
SUPERPOSITION FOR LAMBDA-FREE HIGHER-ORDER LOGIC

ALEXANDER BENTKAMP^a, JASMIN BLANCHETTE^a, SIMON CRUANES^b,
AND UWE WALDMANN^c

^a Vrije Universiteit Amsterdam, Department of Computer Science, De Boelelaan 1111, 1081 HV
Amsterdam, The Netherlands
e-mail address: a.bentkamp@vu.nl
e-mail address: j.c.blanchette@vu.nl

^b Aesthetic Integration, 600 Congress Ave., Austin, Texas, 78701, USA
e-mail address: simon@imandra.ai

^c Max-Planck-Institut für Informatik, Saarland Informatics Campus, Saarbrücken, Germany
e-mail address: uwe@mpi-inf.mpg.de

ABSTRACT. We introduce refutationally complete superposition calculi for intentional and extensional clausal λ -free higher-order logic, two formalisms that allow partial application and applied variables. The calculi are parameterized by a term order that need not be fully monotonic, making it possible to employ the λ -free higher-order lexicographic path and Knuth–Bendix orders. We implemented the calculi in the Zipperposition prover and evaluated them on Isabelle/HOL and TPTP benchmarks. They appear promising as a stepping stone towards complete, highly efficient automatic theorem provers for full higher-order logic.

1. INTRODUCTION

Superposition is a highly successful calculus for reasoning about first-order logic with equality. We are interested in *graceful* generalizations to higher-order logic: calculi that, as much as possible, coincide with standard superposition on first-order problems and that scale up to arbitrary higher-order problems.

As a stepping stone towards full higher-order logic, in this article we restrict our attention to a clausal λ -free fragment of polymorphic higher-order logic that supports partial application and application of variables (Section 2). This formalism is expressive enough to permit the axiomatization of higher-order combinators such as $\mathbf{pow} : \Pi\alpha. \mathbf{nat} \rightarrow (\alpha \rightarrow \alpha) \rightarrow \alpha \rightarrow \alpha$ (intended to denote the iterated application $h^n x$):

$$\mathbf{pow}\langle\alpha\rangle \mathbf{Zero} h \approx \mathbf{id}\langle\alpha\rangle \qquad \mathbf{pow}\langle\alpha\rangle (\mathbf{Succ} n) h x \approx h (\mathbf{pow}\langle\alpha\rangle n h x)$$

Conventionally, functions are applied without parentheses and commas, and variables are italicized. Notice the variable number of arguments to $\mathbf{pow}\langle\alpha\rangle$ and the application of h . The

Key words and phrases: superposition calculus, Boolean-free lambda-free higher-order logic, refutational completeness.

Extended version of Bentkamp et al., “Superposition for lambda-free higher-order logic” [BBCW18].

expressiveness of full higher-order logic can be recovered by introducing SK-style combinators to represent λ -abstractions and proxies for the logical symbols [Ker91, Rob70].

A widespread technique to support partial application and application of variables in first-order logic is to make all symbols nullary and to represent application of functions by a distinguished binary symbol $\mathbf{app} : \Pi\alpha, \beta. \mathbf{fun}(\alpha, \beta) \times \alpha \rightarrow \beta$, where \mathbf{fun} is an uninterpreted binary type constructor. Following this scheme, the higher-order term $f(h\ f)$, where $f : \kappa \rightarrow \kappa'$, is translated to $\mathbf{app}(f, \mathbf{app}(h, f))$ —or rather $\mathbf{app}\langle\kappa, \kappa'\rangle(f, \mathbf{app}\langle\mathbf{fun}(\kappa, \kappa'), \kappa\rangle(h, f))$ if we specify the type arguments. We call this the *applicative encoding*. The existence of such a reduction to first-order logic explains why λ -free higher-order terms are also called “applicative first-order terms.” Unlike for full higher-order logic, most general unifiers are unique for our λ -free fragment, just as they are for applicatively encoded first-order terms.

Although the applicative encoding is complete [Ker91] and is employed fruitfully in tools such as HOLyHammer and Sledgehammer [BKPU16], it suffers from a number of weaknesses, all related to its gracelessness. Transforming all the function symbols into constants considerably restricts what can be achieved with term orders; for example, argument tuples cannot easily be compared using different methods for different symbols [Kop12, Section 2.3.1]. In a prover, the encoding also clutters the data structures, slows down the algorithms, and neutralizes the heuristics that look at the terms’ root symbols. But our chief objection is the sheer clumsiness of encodings and their poor integration with interpreted symbols. And they quickly accumulate; for example, using the traditional encoding of polymorphism relying on a distinguished binary function symbol \mathbf{t} [BBPS16, Section 3.3] in conjunction with the applicative encoding, the term $\mathbf{Succ}\ x$ becomes $\mathbf{t}(\mathbf{nat}, \mathbf{app}(\mathbf{t}(\mathbf{fun}(\mathbf{nat}, \mathbf{nat}), \mathbf{Succ}), \mathbf{t}(\mathbf{nat}, x)))$. The term’s simple structure is lost in translation.

Hybrid schemes have been proposed to strengthen the applicative encoding: If a given symbol always occurs with at least k arguments, these can be passed directly [MP08]. However, this relies on a closed-world assumption: that all terms that will ever be compared arise in the initial problem. This noncompositionality conflicts with the need for complete higher-order calculi to synthesize arbitrary terms during proof search [BM14]. As a result, hybrid encodings are not an ideal basis for higher-order automated reasoning.

Instead, we propose to generalize the superposition calculus to *intensional* and *extensional* clausal λ -free higher-order logic. For the extensional version of the logic, the property $(\forall x. h\ x \approx k\ x) \rightarrow h \approx k$ holds for all functions h, k of the same type. For each logic, we present two calculi (Section 3). The intentional calculi perfectly coincide with standard superposition on first-order clauses; the extensional calculi depend on an extra axiom.

Superposition is parameterized by a term order, which is used to prune the search space. If we assume that the term order is a simplification order enjoying totality on ground terms (i.e., terms containing no term or type variables), the standard calculus rules and completeness proof can be lifted verbatim. The only necessary changes concern the basic definitions of terms and substitutions. However, there is one monotonicity property that is hard to obtain unconditionally: *compatibility with arguments*. It states that $s' \succ s$ implies $s' t \succ s t$ for all terms s, s', t such that $s\ t$ and $s'\ t$ are well typed. Blanchette, Waldmann, and colleagues recently introduced graceful generalizations of the lexicographic path order (LPO) [BWW17] and the Knuth–Bendix order (KBO) [BBWW17] with argument coefficients, but they both lack this property. For example, given a KBO with $\mathbf{g} \succ \mathbf{f}$, it may well be that $\mathbf{g}\ \mathbf{a} \prec \mathbf{f}\ \mathbf{a}$ if \mathbf{f} has a large enough multiplier on its argument.

Our superposition calculi are designed to be refutationally complete for such nonmonotonic orders (Section 4). To achieve this, they include an inference rule for argument

congruence, which derives $C \vee s x \approx t x$ from $C \vee s \approx t$. The redundancy criterion is defined in such a way that the larger, derived clause is not subsumed by the premise. In the completeness proof, the most difficult case is the one that normally excludes superposition at or below variables using the induction hypothesis. With nonmonotonicity, this approach no longer works, and we propose two alternatives: Either perform some superposition inferences into higher-order variables or “purify” the clauses to circumvent the issue. We refer to the corresponding calculi as *nonpurifying* and *purifying*.

The calculi are implemented in the Zipperposition prover [Cru17] (Section 5). We evaluate them on first- and higher-order Isabelle/HOL [BN10] and TPTP benchmarks [SSCB12, SBBT09] and compare them with the applicative encoding (Section 6). We find that there is a substantial cost associated with the applicative encoding, that the nonmonotonicity is not particularly expensive, and that the nonpurifying calculi outperform the purifying calculi.

An earlier version of this work was presented at IJCAR 2018 [BBCW18]. This article extends the conference paper with detailed soundness and completeness proofs and more explanations. Because of too weak selection restrictions on the purifying calculi, our claim of refutational completeness in the conference version was not entirely correct. We now strengthened the selection restrictions accordingly. Moreover, we extended the logic with polymorphism, leading to minor modifications to the calculus. We also simplified the presentation of the clausal fragment of the logic that interests us. In particular, we removed mandatory arguments. The redundancy criterion also differs slightly from the conference version. Finally, we updated the empirical evaluation to reflect recent improvements in the Zipperposition prover.

Since the publication of the conference paper, two research groups have built on our work. Bhayat and Reger [BR20] extended our intensional nonpurifying calculus to a version of higher-order logic with combinators. They use it with a nonmonotonic order that orients the defining equations of SK-style combinators. Bentkamp et al. [BBT⁺21] extended our extensional nonpurifying calculus to a calculus on λ -terms. For them, support for nonmonotonic orders is crucial as well because no ground-total simplification order exists for λ -terms up to β -conversion. Based on these two extensions of our calculi, the provers Zipperposition and Vampire achieved first and third place, respectively, in the higher-order category of the 2020 edition of the CADE ATP System Competition (CASC-J10).

2. LOGIC

Our logic is intended as an intermediate step on the way towards full higher-order logic (also called simple type theory) [Chu40, GM93]. Refutational completeness of calculi for higher-order logic is usually stated in terms of Henkin semantics [Hen50, BM14], in which the universes used to interpret functions need only contain the functions that can be expressed as terms. Since the terms of λ -free higher-order logic exclude λ -abstractions, in “ λ -free Henkin semantics” the universes interpreting functions can be even smaller. In that sense, our semantics resemble Henkin prestructures [Lei94, Section 5.4]. In contrast to other higher-order logics [Vää19], there are no comprehension principles, and we disallow nesting of Boolean formulas inside terms.

2.1. Syntax. We fix a set Σ_{ty} of type constructors with arities and a set \mathcal{V}_{ty} of type variables. We require at least one nullary type constructor and a binary type constructor \rightarrow to be present in Σ_{ty} . We inductively define a λ -free higher-order type to be either a type variable

$\alpha \in \mathcal{V}_{\text{ty}}$ or of the form $\kappa(\bar{\tau}_n)$ for an n -ary type constructor $\kappa \in \Sigma_{\text{ty}}$ and types $\bar{\tau}_n$. Here and elsewhere, we write \bar{a}_n or \bar{a} to abbreviate the tuple (a_1, \dots, a_n) or product $a_1 \times \dots \times a_n$, for $n \geq 0$. We write κ for $\kappa()$ and $\tau \rightarrow v$ for $\rightarrow(\tau, v)$. A type declaration is an expression of the form $\Pi \bar{\alpha}_m. \tau$ (or simply τ if $m = 0$), where all type variables occurring in τ belong to $\bar{\alpha}_m$.

We fix a set Σ of symbols with type declarations, written as $f : \Pi \bar{\alpha}_m. \tau$ or f , and a set \mathcal{V} of typed variables, written as $x : \tau$ or x . We require Σ to contain a symbol with type declaration $\Pi \alpha. \alpha$, to ensure that the (Herbrand) domain of every type is nonempty. The sets $(\Sigma_{\text{ty}}, \Sigma)$ form the logic's signature. We reserve the letters s, t, u, v, w for terms and x, y, z for variables and write $: \tau$ to indicate their type. The set of λ -free higher-order terms is defined inductively as follows. Every variable $x : \tau \in \mathcal{V}$ is a term. If $f : \Pi \bar{\alpha}_m. \tau$ is a symbol and \bar{v}_m are types, then $f(\bar{v}_m) : \tau\{\bar{\alpha}_m \mapsto \bar{v}_m\}$ is a term. If $t : \tau \rightarrow v$ and $u : \tau$, then $t u : v$ is a term, called an *application*. Nonapplication terms are called *heads*. Application is left-associative, and correspondingly the function type constructor \rightarrow is right-associative. Using the spine notation [CP03], terms can be decomposed in a unique way as a head t applied to zero or more arguments: $t s_1 \dots s_n$ or $t \bar{s}_n$ (abusing notation). A term is *ground* if it is built without using type or term variables.

Substitution and unification are generalized in the obvious way, without the difficulties caused by λ -abstractions. A substitution has the form $\{\bar{\alpha}_m, \bar{x}_n \mapsto \bar{v}_m, \bar{s}_n\}$, where each x_j has type τ_j and each s_j has type $\tau_j\{\bar{\alpha}_m \mapsto \bar{v}_m\}$, mapping m type variables to m types and n term variables to n terms. A unifier of two terms s and t is a substitution ρ such that $s\rho = t\rho$. A most general unifier $\text{mgu}(s, t)$ of two terms s and t is a unifier σ of s and t such that for every other unifier θ , there exists a substitution ρ such that $\alpha\theta = \alpha\sigma\rho$ and $x\theta = x\sigma\rho$ for all $\alpha \in \mathcal{V}_{\text{ty}}$ and all $x \in \mathcal{V}$. As in first-order logic, the most general unifier is unique up to variable renaming. For example, $\text{mgu}(x \mathbf{b} z, \mathbf{f} \mathbf{a} y \mathbf{c}) = \{x \mapsto \mathbf{f} \mathbf{a}, y \mapsto \mathbf{b}, z \mapsto \mathbf{c}\}$, and $\text{mgu}(y(\mathbf{f} \mathbf{a}), \mathbf{f}(y \mathbf{a})) = \{y \mapsto \mathbf{f}\}$, assuming that the types of the unified subterms are equal.

An equation $s \approx t$ is formally an unordered pair of terms s and t of the same type. A literal is an equation or a negated equation, written $s \not\approx t$. A clause $L_1 \vee \dots \vee L_n$ is a finite multiset of literals L_j . The empty clause is written as \perp .

2.2. Semantics. A *type interpretation* $\mathcal{I}_{\text{ty}} = (\mathcal{U}, \mathcal{J}_{\text{ty}})$ is defined as follows. The set \mathcal{U} is a nonempty collection of nonempty sets, called *universes*. The function \mathcal{J}_{ty} associates a function $\mathcal{J}_{\text{ty}}(\kappa) : \mathcal{U}^n \rightarrow \mathcal{U}$ with each n -ary type constructor κ . A *type valuation* ξ is a function that maps every type variable to a universe. The *denotation* of a type for a type interpretation \mathcal{I}_{ty} and a type valuation ξ is defined by $\llbracket \alpha \rrbracket_{\mathcal{I}_{\text{ty}}}^\xi = \xi(\alpha)$ and $\llbracket \kappa(\bar{\tau}) \rrbracket_{\mathcal{I}_{\text{ty}}}^\xi = \mathcal{J}_{\text{ty}}(\kappa)(\llbracket \bar{\tau} \rrbracket_{\mathcal{I}_{\text{ty}}}^\xi)$. Here and elsewhere, we abuse notation by applying an operation on a tuple when it must be applied elementwise; thus, $\llbracket \bar{\tau}_n \rrbracket_{\mathcal{I}_{\text{ty}}}^\xi$ stands for $\llbracket \tau_1 \rrbracket_{\mathcal{I}_{\text{ty}}}^\xi, \dots, \llbracket \tau_n \rrbracket_{\mathcal{I}_{\text{ty}}}^\xi$.

A type valuation ξ can be extended to be a *valuation* by additionally assigning an element $\xi(x) \in \llbracket \tau \rrbracket_{\mathcal{I}_{\text{ty}}}^\xi$ to each variable $x : \tau$. An *interpretation function* \mathcal{J} for a type interpretation \mathcal{I}_{ty} associates with each symbol $f : \Pi \bar{\alpha}_m. \tau$ and universe tuple $\bar{U}_m \in \mathcal{U}^m$ a value $\mathcal{J}(f, \bar{U}_m) \in \llbracket \tau \rrbracket_{\mathcal{I}_{\text{ty}}}^\xi$, where ξ is the type valuation that maps each α_i to U_i . Loosely following Fitting [Fit02, Section 2.5], an *extension function* \mathcal{E} associates to any pair of universes $U_1, U_2 \in \mathcal{U}$ a function $\mathcal{E}_{U_1, U_2} : \mathcal{J}_{\text{ty}}(\rightarrow)(U_1, U_2) \rightarrow (U_1 \rightarrow U_2)$. Together, a type interpretation, an interpretation function, and an extension function form an *interpretation* $\mathcal{I} = (\mathcal{U}, \mathcal{J}_{\text{ty}}, \mathcal{J}, \mathcal{E})$.

An interpretation is *extensional* if \mathcal{E}_{U_1, U_2} is injective for all U_1, U_2 . Both intensional and extensional logics are widely used for interactive theorem proving; for example, Coq's calculus

of inductive constructions is intensional [BC04], whereas Isabelle/HOL is extensional [NPW02]. The semantics is *standard* if \mathcal{E}_{U_1, U_2} is bijective for all U_1, U_2 .

For an interpretation $\mathcal{I} = (\mathcal{U}, \mathcal{I}_{\text{ty}}, \mathcal{J}, \mathcal{E})$ and a valuation ξ , the denotation of a term is defined as follows: For variables x , let $\llbracket x \rrbracket_{\mathcal{I}}^{\xi} = \xi(x)$. For symbols f , let $\llbracket f \langle \bar{\tau} \rangle \rrbracket_{\mathcal{I}}^{\xi} = \mathcal{J}(f, \llbracket \bar{\tau} \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi})$. For applications $s t$ of a term $s : \tau \rightarrow v$ to a term $t : \tau$, let $U_1 = \llbracket \tau \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi}$, $U_2 = \llbracket v \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi}$, and $\llbracket s t \rrbracket_{\mathcal{I}}^{\xi} = \mathcal{E}_{U_1, U_2}(\llbracket s \rrbracket_{\mathcal{I}}^{\xi})(\llbracket t \rrbracket_{\mathcal{I}}^{\xi})$. If t is a ground term, we also write $\llbracket t \rrbracket_{\mathcal{I}}$ for the denotation of t because it does not depend on the valuation.

An equation $s \approx t$ is true in \mathcal{I} for ξ if $\llbracket s \rrbracket_{\mathcal{I}}^{\xi} = \llbracket t \rrbracket_{\mathcal{I}}^{\xi}$; otherwise, it is false. A disequation $s \not\approx t$ is true if $s \approx t$ is false. A clause is true if at least one of its literals is true. The interpretation \mathcal{I} is a model of a clause C , written $\mathcal{I} \models C$, if C is true in \mathcal{I} for all valuations ξ . It is a model of a set of clauses if it is a model of all contained clauses.

For example, given the signature $(\{\kappa, \rightarrow\}, \{\mathbf{a} : \kappa\})$ and a variable $h : \kappa \rightarrow \kappa$, the clause $h \mathbf{a} \not\approx \mathbf{a}$ has an extensional model with $\mathcal{U} = \{U_1, U_2\}$, $U_1 = \{a, b\}$ ($a \neq b$), $U_2 = \{f\}$, $\mathcal{I}_{\text{ty}}(\kappa) = U_1$, $\mathcal{I}_{\text{ty}}(\rightarrow)(U_1, U_1) = U_2$, $\mathcal{J}(\mathbf{a}) = a$, $\mathcal{E}_{U_1, U_1}(f)(a) = \mathcal{E}_{U_1, U_1}(f)(b) = b$.

3. THE CALCULI

We introduce four versions of the *Boolean-free λ -free higher-order superposition calculus*, articulated along two axes: intentional versus extensional, and nonpurifying versus purifying. To avoid repetitions, our presentation unifies them into a single framework.

3.1. The Inference Rules. To support nonmonotonic term orders, we restrict superposition inferences to *green subterms*, which are defined inductively as follows:

Definition 3.1 (Green subterms and contexts). A term t' is a *green subterm* of t if $t = t'$ or if $t = s \bar{u}$ and t' is a green subterm of u_i for some i . We write $s \langle u \rangle$ to indicate that the subterm u of $s[u]$ is a green subterm. Correspondingly, we call the context $s \langle \rangle$ around a green subterm a *green context*.

By this definition, f and $f \mathbf{a}$ are subterms of $f \mathbf{a} \mathbf{b}$, but not green subterms. The green subterms of $f \mathbf{a} \mathbf{b}$ are \mathbf{a} , \mathbf{b} , and $f \mathbf{a} \mathbf{b}$. Thus, $[\] \mathbf{a} \mathbf{b}$ and $[\] \mathbf{b}$ are not green contexts, but $f [\] \mathbf{b}$, $f \mathbf{a} [\]$, and $[\]$ are.

The calculi are parameterized by a partial order \succ on terms that

- is well founded on ground terms;
- is total on ground terms;
- has the subterm property on ground terms;
- is *compatible with green contexts* on ground terms: $t' \succ t$ implies $s \langle t' \rangle \succ s \langle t \rangle$;
- is stable under grounding substitutions: $t \succ s$ implies $t\theta \succ s\theta$ for all substitutions θ grounding t and s .

The order need not be *compatible with arguments*: $s' \succ s$ need not imply $s' t \succ s t$, even on ground terms. The literal and clause orders are defined from \succ as multiset extensions in the standard way [BG94]. Despite their names, the term, literal, and clause orders need not be transitive on nonground entities.

The λ -free higher-order generalizations of LPO [BWW17] and KBO [BBWW17] as well as EPO [Ben18] fulfill these requirements, with the caveat that they are defined on untyped terms. To use them on polymorphic terms, we can encode type arguments as term arguments and ignore the remaining type information.

Literal selection is supported. The selection function maps each clause C to a subclause of C consisting of negative literals. A literal L is (*strictly*) *eligible* w.r.t. a substitution σ in C if it is selected in C or there are no selected literals in C and $L\sigma$ is (*strictly*) maximal in $C\sigma$. If σ is the identity substitution, we leave it implicit.

The following four rules are common to all four calculi. We regard positive and negative superposition as two cases of the same rule

$$\frac{\overbrace{D' \vee t \approx t'}^D \quad \overbrace{C' \vee s\langle u \rangle \approx s'}^C}{(D' \vee C' \vee s\langle t' \rangle \approx s')\sigma} \text{SUP}$$

where \approx denotes \approx or $\not\approx$ and the following conditions are fulfilled:

1. $\sigma = \text{mgu}(t, u)$;
2. $t\sigma \not\approx t'\sigma$;
3. $s\langle u \rangle\sigma \not\approx s'\sigma$;
4. $t \approx t'$ is strictly eligible w.r.t. σ in D ;
5. $C\sigma \not\approx D\sigma$;
6. $s\langle u \rangle \approx s'$ is eligible w.r.t. σ in C and, if positive, even strictly eligible;
7. the *variable condition* must hold, which is specified individually for each calculus below.

In each calculus, we will define the variable condition to coincide with the condition “ $u \notin \mathcal{V}$ ” if the premises are first-order.

The equality resolution and equality factoring rules are almost identical to their standard counterparts:

$$\frac{\overbrace{C' \vee s \not\approx s'}^C}{C'\sigma} \text{ERES}$$

where

1. $\sigma = \text{mgu}(s, s')$;
2. $s \not\approx s'$ is eligible w.r.t. σ in C .

$$\frac{\overbrace{C' \vee s' \approx t' \vee s \approx t}^C}{(C' \vee t \not\approx t' \vee s \approx t')\sigma} \text{EFACT}$$

where

1. $\sigma = \text{mgu}(s, s')$;
2. $s\sigma \not\approx t\sigma$;
3. $s \approx t$ is eligible w.r.t. σ in C .

The following *argument congruence* rule compensates for the limitation that the superposition rule applies only to green subterms:

$$\frac{\overbrace{C' \vee s \approx s'}^C}{C'\sigma \vee s\sigma \bar{x} \approx s'\sigma \bar{x}} \text{ARGCONG}$$

where

1. $s \approx s'$ is strictly eligible w.r.t. σ in C ;
2. \bar{x} is a nonempty tuple of distinct fresh variables;
3. σ is the most general type substitution that ensures well-typedness of the conclusion.

In particular, if s takes m arguments, there are m ARGCONG conclusions for this literal, for which σ is the identity and \bar{x} is a tuple of $1, \dots, m-1$, or m variables. If the result type of s is a type variable, we have in addition infinitely many ARGCONG conclusions, for

which σ instantiates the type variable in the result type of s with $\bar{\alpha}_k \rightarrow \beta$ for some $k > 0$ and fresh type variables $\bar{\alpha}_k$ and β and for which \bar{x} is a tuple of $m + k$ variables. In practice, the enumeration of the infinitely many conclusions must be interleaved with other inferences via some form of dovetailing.

For the **intensional nonpurifying** calculus, the variable condition of the SUP rule is as follows:

Either $u \notin \mathcal{V}$ or there exists a grounding substitution θ with $t\sigma\theta \succ t'\sigma\theta$ and $C\sigma\theta \prec C\{u \mapsto t'\}\sigma\theta$.

This condition generalizes the standard condition that $u \notin \mathcal{V}$. The two coincide if C is first-order or if the term order is monotonic. In some cases involving nonmonotonicity, the variable condition effectively mandates SUP inferences at variable positions of the right premise, but never below. We will call these inferences *at variables*.

For the **extensional nonpurifying** calculus, the variable condition uses the following definition.

Definition 3.2. A term of the form $x \bar{s}_n$, for $n \geq 0$, *jells* with a literal $t \approx t' \in D$ if $t = \tilde{t} \bar{y}_n$ and $t' = \tilde{t}' \bar{y}_n$ for some terms \tilde{t}, \tilde{t}' and distinct variables \bar{y}_n that do not occur elsewhere in D .

Using the naming convention from Definition 3.2 for \tilde{t}' , the variable condition can be stated as follows:

If u has a variable head x and jells with the literal $t \approx t' \in D$, there must exist a grounding substitution θ with $t\sigma\theta \succ t'\sigma\theta$ and $C\sigma\theta \prec C''\sigma\theta$, where $C'' = C\{x \mapsto \tilde{t}'\}$.

If C is first-order, this amounts to $u \notin \mathcal{V}$. Since the order is compatible with green contexts, the substitution θ can exist only if x occurs applied in C .

Moreover, the extensional nonpurifying calculus has one additional rule, the positive extensionality rule, and one axiom, the extensionality axiom. The rule is

$$\frac{C' \vee s \bar{x} \approx s' \bar{x}}{C' \vee s \approx s'} \text{PosEXT}$$

where

1. \bar{x} is a tuple of distinct variables that do not occur in C' , s , or s'
2. $s \bar{x} \approx s' \bar{x}$ is strictly eligible in the premise.

The extensionality axiom uses a polymorphic Skolem symbol $\text{diff} : \Pi\alpha, \beta. (\alpha \rightarrow \beta)^2 \rightarrow \alpha$ characterized by the axiom

$$x (\text{diff}\langle\alpha, \beta\rangle x y) \not\approx y (\text{diff}\langle\alpha, \beta\rangle x y) \vee x \approx y \quad (\text{EXT})$$

Unlike the nonpurifying calculi, the purifying calculi never perform superposition at variables. Instead, they rely on purification [Bra75, DH86, SS89, SL91] (also called abstraction) to circumvent nonmonotonicity. The idea is to rename apart problematic occurrences of a variable x in a clause to x_1, \dots, x_n and to add *purification literals* $x_1 \not\approx x, \dots, x_n \not\approx x$ to connect the new variables to x . We must then ensure that all clauses are purified, by processing the initial clause set and the conclusion of every inference or simplification.

In the **intensional purifying** calculus, the purification $\text{pure}(C)$ of clause C is defined as the result of the following procedure. Choose a variable x that occurs applied in C and also unapplied in a literal of C that is not of the form $x \not\approx y$. If no such variable exists, terminate. Otherwise, replace all unapplied occurrences of x in C by a fresh variable x'

and add the purification literal $x' \not\approx x$. Then repeat these steps with another variable. The procedure terminates because the number of variables that can be chosen reduces with each step. For example,

$$\mathit{pure}(x \mathbf{a} \approx x \mathbf{b} \vee \mathbf{f} x \approx \mathbf{g} x) = x \mathbf{a} \approx x \mathbf{b} \vee \mathbf{f} x' \approx \mathbf{g} x' \vee x \not\approx x'$$

The variable condition is standard:

The term u is not a variable.

The conclusion C of ARGCONG is changed to $\mathit{pure}(C)$; the other rules preserve purity of their premises.

In the **extensional purifying** calculus, $\mathit{pure}(C)$ is defined as follows. Choose a variable x occurring in green subterms $x \bar{u}$ and $x \bar{v}$ in literals of C that are not of the form $x \not\approx y$, where \bar{u} and \bar{v} are distinct (possibly empty) term tuples. If no such variable exists, terminate. Otherwise, replace all green subterms $x \bar{v}$ with $x' \bar{v}$, where x' is fresh, and add the purification literal $x' \not\approx x$. Then repeat the procedure until no variable fulfilling the requirements is left. The procedure terminates because the number of variables fulfilling the requirements does not increase and with each step, the number of distinct tuples \bar{u} and \bar{v} fulfilling the requirements decreases for the chosen variable x . For example,

$$\mathit{pure}(x \mathbf{a} \approx x \mathbf{b} \vee \mathbf{f} x \approx \mathbf{g} x) = x \mathbf{a} \approx x' \mathbf{b} \vee \mathbf{f} x'' \approx \mathbf{g} x'' \vee x' \not\approx x \vee x'' \not\approx x$$

Like the extensional nonpurifying calculus, this calculus also contains the POEXT rule and axiom (EXT) introduced above. The variable condition is as follows:

Either u has a nonvariable head or u does not jell with the literal $t \approx t' \in D$.

The conclusion E of each rule is changed to $\mathit{pure}(E)$, except for POEXT, which preserves purity.

Finally, we impose further restrictions on literal selection. In the nonpurifying calculi, a literal must not be selected if $x \bar{u}$ is a maximal term of the clause and the literal contains a green subterm $x \bar{v}$ with $\bar{v} \neq \bar{u}$. In the purifying calculi, a literal must not be selected if it contains a variable of functional type. These restrictions are needed in our completeness proof to show that superposition inferences below variables are redundant. For the purifying calculi, the restriction is also needed to avoid inferences whose conclusion is identical to their premise, such as

$$\frac{z \mathbf{a} \approx \mathbf{b} \vee z \not\approx x \vee \mathbf{f} \not\approx x}{z \mathbf{a} \approx \mathbf{b} \vee z' \not\approx \mathbf{f} \vee z' \not\approx z} \text{ERES}$$

where $\mathbf{f} \not\approx x$ is selected. For the nonpurifying calculi, it might be possible to avoid the restriction at the cost of a more elaborate argument.

Example 3.3. We illustrate the different variable condition with a few examples. Consider a left-to-right LPO [BWW17] instance with precedence $\mathbf{h} \succ \mathbf{g} \succ \mathbf{f} \succ \mathbf{c} \succ \mathbf{b} \succ \mathbf{a}$. Given the clauses D and C of a superposition inference in which the grayed subterms are unified, we list below whether each calculus's variable condition is fulfilled (✓) or not.

D	C	intensional nonpurifying	extensional nonpurifying	intensional purifying	extensional purifying
$h \approx g$	$f y \approx c$				
$f(h a) \approx h$	$g y \approx y b$	✓	✓		
$h x \approx f x$	$g(y b) y \approx a$	✓		✓	
$x \approx c \vee h x \approx f x$	$g(y b) y \approx a$	✓	✓	✓	✓

For the purifying calculi, the clauses would actually undergo purification first, but this has no impact on the variable conditions. In the last row, the term $y b$ does not jell with $h x \approx f x \in D$ because x occurs also in the first literal of D .

Remark 3.4. In descriptions of first-order logic with equality, the property $y \approx y' \rightarrow f(\bar{x}, y, \bar{z}) \approx f(\bar{x}, y', \bar{z})$ is often referred to as “function congruence.” It seems natural to use the same label for the higher-order version $t \approx t' \rightarrow s t \approx s t'$ and to call the companion property $s \approx s' \rightarrow s t \approx s' t$ “argument congruence,” whence the name ARGCONG for our inference rule. This nomenclature is far from universal; for example, the Isabelle/HOL theorem *fun_cong* captures argument congruence and *arg_cong* captures function congruence.

3.2. Rationale for the Inference Rules. A key restriction of all four calculi is that they superpose only at green subterms, mirroring the term order’s compatibility with green contexts. The ARGCONG rule then makes it possible to simulate superposition at nongreen subterms. However, in conjunction with the SUP rules, ARGCONG can exhibit an unpleasant behavior, which we call *argument congruence explosion*:

$$\text{ARGCONG} \frac{g \approx f}{g x \approx f x} \quad \text{SUP} \frac{g x \approx f x \quad h a \not\approx b}{f a \not\approx b} \qquad \text{ARGCONG} \frac{g \approx f}{g x y z \approx f x y z} \quad \text{SUP} \frac{g x y z \approx f x y z \quad h a \not\approx b}{f x y a \not\approx b}$$

In both derivation trees, the higher-order variable h is effectively the target of a SUP inference. Such derivations essentially amount to superposition at variable positions (as shown on the left) or even superposition below variable positions (as shown on the right), both of which can be extremely prolific. In standard superposition, the explosion is averted by the condition on the SUP rule that $u \notin \mathcal{V}$. In the extensional purifying calculus, the variable condition tests that either u has a nonvariable head or u does not jell with the literal $t \approx t' \in D$, which prevents derivations such as the above. In the corresponding nonpurifying calculus, some such derivations may need to be performed when the term order exhibits nonmonotonicity for the terms of interest.

In the intensional calculi, the explosion can arise because the variable conditions are weaker. The following example shows that the intensional nonpurifying calculus would be incomplete if we used the variable condition of the extensional nonpurifying calculus.

Example 3.5. Consider a left-to-right LPO [BWW17] instance with precedence $h \succ g \succ f \succ b \succ a$, and consider the following unsatisfiable clause set:

$$h x \approx f x \qquad g(x b) x \approx a \qquad g(f b) h \not\approx a$$

The only possible inference is a SUP inference of the first into the second clause, but the variable condition of the extensional nonpurifying calculus is not met.

It is unclear whether the variable condition of the intensional purifying calculus could be strengthened, but our completeness proof suggests that it cannot.

The variable conditions in the extensional calculi are designed to prevent the argument congruence explosion shown above, but since they consider only the shape of the clauses, they might also block SUP inferences whose side premises do not originate from ARGCONG. This is why we need the POSEXT rule.

Example 3.6. In the following unsatisfiable clause set, the only possible inference from these clauses in the extensional nonpurifying calculus is POSEXT, showing its necessity:

$$g x \approx f x \quad g \not\approx f \quad x (\text{diff}\langle\alpha, \beta\rangle x y) \not\approx y (\text{diff}\langle\alpha, \beta\rangle x y) \vee x \approx y$$

The same argument applies for the purifying calculus with the difference that the third clause must be purified.

Due to nonmonotonicity, for refutational completeness we need either to purify the clauses or to allow some superposition at variable positions, as mandated by the respective variable conditions. Without either of these measures, at least the extensional calculi and presumably also the intensional calculi would be incomplete, as the next example demonstrates.

Example 3.7. Consider the following clause set:

$$\begin{aligned} k (g x) \approx k (x b) \quad k (f (h a) b) \not\approx k (g h) \quad f (h a) \approx h \quad f (h a) x \approx h x \\ x (\text{diff}\langle\alpha, \beta\rangle x y) \not\approx y (\text{diff}\langle\alpha, \beta\rangle x y) \vee x \approx y \end{aligned}$$

Using a left-to-right LPO [BWW17] instance with precedence $k \succ h \succ g \succ f \succ b \succ a$, this clause set is saturated w.r.t. the extensional purifying calculus when omitting purification. It also quickly saturates using the extensional nonpurifying calculus when omitting SUP inferences at variables. By contrast, the intensional calculi derive \perp , even without purification and without SUP inferences at variables, because of the less restrictive variable conditions.

This raises the question as to whether the intensional calculi actually need to purify or to perform SUP inferences at variables. We conjecture that omitting purification and SUP inferences at variables in the intensional calculi is complete when redundant clauses are kept but that it is incomplete in general.

We initially considered inference rules instead of axiom (EXT). However, we did not find a set of inference rules that is complete and leads to fewer inferences than (EXT). We considered the POSEXT rule described above in combination with the following rule:

$$\frac{C \vee s \not\approx t}{C \vee s (\text{sk}\langle\bar{\alpha}\rangle \bar{x}_n) \not\approx t (\text{sk}\langle\bar{\alpha}\rangle \bar{x}_n)} \text{NEGEXT}$$

where sk is a fresh Skolem symbol and $\bar{\alpha}$ and \bar{x}_n are the type and term variables occurring free in the literal $s \not\approx t$. However, these two rules do not suffice for a refutationally complete calculus, as the following example demonstrates:

Example 3.8. Consider the clause set

$$f x \approx a \quad g x \approx a \quad h f \approx b \quad h g \not\approx b$$

Assuming that all four equations are oriented from left to right, this set is saturated w.r.t. the extensional calculi if (EXT) is replaced by NEGEXT; yet it is unsatisfiable in an extensional logic.

Example 3.9. A significant advantage of our calculi over the use of standard superposition on applicatively encoded problems is the flexibility they offer in orienting equations. The following equations provide two definitions of addition on Peano numbers:

$$\begin{array}{ll} \text{add}_L \text{Zero } y \approx y & \text{add}_R x \text{Zero} \approx x \\ \text{add}_L (\text{Succ } x) y \approx \text{add}_L x (\text{Succ } y) & \text{add}_R x (\text{Succ } y) \approx \text{add}_R (\text{Succ } x) y \end{array}$$

Let $\text{add}_L (\text{Succ}^{100} \text{Zero}) n \not\approx \text{add}_R n (\text{Succ}^{100} \text{Zero})$ be the negated conjecture. With LPO, we can use a left-to-right comparison for add_L 's arguments and a right-to-left comparison for add_R 's arguments to orient all four equations from left to right. Then the negated conjecture can be simplified to $\text{Succ}^{100} n \not\approx \text{Succ}^{100} n$ by simplification (demodulation), and \perp can be derived with a single inference. If we use the applicative encoding instead, there is no instance of LPO or KBO that can orient both recursive equations from left to right. For at least one of the two sides of the negated conjecture, simplification is replaced by 100 SUP inferences, which is much less efficient, especially in the presence of additional axioms.

3.3. Soundness. To show the inferences' soundness, we need the substitution lemma for our logic:

Lemma 3.10 (Substitution lemma). *Let $\mathcal{I} = (\mathcal{U}, \mathcal{J}_{\text{ty}}, \mathcal{J}, \mathcal{E})$ be a λ -free higher-order interpretation. Then*

$$\llbracket \tau \rho \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi} = \llbracket \tau \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi'} \text{ and } \llbracket t \rho \rrbracket_{\mathcal{I}}^{\xi} = \llbracket t \rrbracket_{\mathcal{I}}^{\xi'}$$

for all terms t , all types τ , and all substitutions ρ , where $\xi'(\alpha) = \llbracket \alpha \rho \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi}$ for all type variables α and $\xi'(x) = \llbracket x \rho \rrbracket_{\mathcal{I}}^{\xi}$ for all term variables x .

Proof. First, we prove that $\llbracket \tau \rho \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi} = \llbracket \tau \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi'}$ by induction on the structure of τ . If $\tau = \alpha$ is a type variable,

$$\llbracket \alpha \rho \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi} = \xi'(\alpha) = \llbracket \alpha \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi'}$$

If $\tau = \kappa(\bar{v})$ for some type constructor κ and types \bar{v} ,

$$\llbracket \kappa(\bar{v}) \rho \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi} = \mathcal{J}_{\text{ty}}(\kappa)(\llbracket \bar{v} \rho \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi}) \stackrel{\text{IH}}{=} \mathcal{J}_{\text{ty}}(\kappa)(\llbracket \bar{v} \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi'}) = \llbracket \kappa(\bar{v}) \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi'}$$

Next, we prove $\llbracket t \rho \rrbracket_{\mathcal{I}}^{\xi} = \llbracket t \rrbracket_{\mathcal{I}}^{\xi'}$ by structural induction on t . If $t = y$, then by the definition of the denotation of a variable

$$\llbracket y \rho \rrbracket_{\mathcal{I}}^{\xi} = \xi'(y) = \llbracket y \rrbracket_{\mathcal{I}}^{\xi'}$$

If $t = f(\bar{\tau})$, then by the definition of the term denotation

$$\llbracket f(\bar{\tau}) \rho \rrbracket_{\mathcal{I}}^{\xi} = \mathcal{J}(f, \llbracket \bar{\tau} \rho \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi}) = \mathcal{J}(f, \llbracket \bar{\tau} \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi'}) = \llbracket f(\bar{\tau}) \rrbracket_{\mathcal{I}}^{\xi'}$$

If $t = u v$, then by the definition of the term denotation

$$\llbracket (u v) \rho \rrbracket_{\mathcal{I}}^{\xi} = \mathcal{E}_{U_1, U_2}(\llbracket u \rho \rrbracket_{\mathcal{I}}^{\xi})(\llbracket v \rho \rrbracket_{\mathcal{I}}^{\xi}) \stackrel{\text{IH}}{=} \mathcal{E}_{U_1, U_2}(\llbracket u \rrbracket_{\mathcal{I}}^{\xi'})(\llbracket v \rrbracket_{\mathcal{I}}^{\xi'}) = \llbracket u v \rrbracket_{\mathcal{I}}^{\xi'}$$

where u is of type $\tau \rightarrow v$, $U_1 = \llbracket \tau \rho \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi} = \llbracket \tau \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi'}$, and $U_2 = \llbracket v \rho \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi} = \llbracket v \rrbracket_{\mathcal{I}_{\text{ty}}}^{\xi'}$. \square

Lemma 3.11. *If $\mathcal{I} \models C$ for some interpretation \mathcal{I} and some clause C , then $\mathcal{I} \models C\rho$ for all substitutions ρ .*

Proof. We need to show that $C\rho$ is true in \mathcal{I} for all valuations ξ . Given a valuation ξ , define ξ' as in Lemma 3.10. Then, by Lemma 3.10, a literal in $C\rho$ is true in \mathcal{I} for ξ if and only if the corresponding literal in C is true in \mathcal{I} for ξ' . There must be at least one such literal because $\mathcal{I} \models C$ and hence C is in particular true in \mathcal{I} for ξ' . Therefore, $C\rho$ is true in \mathcal{I} for ξ . \square

Theorem 3.12 (Soundness of the intensional calculi). *The inference rules SUP, ERES, EFACT, and ARGCONG are sound (even without the variable condition and the side conditions on order and eligibility).*

Proof. We fix an inference and an interpretation \mathcal{I} that is a model of the premises. We need to show that it is also a model of the conclusion.

From the definition of the denotation of a term, it is obvious that congruence holds at all subterms in our logic. By Lemma 3.11, \mathcal{I} is a model of the σ -instances of the premises as well, where σ is the substitution used for the inference. Fix a valuation ξ . By making case distinctions on the truth in \mathcal{I} under ξ of the literals of the σ -instances of the premises, using the conditions that σ is a unifier, and applying congruence, it follows that the conclusion is also true in \mathcal{I} under ξ . \square

Theorem 3.13 (Soundness of the extensional calculi). *The inference rules SUP, ERES, EFACT, ARGCONG, and POSEXT are sound w.r.t. extensional interpretations (even without the variable condition and the side conditions on order and eligibility).*

Proof. We only need to prove POSEXT sound. For the other rules, we can proceed as in Theorem 3.12. By induction on the length of \bar{x} , it suffices to prove POSEXT sound for one variable x instead of a tuple \bar{x} . We fix an inference and an extensional interpretation \mathcal{I} that is a model of the premise $C' \vee s x \approx s' x$. We need to show that it is also a model of the conclusion $C' \vee s \approx s'$.

Let ξ be a valuation. If C' is true in \mathcal{I} under ξ , the conclusion is clearly true as well. Otherwise C' is false in \mathcal{I} under ξ , and also under $\xi[x \mapsto a]$ for all a because x does not occur in C' . Since the premise is true in \mathcal{I} , $s x = s' x$ must be true in \mathcal{I} under $\xi[x \mapsto a]$ for all a . Hence, for appropriate universes U_1, U_2 , we have $\mathcal{E}_{U_1, U_2}(\llbracket s \rrbracket_{\mathcal{I}}^{\xi[x \mapsto a]})(a) = \llbracket s x \rrbracket_{\mathcal{I}}^{\xi[x \mapsto a]} = \llbracket s' x \rrbracket_{\mathcal{I}}^{\xi[x \mapsto a]} = \mathcal{E}_{U_1, U_2}(\llbracket s' \rrbracket_{\mathcal{I}}^{\xi[x \mapsto a]})(a)$. Since s and s' do not contain x , $\llbracket s \rrbracket_{\mathcal{I}}^{\xi[x \mapsto a]}$ and $\llbracket s' \rrbracket_{\mathcal{I}}^{\xi[x \mapsto a]}$ do not depend on a . Thus, $\mathcal{E}_{U_1, U_2}(\llbracket s \rrbracket_{\mathcal{I}}^{\xi}) = \mathcal{E}_{U_1, U_2}(\llbracket s' \rrbracket_{\mathcal{I}}^{\xi})$. Since \mathcal{I} is extensional, \mathcal{E}_{U_1, U_2} is injective and hence $\llbracket s \rrbracket_{\mathcal{I}}^{\xi} = \llbracket s' \rrbracket_{\mathcal{I}}^{\xi}$. It follows that $s \approx s'$ is true in \mathcal{I} under ξ , and so is the entire conclusion of the inference. \square

A problem expressed in higher-order logic must be transformed into clausal normal form before the calculi can be applied. This process works as in the first-order case, except for skolemization. The issue is that skolemization, when performed naively, is unsound for higher-order logic with a Henkin semantics [Mil87, Section 6], because it introduces new functions that can be used to instantiate variables.

The core of this article is not affected by this because the problems are given in clausal form. For the implementation, we claim soundness only w.r.t. models that satisfy the axiom of choice, which is the semantics mandated by the TPTP THF format [SBBT09]. By contrast, refutational completeness holds w.r.t. arbitrary models as defined above. Alternatively, skolemization can be made sound by introducing mandatory arguments as described by Miller [Mil87, Section 6] and in the conference version of this article [BBCW18].

This issue also affects axiom (EXT) because it contains the Skolem symbol `diff`. As a consequence, (EXT) does not hold in all extensional interpretations. The extensional calculi are thus only sound w.r.t. interpretations in which (EXT) holds. However, we can prove that (EXT) is compatible with our logic:

Theorem 3.14. *Axiom (EXT) is satisfiable.*

Proof. For a given signature, let $(\mathcal{U}, \mathcal{J}_{\text{ty}}, \mathcal{J}, \mathcal{E})$ be an Herbrand interpretation. That is, we define \mathcal{U} to contain the set \mathcal{U}_τ of all terms of type τ for each ground type τ , we define \mathcal{J}_{ty} by $\mathcal{J}_{\text{ty}}(\kappa)(\bar{\tau}) = \kappa(\bar{\tau})$, we define \mathcal{J} by $\mathcal{J}(f, \mathcal{U}_{\bar{\tau}}) = f\langle\bar{\tau}\rangle$, and we define \mathcal{E} by $\mathcal{E}_{U_\tau, U_v}(f)(\mathbf{a}) = f \mathbf{a}$. Then $\mathcal{E}_{U_\tau, U_v}$ is clearly injective and hence \mathcal{I} is extensional. To show that $\mathcal{I} \models (\text{EXT})$, we need to show that (EXT) is true under all valuations. Let ξ be a valuation. If $x \approx y$ is true under ξ , (EXT) is also true. Otherwise $x \approx y$ is false under ξ , and hence $\xi(x) \neq \xi(y)$. Then we have $\llbracket x (\text{diff}\langle\alpha, \beta\rangle x y) \rrbracket_{\mathcal{I}}^\xi = (\xi(x)) (\text{diff}\langle\alpha, \beta\rangle (\xi(x)) (\xi(y))) \neq (\xi(y)) (\text{diff}\langle\alpha, \beta\rangle (\xi(x)) (\xi(y))) = \llbracket y (\text{diff}\langle\alpha, \beta\rangle x y) \rrbracket_{\mathcal{I}}^\xi$. Therefore, $x (\text{diff}\langle\alpha, \beta\rangle x y) \not\approx y (\text{diff}\langle\alpha, \beta\rangle x y)$ is true in \mathcal{I} under ξ and so is (EXT). \square

3.4. The Redundancy Criterion. For our calculi, a redundant clause cannot simply be defined as a clause whose ground instances are entailed by smaller (\prec) ground instances of existing clauses, because this would make all ARGCONG inferences redundant. Our solution is to base the redundancy criterion on a weaker ground logic—ground monomorphic first-order logic—in which argument congruence does not hold. This logic also plays a central role in our refutational completeness proof.

We employ an encoding \mathcal{F} to translate ground λ -free higher-order terms into ground first-order terms. It indexes each symbol occurrence with its type arguments and its term argument count. Thus, $\mathcal{F}(f) = f_0$, $\mathcal{F}(f \mathbf{a}) = f_1(\mathbf{a}_0)$, and $\mathcal{F}(g\langle\kappa\rangle) = \mathbf{g}_0^\kappa$. This is enough to disable argument congruence; for example, $\{f \approx h, f \mathbf{a} \not\approx h \mathbf{a}\}$ is unsatisfiable, whereas its encoding $\{f_0 \approx h_0, f_1(\mathbf{a}_0) \not\approx h_1(\mathbf{a}_0)\}$ is satisfiable. For clauses built from fully applied ground terms, the two logics are isomorphic, as we would expect from a graceful generalization.

Given a λ -free higher-order signature $(\Sigma_{\text{ty}}, \Sigma)$, we define a first-order signature $(\Sigma_{\text{ty}}, \Sigma_{\text{GF}})$ as follows. The type constructors Σ_{ty} are the same in both signatures, but \rightarrow is uninterpreted in first-order logic. For each symbol $f : \prod_{\bar{\alpha}_m} \tau_1 \rightarrow \dots \rightarrow \tau_n \rightarrow \tau$ in Σ , where τ is not functional, we introduce symbols $f_l^{\bar{v}_m} \in \Sigma_{\text{GF}}$ with argument types $\bar{\tau}_l \sigma$ and return type $(\tau_{l+1} \rightarrow \dots \rightarrow \tau_n \rightarrow \tau) \sigma$, where $\sigma = \{\bar{\alpha}_m \mapsto \bar{v}_m\}$, for each tuple of ground types \bar{v}_m and each $l \in \{0, \dots, n\}$.

For example, let $\Sigma = \{\mathbf{a} : \kappa, \mathbf{g} : \kappa \rightarrow \kappa \rightarrow \kappa\}$. The corresponding first-order signature is $\Sigma_{\text{GF}} = \{\mathbf{a}_0 : \kappa, \mathbf{g}_0 : \kappa \rightarrow \kappa \rightarrow \kappa, \mathbf{g}_1 : \kappa \Rightarrow \kappa \rightarrow \kappa, \mathbf{g}_2 : \kappa^2 \Rightarrow \kappa\}$ where $f : \bar{\tau} \Rightarrow v$ denotes a first-order function symbol f with argument types $\bar{\tau}$ and return type v , and \rightarrow is an uninterpreted binary type constructor. The term $\mathcal{F}(g \mathbf{a} \mathbf{a}) = \mathbf{g}_2(\mathbf{a}_0, \mathbf{a}_0)$ has type κ , and $\mathcal{F}(g \mathbf{a}) = \mathbf{g}_1(\mathbf{a}_0)$ has type $\kappa \rightarrow \kappa$.

Thus, we consider three levels of logics: the λ -free higher-order level H over a given signature $(\Sigma_{\text{ty}}, \Sigma)$, the ground λ -free higher-order level GH, corresponding to H's ground fragment, and the ground monomorphic first-order level GF over the signature $(\Sigma_{\text{ty}}, \Sigma_{\text{GF}})$ defined above. We use \mathcal{T}_H , \mathcal{T}_{GH} , and \mathcal{T}_{GF} to denote the respective sets of terms, \mathcal{T}_{y}_H , $\mathcal{T}_{\text{y}}_{\text{GH}}$, and $\mathcal{T}_{\text{y}}_{\text{GF}}$ to denote the respective sets of types, and \mathcal{C}_H , \mathcal{C}_{GH} , and \mathcal{C}_{GF} to denote the respective sets of clauses. In the purifying calculi, we exceptionally let \mathcal{C}_H denote the set of purified clauses. Each of the three levels has an entailment relation \models . A clause set N_1 entails a clause set N_2 , denoted $N_1 \models N_2$, if any model of N_1 is also a model of N_2 . On H and GH, we use λ -free higher-order models for the intensional calculi and extensional λ -free higher-order models for the extensional calculi; on GF, we use first-order models. This machinery may seem excessive, but it is essential to define redundancy of clauses and inferences properly, and it will play an important role in the refutational completeness proof (Section 4).

The three levels are connected by two functions, \mathcal{G} and \mathcal{F} :

Definition 3.15 (Grounding function \mathcal{G} on terms and clauses). The grounding function \mathcal{G} maps terms $t \in \mathcal{T}_H$ to the set of their ground instances—i.e., the set of all $t\theta \in \mathcal{T}_{GH}$ where θ is a substitution. It also maps clauses $C \in \mathcal{C}_H$ to the set of their ground instances—i.e., the set of all $C\theta \in \mathcal{C}_{GH}$ where θ is a substitution.

Definition 3.16 (Encoding \mathcal{F} on terms and clauses). The encoding $\mathcal{F} : \mathcal{T}_{GH} \rightarrow \mathcal{T}_{GF}$ is recursively defined as $\mathcal{F}(f\langle \bar{v}_m \rangle \bar{u}_l) = f_l^{\bar{v}_m}(\mathcal{F}(\bar{u}_l))$. The encoding \mathcal{F} is extended to map from \mathcal{C}_{GH} to \mathcal{C}_{GF} by mapping each literal and each side of a literal individually.

The encoding \mathcal{F} is bijective with inverse \mathcal{F}^{-1} . Using \mathcal{F}^{-1} , the clause order \succ on \mathcal{T}_{GH} can be transferred to \mathcal{T}_{GF} by defining $t \succ s$ as equivalent to $\mathcal{F}^{-1}(t) \succ \mathcal{F}^{-1}(s)$. The property that \succ on clauses is the multiset extension of \succ on literals, which in turn is the multiset extension of \succ on terms, is maintained because \mathcal{F}^{-1} maps the multiset representations elementwise.

Schematically, the three levels are connected as follows:

$$\begin{array}{ccccc} H & \xrightarrow{\mathcal{G}} & GH & \xrightarrow{\mathcal{F}} & GF \\ \text{higher-order} & & \text{ground higher-order} & & \text{ground first-order} \end{array}$$

Crucially, green subterms in \mathcal{T}_{GH} correspond to subterms in \mathcal{T}_{GF} (Lemma 3.17), whereas nongreen subterms in \mathcal{T}_{GH} are not subterms at all in \mathcal{T}_{GF} . For example, \mathbf{a} is a green subterm of $\mathbf{f} \mathbf{a}$, and correspondingly $\mathcal{F}(\mathbf{a}) = \mathbf{a}_0$ is a subterm of $\mathcal{F}(\mathbf{f} \mathbf{a}) = \mathbf{f}_1(\mathbf{a}_0)$. On the other hand, \mathbf{f} is not a green subterm of $\mathbf{f} \mathbf{a}$, and correspondingly $\mathcal{F}(\mathbf{f}) = \mathbf{f}_0$ is not a subterm of $\mathcal{F}(\mathbf{f} \mathbf{a}) = \mathbf{f}_1(\mathbf{a}_0)$.

To state the correspondence between green subterms in \mathcal{T}_{GH} and subterms in \mathcal{T}_{GF} explicitly, we define positions of green subterms as follows. A term $s \in \mathcal{T}_H$ is a green subterm at position ϵ of s . If a term $s \in \mathcal{T}_H$ is a green subterm at position p of u_i for some $1 \leq i \leq n$, then s is a green subterm at position $i.p$ of $\mathbf{f}\langle \bar{\tau} \rangle \bar{u}_n$ and of $x \bar{u}_n$. For \mathcal{T}_{GF} , positions are defined as usual in first-order logic.

Lemma 3.17. *Let $s, t \in \mathcal{T}_{GH}$. We have $\mathcal{F}(t\langle s \rangle_p) = \mathcal{F}(t)[\mathcal{F}(s)]_p$. In other words, s is a green subterm of t at position p if and only if $\mathcal{F}(s)$ is a subterm of $\mathcal{F}(t)$ at position p .*

Proof. By induction on p . If $p = \epsilon$, then $s = t[s]_p$. Hence $\mathcal{F}(t[s]_p) = \mathcal{F}(s) = \mathcal{F}(t)\langle \mathcal{F}(s) \rangle_p$. If $p = i.p'$, then $t[s]_p = \mathbf{f}\langle \bar{\tau} \rangle \bar{u}_n$ with $u_i = u_i[s]_{p'}$. Applying \mathcal{F} , we obtain by the induction hypothesis that $\mathcal{F}(u_i) = \mathcal{F}(u_i)\langle \mathcal{F}(s) \rangle_{p'}$. Therefore, $\mathcal{F}(t[s]_p) = \mathbf{f}_n^{\bar{\tau}}(\mathcal{F}(u_1), \dots, \mathcal{F}(u_{i-1}), \mathcal{F}(u_i)\langle \mathcal{F}(s) \rangle_{p'}, \mathcal{F}(u_{i+1}), \dots, \mathcal{F}(u_n))$. It follows that $\mathcal{F}(t[s]_p) = \mathcal{F}(t)\langle \mathcal{F}(s) \rangle_p$. \square

Corollary 3.18. *Given $s, t \in \mathcal{T}_{GF}$, we have $\mathcal{F}^{-1}(t[s]_p) = \mathcal{F}^{-1}(t)\langle \mathcal{F}^{-1}(s) \rangle_p$.*

Lemma 3.19. *Well-foundedness, totality, compatibility with contexts, and the subterm property hold for \succ on \mathcal{T}_{GF} .*

Proof. COMPATIBILITY WITH CONTEXTS: We must show that $s \succ s'$ implies $t[s]_p \succ t[s']_p$ for terms $t, s, s' \in \mathcal{T}_{GF}$. Assuming $s \succ s'$, we have $\mathcal{F}^{-1}(s) \succ \mathcal{F}^{-1}(s')$. By compatibility with green contexts on \mathcal{T}_{GH} , it follows that $\mathcal{F}^{-1}(t)\langle \mathcal{F}^{-1}(s) \rangle_p \succ \mathcal{F}^{-1}(t)\langle \mathcal{F}^{-1}(s') \rangle_p$. By Corollary 3.18, we have $t[s]_p \succ t[s']_p$.

WELL-FOUNDEDNESS: Assume that there exists an infinite descending chain $t_1 \succ t_2 \succ \dots$ in \mathcal{T}_{GF} . By applying \mathcal{F}^{-1} , we then obtain the chain $\mathcal{F}^{-1}(t_1) \succ \mathcal{F}^{-1}(t_2) \succ \dots$ in \mathcal{T}_{GH} , contradicting well-foundedness on \mathcal{T}_{GH} .

TOTALITY: Let $s, t \in \mathcal{T}_{\text{GF}}$. Then $\mathcal{F}^{-1}(t)$ and $\mathcal{F}^{-1}(s)$ must be comparable by totality in \mathcal{T}_{GH} . Hence, t and s are comparable.

SUBTERM PROPERTY: By Corollary 3.18 and the subterm property on \mathcal{T}_{GH} , we have $\mathcal{F}^{-1}(t[s]_p) = \mathcal{F}^{-1}(t) \langle \mathcal{F}^{-1}(s) \rangle_p \succ \mathcal{F}^{-1}(s)$. Hence, $t[s]_p \succ s$. \square

In standard superposition, redundancy relies on the entailment relation \models on ground clauses. We will define redundancy on GH and H in the same way, but using GF's entailment relation. This notion of redundancy gracefully generalizes the first-order notion, without making all ARGCONG inferences redundant.

The standard redundancy criterion for standard superposition cannot justify subsumption deletion. Following Waldmann et al. [WTRB20], we incorporate subsumption into the redundancy criterion. A clause C *subsumes* D if there exists a substitution σ such that $C\sigma \subseteq D$. A clause C *strictly subsumes* D if C subsumes D but D does not subsume C . Let \sqsupset stand for “is strictly subsumed by”. Using the applicative encoding, it is easy to show that \sqsupset is well founded because strict subsumption is well founded in first-order logic.

We define the sets of redundant clauses w.r.t. a given clause set as follows:

- Given $C \in \mathcal{C}_{\text{GF}}$ and $N \subseteq \mathcal{C}_{\text{GF}}$, let $C \in \text{GFRed}_C(N)$ if $\{D \in N \mid D \prec C\} \models C$.
- Given $C \in \mathcal{C}_{\text{GH}}$ and $N \subseteq \mathcal{C}_{\text{GH}}$, let $C \in \text{GHRed}_C(N)$ if $\mathcal{F}(C) \in \text{GFRed}_C(\mathcal{F}(N))$.
- Given $C \in \mathcal{C}_{\text{H}}$ and $N \subseteq \mathcal{C}_{\text{H}}$, let $C \in \text{HRed}_C(N)$ if for every $D \in \mathcal{G}(C)$, we have $D \in \text{GHRed}_C(\mathcal{G}(N))$ or there exists $C' \in N$ such that $C \sqsupset C'$ and $D \in \mathcal{G}(C')$.

Along with the three levels of logics, we consider three inference systems: HInf , GHInf and GFInf . HInf is one of the four variants of the inference system described in Section 3.1. For uniformity, we regard axiom (EXT) as a premise-free inference rule EXT whose conclusion is (EXT). In the purifying calculi, the conclusion of EXT must be purified. GHInf consists of all SUP, ERES, and EFACT inferences from HInf whose premises and conclusion are ground, a premise-free rule GEXT whose infinitely many conclusions are the ground instances of (EXT), and the following ground variant of ARGCONG:

$$\frac{C' \vee s \approx s'}{C' \vee s \bar{u}_n \approx s' \bar{u}_n} \text{GARGCONG}$$

where $s \approx s'$ is strictly eligible in the premise and \bar{u}_n is a nonempty tuple of ground terms. GFInf contains all SUP, ERES, and EFACT inferences from GHInf translated by \mathcal{F} . It coincides exactly with standard first-order superposition. Given a SUP, ERES, or EFACT inference $\iota \in \text{GHInf}$, let $\mathcal{F}(\iota)$ denote the corresponding inference in GFInf .

Each of the three inference systems is parameterized by a selection function. For HInf , we globally fix a selection function HSel . For GHInf and GFInf , we need to consider different selection functions GHSel and GFSel . We write $\text{GHInf}^{\text{GHSel}}$ for GHInf and $\text{GFInf}^{\text{GFSel}}$ for GFInf to make the dependency on the respective selection functions explicit. Let $\mathcal{G}(\text{HSel})$ denote the set of all selection functions on \mathcal{C}_{GH} such that for each clause in $C \in \mathcal{C}_{\text{GH}}$, there exists a clause $D \in \mathcal{C}_{\text{H}}$ with $C \in \mathcal{G}(D)$ and corresponding selected literals. For each selection function GHSel on \mathcal{C}_{GH} , via the bijection \mathcal{F} , we obtain a corresponding selection function on \mathcal{C}_{GF} , which we denote by $\mathcal{F}(\text{GHSel})$.

Notation 3.20. Given an inference ι , we write $\text{prems}(\iota)$ for the tuple of premises, $\text{mprem}(\iota)$ for the main (i.e., rightmost) premise, $\text{preconcl}(\iota)$ for the conclusion before purification, and $\text{concl}(\iota)$ for the conclusion after purification. For the nonpurifying calculi, $\text{preconcl}(\iota) = \text{concl}(\iota)$ simply denotes the conclusion.

Definition 3.21 (Encoding \mathcal{F} on inferences). Given a SUP, ERES, or EFACT inference $\iota \in GHI\text{nf}$, let $\mathcal{F}(\iota) \in GF\text{Inf}$ denote the inference defined by $\text{prems}(\mathcal{F}(\iota)) = \mathcal{F}(\text{prems}(\iota))$ and $\text{concl}(\mathcal{F}(\iota)) = \mathcal{F}(\text{concl}(\iota))$.

Definition 3.22 (Grounding function \mathcal{G} on inferences). Given a selection function $GHSel \in \mathcal{G}(HSel)$, and a non-POSEXT inference $\iota \in HI\text{nf}$, we define the set $\mathcal{G}^{GHSel}(\iota)$ of ground instances of ι to be all inferences $\iota' \in GHI\text{nf}^{GHSel}$ such that $\text{prems}(\iota') = \text{prems}(\iota)\theta$ and $\text{concl}(\iota') = \text{preconcl}(\iota)\theta$ for some grounding substitution θ . For POSEXT inferences ι , which cannot be grounded, we let $\mathcal{G}^{GHSel}(\iota) = \text{undef}$.

This will map SUP to SUP, EFACT to EFACT, ERES to ERES, EXT to GEXT, and ARGCONG to GARGCONG inferences, but it is also possible that $\mathcal{G}^{GHSel}(\iota)$ is the empty set for some inferences ι .

We define the sets of redundant inferences w.r.t. a given clause set as follows:

- Given $\iota \in GF\text{Inf}^{GF\text{Sel}}$ and $N \subseteq \mathcal{C}_{GF}$, let $\iota \in GF\text{Red}_I^{GF\text{Sel}}(N)$ if $\text{prems}(\iota) \cap GF\text{Red}_C(N) \neq \emptyset$ or $\{D \in N \mid D \prec \text{mprem}(\iota)\} \models \text{concl}(\iota)$.
- Given $\iota \in GHI\text{nf}^{GHSel}$ and $N \subseteq \mathcal{C}_{GH}$, let $\iota \in GH\text{Red}_I^{GHSel}(N)$ if
 - ι is not a GARGCONG or GEXT inference and $\mathcal{F}(\iota) \in GF\text{Red}_I^{\mathcal{F}(GHSel)}(\mathcal{F}(N))$; or
 - the calculus is nonpurifying and ι is a GARGCONG or GEXT inference and $\text{concl}(\iota) \in N \cup GH\text{Red}_C(N)$; or
 - the calculus is purifying and ι is a GARGCONG or GEXT inference and $\mathcal{F}(N) \models \mathcal{F}(\text{concl}(\iota))$.
- Given $\iota \in HI\text{nf}$ and $N \subseteq \mathcal{C}_H$, let $\iota \in H\text{Red}_I(N)$ if
 - ι is not a POSEXT inference and $\mathcal{G}^{GHSel}(\iota) \subseteq GH\text{Red}_I(\mathcal{G}(N))$ for all $GHSel \in \mathcal{G}(HSel)$;
 - or
 - ι is a POSEXT inference and $\mathcal{G}(\text{concl}(\iota)) \subseteq \mathcal{G}(N) \cup GH\text{Red}_C(\mathcal{G}(N))$.

Occasionally, we omit the selection function in the notation when it is irrelevant. A clause set N is *saturated* w.r.t. an inference system and a redundancy criterion (Red_I, Red_C) if every inference from clauses in N is in $Red_I(N)$.

3.5. Simplification Rules. The redundancy criterion ($HRed_I, HRed_C$) is strong enough to support most of the simplification rules implemented in Schulz’s first-order prover E [Sch02, Sections 2.3.1 and 2.3.2], some only with minor adaptations. Deletion of duplicated literals, deletion of resolved literals, syntactic tautology deletion, negative simplify-reflect, and clause subsumption adhere to our redundancy criterion.

Positive simplify-reflect and equality subsumption are supported by our criterion if they are applied in green contexts. Semantic tautology deletion can be applied as well, but we must use the entailment relation of the GF level—i.e., only rewriting in green contexts can be used to establish the entailment. Similarly, rewriting of positive and negative literals (demodulation) can only be applied in green contexts. Moreover, for positive literals, the rewriting clause must be smaller than the rewritten clause—a condition that is also necessary with the standard first-order redundancy criterion but not always fulfilled by Schulz’s rule. As for destructive equality resolution, even in first-order logic the rule cannot be justified with the standard redundancy criterion, and it is unclear whether it preserves refutational completeness.

4. REFUTATIONAL COMPLETENESS

Besides soundness, the most important property of the four calculi introduced in Section 3.1 is refutational completeness. We will prove the static and dynamic refutational completeness of $HInf$ w.r.t. $(HRed_I, HRed_C)$, which is defined as follows:

Definition 4.1 (Static refutational completeness). Let Inf be an inference system and let (Red_I, Red_C) be a redundancy criterion. The inference system Inf is called *statically refutationally complete* w.r.t. (Red_I, Red_C) if we have $N \models \perp$ if and only if $\perp \in N$ for every clause set N that is saturated w.r.t. Inf and Red_I .

Definition 4.2 (Dynamic refutational completeness). Let Inf be an inference system and let (Red_I, Red_C) be a redundancy criterion. Let $(N_i)_i$ be a finite or infinite sequence over sets of clauses. Such a sequence is called a *derivation* if $N_i \setminus N_{i+1} \subseteq Red_C(N_{i+1})$ for all i . It is called *fair* if all Inf -inferences from clauses in the limit inferior $\bigcup_i \bigcap_{j \geq i} N_j$ are contained in $\bigcup_i Red_I(N_i)$. The inference system Inf is called *dynamically refutationally complete* w.r.t. (Red_C, Red_I) if for every fair derivation $(N_i)_i$ such that $N_0 \models \perp$, we have $\perp \in N_i$ for some i .

To circumvent the term order's potential nonmonotonicity, our SUP inference rule only considers green subterms. This is reflected in our proof by the reliance on the GF level introduced in Section 3.4. The equation $g_0 \approx f_0 \in C_{GF}$, which corresponds to the equation $g \approx f \in C_{GH}$, cannot be used directly to rewrite the clause $g_1(a_0) \not\approx f_1(a_0) \in C_{GF}$, which corresponds to $g a \not\approx f a \in C_{GH}$. Instead, we first need to apply ARGCONG to derive $g x \approx f x$, which after grounding and transfer to GF yields $g_1(a_0) \approx f_1(a_0)$. The GF level is a device that enables us to reuse the refutational completeness result for standard (first-order) superposition.

The proof proceeds in three steps, corresponding to the three levels GF, GH, and H introduced in Section 3.4:

- (1) We use Bachmair and Ganzinger's work on the refutational completeness of standard superposition [BG94] to prove the static refutational completeness of $GFIInf$.
- (2) From the first-order model constructed in Bachmair and Ganzinger's proof, we derive a λ -free higher-order model to prove the static refutational completeness of $GHIInf$.
- (3) We use the saturation framework of Waldmann et al. [WTRB20] to lift the static refutational completeness of $GHIInf$ to static and dynamic refutational completeness of $HInf$.

In the first step, since the inference system $GFIInf$ is standard ground superposition, we only need to work around minor differences between Bachmair and Ganzinger's definitions and ours. Given a saturated clause set $N \subseteq C_{GF}$ with $\perp \notin N$, Bachmair and Ganzinger prove refutational completeness by constructing a term rewriting system R_N and showing that it can be viewed as an interpretation that is a model of N . This step is exclusively concerned with ground first-order clauses.

In the second step, we derive refutational completeness of $GHIInf$. Given a saturated clause set $N \subseteq C_{GH}$ with $\perp \notin N$, we use the first-order model $R_{\mathcal{F}(N)}$ of $\mathcal{F}(N)$ constructed in step (1) to derive a clausal higher-order interpretation that is a model of N . Thanks to saturation w.r.t. GARGCONG, the higher-order interpretation can conflate the interpretations of the members $f_0^{\vec{v}}, \dots, f_n^{\vec{v}}$ of a same symbol family. In the extensional calculi, saturation w.r.t. GEXT can be used to show that the constructed interpretation is extensional.

In the third step, we employ the saturation framework by Waldmann et al. [WTRB20], which is largely based on Bachmair and Ganzinger's [BG01], to prove $HInf$ refutationally

complete. Like Bachmair and Ganzinger’s, the saturation framework allows us to prove the static and dynamic refutational completeness of our calculus on the nonground level. On top of that, it allows us to use the redundancy criterion defined in Section 3.4, which supports deletion of subsumed formulas. Moreover, the saturation framework provides completeness theorems for prover architectures, such as the DISCOUNT loop. The main proof obligation the saturation framework leaves to us is that there exist inferences in $HInf$ corresponding to all nonredundant inferences in $GHInf$. For monotone term orders, we can avoid SUP inferences into variables x by exploiting the clause order’s compatibility with contexts: If $t' \prec t$, we have $C\{x \mapsto t'\} \prec C\{x \mapsto t\}$, which allows us to show that SUP inferences into variables are redundant. This technique fails for variables x that occur applied in C , because the order lacks compatibility with arguments. This is why the calculi must either purify clauses to make this line of reasoning work again or perform some SUP inferences into variables.

4.1. The Ground First-Order Level. We use Bachmair and Ganzinger’s results on standard superposition [BG94] to prove GF refutationally complete. In the subsequent steps, we will also make use of specific properties of Bachmair and Ganzinger’s model.

Bachmair and Ganzinger’s logic and inference system differ in some details from GF:

- Bachmair and Ganzinger use untyped first-order logic, whereas GF’s logic is typed. Bachmair and Ganzinger’s proof works verbatim for monomorphic first-order logic as well, but we need to require that the order \succ has the subterm property to show that there exist no critical pairs in the term rewriting system, as observed by Wand [Wan17, Section 3.2.1].
- In their redundancy criterion for clauses, Bachmair and Ganzinger require that a finite subset of $\{D \in N \mid D \prec C\}$ entails C , whereas we require that $\{D \in N \mid D \prec C\}$ entails C . By compactness of first-order logic, the two criteria are equivalent.

Bachmair and Ganzinger prove refutational completeness for nonground clause sets, but we only require the ground result here.

The basis of Bachmair and Ganzinger’s proof is that a term rewriting system R defines an interpretation \mathcal{T}_{GF}/R such that for every ground equation $s \approx t$, we have $\mathcal{T}_{GF}/R \models s \approx t$ if and only if $s \leftrightarrow_R^* t$. Formally, \mathcal{T}_{GF}/R denotes the monomorphic first-order interpretation whose universes U_τ consist of the R -equivalence classes over \mathcal{T}_{GF} containing terms of type τ . The interpretation \mathcal{T}_{GF}/R is term-generated—that is, for every element a of the universe of this interpretation and for any valuation ξ , there exists a ground term t such that $\llbracket t \rrbracket_{\mathcal{T}_{GF}/R}^\xi = a$. To lighten notation, we will write R to refer to both the term rewriting system \bar{R} and the interpretation \mathcal{T}_{GF}/R .

The term rewriting system is constructed as follows. Let $N \subseteq \mathcal{C}_{GF}$. We first define sets of rewrite rules E_N^C and R_N^C for all $C \in N$ by induction on the clause order. Assume that E_N^D has already been defined for all $D \in N$ such that $D \prec C$. Then $R_N^C = \bigcup_{D \prec C} E_N^D$. Let $E_N^C = \{s \rightarrow t\}$ if the following conditions are met:

- (a) $C = C' \vee s \approx t$;
- (b) $s \approx t$ is strictly maximal in C ;
- (c) $s \succ t$;
- (d) C' is false in R_N^C ;
- (e) s is irreducible w.r.t. R_N^C .

Then C is said to *produce* $s \rightarrow t$. Otherwise, $E_N^C = \emptyset$. Finally, $R_N = \bigcup_D E_N^D$.

We call an inference $\iota \in GFInf$ *B&G-redundant* if some $C \in \text{prems}(\iota)$ is true in R_N^C or $\text{concl}(\iota)$ is true in $R_N^{\text{mprem}(\iota)}$. We call a set $N \subseteq C_{GF}$ *B&G-saturated* if all inferences from N are B&G-redundant.

Lemma 4.3. *If $\perp \notin N$ and $N \subseteq C_{GF}$ is saturated w.r.t. $GFInf$ and $GFRed_I$, then N is B&G-saturated.*

Proof. Let $N \subseteq C_{GF}$ be saturated w.r.t. $GFInf$ and $GFRed_I$. To show that N is B&G-saturated, let ι be an inference from N . We need to show that ι is B&G-redundant w.r.t. N . We proceed by well-founded induction on $\text{mprem}(\iota)$ w.r.t. \succ . By the induction hypothesis, for all inferences ι' with $\text{concl}(\iota') \prec \text{mprem}(\iota)$, ι' is B&G-redundant w.r.t. N . By Lemma 5.5 of Bachmair and Ganzinger, ι is B&G-redundant w.r.t. N . \square

Lemma 4.4. *Let $\perp \notin N$ and $N \subseteq C_{GF}$ be saturated w.r.t. $GFInf$ and $GFRed_I$. If $C = C' \vee s \approx t \in N$ produces $s \rightarrow t$, then $s \approx t$ is strictly eligible in C and C' is false in R_N .*

Proof. By Lemma 4.3, N is also B&G-saturated. By condition (d), C' is false in R_N^C . Since $s \succ t$ by condition (c) and s is irreducible w.r.t. R_N^C by condition (e), $s \approx t$ is also false in R_N^C . Hence, C is false in R_N^C . Using this and conditions (a), (b), (c), and (e), we can apply Lemma 4.11 of Bachmair and Ganzinger using N for N and for N' , C for C and for D , s for s , and t for t . Part (ii) of that lemma shows that $s \approx t$ is strictly eligible in C , and part (iv) shows that C' is false in R_N . \square

Theorem 4.5 (Ground first-order static refutational completeness). *The inference system $GFInf$ is statically refutationally complete w.r.t. $(GFRed_I, GFRed_C)$. More precisely, if $N \subseteq C_{GF}$ is a clause set saturated w.r.t. $GFInf$ and $GFRed_I$ such that $\perp \notin N$, then R_N is a model of N .*

Proof. By Lemma 4.3, N is also B&G-saturated. It follows that R_N is a model of N , as shown in the proof of Theorem 4.14 of Bachmair and Ganzinger. \square

4.2. The Ground Higher-Order Level. In this subsection, let $GHSel$ be a selection function on C_{GH} , and let $N \subseteq C_{GH}$ with $\perp \notin N$ be a clause set saturated w.r.t. $GHIInf^{GHSel}$ and $GHRed_I^{GHSel}$. Clearly, $\mathcal{F}(N)$ is then saturated w.r.t. $GFInf^{\mathcal{F}(GHSel)}$ and $GFRed_I^{\mathcal{F}(GHSel)}$.

We abbreviate $R_{\mathcal{F}(N)}$ as R . From R , we construct a model \mathcal{I}^{GH} of N . The key properties enabling us to perform this construction are that R is term-generated and that the interpretations of the members $f_0^{\vec{v}}, \dots, f_n^{\vec{v}}$ of a same symbol family behave in the same way thanks to the ARGCONG rule.

Lemma 4.6 (Argument congruence). *For terms $s, t, u \in \mathcal{T}_{GH}$, if $\llbracket \mathcal{F}(s) \rrbracket_R = \llbracket \mathcal{F}(t) \rrbracket_R$, then $\llbracket \mathcal{F}(s u) \rrbracket_R = \llbracket \mathcal{F}(t u) \rrbracket_R$.*

Proof. What we want to show is equivalent to

$$\mathcal{F}(s) \leftrightarrow_R^* \mathcal{F}(t) \text{ implies } \mathcal{F}(s u) \leftrightarrow_R^* \mathcal{F}(t u)$$

By induction on the number of rewrite steps and due to symmetry, it suffices to show that

$$\mathcal{F}(s) \rightarrow_R \mathcal{F}(t) \text{ implies } \mathcal{F}(s u) \leftrightarrow_R^* \mathcal{F}(t u)$$

If the rewrite step from $\mathcal{F}(s)$ is below the top level, this is obvious because there is a corresponding rewrite step from $\mathcal{F}(s u)$. If it is at the top level, $\mathcal{F}(s) \rightarrow \mathcal{F}(t)$ must be rule of R . This rule must originate from a productive clause of the form $\mathcal{F}(C) = \mathcal{F}(C' \vee s \approx t)$. By

Lemma 4.4, $\mathcal{F}(s \approx t)$ is strictly eligible in $\mathcal{F}(C)$ w.r.t. $\mathcal{F}(GHSel)$, and hence $s \approx t$ is strictly eligible in C w.r.t. $GHSel$. Moreover, s and t have functional type. Thus, the following GARGCONG inference ι is applicable:

$$\frac{C' \vee s \approx t}{C' \vee s u \approx t u} \text{GARGCONG}$$

By saturation, ι is redundant w.r.t. N —i.e., we have $\text{concl}(\iota) \in N \cup GHRed_C(N)$ (for the nonpurifying calculi) or $\mathcal{F}(N) \models \text{concl}(\iota)$ (for the purifying calculi). In both cases, by Theorem 4.5, $\mathcal{F}(\text{concl}(\iota))$ is then true in R . By Lemma 4.4, $\mathcal{F}(C')$ is false in R . Therefore, $\mathcal{F}(s u \approx t u)$ must be true in R . \square

Lemma 4.7. *For terms $s, t, u, v \in \mathcal{T}_{\text{GH}}$, if $\llbracket \mathcal{F}(s) \rrbracket_R = \llbracket \mathcal{F}(t) \rrbracket_R$ and $\llbracket \mathcal{F}(u) \rrbracket_R = \llbracket \mathcal{F}(v) \rrbracket_R$, then $\llbracket \mathcal{F}(s u) \rrbracket_R = \llbracket \mathcal{F}(t v) \rrbracket_R$.*

Proof. By Lemma 4.6, we have $\llbracket \mathcal{F}(s u) \rrbracket_R = \llbracket \mathcal{F}(t u) \rrbracket_R$. It remains to show that $\llbracket \mathcal{F}(t u) \rrbracket_R = \llbracket \mathcal{F}(t v) \rrbracket_R$. Since t is ground, it must be of the form $f\langle \bar{v}_m \rangle \bar{t}_n$. The interpretation R defined above is an interpretation (U, \mathcal{J}) in monomorphic first-order logic. Then

$$\llbracket \mathcal{F}(t u) \rrbracket_R = \mathcal{J}(f_{n+1}^{\bar{v}_m})(\llbracket \mathcal{F}(\bar{t}_n) \rrbracket_R, \llbracket \mathcal{F}(u) \rrbracket_R) = \mathcal{J}(f_{n+1}^{\bar{v}_m})(\llbracket \mathcal{F}(\bar{t}_n) \rrbracket_R, \llbracket \mathcal{F}(v) \rrbracket_R) = \llbracket \mathcal{F}(t v) \rrbracket_R \quad \square$$

Definition 4.8. Define a higher-order interpretation $\mathcal{I}^{\text{GH}} = (\mathcal{U}^{\text{GH}}, \mathcal{J}_{\text{ty}}^{\text{GH}}, \mathcal{J}^{\text{GH}}, \mathcal{E}^{\text{GH}})$ as follows. The interpretation R defined above is an interpretation (U, \mathcal{J}) in monomorphic first-order logic, where U_τ is its universe for type τ , and \mathcal{J} is its interpretation function. Let $\mathcal{U}^{\text{GH}} = \{U_\tau \mid \tau \text{ is a ground type}\}$. Let $\mathcal{J}_{\text{ty}}^{\text{GH}}(\kappa)(U_{\bar{\tau}}) = U_{\kappa(\bar{\tau})}$ for all type constructors κ and type tuples $\bar{\tau}$. Let $\mathcal{J}^{\text{GH}}(f, U_{\bar{\tau}}) = \mathcal{J}(f_{\bar{\tau}})$.

Since R is term-generated, for every $a \in U_{\tau \rightarrow v}$ and $b \in U_\tau$, there exist ground terms $s : \tau \rightarrow v$ and $u : \tau$ such that $\llbracket \mathcal{F}(s) \rrbracket_R = a$ and $\llbracket \mathcal{F}(u) \rrbracket_R = b$. We define \mathcal{E}^{GH} by $\mathcal{E}_{U_\tau, U_v}^{\text{GH}}(a)(b) = \llbracket \mathcal{F}(s u) \rrbracket_R$ for any term u . By Lemma 4.7, this definition is independent of the choice of s and u .

Lemma 4.9 (Model transfer to GH). *\mathcal{I}^{GH} is a model of N . In the extensional calculi, \mathcal{I}^{GH} is an extensional model of N .*

Proof. We first prove by induction on terms $t \in \mathcal{T}_{\text{GH}}$ that $\llbracket t \rrbracket_{\mathcal{I}^{\text{GH}}} = \llbracket \mathcal{F}(t) \rrbracket_R$. Let $t \in \mathcal{T}_{\text{GH}}$, and assume as the induction hypothesis that $\llbracket u \rrbracket_{\mathcal{I}^{\text{GH}}} = \llbracket \mathcal{F}(u) \rrbracket_R$ for all subterms u of t . If t is of the form $f\langle \bar{v} \rangle$, then

$$\llbracket t \rrbracket_{\mathcal{I}^{\text{GH}}} = \mathcal{J}^{\text{GH}}(f, U_{\bar{v}}) = \mathcal{J}(f_{\bar{v}}) = \llbracket f_{\bar{v}} \rrbracket_R = \llbracket \mathcal{F}(f\langle \bar{v} \rangle) \rrbracket_R = \llbracket \mathcal{F}(t) \rrbracket_R$$

If t is an application $t = t_1 t_2$, where t_1 is of type $\tau \rightarrow v$, then, using the definition of the term denotation and of \mathcal{E}^{GH} , we have

$$\llbracket t_1 t_2 \rrbracket_{\mathcal{I}^{\text{GH}}} = \mathcal{E}_{U_\tau, U_v}^{\text{GH}}(\llbracket t_1 \rrbracket_{\mathcal{I}^{\text{GH}}})(\llbracket t_2 \rrbracket_{\mathcal{I}^{\text{GH}}}) \stackrel{\text{IH}}{=} \mathcal{E}_{U_\tau, U_v}^{\text{GH}}(\llbracket \mathcal{F}(t_1) \rrbracket_R)(\llbracket \mathcal{F}(t_2) \rrbracket_R) = \llbracket \mathcal{F}(t_1 t_2) \rrbracket_R$$

So we have shown that $\llbracket t \rrbracket_{\mathcal{I}^{\text{GH}}} = \llbracket \mathcal{F}(t) \rrbracket_R$ for all terms t . It follows that a ground equation $s \approx t$ is true in \mathcal{I}^{GH} if and only if $\mathcal{F}(s \approx t)$ is true in R . Hence a ground clause C is true in \mathcal{I}^{GH} if and only if $\mathcal{F}(C)$ is true in R . By Theorem 4.5, R is a model of $\mathcal{F}(N)$. Thus, \mathcal{I}^{GH} is a model of N .

For the extensional calculi, it remains to show that \mathcal{I}^{GH} is extensional—i.e., we have to show that for all τ and v and all $a, b \in U_{\tau \rightarrow v}$, if $a \neq b$, then $\mathcal{E}^{\text{GH}}(a) \neq \mathcal{E}^{\text{GH}}(b)$. Since R is term-generated, there are terms $s, t \in \mathcal{T}_{\text{GF}}$ such that $\llbracket s \rrbracket_R = a$ and $\llbracket t \rrbracket_R = b$. By what

we have shown above, it follows that $\llbracket s' \rrbracket_{\mathcal{I}^{\text{GH}}} = a$ and $\llbracket t' \rrbracket_{\mathcal{I}^{\text{GH}}} = b$ for $s' = \mathcal{F}^{-1}(s)$ and $t' = \mathcal{F}^{-1}(t)$.

Since N is saturated, the GEXT inference that generates the clause

$$C = s' (\text{diff}\langle \tau, v \rangle s' t') \not\approx t' (\text{diff}\langle \tau, v \rangle s' t') \vee s' \approx t'$$

is redundant—i.e., $C \in N \cup \text{GHRed}_C(N)$ (in the nonpurifying calculi) or $\mathcal{F}(N) \models \mathcal{F}(C)$ (in the purifying calculi). In both cases, it follows that $R \models \mathcal{F}(C)$ by Theorem 4.5 and thus $\mathcal{I}^{\text{GH}} \models C$ by what we have shown above. The second literal of C is false in \mathcal{I}^{GH} because $\llbracket s' \rrbracket_{\mathcal{I}^{\text{GH}}} = a \neq b = \llbracket t' \rrbracket_{\mathcal{I}^{\text{GH}}}$. So the first literal of C must be true in \mathcal{I}^{GH} and thus

$$\begin{aligned} \mathcal{E}^{\text{GH}}(a)(\llbracket \text{diff}\langle \tau, v \rangle s' t' \rrbracket_{\mathcal{I}^{\text{GH}}}) &= \llbracket s' (\text{diff}\langle \tau, v \rangle s' t') \rrbracket_{\mathcal{I}^{\text{GH}}} \\ &\neq \llbracket t' (\text{diff}\langle \tau, v \rangle s' t') \rrbracket_{\mathcal{I}^{\text{GH}}} = \mathcal{E}^{\text{GH}}(b)(\llbracket \text{diff}\langle \tau, v \rangle s' t' \rrbracket_{\mathcal{I}^{\text{GH}}}) \end{aligned}$$

It follows that $\mathcal{E}^{\text{GH}}(a) \neq \mathcal{E}^{\text{GH}}(b)$. \square

We summarize the results of this subsection in the following theorem:

Theorem 4.10 (Ground static refutational completeness). *Let GHSel be a selection function on \mathcal{C}_{GH} . Then the inference system $\text{GHInf}^{\text{GHSel}}$ is statically refutationally complete w.r.t. $(\text{GHRed}_1, \text{GHRed}_C)$. That means, if $N \subseteq \mathcal{C}_{\text{GH}}$ is saturated w.r.t. $\text{GHInf}^{\text{GHSel}}$ and $\text{GHRed}_1^{\text{GHSel}}$, then $N \models \perp$ if and only if $\perp \in N$.*

The construction of \mathcal{I}^{GH} relies on the specific properties of R . It would not work with an arbitrary first-order interpretation. Transforming a λ -free higher-order interpretation into a first-order interpretation is easier:

Lemma 4.11. *Given an interpretation \mathcal{I} on GH, there exists an interpretation \mathcal{I}^{GF} on GF such that for any clause $C \in \mathcal{C}_{\text{GH}}$ the truth values of C in \mathcal{I} and of $\mathcal{F}(C)$ in \mathcal{I}^{GF} coincide.*

Proof. Let $\mathcal{I} = (\mathcal{U}, \mathcal{I}_{\text{ty}}, \mathcal{J}, \mathcal{E})$ be a λ -free higher-order interpretation. Let $\mathcal{U}_{\tau}^{\text{GF}} = \llbracket \tau \rrbracket_{\mathcal{I}_{\text{ty}}}$ be the first-order type universe for the ground type τ . For a symbol $f_l^{\bar{v}_m} \in \Sigma_{\text{GF}}$ and universe elements \bar{a}_l , let $\mathcal{J}^{\text{GF}}(f_l^{\bar{v}_m})(\bar{a}_l) = \llbracket f\langle \bar{v}_m \rangle \bar{x}_l \rrbracket_{\mathcal{I}}^{\{\bar{x}_l \mapsto \bar{a}_l\}}$. This defines an interpretation $\mathcal{I}^{\text{GF}} = (\mathcal{U}^{\text{GF}}, \mathcal{J}^{\text{GF}})$ on GF.

We need to show that for any $C \in \mathcal{C}_{\text{GH}}$, $\mathcal{I} \models C$ if and only if $\mathcal{I}^{\text{GF}} \models \mathcal{F}(C)$. It suffices to show that $\llbracket t \rrbracket_{\mathcal{I}} = \llbracket \mathcal{F}(t) \rrbracket_{\mathcal{I}^{\text{GF}}}$ for all terms $t \in \mathcal{T}_{\text{GH}}$. We prove this by induction on t . Since t is ground, it must be of the form $f\langle \bar{v}_m \rangle \bar{s}_l$. Then $\mathcal{F}(t) = f_l^{\bar{v}_m}(\mathcal{F}(\bar{s}_l))$ and hence

$$\llbracket \mathcal{F}(t) \rrbracket_{\mathcal{I}^{\text{GF}}} = \mathcal{J}^{\text{GF}}(f_l^{\bar{v}_m})(\llbracket \mathcal{F}(\bar{s}_l) \rrbracket_{\mathcal{I}^{\text{GF}}}) \stackrel{\text{IH}}{=} \mathcal{J}^{\text{GF}}(f_l^{\bar{v}_m})(\llbracket \bar{s}_l \rrbracket_{\mathcal{I}}) = \llbracket f\langle \bar{v}_m \rangle \bar{s}_l \rrbracket_{\mathcal{I}} = \llbracket t \rrbracket_{\mathcal{I}}$$

using the definition of \mathcal{J}^{GF} and Lemma 3.10 for the third step. \square

4.3. The Nonground Higher-Order Level. To lift the result to the nonground level, we employ the saturation framework of Waldmann et al. [WTRB20]. It is easy to see that the entailment relation \models on GH is a consequence relation in the sense of the framework. It remains to show that our redundancy criterion on GH is a redundancy criterion in the sense of the framework and that \mathcal{G} is a grounding function in the sense of the framework:

Lemma 4.12. *The redundancy criterion for GH is a redundancy criterion in the sense of the framework.*

Proof. We must prove the conditions (R1) to (R4) defined by Waldmann et al., which, adapted to our context, state the following for all clause sets $N, N' \subseteq \mathcal{C}_{\text{GH}}$:

- (R1) if $N \models \perp$, then $N \setminus \mathit{GHRed}_C(N) \models \perp$;
- (R2) if $N \subseteq N'$, then $\mathit{GHRed}_C(N) \subseteq \mathit{GHRed}_C(N')$ and $\mathit{GHRed}_I(N) \subseteq \mathit{GHRed}_I(N')$;
- (R3) if $N' \subseteq \mathit{GHRed}_C(N)$, then $\mathit{GHRed}_C(N) \subseteq \mathit{GHRed}_C(N \setminus N')$ and $\mathit{GHRed}_I(N) \subseteq \mathit{GHRed}_I(N \setminus N')$;
- (R4) if $\iota \in \mathit{GHInf}$ and $\mathit{concl}(\iota) \in N$, then $\iota \in \mathit{GHRed}_I(N)$.

(R1) It suffices to show that $N \setminus \mathit{GHRed}_C(N) \models N$ for $N \subseteq \mathcal{C}_{\text{GH}}$. Let \mathcal{I} be a model of $N \setminus \mathit{GHRed}_C(N)$. In the extensional calculi, let \mathcal{I} be extensional. Then by Lemma 4.11, there exists a model \mathcal{I}^{GF} of $\mathcal{F}(N \setminus \mathit{GHRed}_C(N)) = \mathcal{F}(N) \setminus \mathit{GFRed}_C(\mathcal{F}(N))$. By Lemma 5.2 of Bachmair and Ganzinger, this is also a model of $\mathcal{F}(N)$. By Lemma 4.11, it follows that $\mathcal{I} \models N$.

(R2) By Lemma 5.6(i) of Bachmair and Ganzinger, this holds on GF. For clauses and all inferences except GARGCONG and GEXT, this implies that it holds on GH as well because \mathcal{F} is a redundancy-preserving bijection between \mathcal{C}_{GH} and \mathcal{C}_{GF} and between these inferences. For GARGCONG and GEXT inferences, it holds because it holds on clauses.

(R3) The proof is analogous to (R2), with Lemma 5.6(ii) of Bachmair and Ganzinger instead of Lemma 5.6(i).

(R4) We must show that for all inferences with $\mathit{concl}(\iota) \in N$, we have $\iota \in \mathit{GHRed}_I(N)$. Since $\mathit{concl}(\iota) \prec \mathit{mprem}(\iota)$ for all $\iota \in \mathit{GFInf}$, this holds on GF. For all inferences except GARGCONG and GEXT, since \mathcal{F} is a bijection preserving redundancy, it follows that it also holds also on GH. For GARGCONG and GEXT inferences, it holds by definition. \square

Lemma 4.13. *The grounding functions $\mathcal{G}^{\mathit{GHSel}}$ for $\mathit{GHSel} \in \mathcal{G}(\mathit{HSel})$ are grounding functions in the sense of the framework.*

Proof. We must prove the conditions (G1) to (G3) defined by Waldmann et al., which, adapted to our context, state the following:

- (G1) $\mathcal{G}(\perp) = \{\perp\}$;
- (G2) for every $C \in \mathcal{C}_{\text{GH}}$, if $\perp \in \mathcal{G}(C)$, then $C = \perp$;
- (G3) for every $\iota \in \mathit{HInf}$, if $\mathcal{G}^{\mathit{GHSel}}(\iota) \neq \mathit{undef}$, then $\mathcal{G}^{\mathit{GHSel}}(\iota) \subseteq \mathit{GHRed}_I^{\mathit{GHSel}}(\mathcal{G}(\mathit{concl}(\iota)))$.

Clearly, $C = \perp$ if and only if $\perp \in \mathcal{G}(C)$ if and only if $\mathcal{G}(C) = \{\perp\}$, proving (G1) and (G2). For (G3), we have to show for all non-POSEXT inferences $\iota \in \mathit{HInf}$ that $\mathcal{G}^{\mathit{GHSel}}(\iota) \subseteq \mathit{GHRed}_I^{\mathit{GHSel}}(\mathcal{G}(\mathit{concl}(\iota)))$. Let $\iota \in \mathit{HInf}$ and $\iota' \in \mathcal{G}^{\mathit{GHSel}}(\iota)$. By the definition of $\mathcal{G}^{\mathit{GHSel}}$, there exists a grounding substitution θ such that $\mathit{prems}(\iota') = \mathit{prems}(\iota)\theta$ and $\mathit{concl}(\iota') = \mathit{preconcl}(\iota)\theta$. We want to show that $\iota' \in \mathit{GHRed}_I^{\mathit{GHSel}}(\mathcal{G}(\mathit{concl}(\iota)))$.

If ι' is not an GARGCONG or GEXT inference, by the definition of inference redundancy, it suffices to show that $\{D \in \mathcal{F}(\mathcal{G}(\mathit{concl}(\iota))) \mid D \prec \mathit{mprem}(\mathcal{F}(\iota'))\} \models \mathit{concl}(\mathcal{F}(\iota'))$. We define a substitution θ' that extends θ to all variables in $\mathit{concl}(\iota)$. Due to purification, the clause $\mathit{concl}(\iota)$ may contain variables not present in $\mathit{preconcl}(\iota)$. For each such variable x' , let x be the variable in $\mathit{preconcl}(\iota)$ that x' stems from and define $x'\theta' = x\theta$. Then the clause $\mathcal{F}(\mathit{concl}(\iota)\theta')$ differs from the clause $\mathcal{F}(\mathit{concl}(\iota')) = \mathcal{F}(\mathit{preconcl}(\iota)\theta')$ only in some additional grounded purification literals, which all have the form $t \not\approx t$ and are thus trivially false in any interpretation. Hence, $\mathcal{F}(\mathit{concl}(\iota)\theta') \models \mathcal{F}(\mathit{concl}(\iota'))$. Since one of the variables of a purification literal must appear applied in the clause, for each grounded purification literal $t \not\approx t$ the term t must be smaller than the maximal term of the clause $\mathcal{F}(\mathit{concl}(\iota'))$.

If no literals are selected in $\mathit{mprem}(\mathcal{F}(\iota'))$, inspection of the rules in GFInf shows that $\mathcal{F}(\mathit{concl}(\iota)\theta') \prec \mathit{mprem}(\mathcal{F}(\iota'))$. Otherwise, ι' can only be an ERES inference or a SUP

inference into a negative literal. If it is an ERES inference, due to the selection restrictions, the substitution σ used in ι is the identity for all variables of functional type. Therefore, applying σ cannot trigger any purification and hence $\mathcal{F}(\text{concl}(\iota)\theta') = \mathcal{F}(\text{preconcl}(\iota)\theta') \prec \text{mprem}(\mathcal{F}(\iota'))$. If ι' is a SUP inference into a negative literal, due to the selection restrictions, the substitution $\sigma = \text{mgu}(t, u)$ used in ι is the identity for all variables of functional type that stem from the main premise. Therefore the variables from the main premise C need not be purified. The variables from the side premise D might need to be purified, yielding purification literals of the form $x \approx y$ where $x\theta' = y\theta'$. Then x or y must appear applied in D and hence $x\theta'$ is smaller than $t\theta'$. Again, it follows that $\mathcal{F}(\text{concl}(\iota)\theta') \prec \text{mprem}(\mathcal{F}(\iota'))$.

This proves $\{D \in \mathcal{F}(\mathcal{G}(\text{concl}(\iota))) \mid D \prec \text{mprem}(\mathcal{F}(\iota'))\} \models \text{concl}(\mathcal{F}(\iota'))$.

In the nonpurifying calculi, if ι' is an GARGCONG or GEXT inference, it suffices to show that $\text{concl}(\iota') \in \mathcal{G}(\text{concl}(\iota))$. This holds because $\text{concl}(\iota') = \text{preconcl}(\iota)\theta = \text{concl}(\iota)\theta$. In the purifying calculi, if ι' is an GARGCONG or GEXT inference, we must show that $\mathcal{F}(\mathcal{G}(\text{concl}(\iota))) \models \mathcal{F}(\text{concl}(\iota'))$. Defining θ' as above, we have $\mathcal{F}(\text{concl}(\iota)\theta') \models \mathcal{F}(\text{concl}(\iota'))$, as desired. \square

To lift the completeness result of the previous section to the nonground calculus $HInf$, we employ Theorem 14 of Waldmann et al., which, adapted to our context, is stated as follows. The theorem uses the notation $Inf(N)$ to denote the set of Inf -inferences whose premises are in N , for an inference system Inf and a clause set N . Moreover, it uses the Herbrand entailment $\models_{\mathcal{G}}$ on \mathcal{C}_H , which is defined as $N_1 \models_{\mathcal{G}} N_2$ if $\mathcal{G}(N_1) \models \mathcal{G}(N_2)$.

Theorem 4.14 (Lifting theorem). *If $GHIInf^{GHSel}$ is statically refutationally complete w.r.t. $(GHRed_1^{GHSel}, GHRed_C)$ for every $GHSel \in \mathcal{G}(HSel)$, and if for every $N \subseteq \mathcal{C}_H$ that is saturated w.r.t. $HInf$ and $HRed_1$ there exists a $GHSel \in \mathcal{G}(HSel)$ such that $GHIInf^{GHSel}(\mathcal{G}(N)) \subseteq \mathcal{G}^{GHSel}(HInf(N)) \cup GHRed_1^{GHSel}(\mathcal{G}(N))$, then $HInf$ is statically refutationally complete w.r.t. $(HRed_1, HRed_C)$ and $\models_{\mathcal{G}}$.*

Proof. This is essentially Theorem 14 of Waldmann et al. We take H for \mathbf{F} , $\mathcal{G}H$ for \mathbf{G} , and $\mathcal{G}(HSel)$ for Q . It is easy to see that the entailment relation \models on $\mathcal{G}H$ is a consequence relation in the sense of the framework. By Lemma 4.12 and 4.13, $(GHRed_1^{GHSel}, GHRed_C)$ is a redundancy criterion in the sense of the framework, and \mathcal{G}^{GHSel} are grounding functions in the sense of the framework, for all $GHSel \in \mathcal{G}(HSel)$. The redundancy criterion $(HRed_1, HRed_C)$ matches exactly the intersected lifted redundancy criterion $Red^{\cap \mathcal{G}, \square}$ of Waldmann et al. Theorem 14 of Waldmann et al. states the theorem only for $\square = \emptyset$. By Lemma 16 of Waldmann et al., it also holds if $\square \neq \emptyset$. \square

Let $N \subseteq \mathcal{C}_H$ be a clause set saturated w.r.t. $HInf$ and $HRed_1$. We assume that $HSel$ fulfills the selection restrictions introduced in Section 3.1. For the above theorem to apply, we need to show that there exists a selection function $GHSel \in \mathcal{G}(HSel)$ such that all inferences $\iota \in GHIInf^{GHSel}$ with $\text{prems}(\iota) \in \mathcal{G}(N)$ are liftable or redundant. Here, by *liftable*, we mean that ι is a \mathcal{G}^{GHSel} -ground instance of a $HInf$ -inference from N ; by *redundant*, we mean that $\iota \in GHRed_1^{GHSel}(\mathcal{G}(N))$.

To choose the right selection function $GHSel \in \mathcal{G}(HSel)$, we observe that each ground clause $C \in \mathcal{G}(N)$ must have at least one corresponding clause $D \in N$ such that C is a ground instance of D . We choose one of them for each $C \in \mathcal{G}(N)$, which we denote by $\mathcal{G}^{-1}(C)$. Then let $GHSel$ select those literals in C that correspond to the literals selected by $HSel$ in

$\mathcal{G}^{-1}(C)$. Given this selection function $GHSel$, we can show that all inferences from $\mathcal{G}(N)$ are liftable or redundant.

All non-SUP inferences in $GHIInf$ are liftable (Lemma 4.16). For SUP, some inferences are liftable (Lemma 4.17) and some are redundant (Lemma 4.18). As in standard superposition, SUP inferences into positions below variables are redundant. The variable condition of each of the four calculi is designed to cover the nonredundant SUP inferences into positions of variable-headed terms, which makes these inferences liftable.

Lemma 4.15. *Let σ be the most general unifier of s and s' . Let θ be an arbitrary unifier of s and s' . Then $\sigma\theta = \theta$.*

Proof. Like in first-order logic, we can assume that σ is idempotent without loss of generality [Fit96, Corollary 7.2.11]. Since σ is most general, there exists a substitution ρ such that $\sigma\rho = \theta$. Therefore, by idempotence, $\sigma\theta = \sigma\sigma\rho = \sigma\rho = \theta$. \square

Lemma 4.16 (Lifting of ERES, EFACT, GARGCONG, and GEXT). *All ERES, EFACT, GARGCONG, and GEXT inferences are liftable.*

Proof. ERES: Let $\iota \in GHIInf^{GHSel}$ be an ERES inference with $prems(\iota) \in \mathcal{G}(N)$. Then ι is of the form

$$\frac{C\theta = C'\theta \vee s\theta \not\approx s'\theta}{C'\theta} \text{ERES}$$

where $\mathcal{G}^{-1}(C\theta) = C = C' \vee s \not\approx s'$ and the literal $s\theta \not\approx s'\theta$ is eligible w.r.t. $GHSel$. Let $\sigma = \text{mgu}(s, s')$. It follows that $s \not\approx s'$ is eligible in C w.r.t. σ and $HSel$. Moreover, $s\theta$ and $s'\theta$ are unifiable and ground, and therefore $s\theta = s'\theta$. Thus, the following inference $\iota' \in HIInf$ is applicable:

$$\frac{C' \vee s \not\approx s'}{\text{pure}(C'\sigma)} \text{ERES}$$

(where *pure* is the identity in the nonpurifying calculi). By Lemma 4.15, we have $C'\sigma\theta = C'\theta$. Therefore, ι is the θ -ground instance of ι' and is therefore liftable.

EFACT: Analogously, if $\iota \in GHIInf^{GHSel}$ is an EFACT inference with $prems(\iota) \in \mathcal{G}(N)$, then ι is of the form

$$\frac{C\theta = C'\theta \vee s'\theta \approx t'\theta \vee s\theta \approx t\theta}{C'\theta \vee t\theta \not\approx t'\theta \vee s\theta \approx t\theta} \text{EFACT}$$

where $\mathcal{G}^{-1}(C\theta) = C = C' \vee s' \approx t' \vee s \approx t$, the literal $s\theta \approx t\theta$ is eligible in C w.r.t. $GHSel$, and $s\theta \not\approx t\theta$. Let $\sigma = \text{mgu}(s, s')$. Hence, $s \approx t$ is eligible in C w.r.t. σ and $HSel$. We have $s \not\approx t$. Moreover, $s\theta$ and $s'\theta$ are unifiable and ground. Hence, $s\theta = s'\theta$. Thus, the following inference $\iota' \in HIInf$ is applicable:

$$\frac{C' \vee s' \approx t' \vee s \approx t}{\text{pure}((C' \vee t \not\approx t' \vee s \approx t')\sigma)} \text{EFACT}$$

By Lemma 4.15, we have $preconcl(\iota')\theta = concl(\iota)$. Hence, ι is the θ -ground instance of ι' and is therefore liftable.

GARGCONG: Let $\iota \in GHIInf^{GHSel}$ be a GARGCONG inference with $prems(\iota) \in \mathcal{G}(N)$. Then ι is of the form

$$\frac{C\theta = C'\theta \vee s\theta \approx s'\theta}{C'\theta \vee s\theta \bar{u}_n \approx s'\theta \bar{u}_n} \text{GARGCONG}$$

where $\mathcal{G}^{-1}(C\theta) = C = C' \vee s \approx s'$, the literal $s\theta \approx s'\theta$ is strictly eligible w.r.t. $GHSel$, and $s\theta$ and $s'\theta$ are of functional type. It follows that s and s' have either a functional or a polymorphic type. Let σ be the most general substitution such that $s\sigma$ and $s'\sigma$ take n arguments. Then $s \not\approx s'$ is eligible in C w.r.t. σ and $HSel$. Hence the following inference $\iota' \in HInf$ is applicable:

$$\frac{C' \vee s \approx s'}{\text{pure}(C'\sigma \vee s\sigma \bar{x}_n \approx s'\sigma \bar{x}_n)} \text{ARGCONG}$$

Then ι is a ground instance of ι' and is therefore liftable.

GEXT: The conclusion of a GEXT inference in $GHInf$ is by definition a ground instance of the conclusion of the EXT inference in $HInf$ before purification. Hence, the GEXT inference is a ground instance of the EXT inference. Therefore it is liftable. \square

Lemma 4.17 (Lifting of SUP). *Let $\iota \in GHInf^{GHSel}$ be a SUP inference*

$$\frac{D'\theta \vee t\theta \approx t'\theta \quad C'\theta \vee s\theta \langle t\theta \rangle_p \approx s'\theta}{D'\theta \vee C'\theta \vee s\theta \langle t'\theta \rangle_p \approx s'\theta} \text{SUP}$$

where $\mathcal{G}^{-1}(D\theta) = D = D' \vee t \approx t'$ and $\mathcal{G}^{-1}(C\theta) = C = C' \vee s \approx s'$. Suppose that the position p exists as a green subterm in s . Let u be the green subterm of s at that position and $\sigma = \text{mgu}(t, u)$ (which exists since θ is a unifier). If the variable condition holds for C , t , t' , u , and σ , then ι is liftable.

Proof. The inference conditions of ι can be lifted to D and C . That $t\theta \approx t'\theta$ is strictly eligible in $D\theta$ w.r.t. $GHSel$ implies that $t \approx t'$ is strictly eligible in D w.r.t. σ and $HSel$. If $s\theta \approx s'\theta$ is (strictly) eligible in $C\theta$ w.r.t. $GHSel$, then $s \approx s'$ is (strictly) eligible in C w.r.t. σ and $HSel$. Moreover, $D\theta \not\approx C\theta$ implies $D \not\approx C$, $t\theta \not\approx t'\theta$ implies $t \not\approx t'$, and $s\theta \not\approx s'\theta$ implies $s \not\approx s'$.

By assumption, p is a position of s and the variable condition holds. Thus, the following inference $\iota' \in HInf$ is applicable:

$$\frac{D' \vee t \approx t' \quad C' \vee s \langle u \rangle_p \approx s'}{\text{pure}((D' \vee C' \vee s \langle t' \rangle_p \approx s')\sigma)} \text{SUP}$$

By Lemma 4.15, we have $(\text{preconcl}(\iota'))\theta = \text{concl}(\iota)$. Hence, ι is the θ -ground instance of ι' and is therefore liftable. \square

The other SUP inferences might not be liftable, but they are redundant:

Lemma 4.18. *Let $\iota \in GHInf^{GHSel}$ be a SUP inference*

$$\frac{D'\theta \vee t\theta \approx t'\theta \quad C'\theta \vee s\theta \langle t\theta \rangle_p \approx s'\theta}{D'\theta \vee C'\theta \vee s\theta \langle t'\theta \rangle_p \approx s'\theta} \text{SUP}$$

where $\mathcal{G}^{-1}(D\theta) = D = D' \vee t \approx t'$ and $\mathcal{G}^{-1}(C\theta) = C = C' \vee s \approx s'$. Suppose that Lemma 4.17 does not apply. This could be either because the position p is below a variable in s or because the variable condition does not hold. Then $\iota \in GHRed_1^{GHSel}(\mathcal{G}(N))$.

Proof. By the definition of $GHRed_1$, to show $\iota \in GHRed_1^{GHSel}(\mathcal{G}(N))$, it suffices to prove that $\{E \in \mathcal{F}(\mathcal{G}(N)) \mid E \prec \mathcal{F}(C\theta)\} \models \mathcal{F}(\text{concl}(\iota))$. Let \mathcal{I} be a first-order model of all $E \in \mathcal{F}(\mathcal{G}(N))$ with $E \prec \mathcal{F}(C\theta)$. We must show that $\mathcal{I} \models \mathcal{F}(\text{concl}(\iota))$. If $\mathcal{I} \models \mathcal{F}(D'\theta)$, this is obvious. So we further assume that $\mathcal{I} \not\models \mathcal{F}(D'\theta)$. Since $D\theta \prec C\theta$ by the SUP inference

conditions, it follows that $\mathcal{I} \models \mathcal{F}(t\theta \approx t'\theta)$. By congruence, it suffices to show $\mathcal{I} \models \mathcal{F}(C\theta)$. We proceed by a case distinction on the two possible reasons why Lemma 4.17 does not apply:

CASE 1: The position p is below a variable in s . Then $t\theta$ is a proper green subterm of $x\theta$ and hence a green subterm of $x\theta \bar{w}$ for any arguments \bar{w} . Let v be the term that we obtain by replacing $t\theta$ by $t'\theta$ in $x\theta$ at the relevant position. It follows from our assumptions about \mathcal{I} that $\mathcal{I} \models \mathcal{F}(t\theta \approx t'\theta)$, and by congruence, $\mathcal{I} \models \mathcal{F}(x\theta \bar{w} \approx v \bar{w})$ for any arguments \bar{w} . Hence, $\mathcal{I} \models \mathcal{F}(C\theta)$ if and only if $\mathcal{I} \models \mathcal{F}(C\{x \mapsto v\}\theta)$. By the inference conditions we have $t\theta \succ t'\theta$, which implies $\mathcal{F}(C\theta) \succ \mathcal{F}(C\{x \mapsto v\}\theta)$ by compatibility with green contexts. Therefore, we have $\mathcal{I} \models \mathcal{F}(C\{x \mapsto v\}\theta)$ and hence $\mathcal{I} \models \mathcal{F}(C\theta)$.

CASE 2: The variable condition does not hold. In the extensional calculi, it follows that u has a variable head and jells with $t \approx t'$. By Definition 3.2, this means that u , t , and t' have the following form: $u = x \bar{v}_n$ for some variable x and a tuple of terms \bar{v}_n of length $n \geq 0$; $t = \tilde{t} \bar{x}_n$ and $t' = \tilde{t}' \bar{x}_n$, where \bar{x}_n are variables that do not occur elsewhere in D .

For the intensional calculi, we have $u \in \mathcal{V}$. Thus, u , t , and t' can be written in the same form as described above for the extensional calculi, with $n = 0$.

CASE 2.1 (PURIFYING CALCULI): First, we assume that x occurs only with arguments \bar{v}_n in C . For the intensional calculus, this must be the case because $n = 0$ and hence x can only occur without arguments by the definition of *pure* and the literal selection restriction. Define a substitution θ' by $x\theta' = \tilde{t}'\theta$ and $y\theta' = y\theta$ for other variables y . Since $t\theta \succ t'\theta$ by the inference conditions, we have $C\theta \succ C\theta'$. Moreover, $C\theta' \in \mathcal{G}(N)$. Then $\mathcal{I} \models \mathcal{F}(C\theta)$ by congruence, because $\mathcal{I} \models \mathcal{F}(C\theta')$ and $\mathcal{I} \models \mathcal{F}(t\theta \approx t'\theta)$.

Now we assume that x occurs with arguments other than \bar{v}_n in C . This can only happen in the extensional calculus and by the selection restrictions, $s\theta \approx s'\theta$ must not be selected in $C\theta$. Therefore, $s\theta$ is the maximal term in $C\theta$. Then $s \neq x$ and hence $\bar{v}_n \neq \varepsilon$ because otherwise $s\theta = x\theta$ would be smaller than the applied occurrence of $x\theta$ in $C\theta$.

Define a substitution θ'' such that $x\theta'' = \tilde{t}'\theta$, $y\theta'' = \tilde{t}'\theta$ for other variables y if $y\theta = s\theta$ and C contains the literal $x \not\approx y$, and $y\theta'' = y\theta$ otherwise.

We show that $C\theta \succ C\theta''$ by proving that no literal of $C\theta''$ is larger than the maximal literal $s\theta \approx s'\theta$ of $C\theta$ and that $s\theta \approx s'\theta$ appears more often in $C\theta$ than in $C\theta''$.

For a literal of the form $x \not\approx y$, we have $x\theta'' \prec s\theta$ and $y\theta'' \prec s\theta$. For literals that are not of this form, by the definition of *pure* in the extensional calculus, x occurs always with arguments \bar{v}_n . Hence these literals are equal or smaller in $C\theta''$ than in $C\theta$, because $x\theta'' \bar{v}_n \prec x\theta \bar{v}_n$ and $y\theta'' \preceq y\theta$. Therefore, no literal of $C\theta''$ is larger than the maximal literal $s\theta \approx s'\theta$ of $C\theta$. Moreover, these inequalities show that every occurrence of $s\theta \approx s'\theta$ in $C\theta''$ corresponds to an occurrence of $s\theta \approx s'\theta$ in $C\theta$ that corresponds to a literal in C without the variable x . Since at least one occurrence of $s\theta \approx s'\theta$ in $C\theta$ corresponds to a literal in C containing x , $s\theta \approx s'\theta$ appears more often in $C\theta$ than in $C\theta''$. This concludes the argument that $C\theta \succ C\theta''$. It follows that $\mathcal{I} \models \mathcal{F}(C\theta'')$.

We need to show that $\mathcal{I} \models \mathcal{F}(C\theta)$. There is a POEXT inference from D to $D' \vee \tilde{t} \approx \tilde{t}'$. This inference is in $HRed_1(N)$ because N is saturated. Therefore, $D'\theta \vee \tilde{t}\theta \approx \tilde{t}'\theta$ is in $\mathcal{G}(N) \cup GHRed_C(\mathcal{G}(N))$. It follows that $\mathcal{I} \models \mathcal{F}(D'\theta \vee \tilde{t}\theta \approx \tilde{t}'\theta)$ because this clause is smaller than $\mathcal{F}(D'\theta)$ and hence smaller than $\mathcal{F}(C\theta)$. Since $\mathcal{F}(D'\theta)$ is false in \mathcal{I} , we have $\mathcal{I} \models \mathcal{F}(\tilde{t}\theta \approx \tilde{t}'\theta)$.

For every literal of the form $x \not\approx y$, where $y\theta = s\theta$, the variable y can only occur without arguments in C because of the maximality of $s\theta$. We distinguish two cases. If for every

literal of the form $x \not\approx y$ where $y\theta = s\theta$, we have $\mathcal{I} \models \mathcal{F}(y\theta'' \approx y\theta)$, then $\mathcal{I} \models \mathcal{F}(C\theta)$ by congruence. If for some literal of the form $x \not\approx y$ where $y\theta = s\theta$, we have $\mathcal{I} \models \mathcal{F}(y\theta'' \not\approx y\theta)$, then $\mathcal{I} \models \mathcal{F}(y\theta \not\approx x\theta)$ because $y\theta'' = \tilde{t}'\theta$, $\mathcal{I} \models \mathcal{F}(\tilde{t}'\theta \approx \tilde{t}\theta)$, and $\tilde{t}\theta = x\theta$. Hence a literal of $\mathcal{F}(C\theta)$ is true in \mathcal{I} and therefore $\mathcal{I} \models C\theta$.

CASE 2.2 (NONPURIFYING CALCULI): Since the variable condition does not hold, we have $C\theta \succeq C''\theta$, where $C'' = C\{x \mapsto \tilde{t}'\}$. We cannot have $C\theta = C''\theta$ because $x\theta = \tilde{t}\theta \neq \tilde{t}'\theta$ and x occurs in C . Hence, we have $C\theta \succ C''\theta$.

By the definition of \mathcal{I} , $C\theta \succ C''\theta$ implies $\mathcal{I} \models \mathcal{F}(C''\theta)$. We will use equalities that are true in \mathcal{I} to rewrite $\mathcal{F}(C\theta)$ into $\mathcal{F}(C''\theta)$, which implies $\mathcal{I} \models \mathcal{F}(C\theta)$ by congruence.

By saturation of N , for any well-typed m -tuple of fresh variables \bar{z} , we can use a `POSEXT` with premise D (if $n > m$) or `ARGCONG` inference with premise D (if $n < m$) or using D itself (if $n = m$) to show that $\mathcal{G}(D' \vee \tilde{t} \bar{z} \approx \tilde{t}' \bar{z}) \subseteq \mathcal{G}(N) \cup \text{GHRed}_C(\mathcal{G}(N))$. Hence, $D'\theta \vee \tilde{t}\theta \bar{u} \approx \tilde{t}'\theta \bar{u}$ is in $\mathcal{G}(N) \cup \text{GHRed}_C(\mathcal{G}(N))$ for any ground arguments \bar{u} .

We observe that whenever $\tilde{t}\theta \bar{u}$ and $\tilde{t}'\theta \bar{u}$ are smaller than the maximal term of $C\theta$ for some arguments \bar{u} , we have

$$\mathcal{I} \models \mathcal{F}(\tilde{t}\theta \bar{u}) \approx \mathcal{F}(\tilde{t}'\theta \bar{u}) \quad (\dagger)$$

To show this, we assume that $\tilde{t}\theta \bar{u}$ and $\tilde{t}'\theta \bar{u}$ are smaller than the maximal term of $C\theta$ and we distinguish two cases: If $t\theta$ is smaller than the maximal term of $C\theta$, all terms in $D'\theta$ are smaller than the maximal term of $C\theta$ and hence $D'\theta \vee \tilde{t}\theta \bar{u} \approx \tilde{t}'\theta \bar{u} \prec C\theta$. If, on the other hand, $t\theta$ is equal to the maximal term of $C\theta$, $\tilde{t}\theta \bar{u}$ and $\tilde{t}'\theta \bar{u}$ are smaller than $t\theta$. Hence $\tilde{t}\theta \bar{u} \approx \tilde{t}'\theta \bar{u} \prec t\theta \approx t'\theta$ and $D'\theta \vee \tilde{t}\theta \bar{u} \approx \tilde{t}'\theta \bar{u} \prec D\theta \prec C\theta$. In both cases, since $\mathcal{F}(D'\theta)$ is false in \mathcal{I} by assumption, $\mathcal{I} \models \mathcal{F}(\tilde{t}\theta \bar{u}) \approx \mathcal{F}(\tilde{t}'\theta \bar{u})$.

We proceed by a case distinction on whether $s\theta$ appears in a selected or in a maximal literal of $C\theta$. In both cases we provide an algorithm that establishes the equivalence of $C\theta$ and $C''\theta$ via rewriting using (\dagger) . This might seem trivial at first sight, but we can only use the equations (\dagger) if $\tilde{t}\theta \bar{u}$ and $\tilde{t}'\theta \bar{u}$ are smaller than the maximal term of $C\theta$. Moreover, \bar{u} might itself contain positions where we want to rewrite, so the order of rewriting matters.

CASE 2.2.1: $s\theta$ is the maximal side of a maximal literal of $C\theta$. Then, since $C\theta \succ C''\theta$, every term in $C\theta$ and in $C''\theta$ is smaller than or equal to $s\theta$. Let C_0 and \tilde{C}_0 be the clauses resulting from rewriting $\mathcal{F}(t\theta) \rightarrow \mathcal{F}(t'\theta)$ wherever possible in $\mathcal{F}(C\theta)$ and $\mathcal{F}(C''\theta)$, respectively. Since $\mathcal{F}(t\theta)$ is a subterm of $\mathcal{F}(s\theta)$, now every term in C_0 and \tilde{C}_0 is strictly smaller than $\mathcal{F}(s\theta)$.

We define C_1, C_2, \dots inductively as follows: Given C_i , choose a subterm of the form $\mathcal{F}(\tilde{t}\theta \bar{u})$ where $\tilde{t}\theta \bar{u} \succ \tilde{t}'\theta \bar{u}$ or of the form $\mathcal{F}(\tilde{t}'\theta \bar{u})$ where $\tilde{t}'\theta \bar{u} \succ \tilde{t}\theta \bar{u}$. Let C_{i+1} be the clause resulting from rewriting that subterm $\mathcal{F}(\tilde{t}\theta \bar{u})$ to $\mathcal{F}(\tilde{t}'\theta \bar{u})$ or that subterm $\mathcal{F}(\tilde{t}'\theta \bar{u})$ to $\mathcal{F}(\tilde{t}\theta \bar{u})$ in C_i , depending on which term was chosen. Analogously, we define $\tilde{C}_1, \tilde{C}_2, \dots$ by applying the same algorithm to \tilde{C}_0 . In both cases, the process terminates because \succ is compatible with green contexts and well founded. Let C_* and \tilde{C}_* be the respective final clauses.

The algorithm preserves the invariant that every term in C_i and \tilde{C}_i is strictly smaller than $s\theta$. By congruence via (\dagger) , applied at every step of the algorithm, we know that C_* and $\mathcal{F}(C\theta)$ are equivalent in \mathcal{I} and that \tilde{C}_* and $\mathcal{F}(C''\theta)$ are equivalent in \mathcal{I} as well.

We show that $C_* = \tilde{C}_*$. Assume that $C_* \neq \tilde{C}_*$. The algorithm preserves a second invariant, namely that $\mathcal{F}^{-1}(C_i)$ and $\mathcal{F}^{-1}(\tilde{C}_j)$ are equal except for positions where one contains $\tilde{t}\theta$ and the other one contains $\tilde{t}'\theta$. Consider a deepest position where $\mathcal{F}^{-1}(C_*)$ and $\mathcal{F}^{-1}(\tilde{C}_*)$ are different. The respective position in C_* and \tilde{C}_* then contains $\mathcal{F}(\tilde{t}\theta \bar{u})$ and $\mathcal{F}(\tilde{t}'\theta \bar{u})$ or vice versa. The arguments \bar{u} must be equal because we consider a deepest position. But then

$\tilde{t}\theta \bar{u} \succ \tilde{t}'\theta \bar{u}$ or $\tilde{t}\theta \bar{u} \prec \tilde{t}'\theta \bar{u}$, which is impossible since the algorithm terminated in C_* and \tilde{C}_* . This shows that $C_* = \tilde{C}_*$. Hence $\mathcal{F}(C\theta)$ and $\mathcal{F}(C''\theta)$ are equivalent, which proves $\mathcal{I} \models \mathcal{F}(C\theta)$.

CASE 2.2.2: $s\theta$ is the maximal side of a selected literal of $C\theta$. Then, by the selection restrictions, x cannot be the head of a maximal literal of C .

At every position where $x \bar{u}$ occurs in C with some (or no) arguments \bar{u} , we rewrite $(\tilde{t} \bar{u})\theta$ to $(\tilde{t}' \bar{u})\theta$ in $C\theta$ if $(\tilde{t} \bar{u})\theta \succ (\tilde{t}' \bar{u})\theta$. We start with the innermost occurrences of x , so that the order of the two terms at one step does not change by later rewriting. Analogously, at every position where $x \bar{u}$ occurs in C with some (or no) arguments \bar{u} , we rewrite $(\tilde{t}' \bar{u})\theta$ to $(\tilde{t} \bar{u})\theta$ in $C''\theta$ if $(\tilde{t}' \bar{u})\theta \succ (\tilde{t} \bar{u})\theta$, again starting with the innermost occurrences.

We never rewrite at the top level of the maximal term of $C\theta$ or $C''\theta$ because x cannot be the head of a maximal literal of C . The two resulting clauses are identical because $C\theta$ and $C''\theta$ only differ at positions where x occurs in C . The rewritten terms are all smaller than the maximal term of $C\theta$. With (\dagger) , this implies that $\mathcal{I} \models \mathcal{F}(C\theta)$ by congruence. \square

With these properties of our inference systems in place, the saturation framework's lifting theorem (Theorem 4.14) guarantees static and dynamic refutational completeness of $HInf$ w.r.t. $HRed_1$. However, this theorem gives us refutational completeness w.r.t. the Herbrand entailment $\models_{\mathcal{G}}$, defined as $N_1 \models_{\mathcal{G}} N_2$ if $\mathcal{G}(N_1) \models \mathcal{G}(N_2)$, whereas our semantics is Tarski entailment \models , defined as $N_1 \models N_2$ if any model of N_1 is a model of N_2 . The following lemma repairs this mismatch:

Lemma 4.19. *For $N \subseteq \mathcal{C}_H$, we have $N \models_{\mathcal{G}} \perp$ if and only if $N \models \perp$.*

Proof. By Lemma 3.11, any model of N is also a model of $\mathcal{G}(N)$ —i.e., $N \not\models \perp$ implies $N \not\models_{\mathcal{G}} \perp$. For the other direction, we need to show that $N \not\models_{\mathcal{G}} \perp$ implies $N \not\models \perp$. Assume that $N \not\models_{\mathcal{G}} \perp$ —i.e., $\mathcal{G}(N) \not\models \perp$. Then there is a model \mathcal{I} of $\mathcal{G}(N)$. We must show that there exists a model of N —i.e., $N \not\models \perp$. Let \mathcal{I}' be an interpretation derived from \mathcal{I} by removing all universes that are not the denotation of a type in $\mathcal{T}_{\mathcal{G}H}$ and removing all domain elements that are not the denotation of a term in $\mathcal{T}_{\mathcal{G}H}$, making \mathcal{I}' term-generated. Clearly, in our clausal logic, this leaves the denotations of terms and the truth of ground clauses unchanged. Thus, $\mathcal{I}' \models \mathcal{G}(N)$. We will show that $\mathcal{I}' \models N$. Let $C \in N$. We want to show that C is true in \mathcal{I}' for all valuations ξ . Fix a valuation ξ . By construction, for each variable x , there exists a ground term s_x such that $\llbracket s_x \rrbracket_{\mathcal{I}'} = \xi(x)$. Let ρ be the substitution that maps every free variable x in C to s_x . Then $\xi(x) = \llbracket s_x \rrbracket_{\mathcal{I}'} = \llbracket x\rho \rrbracket_{\mathcal{I}'}$ for all x . By treating the type variables of C in the same way, we can also achieve that $\xi(\alpha) = \llbracket \alpha\rho \rrbracket_{\mathcal{I}'}$ for all α . By Lemma 3.10, $\llbracket t\rho \rrbracket_{\mathcal{I}'} = \llbracket t \rrbracket_{\mathcal{I}'}^{\xi}$ for all terms t and $\llbracket \tau\rho \rrbracket_{\mathcal{I}'} = \llbracket \tau \rrbracket_{\mathcal{I}'}^{\xi}$ for all types τ . Hence, $C\rho$ and C have the same truth value in \mathcal{I}' for ξ . Since $\mathcal{I}' \models \mathcal{G}(N)$, $C\rho$ is true in \mathcal{I}' and thus C is true in \mathcal{I}' as well. \square

Theorem 4.20 (Static refutational completeness). *The inference system $HInf$ is statically refutationally complete w.r.t. $(HRed_1, HRed_C)$. That means, if $N \subseteq \mathcal{C}_H$ is a clause set saturated w.r.t. $HInf$ and $HRed_1$, then we have $N \models \perp$ if and only if $\perp \in N$.*

Proof. We apply Theorem 4.14. By Theorem 4.10, $GHInf^{GHSel}$ is statically refutationally complete for all $GHSel \in \mathcal{G}(HSel)$. By Lemmas 4.16, 4.17, and 4.18, for every saturated $N \subseteq \mathcal{C}_H$, there exists a selection function $GHSel \in \mathcal{G}(HSel)$ such that all inferences $\iota \in GHInf^{GHSel}$ with $prems(\iota) \in \mathcal{G}(N)$ either are \mathcal{G}^{GHSel} -ground instances of $HInf$ -inferences from N or belong to $GHRed_1^{GHSel}(\mathcal{G}(N))$.

Theorem 4.14 implies that if $N \subseteq C_H$ is a clause set saturated w.r.t. $HInf$ and $HRed_I$, then $N \models_G \perp$ if and only if $\perp \in N$. By Lemma 4.19, this also holds for the Tarski entailment \models . That is, if $N \subseteq C_H$ is a clause set saturated w.r.t. $HInf$ and $HRed_I$, then $N \models \perp$ if and only if $\perp \in N$. \square

From static refutational completeness, we can easily derive dynamic refutational completeness.

Theorem 4.21 (Dynamic refutational completeness). *The inference system $HInf$ is dynamically refutationally complete w.r.t. $(HRed_I, HRed_C)$, as defined in Definition 4.2.*

Proof. By Theorem 17 of Waldmann et al., this follows from Theorem 4.20 and Lemma 4.19. \square

5. IMPLEMENTATION

Zipperposition [Cru15, Cru17] is an open source superposition-based theorem prover written in OCaml.¹ It was initially designed for polymorphic first-order logic with equality, as embodied by TPTP TF1 [BP13]. We will refer to this implementation as Zipperposition’s first-order mode. Later, Cruanes extended the prover with a pragmatic higher-order mode with support for λ -abstractions and extensionality, without any completeness guarantees. We have now also implemented complete λ -free higher-order modes based on the four calculi described in this article.

The pragmatic higher-order mode provided a convenient basis to implement our calculi. It employs higher-order term and type representations and orders. Its ad hoc calculus extensions are similar to our calculi. They include an ARGCONG-like rule and a POSEXT-like rule, and SUP inferences are performed only at green subterms. One of the bugs we found during our implementation work occurred because argument positions shift when applied variables are instantiated. We resolved this by numbering argument positions in terms from right to left.

To implement the λ -free higher-order mode, we restricted the unification algorithm to non- λ -abstractions. To satisfy the requirements on selection, we avoid selecting literals that contain higher-order variables. To comply with our redundancy notion, we disabled rewriting of nongreen subterms. To improve term indexing of higher-order terms, we replaced the imperfect discrimination trees by fingerprint indices [Sch12]. To speed up the computation of the SUP conditions, we omit the condition $C\sigma \not\prec D\sigma$ in the implementation, at the cost of performing some additional inferences.

For the purifying calculi, we implemented purification as a simplification rule. This ensures that it is applied aggressively on all clauses, whether initial clauses from the problem or clauses produced during saturation, before any inferences are performed.

For the nonpurifying calculi, we added the possibility to perform SUP inferences at variable positions. This means that variables must be indexed as well. In addition, we modified the variable condition. Depending on the term ordering, it may be expensive or even impossible to decide whether there exists a grounding substitution θ with $t\sigma\theta \succ t'\sigma\theta$ and $C\sigma\theta \prec C''\sigma\theta$. We overapproximate the condition as follows: (1) check whether x appears with different arguments in the clause C ; (2) use a term-order-specific algorithm to determine whether there might exist a grounding substitution θ and terms \bar{u} such that $t\sigma\theta \succ t'\sigma\theta$

¹<https://github.com/sneeuwballen/zipperposition>

and $t\sigma\bar{u} \prec t'\sigma\bar{u}$; and (3) check whether $C\sigma \not\prec C''\sigma$. If these three conditions apply, we conclude that there might exist a ground substitution θ witnessing nonmonotonicity.

For the extensional calculi, we add axiom (EXT) to the clause set. To curb the explosion associated with extensionality, this axiom and all clauses derived from it are penalized by the clause selection heuristic. We also added the NEGEXT rule described in Section 3.2, which resembles Vampire’s extensionality resolution rule [GKKV14].

The ARGCONG rule can have infinitely many conclusions on polymorphic clauses. To capture this in the implementation, we store these infinite sequences of conclusions in the form of finite instructions of how to obtain the actual clauses. In addition to the usual active and passive set of the DISCOUNT loop architecture [ADF95], we use a set of scheduled inferences that stores these instructions. We visit the scheduled inferences in this additional set and the clauses in the passive set fairly to achieve dynamic completeness of our prover architecture. Waldmann et al. [WTRB20, Example 34] and Bentkamp et al. [BBT⁺21, Section 6] describe this architecture in more detail.

Using Zipperposition, we can quantify the disadvantage of the applicative encoding on Example 3.9. A well-chosen KBO instance with argument coefficients allows Zipperposition to derive \perp in 4 iterations of the prover’s main loop and 0.03 s. KBO or LPO with default settings needs 203 iterations and 0.4 s, whereas KBO or LPO on the applicatively encoded problem needs 203 iterations and more than 1 s due to the larger terms.

6. EVALUATION

We evaluated Zipperposition’s implementation of our four calculi on Sledgehammer-generated Isabelle/HOL benchmarks [BN10] and on benchmarks from the TPTP library [SSCB12, SBBT09]. Our experimental data is available online.² We used the development version of Zipperposition, revision 2031e216.³

The Sledgehammer benchmarks, corresponding to Isabelle’s Judgment Day suite, were regenerated to target clausal λ -free higher-order logic. They comprise 2506 problems in total, divided in two groups based on the number of Isabelle facts (lemmas, definitions, etc.) selected for inclusion in each problem: either 16 facts (SH16) or 256 facts (SH256). The problems were generated by encoding λ -expressions as λ -lifted supercombinators [MP08].

From the TPTP library, we collected 708 first-order problems in TFF format and 717 higher-order problems in THF format, both groups containing both monomorphic and polymorphic problems. We excluded all problems that contain interpreted arithmetic symbols, the symbols (@@+), (@@-), (@+), (@-), (&), or tuples, as well as the SYN000 problems, which are only intended to test the parser, and problems whose clausal normal form takes longer than 15 s to compute or falls outside the λ -free fragment described in Section 2.

We want to answer the following questions:

- (1) What is the overhead of our calculi on first-order benchmarks?
- (2) Do the calculi outperform the applicative encoding?
- (3) Do the purifying or the nonpurifying calculi achieve better results?
- (4) What is the cost of nonmonotonicity?

Since the present work is only a stepping stone towards a prover for full higher-order logic, it would be misleading to compare this prototype with state-of-the-art higher-order provers

²<https://doi.org/10.5281/zenodo.3992618>

³<https://github.com/sneeuwballen/zipperposition/tree/2031e216c1941acd76187882a073e8f1e533>

that support a stronger logic. Many of the higher-order problems in the TPTP library are in fact satisfiable for our λ -free logic, even though they may be unsatisfiable for full higher-order logic and labeled as such in the TPTP.

To answer question (1) and (3), we ran Zipperposition’s first-order mode on the first-order benchmarks and the purifying and nonpurifying modes on all benchmarks. To answer question (2), we implemented an applicative encoding mode in Zipperposition, which performs the applicative encoding after the clausal normal form transformation and then proceeds with Zipperposition’s first-order mode. The encoding makes all function symbols nullary and replaces all applications with a polymorphic binary `app` symbol.

We instantiated all four calculi with three term orders: LPO [BWW17], KBO [BBWW17] (without argument coefficients), and EPO [Ben18]. Among these, LPO is the only nonmonotonic order and therefore the most relevant option to evaluate our calculi, which are designed to cope with nonmonotonicity. To answer question (4), we also include the monotone orders KBO and EPO. EPO is an order designed to resemble LPO while fulfilling the requirements of a ground-total simplification order on λ -free terms. KBO and EPO serve as a yardstick to assess the cost of nonmonotonicity. With these monotone orders, no superposition inferences at variables are necessary and thus the nonpurifying calculi become a straightforward generalization of the standard superposition calculus with the caveat that it may be more efficient to superpose at nongreen subterms directly instead of relying on the `ARGCONG` rule. On first-order benchmarks and in the applicative encoding mode, all three orders are monotone because they are monotone on first-order terms.

Figure 1 summarizes, for the intensional calculi, the number of solved satisfiable and unsatisfiable problems within 180s, and the time taken to show unsatisfiability. Figure 2 presents the corresponding data for the extensional calculi. The average time is computed over the problems that all configurations for the respective benchmark set and term order found to be unsatisfiable within the time limit. For each combination of benchmark set and term order, the best result is highlighted in bold. The evaluation was carried out on StarExec Iowa [SST14] using Intel Xeon E5-2609 0 CPUs clocked at 2.40 GHz.

The experimental results on the TFF part of the TPTP library confirm that our calculi handle the vast majority of problems that are solvable in first-order mode gracefully, and thus that the overhead is minimal, answering question (1). On first-order problems, the calculi are occasionally at variance with the first-order mode, due to the interaction of `ARGCONG` with polymorphic types and due to the extensionality axiom (`EXT`). In contrast, the applicative encoding is comparatively inefficient on problems that are already first-order. For LPO, the success rate decreases by around 15%, and the average time to show unsatisfiability triples.

The SH16 benchmarks consist mostly of small higher-order problems. The small number of axioms benefits the applicative encoding enough to outperform the purifying calculi but not the nonpurifying ones. The SH256 benchmarks are also higher-order but much larger. Such problems are underrepresented in the TPTP library. On these, our calculi clearly outperform the applicative encoding, answering question (2) decisively. This is hardly surprising given that the proving effort is dominated by first-order reasoning, which they can perform gracefully.

The THF benchmarks generally require more sophisticated higher-order reasoning than the Sledgehammer benchmarks, as observed by Sultana, Blanchette, and Paulson [SBP13, Section 5]. On these benchmarks, the empirical results are less clear; the applicative encoding and our calculi are roughly neck-and-neck. The nonpurifying calculi detect unsatisfiability slightly more frequently, whereas the applicative encoding tends to find more saturations.

		# sat			# unsat			∅ time		
		LPO	KBO	EPO	LPO	KBO	EPO	LPO	KBO	EPO
SH16	applicative encoding	111	189	65	373	382	157	0.9	1.2	10.7
	nonpurifying calculus	136	165	133	383	385	381	0.4	0.3	0.0
	purifying calculus	82	98	82	363	363	355	1.3	2.0	0.0
SH256	applicative encoding	1	1	1	471	488	36	9.4	8.7	63.8
	nonpurifying calculus	1	1	1	543	554	498	2.3	2.3	0.1
	purifying calculus	1	1	1	523	528	484	2.6	3.4	0.5
TFF	first-order mode	0	0	0	212	229	107	1.9	2.3	1.5
	applicative encoding	0	0	0	180	205	21	7.0	10.0	4.6
	nonpurifying calculus	0	0	0	210	229	105	1.9	2.4	1.5
	purifying calculus	0	0	0	211	229	105	2.1	2.6	1.6
THF	applicative encoding	127	115	111	523	522	428	0.9	0.6	0.8
	nonpurifying calculus	111	114	112	529	527	516	0.3	0.3	0.0
	purifying calculus	108	109	108	528	526	514	0.3	0.5	0.0

FIGURE 1. Evaluation of the intensional calculi

		# sat			# unsat			∅ time		
		LPO	KBO	EPO	LPO	KBO	EPO	LPO	KBO	EPO
SH16	applicative encoding	79	152	48	379	386	157	1.2	1.3	11.4
	nonpurifying calculus	103	131	95	386	393	387	0.4	0.1	0.0
	purifying calculus	32	57	32	367	365	363	2.0	1.7	0.0
SH256	applicative encoding	1	1	1	462	486	36	7.5	9.4	63.8
	nonpurifying calculus	1	1	1	548	572	504	1.9	2.1	0.1
	purifying calculus	1	1	1	512	529	482	2.2	5.0	0.1
TFF	first-order mode	0	0	0	212	229	107	1.9	2.5	1.5
	applicative encoding	0	0	0	178	202	21	7.9	11.7	4.7
	nonpurifying calculus	0	0	0	207	229	106	2.1	3.0	1.5
	purifying calculus	0	0	0	210	229	105	2.2	3.2	1.6
THF	applicative encoding	108	109	105	526	527	436	0.9	0.6	1.1
	nonpurifying calculus	106	108	107	539	535	526	0.3	0.3	0.0
	purifying calculus	96	97	96	530	529	519	0.3	0.6	0.0

FIGURE 2. Evaluation of the extensional calculi

It seems that, due to the large amount of higher-order reasoning necessary to solve TPTP problems, the advantage of our calculi on the first-order parts of the derivation is not a decisive factor on these benchmarks.

Concerning question (3), the nonpurifying calculi outperform their purifying relatives across all benchmarks. The raw data show that on most benchmark sets, the problems solved by the nonpurifying calculi are almost a superset of the problems solved by the purifying calculi. Only on the SH256 benchmarks, the purifying calculi can solve a fair number of problems that the nonpurifying calculi cannot solve (11 problems for the intensional calculi with LPO and 9 problems for the extensional calculi with LPO).

KBO tends to have a slight advantage over LPO on all benchmark sets. But the gap between KBO and LPO is not larger on the higher-order benchmarks than on TFF. Since LPO is monotonic on first-order terms but nonmonotonic on higher-order terms, whereas KBO is monotonic on both, the best answer we can give to question (4) is that no substantial cost seems to be associated with nonmonotonicity. In particular, for the nonpurifying calculi, the additional superposition inferences at variables necessary with LPO do not have a negative impact on the overall performance. EPO generally performs worse than the other two orders, with the exception of the nonpurifying calculus on SH16 benchmarks, where it is roughly neck-and-neck with LPO. This suggests that for small, mildly higher-order problems, EPO can be a viable LPO-like complement to KBO if one considers the effort to implement our calculi too high.

7. DISCUSSION AND RELATED WORK

Our calculi join a long list of extensions and refinements of superposition. Among the most closely related is Peltier’s [Pel16] Isabelle/HOL formalization of the refutational completeness of a superposition calculus that operates on λ -free higher-order terms and that is parameterized by a monotonic term order. Extensions with polymorphism and induction, independently developed by Cruanes [Cru15, Cru17] and Wand [Wan17], contribute to increasing the power of automatic provers. Detection of inconsistencies in axioms, as suggested by Schulz et al. [SSUP17], is important for large axiomatizations.

Also of interest is Bofill and Rubio’s [BR13] integration of nonmonotonic orders in ordered paramodulation, a precursor of superposition. Their work is a veritable tour de force, but it is also highly complicated and restricted to ordered paramodulation. Lack of compatibility with arguments being a mild form of nonmonotonicity, it seemed preferable to start with superposition, enrich it with an ARGCONG rule, and tune the side conditions until we obtained a complete calculus.

Most complications can be avoided by using a monotonic order such as KBO without argument coefficients. However, coefficients can be useful to help achieve compatibility with β -reduction. For example, the term $\lambda x. x + x$ could be treated as a constant with a coefficient of 2 on its argument and a heavy weight to ensure $(\lambda x. x + x) y \succ y + y$. Although they do not use argument coefficients, the recently developed λ -superposition calculus by Bentkamp et al. [BBT⁺21] and combinatory superposition calculus by Bhayat and Reger [BR20] need a nonmonotonic order to cope with β -reduction. They are modeled after our extensional nonpurifying and intensional nonpurifying calculi, respectively.

Many researchers have proposed or used encodings of higher-order logic constructs into first-order logic, including Robinson [Rob70], Kerber [Ker91], Dougherty [Dou93], Dowek et al. [DHK95], Hurd [Hur03], Meng and Paulson [MP08], Obermeyer [Obe09], and Czajka [Cza16]. Encodings of types, such as those by Bobot and Paskevich [BP11] and Blanchette et al. [BBPS16], are also crucial to obtain a sound encoding of higher-order

logic. These ideas are implemented in proof assistant tools such as HOLyHammer and Sledgehammer [BKPU16].

In the term rewriting community, λ -free higher-order logic is known as applicative first-order logic. First-order rewrite techniques can be applied to this logic via the applicative encoding. However, there are similar drawbacks as in theorem proving to having **app** as the only nonnullary symbol. Hirokawa et al. [HMZ13] propose a technique that resembles our mapping \mathcal{F} to avoid these drawbacks.

Another line of research has focused on the development of automated proof procedures for higher-order logic. Robinson’s [Rob69], Andrews’s [And71], and Huet’s [Hue73] pioneering work stands out. Andrews [And01] and Benzmüller and Miller [BM14] provide excellent surveys. The competitive higher-order automatic theorem provers include LEO-II [BPTF08] (based on RUE resolution), Satallax [Bro12] (based on a tableau calculus and a SAT solver), agsyHOL [Lin14] (based on a focused sequent calculus and a generic narrowing engine), Leo-III [SB18] (based on a pragmatic higher-order version of ordered paramodulation with no completeness guarantees), CVC4 and veriT [BREO⁺19] (both based on satisfiability modulo theories), and Vampire [BR19, BR20] (based on superposition and SK-style combinators). The Isabelle proof assistant [NPW02] (which includes a tableau reasoner and a rewriting engine) and its Sledgehammer subsystem have also participated in the higher-order division of the CADE ATP System Competition.

Zipperposition is a convenient vehicle for experimenting and prototyping because it is easier to understand and modify than highly-optimized C or C++ provers. Our middle-term goal is to design higher-order superposition calculi, implement them in state-of-the-art provers such as E [Sch13], SPASS [WDF⁺09], and Vampire [KV13], and integrate these in proof assistants to provide a high level of automation. With its stratified architecture, Otter- λ [Bee04] is perhaps the closest to what we are aiming at, but it is limited to second-order logic and offers no completeness guarantees. As a first step, Vukmirović, Blanchette, Cruanes, and Schulz [VBCS19] have generalized E’s data structures and algorithms to clausal λ -free higher-order logic, assuming a monotonic KBO [BBWW17].

8. CONCLUSION

We presented four superposition calculi for intensional and extensional clausal λ -free higher-order logic and proved them refutationally complete. The calculi nicely generalize standard superposition and are compatible with our λ -free higher-order LPO and KBO. Especially on large problems, our experiments confirm what one would naturally expect: that native support for partial application and applied variables outperforms the applicative encoding.

The new calculi reduce the gap between proof assistants based on higher-order logic and superposition provers. We can use them to reason about arbitrary higher-order problems by axiomatizing suitable combinators. But perhaps more importantly, they appear promising as a stepping stone towards complete, highly efficient automatic theorem provers for full higher-order logic. Indeed, the subsequent work by Bentkamp et al. [BBT⁺21], which introduces support for λ -expressions, and Bhayat and Reger [BR20], which works with SK-style combinators, is largely based on our nonpurifying calculi.

Acknowledgment. We are grateful to the maintainers of StarExec for letting us use their service. We want to thank Petar Vukmirović for his work on Zipperposition and for his advice. We thank Christoph Benzmüller, Sander Dahmen, Johannes Hölzl, Anders Schlichtkrull,

Stephan Schulz, Alexander Steen, Geoff Sutcliffe, Andrei Voronkov, Daniel Wand, Christoph Weidenbach, and the participants in the 2017 Dagstuhl Seminar on Deduction beyond First-Order Logic for stimulating discussions. We also want to thank Christoph Benzmüller, Ahmed Bhayat, Wan Fokkink, Alexander Steen, Mark Summerfield, Sophie Tourret, and the anonymous reviewers for suggesting several textual improvements. Bentkamp and Blanchette’s research has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 713999, Matryoshka).

REFERENCES

- [ADF95] Jürgen Avenhaus, Jörg Denzinger, and Matthias Fuchs. DISCOUNT: A system for distributed equational deduction. In Jieh Hsiang, editor, *RTA-95*, volume 914 of *LNCS*, pages 397–402. Springer, 1995.
- [And71] Peter B. Andrews. Resolution in type theory. *J. Symb. Log.*, 36(3):414–432, 1971.
- [And01] Peter B. Andrews. Classical type theory. In John Alan Robinson and Andrei Voronkov, editors, *Handbook of Automated Reasoning*, volume II, pages 965–1007. Elsevier and MIT Press, 2001.
- [BBCW18] Alexander Bentkamp, Jasmin Christian Blanchette, Simon Cruanes, and Uwe Waldmann. Superposition for lambda-free higher-order logic. In Didier Galmiche, Stephan Schulz, and Roberto Sebastiani, editors, *IJCAR 2018*, volume 10900 of *LNCS*, pages 28–46. Springer, 2018.
- [BBPS16] Jasmin Christian Blanchette, Sascha Böhme, Andrei Popescu, and Nicholas Smallbone. Encoding monomorphic and polymorphic types. *Log. Meth. Comput. Sci.*, 12(4:13):1–52, 2016.
- [BBT⁺21] Alexander Bentkamp, Jasmin Blanchette, Sophie Tourret, Petar Vukmirović, and Uwe Waldmann. Superposition with lambdas. *J. Autom. Reason.*, 2021.
- [BBWW17] Heiko Becker, Jasmin Christian Blanchette, Uwe Waldmann, and Daniel Wand. A transfinite Knuth–Bendix order for lambda-free higher-order terms. In Leonardo de Moura, editor, *CADE-26*, volume 10395 of *LNCS*, pages 432–453. Springer, 2017.
- [BC04] Yves Bertot and Pierre Castéran. *Interactive Theorem Proving and Program Development: Coq’Art: The Calculus of Inductive Constructions*. Texts in Theoretical Computer Science. Springer, 2004.
- [Bee04] Michael Beeson. Lambda logic. In David A. Basin and Michaël Rusinowitch, editors, *IJCAR 2004*, volume 3097 of *LNCS*, pages 460–474. Springer, 2004.
- [Ben18] Alexander Bentkamp. Formalization of the embedding path order for lambda-free higher-order terms. *Archive of Formal Proofs*, 2018. http://isa-afp.org/entries/Lambda_Free_EPO.html.
- [BG94] Leo Bachmair and Harald Ganzinger. Rewrite-based equational theorem proving with selection and simplification. *J. Log. Comput.*, 4(3):217–247, 1994.
- [BG01] Leo Bachmair and Harald Ganzinger. Resolution theorem proving. In John Alan Robinson and Andrei Voronkov, editors, *Handbook of Automated Reasoning*, volume I, pages 19–99. Elsevier and MIT Press, 2001.
- [BKPU16] Jasmin Christian Blanchette, Cezary Kaliszyk, Lawrence C. Paulson, and Josef Urban. Hammering towards QED. *J. Formaliz. Reas.*, 9(1):101–148, 2016.
- [BM14] Christoph Benzmüller and Dale Miller. Automation of higher-order logic. In Jörg H. Siekmann, editor, *Computational Logic*, volume 9 of *Handbook of the History of Logic*, pages 215–254. Elsevier, 2014.
- [BN10] Sascha Böhme and Tobias Nipkow. Sledgehammer: Judgement Day. In Jürgen Giesl and Reiner Hähnle, editors, *IJCAR 2010*, volume 6173 of *LNCS*, pages 107–121. Springer, 2010.
- [BP11] François Bobot and Andrei Paskevich. Expressing polymorphic types in a many-sorted language. In Cesare Tinelli and Viorica Sofronie-Stokkermans, editors, *FroCoS 2011*, volume 6989 of *LNCS*, pages 87–102. Springer, 2011.
- [BP13] Jasmin Christian Blanchette and Andrei Paskevich. TFF1: The TPTP typed first-order form with rank-1 polymorphism. In Maria Paola Bonacina, editor, *CADE-24*, volume 7898 of *LNCS*, pages 414–420. Springer, 2013.
- [BPTF08] Christoph Benzmüller, Lawrence C. Paulson, Frank Theiss, and Arnaud Fietzke. LEO-II—A cooperative automatic theorem prover for higher-order logic. In Alessandro Armando, Peter

- Baumgartner, and Gilles Dowek, editors, *IJCAR 2008*, volume 5195 of *LNCS*, pages 162–170. Springer, 2008.
- [BR13] Miquel Bofill and Albert Rubio. Paramodulation with non-monotonic orderings and simplification. *J. Autom. Reason.*, 50(1):51–98, 2013.
- [BR19] Ahmed Bhayat and Giles Reger. Restricted combinatory unification. In Pascal Fontaine, editor, *CADE-27*, volume 11716 of *LNCS*, pages 74–93. Springer, 2019.
- [BR20] Ahmed Bhayat and Giles Reger. A combinator-based superposition calculus for higher-order logic. In Nicolas Peltier and Viorica Sofronie-Stokkermans, editors, *IJCAR 2020, Part I*, volume 12166 of *LNCS*, pages 278–296. Springer, 2020.
- [Bra75] D. Brand. Proving theorems with the modification method. *SIAM J. Comput.*, 4:412–430, 1975.
- [BREO⁺19] Haniel Barbosa, Andrew Reynolds, Daniel El Ouraoui, Cesare Tinelli, and Clark Barrett. Extending SMT solvers to higher-order logic. In Pascal Fontaine, editor, *CADE-27*, volume 11716 of *LNCS*, pages 35–54. Springer, 2019.
- [Bro12] Chad E. Brown. Satallax: An automatic higher-order prover. In Bernhard Gramlich, Dale Miller, and Uli Sattler, editors, *IJCAR 2012*, volume 7364 of *LNCS*, pages 111–117. Springer, 2012.
- [BWW17] Jasmin Christian Blanchette, Uwe Waldmann, and Daniel Wand. A lambda-free higher-order recursive path order. In Javier Esparza and Andrzej S. Murawski, editors, *FoSSaCS 2017*, volume 10203 of *LNCS*, pages 461–479. Springer, 2017.
- [Chu40] Alonzo Church. A formulation of the simple theory of types. *J. Symb. Log.*, 5(2):56–68, 1940.
- [CP03] Iliano Cervesato and Frank Pfenning. A linear spine calculus. *J. Log. Comput.*, 13(5):639–688, 2003.
- [Cru15] Simon Cruanes. *Extending Superposition with Integer Arithmetic, Structural Induction, and Beyond*. Ph.D. thesis, École polytechnique, 2015.
- [Cru17] Simon Cruanes. Superposition with structural induction. In Clare Dixon and Marcelo Finger, editors, *FroCoS 2017*, volume 10483 of *LNCS*, pages 172–188. Springer, 2017.
- [Cza16] Łukasz Czajka. Improving automation in interactive theorem provers by efficient encoding of lambda-abstractions. In Jeremy Avigad and Adam Chlipala, editors, *CPP 2016*, pages 49–57. ACM, 2016.
- [DH86] Vincent J. Digricoli and Malcolm C. Harrison. Equality-based binary resolution. *J. ACM*, 33(2):253–289, 1986.
- [DHK95] Gilles Dowek, Thérèse Hardin, and Claude Kirchner. Higher-order unification via explicit substitutions (extended abstract). In *LICS '95*, pages 366–374. IEEE Computer Society, 1995.
- [Dou93] Daniel J. Dougherty. Higher-order unification via combinators. *Theor. Comput. Sci.*, 114(2):273–298, 1993.
- [Fit96] Melvin Fitting. *First-Order Logic and Automated Theorem Proving*. Springer-Verlag, 2nd edition, 1996.
- [Fit02] Melvin Fitting. *Types, Tableaux, and Gödel's God*. Kluwer, 2002.
- [GKKV14] Ashutosh Gupta, Laura Kovács, Bernhard Kragl, and Andrei Voronkov. Extensional crisis and proving identity. In Franck Cassez and Jean-François Raskin, editors, *ATVA 2014*, volume 8837 of *LNCS*, pages 185–200. Springer, 2014.
- [GM93] M. J. C. Gordon and T. F. Melham, editors. *Introduction to HOL: A Theorem Proving Environment for Higher Order Logic*. Cambridge University Press, 1993.
- [Hen50] Leon Henkin. Completeness in the theory of types. *J. Symb. Log.*, 15(2):81–91, 1950.
- [HMZ13] Nao Hirokawa, Aart Middeldorp, and Harald Zankl. Uncurrying for termination and complexity. *J. Autom. Reasoning*, 50(3):279–315, 2013.
- [Hue73] Gérard P. Huet. A mechanization of type theory. In Nils J. Nilsson, editor, *IJCAI-73*, pages 139–146. William Kaufmann, 1973.
- [Hur03] Joe Hurd. First-order proof tactics in higher-order logic theorem provers. In Myla Archer, Ben Di Vito, and César Muñoz, editors, *Design and Application of Strategies/Tactics in Higher Order Logics*, NASA Technical Reports, pages 56–68, 2003.
- [Ker91] Manfred Kerber. How to prove higher order theorems in first order logic. In John Mylopoulos and Raymond Reiter, editors, *IJCAI-91*, pages 137–142. Morgan Kaufmann, 1991.
- [Kop12] Cynthia Kop. *Higher Order Termination: Automatable Techniques for Proving Termination of Higher-Order Term Rewriting Systems*. Ph.D. thesis, Vrije Universiteit Amsterdam, 2012.

- [KV13] Laura Kovács and Andrei Voronkov. First-order theorem proving and Vampire. In Natasha Sharygina and Helmut Veith, editors, *CAV 2013*, volume 8044 of *LNCS*, pages 1–35. Springer, 2013.
- [Lei94] Daniel Leivant. Higher order logic. In Dov M. Gabbay, Christopher J. Hogger, J. A. Robinson, and Jörg H. Siekmann, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 2, Deduction Methodologies*, pages 229–322. Oxford University Press, 1994.
- [Lin14] Fredrik Lindblad. A focused sequent calculus for higher-order logic. In Stéphane Demri, Deepak Kapur, and Christoph Weidenbach, editors, *IJCAR 2014*, volume 8562 of *LNCS*, pages 61–75. Springer, 2014.
- [Mil87] Dale A. Miller. A compact representation of proofs. *Studia Logica*, 46(4):347–370, 1987.
- [MP08] Jia Meng and Lawrence C. Paulson. Translating higher-order clauses to first-order clauses. *J. Autom. Reason.*, 40(1):35–60, 2008.
- [NPW02] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL: A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002.
- [Obe09] Fritz H. Obermeyer. *Automated Equational Reasoning in Nondeterministic λ -Calculi Modulo Theories \mathcal{H}^** . Ph.D. thesis, Carnegie Mellon University, 2009.
- [Pel16] Nicolas Peltier. A variant of the superposition calculus. *Archive of Formal Proofs*, 2016.
- [Rob69] J.A. Robinson. Mechanizing higher order logic. In B. Meltzer and D. Michie, editors, *Machine Intelligence*, volume 4, pages 151–170. Edinburgh University Press, 1969.
- [Rob70] J.A. Robinson. A note on mechanizing higher order logic. In B. Meltzer and D. Michie, editors, *Machine Intelligence*, volume 5, pages 121–135. Edinburgh University Press, 1970.
- [SB18] Alexander Steen and Christoph Benzmüller. The higher-order prover Leo-III. In Didier Galmiche, Stephan Schulz, and Roberto Sebastiani, editors, *IJCAR 2018*, volume 10900 of *LNCS*, pages 108–116. Springer, 2018.
- [SBBT09] Geoff Sutcliffe, Christoph Benzmüller, Chad E. Brown, and Frank Theiss. Progress in the development of automated theorem proving for higher-order logic. In Renate A. Schmidt, editor, *CADE-22*, volume 5663 of *LNCS*, pages 116–130. Springer, 2009.
- [SBP13] Nik Sultana, Jasmin Christian Blanchette, and Lawrence C. Paulson. LEO-II and Satallax on the Sledgehammer test bench. *J. Applied Logic*, 11(1):91–102, 2013.
- [Sch02] Stephan Schulz. E - a brainiac theorem prover. *AI Commun.*, 15(2-3):111–126, 2002.
- [Sch12] Stephan Schulz. Fingerprint indexing for paramodulation and rewriting. In Bernhard Gramlich, Dale Miller, and Uli Sattler, editors, *IJCAR 2012*, volume 7364 of *LNCS*, pages 477–483. Springer, 2012.
- [Sch13] Stephan Schulz. System description: E 1.8. In Kenneth L. McMillan, Aart Middeldorp, and Andrei Voronkov, editors, *LPAR-19*, volume 8312 of *LNCS*, pages 735–743. Springer, 2013.
- [SL91] Wayne Snyder and Christopher Lynch. Goal directed strategies for paramodulation. In Ronald V. Book, editor, *RTA-91*, volume 488 of *LNCS*, pages 150–161. Springer, 1991.
- [SS89] Manfred Schmidt-Schauß. Unification in a combination of arbitrary disjoint equational theories. *J. Symb. Comput.*, 8:51–99, 1989.
- [SSCB12] Geoff Sutcliffe, Stephan Schulz, Koen Claessen, and Peter Baumgartner. The TPTP typed first-order form with arithmetic. In Nikolaj Bjørner and Andrei Voronkov, editors, *LPAR-18*, volume 7180 of *LNCS*, pages 406–419. Springer, 2012.
- [SST14] Aaron Stump, Geoff Sutcliffe, and Cesare Tinelli. StarExec: A cross-community infrastructure for logic solving. In Stéphane Demri, Deepak Kapur, and Christoph Weidenbach, editors, *IJCAR 2014*, volume 8562 of *LNCS*, pages 367–373. Springer, 2014.
- [SSUP17] Stephan Schulz, Geoff Sutcliffe, Josef Urban, and Adam Pease. Detecting inconsistencies in large first-order knowledge bases. In Leonardo de Moura, editor, *CADE-26*, volume 10395 of *LNCS*, pages 310–325. Springer, 2017.
- [Vää19] Jouko Väänänen. Second-order and higher-order logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2019 edition, 2019.
- [VBCS19] Petar Vukmirović, Jasmin Christian Blanchette, Simon Cruanes, and Stephan Schulz. Extending a brainiac prover to lambda-free higher-order logic. In Tomas Vojnar and Lijun Zhang, editors, *TACAS 2019*, volume 11427 of *LNCS*, pages 192–210. Springer, 2019.

- [Wan17] Daniel Wand. *Superposition: Types and Polymorphism*. Ph.D. thesis, Universität des Saarlandes, 2017.
- [WDF⁺09] Christoph Weidenbach, Dilyana Dimova, Arnaud Fietzke, Rohit Kumar, Martin Suda, and Patrick Wischnewski. SPASS version 3.5. In Renate A. Schmidt, editor, *CADE-22*, volume 5663 of *LNCS*, pages 140–145. Springer, 2009.
- [WTRB20] Uwe Waldmann, Sophie Tourret, Simon Robillard, and Jasmin Blanchette. A comprehensive framework for saturation theorem proving. In Nicolas Peltier and Viorica Sofronie-Stokkermans, editors, *IJCAR 2020, Part I*, volume 12166 of *LNCS*, pages 316–334. Springer, 2020.