The van Benthem Characterisation Theorem for Descriptive Models

MSc Thesis *(Afstudeerscriptie)*

written by

Tim Henke
(born July 21st, 1996 in Amsterdam, The Netherlands)

under the supervision of **dr. Nick Bezhanishvili**, and submitted to the Board of Examiners in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:	**Members of the Thesis Committee:**
July 2, 2019	prof. dr. Yde Venema (chair)
	dr. Nick Bezhanishvili (supervisor)
	prof. dr. Johan van Benthem
	dr. Alexandru Baltag

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

This thesis investigate the modal and first-order model theory of the class of models over descriptive general frames. Descriptive general frames are Stone spaces with a suitable relation over which every modal logic is complete. The main theorem of this thesis is the van Benthem Characterisation Theorem for the class of descriptive general models. Moreover, a model-theoretic analysis is given to prove that many important results from classical model theory, including the Compactness Theorem for first-order logic and the upward Löwenheim-Skolem Theorem, fail on the class of descriptive general models. The main tool developed in this thesis is the descriptive unravelling, a version of the unravelling tree that is modified to be descriptive. A careful analysis of the operation is provided and three isomorphic constructions are given: a construction through duality theorems, a construction through a topological toolkit based on nets that is also developed, and an explicit construction in terms of finite and infinite paths.

Title: The van Benthem Characterisation Theorem for Descriptive Models
Author: Tim Henke BSc
Date of public defense: July 2, 2019

Institute for Logic, Language and Computation
Universiteit van Amsterdam
Science Park 904, 1098 XH Amsterdam
`https://www.illc.uva.nl/`

Acknowledgements

First and foremost, I would like to thank dr. Nick Bezhanishvili for being my supervisor, reading so many draft versions of the thesis and providing an immense number of comments. He has given me excellent guidance and thanks to him I have enjoyed this project endlessly.
Moreover, I am very grateful to professor Ian Hodkinson for his enormously helpful comments and his extremely sharp eye for errors. His advice and his suggestions have made large portions of the thesis much better.
Next, my thanks goes out to Serge Hartog, Annelene Schulze and Anna Sisák for their continuous support and company during the writing process. My thanks especially to Serge, who provided a much-needed stylistic eye to the thesis.
Finally, I would like to thank the thesis committee for agreeing to read and review my thesis.

Contents

1 Introduction **1**

2 Modal logic, general frames, and bisimulations **5**

2.1 The language of modal logic 5
2.1.1 Syntax and semantics of modal logic 5
2.1.2 Finitary modal logic 7
2.2 General frames 9
2.3 Topology of general frames 11
2.4 Algebraic duality 16
2.4.1 Stone duality 16
2.4.2 Jónsson-Tarski duality 17
2.5 Bisimulations 18
2.5.1 Kripke and Vietoris bisimulations 18
2.5.2 Finite bisimulations 20
2.6 Nets as topological toolkit 22
2.6.1 Nets: basic definitions 23
2.6.2 Topological construction of descriptive completion 28

3 Model theory of descriptive models **35**

3.1 Compactness theorems for descriptive models 35
3.2 Model-theoretic methods for descriptive models 40
3.3 Failure of classical model-theoretic results on descriptive models 43

4 The van Benthem Characterisation Theorem for descriptive models **49**

4.1 The classical van Benthem Characterisation Theorem 49
4.2 Modifying the proof structure 51
4.2.1 Proof structure of classical van Benthem Characterisation Theorem 51
4.2.2 Proof structure of the van Benthem Characterisation Theorem for finite models 52
4.2.3 Proof structure of the van Benthem Characterisation Theorem for descriptive models 55
4.3 The van Benthem Characterisation Theorem for descriptive models 57
4.3.1 Unravelling for descriptive frames 57
4.3.2 The descriptive unravelling 60
4.3.3 Preservation under finite bisimulations 63

5 Further analysis of the descriptive unravelling **69**

5.1 The $\mathfrak{g}_{-\infty}$-operation 69

5.2 The $\mathfrak{g}_{-\infty}$-frame as limit points 73
5.3 The structure of the $\mathfrak{g}_{-\infty}$-frame 78
5.4 A review of the descriptive unravelling 79

6 Conclusions and future work **81**
6.1 Conclusion 81
6.2 Directions for future research 81

Bibliography **83**

1 Introduction

Modal logic extends propositional logic with modal operators that lend themselves to a wide variety of interpretations. Among the diversity, modal logic is used to represent knowledge and belief (known as epistemic modalities), possibilities and necessities (known as alethic modalities), duties and permissions (deontic modalities), for statements about the future and past (known as temporal modalities) and many more applications [8, 11, 26, 4]. The language and its abstractions have also been found widely applicable in theoretical computer science and artificial intelligence [8, 29, 30, 36].

The conventional structures on which modal logic is interpreted are Kripke models. Bisimulations are relations on Kripke models that are indistinguishable by modal logic, in the sense that the truth of modal formulae is preserved under these relations. The celebrated van Benthem Characterisation Theorem states that in fact, any first-order formula that is invariant under bisimulations must be equivalent to a modal formula [1, 8]. This is often formulated more succinctly by saying that modal logic is the bisimulation-invariant fragment of first-order logic. The theorem has inspired many generalisations and alternative versions, including a similar characterisation for intuitionistic logic [31, 32], neighbourhood models [21], and numerous coalgebraic generalisations [34]. Another famous result is the Janin-Walukiewicz Theorem [23], showing that the modal μ-calculus is the bisimulation-invariant fragment of monadic second-order logic.

A result that is of particular importance to this thesis was given in [33], showing that over finite models, too, modal logic is the bisimulation-invariant fragment of first-order logic. The reason why this is particularly remarkable is that the Compactness Theorem of first-order logic features prominently in the proof of the original van Benthem Characterisation Theorem, while the class of finite models crucially lacks the compactness property. The result for finite models has in turn been generalised to multiple other classes lacking the compactness property, including the classes of rooted finite models, rooted transitive models, and in well-founded transitive models [14].

The main result of this thesis will be the van Benthem Characterisation Theorem for the class of models over descriptive general frames. These are totally separated, compact topological spaces, commonly known as Stone spaces, with a suitable relation, whose models have valuations restricted to clopen sets.

Descriptive general frames are an important class of frames with respect to which every modal logic is complete. This is due to a topological duality known as Jónsson-Tarski duality [25], establishing a dual equivalence between the category of descriptive general frames and the category of modal algebras. This duality makes them a crucial tool to study modal algebras.

Apart from this, descriptive general frames are a natural generalisation of finite frames. Topological compactness is a common generalisation of finiteness and the classes of finite and descriptive general frames are both closed under many of finite disjoint unions and p-morphic images, they both are not closed under infinite disjoint unions and they both have the Hennessy-Milner property (see [5, Corollary 3.10]). Moreover, in this thesis, it will be shown that they have a similar model theory.

Finally, descriptive general frames can be viewed as coalgebras for the Vietoris functor on the category of Stone spaces [27], just like Kripke frames are coalgebras for the powerset functor and finite Kripke frames are coalgebras for the powerset functor restricted to finite sets. As far as we are aware, the result in this thesis is the first van Benthem-style characterisation theorem for coalgebras over a category of topological spaces. This may pave the way for generalisations of the theorem for coalgebras of other Vietoris-like functors. These potential generalisations are discussed in Chapter 6.

The following will be a detailed description of the contents of the thesis.

In Chapter 2, descriptive general frames and their models are discussed and analysed, providing background theory and an important topological toolkit. Algebraic duality for general frames is presented and a "descriptive completion"-operation is introduced to turn every general frame into a descriptive general frame. Moreover, the topology of general frames is examined and nets are discussed as a useful tool with which to grapple with descriptive general frames, including two topological constructions of the descriptive completion.

Chapter 3 focuses on the model theory of descriptive models and proves the failure of a number of important results from classical model theory, like the Compactness Theorem for first-order logic and the upward Löwenheim-Skolem Theorem, on the class of descriptive general models. The topological tools presented in Chapter 2 are used to develop a lemma that provides the failure of the different results.

Inspired by the model-theoretic similarities between the classes of finite and descriptive models, Chapter 4 uses the approach from [33] to prove the main result of the thesis: the van Benthem Characterisation Theorem for descriptive models. The first section is dedicated to providing a proof for the classical van Benthem Characterisation Theorem. The section after that discusses the role of the Compactness Theorem for first-order logic in the proof of the classical van Benthem Characterisation Theorem, and how the approach in [33] gets around this. Then a preview will be given for how this approach should be modified for descriptive models, after which the third section produces the complete and detailed proof. The central tool in the proof is the "descriptive unravelling", the descriptive completion of the unravelling tree commonly seen and used in modal logic. Roughly speaking, we add new points to the standard unravelling at "infinite distance" from the points of unravelling, yet guaranteeing that the new model is descriptive. Proving a number of important qualities of this construction then allows for its use in the proof through an Ehrenfeucht-Fraïssé argument.

Chapter 5 is dedicated to a careful examination of the descriptive unravelling, producing an explicit construction of the operation in terms of finite and infinite paths on the original frame and a visualisable classification of the possible shapes that it may take. The equivalent topological construction of the descriptive completion is used to examine what exactly changes when the unravelling of a descriptive general frame is compactified in the descriptive completion and the topological toolkit from Chapter 2 is used to give an explicit isomorphic construction to better understand the structure of the descriptive unravelling. The final result of the chapter is a concrete classification of the possible relational structures of the descriptive unravelling, and it is concluded with a discussion on the uses of the different constructions for the descriptive unravelling.

Finally, Chapter 6 discusses future research directions in which to take these results and how the tools developed in the thesis might be applied in different contexts.

The main original contributions of this thesis are:

- The van Benthem Characterisation Theorem for descriptive models;
- The descriptive unravelling as a tool to understand bisimulation-invariant phenomena on descriptive models;
- A topological toolkit, specifically using nets, to understand descriptive general frames better and two alternative constructions of the descriptive unravelling;
- Failure of a number of model-theoretic results on the class of descriptive models, including specifically the Compactness Theorem for first-order logic, the upward Löwenheim-Skolem Theorem, and the Beth Definability Theorem.

2 Modal logic, general frames, and bisimulations

2.1 The language of modal logic

As stated in the introduction, this thesis will revolve around the language of modal logic, which extends propositional logic with additional "modal operators".

2.1.1 Syntax and semantics of modal logic

In this thesis, only the uni-modal language will be considered. This choice is made mainly to improve legibility.

Definition 2.1. The *modal language* over propositional variables from a set P of any cardinality is defined as all formulae that may be constructed in the recursive schema

$$\varphi \quad ::= \quad p \;\big|\; \top \;\big|\; \bot \;\big|\; \neg\varphi \;\big|\; \varphi \wedge \varphi \;\big|\; \varphi \vee \varphi \;\big|\; \Diamond\varphi \;\big|\; \Box\varphi$$

where $p \in \mathsf{P}$. This set of formulae will be referred to as $\mathtt{ML}(\mathsf{P})$ or $\mathtt{ML}$, if the set of propositional variables is clear from context. ◀

Modal formulae are interpreted on relational structures through the standard Kripke semantics.

Definition 2.2. A *Kripke frame* is a pair $\mathfrak{F} = (W, R)$ of a non-empty set W and a relation $R \subseteq W \times W$. A *Kripke model* over a set of propositional variables P is a triple $\mathfrak{M} = (W, R, V)$ where (W, R) is a Kripke frame and $V : \mathsf{P} \to \mathscr{P}(W)$ is a function. The class of all Kripke models will be denoted by $\mathscr{K}$. Sometimes Rwv or wRv will be used as alternative notation for $(w, v) \in R$. ◀

Definition 2.3. Let (W, R) be a Kripke frame and let $X \subseteq W$. Define

$$\begin{aligned} \langle R\rangle X &:= \{w \in W \mid \exists x \in X : Rwx\}; \\ [R]X &:= \{w \in W \mid \forall v \in W : Rwv \implies v \in X\}. \end{aligned}$$

These are dual notions in the sense that $\langle R\rangle X = ([R]X^c)^c$ and therefore also $[R]X = (\langle R\rangle X^c)^c$, where $X^c := W \setminus X$ is the set-theoretic complement. ◀

Definition 2.4. For a Kripke model $\mathfrak{M} = (W, R, V)$ the *truth function* $[\![\cdot]\!]^{\mathfrak{M}}$ from $\mathtt{ML}(\mathsf{P})$ to $\mathscr{P}(W)$ is defined through recursion on the complexity of formulae by

$$\begin{aligned} [\![\top]\!]^{\mathfrak{M}} &:= W; & [\![\bot]\!]^{\mathfrak{M}} &:= \emptyset; \\ [\![p]\!]^{\mathfrak{M}} &:= V(p) & [\![\neg\varphi]\!]^{\mathfrak{M}} &:= W \setminus [\![\varphi]\!]^{\mathfrak{M}}; \\ [\![\varphi \wedge \psi]\!]^{\mathfrak{M}} &:= [\![\varphi]\!]^{\mathfrak{M}} \cap [\![\psi]\!]^{\mathfrak{M}}; & [\![\varphi \vee \psi]\!]^{\mathfrak{M}} &:= [\![\varphi]\!]^{\mathfrak{M}} \cup [\![\psi]\!]^{\mathfrak{M}}; \\ [\![\Diamond\varphi]\!]^{\mathfrak{M}} &:= \langle R\rangle[\![\varphi]\!]^{\mathfrak{M}}; & [\![\Box\varphi]\!]^{\mathfrak{M}} &:= [R][\![\varphi]\!]^{\mathfrak{M}}. \end{aligned}$$

If $\mathfrak{M} = (W, R, V)$ is a model and $\varphi \in \mathtt{ML}(\mathsf{P})$, then for each point $w \in W$ the formula φ is said to be *satisfied* at w or at $\mathfrak{M}, w$ if $w \in [\![\varphi]\!]^{\mathfrak{M}}$. ◀

Convention 2.5. The standard models on which to evaluate modal formulae are pointed Kripke models. That is, pairs of the form $(\mathfrak{M}, w)$ where $\mathfrak{M} = (W, R, V)$ is a Kripke model and $w \in W$ is a point. The brackets will be included only if omitting them is ambiguous. Then $\mathfrak{M}, w \Vdash \varphi$ is understood to mean that $\mathfrak{M}, w$ satisfies φ, that is $w \in [\![\varphi]\!]^{\mathfrak{M}}$. Contrast this with the symbol for semantic consequence for first-order logic ($\mathtt{FOL}$), where a model $\mathfrak{A}$ validating a formula α is written as $\mathfrak{A} \models \alpha$.

The modal theory of a pointed model $\mathfrak{M}, w$ is the set of modal formulae satisfied by it and is written as $\mathrm{Th}_{\mathtt{ML}}(\mathfrak{M}, w)$. The equality $\mathrm{Th}_{\mathtt{ML}}(\mathfrak{M}, w) = \mathrm{Th}_{\mathtt{ML}}(\mathfrak{N}, v)$ is abbreviated as $\mathfrak{M}, w \leftrightsquigarrow \mathfrak{N}, v$, which is to be contrasted with the symbol for first-order equivalence, $\mathfrak{A} \equiv^{\mathtt{FOL}} \mathfrak{B}$, where the label of $\mathtt{FOL}$ is left explicit to avoid confusion.

Finally, for a given language $\mathcal{L}$, and a formula ξ (not necessarily in the same language), one can write $\mathrm{Th}^{\mathcal{C}}_{\mathcal{L}}(\xi)$ for the set of all formulae ζ in $\mathcal{L}$ such that if $\mathfrak{A}$ is a model in $\mathcal{C}$ that validates ξ, then $\mathfrak{A}$ also validates ζ. Observe that this makes sense only if ξ and ζ have the same type of models, and $\mathcal{C}$ consists of such models. If the class $\mathcal{C}$ is obvious from context, it may be omitted from the notation.

For a final notational convention, if $\xi \in \mathrm{Th}^{\mathcal{C}}_{\mathcal{L}}(\zeta)$ and $\zeta \in \mathrm{Th}^{\mathcal{C}}_{\mathcal{L}'}(\xi)$, then one can write $\xi \equiv^{\mathcal{C}} \zeta$, where again $\mathcal{C}$ may be omitted if it is redundant.

With this notion of equivalence of formulae from different languages, it turns out that $\mathtt{ML}$ can be thought of as a sublanguage of first-order language, $\mathtt{FOL}$. There exists a standard inclusion of the former into the latter, known as the standard translation.

Definition 2.6. Given two <u>distinct</u> variables x_0 and x_1 in $\mathtt{FOL}$, the *standard translation* $\mathrm{ST}_{x_i} : \mathtt{ML} \to \mathtt{FOL}$ for $i \in \{0, 1\}$ is defined through the joint recursive schema

$$\begin{aligned} &\mathrm{ST}_{x_i}(p) = Px_i; && \mathrm{ST}_{x_i}(\neg\varphi) = \neg\,\mathrm{ST}_{x_i}(\varphi); \\ &\mathrm{ST}_{x_i}(\varphi \vee \psi) = \mathrm{ST}_{x_i}(\varphi) \vee \mathrm{ST}_{x_i}(\psi); && \mathrm{ST}_{x_i}(\varphi \wedge \psi) = \mathrm{ST}_{x_i}(\varphi) \wedge \mathrm{ST}_{x_i}(\psi); \\ &\mathrm{ST}_{x_i}(\Diamond\varphi) = \exists x_{1-i}[Rx_ix_{1-i} \wedge \mathrm{ST}_{x_{1-i}}(\varphi)]; && \mathrm{ST}_{x_i}(\Box\varphi) = \forall x_{1-i}[Rx_ix_{1-i} \to \mathrm{ST}_{x_{1-i}}(\varphi)]. \end{aligned}$$

The variables x_0 and x_1 are also often replaced by x and y. ◀

The proposition below will show why this is a meaningful translation.

Proposition 2.7. For any pointed model $\mathfrak{M}, w$ the two formulae $\varphi \in \mathtt{ML}$ and $\mathrm{ST}_{x_0}(\varphi)$ are equivalent in the previously described sense:

$$\mathfrak{M}, w \Vdash \varphi \iff \mathfrak{M} \models \mathrm{ST}_{x_0}(\varphi)\left[{}^{w}\!/_{x_0}\right].$$

Proof. The proof is a simple induction on the complexity of formulae. $\square$

Borrowing notation from topology and writing $\mathscr{K}_*$ for the class of pointed Kripke models, and using Convention 2.5, this becomes $\varphi \equiv^{\mathscr{K}_*} \mathrm{ST}_{x_0}(\varphi)$, where first-order formulae in one free variable are interpreted on pointed Kripke models in the obvious way.

2.1.2 Finitary modal logic

A pointed Kripke model $\mathfrak{M}, w$ validates a modal formula such as $\overbrace{\Diamond \cdots \Diamond}^{m \text{ times}} \varphi$ (usually abbreviated to $\Diamond^m \varphi$) if and only if there exists an R-path from w in $\mathfrak{M}$ of length m at the end of which φ is satisfied. As there is no restriction on the size of m, sets of modal formulae can be used to characterise behaviour of all points that can be reached from w, in the entire connected component.

While this expressive power is extremely useful, it hinders the analysis of the language. In particular, as will be demonstrated in Proposition 2.12, it means the language contains infinitely many non-equivalent formulae. However, the modal language can be decomposed in finitary sublanguages that are more well-behaved. This tool will be vital to proof of the main theorem of this thesis.

Definition 2.8. The *finite modal language up to depth n* over propositional variables P, written as $\mathtt{ML}_n(\mathsf{P})$, will be constructed by the recursive schema

$$\varphi \quad ::= \quad p \mid \top \mid \bot \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \Diamond\psi \mid \Box\psi$$

where $p \in \mathsf{P}$ and $\psi \in \mathtt{ML}_k(\mathsf{P})$ for $k < n$. Formulae in $\mathtt{ML}_n(\mathsf{P}) \setminus \bigcup_{i<n} \mathtt{ML}_i(\mathsf{P})$ are said to have modal depth n. ◀

As the next proposition will show, this is indeed a decomposition of the full language.

Proposition 2.9. For any set of propositional variables P the equality

$$\mathtt{ML}(\mathsf{P}) = \bigcup_{n \in \mathbb{N}} \mathtt{ML}_n(\mathsf{P}) \tag{2.1}$$

holds.

Proof. Induction on n immediately gives that $\mathtt{ML}_n(\mathsf{P}) \subseteq \mathtt{ML}(\mathsf{P})$ for all n. Induction on the complexity of formulae in $\mathtt{ML}$ gives the other inclusion. $\square$

All these languages contain infinitely many syntactically different formulae.

Corollary 2.10. $|\mathtt{ML}(\mathsf{P})| = |\mathtt{ML}_n(\mathsf{P})| = \aleph_0 \cdot |\mathsf{P}|$.

Proof. The second equality is given by induction on n. There are $\aleph_0 \cdot \kappa$ Boolean expressions over κ propositional variables, so $|\mathtt{ML}_0(\mathsf{P})| = \aleph_0 \cdot |\mathsf{P}|$. Viewing $\mathtt{ML}_{n+1}(\mathsf{P})$ as Boolean expressions over P and $\{\heartsuit\varphi \mid \varphi \in \mathtt{ML}_n, \heartsuit \in \{\Diamond, \Box\}\}$, the cardinality is given by

$$|\mathtt{ML}_{n+1}| = \aleph_0 \cdot (|\mathtt{ML}_n(\mathsf{P})| + |\mathsf{P}|) = \aleph_0 \cdot (\aleph_0 \cdot |\mathsf{P}| + |\mathsf{P}|) = \aleph_0^2 \cdot |\mathsf{P}| = \aleph_0 \cdot |\mathsf{P}|$$

by the induction hypothesis. Moreover,

$$\aleph_0 \cdot |\mathsf{P}| = |\mathtt{ML}_0(\mathsf{P})| \leq |\mathtt{ML}(\mathsf{P})| = \left|\bigcup_{n\in\mathbb{N}} \mathtt{ML}_n(\mathsf{P})\right| \leq \sum_{n\in\mathbb{N}} |\mathtt{ML}_n(\mathsf{P})| = \aleph_0^2 \cdot |\mathsf{P}| = \aleph_0 \cdot |\mathsf{P}|,$$

concluding the proof. □

However, while there are infinitely many syntactically different formulae in these languages, the next proposition will show that the finitary modal languages only have finitely many formulae up to equivalence. This means that theories in these languages can be reduced to a single formula, providing much stronger results.

Proposition 2.11. If P is a finite set of propositional variables, and $n \in \mathbb{N}$ is a finite, natural number, then $\mathtt{ML}_n(\mathsf{P})$ contains only finitely many formulae up to equivalence as in Convention 2.5. That is, there is a finite set of formulae $\{\varphi_0, \ldots, \varphi_k\} \subseteq \mathtt{ML}_n(\mathsf{P})$ with the property that for any $\varphi \in \mathtt{ML}_n(\mathsf{P})$ there is a non-negative $i \leq k$ such that $\varphi \equiv \varphi_i$.

Proof. By induction on n. Suppose that the statement is true for all $\ell < n$. Observe that $\mathtt{ML}_n(\mathsf{P})$ consists of all Boolean polynomials in $\mathsf{P} \cup \{\Diamond\psi, \Box\psi \mid \psi \in \mathtt{ML}_\ell(\mathsf{P}), \ell < n\}$. By assumption and the induction hypothesis, this is finite up to equivalence, and since there are only finitely many Boolean expressions up to equivalence over a finite set, it follows that $\mathtt{ML}_n(\mathsf{P})$ is finite up to equivalence. □

Contrast this with the original modal language, which does not have this property.

Proposition 2.12. The language $\mathtt{ML}(\mathsf{P})$ contains $\aleph_0 \cdot |\mathsf{P}|$ many formulae that are pairwise non-equivalent over the class of pointed Kripke models.

Proof. Corollary 2.10 implies immediately that there are at most $\aleph_0 \cdot |\mathsf{P}|$ many such formulae. To see that there are at least so many non-equivalent formulae, note first that the formulae p and q are evidently non-equivalent for distinct $p, q \in \mathsf{P}$. Secondly, any model $\mathfrak{M}$ based on $\mathbb{N}$ with the predecessor relation, illustrated in Figure 2.1, has the property that $\mathfrak{M}, n \Vdash \Diamond^n\top$, but $\mathfrak{M}, n \nVdash \Diamond^{n+1}\top$, implying that $\{\Diamond^n\top\}_{n\in\mathbb{N}}$ are all non-equivalent. This guarantees that there are at least $\max\{\aleph_0, |\mathsf{P}|\} = \aleph_0 \cdot |\mathsf{P}|$ many non-equivalent formulae in $\mathtt{ML}$. □

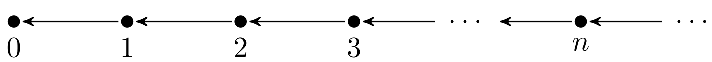

Figure 2.1: The frame given by $\mathbb{N}$ equipped with the predecessor relation.

Equivalence on these finitary languages will come back often in this thesis.

Definition 2.13. Let $\ell \in \mathbb{N}$. Two pointed models $\mathfrak{M}, w$ and $\mathfrak{N}, v$ will be called *modally ℓ-equivalent* if for all $\varphi \in \mathtt{ML}_\ell(\mathsf{P})$:

$$\mathfrak{M}, w \Vdash \varphi \text{ if and only if } \mathfrak{N}, v \Vdash \varphi.$$

This is denoted by $\mathfrak{M}, w \leftrightsquigarrow_\ell \mathfrak{N}, v$ or $\mathfrak{M}, w \equiv_\ell^{\mathtt{ML}} \mathfrak{N}, v$ over P. If the set of propositional variables is clear from context, it may be omitted. ◀

2.2 General frames

Investigating the class of models based on frames with a fixed set of graph-theoretic or relational properties has yielded a rich theory of normal modal logics [8, 11, 26]. These are sets of modal formulae closed under reasonable rules and axioms, see also [8, Def. 1.39] for more details. Typical soundness and completeness results in modal logic state that a normal modal logic Λ is sound and complete with respect to all models based Kripke frames in a fixed class.

There are some normal modal logics, however, which are not sound and complete with respect to any class of Kripke frames [9, 2, 38]. In fact, some are not even sound with respect to any Kripke frame [3]. In these cases, requiring validity on all models based on frame is too strict. A weaker requirement is to demand validity on all models whose valuations are restricted to a given collection of "admissible sets".

Definition 2.14. A triplet $\mathfrak{g} = (W, R, A)$ is a *general frame* if (W, R) is a Kripke frame and A is a field of sets[1] over W that is closed under the operation $\langle R \rangle$ (i.e. for every $a \in A$ also $\langle R \rangle a \in A$). The underlying frame of $\mathfrak{g}$ will be denoted by $\mathfrak{g}_{\#} := (W, R)$.

A quadruplet $\mathfrak{m} = (W, R, A, V)$ is a *general model* based on $\mathfrak{g} = (W, R, A)$ over a set of propositional variables P if $V : \mathsf{P} \to A$ is a function. The definition of $[\![\cdot]\!]^{\mathfrak{m}}$ is the same for general models as it is for Kripke models. ◀

Now soundness and completeness with respect to a class of general frames is validity on all models whose valuations are contained in the collection of admissible sets. With this notion of soundness and completeness, it turns out that in fact all normal modal logics are sound and complete with respect to some class of general frames.

To motivate the closure conditions on the collection of admissible sets, the next proposition will show that the truth set of any formula is in fact an admissible set.

Proposition 2.15. Let $\mathfrak{m} = (W, R, A, V)$ be a general model and $\varphi \in \mathtt{ML}$ a formula. Then $[\![\varphi]\!]^{\mathfrak{m}} \in A$.

Proof. By induction on the complexity of φ. □

When algebraic notions are introduced in Section 2.4, it will become clear that some specific types of general frames are of special importance. It will be useful to define a few important classes of general frames that are relevant to this thesis.

Definition 2.16. Let $\mathfrak{g} = (W, R, A)$ be a general frame. Then

- $\mathfrak{g}$ is called *differentiated* if for all distinct $w, v \in W$ there exists an $a \in A$ such that $v \in a$ but $w \notin a$;
- $\mathfrak{g}$ is *tight* if for all $w, v \in W$ with $(w, v) \notin R$ there exists an $a \in A$ such that $v \in a$ and $w \notin \langle R \rangle a$;

[1] That is, a non-empty collection of sets that is closed under union, intersection and complementation.

- $\mathfrak{g}$ is called *compact* if all collections $\mathcal{A} \subseteq A$ with empty intersection have a finite subcollection $\mathcal{A}_0 \subseteq \mathcal{A}$ with empty intersection.
- $\mathfrak{g}$ is called (*pre*)*image-compact* if for any point $w \in W$, a collection $\mathcal{A} \subseteq A$ whose intersection is disjoint from $R[w]$ ($R^{-1}[w] = \langle R \rangle \{w\}$), the set of successors (predecessors) of w, has a finite subcollection $\mathcal{A}_0 \subseteq \mathcal{A}$ whose intersection is also disjoint from $R[w]$ ($R^{-1}[w]$).
- $\mathfrak{g}$ is called *descriptive* if it is differentiated, tight, and compact. The class of models based on descriptive frames will be denoted by $\mathfrak{D}$.

Models based on such a general frame are called *differentiated, tight, compact,* or *descriptive models.* ◀

Image-compact frames and models will show up regularly in this thesis.

Remark 2.17. From the De Morgan laws and the fact that A is closed under complementation, it is easy to see that a general frame (W, R, A) is compact if and only if each collection $\mathcal{A} \subseteq A$ such that $\bigcup \mathcal{A} = W$ has a finite subcollection $\mathcal{A}_0 \subseteq \mathcal{A}$ with $\bigcup \mathcal{A}_0 = W$. ◀

Example 2.18. To perhaps clarify some of these notions, it might be useful to consider some examples and counterexamples.

- The most trivial example of a descriptive frame are the unique general frame consisting of a single reflexive point and the unique general frame of a single irreflexive point. These are unique, because there is only one field of sets on a singleton set.
- A less trivial common example of a descriptive frame is $(\omega+1, \geq)$ or $(\omega+1, \leq)$ with admissible sets all finite subsets of ω and their complements in $\omega + 1 = \mathbb{N} \cup \{\omega\}$. The properties are easily checked. They are schematically illustrated in Figure 2.2a and 2.2b. Note that not all arrows in $\mathbb{N}$ are displayed to avoid clutter.
- An example of a compact and differentiated frame that is not tight is $(\omega+1, >)$, as schematically displayed in Figure 2.2c (again with only successor arrows in $\mathbb{N}$), with the same admissible sets. This is not tight, because any admissible set containing ω also contains finite numbers by construction, so that ω is again in the $>$-preimage. In Section 2.6.1, it will turn out that there is a much more intuitive reason why this frame is not tight. The frame $(\omega + 1, <)$ with these admissible sets is not a general frame, because the $<$-preimage of the whole set is $\mathbb{N}$, which is not admissible.
- For a non-compact frame that is tight and differentiated, one can take any infinite frame with the powerset as collection of admissible sets.
- To get a non-differentiated frame, one can take any (differentiated) general frame, pick an arbitrary point and copy it, where both copies inherit the relations of the original and each admissible set gets both copies if and only if it had the original point, otherwise it contains neither. This operation preserves both tightness and compactness, but the result will never be differentiated, because the two copies cannot be distinguished.

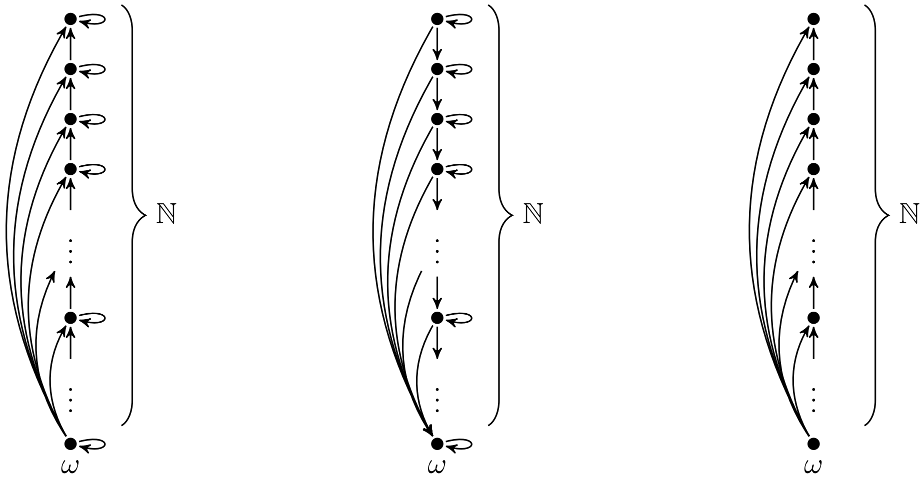

(a) Schematic illustration of $(\omega + 1, \geq)$.

(b) Schematic illustration of $(\omega + 1, \leq)$.

(c) Schematic illustration of $(\omega + 1, >)$.

Figure 2.2: Schematic illustrations of $\omega + 1$ with relations $\geq$, $\leq$ and $>$ respectively. Note that only the successor relations are drawn in the points from $\mathbb{N}$ to avoid cluttering. However, the relation intended is the total order.

2.3 Topology of general frames

An extremely useful lens through which to view general frames is that of topology. General frames can be equipped with an extremely natural topology which provides an immense toolbox for analysing them. References for this subject can be found in [8, 11]. For the sake of keeping this thesis self-contained, all relevant theory is formally stated in this section.

Definition 2.19. Let $\mathfrak{g} = (W, R, A)$ be a general frame. Then the topological space $\mathcal{T}(\mathfrak{g})$ with universe W and topology whose basis[2] is A will be called the $\mathfrak{g}$-*space*. ◀

Note that A is closed under intersections, so that it is in fact a basis. Moreover, because A is closed under complementation, this is in fact a basis of clopens (sets that are both closed and open). It turns out that this topological view of general frames is very powerful, and it will be used extensively throughout the thesis. The distinction between a general frame $\mathfrak{g}$ and the associated $\mathfrak{g}$-space will not always be made explicitly, and topological properties of general frames are to be understood as properties of the associated space.

To understand why the topology is a meaningful construction on a general frame, it is instructive to turn to the morphisms. The morphisms on topological spaces are called continuous functions.

[2]Recall that a basis is a collection of sets $\mathcal{B}$ such that for each $x \in B_0 \cap B_1$ with $B_0, B_1 \in \mathcal{B}$ there is a B_2 with $x \in B_2 \subseteq B_0 \cap B_1$. The topology generated by a basis $\mathcal{B}$ is given by the collection of all unions of subsets of $\mathcal{B}$.

Definition 2.20. Let $\mathcal{X} = (X, \mathcal{T})$ And $\mathcal{X}' = (X', \mathcal{T}')$ be two topological spaces. A function $f : X \to X'$ is called *continuous* if for each open $U \in \mathcal{T}'$ its inverse image under f is open, i.e. $f^{-1}[U] \in \mathcal{T}$. ◀

The connection with general frames is then understood through the action of continuous functions on the admissible sets.

Proposition 2.21. Let $\mathfrak{g} = (W, R, A)$ and $\mathfrak{g}' = (W', R', A')$ be two general frames. If $f : W \to W'$ generates a morphism of fields of sets $\mathscr{P}^{\mathrm{op}} f : A' \to A$ sending a to $f^{-1}[a]$, then f is continuous as map from $\mathcal{T}(\mathfrak{g})$ to $\mathcal{T}(\mathfrak{g}')$. If A is also the collection of all clopen sets on $\mathcal{T}(\mathfrak{g})$, then the converse implication holds, too.

Proof. If $U = \bigcup \mathcal{A}$ is open for $\mathcal{A} \subseteq A'$, then

$$f^{-1}[U] = f^{-1}\left[\bigcup \mathcal{A}\right] = \bigcup \{f^{-1}[a] \mid a \in \mathcal{A}\}$$

is open, as $\{f^{-1}[a] \mid a \in \mathcal{A}\} \subseteq A$ by assumption.

Moreover, f is continuous and $a \in A'$, then a and a^c are both open by construction of the topology so that $f^{-1}[a]$ and $f^{-1}[a^c] = f^{-1}[a]^c$ are both open, implying that $f^{-1}[a]$ is clopen. If A is the collection of all clopens, this means $\mathscr{P}^{\mathrm{op}} f$ maps to A and thus that is a morphism of fields of sets. □

This topological perspective provides a good reason for redefining classes of general frames in topological terms.

Definition 2.22. A topological space is called *totally separated* if every two distinct points have disjoint, clopen neighbourhoods. ◀

Proposition 2.23. Let $\mathfrak{g}$ be a general frame. Then $\mathfrak{g}$ is differentiated if and only if the $\mathfrak{g}$-space is totally separated.

Proof. The implication from left to right is obvious. For the implication from right to left, consider two distinct points, w and v. There must be a clopen set C such that $w \in C$ and $v \notin C$. A clopen set is in particular open, so that there is a basis element $a \subseteq C$ with $w \in a$. From $v \notin C$ we infer $v \notin a$ so that a is an admissible set containing one but not the other, proving differentiatedness. □

Definition 2.24. A topological space is called *Hausdorff* if every two distinct points have disjoint open neighbourhoods. ◀

Corollary 2.25. If a general frame $\mathfrak{g}$ is differentiated, then the $\mathfrak{g}$-space is Hausdorff.

Proof. This follows immediately from Proposition 2.23 and the fact that clopen sets and their complements are open. □

Definition 2.26. Let $\mathcal{X} = (X, \mathcal{T})$ be a topological space. The space $\mathcal{X}$ is said to be *compact* if for every collection of open sets $\mathcal{U} \subseteq \mathcal{T}$ such that $\bigcup \mathcal{U} = X$ there is a finite subcover $\mathcal{U}_0 \subseteq \mathcal{U}$ such that $\bigcup \mathcal{U}_0 = X$. A subset $C \subseteq X$ is said to be compact if the subspace it induces is compact. ◀

Proposition 2.27. Let $\mathfrak{g} = (W, R, A)$ be a general frame. Then $\mathfrak{g}$ is compact if and only if the $\mathfrak{g}$-space is compact in the topological sense.

Proof. For the left to right implication, $\mathcal{U}$ is an open cover of W. Then each element $U \in \mathcal{U}$ is a union of basis elements $U = \bigcup \mathcal{A}_U$ for some $\mathcal{A}_U \subseteq A$. Consider the collection

$$\mathcal{A} := \bigcup \{\mathcal{A}_U \mid U \in \mathcal{U}\} \subseteq A.$$

Then evidently $\bigcup \mathcal{A} = \bigcup \mathcal{U} = W$. By compactness, there is a finite subset $\mathcal{A}_0 \subseteq \mathcal{A} = \bigcup \{\mathcal{A}_U \mid U \in \mathcal{U}\}$ that covers W. So for each element $a \in \mathcal{A}_0$, there is a $U_a \in \mathcal{U}$ such that $a \in \mathcal{A}_U$, implying that $a \subseteq U_a$. From finite choice, this gives a finite subcover $\mathcal{U}_0 = \{U_a \mid a \in \mathcal{A}_0\} \subseteq \mathcal{U}$.

For the right to left implication, it is sufficient to note that all $a \in A$ are open. □

Corollary 2.28. A general frame $\mathfrak{g} = (W, R, A)$ is (pre)image-compact if and only if for each $w \in W$ the subset $R[w]$ $(R^{-1}[w])$ is compact as subset of the $\mathfrak{g}$-space.

Proof. This follows immediately from Proposition 2.27 when viewing $R[w]$ $(R^{-1}[w])$ with the restricted relation and admissible sets as a general frame. □

Definition 2.29 ([13, 24])**.** A topological space is called a *Stone space* if it satisfies one of the following equivalent properties:

- It is compact and totally separated.
- It is compact, Hausdorff, and has a basis of clopen sets. ◀

Remark 2.30. As can be seen immediately from Propositions 2.23 and 2.27, a general frame $\mathfrak{g}$ is differentiated and compact if and only if the $\mathfrak{g}$-space is a Stone space. ◀

This topological terminology provides the opportunity to talk about subsets of W being compact.

For sufficiently well-behaved topological spaces, there is a convenient correspondence between closedness and compactness of subsets.

Lemma 2.31. If $\mathcal{X} = (X, \mathcal{T})$ is a Hausdorff space, then any compact subset $C \subseteq X$ is closed. If a (not necessarily Hausdorff) topological space $\mathcal{X}$ is compact, then any closed subset is also compact.

This is a well-known result in topology, but the proof will be stated briefly in an effort to keep the thesis more self-contained.

Proof of 2.31. Let C be compact. By the Hausdorff property, for every $x \in C$ and $y \notin C$, there exist disjoint open neighbourhoods U_x^y and V_y^x such that $x \in U_x^y$ and $y \in V_y^x$. Fix a $y \notin C$ and consider $\mathcal{U}_y := \{U_x^y \mid x \in C\}$. As $x \in U_x^y$, it follows that $\bigcup \mathcal{U}_y \supset C$, so $\mathcal{U}_y$ is an open cover of C. By compactness, it has a finite subcover $\{U_{x_1}^y, \ldots, U_{x_n}^y\}$. This implies that $V_y := V_y^{x_1} \cap \cdots \cap V_y^{x_n}$ is disjoint from $U_{x_1}^y \cup \cdots \cup U_{x_n}^y \supseteq C$, so that y has an open neighbourhood V_y disjoint from C. As $y \in V_y^x$ for all y, x, it follows that

$$X \setminus C = \bigcup_{y \in X \setminus C} V_y$$

must be open, making C closed.

If $\mathcal{X}$ is compact and $C \subseteq X$ is a closed subset, then any open cover of C can be turned into an open cover of $\mathcal{X}$ by adding $X \setminus C$. Compactness of X means this has a finite subcover, which must then also cover C. As clearly $X \setminus C$ does not cover C, this must be a finite subcover of the original cover of C. □

As such, in compact Hausdorff spaces, closed and compact subsets are interchangeable.

Lemma 2.32. If $\mathfrak{g} = (W, R, A)$ is a tight general frame, then $\mathfrak{g}$ is preimage-closed.

Proof. Let $w \in W$ be a point. Then

$$R^{-1}[x] = \bigcap_{\substack{a \in A \\ x \in a}} \langle R \rangle a.$$

The left to right inclusion is obvious, since $\{x\} \subseteq a \implies R^{-1}[x] = \langle R \rangle \{x\} \subseteq \langle R \rangle a$. Tightness provides the other inclusion through contraposition: suppose that $y \notin R^{-1}[x]$, then there exists an $a \in A$ such that $x \in a$ but $y \notin \langle R \rangle a \supseteq \bigcap \{\langle R \rangle b \mid b \in A, x \in b\}$. □

Proposition 2.33. If $\mathfrak{g}$ is a compact and tight general frame, then $\mathfrak{g}$ is preimage-compact. In particular, all descriptive frames are preimage-compact.

Proof. Let w be a point. Then $R^{-1}[w]$ is closed by the above lemma, implying that it is compact, because a closed subset of a compact space is again compact by Lemma 2.31. □

Lemma 2.34. If (W, R, A) is a descriptive frame and $X \subseteq W$ is closed, or equivalently compact, then $R[X]$ is closed, and therefore compact.

Proof. Note that descriptive frames are compact and Hausdorff by Proposition 2.27 and Corollary 2.25, so that Lemma 2.31 shows that closed and compact are the same. It is thus sufficient to show that the R-image of a closed set is closed. To this end let X be closed. The following argument will show that $R[X]$ is equal to its closure, that is

$$R[X] = \overline{R[X]} := \bigcap \{a \in A \mid R[X] \subseteq a\}.$$

As all sets are included in their closure, the left to right inclusion is trivial. For the other inclusion, assume that $w \notin R[X]$. Let $x \in X$. By tightness, there exists an $a_x \in A$ such that $w \in a_x$ but $x \notin \langle R \rangle a_x$. In particular, $x \in W \setminus \langle R \rangle a_x = [R] a_x^c$. Hence there is an open cover of X given by $\{[R]a_x^c \mid x \in X\}$. Because X is compact, there is a finite subcover $\{[R]a_{x_i}^c\}_{i=1}^n$ for some n. Then

$$X \subseteq \bigcup_{i=1}^{n} [R]a_{x_i}^c \subseteq [R] \bigcup_{i=1}^{n} a_{x_i}^c \implies R[X] \subseteq \bigcup_{i=1}^{n} a_{x_i}^c.$$

Note that the latter is a finite union of clopen sets, hence clopen. As such, it follows that also $\overline{R[X]} \subseteq \bigcup_{i=1}^n a_{x_i}^c$. Since $w \notin a_x^c$ for any $x \in X$, it follows that $w \notin \overline{R[X]}$. □

Proposition 2.35. Descriptive frames are image-compact.

Proof. Note that singletons are finite, hence compact. Then Lemma 2.34 gives the conclusion. □

In fact, this is an exact characterisation of descriptive frames.

Definition 2.36 ([7, 27])**.** A *modal space* is a triple $(X, R, \mathcal{T})$ such that $(X, \mathcal{T})$ is a Stone space, R is image-compact and the R-preimage set of clopen sets is clopen. If R is a partial order, then this is also called an *Esakia space.* ◀

Proposition 2.37. There is a one-to-one correspondence between descriptive frames and modal spaces given by taking the topology generated by the admissible sets on the one hand, and taking the clopen sets of the topology in the other direction.

That is, if $\mathcal{T}_A$ is the topology generated by a basis A and $\mathbf{Clop}(\mathcal{T})$ is the clopens of a topology $\mathcal{T}$ then sending a descriptive frame (W, R, A) to $(W, R, \mathcal{T}_A)$ and sending a modal space $(X, R, \mathcal{T})$ to $(X, R, \mathbf{Clop}(\mathcal{T}))$ provide a one-to-one correspondence in the sense that $\mathbf{Clop}(\mathcal{T}_A) = A$ and $\mathcal{T}_{\mathbf{Clop}(\mathcal{T})} = \mathcal{T}$.

Proof. Given a descriptive frame $\mathfrak{g} = (W, R, A)$, the $\mathfrak{g}$-space is a Stone space from Remark 2.30, the R-image of a point is compact by Lemma 2.34, and the preimage of a clopen set is clopen. The last claim follows because any clopen set is an admissible set, since closed subsets of a compact space are compact, and being open they must be a union of basis elements. As such, any clopen set is a finite union of admissible sets, making it an admissible set, which are clopen and closed under R-preimages.

On the other hand, if $(X, \mathcal{T}, R)$ is a modal space and $A := \mathbf{Clop}(X, \mathcal{T})$ is the set of clopens in $(X, \mathcal{T})$, then closure of A under R-preimage guarantees that (X, R, A) is a general frame. Clearly, it is compact and differentiated, and if $(w, v) \notin R$, then $v \notin R[w]$, which is compact by hypothesis. Moreover, v can be separated from any point $t \in R[w]$ by a clopen set $C_t \ni t$ as $(X, \mathcal{T})$ is totally separated. By compactness, this has a finite subcover whose union then produces a clopen set C containing all of $R[w]$ but not v. As $v \in X \setminus C$ but $R[w] \subseteq C$, tightness follows.

Finally, note that by definition Stone spaces have a basis of clopens, so $\mathcal{T}_{\mathbf{Clop}(\mathcal{T})} = \mathcal{T}$ and note that any clopen set in $\mathcal{T}_A$ is open, so a union of elements in A, and a closed subset of a compact space, so compact. This means that it must be a finite union of those elements in A. A finite union of admissible sets is admissible, so $\mathbf{Clop}(\mathcal{T}_A) = A$. □

2.4 Algebraic duality

The previous section ended with Proposition 2.37, which showed that there is a meaningful correspondence between descriptive frames and modal spaces. This leads back to an important motivation for considering descriptive frames. Stone duality is a celebrated contravariant category equivalence between Stone spaces and Boolean algebras, which are the algebraic models of classical propositional logic.

2.4.1 Stone duality

An extraordinarily rich theory exists around topological duality of algebraic structures, only a small portion of which will be relevant to this thesis. In the interest of brevity, only the bare necessities of this theory will be mentioned. First, there is the classic Stone duality, establishing a dual equivalence between the category of Boolean algebras and the category of Stone spaces.

Definition 2.38. The category **BA** of Boolean algebras has an object class consisting of all Boolean algebras and as morphisms all maps preserving the meet, join and complementation operations. ◀

The next relevant category is that of the Stone spaces mentioned above.

Definition 2.39. The category **Stone** of Stone spaces consists of all totally separated, compact topological spaces (see Definition 2.29) with corresponding morphisms the collection of all continuous maps between them. ◀

It turns out that these two categories can then be linked by a contravariant functor.

Definition 2.40. The **Spec**-*functor*, denoted **Spec** or $(-)_*$, goes from **BA** to **Stone** by mapping a Boolean algebra $\mathcal{B}$ to the Stone space $\mathbf{Spec}(\mathcal{B}) = \mathcal{B}_*$ on Uf $\mathcal{B}$, the set of ultrafilters[3] on $\mathcal{B}$, with a topology generated by all sets $\widehat{b} := \{F \in \mathrm{Uf}\ \mathcal{B} \mid b \in F\}$ for b in $\mathcal{B}$. A morphism $f \in \mathrm{Hom}_{\mathbf{BA}}(\mathcal{B}_1, \mathcal{B}_2)$ is mapped to $\mathbf{Spec}(f) = f_*$ defined by $f_* := \mathscr{P}^{\mathrm{op}} f : \mathbf{Spec}(\mathcal{B}_2) \to \mathbf{Spec}(\mathcal{B}_1)$ with $f_*(F) = f^{-1}[F]$. ◀

Definition 2.41. The **Clop**-functor, written as $(-)^*$ maps objects $\mathcal{X}$ in **Stone** to the field of sets of clopens, $\mathcal{X}_* := (\mathbf{Clop}(\mathcal{X}), \cap, \cup, \cdot^c)$ in $\mathsf{ob}\,\mathbf{BA}$. Morphisms $f : \mathcal{X} \to \mathcal{Y}$ in **Stone**, being continuous functions between the topological spaces, are mapped to $f^* := \mathscr{P}^{\mathrm{op}} f$, so that again any clopen A in $\mathcal{Y}$ is mapped to $f^*(C) = f^{-1}[C]$. ◀

Stone duality is the dual category equivalence established by these two functors.

Theorem 2.42 (Stone duality [8, 24, 35])**.** *The functors* **Spec** *and* **Clop** *establish a dual equivalence between* **BA** *and* **Stone**.

[3] An ultrafilter of a Boolean algebra is an upwards subset F of the universe that is closed under meet and contains b or its complement for each b.

2.4.2 Jónsson-Tarski duality

Stone duality is important for this thesis, because a similar duality holds for the category of modal algebras and the category of descriptive frames. As Proposition 2.37 shows, descriptive frames are simply Stone spaces with extra structure. Similarly, modal algebras are Boolean algebras with extra structure, as seen in the next definition.

Definition 2.43. A *modal algebra* is a pair $\mathcal{M} = (\mathcal{B}, \Diamond)$ where $\mathcal{B} = (B, \vee, \wedge, \neg, 0, 1)$ is a Boolean algebra, and $\Diamond : B \to B$ is an operator satisfying $\Diamond 0 = 0$ and $\Diamond(a \vee b) = \Diamond a \vee \Diamond b$.

The category **MA** has as objects all modal algebras and as morphisms the Boolean morphisms f such that $f(\Diamond a) = \Diamond f(a)$ for all a. ◀

Definition 2.44. The category **DF** is the category with all descriptive frames as objects, and as morphisms the continuous[4] bounded morphisms.[5] ◀

The functors **Spec** and **Clop** can be extended to incorporate these expansions of the structures.

Definition 2.45. Let $\mathcal{M} = (\mathcal{B}, \Diamond)$ be a modal algebra. Then let $\mathcal{M}_* = (W, R, A)$ be the descriptive frame with $(W, A) = \mathbf{Spec}\mathcal{B}$, and R a relation on the ultrafilters given by FRF' if and only if for every $b \in F'$ also $\Diamond b \in F$. The morphisms are mapped in the same way as under Stone duality. ◀

Remark 2.46. Note that the definition of the relation, given by $(F, F') \in R \iff \forall b \in B[b \in F' \implies \Diamond b \in F]$ is equivalent to the definition $(F, F') \in R \iff \forall b \in B[\Box b \in F \implies b \in F']$. After all,

$$\begin{aligned}\big(b \in F' \implies \Diamond b \in F\big) &\iff \big(\Diamond b \notin F \implies b \notin F'\big) \iff \big(\neg\Diamond b \in F \implies \neg b \in F'\big)\\ &\iff \big(\Box\neg b \in F \implies \neg b \in F'\big).\end{aligned}$$

Quantifying over all b and using that $\neg\neg b = b$ in modal algebras then implies that these two universal statements are equivalent. ◀

Definition 2.47. Let $\mathfrak{g} = (W, R, A)$ be a descriptive frame. Then the associated modal algebra is $\mathfrak{g}^* = (A, \langle R \rangle)$. The morphisms are again mapped as under Stone duality. ◀

Remark 2.48. For both of these functors, note that when converting the descriptive frames back to Stone spaces as per Proposition 2.37 produces again the **Spec**- and **Clop**-functor from the Stone duality. ◀

Theorem 2.49 (Jónsson-Tarski duality [8, 25, 26])**.** *These two functors present a dual category equivalence between* **MA** *and* **DF**.

[4] Recall Definition 2.20

[5] A bounded morphism from between Kripke frames $\mathfrak{F} = (W, R)$ and $\mathfrak{F}' = (W', R')$ is a map $f : W \to W'$ with the property that if wRv then $f(w)Rf(v)$ and if $f(w)R'v'$ then there is a $v \in W$ such that wRv and $f(v) = v'$.

An interesting and vitally important note is that nothing prevents the use of the $\cdot^*$-functor to be restricted to descriptive frames. For the entire class of general frames, one may define the exact same operation. Combining it with the other functor results in an operation that turns every general frame into a descriptive frame. This may be called the *double dual*, the *dual completion*, or the *descriptive completion* of the general frame.

Definition 2.50. Let $\mathfrak{g} = (W, R, A)$ be a general frame. Then its *descriptive completion* is $(\mathfrak{g}^*)_* =: (W_*, R_*, \widehat{A})$, where $\widehat{A} = \{\widehat{a} \mid a \in A\}$. This is always a descriptive frame. If $\mathfrak{m} = (\mathfrak{g}, V)$ is a general model, there is an induced descriptive model $\mathfrak{m}_* := ((\mathfrak{g}^*)_*, V_*)$ where V_* is a valuation such that if originally $V(p) = a$, then $V_*(p) = \widehat{a}$. This will be called its *completed model.* ◀

In the special case that $\mathfrak{g}$ is a Kripke frame, meaning that $A = \mathscr{P}(W)$ is the powerset, then $(\mathfrak{g}^*)_*$ is known as the *ultrafilter extension* of $\mathfrak{g}$, see also [8, Section 2.5].

Proposition 2.51. Let $\mathfrak{g} = (W, R, A)$ be a general frame. There is a canonical relation-preserving map (crucially not necessarily a bounded morphism) from W to W_* given by sending a world w to the ultrafilter generated by w, given by $F_w := \{a \in A \mid w \in a\}$. If $\mathfrak{g}$ is tight, then two related points in the image were also related in $\mathfrak{g}$. If $\mathfrak{g}$ is differentiated, this morphism is an embedding, denoted by $\iota_{\mathfrak{g}}$ or ι if the relevant general frame is clear from context.

Proof. Suppose that v is a successor of w. Then in particular, if $v \in a$, it follows that $w \in \langle R \rangle a$. From this, we infer that by definition F_v is a successor of F_w. If in addition $\mathfrak{g}$ is tight and v is not a successor, then there exists an a such that $v \in a$ but $w \notin \langle R \rangle a$ from which it follows that F_v is also not a successor of F_w. Finally, if $\mathfrak{g}$ is differentiated, then for any two distinct points w and v there exists an admissible set $a \in A$ such that $w \in a$ but $v \notin a$, so that $a \in F_w$ and $a \notin F_v$, implying $F_w \neq F_v$. □

2.5 Bisimulations

An important notion in modal logic that is central to main theorem of this thesis is the concept of bisimulations. Intuitively, if two pointed models $\mathfrak{M}, w$ and $\mathfrak{N}, v$ are bisimilar, they are impossible to distinguish by walking from w and v through the models and looking only at the propositional variables that are true at the points reached.

2.5.1 Kripke and Vietoris bisimulations

Definition 2.52. Let $\mathfrak{M}$ and $\mathfrak{M}'$ be (general) models on frames (W, R) and (W', R') with valuations V and V'. A relation $Z \subseteq W \times W'$ is a called (*Kripke*) *bisimulation* if for all $(w, w') \in Z$ the following three conditions hold:

(prop) If $p \in \mathsf{P}$ then $w \in V(p) \iff w' \in V'(p)$;

(forth) If $v \in R[w]$ then there exists a $v' \in R'[w']$ such that $(v, v') \in Z$;

(back) If $v' \in R'[w']$ then there exists a $v \in R[w]$ such that $(v, v') \in Z$.

Another way to think of these conditions is in terms of disjointness conditions:

(prop) If p is a propositional variable, then

$$Z \cap (V(p) \times V'(p)^c) = Z \cap (V(p)^c \times V'(p)) = \emptyset;$$

(forth) For all sets $X \subseteq W$ and $Y \subseteq W'$

$$X \times Y \cap Z = \emptyset \implies \Big(\langle R \rangle X \times [R'] Y \Big) \cap Z = \emptyset;$$

(back) For all sets $X \subseteq W$ and $Y \subseteq W'$

$$X \times Y \cap Z = \emptyset \implies \Big([R] X \times \langle R' \rangle Y \Big) \cap Z = \emptyset.$$

These conditions can easily be seen to be equivalent and will turn out to be a useful way to look at these relations. ◀

As will be shown in Corollary 2.60, two points linked by a bisimulation satisfy the same modal formulae. This is key to their importance in modal logic.

Definition 2.53 ([5, 17])**.** Let $\mathfrak{m}$ and $\mathfrak{m}'$ be general models and Z a Kripke bisimulation between them. If Z is closed in the product topology of the two associated topological spaces, then it is called a *Vietoris bisimulation.* ◀

If two points are linked by a Kripke bisimulation this is denoted by $\mathfrak{m}, w \underline{\leftrightarrow} \mathfrak{m}', w'$ and if they are linked by a Vietoris bisimulation this will be written as $\mathfrak{m}, w \leftrightharpoons \mathfrak{m}', w'$.

Discussion 2.54. This definition is motivated by a coalgebraic perspective on descriptive frames, in which it is the coalgebraic definition of a bisimulation on descriptive frames as coalgebras for the Vietoris functor. However, on descriptive frames this is equivalent to the definition provided above. A detailed exploration of the coalgebraic notion of descriptive frames can be found in [5, 27, 37] and an extensive treatment of Vietoris bisimulations, including this equivalence, can be found in [5, 17].

Proposition 2.55. [5, theorem 5.2] If $(\mathfrak{F}, A, V)$ and $(\mathfrak{F}', A', V')$ are image-compact general models and Z is a Kripke bisimulation between $(\mathfrak{F}, V)$ and $(\mathfrak{F}', V')$, then its closure in the product topology, $\overline{Z}$, is a Vietoris bisimulation.

Proof. To see that $\overline{Z}$ is a Kripke bisimulation, consider the disjointness conditions.

Suppose w and w' disagree on a propositional variable p. Then without loss of generality $(w, w') \in V(p) \times V'(p)^c$ is an open neighbourhood disjoint from Z, from which it follows that $(w, w') \notin \overline{Z}$.

For the forth condition, the approach will be different from [5] and more similar to the argument for the propositional condition. Suppose that Rwv and $(v, v') \notin \overline{Z}$ for all $v' \in R'[w']$. Then for each $v' \in R'[w']$ there are $a_{v'} \in A$ and $b_{v'} \in A'$ such that $(v, v') \in a_{v'} \times b_{v'}$ and $a_{v'} \times b_{v'} \cap Z = \emptyset$, because the pair is not in the closure.

Then $\mathcal{U} = \{b_{v'} \mid v' \in R'[w']\}$ is an open cover of $R'[w']$, which was compact per assumption. So there is a finite subcover. Define b to be the finite union of that subcover and a the intersection of the corresponding $a_{v'}$. Then $a \times b \supseteq \{v\} \times R'[w']$ is a product of clopens that is disjoint from Z and thus $(\langle R\rangle a) \times ([R']b) \ni (w, w')$ is an open neighbourhood that is disjoint from Z. So $(w, w') \notin \overline{Z}$. The back condition is proven identically. $\square$

Corollary 2.56. If $\mathfrak{m}, w$ and $\mathfrak{n}, v$ are two pointed image-compact general models, then the following are equivalent:

1. $\mathfrak{m}, w \,\underline{\leftrightarrow}\, \mathfrak{n}, v$;

2. $\mathfrak{m}, w \leftrightharpoons \mathfrak{n}, v$.

Proof. The implication (2) $\implies$ (1) is immediate from the fact that all Vietoris bisimulations are a forteriori Kripke bisimulations. The implication (1) $\implies$ (2) follows from proposition 2.55 and the fact that $Z \subseteq \overline{Z}$. $\square$

2.5.2 Finite bisimulations

Crucial for this thesis will be the notion of finite (approximations to) bisimulations. Like the finitary modal language, the finite number of steps involved makes them much easier to deal with.

Definition 2.57. Let $k \in \mathbb{N}$ be a natural number, P a set of propositional variables and $\mathfrak{M}$ and $\mathfrak{M}'$ be two (general) models with frames (W, R) and (W', R') and valuations V and V' over P respectively. Then a k-bisimulation over P is a $\subseteq$-decreasing $(k+1)$-sequence $(Z_\ell)_{0 \le \ell \le k}$ of relations $Z_\ell \subseteq W \times W'$ such that for all natural numbers $\ell \le k$ and $(w, w') \in Z_\ell$:

(prop) If $p \in \mathsf{P}$ then $w \in V(p)$ if and only if $w' \in V'(p)$;

(forth) If $v \in R[w]$ then there exists a $v' \in R'[w']$ such that for all non-negative $m < \ell$ the pair (v, v') is an element of Z_m;

(back) If $v' \in R'[w']$ then there exists a $v \in R[w]$ such that for all non-negative $m < \ell$ the pair (v, v') is an element of Z_m.

Once more, these could all be rephrased in terms of disjointness conditions:

(prop) If p is a propositional variable, then

$$(V(p) \times V'(p)^c) \cap \bigcup_{0 \le \ell \le k} Z_\ell = (V(p)^c \times V'(p)) \cap \bigcup_{0 \le \ell \le k} Z_\ell = \emptyset;$$

(forth) For all sets $X \subseteq W$ and $Y \subseteq W'$ and $\ell \le k$

$$X \times Y \cap \bigcap_{0 \le m < \ell} Z_\ell = \emptyset \implies \Big(\langle R\rangle X \times [R']Y\Big) \cap Z_\ell = \emptyset;$$

(back) For all sets $X \subseteq W$ and $Y \subseteq W'$ and $\ell \leq k$

$$X \times Y \cap \bigcap_{0 \leq m < \ell} Z_m = \emptyset \implies \Big([R]X \times \langle R' \rangle Y\Big) \cap Z_\ell = \emptyset.$$

One could also define finite Vietoris bisimulations when all Z_ℓ are closed, but this notion will not be useful for this thesis.

If there exists an k-bisimulation $(Z_\ell)_{0 \leq \ell \leq k}$ with $(w, v) \in Z_k$, this is denoted by $\mathfrak{M}, w \underline{\leftrightarrow}_k \mathfrak{N}, v$. ◀

Remark 2.58. If Z is a bisimulation, then for any $k \in \mathbb{N}$ the sequence $(Z)_{0 \leq \ell \leq k}$ is a finitary bisimulation. As such, $\mathfrak{M}, w \underline{\leftrightarrow} \mathfrak{N}, v$ implies $\mathfrak{M}, w \underline{\leftrightarrow}_k \mathfrak{N}, v$ for any k. ◀

Lemma 2.59. Let $k \in \mathbb{N}$ and $\mathfrak{M}, w$ and $\mathfrak{N}, v$ be two pointed (general) models and P a set of propositional variables. Then

$$\text{If } \mathfrak{M}, w \underline{\leftrightarrow}_k \mathfrak{N}, v \text{ over } \mathsf{P} \text{ then } \mathfrak{M}, w \leftrightsquigarrow_k \mathfrak{N}, v \text{ over } \mathsf{P}$$

and if P is finite, then the converse implication holds, too.

Proof. Let $\mathfrak{M} = (W, R, V)$ and $\mathfrak{N} = (W', R', V')$ be two Kripke models. The statement follows by induction on k. For the induction hypothesis, suppose that it is true for all $\ell < k$. Assume that $\mathfrak{M}, w \underline{\leftrightarrow}_k \mathfrak{N}, v$ and let $p \in \mathsf{P}$. Then there is a $(k+1)$-sequence $(Z_\ell)_{0 \leq \ell \leq k}$ with $(w, v) \in Z_k$ satisfying the relevant conditions, implying in particular

$$\mathfrak{M}, w \Vdash p \iff w \in V(p) \iff v \in V'(p) \iff \mathfrak{N}, v \Vdash p.$$

Now suppose $\Diamond\psi, \Box\psi \in \mathtt{ML}_k(\mathsf{P})$, for $\psi \in \mathtt{ML}_\ell(\mathsf{P})$ with $\ell < k$. Remark that this assumption is vacuous if $k = 0$. Then $\mathfrak{M}, w \Vdash \Diamond\psi$ if and only if there is a $\widetilde{w} \in R[w]$ such that $\mathfrak{M}, \widetilde{w} \Vdash \psi$. Then there must exist a $\widetilde{v} \in R'[v]$ such that $(\widetilde{w}, \widetilde{v}) \in Z_\ell$ so that $\mathfrak{M}, \widetilde{w} \underline{\leftrightarrow}_\ell \mathfrak{N}, \widetilde{v}$ via $(Z_m)_{0 \leq m \leq \ell}$, from which the induction hypothesis gives $\mathfrak{N}, \widetilde{v} \Vdash \psi$ and thus $\mathfrak{M}, v \Vdash \Diamond\psi$. The other direction is completely symmetric and the case for $\Box\psi$ can be seen immediately from the duality $\Box = \neg\Diamond\neg$.

So $\mathfrak{M}, w$ and $\mathfrak{N}, v$ agree on all propositional variables and on formulae with a box or diamond, implying that they agree on all Boolean combinations, which comprise $\mathtt{ML}_k(\mathsf{P})$ in its entirety, so that $\mathfrak{M}, w \leftrightsquigarrow_k \mathfrak{N}, v$.

For the final addition, suppose that P is finite. Once again we perform induction on k. Let the induction hypothesis state that the right to left implication holds for all $\ell < k$. From proposition 2.11 it follows then that all $\mathtt{ML}_\ell(\mathsf{P})$ are finite up to equivalence. Choose a set of non-equivalent representations $\{\varphi_0, \dots, \varphi_n\}$ such that each of the formulae in $\bigcup_{\ell < k} \mathtt{ML}_\ell(\mathsf{P})$ is equivalent to one of them and consider $(\leftrightsquigarrow_\ell)_{0 \leq \ell \leq k}$ as finitary bisimulation. Suppose that $\mathfrak{M}, w \leftrightsquigarrow_k \mathfrak{N}, v$. Then certainly

$$w \in V(p) \iff \mathfrak{M}, w \Vdash p \iff \mathfrak{N}, v \Vdash p \iff v \in V'(p).$$

For the forth condition, suppose that $\widetilde{w}$ is a successor of w. Consider then $I \subseteq \{0, \dots, n\}$ such that $\mathfrak{M}, \widetilde{w} \Vdash \varphi_i \iff i \in I$. Then in particular, since $\widetilde{w}$ is a successor of w, it follows that $\mathfrak{M}, w \Vdash \Diamond\left(\bigwedge_{i \in I} \varphi_i \wedge \bigwedge_{i \notin I} \neg\varphi_i\right)$, which is a formula in $\mathtt{ML}_k(\mathsf{P})$ by construction, so that $\mathfrak{N}, v \Vdash \Diamond\left(\bigwedge_{i \in I} \varphi_i \wedge \bigwedge_{i \notin I} \neg\varphi_i\right)$. As such, v has a successor $\widetilde{v}$ with $\mathfrak{N}, \widetilde{v} \Vdash \varphi_i \iff i \in I$. These were representatives of all formulae in $\bigcup_{\ell<k} \mathtt{ML}_\ell(\mathsf{P})$, so that $\widetilde{w} \leftrightsquigarrow_\ell \widetilde{v}$ for all $\ell < k$. The back condition is checked identically. This implies that $(\leftrightsquigarrow_\ell)_{0 \leq \ell \leq k}$ is indeed a finitary bisimulation. So $\mathfrak{M}, w \underline{\leftrightarrow}_k \mathfrak{N}, v$. □

Corollary 2.60. For any two pointed (general) models $\mathfrak{M}, w$ and $\mathfrak{N}, v$ and set P of propositional variables

$$\mathfrak{M}, w \underline{\leftrightarrow} \mathfrak{N}, v \text{ over } \mathsf{P} \implies \mathfrak{M}, w \leftrightsquigarrow \mathfrak{N}, v \text{ over } \mathsf{P}.$$

Proof. This follows immediately from Lemma 2.59 and Remark 2.58. □

An important note to make is that these are equivalence relations, in particular transitive.

Lemma 2.61. Assume that $\mathfrak{M}_0, w_0 \underline{\leftrightarrow}_\ell \mathfrak{M}_1, w_1$ and $\mathfrak{M}_1, w_1 \underline{\leftrightarrow}_k \mathfrak{M}_2, w_2$ for $\ell \leq k$. Then $\mathfrak{M}_0, w_0 \underline{\leftrightarrow}_\ell \mathfrak{M}_2, w_2$. In particular, if $\mathfrak{M}_0, w_0 \underline{\leftrightarrow} \mathfrak{M}_1, w_1$ and $\mathfrak{M}_0, w_0 \underline{\leftrightarrow}_\ell \mathfrak{N}, v$, then $\mathfrak{M}_1, w_1 \underline{\leftrightarrow}_\ell \mathfrak{N}, v$.

Proof. Let $(Z_m)_{0 \leq m \leq \ell}$ be an ℓ-bisimulation between $\mathfrak{M}_0, w_0$ and $\mathfrak{M}_1, w_1$ and let $(\widetilde{Z}_n)_{0 \leq n \leq k}$ be a k-bisimulation between $\mathfrak{M}_1, w_1$ and $\mathfrak{M}_2, w_2$. Then consider $(Z_m; \widetilde{Z}_m)_{0 \leq m \leq \ell}$ as an ℓ-bisimulation between $\mathfrak{M}_0, w_0$ and $\mathfrak{M}_2, w_2$, where ; denotes composition of relations. Suppose that $(v_0, v_2) \in Z_m; \widetilde{Z}_m$ for some $m \leq \ell$. Then there exists a v_1 such that $v_0 Z_m v_1$ and $v_1 \widetilde{Z}_m v_2$. Consequently, v_0 satisfies the same propositional variables as v_1, which in turn satisfies the same propositional variables as v_2, verifying that v_0 and v_2 satisfy the same propositional variables.

For the forth condition, if v_0 has a successor x_0, then there is a successor x_1 of v_1 such that $x_0 Z_n x_1$ for all $n < m$. Then from $v_1 \widetilde{Z}_m v_2$ it follows that there is a successor x_2 of v_2 such that $x_1 \widetilde{Z}_n x_2$ for all $n < m$. Therefore $(x_0, x_2) \in (Z_n; \widetilde{Z}_n)$ for all $n < m$. The back condition is identical. The final observation then follows from Remark 2.58. □

2.6 Nets as topological toolkit

This section will develop nets as a tool to analyse general frames more carefully from a topological perspective. Nets are a fairly advanced tool in general topology and are not extremely widely used because they do not always behave as one might expect. However, they will prove a strong and intuitive tool for the purposes of this thesis, remaining very well-behaved for the purposes of this thesis. Mostly, they will be used to give the reader a feeling for the behaviour of general frames.

This means that they are not crucial for an understanding of the main results of this thesis, and a reader not interested in this intricate topological development can skip this section without concern for their understanding of the main theorem of the thesis. Care was taken to avoid including them in the most important lines of reasoning. However, readers with a background in topology could consider the section a formal framework for intuitions they may already have or develop along the section and as an potential tool for further research.

In the Chapter 3 they are used to demonstrate failures of model-theoretic theorems on the class of descriptive models, necessitating different techniques in Chapter 4. However, in an effort not to alienate readers with a weaker background in topology, alternative reasoning is presented, but these proofs will be much less intuitive.

They will also be a vital tool in Chapter 5, where the contributions of this thesis will be developed further in preparation for future research. These proofs will be supplied with alternative reasoning.

2.6.1 Nets: basic definitions

This section will explore topological notions of convergence and continuity and tie this back to general frames. The concepts introduced here will also come back in Section 2.6.2 as an alternative and hopefully enlightening point of view. The definitions and basic results will mostly be following [10].

Definition 2.62. Let $\mathcal{X} = (X, \mathcal{T})$ be a topological space, α an ordinal, $x \in X$ and $(x_\beta)_{\beta\in\alpha}$ be an α-sequence in X. Then the sequence *converges to* x if for every open neighbourhood $U \ni x$ there is a $\gamma \in \alpha$ such that $\beta \geq \gamma$ implies $x_\beta \in U$. This is often abbreviated to saying that x_β is *eventually in* U. The sequence is then called *convergent*. For the set of points to which a sequence $(x_\beta)_{\beta\in\alpha}$ converges, one writes $\lim_{\beta\in\alpha} x_\beta$. If this is a singleton $\{x\}$, one may also write

$$\lim_{\beta\in\alpha} x_\beta = x$$

and call x the *limit* of the sequence. ◀

Sequences are extremely useful for well-behaved topological spaces like metrisable spaces, but almost any interesting result for sequences requires the space to be at least first-countable. However, there exists a more general notion: nets. Instead of a well-order, they have a directed order. This makes them slightly less intuitive, but greatly increases their expressive power.

Definition 2.63. Let $\mathbf{D} = (D, \prec)$ be a directed order[6]. A $\mathbf{D}$-*net* is a D-labeled set $(x_d)_{d\in D}$. A net $(x_d)_{d\in D}$ is said to be *eventually in* a set Y if there is a $d_0 \in D$ such that $d \succeq d_0$ implies $x_d \in Y$. Moreover, if for every $d_0 \in D$ there is a $d \succeq d_0$ such that $x_d \in Y$, then $(x_d)_{d\in D}$ is said to be *frequently in* Y.

If $\mathbf{\Delta} = (\Delta, \lhd)$ and $\mathbf{D} = (D, \prec)$ are directed orders and $(x_d)_{d\in D}$ is a net, then a *subnet* of $(x_d)_{d\in D}$ is characterised by a cofinal and monotone[7] map $h : \Delta \to D$ and given by $(x_{h(\delta)})_{\delta\in\Delta}$.

Letters used for nets are ν or n. ◀

[6] A directed order is a preorder in which every two elements have a common upper bound.

[7] Recall that cofinal means that for every $d \in D$ there is a $\delta \in \Delta$ such that $h(\delta) \geq d$ and monotone means that for every two δ, δ' if $\delta \lhd \delta'$ then $h(\delta) \preceq h(\delta')$.

Subnets should be treated with some care, as they may occasionally prove counter-intuitive. An important detail to keep in mind is that a subnet of a sequence is not necessarily a sequence.

Remark 2.64. A net $(x_d)_{d\in D}$ on X is frequently in a set $Y \subseteq$ if and only if it is not eventually in Y^c. ◀

Definition 2.65. A net $(x_d)_{d\in D}$ *converges* to a point x if for every open neighbourhood U of x, the net $(x_d)_{d\in D}$ is eventually in U. If the limit point is unique, one writes

$$\lim_{d\in D} x_d = x$$

and x is called the *limit point* of $(x_d)_{d\in D}$. ◀

Remark 2.66. Every net in a Hausdorff space (recall Definition 2.24) converges to at most one point. After all, any two distinct points have disjoint neighbourhoods U and V, and a net cannot eventually be in both. After all, if x_d is in U for $d \succeq d_0$ and $x_d \in V$ for $d \succeq d_1$, then for d_2 a common upper bound for d_0 and d_1, the element $x_{d_2} \in U \cap V = \emptyset$, which is a contradiction.

In particular, by Corollary 2.25, every net on the $\mathfrak{g}$-space of a differentiated general frame $\mathfrak{g}$ must converge to at most one point. ◀

It turns out that nets capture many important topological concepts. In fact, nets can be used to identify open and closed sets.

Proposition 2.67. [10, Proposition 6.7] Let $\mathcal{X}$ be a topological space. Then the following hold for all subsets Y of the carrier:

1. Y is closed if and only if every convergent net that is eventually in Y has its limit points in Y;

2. Y is open if and only if every convergent net with limit points contained in Y is eventually in Y.

It should be obvious from this result that nets have great expressive power. For example, one may express topological compactness in the terms of nets.

Proposition 2.68. [10, Theorem 7.14] A topological space $\mathcal{X} = (X, \mathcal{T})$ is compact if and only if all nets on $\mathcal{X}$ have convergent subnets.

Another important notion expressible in terms of nets is the continuity of functions.

Proposition 2.69. [10, Proposition 6.6] Let $f : \mathcal{X} \to \mathcal{Y}$ be a function. Then the following are equivalent:

1. f is continuous (recall Definition 2.20);

2. For every net $(x_d)_{d\in D}$ converging to x, the net $(f(x_d))_{d\in D}$ converges to $f(x)$.

Inspired by this, one may define continuity for relations, too. However, there are two plausible definitions, both important corresponding notions in point-set topology and general frames. These will be *full* and *functional continuity*. In the most general context, neither one is stronger than the other.

Definition 2.70. Two nets $(x_d)_{d\in D}$ and $(y_\delta)_{\delta\in\Delta}$ are said to be *frequently (R-)related* if for all $d_0 \in D$ and $\delta_0 \in \Delta$ there are d and δ greater than d_0 and δ_0 such that Rx_dy_δ. ◀

An equivalent, but less symmetric way of saying this is that for any $d_0 \in D$ the net $(y_\delta)_{\delta\in\Delta}$ is frequently in $R[\{x_d\}_{d\succeq d_0}]$.

Definition 2.71. Let $\mathcal{X}$ and $\mathcal{Y}$ be topological spaces on carriers X and Y, and let $R \subseteq X \times Y$ be a relation. Then R is called *fully continuous* if for any two frequently R-related nets $(x_d)_{d\in D}$ and $(y_\delta)_{\delta\in\Delta}$ converging on $\mathcal{X}$ and $\mathcal{Y}$, any two respective limit points x and y are R-related.

We say moreover that R is *functionally continuous* if for any two any two frequently related nets $(x_d)_{d\in D}$ and $(y_\delta)_{\delta\in\Delta}$ on $\mathcal{X}$ and $\mathcal{Y}$ respectively, if $(x_d)_{d\in D}$ converges to x, then there must be a limit point y of $(y_\delta)_{\delta\in\Delta}$ such that Rxy. That is, instead of demanding that all limit points are related, each limit point of $(x_d)_{d\in D}$ has to be related to at least one limit point of $(y_\delta)_{\delta\in\Delta}$. ◀

These distinct notions of continuity are comparable for sufficiently separated spaces: in a Hausdorff space, there is at most one limit point and functional continuity becomes strictly stronger than full continuity. In the context of Hausdorff spaces, then, it is enough to speak only of weak and strong continuity.

The next proposition justifies the name functionally continuous.

Proposition 2.72. A function is continuous if and only if the associated relation is functionally continuous.

Proof. Immediate from Proposition 2.69. □

The next obvious step is to relate this topological material to general frames. It turns out that continuity of relations fits corresponds to a previously discussed concept.

Proposition 2.73. Let $\mathfrak{g} = (W, R, A)$ be a general frame. Then R is fully continuous on the $\mathfrak{g}$-space if and only if $\mathfrak{g}$ is tight.

Proof. For the implication from right to left, assume the frame is tight, and let $(x_d)_{d\in D}$ and $(y_\delta)_{\delta\in\Delta}$ converge to x and y respectively. Suppose for contraposition that $(x, y) \notin R$. From tightness, there must exist an $a \in A$ such that $y \in a$ but $x \notin \langle R\rangle a$.

Since admissible sets are clopen, $(y_\delta)_{\delta\in\Delta}$ is eventually in a, but $(x_d)_{d\in D}$ is eventually in $W \setminus \langle R\rangle a = [R]a^c$. So there exist d_0 and δ_0 such that $d \succeq d_0$ implies $x_d \in [R]a^c$ and $\delta \succeq \delta_0$ implies $y_\delta \in a$. As such, for sufficiently large d and δ, x_d and y_δ are not related, so that the nets are certainly not frequently related. Continuity of R then follows from contraposition.

For the other direction, suppose that the frame is not tight. Then there is a pair $(w, v) \notin R$ such that for each $a \ni v$ also $w \in \langle R \rangle a$. Now consider the directed orders D_w and D_v defined by

$$D_w := \{(x, b) \mid x, w \in b \in A\}$$

ordered by $(x, a) \succeq (y, b)$ if and only if $a \subseteq b$ and similar for D_v. They naturally define nets by projection on the first coordinate. Note that these nets converge by construction to w and v. Moreover, let $(x, b) \in D_w$ and let $(y, a) \in D_v$. Then $w \in \langle R \rangle a$, so $(w, \langle R \rangle a) \in D_w$. Thus there must be a $(z, b \cap \langle R \rangle a) \in D_w$. This means in particular that there is a $z' \in a$ such that zRz'. As $(z', a) \succeq (y, a)$, this demonstrates that the two nets are frequently related.

Therefore, since the assumption was that $(w, v) \notin R$, the relation cannot be fully continuous. $\square$

This provides a new characterisation of modal spaces from Definition 2.36.

Proposition 2.74. A triple $(X, R, \mathcal{T})$ is a modal space if and only if $(X, \mathcal{T})$ is a Stone space, R is fully continuous on it, and the inverse image under R preserves clopen sets.

Proof. This is immediate from Propositions 2.73 and 2.37. The preservation of clopen sets under inverse R-images is needed to ensure that the basis of clopens can be a collection of admissible sets that is closed under $\langle R \rangle$. $\square$

As an illustration of the usefulness of nets in this context, the following will be a visual rewording of Example 2.18.

Example 2.75. Recall from Example 2.18 that the frame $(\omega, \geq)$ (see Figure 2.3a) with collection of admissible sets consisting of finite subsets of $\mathbb{N}$ and their complements in $\mathbb{N} \cup \{\omega\}$ was a tight (in fact descriptive) frame, while $(\omega, >)$ (see Figure 2.3b) with the same admissible sets was not. With Proposition 2.73 this can be seen visually in two ways.

Note that in the induced topology, the sequence $(n)_{n\in\mathbb{N}}$ converges to ω. After all, any open neighbourhood of ω must contain an admissible set with ω as element. Any such admissible set must be cofinite in $\mathbb{N}$ so that $(n)_{n\in\mathbb{N}}$ is eventually in this admissible set and hence in the open neighbourhood.

Noting that also the constant net at ω (over any directed set) must converge to ω, it follows that for any tight relation R with $R\omega n$ for any n the sequence $(n)_{n\in\mathbb{N}}$ is eventually (and thus certainly frequently) related to the constant sequence at ω. By Proposition 2.73, tightness of R then implies that their limits, both ω, must be related as well, so that ω is reflexive. This argument is visually illustrated in Figure 2.3c.

Alternatively, let R be a tight relation such that $(n + 1, n) \in R$ for all n. Then the same trick can be applied with both sequences equal to $(n)_{n\in\mathbb{N}}$. For any numbers N and N', the components $n = \max\{N, N'\}$ and $n + 1$ are both larger than N and N', and $(n + 1)Rn$. Thus the sequences are frequently related, so that, again from tightness, the limit, ω, must be related to itself. This argument is visually illustrated in Figure 2.3d.

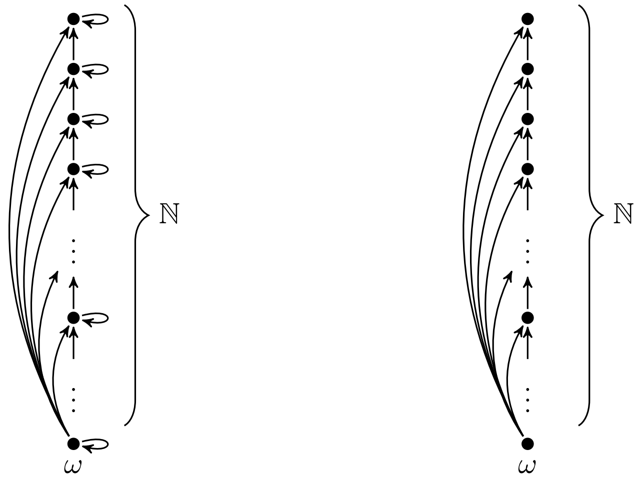

(a) Schematic illustration of $(\omega + 1, \geq)$.

(b) Schematic illustration of $(\omega + 1, >)$.

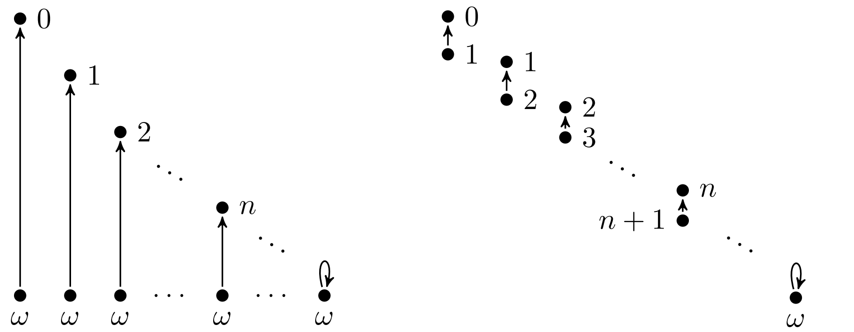

(c) Illustration of how the convergence of the sequence $(0, 1, 2, \dots)$ to ω implies any tight frame on this topology in which ω sees all natural numbers must have ω reflexive.

(d) Illustration of a similar convergence conclusion, this time showing that any relation including the successor relation must have ω reflexive.

Figure 2.3: Visual illustrations of convergences on the structure $\omega+1$ of natural numbers to ω showing that any tight relation on this structure containing either (ω, n) for all n or $(n+1, n)$ for all n must have (ω, ω) as well.

Visual reasoning of this type is helpful in developing intuitions of the behaviour of descriptive frames. It is useful both in understanding specific structures and classes of structures. In the latter case care must be taken that nets do not always behave exactly like sequences but thinking in terms of sequences is still in all cases a helpful guide.

2.6.2 Topological construction of descriptive completion

The suggestive use of the name descriptive completion from Definition 2.50 implies that in some sense it undergoes the minimal number of adjustments to become a descriptive frame. There is a topological sense in which this is exactly right. Each "gap" that undermines compactness is filled and tightness and differentiatedness are enforced with minimal adjustments.

To see this first requires some definitions.

Definition 2.76. Let $(x_d)_{d\in D}$ and $(y_\delta)_{\delta\in\Delta}$ be two nets on a set X. Then they are *equivalent nets* (denoted by $(x_d)_{d\in D} \sim (y_\delta)_{\delta\in\Delta}$) if for every $Y \subseteq X$ the net $(x_d)_{d\in D}$ is eventually in Y if and only if $(y_\delta)_{\delta\in\Delta}$ is eventually in Y. ◀

Definition 2.77. A net $(x_d)_{d\in D}$ on a set X is called *universal* or *an ultranet* if for every $Y \subseteq X$ either $(x_d)_{d\in D}$ is eventually in Y or it is eventually in $X \setminus Y$. ◀

The name ultranet is of course suggestive of ultrafilters. It turns out that there is a direct correspondence to be found.

Proposition 2.78. [16, Theorem 1.6.13] Let $\mathcal{X}$ be a topological space on a carrier X.

1. Each net on X canonically generates a filter on $\mathscr{P}(X)$ such that two nets are assigned the same filter if and only if they are equivalent as in Definition 2.76.
2. Each ultranet in this way canonically generates an ultrafilter on $\mathscr{P}(X)$.
3. Each filter F has a natural associated net that generates it.
4. There is a bijection between the set of equivalence classes of nets of size at most $|X \times \mathscr{P}(X)|$ and the set of filters of $\mathscr{P}(X)$.
5. There is a bijection between the set of equivalence classes of universal nets of size at most $|X \times \mathscr{P}(X)|$ and the set of ultrafilters of $\mathscr{P}(X)$.

Proof. These items can be proven one by one.

1. Let $\nu = (x_d)_{d\in D}$ be a net. Then define its generated filter as
$$F_n := \{Y \in \mathscr{P}(X) \mid \nu \text{ eventually in } Y\}.$$
This is a filter. After all, if ν is eventually in $Y \subseteq Y'$, then obviously ν is also eventually in Y'. Moreover, if ν is eventually in Y_0 and eventually in Y_1, then there exist d_0 and d_1 such that $d \succeq d_0 \implies x_d \in Y_0$ and $d \succeq d_1 \implies x_d \in Y_1$. As D is directed, there is a d_2 larger than both d_0 and d_1. Now if $d \succeq d_2$, then $d \succeq d_0$ and $d \succeq d_1$, so $x_d \in Y_0 \cap Y_1$. So ν is eventually in $Y_0 \cap Y_1$.
Moreover, if $F_\nu = F_{\nu'}$, then for every set $Y \in \mathscr{P}(X)$, the net ν is eventually in Y if and only if ν' is eventually in Y, meaning exactly that they are equivalent.
2. This follows immediately from the definition of universal nets.

3. Let F be a filter. Construct the net

$$\begin{aligned} D_F &:= \{(x,Y) \mid x \in Y \in F\}; \\ (x,Y) \preceq (x',Y') &:\iff Y \supseteq Y'; \\ x_d &:= \pi_0(d) && \text{where } d \in D_F \text{ and } \pi_0(x,Y) = x; \\ n_F &:= (x_d)_{d \in D_F} && (2.2) \end{aligned}$$

   Let $Y \in F$. Then Y is non-empty, so that there is an $x \in Y$. Let $d = (x',Y') \succeq (x,Y) \in D_F$. Then $x_d = x' \in Y' \subseteq Y$, so n_F is eventually in Y.
   Moreover, suppose that n_F is eventually in Y. Then there is a $d_0 = (x_0, Y_0)$ such that $d \succeq d_0 \implies x_d \in Y$. In particular, for any $y \in Y_0$, it follows that $d := (y, Y_0) \succeq (x_0, Y_0)$, so that $y = x_d \in Y$. Thus $Y \supseteq Y_0 \in F$, meaning $Y \in F$. Thus n_F generates F.

4. Injectivity of the map given by the generation is immediate from item 1. Surjectivity follows from the fact that $D_F \subseteq X \times F \subseteq X \times \mathscr{P}(X)$, so that n_F satisfies the size requirement.

5. This follows immediately from items 2 and 4, combinded with the fact that n_F is by construction a universal net for each ultrafilter F. □

The connection between universal nets and the descriptive completion will emerge from this correspondence. The filters used to construct the descriptive completion should in some way be coupled to the nets. It is important to first investigate the connection to compactness.

Proposition 2.79. [10, Theorem 7.14] Let $\mathcal{X}$ be a topological space. The following are equivalent.

1. $\mathcal{X}$ is compact;

2. Every universal net is convergent.

Proof. Towards a contradiction for the implication (1) $\implies$ (2), assume that $\mathcal{X}$ is compact but $(x_d)_{d \in D}$ is a universal net that is not convergent. Then for every x there exists an open neighbourhood U_x such that $(x_d)_{d \in D}$ is not eventually in U_x, meaning $(x_d)_{d \in D}$ is eventually in U_x^c.

The collection $\{U_x \mid x \in X\}$ is an open cover and hence has a finite subcover, $\{U_{x_1}, \dots, U_{x_n}\}$. Since $(x_d)_{d \in D}$ is eventually in $U_{x_i}^c$ for all i, it follows that

$$(x_d)_{d \in D} \text{ is eventually in } U_{x_1}^c \cap \dots \cap U_{x_n}^c = \Big(U_{x_1} \cup \dots \cup U_{x_n}\Big)^c = X^c = \emptyset,$$

an obvious contradiction.

Next, assume that every universal net is convergent and let $\mathcal{C}$ be a collection of closed sets with the finite intersection property. Then $\mathcal{C}$ can be extended to an ultrafilter F on $\mathscr{P}(X)$. Consider the directed set $\{(x,Y) \mid x \in Y \in F\}$ with preorder

$$Y \subseteq Y' \implies (x,Y) \succeq (y,Y') \qquad \forall x \in Y, y \in Y'.$$

Note that this is directed by virtue of the finite intersection property. By Proposition 2.78, this is a universal net. Hence it converges to some point x_0.

For any $C \in \mathcal{C}$, the subnet given by $\{(x, Y) \mid Y \subseteq C\}$ is contained entirely in C and non-empty due to C being non-empty. Since a subnet of a convergent net converges to the same points, Proposition 2.67 implies that $x_0 \in C$. Since this was true for all C, it follows that $x_0 \in \bigcap \mathcal{C} \neq \emptyset$. So any collection of closed sets with the finite intersection property has a non-empty intersection, implying compactness. □

Using this, we see that universal nets have a distinctive connection to the properties associated with descriptive frames.

- A general frame $\mathfrak{g}$ is differentiated if and only if every universal net on the $\mathfrak{g}$-space has at most one limit point.
- A general frame $\mathfrak{g}$ is compact if and only if every universal net on the $\mathfrak{g}$-space has at least one limit point.
- A general frame $\mathfrak{g}$ is tight if and only if every two universal nets $(x_d)_{d \in D}$ and $(y_\delta)_{\delta \in \Delta}$ converge to R-related points if they are frequently related.

This inspires another notion of a descriptive completion. One can compactify a frame by adding a limit for each universal net, differentiate a frame by identifying limit sets to a single point and tighten a frame by adding relations between limit points of eventually related universal nets.

However, universal nets are not the right notion. As Proposition 2.78 shows, universal nets correspond to ultrafilters on the powerset algebra, not on the admissible sets. A consequence of this is that non-equivalent ultranets, according to Definition 2.76, may converge to the same points. As such, they do not pose the right notion to construct the descriptive completion, but they do hint at one.

To accomplish such a construction, the notion of universal nets must be weakened. On topological spaces with a basis of clopens, the property that makes ultranets so powerful, being eventually in all subsets or their complements, can be reduced to eventually being in all “topologically relevant” sets or their complements.

Definition 2.80. Let $\mathcal{X}$ be a topological space with a basis $\mathcal{B}$ of clopens. A net $(x_d)_{d \in D}$ is *semi-universal* if for each $B \in \mathcal{B}$ either $(x_d)_{d \in D}$ is eventually in B or it is eventually in B^c. ◀

For many common topological spaces, this notion makes very little sense, as it is quite rare for clopen sets to generate the topology. However, in the context of general frames, this is always guaranteed. As one would hope, this captures the topology in the same way.

Proposition 2.81. Let $\mathcal{X} = (X, \mathcal{T})$ be a topological space with a basis $\mathcal{B}$ of clopens. Then the following are equivalent

1. $\mathcal{X}$ is compact;

2. Any universal net on $\mathcal{X}$ is convergent;

3. Any semi-universal net on $\mathcal{X}$ is convergent.

Proof. The equivalence (1) $\iff$ (2) is given by Proposition 2.79, and the implication (3) $\implies$ (2) is obvious, as any universal net is also semi-universal. To complete the equivalence, the implication (1) $\implies$ (3) is sufficient. The proof is essentially identical to the one for Proposition 2.79.

Suppose for contradiction that $\mathcal{X}$ is compact and there is a semi-universal net $(x_d)_{d\in D}$ on $\mathcal{X}$ that is not convergent. For each point $p \in X$, there must then be an open neighbourhood U_p of p such $(x_d)_{d\in D}$ is not eventually in U_p. Because $\mathcal{B}$ was a basis of clopens, there is a clopen $B_p \in \mathcal{B}$ such that $p \in B_p \subseteq U_p$, so that $(x_d)_{d\in D}$ is not eventually in B_p, meaning by definition it is eventually in B_p^c. As $p \in B_p$, it follows that $\{B_p \mid p \in X\}$ is an open cover of $\mathcal{X}$ so that compactness gives a finite subcover $\{B_{p_1}, \ldots, B_{p_n}\}$. As $(x_d)_{d\in D}$ is eventually in $B_{p_i}^c$ for each $i \leq n$, it follows that $(x_d)_{d\in D}$ is eventually in $B_{x_1}^c \cap \cdots \cap B_{x_n}^c = (B_{x_1} \cup \cdots \cup B_{x_n})^c = X^c = \emptyset$, which is a contradiction. $\square$

Comparable to the equivalence in Definition 2.76, there is a notion of semi-equivalence that is particularly important for semi-universal nets.

Definition 2.82. Let $\mathcal{X} = (X, \mathcal{T})$ be a topological space generated by a basis of clopens $\mathcal{B}$. Two nets $(x_d)_{d\in D}$ and $(y_\delta)_{\delta\in\Delta}$ are semi-equivalent with respect to $\mathcal{B}$ if for every $B \in \mathcal{B}$ the net $(x_d)_{d\in D}$ is eventually in B if and only if $(y_\delta)_{\delta\in\Delta}$ is eventually in B. This will be denoted by $(x_d)_{d\in D} \sim_{\mathcal{B}} (y_\delta)_{\delta\in\Delta}$. The equivalence class of $(x_d)_{d\in D}$ is then written as $[(x_d)_{d\in D}]_{\mathcal{B}}$. ◀

The advantage of this equivalence relation is that on these spaces the superfluous non-topological information present in ultranets up to equivalence that allowed non-equivalent nets to converge to the same point, is no longer present in semi-universal nets up to semi-equivalence.

Proposition 2.83. Let $\mathcal{X}$ be a topological space with a basis $\mathcal{B}$ of clopens and let $(x_d)_{d\in D}$ be a net and x a point. Then $(x_d)_{d\in D}$ converges to x if and only if for every $B \in \mathcal{B}$

$$(x_d)_{d\in D} \text{ is eventually in } B \text{ if and only if } x \in B$$

Proof. The right to left implication is immediate from the definition. The left to right implication follows quite swiftly from contraposition, because if $x \notin B$, then $x \in B^c$, which is also clopen, so that $(x_d)_{d\in D}$ is eventually in B^c, and thus certainly not eventually in B. $\square$

This means that on these spaces, there is an intricate link between semi-equivalence and convergence.

Corollary 2.84. Let $\mathcal{X}$ be a topological space with a basis of $\mathcal{B}$ clopens and let $(x_d)_{d\in D}$ and $(y_\delta)_{\delta\in\Delta}$ be convergent nets. Then they are semi-equivalent with respect to $\mathcal{B}$ if and only if they converge to the same points.

Proof. The left to right direction is immediate from the definition. For the other direction, let x be a shared limit point. Then for each clopen basis element

$$(x_d)_{d\in D} \text{ is eventually in } B \iff x \in B \iff (y_\delta)_{\delta\in\Delta} \text{ is eventually in } B,$$

where both equivalences are given by Proposition 2.83. □

Like universal nets up to equivalence had a natural correspondence to ultrafilters on the powerset algebra in Proposition 2.78, so do semi-universal nets up to equivalence have a natural correspondence with ultrafilters on admissible sets.

Remark 2.85. It is not hard to see that there is a correspondence between semi-universal nets up to semi-equivalence are in the same correspondence to ultrafilters on the admissible sets as universal nets were with ultrafilters on $\mathscr{P}(W)$. This will be made explicit in Theorem 2.88. ◀

The upshot of Proposition 2.81 and Corollary 2.84 is that semi-universal nets up to semi-equivalence capture exactly all failures of compactness on topological spaces with a clopen basis, like $\mathfrak{g}$-spaces.

As Remark 2.66 suggested how nets capture differentiatedness of $\mathfrak{g}$-spaces as the uniqueness of limit points and Proposition 2.73 showed that tightness is captured by nets through the relation of limit points of frequently related nets, this means that a general frame being descriptive is wholly captured by semi-universal nets. Consequently, semi-universal nets can be used to construct a new descriptive completion of general frames by adding, identifying, and relating limit points.

The construction should put to one's mind the construction of the real numbers from Cauchy sequences of rational numbers. The two are not the same, but they resemble one another.

Definition 2.86. The *topological descriptive completion* $\mathfrak{g}_\bullet$ *of* $\mathfrak{g} = (W, R, A)$ is the general frame on all semi-universal nets on W up to semi-equivalence, R-relating all frequently related sets and with admissible sets $\widetilde{a}$ for each $a \in A$ consisting of all semi-universal nets eventually in a.

Formally, the universe $W_\bullet$ is the set of all semi-universal nets on the $\mathfrak{g}$-space up to semi-equivalence, written $[n]_A$ for all such nets n over directed sets of size at most $|W \times A|$.[8] The relation $R_\bullet$ is given by relating equivalent classes of eventually related nets and the collection of admissible sets $A_\bullet$ is given by $\{\widetilde{a} \mid a \in A\}$ where $\widetilde{a}$ is the collection of equivalence classes of universal nets that are eventually in a. ◀

As should be expected, this is a general frame.

Proposition 2.87. If $\mathfrak{g}$ is a general frame, then $\mathfrak{g}_\bullet$ is also a general frame.

Proof. If A is the collection of admissible sets of $\mathfrak{g}$, then observe that $\widetilde{a} \cup \widetilde{b} = \widetilde{a \cup b}$, as being eventually in a or eventually in b implies being eventually in $a \cup b$. Moreover, by definition of semi-universal nets, they are either eventually in a or a^c for each $a \in A$.

[8]The size restriction is imposed to avoid set-theoretic problems.

This means that $\widetilde{a}^c = \widetilde{a^c}$. This shows, together with the De Morgan laws that $A_\bullet$ is a field of sets.

To see that A is closed under $\langle R_\bullet \rangle$, see that for a semi-universal net n the equivalence class $[n]_A \in \langle R_\bullet \rangle \widetilde{a}$ if and only if $n \sim_A n'$ frequently related to ν eventually in a, which is true if and only if n' is frequently in $\langle R \rangle a$. As n' is semi-universal, this means that n' is eventually in $\langle R \rangle a$, since it is not eventually in $(\langle R \rangle a)^c$. As $n \sim_A n'$ were equivalent, this means n is eventually in $\langle R \rangle a$. This demonstrates that $\langle R_\bullet \rangle \widetilde{a} = \widetilde{\langle R \rangle a}$, proving the proposition. $\square$

The connection from Remark 2.85 suggests a connection between the topological descriptive completion from Definition 2.86 and the descriptive completion from Definition 2.50.

Theorem 2.88. *Let $\mathfrak{g}$ be a general frame. Then*

$$\mathfrak{g}_\bullet \cong (\mathfrak{g}^*)_*$$

as general frames. Consequently, $\mathfrak{g}_\bullet$ is descriptive.

Proof. As announced, the proof will use the correspondence from Proposition 2.78. That is, consider the map from $W_\bullet$ to W_* that sends an equivalence class of nets $[n]_A$ to the filter $F_n \cap A$ containing all addmissible sets in which n eventually lands. This is well-defined and injective by definition of semi-equivalence. Moreover, it is surjective, because for any ultrafilter $F \subseteq A$ one can construct like in Formula 2.2 the filter $n_F = (x_d)_{d \in D_F}$ over the directed order $\mathbf{D}_F = (D_F, \preceq)$ given by

$$\begin{aligned} D_F &:= \{(x, a) \mid a \in F \text{ and } x \in a\}; \\ (x, a) &\preceq (x', a') \text{ if and only if } a \supseteq a'; \\ x_{(y,a)} &:= y. \end{aligned}$$

By construction, for any non-empty $a \in F$, there is an $x \in a$ such that if $(y, a') \succeq (x, a)$ if and only if $x_{(y,a')} = y \in a' \subseteq a$, implying that n_F is eventually in a for each $a \in F$. Consequently, n_F is mapped to F, proving surjectivity.

Furthermore, if $[n]_A$ is mapped to $F \subseteq A$, then

$$[n]_A \in \widetilde{a} \text{ if and only if } n \text{ eventually in } a \iff a \in F \iff F \in \widehat{a}.$$

To see that the relation is preserved, suppose that n is a semi-universal net that is frequently related to a semi-universal n' and their equivalence classes are mapped to F and F' in $\mathrm{Uf}A$ respectively. Let $a \in F'$. Then n' is eventually in a, so that n is frequently in $\langle R \rangle a$. By semi-universality, it is thus eventually in $\langle R \rangle a$, implying that $\langle R \rangle a \in F$, showing that F' is an R_*-successor of F.

Finally, suppose that F' is an R_*-successor of F. Then consider the nets $n := n_F$ and $n' := n_{F'}$ as defined above. Then consider any two $(x, a) \in D_F$ and $(y, b) \in D_{F'}$ for $x \in a \in F$ and $y \in b \in F'$. By the R_*-relation, $\langle R \rangle b \in F$, so that $\langle R \rangle b \cap a \in F$ is non-empty. Therefore, there exist $(z_0, \langle R \rangle b \cap a) \succeq (x, a)$ and $(z_1, b) \succeq (y, b)$ such that $z_0 R z_1$. Therefore, n and n' are frequently related, and their equivalence classes thus $R_\bullet$-related. As $(\mathfrak{g}^*)_*$ is descriptive, $\mathfrak{g}_\bullet$ is as well. $\square$

3 Model theory of descriptive models

As the classes of descriptive frames and descriptive models are key to this thesis, it is worth investigating their model theory. In particular, the main theorem of this thesis, the van Benthem Characterisation Theorem for descriptive models, will require a fundamentally different proof from the classical theorem, because, as will be shown in Theorem 3.5, the Compactness Theorem of first-order logic fails on the class of descriptive models.

Apart from considerations of compactness, the model theory of descriptive models will be examined and compared to finite model theory. Some results from classical model theory will be discussed and investigated for the class of descriptive models.

3.1 Compactness theorems for descriptive models

Having discussed two notions of compactness in the previous chapter, both for Boolean algebras and topologies, there is a third notion of compactness that will be of interest for later reference.

Definition 3.1. Let $\mathcal{L}$ be a language and $\mathcal{C}$ be a class of $\mathcal{L}$-models. We say that $\mathcal{C}$ is *compact* over $\mathcal{L}$ if any set of formulae $\Phi \subseteq \mathcal{L}$ is satisfiable in $\mathcal{C}$ if and only if every finite subset of Φ is satisfiable in $\mathcal{C}$. When the language is left unspecified, it is understood to be the first-order language. In such a case, one might say that $\mathcal{C}$ has *the compactness property.* ◀

Theorem 3.2 (Compactness Theorem for first-order logic)**.** *Let* $\mathtt{FOL}(\mathcal{R}, \mathcal{F})$ *be a first-order language with relation symbols (and predicates)* $\mathcal{R}$ *and function symbols* $\mathcal{F}$*. Let* $\mathcal{C}$ *denote its class of models. Then* $\mathcal{C}$ *is compact over* $\mathtt{FOL}$.

This famous theorem often goes by the name of *The Compactness Theorem.* It has an immediate consequence for modal logic.

Corollary 3.3. The class $\mathscr{K}$ of all Kripke models with propositional variables P is compact over $\mathtt{ML}(\mathsf{P})$.

Proof. This is immediate when one recalls that $\mathtt{ML}(\mathsf{P})$ can be seen as a subset of $\mathtt{FOL}(\mathsf{P} \cup \{R\}, \emptyset)$ (through the standard translation from Definition 2.6) and sees that Kripke models over P are exactly the appropriate models for this language. □

This argument shows that the compactness property over $\mathtt{FOL}$ implies compactness over $\mathtt{ML}$. The other direction does not always hold.

Example 3.4. Consider the class consisting only of the model $\mathfrak{M} := (\mathbb{N}, <, V)$ with empty valuation. The following will show that class is compact over `ML` but not over `FOL`.

To see that it is not compact over `FOL`, let x_0 and x_1 be two distinct variables and define recursively a formula $\alpha_n(x)$ stating that x has at least n predecessors:

$$\alpha_0(x_i) := \top;$$
$$\alpha_{n+1}(x_i) := \exists x_{1-i}[Rx_{1-i}x_i \wedge \alpha_n(x_{1-i})].$$

It is easy to see that $(\mathbb{N}, <) \models \alpha_n[m] \iff m \geq n$. The set $\{\alpha_n(x_0)\}_{n\in\mathbb{N}}$ is therefore finitely satisfiable, but it is not satisfiable.

However, $\mathfrak{M}, n \underline{\leftrightarrow} \mathfrak{M}, m$ for all $n, m \in \mathbb{N}$, because the relation only points upwards. Therefore, any finitely satisfiable set $\Phi \subseteq$ `ML` is finitely satisfied in $\mathfrak{M}, 0$, because bisimilar states are modally equivalent. Then every formula $\varphi \in \Phi$ is satisfied in $\mathfrak{M}, 0$. Thus Φ is satisfied in $\mathfrak{M}, 0$. This implies that the class consisting of $\mathfrak{M}$ is compact over `ML`. ◀

Compactness in this sense is a powerful model-theoretic tool. As will be shown in Section 4.1, the classical proof for the van Benthem Characterisation Theorem on the class of all Kripke models relies heavily on Theorem 3.2, proving two powerful lemmas. In contrast, the next theorem will show that the class of descriptive models is not compact over `FOL`. This will prevent the proof of the classical van Benthem Characterisation Theorem to apply directly to the class of descriptive models as well, and necessitate different techniques.

Theorem 3.5. *The class of descriptive models is not compact over* `FOL`*.*

The proof of this theorem requires a lemma that will help reduce the model theory of descriptive models to the model theory of finite models.

Lemma 3.6. An infinite, irreflexive, linear order cannot be given the structure of a descrtiptive frame. Therefore, the subclass of finite, irreflexive, linear orders of the class of descriptive models can be defined by a single first-order sentence.

There will be two proofs given for this lemma. The first proof will make use of nets. An attempt will be made to give an intuition for the result. This proof build on intuitions discussed in Example 2.75 and Figure 2.3.

An infinite, irreflexive, and linear order must have $(\mathbb{N}, <)$ or $(\mathbb{N}, >)$ as a suborder. Intuitively, such a sequence of points in $\mathbb{N}$ must have some limit point ω similar to the one in Figure 2.3 for the frame to remain compact. As illustrated visually in Figure 2.3, tightness, or continuity of the relation, then forces the limit point to be reflexive, contradicting the assumption of irreflexivity.

The second proof will instead operate in basic point-set topology to make it accessible. It is inspired by the first proof, but the nets are removed and the proof is presented in the style of a typical point-set approach.

First proof of Lemma 3.6. Suppose that $\mathfrak{g} = (W, R, A)$ is a descriptive structure on an infinite, irreflexive, linear order (W, R). A standard Bolzano-Weierstrass argument will guarantee it must have an infinite R-increasing sequence or an infinite R-decreasing sequence. After all, let $(s_n)_{n\in\mathbb{N}}$ be an infinite subsequence. If there is an infinite subsequence of dominant elements (elements x_n with Rx_mx_n for all $m > n$), then this subsequence of dominants is an R-descreasing subsequence. If there are only finitely many such elements, that means that for every x_n with n sufficiently large, there is an $m > n$ such that Rx_nx_m, so that the Axiom of (Dependent) Choice guarantees an increasing subsequence.

Assume now that $(x_n)_{n\in\mathbb{N}}$ is such an infinite R-sequence, either increasing or decreasing. From topological compactness, it must have a convergent subnet $(y_d)_{d\in D}$ by Proposition 2.68. Since $n > m$ implies x_mRx_n for the increasing case and x_nRx_m for the decreasing case, it follows that $d \succ d'$ implies $y_{d'}Ry_d$ or it implies $y_dRy_{d'}$. In particular, this establishes that for any $d_0 \in D$ and $d \succ d_0$ the points y_{d_0} and y_d are R-related. As $(D, \succeq)$ is cofinal in $\mathbb{N}$ and thus unbounded in itself, $(y_d)_{d\in D}$ must be frequently related to itself. Because $\mathfrak{m}$ was tight per assumption, R is fully continuous from Proposition 2.73 and hence the limit point of $(y_d)_{d\in D}$ is related to itself. This is in contradiction with the assumption that $\mathfrak{g}$ is irreflexive.

Thus no infinite irreflexive linear order can be given the structure of a descriptive frame. The irreflexive linear orders are definable in the single first-order sentence λ given by

$$\lambda := \qquad \underbrace{\forall x \forall y[(x \equiv y \vee Rxy) \leftrightarrow \neg Ryx]}_{\text{total, irreflexive, and antisymmetric}} \wedge \underbrace{\forall x \forall y \forall z[(Rxy \wedge Ryz) \to Rxz]}_{\text{transitive}}. \tag{3.1}$$

Since all finite irreflexive linear orders can be given the structure of a descriptive frame (with the powerset as a collection of admissible sets) the subclass of finite, irreflexive linear orders in the class of descriptive models can be defined by a single first-order sentence. □

Second proof of Lemma 3.6. Suppose again $\mathfrak{g} = (W, R, A)$ is a descriptive frame such that (W, R) is an infinite, irreflexive linear order. The sets $W \supseteq R[W] \supseteq \cdots \supseteq R^n[W] \supseteq \ldots$ must be non-empty for all n, otherwise $\mathfrak{g}$ would be a finite chain. As such, $C := \bigcap_n R^n[W]$ is closed and non-empty, because the $R^n[W]$ is closed for all n by Lemma 2.34 and the space is topologically compact by assumption. By tightness and irreflexivity, for each $x \in C$ there must be a clopen set a_x such that $x \in a_x$ but $x \notin \langle R\rangle a_x$. As then

$$C \cap \bigcap_{x\in C} \langle R\rangle a_x = \emptyset,$$

topological compactness of $\mathfrak{g}$ gives a finite set $\{x_1, \ldots, x_k\}$ such that the finite intersection $C \cap \langle R\rangle a_{x_1} \cap \cdots \cap \langle R\rangle a_{x_k}$ is empty. Since the x_i are linearly ordered, there must be a least one, say x_1, so that $\langle R\rangle a_{x_1} \subseteq \langle R\rangle a_{x_i}$ for all i by transitivity, yielding $C \cap \langle R\rangle a_{x_1} = \emptyset$.

The fact that $\langle R\rangle a_{x_1}$ is clopen, hence compact, and the fact that $C = \bigcap_n R^n[W]$, implies that there is an $n \in \mathbb{N}$ such that $R^n[W] \cap \langle R\rangle a_{x_1} = \emptyset$, so that $x_1 \notin R^{n+1}[W]$, which contradicts $x_1 \in C$. Thus no such descriptive frame exists.

Again, all finite, irreflexive linear orders are can be made into descriptive frames, so Formula 3.1 defines the subclass of finite, irreflexive linear orders. □

This lemma gives Theorem 3.5 almost immediately.

Proof of Theorem 3.5. From Lemma 3.6, the compactness theorem fails almost immediately. Taking λ to be Formula 3.1 and letting φ_n denote the existence of at least n elements,

$$\varphi_n := \qquad \exists x_1 \cdots \exists x_n \bigwedge_{1 \leq i < j \leq n} \neg x_i \equiv x_j; \qquad\qquad n \in \mathbb{N}$$

Lemma 3.6 implies directly that $\{\lambda\} \cup \{\varphi_n\}_{n \in \mathbb{N}}$ is not satisfiable on the class of descriptive models. However, every finite subset is satisfied on a finite, irreflexive linear order, which can always be given the structure of a descriptive model with the powerset as collection of admissible sets. ☐

As mentioned earlier, while compactness over `FOL` implies compactness over `ML`, the other direction does not hold. The following is another such example.

Theorem 3.7. *The class of descriptive models is compact over* $\mathtt{ML}(\mathsf{P})$.

This is proven through the simple observation that the model often known as the *canonical model* is a descriptive model. This model is usually constructed to infer completeness, but in this case it is reformulated to provide compactness. The construction is subtly different to avoid having to include deduction systems in the thesis, which is why the details have been included. The reader could compare this to [8, Section 4.2].

Definition 3.8. The *canonical model* and *canonical general frame* over the class $\mathscr{K}$ of Kripke frames are constructed as follows:

$$\begin{aligned}
W_{\mathscr{K}} &:= \{T \subseteq \mathtt{ML}(\mathsf{P}) \mid T \cup \{\varphi\} \text{ satisfiable on } \mathscr{K} \iff \varphi \in T\}; \\
R_{\mathscr{K}} &:= \{(T, T') \mid \forall \varphi \in \mathtt{ML}(\mathsf{P}) : \varphi \in T' \implies \Diamond\varphi \in T\}; \\
V_{\mathscr{K}}(p) &:= \{T \mid p \in T\}; \\
\mathfrak{M}_{\mathscr{K}} &:= (W_{\mathscr{K}}, R_{\mathscr{K}}, V_{\mathscr{K}}); \\
A_{\mathscr{K}} &:= \{[\![\varphi]\!]^{\mathfrak{M}_{\mathscr{K}}} \mid \varphi \in \mathtt{ML}\}; \\
\mathfrak{g}_{\mathscr{K}} &:= (W_{\mathscr{K}}, R_{\mathscr{K}}, A_{\mathscr{K}}).
\end{aligned}$$

The sets in $W_{\mathscr{K}}$ will be called *maximally satisfiable.* Compare [8, Definition 4.18]. ◀

Lemma 3.9 (Truth Lemma)**.** For each $T \in W_{\mathscr{K}}$ the equivalence

$$\varphi \in T \iff \mathfrak{M}_{\mathscr{K}}, T \Vdash \varphi$$

holds. Compare [8, Lemma 4.21].

This is the standard Truth Lemma, but is again included and proven because the construction is for satisfiability and not for consistency.

Proof of Lemma 3.9. Notice that for any $\varphi, \psi \in \mathtt{ML}$:

- $\varphi \in T \iff \neg\varphi \notin T$. After all, they are not simultaneously satisfiable but any instance satisfying T must have one or the other, so that maximality gives that one of them must be in T.

- $\varphi \vee \psi \in T \iff \varphi \in T$ or $\psi \in T$. This is immediate, because any instance of one side must be an instance of the other, so that maximality of T is sufficient.

- $\varphi \wedge \psi \in T \iff \varphi \in T$ and $\psi \in T$. Again, any instance of $\varphi \wedge \psi$ is an instance of φ and of ψ and vice versa, from which maximality suffices.

A simple induction on the complexity of formulae will then show that the modal theory of T is exactly T. That is, $\mathrm{Th}_{\mathtt{ML}}(\mathfrak{M}_{\mathscr{K}}, T) = T$.

This is trivial for the propositional variables and the above observations handle the Boolean cases. For the modal case, if $\mathfrak{M}_{\mathscr{K}}, T \Vdash \Diamond\varphi$, then there is an $R_{\mathscr{K}}$-successor T' of T with $\mathfrak{M}_{\mathscr{K}}, T' \Vdash \varphi$, meaning $\varphi \in T'$ by the induction hypothesis and the construction of $R_{\mathscr{K}}$ then gives that $\Diamond\varphi \in T$.

In the other direction, if $\Diamond\varphi \in T$, a pointed model $\mathfrak{M}, w$ satisfying T must have a successor v such that $\mathfrak{M}, v \Vdash \varphi$. But then for each $\psi \in \mathrm{Th}_{\mathtt{ML}}(\mathfrak{M}, v)$ we see that $\Diamond\psi \in T$. Since $\mathrm{Th}_{\mathtt{ML}}(\mathfrak{M}, v)$ is clearly maximal, it follows that $\mathrm{Th}_{\mathtt{ML}}(\mathfrak{M}, v)$ is an $R_{\mathscr{K}}$-successor of T so that $\mathfrak{M}_{\mathscr{K}}, T \Vdash \Diamond\varphi$. □

Lemma 3.10. The canonical general frame is descriptive. Compare [8, Proposition 5.69].

Proof. It is differentiated. After all, suppose $T \neq T'$. Then there is at least one $\varphi \in T$ with $\varphi \notin T'$. Thus $T \in [\![\varphi]\!]^{\mathfrak{M}_{\mathscr{K}}}$ but $T' \notin [\![\varphi]\!]^{\mathfrak{M}_{\mathscr{K}}}$ by the Truth Lemma.

It is tight. Suppose that for each $[\![\varphi]\!]^{\mathfrak{M}_{\mathscr{K}}} \in A_{\mathscr{K}}$ and two maximally satisfiable theories T and T' if $T' \in [\![\varphi]\!]^{\mathfrak{M}_{\mathscr{K}}}$ then $T \in \langle R_{\mathscr{K}} \rangle [\![\varphi]\!]^{\mathfrak{M}_{\mathscr{K}}} = [\![\Diamond\varphi]\!]^{\mathfrak{M}_{\mathscr{K}}}$. By the Truth Lemma, this means that $\varphi \in T' \implies \Diamond\varphi \in T$. So $(T, T') \in R_{\mathscr{K}}$.

Finally, it is compact. This is a consequence of compactness of $\mathscr{K}$ over `ML`. Suppose that $\mathcal{A} \subseteq A_{\mathscr{K}}$ has the finite intersection property. Write $T = \{\varphi \in \mathtt{ML} \mid [\![\varphi]\!]^{\mathfrak{M}_{\mathscr{K}}} \in \mathcal{A}\}$. Let $T_0 \subseteq T$ be finite. Then $\{[\![\varphi]\!]^{\mathfrak{M}_{\mathscr{K}}} \mid \varphi \in T_0\} \subseteq \mathcal{A}$ is a finite subset and thus has a non-empty intersection per assumption. Thus T_0 is satisfiable in $\mathfrak{M}_{\mathscr{K}}$. Therefore T is finitely satisfiable. By compactness of $\mathscr{K}$ over `ML`, it follows that T is satisfiable. So there is some pointed model $\mathfrak{M}, w$ with $T \subseteq \mathrm{Th}_{\mathtt{ML}}(\mathfrak{M}, w) =: \overline{T}$. The latter is a maximally satisfiable set, so the Truth lemma gives $\mathfrak{M}_{\mathscr{K}}, \overline{T} \Vdash \varphi$ for all $\varphi \in T$, yielding $\overline{T} \in \bigcap \mathcal{A}$. □

Proof of Theorem 3.7. Let $T \subseteq \mathtt{ML}$ be a modal theory that is finitely satisfiable on descriptive models. Then it is clearly also finitely satisfiable on Kripke models. By compactness of $\mathscr{K}$ over `ML`, the theory T must then also be satisfiable in a Kripke model, say $\mathfrak{M}, w \Vdash T$. Write $\overline{T} := \mathrm{Th}_{\mathtt{ML}}(\mathfrak{M}, w)$. Then $\overline{T}$ is a maximally satisfiable set of modal formulae. By the Truth Lemma, $\mathfrak{M}_{\mathscr{K}}, \overline{T} \Vdash T$. As $\mathfrak{M}_{\mathscr{K}}$ is a descriptive model by Lemma 3.10, T is thus satisfiable on a descriptive model. □

Exactly the same proof gives compactness over `ML` for a number of classes.

Corollary 3.11. Any subclass of $\mathscr{K}$ that includes the canonical model is compact over `ML`. In particular,

- The class of image-compact general models,
- The class of tight models,
- The class of differentiated models,
- The class of compact models,
- Any intersection of the above,

are all compact over `ML`.

3.2 Model-theoretic methods for descriptive models

The compactness theorem is a vital tool of classical model theory and classes lacking the compactness property require fundamentally different techniques for model-theoretic results. By far the most important such class is the class of finite models, and extensive research has been done in this field [15, 19, 28].

Theorem 3.5 shows that, like the class of finite models, the class of descriptive models does not have the compactness property over `FOL`. This means that similar model-theoretic exploration could be conducted for the class of descriptive models.

To do this, it is useful to first consider the model-theoretic tools used in finite model theory and see how they can be applied and extended to the class of descriptive models.

An important tool for finite model theory has been the Ehrenfeucht-Fraïssé method. This is a game-theoretic technique that has been key to many non-expressibility results in finite model theory. It will be important to this thesis, so it will be lined out carefully here. The presentation will follow [19].

Definition 3.12. Let n be a natural number and $\mathfrak{A}$ and $\mathfrak{B}$ be two first-order structures of the same type. The Ehrenfeucht-Fraïssé game of n moves on $\mathfrak{A}$ and $\mathfrak{B}$ is defined as follows:

- The game is played between two players, commonly referred to as Spoiler and Duplicator or $\forall$ and $\exists$.
- The positions that may occur in the game are of two types:
  1. k-tuples of pairs (a, b) of elements a in $\mathfrak{A}$ and b in $\mathfrak{B}$ for $0 \leq k \leq n$,
  2. k-tuples of which the first $k-1$ components are pairs like above and the final component is a single element either in $\mathfrak{A}$ or in $\mathfrak{B}$ for $1 \leq k \leq n$.
- In positions of type 1 of length strictly less than n, Spoiler may move to any position of type 2 that extends the current position. That is, in a position $((a_1, b_1), \ldots, (a_k, b_k))$, Spoiler may move to any position of the form $((a_1, b_1), \ldots, (a_k, b_k), a_{k+1})$ or of the form $((a_1, b_1), \ldots, (a_k, b_k), b_{k+1})$.

- In positions of type 2, Duplicator may move to any position of type 1 that completes the final component to a pair. That is, in any position $((a_1, b_1), \dots, (a_k, b_k), a_{k+1})$, Duplicator may play any move of the form $((a_1, b_1), \dots, (a_k, b_k), (a_{k+1}, b_{k+1}))$, and in any position of the form $((a_1, b_1), \dots, (a_k, b_k), b_{k+1})$, Duplicator may then move to any position of the form $((a_1, b_1), \dots, (a_k, b_k), (a_{k+1}, b_{k+1}))$.

- When the game reaches a position of type 1 and length n, the game stops. Then Duplicator wins if and only if the final position $((a_i, b_i))_{1 \leq i \leq n}$ generates a function $\{(a_i, b_i)\}_{1 \leq i \leq n}$ that is a *local isomorphism*. That is, if the substructure $\mathfrak{A} \restriction \{a_i\}_{i=1}^n$ of $\mathfrak{A}$ generated by the a_i is isomorphic to the $\mathfrak{B} \restriction \{b_i\}_{i=1}^n$ by mapping a_i to b_i. Substructure generation is formalised in Definition 3.13.

The initial position is always the empty tuple. ◀

It will be instructive to formalise the substructure construction necessary for the winning condition.

Definition 3.13. Let $\mathfrak{A}$ be a first-order structure and S a subset of its universe. Write $\mathcal{F}$ for its collection of functions and constants and $a_f \in \mathbb{N}$ for the arity of $f \in \mathcal{F}$. The *universe generated by* S, written $\mathrm{Gen}(S)$, is recursively defined by

$$\begin{aligned} \mathrm{Gen}_0(S) &:= S; \\ \mathrm{Gen}_{n+1}(S) &:= \mathrm{Gen}_n(S) \cup \{f(s_1, \dots, s_{a_f}) \mid s_1, \dots, s_{a_f} \in \mathrm{Gen}_n(S) \text{ and } f \in \mathcal{F}\}; \\ \mathrm{Gen}(S) &:= \bigcup_{n \in \mathbb{N}} \mathrm{Gen}_n(S). \end{aligned}$$

The *substructure generated by* S, written $\mathfrak{A} \restriction S$, is then defined to be the structure with universe $\mathrm{Gen}(S)$ and predicates and functions given by restricting those from $\mathfrak{A}$. ◀

Example 3.14. Two quick examples of the substructure generation might be useful.

- For structures without function symbols, the substructure generated by a set is simply the predicates restricted to that set. For example, the substructure of $\mathbb{Z}$ with the successor relation generated by $\{0, 1, 2, 4\}$ is the structure displayed in Figure 3.1.

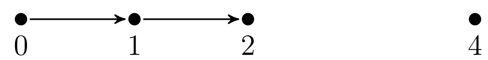

Figure 3.1: The substructure of $\mathbb{Z}$ with the successor relation generated by $\{0, 1, 2, 4\}$.

- For the structure $\mathcal{N}$ with universe $\mathbb{N}$, constant 0 and the successor function, the substructure generated by the empty set is the whole structure. That is, $\mathcal{N} \restriction \emptyset = \mathcal{N}$. Following the definitions, $\mathrm{Gen}_1(\emptyset) = \{0\}$ and by induction $\mathrm{Gen}_{n+1}(\emptyset) = \{0, 1, \dots, n\}$.

Discussion 3.15. The analysis of Ehrenfeucht-Fraïssé games, as with all game-theoretic methods in model theory, focuses on *winning strategies*. Winning strategies are formalised in [19, Definition 2.3.4], but will not be explicitly defined in this thesis. It is our view that the value of game-theoretic methods lies in their intuitive nature and that working with strategy-defining functions undermines this intuition. Instead, the strategies will be carefully verbally explained, trusting that any reader desiring a formal winning strategy can transform these explanations into precise functions.

The Ehrenfeucht-Fraïssé method can be used to prove negative results about the expressivity of first-order logic. That is, it can be used to show that certain structures cannot be distinguished by the language or by parts of the language. An Ehrenfeucht-Fraïssé game of length n can be used to prove equivalence up to "quantifier depth" n.

Definition 3.16. Let α be a first-order formula. The *quantifier-depth* of α, denoted as $q(\alpha)$ is given recursively by

$$\begin{aligned} q(t_1 \approx t_2) &= q(Rt_1, \dots, t_n) = 0; && \text{for any } n\text{-ary relation } R \text{ and terms } t_i \\ q(\alpha \wedge \beta) &= q(\alpha \vee \beta) = \max\{q(\alpha), q(\beta)\}; \\ q(\neg\alpha) &= q(\alpha); \\ q(\exists x\alpha) &= q(\forall x\alpha) = q(\alpha) + 1. \end{aligned}$$

Write $\mathtt{FOL}_n := \{\varphi \in \mathtt{FOL} \mid q(\varphi) \leq n\}$. Two structures $\mathfrak{A}$ and $\mathfrak{B}$ are said to be n-equivalent, notation $\mathfrak{A} \equiv_n^{\mathtt{FOL}} \mathfrak{B}$, if any $\varphi \in \mathtt{FOL}_n$ has $\mathfrak{A} \models \varphi$ if and only if $\mathfrak{B} \models \varphi$. ◀

This definition allows for the formulation of the Ehrenfeucht-Fraïssé Theorem, which demonstrates the power of these games.

Theorem 3.17 (Ehrenfeucht-Fraïssé)**.** *Let n be a natural number and $\mathfrak{A}$ and $\mathfrak{B}$ two first-order structures. If Duplicator has a winning strategy for the Ehrenfeucht-Fraïssé game of n moves, then $\mathfrak{A} \equiv_n^{\mathtt{FOL}} \mathfrak{B}$.*

This method is so powerful because it reduces the abstract and hard to grasp notion of n-equivalence to intuitive reasoning about the similarity between two structures. Arguments about such equivalence usually require complicated inductions on the complexity of formulae, while winning strategies for games can often be formulated intuitively as a reaction plan for Duplicator.

An important application of the Ehrenfeucht-Fraïssé method is Hanf's Lemma, a result showing that the first-order theory of a Kripke model depends only on the local behaviours. The exact meaning of local is given by the notion of the Gaifman-neighbourhood.

Definition 3.18. Let $\mathfrak{F} = (W, R)$ be a frame, $\mathfrak{M} = (\mathfrak{F}, V)$ a model and $S \subseteq W$. Then the *Gaifman neighbourhood of size ℓ of S* is defined recursively by

$$\begin{aligned} N_0^{\mathfrak{F}}(S) &= S; \\ N_{\ell+1}^{\mathfrak{F}}(S) &= N_\ell^{\mathfrak{F}}(S) \cup \langle R \rangle N_\ell^{\mathfrak{F}}(S) \cup R\big[N_\ell^{\mathfrak{F}}(S)\big]; \\ \mathcal{N}_\ell^{\mathfrak{F}}(S) &= (N_\ell^{\mathfrak{F}}(S), R \restriction N_\ell^{\mathfrak{F}}(S)) && \text{where } R \restriction A := R \cap A \times A; \\ \mathcal{N}_\ell^{\mathfrak{M}}(S) &= (\mathcal{N}_\ell^{\mathfrak{F}}(S), V \cap N_\ell^{\mathfrak{F}}(S)) && \text{where } V \cap N_\ell^{\mathfrak{F}}(S) := \Big(p \mapsto V(p) \cap N_\ell^{\mathfrak{F}}(S)\Big). \end{aligned}$$

That is, the ℓ-neighbourhood of S is the set of points that can be reached in ℓ steps along or against R. The letter N is used for the set, and $\mathscr{N}$ is used for the subframe and the submodel. ◀

Theorem 3.19 (Hanf's Lemma [20], presented after [15]). *Let $\mathfrak{M}$ and $\mathfrak{N}$ be Kripke models and let $n \in \mathbb{N}$. Define $\ell = 3^n$. Suppose that the cardinality of the ℓ-neighbourhood of each point in $\mathfrak{M}$ and $\mathfrak{N}$ is bounded from above by $m \in \mathbb{N}$ and for each Kripke model $\mathbb{M}$ one of the following two conditions hold:*

1. *$\mathfrak{M}$ and $\mathfrak{N}$ have the same number of points whose ℓ-Gaifman-neighbourhood is isomorphic to $\mathbb{M}$;*
2. *$\mathfrak{M}$ and $\mathfrak{N}$ both have at least $m \cdot n$ points whose ℓ-Gaifman-neighbourhood is isomorphic to $\mathbb{M}$.*

Then $\mathfrak{M} \equiv_n^{\mathtt{FOL}} \mathfrak{N}$.

With these tools in place some results from classical modeltheory can be investigated.

3.3 Failure of classical model-theoretic results on descriptive models

An important consequence of the failure of the compactness theorem on the class of descriptive models is that standard syntactical deduction systems cannot adequately describe semantic consequence relations on the class of descriptive models.

Theorem 3.20 (Failure of Completeness Theorem). *No first-order deduction system in which the inference rules have a finite number of formulae as input is sound and complete for the class of descriptive models.*

Proof. If it were, then by completeness any semantic consequence $\Psi \models \psi$ could be inferred in this deduction system. If the rules of the deduction system have a finite input, then each step can only use a finite number of previous steps, implying that a derivation of finitely many steps can only use finitely many premises. This means that a ψ can also be deduced from a finite subset $\Psi_0 \subseteq \Psi$ in this system. But then $\Psi_0 \models \psi$ by soundness.

However, the proof of Theorem 3.5 gave a set of formulae Φ such that $\Phi \models \bot$, but no finite subset $\Phi_0 \subseteq \Phi$ entails $\bot$. □

As announced in the introduction, there is a strong similarity between finite model theory and the model theory of descriptive models. As the Compactness Theorem for first-order logic fails on both, many theorems from classical model theory that are proven through compactness fail on these classes.

One such result is an important theorem in classical model theory known as Beth's Definability Theorem. This theorem involves the definability of relations added to a language. The treatment here will follow [15, Section 3.5] for finite model theory.

Definition 3.21. Consider the first-order language $\texttt{FOL}(\mathcal{R},\mathcal{F})$ with a set of relation symbols $\mathcal{R}$ and function symbols $\mathcal{F}$, and let $\mathcal{C}$ be some class of models for this signature. Let $R \notin \mathcal{R}$ be an n-ary relation symbol and $\varphi \in \texttt{FOL}(\mathcal{R} \sqcup \{R\}, \mathcal{F})$ be a formula, where $\sqcup$ denotes disjoint union.

The formula φ *implicitly defines R over $\mathcal{C}$* if for each structure $\mathfrak{A}$ over $\mathcal{R}$ and $\mathcal{F}$, there is a at most one relation $R^{\mathfrak{A}}$ such that if $(\mathfrak{A}, R^{\mathfrak{A}})$ expands $\mathfrak{A}$ to a structure over $\mathcal{R} \sqcup \{R\}$ and $\mathcal{F}$, then $(\mathfrak{A}, R^{\mathfrak{A}}) \models \varphi$.

If $\varphi \in \texttt{FOL}(\mathcal{R} \sqcup \{R\}, \mathcal{F})$ implicitly defines the relation R, then formula $\psi(x_1, \ldots, x_n) \in \texttt{FOL}(\mathcal{R},\mathcal{F})$ *explicitly defines R relative to φ over $\mathcal{C}$* if

$$\varphi \models_{\mathcal{C}} \forall x_1 \cdots \forall x_n [Rx_1 \ldots x_n \leftrightarrow \psi(x_1, \ldots, x_n)].$$

A class $\mathcal{C}$ has the *Beth property* if these two notions of definability coincide. ◀

Theorem 3.22 (Beth's Definability Theorem). *Let $\mathcal{R}$ and $\mathcal{F}$ be collections of relation symbols and function symbols respectively, and suppose that $R \notin \mathcal{R}$ is an n-ary relation symbol. Let $\mathcal{C}$ be the class of all models over this signature. If $\varphi \in \texttt{FOL}(\mathcal{R} \sqcup \{R\}, \mathcal{F})$ implicitly defines R over $\mathcal{C}$, then there exists a formula $\psi(x_1, \ldots, x_n) \in \texttt{FOL}(\mathcal{R},\mathcal{F})$ that explicitly defines R relative to φ over $\mathcal{C}$.*

In [15, Proposition 3.5.1], it is shown that this theorem does not hold when $\mathcal{C}$ is instead taken to be the class of finite models. Following the strategy of this proof shows that the theorem also fails for the class of descriptive models. Lemma 3.6 will allow the proof from [15] to be transferred to the class of descriptive models.

Theorem 3.23 (Failure of Beth Definability Theorem on descriptive models). *Let R be a binary relation symbol and let P be a unary predicate symbol. Then there exists a formula φ that implicitly defines P over the class $\mathfrak{D}$ of descriptive models such that there is no formula $\psi \in \texttt{FOL}(\{R, P\}, \emptyset)$ that explicitly defines P relative to φ over $\mathfrak{D}$.*

Proof. Let λ be Formula 3.1 defining the finite, irreflexive linear orders over the class of descriptive models as per Lemma 3.6. Then take φ to be the formula saying that R is an irreflexive linear order, the R-minimum does not have P and any two immediate successors disagree on P. Formally, φ is

$$\lambda \wedge (\underbrace{\forall x[\neg\exists y\; Ryx \to \neg Px]}_{\text{minimum does not have } P} \wedge \underbrace{\forall x \forall y[\neg\exists z[(Rxz \wedge Rzy) \vee (Rzx \wedge Ryz)] \to (Px \leftrightarrow \neg Px)])}_{\text{successors disagree on } P}.$$

The models of λ are precisely finite linear orders that assign P to exactly the even points.

Suppose now that there were a formula $\psi(x) \in \texttt{FOL}(\{R\}, \emptyset)$ that explicitly defined P. Then the formula ε stating that R is an irreflexive linear order and the maximum has the property ψ, given explicitly by

$$\lambda \wedge \forall x[\neg\exists y\; Rxy \to \psi(x)]$$

would be a formula in $\texttt{FOL}(\{R\}, \emptyset)$ that characterises exactly the finite, even, linear orders. With the Ehrenfeucht-Fraïssé method, this will be shown to be impossible.

Any two finite linear orders $\mathfrak{A}$ and $\mathfrak{B}$ of size larger than 2^{n+1} are n-equivalent. Finite linear orders are discrete, and one can thus speak reasonably about distance. The winning strategy for Duplicator is as follows: for a position $((a_1, b_1), (a_2, b_2), \ldots, (a_{k-1}, b_{k-1}), a_k)$, if a_k is at distance strictly less than 2^{n-k} from the maximum, minimum, or another a_i, then place it on the same side at the same distance from the corresponding point in $\mathfrak{B}$. Otherwise, place it between the b_i as it is between the corresponding a_i at distances all greater than 2^{n-k}. The response to a position $((a_1, b_1), (a_2, b_2), \ldots, (a_{k-1}, b_{k-1}), b_k)$ is symmetric.

Inductively, one can see that this is well-defined, because by the triangle inequality, if a_k is at distance less than 2^{n-k} from both a_i and a_j for $i, j < k$, then a_i is at distance less than $2^{n-k} + 2^{n-k} = 2 \cdot 2^{n-k} = 2^{n-(k-1)}$ from a_j. So b_i is at the same distance from b_j by the induction hypothesis, so that b_k can be chosen at the corresponding distances from both.

Moreover, if there is no point of distance at least 2^{n-k} from both b_i and b_j, then similarly they are at distance at most $2^{n-(k-1)} - 1$ for $i, j < k$. Inductively, a_i and a_j are at the same distance, so that a_k cannot have been placed at a distance at least 2^{n-k} from both.

This induction works for models larger than 2^{n+1}, because the induction basis, that the minimum and the maximum are at distance at least 2^n, is then true.

After n moves, this ensures that $a_i R a_j$ if and only if $b_i R b_j$. Because there are no function symbols, this is a local isomorphism, winning the game for Duplicator. $\square$

The same technique can be used to show the failure of the Craig Interpolation Theorem.

Theorem 3.24 (Craig Interpolation Theorem [12])**.** *Let $\varphi \in \mathtt{FOL}(\mathcal{R}_0, \mathcal{F}_0)$ and $\psi \in \mathtt{FOL}(\mathcal{R}_1, \mathcal{F}_1)$ be formulae and let for i equal to 0 or to 1, the class $\mathcal{C}_i$ contain all structures over $\mathcal{R}_i$ and $\mathcal{F}_i$ and $\mathcal{C}_2$ the class of all structures over $\mathcal{R}_0 \cup \mathcal{R}_1$ and $\mathcal{F}_0 \cup \mathcal{F}_1$. If $\varphi \models_{\mathcal{C}_2} \psi$ then there exists a $\theta \in \mathtt{FOL}(\mathcal{R} \cap \mathcal{R}', \mathcal{F} \cap \mathcal{F}')$ such that $\varphi \models_{\mathcal{C}_0} \theta$ and $\theta \models_{\mathcal{C}_1} \psi$.*

The Craig Interpolation Theorem fails on the class of descriptive models for essentially the same reason that the Beth Definability Theorem fails. Usually, the Beth Definability Theorem is stated as a consequence of the Craig Interpolation Theorem, so that the failure of former immediately implies failure of the latter, but a direct proof will be informative.

Theorem 3.25 (Failure of Craig Interpolation Theorem on descriptive models)**.** *Let R be a binary relation symbol, let P_0, P_1 be unary predicate symbols, and let $\mathscr{D}_0, \mathscr{D}_1$, and $\mathscr{D}_{01}$ denote the class of descriptive models over predicate sets $\{P_0\}$, $\{P_1\}$ and $\{P_0, P_1\}$ respectively. Then there exist formulae $\varphi \in \mathtt{FOL}(\{R, P_0\}, \emptyset)$ and $\psi \in \mathtt{FOL}(\{R, P_1\}, \emptyset)$ such that $\varphi \models_{\mathscr{D}_{01}} \psi$, but there is no formula $\theta \in \mathtt{FOL}(\{R\}, \emptyset)$ such that $\varphi \models_{\mathscr{D}_0} \theta$ and $\theta \models_{\mathscr{D}_1} \psi$.*

Proof. Let φ state that R is an irreflexive linear order, P_0 occurs only on the even points like in the proof of Theorem 3.23, as well as that the maximum satisfies P_0. Let ψ state that R is an irreflexive linear order and if P_1 occurs only on the even points, then the maximum has P_1.

The formula φ is true on the finite, even, irreflexive linear orders with P_0 on the even points. Similarly, for ψ to be true, a structure must be a finite, irreflexive linear order with P_1 true on an odd point or false on an even point, or it must be even. From this, it is easily seen that $\varphi \models_{\mathscr{D}_{01}} \psi$.

Suppose now for contradiction that there is an interpolant $\theta \in \mathtt{FOL}(\{R\}, \emptyset)$ such that $\varphi \models_{\mathscr{D}_0} \theta$ and $\theta \models_{\mathscr{D}_1} \psi$. Then in particular θ must be true on all finite, even, irreflexive linear R-orders. By the Ehrenfeucht-Fraïssé previously presented, this means that θ is true on all sufficiently large finite, irreflexive linear R-orders. But an odd linear order among these can be expanded with a predicate P_1 on the even points, from which it follows that $\theta \not\models_{\mathscr{D}_1} \psi$. This is a contradiction. □

These two results use Lemma 3.6 to reduce a model-theoretic problem for the class of descriptive models to the corresponding result in finite model theory. This reasoning can resolve model-theoretic problems on descriptive models that have an analogue on finite models, but does not help when considering results that have no sensible corresponding statement in finite model theory.

For an example of such results, consider the celebrated upward Löwenheim-Skolem Theorem. It states that any first-order theory T with an infinite model has arbitrarily large models.

Theorem 3.26 (Upward Löwenheim-Skolem Theorem)**.** *Let T be a set of first-order formulae that is satisfiable on an infinite model $\mathfrak{A}$. Then for every cardinal κ larger than the size of $\mathfrak{A}$ there exists a model $\mathfrak{B}$ of T with cardinality at least κ.*

Proof. Introduce to the language a set of constants c_α for any $\alpha \in \kappa$ and define a theory

$$T' := T \cup \{\neg c_\alpha \equiv c_\beta \mid \alpha, \beta \in \kappa\}.$$

This theory is finitely satisfiable on $\mathfrak{A}$ since it is infinite and hence satisfiable by the compactness theorem for first-order logic. A model of T' is a model of T and has at least size κ. □

The upward Löwenheim-Skolem theorem has no meaningful interpretation in finite model theory, because the hypothesis of the implication is trivially false. However, as the theorem is an immediate consequence of the compactness theorem, it is natural to wonder if it holds on the class of descriptive models on which this the compactness property fails. Indeed this class turns out to lack the upward Löwenheim-Skolem property. In fact, it even fails when the theorem is restricted to speak only about formulae.

Theorem 3.27 (Failure of Upward Löwenheim-Skolem Theorem for descriptive models)**.** *There exists a first-order formula that is satisfiable on an infinite descriptive model, but not satisfiable on any uncountable descriptive model.*

Proof. Let φ be a formula encoding a linear order with a reflexive minimum and no other reflexive points. That is, φ is the formula

$$\begin{aligned}
&\forall x \forall y [x \equiv y \vee (Rxy \wedge \neg Ryx) \vee (Ryx \wedge \neg Rxy)] && \text{(antisymmetric \& total)} \\
\wedge\ &\forall x \forall y \forall z [(Rxy \wedge Ryz) \to Rxz] && \text{(transitive)} \\
\wedge\ &\forall x \forall y [Rxx \wedge Ryy \to x \equiv y] && \text{(at most one reflexive point)} \\
\wedge\ &\exists x [Rxx \wedge \forall y\ Rxy] && \text{(there is a reflexive minimum).}
\end{aligned}$$

First, consider the Kripke frame $(\omega + 1, R)$ with Rab if and only if $a > b$ or $a = b = \omega$, equipped with the collection of admissible sets given by the finite subsets of $\mathbb{N}$ and cofinite subsets of $\mathbb{N}$ with ω. See Figure 3.2 for the underlying Kripke frame. This is easily checked to be a descriptive frame.Then any model based on this frame satisfies φ, because the relation is a linear order, has a reflexive minimum and no other reflexive points.

Now suppose that $\mathfrak{m}$ is a descriptive model of cardinality higher than $\aleph_0$ whose underlying frame is a linear order with exactly one reflexive point that is also its minimum. By the pigeon hole principle, there is at least one point w that is not the minimum with infinitely many successors. Then $R[w]$ is a compact subspace (recall Lemma 2.34) and hence induces another infinite descriptive model. By Lemma 3.6, it follows that $R[w]$ contains another reflexive point, which means that $\mathfrak{m}$ contains at least two reflexive points, implying that $\mathfrak{m} \not\models \varphi$. $\square$

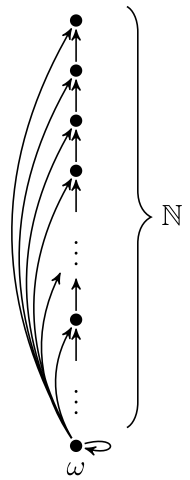

Figure 3.2: Schematic illustration of $(\omega+1, R)$ for R the $\geq$-relation without reflexivity on the natural numbers. Note that not all arrows between the natural numbers are drawn to avoid cluttering, but the relation should be understood to be transitive.

4 The van Benthem Characterisation Theorem for descriptive models

This chapter will prove the main result of this thesis: the van Benthem Characterisation Theorem for the class of descriptive models. The first section will briefly discuss the classical van Benthem Characterisation Theorem. The section after that will discuss the proof strategies for the contexts of Kripke models, finite models and descriptive models. In particular, the failure of the Compactness Theorem for first-order logic will be discussed. It will outline roughly the techniques used in [33] for the class of finite models to avoid the use of the Compactness Theorem and how these will be modified to obtain the result for descriptive models. Finally, the third section of the chapter will develop the constructions announced in the second section and use them to prove the main result.

4.1 The classical van Benthem Characterisation Theorem

This section is intended to provide the reader with insight into the classical van Benthem Characterisation Theorem and the techniques used to prove it. It will not provide a complete proof, but will instead focus on what properties of the class of Kripke models are used to show an explicit contrast with later incarnations of this theorem and their proofs.

Theorem 4.1. ***The classical van Benthem Characterisation Theorem***
Let α be a first-order formula in one free variable and let $\mathscr{K}_*$ denote the class of all pointed Kripke models. Then the following are equivalent:

1. *For all pointed models $\mathfrak{M}, w$ and $\mathfrak{N}, v$ in $\mathscr{K}_*$,*

$$\text{if } \mathfrak{M}, w \underline{\leftrightarrow} \mathfrak{N}, v \text{ then } \mathfrak{M} \models \alpha[w] \iff \mathfrak{N} \models \alpha[v];$$

2. *There exists a modal formula φ such that for all pointed models $\mathfrak{M}, w$ in $\mathscr{K}_*$*

$$\mathfrak{M}, w \Vdash \varphi \text{ if and only if } \mathfrak{M} \models \alpha[w].$$

The first statement will in future be abbreviated by saying that the first-order formula α is (Kripke-)bisimulation-invariant over $\mathscr{K}$. Observe that the implication $2 \implies 1$ is both common knowledge and follows from Corollary 2.60. All that is left is to prove the other implication.

The proof illustrated here will be a carefully presented version of the one given in [8, Section 2.6,p. 103] with explicit comments made on matters of relevance to the later results.

The proof follows from two lemmas. The first lemma requires a fair amount of preceding theory on saturation, which will not be useful in the theory to come, and as such will remain unproven.

Lemma 4.2. Let $\mathfrak{M}, w$ and $\mathfrak{N}, v$ be two modally equivalent pointed Kripke models. Then there exist elementary extensions[1] $\mathfrak{M}^*, w^*$ of $\mathfrak{M}, w$ and $\mathfrak{N}^*, v^*$ of $\mathfrak{N}, v$ such that $\mathfrak{M}^*, w^* \underline{\leftrightarrow} \mathfrak{N}^*, v^*$.

This lemma and the relevant theory can be found in [8, Section 2.6, p. 101] as the "Detour lemma". However, the techniques of the proof cannot be transferred to the case of descriptive models due to the heavy reliance on the Compactness Theorem for first-order logic (see Theorem 3.2), so it will not be helpful to repeat it here. Instead, it is stated without proof. All that matters for this discussion is that the proof uses the compactness property of $\mathscr{K}$ over first-order logic.

The next step towards Theorem 4.1, however, will be instructive to later incarnations of the van Benthem Characterisation Theorem. As such, the proof, using Lemma 4.2, will be spelled out explicitly.

Lemma 4.3. Let α be a first order formula in one free variable and let $\mathscr{K}_*$ be the class of all pointed Kripke models. Suppose that α is bisimulation-invariant over $\mathscr{K}_*$. Then $\mathrm{Th}_{\mathtt{ML}}^{\mathscr{K}_*}(\alpha) \models_{\mathscr{K}_*} \alpha$.

Proof. Write $T := \mathrm{Th}_{\mathtt{ML}}^{\mathscr{K}_*}(\alpha) = \{\varphi \in \mathtt{ML} : \alpha \models_{\mathscr{K}_*} \varphi\}$ for the modal theory of α and assume $\mathfrak{M}, w$ satisfies T. Then in particular $\mathrm{Th}_{\mathtt{ML}}(\mathfrak{M}, w) \cup \{\alpha\}$ is satisfiable. After all, suppose it is not, then by the Compactness Theorem for first-order logic there would be a finite subset $T_0 \subseteq \mathrm{Th}_{\mathtt{ML}}(\mathfrak{M}, w)$ such that $T_0 \cup \{\alpha\}$ is not satisfiable. In particular, $\alpha \models_{\mathscr{K}_*} \neg \bigwedge T_0$ and hence $\neg \bigwedge T_0 \in T$, implying $\mathfrak{M}, w \Vdash \neg \bigwedge T_0$ in obvious contradiction to $T_0 \subseteq \mathrm{Th}_{\mathtt{ML}}(\mathfrak{M}, w)$. So there exists a pointed model $\mathfrak{N}, v$ that is modally equivalent to $\mathfrak{M}, w$ and which satisfies α. Using Lemma 4.2, the pointed Kripke models $\mathfrak{N}, v$ and $\mathfrak{M}, w$ have bisimilar elementary extensions $\mathfrak{N}^*, v^*$ and $\mathfrak{M}^*, w^*$ respectively. In particular, $\mathfrak{N}^* \models \alpha[v^*]$ using the elementary embedding, and bisimilation-invariance of α gives that $\mathfrak{M}^* \models \alpha[w^*]$ yielding $\mathfrak{M} \models \alpha[w]$ by the elementary embedding again. It follows that $\mathrm{Th}_{\mathtt{ML}}^{\mathscr{K}_*}(\alpha) \models_{\mathscr{K}_*} \alpha$. □

From this lemma, the Compactness Theorem implies the van Benthem Characterisation Theorem almost immediately.

Proof of Theorem 4.1. Suppose that α is a first order formula and bisimulation-invariant over the class $\mathscr{K}_*$ of all models. Then Lemma 4.3 gives that $\mathrm{Th}_{\mathtt{ML}}^{\mathscr{K}_*}(\alpha) \models_{\mathscr{K}_*} \alpha$. Compactness now immediately tells us that there is a finite subset $T_0 \subseteq \mathrm{Th}_{\mathtt{ML}}^{\mathscr{K}_*}(\alpha)$ such that

[1] A first-order structure $\mathfrak{A}^*$ is an elementary extension of $\mathfrak{A}$ if there exists an injective homomorphism $h : \mathfrak{A} \to \mathfrak{A}^*$ such that $\mathfrak{A} \models \varphi(a_1, \ldots, a_n)$ if and only if $\mathfrak{A}^* \models \varphi(h(a_1), \ldots, h(a_n))$ for all first-order formulae φ. The map h is called an elementary embedding.

$\bigwedge T_0 \models_{\mathscr{K}_*} \alpha$. Since also $\alpha \models_{\mathscr{K}_*} T$ by construction, the modal formula $\bigwedge T_0$ is equivalent to α on $\mathscr{K}_*$. $\square$

4.2 Modifying the proof structure

The proof from the previous section relies heavily on the Compactness Theorem for first-order logic. Three separate steps in the argument make use of the theorem. However, Theorem 3.5 from the previous chapter stated that the class of descriptive models does not have the compactness property over first-order logic. This means that the techniques used to prove the classical theorem cannot be transferred. A new proof for this class should thus be based on a careful analysis of the structure of the proof and this structure must be modified to not rely on the Compactness Theorem.

In [33], this is done for the class of finite models, which also lacks the compactness property. The approach taken in this thesis will be very close to the one taken there. As such, the proof structure from [33] will be explained, and the idea of proof of the main theorem of this thesis will be announced. The similarities and differences between these approaches will be made explicit.

4.2.1 Proof structure of classical van Benthem Characterisation Theorem

Consider once more the structure of the proof presented for Theorem 4.1. The first lemma that was used was the Detour Lemma. It can be visualised by the following diagram

$$\begin{array}{ccc} \mathfrak{M}, w & \equiv^{\mathtt{ML}} & \mathfrak{N}, v \\ \mathtt{FOL}\,\big\| & & \mathtt{FOL}\,\big\| \\ \mathfrak{M}^*, w^* & \underline{\leftrightarrow} & \mathfrak{N}^*, v^* \end{array} \tag{4.1}$$

where the top line is assumed and the existence of the rest is inferred. While the first order formula α cannot a priori be "transferred" through the relation of modal equivalence, it is preserved under first-order equivalence and, by assumption, bisimulation. This allows nonetheless the conclusion that if $\mathfrak{M} \models \alpha[w]$, then $\mathfrak{N} \models \alpha[v]$ by taking the alternative route through the diagram.

From there, the important steps are Lemma 4.3 and the finishing remarks below that. The former implies that any formula satisfying the requirements of Theorem 4.1 is, at least, modally expressible. That is, there is some set of modal formulae that is exactly equivalent to the formula in question. The remarks that follow then reduce this theory to a single formula.

These three steps are exactly the separate uses of the Compactness Theorem.

- The proof of Lemma 4.2, giving the existence of diagram 4.1, requires a claim that follows from the compactness property.

- If $\mathfrak{M}, w$ satisfies the modal theory of α, the satisfiability of $\mathrm{Th}_{\mathtt{ML}}(\mathfrak{M}, w) \cup \{\alpha\}$ requires the compactness property.
- The final conclusion, that the modal theory of α can be reduced to a single formula, also requires the Compactness Theorem for first-order logic over the class of Kripke models.

When considering classes without the compactness property, these claims have to be replaced by ones that do not rely on the Compactness Theorem.

4.2.2 Proof structure of the van Benthem Characterisation Theorem for finite models

A vital inspiration to this thesis is Lemma 4 from [33], proving the van Benthem Characterisation Theorem for finite models.

Theorem 4.4 (Van Benthem Characterisation Theorem for finite models [33])**.** *Let α be a first order formula in one free variable. Then the following are equivalent:*

1. *There exists a modal formula φ such that for all finite pointed Kripke models $\mathtt{m}, w$*
$$\mathtt{m}, w \Vdash \varphi \text{ if and only if } \mathtt{m} \models \alpha[w].$$
2. *For any two finite pointed Kripke models $\mathtt{m}, w$ and $\mathtt{n}, v$*
$$\text{If } \mathtt{m}, w \mathrel{\underline{\leftrightarrow}} \mathtt{n}, v \text{ then } \mathtt{m} \models \alpha[w] \text{ if and only if } \mathtt{n} \models \alpha[v].$$

The following will be a sketch of the proof. Although the presentation will not be exactly the same as in [33], it will be equivalent.

Again, the implication $1 \implies 2$ is immediate from Corollary 2.60 and only the converse implication requires thought. The proof in [33] avoids the usage of the Compactness Theorem by considering a language that is finite up to equivalence. To this end, fix a first-order formula α in one free variable that is invariant under bisimulations between finite models. Recall from Proposition 2.11 that for a finite set of propositional variables P_0 and a natural number $\ell \in \mathbb{N}$ the language $\mathtt{ML}_\ell(\mathsf{P}_0)$ has this property. Choosing this finite set and number appropriately based on the first-order formula α that is under consideration then allows the author to reduce its finite modal theory to a single formula that can be shown to be equivalent to α.

It stands to reason that the finite set of predicates used to make sure the modal language under consideration is finite up to equivalence, will be the set of all predicates occurring in α. This is made precise in the next definition.

Definition 4.5. Let α be a first-order formula. The *predicate set* of α, denoted $\mathbf{P}(\alpha)$ is defined recursively as

$$\mathbf{P}(Px) = \{P\}; \qquad \mathbf{P}(Rx_1x_2\ldots x_n) = \mathbf{P}(x \approx y) = \emptyset;$$
$$\mathbf{P}(\alpha \vee \beta) = \mathbf{P}(\alpha \wedge \beta) = \mathbf{P}(\alpha) \cup \mathbf{P}(\beta); \qquad \mathbf{P}(\forall x\alpha) = \mathbf{P}(\exists x\alpha) = \mathbf{P}(\neg\alpha) = \mathbf{P}(\alpha)$$

◀

Then $\mathtt{ML}_\ell(\mathsf{P}_0)$ for $\mathsf{P}_0 = \mathbf{P}(\alpha)$ and $\ell = 3^{q(\alpha)}$ will be the language used in the argument. The reasoning then rests on the following lemma.

Lemma 4.6. If α is a first order formula in one free variable that is invariant under bisimulations over the class of finite models. Then α is also invariant under ℓ-bisimulations over the same class, for $\ell = 3^{q(\alpha)+1}$.

As the finite number of propositional variables makes $\equiv_\ell^{\mathtt{ML}}$ and $\underline{\leftrightarrow}_\ell$ coincide (recall Lemma 2.59), this implies that α is preserved under $\equiv_\ell^{\mathtt{ML}}$. Through the finiteness of the language, this then makes it possible to reduce α to a single modal formula. Towards this lemma, fix finite models $\mathtt{m}$ and $\mathtt{n}$ such that $\mathtt{m}, w \underline{\leftrightarrow} \mathtt{n}, v$ and $\mathtt{m} \models \alpha[w]$. The diagram used to prove that $\mathtt{n} \models \alpha[v]$ will not be of the form 4.1, but instead of the form 4.2

$$\begin{array}{ccc} \mathtt{m}, w & \underline{\leftrightarrow}_\ell & \mathtt{n}, v \\ & & \\ |\!\updownarrow & & |\!\updownarrow \\ & & \\ \mathtt{m}^*, w^* & \equiv_n^{\mathtt{FOL}} & \mathtt{n}^*, v^* \end{array} \tag{4.2}$$

for some yet to be specified models $\mathtt{m}^*$ and $\mathtt{n}^*$. Once again, while the formula α can a priori not be assumed to be preserved under $\underline{\leftrightarrow}_\ell^{\mathtt{ML}}$, it is by assumption preserved under bisimulation and for $n \geq q(\alpha)$ it will be preserved under $\equiv_n^{\mathtt{FOL}}$. Note that full first-order equivalence could never be achieved, because they are not even guaranteed to be modally equivalent.

The relation $\equiv_n^{\mathtt{FOL}}$ will be established through the Ehrenfeucht-Fraïssé method from Section 3.2. The exact construction of $\mathtt{m}^*$ is designed specifically to make the application of this tool possible. These will be treelike structures, which have a much simpler structure, allowing for the application of Ehrenfeucht-Fraïssé arguments.

The obvious candidate for such a tree would be the unravelling tree: the tree of all paths through the model given the obvious structure of linking paths that extend one another, in this thesis denoted by $\mathfrak{T}(\mathtt{m})$. A more exact construction can be found in Definition 4.9, but for the purposes of this discussion the intuition behind the structure is much more important than the details. However, this is not in general a finite structure. Any cycles in the original model will produce arbitrarily long paths and thus an infinite unravelling tree. As the hypothesis of the theorem is only that the first-order formula is preserved under bisimulations between finite models, the tree must be in some sense "pruned".

What this means is that the tree is artificially shortened by at some finite distance away from the starting point, putting the original model back into the tree. A sketch of this process can be found in Figure 4.1. This creates what the author calls a pseudotree that has a treelike structure in the initial segment, but breaks this at some sufficient distance from the starting point. This pseudotree based on $\mathtt{m}$ will for now be called $\mathrm{Fin}\mathfrak{T}(\mathtt{m})$.

Note that a model is always bisimilar to its unravelling, and the construction of the pseudotree guarantees that the unravelling is bisimilar to the final result. Since $\mathtt{m}$ and

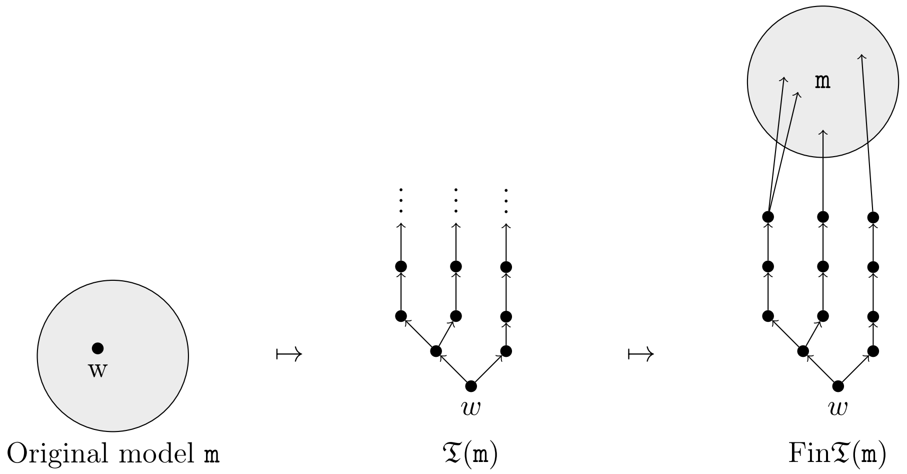

Figure 4.1: Schematic illustration of the pruning process that turns a finite model into a finite pseudotree whose initial structure is treelike.

$\mathrm{Fin}\mathfrak{T}(\mathtt{m})$ are finite, bisimilar structures, bisimulation-invariance of α on finite structures means validity of $\alpha[w]$ in $\mathtt{m}$ guarantees validity of $\alpha[w]$ in $\mathrm{Fin}\mathfrak{T}(\mathtt{m})$. The exact same construction can be performed on $\mathtt{n}$, so that diagram 4.2 can be updated to

$$\begin{array}{ccccccc} \mathfrak{T}(\mathtt{m}), w & \underline{\leftrightarrow} & \mathtt{m}, w & \underline{\leftrightarrow}_\ell & \mathtt{n}, v & \underline{\leftrightarrow} & \mathfrak{T}(\mathtt{n}), v \\ \\ |\!\updownarrow & & |\!\updownarrow & & |\!\updownarrow & & |\!\updownarrow \\ \\ \mathrm{Fin}\mathfrak{T}(\mathtt{m}), w & \underline{\leftrightarrow} & \mathtt{m}^*, w^* & \equiv_n^{\mathtt{FOL}} & \mathtt{n}^*, v^* & \underline{\leftrightarrow} & \mathrm{Fin}\mathfrak{T}(\mathtt{n}), v \end{array} \tag{4.3}$$

with a still unspecified construction $\mathtt{m}^*$ and $\mathtt{n}^*$. Recall that the aim was to perform an Ehrenfeucht-Fraïssé argument on these two models.

From the fact that $\mathtt{m}, w \,\underline{\leftrightarrow}_\ell\, \mathtt{n}, v$ and Lemma 2.61, this finitary bisimulation may be carried to the pseudotrees to give $\mathrm{Fin}\mathfrak{T}(\mathtt{m}), w \,\underline{\leftrightarrow}_\ell\, \mathrm{Fin}\mathfrak{T}(\mathtt{n}), v$. This means that up to the exact number of successors, the first ℓ layers of the treelike part of the pseudotrees will be isomorphic. To illustrate this point, consider the following: if $\mathfrak{M}, w \,\underline{\leftrightarrow}_1\, \mathfrak{N}, v$, then any successor of w must have a matching successor of v with the same propositional type and vice versa. There may of course be more, but there must be at least one. We could, then, in principle, duplicate the points until they had the same number of each type, at which time their successor sets should be isomorphic. The analogous statement is true for $\underline{\leftrightarrow}_\ell$.

By appropriately duplicating points, the pseudotrees can be modified to $\widetilde{\mathrm{Fin}}\mathfrak{T}(\mathtt{m})$ and $\widetilde{\mathrm{Fin}}\mathfrak{T}(\mathtt{n})$ with isomorphic ℓ-neighbourhoods around w and v, while still remaining finite.

For ℓ sufficiently large compared to n, the Ehrenfeucht-Fraïssé game of n moves cannot establish that the non-isomorphic parts are actually in the same connected component.

If there are not enough moves to establish the path, the connection cannot be expressible with n quantifiers. That means intuitively that, for the purposes of Ehrenfeucht-Fraïssé games, the model $\mathtt{M} := \widetilde{\mathrm{Fin}}\mathfrak{T}(\mathtt{m}) \uplus \widetilde{\mathrm{Fin}}\mathfrak{T}(\mathtt{n})$ can be thought of as the disjoint unions of a neighbourhood of w, an isomorphic neighbourhood of v and some points away from both. This establishes the final diagram

$$\begin{array}{ccccccc}
\mathfrak{T}(\mathtt{m}),w & \underline{\leftrightarrow} & \mathtt{m},w & \underline{\leftrightarrow}_\ell & \mathtt{n},v & \underline{\leftrightarrow} & \mathfrak{T}(\mathtt{n}),v \\
|\updownarrow & & & & & & |\updownarrow \\
\mathrm{Fin}\mathfrak{T}(\mathtt{m}),w & & |\updownarrow & & |\updownarrow & & \mathrm{Fin}\mathfrak{T}(\mathtt{n}),v \\
|\updownarrow & & & & & & |\updownarrow \\
\widetilde{\mathrm{Fin}}\mathfrak{T}(\mathtt{m}),w & \underline{\leftrightarrow} & \mathtt{M},w & \equiv_n^{\mathtt{FOL}} & \mathtt{M},v & \underline{\leftrightarrow} & \widetilde{\mathrm{Fin}}\mathfrak{T}(\mathtt{n}),v
\end{array} \tag{4.4}$$

It now follows that, since $\mathtt{m} \models \alpha[w]$, also $\mathtt{M} \models \alpha[w]$, from which equivalence up to quantifier depth $n \geq q(\alpha)$ implies $\mathtt{M} \models \alpha[v]$ as well, so that finally $\mathtt{n} \models \alpha[v]$ by bisimulation-invariance on the class of finite models. Note that, while there are bisimulations in the diagram that are not between finite models, the only bisimulations relevant to this argument are between finite models.

4.2.3 Proof structure of the van Benthem Characterisation Theorem for descriptive models

Now we formulate for the first time the van Benthem Characterisation Theorem for descriptive models.

Theorem 4.7 (Main theorem: van Benthem Characterisation Theorem for descriptive models)**.** *Let α be a first order formula in one free variable. Then the following are equivalent:*

1. *There exists a modal formula φ such that for all pointed, descriptive models $\mathfrak{m}, w$*

$$\mathfrak{m}, w \Vdash \varphi \textit{ if and only if } \mathfrak{m} \models \alpha[w].$$

2. *If $\mathfrak{m}, w$ and $\mathfrak{n}, v$ are two Kripke-bisimilar pointed, descriptive models, then*

$$\mathfrak{m} \models \alpha[w] \iff \mathfrak{n} \models \alpha[v];$$

3. *If $\mathfrak{m}, w$ and $\mathfrak{n}, v$ are two Vietoris-bisimilar pointed, descriptive models, then*

$$\mathfrak{m} \models \alpha[w] \iff \mathfrak{n} \models \alpha[v].$$

This section should be read as a preview sketch of how Section 4.3 will prove this theorem, mentioning explicitly variations on and similarities to the proof for finite models as discussed in the previous section.

The equivalence $2 \iff 3$ is immediate from Corollary 2.56. Again, the implication $1 \implies 2$ is immediate from Corollary 2.60. The only implication that requires work is $2 \implies 1$. Fix a first-order formula α satisfying 2 and consider again the language $\mathtt{ML}_\ell(\mathbf{P}(\alpha))$ for $\ell = 3^{q(\alpha)+1}$. Again the argument relies on a lemma similar to 4.6.

Lemma 4.8. If α is a first order formula in one free variable that is invariant under bisimulations over the class of descriptive models. Then α is also invariant under ℓ-bisimulations over the same class, for $\ell = 3^{q(\alpha)+1}$.

After proving this lemma, the argumentation is identical to the argumentation from the previous subsection. For the lemma, fix descriptive models $\mathfrak{m}$ and $\mathfrak{n}$ with $\mathfrak{m}, w \mathrel{\underline{\leftrightarrow}_\ell} \mathfrak{n}, v$ and $\mathfrak{m} \models \alpha[w]$.

Proving the lemma for descriptive models will be very similar to the approach for finite models. Again, the unravelling tree will be taken and, although it is not in general descriptive, it will be modified to become descriptive. Moreover, a duplication process will be undertaken to turn ℓ-bisimilarity into a isomorphism of neighbourhoods. In this case, the duplication is an indiscriminate copying process making κ many copies of each point with κ an infinite cardinal at least larger than the universes of both $\mathfrak{m}$ and $\mathfrak{n}$. They will be denoted by $\mathfrak{m}^{\otimes\kappa}$ and $\mathfrak{n}^{\otimes\kappa}$, whose precise meaning will be explained later. Then the disjoint union will enable an Ehrenfeucht-Fraïssé argument that finishes the argument. A difference, however, will be the order of these operations. This is not vital, but simply makes matters easier.

$$\begin{array}{ccccccc}
\mathfrak{m}^{\otimes\kappa}, (w,0) & \underline{\leftrightarrow} & \mathfrak{m}, w & \underline{\leftrightarrow}_\ell & \mathfrak{n}, v & \underline{\leftrightarrow} & \mathfrak{m}^{\otimes\kappa}, (v,0) \\
\mathrel{|}\!\updownarrow & & & & & & \mathrel{|}\!\updownarrow \\
\mathfrak{T}(\mathfrak{m}^{\otimes\kappa}), ((w,0)) & & \mathrel{|}\!\updownarrow & & \mathrel{|}\!\updownarrow & & \mathfrak{T}(\mathfrak{n}^{\otimes\kappa}), ((v,0)) \\
\mathrel{|}\!\updownarrow & & & & & & \mathrel{|}\!\updownarrow \\
\widehat{\mathfrak{m}^{\otimes\kappa}}, \widehat{w} & \underline{\leftrightarrow} & \widehat{\mathfrak{m}^{\otimes\kappa}} \uplus \widehat{\mathfrak{n}^{\otimes\kappa}}, \widehat{w} & \equiv_n^{\mathtt{FOL}} & \widehat{\mathfrak{m}^{\otimes\kappa}} \uplus \widehat{\mathfrak{n}^{\otimes\kappa}}, \widehat{v} & \underline{\leftrightarrow} & \widehat{\mathfrak{n}^{\otimes\kappa}}, \widehat{v}
\end{array} \tag{4.5}$$

Then again, from $\mathfrak{m} \models \alpha[w]$ and the fact that α is preserved under bisimulations between descriptive frames, it follows that $\widehat{\mathfrak{m}^{\otimes\kappa}} \uplus \widehat{\mathfrak{n}^{\otimes\kappa}} \models \alpha[\widehat{w}]$. First-order equivalence for $n \geq q(\alpha)$ gives $\widehat{\mathfrak{m}^{\otimes\kappa}} \uplus \widehat{\mathfrak{n}^{\otimes\kappa}} \models \alpha[\widehat{v}]$, so that bisimulation-invariance again provides $\mathfrak{n} \models \alpha[v]$.

Noting then that this means α is preserved under $\equiv_\ell^{\mathtt{ML}}$ and $\mathtt{ML}_\ell(\mathbf{P}(\alpha))$ is finite up to equivalence will provide a single modal formula equivalent to α.

The exact details of the consructions have purposefully been left out, because this was meant to illustrate how the proof will be performed. In this way, the motivation for the constructions in the next section will be more obvious.

4.3 The van Benthem Characterisation Theorem for descriptive models

This section will tackle the proof of the van Benthem Characterisation Theorem for descriptive models, stating that modal logic is the Kripke-bisimulation-invariant fragment and the Vietoris-bisimulation-invariant fragment of first-order logic (see Theorem 4.7). This is the main theorem of the thesis.

4.3.1 Unravelling for descriptive frames

Towards proving the van Benthem theorem for descriptive frames, it will be necessary to modify the notion of unravelling trees and unravelling forests to descriptive frames.

Definition 4.9. Let $\mathfrak{F} = (W, R)$ be a frame. Define the *unravelling forest* of $\mathfrak{F}$ as follows:

$$\begin{aligned} W^{\mathfrak{T}} &:= \{(w_i)_{i\leq n} \in W^n \mid n \in \mathbb{N}, \forall i < n : w_i R w_{i+1}\}; \\ R^{\mathfrak{T}} &:= \{((w_i)_{i\leq n}, (w_i)_{i\leq n+1}) : (w_i)_{i\leq n+1} \in W^{\mathfrak{T}}_w\}; \\ \mathfrak{T}(\mathfrak{F}) &= (W^{\mathfrak{T}}, R^{\mathfrak{T}}). \end{aligned}$$

In words, $W^{\mathfrak{T}}$ is the set of all paths in $\mathfrak{F}$, with two paths connected if one extends the other by one step.

For a specific world $w \in W$, write $\mathfrak{T}_w(\mathfrak{F}) = (W^{\mathfrak{T}}_w, R^{\mathfrak{T}}_w)$ for the connected component of $\mathfrak{T}(\mathfrak{F})$ of all paths starting at w.

For the unravelling forest, there exists a surjective bounded morphism $\pi : \mathfrak{T}(\mathfrak{F}) \twoheadrightarrow \mathfrak{F}$ by sending each path to its endpoint.

Write $l((w_i)_{0\leq i\leq n}) = n$ for the length of a path and $\mathcal{I} = \{\vec{w} \in W^{\mathfrak{T}} : l(\vec{w}) = 0\}$ for the set of initial paths. ◀

Now this classical construction must be modified to incorporate the admissible sets of the general frames. The most canonical way of doing that, pulling back the admissible sets through the projection map π, does not offer a useful solution. The resulting frame would not inherit differentiatedness nor tightness. To accomplish this inheritance, $R^{\mathfrak{T}}$-closure and $\mathcal{I}$ will be added.

Definition 4.10. Let $\mathfrak{g} = (\mathfrak{F}, A)$ be a general frame. Define the *unravelling cover* $\mathfrak{T}(\mathfrak{g})$ of $\mathfrak{g}$ to be $\mathfrak{T}(\mathfrak{g}) := (\mathfrak{T}(\mathfrak{F}), A^{\mathfrak{T}})$, the forest with as underlying frame the classical unravelling forest of $\mathfrak{g}_\#$ (recall Definition 2.14 for notation) and with collection of admissible sets $A^{\mathfrak{T}}$, a subalgebra of $(\mathscr{P}(W^{\mathfrak{T}}), \langle R^{\mathfrak{T}}\rangle)$ defined through the surjective bounded morphism

$\pi : \mathfrak{T}(\mathfrak{F}) \twoheadrightarrow \mathfrak{F}$ by the following recursive schema:

$$
\begin{aligned}
& & & \mathcal{I} \in A^{\mathfrak{T}}; & (1)\\
a \in A & & \Longrightarrow & \pi^{-1}[a] \in A^{\mathfrak{T}}; & (2)\\
b \in A^{\mathfrak{T}} & & \Longrightarrow & R^{\mathfrak{T}}[b] \in A^{\mathfrak{T}}; & (3)\\
b, b' \in A^{\mathfrak{T}} & & \Longrightarrow & b \cup b', b \cap b' \in A^{\mathfrak{T}}. & (4)
\end{aligned}
$$

Note the double usage of $\mathfrak{T}$ for both frames and general frames. This should not be causing too much confusion, but it is worth paying attention to. The unravelling subcovering tree at $w \in W$ will be $\mathfrak{T}_w(\mathfrak{g}) = (\mathfrak{T}_w(\mathfrak{F}), A^{\mathfrak{T}}_w)$ with

$$A^{\mathfrak{T}}_w := \{a \cap W^{\mathfrak{T}}_w \mid a \in A^{\mathfrak{T}}\}.$$

Note that this is not necessarily a cover of $\mathfrak{g}$ as it is not necessarily surjective.

In a similar vein, for a general model $\mathfrak{m} = (\mathfrak{g}, V)$, the unravelling cover will be $\mathfrak{T}(\mathfrak{m}) := (\mathfrak{T}(\mathfrak{g}), \mathscr{P}^{\mathrm{op}}\pi \circ V)$. ◀

Despite these extra additions, the resulting frame is always a general frame.

Proposition 4.11. Let $\mathfrak{g} = (W, R, A)$ be a general frame. Then $\mathfrak{T}(\mathfrak{g})$ is a general frame.

Proof. All to be checked is that $A^{\mathfrak{T}}$ is a field of sets and closed under the $\langle R^{\mathfrak{T}}\rangle$-operation. Closure under binary union and intersection is immediate. Closure under complements can by checked by induction on the construction of $A^{\mathfrak{T}}$. The complement of the initial points $\mathcal{I}^c = R^{\mathfrak{T}}[W^{\mathfrak{T}}]$ is the set of points with a predecessor, and is in $A^{\mathfrak{T}}$ because $W^{\mathfrak{T}} = \pi^{-1}[W]$ is in $A^{\mathfrak{T}}$.

Moreover, $(\pi^{-1}[a])^c = \pi^{-1}[a^c] \in A^{\mathfrak{T}}$ and if $b_0, b_1 \in A^{\mathfrak{T}}$ have $b_0^c, b_1^c \in A^{\mathfrak{T}}$, then $(b_0 \cup b_1)^c = b_0^c \cap b_1^c \in A^{\mathfrak{T}}$ and $(b_0 \cap b_1)^c = b_0^c \cup b_1^c \in A^{\mathfrak{T}}$.

The final case, the $R^{\mathfrak{T}}$-image of an admissible set, requires a minor observation about $R^{\mathfrak{T}}$: distinct points have disjoint image-sets. After all, two distinct paths cannot have the same extension. An alternative way of saying this is that each point has at most one $R^{\mathfrak{T}}$-predecessor. As such, we have for all $a \subseteq W^{\mathfrak{T}}$ that

$$(R^{\mathfrak{T}}[a])^c = R^{\mathfrak{T}}[a^c] \cup \mathcal{I} \in A^{\mathfrak{T}},$$

because each point with no predecessor in a either has a its unique predecessor in a^c or has no predecessor.

For closure under $\langle R^{\mathfrak{T}}\rangle$, note that every element in $A^{\mathfrak{T}}$ can be written as a finite union of finite intersections of elements of the form $\pi^{-1}[a]$ for $a \in A$ admissible, $R^{\mathfrak{T}}[b]$ for $b \in A^{\mathfrak{T}}$ or $\mathcal{I}$. This is evident for elements of the forms (1), (2) or (3) and follows by induction on the construction for elements of the form (4), as intersection distributes over union.

As $\langle R^{\mathfrak{T}}\rangle$ distributes over unions, it is sufficient to show closure for finite intersections. By induction on n, it will be shown that for $b_1, \ldots, b_n$ of the forms (1),(2) and (3), the set $\langle R^{\mathfrak{T}}\rangle(b_1 \cap \cdots \cap b_n)$ is in $A^{\mathfrak{T}}$.

For $n = 1$, note that $\langle R^{\mathfrak{T}}\rangle\mathcal{I} = \emptyset = \pi^{-1}[\emptyset] \in A^{\mathfrak{T}}$ and if $a \in A$ then $\langle \mathbb{R}^{\mathfrak{T}}\rangle\pi^{-1}[a] = \pi^{-1}[\langle R\rangle a] \in A^{\mathfrak{T}}$ because π is a bounded morphism.

Moreover, $\langle R^{\mathfrak{T}}\rangle(R^{\mathfrak{T}}[b])$, again because each point has at most one predecessor, is the subset of b given by points with at least one successor. After all, $x \in \langle R^{\mathfrak{T}}\rangle(R^{\mathfrak{T}}[b])$ if and only if there is a y such that $y \in R^{\mathfrak{T}}[b]$ and $xR^{\mathfrak{T}}y$. That is equivalent y being a successor to x and having a predecessor in b, and since predecessors in this frame are unique, this predecessor must be x. So $\langle R^{\mathfrak{T}}\rangle(R^{\mathfrak{T}}[b]) = \langle R^{\mathfrak{T}}\rangle W^{\mathfrak{T}} \cap b = \pi^{-1}[\langle R\rangle W] \cap b \in A^{\mathfrak{T}}$.

Now suppose that $\langle R^{\mathfrak{T}}\rangle(b_1 \cap \cdots \cap b_n)$ when all b_i are of the forms (1),(2) or (3). Consider then $b_0 \cap b_1 \cap \cdots \cap b_n$ for $n \geq 1$. If one of the b_i is of the form $\mathcal{I}$, then it follows immediately that $\langle R^{\mathfrak{T}}\rangle(b_0 \cap b_1 \cap \cdots \cap b_n) \subseteq \langle R^{\mathfrak{T}}\rangle\mathcal{I} = \emptyset$. If one of them, without loss of generality b_0, is of the form $R^{\mathfrak{T}}[b']$, then observe that $x \in \langle R^{\mathfrak{T}}\rangle(R^{\mathfrak{T}}[b'] \cap b_1 \cap \cdots \cap b_n)$ if and only if $\exists y \in R^{\mathfrak{T}}[x] : y \in R^{\mathfrak{T}}[b'] \cap b_1 \cap \cdots \cap b_n$. By the uniquenss of predecessors, this is equivalent to having $x \in b'$ and $x \in \langle R^{\mathfrak{T}}\rangle(b_1 \cap \cdots \cap b_n)$. As the latter was admissible by the induction hypothesis, it follows that $\langle R^{\mathfrak{T}}\rangle(R^{\mathfrak{T}}[b'] \cap b_1 \cap \cdots \cap b_n) = b' \cap \langle R^{\mathfrak{T}}\rangle(b_1 \cap \cdots \cap b_n)$, which is admissible by closure under intersection.

Finally, if none of the b_i are of the form $\mathcal{I}$ or $R^{\mathfrak{T}}[b']$ for some $b' \in A^{\mathfrak{T}}$, then they must all be of the form $b_i = \pi^{-1}[a_i]$, from which it follows that

$$\begin{aligned} \langle R^{\mathfrak{T}}\rangle(b_0 \cap b_1 \cap \cdots \cap b_n) &= \langle R^{\mathfrak{T}}\rangle(\pi^{-1}[a_0] \cap \cdots \cap \pi^{-1}[a_n]) = \langle R^{\mathfrak{T}}\rangle\pi^{-1}[a_0 \cap \cdots \cap a_n] \\ &= \pi^{-1}[\langle R\rangle(a_0 \cap \cdots \cap a_n)] \in A^{\mathfrak{T}}, \end{aligned}$$

because $\langle R\rangle(a_0 \cap \cdots \cap a_n) \in A$. □

Now towards the main result of the thesis, it is useful to find out which properties of the general frame are transferred to its unravelling cover.

Remark 4.12. For all $a \in A^{\mathfrak{T}}$ and $n \in \mathbb{N}$, the set $(R^{\mathfrak{T}})^n[a]$ is admissible. ◀

To get the van Benthem result for descriptive models, it would be reasonable to hope that properties defining descriptive frames are preserved under this construction. As was hinted to above, this construction in fact does preserve differentiatedness and tightness.

Proposition 4.13. If $\mathfrak{g} = (\mathfrak{F}, A)$ is differentiated, then so is $\mathfrak{T}(\mathfrak{g})$.

Proof. Let $\vec{x} = (x_i)_{i \leq n}, \vec{y} = (y_j)_{j \leq m} \in W^{\mathfrak{T}}$ be distinct paths. If $n > m$, then $\vec{y} \notin (R^{\mathfrak{T}})^n[W^{\mathfrak{T}}] \ni \vec{x}$. So assume $n = m$. Then their distinction must mean there is some $k \leq n$ such that $x_k \neq y_k$. Then by differentiatedness of $\mathfrak{g}$ there is an $a \in A$ such that $y_k \notin a \ni x_k$. But then $(y_i)_{i \leq k} \notin \pi^{-1}[a] \ni (x_i)_{i \leq k}$, meaning $\vec{y} \notin (R^{\mathfrak{T}})^{n-k}[\pi^{-1}[a]] \ni \vec{x}$. □

Proposition 4.14. If $\mathfrak{g} = (\mathfrak{F}, A)$ is differentiated, then $\mathfrak{T}(\mathfrak{g})$ is tight.

Proof. Let $(\vec{x}, \vec{y}) \notin R^{\mathfrak{T}}$. If the length $l(\vec{y})$ of $\vec{y}$ is not $l(\vec{x}) + 1$, then $\vec{y} \in (R^{\mathfrak{T}})^{l(\vec{y})}[\mathcal{I}]$, but $\vec{x} \notin \langle R^{\mathfrak{T}}\rangle(R^{\mathfrak{T}})^{l(\vec{y})}[\mathcal{I}] \subseteq (R^{\mathfrak{T}})^{l(\vec{y})-1}[\mathcal{I}]$, because all paths in $(R^{\mathfrak{T}})^n[\mathcal{I}]$ have length n.

If $n := l(\vec{y}) = l(\vec{x}) + 1$, then there must be a $k \leq l(\vec{x})$ such that $x_k \neq y_k$. By assumption, $\mathfrak{g}$ was differentiated, so there exists an $a \in A$ with $x_k \notin a \ni y_k$, meaning $\vec{y} \in (R^{\mathfrak{T}})^{n-k}[\pi^{-1}[a]] \in A^{\mathfrak{T}}$ but $\vec{x} \notin (R^{\mathfrak{T}})^{n-k-1}[\pi^{-1}[a]] \supseteq \langle R^{\mathfrak{T}}\rangle(R^{\mathfrak{T}})^{n-k}[\pi^{-1}[a]]$. □

Although surprising, it may be understandable that tightness is not required, as the structure of the forest itself already separates unconnected points automatically.

Proposition 4.15. Let $\mathfrak{g} = (W, R, A)$ be a general frame and $\mathfrak{T}(\mathfrak{g})$ its unravelling cover. Take $\vec{x} \in W^{\mathfrak{T}}$. Then $\pi \restriction R^{\mathfrak{T}}[\vec{x}] : R^{\mathfrak{T}}[\vec{x}] \to R[\pi(\vec{x})]$ is a homeomorphism between the subspace topologies.

Proof. It is obviously a continuous bijection as $\pi^{-1}[a] \in A^{\mathfrak{T}}$ for all $a \in A$, which is the basis of the topology on $\mathfrak{g}$. To prove continuity of the inverse, it is sufficient to show that all basis elements in the subspace topology of $R^{\mathfrak{T}}[\vec{x}]$ are of the form $\pi^{-1}[a] \cap R^{\mathfrak{T}}[\vec{x}]$.

By induction on the recursion schema for $A^{\mathfrak{T}}$. It is obvious for the restriction of $\pi^{-1}[a]$ and $\mathcal{I}$. If $b \in A^{\mathfrak{T}}$, then

$$R^{\mathfrak{T}}[b] \cap R^{\mathfrak{T}}[\vec{x}] = \begin{cases} \emptyset & \text{if } \vec{x} \notin b, \\ R^{\mathfrak{T}}[\vec{x}] & \text{if } \vec{x} \in b, \end{cases}$$

by using again that distinct points have disjoint $R^{\mathfrak{T}}$-successor sets.

Finally, π^{-1} distributes over intersection and union, completing the induction. □

Corollary 4.16. Let $\mathfrak{g}$ be an image-compact general frame. Then $\mathfrak{T}(\mathfrak{g})$ is image-compact.

Corollary 4.17. Let $\mathfrak{g}$ be a descriptive frame. Then $\mathfrak{T}(\mathfrak{g})$ is a differentiated, tight, and image-compact.

4.3.2 The descriptive unravelling

The unravelling cover of a compact frame is not necessarily compact, as the following example demonstrates.

Example 4.18. Consider the descriptive frame consisting of a single, reflexive point with the only field of sets possible. Its unravelling cover is $\mathbb{N}$ with the finite and cofinite sets as admissible sets, which is not compact.

This shows that while differentiatedness of $\mathfrak{g}$ implies $\mathfrak{T}(\mathfrak{g})$ is differentiated and tight, compactness may not be preserved. In fact, no collection of admissible sets can be constructed with which an unravelling forest of a descriptive frame with arbitrarily long paths is descriptive.

Proposition 4.19. Let $\mathfrak{g} = (\mathfrak{F}, A)$ be a descriptive frame. If the path lengths in $\mathfrak{g}$ are unbounded, then $\mathfrak{T}(\mathfrak{F})$ cannot be made into a descriptive frame.

Proof. For contraposition, let $(\mathfrak{T}(\mathfrak{F}), \mathcal{A})$ be descriptive. Then in particular from Lemma 2.34 and induction it follows that $(R^{\mathfrak{T}})^n[W^{\mathfrak{T}}]$ is closed for all n. Clearly, $\bigcap_{n\in\mathbb{N}} (R^{\mathfrak{T}})^n[W^{\mathfrak{T}}] = \emptyset$, as any paths is of finite length. By compactness, then, there exists an n such that $(R^{\mathfrak{T}})^n[W^{\mathfrak{T}}] = \emptyset$. Therefore, there exist no paths of length n or more in $\mathfrak{g}$. □

Like for the finite models, the unravelling must be modified to become descriptive. In principle, this could be done in the same way as was done in [33]. Putting the original frame at a sufficiently long distance from $\mathcal{I}$ would suffice for a reproduction of the argument. However, for descriptive frames, there exists an alternative construction that will be used in this thesis: through Jónsson-Tarski duality.

Recall from Definition 2.50 that every general frame has a "descriptive completion", its double dual under the functors **Spec** and **Clop**.

Definition 4.20. Let $\mathfrak{g}$ be a general frame. Then let its *compactified unravelling* or *descriptive unravelling* be the descriptive completion of the unravelling cover of $\mathfrak{g}$. Write

$$\widehat{\mathfrak{g}} := ((\mathfrak{T}(\mathfrak{g}))^*)_*$$

to abbreviate. ◀

This will turn out to be a very well-behaved construction for descriptive frames and will be key to the approach taken in this thesis. Moreover, this construction of the descriptive completion should be much more applicable to related future research. In particular, as will be discussed in Chapters 5 and 6, it will likely be conveniently applicable to achieve results in the modal μ-calculus and coalgebraic constructions. For coalgebraic purposes, too, the cleaner functorial construction of the descriptive completion should be much easier to work with. Section 5.1 will also show that has a very well-behaved underlying frame, in an attempt to pave the way for future research.

For now the focus will be on showing that this construction is well-behaved. In particular, the unravelling forest had a number of useful preservation properties as shown in Section 4.3.1, and the following results will show that these are maintained for the descriptive unravelling.

Lemma 4.21. Let $\mathfrak{g} = (W, R, A)$ be image-compact. Let $F_x, F \in \mathrm{Uf}A$, where F_x is the ultrafilter generated by x. Then in the double dual descriptive frame $F_x R_* F$ if and only if $F = F_y$ for some $y \in R[x]$.

Proof. For the implication from right to left, assume that $y \in R[x]$. To prove that $F_x R_* F_y$, let $a \in F_y$. Then $y \in a$. From xRy it follows that $x \in \langle R \rangle a$, yielding $\langle R \rangle a \in F_x$. Since a was arbitrary, this holds for all $a \in F_y$, so that $F_x R_* F_y$.

For the implication from left to right, assume $F_x R_* F$. By definition, if $a \in F$ then $\langle R \rangle a \in F_x$, implying $x \in \langle R \rangle a$. Therefore, there exists an $x' \in a$ such that xRx'.

So for every $a \in F$, we have $a \cap R[x] \neq \emptyset$. Since F is closed under finite intersections, we find that $\{a \cap R[x] : a \in F\}$ has the finite intersection property. By compactness of $R[x]$, the set $\bigcap \{a \cap R[x] : a \in F\} = R[x] \cap \bigcap F$ is non-empty. Therefore, there exists a $y \in R[x]$ such that for all $a \in F$ we have $y \in a$. So $F = F_y$ for some $y \in R[x]$, because it is an ultrafilter. □

Remark 4.22. One might have expected tightness to show up in the proposition above to prove that $F_x R_* F_y \implies xRy$, but tightness is an immediate consequence from image-compactness and differentiability, so it may be reasonable to expect that image-compactness is strong enough to prove something not quite as strong as tightness. ◀

Remark 4.23. While this result may be surprising from the algebraic construction, it is in fact quite natural from the topological construction of the descriptive completion presented in Definition 2.86. After all, if $(y_d)_{d\in D}$ is a semi-universal net that is frequently related to the constant net at x, then it must be frequently in $R[x]$. Then there is a semi-univeral subnet $(z_\delta)_{\delta\in\Delta}$ that is eventually in $R[x]$ and semi-equivalent to $(y_d)_{d\in D}$, because if a net is eventually in Y, then any subnet is also eventually in Y. This is a semi-universal net in a compact space, so that it converges to a point $z \in R[x]$ by Proposition 2.81. By Corollary 2.84, it must be semi-equivalent to a constant at z. ◀

Proposition 4.24. Let $\mathfrak{g}$ be a differentiated and image-compact frame. Then $\mathfrak{g} \rightarrowtail (\mathfrak{g}^*)_*$ is a generated subframe[2] through a topological embedding $\iota_\mathfrak{g} : x \mapsto F_x$. That is, $\mathfrak{g}$ is homeomorphic to its image under $\iota_\mathfrak{g}$.

Proof. Lemma 4.21 gives immediately that it is a bounded morphism. From the fact that $\mathfrak{g}$ is differentiated, it follows that $\iota_\mathfrak{g}$ injective, making $\mathfrak{g}$ a generated subframe. To see that it is a homeomorphism on its image, let a be an admissible set on $\mathfrak{g}$. Then

$$\iota_\mathfrak{g}(x) = F_x \in \widehat{a} \iff a \in F_x \iff x \in a,$$

so that $\iota_\mathfrak{g}$ and $\iota_\mathfrak{g}^{-1}$ preserve the basis elements of the topology, ensuring continuity for both it and its inverse restricted to the image. □

Theorem 4.25. *Let $\mathfrak{g}$ be a descriptive frame. Then $\widehat{\iota} : \mathfrak{T}(\mathfrak{g}) \rightarrowtail \widehat{\mathfrak{g}}$ continuously.*

Proof. Note that descriptive frames are in particular, image-compact and differentiated, so Corollary 4.16 and Proposition 4.13 gives that $\mathfrak{T}(\mathfrak{g})$ is image-compact and differentiated. Proposition 4.24 then gives the theorem. □

In fact, an even stronger claim is true.

Theorem 4.26. *Let $\mathfrak{g}$ be a descriptive frame. Then $\widehat{\mathfrak{g}}_\# = \mathfrak{T}(\mathfrak{g})_\# \uplus \mathfrak{L}$ for some unspecified frame of limit points $\mathfrak{L}$, where the $\#$-operation takes the underlying Kripke frame of a general frame.*

Proof. Let $\mathfrak{g} = (W, R, A)$. From Theorem 4.25, it is sufficient to show that two points in $\widehat{\mathfrak{g}}$ can only be $(R^\mathfrak{T})_*$-related if they are both inside or both outside $\mathfrak{T}(\mathfrak{g})_\#$. Theorem 4.25 implies that if w is in the image of $\widehat{\iota} : \mathfrak{T}(\mathfrak{g}) \rightarrowtail \widehat{\mathfrak{g}}$, then the $(R^\mathfrak{T})_*$-successor set of w is, too. To complete the theorem, the predecessor set has to be, as well. This means that if $F(R^\mathfrak{T})_* F_x$ for some ultrafilter F and the ultrafilter F_x generated by x, then $F = F_y$ for $y \in (R^\mathfrak{T})^{-1}[x]$.

Towards contraposition, assume that $F \neq F_y$ for any $y \in (R^\mathfrak{T})^{-1}[x]$. In $\mathfrak{T}(\mathfrak{g})$, each point has at most one predecessor, so $(R^\mathfrak{T})^{-1}[x] \subseteq \{y\}$ for some y. In particular, this means there exists some $a \in F \subseteq A^\mathfrak{T}$ such that $x \notin R^\mathfrak{T}[a]$, either because it has no predecessor or because $y \notin a$. By construction of $A^\mathfrak{T}$, the set $R^\mathfrak{T}[a]$ is in $A^\mathfrak{T}$, so that $a \subseteq [R^\mathfrak{T}]R^\mathfrak{T}[a] \in F$ by monotonicity of filters. Since $x \notin R^\mathfrak{T}[a]$, also $R^\mathfrak{T}[a] \notin F_x$, implying that $(F, F_x) \notin (R^\mathfrak{T})_*$. □

[2]Generated subframes are given by injective bounded morphisms

As such, there is is an isomorphic copy of $\mathfrak{T}(\mathfrak{g})_\#$ in $\widehat{\mathfrak{g}}_\#$. A topologically flavoured and more intuitive proof of this fact will be given in Corollary 5.12. The next step is to upgrade the descriptive unravelling to a descriptive model.

Corollary 4.27. Let $\mathfrak{m} = (\mathfrak{g}, V)$ be a descriptive model. Define $\widehat{\mathfrak{m}} := (\widehat{\mathfrak{g}}, \widehat{V})$ with $\widehat{V}(p) = \widehat{\pi^{-1}[V(p)]}$. Then $\pi^T;\widehat{\iota}$, where ; denotes composition of the relations, is a Kripke bisimulation, and hence its closure is a Vietoris bisimulation.

Proof. Note that π and $\widehat{\iota}$ are bounded morphisms, so they satisfy the back and forth conditions by construction. The propositional requirements is satisfied because

$$\pi(\vec{x}) \in V(p) \iff \vec{x} \in \pi^{-1}[V(p)] \iff \pi^{-1}[V(p)] \in F_{\vec{x}} \iff F_{\vec{x}} \in \widehat{\pi^{-1}[V(p)]} = \widehat{V}(p).$$

The final remark is then given by Proposition 2.55. □

4.3.3 Preservation under finite bisimulations

The previous section provides a tool with which to show invariance under bisimulation implies invariance under some finite bisimulation. This tool will now be used to achieve this through Ehrenfeucht-Fraïssé methods. To this end, there is a final combinatorial construction that will prove useful: a duplication process. It will be useful to copy points in a manner that preserves preserves compactness.

Definition 4.28. Let A and B be fields of sets over universes X and Y. Write $A \otimes B$ for the field of sets over the universe $X \times Y$ generated by $\{a \times b \mid a \in A, b \in B\}$. ◀

Proposition 4.29. Let A and B be fields of sets generating topological spaces $\mathcal{X}$ and $\mathcal{Y}$ as bases. Then $A \otimes B$ is a basis for the product space $\mathcal{X} \times \mathcal{Y}$.

Proof. To see that $A \otimes B$ generates a finer topology, note that the basis of $\mathcal{X} \times \mathcal{Y}$ consists of all $U \times V$ where U and V are open subsets of $\mathcal{X}$ and $\mathcal{Y}$ respectively. This means that $U = \bigcup_{\alpha \in I} a_\alpha$ and $V = \bigcup_{\beta \in J} b_\beta$ for $a_\alpha \in A$ and $b_\beta \in B$ and index sets I and J. But then it is immediate that

$$U \times V = \left(\bigcup_{\alpha \in I} a_\alpha \right) \times \left(\bigcup_{\beta \in J} b_\beta \right) = \bigcup_{\alpha \in I, \beta \in J} a_\alpha \times b_\beta$$

is generated by $A \otimes B$.

To see that it is coarser, remark that all elements of $A \otimes B$ are products of sets open in $\mathcal{X}$ with open sets of $\mathcal{Y}$. □

Corollary 4.30. Let A and B be compact fields of sets. Then $A \otimes B$ is compact.

Corollary 4.31. Let A and B be differentiated fields of sets. Then $A \otimes B$ is differentiated.

Proof of Corollaries 4.30 and 4.31. This follows immediately from Proposition 4.29 and the fact that the product of two compact spaces is compact and the product of totally separated spaces is totally separated (recall Definition 2.22). □

Definition 4.32. Let $\mathfrak{g} = (W, R, A)$ be a general frame and let $\mathcal{F}$ be a field of sets over a universe X. Define the $\mathcal{F}$*-multiplier* of $\mathfrak{g}$ by

$$\begin{aligned} R \otimes X &:= \{((w, x_1), (v, x_2)) \in (W \times X)^2 \mid (w, v) \in R, x_1, x_2 \in X\}; \\ \mathfrak{g}^{\otimes \mathcal{F}} &:= (W \times X, R \otimes X, A \otimes \mathcal{F}). \end{aligned}$$

If $\mathfrak{m} = (\mathfrak{g}, V)$ is a general model, then define $V(\cdot) \times X$ by $p \mapsto V(\cdot) \times X)$ and

$$\mathfrak{m}^{\otimes \mathcal{F}} := (\mathfrak{g}^{\otimes \mathcal{F}}, V(\cdot) \times X),$$

which will be called the $\mathcal{F}$-multiplier of $\mathfrak{m}$. ◀

Remark 4.33. There is an obvious surjective continuous bounded morphism $\pi_0 : \mathfrak{F}^{\otimes X} \twoheadrightarrow \mathfrak{F}$ given by projection on the first coordinate. ◀

Lemma 4.34. Let $\mathfrak{g} = (W, R, A)$ be a descriptive frame and let $\mathcal{F}$ over X be a compact and differentiated field of sets. Then $\mathfrak{g}^{\otimes \mathcal{F}}$ is a descriptive frame.

Proof. Compactness and differentiatedness follow immediately from Corollaries 4.30 and 4.31. To see tightness, let $((w, x_1), (v, x_2)) \notin R \otimes X$. Then $(w, v) \notin R$. From tightness of $\mathfrak{g}$ follows the existence of $a \in A$ such that $v \in a$ but $w \notin \langle R \rangle a$. Then in particular, $(v, x_2) \in a \times X$, but $(w, x_1) \notin (\langle R \rangle a) \times X$. It is a general frame in the first place, because

$$\begin{aligned} \langle R \otimes X \rangle (a \times X) &= \{\vec{s} \in W \times X \mid \exists \vec{t} \in a \times X : \vec{s}(R \otimes X)\vec{t}\} \\ &= \{(s, x) \in W \times X \mid \exists (t, y) \in a \times X : sRt\} \\ &= \{s \in W \times X \mid \exists t \in a : sRt\} \times X = (\langle R \rangle a) \times X, \end{aligned}$$

completing the proof. □

As mentioned before, this construction will be used to apply Ehrenfeucht-Fraïssé method. Multiple constructions are conceivable, but the approach taken here is adopted for its convenience. It will be inspired by Hanf's Lemma [20, Lemma 2.3], see also Theorem 3.19. Like Hanf's Lemma, it relies on the notion of the Gaifman neighbourhood as introduced in Definition 3.18. Recall that the ℓ-neighbourhood $\mathcal{N}_\ell^{\mathfrak{M}}$ of a point w in a model $\mathfrak{M}$ is the submodel of all points that can be reached from w through a path of length at most ℓ in which each step moves either to an R-successor or R-predecessor.

Lemma 4.35. Let $\mathfrak{M} = X \uplus \biguplus_{\alpha \in I} B_\alpha$ and $\mathfrak{N} = X \uplus \biguplus_{\rho \in J} C_\rho$ be two (general) models such that

- $B_\alpha \cong B_\beta$ for all $\alpha, \beta \in I$;
- $C_\rho \cong C_\pi$ for all $\rho, \pi \in J$;
- I and J are infinite.

Taking $\ell = 3^n$, suppose that for each point w in B_α there exists a point v in a C_ρ such that $\mathcal{N}_\ell^{\mathfrak{M}}(w) \cong \mathcal{N}_\ell^{\mathfrak{N}}(v)$ and vice versa. Then for w_0 in $\mathfrak{M}$ and v_0 in $\mathfrak{N}$

$$\mathcal{N}_\ell^{\mathfrak{M}}(w_0) \cong \mathcal{N}_\ell^{\mathfrak{N}}(v_0) \implies \mathfrak{M}, w_0 \equiv_n \mathfrak{N}, v_0.$$

Proof. The proof uses the Ehrenfeucht-Fraïssé method (recall Definition 3.12 and Theorem 3.17). By induction on the number of rounds played so far, it will be shown that Duplicator can counter any move by Spoiler. To this end, write $\ell(k) := 3^{n-k}$.

More precisely, let the induction hypothesis denote that for any $(a_i)_{i<k}$ and $(b_i)_{i<k}$ such that

$$X \cup \mathscr{N}_{\ell(k)}\Big(\{w_0\} \cup \{a_i\}_{i<k}\Big) \cong X \cup \mathscr{N}_{\ell(k)}\Big(\{v_0\} \cup \{b_i\}_{i<k}\Big)$$
$$\text{where } \begin{cases} a_i \mapsto b_i, \\ w_0 \mapsto v_0, \\ x \mapsto x & \text{if } x \in X. \end{cases}$$

a move a_k can be countered with a move b_k such that the above condition holds with k replaced by $k+1$. In particular, there will be a local isomorphism between the elements chosen. From symmetry, the response of an a_k to a b_k can be given similarly. Inductively performing this until $k = n$ then gives victory for Duplicator.

Suppose that there have been k turns and the inductive hypothesis holds. When Spoiler makes a move a_k, there are two cases to consider:

1. $a_k \in N_{2\cdot\ell(k+1)}\Big(\{w_0\} \cup \{a_i\}_{i<k}\Big) \cup X$,
2. $a_k \notin N_{2\cdot\ell(k+1)}\Big(\{w_0\} \cup \{a_i\}_{i<k}\Big) \cup X$

In the former case, let $\theta : X \cup N_{\ell(k)}(\{a_i\}_{i<k}) \to X \cup N_{\ell(k)}(\{b_i\}_{i<k})$ be the isomorphism assumed in the induction hypothesis and let $b_k = \theta(a_k)$. Observe that $2 \cdot 3^{n-k-1} + 3^{n-k-1} = 3^{n-k}$, so $N_{\ell(k-1)}(\{a_i\}_{i\le k}) \subseteq N_{\ell(k)}(\{a_i\}_{i<k})$ and the same for b. It follows immediately that the restriction of θ is again an isomorphism between these two smaller neighbourhoods.

In the second case, remark that $N_{\ell(k)}(\{b_i\}_{i<k})$ can only intersect at most k connected components. Because J was infinite per assumption, there is a $\rho \in J$ such that $N_{\ell(k)}(\{b_i\}_{i<k}) \cap C_\rho = \emptyset$. Per assumption, there is a π such that

$$\mathscr{N}_{\ell(0)}(a_k) \cong \mathscr{N}_{\ell(0)}(v) \subseteq C_\pi \cong C_\rho$$

for some v. Therefore, there is a b_k in C_ρ (namely the image of v under the isomorphism directly above) such that $\mathscr{N}_{\ell(0)}(a_k) \cong \mathscr{N}_{\ell(0)}(b_k)$. Choosing it, one obtains

$$\begin{aligned} \mathscr{N}_{\ell(k+1)}(a_k) &\cong \mathscr{N}_{\ell(k+1)}(b_k); \\ X \cup \mathscr{N}_{\ell(k+1)}(\{a_i\}_{i<k}) &\cong X \cup \mathscr{N}_{\ell(k+1)}(\{a_i\}_{i<k}) \qquad \text{through restriction of } \theta; \\ X \cup \mathscr{N}_{\ell(k+1)}(\{a_i\}_{i\le k}) &= \mathscr{N}_{\ell(k+1)}(a_k) \sqcup \big(X \cup \mathscr{N}_{\ell(k+1)}(\{a_i\}_{i<k})\big) \\ &\cong \mathscr{N}_{\ell(k+1)}(b_k) \sqcup \big(X \cup \mathscr{N}_{\ell(k+1)}(\{b_i\}_{i<k})\big) \\ &= X \cup \mathscr{N}_{\ell(k+1)}(\{b_i\}_{i\le k}), \end{aligned}$$

where $\sqcup$ denotes disjoint union and the isomorphism on the different disjoint components is preserved, so that it is still the identity on X, and sends the a_i to the b_i.

Symmetry of the models $\mathfrak{M}$ and $\mathfrak{N}$ guarantees that the exact same argument provides a response a_k to a move b_k by Spoiler.

At the end, this gives Duplicator an isomorphism

$$X \cup \mathscr{N}_{\ell(n)}\Big(\{w_0\} \cup \{a_i\}_{i<n}\Big) \cong X \cup \mathscr{N}_{\ell(n)}\Big(\{w_0\} \cup \{a_i\}_{i<n}\Big)$$

which, using $\ell(n) = 1$, can be restricted to a local isomorphism

$$\{w_0\} \cup \{a_i\}_{i<n} \cong \{v_0\} \cup \{b_i\}_{i<n},$$

winning the game for Duplicator. $\square$

This final lemma now permits a proof of Theorem 4.7

Proof of Theorem 4.7. Remark that the equivalence (2) $\iff$ (3) follows immediately from Corollary 2.56. After all, supposing (2) and assuming $\mathfrak{m}, w \leftrightharpoons \mathfrak{n}, v$, it follows from the corollary that $\mathfrak{m}, w \mathrel{\underline{\leftrightarrow}} \mathfrak{n}, v$ so that the supposition of (2) gives agreement on α. Exactly the same argument provides the other implication.

The implication (1) $\implies$ (2), apart from being an extremely well-known result in modal logic, again follows immediately from Corollary 2.60.

The only interesting implication is the implication towards (1). Towards the implication (2) $\implies$ (1), it will turn out to be sufficient that Kripke-bisimulation-invariance implies finite-bisimulation-invariance. Compare Lemmas 4.6 and 4.8.

Let $\alpha(x)$ be a bisimulation-invariant formula. Write $\ell = 2 \cdot 3^{\mathsf{q}(\alpha)}$ and assume that $\mathfrak{m}, w$ and $\mathfrak{n}, v$ are pointed descriptive models with respective universes W and $\widetilde{w}$ such that

- $\mathfrak{m} \models \alpha[w]$;
- $\mathfrak{m}, w \mathrel{\underline{\leftrightarrow}}_\ell \mathfrak{n}, v$ over $\mathsf{P} := \mathbf{P}(\alpha)$.

Let κ be an infinite cardinal greater than those of the universes of $\mathfrak{m}$ and $\mathfrak{n}$. Then the order topology is compact and totally separated on the ordinal $\kappa + 1$. Write K for the field of clopens of this topology. Remark 4.33 gives that $\mathfrak{m}^{\otimes K}, (w, 0) \mathrel{\underline{\leftrightarrow}}_\ell \mathfrak{n}^{\otimes K}, (v, 0)$ and Corollary 4.27 implies

$$\widehat{\mathfrak{m}^{\otimes K}}, F_{((w,0))} \mathrel{\underline{\leftrightarrow}}_\ell \widehat{\mathfrak{n}^{\otimes K}}, F_{((v,0))}.$$

Write $\widehat{w} := F_{((w,0))}$ and $\widehat{v} := F_{((v,0))}$ to avoid unnecessarily complicated expressions. Observe that the bisimulations provided by Remark 4.33 and Corollary 4.27, together with the bisimulation-invariance of α imply that $\widehat{m^{\otimes K}} \models \alpha[\widehat{w}]$. Showing now that $\widehat{\mathfrak{n}^{\otimes K}} \models \alpha[\widehat{v}]$ would, through the same bisimulations, then show $\mathfrak{n} \models \alpha[v]$, as desired.

Following Theorem 4.26, the models $\widehat{\mathfrak{m}^{\otimes K}}$ and $\widehat{\mathfrak{n}^{\otimes K}}$ can be written as

$$\begin{aligned}
\widehat{\mathfrak{m}^{\otimes K}} &= \mathfrak{L} \uplus \mathfrak{T}(\mathfrak{m}^{\otimes K}) \stackrel{*}{=} \mathfrak{L} \uplus \biguplus_{\substack{\alpha\in\kappa+1 \\ w\in W}} \mathfrak{T}_{(w,\alpha)}(\mathfrak{m}^{\otimes K}); \\
\widehat{\mathfrak{n}^{\otimes K}} &= \Lambda \uplus \mathfrak{T}(\mathfrak{n}^{\otimes K}) \stackrel{*}{=} \Lambda \uplus \biguplus_{\substack{\alpha\in\kappa+1 \\ v\in \widetilde{W}}} \mathfrak{T}_{(v,\alpha)}(\mathfrak{n}^{\otimes K}); \\
\widehat{\mathfrak{m}^{\otimes K}} \uplus \widehat{\mathfrak{n}^{\otimes K}} &= (\mathfrak{L} \uplus \Lambda) \uplus \biguplus_{\alpha\in\kappa+1} \Big(\biguplus_{w\in W} \mathfrak{T}_{(w,\alpha)}(\mathfrak{m}^{\otimes K}) \uplus \biguplus_{v\in\widetilde{W}} \mathfrak{T}_{(v,\alpha)}(\mathfrak{n}^{\otimes K}) \Big),
\end{aligned}$$

where the equalities marked by $*$ follow from the fact that two paths can only be related in the tree if one extends the other, and two paths with different starting points cannot extend one another. The submodels $\mathfrak{L}$ and Λ are unspecified frames of limit points.

For any $\alpha, \beta \in \kappa + 1$,

$$\biguplus_{w\in W} \mathfrak{T}_{(w,\alpha)}(\mathfrak{m}^{\otimes K}) \uplus \biguplus_{v\in\widetilde{W}} \mathfrak{T}_{(v,\alpha)}(\mathfrak{n}^{\otimes K}) \cong \biguplus_{w\in W} \mathfrak{T}_{(w,\beta)}(\mathfrak{m}^{\otimes K}) \uplus \biguplus_{v\in\widetilde{W}} \mathfrak{T}_{(v,\beta)}(\mathfrak{n}^{\otimes K})$$

through the map of switching the initial point. Moreover,

$$\begin{array}{ccc}
\widehat{\mathfrak{m}^{\otimes K}}, \widehat{w} \underline{\leftrightarrow} \widehat{\mathfrak{m}^{\otimes K}} \uplus \widehat{\mathfrak{n}^{\otimes K}}, \widehat{w} & \Longrightarrow & \widehat{\mathfrak{m}^{\otimes K}} \uplus \widehat{\mathfrak{n}^{\otimes K}} \models \alpha[\widehat{w}] \\
\text{and} & & \\
\widehat{\mathfrak{n}^{\otimes K}}, \widehat{v} \underline{\leftrightarrow} \widehat{\mathfrak{m}^{\otimes K}} \uplus \widehat{\mathfrak{n}^{\otimes K}}, \widehat{v}. & &
\end{array}$$

Therefore, if $\mathcal{N}_\ell^{\widehat{\mathfrak{m}^{\otimes K}}}(\widehat{w}) \cong \mathcal{N}_\ell^{\widehat{\mathfrak{n}^{\otimes K}}}(\widehat{v})$, applying Lemma 4.35 for

$$\begin{aligned}
\mathfrak{M} = \mathfrak{N} &:= \widehat{\mathfrak{m}^{\otimes K}} \uplus \widehat{\mathfrak{n}^{\otimes K}}; \\
I = J &:= \kappa + 1; \\
B_\alpha = C_\alpha &:= \biguplus_{w\in W} \mathfrak{T}_{(w,\alpha)}(\mathfrak{m}^{\otimes K}) \uplus \biguplus_{v\in\widetilde{W}} \mathfrak{T}_{(v,\alpha)}(\mathfrak{n}^{\otimes K}); \\
X &= \mathfrak{L} \uplus \Lambda
\end{aligned}$$

will prove that

$$\mathfrak{M}, \widehat{w} \equiv_{\mathrm{q}(\alpha)} \mathfrak{N}, \widehat{v}.$$

This implies $\widehat{\mathfrak{m}^{\otimes K}} \uplus \widehat{\mathfrak{n}^{\otimes K}} \models \alpha[\widehat{v}]$, from which the bisimulation above provides $\widehat{\mathfrak{n}^{\otimes K}} \models \alpha[\widehat{v}]$, which in turn shows $\mathfrak{n} \models \alpha[v]$ from the previous bisimulations. This implies that α is preserved under finitary bisimulations. As Lemma 2.59 implies that bisimilarity up to depth ℓ over finitely many propositional variables is identical to modal equivalence up to depth ℓ, this means that α is characterised by some collection of theories in $\mathtt{ML}_\ell$. As there are finitely many such theories, and each can be characterised by a single formula, there is a modal formula of depth up to ℓ equivalent to α, concluding the proof.

This final claim of isomorphism of neighbourhoods will be achieved by showing inductively that there exists a sequence of isomorphisms $(f_i)_{i\leq\ell}$ with

$$
\begin{aligned}
&f_i : \mathcal{N}_i^{\widehat{\mathfrak{m}^{\otimes K}}}(\widehat{w}) \cong \mathcal{N}_i^{\widehat{\mathfrak{n}^{\otimes K}}}(\widehat{v}); && \\
&f_i \restriction \mathcal{N}_j^{\widehat{\mathfrak{m}^{\otimes K}}}(\widehat{w}) = f_j && \text{for } j \leq i; \\
&N_i^{\widehat{\mathfrak{m}^{\otimes K}}}(\widehat{w}) \ni x \leftrightsquigarrow_{\ell-i} f_i(x) \in N_i^{\widehat{\mathfrak{n}^{\otimes K}}}(\widehat{v}) && \text{over } \mathsf{P} = \mathbf{P}(\alpha)
\end{aligned}
$$

Clearly, the map $f_0 : \mathcal{N}_0^{\widehat{\mathfrak{m}^{\otimes K}}}(\widehat{w}) \to \mathcal{N}_0^{\widehat{\mathfrak{n}^{\otimes K}}}(\widehat{v})$ is an isomorphism, since both consist of a single irreflexive point and must satisfy the same propositional variables, because $\mathfrak{m}^{\otimes K}, \widehat{w} \underline{\leftrightarrow}_\ell \mathfrak{n}^{\otimes K}, \widehat{v}$, so they satisfy the same propositional variables, implying that f_0 is an isomorphism. Moreover, through Lemma 2.59, it implies that $\widehat{w} \leftrightsquigarrow_\ell \widehat{v}$.

Now suppose that $f_i : \mathcal{N}_i^{\widehat{\mathfrak{m}^{\otimes K}}}(\widehat{w}) \to \mathcal{N}_i^{\widehat{\mathfrak{n}^{\otimes K}}}(\widehat{v})$ is an isomorphism. Since $\widehat{w}$ is the root of its treelike connected component, if R is the relation on $\widehat{\mathfrak{m}^{\otimes K}}$, then

$$
N_j^{\widehat{\mathfrak{m}^{\otimes K}}}(\widehat{w}) = \bigcup_{k\leq j} R^k[\widehat{w}],
$$

so the isomorphism f_{i+1} only needs to extend f_i on the successors of $R^i[\widehat{w}]$. Because the connected component under consideration is a tree, no two points share successors. Therefore, the successors of each such point may be considered separately.

Let $x \in R^i[\widehat{w}]$. Each theory in $\mathtt{ML}_{\ell-i-1}(\mathsf{P})$ has a single characterising formula, as the language is finite. Write Σ for the set of these formulae. For each $\varphi \in \Sigma$, observe that if $\mathfrak{m}^{\otimes K}, x \Vdash \Diamond\varphi$, then this type occurs κ many times in $R[x]$ and otherwise it occurs 0 times, because if it occurs once, then it must occur in all the K-duplicates. The same reasoning applies to $f_i(x)$, so because they agree on each $\Diamond\varphi$ per assumption, for each theory x and $f_i(x)$ either both have κ many successors with that theory, or none.

It follows that for every $\varphi \in \Sigma$, there exists a bijection g_φ^x between the successors of x satisfying φ and the successors of $f_i(x)$ satisfying φ. Choosing one such g_φ^x for each φ allows the construction of

$$
f_{i+1} = f_i \sqcup \bigsqcup_{\substack{x\in R^i[\{\widehat{w}\}] \\ \varphi\in\Sigma}} g_\varphi^x,
$$

this preserves all new relations and predicates. Moreover, $y \in R[x]$ with $x \in R^i[\widehat{w}]$ has modal theory characterised by $\varphi \in \Sigma$, by construction $y \leftrightsquigarrow_{\ell-i-1} g_\varphi^x(y)$. Since this also held for f_i by induction hypothesis, the inductive condition is satisfied again. $\square$

5 Further analysis of the descriptive unravelling

In terms of results, the key contribution of this thesis is the main theorem, the van Benthem characterisation theorem. The essential technique in this proof, the descriptive unravelling, is based on the more classical unravelling forest or unravelling tree, which has proven an extremely powerful tool for the class of Kripke structures. Because the limit points that disturb the treelike structure of the construction cannot be reached from the treelike points, these constructions behave much alike and this may prove an exceptionally useful tool when one is restricted to descriptive frames.

With this in mind, future research may benefit from this construction. To facilitate this, a careful analysis of the exact structure of the descriptive unravelling is in order. While usually the properties of the tree-like components will be most important, the remaining structure may also be enlightening.

5.1 The $\mathfrak{g}_{-\infty}$-operation

Up to this point, the limit points in the descriptive unravelling have been treated rather like a black box. They were unimportant for the argument and could safely be ignored. However, future research using this technique might be forced to dive into the precise details of these points and require them to be well-behaved.

As it turns out, the limit point structure is conveniently similar in structure to the unravelling cover that is being completed. Specifically, Theorem 5.17 will show that they can be recreated with a similar path-like construction and a comparable recursively defined collection of admissible sets.

In this analysis, the topological toolbox introduced in Section 2.6.1 will prove exceptionally useful. Of course, because they are equivalent as proved in Theorem 2.88, the proof techniques should translate immediately to proofs using only the ultrafilters, but because ultrafilters not generated by elements are notoriously unintuitive and difficult to work with, this will be undesirable. The topological toolbox will provide intuition into the workings of the descriptive completion.

Recall from Theorem 4.26 that $\widehat{\mathfrak{g}}_{\#} = \mathfrak{T}(\mathfrak{g})_{\#} \uplus \mathfrak{L}$ for some unspecified frame of limit points $\mathfrak{L}$. The argument was based exclusively on the properties of filters and the result provided little intuition. This section aims to pinpoint the precise structure of $\mathfrak{L}$. First a new operation on descriptive frames is introduced, after which it is shown that this operation gives the frame of limit points. In the next two sections, this operation is used to understand the exact structure of $\mathfrak{L}$.

It will turn out that the limit points in $\widehat{\mathfrak{g}}$ have the structure of what will be called the $\mathfrak{g}_{-\infty}$-frame. This will be a general frame consisting of all infinitely descending paths in $\mathfrak{g}$ with appropriate relation and collection of admissible sets.

Definition 5.1. Let $\mathfrak{F} = (W, R)$ be a frame. Then define $\mathfrak{F}_{-\infty}$ to be the frame of infinitely descending paths defined by $F_{-\infty} = (W_{-\infty}, R_{-\infty})$, where $W_{-\infty}$ is the set of infinitely descending R-paths, formally defined by

$$W_{-\infty} := \{(w_n)_{n\in\mathbb{N}} \mid w_{n+1} R w_n\}$$

and $R_{-\infty}$ connects paths that extend one another by one step. That is, $(w_n)_{n\in\mathbb{N}}$ is an $R_{-\infty}$-successor of $(v_m)_{m\in\mathbb{N}}$ if and only if $v_n = w_{n+1}$. ◀

Remark 5.2. The convention for these paths is different from the convention for paths in the tree from Definition 4.9. Specifically, these are formulated as R-descending sequences, while in the unravelling tree, they are R-ascending sequences. The latter choice was made to conform to convention for the unravelling tree. The sequences in $W_{-\infty}$ have been chosen to be R-descending to avoid cumbersome notation. In Figure 5.1, a few graphical examples of these paths have been given to clarify the point. ◀

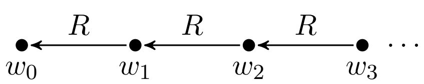

(a) An illustration of a path that becomes a point $(w_n)_{n\in\mathbb{N}}$ in $W_{-\infty}$. Make special note of the direction of R.

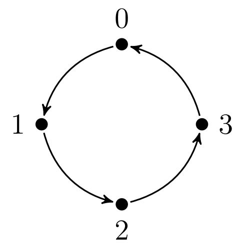

(b) A cycle that would produce a point in $W_{-\infty}$ that would be written as $(0, 3, 2, 1, 0, 3, 2, 1, \dots)$.

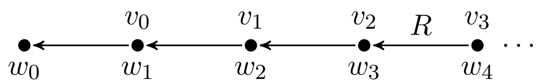

(c) An illustration of two points $\vec{w} = (w_n)_{n\in\mathbb{N}}$ and $\vec{v} = (v_n)_{n\in\mathbb{N}}$ such that $\vec{v} R_{-\infty} \vec{w}$ (that is, the former is an $R_{-\infty}$-successor of the latter).

Figure 5.1: A few figures to illustrate the structure on $\mathfrak{F}_{-\infty}$.

Note that the tail of any sequence $(w_n)_{n\in\mathbb{N}} \in W_{-\infty}$ is still in $W_{-\infty}$. This motivates a convenient notation.

Definition 5.3. Let $\vec{w} := (w_n)_{n\in\mathbb{N}} \in W_{-\infty}$ be an infinitely descending R-chain. Then the k-shift of $\vec{w}$, denoted by $\vec{w}[k]$, is the sequence $(w_{n+k})_{n\in\mathbb{N}} = (w_k, w_{k+1}, w_{k+2}, \dots)$. In this notation, it follows that $\vec{v} R_{-\infty} \vec{w}$ if and only if $\vec{v} = \vec{w}[1]$. See Figure 5.1c. ◀

As with the unravelling forest, there is a natural map from $\mathfrak{F}_{-\infty}$ to $\mathfrak{F}$ sending a path to its endpoint.

Proposition 5.4. Let $\pi_{-\infty} : \mathfrak{F}_{-\infty} \to \mathfrak{F}$ be given by sending an infinite path $(w_n)_{n\in\mathbb{N}}$ to the point w_0. Then $\pi_{-\infty}$ is a bounded morphism.

Proof. Let $\vec{w} = (w_n)_{n\in\mathbb{N}}$ be an $R_{-\infty}$-successor to $\vec{v} = (v_m)_{m\in\mathbb{N}} = (w_{n+1})_{n\in\mathbb{N}}$. Then note in particular that $w_0 = \pi_{-\infty}(\vec{w})$ is a R-successor to $w_1 = v_0 = \pi_{-\infty}(\vec{v})$ because $\vec{w}$ is an infinite R-descending path.

Conversely, if $\vec{w} = (w_n)_{n\in\mathbb{N}}$ is an infinite path and v is a successor of $\pi_{-\infty}(\vec{w})$, then it follows that $\vec{v} := (v, w_0, w_1, \dots)$ is an infinite descending R-path such that $\pi_{-\infty}(\vec{v}) = v$ and $\vec{v}$ is by definition an $R_{-\infty}$-successor to $\vec{w}$. $\square$

This map $\pi_{-\infty}$ also allows for this construction to be extended to general frames by taking the same admissible sets as were used for the unravelling forest.

Definition 5.5. Let $\mathfrak{g} = (\mathfrak{F}, A)$ be a general frame based on a frame $\mathfrak{F} = (W, R)$. Then define $\mathfrak{g}_{-\infty} = (\mathfrak{F}_{-\infty}, A_{-\infty})$ to be the general frame based on $\mathfrak{F}_{-\infty}$ with admissible sets given by the recursion

$$\begin{array}{lcl} a \in A & \Longrightarrow & (\pi_{-\infty})^{-1}[a] \in A_{-\infty}; \\ b \in A_{-\infty} & \Longrightarrow & R_{-\infty}[b] \in A_{-\infty}; \\ b, b' \in A_{-\infty} & \Longrightarrow & b \cup b', b \cap b' \in A_{-\infty}. \end{array}$$

◀

Proposition 5.6. Let $\mathfrak{g}$ be a general frame. Then $\mathfrak{g}_{-\infty}$ is a general frame.

Proof. It should be checked that $A_{-\infty}$ is a field of sets closed under $\langle R_{-\infty}\rangle$. The recursion is effectively identical to the one for $A^{\mathfrak{T}}$, missing only the addition of the set of initial points, since there are no points without predecessor[1] and since in both each point has at most one predecessor, which is the only property used in the proof for Proposition 4.11, the exact same proof provides the theorem here. $\square$

Analogues to Propositions 4.13 and 4.14 for the unravelling forest can be found, showing that $\mathfrak{g}_{-\infty}$ inherits differentiatedness and the combination of differentiatedness and tightness for the same reasons.

Proposition 5.7. Let $\mathfrak{g}$ be a differentiated frame. Then $\mathfrak{g}_{-\infty}$ is differentiated.

Proof. Let $\vec{w} = (w_n)_{n\in\mathbb{N}}$ and $\vec{v} = (v_n)_{n\in\mathbb{N}}$ be distinct infinitely R-descending paths. Then there must be a least $m \in \mathbb{N}$ such that $w_n \neq v_n$. As a consequence of differentiatedness of $\mathfrak{g}$, there must be an admissible a such that $w_n \in a$ but $v_n \notin a$. Then $\vec{w}[m] \in (\pi_{-\infty})^{-1}[a]$ and $\vec{v}[m] \notin (\pi_{-\infty})^{-1}[a]$. It is then clear that $\vec{w} \in (R_{-\infty})^m[(\pi_{-\infty})^{-1}[a]]$ but $\vec{v} \notin (R_{-\infty})^m[(\pi_{-\infty})^{-1}[a]] \in A_{-\infty}$, showing differentiatedness. $\square$

[1] Strictly speaking, they could be added to the recursion, but the set would simply be empty.

Proposition 5.8. Let $\mathfrak{g}$ be a differentiated general frame. Then $\mathfrak{g}_{-\infty}$ is tight.

Proof. Let $\vec{w} = (w_n)_{n\in\mathbb{N}}$ not extend $\vec{v} = (v_m)_{m\in\mathbb{N}}$. Then there must be a $k \in \mathbb{N}$ such that $w_{k+1} \neq v_k$. From differentiatedness of $\mathfrak{g}$, the frame $\mathfrak{g}_{-\infty}$ is differentiated, so there is an admissible a in $\mathfrak{g}_{-\infty}$ such that $\vec{w}[k+1] \in a$ but $\vec{v}[k] \notin a$. It follows that $\vec{w} \in (R_{-\infty})^{k+1}[a]$ but $\vec{v} \notin (R_{-\infty})^k[a] \supseteq \langle R_{-\infty}\rangle(R_{-\infty})^{k+1}[a]$. □

In contrast to the unravelling cover, however, $\mathfrak{g}_{-\infty}$ even inherits compactness. The proof is conveniently short when using nets, as there is a useful characterisation of compactness in terms of ultranets. Recall from Definition 2.77 that an ultranet, or universal net, on a topological space $(X, \mathcal{T})$ is a net that for every subset $Y \subseteq X$ is eventually in Y or in Y^c. An important property of ultranets was formulated in Proposition 2.81, stating that a space is compact if and only if all its ultranets converge, if and only if all its semi-universal nets converge. This chapter will accomplish compactness results with both ultranets or semi-universal nets, depending on which is convenient.

The proof is included in Section 2.6.2 and will not be repeated here. The reason why universal nets will be so convenient, is that any shifted universal net remains universal.

Lemma 5.9. Let $(\vec{x}_d)_{d\in D}$ be a (semi-)universal net in $\mathfrak{g}_{-\infty}$. Then the k-shifted net $(\vec{x}_d[k])_{d\in D}$ is a (semi-)universal net for every $k \in \mathbb{N}$, where $\vec{x}_d[k]$ is the k-shift of the sequence $x_d \in W_{-\infty}$.

Proof. The universal case will be shown by induction on k. The semi-universal case is identical. Suppose $(\vec{x}_d[k])_{d\in D}$ is universal and let $Y \subseteq W_{-\infty}$. Then, by the induction hypothesis, $(\vec{x}_d[k])_{d\in D}$ is eventually in $R_{-\infty}[Y]$ or in $(R_{-\infty}[Y])^c = R_{-\infty}[Y^c]$, where the equality was noted earlier to follow from disjointness of R-image sets. But note that every point $\vec{w}$ in $\mathfrak{g}_{-\infty}$ has a unique predecessor $\vec{w}[1]$, so for any set $X \subseteq W_{-\infty}$ the net $(\vec{x}_d[k])_{d\in D}$ is eventually in $R_{-\infty}[X]$ if and only if $(\vec{x}_d[k+1])_{d\in D}$ is eventually in X, so that $(\vec{x}_d[k+1])_{d\in D}$ is eventually in Y or eventually in Y^c. Because $Y \subseteq W_{-\infty}$ was arbitrary, it follows that $(\vec{x}_d[k+1])_{d\in D}$ is universal. □

From here, the compactness of $\mathfrak{g}_{-\infty}$ of a compact general frame $\mathfrak{g}$ is straightforwardly proven.

Theorem 5.10. *Let $\mathfrak{g}$ be a descriptive frame. Then $\mathfrak{g}_{-\infty}$ is descriptive.*

Proof. For differentiatedness and tightness see Propositions 5.7 and 5.8. All that needs to be proven is compactness.

For the characterisation of compactness from Proposition 2.81, let $(\vec{x}_d)_{d\in D}$ be any ultranet in $\mathfrak{g}_{-\infty}$. To show that it converges, an accumulation point first needs to be found. To do this, every $(\vec{x}_d[n])_{d\in D}$ can be projected under $\pi_{-\infty}$ to $\mathfrak{g}$, where it is still universal. There, compactness guarantees an accumulation point. The sequence of these accumulation points will turn out to be a point in $W_{-\infty}$, which will be an accumulation point to which $(\vec{x}_d)_{d\in D}$ converges. This convergence can be shown through a careful induction on the elements of $A_{-\infty}$.

Then by Lemma 5.9, for any n the net $(\vec{x}_d[n])_{d\in D}$ is universal. As a consequence, for each $n \in \mathbb{N}$ the net $(y^n_d)_{d\in D}$ in $\mathfrak{g}$ defined by $y^n_d = \pi_{-\infty}(x_d[n])$ is universal. After all, let X be a subset of the universe of $\mathfrak{g}$. Then by universality of $(\vec{x}_d[n])_{d\in D}$ it is either eventually in $(\pi_{-\infty})^{-1}[X]$ or eventually in $(\pi_{-\infty})^{-1}[X]^c = (\pi_{-\infty})^{-1}[X^c]$. An immediate consequence is that $y_d^n = \pi_{-\infty}(x_d[n])$ is eventually in X or is eventually in X^c.

The compactness of $\mathfrak{g}$ then implies through Proposition 2.81 that y_d^n is convergent. Let w_n then be the limit point of $(y_d^n)_{d\in D}$ for every n. Note that y_d^n is an R-successor of y_d^{n+1} in $\mathfrak{g}$. After all, the former is the head of $\vec{x}_d[n]$ and the latter of $\vec{x}_d[n+1]$. Since $\vec{x}_d$ is a descending R-sequence, its n-th element, which will be the head of its n-shift $\vec{x}_d[n]$, is a successor of its $(n+1)$-th element, the head of its $n+1$-shift $\vec{x}_d[n+1]$.

As such, for each n and d, the element y_d^n is a successor of y_d^{n+1}. As R is fully continuous from Proposition 2.73 and the assumed tightness, it follows that w_n is an R-successor of w_{n+1}. As a consequence, the sequence $\vec{w} := (w_n)_{n\in\mathbb{N}}$ is an infinitely R-descending sequence and hence a point in $W_{-\infty}$.

An careful induction on the construction of $A_{-\infty}$ will now show for all n simultaneously that $(\vec{x}_d[n])_{d\in D}$ converges to $\vec{w}[n]$. Specifically, it will be shown by induction that for a given $b \in A_{-\infty}$, that for all $n \in \mathbb{N}$, if $\vec{w}[n] \in b$, then $(\vec{x}_d[n])_{d\in D}$ is eventually in b.

- For a $\mathfrak{g}$-admissible set a, suppose that $\vec{w}[n] \in (\pi_{-\infty})^{-1}[a]$ for some n. Then $w_n \in a$ and thus $(y_d^n)_{d\in D}$ is eventually in a, as it converges to w_n by assumption. Then as $y_d^n = \pi_{-\infty}(\vec{x}_d[n])$, it follows that $(\vec{x}_d[n])_{d\in D}$ is eventually in $(\pi_{-\infty})^{-1}[a]$.
- Now suppose for the induction hypothesis that $b \in A_{-\infty}$ has the property that for all $n \in \mathbb{N}$ if $\vec{w}[n] \in b$, then $(\vec{x}_d[n])_{d\in D}$ is eventually b. Let now m be arbitrary and suppose that $\vec{w}[m] \in R_{-\infty}[b]$. By disjointness of image-sets, it follows that $\vec{w}[m+1] \in b$. The induction hypothesis then grants that $(\vec{x}_d[m+1])_{d\in D}$ is eventually in b, so that $(\vec{x}_d[m])_{d\in D} \in R_{-\infty}[(\vec{x}_d[m+1])_{d\in D}] \subseteq R_{-\infty}[b]$.
- The cases $b \cup b'$ and $b \cap b'$ are easily checked by noting that a net is eventually in $b \cup b'$ if and only if it is eventually in b or eventually in b' and a net is eventually in $b \cap b'$ if and only if it is both eventually in b and eventually in b'.

Thus the arbitrary universal net $(\vec{x}_d)_{d\in D}$ converges, from which it follows that $\mathfrak{g}_{-\infty}$ is compact. □

5.2 The $\mathfrak{g}_{-\infty}$-frame as limit points

The infinite paths in $W_{-\infty}$ should remind the reader of the paths in the unravelling tree. Noting how closely the structures resemble one another, it is perhaps unsurprising that they embody the structure of the limit points.

To associate these with one another, it will be very useful to have a concrete grasp of the limit points as a general frame. As a byproduct of the completion from the double application of Jónsson-Tarski duality, it is hard to understand and work with. In its place, the theory developed in Section 2.6.2 will allow the limit points to be viewed as nets, which will turn out to be a sufficiently concrete image to understand it as $\mathfrak{g}_{-\infty}$-frame.

First, it will prove useful to revisit Theorem 4.26. It stated that for a descriptive frame $\mathfrak{g}$, the descriptive unravelling $\widehat{\mathfrak{g}}$ as a Kripke frame was isomorphic to the disjoint union of $\mathfrak{T}(\mathfrak{g})$ and a frame of limit points $\mathfrak{L}$. Together with the continuous embedding $\widehat{\iota} : \mathfrak{T}(\mathfrak{g}) \rightarrowtail \widehat{\mathfrak{g}}$, this makes it possible to characterise the nets that become points in $\mathfrak{L}$ in the $\mathfrak{g}_\bullet$-construction.

Lemma 5.11. Let $\mathfrak{g} = (W, R, A)$ be a descriptive frame and let ν be a universal net on $\mathfrak{T}(\mathfrak{g})$. Then either ν converges to a point $\vec{x}$ in $\mathfrak{T}(\mathfrak{g})$ or it is eventually in $(R^{\mathfrak{T}})^n[W^{\mathfrak{T}}]$ for all $n \in \mathbb{N}$.

Proof. From universality, it follows that for each n either ν is eventually in $(R^{\mathfrak{T}})^n[W^{\mathfrak{T}}]$, or in its complement, $\bigcup_{0\leq k<n}(R^{\mathfrak{T}})^k[\mathcal{I}]$. Suppose that there is an n such that ν is eventually in the latter. This set is compact. After all, Proposition 4.24 states that $\widehat{\iota}$ is an embedding and a bounded morphism, so that as topological spaces

$$\bigcup_{0\leq k<n}(R^{\mathfrak{T}})^k[\mathcal{I}] \cong \widehat{\iota}\left[\bigcup_{0\leq k<n}(R^{\mathfrak{T}})^k[\mathcal{I}]\right] = \bigcup_{0\leq k<n}(R_*)^k[\widehat{\iota}[\mathcal{I}]].$$

Since $\mathcal{I}$ with the subspace topology is homeomorphic to the $\mathfrak{g}$-space by construction, it is compact. By the fact that $\widehat{\iota}$ is an embedding, then, $\widehat{\mathcal{I}} := \widehat{\iota}[\mathcal{I}]$ is compact. As the successor set of a compact set is compact by Lemma 2.34, each $(R_*)^k[\widehat{\mathcal{I}}]$ is compact, so that their finite union is as well.

In a compact space, every ultranet converges by Proposition 2.81, meaning that $(x_d)_{d\in D}$ must converge. As no point of $\mathfrak{T}(\mathfrak{g})$ is in $(R^{\mathfrak{T}})^n[W^{\mathfrak{T}}]$ for all n, this proves the lemma. $\square$

As an application, this also provides an easy topological proof of Theorem 4.26, stating that $\widehat{\mathfrak{g}}_\# \cong (\mathfrak{T}(\mathfrak{g})_\bullet)_\# \cong \mathfrak{T}(\mathfrak{g})_\# \uplus \mathfrak{L}$.

Corollary 5.12. Let $\mathfrak{g}$ be a descriptive frame. Then $\widehat{\mathfrak{g}}_\# = \mathfrak{T}(\mathfrak{g})_\# \uplus \mathfrak{L}$ for some unspecified frame of limit points $\mathfrak{L}$, where $\#$ again takes the underlying Kripke frame.

Proof. It is easy to check that in the topological completion $\mathfrak{T}(\mathfrak{g})_\bullet$, the constant nets reproduce the structure of $\mathfrak{T}(\mathfrak{g})$. Now suppose that $(\vec{x}_d)_{d\in D}$ is a universal net in $\mathfrak{T}(\mathfrak{g})$ and let ν be a constant net of a point $\vec{w}$. It is sufficient to show that if ν is $R_\bullet$-related to $(\vec{x}_d)_{d\in D}$, then $(\vec{x}_d)_{d\in D}$ is equivalent to a constant net on $\mathfrak{T}(\mathfrak{g})$.

For $(x_d)_{d\in D}$ to be $R_\bullet$-related to $\vec{w}$, it must frequently be in $R^{\mathfrak{T}}[\vec{w}]$ or $\langle R^{\mathfrak{T}}\rangle\{\vec{w}\}$. But, if n is the length of the path $\vec{w}$, this means that $(x_d)_{d\in D}$ is not eventually in $(R^{\mathfrak{T}})^{n+2}[W^{\mathfrak{T}}]$, from which it follows that $(x_d)_{d\in D}$ converges by Lemma 5.11. As it converges, there must be a point $\vec{v}$ in $\mathfrak{T}(\mathfrak{g})$ such that for every open set a on $\mathfrak{T}(\mathfrak{g})$ containing $\vec{v}$, it lands eventually in a. As such, it must be equivalent to the constant net at $\vec{v}$. $\square$

This characterisation of the non-convergent nets then allows an exact and explicit construction of the descriptive unravelling.

Definition 5.13. Let $\mathfrak{F} = (W, R)$ be a Kripke frame. Then the *infinitary unravelling forest* $\mathfrak{T}_{-\infty} = (W^{\mathfrak{T}_{-\infty}}, R^{\mathfrak{T}_{-\infty}})$ is given by

$$\mathfrak{T}_{-\infty}(\mathfrak{F}) := \mathfrak{T}(\mathfrak{F}) \uplus \mathfrak{F}_{-\infty}$$

is the frame on all R-paths, finite or infinite, with a final point. ◀

The bounded morphisms $\pi : \mathfrak{T}(\mathfrak{F}) \to \mathfrak{F}$ and $\pi_{-\infty} : \mathfrak{F}_{-\infty} \to \mathfrak{F}$ now give rise to a bounded morphism $\pi_{\mathfrak{T}_{-\infty}} : \mathfrak{T}_{-\infty}(\mathfrak{F}) \to \mathfrak{F}$ by applying π to points in $\mathfrak{T}(\mathfrak{F})$ and $\pi_{-\infty}$ to points in $\mathfrak{F}_{-\infty}$. It is immediate that this is a bounded morphism, since no points from the two frames are related. Just as these two maps gave the unravelling cover and the $\mathfrak{g}_{-\infty}$-frame their admissible sets, this map can also be used to give $\mathfrak{T}_{-\infty}(\mathfrak{F})$ the structure of a general frame.

Definition 5.14. Let $\mathfrak{g} = (\mathfrak{F}, A)$ be a general frame based on a Kripke frame $\mathfrak{F} = (W, R)$. Then the *infinitary unravelling cover* $\mathfrak{T}_{-\infty}(\mathfrak{g}) = (\mathfrak{T}_{-\infty}(\mathfrak{F}), A^{\mathfrak{T}_{-\infty}})$ is the general frame with admissible sets $A^{\mathfrak{T}_{-\infty}}$ recursively defined using $\pi_{\mathfrak{T}_{-\infty}} : \mathfrak{T}_{-\infty}(\mathfrak{F}) \to \mathfrak{F}$ through

$$\begin{array}{lcl} & & \mathcal{I} \in A^{\mathfrak{T}_{-\infty}}; \\ a \in A & \Longrightarrow & (\pi_{\mathfrak{T}_{-\infty}})^{-1}[a] \in A^{\mathfrak{T}_{-\infty}}; \\ b \in A^{\mathfrak{T}_{-\infty}} & \Longrightarrow & R^{\mathfrak{T}_{-\infty}}[b] \in A^{\mathfrak{T}_{-\infty}}; \\ b, b' \in A^{\mathfrak{T}_{-\infty}} & \Longrightarrow & b \cup b', b \cap b' \in A^{\mathfrak{T}_{-\infty}}. \end{array}$$

◀

Proposition 5.15. If $\mathfrak{g}$ is a general frame, then $\mathfrak{T}_{-\infty}(\mathfrak{g})$ is a general frame.

Proof. Once again, the recursion has the exact same structure as the recursion for $A^{\mathfrak{T}}$. Since still every point has at most one predecessor, the proof for Proposition 4.11 may be repeated word for word, replacing only every instance of $\mathfrak{T}$ with $\mathfrak{T}_{-\infty}$. □

The structure of the infinitary unravelling cover is expressly designed to extend and unite the unravelling cover and the $\mathfrak{g}_{-\infty}$-structure.

Lemma 5.16. Let $\mathfrak{g} = (\mathfrak{F}, A)$ be a descriptive frame. Then the maps $\iota_{-\infty} : \mathfrak{F}_{-\infty} \rightarrowtail \mathfrak{T}(\mathfrak{F}) \uplus \mathfrak{F}_{-\infty}$ and $\iota_{\mathfrak{T}} : \mathfrak{T}(\mathfrak{F}) \rightarrowtail \mathfrak{T}(\mathfrak{F}) \uplus \mathfrak{F}_{-\infty}$, when considered as maps between general frames $\iota_{-\infty} : \mathfrak{g}_{-\infty} \to \mathfrak{T}_{-\infty}(\mathfrak{g})$ and $\iota_{\mathfrak{T}} : \mathfrak{T}(\mathfrak{g}) \to \mathfrak{T}_{-\infty}(\mathfrak{g})$ are embeddings in the sense that they are homeomorphisms onto the image.

Proof. Injectivity is immediate. To see that they are homeomorphisms, it is sufficient to see that the admissible sets restricted to $W^{\mathfrak{T}}$ and $W_{-\infty}$ are exactly the collections $A^{\mathfrak{T}}$ and $A_{-\infty}$ respectively. This can be done by an induction on the recursive construction of $\mathfrak{T}_{-\infty}$. The proof is easy, but notationally involved.

For the initial paths, note that $\mathcal{I} \cap W^{\mathfrak{T}} = \mathcal{I}$, and $\mathcal{I} \cap W_{-\infty} = \emptyset$. For the inverse images, $(\pi_{\mathfrak{T}_{-\infty}})^{-1}[a] \cap W^{\mathfrak{T}} = \pi^{-1}[a]$ and $(\pi_{\mathfrak{T}_{-\infty}})^{-1}[a] \cap W_{-\infty} = (\pi_{-\infty})^{-1}[a]$, as π and $\pi_{-\infty}$ are exactly the restrictions of $\pi_{\mathfrak{T}_{-\infty}}$ to $W^{\mathfrak{T}}$ and $W_{-\infty}$ respectively.

Now suppose for the induction hypothesis that for $b, b' \in A^{\mathfrak{T}_{-\infty}}$ the intersections $b \cap W^{\mathfrak{T}}, b' \cap W^{\mathfrak{T}}$ are in $A^{\mathfrak{T}}$ and that $b \cap W_{-\infty}$ and $b' \cap W_{-\infty}$ are in $A_{-\infty}$. Then note that $R^{\mathfrak{T}_{-\infty}}[b] \cap W^{\mathfrak{T}} = R^{\mathfrak{T}_{-\infty}}[b \cap W^{\mathfrak{T}}] = R^{\mathfrak{T}}[b \cap W^{\mathfrak{T}}]$ and $R^{\mathfrak{T}_{-\infty}}[b] \cap W_{-\infty} = R_{-\infty}[b \cap W_{-\infty}]$, because $R^{\mathfrak{T}_{-\infty}}$ does not connect points from the two sets and restricts to their relations.

Moreover, intersecting with $W^{\mathfrak{T}}$ and $W_{-\infty}$ distribute over union and intersection, completing the induction. □

With the concrete structure of the infinitary unravelling cover in place, it can be used to understand the descriptive unravelling.

Theorem 5.17. *Let $\mathfrak{g} = (W, R, A)$ be a descriptive frame. Then the topological descriptive completion is isomorphic to the infinitary unravelling, $\mathfrak{T}(\mathfrak{g})_\bullet \cong \mathfrak{T}_{-\infty}(\mathfrak{g})$ through the map $\Theta : W^{\mathfrak{T}_{-\infty}} \to W_\bullet$ that sends each path to the net of finite paths leading up to it. Specifically,*

$$\Theta : W^{\mathfrak{T}_{-\infty}} \to W_\bullet : \begin{cases} \vec{w} & \mapsto [(\vec{w})]_A \text{ constant net if } \vec{w} \in W^{\mathfrak{T}}, \\ \vec{w} = (w_n)_{n\in\mathbb{N}} & \mapsto [((w_m, \dots, w_0))_{m\in\mathbb{N}}]_A \end{cases}$$

maps finite paths in $W^{\mathfrak{T}}$ to the equivalence class of a constant net at that point and infinite paths to the equivalence class of the net of all finite paths with which the infinite path ends.

The proof of this theorem is extremely involved, not because it is conceptually difficult, but because the isomorphism provided has to be checked on all counts: well-definedness, injectivity, surjectivity, relation-preservation in both directions, and inducing an isomorphism between the collections admissible sets. For the sake of completeness, the proof has been included, but the theorem is much more important and insightful than the proof.

Proof of Theorem 5.17. For $\vec{w} \in W^{\mathfrak{T}}$, this is well-defined, because constant nets are certainly semi-universal. By induction on the construction of $A^{\mathfrak{T}}$ over all simultaneous $\vec{w}$, it can be seen that $\Theta(\vec{w}) = (\vec{x}_m)_{m\in\mathbb{N}}$ is also well-defined for $\vec{w} \in W_{-\infty}$. This is easy to see for the induction basis, as

$$\begin{aligned} \vec{x}_m = (w_m, w_{m-1}, \dots, w_1, w_0) \in \pi^{-1}[a] \iff w_0 \in a \iff w_0 \notin a^c \\ \iff \vec{x}_m = (w_m, w_{m-1}, \dots, w_1, w_0) \notin \pi^{-1}[a^c] = \pi^{-1}[a]^c, \end{aligned}$$

so that $(x_m)_{m\in\mathbb{N}}$ is either constantly (and thus eventually) in $\pi^{-1}[a]$ or constantly in $\pi^{-1}[a]^c$. The case for union and intersection follows immediately from the De Morgan laws.

For the case $R^{\mathfrak{T}}[b]$ with $b \in A^T$, define $(\vec{y}_k)_{k\in\mathbb{N}} = \Theta(\vec{w}[1])$. By the induction hypothesis $(\vec{y}_k)_{k\in\mathbb{N}}$ is eventually in b or eventually in b^c, meaning that $\vec{y}_k \in b$ or $\vec{y}_k \in b^c$ for sufficiently large $k \geq k_0$. But since for $m \geq 1$, $\vec{x}_m = (w_m, w_{m-1}, \dots, w_1, w_0)$ and $\vec{y}_{m-1} = (w_m, \dots, w_1)$, this means that $\vec{x}_m$ is a successor to $\vec{y}_{m-1}$, so that $\vec{x}_m$ is in $R^{\mathfrak{T}}[b]$ for sufficiently large $m \geq k_0 + 1$. The same reasoning applies if $\vec{y}_k$ is eventually in b^c so that then x_m eventually in $R^{\mathfrak{T}}[b^c] \subseteq R^{\mathfrak{T}}[b^c] \cup \mathcal{I} = (R^{\mathfrak{T}}[b])^c$.

To see that Θ is injective, let $\vec{w} = (w_n)_{n\in\mathbb{N}}, \vec{v} = (v_n)_{n\in\mathbb{N}} \in W_{-\infty}$ be distinct, infinitely R-descending paths. Then there must be some natural number $k \in \mathbb{N}$ such that $w_k \neq v_k$. By differentiatedness of $\mathfrak{g}$, there is an $a \in A$ such that $w_k \in a$ but $v_k \notin a$. Writing $\Theta(\vec{w}) = (\vec{x}_m)_{m\in\mathbb{N}}$ and $\Theta(\vec{v}) = (\vec{y}_m)_{m\in\mathbb{N}}$, note that for $m \geq k$ the finite paths $\vec{x}_m = (w_m, \dots, w_k, \dots, w_0)$ must in $(R^{\mathfrak{T}})^k[w_k] \subseteq (R^{\mathfrak{T}})^k[\pi^{-1}[a]]$, while $\vec{y}_m = (v_m, \dots, v_k, \dots, w_0)$

must be in $(R^{\mathfrak{T}})^k[v_k] \subseteq (R^{\mathfrak{T}})^k[\pi^{-1}[a^c]] \subseteq (R^{\mathfrak{T}})^k[\pi^{-1}[a]^c] \cup \mathcal{I} = ((R^{\mathfrak{T}})^k[\pi^{-1}[a]])^c$, so that they are not semi-equivalent.

For surjectivity, the trick from Theorem 5.18 can be applied to use the projection $\pi : \mathfrak{T}(\mathfrak{g}) \to \mathfrak{g}$ to construct from a semi-universal net on $\mathfrak{T}(\mathfrak{g})$ a point in $W^{\mathfrak{T}_{-\infty}}$. So let $(\vec{x}_d)_{d\in D}$ be a semi-universal net on $\mathfrak{T}(\mathfrak{g})$. Recall from Lemma 5.11 that either $(\vec{x}_d)_{d\in D}$ converges or it is eventually in $(R^{\mathfrak{T}})^n[W^{\mathfrak{T}}]$ for all n. If it converges to a point $\vec{x}$, then it is semi-equivalent to the constant path at $\vec{x}$ (see also Corollary 2.84), so that its equivalence class is reached by $\Theta(\vec{x})$.

So assume now it is eventually in $(R^{\mathfrak{T}})^n[W^{\mathfrak{T}}]$ for all $n \in \mathbb{N}$. Fix an n and a $d_n \in D$ such that $d \succeq d_n \implies \vec{x}_d \in (R^{\mathfrak{T}})^n[W^{\mathfrak{T}}]$. In this case, for each such d, if $\vec{x}_d = (w_r^d, \ldots, w_n^d, \ldots, w_0^d)$, it can be shifted by n to become $\vec{x}_d[n] = (w_r^d, \ldots, w_n^d)$. Exactly the arguments given in Lemma 5.9 to show that the shift of an semi-universal net on $\mathfrak{g}_{-\infty}$ was a semi-universal net can be applied to show that $(\vec{x}_d[n])_{d\succeq d_n}$ is a semi-universal net on $\mathfrak{T}(\mathfrak{g})$. Projecting then to $\mathfrak{g}$ by applying π to get a net $(w_n^d)_{d\in D}$ for each n gives again a semi-universal net on $\mathfrak{g}$, because the π-inverse image of admissible sets on $\mathfrak{g}$ is admissible on $\mathfrak{T}(\mathfrak{g})$. By Proposition 2.81 and compactness of $\mathfrak{g}$, each $(w_n^d)_{d\in D}$ must have a limit point w_n.

If $\vec{w} = (w_n)_{n\in\mathbb{N}}$, then $(\vec{x}_d[k])_{d\succeq d_k}$ is in the equivalence class $\Theta(\vec{w}[k])$, which can again be seen from induction on the construction of $A^{\mathfrak{T}}$ for all simultaneous k. If $(\vec{x}_d[k])_{d\in D}$ is eventually in $\pi^{-1}[a]$, then $(w_k^d)_{d\in D}$ is eventually in a, so that Proposition 2.83 implies that its convergence to w_k grants $w_k \in a$. From this, it follows that $(w_m, \ldots, w_k) \in \pi^{-1}[a]$ for all $m \geq k$.

Now since the union and intersection are again trivial, the only interesting case to consider is $R^{\mathfrak{T}}[b]$ for an admissible $b \in A^{\mathfrak{T}}$ in which either both $(\vec{x}_d[k])_{d\succeq d_k}$ and the net of finite paths $((w_m, \ldots, w_k))_{m\in\mathbb{N}}$ are eventually or are both eventually not, for all k. If $(\vec{x}_d[k])_{d\succeq d_k}$ is eventually in $R^{\mathfrak{T}}[b]$, then $(\vec{x}_d[k+1])_{d\succeq d_{k+1}}$ is eventually in b, so that the induction hypothesis grants that $((w_m, \ldots, w_{k+1}))_{m\in\mathbb{N}}$ is eventually in b and thus $((w_m, \ldots, w_k))_{m\in\mathbb{N}}$ is eventually in $R^{\mathfrak{T}}[b]$. Therefore, $(\vec{x}_d[k])_{d\succeq d_k}$ is in the equivalence class $\Theta(\vec{w})$.

The relation is obviously preserved in both directions on $\mathfrak{T}(\mathfrak{g})$, as it is the identity. By injectivity, it follows that $\mathfrak{g}_{-\infty}$ is mapped to $\mathfrak{L}$. Since there are no $R^{\mathfrak{T}_{-\infty}}$-relations between $\mathfrak{T}(\mathfrak{g})$ and $\mathfrak{g}_{-\infty}$ by construction, and no $R_\bullet$-relations between the constant nets and the nets representing the limit points, all that needs to be checked is that $\vec{w}R_{-\infty}\vec{v}$ if and only if $\Theta(\vec{w})R_\bullet\Theta(\vec{v})$ for $\vec{w}, \vec{v} \in W_{-\infty}$. The infinite path $\vec{v}$ is a $R_{-\infty}$-successor of $\vec{w}$ if and only if v_0 is an R-successor of w_0. If it is, then $(w_m, \ldots, w_0)$ and $(v_{m+1}, \ldots, v_0)$ are $R^{\mathfrak{T}}$-related for all m, so $\Theta(\vec{v})$ is an $R_\bullet$-successor of $\Theta(\vec{w})$. If v_0 is not an R-successor of w_0, then by tightness of $\mathfrak{g}$, there exists an $a \ni v_0$ such that $w_0 \notin \langle R\rangle a$. As such, all nets in $\Theta(\vec{v})$ must be eventually in $\pi^{-1}[a]$ and all nets in $\Theta(\vec{w})$ must eventually be in $\pi^{-1}[(\langle R\rangle a)^c] = (\langle R^{\mathfrak{T}}\rangle\pi^{-1}[a])^c$, guaranteeing that $\Theta(\vec{v})$ is not a $R_\bullet$-successor of $\Theta(\vec{w})$.

Finally, it must be checked that Θ maps admissible sets in $A^{\mathfrak{T}_{-\infty}}$ to admissible sets in $A_\bullet$. Under Θ a set $b \in A^{\mathfrak{T}-\infty}$ is mapped to $\widetilde{b \cap W^{\mathfrak{T}}}$, where $\widetilde{a}$ is the set of equivalence classes of semi-universal nets eventually in a. For points $\vec{w} \in W^{\mathfrak{T}}$, it is trivial that $\vec{w} \in \widetilde{b \cap W^{\mathfrak{T}}}$. For points $\vec{w} \in W^{\mathfrak{T}-\infty}$, it can be seen through an inductive argument

$$\vec{w} \in (\pi_{-\infty})^{-1}[a] \iff w_0 \in a \iff \forall m \ (w_m, \ldots, w_0) \in \pi^{-1}[a] \iff \Theta(\vec{w}) \in \widetilde{\pi^{-1}[a]};$$

$$\begin{aligned}\vec{w} \in R^{\mathfrak{T}-\infty}[b] \iff \vec{w}[1] \in b \overset{\text{IH}}{\iff} \Theta(\vec{w}[1]) \in \widetilde{b \cap W^{\mathfrak{T}}} \iff \Theta(\vec{w}) &\in R^{\mathfrak{T}}[\widetilde{b \cap W^{\mathfrak{T}}}] \\ &= \widetilde{R^{\mathfrak{T}-\infty}[b] \cap W^{\mathfrak{T}}},\end{aligned}$$

where use of the induction hypothesis is marked by IH and the fact that Θ preserves the relation is used. The induction step for union and intersection is trivial. Lemma 5.16 then gives that this is surjective on $A_\bullet$ and a map between fields of sets induced by a bijection is always injective. $\square$

5.3 The structure of the $\mathfrak{g}_{-\infty}$-frame

Having established that the $\mathfrak{g}_{-\infty}$-construction provides the limit points of the descriptive unravelling, the next step is to understand what it looks like. It turns out to be remarkably well-behaved.

Theorem 5.18. *Let $\mathfrak{F}$ be a Kripke frame. Then each $R_{-\infty}$-connected component in $\mathfrak{F}_{-\infty}$ is either a non-well-founded tree or an n-cycle for $n \geq 1$ with possibly a tree attached at each point along the cycle. Examples for these two options are illustrated in Figure 5.2.*

Proof of Theorem 5.18. This theorem is a consequence of the fact that each point in $W_{-\infty}$ has exactly one predecessor. If $\vec{w} \in W_{-\infty}$, then its only $R_{-\infty}$-predecessor is $\vec{w}[1]$. This also means immediately that the sequence $(\vec{w}[k])_{k \in \mathbb{N}}$ is an infinite $R_{-\infty}$-descending path, meaning that no connected component may be well-founded.

Now from the fact that each point has a unique $R_{-\infty}$-predecessor, it follows that each point in $W_{-\infty}$ can be part of at most one cycle. After all, any cycle must include its unique predecessor, and its unique predecessor, et cetera. When this process finally returns to the original point, the cycle is complete.

If $\vec{w}$ is in a cycle and $\vec{v}$ is an $R_{-\infty}$-successor of $\vec{w}$ but not in the cycle, then $\vec{v}$ is in no cycle. After all, suppose it is. Then surely its unique predecessor must be in this cycle. But that is $\vec{w}$ and by the previous claim $\vec{w}$ cannot be in two cycles.

Finally, if $\vec{w}$ is not in a cycle, then its generated subframe $\bigcup_{n \in \mathbb{N}} (R_{-\infty})^n[w]$ is a tree. After all, by the uniqueness of $R_{-\infty}$-predecessors, it can never happen that two distinct points share a successor. Moreover, no cycle can happen in the successors of $\vec{w}$, because by the uniqueness of predecessors, $\vec{w}$ must be a part of that cycle, in contradiction with the assumption.

Combining all of the above, it follows that any point must either be reachable through an $R_{-\infty}$-path from a cycle or be part of a non-well-founded tree. In the former case, each point in this cycle must have a successor set containing the next point in the cycle and furthermore some number of points (possibly none) that are the root of a tree that is their generated subframe. This proves the theorem. $\square$

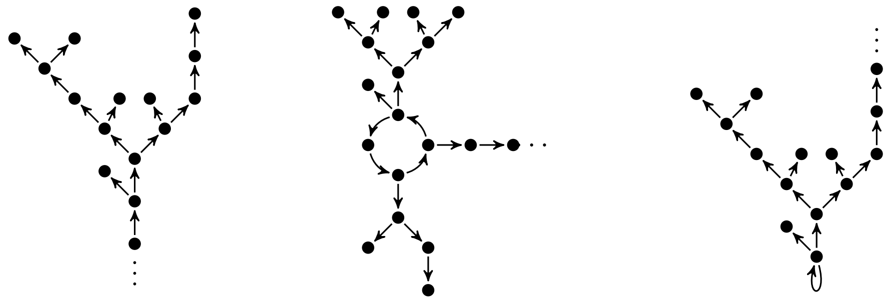

(a) A non-well-founded tree as an example of what an $R_{-\infty}$-connected component of a frame $\mathfrak{F}_{-\infty}$ might look like. Note specifically that the dots are meant to indicate that the downwards chain under the tree goes on forever. The upwards chains may or may not go on forever.

(b) A 4-cycle with at three of the four points a tree-like structure attached, as an example of what an $R_{-\infty}$-connected component might look like. Some of the points have no tree attached, which is also acceptable. Of the trees attached at the cycle, two are finite, and one is an infinite chain.

(c) A final example of what an $R_{-\infty}$-connected component might look like. Note that at the root of this tree is a reflexive point, which can be considered a 1-cycle, which is also allowed. Observe also that some of the branches of the attached tree are finite and some are not.

Figure 5.2: Three examples of the possible structures of $R_{-\infty}$-connected components in $\mathfrak{F}_{-\infty}$. All of them are non-well-founded, either due to an infinite descending chain or because of a cycle.

5.4 A review of the descriptive unravelling

A vital tool and key contribution of this thesis has been the descriptive unravelling construction. Unravelling constructions have proven extremely useful to obtain results about bisimulation-invariant phenomena.

In particular, many recent incarnations of the van Benthem Characterisation Theorem [14, 34, 33], as well as other modal characterisation results like the Janin-Walukiewicz Theorem [23] depend on some unravelling construction. Any future modal characterisation theorems over the class of descriptive models would likely have use for an unravelling construction in this class.

In this thesis, three different constructions of the descriptive unravelling are given. One is constructed algebraically through Jónsson-Tarski duality in Definition 2.50 as the double dual of an appropriate general frame based on the classical unravelling.

Secondly, in Definition 2.86, it is constructed topologically by adding semi-universal nets wherever needed to compactify the space, which turn out to be all semi-universal nets eventually escaping to arbitrarily long paths on the classical unravelling. Finally, the infinitary unravelling construction, given in Definition 5.13, gives an explicit construction of the descriptive unravelling as consisting of all paths on the original frame that are finite in the future and either finite or infinite in the past.

Note that all of these constructions have unique advantages that make them a useful way of looking at the descriptive unravelling. The descriptive unravelling through algebraic methods uses exclusively functorial constructions, because the unravelling tree, even when turned into a general frame, is functorial and the duality used is a composition of two contrapositive functors. This makes it very suitable for coalgebraic analyses.

The infinitary unravelling also has a somewhat coalgebraic flavour, in that it resembles other closure constructions in coalgebraic contexts. Moreover, as seen in the previous section, it lends itself to much more convenient structural analysis. Pinning down the exact structure of the tree is particularly useful for game-theoretic contexts, which means that Ehrenfeucht-Fraïssé-type arguments become more applicable, and languages with game-theoretic semantics, like automata-languages, are more easily analysed on this construction.

Finally, the construction through nets is designed explicitly to make topological analyses more convenient. In particular, frames of which the topology is well-understood will allow for an effective analysis of the nets on the frame, which in turn facilitates the analysis of the descriptive completion.

As Theorems 2.88 and 5.17 show these three structures are isomorphic, anyone wishing to analyse the descriptive unravelling can switch back and forth between these constructions depending on whichever one is more convenient, possibly linking properties derived from one analysis to properties derived in the other. These three perspectives on the unravelling are a key contribution of this thesis.

6 Conclusions and future work

6.1 Conclusion

This thesis investigates the modal and first-order model theory of descriptive models. The main theorem was the van Benthem Characterisation Theorem for descriptive models (Theorem 4.7). It states that any first-order formula in one free variable that is invariant under (Vietoris) bisimulations, is in fact equivalent to a modal formula on the class of descriptive models. Chapter 4 is dedicated entirely to a proof of this theorem.

Chapter 3 investigates the model theory of descriptive models and demonstrates the failure of a number of classical model-theoretic results, including the Compactness Theorem for first-order logic and the upward Löwenheim-Skolem Theorem, on the class of descriptive models. This necessitated the use of methods from finite model theory to achieve the van Benthem Characterisation Theorem, especially the Ehrenfeucht-Fraïssé method.

The main tool developed in order to permit these arguments is the "descriptive unravelling" of descriptive frame $\mathfrak{g}$, written as $\widehat{\mathfrak{g}}$, from Definition 4.9, adapting the commonly used tool of the unravelling tree to the context of descriptive models. Using Jónsson-Tarski duality, the regular unravelling can be compactified into a descriptive frame to obtain the descriptive unravelling.

Chapter 2 provides a topological toolkit of nets fine-tuned especially for topological spaces associated with general frames, to understand such compactifications better. As the precise structure of the descriptive unravelling is key to the proof of the main theorem, much attention has been given to applying these topological tools to the descriptive unravelling, achieving an isomorphic topological construction.

Chapter 5 provides a third, explicit, isomorphic construction, which is used to characterise exactly the structure of the descriptive unravelling. A brief elaboration is given on the three constructions and their potential uses in future research.

6.2 Directions for future research

This final section is dedicated to a discussion of potential directions for future work. As discussed in the introduction, many generalisations and adaptations exist of the classical van Benthem Characterisation Theorem. Directions that will be mentioned are a Janin-Walukiewicz-type result for descriptive models, coalgebraic generalisations, intuitionistic van Benthem-like theorems on Esakia spaces and a characterisation result for neighbourhood models. Especially the first two will be discussed elaborately.

Deserving special mention is the Janin-Walukiewicz Theorem [23]. This theorem is an analogue of the van Benthem Characterisation Theorem for the modal μ-calculus. It states that the modal μ-calculus is the bisimulation-invariant fragment of monadic first-order logic. The proof of this theorem also relies centrally on unravelling trees and the convenient properties of trees for game-theoretic semantics. In [6], a subclass of descriptive frames is designed to allow interpretation of the modal μ-calculus. To adapt this theorem to this class of descriptive models, one might use the descriptive unravelling to achieve similar results. As was shown in Chapter 5, the shape of descriptive unravellings is quite well-behaved and they, too, will likely satisfy desirable properties for game-theoretic semantics.

The coalgebraic generalisations of the van Benthem Characterisation Theorem for finite frames are inspired by the view of Kripke frames as coalgebras for the powerset functor on the category of sets. Similarly, the finite Kripke frames can be viewed as coalgebras for the powerset functor in the category of finite sets. This then leads to a strategy based on the pseudotrees as in [33] to obtain the van Benthem Characterisation Theorem for finite supported coalgebras. As was shown in this thesis, descriptive frames are model-theoretically highly similar to finite frames. As descriptive frames are coalgebras for the Vietoris functor on the category of Stone spaces, one could consider combining the constructions from this thesis with the approach used in [34]. There again, a type of pseudotree is introduced to apply arguments from finite model theory to obtain the result. Replacing the pseudotree with the descriptive unravelling could give way to a similar result for alternative Vietoris-like coalgebras on Stone spaces.

In [31, 32], the van Benthem Characterisation Theorem is treated for intuitionistic frames. Descriptive intuitionistic frames are known as Esakia spaces [18] and with the techniques from this thesis, one could pursue modal characterisation theorems on these classes.

Finally, neighbourhood structures have been given the structure of Stone coalgebras in [22], so that the constructions presented here may be combined with the approach from [21] to achieve a modal characterisation result on these neighbourhood structures over Stone spaces.

Bibliography

[1] J. v. Benthem. *Modal correspondence theory.* PhD thesis, University of Amsterdam, 1976.

[2] J. v. Benthem. Two simple incomplete modal logics. *Theoria*, 44(1):25–37, 1978.

[3] J. v. Benthem. Syntactic aspects of modal incompleteness theorems. *Theoria*, 45(2):63–77, 1979.

[4] J. v. Benthem. Modal logic for open minds. 2010.

[5] N. Bezhanishvili, G. Fontaine, and Y. Venema. Vietoris bisimulations. *Journal of Logic and Computation*, 20(5):1017–1040, 2010.

[6] N. Bezhanishvili and I. Hodkinson. Sahlqvist theorem for modal fixed point logic. *Theoret. Comput. Sci.*, 424:1–19, 2012.

[7] N. Bezhanishvili and A. Kurz. Free modal algebras: A coalgebraic perspective. In *International Conference on Algebra and Coalgebra in Computer Science*, pages 143–157. Springer, 2007.

[8] P. Blackburn, M. De Rijke, and Y. Venema. Modal logic (cambridge tracts in theoretical computer science), 2002.

[9] G. Boolos and G. Sambin. An incomplete system of modal logic. *Journal of Philosophical logic*, 14(4):351–358, 1985.

[10] G. E. Bredon. *Topology and geometry*, volume 139. Springer Science & Business Media, 2013.

[11] A. Chagrov and M. Zakharyaschev. *Modal logic*, volume 35 of *Oxford Logic Guides.* The Clarendon Press, Oxford University Press, New York, 1997. Oxford Science Publications.

[12] W. Craig. Three uses of the herbrand-gentzen theorem in relating model theory and proof theory. *The Journal of Symbolic Logic*, 22(3):269–285, 1957.

[13] B. Davey and H. Priestley. *Introduction to lattices and order.* Cambridge university press, 2002.

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

[15] H.-D. Ebbinghaus and J. Flum. *Finite model theory.* Springer Science & Business Media, 2005.

[16] R. Engelking. *General topology.* Heldermann, 1989.

[17] S. Enqvist and S. Sourabh. Bisimulations for coalgebras on Stone spaces. *Journal of Logic and Computation*, 28(6):991–1010, 2018.

[18] L. L. Esakia. Topological Kripke models. *Dokl. Akad. Nauk SSSR*, 214:298–301, 1974.

[19] E. Grädel, P. G. Kolaitis, L. Libkin, M. Marx, J. Spencer, M. Y. Vardi, Y. Venema, and S. Weinstein. *Finite Model Theory and its applications.* Springer Science & Business Media, 2007.

[20] W. Hanf. Model-theoretic methods in the study of elementary logic. In *The theory of models*, pages 132–145. Elsevier, 2014.

[21] H. Hansen, C. Kupke, and E. Pacuit. Neighbourhood Structures: Bisimilarity and Basic Model Theory. *Logical Methods in Computer Science*, 5(2), 2009.

[22] H. H. Hansen and C. Kupke. A coalgebraic perspective on monotone modal logic. *Electronic Notes in Theoretical Computer Science*, 106:121–143, 2004.

[23] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic. In *International Conference on Concurrency Theory*, pages 263–277. Springer, 1996.

[24] P. T. Johnstone. *Stone spaces*, volume 3. Cambridge university press, 1986.

[25] B. Jónsson and A. Tarski. Boolean algebras with operators. Part I. *American journal of mathematics*, 73(4):891–939, 1951.

[26] M. Kracht. *Tools and techniques in modal logic*, volume 97. Elsevier Amsterdam, 1999.

[27] C. Kupke, A. Kurz, and Y. Venema. Stone coalgebras. *Theoretical Computer Science*, 327(1-2):109–134, 2004.

[28] L. Libkin. *Elements of finite model theory.* Springer Science & Business Media, 2013.

[29] J. McCarthy. Epistemological problems of artificial intelligence. In *Readings in artificial intelligence*, pages 459–465. Elsevier, 1981.

[30] J. Minker. *Logic-based artificial intelligence*, volume 597. Springer Science & Business Media, 2012.

[31] G. Olkhovikov. Model-theoretic characterization of intuitionistic predicate formulas. *J. Log. Comput.*, 24(4):809–829, 2014.

[32] G. Olkhovikov. On generalized Van Benthem-type characterizations. *Annals of Pure and Applied Logic*, 168(9):1643–1691, 2017.

[33] E. Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6(4):427–439, 1997.

[34] L. Schröder, D. Pattinson, and T. Litak. A van Benthem/Rosen theorem for coalgebraic predicate logic. *Journal of Logic and Computation*, 27(3):749–773, 2017.

[35] M. H. Stone. The theory of representation for Boolean algebras. *Transactions of the American Mathematical Society*, 40(1):37–111, 1936.

[36] R. Turner et al. *Logics for artificial intelligence.* E. Horwood Chichester, 1984.

[37] Y. Venema and J. Vosmaer. Modal logic and the Vietoris functor. In *Leo Esakia on Duality in Modal and Intuitionistic Logics*, pages 119–153. Springer, 2014.

[38] F. Wolter et al. A counterexample in tense logic. *Notre Dame Journal of Formal Logic*, 37(2):167–173, 1996.