Intuitionistic μ-calculus with the Lewis arrow

TABLEAUX 2025, Reykjavik University

Bahareh Afshari & Lide Grotenhuis

27 Sep 2025

University of Gothenburg & University of Amsterdam

Introduction

Introduction: the modal μ-calculus

Introduction: the modal μ-calculus

Fixpoint logics contain formulas ψ that satisfy

$$\psi \equiv \varphi(\psi).$$

for some formula φ.

Introduction: the modal μ-calculus

Fixpoint logics contain formulas ψ that satisfy

$$\psi \equiv \varphi(\psi).$$

for some formula φ.

Example. "A until B" in temporal logic: $A \mathrel{\mathsf{U}} B \equiv B \vee (A \wedge \mathsf{X}(A \mathrel{\mathsf{U}} B))$.

Introduction: the modal μ-calculus

Fixpoint logics contain formulas ψ that satisfy

$$\psi \equiv \varphi(\psi).$$

for some formula φ.

Example. "A until B" in temporal logic: $A \mathsf{U} B \equiv B \vee (A \wedge \mathsf{X}(A \mathsf{U} B))$.

The modal μ-calculus extends basic modal logic by explicit fixpoint operators μ and ν:

- $\mu X.\varphi(X)$ denotes the least fixpoint of $\varphi(X)$,
- $\nu X.\varphi(X)$ denotes the greatest fixpoint of $\varphi(X)$.

Introduction: the modal μ-calculus

Fixpoint logics contain formulas ψ that satisfy

$$\psi \equiv \varphi(\psi).$$

for some formula φ.

Example. "A until B" in temporal logic: $A \, \mathsf{U} \, B \equiv B \vee (A \wedge \mathsf{X}(A \, \mathsf{U} \, B))$.

The modal μ-calculus extends basic modal logic by explicit fixpoint operators μ and ν:

- $\mu X.\varphi(X)$ denotes the least fixpoint of $\varphi(X)$,
- $\nu X.\varphi(X)$ denotes the greatest fixpoint of $\varphi(X)$.

The resulting logic is very expressive yet well-behaved: decidable, small model property, finitary axiomatization, uniform interpolation...

Introduction: the modal μ-calculus

Fixpoint logics contain formulas ψ that satisfy

$$\psi \equiv \varphi(\psi).$$

for some formula φ.

Example. "A until B" in temporal logic: $A \,\mathsf{U}\, B \equiv B \vee (A \wedge \mathsf{X}(A \,\mathsf{U}\, B))$.

The modal μ-calculus extends basic modal logic by explicit fixpoint operators μ and ν:

- $\mu X.\varphi(X)$ denotes the least fixpoint of $\varphi(X)$,
- $\nu X.\varphi(X)$ denotes the greatest fixpoint of $\varphi(X)$.

The resulting logic is very expressive yet well-behaved: decidable, small model property, finitary axiomatization, uniform interpolation...

Non-wellfounded and cyclic proof systems provide natural syntactic characterisations of the modal μ-calculus and its fragments.

Introduction: intuitionistic fixpoints

Introduction: intuitionistic fixpoints

Modal fixpoint logics over an intuitionistic propositional base are gaining attention:

Modal fixpoint logics over an intuitionistic propositional base are gaining attention:

- Intuitionistic linear-time temporal logic (Balbiani, Boudou, Diégues & Fernández-Duque, 2019, 2022);

Introduction: intuitionistic fixpoints

Modal fixpoint logics over an intuitionistic propositional base are gaining attention:

- Intuitionistic linear-time temporal logic (Balbiani, Boudou, Diégues & Fernández-Duque, 2019, 2022);
- Intuitionistic common-knowledge logic (Jäger & Marti, 2016);

Introduction: intuitionistic fixpoints

Modal fixpoint logics over an intuitionistic propositional base are gaining attention:

- Intuitionistic linear-time temporal logic (Balbiani, Boudou, Diégues & Fernández-Duque, 2019, 2022);
- Intuitionistic common-knowledge logic (Jäger & Marti, 2016);
- Intuitionistic Gödel-Löb logic (Das, van der Giessen, Marin, 2023).

Introduction: intuitionistic fixpoints

Modal fixpoint logics over an intuitionistic propositional base are gaining attention:

- Intuitionistic linear-time temporal logic (Balbiani, Boudou, Diégues & Fernández-Duque, 2019, 2022);
- Intuitionistic common-knowledge logic (Jäger & Marti, 2016);
- Intuitionistic Gödel-Löb logic (Das, van der Giessen, Marin, 2023).

Introduction: intuitionistic fixpoints

Modal fixpoint logics over an intuitionistic propositional base are gaining attention:

- Intuitionistic linear-time temporal logic (Balbiani, Boudou, Diégues & Fernández-Duque, 2019, 2022);
- Intuitionistic common-knowledge logic (Jäger & Marti, 2016);
- Intuitionistic Gödel-Löb logic (Das, van der Giessen, Marin, 2023).

Our aim. Develop a general framework and proof-theoretic techniques for studying intuitionistic modal fixpoint logics.

Introduction: intuitionistic fixpoints

Modal fixpoint logics over an intuitionistic propositional base are gaining attention:

- Intuitionistic linear-time temporal logic (Balbiani, Boudou, Diégues & Fernández-Duque, 2019, 2022);
- Intuitionistic common-knowledge logic (Jäger & Marti, 2016);
- Intuitionistic Gödel-Löb logic (Das, van der Giessen, Marin, 2023).

Our aim. Develop a general framework and proof-theoretic techniques for studying intuitionistic modal fixpoint logics.

In earlier work, we (Afshari, G., Leigh & Zenger) provided proof systems for:

1. intuitionistic linear-time temporal logic (2023);
2. intuitionistic modal logic with the master modality (2024).

Introduction: intuitionistic fixpoints

Modal fixpoint logics over an intuitionistic propositional base are gaining attention:

- Intuitionistic linear-time temporal logic (Balbiani, Boudou, Diégues & Fernández-Duque, 2019, 2022);
- Intuitionistic common-knowledge logic (Jäger & Marti, 2016);
- Intuitionistic Gödel-Löb logic (Das, van der Giessen, Marin, 2023).

Our aim. Develop a general framework and proof-theoretic techniques for studying intuitionistic modal fixpoint logics.

In earlier work, we (Afshari, G., Leigh & Zenger) provided proof systems for:

1. intuitionistic linear-time temporal logic (2023);
2. intuitionistic modal logic with the master modality (2024).

Current work. We study an intuitionistic version of the modal μ-calculus with the Lewis arrow (a generalisation of the modal $\Box$). We provide game semantics and a non-wellfounded analytic proof system.

Introduction: the Lewis arrow

Introduction: the Lewis arrow

Dissatisfied with material implication, Lewis (1914,1932) introduced several axiom systems (S1-S5) meant to formalize strict implication:

$$\varphi \prec\!\!\!\!\!- \psi \qquad\qquad \text{“}\psi \text{ can be inferred from } \varphi\text{.”}$$

Introduction: the Lewis arrow

Dissatisfied with material implication, Lewis (1914,1932) introduced several axiom systems (S1-S5) meant to formalize strict implication:

$$\varphi \strictif \psi \qquad\qquad \text{“}\psi \text{ can be inferred from } \varphi\text{.”}$$

For the modern modal logician:

$$\varphi \strictif \psi := \Box(\varphi \to \psi).$$

So $\Box\varphi \equiv \top \strictif \varphi$.

Dissatisfied with material implication, Lewis (1914,1932) introduced several axiom systems (S1-S5) meant to formalize strict implication:

$\varphi \dashv\!\!3\, \psi$ "ψ can be inferred from φ."

For the modern modal logician:

$$\varphi \dashv\!\!3\, \psi := \Box(\varphi \to \psi).$$

So $\Box\varphi \equiv \top \dashv\!\!3\, \varphi$.

In an intuitionistic setting, $\dashv\!\!3$ is not interdefinable with $\Box$, as was observed in the study of intuitionistic provability logic (Iemhof 2003, Litak & Visser 2017).

Contents

The logic iL_μ

Syntax

Syntax

Fix some set Prop of propositions/variables. Formulas of iL_μ are given by the grammar:

$$\varphi, \psi ::= \bot \mid \top \mid P \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \varphi \to \psi \mid \varphi \multimap \psi \mid \mu X.\varphi \mid \nu X.\varphi$$

with $P, X \in$ Prop and X (weakly) positive in φ. We define $\Box\varphi := \top \multimap \varphi$.

Note: μ and ν are not interdefinable in the intuitionistic setting.

Syntax

Fix some set Prop of propositions/variables. Formulas of iL_μ are given by the grammar:

$$\varphi, \psi ::= \bot \mid \top \mid P \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \varphi \rightarrow \psi \mid \varphi \multimap \psi \mid \mu X.\varphi \mid \nu X.\varphi$$

with $P, X \in$ Prop and X (weakly) positive in φ. We define $\Box\varphi := \top \multimap \varphi$.

Note: μ and ν are not interdefinable in the intuitionistic setting.

We consider formulas φ that are clean: each bound variable X belongs to a unique subformula $\eta X.\psi_X$ of φ.

Syntax

Fix some set Prop of propositions/variables. Formulas of iL_μ are given by the grammar:

$$\varphi, \psi ::= \bot \mid \top \mid P \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \varphi \to \psi \mid \varphi \prec \psi \mid \mu X.\varphi \mid \nu X.\varphi$$

with $P, X \in$ Prop and X (weakly) positive in φ. We define $\Box\varphi := \top \prec \varphi$.

Note: μ and ν are not interdefinable in the intuitionistic setting.

We consider formulas φ that are clean: each bound variable X belongs to a unique subformula $\eta X.\psi_X$ of φ.

Moreover, to keep track of negative/positive formula occurrences, we will consider polarised (sub)formulas φ^p with $p \in \{+, -\}$.

$$Sub((\varphi_1 \star \varphi_2)^p) := \{\varphi_1^{-p}, \varphi_2^p\} \cup Sub(\varphi_1^{-p}) \cup Sub(\varphi_2^p) \qquad \text{if } \star \in \{\to, \prec\}.$$

Relational semantics: bi-relational models

Formulas are evaluated in bi-relational Kripke models $M = (W, \leq, R, V)$, where

1. $\leq$ is a partial order (the intuitionistic relation),
2. $R \subseteq W^2$ (the modal relation),
3. if $w \leq vRu$ then wRu (triangle confluence).

Relational semantics: bi-relational models

Formulas are evaluated in bi-relational Kripke models $M = (W, \leq, R, V)$, where

1. $\leq$ is a partial order (the intuitionistic relation),
2. $R \subseteq W^2$ (the modal relation),
3. if $w \leq vRu$ then wRu (triangle confluence).

The truth relation for $\to$, $\multimap$ and the fixpoint operators is defined by

$$\begin{array}{lll} M, s \models \varphi \to \psi & \text{iff} & \text{for all } t \geq s \text{ if } M, t \models \varphi, \text{ then } M, t \models \psi, \\ M, s \models \varphi \multimap \psi & \text{iff} & \text{for all } sRt \text{ if } M, t \models \varphi, \text{ then } M, t \models \psi, \\ M, s \models \mu X.\varphi & \text{iff} & s \in LFP(\varphi_X^M), \\ M, s \models \nu X.\varphi & \text{iff} & s \in GFP(\varphi_X^M), \end{array}$$

where $\varphi_X^M : \mathcal{P}(W) \to \mathcal{P}(W)$ is the function given by $S \mapsto [\![\varphi]\!]^M_{X \mapsto S}$.

(Algebraic) semantics of iL_μ: monotonicity and confluence properties

A key property of intuitionistic Kripke semantics is monotonicity: if $v \geq w$ and $w \models \varphi$, then $v \models \varphi$.

(Algebraic) semantics of iL_μ: monotonicity and confluence properties

A key property of intuitionistic Kripke semantics is monotonicity: if $v \geq w$ and $w \models \varphi$, then $v \models \varphi$.

To obtain monotonicity for iL_μ, models need to satisfy triangle confluence. For $\Box$-formulas, forth-down confluence suffices.

$$
\begin{tikzcd}
v \arrow[r, "R"] & u \\
w \arrow[u, "\leq"] \arrow[ur, dashed, "R"'] &
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
v \arrow[r, "R"] & u \\
w \arrow[u, "\leq"] \arrow[r, dashed, "R"'] & \cdot \arrow[u, dashed, "\leq"']
\end{tikzcd}
$$

Figure 1: Triangle confluence (left) and Forth-down confluence (right)

(Algebraic) semantics of iL_μ: monotonicity and confluence properties

A key property of intuitionistic Kripke semantics is monotonicity: if $v \geq w$ and $w \models \varphi$, then $v \models \varphi$.

To obtain monotonicity for iL_μ, models need to satisfy triangle confluence. For $\Box$-formulas, forth-down confluence suffices.

$$\begin{tikzcd}
v \arrow[r, "R"] & u \\
w \arrow[u, "\leq"] \arrow[ur, dashed, "R"'] &
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
v \arrow[r, "R"] & u \\
w \arrow[u, "\leq"] \arrow[r, dashed, "R"'] & \cdot \arrow[u, dashed, "\leq"']
\end{tikzcd}$$

Figure 1: Triangle confluence (left) and Forth-down confluence (right)

Lemma

A $\Box$-formula φ is valid on all forth-down confluent models iff it is valid on all triangle confluent models.

(Algebraic) semantics of iL_μ: monotonicity and confluence properties

A key property of intuitionistic Kripke semantics is monotonicity: if $v \geq w$ and $w \models \varphi$, then $v \models \varphi$.

To obtain monotonicity for iL_μ, models need to satisfy triangle confluence. For $\Box$-formulas, forth-down confluence suffices.

$$\begin{tikzcd}
v \arrow[r, "R"] & u \\
w \arrow[u, "\leq"] \arrow[ur, dashed, "R"'] &
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
v \arrow[r, "R"] & u \\
w \arrow[u, "\leq"] \arrow[r, dashed, "R"'] & \cdot \arrow[u, dashed, "\leq"']
\end{tikzcd}$$

Figure 1: Triangle confluence (left) and Forth-down confluence (right)

Lemma

A $\Box$-formula φ is valid on all forth-down confluent models iff it is valid on all triangle confluent models.

As $-\!\!\!\!\ni$-formulas are not monotone for the weaker condition, we obtain that $-\!\!\!\!\ni$ indeed cannot be expressed in terms of $\Box$.

Game semantics for iL_μ

Game semantics for iL_μ: the evaluation game

Given a model $M = (W, \leq, R, V)$ and clean formula ψ, we define an evaluation game $\mathcal{E}(\varphi, M)$ between $\exists$ and $\forall$.

Given a model $M = (W, \leq, R, V)$ and clean formula ψ, we define an evaluation game $\mathcal{E}(\varphi, M)$ between $\exists$ and $\forall$.

Intuition: at position (ψ^+, s), $\exists$ wants to show that $M, s \models \psi$, while $\forall$ wants to show the converse.

Game semantics for iL_μ: the evaluation game

Given a model $M = (W, \leq, R, V)$ and clean formula ψ, we define an evaluation game $\mathcal{E}(\varphi, M)$ between $\exists$ and $\forall$.

Intuition: at position (ψ^+, s), $\exists$ wants to show that $M, s \models \psi$, while $\forall$ wants to show the converse.

Position	Player	Admissible moves
(P^+, s), $P \notin BV(\psi)$, $s \in V(P)$	$\forall$	$\emptyset$
(P^+, s), $P \notin BV(\psi)$, $s \notin V(P)$	$\exists$	$\emptyset$
$(\varphi_1 \wedge \varphi_2^+, s)$	$\forall$	$\{(\varphi_i^+, s) : i = 1, 2\}$
$(\varphi_1 \vee \varphi_2^+, s)$	$\exists$	$\{(\varphi_i^+, s) : i = 1, 2\}$
$(\varphi_1 \to \varphi_2^+, s)$	$\forall$	$\{(\varphi_1 \to \varphi_2^+, s, t) : s \leq t\}$
$(\varphi_1 \to \varphi_2^+, s, t)$	$\exists$	$\{(\varphi_1^-, t), (\varphi_2^+, t)\}$
$(\varphi_1 \prec \varphi_2^+, s)$	$\forall$	$\{(\varphi_1 \prec \varphi_2^+, s, t) : sRt\}$
$(\varphi_1 \prec \varphi_2^+, s, t)$	$\exists$	$\{(\varphi_1^-, t), (\varphi_2^+, t)\}$
$((\eta X.\psi_X)^p, s)$	-	$\{(\psi_X^p, s)\}$
(X^p, s), $X \in BV(\psi)$	-	$\{(\psi_X^p, s)\}$

Game semantics for iL_μ: the evaluation game

Given a model $M = (W, \leq, R, V)$ and clean formula ψ, we define an evaluation game $\mathcal{E}(\varphi, M)$ between $\exists$ and $\forall$.

Intuition: at position (ψ^+, s), $\exists$ wants to show that $M, s \models \psi$, while $\forall$ wants to show the converse.

Position	Player	Admissible moves
(P^+, s), $P \notin BV(\psi)$, $s \in V(P)$	$\forall$	$\emptyset$
(P^+, s), $P \notin BV(\psi)$, $s \notin V(P)$	$\exists$	$\emptyset$
$(\varphi_1 \wedge \varphi_2^+, s)$	$\forall$	$\{(\varphi_i^+, s) : i = 1, 2\}$
$(\varphi_1 \vee \varphi_2^+, s)$	$\exists$	$\{(\varphi_i^+, s) : i = 1, 2\}$
$(\varphi_1 \to \varphi_2^+, s)$	$\forall$	$\{(\varphi_1 \to \varphi_2^+, s, t) : s \leq t\}$
$(\varphi_1 \to \varphi_2^+, s, t)$	$\exists$	$\{(\varphi_1^-, t), (\varphi_2^+, t)\}$
$(\varphi_1 \multimap \varphi_2^+, s)$	$\forall$	$\{(\varphi_1 \multimap \varphi_2^+, s, t) : sRt\}$
$(\varphi_1 \multimap \varphi_2^+, s, t)$	$\exists$	$\{(\varphi_1^-, t), (\varphi_2^+, t)\}$
$((\eta X.\psi_X)^p, s)$	-	$\{(\psi_X^p, s)\}$
(X^p, s), $X \in BV(\psi)$	-	$\{(\psi_X^p, s)\}$

For negative positions (ψ^-, s) swap the roles of $\exists$ and $\forall$.

Game semantics for iL_μ: winning conditions and adequacy

Game semantics for iL_μ: winning conditions and adequacy

We write $\mathcal{E}(\varphi, M)@q$ for the evaluation game with starting position q. A play of $\mathcal{E}(\varphi, M)@q$ is either infinite or ends in a position with no admissible moves. Finite plays are lost by the player who got stuck.

Game semantics for iL_μ: winning conditions and adequacy

We write $\mathcal{E}(\varphi, M)@q$ for the evaluation game with starting position q. A play of $\mathcal{E}(\varphi, M)@q$ is either infinite or ends in a position with no admissible moves. Finite plays are lost by the player who got stuck.

Who wins an infinite play?

Lemma

Let π an infinite play of $\mathcal{E}(\varphi, M)@(\varphi^+, s)$. Then there is a unique, outermost $X_\pi \in BV(\varphi)$ occurring infinitely often in π. Moreover, there is a unique polarity p_π such that $X_\pi^{p_\pi}$ occurs infinitely often in π.

We write $\mathcal{E}(\varphi, M)@q$ for the evaluation game with starting position q. A play of $\mathcal{E}(\varphi, M)@q$ is either infinite or ends in a position with no admissible moves. Finite plays are lost by the player who got stuck.

Who wins an infinite play?

Lemma

Let π an infinite play of $\mathcal{E}(\varphi, M)@(\varphi^+, s)$. Then there is a unique, outermost $X_\pi \in BV(\varphi)$ occurring infinitely often in π. Moreover, there is a unique polarity p_π such that $X_\pi^{p_\pi}$ occurs infinitely often in π.

Recall that every bound variable is bound by either μ or ν. The infinite play π is won by $\exists$ iff X_π is a negative μ-variable or a positive ν-variable.

Game semantics for iL_μ: winning conditions and adequacy

We write $\mathcal{E}(\varphi, M)@q$ for the evaluation game with starting position q. A play of $\mathcal{E}(\varphi, M)@q$ is either infinite or ends in a position with no admissible moves. Finite plays are lost by the player who got stuck.

Who wins an infinite play?

Lemma

Let π an infinite play of $\mathcal{E}(\varphi, M)@(\varphi^+, s)$. Then there is a unique, outermost $X_\pi \in BV(\varphi)$ occurring infinitely often in π. Moreover, there is a unique polarity p_π such that $X_\pi^{p_\pi}$ occurs infinitely often in π.

Recall that every bound variable is bound by either μ or ν. The infinite play π is won by $\exists$ iff X_π is a negative μ-variable or a positive ν-variable.

Theorem (Adequacy of the Game Semantics)

For any clean formula φ and pointed model (M, s), we have

$$M, s \models \varphi \text{ iff } \exists \text{ has a (positional) winning strategy in } \mathcal{E}(\varphi, M)@(\varphi^+, s).$$

Guardedness

Guardedness

We call a variable X guarded in φ if every occurrence of X in φ is in the scope of some $\mathrel{-\!\!3}$-operator. A formula φ is guarded if for every subformula $\eta X.\psi$ of φ, X is guarded in ψ.

Guardedness

We call a variable X guarded in φ if every occurrence of X in φ is in the scope of some $\strictif$-operator. A formula φ is guarded if for every subformula $\eta X.\psi$ of φ, X is guarded in ψ.

Theorem

Every iL_μ*-formula* φ *is equivalent to a guarded one.*

We call a variable X guarded in φ if every occurrence of X in φ is in the scope of some $\multimap$-operator. A formula φ is guarded if for every subformula $\eta X.\psi$ of φ, X is guarded in ψ.

Theorem

Every iL_μ*-formula* φ *is equivalent to a guarded one.*

Proof sketch: By induction on φ. For the fixpoint case $\eta X.\psi$, use Ruitenburg's theorem for IPC:

Theorem (Ruitenburg, 1984)

Let φ *be a formula of* IPC and X a propositional letter such that X is positive in φ. Define $\varphi^0_X := X$ and $\varphi^{n+1}_X := \varphi[\varphi^n_X/X]$. Then there exists an N such that $\varphi^N_X \equiv \varphi^{N+1}_X$.

A non-wellfounded proof system for iL_μ

A non-wellfounded proof system for iL_μ: the propositional rules

A non-wellfounded proof system for iL_μ: the propositional rules

We define a sequent as a finite set of polarised formulas. We let $\Gamma \Rightarrow \Delta$ denote $\{\varphi^+ : \varphi \in \Gamma\} \cup \{\varphi^- : \varphi \in \Delta\}$ and its interpretation is given by

$$\bigwedge \Gamma \to \bigvee \Delta.$$

A non-wellfounded proof system for iL_μ: the propositional rules

We define a sequent as a finite set of polarised formulas. We let $\Gamma \Rightarrow \Delta$ denote $\{\varphi^+ : \varphi \in \Gamma\} \cup \{\varphi^- : \varphi \in \Delta\}$ and its interpretation is given by

$$\bigwedge \Gamma \to \bigvee \Delta.$$

For the propositional rules, we use standard multi-conclusion rules for IPC.

$$\frac{}{\Gamma, A \Rightarrow A, \Delta}\ \mathrm{id} \qquad \frac{}{\Gamma, \bot \Rightarrow \Delta}\ \bot$$

$$\frac{\Gamma, A, B \Rightarrow \Delta}{\Gamma, A \wedge B \Rightarrow \Delta}\ \wedge\mathrm{L} \qquad \frac{\Gamma \Rightarrow A, \Delta \quad \Gamma \Rightarrow B, \Delta}{\Gamma \Rightarrow A \wedge B, \Delta}\ \wedge\mathrm{R}$$

$$\frac{\Gamma, A \Rightarrow \Delta \quad \Gamma, B \Rightarrow \Delta}{\Gamma, A \vee B \Rightarrow \Delta}\ \vee\mathrm{L} \qquad \frac{\Gamma \Rightarrow A, B, \Delta}{\Gamma \Rightarrow A \vee B, \Delta}\ \vee\mathrm{R}$$

$$\frac{\Gamma, A \to B \Rightarrow A, \Delta \quad \Gamma, B \Rightarrow \Delta}{\Gamma, A \to B \Rightarrow \Delta}\ {\to}\mathrm{L} \qquad \frac{\Gamma, A \Rightarrow B}{\Gamma \Rightarrow A \to B, \Delta}\ {\to}\mathrm{R}$$

A non-wellfounded proof system for iL_μ: the modal rule

Consider the following sound rule for the modality $\multimap$:

$$\dfrac{A \Rightarrow B, C \qquad D, A \Rightarrow B}{\Gamma, C \multimap D \Rightarrow A \multimap B, \Delta}\ \multimap_1$$

A non-wellfounded proof system for iL_μ: the modal rule

Consider the following sound rule for the modality $\strictif$:

$$\frac{A \Rightarrow B, C \qquad D, A \Rightarrow B}{\Gamma, C \strictif D \Rightarrow A \strictif B, \Delta}\ \strictif_1$$

For completeness, we generalize it to the following:

$$\frac{\{\mathcal{D}_j, A \Rightarrow B, \mathcal{C}_j\}_{j \leq 2^n}}{\Gamma, \{C_i \strictif D_i\}_{i \leq n} \Rightarrow A \strictif B, \Delta}\ \strictif_n$$

where $n \geq 0$, and the sets $\mathcal{D}_1, \ldots, \mathcal{D}_{2^n}$ and $\mathcal{C}_1, \ldots, \mathcal{C}_{2^n}$ enumerate the subsets of $\{D_1, \ldots, D_n\}$ and $\{C_1, \ldots, C_n\}$, respectively, such that

$$D_i \in \mathcal{D}_j \text{ if and only if } C_i \notin \mathcal{C}_j.$$

A non-wellfounded proof system for iL_μ: the fixpoint rules

A non-wellfounded proof system for iL_μ: the fixpoint rules

For $\eta \in \{\mu, \nu\}$, we have the fixpoint rules:

$$\frac{\Gamma, \psi \Rightarrow \Delta}{\Gamma, \eta X.\psi \Rightarrow \Delta}\ \eta\mathsf{L} \quad \frac{\Gamma \Rightarrow \psi, \Delta}{\Gamma \Rightarrow \eta X.\psi, \Delta}\ \eta\mathsf{R}$$

$$\frac{\Gamma, \psi_X \Rightarrow \Delta}{\Gamma, X \Rightarrow \Delta}\ X\mathsf{L} \qquad \frac{\Gamma \Rightarrow \psi_X, \Delta}{\Gamma \Rightarrow X, \Delta}\ X\mathsf{R}$$

We work in the context of a clean formula φ, so each bound variable $X \in BV(\varphi)$ has an associated fixpoint formula ψ_X.

A non-wellfounded proof system for iL_μ: the fixpoint rules

For $\eta \in \{\mu, \nu\}$, we have the fixpoint rules:

$$\frac{\Gamma, \psi \Rightarrow \Delta}{\Gamma, \eta X.\psi \Rightarrow \Delta}\ \eta\mathsf{L} \qquad \frac{\Gamma \Rightarrow \psi, \Delta}{\Gamma \Rightarrow \eta X.\psi, \Delta}\ \eta\mathsf{R}$$

$$\frac{\Gamma, \psi_X \Rightarrow \Delta}{\Gamma, X \Rightarrow \Delta}\ X\mathsf{L} \qquad \frac{\Gamma \Rightarrow \psi_X, \Delta}{\Gamma \Rightarrow X, \Delta}\ X\mathsf{R}$$

We work in the context of a clean formula φ, so each bound variable $X \in BV(\varphi)$ has an associated fixpoint formula ψ_X.

This concludes the rules of nwIL_μ.

A derivation T in nwIL_μ is a finite or infinite tree labelled according to the rules of nwIL_μ.

A non-wellfounded proof system for iL_μ: derivations and proofs

A derivation T in nwIL_μ is a finite or infinite tree labelled according to the rules of nwIL_μ.

Given a path ρ through T, a trace on ρ is a sequence $(\varphi_i^{p_i})_i$ of polarised formulas following the principal-residual relation.

$$\Gamma', \psi' \Rightarrow \Delta'$$

$$\Gamma, \psi \Rightarrow \Delta$$

A non-wellfounded proof system for iL_μ: derivations and proofs

A derivation T in nwIL_μ is a finite or infinite tree labelled according to the rules of nwIL_μ.

Given a path ρ through T, a trace on ρ is a sequence $(\varphi_i^{p_i})_i$ of polarised formulas following the principal-residual relation.

$$\Gamma', \psi' \Rightarrow \Delta'$$

$$\Gamma, \psi \Rightarrow \Delta$$

Each (non-stagnating) trace has a unique outermost bound variable X that occurs infinitely often and has a well-defined polarity.

A non-wellfounded proof system for iL_μ: derivations and proofs

A derivation T in nwIL_μ is a finite or infinite tree labelled according to the rules of nwIL_μ.

Given a path ρ through T, a trace on ρ is a sequence $(\varphi_i^{p_i})_i$ of polarised formulas following the principal-residual relation.

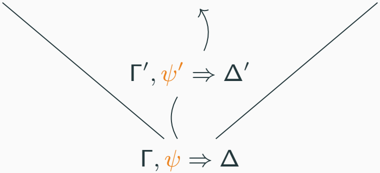

Each (non-stagnating) trace has a unique outermost bound variable X that occurs infinitely often and has a well-defined polarity.

A derivation is a proof in nwIL_μ if every infinite path of T has either a negative μ-trace or a positive ν-trace.

A non-wellfounded proof system for iL_μ: soundness and completeness

Theorem

If φ is provable in nwIL_μ then it is valid on triangle models.

A non-wellfounded proof system for iL_μ: soundness and completeness

Theorem

If φ is provable in nwIL_μ then it is valid on triangle models.

Theorem

Every guarded formula valid on triangle models is provable in nwIL_μ.

A non-wellfounded proof system for iL_μ: soundness and completeness

Theorem

If φ is provable in nwIL_μ then it is valid on triangle models.

Theorem

Every guarded formula valid on triangle models is provable in nwIL_μ.

Proof sketch:

- Given a sequent σ, construct an ω-regular validity game between Prover and Refuter (à la Niwiński & Walukiewicz, 1996).

A non-wellfounded proof system for iL_μ: soundness and completeness

Theorem

If φ is provable in nwIL_μ then it is valid on triangle models.

Theorem

Every guarded formula valid on triangle models is provable in nwIL_μ.

Proof sketch:

- Given a sequent σ, construct an ω-regular validity game between Prover and Refuter (à la Niwiński & Walukiewicz, 1996).
- A winning strategy for Prover corresponds to a (regular) proof of σ.

A non-wellfounded proof system for iL_μ: soundness and completeness

Theorem

If φ is provable in nwIL_μ then it is valid on triangle models.

Theorem

Every guarded formula valid on triangle models is provable in nwIL_μ.

Proof sketch:

- Given a sequent σ, construct an ω-regular validity game between Prover and Refuter (à la Niwiński & Walukiewicz, 1996).
- A winning strategy for Prover corresponds to a (regular) proof of σ.
- A winning strategy for Refuter induces a (pre)countermodel M for σ.

A non-wellfounded proof system for iL_μ: soundness and completeness

Theorem

If φ is provable in $nwIL_\mu$ then it is valid on triangle models.

Theorem

Every guarded formula valid on triangle models is provable in $nwIL_\mu$.

Proof sketch:

- Given a sequent σ, construct an ω-regular validity game between Prover and Refuter (à la Niwiński & Walukiewicz, 1996).
- A winning strategy for Prover corresponds to a (regular) proof of σ.
- A winning strategy for Refuter induces a (pre)countermodel M for σ.
- We make M satisfy triangle confluence by replacing the modal relation R by the composition $\leq; R$. This does not break monotonicity of the valuation nor falsification of φ in M.

Conclusion

Conclusion

Conclusion

The logic iL_μ is an intuitionistic version of the modal μ-calculus with an expressive universal modality that has many desirable properties:

Conclusion

The logic iL_μ is an intuitionistic version of the modal μ-calculus with an expressive universal modality that has many desirable properties:

- Elegant game semantics;

Conclusion

The logic iL_μ is an intuitionistic version of the modal μ-calculus with an expressive universal modality that has many desirable properties:

- Elegant game semantics;
- Guardedness property;

Conclusion

The logic iL_μ is an intuitionistic version of the modal μ-calculus with an expressive universal modality that has many desirable properties:

- Elegant game semantics;
- Guardedness property;
- A guarded complete, cut-free, non-wellfounded proof system;

Conclusion

The logic iL_μ is an intuitionistic version of the modal μ-calculus with an expressive universal modality that has many desirable properties:

- Elegant game semantics;
- Guardedness property;
- A guarded complete, cut-free, non-wellfounded proof system;
- Regular proofs and decidability (as the validity game is ω-regular).

Conclusion

The logic iL_μ is an intuitionistic version of the modal μ-calculus with an expressive universal modality that has many desirable properties:

- Elegant game semantics;
- Guardedness property;
- A guarded complete, cut-free, non-wellfounded proof system;
- Regular proofs and decidability (as the validity game is ω-regular).

Conclusion

The logic iL_μ is an intuitionistic version of the modal μ-calculus with an expressive universal modality that has many desirable properties:

- Elegant game semantics;
- Guardedness property;
- A guarded complete, cut-free, non-wellfounded proof system;
- Regular proofs and decidability (as the validity game is ω-regular).

The latter also means we can obtain a cyclic proof system with a local soundness condition.

Conclusion

The logic iL_μ is an intuitionistic version of the modal μ-calculus with an expressive universal modality that has many desirable properties:

- Elegant game semantics;
- Guardedness property;
- A guarded complete, cut-free, non-wellfounded proof system;
- Regular proofs and decidability (as the validity game is ω-regular).

The latter also means we can obtain a cyclic proof system with a local soundness condition.

Directions for future work:

Conclusion

The logic iL_μ is an intuitionistic version of the modal μ-calculus with an expressive universal modality that has many desirable properties:

- Elegant game semantics;
- Guardedness property;
- A guarded complete, cut-free, non-wellfounded proof system;
- Regular proofs and decidability (as the validity game is ω-regular).

The latter also means we can obtain a cyclic proof system with a local soundness condition.

Directions for future work:

- Completeness for unguarded formulas;

Conclusion

The logic iL_μ is an intuitionistic version of the modal μ-calculus with an expressive universal modality that has many desirable properties:

- Elegant game semantics;
- Guardedness property;
- A guarded complete, cut-free, non-wellfounded proof system;
- Regular proofs and decidability (as the validity game is ω-regular).

The latter also means we can obtain a cyclic proof system with a local soundness condition.

Directions for future work:

- Completeness for unguarded formulas;
- Interpolation (might be difficult with the Lewis arrow);

Conclusion

The logic iL_μ is an intuitionistic version of the modal μ-calculus with an expressive universal modality that has many desirable properties:

- Elegant game semantics;
- Guardedness property;
- A guarded complete, cut-free, non-wellfounded proof system;
- Regular proofs and decidability (as the validity game is ω-regular).

The latter also means we can obtain a cyclic proof system with a local soundness condition.

Directions for future work:

- Completeness for unguarded formulas;
- Interpolation (might be difficult with the Lewis arrow);
- A finite Hilbert axiomatisation?

Conclusion

The logic iL_μ is an intuitionistic version of the modal μ-calculus with an expressive universal modality that has many desirable properties:

- Elegant game semantics;
- Guardedness property;
- A guarded complete, cut-free, non-wellfounded proof system;
- Regular proofs and decidability (as the validity game is ω-regular).

The latter also means we can obtain a cyclic proof system with a local soundness condition.

Directions for future work:

- Completeness for unguarded formulas;
- Interpolation (might be difficult with the Lewis arrow);
- A finite Hilbert axiomatisation?
- Adding diamonds?

Conclusion

The logic iL_μ is an intuitionistic version of the modal μ-calculus with an expressive universal modality that has many desirable properties:

- Elegant game semantics;
- Guardedness property;
- A guarded complete, cut-free, non-wellfounded proof system;
- Regular proofs and decidability (as the validity game is ω-regular).

The latter also means we can obtain a cyclic proof system with a local soundness condition.

Directions for future work:

- Completeness for unguarded formulas;
- Interpolation (might be difficult with the Lewis arrow);
- A finite Hilbert axiomatisation?
- Adding diamonds?

References

- Balbiani, P., Boudou, J., Diéguez, M. & Fernández-Duque, D. (2019). Intuitionistic linear temporal logics.
- Das, A., van der Giessen, I. & Marin, S. (2023). Intuitionistic Gödel-Löb logic, à la Simpson: Labelled systems and birelational semantics.
- Iemhoff, R., De Jongh, D., & Zhou, C. (2005). Properties of intuitionistic provability and preservativity logics.
- Jäger, G. & M. Marti (2016). Intuitionistic common knowledge or belief.
- Litak, T., & Visser, A. (2018). Lewis meets Brouwer: constructive strict implication.
- Ruitenburg, W. (1984). On the period of sequences (an(p)) in intuitionistic propositional calculus.
- Venema, Y. (2024). Lectures on the modal μ-calculus.