

Unified Classical Logic Completeness

A Coinductive Pearl

Jasmin Christian Blanchette¹, Andrei Popescu^{1,2}, and Dmitriy Traytel¹

¹ Fakultät für Informatik, Technische Universität München, Germany

² Institute of Mathematics Simion Stoilow of the Romanian Academy, Bucharest, Romania

Abstract. Codatatypes are absent from many programming and specification languages. We make a case for their importance by revisiting a classical result: the completeness theorem for first-order logic established through a Gentzen system. The core of the proof establishes an abstract property of possibly infinite derivation trees, independently of the concrete syntax or inference rules. This separation of concerns simplifies the presentation. The abstract proof can be instantiated for a wide range of Gentzen and tableau systems as well as various flavors of first-order logic. The corresponding Isabelle/HOL formalization demonstrates the recently introduced support for codatatypes and the Haskell code generator.

1 Introduction

Gödel’s completeness theorem [12] is a major result about first-order logic (FOL). It forms the basis of results and techniques in various areas, including mathematical logic, automated deduction, and program verification. It can be stated as follows: If a set of formulas is syntactically consistent (i.e., no contradiction arises from it), then it has a model. The theorem enjoys many accounts in the literature that generalize and simplify the original proof; indeed, a textbook on mathematical logic would be incomplete without a proof of this fundamental theorem.

Formal logic has always been a battleground between semantic and syntactic methods. Generally, mathematicians belong to the semantic school, whereas computer scientists tend to take the other side of the argument. The completeness theorem, which combines syntax and semantics, is also disputed, with the result that each school has its own proof. In his review of Gallier’s *Logic for Computer Science* [11], Pfenning notes the following [29]:

All too often, proof-theoretic methods are neglected in favor of shorter, and superficially more elegant semantic arguments. [In contrast, in Gallier’s book] the treatment of the proof theory of the Gentzen system is oriented towards computation with proofs. For example, a pseudo-Pascal version of a complete search procedure for first-order cut-free Gentzen proofs is presented.

In the context of completeness, the “superficially more elegant semantic arguments” are proofs that rely on Hilbert systems. These systems have several axioms but only one or two deduction rules, providing minimal support for presenting the structure of proofs or for modeling proof search. A proof of completeness based on Hilbert systems follows the *Henkin style*: It employs a heavy bureaucratic apparatus to establish facts about deduction and conservative language extensions, culminating in a nonconstructive step:

an application of Zorn’s lemma to extend any syntactically consistent set of formulas to a maximally consistent one, from which a model is produced.

In contrast, a proof of completeness based on more elaborate Gentzen or tableau systems follows the *Beth–Hintikka style* [20]. It performs a search that builds either a finite deduction tree yielding a proof (or refutation, depending on the system) or an infinite tree from which a countermodel (or model) can be extracted. Such completeness proofs have an intuitive content that stresses the tension of the argument: The deduction system systematically tries to prove the goal; a failure yields, at the limit, a countermodel.

The intuitive appeal of the Beth–Hintikka approach comes at a price: It requires reasoning about infinite derivation trees and infinite paths. Unfortunately, convenient means to reason about infinite (or lazy) data structures are lacking in mainstream mathematics. In textbooks, at best the trees are defined rigorously (e.g., as prefix-closed sets), but the reasoning relies on the intuitive notion of trees, as Gallier does. One could argue that trees are intuitive and do not need a formal treatment, but the same holds for the syntax of formulas, which is treated very rigorously in most of the textbooks.

This paper presents a rigorous Beth–Hintikka-style proof of the completeness theorem, based on a Gentzen system. The potentially infinite trees are captured by codatatypes (also called coinductive datatypes or final coalgebras) [18]. Another novel aspect of the proof is its modularity: The core tree construction argument is isolated from the proof system and concrete formula syntax (Section 3). The abstract proof can be instantiated for a wide range of Gentzen and tableau systems as well as various flavors of FOL (Sections 4 and 5). This modularization replaces the textbook proofs by analogy. The core of the argument amounts to reasoning about a lazy functional program.

The proof is formalized in Isabelle/HOL [28] (Section 6). The tree construction makes use of a new definitional package for codatatypes [5], which automates the derivation of characteristic theorems from specifications of the constructors. Through Isabelle’s code generator [14], the corecursive construction gives rise to a Haskell program that implements a semidecision procedure for validity instantiable with various proof systems, yielding verified sound and complete provers.

Conventions. Isabelle/HOL is a proof assistant based on classical higher-order logic (HOL) with Hilbert choice, the axiom of infinity, and rank-1 polymorphism. It is the logic of Gordon’s original HOL system and of its many successors [13]. HOL notations are a mixture of functional programming and mathematics. We refer to Nipkow and Klein [27] for a modern introduction. In this paper, the logic is viewed as a framework for expressing mathematics, much like set theory is employed by working mathematicians. In keeping with the standard semantics of HOL, types α are identified with sets.

2 A Gentzen System for First-Order Logic

We fix a first-order language: a countably infinite set var of variables x, y, z and countable sets fsym and psym of function symbols f and predicate symbols p together with an assignment $\text{ar} : \text{fsym} \uplus \text{psym} \rightarrow \text{nat}$ of numeric arities. Terms $t \in \text{term}$ are symbolic expressions built inductively from variables by application of function symbols $f \in \text{fsym}$ to tuples of arguments whose lengths respect the arities: $f(t_1, \dots, t_{\text{ar } f})$. Atoms $a \in \text{atom}$ are expressions of the form $p(t_1, \dots, t_{\text{ar } p})$, where $p \in \text{psym}$ and $t_1, \dots, t_{\text{ar } p} \in \text{term}$.

Formulas φ, ψ may be atoms, negations, conjunctions, or universal quantifications. They are defined as follows:

datatype fmla = Atm atom | Neg fmla | Conj fmla fmla | All var fmla

A structure $\mathcal{S} = (S, (F_f)_{f \in \text{fsym}}, (P_p)_{p \in \text{psym}})$ for the given language consists of a carrier set S , together with a function $F_f : S^n \rightarrow S$ for each n -ary $f \in \text{fsym}$ and a predicate $P_p : S^n \rightarrow \text{bool}$ for each n -ary $p \in \text{psym}$. The notions of interpretation of a term t and satisfaction of a formula φ by a structure $\mathcal{S}$ with respect to a variable valuation $\xi : \text{var} \rightarrow S$ are defined in the standard way. For terms:

$$\llbracket x \rrbracket_{\xi}^{\mathcal{S}} = \xi x \quad \llbracket f(t_1, \dots, t_n) \rrbracket_{\xi}^{\mathcal{S}} = F_f(\llbracket t_1 \rrbracket_{\xi}^{\mathcal{S}}, \dots, \llbracket t_n \rrbracket_{\xi}^{\mathcal{S}})$$

For atoms: $\mathcal{S} \models_{\xi} p(t_1, \dots, t_n)$ iff $P_p(\llbracket t_1 \rrbracket_{\xi}^{\mathcal{S}}, \dots, \llbracket t_n \rrbracket_{\xi}^{\mathcal{S}})$. For formulas:

$$\begin{aligned} \mathcal{S} \models_{\xi} \text{Atm } a & \text{ iff } \mathcal{S} \models_{\xi} a & \mathcal{S} \models_{\xi} \text{Conj } \varphi \psi & \text{ iff } \mathcal{S} \models_{\xi} \varphi \wedge \mathcal{S} \models_{\xi} \psi \\ \mathcal{S} \models_{\xi} \text{Neg } \varphi & \text{ iff } \mathcal{S} \not\models_{\xi} \varphi & \mathcal{S} \models_{\xi} \text{All } x \varphi & \text{ iff } \forall a \in S. \mathcal{S} \models_{\xi[x \leftarrow a]} \varphi \end{aligned}$$

The following substitution lemma relates the notions of satisfaction and capture-avoiding substitution $\varphi[t/x]$ of a term t for a variable x in a formula φ :

Lemma 1. $\mathcal{S} \models_{\xi} \varphi[t/x]$ iff $\mathcal{S} \models_{\xi[x \leftarrow \llbracket t \rrbracket_{\xi}^{\mathcal{S}}]} \varphi$.

A sequent is a pair $\Gamma \triangleright \Delta$ of finite formula sets. Satisfaction is extended to sequents: $\mathcal{S} \models_{\xi} \Gamma \triangleright \Delta$ iff $(\forall \varphi \in \Gamma. \mathcal{S} \models_{\xi} \varphi) \Rightarrow (\exists \psi \in \Delta. \mathcal{S} \models_{\xi} \psi)$.

The proof system on sequents is defined inductively as follows, where the notation Γ, φ abbreviates the set $\Gamma \cup \{\varphi\}$:

$$\begin{array}{c} \frac{}{\Gamma, \text{Atm } a \triangleright \Delta, \text{Atm } a} \text{AX} \quad \frac{\Gamma \triangleright \Delta, \varphi}{\Gamma, \text{Neg } \varphi \triangleright \Delta} \text{NEGL} \quad \frac{\Gamma, \varphi \triangleright \Delta}{\Gamma \triangleright \Delta, \text{Neg } \varphi} \text{NEGR} \\[10pt] \frac{\Gamma, \varphi, \psi \triangleright \Delta}{\Gamma, \text{Conj } \varphi \psi \triangleright \Delta} \text{CONJL} \quad \frac{\Gamma \triangleright \Delta, \varphi \quad \Gamma \triangleright \Delta, \psi}{\Gamma \triangleright \Delta, \text{Conj } \varphi \psi} \text{CONJR} \\[10pt] \frac{\Gamma, \text{All } x \varphi, \varphi[t/x] \triangleright \Delta}{\Gamma, \text{All } x \varphi \triangleright \Delta} \text{ALLL} \quad \frac{\Gamma \triangleright \Delta, \varphi[y/x]}{\Gamma \triangleright \Delta, \text{All } x \varphi} \text{ALLR} \quad (y \text{ fresh}) \end{array}$$

The rules are applied from bottom to top. One chooses a formula from either side of the sequent, the eigenformula, and applies a rule according to the topmost connective or quantifier. For a given choice of eigenformula, at most one rule is applicable. The aim of applying the rules is to prove the sequent by building a finite derivation tree whose branches are closed by an axiom (AX). The completeness theorem states that any sequent $\Gamma \triangleright \Delta$ either is provable (denoted by $\vdash$) or has a countermodel, i.e., a structure $\mathcal{S}$ and a valuation ξ that falsify it: $\vdash \Gamma \triangleright \Delta \vee (\exists \mathcal{S}, \xi. \mathcal{S} \not\models_{\xi} \Gamma \triangleright \Delta)$.

3 Abstract Completeness

The proof of the completeness theorem is divided in two parts. The first part, described in this section, focuses on the core of the completeness argument in an abstract, syntax-free manner. This level captures the tension between the existence of a proof or of an abstract notion of countermodel; the latter is introduced through what we call an escape

path—an infinite sequence of rule applications that “escapes” the proof attempt. The tension is distilled in a completeness result: Either there exists a finite derivation tree or there exists an infinite derivation tree with a suitable escape path. The second part maps the abstract escape path to a concrete, proof-system-specific countermodel. Section 4 performs this connection for the Gentzen system presented in Section 2.

Rule Systems. We abstract away the syntax of formulas and sequents and the specific rules of the proof system. We fix countable sets state and rule for states and rules. A state represents a formal statement in the logic. We assume that the meaning of the rules is given by an effect relation $\text{eff} : \text{rule} \rightarrow \text{state} \rightarrow \text{state fset} \rightarrow \text{bool}$, where $\alpha \text{ fset}$ denotes the set of finite subsets of α . The reading of $\text{eff } r \ s \ ss$ is as follows: Starting from state s , applying rule r expands s into the states ss . The triple $\mathcal{R} = (\text{state}, \text{rule}, \text{eff})$ forms a *rule system*.

Example 1. The Gentzen system from Section 2 can be presented as a rule system. The set state is the set of sequents, and rule consists of the following: a rule Ax_a for each atom a ; rules NEGL_φ and NEGR_φ for each formula φ ; rules $\text{CONJL}_{\varphi,\psi}$ and $\text{CONJR}_{\varphi,\psi}$ for each pair of formulas φ and ψ ; a rule $\text{ALLL}_{x,\varphi,t}$ for each variable x , formula φ , and term t ; and a rule $\text{ALLR}_{x,\varphi}$ for each variable x and formula φ .

The eigenformula is part of the rule. Hence we have a countably infinite number of rules. The effect is defined as follows, where semicolons (;) separate set elements:

$\text{eff } \text{Ax}_a \ (\Gamma, \text{Atm } a \triangleright \Delta, \text{Atm } a) \ \emptyset$
 $\text{eff } \text{NEGR}_\varphi \ (\Gamma \triangleright \Delta, \text{Neg } \varphi) \ \{\Gamma, \varphi \triangleright \Delta\}$
 $\text{eff } \text{NEGL}_\varphi \ (\Gamma, \text{Neg } \varphi \triangleright \Delta) \ \{\Gamma \triangleright \Delta, \varphi\}$
 $\text{eff } \text{CONJL}_{\varphi,\psi} \ (\Gamma, \text{Conj } \varphi \ \psi \triangleright \Delta) \ \{\Gamma, \varphi, \psi \triangleright \Delta\}$
 $\text{eff } \text{CONJR}_{\varphi,\psi} \ (\Gamma \triangleright \Delta, \text{Conj } \varphi \ \psi) \ \{\Gamma \triangleright \Delta, \varphi; \Gamma \triangleright \Delta, \psi\}$
 $\text{eff } \text{ALLL}_{x,\varphi,t} \ (\Gamma, \text{All } x \ \varphi \triangleright \Delta) \ \{\Gamma, \text{All } x \ \varphi, \varphi[t/x] \triangleright \Delta\}$
 $\text{eff } \text{ALLR}_{x,\varphi} \ (\Gamma \triangleright \Delta, \text{All } x \ \varphi) \ \{\Gamma \triangleright \Delta, \varphi[y/x]\}$ where y is fresh for Γ and $\text{All } x \ \varphi$

Derivation Trees. Possibly infinite trees are represented by the following codatatype:

codatatype $\alpha \text{ tree} = \text{Node} \ (\text{lab} : \alpha) \ (\text{sub} : (\alpha \text{ tree}) \text{ fset})$

This definition introduces a constructor $\text{Node} : \alpha \rightarrow (\alpha \text{ tree}) \text{ fset} \rightarrow \alpha \text{ tree}$ and two selectors $\text{lab} : \alpha \text{ tree} \rightarrow \alpha$, $\text{sub} : \alpha \text{ tree} \rightarrow (\alpha \text{ tree}) \text{ fset}$. Trees have the form $\text{Node } a \ Ts$, where a is the tree’s *label* and Ts is the finite set of its (immediate) *subtrees*. The **codatatype** keyword indicates that, unlike for inductive datatypes, this tree formation rule may be applied an infinite number of times.

A *step* combines the current state and the rule to be applied: $\text{step} = \text{state} \times \text{rule}$. Derivation trees are defined as trees labeled by steps, $\text{dtree} = \text{step tree}$, in which the root’s label (s, r) represents the proved goal s and the first (backward) applied rule r . The well-formed derivation trees are captured by the predicate $\text{wf} : \text{dtree} \rightarrow \text{bool}$ defined by the coinductive rule

$$\frac{\text{eff } r \ s \ (\text{image } (\text{fst} \circ \text{lab}) \ Ts) \quad \forall T \in Ts. \ \text{wf } T}{\text{wf } (\text{Node } (s, r) \ Ts)} \text{WF}$$

(Double lines distinguish coinductive rules from their inductive counterparts.) Thus, the predicate wf is the greatest (weakest) solution to

$$\text{wf}(\text{Node}(s, r) Ts) \Leftrightarrow \text{eff } r \text{ s } (\text{image}(\text{fst} \circ \text{lab}) Ts) \wedge (\forall T \in Ts. \text{wf } T)$$

The term $\text{image } f A$ denotes the image of set A through function f , and fst is the left projection operator (i.e., $\text{fst}(x, y) = x$).

The first assumption requires that the rule r from the root be applied to obtain the subtrees' labels. The second assumption requires that wellformedness hold for the immediate subtrees. The coinductive nature of the definition ensures that these properties hold for arbitrarily deep subtrees of T , even if T has infinite paths.

Proofs. The finite derivation trees—the trees that would result from an inductive datatype definition with the same constructors—can be carved out of the codatatype dtree using the predicate finite defined inductively (i.e., as a least fixpoint) by the rule

$$\frac{\forall T \in Ts. \text{finite } T}{\text{finite}(\text{Node}(s, r) Ts)} \text{FIN}$$

A *proof* of a state s is a finite well-formed derivation tree with the state s at its root. An infinite well-formed derivation tree represents a failed proof attempt.

Example 2. Given the instantiation of Example 1, Figure 1 shows a finite derivation tree for the sequent $\text{All } x (p(x)) \triangleright \text{Conj } (p(y)) (p(z))$ written using the familiar syntax for logical symbols. Figure 2 shows an infinite tree for the same sequent.

Escape Paths. An infinite path in a derivation tree can be regarded as a way to “escape” the proof. To represent infinite paths independently of trees, we introduce the codatatype of streams over a type α with the constructor SCons and the selectors shead and stail :

$$\text{codatatype } \alpha \text{ stream} = \text{SCons}(\text{shead}: \alpha) (\text{stail}: \alpha \text{ stream})$$

$$\frac{\frac{\frac{}{\forall x. p(x), p(y) \triangleright p(y)} \text{AX}_{p(y)}}{\forall x. p(x) \triangleright p(y)} \text{ALLL}_{x, p(x), y} \quad \frac{\frac{\frac{}{\forall x. p(x), p(z) \triangleright p(z)} \text{AX}_{p(z)}}{\forall x. p(x) \triangleright p(z)} \text{ALLL}_{x, p(x), z}}{\forall x. p(x) \triangleright p(y) \wedge p(z)} \text{CONJR}_{p(y), p(z)}$$

Fig. 1. A proof

$$\frac{\frac{\frac{}{\forall x. p(x), p(y) \triangleright p(y)} \text{AX}_{p(y)}}{\forall x. p(x) \triangleright p(y)} \text{ALLL}_{x, p(x), y} \quad \frac{\frac{\frac{}{\forall x. p(x), p(y) \triangleright p(z)} \text{ALLL}_{x, p(x), y}}{\forall x. p(x) \triangleright p(z)} \text{ALLL}_{x, p(x), y}}{\forall x. p(x) \triangleright p(y) \wedge p(z)} \text{CONJR}_{p(y), p(z)}$$

Fig. 2. A failed proof attempt

The coinductive predicate $\text{ipath} : \text{dtree} \rightarrow \text{step stream} \rightarrow \text{bool}$ ascertains whether a stream of steps is an infinite path in a tree:

$$\frac{T \in Ts \quad \text{ipath } T \sigma}{\text{ipath } (\text{Node } (s, r) Ts) (\text{SCons } (s, r) \sigma)} \text{IPATH}$$

An *escape path* is a stream of steps that can form an infinite path in a derivation tree. It is defined coinductively as the predicate $\text{epath} : \text{step stream} \rightarrow \text{bool}$, which requires that every element in the given stream be obtained by applying an existing rule and choosing one of the resulting states:

$$\frac{\text{eff } r \ s \ ss \quad s' \in ss \quad \text{epath } (\text{SCons } (s', r') \sigma)}{\text{epath } (\text{SCons } (s, r) (\text{SCons } (s', r') \sigma))} \text{EPATH}$$

The following lemma is easy to prove by coinduction.

Lemma 2. *For any stream σ and tree T , if $\text{wf } T$ and $\text{ipath } \sigma \ T$, then $\text{epath } \sigma$.*

Example 3. The stream

$$(\forall x. p(x) \triangleright p(y) \wedge p(z)) \cdot (\forall x. p(x) \triangleright p(z)) \cdot (\forall x. p(x), p(y) \triangleright p(z))^\infty$$

where $s \cdot \sigma = \text{SCons } s \ \sigma$ and $s^\infty = s \cdot s \cdot \dots$ is an escape path for the tree of Figure 2.

Since the trees are finitely branching, König's lemma applies. Its proof allows us to study a first simple corecursive definition.

Lemma 3. *If the tree T is infinite, there exists an infinite path σ in T .*

Proof. By the contrapositive of FIN, if $\text{Node } (s, r) \ Ts$ is infinite, there exists an infinite subtree $T \in Ts$. Let $f : \{T \in \text{dtree} \mid \neg \text{finite } T\} \rightarrow \{T \in \text{dtree} \mid \neg \text{finite } T\}$ be a function witnessing this fact—i.e., $f \ T$ is an immediate infinite subtree of T . The desired infinite path $p : \{T \in \text{dtree} \mid \neg \text{finite } T\} \rightarrow \text{step stream}$ can be defined by primitive corecursion over the codatatype of streams: $p \ T = \text{SCons } (\text{lab } T) (p \ (f \ T))$. The predicate $\text{ipath } (p \ T) \ T$ holds by straightforward coinduction on the definition of ipath . $\square$

Countermodel Paths. A countermodel path is a structure that witnesses the unprovability of a state s . Any escape path starting in s is a candidate for a countermodel path, given that it indicates a way to apply the proof rules without reaching any result. For it to be a genuine countermodel path, all possible proofs must have been attempted. More specifically, whenever a rule becomes enabled along the escape path, it is eventually applied later in the sequence. For FOL with sequents as states, such paths can be used to produce actual countermodels by interpreting as true (resp. false) all statements made along the path on the left (resp. right) of the sequents.

A rule r is *enabled* in a state s if it has an effect (i.e., $\exists ss. \text{eff } r \ s \ ss$). This is written $\text{enabled } r \ s$. For any rule r , stream σ , and predicate $P : \alpha \text{ stream} \rightarrow \text{bool}$:

- $\text{taken}_r \ \sigma$ iff r is taken at the start of the stream (i.e., $\text{shead } \sigma = (s, r)$ for some s);
- $\text{enabledAt}_r \ \sigma$ iff r is enabled at the beginning of the stream (i.e., if $\text{shead } \sigma = (s, r')$, then $\text{enabled } r \ s$);

$$\begin{array}{c}
\vdots \\
\hline
\forall x. p(x), p(t_1), p(t_2), p(t_3) \triangleright q(y) \quad \text{ALLL}_{x,p(x),t_4} \\
\hline
\forall x. p(x), p(t_1), p(t_2) \triangleright q(y) \quad \text{ALLL}_{x,p(x),t_3} \\
\hline
\forall x. p(x), p(t_1) \triangleright q(y) \quad \text{ALLL}_{x,p(x),t_2} \\
\hline
\forall x. p(x) \triangleright q(y) \quad \text{ALLL}_{x,p(x),t_1}
\end{array}$$

Fig. 3. A derivation tree with a countermodel path

- $\text{ev } P \sigma$ (“eventually P ”) iff P is true for some suffix of σ ;
- $\text{alw } P \sigma$ (“always P ”) iff P is true for all suffixes of σ .

A stream of steps σ is *saturated* if, at each point, any enabled rule is taken at a later point: $\forall r \in \text{rule}. \text{alw } (\lambda \sigma'. \text{enabledAt}_r \sigma' \Rightarrow \text{ev taken}_r \sigma') \sigma$. A *countermodel path* for a state s is a saturated escape path σ starting at s (i.e., $\text{shead } \sigma = (s, r)$ for some r).

Example 4. The escape path given in Example 3 is not saturated, because the rule $\text{ALLL}_{x,p(x),z}$ is enabled starting from the first position but never taken.

Example 5. The escape path associated with the tree of Figure 3 is a countermodel path for $\forall x. p(x) \triangleright q(y)$, assuming that each possible term occurs infinitely often in the sequence $t_1, t_2, \dots$. The only enabled rules along the path are of the form $\text{ALLL}_{x,p(x),-}$, and each is always eventually taken.

Completeness. For the proof of completeness, we assume that the set of rules satisfies the following properties:

- *Availability:* At each state, at least one rule is enabled (i.e., $\forall s. \exists r. \text{enabled } r \ s$).
- *Persistence:* At each state, if a rule is enabled but not taken, it remains enabled (i.e., $\forall s, r, r', s', ss. \text{enabled } r' \ s \wedge r' \neq r \wedge \text{eff } r \ s \ ss \wedge s' \in \text{set } ss \Rightarrow \text{enabled } r' \ s'$).

(We will later remove the first condition with Theorem 6.) The above conditions are local properties of the rules’ effect, not global properties of the proof system. This makes them easy to verify for particular systems.

Saturation is a stronger condition than the standard properties of fairness and justice [10]. Fairness would require the rules to be continuously enabled to guarantee that they are eventually taken. The property of justice is stronger in that it would require the rules to be enabled infinitely often, but not necessarily continuously. Saturation goes further: If a rule is ever enabled, it will certainly be chosen at a later point. Saturation may seem too strong for the task at hand; however, in the presence of persistence, the notions of fairness, justice, and saturation all coincide.

Theorem 4. *Given a rule system that fulfills availability and persistence, every state admits a proof or a countermodel path.*

Proof. The proof uses the following combinators:

- $\text{stake} : \alpha \text{ stream} \rightarrow \text{nat} \rightarrow \alpha \text{ list}$ maps ρ and n to the list of the first n elements of ρ ;
- $\text{smap} : (\alpha \rightarrow \beta) \rightarrow \alpha \text{ stream} \rightarrow \beta \text{ stream}$ maps f to every element of the stream;

- $\text{nats} : \text{nat stream}$ denotes the stream of natural numbers: $0 \cdot 1 \cdot 2 \cdot 3 \cdot \dots$;
- $\text{flat} : (\alpha \text{ list}) \text{ stream} \rightarrow \alpha \text{ stream}$ maps a stream of finite nonempty lists to the stream obtained by concatenating those lists;
- $\text{sdropWhile} : (\alpha \rightarrow \text{bool}) \rightarrow \alpha \text{ stream} \rightarrow \alpha \text{ stream}$ removes the maximal prefix of elements that fulfill a given predicate from a given stream (or returns an irrelevant default value if the predicate holds for all elements of the stream).

We start by constructing a stream of rules fenum in a fair fashion, so that every rule occurs infinitely often in fenum . Let enum be a stream such that its elements cover the entire set rule (which is required to be countable). Take $\text{fenum} = \text{flat} (\text{smap} (\text{stake enum}) (\text{stail nats}))$. Thus, if $\text{enum} = r_1 \cdot r_2 \cdot r_3 \cdot \dots$, then $\text{fenum} = r_1 \cdot r_1 \cdot r_1 \cdot r_2 \cdot r_1 \cdot r_2 \cdot r_3 \cdot \dots$.

Let s be a state. Using fenum , we build a derivation tree T_0 labeled with s such that all its infinite paths are saturated. Let fair be the subset of rule stream consisting of the fair streams. Clearly, any suffix of an element in fair also belongs to fair . In particular, fenum and all its suffixes belong to fair . Given $\rho \in \text{fair}$ and $s \in \text{state}$, $\text{sdropWhile} (\lambda r. \neg \text{enabled } r \ s) \ \rho$ has the form $\text{SCons } r \ \rho'$, making r the first enabled rule in ρ . Such a rule exists because, by availability, at least one rule is enabled in s and, by fairness, all the rules occur in ρ . Since $\text{enabled } r \ s$, we can pick a state set ss such that $\text{eff } r \ s \ ss$. We define $\text{mkTree} : \text{fair} \rightarrow \text{state} \rightarrow \text{dtree}$ corecursively as $\text{mkTree } \rho \ s = \text{Node } (s, r) (\text{image } (\text{mkTree } \rho') \ ss)$.

We prove that, for all $\rho \in \text{fair}$ and s , the derivation tree $\text{mkTree } \rho \ s$ is well formed and all its infinite paths are saturated. Wellformedness is obvious because at each point the continuation is built starting with the effect of a rule. For saturation, we show that if rule r is enabled at state s and $\text{ipath } (\text{mkTree } \rho \ s) \ \sigma$, then r appears along σ (i.e., there exists a state s' such that (s', r) is in σ). This follows by induction on the position of r in ρ , $\text{pos } r \ \rho$ —formally, the length of the shortest list ρ_0 such that $\rho = \rho_0 @ \text{SCons } r \ _$, where $@$ denotes concatenation. Let r' be the first rule from ρ enabled at state s . If $r = r'$, then $\text{mkTree } \rho \ s$ has label (s, r) already. Otherwise, ρ has the form $\rho_1 @ [r'] @ \rho'$, with r not in ρ_1 , hence $\text{pos } r \ \rho' < \text{pos } r \ \rho$. From the definitions of ipath and mkTree , it follows that $\text{ipath } (\text{mkTree } \rho' \ s') (\text{stail } \sigma)$ holds for some $s' \in ss$ such that $\text{eff } r \ s' \ ss$. By the induction hypothesis, r appears along $\text{stail } \sigma$, hence along σ as desired. In particular, $T_0 = \text{mkTree } \text{fenum} \ s$ is well formed and all its infinite paths are saturated.

Finally, if T_0 is finite, it is the desired finite derivation tree. Otherwise, by Lemma 3 (König) it has an infinite path. This path is necessarily saturated; by Lemma 2, it is the desired countermodel path. $\square$

Theorem 4 captures the abstract essence of arguments from the literature, although this is sometimes hard to grasp under the thick forest of syntactic details and concrete strategies for fair enumeration: A fair tree is constructed, which attempts a proof; in case of failure, the tree exhibits a saturated escape path.

If we are not interested in witnessing the proof attempt closely, Theorem 4 can be established more directly by building the fair path without going through an intermediate fair tree. The key observation is that if a state s has no proof and $\text{eff } r \ s \ ss$, there must exist a state $s' \in ss$ that has no proof. (Otherwise, we could compose the proofs of all s' into a proof of s by applying rule r .) Let $\text{pick } r \ s \ ss$ denote such an s' . We proceed directly to the construction of a saturated escape path as a corecursive predicate

$\text{mkPath} : \text{fair} \rightarrow \{s \in \text{state}. s \text{ has no proof}\} \rightarrow \text{step stream}$ following the same idea as for the previous tree construction: $\text{mkPath } \rho s = \text{SCons } (s, r) (\text{mkPath } \rho' (\text{pick } r s ss))$, where again $\text{SCons } r \rho' = \text{sdropWhile } (\lambda r. \neg \text{enabled } r s) \rho$ and ss is such that $\text{eff } r s ss$. Fairness of $\text{mkPath } \rho s$ follows by a similar argument as before for fairness of the tree.

Omitting the Availability Assumption. The above result assumes availability and persistence. Among these assumptions, persistence is essential: It ensures that the constructed fair path is saturated, meaning that every rule available at any point is eventually applied. Availability can be added later to the system without affecting its behavior by introducing a special “idle” rule.

Lemma 5. *A rule system $\mathcal{R} = (\text{state}, \text{rule}, \text{eff})$ that fulfills persistence can be transformed into an equivalent rule system $\mathcal{R}_{\text{idle}} = (\text{state}, \text{rule}_{\text{idle}}, \text{eff}_{\text{idle}})$ that fulfills both persistence and availability, with $\text{rule}_{\text{idle}} = \text{rule} \cup \{\text{IDLE}\}$ and eff_{idle} behaving like eff on rule and $\text{eff}_{\text{idle}} \text{IDLE } s ss \Leftrightarrow ss = \{s\}$.*

Proof. Availability for the modified system follows from the continuous enabledness of IDLE. Persistence follows from the persistence of the original system together with the property that IDLE is continuously enabled and does not alter the state. The modified system is equivalent to the original one because IDLE does not alter the state. $\square$

Theorem 6. *Given a rule system $\mathcal{R}$ that fulfills persistence, every state admits a proof over $\mathcal{R}$ or a countermodel path over $\mathcal{R}_{\text{idle}}$.*

Proof. We first apply Theorem 4 to the system $\mathcal{R}_{\text{idle}}$ to obtain that every state admits either a proof or a countermodel path, both in this system. And since $\mathcal{R}$ and $\mathcal{R}_{\text{idle}}$ are equivalent, any proof of $\mathcal{R}_{\text{idle}}$ yields one of $\mathcal{R}$. $\square$

4 Concrete Completeness

The abstract completeness proof is parameterized by a rule system. This section concretizes the result for the Gentzen system from Section 2 to derive the standard completeness theorem. Example 1 recast it as a rule system; we must verify that it fulfills persistence and interpret abstract countermodel paths as actual FOL countermodels.

The Gentzen rules are persistent because they preserve the context surrounding the eigenformulas. For example, an application of AX_a (which affects only the atom a) leaves any potential enabledness of $\text{ALL}_{x,\varphi,t}$ (which affects only formulas with All at the top) unchanged; moreover, AX_a does not overlap with AX_b for $a \neq b$. A minor subtlety concerns $\text{ALL}_{x,\varphi}$, which requires the existence of a fresh y ; but since the sequents are finite, we can always find a fresh variable in the infinite set var .

On the other hand, availability does not hold for the proof system; for example, the sequent $p(x) \triangleright q(x)$ has no enabled rule. Hence, we need Theorem 6 and its IDLE rule.

Lemma 7. *If $\Gamma \triangleright \Delta$ admits a countermodel path, there exist a structure $\mathcal{S}$ and a valuation $\xi : \text{var} \rightarrow S$ such that $\mathcal{S} \not\models_{\xi} \Gamma \triangleright \Delta$.*

Proof. Let σ be a countermodel path for $\Gamma \triangleright \Delta$ (i.e., a saturated escape path with $\Gamma \triangleright \Delta$ as the first state). Let $\bar{\Gamma}$ be the union of the left-hand sides of sequents occurring in σ ,

and let $\tilde{\Delta}$ be the union of the corresponding right-hand sides. Clearly, $\Gamma \subseteq \tilde{\Gamma}$ and $\Delta \subseteq \tilde{\Delta}$. The pair $(\tilde{\Gamma}, \tilde{\Delta})$ can be shown to be well behaved with respect to all the connectives and quantifiers in the following sense:

1. For all atoms a , $\text{Atm } a \notin \tilde{\Gamma} \cap \tilde{\Delta}$.
2. If $\text{Neg } \varphi \in \tilde{\Gamma}$, then $\varphi \in \tilde{\Delta}$.
3. If $\text{Neg } \varphi \in \tilde{\Delta}$, then $\varphi \in \tilde{\Gamma}$.
4. If $\text{Conj } \varphi \psi \in \tilde{\Gamma}$, then $\varphi \in \tilde{\Gamma}$ and $\psi \in \tilde{\Gamma}$.
5. If $\text{Conj } \varphi \psi \in \tilde{\Delta}$, then $\varphi \in \tilde{\Delta}$ or $\psi \in \tilde{\Delta}$.
6. If $\text{All } x \varphi \in \tilde{\Gamma}$, then $\varphi[t/x] \in \tilde{\Gamma}$ for all t .
7. If $\text{All } x \varphi \in \tilde{\Delta}$, there exists a variable y such that $\varphi[y/x] \in \tilde{\Delta}$.

These properties follow from the saturation of σ with respect to the corresponding rules. The proofs are routine. For example, if $\text{All } x \varphi \in \tilde{\Gamma}$ and t is a term, $\text{ALLL}_{x,\varphi,t}$ is enabled in σ and hence eventually taken, ensuring that $\varphi[t/x] \in \tilde{\Gamma}$.

We construct the concrete (Herbrand) countermodel $\mathcal{S} = (S, F, P)$ as follows. Let the domain S be the set term, and let ξ be the embedding of variables into terms. For each n -ary f and p and each $t_1, \dots, t_n \in S$, we define $F_f(t_1, \dots, t_n) = f(t_1, \dots, t_n)$ and $P_p(t_1, \dots, t_n) \Leftrightarrow p(t_1, \dots, t_n) \in \tilde{\Gamma}$.

To prove $\mathcal{S} \not\models_{\xi} \Gamma \triangleright \Delta$, it suffices to show that $\forall \varphi \in \tilde{\Gamma}. \mathcal{S} \models_{\xi} \varphi$ and $\forall \varphi \in \tilde{\Delta}. \mathcal{S} \not\models_{\xi} \varphi$. These two facts follow together by induction on the depth of φ . In the base case, if $\text{Atm } a \in \tilde{\Gamma}$, then $\mathcal{S} \models_{\xi} \text{Atm } a$ follows directly from the definition of $\mathcal{S}$; moreover, if $\text{Atm } a \in \tilde{\Delta}$, then by property 1 $\text{Atm } a \notin \tilde{\Gamma}$, hence again $\mathcal{S} \not\models_{\xi} \text{Atm } a$ follows from the definition of $\mathcal{S}$. The only nontrivial inductive case is All , which requires Lemma 1 (substitution). Assume $\text{All } x \varphi \in \tilde{\Gamma}$. By property 6, we have $\varphi[t/x] \in \tilde{\Gamma}$ for any t . Hence, by the induction hypothesis, $\mathcal{S} \models_{\xi} \varphi[t/x]$. By Lemma 1, $\mathcal{S} \models_{\xi[x \mapsto t]} \varphi$ for all t ; that is, $\mathcal{S} \models_{\xi} \text{All } x \varphi$. The second fact, concerning $\tilde{\Delta}$, follows similarly from property 7. $\square$

Theorem 8. *For any sequent $\Gamma \triangleright \Delta$, we have $\vdash \Gamma \triangleright \Delta \vee (\exists \mathcal{S}, \xi. \mathcal{S} \not\models_{\xi} \Gamma \triangleright \Delta)$.*

Proof. From Theorem 6 and Lemma 7. $\square$

The rule ALLL from Section 2 stores, in the left context, a copy of the universal formula $\text{All } x \varphi$ when applied backward. This is crucial for concrete completeness since a fair enumeration should try all the t instances of the universally quantified variable x , which requires availability of $\text{All } x \varphi$ even after its use. If we labeled ALLL as $\text{ALLL}_{x,\varphi}$ instead of $\text{ALLL}_{x,\varphi,t}$, thereby delegating the choice of t to the nondeterminism of eff , the system would still be persistent as required by the abstract completeness proof, but Lemma 7 (and hence concrete completeness) would not hold—more specifically, property 6 from the lemma’s proof would fail.

5 Further Concrete Instances

Theorem 6 is applicable to classical FOL Gentzen systems from the literature, in several variants: with sequent components represented as lists, multisets or sets, one-sided or two-sided, and so on. This includes the systems G' , GCNF' , G , and $G_{=}$ from Gallier [11] and the systems $G1$, $G2$, $G3$, $GS1$, $GS2$, and $GS3$ from Troelstra and Schwichtenberg [37]. Persistence is easy to check. The syntax-independent part of the argument is provided by Theorem 6, while an ad hoc step analogous to Lemma 7 is required to build a concrete countermodel.

Several FOL refutation systems based on tableaux or resolution are instances of the abstract theorem, providing that we read the abstract notion of “proof” as “refutation” and “countermodel” as “model.” Nondestructive tableaux [15]—including those presented in Bell and Machover [1] and in Fitting [9]—are usually persistent when regarded as derivation systems. After an application of Theorem 6, the argument for interpreting the abstract model is similar to that for Gentzen systems (Lemma 7).

Regrettably, abstract completeness is not directly applicable beyond classical logic. It is generally not clear how to extract a specific model from a nonstandard logic from an abstract (proof-theoretic) model. Another issue is that standard sequent systems for nonclassical variations of FOL such as modal or intuitionistic logics do not satisfy persistence. A typical right rule for the modal operator $\Box$ (“must”) is as follows [37]:

$$\frac{\Box \Gamma \triangleright \Diamond \Delta, \varphi}{\Box \Gamma \triangleright \Diamond \Delta, \Box \varphi} \text{ MUSTR}$$

To be applicable, the rule requires that all the formulas in the context surrounding the eigenformula have $\Box$ or $\Diamond$ at the top. Other rules may remove these operators, or introduce formulas that do not have them, thus disabling MUSTR.

Recent work targeted at simplifying completeness arguments [26] organizes modal logics as labeled transition systems, for which Kripke completeness is derived. In the proposed systems, the above rule becomes

$$\frac{\Gamma, w R w' \triangleright \Delta, w' : \varphi}{\Gamma \triangleright \Delta, w : \Box \varphi} \text{ MUSTR}' \quad (w' \text{ fresh})$$

The use of labels for worlds (w, w') and the bookkeeping of the accessibility relation R makes it possible to recast the rule so that only resilient facts are ever assumed about the context. The resulting proof system satisfies persistence, enabling Theorem 6. The Kripke countermodel construction is roughly as for classical FOL Gentzen systems.

6 Formalization and Implementation

The definitions, lemmas, and theorems presented in Sections 2 to 4 are formalized in the proof assistant Isabelle/HOL. The instantiation step of Section 4 is formalized for a richer version of FOL, with sorts and interpreted equality, as required by our motivating application (efficient encodings of sorts in unsorted FOL [4]). The formal development is publicly available [6, 7].

The necessary codatatypes and corecursive definitions are realized using a recently introduced definitional package [5]. The tree codatatype illustrates the support for corecursion through permutative data structures (with non-free constructors) such as finite sets, a feature that is not available in any other proof assistant.

For generating code, we make the additional assumption that the effect relation corresponds to a partial function $\text{eff}' : \text{rule} \rightarrow \text{state} \rightarrow (\text{state fset}) \text{ option}$, where the Isabelle datatype $\alpha \text{ option}$ enriches a copy of α with a special value `None`.³ From this function,

³ In the proof system from Example 1, eff is not deterministic due to the rule ALLR. It can be made deterministic by refining the rule with a systematic choice of the fresh variable y .

we build the relational eff as the partial function’s graph. Isabelle’s code generator [14] can then produce Haskell code for the computable part of our completeness proof—the abstract prover `mkTree`, defined corecursively in the proof of Theorem 4:

```
data Stream a = SCons a (Stream a)
newtype FSet a = FSet [a]
data Tree a = Node a (FSet (Tree a))

fmap f (FSet xs) = FSet (map f xs)

sdropWhile p (SCons a  $\sigma$ ) =
  if p a then sdropWhile p  $\sigma$  else SCons a  $\sigma$ 

mkTree eff  $\rho$  s =
  Node (s, r) (fmap (mkTree eff  $\rho'$ ) (fromJust (eff r s)))
  where SCons r  $\rho'$  = sdropWhile ( $\backslash r \rightarrow$  not (isJust (eff r s)))  $\rho$ 
```

Finite sets are represented as lists. The functions `isJust` : α option $\rightarrow$ bool and `fromJust` : α option $\rightarrow \alpha$ are the Haskell-style discriminator and selector for option. Since the Isabelle formalization is parametric over rule systems (state, rule, eff), the code for `mkTree` explicitly takes `eff` as a parameter.

Although the code generator was not designed with codatatypes in mind, it is general enough to handle them. Internally, it reduces Isabelle specifications to higher-order rewrite systems [24] and generates functional code in Haskell, OCaml, Scala, or Standard ML. Partial correctness is guaranteed regardless of the target language’s evaluation strategy. However, for the guarantee to be non-vacuous for corecursive definitions, one needs a language with a lazy evaluation strategy, such as Haskell.

The verified contract of the program reads as follows: Given an available and persistent rule system (state, rule, eff), a fair rule enumeration ρ , and a state s representing the formula to prove, `mkTree eff  $\rho$  s` either yields a finite derivation tree of s or produces an infinite fair derivation tree whose infinite paths are all countermodel paths. These guarantees involve only partial correctness of ground term evaluation.

The generated code is a generic countermodel-producing semidecision procedure parameterized by the the proof system. Moreover, the fair rule enumeration parameter ρ can be instantiated to various choices that may perform better than the simple scheme described in Section 3.

7 Related Work

This paper joins a series of pearls aimed at reclaiming mathematical concepts and results for coinductive methods, including streams [31, 35], regular expressions [32, 34], and automata [33]. Some developments pass the ultimate test of formalization, usually in Agda and Coq, the codatatype-aware proof assistants par excellence: the sieve of Eratosthenes [3], real number basics [8], and temporal logic for red–blue trees [25].

So why write yet another formalized pearl involving coinduction? First, because we finally could—with the new codatatype package, Isabelle has caught up with its rivals in this area. Second, because, although codatatypes are a good match for the completeness theorem, there seems to be no proof in the literature that takes advantage of this.

While there are many accounts of the completeness theorem for FOL and related logics, most of them favor the more mathematical Henkin style, which obfuscates the rich structure of proof and failure. This preference has a long history. It is positively motivated by the ability to support uncountable languages. More crucially, it is negatively motivated by the lack of rigor perceived in the alternative: “geometric” reasoning about infinite trees. Negri [26] gives a revealing account in the context of modal logic, quoting reviews that were favorable to Kripke’s completeness result [21] but critical of his informal argument based on infinite tableau trees.⁴ Kaplan [19] remarks that “although the author extracts a great deal of information from his tableau constructions, a completely rigorous development along these lines would be extremely tedious.”

A few textbooks venture in a proof-theoretic presentation of completeness, notably Gallier’s [11]. Such a treatment highlights not only the structure, but also the algorithmic content of the proofs. The price is usually a lack of rigor, in particular a gap between the definition of derivation trees and its use in the completeness argument. This lack of rigor should not be taken lightly, as it may lead to serious ambiguities or errors: In the context of a tableau completeness proof development, Hähnle [15] first performs an implicit transition from finite to possibly infinite tableaux, and then claims that tableau chain suprema exist by wrongly invoking Zorn’s lemma [15, Definition 3.16].⁵

The completeness theorem has been mechanized before in proof assistants. Schlöder and Koepke, in Mizar [36], formalize a Henkin-style argument for possibly uncountable languages. Building on an early insight by Krivine [22] concerning the expressibility of the completeness proof in intuitionistic second-order logic, Ilik [17] analyzes Henkin-style arguments for classical and intuitionistic logic with respect to standard and Kripke models and formalizes them in Coq (without employing codatatypes).

At least three proofs were developed using HOL-based provers. Harrison [16], in HOL Light, and Berghofer [2], in Isabelle, formalize Henkin-style arguments. Ridge and Margetson [23, 30], in Isabelle, employ proof trees constructed as graphs of nodes that carry their levels as natural numbers. This last work has the merits of analyzing the computational content of proofs in the style of Gallier [11] and discussing an OCaml implementation. Our formalization relates to this work in a similar way to which our presentation relates to Gallier’s: The newly introduced support for codatatypes and co-recursion in Isabelle provides suitable abstraction mechanisms for reasoning about infinite trees, avoiding boilerplate for tree manipulation based on numeric indexing. Moreover, codatatypes are mapped naturally to Haskell types, allowing Isabelle’s code generator to produce certified Haskell code. Finally, our proof is more abstract and applies to several variants of FOL and beyond.

8 Conclusion

The completeness theorem is a fundamental result about classical logic. Its proof is presented in many variants in the literature. Few of these presentations emphasize the algorithmic content, and none of them uses codatatypes. Gallier’s pseudo-Pascal code

⁴ And Kripke’s degree of rigor in this early paper is not far from today’s state of the art in proof theory; see, e.g., Troelstra and Schwichtenberg [37].

⁵ This is the only error we found in this otherwise excellent chapter on tableaux.

is inspiring, but we prefer “pseudo-Haskell,” i.e., Isabelle/HOL with codatatypes. In our view, coinduction is the key to formulate an account that is both mathematically rigorous and abundant in algorithmic content. The definition of the abstract prover mkTree is stated rigorously, is accessible to functional programmers, and replaces pages of verbose descriptions.

The advantages of machine-checked metatheory are well known from programming language research, where new results are often formalized and proof assistants are used in the classroom. This paper, like its predecessor [4], reported on some steps we have taken to apply the same methods to formal logic and automated reasoning.

Acknowledgment. Tobias Nipkow made this work possible. Mark Summerfield and the anonymous reviewers suggested many textual improvements to earlier versions of this paper. Blanchette is supported by the Deutsche Forschungsgemeinschaft (DFG) project Hardening the Hammer (grant Ni 491/14-1). Popescu is supported by the DFG project Security Type Systems and Deduction (grant Ni 491/13-2) as part of the program Reliably Secure Software Systems (RS³, priority program 1496). Traytel is supported by the DFG program Program and Model Analysis (PUMA, doctorate program 1480). The authors are listed alphabetically.

References

1. Bell, J.L., Machover, M.: A Course in Mathematical Logic. North-Holland (1977)
2. Berghofer, S.: First-order logic according to Fitting. In: Klein, G., Nipkow, T., Paulson, L. (eds.) Archive of Formal Proofs. <http://afp.sf.net/entries/FOL-Fitting.shtml> (2007)
3. Bertot, Y.: Filters on coinductive streams, an application to Eratosthenes’ sieve. In: Urzyczyn, P. (ed.) TLCA 2005. LNCS, vol. 3461, pp. 102–115. Springer (2005)
4. Blanchette, J.C., Popescu, A.: Mechanizing the metatheory of Sledgehammer. In: Fontaine, P., Ringeissen, C., Schmidt, R.A. (eds.) FroCoS 2013. LNAI, vol. 8152, pp. 245–260. Springer (2013)
5. Blanchette, J.C., Hölzl, J., Lochbihler, A., Panny, L., Popescu, A., Traytel, D.: Truly modular (co)datatypes for Isabelle/HOL. In: Klein, G., Gamboa, R. (eds.) ITP 2014. LNCS, Springer (2014)
6. Blanchette, J.C., Popescu, A., Traytel, D.: Abstract completeness. In: Klein, G., Nipkow, T., Paulson, L. (eds.) Archive of Formal Proofs. http://afp.sf.net/entries/Abstract_Completeness.shtml (2014)
7. Blanchette, J.C., Popescu, A., Traytel, D.: Formal development associated with this paper. http://www21.in.tum.de/~traytel/compl_devel.zip (2014)
8. Ciaffaglione, A., Gianantonio, P.D.: A certified, corecursive implementation of exact real numbers. Theor. Comput. Sci. 351(1), 39–51 (2006)
9. Fitting, M.: First-Order Logic and Automated Theorem Proving, Second Edition. Graduate Texts in Computer Science, Springer (1996)
10. Francez, N.: Fairness. Texts and Monographs in Computer Science, Springer (1986)
11. Gallier, J.H.: Logic for Computer Science: Foundations of Automatic Theorem Proving. Computer Science and Technology, Harper & Row (1986)
12. Gödel, K.: Über die Vollständigkeit des Logikkalküls. Ph.D. thesis, Universität Wien (1929)
13. Gordon, M.J.C., Melham, T.F. (eds.): Introduction to HOL: A Theorem Proving Environment for Higher Order Logic. Cambridge University Press (1993)

14. Haftmann, F., Nipkow, T.: Code generation via higher-order rewrite systems. In: Blume, M., Kobayashi, N., Vidal, G. (eds.) FLOPS 2010. LNCS, vol. 6009, pp. 103–117. Springer (2010)
15. Hähnle, R.: Tableaux and related methods. In: Robinson, A., Voronkov, A. (eds.) Handbook of Automated Reasoning, vol. I, pp. 100–178. Elsevier (2001)
16. Harrison, J.: Formalizing basic first order model theory. In: Grundy, J., Newey, M.C. (eds.) TPHOLs '98. LNCS, vol. 1479, pp. 153–170. Springer (1998)
17. Ilik, D.: Constructive Completeness Proofs and Delimited Control. Ph.D. thesis, École Polytechnique (2010)
18. Jacobs, B., Rutten, J.: A tutorial on (co)algebras and (co)induction. Bull. Eur. Assoc. Theor. Comput. Sci. 62, 222–259 (1997)
19. Kaplan, D.: Review of Kripke (1959) [21]. J. Symb. Log. 31, 120–122 (1966)
20. Kleene, S.C.: Mathematical Logic. John Wiley & Sons (1967)
21. Kripke, S.: A completeness theorem in modal logic. J. Symb. Log. 24(1), 1–14 (1959)
22. Krivine, J.L.: Une preuve formelle et intuitionniste du théorème de complétude de la logique classique. Bull. Symb. Log. 2(4), 405–421 (1996)
23. Margetson, J., Ridge, T.: Completeness theorem. In: Klein, G., Nipkow, T., Paulson, L. (eds.) Archive of Formal Proofs. <http://afp.sf.net/entries/Completeness.shtml> (2004)
24. Mayr, R., Nipkow, T.: Higher-order rewrite systems and their confluence. Theor. Comput. Sci. 192(1), 3–29 (1998)
25. Nakata, K., Uustalu, T., Bezem, M.: A proof pearl with the fan theorem and bar induction: Walking through infinite trees with mixed induction and coinduction. In: Yang, H. (ed.) APLAS 2011. LNCS, vol. 7078, pp. 353–368. Springer (2011)
26. Negri, S.: Kripke completeness revisited. In: Primiero, G., Rahman, S. (eds.) Acts of Knowledge: History, Philosophy and Logic: Essays Dedicated to Göran Sundholm, pp. 247–282. College Publications (2009)
27. Nipkow, T., Klein, G.: Concrete Semantics: A Proof Assistant Approach. Springer (to appear), <http://www.in.tum.de/~nipkow/Concrete-Semantics>
28. Nipkow, T., Paulson, L.C., Wenzel, M.: Isabelle/HOL: A Proof Assistant for Higher-Order Logic. LNCS, vol. 2283. Springer (2002)
29. Pfenning, F.: Review of “Jean H. Gallier: Logic for Computer Science, Harper & Row, New York 1986” [11]. J. Symb. Log. 54(1), 288–289 (1989)
30. Ridge, T., Margetson, J.: A mechanically verified, sound and complete theorem prover for first order logic. In: Hurd, J., Melham, T.F. (eds.) TPHOLs 2005. LNCS, vol. 3603, pp. 294–309. Springer (2005)
31. Roşu, G.: Equality of streams is a Π_2^0 -complete problem. In: Reppey, J.H., Lawall, J.L. (eds.) ICFP '06. ACM (2006)
32. Roşu, G.: An effective algorithm for the membership problem for extended regular expressions. In: Seidl, H. (ed.) FoSSaCS 2007. LNCS, vol. 4423, pp. 332–345. Springer (2007)
33. Rutten, J.J.M.M.: Automata and coinduction (an exercise in coalgebra). In: Sangiorgi, D., de Simone, R. (eds.) CONCUR '98. LNCS, vol. 1466, pp. 194–218. Springer (1998)
34. Rutten, J.J.M.M.: Regular expressions revisited: A coinductive approach to streams, automata, and power series. In: Backhouse, R.C., Oliveira, J.N. (eds.) MPC 2000. LNCS, vol. 1837, pp. 100–101. Springer (2000)
35. Rutten, J.J.M.M.: Elements of stream calculus (an extensive exercise in coinduction). Electr. Notes Theor. Comput. Sci. 45, 358–423 (2001)
36. Schlöder, J.J., Koepke, P.: The Gödel completeness theorem for uncountable languages. Formalized Mathematics 20(3), 199–203 (2012)
37. Troelstra, A.S., Schwichtenberg, H.: Basic Proof Theory, Second Edition. Cambridge University Press (2000)