Towards a Model Theory of Ordered Logics: Expressivity and Interpolation

Bartosz Bednarczyk ✉ 🏠 📙
Computational Logic Group, Technische Universität Dresden, Germany
Institute of Computer Science, University of Wrocław, Poland

Reijo Jaakkola ✉ 🏠 📙
Tampere University, Finland

---- **Abstract** ----

We consider the family of guarded and unguarded ordered logics, that constitute a recently redis-covered family of decidable fragments of first-order logic (FO), in which the order of quantification of variables coincides with the order in which those variables appear as arguments of predicates. While the complexities of their satisfiability problems are now well-established, their model theory, however, is poorly understood. Our paper aims to provide some insight into it.

We start by providing suitable notions of bisimulation for ordered logics. We next employ bisimulations to compare the relative expressive power of ordered logics, and to characterise our logics as bisimulation-invariant fragments of FO à la van Benthem.

Afterwards, we study the Craig Interpolation Property (CIP). We refute yet another claim from the infamous work by Purdy, by showing that the fluted and forward fragments do not enjoy CIP. We complement this result by showing that the ordered fragment and the guarded ordered logics enjoy CIP. These positive results rely on novel and quite intricate model constructions, which take full advantage of the "forwardness" of our logics.

1 Introduction

An ongoing research in computational logic has lead to discovery of new decidable fragments of first-order logics (FO) that extend modal and description logics. The main ideas that were proposed in the past involve: restricting the number of variables [8], relativised quantifica-tion [1, 26], restricted use of negation [24], relativised negation [3], one-dimensionality and uniformity [11], separateness [25] and ordered quantification [12, 22]. To compare aforemen-tioned logics, the authors of [1, Section 4.7] proposed a list of desirable meta-properties of logic, which can serve as a yardstick to measure how "nice" a given logic is. We expect a logic L to

(A) be decidable and have the *Finite Model Property* (FMP),

(B) satisfy the Craig Interpolation Property (CIP), *i.e.* for any L-formulae φ, ψ such that $\varphi \models \psi$ there should be an L-formulae χ, called an *interpolant*, that uses only symbols appearing in the common vocabulary of φ and ψ, so that $\varphi \models \chi \models \psi$ holds,

(C) and to satisfy the analog of Łoś-Tarski Preservation Theorem (ŁTPT), *i.e.* any L-formula φ preserved under substructures should be equivalent to some universal L-formula.

It turned out that FO^2 and GF, example logics based on restricted number of variables and relativised quantification, are not "nice" as they do not enjoy CIP [17, Examples 1–2]. In contrast, UNFO and GNFO, the logics based on relativised negation, fulfil the properties (A)–(C), consult: [24, 2, 6]. For one-dimensionality, separateness and ordered quantification we have partial results only.

In this paper we take a closer look at logics enjoying ordered quantification, which have been receiving increasing attention recently [20, 4, 15]. Their syntax can be informally explained as follows. We first require that all variables appearing in formulae are additionally indexed by the quantifier depth and then impose a certain restriction on such numbers in variable sequences in atoms. Assuming that $\alpha(\overline{x})$ is in the scope of the n-th quantifier (but not the $(n+1)$-th), in the fluted fragment $\mathsf{L_{suf}}$ of Quine [22] (resp. in the ordered fragment $\mathsf{L_{pre}}$ by Herzig [12][1]) the tuple $\overline{x}$ is required to be a suffix (resp. a prefix) of the sequence $x_1, x_2, \ldots, x_n$. The forward fragment $\mathsf{L_{inf}}$ [4] is more liberal and allows infixes in place of suffixes or prefixes. An example formula $\varphi \in (\mathsf{L_{suf}} \cap \mathsf{L_{inf}}) \setminus \mathsf{L_{pre}}$ is given below:

1. No student admires every professor.

$$\forall x_1\ (\mathrm{student}(x_1) \to \neg\forall x_2\ (\mathrm{professor}(x_2) \to \mathrm{admires}(x_1, x_2)))$$

2. No lecturer introduces any professor to every student.

$$\forall x_1\ \mathrm{lecturer}(x_1) \to \neg\exists x_2\ [\mathrm{professor}(x_2) \wedge \forall x_3\ (\mathrm{student}(x_3) \to \mathrm{introduce}(x_1, x_2, x_3))]$$

Next, we provide a few coexamples, *i.e.* formulae that, as stated, do not belong to any of $\mathsf{L_{inf}}, \mathsf{L_{pre}}, \mathsf{L_{suf}}$. The blue colour indicates a mismatch in the variable ordering.

1. The relation isPartOf is transitive.

$$\forall x_1\ \forall x_2\ \forall x_3\ \mathrm{isPartOf}(x_1, x_2) \wedge \mathrm{isPartOf}(x_2, x_3) \to \mathrm{isPartOf}(x_1, x_3)$$

2. A narcissist is a person who loves himself.

$$\forall x_1\ \mathrm{narcissist}(x_1) \to \mathrm{person}(x_1) \wedge \mathrm{loves}(x_1, x_1)$$

3. The binary relation hasChild is the inverse of the hasParent relation.

$$\forall x_1\ \forall x_2\ \mathrm{hasParent}(x_1, x_2) \leftrightarrow \mathrm{hasParent}(x_2, x_1)$$

All of $\mathsf{L_{inf}}, \mathsf{L_{suf}}, \mathsf{L_{pre}}$ are decidable and have the Finite Model Property. Their satisfiability problem is, respectively, TOWER-complete for $\mathsf{L_{inf}}$ and $\mathsf{L_{suf}}$, and PSPACE-complete for $\mathsf{L_{pre}}$. Somehow unexpectedly, the TOWER-completeness of $\mathsf{L_{suf}}$ was established only recently by Pratt-Hartmann et al. [20], after pointing out a mistake in the proof of the exponential-size model of $\mathsf{L_{suf}}$ by Purdy [21] and disproving Purdy's claim of NEXPTIME-completeness of $\mathsf{L_{suf}}$. The model theory of $\mathsf{L_{inf}}, \mathsf{L_{suf}}$, and $\mathsf{L_{pre}}$ is, however, poorly understood. The only results that we are aware of are Purdy's claims that $\mathsf{L_{suf}}$ has CIP [21, Thm. 14] and ŁTPT [21, Corr. 17]. But in the light of previously discovered errors, one should treat Purdy's paper with caution.

---

[1] Strictly speaking, the syntax of $\mathsf{L_{pre}}$ is slightly more liberal than the original syntax of the ordered fragment as defined by Herzig, since the syntax of $\mathsf{L_{pre}}$ allows requantifying variables.

1.1 Our results

This paper kick-starts a project of understanding the model theory of *ordered logics*, by which we mean the logics $\mathsf{L_{pre}}, \mathsf{L_{suf}}$, and $\mathsf{L_{inf}}$ as well as their intersections with the guarded fragment GF [1], focusing on the problems mentioned in the introduction.

In Section 3, we design a suitable notion of bisimulations and compare the relative expressive power of ordered logics. Our proofs employ standard model-theoretic constructions like the Compactness Theorem and ω-saturated structures. Next, we investigate CIP in Section 4, which is the main technical contribution of the paper. First, we focus on interpolation for the fluted and the forward fragments. We show that, surprisingly, $\mathsf{L_{inf}}$ and $\mathsf{L_{suf}}$ do not enjoy CIP, refuting yet another claim from the infamous work of Purdy [21, Thm. 14]. Fortunately, other members of the family of ordered logics enjoy CIP, as shown in Sections 4.2–4.3. We stress here that standard techniques for proving CIP, *e.g.* those based on zig-zag products [19, 14, 2, 16], do not seem to work in our case.[2] This forces us to take a different route: we construct models explicitly by specifying types of tuples.

We believe that our proof methods, which are based on novel and intricate model-theoretic constructions, are very general. In particular, we believe that our CIP proof for guarded ordered logics can serve as a useful meta-technique (or even a heuristic) for (dis)proving CIP for fragments of GF. For instance, the proof can be adopted to fragments with CIP, deriving existing results (*e.g.* for the 2-variable GF [14] or the uniform one-dimensional GF [16]) and its failure gives hints why a certain fragment may not have CIP (*e.g.* in the case of full GF).

2 Preliminaries

Henceforth, we employ standard terminology from (finite and classical) model theory [18, 13]. All the logics considered here will be fragments of the first-order logic (FO) over purely-relational equality-free vocabularies, under the usual syntax and semantics.

We fix a countably infinite set of variables $\{x_i \mid i \in \mathbb{N}\}$ and throughout this paper all the formulas use only variables from this set. With $\mathsf{sig}(\varphi)$ we denote the set of relational symbols appearing in φ. We use $\mathsf{ar}(\mathrm{R})$ to denote the arity of R. For a logic L and a signature σ we use $\mathsf{L}[\sigma]$ in place of $\{\varphi \in \mathsf{L} \mid \mathsf{sig}(\varphi) \subseteq \sigma\}$. The k-variable fragment of L (*i.e.* employing only the variables $x_1, x_2, \ldots, x_k$) is denoted L^k. We write $\varphi(\overline{x})$ to indicate that all free variables from φ are members of $\overline{x}$. If $\overline{x}$ contains precisely the free variables of φ, then we will emphasise this separately. Given a structure $\mathfrak{A}$ and $B \subseteq A$, we will use $\mathfrak{A} \restriction B$ to denote the *substructure* of $\mathfrak{A}$ that B induces.

Tuples and subsequences.   An n-tuple is a tuple with n elements. The 0-tuple is denoted with ϵ. We use $\overline{x}_{i \ldots j}$ to denote the $(j{-}i{+}1)$-tuple $x_i, x_{i+1}, \ldots, x_j$. We say that $\overline{x}_{i \ldots j}$ is an infix of a tuple $\overline{x}_{k \ldots l}$ if $k \leq i \leq j \leq l$ holds. If, in addition, $k = i$ (resp. $j = l$) we say that $\overline{x}_{i \ldots j}$ is a prefix (resp. suffix) of $\overline{x}_{k \ldots l}$. We use the word *affix* as a place-holder for the words *pre*fix, *suf*fix or *inf*ix. For a set S, we write $\overline{x} \sqsubseteq S$ iff $x_i \in S$ for all indices $1 \leq i \leq |\overline{x}|$, where $|\overline{x}|$ denotes the length of $\overline{x}$. A tuple $\overline{\mathrm{a}} \sqsubseteq A$ is σ-*live* in $\mathfrak{A}$ if $|\overline{\mathrm{a}}| \leq 1$ or $\overline{\mathrm{a}} \in \mathrm{R}^{\mathfrak{A}}$ for some $\mathrm{R} \in \sigma$.

Logics.   We next introduce the logics $\mathsf{L_{affix}} \in \{\mathsf{L_{pre}}, \mathsf{L_{suf}}, \mathsf{L_{inf}}\}$. We start from $\mathsf{L_{suf}}$, which for technical reasons we need to define separately from $\mathsf{L_{pre}}$ and $\mathsf{L_{inf}}$. For every $n \in \mathbb{N}$, we define the set $\mathsf{L_{suf}}(n)$ as follows:

---

[2] For logics that are closed under negation on the level of formulas, zig-zag constructions seem to work only if the logics are *one-dimensional* and *uniform*, see [16] for more details. None of our logics are one-dimensional nor uniform.

- an atom $\alpha(\overline{x})$ is in $\mathsf{L}_{\mathsf{suf}}(n)$ if $\overline{x}$ is a suffix of $\overline{x}_{1\ldots n}$,
- $\mathsf{L}_{\mathsf{suf}}(n)$ is closed under Boolean connectives $\wedge, \vee, \neg, \rightarrow$,
- if φ is in $\mathsf{L}_{\mathsf{suf}}(n{+}1)$ then $\exists x_{n+1}\, \varphi$ and $\forall x_{n+1}\, \varphi$ are in $\mathsf{L}_{\mathsf{suf}}(n)$.

We put $\mathsf{L}_{\mathsf{suf}} := \mathsf{L}_{\mathsf{suf}}(0)$, which is exclusively composed of sentences.

To define the fragments $\mathsf{L} \in \{\mathsf{L}_{\mathsf{pre}}, \mathsf{L}_{\mathsf{inf}}\}$, for every $n \in \mathbb{N}$ we define the set of $\mathsf{L}(n)$ as follows:

- an atom $\alpha(\overline{x})$ is in $\mathsf{L}_{\mathsf{pre}}(n)$ if $\overline{x}$ is a prefix of $\overline{x}_{1\ldots n}$ and in $\mathsf{L}_{\mathsf{inf}}(n)$ if $\overline{x}$ is an infix of $\overline{x}_{1\ldots n}$.
- if $\varphi \in \mathsf{L}(n_1)$ and $\psi \in \mathsf{L}(n_2)$, then for all $n \geq \max\{n_1, n_2\}$ we have that $\neg \varphi, (\varphi \vee \psi), (\varphi \wedge \psi), (\varphi \rightarrow \psi)$ are in $\mathsf{L}(n)$.
- if φ is in $\mathsf{L}(n{+}1)$ then $\exists x_{n+1}\, \varphi$ and $\forall x_{n+1}\, \varphi$ are in $\mathsf{L}(n)$.

We set $\mathsf{L} := \mathsf{L}(0)$, which is exclusively composed of sentences. We stress that in contrast to $\mathsf{L}_{\mathsf{suf}}$, the logics $\mathsf{L} \in \{\mathsf{L}_{\mathsf{pre}}, \mathsf{L}_{\mathsf{inf}}\}$ allow us to requantify variables. We recommend the reader to employ the above definition to show that $\forall x_1 \forall x_2 \forall x_3 (\mathrm{R}(x_1 x_2 x_3) \rightarrow (\mathrm{A}(x_1) \wedge \exists x_2 \exists x_3 \mathrm{S}(x_1 x_2 x_3))) \in \mathsf{L}_{\mathsf{inf}}$.

Notice that if $\varphi(\overline{x}) \in \mathsf{L}_{\mathsf{affix}}(n)$, where $\overline{x}$ lists all the free variables of φ in order (with respect to their indices), then $\overline{x}$ is an affix of the tuple $\overline{x}_{1\ldots n}$. The logics $\mathsf{L}_{\mathsf{pre}}, \mathsf{L}_{\mathsf{suf}}$, and $\mathsf{L}_{\mathsf{inf}}$ were studied under the names of ordered [12], fluted [22], and forward [4] fragments. The *guarded* counterparts $\mathsf{G}_{\mathsf{affix}}$ of $\mathsf{L}_{\mathsf{affix}}$, are defined as the intersection of $\mathsf{L}_{\mathsf{affix}}$ and the guarded fragment GF [1], *i.e.* by imposing that blocks of quantifiers are relativised by atoms (recalled below). Abusing notation, we speak about all these logics collectively as *ordered logics*.

For reader's convenience we recall that GF is the smallest fragment of FO such that:

- Every atomic formula is in GF;
- GF is closed under boolean connectives $\wedge, \vee, \neg, \rightarrow$;
- If $\varphi(\overline{x}, \overline{y})$ is in GF and $\alpha(\overline{x}, \overline{y})$ is an atom containing all free variables of φ then both $\forall \overline{y}\, (\alpha(\overline{x}, \overline{y}) \rightarrow \varphi(\overline{x}, \overline{y}))$ and $\exists \overline{y}\, (\alpha(\overline{x}, \overline{y}) \wedge \varphi(\overline{x}, \overline{y}))$ are in GF;
- If $\varphi(x)$ has only a single free-variable x, then $\forall x\, \varphi$ and $\exists x\, \varphi$ are in GF.

The atoms α, appearing in the 3rd item of the above definition is called a *guard*.

For a finite signature σ and $n \in \mathbb{N}$, a (σ, n)-*affix-type* is a conjunction of atoms with n free variables $\overline{x}_{1\ldots n}$, in which for every $\mathrm{R} \in \sigma$ and every affix $\overline{x}_{l\ldots k}$ of $\overline{x}_{1\ldots n}$, of length $\mathsf{ar}(\mathrm{R})$, exactly one of $\mathrm{R}(\overline{x}_{l\ldots k}), \neg \mathrm{R}(\overline{x}_{l\ldots k})$ appears as a conjunct. For a σ-structure $\mathfrak{A}$ and a tuple $\overline{\mathsf{a}} \mathrel{\sqsubseteq} A$ with $\mathsf{tp}^{\mathsf{L}_{\mathsf{affix}}[\sigma]}_{\mathfrak{A}}(\overline{\mathsf{a}})$ we denote the *unique* $(\sigma, |\overline{\mathsf{a}}|)$-affix-type realised by $\overline{\mathsf{a}}$ in $\mathfrak{A}$.

2.1 Model Checking

Before jumping into the main part of the paper, we would like to point out some results on the combined complexity of model checking problems of ordered logics, since these seem to be missing from the literature. In what follows we will employ the *matrix encoding of structure*, that is a standard encoding in finite model theory [18, p. 88]. Given a $\{\mathrm{R}_1, \ldots, \mathrm{R}_m\}$-structure $\mathfrak{A}$ with a linearly-ordered domain A, by its *matrix encoding* we mean a binary string $\mathrm{menc}(\mathfrak{A}) := 0^n 1\,\mathrm{menc}(\mathrm{R}_1) \ldots \mathrm{menc}(\mathrm{R}_m)$, where $\mathrm{menc}(\mathrm{R}_i)$ is a binary sequence of length $|A|^{\mathsf{ar}(\mathrm{R}_i)}$, in which the j-th bit is 1 iff the j-th tuple in the lexicographic ordering of $|A|^{\mathsf{ar}(\mathrm{R}_i)}$ belongs to $\mathrm{R}_i^{\mathfrak{A}}$.

The following theorem collects our complexity results. We have not tried to optimise the upper bounds for $\mathsf{G}_{\mathsf{pre}}$ and $\mathsf{L}_{\mathsf{pre}}$: it is quite possible that they can be improved further.

▶ **Theorem 1.** *Under the matrix encoding of structures, the combined complexity of the model-checking problem for a logic L is*

1. *decidable in PTIME for $\mathsf{G}_{\mathsf{pre}}$ and $\mathsf{L}_{\mathsf{pre}}$,*
2. *PTIME-complete for $\mathsf{L} \in \{\mathsf{G}_{\mathsf{suf}}, \mathsf{G}_{\mathsf{inf}}, \mathsf{L}_{\mathsf{suf}}\}$, and*
3. *PSPACE-complete for $\mathsf{L} = \mathsf{L}_{\mathsf{inf}}$.*

Proof. The upper bound for $\mathsf{G}_{\mathsf{pre}}$ follows from the second item while the upper bound for $\mathsf{L}_{\mathsf{pre}}$ is proved in [5, App. A.1]. For the second item, the lower bound follows for all of the logics from the fact that they embed standard modal logic, for which the combined complexity is PTIME-complete [10, Cor. 3.1.7]. For $\mathsf{G}_{\mathsf{suf}}$ and $\mathsf{G}_{\mathsf{inf}}$ matching upper bounds follow from the fact that the combined complexity of the guarded fragment is PTIME-complete, while for $\mathsf{L}_{\mathsf{suf}}$ the matching upper bound is proved in [5, App. A.1]. Finally, for the third item, the upper bound follows from the fact that the combined complexity of FO is PSPACE-complete [7], while the matching lower bound follows from the fact that $\mathsf{L}_{\mathsf{inf}}$ contains monadic FO, for which the combined complexity of model-checking is PSPACE-complete [18, p. 99]. ◀

The matrix encoding is not the only natural way of encoding models. Another option would be to use the *list/database encoding* of models, where one essentially encodes relations by listing the tuples that they contain, as opposed to describing their adjacency matrices. It is easy to see that, if there is no bound on the arities of the relation symbols, then the list encoding of a model can be exponentially more succinct than its matrix encoding. Our proofs for the upper bounds of $\mathsf{L}_{\mathsf{pre}}$ and $\mathsf{L}_{\mathsf{suf}}$ are heavily dependent on the fact that we are using the matrix encoding of models, and hence it is conceivable that the complexities are higher if we are using list encoding.[3] We leave the related investigations as a very interesting future research direction.

3 Expressive power

We study the relative expressive power of ordered logics with a suitable notion of bisimulations.

▶ **Definition 2.** *A non-empty set* $\mathcal{Z} \subseteq \bigcup_{n<\omega}(A^n \times B^n)$ *is a* $\mathsf{L}_{\mathsf{affix}}[\sigma]$-*bisimulation between pointed structures* $\mathfrak{A}, \overline{\mathrm{a}}$ *and* $\mathfrak{B}, \overline{\mathrm{b}}$, *where* $|\overline{\mathrm{a}}| = |\overline{\mathrm{b}}|$, *if and only if* $(\overline{\mathrm{a}}, \overline{\mathrm{b}}) \in \mathcal{Z}$ *and for all* $(\overline{\mathrm{c}}, \overline{\mathrm{d}}) \in \mathcal{Z}$ *the following conditions hold:*

(atomic harmony) $\mathsf{tp}_{\mathfrak{A}}^{\mathsf{L}_{\mathsf{affix}}[\sigma]}(\overline{\mathrm{c}}) = \mathsf{tp}_{\mathfrak{B}}^{\mathsf{L}_{\mathsf{affix}}[\sigma]}(\overline{\mathrm{d}})$.

(forth) *For a (possibly empty) affix* $\overline{\mathrm{c}}_{i\ldots j}$ *of* $\overline{\mathrm{c}}$ *and* $\mathrm{e} \in A$ *there is* $\mathrm{f} \in B$ *s.t.* $(\overline{\mathrm{c}}_{i\ldots j}\mathrm{e}, \overline{\mathrm{d}}_{i\ldots j}\mathrm{f}) \in \mathcal{Z}$.

(back) *For a (possibly empty) affix* $\overline{\mathrm{d}}_{i\ldots j}$ *of* $\overline{\mathrm{d}}$ *and* $\mathrm{f} \in B$ *there is* $\mathrm{e} \in A$ *s.t.* $(\overline{\mathrm{c}}_{i\ldots j}\mathrm{d}, \overline{\mathrm{d}}_{i\ldots j}\mathrm{f}) \in \mathcal{Z}$.

For $\mathsf{G}_{\mathsf{affix}}$, we replace the conditions (forth), (back) by their guarded counterparts:

(gforth) For a (possibly empty) affix $\overline{\mathrm{c}}_{i\ldots j}$ of $\overline{\mathrm{c}}$ and a σ-live tuple $\overline{\mathrm{e}}$ in $\mathfrak{A}$ such that $\overline{\mathrm{c}}_{i\ldots j} = \overline{\mathrm{e}}_{1\ldots j-i+1}$ there is a σ-live tuple $\overline{\mathrm{f}}$ with $\overline{\mathrm{d}}_{i\ldots j} = \overline{\mathrm{f}}_{1\ldots j-i+1}$ and $(\overline{\mathrm{e}}, \overline{\mathrm{f}}) \in \mathcal{Z}$,

(gback) For a (possibly empty) affix $\overline{\mathrm{d}}_{i\ldots j}$ of $\overline{\mathrm{d}}$ and a σ-live tuple $\overline{\mathrm{f}}$ in $\mathfrak{B}$ such that $\overline{\mathrm{d}}_{i\ldots j} = \overline{\mathrm{f}}_{1\ldots j-i+1}$ there is a σ-live tuple $\overline{\mathrm{e}}$ with $\overline{\mathrm{c}}_{i\ldots j} = \overline{\mathrm{e}}_{1\ldots j-i+1}$ and $(\overline{\mathrm{e}}, \overline{\mathrm{f}}) \in \mathcal{Z}$,

For a logic L and a finite signature σ, we write $\mathfrak{A} \equiv_{\mathsf{L}[\sigma]} \mathfrak{B}$ if $\mathfrak{A}$ and $\mathfrak{B}$ satisfy the same $\mathsf{L}[\sigma]$-sentences, and we write $\mathfrak{A} \sim_{\mathsf{L}[\sigma]} \mathfrak{B}$ if there is an $\mathsf{L}[\sigma]$-bisimulation between $\mathfrak{A}$ and $\mathfrak{B}$. If $|\overline{\mathrm{a}}| = |\overline{\mathrm{b}}|$, we use $\mathfrak{A}, \overline{\mathrm{a}} \equiv_{\mathsf{L}[\sigma]} \mathfrak{B}, \overline{\mathrm{b}}$ to denote that for every (possibly empty) affix $\overline{\mathrm{a}}_{i\ldots j}$ of $\overline{\mathrm{a}}$ and $\varphi(\overline{x}_{i\ldots j}) \in \mathsf{L}[\sigma]$, where $\overline{x}_{i\ldots j}$ is an affix of $(x_1, \ldots, x_n)$, we have that $\mathfrak{A} \models \varphi(\overline{\mathrm{a}}_{i\ldots j})$ if and only if $\mathfrak{B} \models \varphi(\overline{\mathrm{b}}_{i\ldots j})$. For the next lemma consult [5, App. B.1]

▶ **Lemma 3.** *Let* $\mathsf{L} \in \{\mathsf{L}_{\mathsf{affix}}, \mathsf{G}_{\mathsf{affix}}\}$. *Then* $\mathfrak{A}, \overline{\mathrm{a}} \sim_{\mathsf{L}[\sigma]} \mathfrak{B}, \overline{\mathrm{b}}$ *implies* $\mathfrak{A}, \overline{\mathrm{a}} \equiv_{\mathsf{L}[\sigma]} \mathfrak{B}, \overline{\mathrm{b}}$. *The converse holds over* ω-*saturated* $\mathfrak{A}$ *and* $\mathfrak{B}$.

---

[3] They can not decrease, because a list encoding of a model can always be constructed efficiently from its matrix encoding.

A logic L_2 is *at least as expressive as* a logic L_1 (written $L_1 \preceq L_2$) if for all $\varphi \in L_1$ there is a $\psi \in L_2$ such that $\varphi \equiv \psi$. We write $L_1 \approx L_2$ iff $L_1 \preceq L_2$ and $L_2 \preceq L_1$. In case $L_1 \not\preceq L_2$ and $L_2 \not\preceq L_1$ we call L_1 and L_2 incomparable. Lastly, $L_1 \prec L_2$ denotes that L_2 is *strictly more expressive* than L_1, *i.e.* $L_1 \preceq L_2$ and $L_1 \not\approx L_2$. Note that, by definition, all the considered fragments L satisfy $L \preceq L_{\mathsf{inf}}$ and $L \prec \mathsf{FO}$ (every such L is decidable). Moreover, $\mathsf{G_{affix}} \prec \mathsf{L_{affix}}$ is a consequence of $\forall x_1 \forall x_2 R(x_1, x_2)$ not being $\mathsf{GF}[\{R\}]$-definable (which is well-known and follows from the fact that GF has the tree-model property). Our results are as follows:

▶ **Theorem 4.** *(a)* $\mathsf{L_{pre}} \prec \mathsf{L_{suf}} \approx \mathsf{L_{inf}} \prec \mathsf{FO}$, *(b)* $\mathsf{G_{affix}} \prec \mathsf{L_{affix}}$ *for all affixes, (c)* $\mathsf{G_{suf}} \prec \mathsf{G_{inf}}$, *(d)* $\mathsf{G_{pre}} \prec \mathsf{G_{inf}}$, *and (e) otherwise the logics are incomparable.*

Proof. Full proofs are in [5, App. B.2]. The relationships between different logics with separating examples (omitting trivial examples due to guardedness) are depicted below. With φ_{pre} we denote the formula $\forall x_1 x_2 x_3 \ R(x_1 x_2 x_3) \to S(x_1 x_2)$, while φ_{suf} denotes $\forall x_1 x_2 x_3 \ R(x_1 x_2 x_3) \to T(x_2 x_3)$. Solid (resp. dashed) arrows from L_1 to L_2 denote that $L_1 \prec L_2$ holds (resp. that the logics are incomparable).
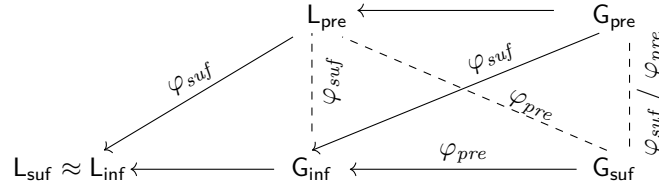
The equi-expressivity of $\mathsf{L_{inf}}$ and $\mathsf{L_{suf}}$ is an easy observation: we turn each maximally nested subformulae into DNF and push the atoms violating the definition of $\mathsf{L_{affix}}$ outside.   ◀

Knowing the relative expressive power of our logics, we would like to characterise them as bisimulation-invariant fragments of FO, as was done with other decidable logics, see *e.g.* [9]. Given a formula $\varphi(\overline{x}) \in L$, we say that it is $\sim_L$-invariant iff for all $\mathfrak{A}, \overline{a} \sim_L^{\mathsf{sig}(\varphi)} \mathfrak{B}, \overline{b}$ we have $\mathfrak{A} \models \varphi(\overline{a}) \Leftrightarrow \mathfrak{B} \models \varphi(\overline{b})$. L is $\sim_L$-invariant iff all its formulae are $\sim_L$-invariant. We will next show that $\mathsf{L_{affix}}$ (resp. $\mathsf{G_{affix}}$) are exactly the $\sim_{\mathsf{L_{affix}}}$- (resp. $\sim_{\mathsf{G_{affix}}}$-) invariant fragments of FO. This confirms that our notion of bisimulation is the right one.

▶ **Theorem 5.** *Let* $L \in \{\mathsf{L_{affix}}, \mathsf{G_{affix}}\}$ *and let* $\varphi(\overline{x})$ *be a* $\sim_L$-*invariant* FO *formula. Then there exists a formula* $\psi(\overline{x})$ *in* L *which is equivalent with* $\varphi(\overline{x})$.

Proof. We follow standard proof methods, see *e.g.* [2, Thm. 3.2]. Suppose $\varphi(x_1, \ldots, x_n) \in \mathsf{FO}$ is $\sim_L$-invariant, where $\overline{x} = (x_1, \ldots, x_n)$ enumerates precisely the set of free variables of φ. The case when φ is unsatisfiable φ is trivial, thus assume otherwise. Consider the set $\Gamma := \{\chi(\overline{x}_{i \ldots j}) \in L \mid \varphi(\overline{x}) \models \chi(\overline{x}_{i \ldots j})\}$. Clearly $\varphi(\overline{x}) \models \Gamma$. Since FO is compact, it suffices to show that $\Gamma \models \varphi(\overline{x})$. Let $\mathfrak{A}$ be a structure and $\overline{a} \in A^n$ so that $\mathfrak{A} \models \chi(\overline{a}_{i \ldots j})$, for every $\chi(\overline{x}_{i \ldots j}) \in \Gamma$. Next, consider the set $\Sigma := \{\chi(\overline{x}_{i \ldots j}) \in L \mid \mathfrak{A} \models \chi(\overline{a}_{i \ldots j})\}$. Again, by compactness of FO we can show that $\Sigma \cup \{\varphi\}$ is consistent. Take a structure $\mathfrak{B}$ and $\overline{b} \in B^n$ so that $\mathfrak{B} \models \varphi(\overline{b})$ and $\mathfrak{B} \models \chi(\overline{b}_{i \ldots j})$, for every $\chi(\overline{x}_{i \ldots j}) \in \Sigma$. Observe that by construction $\mathfrak{A}, \overline{a} \equiv_L \mathfrak{B}, \overline{b}$. Replacing $\mathfrak{A}$ and $\mathfrak{B}$ with their ω-saturated elementary extensions $\hat{\mathfrak{A}}$ and $\hat{\mathfrak{B}}$, we know by Lemma 3 that $\hat{\mathfrak{A}}, \overline{a} \sim_L \hat{\mathfrak{B}}, \overline{b}$. Chasing the resulting diagram we get $\mathfrak{A} \models \varphi(\overline{a})$.   ◀

4   Craig Interpolation

Recall that the Craig Interpolation Property (CIP) for a logic L states that if $\varphi(\overline{x}) \models \psi(\overline{x})$ holds (with φ and ψ having the same free variables), then there is a $\chi(\overline{x}) \in L[\mathsf{sig}(\varphi) \cap \mathsf{sig}(\psi)]$ (an L-*interpolant*) such that $\varphi(\overline{x}) \models \chi(\overline{x})$ and $\chi(\overline{x}) \models \psi(\overline{x})$ hold. We always assume that both φ and ψ are satisfiable, otherwise we can take $\bot$ as a trivial interpolant.

To reason about interpolants we employ the notion of *joint consistency* [23]. We say that L-formulae $\varphi(\overline{x}_{1\ldots n})$ and $\psi(\overline{x}_{1\ldots n})$ (having exactly $\overline{x}_{1\ldots n}$ free) are *jointly-$L[\tau]$-consistent* (or just *jointly consistent* in case $\tau := \mathsf{sig}(\varphi) \cap \mathsf{sig}(\psi)$ and L are known from the context), if there are structures $\mathfrak{A} \models \varphi(\overline{\mathrm{a}})$ and $\mathfrak{B} \models \psi(\overline{\mathrm{b}})$ such that $\mathfrak{A}, \overline{\mathrm{a}} \sim_{\mathsf{L}[\tau]} \mathfrak{B}, \overline{\mathrm{b}}$. The next lemma is classic and links joint consistency and interpolation: see [5, App. C.1].

▶ **Lemma 6.** *Let* $\mathsf{L} \subseteq \mathsf{FO}$, *and let* $\varphi(\overline{x}), \psi(\overline{x}) \in \mathsf{L}$ *with* $\tau := \mathsf{sig}(\varphi) \cap \mathsf{sig}(\psi)$. *Then* $\varphi(\overline{x})$ *and* $\neg\psi(\overline{x})$ *are jointly consistent iff there is no* $\mathsf{L}[\tau]$-*interpolant for* $\varphi(\overline{x}) \models \psi(\overline{x})$.

We simplify the reasoning about ordered logics by employing suitable normal forms. We say that a formula $\varphi(\overline{x})$ from[4] $\mathsf{L}_{\mathsf{pre}}$ (resp. from $\mathsf{G}_{\mathsf{affix}}$) is in *normal form* if it has the shape:

(NForm-$\mathsf{L}_{\mathsf{pre}}$) $\mathrm{H}(\overline{x}) \wedge \bigwedge_{i=1}^{s} \forall \overline{x}_{1\ldots\ell_i} (\alpha_i \to \exists x_{\ell_i+1} \beta_i) \wedge \bigwedge_{j=1}^{t} \forall \overline{x}_{1\ldots\ell_j} (\alpha_j \to \forall x_{\ell_j+1} \beta_j)$,

(NForm-$\mathsf{G}_{\mathsf{affix}}$) $\mathrm{H}(\overline{x}) \wedge \bigwedge_{i=1}^{s} \forall \overline{x}_{1\ldots\ell_i} (\mathrm{R}_i(\overline{x}_{1\ldots\ell_i}) \to \exists \overline{x}_{\ell_i+1\ldots\ell_i+k_i} (\mathrm{S}_i(\overline{x}_{1\ldots\ell_i+k_i}) \wedge \psi_i(\overline{x}_{1\ldots\ell_i+k_i}))) \wedge \bigwedge_{j=1}^{t} \forall \overline{x}_{1\ldots\ell_j} (\mathrm{R}_j(\overline{x}_{1\ldots\ell_j}) \to \psi_j(\overline{x}_{1\ldots\ell_j}) \to \forall \overline{x}_{\ell_j+1\ldots\ell_j'} (\mathrm{T}_j(\overline{x}_{1\ldots\ell_j'}) \to \psi_j'(\overline{x}_{1\ldots\ell_j'})))$,

where $\alpha_i, \alpha_j, \beta_i$ and β_j are quantifier-free $\mathsf{L}_{\mathsf{pre}}$-formulae, $\mathrm{R}_i, \mathrm{R}_j, \mathrm{T}_j$ and H are relational symbols, and ψ_i, ψ_j and ψ_j' are $\mathsf{G}_{\mathsf{affix}}$-formulae. The symbol H is called the *head* of $\varphi(\overline{x})$. We will often speak about *existential/universal requirements* of a formula, meaning the appropriate subformulae with the maximal quantifier prefix $\forall^* \exists^*$ and $\forall^*$. In aforementioned normal forms we implicitly allow various parameters to be zero, *e.g.* in subformulae of the form $\forall \overline{x}_{1\ldots\ell_i} (\alpha_i \to \exists x_{\ell_i+1} \beta_i)$ we allow $\ell_i = 0$, and we agree that the result is $\exists x_{\ell_i+1} \beta_i$.

The following lemma can be shown using standard renaming techniques, in complete analogy to [4, 15], with a minor (but technically tedious) modification in the case of $\mathsf{G}_{\mathsf{suf}}$, see [5, App. C.2].

▶ **Lemma 7.** *Let* $\mathsf{L} \in \{\mathsf{L}_{\mathsf{pre}}, \mathsf{G}_{\mathsf{affix}}\}$, *and take* $\varphi(\overline{x}), \psi(\overline{x}) \in \mathsf{L}$. *Suppose that there are models* $\mathfrak{A}$ *and* $\mathfrak{B}$ *such that* $\mathfrak{A} \models \varphi(\overline{\mathrm{a}})$, $\mathfrak{B} \models \psi(\overline{\mathrm{b}})$ *and* $\mathfrak{A}, \overline{\mathrm{a}} \sim_{\mathsf{L}[\tau]} \mathfrak{B}, \overline{\mathrm{b}}$, *where* $\tau = \mathsf{sig}(\varphi) \cap \mathsf{sig}(\psi)$. *Then there exist formulae* $\varphi'(\overline{x}), \psi'(\overline{x}) \in \mathsf{L}$ *in normal form and extensions* $\mathfrak{A}'$ *and* $\mathfrak{B}'$ *of* $\mathfrak{A}$ *and* $\mathfrak{B}$ *respectively, such that (i)* $\varphi'(\overline{x})$ *and* $\psi'(\overline{x})$ *have the same head* H, *(ii)* $\mathsf{sig}(\varphi') \cap \mathsf{sig}(\psi') = \tau \cup \{\mathrm{H}\}$, *(iii)* $\varphi'(\overline{x}) \models \varphi(\overline{x})$ *and* $\psi'(\overline{x}) \models \psi(\overline{x})$, *and (iii)* $(\mathfrak{A}', \overline{\mathrm{a}}) \sim_{\mathsf{L}[\tau \cup \{\mathrm{H}\}]} (\mathfrak{B}', \overline{\mathrm{b}})$ *holds.*

The following lemma is a useful tool when dealing with interpolation, allowing us to switch our attention to a certain satisfiability problem. Its proof is routine, consult [5, App. C.3].

▶ **Lemma 8.** *Let* $\mathsf{L} \in \{\mathsf{L}_{\mathsf{pre}}, \mathsf{G}_{\mathsf{affix}}\}$. *If for any jointly-consistent* L-*formulae* $\varphi(\overline{x}), \psi(\overline{x})$ *in normal forms from Lemma 7 with the same head, there is* $\mathfrak{A} \models \varphi(\overline{x}) \wedge \psi(\overline{x})$, *then* L *has CIP.*

4.1 Disproving CIP in $\mathsf{L}_{\mathsf{inf}}$ and $\mathsf{L}_{\mathsf{suf}}$

We start our investigation of CIP for $\mathsf{L}_{\mathsf{affix}}$ and $\mathsf{G}_{\mathsf{affix}}$ by further discrediting the infamous work of Purdy [21]. We prove, in stark contrast to [21, Thm. 14], that $\mathsf{L}_{\mathsf{suf}}$ does not have CIP.

▶ **Theorem 9.** $\mathsf{L}_{\mathsf{inf}}$ *and* $\mathsf{L}_{\mathsf{suf}}$ *do not have CIP. More specifically, there are* $\mathsf{L}_{\mathsf{suf}}^2$-*sentences* φ, ψ *with* $\varphi \models \psi$ *but without any* $\mathsf{L}_{\mathsf{inf}}[\mathsf{sig}(\varphi) \cap \mathsf{sig}(\psi)]$-*interpolant.*

Proof. Consider the following $\mathsf{L}_{\mathsf{inf}}^3$-sentences φ and ψ, presented respectively below:

$\forall \overline{x}_{1\ldots 3} [(\mathrm{R}(x_1, x_2) \wedge \mathrm{R}(x_2, x_3)) \to (\mathrm{P}_1(x_1) \wedge \mathrm{P}_2(x_3))] \wedge \forall x_1 \forall x_2 [(\mathrm{P}_1(x_1) \wedge \mathrm{P}_2(x_2)) \to \mathrm{R}(x_1, x_2)]$

$\exists \overline{x}_{1\ldots 3} [\mathrm{R}(x_1, x_2) \wedge \mathrm{R}(x_2, x_3) \wedge \mathrm{Q}_1(x_1) \wedge \mathrm{Q}_2(x_3)] \wedge \forall x_1 \forall x_2 [(\mathrm{Q}_1(x_1) \wedge \mathrm{Q}_2(x_2)) \to \neg\mathrm{R}(x_1, x_2)]$,

---

[4] To avoid notational glitter we will be a bit careless when dealing with formulae with free-variables.

with $\mathfrak{A} \models \varphi$ iff $(R^{\mathfrak{A}} \circ R^{\mathfrak{A}}) \subseteq P_1^{\mathfrak{A}} \times P_2^{\mathfrak{A}}$ and $P_1^{\mathfrak{A}} \times P_2^{\mathfrak{A}} \subseteq R^{\mathfrak{A}}$, and $\mathfrak{B} \models \psi$ iff $Q_1^{\mathfrak{B}} \times Q_2^{\mathfrak{B}} \subseteq B^2 \setminus R^{\mathfrak{B}}$ and there are $(a, b), (b, c) \in R^{\mathfrak{B}}$ with $a \in Q_1^{\mathfrak{B}}$ and $c \in Q_2^{\mathfrak{B}}$. Observe that $\varphi \models \neg\psi$, since φ entails transitivity of R, while ψ entails that this is not the case. But φ and ψ are jointly-$\mathsf{L}_{\mathsf{inf}}[\{R\}]$-consistent (it suffices to take $\mathfrak{A}$ and $\mathfrak{B}$ depicted below, *cf.* [5, App. C.5].
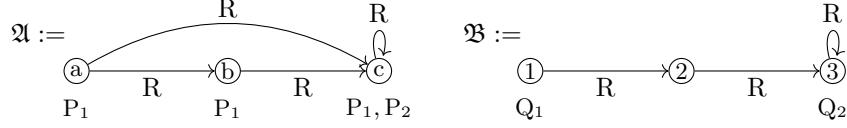
Hence, by Lemma 6 there is no $\mathsf{L}_{\mathsf{inf}}[\{R\}]$-interpolant for $\varphi \models \neg\psi$. By slightly obfuscating φ and ψ (*i.e.* by shifting quantifiers and introducing a unary symbol to get rid of the third variable) we can make our counterexample formulae to be in $\mathsf{L}_{\mathsf{suf}}^2$; consult [5, App. C.6]. ◄

We left open the question whether $\mathsf{L}_{\mathsf{suf}}$ and $\mathsf{L}_{\mathsf{inf}}$ have the *Projective Beth Definability Property* (PBDP). We conjuncture that the answer is *no*, but we haven't found a suitable example yet.

4.2 Restoring CIP in $\mathsf{L}_{\mathsf{pre}}$

Even though $\mathsf{L}_{\mathsf{suf}}$ and $\mathsf{L}_{\mathsf{inf}}$ fail to have CIP, it turns out that $\mathsf{L}_{\mathsf{pre}}$ still has it. To prove interpolation for $\mathsf{L}_{\mathsf{pre}}$, we are going to construct a model for two jointly consistent $\mathsf{L}_{\mathsf{pre}}$ formulae $\varphi(\overline{x})$ and $\psi(\overline{x})$. However, rather than modifying existing amalgamation-based arguments used, for instance, in [19, 14, 2], we will construct our model explicitly by specifying prefix-types for tuples. We feel that our approach, which is more direct in nature than other arguments found in the literature, could potentially be useful also in other contexts.

Take $\varphi(\overline{x})$ and $\psi(\overline{x})$ in normal form (NForm-$\mathsf{L}_{\mathsf{pre}}$) satisfying the premise of Lemma 8. Hence, there are structures $\mathfrak{A}$ and $\mathfrak{B}$ and tuples $\overline{a} \in A^k$ and $\overline{b} \in B^k$ such that that $\mathfrak{A} \models \varphi(\overline{a})$, $\mathfrak{B} \models \psi(\overline{b})$ and $(\mathfrak{A}, \overline{a}) \sim_{\mathsf{L}_{\mathsf{pre}}[\sigma]} (\mathfrak{B}, \overline{b})$, where $\sigma := \mathsf{sig}(\varphi) \cap \mathsf{sig}(\psi)$. Let $\tau := \mathsf{sig}(\varphi) \cup \mathsf{sig}(\psi)$.

We will define a sequence of τ-structures $\mathfrak{U}_1 \leq \ldots \leq \mathfrak{U}_M := \mathfrak{U}$, where $M = \max\{\mathsf{ar}(R) \mid R \in \tau\}$, satisfying the following inductive assumptions: (i) $U_i = \mathbb{N}$, (ii) the interpretation of symbols from τ of arity $> i$ is empty, and (iii) for any i-tuple $\overline{c}$ in $\mathfrak{U}_i$ there are i-tuples $\overline{d}$ in $\mathfrak{A}$ and $\overline{e}$ in $\mathfrak{B}$ so that $(\mathfrak{A}, \overline{d}) \sim_{\mathsf{L}_{\mathsf{pre}}[\sigma]} (\mathfrak{B}, \overline{e})$ and $\mathsf{tp}_{\mathfrak{U}_i}^{\mathsf{L}_{\mathsf{pre}}[\tau]}(\overline{c}) = \mathsf{tp}_{\mathfrak{A}}^{\mathsf{L}_{\mathsf{pre}}[\mathsf{sig}(\varphi)]}(\overline{d}) \cup \mathsf{tp}_{\mathfrak{B}}^{\mathsf{L}_{\mathsf{pre}}[\mathsf{sig}(\psi)]}(\overline{e})$ hold. The last condition guarantees that no tuple $\overline{c}$ of $\mathfrak{U}_i$ violates the universal requirements of φ and ψ, since otherwise the corresponding tuple would violate them, contradicting modelhood of $\mathfrak{A}$ or $\mathfrak{B}$.[5]

For the inductive base, take $\mathfrak{U}_1$ with domain $\mathbb{N}$ and empty interpretation of symbols from τ. Our goal is to realise each $(\mathsf{sig}(\varphi), 1)$-prefix-type, which is realised in $\mathfrak{A}$ and $\mathfrak{B}$, in $\mathfrak{U}_1$ in a careful way, suggested by the inductive assumption. Let t be a $(\mathsf{sig}(\varphi), 1)$-prefix type realised in $\mathfrak{A}$ and let $c \in A$ be some element witnessing it. Since $(\mathfrak{A}, \overline{a}) \sim_{\mathsf{L}_{\mathsf{pre}}[\sigma]} (\mathfrak{B}, \overline{b})$ holds, there exists an element d of $\mathfrak{B}$ so that $\mathsf{tp}_{\mathfrak{A}}^{\mathsf{L}_{\mathsf{pre}}[\sigma]}(c) = \mathsf{tp}_{\mathfrak{B}}^{\mathsf{L}_{\mathsf{pre}}[\sigma]}(d)$. Now we will assign the $(\tau, 1)$-prefix-type $\mathsf{tp}_{\mathfrak{A}}^{\mathsf{L}_{\mathsf{pre}}[\mathsf{sig}(\varphi)]}(c) \cup \mathsf{tp}_{\mathfrak{B}}^{\mathsf{L}_{\mathsf{pre}}[\mathsf{sig}(\psi)]}(d)$ to some element e of $\mathfrak{U}_1$, for which we have not yet assigned a $(\tau, 1)$-prefix-type. For the remaining elements of $\mathfrak{U}_1$, having no $(\tau, 1)$-prefix-type assigned, we assign any of the previously realised types.

---

[5] We note that this claim no longer holds if $\mathsf{L}_{\mathsf{pre}}$ is replaced by either $\mathsf{L}_{\mathsf{suf}}$ or $\mathsf{L}_{\mathsf{affix}}$, which is why the forthcoming construction does not work for these logics.

Suppose then that $\mathfrak{U}_k$ is defined. To define $\mathfrak{U}_{k+1}$, we will start by providing witnesses for the existential requirements of φ and ψ; since the two cases are rather analogous, we will restrict our attention to the former case. Consider an existential requirement $\varphi_i^\exists$ of $\varphi(\overline{x})$ and let $\overline{e} \in U_k^k$ be a k-tuple so that $\mathfrak{U} \models \alpha_i(\overline{e})$. By construction, there exists a tuple $\overline{a} \in A^k$ witnessing $\mathsf{tp}_{\mathfrak{U}_k}^{\mathsf{L}_{\mathsf{pre}}[\mathsf{sig}(\varphi)]}(\overline{e}) = \mathsf{tp}_{\mathfrak{A}}^{\mathsf{L}_{\mathsf{pre}}[\mathsf{sig}(\varphi)]}(\overline{a})$. Since $\mathfrak{A} \models \varphi_i^\exists$, there exists an element $c \in A$ so that $\mathfrak{A} \models \beta_i(\overline{a}, c)$. Due to $(\mathfrak{A}, \overline{a}) \sim_{\mathsf{L}_{\mathsf{pre}}[\sigma]} (\mathfrak{B}, \overline{b})$, we know that there exists an element $d \in B$ satisfying $\mathsf{tp}_{\mathfrak{A}}^{\mathsf{L}_{\mathsf{pre}}[\sigma]}(\overline{a}, c) = \mathsf{tp}_{\mathfrak{B}}^{\mathsf{L}_{\mathsf{pre}}[\sigma]}(\overline{b}, d)$. Now we pick an element $f \in U$ for which we have not yet assigned a $(\tau, k+1)$-prefix-type for the tuple $(\overline{e}, f)$ (recall that the domain of our model is $\mathbb{N}$, so such an element always exists). We assign the following $(\tau, k+1)$-prefix-type to the tuple $(\overline{e}, f)$: $\mathsf{tp}_{\mathfrak{A}}^{\mathsf{L}_{\mathsf{pre}}[\mathsf{sig}(\varphi)]}((\overline{a}, c)) \cup \mathsf{tp}_{\mathfrak{B}}^{\mathsf{L}_{\mathsf{pre}}[\mathsf{sig}(\psi)]}((\overline{b}, d))$. Note that the assigned $(\tau, k+1)$-prefix-type is consistent with the (τ, k)-prefix-type that we assigned to $\overline{e}$. Having assigned witnesses to relevant existential requirements of $\varphi(\overline{x})$ and $\psi(\overline{x})$, there are still $(k+1)$-tuples of elements of $\mathfrak{U}$ for which we have not yet assigned a $(\tau, k+1)$-prefix-type. For those tuples we will assign any $(\tau, k+1)$-prefix-type that we have already assigned to some other $(k+1)$-tuple of elements of $\mathfrak{U}_{k+1}$. This completes the construction of $\mathfrak{U}_{k+1}$.

By construction, it is clear that there exists a tuple $\overline{e}$ of elements of $\mathfrak{U}$ so that $\overline{e} \in H^{\mathfrak{U}}$; in particular, $\mathfrak{U} \models \varphi(\overline{e}) \wedge \psi(\overline{e})$ holds. Thus, by Lemma 8 we conclude:

▶ **Theorem 10.** $\mathsf{L}_{\mathsf{pre}}$ *enjoys the Craig Interpolation Property.*

4.3 Restoring CIP in guarded logics

Finally we turn our attention to the logics $\mathsf{G}_{\mathsf{affix}}$ and present the main contribution of the paper. It will be convenient to employ suitable tree-like models. Intuitively, HATs [4] are just trees in which relations connect elements but only in a level-by-level ascending order; see Figure 1. HAHs are collections of HATs.

▶ **Definition 11.** *A structure* $\mathfrak{T}$ *is a* higher-arity tree *(HAT) if its domain is a prefix-closed subset of sequences from* $\mathbb{N}^*$ *and for all relation symbols* R *we have that* $(d_1, \ldots, d_k) = \overline{d} \in R^{\mathfrak{T}}$ *implies that for each index* $i < k$ *there exists a number* n_i *such that* $d_{i+1} = d_i \cdot n_i$, *where* $d_i \cdot n_i$ *means that the element* n_i *is appended to the sequence* d_i. *A structure* $\mathfrak{H}$ *is a* higher-arity hedge *(HAH) if* $\mathfrak{H}$ *becomes a HAT if extended by a single element* ε.
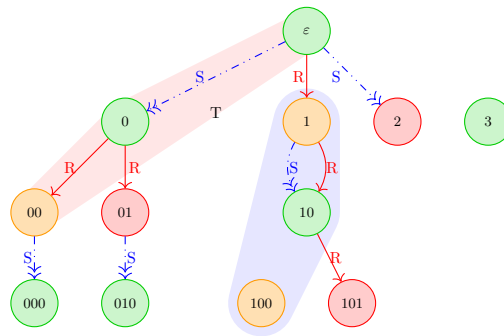
▶ **Figure 1** An example HAT $\mathfrak{T}$. All relations go down lvl-by-lvl. The red area means $(\varepsilon, 0, 00) \in T^{\mathfrak{T}}$.

By a subtree of a HAT $\mathfrak{T}$ rooted at an element d we mean a substructure of $\mathfrak{T}$ with the domain composed of all elements of the form dw for a possibly empty word w. Note that such a subtree is also a HAT after an obvious renaming.

We are going to employ the following lemma, stating that for our purposes we can focus on tree-like models only. Its proof relies on the suitable notion of unravelling, see [5, App. C.4].

▶ **Lemma 12.** *Let a logic* L *be any of* G_{affix}, φ, ψ *be* L-*formulae and* $\sigma := \text{sig}(\varphi) \cap \text{sig}(\psi)$ *containing the predicate* H. *Assume that models* $\mathfrak{A} \models \varphi(\overline{a}), \mathfrak{B} \models \psi(\overline{b})$ *are given such that* $\overline{a} \in H^{\mathfrak{A}}, \overline{b} \in H^{\mathfrak{B}}$ *and* $(\mathfrak{A}, \overline{a}) \sim_{G_{\text{affix}}[\sigma]} (\mathfrak{B}, \overline{b})$ *hold. Then there are HAH models* $\mathfrak{T}_{\mathfrak{A}} \models \varphi(\overline{c}), \mathfrak{T}_{\mathfrak{B}} \models \psi(\overline{d})$ *satisfying* $\overline{c} \in H^{\mathfrak{T}_{\mathfrak{A}}}, \overline{d} \in H^{\mathfrak{T}_{\mathfrak{B}}}$ *and* $(\mathfrak{T}_{\mathfrak{A}}, \overline{c}) \sim_{G_{\text{affix}}[\sigma]} (\mathfrak{T}_{\mathfrak{B}}, \overline{d})$.

Take $\varphi(\overline{x})$, $\psi(\overline{x})$ in form (NForm-G_{affix}) with the same head, satisfying the premise of Lemma 8. We have structures $\mathfrak{A}$ and $\mathfrak{B}$ and tuples $\overline{a} \in A^k$ and $\overline{b} \in B^k$ so that $\mathfrak{A} \models \varphi(\overline{a}), \mathfrak{B} \models \psi(\overline{b})$ and $(\mathfrak{A}, \overline{a}) \sim_{G_{\text{affix}}[\sigma]} (\mathfrak{B}, \overline{b})$, where $\sigma := \text{sig}(\varphi) \cap \text{sig}(\psi)$ and $\tau := \text{sig}(\varphi) \cup \text{sig}(\psi)$. Using Lemma 12 we can assume that $\mathfrak{A}$ and $\mathfrak{B}$ are τ-HAHs. As done before, we aim at constructing a τ-structure being a model of both φ and ψ. To do this, we will construct a growing sequence of τ-HAHs $\mathfrak{U}_0 := \mathfrak{A} \leq \mathfrak{U}_1 \leq \ldots \leq \mathfrak{U}_n \leq \ldots$, whose limit $\mathfrak{U}$ will be a model of $\varphi \wedge \psi$. For simplicity, let us employ the following naming scheme. A tuple $\overline{d}$ from $\mathfrak{U}_n$ ($\overline{d} \sqsubseteq U_n$) is called (a) *n-fresh* if $\overline{d} \sqsubseteq U_{n-1}$ and *n-aged* otherwise, (b) *maximal* if its not an affix of any different σ-live tuple.

A high-level idea of the construction of the sequence $\mathfrak{U}_i$, obfuscated by many challenging technical details, is as follows. Starting from $\mathfrak{A}$ we inductively "complete" types of all σ-live tuples to become proper τ-live tuples. This will help, if done carefully and in a bisimilarity-preserving way, the structure $\mathfrak{U}_i$ to fulfil the universal constraints of φ and ψ, but may introduce tuples without witnesses for the existential constraints. Hence, after each "completion" phase, we will "repair" the obtained structure by "copying" some substructures of $\mathfrak{A}$ and $\mathfrak{B}$ and "gluing" them on existing witness-lacking tuples (providing the required witnesses).

During the construction we will make sure that for every n-aged $\text{sig}(\varphi)$-live (resp. $\text{sig}(\psi)$-live) k-tuple in $\mathfrak{U}$, there exists a k-tuple in $\mathfrak{A}$ (resp. in $\mathfrak{B}$) having equal $(\text{sig}(\varphi), k)$-affix-type (resp. $(\text{sig}(\psi), k)$-affix-type). This will be controlled by means of partial *witness functions* $\text{wit}_{\mathfrak{A}} : U_n \to A, \text{wit}_{\mathfrak{B}} : U_n \to B$, intuitively pinpointing from where a tuple in $\mathfrak{U}$ originated from. To make the construction work, the witness function will fulfil several technical criteria, that are listed below. Conditions (a) and (b) speak about the compatibility of types between a tuple and its witness tuple; this guarantees that no tuple from $\mathfrak{U}_n$ violate the universal requirements of φ and ψ. Conditions (c)–(d) guarantees the satisfaction of the existential requirements of φ and ψ (condition (c) takes care of "local" requirements while (d) handles the "global" ones). Formally, for every n-aged $\overline{c}$ from $\mathfrak{U}_n$ we have that:

(a) If $\overline{c}$ is σ-live then both $\overline{d} := \text{wit}_{\mathfrak{A}}(\overline{c})$ and $\overline{e} := \text{wit}_{\mathfrak{B}}(\overline{c})$ are defined, $(\mathfrak{A}, \overline{d}) \sim_{G_{\text{affix}}[\sigma]}$ $(\mathfrak{B}, \overline{e})$ holds and $\text{tp}_{\mathfrak{U}_n}^{G_{\text{affix}}[\tau]}(\overline{c})$ is equal to $\text{tp}_{\mathfrak{A}}^{G_{\text{affix}}[\text{sig}(\varphi)]}(\overline{d}) \cup \text{tp}_{\mathfrak{B}}^{G_{\text{affix}}[\text{sig}(\psi)]}(\overline{e})$.

(b) If $\overline{c}$ is not σ-live but is $\text{sig}(\varphi)$-live (resp. $\text{sig}(\psi)$-live), then $\overline{d} := \text{wit}_{\mathfrak{A}}(\overline{c})$ (resp. $\overline{d} := \text{wit}_{\mathfrak{B}}(\overline{c})$) is defined, and $\text{tp}_{\mathfrak{U}_n}^{G_{\text{affix}}[\text{sig}(\varphi)]}(\overline{c})$ is equal to $\text{tp}_{\mathfrak{A}}^{G_{\text{affix}}[\text{sig}(\varphi)]}(\overline{d})$ (resp. $\text{tp}_{\mathfrak{B}}^{G_{\text{affix}}[\text{sig}(\psi)]}(\overline{d})$).

(c) if $\overline{c}$ is $\text{sig}(\varphi)$-live (resp. $\text{sig}(\psi)$-live), then for every existential requirement $\lambda := R_i(\overline{x}_{1...\ell_i}) \to \exists \overline{x}_{\ell_i...\ell_i+k_i}(S_i(\overline{x}_{1...i+k_i}) \wedge \theta_i(\overline{x}_{1...\ell_i+k_i}))$ from φ (resp. from ψ) with $\overline{c}$ satisfying the premise of λ, there is a tuple $\overline{d}$ in $\mathfrak{U}_n$ so that $\overline{c}\overline{d}$ satisfies the conclusion of λ.

(d) For every $\text{sig}(\varphi)$-live (resp. $\text{sig}(\psi)$-live) tuple $\overline{d}$ from $\mathfrak{A}$ (resp. from $\mathfrak{B}$) there is a tuple $\overline{e}$ in $\mathfrak{U}_1$ such that $\text{tp}_{\mathfrak{U}_n}^{G_{\text{affix}}[\text{sig}(\varphi)]}(\overline{e}) = \text{tp}_{\mathfrak{A}}^{G_{\text{affix}}[\text{sig}(\varphi)]}(\overline{d})$ (resp. $\text{tp}_{\mathfrak{U}_n}^{G_{\text{affix}}[\text{sig}(\varphi)]}(\overline{e}) = \text{tp}_{\mathfrak{B}}^{G_{\text{affix}}[\text{sig}(\psi)]}(\overline{d})$).

While the following property is not necessary to guarantee that the limit $\mathfrak{U}$ is a model of $\varphi \wedge \psi$, it plays an important technical role in the construction:

(e) If $\overline{d}$ is an n-fresh σ-live tuple such that either $\text{wit}_{\mathfrak{A}}(\overline{d})$ or $\text{wit}_{\mathfrak{B}}(\overline{d})$ is undefined, then for every prefix $\overline{d}_{1...k}$ of $\overline{d}$ that is contained in U_{n-1}, meaning that $\overline{d}_{1...k} \sqsubseteq U_{n-1}$, there exists an n-aged σ-live tuple $\overline{c}$ which contains $\overline{d}_{1...k}$ as its affix.

Using conditions (a)–(d) it follows that $\mathfrak{U} \models \varphi \wedge \psi$, allowing us to conclude (by Lemma 8):

▶ **Theorem 13.** $G_{\text{inf}}, G_{\text{suf}}$ *and* G_{pre} *enjoy the Craig Interpolation Property.*

We will now move on to the construction of $\mathfrak{U}$, described below. We start from the crucial, aforementioned notions of *completions* and *repairs*. Intuitively the *completion* just "completes a type of a tuple" in a bisimulation-preserving way, taking all symbols of τ into account. *Repair* simply "plugs in" certain subtrees from $\mathfrak{A}$ or $\mathfrak{B}$ into $\mathfrak{U}$, providing missing witnesses.

▶ **Definition 14** (completion). *Let $(\mathfrak{T}, \overline{d})$ be a pointed τ-HAH, where $\overline{d}$ is σ-live with $\mathsf{wit}_{\mathfrak{A}}, \mathsf{wit}_{\mathfrak{B}}$ defined. The $\overline{d}$-completion of $\mathfrak{T}$ is obtained from $\mathfrak{T}$ by redefining interpretation of symbols from τ in a min. way so that $\mathsf{tp}_{\mathfrak{T}}^{\mathsf{G}_{\mathsf{affix}}[\tau]}(\overline{d})$ equals $\mathsf{tp}_{\mathfrak{A}}^{\mathsf{G}_{\mathsf{affix}}[\mathsf{sig}(\varphi)]}(\mathsf{wit}_{\mathfrak{A}}(\overline{d})) \cup \mathsf{tp}_{\mathfrak{B}}^{\mathsf{G}_{\mathsf{affix}}[\mathsf{sig}(\psi)]}(\mathsf{wit}_{\mathfrak{B}}(\overline{d}))$.*

▶ **Definition 15** (repair). *Let $(\mathfrak{T}, \overline{c})$ be a pointed τ-HAH with only $\overline{d} := \mathsf{wit}_{\mathfrak{A}}(\overline{c})$ defined, where $\overline{c}$ is σ-live in $\mathfrak{T}$. Suppose also that there is a tuple $\overline{e}$ in $\mathfrak{B}$ such that $(\mathfrak{A}, \overline{d}) \sim_{\mathsf{G}_{\mathsf{affix}}[\sigma]} (\mathfrak{B}, \overline{e})$ holds. The $(\mathfrak{B}, \overline{e})$-repair of $\overline{c}$ is a τ-HAH $\mathfrak{T}'$ obtained from $\mathfrak{T}$ in the following five steps:*
1. *Let $\mathfrak{B}_0$ be the subtree of $\mathfrak{B}$ rooted at the first element of $\overline{e}$.*
2. *Take $\mathfrak{T}'$ to be the union of $\mathfrak{T}$ and $\mathfrak{B}_0$ without $\overline{e}$.*
3. *$\mathfrak{T}'$ will contain $\mathfrak{T}$ as a substructure.*
4. *By identifying $\overline{c}$ with $\overline{e}$, we interpret the relation symbols for tuples of elements of $\mathfrak{T}' \upharpoonright B_0$ in such a way that the resulting substructure of $\mathfrak{T}'$ is isomorphic with $\mathfrak{B}_0$.*
5. *We set $\mathsf{wit}_{\mathfrak{B}}$ on freshly added elements to be the identity on $\mathfrak{B}_0$.*
The substructure $\mathfrak{T}' \upharpoonright (B_0 \cup \overline{c})$ is called a $\overline{c}$-component of $\mathfrak{T}'$. $\mathfrak{T}'$ becomes a HAH after a routine renaming. We define $(\mathfrak{A}, \overline{d})$-repair of $\overline{c}$ analogously.
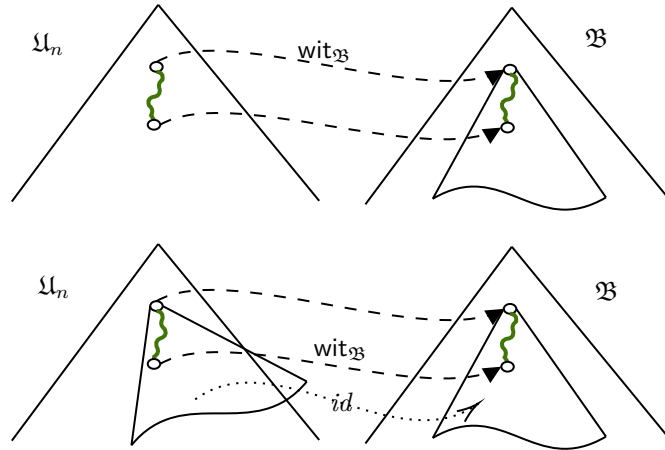
■ **Figure 2** An example structure $\mathfrak{U}_n$ before and after we performed a "$\mathfrak{B}$"-repair.

We proceed with the base of induction, setting first $\mathfrak{U}_0$ to be $\mathfrak{A}$. It will be four-fold.

Base case: Step I. We set up $\mathsf{wit}_{\mathfrak{A}}$ and $\mathsf{wit}_{\mathfrak{B}}$ functions. For $\mathsf{wit}_{\mathfrak{A}}$ we will simply take the identity function. To define $\mathsf{wit}_{\mathfrak{B}}$, we intuitively proceed by traversing $\mathfrak{U}_0$ from top to bottom. More precisely, let $L_0^{\mathfrak{A}}$ denote the set of all maximal σ-live tuples in $\mathfrak{U}_0$. Letting $<_{lex}$ denote the lexicographic ordering of $\mathbb{N}^*$, we construct a well-founded linear ordering $\prec$ on $L_0^{\mathfrak{A}}$ as follows: $\overline{c} \prec \overline{d}$ iff there is an $i \leq \min\{|\overline{c}|, |\overline{d}|\}$ such that $c_i <_{lex} d_i$ and $c_j = d_j$ for every $j <_{lex} i$ (note that if there is no such i, then the tuples are equal due to maximality). One can show that $\overline{c} \prec \overline{d}$ implies that ($\heartsuit$): if $\overline{c}$ and $\overline{d}$ share some elements, then there exists i, j and k such that $\overline{c}_{i \ldots j} = \overline{d}_{1 \ldots k}$ and none of the elements d_ℓ, for $\ell > k$, occur in $\overline{c}$. To prove this, one needs to simply show that if d_k occurs in $\overline{c}$, then $(d_1, \ldots, d_k)$ is an affix of $\overline{c}$ (the proof goes via careful inspection of the definition of HAHs, *cf.* [5, App. C.7]).

We define $\mathsf{wit}_\mathfrak{B}$ inductively w.r.t $\prec$. Consider a maximal σ-live tuple $\overline{\mathrm{d}}$ and suppose that we have defined $\mathsf{wit}_\mathfrak{B}$ for all the σ-live tuples $\overline{\mathrm{c}} \prec \overline{\mathrm{d}}$. There are two cases to consider.

- There exists a tuple $\overline{\mathrm{c}} \prec \overline{\mathrm{d}}$ sharing at least one element with $\overline{\mathrm{d}}$. By $(\heartsuit)$, for every such tuple $\overline{\mathrm{c}}$ there are i, j and k so that $\overline{\mathrm{c}}_{i\ldots j} = \overline{\mathrm{d}}_{1\ldots k}$ and none of the elements d_ℓ, for $\ell > k$, occur in $\overline{\mathrm{c}}$. Let $\overline{\mathrm{c}}$ be the tuple for which the corresponding value k is the largest. Since $\overline{\mathrm{c}}$ is σ-live, by induction hypothesis there exists some $\overline{\mathrm{e}} \sqsubseteq B$ such that $(\mathfrak{A}, \overline{\mathrm{c}}) \sim_{\mathsf{G}_{\mathsf{affix}}[\sigma]} (\mathfrak{B}, \overline{\mathrm{e}})$ holds. Thus there exists some $\overline{\mathrm{f}} \sqsubseteq B$ such that $\overline{\mathrm{f}}_{i\ldots j} = \overline{\mathrm{e}}_{i\ldots j}$ and $(\mathfrak{A}, \overline{\mathrm{d}}) \sim_{\mathsf{G}_{\mathsf{affix}}[\sigma]} (\mathfrak{B}, \overline{\mathrm{f}})$. We now extend $\mathsf{wit}_\mathfrak{B}$ in such a way that $\mathsf{wit}_\mathfrak{B}(\overline{\mathrm{d}}) = \overline{\mathrm{f}}$.

- Otherwise $\overline{\mathrm{c}}$ and $\overline{\mathrm{d}}$ do not share any elements. Since $\overline{\mathrm{d}}$ is σ-live and $(\mathfrak{A}, \overline{\mathrm{a}}) \sim_{\mathsf{G}_{\mathsf{affix}}[\sigma]} (\mathfrak{B}, \overline{\mathrm{b}})$, there exists some $\overline{\mathrm{e}} \sqsubseteq B$ such that $(\mathfrak{A}, \overline{\mathrm{d}}) \sim_{\mathsf{G}_{\mathsf{affix}}[\sigma]} (\mathfrak{B}, \overline{\mathrm{e}})$. We then simply extend $\mathsf{wit}_\mathfrak{B}$ in such a way that $\mathsf{wit}_\mathfrak{B}(\overline{\mathrm{d}}) = \overline{\mathrm{e}}$.

The resulting mapping $\mathsf{wit}_\mathfrak{B}$. This finishes Step I.

Base case: Step II. We next complete types of all fresh (= all in this case) σ-live tuples of $\mathfrak{U}_0$. Take any maximal σ-live tuple $\overline{\mathrm{d}}$ from $\mathfrak{U}_0$ and perform the $\overline{\mathrm{d}}$-completion of $\mathfrak{U}_0$. It is easy to see that this process is conflict-free in the following sense: there is no tuple $\overline{\mathrm{c}}$ and $R \in \mathsf{sig}(\psi)$ so that we end up specifying both $\overline{\mathrm{c}} \in R^{\mathfrak{U}_0}$ and $\overline{\mathrm{c}} \notin R^{\mathfrak{U}_0}$. First, if two maximal σ-live tuples $\overline{\mathrm{d}}$ and $\overline{\mathrm{e}}$ have affixes $\overline{\mathrm{d}}_{i\ldots j}$ and $\overline{\mathrm{e}}_{k\ldots\ell}$ such that $\overline{\mathrm{d}}_{i\ldots j} = \overline{\mathrm{e}}_{k\ldots\ell}$, then we know that $\mathsf{wit}_\mathfrak{B}(\overline{\mathrm{d}}_{i\ldots j}) = \mathsf{wit}_\mathfrak{B}(\overline{\mathrm{e}}_{k\ldots\ell})$, and thus there are no conflicts in the "intersections" of σ-live tuples. Second, by construction $\mathsf{tp}_\mathfrak{A}^{\mathsf{G}_{\mathsf{affix}}[\sigma]}(\overline{\mathrm{d}}) = \mathsf{tp}_\mathfrak{B}^{\mathsf{G}_{\mathsf{affix}}[\sigma]}(\mathsf{wit}_\mathfrak{B}(\overline{\mathrm{d}}))$ holds for all maximal σ-live tuple $\overline{\mathrm{d}}$, and hence the σ-infix-types that we assigned to σ-live tuples are indeed types, *i.e.* they are consistent. Thus our process is conflict-free.

Note that our structure satisfies now conditions (a) and (b).

Base case: Step III. We finish the base case by providing witnesses for fresh $\mathsf{sig}(\psi)$-tuples via repairs. Recall that $L_0^\mathfrak{A}$ denotes the set of all maximal σ-live tuples in $\mathfrak{U}_0^\mathfrak{A}$. For each $\overline{\mathrm{d}} \in L_0^\mathfrak{A}$ we perform the $(\mathfrak{B}, \mathsf{wit}_\mathfrak{B}(\overline{\mathrm{d}}))$-repair of $\overline{\mathrm{d}}$; the resulting structure will be taken to be $\mathfrak{U}_1$. Note that now every $\mathsf{sig}(\psi)$-live tuple in $\mathfrak{U}_1$ has its witnesses for the existential requirements, but there may be new $\mathsf{sig}(\varphi)$-live tuples without them. Moreover, $\mathsf{wit}_\mathfrak{B}$ is defined for all freshly added elements, but $\mathsf{wit}_\mathfrak{A}$ is not. Furthermore, we note that $\mathfrak{U}_1$ now satisfies condition (e), since all the 1-fresh live tuples for which $\mathsf{wit}_\mathfrak{A}$ is not defined are present in the subtrees that we attached to $\mathfrak{U}_0$ during the repair, which is done only at (maximal) σ-live tuples.

Base case: Step IV. It could be the case that the structure $\mathfrak{U}_1$ produced in the previous step violates (d), due to the lack of realisation of a certain type from $\mathfrak{B}$. Thus, as an extra precaution, unique to the base case, we add a disjoint copy of $\mathfrak{B}$ to $\mathfrak{U}_1$ and define $\mathsf{wit}_\mathfrak{B}$ for it to be the identity. Note that now (d) will be satisfied in any extension of $\mathfrak{U}_1$.

Inductive step. The inductive step is analogous to Steps I-III from the base case, hence we keep its description short. Assume that $\mathfrak{U}_n$ is defined and that in the previous step of the construction we employed $\mathfrak{B}$-repairs (the case of $\mathfrak{A}$-repairs is symmetric). Given a component $\mathfrak{C}$ that was created during such a repair, we let $L_n^\mathfrak{C}$ denote the set of all maximal n-fresh σ-live tuples in $\mathfrak{C}$. Since $\mathfrak{C}$ is essentially a HAT (up to renaming), we can again define a well-founded linear order $\prec$ on $L_n^\mathfrak{C}$ in the same way as we did in the base case for $L_0^\mathfrak{A}$. As in the base case, we then define missing values of $\mathsf{wit}_\mathfrak{A}$ for elements of $\mathfrak{C}$ inductively w.r.t $\prec$.

Observe that some of the tuples in $L_n^\mathfrak{C}$ might contain a proper prefix of elements of U_{n-1}. In the case of $\mathsf{G}_{\mathsf{suf}}$ these tuples do not cause any problems to us, because suffix-types do not impose any constraints on proper prefixes. In the cases of $\mathsf{G}_{\mathsf{pre}}$ and $\mathsf{G}_{\mathsf{inf}}$ we handle these

tuples by using the fact that $\mathfrak{U}_n$ satisfies condition (e) as follows. Let $\overline{d} \in L_n^{\mathfrak{C}}$ be such a tuple and let k be the largest index so that $\overline{d}_{1\ldots k} \sqsubseteq U_{n-1}$. Using condition (e), we know that there exists a σ-live n-aged tuple $\overline{c}$ so that $\overline{c}_{i\ldots j} = \overline{d}_{1\ldots k}$, for some i and j. Employing condition (a), we know that $\overline{e} := \mathsf{wit}_{\mathfrak{A}}(\overline{c})$, $\overline{f} := \mathsf{wit}_{\mathfrak{B}}(\overline{c})$ and $\overline{h} := \mathsf{wit}_{\mathfrak{B}}(\overline{d})$ are defined and that $(\mathfrak{A}, \overline{e}) \sim_{\mathsf{G}_{\mathsf{affix}}[\sigma]} (\mathfrak{B}, \overline{f})$. Thus there exists a σ-live tuple $\overline{g} \sqsubseteq A$ such that $\overline{f}_{i\ldots j} = \overline{g}_{i\ldots j}$ and $(\mathfrak{A}, \overline{g}) \sim_{\mathsf{G}_{\mathsf{affix}}[\sigma]} (\mathfrak{B}, \overline{h})$. We now extend $\mathsf{wit}_{\mathfrak{A}}$ in such a way that $\mathsf{wit}_{\mathfrak{A}}(\overline{d}) = \overline{g}$.

The above procedure is repeated for all components $\mathfrak{C}$ that were introduced during the previous repair. Having defined $\mathsf{wit}_{\mathfrak{A}}$ for all the elements, we perform a completion that works for exactly the same reasons as described before. Finally, letting L_n denote the set of all maximal n-fresh σ-live tuples, we perform repair of every tuple in L_n, which results in a model that we select as $\mathfrak{U}_{n+1}$. We stress that every $\mathsf{sig}(\varphi)$-live tuple in $\mathfrak{U}_{n+1}$ has its witnesses for the existential requirements, but there can now be new $\mathsf{sig}(\psi)$-live tuples without them. We also emphasise that $\mathsf{wit}_{\mathfrak{A}}$ is defined for all the new elements but $\mathsf{wit}_{\mathfrak{B}}$ might not be. This concludes the inductive step and hence, also the construction of $\mathfrak{U}$ and the proof of Theorem 13.

▶ Remark 16. The presented model construction is quite generic. Indeed, the only part of the construction which is really specific to $\mathsf{G}_{\mathsf{affix}}$ is the first step of the construction, namely the part where we define inductively the values of witness functions. We expect that the presented technique can be easily adapted to other logics, especially to other fragments of the guarded fragments. For instance, we believe that our technique can be adjusted, *e.g.* to the case of the two-variable GF from [14] as well as to the uniform one-dimensional GF from [16].

5 Conclusions

In this paper kick-started a project of understanding the model theory of the family of guarded and unguarded ordered logics. We first investigated the relative expressive power of ordered logics by means of suitable bisimulations. Afterwards, we proceed with the Craig Interpolation Property (CIP) showing that (i) the fluted and the forward fragments do not enjoy CIP, (ii) while the other logics that we consider enjoy it. The fact that the fluted fragment does not posses CIP was quite unexpected in the light of already existing claims for the contrary [21, Thm. 14]. For the other logics we proposed a novel model-theoretic "complete-and-repair" method of creating a model out of two bisimilar forest-like structures.

There are several interesting future work directions.

1. One example is to investigate the Łoś-Tarski Preservation Theorem as well as other preservation theorems. While we think that we already have a working construction for guarded ordered logics, the status of ŁTPT holding for $\mathsf{L}_{\mathsf{pre}}, \mathsf{L}_{\mathsf{suf}}$, and $\mathsf{L}_{\mathsf{inf}}$ is not clear.[6]

2. Another work direction is to take a look at on effective interpolation, similarly to what has been proposed in [6] as well as on the interpolant existence problem for $\mathsf{L}_{\mathsf{inf}}$ and $\mathsf{L}_{\mathsf{suf}}$, as done in [17]. Preliminary results were obtained. It is also interesting whether the guarded ordered logics enjoy stronger versions of interpolations, *e.g.* Lyndon's interpolation or Otto's interpolation. We are quite optimistic about it.

3. What is the complexity of the model checking problem for ordered logics, if we use list encoding to encode our structures?

---

[6] Purdy provides a "proof" in [21] that $\mathsf{L}_{\mathsf{suf}}$ has ŁTPT. However, his "proof" is sketchy and lacks sufficient mathematical arguments required to verify its correctness. In the light of our discovery of yet another false claim from [21], we believe that it is safe to assume that ŁTPT for $\mathsf{L}_{\mathsf{suf}}$ *is open*.

We are also actively working on the finitary versions of van Benthem theorem for the forward guarded fragment as well as the Lindström-style characterisation theorems. This is an ongoing work of Benno Fünfstück, a master student at TU Dresden, under the supervision of B. Bednarczyk.

───── **References** ─────

1   Hajnal Andréka, István Németi, and Johan van Benthem. Modal Languages and Bounded Fragments of Predicate Logic. *J. Philos. Log.*, 1998.

2   Vince Bárány, Michael Benedikt, and Balder ten Cate. Some Model Theory of Guarded Negation. *J. Symb. Log.*, 2018.

3   Vince Bárány, Balder ten Cate, and Luc Segoufin. Guarded Negation. *J. ACM*, 2015.

4   Bartosz Bednarczyk. Exploiting Forwardness: Satisfiability and Query-Entailment in Forward Guarded Fragment. In *JELIA*, 2021.

5   Bartosz Bednarczyk and Reijo Jaakkola. Towards a model theory of ordered logics: Expressivity and interpolation (extended version), arXiV 2022. `doi:10.48550/ARXIV.2206.11751`.

6   Michael Benedikt, Balder ten Cate, and Michael Vanden Boom. Effective Interpolation and Preservation in Guarded Logics. *ACM Trans. Comput. Log.*, 2016.

7   Dietmar Berwanger and Erich Graedel. Games and Model Checking for Guarded Logics. In *LPAR 2001*, 2001.

8   Erich Grädel, Phokion G. Kolaitis, and Moshe Y. Vardi. On the decision problem for two-variable first-order logic. *Bull. Symb. Log.*, 1997.

9   Erich Grädel and Martin Otto. The Freedoms of (Guarded) Bisimulation. In Alexandru Baltag and Sonja Smets, editors, *Johan van Benthem on Logic and Information Dynamics*. Springer, 2014.

10 Erich Graedel, Phokion Kolaitis, Leonid Libkin, Maarten Marx, Joel Spencer, Moshe Y. Vardi, Yde Venema, and Scott Weinstein. *Finite Model Theory and Its Applications*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2007.

11 Lauri Hella and Antti Kuusisto. One-dimensional Fragment of First-order Logic. In *AIML 2014*, 2014.

12 Andreas Herzig. A New Decidable Fragment of First Order Logic. In *Third Logical Biennial, Summer School and Conference in Honour of S. C. Kleene*, 1990.

13 Wilfrid Hodges. *A Shorter Model Theory*. Cambridge University Press, 1997.

14 Eva Hoogland, Maarten Marx, and Martin Otto. Beth Definability for the Guarded Fragment. In *LPAR*, 1999.

15 Reijo Jaakkola. Ordered Fragments of First-Order Logic. In *MFCS*, 2021.

16 Reijo Jaakkola. Uniform Guarded Fragments. In *FOSSACS*, 2022.

17 Jean Christoph Jung and Frank Wolter. Living without Beth and Craig: Definitions and Interpolants in the Guarded and Two-Variable Fragments. In *LICS*, 2021.

18 Leonid Libkin. *Elements of Finite Model Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2004. `doi:10.1007/978-3-662-07003-1`.

19 Maarten Marx. *Algebraic relativization and arrow logic*. University of Amsterdam, 1995.

20 Ian Pratt-Hartmann, Wieslaw Szwast, and Lidia Tendera. The Fluted Fragment Revisited. *J. Symb. Log.*, 2019.

21 William C. Purdy. Complexity and Nicety of Fluted Logic. *Stud Logica*, 2002.

22 Willard Quine. *The Ways of Paradox and Other Essays, Revised Edition*. Harvard University Press, 1976.

23 Abraham Robinson. A Result on Consistency and Its Application to the Theory of Definition. *Journal of Symbolic Logic*, 1960.

24 Luc Segoufin and Balder ten Cate. Unary negation. *Log. Methods Comput. Sci.*, 9(3), 2013.

25 Thomas Sturm, Marco Voigt, and Christoph Weidenbach. Deciding First-Order Satisfiability when Universal and Existential Variables are Separated. In *LICS 2016*, 2016.

26 Johan van Benthem. Modal Foundations for Predicate Logic. *Log. J. IGPL*, 1997.