Graded modal logic and counting bisimulation

Martin Otto, September 2019

Abstract: This note sketches the extension of the basic characterisation theorems as the bisimulation-invariant fragment of first-order logic to modal logic with graded modalities and matching adaptation of bisimulation. The result is closely related to unpublished work in the diploma thesis of Rebecca Lukas [5], the presentation newly adapted. We focus on showing *expressive completeness* of graded multi-modal logic for those first-order properties of pointed Kripke structures that are preserved under counting bisimulation equivalence among all or among just all finite pointed Kripke structures, similar to the treatment in [6, 4]. The classical version of the characterisation was treated in [2].

1 Modal back-and-forth with costed counting

We write ML for basic modal logic, typically considering its multi-modal format with finitely many modalities indexed as $(\Diamond_i)$ with duals $(\Box_i)$, where we think of the indices i as ranging over a finite set if *agents*, and finitely many basic propositions indexed as (p_j). In the mono-modal setting with just one agent and we write $\Diamond$ and $\Box$ for the associated modalities. Intended models are Kripke structures $\mathfrak{M} = (W, (E_i), (P_j))$ on sets $W \neq \emptyset$ of *possible worlds*, with accessibility relations $E_i = E_i^{\mathfrak{M}} \subseteq W \times W$ as interpretations of binary relation symbols E_i, one for each agent i, and unary predicates $P_j = P_j^{\mathfrak{M}} \subseteq W$ as interpretations of unary relation symbols P_j, one for each basic proposition j. [1] Formulae are evaluated in *pointed* Kripke structures $\mathfrak{M}, w$ with a distinguished world $w \in W$. For any world $u \in W$ we write $E_i[u]$ for the set of its immediate E_i-successors

$$E_i[u] = \{v \in W : (u, v) \in E_i\},$$

which supports the usual interpretation of $\Diamond_i$ according to the inductive clause

$$\mathfrak{M}, w \models \Diamond_i \varphi \ \text{ if } \ \{u \in E_i[w] : \mathfrak{M}, u \models \varphi\} \neq \emptyset.$$

The counting extension of ML, which we denote by CML, has *graded modalities* $\Diamond_i^{\geq k}$ for every $k \in \mathbb{N} \setminus \{0\}$ instead of just $\Diamond_i$, whose semantics captures that there are *at least k successors* w.r.t. accessibility relation E_i such that ...

The defining clause for the semantics of $\Diamond^{\geq k}$ is that

$$\mathfrak{M}, w \models \Diamond^{\geq k} \varphi \ \text{ if } \ |\{u \in E_i[w] : \mathfrak{M}, u \models \varphi\}| \geq k.$$

Using $[\![\varphi]\!]^{\mathfrak{M}}$ to denote the extension of (the property defined by) φ in $\mathfrak{M}$, $[\![\varphi]\!]^{\mathfrak{M}} := \{u \in W : \mathfrak{M}, u \models \varphi\}$, the above clause can be rewritten as

$$[\![\Diamond^{\geq k} \varphi]\!]^{\mathfrak{M}} = \{u \in W : |E_i[u] \cap [\![\varphi]\!]^{\mathfrak{M}}| \geq k\}.$$

---

[1] We drop superscripts to reference $\mathfrak{M}$ whenever this seems uncritical.

In particular, $\Diamond_i^{\geqslant 1}$ is just $\Diamond_i$ so that obviously $\mathsf{ML} \subseteq \mathsf{CML}$ in expressive power. On the other hand, the usual standard translation of ML into first-order logic FO to establish $\mathsf{ML} \subseteq \mathsf{FO}$ easily extends to CML (at the expense of more first-order variables for the parametrisation of that many distinct E_i-successors, or alternatively with the use of first-order counting quantifiers $\exists^{\geqslant k}$ in a two-variable first-order setting), so that

$$\mathsf{ML} \subseteq \mathsf{CML} \subseteq \mathsf{FO},$$

and these inclusions are easily seen to be strict. For FO-formulae we denote the usual quantifier rank as in $\mathrm{qr}(\varphi)$. The fragments $\mathsf{FO}_q := \{\varphi \in \mathsf{FO} : \mathrm{qr}(\varphi) \leqslant q\}$ and the induced approximations $\equiv^{\mathsf{FO}_q}$ to elementay equivalence are defined as usual.

Just as the semantics of basic modal logic ML is invariant under bisimulation equivalence, so graded modal logic CML is invariant under the natural refinement to *graded bisimulation equivalence*. We give a static definition of graded bisimulation *relations* before discussing the more dynamic intuition based on the associated back-and-forth (b&f) game underlying this natural notion of structural equivalence.

Definition 1.1. A non-empty binary relation $Z \subseteq W \times W'$ between the universes of two Kripke structures $\mathfrak{M} = (W, (E_i), (P_j))$ and $\mathfrak{M} = (W', (E'_i), (P'_j))$ is a *graded bisimulation relation* if it satisfies the following conditions:

– *atom equivalence*: for all $(u, u') \in Z$ and j, $u \in P_j \Leftrightarrow u' \in P'_j$.
– *graded forth*: for all $(u, u') \in Z$ and i, for all $k \geqslant 1$: for pairwise distinct $v_1, \ldots, v_k \in E_i[u]$ there are pairwise distinct $v'_1, \ldots, v'_k \in E'_i[u']$ with $(v_1, v'_1), \ldots, (v_k, v'_k) \in Z$.
– *graded back*: for all $(u, u') \in Z$ and i, for all $k \geqslant 1$: for pairwise distinct $v'_1, \ldots, v'_k \in E'_i[u']$ there are pairwise distinct $v_1, \ldots, v_k \in E_i[u]$ with $(v_1, v'_1), \ldots, (v_k, v'_k) \in Z$.

We write $\mathfrak{M}, w \sim_{\#} \mathfrak{M}', w'$, and say that these pointed Kripke structures are *graded bisimilar* if there is such a graded bisimulation relation with $(w, w') \in Z$.

The associated graded bisimulation game is played between two players, $\mathbf{I}$ and $\mathbf{II}$, over the two structures $\mathfrak{M}, \mathfrak{M}'$. Positions are pairings $(u, u') \in W \times W'$ and a single round played from this position allows $\mathbf{I}$ to challenge $\mathbf{II}$ in the following exchange of moves

– $\mathbf{I}$ chooses, for some i, a *finite* non-empty subset of one of $E_i[u]$ or $E'_i[u']$;
  $\mathbf{II}$ must respond with a matching subset of $E'_i[u']$ or $E_i[u]$ on the opposite side of the *same finite size*;
– $\mathbf{I}$ picks a world in the set proposed by $\mathbf{II}$;
  $\mathbf{II}$ must respond by a choice of a matching world in the set proposed by $\mathbf{I}$.

The pairing between the worlds chosen during the second stage of this round is the new position in the game.

Either player loses in this round if stuck (for a choice or response), and the second player also loses as soon as the position violates atom equivalence.

It is not hard to show that a graded bisimulation relation Z as in Definition 1.1 *is* an extensional formalisation of a *winning strategy* in the unbounded (i.e. infinite) graded bisimulation game, which is good for any position $(u, u') \in Z$. Atom equivalence for Z ensures that $\mathbf{II}$ cannot have lost in a position covered by Z; and the b&f conditions for Z make sure that $\mathbf{II}$ can respond to any challenge by $\mathbf{I}$ without leaving cover by Z. The winning strategy embodied in Z is in general non-deterministic in allowing $\mathbf{II}$ several choices in accordance with Z.

Game views of the concept of bisimulations typically suggest the adequate *finite approximations* to the a priori essentially infinitary closure condition implicit in he definition of bisimulation relations: the natural notion of ℓ-bisimulation captures the existence of a winning strategy for $\mathbf{II}$ in the ℓ-*round game*. It typically provides

exactly the right levels of b&f equivalence to match finite levels of logical equivalence based on a gradation by quantifier rank or nesting depth of (modal) quantifiers. This approach generates all kind of natural variants of the classical Ehrenfeucht–Fraïssé connection, which produce precise matches between corresponding levels of b&f structural equivalences and logical equivalences.[2] For basic modal logic ML as well as for its graded variant CML, the natural notion of ℓ-bisimulation $\sim^\ell$ or $\sim^\ell_\#$ lies a limitation to ℓ rounds, which for these logics also corresponds to a limitation of the structural exploration to depth ℓ from the initial worlds – which on the logics' side in turns matches the natural notions of nesting depths for modal operators. This modal nesting depth $\mathrm{nd}(\varphi)$ is defined for CML just as for ML, with the crucial clause in the inductive definition that $\mathrm{nd}(\Diamond^{\geqslant k}\varphi) = \mathrm{nd}(\varphi) + 1$.

Definition 1.2. [graded ℓ-bisimilarity]
Pointed Kripke structures $\mathfrak{M}, w$ and $\mathfrak{M}', w'$ are *graded ℓ-bisimilar* if player **II** has a winning strategy in the ℓ-round graded bisimulation game on $\mathfrak{M}, \mathfrak{M}'$ starting from (w, w'); notation: $\mathfrak{M}, w \sim^\ell_\# \mathfrak{M}', w'$.

In the following we write CML_ℓ for the fragment of formulae of modal nesting depth up to ℓ, and $\equiv^{\mathsf{CML}_\ell}$ for the notion of equivalence between pointed Kripke structures that is induced by indistinguishability in CML_ℓ.

Lemma 1.3 (graded modal Ehrenfeucht–Fraïssé theorem)**.**
For any two finitely branching pointed Kripke structures $\mathfrak{M}, w$ and $\mathfrak{M}', w'$ over a common finite modal signature, and for all $\ell \in \mathbb{N}$:

$$\mathfrak{M}, w \sim^\ell_\# \mathfrak{M}', w' \quad \Longleftrightarrow \quad \mathfrak{M}, w \equiv^{\mathsf{CML}_\ell} \mathfrak{M}', w'$$

The proof is as usual, but it is worth noting that the corresponding modal Ehrenfeucht–Fraïssé theorem, i.e. the variant for ML rather than CML, is good across all pointed Kripke structures, finite branching or not. (It requires the same restriction to finite signatures, though.) The source of this difference is that the cardinality of the set of E_i-successors can be pinpointed precisely by a single CML_1-formula if finite, but requires an infinite collection of formulae (in CML_1) to be characterised as infinite.

Example 1.4. The following two depth 2 rooted tree frames, in a modal signature with a single modality, are indistinguishable in CML, but not $\sim^2_\#$-equivalent:

– $\mathfrak{M}, w$ has of a root node w with children $(u_i)_{i \geqslant 1}$,
   where each u_i has precisely i many children;
– $\mathfrak{M}', w'$ has a root node w' with children $(u'_i)_{i \in \mathbb{N}}$,
   where each u'_i for $i \geqslant 1$ has precisely i many children
   and the extra u_0 has infinitely many children.

They are $\sim^1_\#$-equivalent though.

The reason behind this example is that the existence of a successor that has infinitely many successors is not expressible in CML (nor is it expressible in FO, which precludes easy remedies). This problem does not occur in restriction to finitely branching structures. While such restrictions and problems related to definability also occur for basic modal logic ML, and are familiar from classical modal theory in connection with Hennessy-Milner phenomena, they here already strike at the level of $\sim^\ell_\#$ not just at the level of full $\sim_\#$.

The underlying distinction that really matters for our concerns lies with the index of the associated equivalence relations on the class of all (or even just all

[2]Recall how the classical Ehrenfeucht–Fraïssé matches m-isomorphy, i.e. m-round pebble game b&f equivalence, with m-elemenatry equivalence.

finite) pointed Kripke structures. The index of $\sim^\ell$ and $\equiv^{\mathsf{ML}_\ell}$ is finite as long as the underlying signature is finite, for each $\ell \in \mathbb{N}$. Contrast this with $\sim^1_\#$, even for just a single modality and without propositions. Here we find that the formulae $\Diamond^{\geqslant k}\top$ are pairwise inequivalent, and the family of pointed Kripke frames consisting of a root node with precisely $k-1$ immediate E-successors are pairwise $\sim^1_\#$-inequivalent.

To apply a finer gradation than the straightforward gradation by just modal nesting depth, we account for the *cost of counting*. Obviously this is something to take into account in relation to FO where any natural rendering of counting has a cost in terms of quantification.

Definition 1.5. [c-graded ℓ-bisimilarity]
For $c \in \mathbb{N}$ consider the ℓ-round graded bisimulation game with the additional constraint that the sets s/s' of successor worlds chosen by $\mathbf{I}$, and consequently the responses s'/s given by $\mathbf{II}$, must be of size $\leqslant c$. Call this the *c-graded ℓ-round game*. $\mathfrak{M}, w$ and $\mathfrak{M}', w'$ are *c-graded ℓ-bisimilar*, $\mathfrak{M}, w \sim^{c,\ell}_\# \mathfrak{M}', w'$, if $\mathbf{II}$ has a winning strategy in this c-graded ℓ-round game starting from (w, w').

To match the constraint syntactically, in terms of the grades k on modalities $\Diamond^{\geqslant k}$, we define fragments $\mathsf{CML}_{c,\ell}$. Essentially we refine the primary gradation of CML in terms of nesting depth ℓ, which we have in $\mathsf{CML} = \bigcup_\ell \mathsf{CML}_\ell$.

Let the (counting) *rank*, $\mathrm{rk}_\#(\varphi) \in \mathbb{N}$, of formulae $\varphi \in \mathsf{CML}$ be defined by induction according to

- $\mathrm{rk}_\#(\varphi) := 0$ for propositional formulae;
- $\mathrm{rk}_\#(\neg\varphi) = \mathrm{rk}_\#(\varphi)$;
- $\mathrm{rk}_\#(\varphi_1 * \varphi_2) = \max(\mathrm{rk}_\#(\varphi_1), \mathrm{rk}_\#(\varphi_2))$ for $* = \wedge, \vee$;
- $\mathrm{rk}_\#(\Diamond^{\geqslant k}_i \varphi) := \max(k, \mathrm{rk}_\#(\varphi))$.

The fragment $\mathsf{CML}_{c,\ell}$ comprises all CML-formulae φ that have counting rank $\mathrm{rk}_\#(\varphi) \leqslant c$ and nesting depth $\mathrm{nd}(\varphi) \leqslant \ell$. Correspondingly we define

$$\mathfrak{M}, w \equiv^{\mathsf{CML}_{c,\ell}} \mathfrak{M}', w'$$

as indistinguishability at level c, ℓ: $\mathfrak{M}, w \models \varphi \Leftrightarrow \mathfrak{M}', w' \models \varphi$ for all $\varphi \in \mathsf{CML}_{c,\ell}$. Obviously $\mathsf{CML}_{c,\ell} \subseteq \mathsf{CML}_\ell$ for every c, and $\equiv^{\mathsf{CML}_{c,\ell}}$ can only be coarser than $\equiv^{\mathsf{CML}_\ell}$ while leading to the same common refinement $\equiv^{\mathsf{CML}}$ in the limit as

$$\bigcup_{c,\ell} \mathsf{CML}_{c,\ell} = \mathsf{CML} = \bigcup_\ell \mathsf{CML}_\ell.$$

Crucially, for finite modal signatures, $\mathsf{CML}_{c,\ell}$ is finite up to logical equivalence, for every combination of $c, \ell \in \mathbb{N}$. This is obvious at level $\ell = 0$ (by finiteness of the set of basic propositions in the signature). For fixed c, the relevant $\Diamond^{\geqslant k}_i$-constituents at nesting depth level $\ell + 1$ are constrained to finitely many choices for i (by finiteness of the set of agents in the signature). Inductively, the finitely many representatives of any constituent subformulae at counting rank level c at nesting depth ℓ can only contribute to finitely many disjunctive normal forms involving admissible $\Diamond^{\geqslant k}_i$-applications at counting rank level c and nesting depth $\ell + 1$.

Lemma 1.6 (c-graded modal Ehrenfeucht–Fraïssé theorem)**.**
For any two pointed Kripke structures $\mathfrak{M}, w$ and $\mathfrak{M}', w'$ over a common finite modal signature, and for all $\ell \in \mathbb{N}$:

$$\mathfrak{M}, w \sim^{c,\ell}_\# \mathfrak{M}', w' \quad \Longleftrightarrow \quad \mathfrak{M}, w \equiv^{\mathsf{CML}_{c,\ell}} \mathfrak{M}', w'$$

In comparison with Lemma 1.3 the restriction to finitely branching structures has been lifted. We present a detailed proof, which retraces and implicitly reviews

the key idea also for Lemma 1.3 and any other natural variant of the Ehrenfeucht–Fraïssé connection.

In the following we write $[\mathfrak{M}, w]_{c,\ell}$ for the $\sim_{\#}^{c,\ell}$-equivalence class of a pointed Kripke structure $\mathfrak{M}, w$,

$$[\mathfrak{M}, w]_{c,\ell} := \{\mathfrak{M}', u' \colon \mathfrak{M}', u' \sim_{\#}^{c,\ell} \mathfrak{M}, u\},$$

and, relative to a fixed Kripke structure $\mathfrak{M}$, just $[u]_{c,\ell}$ for the $\sim_{\#}^{c,\ell}$-equivalence classes of its worlds $u \in W$:

$$[u]_{c,\ell} := \{u' \in W \colon \mathfrak{M}, u' \sim_{\#}^{c,\ell} \mathfrak{M}, u\}.$$

Proof of Lemma 1.6. To establish the implication from left to right in the claim of the lemma we show that $\mathrm{rk}_{\#}(\varphi) \leqslant c$, $\mathrm{nd}(\varphi) \leqslant \ell$ and $\mathfrak{M}, w \sim_{\#}^{c,\ell} \mathfrak{M}', w'$ imply that

$$\mathfrak{M}, w \models \varphi \;\Rightarrow\; \mathfrak{M}', w' \models \varphi.$$

This is shown by syntactic induction on $\varphi \in \mathsf{CML}$. The claim is trivial for propositional φ and obviously compatible with boolean connectives. For $\varphi = \Diamond_i^{\geqslant k} \psi$, the definition of counting rank guarantees that $\mathrm{rk}_{\#}(\varphi) \geqslant k, \mathrm{rk}_{\#}(\psi)$.

So $\mathfrak{M}, w \sim_{\#}^{c,\ell} \mathfrak{M}', w'$ for $\ell \geqslant \mathrm{nd}(\varphi)$ and $c \geqslant \mathrm{rk}_{\#}(\varphi)$ implies that **II** has, in the first round, a response if **I** chooses a set $s \subseteq E_i[w] \cap \llbracket \psi \rrbracket^{\mathfrak{M}}$ of size k (such s exists since $\mathfrak{M}, w \models \Diamond^{\geqslant k} \psi$); for this response, say $s' \subseteq E_i[w']$ to be adequate, **II** must have, for every $u' \in s'$ (that **I** could choose in the second part of the first round) some response $u \in s$ such that $\mathfrak{M}, u \sim_{\#}^{c,\ell-1} \mathfrak{M}', u'$. The inductive hypothesis for ψ implies that $s' \subseteq \llbracket \psi \rrbracket^{\mathfrak{M}'}$ as $\mathrm{nd}(\psi) \leqslant \ell - 1$ and $\mathrm{rk}_{\#}(\psi) \leqslant c$. Thus $\mathfrak{M}', w' \models \varphi$.

For the converse implication, we define characteristic formulae $\chi_{\mathfrak{M}, w}^{\ell}$ by induction on $\ell \in \mathbb{N}$ as follows:

$$
\begin{aligned}
\chi_{\mathfrak{M}, w}^{c,0} \quad &:= \quad \bigwedge \{p_j \colon w \in P_j^{\mathfrak{M}}\} \wedge \bigwedge \{\neg p_j \colon w \notin P_j^{\mathfrak{M}}\} \\
\chi_{\mathfrak{M}, w}^{c,\ell+1} \quad &:= \quad \chi_{\mathfrak{M}, w}^{c,\ell} \\
&\quad \wedge \bigwedge \{\; \Diamond_i^{\geqslant k} \chi_{\mathfrak{M}, u}^{c,\ell} \colon \; u \in E_i[w], |E_i[w] \cap [u]_{c,\ell}| \geqslant k, k \leqslant c\} \quad [\textit{forth}] \\
&\quad \wedge \bigwedge \{\neg \Diamond_i^{\geqslant k} \chi_{\mathfrak{M}, u}^{c,\ell} \colon \; u \in E_i[w], |E_i[w] \cap [u]_{c,\ell}| < k \leqslant c\} \qquad [\textit{back}]
\end{aligned}
$$

Inductively it is clear that $\chi_{-}^{c,\ell} \in \mathsf{CML}_{c,\ell}$ since the the formally unbounded conjunctions are finite up to logical equivalence: we now have this, without any assumption on branching degree, as all $\mathsf{CML}_{c,\ell}$ are finite up to logical equivalence. We show that $\chi_{\mathfrak{M}, w}^{c,\ell}$ defines the $\sim_{\#}^{c,\ell}$-equivalence class $[\mathfrak{M}, w]_{c,\ell}$ of $\mathfrak{M}, w$. Obviously $\mathfrak{M}, w \models \chi_{\mathfrak{M}, w}^{c,\ell}$; it remains to show that

$$\mathfrak{M}', w' \models \chi_{\mathfrak{M}, w}^{c,\ell} \;\Rightarrow\; \mathfrak{M}, w \sim_{\#}^{c,\ell} \mathfrak{M}', w'.$$

This is clear for $\ell = 0$, and we show the claim by inductioin on ℓ for fixed c. Assuming the claim at level ℓ, we know that, across all pointed Kripke structures of the relevant (finite) signature, there are finitely many $\sim_{\#}^{c,\ell}$-types t, each defined by some characteristic formula $\chi_t^{c,\ell}$. To establish the claim at level $\ell + 1$ we look at the first round in the c-graded $(\ell + 1)$-round game on $\mathfrak{M}$ and $\mathfrak{M}'$ starting from (w, w') and give good strategy advice for **II**. We distinguish two cases, depending on the challenge played by **I**, and invoking either the *forth* or the *back* constituents in $\chi_{\mathfrak{M}, w}^{c,\ell+1}$.

If **I**'s choice is some $s \subseteq E_i[w]$ of size $|s| \leqslant c$, we partition s into its disjoint constituents

$$s_t := s \cap \llbracket \chi_t^{c,\ell} \rrbracket^{\mathfrak{M}}$$

where t ranges over the finitely many, mutually exclusive, $\sim_{\#}^{c,\ell}$-types. As $s_t \subseteq s$ and $k_t := |s_t| \leqslant c$, $\mathfrak{M}', w' \models \chi_{\mathfrak{M},w}^{c,\ell+1}$ implies in particular that, for each t,

$$\mathfrak{M}', w' \models \lozenge^{\geqslant k_t} \chi_t^{c,\ell}$$

so that $E_i'[w']$ comprises a subset s_t' of size $|s_t'| = k_t = |s_t|$ of worlds in $[\![\chi_t^{c,\ell}]\!]^{\mathfrak{M}'}$. Therefore **II** can respond with $s' := \bigcup_t s_t' \subseteq E_i'[w']$ to make sure that any choice of $u' \in s'$ by **I** allows her a matching choice of some $u \in s$, viz. picking $u \in s_t$ if $u' \in s_t'$. That guarantees $\mathfrak{M}, u \sim_{\#}^{c,\ell} \mathfrak{M}', u'$, and hence a win in the remaining game.

If **I**'s choice is some $s' \subseteq E_i'[w']$ of size $|s| \leqslant c$, we analogously partition s' into its disjoint constituents

$$s_t' := s' \cap [\![\chi_t^{c,\ell}]\!]^{\mathfrak{M}'}$$

and let $k_t := |s_t'| \leqslant c$. Now $\mathfrak{M}', w' \models \chi_{\mathfrak{M},w}^{c,\ell+1}$ implies that $\mathfrak{M}, w \models \lozenge_i^{\geqslant k_t} \chi_t^{c,\ell}$: otherwise the negation of this formula would be a conjunct in the *back*-part of $\chi_{\mathfrak{M},w}^{c,\ell+1}$ whereas $\mathfrak{M}', w' \models \lozenge_i^{\geqslant k_t} \chi_t^{c,\ell}$. Analogous to the above, the union of size k_t-subsets of $s_t \subseteq E_i[w] \cap [\![\chi_t^{c,\ell}]\!]^{\mathfrak{M}}$ serves as a safe response for **II**. $\qquad\square$

Corollary 1.7. *Across all (finite or infinite) pointed Kripke structures over any fixed finite modal signature, every formula $\varphi \in \mathsf{CML}_{c,\ell}$ is logically equivalent to a disjunction over mutually incompatible characteristic formulae $\chi_t^{c,\ell}$ from the finite collection of such characteristic formulae $\chi_t^{c,\ell}$ for all $\sim_{\#}^{c,\ell}$-types t. In particular, $\equiv^{\mathsf{CML}_{c,\ell}}$ and $\sim_{\#}^{c,\ell}$ have finite index for all $c, \ell \in \mathbb{N}$ over the class of all pointed Kripke structures of fixed finite modal signature.*

Corollary 1.8. *The following are equivalent for any class $\mathcal{C}$ of pointed Kripke structures over a finite modal signature and $c, \ell \in \mathbb{N}$:*

   (i) $\mathcal{C}$ is definable by a formula $\varphi \in \mathsf{CML}_{c,\ell}$, i.e.
       $\mathcal{C} = \mathrm{Mod}(\varphi)$ for some $\varphi \in \mathsf{CML}$ of $\mathrm{rk}_{\#}(\varphi) \leqslant c$ and $\mathrm{nd}(\varphi) \leqslant \ell$;

   (ii) $\mathcal{C}$ is closed under $\sim_{\#}^{c,\ell}$.

The same is true of any class $\mathcal{C}$ of finite *pointed Kripke structures over a finite modal signature: $\mathcal{C} = \mathrm{FMod}(\varphi)$ for some $\varphi \in \mathsf{CML}_{c,\ell}$ if, and only if, $\mathcal{C}$ is closed under $\sim_{\#}^{c,\ell}$ within the class of all finite pointed Kripke structures over the same finite signature.*

So, both in the sense of classical and in the sense of finite model theory, definability in CML is equivalent with $\sim_{\#}^{c,\ell}$-invariance for some finite level $c, \ell \in \mathbb{N}$.

Just as for basic modal logic ML in relation to ordinary bisimulation $\sim$ and its finite approximations $\sim^\ell$, this does not immediately relate definability to *bisimulation invariance*. The missing link in both cases is that invariance under the proper *infinitary* bisimulation equivalence ($\sim$ or $\sim_{\#}$) is a much weaker (!) condition than invariance under any one of its much coarser finite approximations ($\sim^\ell$ for $\sim$, or either gradation $\sim_{\#}^\ell$ or $\sim_{\#}^{c,\ell}$ for $\sim_{\#}$). Even in restriction to just finite structures $\sim$ is stronger even than the common refinement of all the finite levels $\sim^\ell$, and the same can be shown for $\sim_{\#}$ in relation to both $\sim_{\#}^\ell$ and $\sim_{\#}^{c,\ell}$.

Just as in the case of basic modal logic and $\sim$, the crux for a characterisation of CML as the $\sim_{\#}$-invariant fragment of FO therefore lies in establishing that for $\varphi \in \mathsf{FO}$, $\sim_{\#}$-invariance implies $\sim_{\#}^{c,\ell}$-invariance for some $c, \ell \in \mathbb{N}$. The crucial difference in these matters arises from the fact that the corresponding assertion for $\sim_{\#}^\ell$ would not suffice: there is no obvious analogue for Corollary 1.8 with $\sim_{\#}^{c,\ell}$ replaced by $\sim_{\#}^\ell$.

2 Characterisation theorems through upgrading

We separately state the two versions of the desired characterisation theorem, one for the setting of classical model theory and one for finite model theory.

Theorem 2.1. *The following are equivalent for any $\varphi(x) \in \mathsf{FO}$ in (the first-order counterpart of) a modal signature:*

  (i) φ is invariant under $\sim_\#$ within the class of all pointed Kripke structures over the modal signature of φ;

 (ii) φ is expressible in graded modal logic CML: $\varphi \equiv \varphi'$ for some $\varphi \in \mathsf{CML}$.

Theorem 2.2. *The following are equivalent for any $\varphi(x) \in \mathsf{FO}$ in (the first-order counterpart of) a modal signature:*

  (i) φ is invariant under $\sim_\#$ within the class of all finite pointed Kripke structures over the modal signature of φ;

 (ii) over the class of all finite pointed Kripke structures φ is expressible in graded modal logic CML: $\varphi \equiv_{\mathrm{fin}} \varphi'$ for some $\varphi \in \mathsf{CML}$.

In both cases, the implications from (ii) to (i) are direct consequences of the Ehrenfeucht–Fraïssé analysis above. In both cases the essential part is the *expressive completeness* claim of CML for $\sim_\#$-invariant first-order properties. And in both cases this claim is equivalent with the assertion that for first-order definable properties of (finite) pointed Kripke structures over a finite modal signature

$$(\ddagger) \qquad \sim_\#\text{-invariance} \quad \text{implies} \quad \sim_\#^{,c,\ell}\text{-invariance for some } c, \ell.$$

As I argued elsewhere in connection with the analogue for ordinary bisimulation $\sim$ and basic modal logic ML, this can be seen as a special *compactness property* for FO, and one that unlike full compactness does not fail in restriction to finite models.

Just like the very similar technique originally proposed for plain ML and $\sim$ in [6, 4], and later extended to several variants e.g. in [7, 1], the approach via an upgrading argument for ($\ddagger$) has the advantage of establishing ($\ddagger$) simultaneously for the classical and the finite model theory version.

First-order logic is known to satisfy strong locality properties over relational structures, as expressed in Gaifman's theorem (cf. eg. [3]). Recall that a first-order formula $\varphi(x)$ (in a single free varaible x) is ℓ-*local* if its semantics only concerns the ℓ-neighbourhood $N^\ell(w)$ of the element w assigned to x in a relational structure $\mathfrak{M}$. Since we are here dealing with relational structures in a modal signature, the ℓ-neighbourhood $N^\ell(w)$ of a world w in $\mathfrak{M} = (W, (E_i^{\mathfrak{M}}), (P_j^{\mathfrak{M}}))$ consists of all worlds $v \in W$ that are at graph distance up to ℓ from w in the *undirected* graph induced by the union of the symmetrisations of all the $E_i^{\mathfrak{M}}$. In this context $\varphi(x)$ is ℓ-local if, for all pointed Kripke structures $\mathfrak{M}, w$

$$\mathfrak{M}, w \models \varphi \;\Leftrightarrow\; \mathfrak{M} \restriction N^\ell(w), w \models \varphi,$$

where $\mathfrak{M} \restriction N^\ell(w)$ refers to the substructure of $\mathfrak{M}$ that is induced on $N^\ell(w) \subseteq W$.

In preparation for the upgrading argument towards ($\ddagger$), consider the evaluation of any particular ℓ-local first-order formula $\varphi^\ell(x)$ at some world w in $\mathfrak{M}$ for which the induced substructure $\mathfrak{M} \restriction N^\ell(w)$ is *tree-like* in the sense that the symmetrisations of the $E_i^{\mathfrak{M}}$

  – are pairwise disjoint, and

  – their union is acyclic

in restriction to $N^\ell(w)$. In particular the undirected graph induced by the union of the accessibility relations E_i in restriction to $N^\ell(w)$ is an undirected tree of depth at most ℓ from its root w, in the graph-theoretic sense.

We call $\mathfrak{M}, w$ *rooted tree-like to depth* ℓ if $\mathfrak{M} \restriction N^\ell(w), w$ is tree-like and the E_i themselves happen to be uniformly directed away from the root w in $\mathfrak{M} \restriction N^\ell(w), w$.

For rooted tree-like $\mathfrak{M} \restriction N^\ell(w), w$, whether or not

$$\mathfrak{M} \restriction N^\ell(w), w \models \varphi^\ell(x)$$

is fully determined by the $\sim_\#^\ell$-type of $\mathfrak{M}, w$ and even by its $\sim_\#^{c,\ell}$-type provided c is chosen large enough w.r.t. the quantifier rank of $q = \mathrm{qr}(\varphi)$ and the number of $\equiv^{\mathsf{FO}_q}$-types (cf. Lemma 2.4 below). The condition that all E_i be directed away from the root is essential because no degree of $\sim_\#^\ell$-equivalence can control in-degrees for E_i-edges, while FO can; and the same goes for disjointness of the E_i and absence of cycles.

Lemma 2.3. *Any $\varphi(x) \in \mathsf{FO}_q$ that is invariant under $\sim_\#$ on all (or just all finite) pointed Kripke structures, is ℓ-local for $\ell = 2^q - 1$ in restriction to all (or just all finite) pointed Kripke structures that are pointed tree-like to depth ℓ. For $\mathfrak{M} \restriction N^\ell(w), w$ is rooted tree-like:*

$$\mathfrak{M}, w \models \varphi \quad \Leftrightarrow \quad \mathfrak{M} \restriction N^\ell(w), w \models \varphi.$$

The proof is strictly analogous to that given for $\sim$-invariant φ in [6, 4]. Essentially, the invariance of φ under *disjoint unions* with other component structures (a fundamental consequence of graded as well as plain bisimulation invariance) reduces the claim to an elementary b&f argument for the q-round Ehrenfeucht–Fraïssé game on

$$\mathfrak{M}', w \quad := \quad q \otimes \mathfrak{M} \quad \oplus \quad \mathfrak{M}, w \quad \oplus \quad q \otimes \mathfrak{M} \restriction N^\ell(w)$$
$$\text{versus} \quad \mathfrak{M}'', w \quad := \quad q \otimes \mathfrak{M} \quad \oplus \quad \mathfrak{M} \restriction N^\ell(w), w \quad \oplus \quad q \otimes \mathfrak{M} \restriction N^\ell(w)$$

where $\oplus$ stands for disjoint union and $q \otimes \mathfrak{M}$ for the disjoint union of q many copies of $\mathfrak{M}$.

Unlike the treatment for basic ML cited above, there is not the immediate short-cut that ℓ-locality together with $\sim_\#$-invariance would imply $\sim_\#^{c,\ell}$-invariance as required for ($\ddagger$); and the obvious analogue for $\sim_\#^\ell$ is not good enough as it does not feed directly into Corollary 1.8; this is what makes the situation slightly more interesting. Rather, we need the following simple fact about FO on tree structures. The proof is again by a standard b&f argument about the classical q-round Ehrenfeucht–Fraïssé game for FO to establish $\simeq_q$.

Lemma 2.4. *For any fixed finite signature and $q, \ell \in \mathbb{N}$ there is $c \in \mathbb{N}$ s.t. for any two rooted tree-like structures $\mathfrak{M} \restriction N^\ell(w), w$ and $\mathfrak{M}' \restriction N^\ell(w'), w'$:*

$$\mathfrak{M} \restriction N^\ell(w), w \sim_\#^{c,\ell} \mathfrak{M}' \restriction N^\ell(w'), w' \quad \Rightarrow \quad \mathfrak{M} \restriction N^\ell(w), w \equiv^{\mathsf{FO}_q} \mathfrak{M}' \restriction N^\ell(w'), w',$$

i.e. for structures that are rooted tree-like to depth ℓ, the $\sim_\#^{c,\ell}$-type fully determines the FO_q-type of their restrictions to depth ℓ, and hence the truth value of any ℓ-local $\varphi(x) \in \mathsf{FO}_q$.

Lemma 2.5 (upgrading lemma).
For $\varphi(x) \in \mathsf{FO}$ that is invariant under $\sim_\#$ over the class of all (or just all finite) pointed Kripke structures, there are $c, \ell \in \mathbb{N}$ s.t. $\varphi(x)$ is invariant under $\sim_\#^{c,\ell}$ over the class of all (or just all finite) pointed Kripke structures. Just as for basic modal logic, the optimal choice for ℓ is $\ell = 2^q - 1$ where q is the quantifier rank of φ.

This then establishes ($\ddagger$) and yields characterisation theorems.

Proof of the lemma. Let $q := \mathrm{qr}(\varphi)$ and $\ell := 2^q - 1$. For pointed Kripke structures $\mathfrak{M}, w$ and $\mathfrak{M}', w'$ over the finite signature of φ, let $\mathfrak{M}^*$ and $\mathfrak{M}'^*$ be the partial directed tree unravellings to depth $\ell + 1$ from roots w and w', merged with copies of the original structures for the continuation beyond that depth; this preserves $\sim_{\#}$ and, where relevant, finiteness of the models in question. Let c be as in Lemma 2.4 for this situation, and assume now that $\mathfrak{M}, w$ and $\mathfrak{M}', w'$, and hence also $\mathfrak{M}^*, w$ and $\mathfrak{M}'^*, w'$, are $\sim_{\#}^{c,\ell}$-equivalent for these values. By Lemma 2.3, $\varphi(x)$ is ℓ-local on $\mathfrak{M}^*, w$ and $\mathfrak{M}'^*, w'$, and by Lemma 2.4, $\mathfrak{M}^* \upharpoonright N^\ell(w), w$ and $\mathfrak{M}'^* \upharpoonright N^\ell(w'), w'$ agree on $\varphi(x)$ since they are FO_q-equivalent. So

$$
\begin{array}{llll}
\mathfrak{M}, w \models \varphi(x) & \Leftrightarrow & \mathfrak{M}^*, w \models \varphi(x) & \text{(by } \sim_{\#}\text{-invariance)} \\
& \Leftrightarrow & \mathfrak{M}^* \upharpoonright N^\ell(w), w \models \varphi(x) & \text{(by } \ell\text{-locality, Lemma 2.3)} \\
& \Leftrightarrow & \mathfrak{M}'^* \upharpoonright N^\ell(w'), w' \models \varphi(x) & \text{(by choice of } c\text{, Lemma 2.4)} \\
& \Leftrightarrow & \mathfrak{M}'^*, w' \models \varphi(x) & \text{(by } \ell\text{-locality, Lemma 2.3)} \\
& \Leftrightarrow & \mathfrak{M}', w' \models \varphi(x) & \text{(by } \sim_{\#}\text{-invariance)}
\end{array}
$$

which establishes $\sim_{\#}^{c,\ell}$-invariance of $\varphi(x)$. $\qquad\square$

Remarks. The classical characterisation of CML of Theorem 2.1 was obtained, in close analogy with van Benthem's classical, compactness- and saturation-based treatment for plain ML in [9], by de Rijke in [2]. The finite model theory version in Theorem 2.2 is there stated as an open problem, without reference to the finite model theory analogue for plain ML due to Rosen [8]. I am not aware of published work filling that gap, and here took the opportunity to adapt the combined and very elementary treatment of the van Benthem–Rosen results from [6, 4] to CML in the most straightforward manner. In following the well established route through an upgrading argument that can be made to work in both classical and finite model theory context, it also retains much of the variability that this approach has shown in applications to other formats for the underlying modal core logic. In particular, natural analogous characterisations obtain for some variants involving global two-way modalities with counting thresholds, and possibly restrictions to special classes of frames, similar to some of the corresponding modification treated e.g. in [7, 1]. Some such were already treated in Rebecca Lukas' diploma thesis [5].

References

[1] A. Dawar and M. Otto. Modal characterisation theorems over special classes of frames. *Annals of Pure and Applied Logic*, 161:1–42, 2009.

[2] M. de Rijke. A note on graded modal logic. *Studia Logica*, 64:271–283, 2000.

[3] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 2nd edition, 1999.

[4] V. Goranko and M. Otto. Model theory of modal logic. In P. Blackburn, J. van Benthem, and F. Wolter, editors, *Handbook of Modal Logic*, pages 249–329. Elsevier, 2007.

[5] R. Lukas. Modallogische Charakterisierungssätze über speziellen Klassen endlicher Strukturen. Diploma thesis, Department of Mathematics, TU Darmstadt, 72 pages, supervisor: M. Otto, 2007.

[6] M. Otto. Elementary proof of the van Benthem–Rosen characterisation theorem. Technical Report 2342, Fachbereich Mathematik, Technische Universität Darmstadt, 2004.

[7] M. Otto. Modal and guarded characterisation theorems over finite transition systems. *Annals of Pure and Applied Logic*, 130:173–205, 2004.

[8] E. Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6:427–439, 1997.

[9] J. van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, Napoli, 1983.