(A, \alpha), (B, \beta)$ (i.e. fuzzy relational models $\mathcal{A}, \mathcal{B}$) is a map $f : A \to B$ such that $\mathrm{G} f\alpha = \beta f$. Explicitly, this means that f is a *bounded morphism*, i.e. $p^{\mathcal{B}}(f(a)) = p^{\mathcal{A}}(a)$ for all atoms p and all $a \in A$, and $R^{\mathcal{B}}(f(a), b) = \bigvee_{f(a')=b} R^{\mathcal{A}}(a, a')$ for every $b \in B$. For models $\mathcal{A}$ and $\mathcal{B}$, we define their *disjoint union* $\mathcal{A} + \mathcal{B}$ as the model with domain $A + B$ (disjoint union of sets), $p^{\mathcal{A}+\mathcal{B}}(c) = p^{\mathcal{A}}(c)$ for $c \in A$ and $p^{\mathcal{A}+\mathcal{B}}(c) = p^{\mathcal{B}}(c)$ otherwise, and $R^{\mathcal{A}+\mathcal{B}}(c, c') = R^{\mathcal{A}}(c, c')$ if $c, c' \in A$, $R^{\mathcal{A}+\mathcal{B}}(c, c') = R^{\mathcal{B}}(c, c')$ if $c, c' \in B$, $R^{\mathcal{A}+\mathcal{B}}(c, c') = 0$ otherwise. This is precisely the categorical coproduct of $\mathcal{A}$ and $\mathcal{B}$ as G-coalgebras; in particular, the injection maps $\mathcal{A} \to \mathcal{A} + \mathcal{B}$ and $\mathcal{B} \to \mathcal{A} + \mathcal{B}$ are bounded morphisms.

3 Pseudometric Spaces

We recall some basics on pseudometric spaces, which differ from metric spaces in that distinct points can have distance 0:

Definition 3.1 (Pseudometric space, non-expansive maps)**.** Given a non-empty set X, a *(bounded) pseudometric on* X is a function $d : X \times X \to [0, 1]$ such that for all $x, y, z \in X$, the following axioms hold: $d(x, x) = 0$ (*reflexivity*), $d(x, y) = d(y, x)$ (*symmetry*), $d(x, z) \leq d(x, y) + d(y, z)$ (*triangle inequality*). If additionally $d(x, y) = 0$ implies $x = y$, then d is a *metric*. A *(pseudo)metric space* is a pair (X, d) where X is a set and d is a (pseudo)metric on X. The *diameter* of $A \subseteq X$ is $\bigvee_{x,y \in A} d(x, y)$. We equip the unit interval $[0, 1]$ with the standard Euclidean distance d_e,

$$d_e(x, y) = |x - y|.$$

A function $f : X \to Y$ between pseudometric spaces (X, d_1), (Y, d_2) is *non-expansive* if $d_2 \circ (f \times f) \leq d_1$, i.e. $d_2(f(x), f(y)) \leq d_1(x, y)$ for all x, y. We then write

$$f : (X, d_1) \to_1 (Y, d_2).$$

The space of non-expansive functions $(X, d_1) \to_1 (Y, d_2)$ is equipped with the *supremum (pseudo)metric* d_∞ defined by

$$d_\infty(f, g) = \sup_{x \in X} d_2(f(x), g(x))$$

In the special case $(Y, d_2) = ([0, 1], d_e)$, we will also denote $d_\infty(f, g)$ as $\|f - g\|_\infty$.

As usual we denote by $B_\epsilon(a) = \{x \in X \mid d(a, x) \leq \epsilon\}$ the *ball* of radius ϵ around a in (X, d). The space (X, d) is *totally bounded* if for every $\epsilon > 0$ there exists a finite ϵ-*cover*, i.e. finitely many elements $a_1, \dots, a_n \in X$ such that $X = \bigcup_{i=1}^{n} B_\epsilon(a_i)$.

Recall that a metric space is compact iff it is complete and totally bounded.

Given a fuzzy relational model $\mathcal{A}$, we extend the semantics of $\Diamond$ to arbitrary functions $f : A \to [0, 1]$ by

$$\Diamond f : A \to [0, 1], \quad (\Diamond f)(a) = \bigvee_{a' \in A} R^{\mathcal{A}}(a, a') \wedge f(a').$$

Lemma 3.1. *The map $f \mapsto \Diamond f$ is non-expansive.*

4 Behavioural Distance and Bisimulation Games

We proceed to define our notion of behavioural distance for fuzzy relational models. We opt for a game-based definition as the basic notion, and relate it to logical distance, showing that fuzzy modal logic is non-expansive w.r.t. behavioural distance. In Section 5 we will give an equivalent characterization in terms of fixed points, and show that all distances coincide at finite depth. Following ideas used in probabilistic bisimulation metrics [13], we use bisimulation games that have crisp outcomes but are parametrized over a maximal allowed deviation; we will then define the distance as the least parameter for which duplicator wins.

Definition 4.1. Let $\mathcal{A}, \mathcal{B}$ be fuzzy relational models, and let $a_0 \in A$, $b_0 \in B$. The *ϵ-bisimulation game for $\mathcal{A}, a_0$ and $\mathcal{B}, b_0$* (or just for a_0, b_0) played by S (*spoiler*) and D (*duplicator*) is given as follows.

- *Configurations:* pairs $(a, b) \in A \times B$ of states.
- *Initial configuration:* (a_0, b_0).
- *Moves:* Player S needs to pick a new state in one of the models $\mathcal{A}$ or $\mathcal{B}$, say $a' \in A$, such that $R^{\mathcal{A}}(a, a') > \epsilon$, and then D needs to pick a state in the other model, say $b' \in B$, such that $R^{\mathcal{B}}(b, b') \geq R^{\mathcal{A}}(a, a') - \epsilon$. The new configuration is then (a', b').
- *Winning condition:* Any player who needs to move but cannot, loses. Player D additionally needs to maintain the following winning condition *before* every round: For every $p \in \mathrm{At}$, $|p^{\mathcal{A}}(a) - p^{\mathcal{B}}(b)| \leq \epsilon$.

There are two variants of the game, the unrestricted game in which D wins infinite plays, and the *depth-n ϵ-bisimulation game*, which is restricted to n rounds, meaning that D wins after n rounds have been played.

Remark 4.2. Note that, since the invariant only needs to hold before every round actually played, D always wins the depth-0 game regardless of a_0 and b_0.

The usual composition lemma for bisimulations then takes the following form:

Lemma 4.3. *Let $\mathcal{A}, \mathcal{B}, C$ be models and $a_0 \in A$, $b_0 \in B$, $c_0 \in C$ such that D wins the ϵ-bisimulation game for (a_0, b_0) and the δ-bisimulation game for (b_0, c_0). Then D also wins the $(\epsilon + \delta)$-bisimulation game for (a_0, c_0). The same holds for the corresponding depth-n bisimulation games.*

As indicated above, we then obtain a notion of behavioural distance by taking infima:

Definition 4.4 (Behavioural distance)**.** Let $\mathcal{A}, a_0$ and $\mathcal{B}, b_0$ be as in Definition 4.1. The *behavioural distance* $d^G(a_0, b_0)$ of a_0 and b_0 is the infimum over all ϵ such that D wins the ϵ-bisimulation game for a_0 and b_0. The *depth-n behavioural distance* $d_n^G(a_0, b_0)$ of a_0 and b_0 is defined analogously, using the depth-n bisimulation game.

This definition is justified by the following lemma, which follows from Lemma 4.3:

Lemma 4.5. *The behavioural distance d^G and all depth-n behavioural distances d^G_n are pseudometrics.*

Remark 4.6. We emphasize that $d^G(a,b) = 0$ does not in general imply that D wins the 0-bisimulation game on a, b. In this sense, the notion of ϵ-bisimulation is thus what enables us to avoid restricting to models that are *witnessed* [21] in the sense that all suprema appearing in the evaluation of existential quantifiers are actually maxima.

We have the expected relationship between the various behavioural pseudometrics:

Lemma 4.7. *For all models $\mathcal{A}, \mathcal{B}$, states $a \in A$, $b \in B$, and $n \geq m \geq 0$, we have*

$$d^G_m(a,b) \leq d^G_n(a,b) \leq d^G(a,b).$$

As usual, behavioural equivalence is invariant under coalgebra morphisms; this can now be phrased as follows:

Lemma 4.8. *Let $\mathcal{A}, \mathcal{B}$ be fuzzy relational models, and let $f : \mathcal{A} \to \mathcal{B}$ be a bounded morphism. Then for every $a \in A$, $d^G(a, f(a)) = 0$.*

Proof (sketch). Player D wins the depth-n ϵ-bisimulation game for every $\epsilon > 0$. □

Since coproduct injections are bounded morphisms, a special case is

Lemma 4.9. *Given models $\mathcal{A}, \mathcal{B}$ and $a \in A$, the state a in $\mathcal{A}$ and the corresponding state a in $\mathcal{A} + \mathcal{B}$ have behavioural distance 0.*

Behavioural distance determines our notion of bisimulation invariance, which we take to mean non-expansivity w.r.t. behavioural distance. To match this with the standard notion, interpret classical crisp bisimilarity as a discrete pseudometric d assigning distance 0 to pairs of bisimilar states and 1 to non-bisimilar ones, and similarly interpret crisp predicates P as maps into $\{0, 1\}$; then P is bisimulation-invariant in the usual sense iff P is non-expansive w.r.t. d. Formal definitions for the fuzzy setting are as follows.

Definition 4.10 (Bisimulation-invariant formulas and predicates)**.** A formula ϕ (either in fuzzy modal logic or in fuzzy FOL, with a single free variable) is *bisimulation-invariant* if for all models $\mathcal{A}, \mathcal{B}$ and all states $a \in A$, $b \in B$,

$$|\phi(a) - \phi(b)| \leq d^G(a,b).$$

Similarly, given a model $\mathcal{A}$, a *(fuzzy) state predicate* on A, i.e. a function $P\colon A \to [0,1]$, is *bisimulation-invariant* if P is non-expansive w.r.t. the bisimulation distance d^G. In both cases, *depth-n bisimulation invariance* is defined in the same way using depth-n behavioural distance.

As expected, Zadeh fuzzy modal logic is bisimulation invariant; more precisely:

Lemma 4.11 (Bisimulation invariance)**.** *Every fuzzy modal formula of rank at most n is depth-n bisimulation-invariant.*

In particular, for every rank-n modal formula ϕ and every fuzzy relational model $\mathcal{A}$, the evaluation map $\phi_{\mathcal{A}} : A \to [0,1]$ is a non-expansive map $(A, d^G_n) \to_1 ([0,1], d_e)$. A forteriori (Lemma 4.7), *every fuzzy modal formula ϕ is bisimulation-invariant*, i.e. $\phi(\cdot)$ is non-expansive w.r.t. (unbounded-depth) behavioural distance d^G.

Example 4.12. The formula $R(x,x)$ in fuzzy FOL fails to be bisimulation-invariant (compare a loop with an infinite chain), and is therefore neither expressible nor approximable by fuzzy modal formulas.

Definition 4.13 (Logical distance)**.** We further define *logical distances* d^L (w.r.t. all modal formulas) and d^L_n (w.r.t. modal formulas of rank at most n) by

$$d^L(a,b) = \bigvee_{\phi \text{ modal}} |\phi(a) - \phi(b)|,$$
$$d^L_n(a,b) = \bigvee_{\mathrm{rk}(\phi)\leq n} |\phi(a) - \phi(b)|.$$

We clearly have

$$d^L_m(a,b) \leq d^L_n(a,b) \leq d^L(a,b) \quad \text{for } n \geq m \geq 0,$$

as well as

$$d^L(a,b) = \bigvee_{n\geq 0} d^L_n(a,b). \tag{1}$$

Using (1) and Lemma 4.7, we can then rephrase bisimulation invariance (Lemma 4.11) as

Lemma 4.14. *For models $\mathcal{A}, \mathcal{B}$, states $a \in A$, $b \in B$, and $n \geq 0$, we have*

$$d^L_n(a,b) \leq d^G_n(a,b) \quad \textit{and} \quad d^L(a,b) \leq d^G(a,b).$$

Remark 4.15. Under Łukasiewicz semantics (Remark 2.1), non-expansivity clearly breaks; e.g. if a and b are states without successors in models $\mathcal{A}$ and $\mathcal{B}$, respectively, such that $p^{\mathcal{A}}(a) = 0.9$, $p^{\mathcal{A}}(b) = 0.8$, and a and b agree on all other atoms, then $d^G(a,b) = 0.1$ but $|\phi(a) - \phi(b)| = 0.2$ for the formula $\phi = p \wedge p$, since under Łukasiewicz semantics, $\phi(a) = p^{\mathcal{A}}(a) + p^{\mathcal{A}}(a) - 1 = 0.8$ and $\phi(b) = p^{\mathcal{B}}(b) + p^{\mathcal{B}}(b) - 1 = 0.6$. See also a similar example in [17]. For a treatment of Łukasiewicz fuzzy modal logic, one would thus need to replace non-expansivity with Lipschitz continuity (see also [39]). Additional problems, however, arise with logical distance: Defining a logical distance for Łukasiewicz modal logic in analogy to the above definition of d^L gives a discrete pseudometric. The reason is that small behavioural differences between models can be amplified arbitrarily in Łukasiewicz logic using conjunction, as illustrated precisely by the above example (where we could also use $p \wedge p \wedge p$ etc.). The statement of a van Benthem theorem for Łukasiewicz modal logic would thus presumably become quite complicated, e.g. would need to stratify over Lipschitz constants.

We now launch into the proof of our target result, which states that *every bisimulation-invariant fuzzy first-order property can be approximated by fuzzy modal formulas*, a converse to bisimulation-invariance of fuzzy modal formulas. As already indicated, we follow a proof strategy established for the classical setting by Otto [33]: We show that

- every bisimulation-invariant fuzzy first-order property is ℓ*-local* for some ℓ, w.r.t. a suitable notion of Gaifman distance (Section 7);
- every ℓ-local bisimulation-invariant fuzzy first-order property is already depth-n bisimulation-invariant for some n (Section 8); and
- every depth-n bisimulation-invariant fuzzy state property is approximable by fuzzy modal formulas of rank at most n (Section 5).

We begin with the last step of this program.

5 Modal Approximation at Finite Depth

Having seen game-based and logical behavioural distances d^L, d^G in the previous section, we proceed to introduce a third, fixed-point based definition, and then show that all three distances agree at finite depth. This happens in a large simultaneous induction, in which we also prove that *every depth-n bisimulation-invariant fuzzy state property is approximable by modal formulas of rank n*. As indicated in the introduction, this is in fact the technical core of the paper. This is in sharp contrast with the classical setting, where the corresponding statement – every depth-n bisimulation-invariant crisp state property is expressible by a crisp modal formula – is completely straightforward.

Assumption 5.1. As usual in proofs of van Benthem type results, we assume from now on that $\mathsf{At} = \{p_1, \dots, p_k\}$ is finite. This is w.l.o.g. for purposes of the proof of our main result, as we will aim to show modal approximability of a given formula, so only finitely many atoms are relevant. Note that, e.g., Theorem 5.3.2 (total boundedness of finite-depth behavioural distance) will presumably not hold without this assumption.

The fixed-point definition of behavioural distance is based on the *Kantorovich lifting* [4]. We first define an *evaluation function* $ev\colon \mathsf{F}[0,1] \to [0,1]$ (recall from Section 2 that $\mathsf{F}X = [0,1]^X$ and $\mathsf{G}X = [0,1]^{\mathsf{At}} \times \mathsf{F}X$) by

$$ev(g) = \textstyle\bigvee_{s\in[0,1]}(g(s) \wedge s) \quad \text{for } g\colon [0,1] \to [0,1].$$

Given a pseudometric space (X, d), we define the *Kantorovich pseudometrics* d^{F} on $\mathsf{F}X$ and d^{G} on $\mathsf{G}X$, respectively, by

$$d^{\mathsf{F}}(g_1, g_2) = \bigvee_{f:(X,d)\to_1([0,1],d_e)} |ev \circ \mathsf{F}f(g_1) - ev \circ \mathsf{F}f(g_2)|$$

$$d^{\mathsf{G}}((r_1, g_1), (r_2, g_2)) = d^{\mathsf{F}}(g_1, g_2) \vee \textstyle\bigvee_{p\in\mathsf{At}} |r_1(p) - r_2(p)|$$

for $r_i \in [0,1]^{\mathsf{At}}$ and $g_i \in \mathsf{F}X$. It follows from general results on lifting metrics along functors [4] that d^{F} and d^{G} are indeed pseudometrics.

Given a fuzzy relational model $\mathcal{A} = (A, (p^{\mathcal{A}})_{p\in\mathsf{At}}, R^{\mathcal{A}})$, viewed as a coalgebra $\alpha\colon A \to \mathsf{G}A$ as discussed in Section 2, we can inductively define a sequence of pseudometrics $(d_n^K)_{n\geq 0}$ on A via the Kantorovich lifting:

$$d_0^K(a,b) = 0 \quad \text{and} \quad d_{n+1}^K = (d_n^K)^{\mathsf{G}} \circ (\alpha \times \alpha),$$

that is, expanding definitions,

$$d_{n+1}^K(a,b) = \bigvee_{p\in\mathsf{At}} |p(a) - p(b)| \vee \bigvee_{f:(A,d_n^K)\to_1([0,1],d_e)} |(\Diamond f)(a) - (\Diamond f)(b)|.$$

(The d_n^K can be seen as approximants of a fixed point, which is not itself needed here.)

Example 5.2. To illustrate the three forms of behavioural distance (logical, game-based, and via the Kantorovich lifting), we use the following model $\mathcal{A}$ with one propositional atom p. In the diagram below, each state has the form $x[p : p^{\mathcal{A}}(x)]$, and transitions from x to y are labelled with their truth values $R^{\mathcal{A}}(x,y)$.

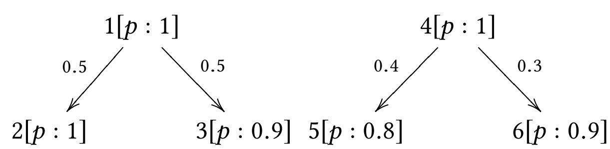

Clearly, it suffices to look at depth 2 in this example. The game-based distance of 1, 4 is $d^G(1,4) = d_2^G(1,4) = 0.2$. To see this, first note that D has a winning strategy for $\epsilon = 0.2$: Player S may pick any transition, and D then always has a transition available as a reply; irrespective of their choices, they end up in a pair of states with values of p differing by at most 0.2, and then S needs to move but cannot. The situation is different for $\epsilon < 0.2$: In this case, S can take the transition from 1 to 2, which D must answer by going from 4 to 5, since $R^{\mathcal{A}}(4,6) \ngeq R^{\mathcal{A}}(1,2) - \epsilon$. But $|p^{\mathcal{A}}(2) - p^{\mathcal{A}}(5)| = 0.2 > \epsilon$, so S wins.

This distance is witnessed by the formula $\phi = \Diamond(p \ominus 0.5)$:

$$\phi(1) = (0.5 \wedge (p^{\mathcal{A}}(2) \ominus 0.5)) \vee (0.5 \wedge (p^{\mathcal{A}}(3) \ominus 0.5)) = 0.5$$

$$\phi(4) = (0.4 \wedge (p^{\mathcal{A}}(5) \ominus 0.5)) \vee (0.3 \wedge (p^{\mathcal{A}}(6) \ominus 0.5)) = 0.3,$$

so $d^L(1,4) = d_2^L(1,4) = 0.2$ (recall $\mathsf{rk}(\phi) = 2$) by Lemma 4.14. Note that $\Diamond p$ would only yield a difference of 0.1.

As to the Kantorovich distances, we have $d_0^K(1,4) = 0$, so that the f over which the supremum in the definition of $d_1^K(1,4)$ is taken are all constant; it is then easily seen that $d_1^K(1,4) = 0.1$, the difference between the maximal transition degree from 1 (0.5) and that from 4 (0.4). The function corresponding to $p \ominus 0.5$ then serves as a witness of the behavioural distance at depth 2, so that $d_2^K(1,4) \geq 0.2$; one can check that in fact $d_2^K(1,4) = 0.2$.

The main result proved in this section is then the following theorem, which as indicated above states in particular that the definitions of behavioural distance coincide at finite depth and that the modal formulas lie dense in the non-expansive state properties:

Theorem 5.3. *Let $\mathcal{A}$ be a fuzzy relational model. Then the following holds for all $n \geq 0$.*

1. $d_n^G = d_n^L = d_n^K =: d_n$ *on* $\mathcal{A}$.
2. *The pseudometric space* (A, d_n) *is totally bounded.*
3. *The modal formulas of rank at most n form a dense subset of the space* $(A, d_n) \rightarrow_1 ([0,1], d_e)$.

(We note that the equality $d_n^L = d_n^G$ is effectively the finite-depth part of a Hennessy-Milner property; the infinite-depth version will, of course, hold only under finite branching. This contrasts somewhat with the probabilistic case [48].)

Proof (sketch). We prove all claims simultaneously by induction on n. The base case $n = 0$ is trivial. The proof of the induction step is split over a number of lemmas proved next:

- Item 1 is proved in Lemmas 5.4 and 5.5.
- Item 2 is proved in Lemma 5.7.
- Item 3 is proved in Lemma 5.9. □

For the remainder of this section, we fix a model $\mathcal{A}$ as in Theorem 5.3 and $n > 0$, and assume as inductive hypothesis that all claims in Theorem 5.3 already hold for all $n' < n$.

Lemma 5.4. *We have* $d_n^L = d_n^K$ *on* $\mathcal{A}$.

Proof (sketch). Let $a, b \in A$ and put $F := (A, d_{n-1}) \rightarrow_1 ([0,1], d_e)$. By Lemma 3.1, the map

$$H\colon (F, d_\infty) \rightarrow ([0,1], d_e), f \mapsto |(\Diamond f)(a) - (\Diamond f)(b)|$$

is continuous. Since by the induction hypothesis, $\mathcal{L}_{n-1}$ is dense in F, it follows that $H[\mathcal{L}_{n-1}]$ is dense in $H[F]$. Thus,

$$\begin{aligned} d_n^K(a,b) &= \bigvee_{p \in \mathrm{At}} |p(a) - p(b)| \vee \bigvee_{f:(A, d_{n-1}) \rightarrow_1 ([0,1], d_e)} |(\Diamond f)(a) - (\Diamond f)(b)| \\ &= \bigvee_{p \in \mathrm{At}} |p(a) - p(b)| \vee \bigvee_{\mathrm{rk}\phi \leq n-1} |(\Diamond \phi)(a) - (\Diamond \phi)(b)| \\ &= \bigvee_{\mathrm{rk}\phi \leq n} |\phi(a) - \phi(b)| = d_n^L(a,b). \end{aligned}$$ □

Lemma 5.5. *We have* $d_n^G = d_n^K$ *on* $\mathcal{A}$.

Proof (sketch). First, $d_n^K(a,b) = d_n^L(a,b) \leq d_n^G(a,b)$ for all a, b by Lemmas 5.4 and 4.11.

In the other direction, if $d_n^K(a,b) \leq \epsilon$, we need to show that D wins the $(\epsilon + \delta)$-game on (a,b) for all $\delta > 0$. W.l.o.g. S moves from a to some a'. We can instantiate the function f in the definition of d_n^K as

$$f\colon (A, d_{n-1}) \rightarrow_1 ([0,1], d_e), f(b') = R(a, a') \ominus d_{n-1}(a', b').$$

A winning reply for D can now be extracted by taking a state b' that approximates the supremum in $(\Diamond f)(b)$ sufficiently closely. One checks that b' is a legal move and that $d_{n-1}(a', b') < \epsilon + \delta$, so D wins. □

Having shown that the pseudometrics d_n^L, d_n^G, d_n^K coincide, we will now use d_n to denote any of them, as indicated in Theorem 5.3.

The next lemma is a version of the Arzelà-Ascoli theorem for total boundedness instead of compactness and non-expansive instead of continuous functions; that is, we impose weaker assumptions on the space but stronger assumptions on the functions.

Lemma 5.6. *Let* $(X, d_1), (Y, d_2)$ *be totally bounded pseudometric spaces. Then the space* $(X, d_1) \rightarrow_1 (Y, d_2)$*, equipped with the supremum pseudometric, is totally bounded.*

The following lemma, the inductive step for Theorem 5.3.2, then guarantees that our variant of Arzelà-Ascoli will actually apply to (A, d_n) in the next round of the induction.

Lemma 5.7. (A, d_n) *is a totally bounded pseudometric space.*

Proof (sketch). Put $F := (A, d_{n-1}) \rightarrow_1 ([0,1], d_e)$ and let $\epsilon > 0$. By Lemma 5.6, F is totally bounded, so as $\mathcal{L}_{n-1}$ is dense in F, there exists a finite $\frac{\epsilon}{6}$-cover of F consisting of formulas $\phi_1, \dots, \phi_m \in \mathcal{L}_{n-1}$. One can now show that the map

$$\begin{aligned} I\colon A &\rightarrow [0,1]^{k+m} \\ a &\mapsto (p_1(a), \dots, p_k(a), (\Diamond \phi_1)(a), \dots, (\Diamond \phi_m)(a)) \end{aligned}$$

is an $\frac{\epsilon}{3}$-*isometry*, i.e. for all $a, b \in A$,

$$|d_n(a,b) - \|I(a) - I(b)\|_\infty| \leq \tfrac{\epsilon}{3}.$$

Using pre-images under I and a simple triangle inequality argument, we can then convert a finite $\frac{\epsilon}{3}$-cover of the compact space $([0,1]^{k+m}, d_\infty)$ into a finite ϵ-cover of (A, d_n). □

We next prove a variant of the lattice version of the Stone-Weierstraß theorem (e.g. [3, Lemma A.7.2]). Again, we only assume the space to be totally bounded instead of compact but require functions to be non-expansive rather than only continuous. (A Stone-Weierstraß argument appears also in a probabilistic Hennessy-Milner result [48]).

Lemma 5.8. *Let* (X, d) *be a totally bounded pseudometric space, and let L be a subset of* $F := (X, d) \rightarrow_1 ([0,1], d_e)$ *such that* $f_1, f_2 \in L$ *implies* $\min(f_1, f_2), \max(f_1, f_2) \in L$. *If each* $f \in F$ *can be approximated at each pair of points by functions in L, then L is dense in F.*

Lemma 5.9. *The modal formulas of rank at most n form a dense subset of the space* $(A, d_n) \rightarrow_1 ([0,1], d_e)$.

Proof (sketch). We proceed as in [48], applying Lemma 5.8 to $\mathcal{L}_n$:

Given a function $f\colon (A, d_n) \to_1 ([0,1], d_e)$ and points $a, b \in A$, a formula ϕ approximating f at a and b can be constructed as follows: Let $\psi \in \mathcal{L}_n$ be such that $|\psi(a) - \psi(b)|$ approximates $|f(a) - f(b)|$ (such a ψ exists by non-expansivity of f). Then ϕ is defined from ψ by means of modified subtraction $\ominus$, which preserves the rank of formulas. □

This concludes the proof of Theorem 5.3. The theorem still leaves one loose end: The modal formulas that approximate a given depth-n bisimulation-invariant state property on a model $\mathcal{A}$ might depend on $\mathcal{A}$. We eliminate this dependency in the next section, using the final chain construction.

6 The Final Chain

The *final chain* [2, 6] of the functor G is a sequence of sets F_k that represent all the possible depth-k behaviours. It is constructed as follows. We take F_0 to be a singleton $F_0 = \{*\}$ (reflecting that all states are equivalent at depth 0), and

$$F_{n+1} = \mathsf{G}F_n = [0,1]^{\mathsf{At}} \times \mathsf{F}F_n.$$

Given a model $\mathcal{A}$, seen as a coalgebra $\alpha\colon A \to \mathsf{G}A$, we can now define a sequence of projections $\pi_n : A \to F_n$, to be thought of as mapping states to their depth-n behaviours, by

$$\pi_0 = ! \quad \text{and} \quad \pi_{n+1} = \mathsf{G}\pi_n \circ \alpha,$$

where ! denotes the unique map $A \to F_0$. Explicitly, π_{n+1} is thus defined by

$$\pi_{n+1}(a) = (\lambda p. p^{\mathcal{A}}(a), \lambda y. \textstyle\bigvee_{\pi_n(a')=y} R^{\mathcal{A}}(a, a')). \qquad (2)$$

We next build a model $\mathcal{F}$ realizing all finite-depth behaviours by taking the union $F = \bigcup_{k\in\mathbb{N}} F_k$ (automatically disjoint). We define the model structure on F by letting every element behave as it claims to: For $(h, g) \in F_{k+1} = [0,1]^{\mathsf{At}} \times \mathsf{F}F_k$ and $y \in F$, we put $p^{\mathcal{F}}(h, g) = h(p)$ for $p \in \mathsf{At}$ and

$$R^{\mathcal{F}}((h, g), y) = g(y) \quad \text{if } y \in F_k,$$

and $R^{\mathcal{F}}((h, g), y) = 0$ otherwise. For $* \in F_0$, we just put $p^{\mathcal{F}}(*) = R^{\mathcal{F}}(*, y) = 0$.

In the proof of our main result (Theorem 8.4), the following lemma will allow us to choose approximating modal formulas uniformly across models.

Lemma 6.1. *Let $\mathcal{A}$ be a model. Then $d_n^G(a, \pi_n(a)) = 0$ for all $a \in A$.*

Proof (sketch). Player D wins the depth-n ϵ-bisimulation game for every $\epsilon > 0$ by maintaining the invariant that in round i, the configuration has the form $(a', \pi_{n-i}(a'))$ for some $a' \in A$. One sees from (2) that this invariant implies the winning condition and can actually be maintained by D. □

7 Locality

We proceed to show that every bisimulation-invariant formula of fuzzy FOL is *local*. To this end, we introduce a notion of Gaifman distance in fuzzy models, as well as a variant of Ehrenfeucht-Fraïssé games. The requisite notions of Gaifman graph and neighbourhood, as well as the definition of locality, are, maybe unexpectedly, fairly crisp. This is technically owed to the fact that unlike continuity, non-expansivity does not go well with chains of ϵ-estimates.

Definition 7.1. Let $\mathcal{A}$ be a fuzzy relational model.

- The *Gaifman graph of* $\mathcal{A}$ is an undirected graph with set A of nodes and an edge $\{a, b\}$ for every $a, b \in A$ such that $R(a, b) > 0$ or $R(a, b) > 0$.
- For every $a, b \in A$, the *Gaifman distance* $D(a, b) \in \mathbb{N} \cup \{\infty\}$ is the minimal length (i.e. number of edges) of a path between a and b in the Gaifman graph.
- For $a \in A$ and $\ell \in \mathbb{N}$, the *neighbourhood* of a with radius ℓ is the set $U^\ell(a)$ given by

$$U^\ell(a) = \{b \in A \mid D(a, b) \le \ell\}.$$

For $\bar{a} = (a_1, \dots, a_n)$, we put $U^\ell(\bar{a}) = \bigcup_{i\le n} U^\ell(a_i)$.

Definition 7.2. Let $\mathcal{A}$ be a fuzzy relational model and $U \subseteq A$. The *restriction* $\mathcal{A}|_U$ of $\mathcal{A}$ to U is the fuzzy relational model $(U, (p^{\mathcal{A}|_U})_{p\in\mathsf{At}}, R^{\mathcal{A}|_U})$ with $p^{\mathcal{A}|_U}(a) = p^{\mathcal{A}}(a)$ and $R^{\mathcal{A}|_U}(a, b) = R^{\mathcal{A}}(a, b)$ for $a, b \in U$. If $U = U^\ell(\bar{a})$ for some vector $\bar{a}$ over A, we also write $\mathcal{A}^\ell_{\bar{a}} := \mathcal{A}|_{U^\ell(\bar{a})}$.

As indicated, the ensuing notion of locality is on-the-nose:

Definition 7.3. A formula ϕ is *ℓ-local* for $\ell \in \mathbb{N}$ if

$$\phi_{\mathcal{A}}(a) = \phi_{\mathcal{A}^\ell_a}(a)$$

for every fuzzy relational model $\mathcal{A}$ and every $a \in A$.

It is easy to see that depth-k behaviour depends only on k-neighbourhoods, i.e.

Lemma 7.4. *For any model $\mathcal{A}$, $a_0 \in A$, and $k > 0$, D wins the depth-k 0-bisimulation game for $\mathcal{A}, a_0$ and $\mathcal{A}^k_{a_0}, a_0$.*

In combination with Lemma 4.11, we obtain

Corollary 7.5. *Every fuzzy modal formula of rank at most k is k-local.*

To establish the desired locality result, we employ Ehrenfeucht-Fraïssé games, introduced next. We phrase Ehrenfeucht-Fraïssé equivalence in terms of a pseudometric, in line with our treatment of behavioural distance, as this is the right way of measuring equivalence w.r.t. fuzzy FOL; in the further technical development, we will actually need only the case with deviation $\epsilon = 0$.

Definition 7.6. Let $\mathcal{A}, \mathcal{B}$ be fuzzy relational models, and let $\bar{a}_0$ and $\bar{b}_0$ be vectors of equal length over A and B respectively. The *ϵ-Ehrenfeucht-Fraïssé game for $\mathcal{A}, \bar{a}_0$ and $\mathcal{B}, \bar{b}_0$* played by S (*spoiler*) and D (*duplicator*) is given as follows.

- *Configurations:* pairs $(\bar{a}, \bar{b})$ of vectors $\bar{a}$ over A and $\bar{b}$ over B.
- *Initial configuration:* $(\bar{a}_0, \bar{b}_0)$.
- *Moves:* Player S needs to pick a new state in one of the models, say $a \in A$, and then D needs to pick a state in the other model, say $b \in B$. The new configuration is then $(\bar{a}a, \bar{b}b)$.
- *Winning condition:* Any player who needs to move but cannot, loses. Player D additionally needs to maintain the condition that $(\bar{a}, \bar{b})$ is a *partial isomorphism up to* ϵ: For all $0 \le i, j \le n$:
  - $a_i = a_j \iff b_i = b_j$
  - $|p^{\mathcal{A}}(a_i) - p^{\mathcal{B}}(b_i)| \le \epsilon$ for all $p \in \mathsf{At}$
  - $|R^{\mathcal{A}}(a_i, a_j) - R^{\mathcal{B}}(b_i, b_j)| \le \epsilon$.

Here, we need only the *n-round ϵ-Ehrenfeucht-Fraïssé game*, which as the name indicates is played for at most n rounds, and D wins after n rounds have been played.

In analogy to the classical setup, fuzzy FOL is invariant under Ehrenfeucht-Fraïssé equivalence in the sense that formula evaluation is non-expansive:

Lemma 7.7 (Ehrenfeucht-Fraïssé invariance)**.** *Let $\mathcal{A}, \mathcal{B}$ be fuzzy relational models and $\bar{a}_0, \bar{b}_0$ vectors of length m over A and B, respectively. If D wins the n-round ϵ-Ehrenfeucht-Fraïssé game on $\bar{a}_0, \bar{b}_0$, then for every first-order formula ϕ with at most m free variables and* $\mathrm{qr}(\phi) \le n$*,*

$$|\phi(\bar{a}_0) - \phi(\bar{b}_0)| \le \epsilon.$$

Lemma 7.8. *Let ϕ be a bisimulation-invariant formula of fuzzy FOL with quantifier rank* $\mathrm{qr}(\phi) \le n$*. Then ϕ is k-local for $k = 3^n$.*

Proof (sketch). Let $\mathcal{A}$ be a model, $a_0 \in A$. Define models $\mathcal{B}$ and C by extending both $\mathcal{A}$ and $\mathcal{A}^k_{a_0}$ by n disjoint copies of both $\mathcal{A}$ and $\mathcal{A}^k_{a_0}$ each. By Lemmas 4.9 and 7.7, it suffices to show that D wins the 0-Ehrenfeucht-Fraïssé game for $\mathcal{B}, a_0$ and C, a_0. Indeed, D wins by maintaining the following invariant, where we put $k_i = 3^{n-i}$ for $0 \le i \le n$:

> If $(\bar{b}, \bar{c}) = ((b_0, \dots, b_i), (c_0, \dots, c_i))$ is the current configuration, then there is an isomorphism between $\mathcal{B}^{k_i}_{\bar{b}}$ and $C^{k_i}_{\bar{c}}$ mapping each b_j to c_j. ◻

8 A Fuzzy van Benthem Theorem

It remains only to establish the implication from locality and bisimulation-invariance to finite-depth bisimulation invariance, using a standard unravelling construction, to finish the proof of our main result.

Definition 8.1. The *unravelling* $\mathcal{A}^*$ of a model $\mathcal{A}$ is the model with set A^+ (non-empty lists over A) of states and

$$p^{\mathcal{A}^*}(\bar{a}) = p^{\mathcal{A}}(\pi(\bar{a})),$$
$$R^{\mathcal{A}^*}(\bar{a}, \bar{a}a) = R^{\mathcal{A}}(\pi(\bar{a}), a),$$

for $\bar{a} \in A^+, a \in A$, where $\pi \colon A^+ \to A$ projects to the last element and all other values of $R^{\mathcal{A}^*}$ are 0.

Lemma 8.2. *For any model $\mathcal{A}$ and $a_0 \in A$, D wins the 0-bisimulation game for $\mathcal{A}, a_0$ and $\mathcal{A}^*, a_0$.*

The following lemma then completes the last step in our program as laid out in Section 4.

Lemma 8.3. *Let ϕ be bisimulation-invariant and k-local. Then ϕ is depth-$(k+1)$ bisimulation-invariant.*

Proof (sketch). Use locality and unravelling (Lemma 8.2) to reduce to tree models of depth k, and then exploit that in such models, winning the depth-$(k+1)$ ϵ-bisimulation game entails winning the unrestriced game. ◻

We finally state our main result:

Theorem 8.4 (Fuzzy van Benthem theorem)**.** *Let ϕ be a formula of fuzzy FOL with one free variable and* $\mathrm{qr}(\phi) = n$*. If ϕ is bisimulation-invariant, then ϕ can by approximated by fuzzy modal formulas of rank at most $3^n + 1$, uniformly over all models; that is: For every $\epsilon > 0$ there exists a fuzzy modal formula ϕ_ϵ such that for every fuzzy relational model $\mathcal{A}$ and every $a \in \mathcal{A}$, $|\phi(a) - \phi_\epsilon(a)| \le \epsilon$.*

Proof (sketch). By Lemmas 7.8 and 8.3, ϕ is depth-k bisimulation-invariant for $k = 3^n + 1$. By Theorem 5.3, ϕ can be modally approximated in rank k on the model $\mathcal{F}$ constructed from the final chain in Section 6. The claim then follows by Lemma 6.1. ◻

Remark 8.5. We leave the Rosen version of the characterization theorem, i.e. whether Theorem 8.4 remains true over finite models, as an open problem. As in the classical case, the unravelling construction is easily made to preserve finite models by using *partial unravelling* up to the locality distance. However, the model construction from the final chain in Section 6 and in fact already the stages of the final chain are infinite, so cannot be used in this version. We thus do obtain a local version of the Rosen theorem, stating that on a fixed finite model, every first-order formula that is bisimulation-invariant over finite models can be approximated by modal formulas. However, it is unclear whether the approximation then works uniformly over models, as in Theorem 8.4.

9 Conclusions

We have established a fuzzy analogue of the classical van Benthem theorem: Every fuzzy first-order formula that is bisimulation-invariant in the sense that its evaluation map is non-expansive w.r.t. a natural notion of behavioural distance can be approximated by fuzzy modal formulas. To our knowledge this is the first modal characterization result of this type for any multi-valued modal logic. We do point out that we leave a nagging open problem: We currently do not know whether the result can be sharpened to claim that every bisimulation-invariant fuzzy first-order formula is in fact *equivalent* to a fuzzy modal formula. This contrasts with the actual technical core of our argument: The key step in our proof is to show that *every* state property that is non-expansive w.r.t. *depth-n* behavioural distance can be approximated, uniformly across models, by fuzzy modal formulas of rank n, a result that certainly cannot be improved to on-the-nose modal expressibility.

Further issues for future research include the question whether our main result has a Rosen variant, i.e. holds also over finite models, and coverage of other semantics of the propositional operators, in particular Łukasiewicz logic. We also aim to extend the modal characterization theorem to further multi-valued logics, such as [0, 1]-valued probabilistic modal logics [48], ideally at a coalgebraic level of generality.

References

[1] S. Abriola, M. Descotte, and S. Figueira. Model theory of XPath on data trees. Part II: Binary bisimulation and definability. *Inf. Comput.*, 255:195–223, 2017.

[2] J. Adámek and V. Koubek. On the greatest fixed point of a set functor. *Theor. Comput. Sci.*, 150:57–75, 1995.

[3] R. Ash. *Real Analysis and Probability.* Academic Press, 1972.

[4] P. Baldan, F. Bonchi, H. Kerstan, and B. König. Behavioral metrics via functor lifting. In *Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2014*, vol. 29 of *LIPIcs*, pp. 403–415. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2014.

[5] B. Balle, P. Gourdeau, and P. Panangaden. Bisimulation metrics for weighted automata. In *Automata, Languages, and Programming, ICALP 2017*, vol. 80 of *LIPIcs*, pp. 103:1–14. Schloss Dagstuhl – Leibniz Center for Informatics, 2017.

[6] M. Barr. Terminal coalgebras in well-founded set theory. *Theor. Comput. Sci.*, 114:299–315, 1993.

[7] G. Bruns and P. Godefroid. Model checking with multi-valued logics. In *Automata, Languages and Programming, ICALP 2004*, vol. 3142 of *LNCS*, pp. 281–293. Springer, 2004.

[8] Y. Cao, S. Sun, H. Wang, and G. Chen. A behavioural distance for fuzzy-transition systems. *IEEE Trans. Fuzzy Sys.*, 21:735–747, 2013.

[9] F. Carreiro. PDL is the bisimulation-invariant fragment of weak chain logic. In *Logic in Computer Science, LICS 2015*, pp. 341–352. IEEE, 2015.

[10] C. Chang and H. Keisler. *Continuous model theory.* Princeton University Press, 1966.

[11] A. Dawar and M. Otto. Modal characterisation theorems over special classes of frames. In *Logic in Computer Science, LICS 2005*, pp. 21–30. IEEE, 2005.

[12] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for labelled Markov processes. *Theor. Comput. Sci.*, 318:323–354, 2004.

[13] J. Desharnais, F. Laviolette, and M. Tracol. Approximate analysis of probabilistic processes: Logic, simulation and games. In *Quantitative Evaluation of Systems, (QEST 2008)*, pp. 264–273. IEEE, 2008.

[14] P. Eleftheriou, C. Koutras, and C. Nomikos. Notions of bisimulation for Heyting-valued modal languages. *J. Log. Comput.*, 22:213–235, 2012.

[15] S. Enqvist, F. Seifan, and Y. Venema. Monadic second-order logic and bisimulation invariance for coalgebras. In *Logic in Computer Science, LICS 2015*. IEEE, 2015.

[16] T. Fan. Fuzzy bisimulation for Gödel modal logic. *IEEE Trans. Fuzzy Sys.*, 23:2387–2396, Dec 2015.

[17] T. Fan and C. Liau. Logical characterizations of regular equivalence in weighted social networks. *Artif. Intell.*, 214:66–88, 2014.

[18] D. Figueira, S. Figueira, and C. Areces. Model theory of XPath on data trees. Part I: Bisimulation and characterization. *J. Artif. Intell. Res. (JAIR)*, 53:271–314, 2015.

[19] M. Fitting. Many-valued modal logics. *Fund. Inform.*, 15:235–254, 1991.

[20] P. Hájek. *Metamathematics of fuzzy logic.* Springer, 1998.

[21] P. Hájek. Making fuzzy description logic more general. *Fuzzy Sets Sys.*, 154:1–15, 2005.

[22] H. Hansen, C. Kupke, and E. Pacuit. Neighbourhood structures: Bisimilarity and basic model theory. *Log. Meth. Comput. Sci.*, 5(2), 2009.

[23] D. Janin and I. Walukiewicz. Automata for the modal μ-calculus and related results. In *Mathematical Foundations of Computer Science, MFCS 1995*, vol. 969 of *LNCS*, pp. 552–562. Springer, 1995.

[24] J. Kontinen, J. Müller, H. Schnoor, and H. Vollmer. A van Benthem theorem for modal team semantics. In *Computer Science Logic, CSL 2015*, vol. 41 of *LIPIcs*, pp. 277–291. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.

[25] S. Kundu and J. Chen. Fuzzy logic or Łukasiewicz logic: A clarification. *Fuzzy Sets Sys.*, 95:369 – 379, 1998.

[26] O. Kupferman and Y. Lustig. Latticed simulation relations and games. In *Automated Technology for Verification and Analysis, ATVA 2007*, vol. 4762 of *LNCS*, pp. 316–330. Springer, 2007.

[27] J. Łukasiewicz and A. Tarski. Untersuchungen über den Aussagenkalkül. *Comptes rendus des séances de la Société des Sciences et des Lettres de Varsovie, Classe III*, 23:1–21, 1930.

[28] T. Łukasiewicz and U. Straccia. Managing uncertainty and vagueness in description logics for the semantic web. *J. Web Sem.*, 6:291–308, 2008.

[29] C. Morgan. Local and global operators and many-valued modal logics. *Notre Dame J. Formal Log.*, 20:401–411, 1979.

[30] O. Morikawa. Some modal logics based on a three-valued logic. *Notre Dame J. Formal Log.*, 30:130–137, 1988.

[31] V. Novák. First-order fuzzy logic. *Stud. Log.*, 46:87–109, 1987.

[32] P. Ostermann. Many-valued modal propositional calculi. *Math. Log. Q.*, 34:343–354, 1988.

[33] M. Otto. Elementary proof of the van Benthem-Rosen characterisation theorem. Technical Report 2342, Dept. Math., TU Darmstadt, 2004.

[34] E. Rosen. Modal logic over finite structures. *J. Log. Lang. Inf.*, 6:427–439, 1997.

[35] J. Rosser and A. Turquette. *Many-Valued Logics.* North Holland, 1952.

[36] J. Rutten. Universal coalgebra: A theory of systems. *Theor. Comput. Sci.*, 249:3–80, 2000.

[37] D. Sánchez and A. Tettamanzi. Reasoning and quantification in fuzzy description logics. In *Fuzzy Logic and Applications, WILF 2005*, vol. 3849 of *LNCS*, pp. 81–88. Springer, 2006.

[38] P. Schotch. Fuzzy modal logic. In *International Symposium on Multiple-Valued Logic, ISMV 1975*, pp. 176–181. IEEE, 1976.

[39] L. Schröder and D. Pattinson. Description logics and fuzzy probability. In *Int. Joint Conf. Artificial Intelligence, IJCAI 2011*, pp. 1075–1081. AAAI, 2011.

[40] L. Schröder, D. Pattinson, and T. Litak. A van Benthem/Rosen theorem for coalgebraic predicate logic. *J. Log. Comput.*, 27:749–773, 2017.

[41] K. Segerberg. Some modal logics based on a three-valued logic. *Theoria*, 33:53–71, 1967.

[42] U. Straccia. A fuzzy description logic. In *Artificial Intelligence, AAAI 1998*, pp. 594–599. AAAI Press / MIT Press, 1998.

[43] H. Sturm and F. Wolter. First-order expressivity for S5-models: Modal vs. two-sorted languages. *J. Philos. Log.*, 30:571–591, 2001.

[44] B. ten Cate, G. Fontaine, and T. Litak. Some modal aspects of XPath. *J. Appl. Non-Classical Log.*, 20:139–171, 2010.

[45] C. Tresp and R. Molitor. A description logic for vague knowledge. In *Europ. Conf. Artificial Intelligence, ECAI 1998*, pp. 361–365. Wiley, 1998.

[46] J. van Benthem. *Modal Correspondence Theory*. PhD thesis, Universiteit van Amsterdam, 1976.

[47] J. van Benthem. Correspondence theory. In D. Gabbay and F. Guenthner, eds., *Handbook of Philosophical Logic*, vol. 3, pp. 325–408. Springer, 2001.

[48] F. van Breugel and J. Worrell. A behavioural pseudometric for probabilistic transition systems. *Theor. Comput. Sci.*, 331:115–142, 2005.

[49] W. Wee and K. Fu. A formulation of fuzzy automata and its application as a model of learning systems. *IEEE Trans. Sys. Sci. Cyb.*, 5:215–223, 1969.

[50] P. Wild and L. Schröder. A characterization theorem for a modal description logic. In *Int. Joint Conf. Artificial Intelligence, IJCAI 2017*, pp. 1304–1310. ijcai.org, 2017.

[51] J. Yen. Generalizing term subsumption languages to fuzzy logic. In *Int. Joint Conf. on Artificial Intelligence, IJCAI 1991*, pp. 472–477. Morgan Kaufmann, 1991.

[52] L. Zadeh. Fuzzy sets. *Inf. Control*, 8:338–353, 1965.

A Omitted Proofs

Proof of Lemma 3.1

Let $\|f-g\|_\infty \leq \epsilon$. We need to show $\|\Diamond f - \Diamond g\|_\infty \leq \epsilon$, so let $a \in A$ and we need to show $|(\Diamond f)(a) - (\Diamond g)(a)| \leq \epsilon$.

Let $a' \in A$. Then $|f(a') - g(a')| \leq \epsilon$ by assumption. Thus, also $|R^{\mathcal{A}}(a,a') \wedge f(a') - R^{\mathcal{A}}(a,a') \wedge g(a')| \leq \epsilon$. Now we may take the supremum over all $a' \in A$ and obtain $|(\Diamond f)(a) - (\Diamond g)(a)| \leq \epsilon$, as desired. □

Proof of Lemma 4.3

A winning strategy for D consists in following the two existing winning strategies in parallel. In detail, D maintains the following invariant on configurations (a,c): There exists a state b such that (a,b) and (b,c) are winning positions for D in the ϵ-game on $\mathcal{A},\mathcal{B}$ and in the δ-game on $\mathcal{B},\mathcal{C}$, respectively. By the triangle inequality, (a,c) then satisfies the winning condition. It remains to show that D can maintain the invariant. Suppose that S makes a move from a to some a' (the case where S moves in $\mathcal{C}$ is entirely symmetric). Let b' be D's reply to a' in the first game. Then b' is a valid move for S in the second game, since $R(b,b') \geq R(a,a') - \epsilon > \epsilon + \delta - \epsilon = \delta$. So D has a reply c' to b' as a move by S in the second game. and c' is also a valid move for D in the $(\epsilon+\delta)$-game, since $R(c,c') \geq R(b,b') - \delta \geq R(a,a') - (\epsilon+\delta)$. Of course, the new configuration (a',c') then satisfies the invariant, as witnessed by b'. □

Proof of Lemma 4.5

Let $\mathcal{A}$ be a model. For $a \in A$, D wins the 0-bisimulation game on a, a by copying every move of S. Thus $d_n^G(a,a) = 0$. The symmetry of d^G follows from that of Definition 4.1. The triangle inequality follows from Lemma 4.3. □

Proof of Lemma 4.7

A winning strategy for D in the ϵ-bisimulation game wins also the depth-n ϵ-bisimulation game, showing that $d_n^G(a,b) \leq d^G(a,b)$; the other inequality is shown in the same way. □

Proof of Lemma 4.8

Let $\epsilon > 0$; we show that D wins the ϵ-bisimulation game on a and $f(a)$. The winning strategy is given by maintaining the invariant that configurations are of the form $(a', f(a'))$ with $a' \in A$. This invariant holds initially, and guarantees that D wins because f preserves truth values of propositional atoms. It remains to show that D can maintain the invariant. So suppose that from a configuration $(a', f(a'))$, S moves in $\mathcal{A}$ along an edge $R^{\mathcal{A}}(a',a'') > \epsilon$. By the definition of bounded morphisms, $R^{\mathcal{B}}(f(a'), f(a'')) \geq R^{\mathcal{A}}(a',a'')$, so D can answer with $f(a'')$, maintaining the invariant. The remaining case is that S moves within $\mathcal{B}$ to some state b such that $R^{\mathcal{B}}(f(a'), b) > \epsilon$. Since $R^{\mathcal{B}}(f(a'), b) = \bigvee_{f(a'')=b} R^{\mathcal{A}}(a',a'')$, there is $a'' \in A$ such that $f(a'') = b$

and $R^{\mathcal{A}}(a', a'') \geq R^{\mathcal{B}}(f(a'), b) - \epsilon$; then D can move to a'', maintaining the invariant. □

Full Proof of Lemma 4.11

Suppose that D wins the ϵ-bisimulation-game for (a, b). We show that $|\phi(a) - \phi(b)| \leq \epsilon$ for all $\phi \in \mathcal{L}_n$ by induction on ϕ. The case for $c \in \mathbb{Q}$ is trivial and the case for propositional atoms p follows from the winning condition.

The inductive cases $\phi \ominus c$, $\neg\phi$, and $\phi \wedge \psi$ are as follows:

$$\begin{aligned}
|(\phi \ominus c)(a) - (\phi \ominus c)(b)| &= |(\phi(a) - c) \wedge 0 - (\phi(b) - c) \wedge 0| \\
&\leq |\phi(a) - \phi(b)| \leq \epsilon \\
|(\neg\phi)(a) - (\neg\phi)(b)| &= |(1 - \phi(a)) - (1 - \phi(b))| \\
&= |\phi(b) - \phi(a)| \leq \epsilon \\
|(\phi \wedge \psi)(a) - (\phi \wedge \psi)(b)| &= |\phi(a) \wedge \psi(a) - \phi(b) \wedge \psi(b)| \\
&\leq |\phi(a) - \phi(b)| \vee |\psi(a) - \psi(b)| \\
&\leq \epsilon,
\end{aligned}$$

in each case using the inductive hypothesis in the last step.

Finally, we treat the case for the modality $\Diamond$: By symmetry, it suffices to show that $(\Diamond\phi)(b) \geq (\Diamond\phi)(a) - \epsilon$. If $(\Diamond\phi)(a) \leq \epsilon$, then this follows immediately, so assume $(\Diamond\phi)(a) > \epsilon$. Let $\delta > 0$; then there is a' such that $R(a, a') > \epsilon$ and $(\Diamond\phi)(a) - (R(a, a') \wedge \phi(a')) < \delta$. Let b' be D's winning answer to S's move a'. Then by induction, $|\phi(a') - \phi(b')| \leq \epsilon$, and moreover $R(b, b') \geq R(a, a') - \epsilon$ by the rules of the game. Thus,

$$\begin{aligned}
(\Diamond\phi)(b) \geq R(b, b') \wedge \phi(b') &\geq (R(a, a') - \epsilon) \wedge (\phi(a') - \epsilon) \\
&= (R(a, a') \wedge \phi(a')) - \epsilon > (\Diamond\phi)(a) - \epsilon - \delta.
\end{aligned}$$

Since this holds for every δ, it follows that $(\Diamond\phi))(b) \geq (\Diamond\phi))(a) - \epsilon$. □

Proof Details for Theorem 5.3

We elaborate details for the base case $n = 0$. For Item 1, we show that all distances are 0 for all pairs $(a, b) \in A \times A$. We have $d_0^G(a, b) = 0$ by Remark 4.2. To see that $d_0^L(a, b) = 0$, just note that the only modal formulas of rank 0 are the constants $c \in \mathbb{Q}$ and Boolean combinations thereof. Finally, $d_0^K(a, b) = 0$ by definition.

Item 2 then follows trivially from the fact that d_0 vanishes.

For Item 3, since d_0 is the zero pseudometric, the space $(A, d_0) \to_1 ([0, 1], d_e)$ is just the space of constant functions. The claim then follows, since the modal formulas correspond to the rational-valued constant functions, and $[0, 1] \cap \mathbb{Q}$ is a dense subset of $[0, 1]$. □

Full Proof of Lemma 5.4

Let $a, b \in A$, put $F := (A, d_{n-1}) \to_1 ([0, 1], d_e)$, and define the map

$$H : (F, d_\infty) \to ([0, 1], d_e), f \mapsto |(\Diamond f)(a) - (\Diamond f)(b)|.$$

This map is continuous because of Lemma 3.1, and using that function evaluation, subtraction, and taking the absolute value are continuous operations.

By the induction hypothesis, $\mathcal{L}_{n-1}$ is dense in F, so $H[\mathcal{L}_{n-1}]$ is also dense in $H[F]$, and $\bigvee H[\mathcal{L}_{n-1}] = \bigvee H[F]$. Now:

$$\begin{aligned}
d_n^K(a, b) &= \bigvee_{p \in \mathrm{At}} |p(a) - p(b)| \vee \bigvee_{f : (A, d_{n-1}) \to_1 ([0,1], d_e)} |(\Diamond f)(a) - (\Diamond f)(b)| \\
&= \bigvee_{p \in \mathrm{At}} |p(a) - p(b)| \vee \bigvee_{\mathrm{rk}\phi \leq n-1} |(\Diamond\phi)(a) - (\Diamond\phi)(b)| \\
&= \bigvee_{\mathrm{rk}\phi \leq n} |\phi(a) - \phi(b)| = d_n^L(a, b).
\end{aligned}$$

In the second to last step, we have used the fact that $\mathcal{L}_n$ is the set of Boolean combinations of formulas $p \in \mathrm{At}$ and $\Diamond\phi$ with $\phi \in \mathcal{L}_{n-1}$. So the "$\leq$" part of this step is clearly true, and the "$\geq$" part also holds by a simple induction over the Boolean combinations, using that these are non-expansive, as shown by the same calculations as in the proof of Lemma 4.11. □

Proof Details for Lemma 5.5

Let $a, b \in A$ and $d_n^L(a, b) \leq \epsilon$. We need to show $d_n^G(a, b) \leq \epsilon$, so it suffices to show that D wins the depth-n $(\epsilon + \delta)$-bisimulation game for every $\delta > 0$. The winning condition for the configuration (a, b) is satisfied by assumption, so now suppose S makes the first move from a to a'. We now consider the function

$$f : (A, d_{n-1}) \to_1 ([0, 1], d_e), f(b') = R(a, a') \ominus d_{n-1}(a', b')$$

f is non-expansive because it is composed from non-expansive functions: the map $x \mapsto c \ominus x$ is non-expansive for every $c \in \mathbb{R}$, and the map $b' \mapsto d_{n-1}(a', b')$ is non-expansive by the triangle inequality.

Now,

$$|(\Diamond f)(a) - (\Diamond f)(b)| \leq d_n^K(a, b) = d_n^L(a, b) \leq \epsilon$$

by Lemma 5.4, so, as $R(a, a') > \epsilon$ by the rules of the game,

$$(\Diamond f)(b) \geq (\Diamond f)(a) - \epsilon \geq R(a, a') \wedge f(a') - \epsilon = R(a, a') - \epsilon > 0.$$

This means there exists some $b' \in A$ such that $f(b') > 0$ and

$$R(a, a') - \epsilon \leq (\Diamond f)(b) \leq R(b, b') \wedge f(b') + \tfrac{\delta}{2}.$$

Rearranging this, we get

$$\begin{aligned}
\epsilon + \tfrac{\delta}{2} &\geq R(a, a') - (R(b, b') \wedge f(b')) \\
&= (R(a, a') - R(b, b')) \vee (R(a, a') - f(b')).
\end{aligned}$$

So first, $R(b, b') \geq R(a, a') - (\epsilon + \frac{\delta}{2})$, which means that b' is a legal reply for D. And second, since $f(b') > 0$,

$$\begin{aligned}
\epsilon + \tfrac{\delta}{2} &\geq R(a, a') - f(b') \\
&= R(a, a') - (R(a, a') - d_{n-1}(a', b')) = d_{n-1}(a', b').
\end{aligned}$$

So the configuration reached after the first round of the game is (a', b'), and D has a winning strategy for the $(\epsilon +$

$\frac{\delta}{2}+\gamma$)-game for every $\gamma > 0$, in particular D wins the $(\epsilon+\delta)$-game. An analogous argument can be used if S makes a move from b to some b' instead. ◻

Proof of Lemma 5.6

Put $F := (X, d_1) \to_1 (Y, d_2)$. Let $\epsilon > 0$. We need to find a finite cover of $(X, d_1) \to_1 (Y, d_2)$ by sets of diameter at most ϵ.

Since (X, d_1) and (Y, d_2) are totally bounded, there exist finite $\frac{\epsilon}{4}$-covers $x_1, \ldots, x_n$ of X and $y_1, \ldots, y_k$ of Y.

Now consider the set Φ of functions $\rho\colon \{1, \ldots, n\} \to \{1, \ldots, k\}$, and for every $\rho \in \Phi$ let

$$F_\rho = \{f \in F \mid f(x_i) \in B_{\frac{\epsilon}{4}}(y_{\rho(i)}) \text{ for all } 1 \le i \le n\}.$$

Then clearly $F = \bigcup_{\rho\in\Phi} F_\rho$, so it remains to show that each F_ρ has diameter at most ϵ.

So let $f, g \in F_\rho$, and consider some $x \in X$. There exists some i such that $x \in B_{\frac{\epsilon}{4}}(x_i)$. Now, by non-expansivity of f and g, and the definition of F_ρ,

$$\begin{aligned} d_2(f(x), g(x)) &\le d_2(f(x), f(x_i)) + d_2(f(x_i), y_{\rho(i)}) \\ &\quad + d_2(y_{\rho(i)}, g(x_i)) + d_2(g(x_i), g(x)) \\ &\le d_1(x, x_i) + \tfrac{\epsilon}{4} + \tfrac{\epsilon}{4} + d_1(x_i, x) \\ &\le \tfrac{\epsilon}{4} + \tfrac{\epsilon}{4} + \tfrac{\epsilon}{4} + \tfrac{\epsilon}{4} = \epsilon. \end{aligned}$$

As f, g, and x were chosen arbitrarily, the diameter of F_ρ is at most ϵ. ◻

Proof Details for Lemma 5.7

We show that the map

$$\begin{aligned} I\colon A &\to [0,1]^{k+m} \\ a &\mapsto (p_1(a), \ldots, p_k(a), (\Diamond\phi_1)(a), \ldots, (\Diamond\phi_m)(a)) \end{aligned}$$

is indeed an $\frac{\epsilon}{3}$-isometry. Let $a, b \in A$. We need to show that

$$|d_n(a,b) - \|I(a) - I(b)\|_\infty| \le \tfrac{\epsilon}{3}.$$

Let $f \in F$ and choose ϕ_i such that $\|f - \phi_i\|_\infty \le \frac{\epsilon}{6}$. Then also $\|\Diamond f - \Diamond\phi_i\|_\infty \le \frac{\epsilon}{6}$, by Lemma 3.1. By the triangle inequality, it follows that

$$\big||(\Diamond f)(a) - (\Diamond f)(b)| - |(\Diamond\phi_i)(a) - (\Diamond\phi_i)(b)|\big| \le \tfrac{\epsilon}{3},$$

so, taking the supremum over all $f \in F$:

$$\Big|\bigvee_{f\in F} |(\Diamond f)(a) - (\Diamond f)(b)| - \bigvee_{i\le m} |(\Diamond\phi_i)(a) - (\Diamond\phi_i)(b)|\Big| \le \tfrac{\epsilon}{3}.$$

Recall from Assumption 5.1 that At $= \{p_1, \ldots, p_k\}$ is finite; then

$$\Big|d_n^K(a,b) - \bigvee_{i\le k} |p_i(a) - p_i(b)| \vee \bigvee_{i\le m} |(\Diamond\phi_i)(a) - (\Diamond\phi_i)(b)|\Big| \le \tfrac{\epsilon}{3},$$

which is what we needed to show.

It remains to give a finite ϵ-cover of (A, d_n). As $[0,1]^{k+m}$ is compact under the supremum metric, it has a finite $\frac{\epsilon}{3}$-cover $v_1, \ldots, v_p$. Then the pre-image of each ball $B_{\frac{\epsilon}{3}}(v_i)$ has diameter at most ϵ: for any $a, b \in I^{-1}[B_{\frac{\epsilon}{3}}(v_i)]$,

$$\begin{aligned} d_n(a,b) &\le \|I(a) - I(b)\|_\infty + \tfrac{\epsilon}{3} \\ &\le \|I(a) - v_i\|_\infty + \|v_i - I(b)\|_\infty + \tfrac{\epsilon}{3} \le \epsilon. \end{aligned}$$

So a finite ϵ-cover of (A, d_n) arises by taking one element from each (non-empty) $I^{-1}[B_{\frac{\epsilon}{3}}(v_i)]$. ◻

Proof of Lemma 5.8

Let $f \in F$ and $\epsilon > 0$. We need to find some $f_\epsilon \in L$ such that $\|f - f_\epsilon\|_\infty \le \epsilon$.

By total boundedness, there exists an $\frac{\epsilon}{4}$-cover $x_1, \ldots, x_n$ of (X, d). By assumption, for every $i, j \in \{1, \ldots, n\}$ there exists some $f_{ij} \in L$ such that $|f(x_i) - f_{ij}(x_i)| \le \frac{\epsilon}{2}$ and $|f(x_j) - f_{ij}(x_j)| \le \frac{\epsilon}{2}$. Now define $f_\epsilon = \bigvee_{i\le n}\bigwedge_{j\le n} f_{ij} \in L$. Then, for any $x \in X$ there exists some k such that $d(x_k, x) \le \frac{\epsilon}{4}$ and thus:

$$\begin{aligned} f_\epsilon(x) &= \bigvee_{i\le n}\bigwedge_{j\le n} f_{ij}(x) \le \bigvee_{i\le n} f_{ik}(x) \le \bigvee_{i\le n} f_{ik}(x_k) + \tfrac{\epsilon}{4} \\ &\le \bigvee_{i\le n} f(x_k) + \tfrac{3\epsilon}{4} = f(x_k) + \tfrac{3\epsilon}{4} \le f(x) + \epsilon, \end{aligned}$$

and, symmetrically:

$$\begin{aligned} f_\epsilon(x) &= \bigvee_{i\le n}\bigwedge_{j\le n} f_{ij}(x) \ge \bigwedge_{j\le n} f_{kj}(x) \ge \bigwedge_{j\le n} f_{kj}(x_k) - \tfrac{\epsilon}{4} \\ &\ge \bigwedge_{j\le n} f(x_k) - \tfrac{3\epsilon}{4} = f(x_k) - \tfrac{3\epsilon}{4} \ge f(x) - \epsilon, \end{aligned}$$

where we have used non-expansivity of f and the f_{ij} as well as the originally assumed property of the f_{ij}. ◻

Proof Details for Lemma 5.9

Lemma 5.8 can be applied because (A, d_n) is totally bounded by Lemma 5.7, and because the set $\mathcal{L}_n$ is clearly closed under $\wedge$ and $\vee$.

Given a function $f\colon (A, d_n) \to_1 ([0,1], d_e)$, $a, b \in A$ and $\epsilon > 0$, we need to find $\phi \in \mathcal{L}_n$ such that $|f(a) - \phi(a)| \le \epsilon$ and $|f(b) - \phi(b)| \le \epsilon$.

W.l.o.g. $f(a) \ge f(b)$ (otherwise we can pass to $1 - f$ and negate the resulting formula). Now put $\Delta = f(a) - f(b)$. Then $\Delta \le d_n(a,b)$ by non-expansivity of f. Since $d_n = d_n^L$, there exists $\psi \in \mathcal{L}_n$ such that $\Delta - \frac{\epsilon}{2} \le \psi(a) - \psi(b)$. Let $u, v, w \in \mathbb{Q} \cap [0,1]$ such that

$$\begin{aligned} \psi(b) - \tfrac{\epsilon}{2} &\le u \le \psi(b) \\ \Delta - \tfrac{\epsilon}{2} &\le v \le \Delta \\ f(b) &\le w \le f(b) + \tfrac{\epsilon}{2}. \end{aligned}$$

Put $\phi = \neg(\neg((\psi \ominus u) \wedge v) \ominus w)$. Then ϕ approximates f at a and b:

$$\begin{aligned} f(a) - \tfrac{\epsilon}{2} &\le \phi(a) \le f(a) + \tfrac{\epsilon}{2} \\ f(b) &\le \phi(b) \le f(b) + \epsilon. \end{aligned}$$

The detailed calculations for the above inequalities follow. Evaluating subformulas at a gives:

$$\psi(a) - \psi(b) \le (\psi \ominus u)(a) \le \psi(a) - \psi(b) + \tfrac{\epsilon}{2}$$
$$\Delta - \tfrac{\epsilon}{2} \le ((\psi \ominus u) \wedge v)(a) \le \Delta$$
$$1 - \Delta \le (\neg((\psi \ominus u) \wedge v))(a) \le 1 - \Delta + \tfrac{\epsilon}{2}$$
$$1 - f(a) - \tfrac{\epsilon}{2} \le (\neg((\psi \ominus u) \wedge v) \ominus w)(a) \le 1 - f(a) + \tfrac{\epsilon}{2}$$
$$f(a) - \tfrac{\epsilon}{2} \le \phi(a) \le f(a) + \tfrac{\epsilon}{2}.$$

Evaluating subformulas at b gives:

$$0 \le (\psi \ominus u)(b) \le 0 + \tfrac{\epsilon}{2}$$
$$0 \le ((\psi \ominus u) \wedge v)(b) \le 0 + \tfrac{\epsilon}{2}$$
$$1 - \tfrac{\epsilon}{2} \le (\neg((\psi \ominus u) \wedge v))(b) \le 1$$
$$1 - f(b) - \epsilon \le (\neg((\psi \ominus u) \wedge v) \ominus w)(b) \le 1 - f(b)$$
$$f(b) \le \phi(b) \le f(b) + \epsilon.$$

□

Full Proof of Lemma 6.1

We show by induction on n that D wins the depth-n ϵ-bisimulation game for every $\epsilon > 0$. The base case is trivial since the depth-0 game is an immediate win; we proceed with the inductive step from n to $n+1$. We need to show that D wins the depth-$(n+1)$ ϵ-bisimulation game for $\mathcal{A}, a$ and $\mathcal{F}, \pi_{n+1}(a)$. By the explicit definition (2) of π_{n+1}, it is immediate that the winning condition holds in the initial configuration.

If S makes the first move from a to some $a' \in A$, then D can reply with $\pi_n(a')$, since by (2),

$$R^{\mathcal{F}}(\pi_{n+1}(a), \pi_n(a')) = \bigvee\nolimits_{\pi_n(a'')=\pi_n(a')} R^{\mathcal{A}}(a, a'') \ge R^{\mathcal{A}}(a, a').$$

By induction, $d_n^G(a', \pi_n(a')) = 0$, so D wins.

If instead, S makes the first move from $\pi_{n+1}(a)$ to some $y \in F$, then $R^{\mathcal{F}}(\pi_{n+1}(a), y) > 0$ by the rules of the game, so $y \in F_n$ by construction of $R^{\mathcal{F}}$. By (2),

$$R^{\mathcal{F}}(\pi_{n+1}(a), y) = \bigvee\nolimits_{\pi_n(a')=y} R^{\mathcal{A}}(a, a')).$$

Thus, D can pick $a' \in A$ with $\pi_n(a') = y$ such that

$$R^{\mathcal{A}}(a, a') \ge R^{\mathcal{F}}(\pi_{n+1}(a), y) - \epsilon.$$

By induction, $d_n^G(a', y) = 0$, so D wins. □

Proof of Lemma 7.4

D wins the game by copying every move S makes. By the definition of $R^{\mathcal{A}^k_{a_0}}$ and the $p^{\mathcal{A}^k_{a_0}}$ it is clear that such a strategy is winning as long as the game never leaves the neighbourhood $U^k(a_0)$. By the rules of the game, S can only ever move along positive edges of the model, so if the configuration after round i is (a_i, a_i), it must hold that $D(a_0, a_i) \le i$ and therefore $a_i \in U^k(a_0)$. □

Proof of Lemma 7.7

Induction on ϕ. The cases for equality, propositional atoms and the relation symbol R follow from the winning condition. The Boolean cases are proved just as in Lemma 4.11. The remaining case is that of existential quantification:

Let $(\bar{a}, \bar{b})$ be the current configuration. Now let $\delta > 0$, let a be such that

$$(\exists x.\, \phi)(\bar{a}) - \phi(\bar{a}a) < \delta,$$

and let b be D's winning answer to S's move a. Then by induction, $|\phi(\bar{a}a) - \phi(\bar{b}b)| \le \epsilon$. Thus,

$$(\exists x.\, \phi)(\bar{b}) \ge \phi(\bar{b}b) \ge \phi(\bar{a}a) - \epsilon > (\exists x.\, \phi)(\bar{a}) - \epsilon - \delta.$$

Since $\delta > 0$ was arbitrary, it follows that

$$(\exists x.\, \phi)(\bar{b}) \ge (\exists x.\, \phi)(\bar{a}) - \epsilon.$$

We show symmetrically that $(\exists x.\, \phi)(\bar{a}) \ge (\exists x.\, \phi)(\bar{b}) - \epsilon$, that is, $|(\exists x.\, \phi)(\bar{a}) - (\exists x.\, \phi)(\bar{b})| \le \epsilon$ as required. □

Proof Details for Lemma 7.8

Recall that D needs to maintain the following invariant:

> If $(\bar{b}, \bar{c}) = ((b_0, \ldots, b_i), (c_0, \ldots, c_i))$ is the current configuration then there is an isomorphism between $\mathcal{B}^{k_i}_{\bar{b}}$ and $C^{k_i}_{\bar{c}}$ mapping each b_j to c_j.

This invariant clearly holds at the beginning of the game: the initial configuration is (a_0, a_0), and $k_0 = k$, so the two models in the invariant are both isomorphic to $\mathcal{A}^k_{a_0}$ and the isomorphism between them maps a_0 to itself.

The invariant also implies the winning condition for D, i.e. that the current configuration is a partial isomorphism up to 0. This is because the isomorphism from the invariant maps each b_j to the corresponding c_j.

It remains to show that D has a way to maintain the invariant. Suppose that $i < n$ and the current configuration is as in the invariant.

First, suppose that S picks $b \in U^{2k_{i+1}}(\bar{b})$. Then D picks a reply c according to the isomorphism. By the triangle inequality for Gaifman distance, $U^{k_{i+1}}(b) \subseteq U^{k_i}(\bar{b})$ (since $2k_{i+1} + k_{i+1} = 3k_{i+1} = k_i$), and thus also $U^{k_{i+1}}(c) \subseteq U^{k_i}(\bar{c})$ by isomorphism. This implies that the domain $U^{k_{i+1}}(\bar{b}b)$ and range $U^{k_{i+1}}(\bar{c}c)$ of the presumptive new isomorphism are contained in the domain and range of the old one. So the new isomorphism can be taken to be the restriction of the old isomorphism to the new domain and range. The case where S picks a new state $c \in U^{2k_{i+1}}(\bar{c})$ is entirely symmetric.

Otherwise, suppose S picks some b in $\mathcal{B}$ with $b \notin U^{2k_{i+1}}(\bar{b})$. Then, by the triangle inequality for Gaifman distance, $U^{k_{i+1}}(\bar{b}) \cap U^{k_{i+1}}(b) = \emptyset$. In this case, D picks as his reply c the copy of b in a fresh copy of either $\mathcal{A}$ or $\mathcal{A}^k_{a_0}$ (i.e. one that has not been played to in the previous rounds). Such a fresh copy is always available, because at most one of them gets visited in each round. Then the radius-k_{i+1} neighbourhoods of b and c are isomorphic because b and c are the same element in isomorphic copies of either $\mathcal{A}$ or $\mathcal{A}^k_{a_0}$. The radius-k_{i+1} neighbourhoods of $\bar{b}$ and $\bar{c}$ are also isomorphic, by restriction of the old isomorphism. We thus have two isomorphisms with disjoint domains and ranges, which we combine to form the requested new isomorphism. Again, the case where S plays in C instead is symmetric. □

Proof of Lemma 8.2

A winning strategy for D is given by $\pi\colon A^+ \to A$, i.e. projection to the last element. More precisely, D wins by maintaining the invariant that the current configuration is of the form $(a, \bar{a})$ with $\pi(\bar{a}) = a$. By definition of $p^{\mathcal{A}^*}$ the invariant implies the winning condition. If S moves from a to some a', then D can reply with a move from $\bar{a}$ to $\bar{a}a'$, which is legal by definition of $R^{\mathcal{A}^*}$. The situation is symmetric if S makes a move in $\mathcal{A}^*$ instead. □

Full Proof of Lemma 8.3

Let $d^G_{k+1}(a, b) < \epsilon$; we show that $|\phi_{\mathcal{A}}(a) - \phi_{\mathcal{B}}(b)| \leq \epsilon$, which proves the claim. By assumption, D wins the ϵ-bisimulation-game for $\mathcal{A}, a$ and $\mathcal{B}, b$. By Lemmas 4.3, 7.4 and 8.2, D also wins the depth-$(k+1)$ ϵ-bisimulation game for $(\mathcal{A}^*)^k_a, a$ and $(\mathcal{B}^*)^k_b, b$.

The models $(\mathcal{A}^*)^k_a$ and $(\mathcal{B}^*)^k_b$ both have the shape of trees of depth k, so for every $0 \leq i \leq k$, before the start of round $i + 1$ of the above game, the two states on either side of the current configuration are nodes at distance i from the root of their tree (i.e. a or b). In particular, if round $k + 1$ needs to be played, then S has no legal move, because the current configuration consists of two leaf nodes.

Using this observation, we conclude that D's winning strategy for the depth-$(k + 1)$ game is in fact also a winning strategy for the unbounded ϵ-bisimulation game, so $|\phi_{(\mathcal{A}^*)^k_a}(a) - \phi_{(\mathcal{B}^*)^k_b}(b)| \leq \epsilon$, by bisimulation invariance of ϕ.

By locality and bisimulation invariance of ϕ, and again Lemma 8.2, we have $\phi_{(\mathcal{A}^*)^k_a}(a) = \phi_{(\mathcal{A}^*)}(a) = \phi_{\mathcal{A}}(a)$ as well as $\phi_{(\mathcal{B}^*)^k_b}(b) = \phi_{(\mathcal{B}^*)}(b) = \phi_{\mathcal{B}}(b)$. Thus $|\phi_{\mathcal{A}}(a) - \phi_{\mathcal{B}}(b)| \leq \epsilon$, as claimed. □

Full Proof of Theorem 8.4

By Lemmas 7.8 and 8.3, ϕ is depth-k bisimulation-invariant for $k = 3^n + 1$. By Theorem 5.3, ϕ can be modally approximated on the model $\mathcal{F}$ constructed from the final chain in Section 6, i.e. for every $\epsilon > 0$ there exists a modal formula ϕ_ϵ of rank at most k such that for every $x \in F$, $|\phi(x)-\phi_\epsilon(x)| \leq \epsilon$. Now let $\mathcal{A}$ be a fuzzy relational model and $a \in A$. By Lemma 6.1, $\phi(a) = \phi(\pi_k(a))$ and $\phi_\epsilon(a) = \phi_\epsilon(\pi_k(a))$ (where π_k is the projection into the final chain), so we obtain $|\phi(a) - \phi_\epsilon(a)| \leq \epsilon$, as required. □

Details for Remark 8.5

In the version of Lemma 8.2 where $\mathcal{A}^*$ is the partial unravelling instead, D wins with a similar, but slightly more complicated invariant: the current configuration is either of the form $(a, \bar{a})$ with $\pi(\bar{a}) = a$ or it is a pair of two equal states, the second being from one of the disjoint copies of $\mathcal{A}$. D can maintain this invariant for the first $k + 1$ rounds just as before, and after that can copy S's moves indefinitely because the game is now played between identical models.