A Sound and Complete Axiomatisation for Intuitionistic Linear Temporal Logic

David Fernández-Duque[1], **Brett McLean**[2], **Lukas Zenger**[3]
[1]Department of Philosophy, University of Barcelona
[2]Department of Mathematics WE16, Ghent University
[3]Institute of Computer Science, University of Bern
fernandez-duque@ub.edu, brett.mclean@ugent.be, lukas.zenger@unibe.ch

Abstract

Intuitionistic linear temporal logic (iLTL) has been studied since at least the 1990s, with renewed interest in the last decade. It enjoys natural semantics over intuitionistic Kripke frames equipped with an order-preserving function representing the temporal dynamics, known as *expanding models*. This leads to a logic that is known to be decidable but whose axiomatisation has long remained open.

We propose an extension of iLTL with the co-implication connective of Heyting–Brouwer logic and call it *bi-intuitionistic linear temporal logic* (biLTL). We establish that this extension is still decidable for the class of expanding models. We moreover give a sound and complete Hilbert-style calculus for it, the first for any logic extending iLTL. As a corollary, the topological semantics for intuitionistic propositional logic cannot be extended to a topological semantics for Heyting–Brouwer logic, which thus establishes co-implication as a distinctive feature of the Kripke semantics for bi-intuitionistic logic.

1 Introduction

Constructive logics (Dalen 1986; Mints 2000) and temporal logics (Pnueli 1977) have long ago established their utility for modelling various aspects of computation, making it clear that a well-behaved combination of the two would be of the utmost value. In fact, there are two seemingly unrelated applications that independently led to the development of intuitionistic temporal logics: the first involves an extension of the Curry–Howard correspondence (Howard 1980) to account for user interaction in software (Kamide and Wansing 2010), and the second to logical modelling and automated theorem proving in the field of dynamical systems (Fernández-Duque 2018). In the latter, intuitionistic logic is interpreted according to its topological semantics over a typically infinite space X (where X may be e.g. the real line or the plane) and the temporal dynamics is modelled via a function $f\colon X \to X$, which, in order to interact well with the intuitionistic semantics, is assumed continuous. The pair (X, f) is a *dynamical system*, and may be used to model change over time in disciplines as diverse as biology, economics, and theoretical computer science. In addition to propositional variables and connectives, intuitionistic linear temporal logic (iLTL) includes the temporal modalities $\bigcirc$, $\Diamond$, and $\Box$. These are interpreted using the function f: $\bigcirc$ is read as 'next' and $\bigcirc\varphi$ holds at x if φ holds at $f(x)$, $\Diamond$ is read as 'eventually' and $\Diamond\varphi$ holds at x if φ holds at $f^n(x)$ for some n, and $\Box$ is read as 'henceforth' and $\Box\varphi$ holds at x if x has a neighbourhood U such that for every $y \in U$ and every natural number n, φ holds at $f^n(y)$.

Since intuitionistic logic also has a Kripke semantics based on partial orders, a bird's-eye-view representation of X may be provided using a finite Kripke model, allowing us to employ tools from *knowledge representation and reasoning* (KRR) for reasoning about topological dynamics.[1] This idea, while fruitful, does come with a caveat—namely, it is known that there are formulas valid over the class of Kripke models that are not valid over the class of topological models, such as the Rodríguez–Vidal formula **RV** $:= \Box(p \vee q) \to \Diamond p \vee \Box q$ (see Example 2).

This has led to a separate study of intuitionistic linear temporal logic over Kripke models, with a complete axiomatisation remaining a challenging open problem (Balbiani et al. 2020). The situation mirrors that of dynamic topological logic (DTL)—the 'classical' precursor of intuitionistic linear temporal logic (Artëmov, Davoren, and Nerode 1997)—where an axiomatisation for the DTL of Kripke models was proposed (Kremer and Mints 2007) but was later proven incomplete (Fernández-Duque 2014), with only an infinitary axiomatisation ever being found more than a decade later (Chopoghloo and Moniri 2022). The logic iLTL was proposed as an alternative to DTL in large part due to the fact that iLTL is decidable (Fernández-Duque 2018) but DTL is not (Konev et al. 2006). But the problem of finding a complete axiomatisation has proven equally elusive (once again with the exception of an infinitary axiomatisation (Chopoghloo and Moniri 2021)). In fact, the axiomatisation we provide in this paper relies on a novel insight that should also shed light on axiomatising DTL. Namely, in order to properly reason about Kripke models in a dynamic setting, it is necessary to look not only 'upward', but also 'downward'!

Intuitionistic logic differs from classical logic in how implication is treated. Classically, $\varphi \to \psi$ is read as material implication: either φ is false or ψ is true. This is often considered to be at odds with a natural language reading of im-

[1]Intuitionistic Kripke models may be seen as a special case of topological models, based on *Aleksandroff topologies* (Aleksandroff 1937).

plication in which $\varphi \to \psi$ means something like 'If φ *were* true, then ψ would also be true'. In the Kripke semantics of intuitionistic logic, this means that any 'possible world' where φ is true also makes ψ true, where a 'possible world' is any world lying above the current world with respect to a given partial order, i.e. models are provided with a partial order $\leq$ and $\varphi \to \psi$ holds at a world w if the material implication holds at every $v \geq w$.

In the standard intuitionistic language, worlds can only ever see 'upward': if, say, $v < w$, then v is in no way involved in the evaluation of formulas at w. This situation is 'remedied' by co-implication, where $\varphi \prec \psi$ holds at w if there is $v \leq w$ that makes φ true but ψ false. Intuitionistic logic expanded with co-implication—known as *bi-intuitionistic* or *Heyting–Brouwer* logic—was first extensively studied by Rauszer in the 1970s (Rauszer 1974a; Rauszer 1974b; Rauszer 1977). A deductive calculus for propositional bi-intuitionistic was proven sound and complete in (Rauszer 1974b), although (Drobyshevich, Odintsov, and Wansing 2022) note that an equivalent system was presented as early as (Moisil 1942), albeit without any discussion of soundness or completeness. See (Drobyshevich, Odintsov, and Wansing 2022) for many more references on the study and use of bi-intuitionistic logic.

There are various reasons to consider co-implication in intuitionistic reasoning. It is argued in (Fernández-Duque, McLean, and Zenger 2023) that for KRR tasks, where iLTL allows reasoning about *acquiring* resources in a temporal context, adding co-implication allows also reasoning about *losing* or *relinquishing* resources. Applicable areas therefore include *expert systems*, *automated planning*, and *temporal logic programming*.

Example 1. *In sensitive healthcare applications such as personalised cancer treatment, decisions must be made by considering various factors, including the type and stage of cancer and the patient's genetic profile. KRR systems can aid in integrating such data from various sources, possibly including some that are more reliable than others or even mutually contradictory.*

A KRR system might analyse a patient's data and predict that chemotherapy should shrink the tumour significantly. However, some of the data is inconclusive and, should it be incorrect, there could be a substantial chance of severe side effects.

In order to model this within our temporal logic framework, let us use the atom c to mean that chemotherapy is applied, r that the tumour is reduced and h that the patient is healthy. The expression $(c \to \Diamond r) \wedge (c \prec \Box h)$ then expresses the situation we have described: the intuitionistic implication tells us that, according to our current state of knowledge, chemotherapy will produce a reduction in the tumour, but co-implication tells us that should we relinquish some of this information, there is a risk of the patient becoming unhealthy. By using bi-intuitionistic logic as a basis for temporal reasoning, such considerations regarding uncertainty are directly built directly into the core of our KRR framework.

In applications such as the above, information is typically discrete and hence best modelled using Kripke semantics. Whether some natural topological analogue to $\prec$ can be defined remains open, but as we will see, the existence of co-implication satisfying propositional Heyting–Brouwer logic is a defining feature of Kripke, as opposed to the more general topological, models.

This is reinforced by our axiomatisation (Section 3). Formulas such as $\Box(p \vee q) \to \Diamond p \vee \Box q$ are nowhere to be seen (although derivable; see Example 4). Instead, by simply combining the natural axioms and rules for iLTL with standard axioms for $\prec$, our axioms 'magically' become complete. In technical parlance, our axiomatisation is not conservative over its $\prec$-free fragment. This might seem surprising, but a possible explanation is that interactions between $\prec$ and the unbounded temporal modalities implicitly arise when axioms for unbounded temporal modalities are applied to formulas that involve $\prec$. As an unexpected corollary, no notion of co-implication is definable in a topological setting validating bi-intuitionistic logic (Corollary 3).

2 Syntax and Semantics

Let us set up our formal system for bi-intuitionistic temporal logic. For this, we first fix a countable set of atomic variables *Prop*.

Definition 1. *The language $\mathcal{L}_{\mathsf{biLTL}}$ of bi-intuitionistic linear temporal logic is defined by the following grammar in Backus–Naur form (where $p \in Prop$):*

$$\varphi, \psi ::= p \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \varphi \to \psi \mid \varphi \prec \psi \mid \bigcirc\varphi \mid \Diamond\varphi \mid \Box\varphi$$

The connective $\prec$ is called **co-implication**. The modalities $\bigcirc$, $\Diamond$ and $\Box$ are read as 'next', 'eventually' and 'henceforth', respectively. We denote by $\mathcal{L}_{\bigcirc}$ the $\Diamond$- and $\Box$-free fragment of $\mathcal{L}_{\mathsf{biLTL}}$.

Define inductively $\bigcirc^0\varphi := \varphi$ and $\bigcirc^{n+1}\varphi := \bigcirc\bigcirc^n\varphi$. Furthermore, define $\top := p \to p$ and $\bot := p \prec p$ (where p is an arbitrary variable), $\neg\varphi := \varphi \to \bot$, and $\sim \varphi := \top \prec \varphi$. We call $\neg$ **strong negation** and $\sim$ **weak negation**.

Definition 2. *An **expanding model** is a tuple $\mathcal{M} = (W, \leq, f, V)$, where*

- *W is a set whose elements are called worlds.*
- *$\leq$ is a partial order on W.*
- *$f : W \to W$ is an order-preserving function:*

$$w \leq v \implies f(w) \leq f(v).$$

- *$V : W \to \mathcal{P}(Prop)$ is a **valuation** function that is monotone in $\leq$:*

$$w \leq v \implies V(w) \subseteq V(v).$$

For the remainder of this paper we will usually call expanding models simply *models*.

The satisfaction relation $\mathcal{M}, w \models \varphi$ between worlds of a model and formulas is defined by structural induction on φ as follows.

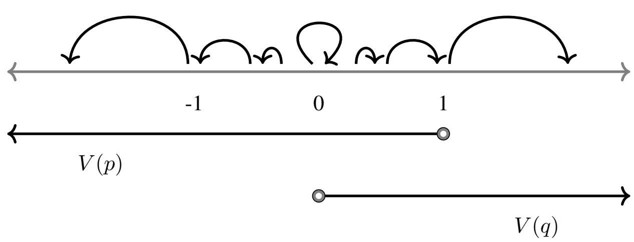

Figure 1: A model based on the real line.

$$
\begin{array}{lcl}
\mathcal{M}, w \models p & \Leftrightarrow & p \in V(w) \\
\mathcal{M}, w \models \varphi \wedge \psi & \Leftrightarrow & \mathcal{M}, w \models \varphi \text{ and } \mathcal{M}, w \models \psi \\
\mathcal{M}, w \models \varphi \vee \psi & \Leftrightarrow & \mathcal{M}, w \models \varphi \text{ or } \mathcal{M}, w \models \psi \\
\mathcal{M}, w \models \varphi \to \psi & \Leftrightarrow & \forall v \geq w\ (\mathcal{M}, v \models \varphi \\
 & & \quad \Rightarrow \mathcal{M}, v \models \psi) \\
\mathcal{M}, w \models \varphi \prec \psi & \Leftrightarrow & \exists v \leq w : \mathcal{M}, v \models \varphi \\
 & & \quad \text{and } \mathcal{M}, v \not\models \psi \\
\mathcal{M}, w \models \bigcirc\varphi & \Leftrightarrow & \mathcal{M}, f(w) \models \varphi \\
\mathcal{M}, w \models \Diamond\varphi & \Leftrightarrow & \exists n \in \mathbb{N} : \mathcal{M}, f^n(w) \models \varphi \\
\mathcal{M}, w \models \Box\varphi & \Leftrightarrow & \forall n \in \mathbb{N} : \mathcal{M}, f^n(w) \models \varphi
\end{array}
$$

If $\mathcal{M}, w \models \varphi$ we say that φ is **true** at w. A formula φ is **satisfiable** if there exists an expanding model $\mathcal{M}$ and a world w such that $\mathcal{M}, w \models \varphi$. A formula φ is **falsifiable** if there exists an expanding model $\mathcal{M}$ and a world w such that $\mathcal{M}, w \not\models \varphi$. A formula φ is **valid** if it is not falsifiable. Given a set of formulas Γ, a model $\mathcal{M}$, and a world w, we write $\mathcal{M}, w \models \Gamma$ if $\mathcal{M}, w \models \varphi$ for all $\varphi \in \Gamma$. A formula φ is a **(local) semantic consequence** of Γ, written $\Gamma \models \varphi$, if for each model $\mathcal{M}$ and world w, if $\mathcal{M}, w \models \Gamma$, then $\mathcal{M}, w \models \varphi$.

The proof the following lemma is by a straightforward induction on the structure of φ, where the order-preservation of f is used for the temporal cases.

Lemma 1 (monotonicity)**.** *Let $\mathcal{M} = (W, \leq, f, V)$ be an expanding model, $w \in W$, and φ be a formula. If $\mathcal{M}, w \models \varphi$ and $w \leq v$, then $\mathcal{M}, v \models \varphi$.*

We will not review *topological semantics* in detail here (see (Boudou et al. 2021) for formal definitions), but an example will be instructive.

Example 2. *Figure 1 shows that the formula **RV** $= \Box(p \vee q) \to \Diamond p \vee \Box q$ is not valid over the real line. For this discussion, it suffices to note that for a formula φ to be true at a point $x \in \mathbb{R}$, topological semantics requires that φ be true in a neighbourhood of x, i.e. in every point of some interval $(x - \varepsilon, x + \varepsilon)$.*

Define a model $\mathcal{M} = (\mathbb{R}, \leq, f, V)$ on $\mathbb{R}$, with the usual ordering and $f(x) = 2x$, $V(p) = (-\infty, 1)$, and $V(q) = (0, \infty)$. Clearly $p \vee q$ is true on all of $\mathbb{R}$, so $\Box(p \vee q)$ is true on $\mathbb{R}$ as well.

*Let us see that $\mathcal{M}, 0 \not\models$ **RV**. Since $\mathcal{M}, 0 \models \Box(p \vee q)$, it suffices to show that $\mathcal{M}, 0 \not\models \Box p \vee \Diamond q$. It is clear that $\mathcal{M}, 0 \not\models \Diamond q$ simply because $f^n(0) = 0 \notin V(q)$ for all n. Meanwhile, we cannot have $\mathcal{M}, 0 \models \Box p$ since for every $x > 0$ there is n with $f^n(x) > 1$, and hence $\mathcal{M}, x \not\models p$, which in turn implies that there can be no neighbourhood of 0 satisfying $\Box p$, and thus, by the topological semantics, $\mathcal{M}, 0 \not\models \Box p$. We conclude that $\mathcal{M}, 0 \not\models$ **RV**.*

3 A Hilbert-style Proof System

This section introduces the Hilbert-style proof system biLTL that captures $\mathcal{L}_{\mathsf{biLTL}}$-validities over expanding models. The system consists of the following axioms:

biInt	all bi-intuitionistic tautologies
D	$\neg\bigcirc\bot$
Dist	$\bigcirc(\varphi \vee \psi) \to (\bigcirc\varphi \vee \bigcirc\psi)$
K	$\bigcirc(\varphi \to \psi) \to (\bigcirc\varphi \to \bigcirc\psi)$
Fix$_\Diamond$	$\varphi \vee \bigcirc\Diamond\varphi \to \Diamond\varphi$
Fix$_\Box$	$\Box\varphi \to \varphi \wedge \bigcirc\Box\varphi$

and the following rules:

Sub	substitutions	**MP**	$\dfrac{\varphi \quad \varphi \to \psi}{\psi}$
Nec	$\dfrac{\varphi}{\bigcirc\varphi}$	**DN**	$\dfrac{\varphi}{\neg\sim\varphi}$
Mon$_\Diamond$	$\dfrac{\varphi \to \psi}{\Diamond\varphi \to \Diamond\psi}$	**Mon$_\Box$**	$\dfrac{\varphi \to \psi}{\Box\varphi \to \Box\psi}$
Ind$_\Diamond$	$\dfrac{\bigcirc\varphi \to \varphi}{\Diamond\varphi \to \varphi}$	**Ind$_\Box$**	$\dfrac{\varphi \to \bigcirc\varphi}{\varphi \to \Box\varphi}$

The system $\mathsf{biLTL}_\bigcirc$ is the restriction of the system biLTL to $\mathcal{L}_\bigcirc$ (i.e., only formulas of $\mathcal{L}_\bigcirc$ may appear in $\mathsf{biLTL}_\bigcirc$ derivations). We write $\vdash \varphi$ if there exists a proof of φ; whether derivability is in $\mathsf{biLTL}_\bigcirc$ or biLTL will be clear from context. A **proof with assumptions in** Γ is defined as usual, with the restriction that only the rule **MP** may be applied to formulas φ for which $\not\vdash \varphi$ holds. We write $\Gamma \vdash \varphi$ if φ is derivable with assumptions in Γ.

(Goré and Shillito 2020) prove a deduction theorem for bi-intuitionistic logic by a standard induction on the length of proofs. It is straightforward to extend their result to the language $\mathcal{L}_\bigcirc$.

Theorem 1 (Deduction Theorem)**.** *For any set of $\mathcal{L}_\bigcirc$-formulas Γ and any $\mathcal{L}_\bigcirc$-formulas φ and ψ, we have that $\Gamma, \varphi \vdash \psi$ if and only if $\Gamma \vdash \varphi \to \psi$.*

Furthermore, the following lemma holds (see (Goré and Shillito 2020), Proposition 7.2).

Lemma 2. *For arbitrary formulas in $\mathcal{L}_{\mathsf{biLTL}}$ the following hold:*

1. $\vdash \varphi \to (\psi \vee \chi) \iff \vdash (\varphi \prec \psi) \to \chi$.
2. *If* $\vdash \varphi \to \varphi'$ *and* $\vdash \psi' \to \psi$, *then* $\vdash (\varphi \prec \psi) \to (\varphi' \prec \psi')$.

We conclude this section with two instructive examples that show how the calculus interacts with co-implication.

Example 3. *The formula $(\bigcirc\varphi \prec \bigcirc\psi) \to \bigcirc(\varphi \prec \psi)$ is derivable in $\mathsf{biLTL}_\bigcirc$. To see this, we use the bi-intuitionistic tautology $\varphi \to (\psi \vee (\varphi \prec \psi))$. Applying necessitation and distribution, we see that $\bigcirc\varphi \to (\bigcirc\psi \vee \bigcirc(\varphi \prec \psi))$ is derivable. By Lemma 2, this shows that $(\bigcirc\varphi \prec \bigcirc\psi) \to \bigcirc(\varphi \prec \psi)$ is derivable as well.*

Example 4. *The formula **RV**, i.e. $\Box(\varphi \vee \psi) \to \Diamond\varphi \vee \Box\psi$, is derivable in biLTL. To see this, note that $\Box(\varphi \vee \psi) \to (\Diamond\varphi \vee (\Box(\varphi \vee \psi) \prec \Diamond\varphi))$ is a substitution instance of a bi-intuitionistic tautology, so it suffices to check that $(\Box(\varphi \vee \psi) \prec \Diamond\varphi) \to \Box\psi$ is derivable. Since $\Box(\varphi \vee \psi) \to \bigcirc\Box(\varphi \vee \psi)$ and $\bigcirc\Diamond\varphi \to \Diamond\varphi$ are derivable, so is $(\Box(\varphi \vee \psi) \prec \Diamond\varphi) \to (\bigcirc\Box(\varphi \vee \psi) \prec \bigcirc\Diamond\varphi))$ by Lemma 2. Hence by*

Example 3, $(\Box(\varphi \vee \psi) \prec \Diamond\varphi) \to \bigcirc(\Box(\varphi \vee \psi) \prec \Diamond\varphi)$ is derivable. By ***Ind***$_\Box$*, we obtain $(\Box(\varphi \vee \psi) \prec \Diamond\varphi) \to \Box(\Box(\varphi\vee\psi)\prec\Diamond\varphi)$, and thus it suffices to show that $(\Box(\varphi\vee\psi) \prec \Diamond\varphi) \to \psi$ is derivable.*

Now, since $\Box(\varphi\vee\psi) \to \varphi\vee\psi$ and $\varphi \to \Diamond\varphi$ are derivable, so is $(\Box(\varphi\vee\psi)\prec\Diamond\varphi) \to ((\varphi\vee\psi)\prec\varphi)$. But $((\varphi\vee\psi)\prec\varphi) \to \psi$ is a bi-intuitionistic tautology; hence $(\Box(\varphi\vee\psi)\prec\Diamond\varphi) \to \psi$ is derivable, as desired.

4 Soundness of biLTL and Completeness of biLTL$_\bigcirc$

In this section we first remark that biLTL is sound (with respect to the class of expanding models), i.e. if a $\mathcal{L}_{\mathsf{biLTL}}$-formula φ is biLTL-provable, then it is valid. Afterwards we show that biLTL$_\bigcirc$ is complete, i.e. if a $\mathcal{L}_\bigcirc$-formula φ is valid, then it is biLTL$_\bigcirc$-provable.

Lemma 3 (Boudou et al. 2021)**.** *The axioms of* biLTL *are valid over the class of expanding models and the rules preserve validity.*

Using Lemma 3 and induction on the length of proofs, we obtain soundness of biLTL.

Theorem 2. *If a formula φ is* biLTL*-provable, then φ is valid over the class of expanding models.*

For completeness of biLTL$_\bigcirc$, we employ a canonical model construction. As the canonical model is also used later on in the completeness proof of biLTL, the following definitions and lemmas apply to both $\mathcal{L}_{\mathsf{biLTL}}$ and $\mathcal{L}_\bigcirc$.

Definition 3. *A* ***prime theory*** *is a set of $\mathcal{L}$-formulas Γ, where $\mathcal{L}$ is either $\mathcal{L}_{\mathsf{biLTL}}$ or $\mathcal{L}_\bigcirc$, such that the following hold:*

1. *Γ is deductively closed: if $\Gamma \vdash \varphi$, then $\varphi \in \Gamma$;*
2. *Γ satisfies the disjunction property: if $\varphi \vee \psi \in \Gamma$, then $\varphi \in \Gamma$ or $\psi \in \Gamma$, and*
3. *Γ is consistent: $\Gamma \nvdash \bot$.*

Given a set of formulas Γ, define

$$\bigcirc^{-1}\Gamma \coloneqq \{\varphi \mid \bigcirc\varphi \in \Gamma\}.$$

Lemma 4. *If Γ is a prime theory, then $\bigcirc^{-1}\Gamma$ is a prime theory as well.*

It is the axiom **Dist** that ensures $\bigcirc^{-1}\Gamma$ satisfies the disjunction property.

We are now ready to define the canonical models for $\mathcal{L}_{\mathsf{biLTL}}$ and $\mathcal{L}_\bigcirc$.

Definition 4. *Let $\mathcal{L}$ be either $\mathcal{L}_\bigcirc$ or $\mathcal{L}_{\mathsf{biLTL}}$. The* ***canonical model*** *for $\mathcal{L}$ is defined to be $\mathcal{M}_{\rm c} = (W_{\rm c}, \leq_{\rm c}, f_{\rm c}, V_{\rm c})$ where*

- *$W_{\rm c} = \{\Gamma \subseteq \mathcal{L} \mid \Gamma$ is a prime theory$\}$,*
- *$\Gamma \leq_{\rm c} \Gamma' \iff \Gamma \subseteq \Gamma'$,*
- *$f_{\rm c}(\Gamma) = \bigcirc^{-1}\Gamma$,*
- *$V_{\rm c}(\Gamma) = \{p \in Prop \mid p \in \Gamma\}$.*

The proof of the following lemma is standard and omitted (see (Boudou, Diéguez, and Fernández-Duque 2017; Boudou, Diéguez, and Fernández-Duque 2022)).

Lemma 5. *The canonical model for either $\mathcal{L}_\bigcirc$ or $\mathcal{L}_{\mathsf{biLTL}}$ is an expanding model.*

For the remainder of this section we work exclusively in the language $\mathcal{L}_\bigcirc$ and show that biLTL$_\bigcirc$ is complete. The following lemma establishes that every consistent set of $\mathcal{L}_\bigcirc$-formulas can be extended to a prime theory. The proof is standard, see e.g. (Goré and Shillito 2020).

Lemma 6 (Lindenbaum lemma)**.** *Suppose $\Gamma \nvdash \chi$. Then there exists a prime theory Δ with $\Gamma \subseteq \Delta$ and $\Delta \nvdash \chi$.*

Lemma 7 (truth lemma)**.** *Let $\mathcal{M}_{\rm c}$ be the canonical model for $\mathcal{L}_\bigcirc$. For every $\Gamma \in W_{\rm c}$ and any $\mathcal{L}_\bigcirc$-formula φ, it holds that*

$$\varphi \in \Gamma \iff \mathcal{M}_{\rm c}, \Gamma \models \varphi.$$

Theorem 3. *If a $\mathcal{L}_\bigcirc$-formula φ is valid over the class of expanding models, then φ is* biLTL$_\bigcirc$*-provable.*

Proof. Suppose φ is not biLTL$_\bigcirc$-provable, i.e. $\emptyset \nvdash \varphi$. By Lemma 6 there exists a prime theory Γ with $\Gamma \nvdash \varphi$. Hence $\varphi \notin \Gamma$, and so by Lemma 7, we have $\mathcal{M}_{\rm c}, \Gamma \not\models \varphi$. We conclude that φ is not valid. □

5 Proof Strategy for the Full Language

In the remainder of the article, we prove that full biLTL is complete for the class of expanding models. Unlike for biLTL$_\bigcirc$, we do not have a truth lemma for the canonical model, since it may be for example that $\Diamond\varphi \in \Gamma$, but there is no n such that $\varphi \in f^n_{\rm c}(\Gamma)$.[2]

A similar situation occurs for classical LTL, but one can then pass to a filtration ${}^{\mathcal{M}_{\rm c}}\!/_{\Sigma}$ of $\mathcal{M}_{\rm c}$, i.e. the quotient of $\mathcal{M}_{\rm c}$ modulo the equivalence relation $\Gamma \sim \Gamma' \iff \Gamma \cap \Sigma = \Gamma' \cap \Sigma$. Assuming Σ is finite, the equivalence class of each prime theory Γ is determined by its *characteristic formula* $\chi(\Gamma) \coloneqq \bigwedge(\Gamma \cap \Sigma)$. The filtrated model *does* respect the semantics of $\Diamond$. More precisely, ${}^{\mathcal{M}_{\rm c}}\!/_{\Sigma}$ satisfies a version of the truth lemma restricted to formulas of Σ. The tradeoff is that ${}^{\mathcal{M}_{\rm c}}\!/_{\Sigma}$ is no longer equipped with a *function*, as the quotient may assign more than one temporal successor to a single prime theory, since $\Gamma \sim \Gamma'$ does not imply $f_{\rm c}(\Gamma) \sim f_{\rm c}(\Gamma')$. However, this is not a problem, since in a later phase one can choose a path $\Gamma_0, \Gamma_1, \Gamma_2, \ldots$ that constitutes a genuine LTL model. In particular, if φ is not derivable, we can choose Σ to be the set of subformulas of φ and their negations and Γ_0 so that $\varphi \notin \Gamma_0$, thereby obtaining a model falsifying φ.

We wish to adapt this strategy, but there is an issue: filtration in general does not conserve order-preservation of the temporal dynamics (i.e. $w \leq v$ implies $f(w) \leq f(v)$), so we must define ${}^{\mathcal{M}_{\rm c}}\!/_{\Sigma}$ differently. This structure should be a *quasimodel*, which is similar to a model except that the temporal transition *function* is replaced by a non-deterministic *relation*. Each point in a quasimodel is assigned a *type*, which is similar to a prime theory except that a type only decides a finite set of formulas; i.e., a type is a pair $\Phi = (\Phi^+, \Phi^-)$ of (usually) finite sets of formulas for which a 'truth lemma' should hold. Quasimodels are designed so that they can be 'unwound' into a genuine model, much like

[2]This is because $\Diamond\varphi \vdash \bigvee_{i<n} \bigcirc^i\varphi$ is not derivable for any specific n, and derivations are finite. Hence it is possible for $\Diamond\varphi$ to hold but each individual $\bigcirc^n\varphi$ to fail in a prime theory.

for the filtrated model of classical LTL. Types, quasimodels, and the more general *labelled systems* are introduced in Section 6, and Section 7 describes the unwinding procedure.

To construct $\mathcal{M}_c/\Sigma$ in the bi-intuitionistic setting, we first construct a structure $\mathbb{U}_\Sigma$ that is finite but still 'too large', as it may contain points that do not correspond to any prime theory. The structure $\mathbb{U}_\Sigma$ is defined in (Fernández-Duque, McLean, and Zenger 2023) in a general modal setting, but we readily apply it to biLTL by simply regarding $\bigcirc$ as a modal operator. While a full construction is out of the scope of this paper, roughly speaking $\mathbb{U}_\Sigma$ consists of the set of all finite, 'acyclic' posets with points labelled by types and bounded in size by some large enough natural number.

In a standard filtration, each prime theory maps to a single equivalence class. Here, however, we have a binary relation E_* between $\mathbb{U}_\Sigma$ and $\mathcal{M}_c$, where there are one or more points of $\mathbb{U}_\Sigma$ linked to each prime theory in $\mathcal{M}_c$. The relation E_* is what we call a *dynamic simulation* (Definition 15), and further is 'exhaustive'. Then $\mathcal{M}_c/\Sigma$ is defined to be the restriction of $\mathbb{U}_\Sigma$ to the domain of E_*. The relation E_* can be thought of as a relational version of the filtration quotient.

Much as the characteristic formula $\chi(\Gamma)$ determines the equivalence class of Γ in the classical setting, we can characterise which points of $\mathbb{U}_\Sigma$ are E_*-related to a given prime theory Γ using *simulation formulas* (Definition 17). Unlike the classical setting, we need two distinct formulas, χ^+ and χ^-, to capture respectively the 'positive' and 'negative' information determining a simulation. Section 8 discusses (dynamic) simulations, while Section 9 defines the formulas χ^+ and χ^- and establishes their basic properties. These simulation formulas enable us to prove that $\mathcal{M}_c/\Sigma$ is indeed a quasimodel (Corollary 1).

At the end of Section 10 we put all the ingredients together to show that biLTL is indeed complete for the class of expanding posets: the argument is that if φ is not derivable then we can find a prime theory Γ with $\varphi \in \Gamma^-$. By choosing a point w of $\mathbb{U}_\Sigma$ with $w \; E_* \; \Gamma$, we see that φ is falsified on $\mathcal{M}_c/\Sigma$. By applying the unwinding procedure to $\mathcal{M}_c/\Sigma$, we obtain a genuine model falsifying φ. Thus every formula that is not derivable can be falsified in some expanding model, i.e. biLTL is complete, our main result (Theorem 5).

6 Types, Labelled Posets, and Quasimodels

From now on, Σ denotes a set of $\mathcal{L}_{\mathsf{biLTL}}$-formulas closed under subformulas.

Definition 5. *Let $\Phi^+, \Phi^- \subseteq \Sigma$. A Σ-**type** is a pair $\Phi = (\Phi^+, \Phi^-)$ of disjoint subsets of Σ with the following properties:*

$\wedge^+$. *If $\varphi \wedge \psi \in \Phi^+$, then $\varphi, \psi \in \Phi^+$.*

$\wedge^-$. *If $\varphi \wedge \psi \in \Phi^-$, then $\varphi \in \Phi^-$ or $\psi \in \Phi^-$.*

$\vee^+$. *If $\varphi \vee \psi \in \Phi^+$, then $\varphi \in \Phi^+$ or $\psi \in \Phi^+$.*

$\vee^-$. *If $\varphi \vee \psi \in \Phi^-$, then $\varphi, \psi \in \Phi^-$.*

$\to^+$. *If $\varphi \to \psi \in \Phi^+$, then $\varphi \in \Phi^-$ or $\psi \in \Phi^+$.*

$\to^-$. *If $\varphi \to \psi \in \Phi^-$, then $\psi \in \Phi^-$.*

$\prec^+$. *If $\varphi \prec \psi \in \Phi^+$, then $\varphi \in \Phi^+$.*

$\prec^-$. *If $\varphi \prec \psi \in \Phi^-$, then $\varphi \in \Phi^-$ or $\psi \in \Phi^+$.*

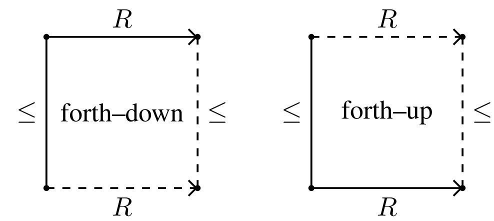

Figure 2: Forward confluence conditions

$\Box^+$. *If $\Box\varphi \in \Phi^+$, then $\varphi \in \Phi^+$.*

$\Diamond^-$. *If $\Diamond\varphi \in \Phi^-$, then $\varphi \in \Phi^-$.*

It is not necessary that $\Phi^+ \cup \Phi^- = \Sigma$. Thus our types are 'partial'.[3] *The set of all Σ-types is denoted by T_Σ.*

To compare types, we define two partial orders on T_Σ:

1. $\Phi \leq_{\mathrm{T}} \Psi$ if and only if $\Phi^+ \subseteq \Psi^+$ and $\Psi^- \subseteq \Phi^-$ (corresponding to the intuitionistic partial order).
2. $\Phi \subseteq_{\mathrm{T}} \Psi$ if and only if $\Phi^+ \subseteq \Psi^+$ and $\Phi^- \subseteq \Psi^-$ (so Ψ 'asserts' more than Φ, both positively and negatively).

Definition 6. *Let Φ be a Σ-type.*

1. *A formula $\varphi \to \psi$ is a **defect** of Φ if $\varphi \to \psi \in \Phi^-$, but $\varphi \notin \Phi^+$.*
2. *A formula $\varphi \prec \psi$ is a defect of Φ if $\varphi \prec \psi \in \Phi^+$, but $\psi \notin \Phi^-$.*

The set of all defects of Φ is denoted by $\delta\Phi$.

In the following, we define labelled posets and quasimodels. We first define labelled posets, which are partial orders whose nodes are labelled by types.

Definition 7. *A Σ-**labelled poset** is a tuple $\mathcal{X} = (X, \leq, \ell)$ where $(X \; \leq)$ is a partial order, and $\ell : X \to \mathrm{T}_\Sigma$ is a labelling function such that the following hold.*

1. *If $x \leq y$, then $\ell(x) \leq_{\mathrm{T}} \ell(y)$.*
2. *If $\varphi \to \psi \in \delta\ell(x)$, then there exists $y \geq x$ with $\varphi \in \ell(y)^+$ and $\psi \in \ell(y)^-$.*
3. *If $\varphi \prec \psi \in \delta\ell(x)$, then there exists $y \leq x$ with $\varphi \in \ell(y)^+$ and $\psi \in \ell(y)^-$.*

If $\chi \in \delta(\ell(x))$ and y is the world from the above definition, then we say that the defect χ is **resolved** at y. We will usually assume that Σ is finite, and thus that labelled posets are labelled with finite types. Given a labelled poset $\mathcal{X} = (X, \leq, \ell)$, a relation $R \subseteq X \times X$ is called **forward confluent** if it satisfies the following two properties (see Figure 2).

Forth–up: If $x \leq x'$ and $x \; R \; y$ then $\exists y' \geq y$ with $x' \; R \; y'$.

Forth–down: If $x \leq x'$ and $x' \; R \; y'$ then $\exists y \leq y'$ with $x \; R \; y$.

Definition 8. *Let Φ, Ψ be Σ-types. The pair (Φ, Ψ) is called **sensible** if the following conditions hold.*

[3]The need to consider partial types will not be too evident in the current work, but it is needed to import results from (Fernández-Duque, McLean, and Zenger 2023). See Footnote 5 for a pointer to why they are needed.

1. *If* $\bigcirc\varphi \in \Phi^+$, *then* $\varphi \in \Psi^+$.
2. *If* $\bigcirc\varphi \in \Phi^-$, *then* $\varphi \in \Psi^-$.
3. *If* $\Diamond\varphi \in \Phi^+$, *then* $\varphi \in \Phi^+$ *or* $\Diamond\varphi \in \Psi^+$.
4. *If* $\Diamond\varphi \in \Phi^-$, *then* $\varphi \in \Phi^-$ *and* $\Diamond\varphi \in \Psi^-$.
5. *If* $\Box\varphi \in \Phi^+$, *then* $\varphi \in \Phi^+$ *and* $\Box\varphi \in \Psi^+$.
6. *If* $\Box\varphi \in \Phi^-$, *then* $\varphi \in \Phi^-$ *or* $\Box\varphi \in \Psi^-$.

Given a Σ*-labelled poset* $\mathcal{X} = (X, \leq, \ell)$*, a pair* $(x, y) \in X \times X$ *is called* ***sensible*** *if* $(\ell(x), \ell(y))$ *is sensible. A relation* $R \subseteq X \times X$ *is called sensible if every pair* $(x, y) \in R$ *is sensible.*

Definition 9. *Given a* Σ*-labelled poset* $\mathcal{X} = (X, \leq, \ell)$*, a sensible relation* $R \subseteq X \times X$ *is called* ω***-sensible*** *if the following hold.*

1. *If* $\Diamond\varphi \in \ell(x)^+$, *then there are* $n \in \mathbb{N}$ *and* $y \in X$ *such that* $x\ R^n\ y$ *and* $\varphi \in \ell(y)^+$.
2. *If* $\Box\varphi \in \ell(x)^-$, *then there are* $n \in \mathbb{N}$ *and* $y \in X$ *such that* $x\ R^n\ y$ *and* $\varphi \in \ell(y)^-$.

Recall: a relation $R \subseteq X \times Y$ is **total** if for each $x \in X$ there exists $y \in Y$ with $x\ R\ y$. If in addition $X = Y$, then we say R is **serial**.

Definition 10. *A* Σ***-labelled system*** *is a tuple* $\mathcal{X} = (X, \leq, \ell, R)$ *consisting of a labelled poset equipped with a forward-confluent sensible relation* $R \subseteq X \times X$*. If moreover* R *is serial and* ω*-sensible, then* $\mathcal{X}$ *is a* Σ***-quasimodel****.*

We may write simply *labelled system* or *quasimodel* when Σ is clear from context. A formula φ is **falsified** at world x of a Σ-quasimodel $\mathcal{X} = (X, \leq, \ell, R)$ if $\varphi \in \ell(x)^-$, and **satisfied** if $\varphi \in \ell(x)^+$. A formula φ is falsifiable over the class of Σ-quasimodels if there exists a Σ-quasimodel $\mathcal{X} = (X, \leq, \ell, R)$ and a world $x \in X$ such that φ is falsified at x. Note that it is possible that a formula $\varphi \in \Sigma$ is neither satisfied nor falsified at a world $x \in X$.

Observe that every expanding model can be regarded as a Σ-quasimodel by simply labelling each world with those formulas in Σ that are true or false respectively. Thus we obtain the following result.

Lemma 8. *If* $\varphi \in \Sigma$ *is falsifiable over the class of expanding models, then* φ *is falsifiable over the class of* Σ*-quasimodels.*

The converse of Lemma 8 is also true, but establishing it requires some work. This will be done in the next section. However, we can already state the result for a particular subclass of quasimodels.

Definition 11. *A* Σ*-quasimodel* $\mathcal{X} = (X, \leq, \ell, R)$ *is* ***functional*** *if* $R \subseteq X \times X$ *is a function.*

Lemma 9. *If a formula is falsifiable over the class of functional* Σ*-quasimodels, then it is falsifiable over the class of expanding models.*

Proof. Let $\mathcal{X} = (X, \leq, \ell, f)$ be a functional Σ-quasimodel, $x \in X$, and φ a formula such that $\varphi \in \ell(x)^-$. Define $\mathcal{M} := (X, \leq, f, V)$ where $V(y) := \ell(y)^+ \cap Prop$. It is routine to check that $\mathcal{M}$ is an expanding model, and $\mathcal{M}, x \not\models \varphi$. $\square$

7 From Quasimodels to Expanding Models

It will be useful to observe that forward confluence can be iterated, thereby yielding the following variant for finite paths.

Lemma 10. *Let* $\mathcal{X} = (X, \leq, \ell, R)$ *be a quasimodel. Suppose that* $w_0\ R\ w_1\ R \ldots R\ w_n$*.*

- *If* $w_0 \leq u_0$ *then there exist* $u_0\ R\ u_1\ R \ldots R\ u_n$ *such that* $w_i \leq u_i$ *for all* $i \leq n$*.*
- *If* $u_0 \leq w_0$ *then there exist* $u_0\ R\ u_1\ R \ldots R\ u_n$ *such that* $u_i \leq w_i$ *for all* $i \leq n$*.*

Proof. Inductively find u_i using the forward confluence of R. $\square$

For the remainder of this section let $\mathcal{X} = (X, \leq, \ell, R)$ be a fixed Σ-quasimodel. Suppose that $\mathcal{X}$ falsifies some formula $\varphi \in \Sigma$. We are going to show how to construct from $\mathcal{X}$ a *functional* Σ-quasimodel falsifying φ. Combining this construction with Lemma 9 then yields the construction of an expanding model falsifying φ.

Given a partial function f we write $\exists f(s)$ if $f(s)$ is defined and $\nexists f(s)$ otherwise.

For a partial order $\leq$, the element y **covers** x if $x \leq y$ and there is no $x < w < y$. We say that $\leq$ is **acyclic** if the *undirected* graph induced by its covering relation is acyclic.

Definition 12. *An* $\mathcal{X}$***-induced structure*** *is a tuple* $\mathcal{I} = (I, \leq_I, \ell_I, f_I)$ *together with a map* $\pi \colon I \to X$ *where:*

1. I *is finite,*
2. $\leq_I$ *is acyclic and if* $w \leq_I v$ *then* $\pi(w) \leq \pi(v)$*,*
3. $\ell_I = \ell_X \circ \pi$*, and*
4. $f_I : I \to I$ *is a partial function such that:*
   (a) *If* $\exists f_I(x)$*, then* $\pi(x)\ R\ \pi(f_I(x))$*.*
   (b) *If* $x \leq_I y$ *then* $\exists f(x) \iff \exists f(y)$*.*
   (c) *If* $x \leq_I y$ *and* $\exists f(x)$*, then* $f_I(x) \leq_I f_I(y)$*.*
   (d) *For each* $x \in I$ *there is a maximal* k *such that* $\exists f^k(x)$*.*

It is instructive to view induced structures as being 'temporally stratified'. In view of (4d) and the assumption that I is finite, there is a maximal k such that $f^k(x)$ is defined for *any* $x \in I$, and hence we may define W_i to be the set of $x \in I$ such that $f^{k-i}(x)$ is defined but $f^{k-i+1}(x)$ is not. This partitions I into sets $W_0, \ldots, W_k$, and it is easy to see that $x \leq_I y$ implies that $x, y \in W_i$ for some i, and moreover $f[W_i] \subseteq W_{i+1}$ for all i.

A *defect* of an $\mathcal{X}$-induced structure records that a claim made by its labelling ℓ_I lacks a witness.

Definition 13. *Let* $\mathcal{I}$ *be an* $\mathcal{X}$*-induced structure.*

1. *A* $\to$***-defect*** *is a pair* $(x, \varphi \to \psi)$ *where* $x \in I$ *and* $\varphi \to \psi \in \ell_I(x)^-$*, but there is no* $y \geq_I x$ *with* $\varphi \in \ell_I(y)^+$ *and* $\psi \in \ell_I(y)^-$*.*
2. *A* $\prec$*-defect is a pair* $(x, \varphi \prec \psi)$ *where* $x \in I$ *and* $\varphi \prec \psi \in \ell_I(x)^+$*, but there is no* $y \leq_I x$ *with* $\varphi \in \ell_I(y)^+$ *and* $\psi \in \ell_I(y)^-$*.*
3. *A* $\bigcirc$*-defect is a world* $x \in I$ *with* $\nexists f_I(x)$*.*
4. *A* $\Diamond$*-defect is a pair* $(x, \Diamond\varphi)$ *where* $x \in I$*,* $\nexists f_I(x)$*, and* $\Diamond\varphi \in \ell_I(x)^+$*, but* $\varphi \notin \ell_I(x)^+$*.*

5. *A $\Box$-defect is a pair $(x, \Box\varphi)$ where $x \in I$, $\nexists f_I(x)$, and $\Box\varphi \in \ell_I(x)^-$, but $\varphi \notin \ell_I(x)^-$.*

Let $x \in X$ be such that $\varphi \in \ell(x)^-$. We build a functional Σ-quasimodel falsifying φ in stages. We start with an $\mathcal{X}$-induced structure $\mathcal{I}_0$ consisting of a single world and then construct in the step $n+1$ an $\mathcal{X}$-induced structure $\mathcal{I}_{n+1}$ extending $\mathcal{I}_n$. We make use of a first-in-first-out queue D that stores the defects of the current $\mathcal{X}$-induced structure. Observe that for any $\mathcal{X}$-induced structure, the set of defects of said structure is always finite (since the structure and Σ are finite) and non-empty (due to $\bigcirc$-defects). The $\mathcal{X}$-induced structure $\mathcal{I}_n$ is defined by induction on n as follows.

For the base case, define $\mathcal{I}_0 = (I_0, \leq_0, \ell_0, f_0)$, where $I_0 = \{x'\}$ (it is not important what x' is), $\leq_0 = \{(x', x')\}$, $\ell_0(x') = \ell(x)$, $f_0 = \emptyset$, and $\pi_0(x') = x$. It is straightforward to check that $(\mathcal{I}_0, \pi_0)$ is an $\mathcal{X}$-induced structure. Initialise D with all defects of $\mathcal{I}_0$ in arbitrary order.

For the inductive step, suppose we have defined $\mathcal{I}_n = (I_n, \leq_n, \ell_n, f_n)$ and π_n, and shown that $(\mathcal{I}_n, \pi_n)$ is an $\mathcal{X}$-induced structure. By inductive hypothesis, D currently stores all defects of $\mathcal{I}_n$. We first show how to define $(\mathcal{I}_{n+1}, \pi_{n+1})$ and then how to update the queue D. We start by setting $I_{n+1} = I_n$. We only treat defects for $\bigcirc$, $\rightarrow$, and $\Diamond$; other cases are similar.

($\bigcirc$-DEFECTS) Suppose the defect at the head of D is a $\bigcirc$-defect $y \in I_n$. Choose any $u \in X$ with $\pi_n(y)\ R\ u$. Add a new point u' to I_{n+1} and define $f_{n+1}(y) = u'$, $\ell_{n+1}(u') = \ell(u)$, and $\pi_{n+1}(u') = u$. We extend f_{n+1} to the connected component of y by adding new worlds, working first 'bottom up' starting with worlds covering y. If y' covers y, use forward confluence to find $z \in X$ with $\pi_n(y')\ R\ z$. Add z' to $\mathcal{I}_{n+1}$ and define $f_{n+1}(y') = z'$, $u' \leq_{n+1} z'$ and close $\leq_{n+1}$ under transitivity and reflexivity,[4], $\pi_{n+1}(z') = z$, and $\ell_{n+1}(z') = \ell(z)$. Then for y'' covering such y', use forward confluence again (relative to y') to define $f_{n+1}(y'')$, and so on. Next repeat the process for those worlds below $\{y' \mid y' \geq y\}$ where f_{n+1} is not yet defined, this time working 'top down'. Continue alternating between 'bottom up' and 'top down' until f_{n+1} is defined on the connected component of y.

This process terminates because the new points are not in the connected component of y, which is finite. Thus if the connected component of y in $\mathcal{I}_n$ is of size m, then m new points are added.

($\rightarrow$-DEFECTS) Suppose the defect at the head of D is a $\rightarrow$-defect $(y, \psi \rightarrow \chi)$. Then $\psi \rightarrow \chi \in \ell_n(y)^-$, but there is no $z' \in I_n$ with $y \leq_n z'$, and $\varphi \in \ell_n(z')^+$ and $\psi \in \ell_n(z')^-$. As $\mathcal{X}$ is a quasimodel, there exists $\pi_n(y) \leq z \in X$ with $\psi \in \ell(z)^+$ and $\chi \in \ell(z)^-$. Let k be maximal such that $f_n^k(y)$ is defined. Using Lemma 10, find $z = z_0\ R\ z_1\ R \ldots R\ z_k$ with $\pi_n(f_n^i(y)) \leq z_i$. Then add points $z_0', \ldots, z_k'$ to I_{n+1} and extend π_n, $\leq_n$, f_n and ℓ_n by setting $\pi_{n+1}(z_i') = z_i$, $f_n^i(y) \leq_{n+1} z_i'$, $f_{n+1}(z_i') = z_{i+1}'$ and $\ell_{n+1}(z_i') = \ell(z_i)$.

($\Diamond$-DEFECTS) Suppose the defect at the head of D is a $\Diamond$-defect $(y, \Diamond\psi)$. Then $\nexists f_n(y)$ and $\Diamond\psi \in \ell_n(y)^+$, but $\psi \notin \ell_n(y)^+$. As $\mathcal{X}$ is a quasimodel we find $u_1, \ldots, u_n \in X$ with $\pi_n(y)\ R\ u_1\ R\ u_2\ R \ldots R\ u_n$ and $\psi \in \ell(u_n)^+$. We add worlds $u_1', \ldots, u_n'$ to I_{n+1} with $\pi_{n+1}(u_i') = u_i$, $f_{n+1}(y) = u_1'$ and $f_{n+1}(u_i') = u_{i+1}'$, and $\ell_{n+1}(u_i') = \ell(u_i)$. Then we proceed as in the case of a $\bigcirc$-defect to define f_{n+1} on the connected component of y, and proceed inductively to define f_{n+1} on the connected component of each u_i'. In this case, we must add n-many components for some natural number n. Hence, the construction for 'next'-defects must be repeated n-many times. Thus the termination of this process is proven by induction on n, with a secondary induction on the number of worlds in a component as in the $\bigcirc$-defect case.

We have shown how to construct $(\mathcal{I}_{n+1}, \pi_{n+1})$ from $(\mathcal{I}_n, \pi_n)$. Next we show how to update the queue D. First, delete every defect from D that has been resolved in the construction of $(\mathcal{I}_{n+1}, \pi_{n+1})$ (observe that in each of the above cases it is possible that multiple defects have been resolved at once). Then we rewrite each remaining defect as follows. If the remaining defect is a $\rightarrow$-defect or a $\prec$-defect we don't change anything. If it is a $\Diamond$-defect $(y, \Diamond\psi)$ we check whether $\nexists f_{n+1}(y)$ holds. If it does we do not change the defect. Otherwise there are $u_1, \ldots, u_k$ with $f_{n+1}(y) = u_1, f_{n+1}(u_1) = u_2, \ldots, f_{n+1}(u_{k-1}) = u_k$ and $\nexists f_{n+1}(u_k)$. By assumption $\Diamond\psi \in \ell_{n+1}(u_k)^+$ and $\psi \notin \ell_{n+1}(u_k)^+$. Thus overwrite $(y, \Diamond\psi)$ with $(u_k, \Diamond\psi)$. The $\Box$-defects and seriality defects are overwritten in the same way. Finally, add all *new* defects of $\mathcal{I}_{n+1}$ to the tail of the queue.

By induction on n, each $\mathcal{I}_n$ is an $\mathcal{X}$-induced structure. Furthermore, by construction, each $\mathcal{I}_{n+1}$ contains $\mathcal{I}_n$ as a substructure. Define the structure $\mathcal{I}_\omega$ to be the limit of the sequence $(\mathcal{I}_n)_{n \in \mathbb{N}}$. More formally, define $(\mathcal{I}_\omega := (I_\omega, \leq_\omega, \ell_\omega, f_\omega), \pi_\omega)$ where

$$\lambda_\omega = \bigcup_{n \in \mathbb{N}} \lambda_n$$

for $\lambda \in \{I, \leq, \ell, f, \pi\}$. Observe that $x' \in I_\omega$ with $\pi(x') = x$ and therefore $\varphi \in \ell_\omega(x')^-$. Thus $\mathcal{I}_\omega$ falsifies φ. Our construction guarantees that we obtained a Σ-labelled quasimodel that is functional.

Lemma 11. *$\mathcal{I}_\omega$ is a functional Σ-quasimodel falsifying φ.*

Theorem 4. *A formula φ is falsifiable over the class of expanding models if and only if φ is falsifiable over the class of Σ-quasimodels.*

Proof. The left-to-right direction is Lemma 8. For the right-to-left direction, suppose φ is falsifiable over the class of Σ-quasimodels. Hence there exists a Σ-quasimodel $\mathcal{X} = (X, \leq, \ell, R)$ and $x \in X$ with $\varphi \in \ell(x)^-$. By Lemma 11 there exists a functional Σ-quasimodel falsifying φ. So by Lemma 9 there exists an expanding model falsifying φ. Thus φ is falsifiable over the class of expanding models. $\Box$

8 Simulations

A key ingredient in our completeness proof will be to relate worlds in a finite quasimodel to prime theories in the canonical model. Unlike in a filtration quotient, this relation will

[4]We will always close $\leq_{n+1}$ under transitivity and reflexivity and will not mention it in the following items.

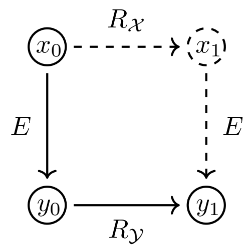

Figure 3: The above diagram can always be completed if $E \subseteq X \times Y$ is a dynamic simulation.

not be a function, but rather given by a simulation, as defined next.

Definition 14. *Let $\Sigma \subseteq \Delta \subseteq \mathcal{L}_{\mathsf{biLTL}}$ be subformula closed, and let $\mathcal{X} = (X, \leq_{\mathcal{X}}, \ell_{\mathcal{X}})$ and $\mathcal{Y} = (Y, \leq_{\mathcal{Y}}, \ell_{\mathcal{Y}})$ be Σ-labelled and Δ-labelled posets respectively. A binary relation $E \subseteq X \times Y$ is a **simulation** if the following hold:*

1. *If $x\ E\ y$, then $\ell_{\mathcal{X}}(x) \subseteq_{\mathrm{T}} \ell_{\mathcal{Y}}(y)$.*
2. *If $x' \geq_{\mathcal{X}} x\ E\ y$, then there exists $y' \in Y$ such that $x'\ E\ y' \geq_{\mathcal{Y}} y$.*
3. *If $x' \leq_{\mathcal{X}} x\ E\ y$, then there exists $y' \in Y$ such that $x'\ E\ y' \leq_{\mathcal{Y}} y$.*

If there exists a simulation E such that $x\ E\ y$, then we write $(\mathcal{X}, x) \rightharpoonup (\mathcal{Y}, y)$.

Lemma 12. *Let $\mathcal{X}$, $\mathcal{Y}$ be labelled systems and $E \subseteq X \times Y$ a simulation. Then $\mathcal{X}\upharpoonright_{E^{-1}[Y]}$ is a labelled system.*

Definition 15. *Let $\mathcal{X} = (X, \leq_{\mathcal{X}}, \ell_{\mathcal{X}}, R_{\mathcal{X}})$ and $\mathcal{Y} = (Y, \leq_{\mathcal{Y}}, \ell_{\mathcal{Y}}, R_{\mathcal{Y}})$ be labelled systems. A **dynamic simulation** between X and Y is a simulation $E \subseteq X \times Y$ satisfying the 'back' condition for R: namely, if $x\ E\ y\ R_{\mathcal{Y}}\ y'$ then there exists x' such that $x\ R_{\mathcal{X}}\ x'\ E\ y'$ (see Figure 3).*

In a more general modal setting, (Fernández-Duque, McLean, and Zenger 2023) construct a 'universal' finite structure, which we denote $\mathbb{U}_\Sigma$, for a given set of formulas Σ, with the property that given *any* Σ-labelled system $\mathcal{M}$, there is a dynamic simulation E_* between $\mathbb{U}_\Sigma$ and $\mathcal{M}$, such that for every world x of $\mathcal{M}$ there exists a world w of $\mathbb{U}_\Sigma$ with $w\ E_*\ x$ and $\ell_{\mathbb{U}}(w) = \ell_{\mathcal{M}}(x)$. We call such an E_* an **exhaustive simulation**. We can then let $\mathcal{M}$ be the canonical model and consider the restriction of $\mathbb{U}_\Sigma$ to the domain of E_*, which by Lemma 12 will be a labelled system. For this, we identify a prime theory Γ with the $\mathcal{L}_{\mathsf{biLTL}}$-type $(\Gamma, \mathcal{L}_{\mathsf{biLTL}} \setminus \Gamma)$; types of this form are **complete**. We obtain the following structure, which plays the role of a 'filtration' in our completeness proof, although we remark that this is not a true filtration in the standard sense, as filtrations do not interact well with confluence properties.

Proposition 1. *Let $\Sigma \subseteq \mathcal{L}_{\mathsf{biLTL}}$ be finite and closed under subformulas. Then there exists a finite, acyclic[5] Σ-labelled system ${}^{\mathcal{M}_{\mathrm{c}}}/_{\Sigma}$ and a total, exhaustive dynamic simulation $E_* \subseteq {}^{\mathcal{M}_{\mathrm{c}}}/_{\Sigma} \times W_{\mathrm{c}}$ (recall that W_{c} is the set of worlds of $\mathcal{M}_{\mathrm{c}}$). Specifically, E_* is the union of all simulations between the two structures.*

In fact, ${}^{\mathcal{M}_{\mathrm{c}}}/_{\Sigma}$ is a quasimodel—a key ingredient in our completeness proof. However, establishing this will require the use of *simulation formulas*, as defined in the next section.

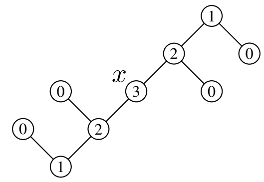

Figure 4: Example x-induced tree $\mathsf{T}(x)$, with heights

9 Simulation Formulas

As before, $\Sigma \subseteq \mathcal{L}_{\mathsf{biLTL}}$ is assumed to be finite and closed under subformulas. Let $\mathcal{X} = (X, \leq_{\mathcal{X}}, \ell_{\mathcal{X}})$ be a Σ-labelled poset, and $x, y \in X$. A **zigzag path** from x to y is a finite sequence $(\rho(i))_{i \leq n}$ of *distinct* worlds, such that $\rho(0) = x$, $\rho(n) = y$, and for all $0 \leq i < n$ either $\rho(i)$ covers $\rho(i+1)$ or $\rho(i+1)$ covers $\rho(i)$. We may also denote a zigzag path by $(\rho(0), \rho(1), \ldots, \rho(n))$. If $\rho = (\rho(i))_{i \leq n}$, the **length** $|\rho|$ of ρ is n. Let

$$\mathsf{ZZP}(x) \coloneqq \{\rho \mid \rho \text{ is a zigzag path starting at } x\}.$$

Definition 16. *Let $\mathcal{X} = (X, \leq_{\mathcal{X}}, \ell_{\mathcal{X}})$ be a finite, acyclic Σ-labelled poset and $x \in X$. The x-**induced tree** is defined as $\mathsf{T}(x) \coloneqq (\mathsf{ZZP}(x), \sqsubset)$, where $\rho \sqsubset \rho'$ if and only if ρ is a proper initial segment of ρ' (see Figure 4).*

Observe that $\mathsf{T}(x)$ is a *finite* tree with the path (x) as *root*. Given $\rho \sqsubseteq \rho'$, we write $\rho' - \rho$ for the suffix of ρ' after ρ. Moreover, we write $\rho \nearrow \rho'$ if each element in $\rho' - \rho$ *covers* its predecessor in ρ', and $\rho \searrow \rho'$ if each element in $\rho' - \rho$ *is covered* by its predecessor in ρ'. For $\rho \in \mathsf{ZZP}(x)$, the **height** of ρ is defined as $h(\rho) \coloneqq \max\{|\rho' - \rho| \mid \rho \sqsubseteq \rho'\}$.

We now define, for x in a finite, acyclic labelled poset, the simulation formulas $\chi^+(x)$ and $\chi^-(x)$, which together encode all worlds accessible from x via a zigzag path. Therefore, satisfying or falsifying $\chi^+(x)$ or $\chi^-(x)$ respectively at some world y of a labelled poset is equivalent to the existence of a simulation involving x and y: see Proposition 2.

We define $\chi^+(x)$ and $\chi^-(x)$ by working 'outside-in', i.e. recursively from the leaves of $\mathsf{T}(x)$ to the root, exploiting the following.

- (i) By *asserting* a formula $\varphi \prec\!\!- \psi$ we can express that there is a world *below* where φ holds and ψ does not.
- (ii) By *denying* a formula $\varphi \to \psi$ we can express that there is a world *above* where φ holds and ψ does not.

We begin by defining for each path ρ in $\mathsf{T}(x)$ different from (x) a formula φ_ρ. The simulation formulas are then composed from these formulas φ_ρ.

Recall that by convention $\bigwedge \emptyset \coloneqq \top$ and $\bigvee \emptyset \coloneqq \bot$.

[5]It is essential that ${}^{\mathcal{M}_{\mathrm{c}}}/_{\Sigma}$ be partially typed in order for it to be both finite and acyclic. Otherwise, the combination of $\to$ and $\prec\!\!-$ can force the existence of infinite zigzag paths which can only be made finite by creating a cycle.

Definition 17. *Let $\mathcal{X} = (X, \leq_{\mathcal{X}}, \ell_{\mathcal{X}})$ be a finite, acyclic Σ-labelled poset, and $x \in X$. For each $\rho = (\rho(0), \ldots, \rho(n)) \in \mathsf{T}(x)$ with $|\rho| > 0$ define the formula φ_ρ by induction on $h(\rho)$. Suppose φ'_ρ has been defined for each ρ' with $h(\rho') < h(\rho)$.*

1. *If $\rho(n-1) >_{\mathcal{X}} \rho(n)$, define*

$$\varphi_\rho := (\bigwedge \ell_{\mathcal{X}}(\rho(n))^+ \wedge \bigwedge_{\rho' \sqsupset \rho \colon \rho \searrow \rho'} \varphi_{\rho'}) \prec (\bigvee \ell_{\mathcal{X}}(\rho(n))^- \vee \bigvee_{\rho' \sqsupset \rho \colon \rho \nearrow \rho'} \varphi_{\rho'})$$

2. *If $\rho(n-1) <_{\mathcal{X}} \rho(n)$, define*

$$\varphi_\rho := (\bigwedge \ell_{\mathcal{X}}(\rho(n))^+ \wedge \bigwedge_{\rho' \sqsupset \rho \colon \rho \searrow \rho'} \varphi_{\rho'}) \to (\bigvee \ell_{\mathcal{X}}(\rho(n))^- \vee \bigvee_{\rho' \sqsupset \rho \colon \rho \nearrow \rho'} \varphi_{\rho'})$$

Then define $\chi^+(x)$ and $\chi^-(x)$ as follows.

$$\chi^+(x) := (\bigwedge \ell_{\mathcal{X}}(x)^+ \wedge \bigwedge_{\rho \sqsupset (x) \colon (x) \searrow \rho} \varphi_\rho) \prec (\bigvee \ell_{\mathcal{X}}(x)^- \vee \bigvee_{\rho \sqsupset (x) \colon (x) \nearrow \rho} \varphi_\rho)$$

$$\chi^-(x) := (\bigwedge \ell_{\mathcal{X}}(x)^+ \wedge \bigwedge_{\rho \sqsupset (x) \colon (x) \searrow \rho} \varphi_\rho) \to (\bigvee \ell_{\mathcal{X}}(x)^- \vee \bigvee_{\rho \sqsupset (x) \colon (x) \nearrow \rho} \varphi_\rho)$$

Recall that $\mathcal{M}_c = (W_c, \leq_c, f_c, V_c)$ is the canonical model.

Proposition 2. *Let $\mathcal{M}_c/\Sigma = (U, \leq, R, \ell)$ and $E_* \subseteq U \times W_c$ be the total, exhaustive simulation provided by Proposition 1. Let $w \in U$ and $\Gamma \in W_c$. The following hold.*

1. *$\chi^+(w) \in \Gamma$ if and only if there exists $\Delta \in W_c$ with $\Delta \leq_c \Gamma$ such that $w \; E_* \; \Delta$.*
2. *$\chi^-(w) \in \mathcal{L}_{\mathsf{biLTL}} \setminus \Gamma$ if and only if there exists $\Delta \in W_c$ with $\Gamma \leq_c \Delta$ such that $w \; E_* \; \Delta$.*

Next we establish some biLTL-derivable properties of χ^+ and χ^-. We begin with the former. These properties are established by using Proposition 2 to see that they are present in *every* prime theory in the canonical model and so derivable. As before, $\mathcal{M}_c/\Sigma = (U, \leq, R, \ell)$; then the reflexive transitive closure of R is denoted R^*.

Proposition 3. *Given $w \in U$ and $\psi \in \Sigma$:*

1. *If $\psi \in \ell(w)^-$, then $\vdash \chi^+(w) \to (\chi^+(w) \prec \psi)$.*
2. *If $\psi \in \ell(w)^+$, then $\vdash \chi^+(w) \to \psi$.*
3. *$\vdash \chi^+(w) \to \bigcirc \bigvee_{wRv} \chi^+(v)$.*

Item 2, for example, follows from the fact that if $\chi^+(w) \in \Gamma$ then $w \; E_* \; \Delta$ for some $\Delta \leq_c \Gamma$, which by the definition of simulations implies that any $\psi \in \ell(w)^+$ must belong to Δ and hence to Γ. Item 3 follows by similar reasoning, using the fact that E_* is dynamic.

The formula χ^- behaves 'dually', as follows.

Proposition 4. *Given $w \in U$ and $\psi \in \Sigma$:*

1. *If $\psi \in \ell(w)^-$, then $\vdash \psi \to \chi^-(w)$.*
2. *If $\psi \in \ell(w)^+$, then $\vdash (\psi \to \chi^-(w)) \to \chi^-(w)$.*
3. *$\vdash \bigcirc \bigwedge_{wRv} \chi^-(v) \to \chi^-(w)$.*

10 Completeness

The simulation formulas $\chi^\pm$ are fundamental in our completeness proof. Specifically, we will use them to show that $\mathcal{M}_c/\Sigma$ is ω-sensible and hence a quasimodel. Since validity over the class of quasimodels is equivalent to validity over the class of expanding models by Theorem 4, completeness will follow. The following lemma is the first step towards establishing ω-sensibility. As above, we write $\mathcal{M}_c/\Sigma = (U, \leq, R, \ell)$ and R^* for the reflexive transitive closure of R, and $E_* \subseteq U \times W_c$ is a total, exhaustive simulation. The following readily follows from Proposition 3 and Proposition 4.

Lemma 13. *If $\Sigma \subseteq \mathcal{L}_{\mathsf{biLTL}}$ is finite and closed under subformulas, and $w \in U$, then:*

1. *$\vdash \bigvee_{wR^*v} \chi^+(v) \to \bigcirc \bigvee_{wR^*v} \chi^+(v)$,*
2. *$\vdash \bigcirc \bigwedge_{wR^*v} \chi^-(v) \to \bigwedge_{wR^*v} \chi^-(v)$.*

In order to complete our proof that $\mathcal{M}_c/\Sigma$ is ω-sensible, it suffices to apply the induction rules $\mathbf{Ind}_\Box$ and $\mathbf{Ind}_\Diamond$ of our calculus to the formulas of Lemma 13.

Proposition 5.

1. *If $w \in U$ and $\Diamond\psi \in \ell(w)^+$, then there exists $v \in R^*(w)$ such that $\psi \in \ell(v)^+$.*
2. *If $w \in U$ and $\Box\psi \in \ell(w)^-$, then there exists $v \in R^*(w)$ such that $\psi \in \ell(v)^-$.*

Proof. We treat only the first item, as the second is analogous, using the respective rules for $\Box$. Towards a contradiction, assume that $w \in U$ and $\Diamond\psi \in \ell(w)^+$, but for all $v \in R^*(w)$, we have $\psi \in \ell(v)^-$. By Lemma 13, $\vdash \bigcirc \bigwedge_{wR^*v} \chi^-(v) \to \bigwedge_{wR^*v} \chi^-(v)$. By the $\mathbf{Ind}_\Diamond$ rule, $\vdash \Diamond \bigwedge_{wR^*v} \chi^-(v) \to \bigwedge_{wR^*v} \chi^-(v)$; in particular,

$$\vdash \Diamond \bigwedge_{wR^*v} \chi^-(v) \to \chi^-(w). \tag{1}$$

Now let $v \in R^*(w)$. By Proposition 4.1 and the assumption that $\psi \in \ell(v)^-$, we have that $\vdash \psi \to \chi^-(v)$, and since v was arbitrary, $\vdash \psi \to \bigwedge_{wR^*v} \chi^-(v)$. Using $\mathbf{Mon}_\Diamond$, we further have that $\vdash \Diamond\psi \to \Diamond \bigwedge_{wR^*v} \chi^-(v)$. This, along with (1), shows that $\vdash \Diamond\psi \to \chi^-(w)$. However, by Proposition 4.2 and our assumption that $\Diamond\psi \in \ell(w)^+$, we have that $\vdash (\Diamond\psi \to \chi^-(w)) \to \chi^-(w)$. Hence by modus ponens we obtain $\vdash \chi^-(w)$. Choosing $\Gamma \in W_c$ such that

$w \; E_* \; \Gamma$, Proposition 2.2 yields $\chi^-(w) \notin \Gamma$, but this contradicts $\vdash \chi^-(w)$. We conclude that there is $v \in R^*(w)$ with $\psi \in \ell(v)^+$, as needed. □

Corollary 1. *If $\Sigma \subseteq \mathcal{L}_{\mathsf{biLTL}}$ is finite and closed under subformulas, then $\mathcal{M}_c/\Sigma$ is a quasimodel.*

Proof. By Proposition 1, $\mathcal{M}_c/\Sigma$ is a labelled system (and serial, since $\mathcal{M}_c$ is), while by Proposition 5, R is ω-sensible. So by Definition 10, $\mathcal{M}_c/\Sigma$ is a quasimodel. □

We are now ready to prove that our calculus is complete.

Theorem 5. *Given $\varphi \in \mathcal{L}_{\mathsf{biLTL}}$, the following are equivalent: (i) $\mathsf{biLTL} \vdash \varphi$, (ii) φ is valid over the class of expanding models, (iii) φ is valid over the class of finite quasimodels.*

Proof. That (i) implies (ii) is Theorem 2 and that (ii) implies (iii) is Theorem 4. We show that (iii) implies (i) by contrapositive. Suppose φ is an unprovable formula and let Σ be the set of subformulas of φ. Since φ is unprovable, there exists $\Gamma \in W_c$ with $\varphi \notin \Gamma$. Since E_* is exhaustive, there is $w \in U$ such that $\varphi \in \ell(w)^-$ and $w \; E_* \; \Gamma$. Hence w is a point in a finite quasimodel falsifying φ. □

Corollary 2. *Derivability in* biLTL *is decidable.*

This follows from the fact that biLTL is axiomatisable and has a finite quasimodel property (with a computable bound on the 'finite'). We also obtain a second, unexpected corollary: $\prec$ cannot be extended to the class of topological models while validating bi-intuitionistic logic, since as verified in (Boudou et al. 2021), this would mean that the class of *dynamic* topological models would validate biLTL, and hence validate **RV**, which we know by Example 4 not to be the case.

Corollary 3. *Suppose that $\prec_{\mathrm{top}}$ assigns to each topological space (X, τ) a binary operation $\tau \times \tau \to \tau$. (Here, τ is the collection of opens.) Consider the semantics that combines standard topological semantics for intuitionistic propositional logic with $\prec_{\mathrm{top}}$ semantics for $\prec$. Then the class of topological spaces does not validate propositional Heyting–Brouwer logic.*

11 Concluding Remarks

We have solved the problem of axiomatising intuitionistic linear temporal logic over Kripke models via a rather unexpected method: by incorporating co-implication into our formal language, the standard axioms automatically become complete. This paves the road for a purely proof-theoretic analysis of intuitionistic temporal logics with multiple natural lines of inquiry, including the existence of cut-free calculi, interpolants, and automated deduction.

Acknowledgements

This work was funded by the SNSF–FWO Lead Agency Grant 200021L 196176 (SNSF)/G0E2121N (FWO). Brett McLean was also funded by FWO Senior Postdoctoral Fellowship 1280024N.

References

Aleksandroff, P. 1937. Diskrete Räume. *Matematicheskii Sbornik* 2:501–518.

Artëmov, S. N.; Davoren, J. M.; and Nerode, A. 1997. Topological semantics for hybrid systems. In Adian, S. I., and Nerode, A., eds., *Logical Foundations of Computer Science, 4th International Symposium, LFCS'97, Yaroslavl, Russia, July 6-12, 1997, Proceedings*, volume 1234 of *Lecture Notes in Computer Science*, 1–8. Springer.

Balbiani, P.; Boudou, J.; Diéguez, M.; and Fernández-Duque, D. 2020. Intuitionistic linear temporal logics. *ACM Trans. Comput. Log.* 21(2):14:1–14:32.

Boudou, J.; Diéguez, M.; Fernández-Duque, D.; and Kremer, P. 2021. Exploring the jungle of intuitionistic temporal logics. *Theory Pract. Log. Program.* 21(4):459–492.

Boudou, J.; Diéguez, M.; and Fernández-Duque, D. 2017. A decidable intuitionistic temporal logic. In *26th EACSL Annual Conference on Computer Science Logic (CSL)*, 14:1–14:17. Stockholm, Sweden: Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.

Boudou, J.; Diéguez, M.; and Fernández-Duque, D. 2022. Complete intuitionistic temporal logics for topological dynamics. *J. Symb. Log.* 87(3):995–1022.

Chopoghloo, S., and Moniri, M. 2021. A strongly complete axiomatization of intuitionistic temporal logic. *J. Log. Comput.* 31(7):1640–1659.

Chopoghloo, S., and Moniri, M. 2022. An infinitary axiomatization of dynamic topological logic. *Log. J. IGPL* 30(1):124–142.

Dalen, D. . 1986. Intuitionistic logic. In *Handbook of Philosophical Logic*, volume 166. Dordrecht: Springer Netherlands. 225–339.

Drobyshevich, S.; Odintsov, S.; and Wansing, H. 2022. Moisil's modal logic and related systems. In Bimbó, K., ed., *Relevance logics and other tools for reasoning*, 150–177. Rickmansworth: College publications.

Fernández-Duque, D.; McLean, B.; and Zenger, L. 2023. A family of decidable bi-intuitionistic modal logics. In Marquis, P.; Son, T. C.; and Kern-Isberner, G., eds., *Proceedings of KR 2023, Rhodes, Greece*, 262–271.

Fernández-Duque, D. 2014. Non-finite axiomatizability of dynamic topological logic. *ACM Trans. Comput. Log.* 15(1):4:1–4:18.

Fernández-Duque, D. 2018. The intuitionistic temporal logic of dynamical systems. *Logical Methods in Computer Science* 14(3):1–35.

Goré, R., and Shillito, I. 2020. Bi-intuitionistic logics: A new instance of an old problem. In *Advances in Modal Logic*.

Howard, W. 1980. The formulas-as-types notion of construction. In Seldin, J. P., and Hindley, J. R., eds., *To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus, and Formalism*. Boston, MA: Academic Press. 479–490.

Kamide, N., and Wansing, H. 2010. Combining linear-time temporal logic with constructiveness and paraconsistency. *Journal of Applied Logic* 8(1):33–61.

Konev, B.; Kontchakov, R.; Wolter, F.; and Zakharyaschev, M. 2006. Dynamic topological logics over spaces with continuous functions. In Governatori, G.; Hodkinson, I. M.; and Venema, Y., eds., *Advances in Modal Logic 6, Noosa, Queensland, Australia*, 299–318. College Publications.

Kremer, P., and Mints, G. 2007. Dynamic topological logic. In Aiello, M.; Pratt-Harman, I.; and van Benthem, J., eds., *Handbook of Spatial Logics*. Springer.

Mints, G. 2000. *A Short Introduction to Intuitionistic Logic*. University Series in Mathematics. Norwell, MA, USA: Kluwer Academic Publishers.

Moisil, G. 1942. *Logique modale*. Bucharest. Institutul Regal de Cercetari Ştiinţifice. Publicaţiile. Imprimeria Naţională.

Pnueli, A. 1977. The temporal logic of programs. In *18th IEEE Symposium on the Foundations of Computer Science*, 46–57. Washington, DC, USA: IEEE Computer Society.

Rauszer, C. 1974a. A formalization of the propositional calculus of H-B logic. *Studia Logica: An International Journal for Symbolic Logic* 33(1):23–34.

Rauszer, C. 1974b. Semi-boolean algebras and their applications to intuitionistic logic with dual operations. *Fundamenta Mathematicae* 83:219–249.

Rauszer, C. 1977. Applications of kripke models to Heyting-Brouwer logic. *Studia Logica: An International Journal for Symbolic Logic* 36(1/2):61–71.