Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning
Main Track

# Two-Variable Logic for Hierarchically Partitioned and Ordered Data

**Oskar Fiuk**[1] , **Emanuel Kieroński**[1] , **Vincent Michielini**[2]
[1]Institute of Computer Science, University of Wrocław,
[2]Faculty of Mathematics, Informatics, and Mechanics, Warsaw University
307023@uwr.edu.pl, emanuel.kieronski@cs.uni.wroc.pl, michielini@mimuw.edu.pl

**Abstract**

We study Two-Variable First-Order Logic, FO$^2$, under semantic constraints that model hierarchically structured data. Our first logic extends FO$^2$ with a linear order $<$ and a chain of increasingly coarser equivalence relations $E_1 \subseteq E_2 \subseteq \ldots$. We show that its finite satisfiability problem is NEXPTIME-complete. We also demonstrate that a weaker variant of this logic without the linear order enjoys the exponential model property. Our second logic extends FO$^2$ with a chain of nested total preorders $\preceq_1 \subseteq \preceq_2 \subseteq \ldots$. We prove that its finite satisfiability problem is also NEXPTIME-complete. However, we show that the complexity increases to EXPSPACE-complete once access to the successor relations of the preorders is allowed. Our last result is the undecidability of FO$^2$ with two independent chains of nested equivalence relations.

## 1 Introduction

Hierarchically partitioned data are pervasive in modern computer systems. For example, Geographical Information Services often organise geospatial information using progressively more detailed fields: country, region, state, and city. Similar hierarchies appear in numerous contexts: Data Storage (organised into folders, subfolders, and files), Network Management (addressing schemes like IPv4/IPv6 with subnet hierarchies), Dependency Maintenance (tools for tracking dependencies between modules, libraries, and services).

To model such hierarchical data, we consider domains in which elements are annotated with data values drawn from potentially infinite or very large domains. A key aspect is that these data values can be tested for equality at multiple levels of precision. This is naturally captured by a family of increasingly coarser equivalence relations: two elements are related by the $k$-th equivalence relation if the $k$-th level equality test holds between their associated data values.

In this work, we establish results on the decidability and complexity of satisfiability problems for several variants of the Two-Variable Fragment of First-Order Logic, FO$^2$, extended to support such increasingly coarser equivalence relations. Our goal is to provide a logical framework that is expressive enough to model complex multi-level relationships while retaining desirable computational properties for reasoning. As most of the investigated logics do not enjoy the finite model property, their general and finite satisfiability problems differ. Our primary focus is on finite satisfiability, while the case of general satisfiability is left for future work.

The motivation for studying FO$^2$ stems from its good algorithmic and model-theoretic properties. FO$^2$ combines an NEXPTIME-complete satisfiability problem and the exponential model property (Grädel, Kolaitis, and Vardi 1997) with a reasonable expressive power. In particular, it embeds (via the so-called *standard translation*) many modal, temporal, and description logics (up to $\mathcal{ALCIOH}^{\cap,\neq}$). Also, it is the maximal, in terms of the number of variables, fragment of First-Order Logic with decidable satisfiability problem, as already the Three-Variable Fragment is undecidable (Kahr, Moore, and Wang 1962). In the last few decades, FO$^2$ together with its variations have been extensively studied, and plenty of results have been obtained in various scenarios (cf. Subsection 1.2 on related work). All of this makes FO$^2$ often the first-choice option for various reasoning tasks.

In the following subsections of this introduction, we define our logics and present the obtained results (Subsection 1.1), discuss related work (Subsection 1.2), compare the expressive power of our logics (Subsection 1.3), and outline the technical sections that follow (Subsection 1.4).

### 1.1 Logics of Interest and Our Results

Our underlying formalism is the Two-Variable Fragment of First-Order Logic, FO$^2$, whose formulas may use only the variables $x$ and $y$; any number of unary and binary *common* relation symbols (i.e., with unconstrained interpretations); the equality symbol; and constant symbols, but no function symbols of positive arity. Extensions of FO$^2$ are denoted by listing the *special* symbols (i.e., those with constrained interpretations) in brackets, e.g., FO$^2[<, \mathcal{EQ}^{\subseteq}]$. This notation makes it explicit which additional semantic constraints are imposed on top of the base FO$^2$ syntax.

**Order on domain elements.** Let $\mathcal{EQ}^{\subseteq}$ denote the family of special symbols $E_1, E_2, \ldots$ whose interpretations are constrained to *nested* equivalence relations, that is for every $k \in \mathbb{N}$: (i) the interpretation of $E_k$ is an equivalence relation, and (ii) the interpretation of $E_{k+1}$ is *coarser* than that of $E_k$ (i.e., $xE_ky \rightarrow xE_{k+1}y$). Let $<$ be a special symbol interpreted as a *strict* linear order on domain elements.[1]

[1]Throughout the paper, we will also use the derived non-strict

Our first considered logic is $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$, supporting both a linear order on elements and hierarchical data values.

Potential applications of $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$ can be found in, e.g., temporal verification of multiprocess systems. We provide a motivational example of enforcing isolation policy in an environment of processes, containers, and events.

System execution is modeled as a linear sequence of events ordered by time using the relation $<$. Events are generated by processes, and processes are grouped into containers. The equivalence relations $E_1$ and $E_2$ capture this hierarchy: $xE_1y$ holds when events $x$ and $y$ come from the same process, and $xE_2y$ when they come from possibly distinct processes yet running in the same container. These are naturally nested: $xE_1y \to xE_2y$.

Importantly, the domain of our model consists only of events: processes are implicitly represented as equivalence classes of the relation $E_1$. That is, all events belonging to the same process form a single equivalence class of $E_1$. (We assume that every process reports at least one event, e.g., *spawn event*, to ensure that it has a non-empty class.) Likewise, containers are implicitly represented as equivalence classes of the relation $E_2$. (We assume that in every container at least one process is running, e.g., *root process*.)

Process-level properties (e.g., *sandboxed*, *privileged*) are expressed via class-wise unary predicates. For example:

$$\forall x, y.\ xE_1y \to \big(\mathrm{sandboxed}(x) \leftrightarrow \mathrm{sandboxed}(y)\big)$$

This ensures that all events from the same process share the same *sandboxed* or *not sandboxed* label, even though the logic only quantifies over events.

Using the above described encoding, we can impose the following isolation policy: *"A sandboxed process must not communicate with events outside its container unless an explicit grant was made before from a privileged process."*

This policy can be formalised with a sentence:

$$\forall x_1.\ \big(\mathrm{sandboxed}(x_1) \wedge \varphi_{\text{cross-container-message}}(x_1)\big) \to \big(\exists y_1.\ y_1 < x_1 \wedge x_1E_1y_1 \wedge \varphi_{\text{permission-grant}}(y_1)\big),$$

where $\varphi_{\text{cross-container-message}}(x_1)$ stands for the formula

$$\exists y_2.\ x_1 < y_2 \wedge \neg x_1E_2y_2 \wedge \mathrm{message}(x_1, y_2),$$

and $\varphi_{\text{permission-grant}}(y_1)$ for

$$\exists x_3.\ \mathrm{privileged}(x_3) \wedge x_3 < y_1 \wedge \mathrm{grant}(x_3, y_1).$$

In the above formulas, we annotated variables as $x_1$, $y_1$, $y_2$, $x_3$ for readability to reflect their roles in different processes. However, only two variables $x$ and $y$, with reuse across quantifiers, are sufficient to express this property.

As stated above, all variables refer to events: $x_1$ is a *dispatch message* event from Process 1 (the *sender*) that initiates the communication, and $y_2$ is a *deliver message* event from Process 2 (the *receiver*) that receives the message from Process 1, as indicated by $\mathrm{message}(x_1, y_2)$. The event $y_1$ is an earlier event from the same process as $x_1$, representing the *grant acknowledge* event, marking the point at which Process 1 becomes aware of and is authorised to act on the granted permission. Finally, $x_3$ is a *grant authorise* event from Process 3 (a *privileged admin*) that issues the permission and notifies Process 1, as indicated by $\mathrm{grant}(x_3, y_1)$.

The intended temporal order of events $x_3{<}y_1{<}x_1{<}y_2$ is enforced: the permission is issued before it is acknowledged ($x_3{<}y_1$); the message is sent after the permission being acknowledged ($y_1{<}x_1$) and before being delivered ($x_1{<}y_2$).

Notice also that, to detect that Process 1 and Process 2 are running in distinct containers, we refer to their representative events using $\neg xE_2y$. Since $E_1$ is nested within $E_2$, it follows that if two events are not $E_2$-related, then their respective $E_1$-classes (i.e., processes) must also belong to distinct $E_2$-classes (i.e., containers).

Our first main contribution is establishing the complexity of the finite satisfiability problem for $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$:

**Theorem 1.** *Finitely satisfiable sentences of* $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$ *admit models of exponential size. The finite satisfiability problem for* $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$ *is* NEXPTIME*-complete.*

Notice that $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$ does not include the induced successor of $<$.[2] (Björklund and Bojańczyk 2007) show that adding it leads to undecidability (see: Subsection 1.2).

**Order on data values.** We consider now a different way to incorporate a linear order into our scenario. We trade an order on domain elements for a family of nested linear orders on data values, i.e., on equivalence classes of $E_1, E_2, \ldots$.

Let $\preceq^{\subseteq}$ denote the family of special symbols $\preceq_1, \preceq_2, \ldots$ whose interpretations are constrained to *nested* total preorders, that is for each $k \in \mathbb{N}$: (i) the interpretation of $\preceq_k$ is a total preorder[3], and (ii) the interpretation of $\preceq_k$ is a subrelation of the interpretation of $\preceq_{k+1}$ (i.e. $x \preceq_k y \to x \preceq_{k+1} y$).

For example, interpretations over $\mathbb{N}$ defined by $n \preceq_k m$ iff $\lfloor \frac{n}{10^k} \rfloor \le \lfloor \frac{m}{10^k} \rfloor$ satisfy the above requirements.

Our second logic is $\mathrm{FO}^2[\preceq^{\subseteq}]$, supporting hierarchical data values that can be compared at multiple levels of granularity using less-than, equal, and greater-than comparison tests.

In $\mathrm{FO}^2[\preceq^{\subseteq}]$, we naturally keep nested equivalence relations $E_1, E_2, \ldots$ Their interpretation is now given by the $\preceq_k$-equivalent elements: $xE_ky \leftrightarrow x \preceq_k y \wedge y \preceq_k x$.

Natural applications of $\mathrm{FO}^2[\preceq^{\subseteq}]$ arise in temporal reasoning tasks that involve increasingly fine-grained notions of time. A representative case study is a system supporting atomic transactions. Here, the elements of the structure represent low-level operations. The finest preorder $\preceq_1$ models the underlying timeline, where $\preceq_1$-equivalent elements are treated as occurring in parallel. Transactions are modeled implicitly as $E_2$-equivalence classes, ordered chronologically by $\preceq_2$. Coarser preorders can be used to represent higher-level structures such as $\preceq_3$ for commit order, $\preceq_4$ for versioning, and so on. With this encoding, we can express that, e.g., "*raised exceptions must be handled in the future, yet within the same transaction*" as follows:

$$\forall x.\ \mathrm{exception}(x) \to \exists y.\ x \prec_1 y \wedge xE_2y \wedge \mathrm{handles}(y, x)$$

For $\mathrm{FO}^2[\preceq^{\subseteq}]$, we establish the following theorem:

linear order $\le$, and their analogues for other orders, e.g., $\prec$ and $\preceq$. Notice that in $\mathrm{FO}^2$ $\le$ is definable from $<$ and vice-versa.

[2] Notice that the induced successor is definable in FO, yet not in $\mathrm{FO}^2$, by the formula $\varphi(x, y) := x < y \wedge \forall z.\ (x < z \to y \le z)$.

[3] A *total preorder* is a transitive relation $\preceq$ such that, for every $x, y$, either $x \preceq y$ or $y \preceq x$ holds (in particular, $x \preceq x$ holds).

**Theorem 2.** *Finitely satisfiable sentences of* $\mathrm{FO}^2[\preceq^{\subseteq}]$ *admit models of exponential size. The finite satisfiability problem for* $\mathrm{FO}^2[\preceq^{\subseteq}]$ *is* NEXPTIME*-complete.*

**Adding successors on data values.** We consider now a logic $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$ that enriches the syntax of $\mathrm{FO}^2[\preceq^{\subseteq}]$ by adding, for each $\preceq_k$-symbol, its induced successor predicate $\mathcal{S}_k$, defined as $\mathcal{S}_k(x,y) := x \prec_k y \wedge \forall z.\ \big(x \prec_k z \rightarrow y \preceq_k z\big)$.

Applications of $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$ naturally extend those of $\mathrm{FO}^2[\preceq^{\subseteq}]$; we continue with the transaction-based system scenario. For example, we can express the property: *"If a transaction fails, then in the immediate* $\preceq_1$*-successor time-slot a rollback must occur, using the last available snapshot from the immediate* $\preceq_2$*-predecessor transaction."* Formally:

$$\forall x.\ \mathrm{fail}(x) \rightarrow \big(\exists y.\ \mathrm{rollback}(y) \wedge \mathcal{S}_1(x,y) \wedge xE_2y \wedge$$
$$\exists x.\ \text{last-snapshot}(x) \wedge \mathcal{S}_2(x,y) \wedge \text{restore-to}(y,x)\big)$$

Here, $\mathcal{S}_1(x,y) \wedge xE_2y$ ensures that the rollback follows the failure in the very next $\preceq_1$-time-slot within the same transaction, while $\mathcal{S}_2(x,y)$ ensures that the snapshot belongs to the transaction immediately preceding the one containing the rollback. In particular, this guarantees that the snapshot was created before the failure. Additional axioms can be imposed to refine the scenario further: for instance, we can define last-snapshot as the $\preceq_1$-maximum snapshot within each $E_2$-class, or require that after a rollback the system proceeds directly to the $\preceq_2$-successor transaction.

Now we state the second main contribution of this paper:

**Theorem 3.** *Finitely satisfiable sentences of* $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$ *admit models of doubly exponential size. The finite satisfiability problem for* $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$ *is* EXPSPACE*-complete.*

**Absence of orders.** Both $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$ and $\mathrm{FO}^2[\preceq^{\subseteq}]$ admit sentences that enforce infinite models (e.g., $\forall x.\ \exists y.\ x{<}y$ and $\forall x.\ \exists y.\ x \prec_1 y$). This contrasts with pure $\mathrm{FO}^2$ which enjoys the *finite model property* (i.e., every satisfiable sentence has a finite model). A natural question arises: Can nested equivalence relations alone enforce infinite models, or is this phenomenon solely due to the presence of linear orders?

Let $\mathrm{FO}^2[\mathcal{EQ}^{\subseteq}]$ denote $\mathrm{FO}^2$ extended with the family of nested equivalence relations $E_1, E_2, \ldots$. We answer that $\mathrm{FO}^2[\mathcal{EQ}^{\subseteq}]$ does indeed enjoy the finite model property:

**Theorem 4.** *Satisfiable sentences of* $\mathrm{FO}^2[\mathcal{EQ}^{\subseteq}]$ *admit models of exponential size. The satisfiability and finite satisfiability problems for* $\mathrm{FO}^2[\mathcal{EQ}^{\subseteq}]$ *coincide and are* NEXPTIME*-complete.*

**Undecidability.** A natural candidate to explore next is a logic supporting two independent families of nested equivalence relations $E_1, E_2, \ldots$ and $F_1, F_2, \ldots$. We prove that reasoning in such a logic is undecidable—even in a very restricted setting—when each family has length 2 and the vocabulary, except the four special equivalence symbols, is composed of unary predicates only.

**Theorem 5.** *The satisfiability and finite satisfiability problems are undecidable for the constant-free, equality-free, monadic fragment of* $\mathrm{FO}^2$ *extended with four special symbols* $E_1, E_2, F_1, F_2$*, interpreted as equivalence relations such that* $E_2$ *is coarser than* $E_1$ *and* $F_2$ *is coarser than* $F_1$*.*

## 1.2 Related Work

$\mathrm{FO}^2$ **with equivalence relations.** $\mathrm{FO}^2$ with a single equivalence relation has the exponential model property and an NEXPTIME-complete satisfiability problem (Kieroński and Otto 2012). With two equivalence relations (not necessarily nested), the finite model property is lost; both satisfiability and finite satisfiability are 2-NEXPTIME-complete (Kieroński et al. 2014). With three equivalence relations, $\mathrm{FO}^2$ is undecidable (Kieroński and Otto 2012).

Since $\mathrm{FO}^2$ can express containment between equivalence relations, the work of (Kieroński et al. 2014) established decidability for $\mathrm{FO}^2$ over hierarchical partitions of depth two (with no ordering). However, it does not yield optimal complexity bounds even in this restricted setting and leave the question of the finite model property unanswered.

$\mathrm{FO}^2$ **with linear orders.** The satisfiability and finite satisfiability problems for $\mathrm{FO}^2$ with a single linear order are NEXPTIME-complete (Otto 2001). With three linear orders, both problems become undecidable (Kieroński 2011). The case of two linear orders was studied in (Zeume and Harwath 2016), where finite satisfiability is shown to be decidable in 2-NEXPTIME when one linear order is accessible via both the order and successor predicates, and the other only via one of them. General satisfiability was not addressed.

**Data Words.** $\mathrm{FO}^2$ with a linear order and an equivalence relation has been studied extensively in the context of *data words*. These are structures interpreting unary predicates, an equivalence relation, a linear order, and its induced successor relation. The satisfiability problem for $\mathrm{FO}^2$ on data words is decidable but non-elementary-hard (Bojańczyk et al. 2011). When the successor relation is omitted, the problem becomes NEXPTIME-complete.

Quite close to our setting are data words with nested equivalences, that where considered in (Björklund and Bojańczyk 2007). It is shown there that satisfiability is decidable when the linear order is accessible only via its successor relation, and becomes undecidable as soon as both the order and its successor are accessible. The variant closest to $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$, where only the linear order (and not its successor) is available, was not explored in that work. Notice that data words restrict the common part of the signature to unary symbols, whereas we allow full $\mathrm{FO}^2$ with arbitrary binary relations that may freely interact with the equivalences.

$\mathrm{FO}^2$ **with total preorders.** The EXPSPACE-completeness of the finite satisfiability problem for $\mathrm{FO}^2$ over structures with one total preorder, its induced successor relation, a linear order, and additional unary relations was established in (Schwentick and Zeume 2012). As in the case of data words, no common binary relations are permitted. When two independent total preorders are available, satisfiability becomes undecidable. To the best of our knowledge, the case of a single total preorder combined with arbitrary binary relations, as well as settings with nested total preorders, have not been investigated so far.

$\mathrm{FO}^2$ **over trees.** An alternative perspective on nested equivalence relations is to interpret them as trees. If only $k$ nested equivalence symbols $E_1, \ldots, E_k$ are considered, structures

can be viewed as modeling the leaves of unranked trees of fixed depth $k$, where $E_i(a, b)$ holds if $a$ and $b$ share a common ancestor at depth $k-i$. Importantly, only leaf nodes constitute the domain; internal tree nodes serve an auxiliary role and are not part of the universe. This interpretation is different from standard two-variable logics over trees, cf. (Bojańczyk et al. 2009; Charatonik and Witkowski 2016; Benaim et al. 2016), where all tree nodes belong to the domain and the structure is accessed via navigational predicates such as *parent*, *child*, *descendant*, etc.

**Other logics.** Nested equivalence relations can be simulated in description logics such as $\mathcal{SHI}$ and $\mathcal{SHOI}$, which include transitive roles ($\mathcal{S}$), inverse roles ($\mathcal{I}$), role hierarchies ($\mathcal{H}$), and, possibly, nominals ($\mathcal{O}$). These logics have EXPTIME-complete satisfiability problems. However, interactions between binary relations (roles) are limited to role hierarchies, meaning one can only express containment between relations (these can be equivalences or common relations). Moreover, these logics do not include linear orderings.

Among first-order fragments, the Unary Negation Fragment, UNFO (ten Cate and Segoufin 2013), is particularly worth to mention. This logic restricts negation to subformulas with at most one free variable. Its extension capturing $\mathcal{SHOI}$, denoted UNFO+$\mathcal{SOH}$, is decidable and 2-EXPTIME-complete (Danielski and Kieroński 2019). Remarkably, UNFO+$\mathcal{SOH}$ enables reasoning over arbitrarily many independent families of nested equivalence relations.

### 1.3 Expressivity of $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$ and $\mathrm{FO}^2[\preceq^{\subseteq}]$

Naturally, $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$ is more expressive than $\mathrm{FO}^2[\preceq^{\subseteq}]$, yet both formalisms are incomparable with $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$. In $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$ and $\mathrm{FO}^2[\preceq^{\subseteq}]$, every $E_k$-class necessarily forms an interval with respect to a linear order induced from the total preorder $\preceq_1$ by resolving ties arbitrarily. In contrast, this interval property cannot be expressed in $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$. On the other hand, $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$ can enforce that certain classes do not form intervals:

$$\forall x.\, \exists y.\, x < y \wedge \neg x E_1 y \wedge \big(P(x) \leftrightarrow \neg P(y)\big)$$
$$\forall x, y.\, \big(P(x) \wedge P(y)\big) \rightarrow x E_1 y$$

### 1.4 Outline of Technical Sections

Section 2 introduces the necessary notions and definitions. Section 3 proves Theorem 1. Section 4 sketches the proofs of Theorems 2 and 4, as these are similar to that of Theorem 1; more details are in the full version. Theorem 3 is established in Section 5, with certain technical details deferred to the full version. Section 6 presents the undecidability result of Theorem 5. Finally, Section 7 concludes the paper.

## 2 Preliminaries

### 2.1 Notation and Conventions

We denote the set of natural numbers including 0 by $\mathbb{N}$. For $k \in \mathbb{N}$, the notation $[k]$ stands for the set $\{1, \ldots, k\}$, with the convention that $[0] = \varnothing$. More generally, we use interval notation $[a, b] \subseteq \mathbb{N}$ to denote the set $\{a, a+1, \ldots, b\}$ whenever $a \leq b$, and the empty set $\varnothing$ whenever $a > b$. If $E$ is an equivalence relation on a set $A$, then $E[a]$ denotes the equivalence class of an element $a \in A$, and $B/E$ denotes the quotient set via $E$ of a subset $B \subseteq A$.

A *signature* $\sigma$ is a finite set of symbols, partitioned as $\sigma = \mathrm{Cons} \cup \mathrm{Rels}$, where Cons is the set of constant symbols and Rels is the set of relation symbols (including special symbols such as $<$, $E_1$, $E_2$, etc.). Every relation symbol has associated arity. We do not allow function symbols of positive arity. The *signature of a formula* is the finite set of relation and constant symbols that appear in the formula.

The *size* (or *length*) of a formula $\varphi$, denoted $|\varphi|$, is defined as the total number of symbols it contains, where each occurrence of a symbol—be it a variable, relation symbol, or constant—contributes 1 to the count.

We use Fraktur letters such as $\mathfrak{A}, \mathfrak{B}, \ldots$ to denote structures, and the corresponding Roman letters $A, B, \ldots$ for their domains. A structure $\mathfrak{A}$ over a signature $\sigma$ interprets the symbols in $\sigma$: a relation symbol $R$ as a relation $R^{\mathfrak{A}} \subseteq A^k$ with $k$ denoting the arity of $R$; and a constant symbol $c$ as an element $c^{\mathfrak{A}} \in A$. If $B \subseteq A$, we write $\mathfrak{A} \restriction B$ for the *restriction* of $\mathfrak{A}$ to the subdomain $B$. The *size* of a structure is the cardinality of its domain. Elements of structures are typically denoted by $a, b, \ldots$; variables by $x, y$, possibly with decorations. We write $\varphi(\bar{x})$ to indicate that all free variables of the formula $\varphi$ are contained in the tuple $\bar{x}$.

An (atomic) 1-*type* over a signature $\sigma$ is a maximal consistent set of literals involving only the variable $x$. Similarly, an (atomic) 2-*type* over $\sigma$ is a maximal consistent set of literals over the variables $x$ and $y$; in particular, a 2-type naturally determines the 1-types of both $x$ and $y$. We use the symbol $\alpha$ (possibly with decorations) to denote 1-types.

Let $\mathfrak{A}$ be a structure. For any $a \in A$, we write $\mathrm{tp}^{\mathfrak{A}}[a]$ to denote the unique 1-type *realised* by $a$ in $\mathfrak{A}$, that is the 1-type $\alpha$ such that $\mathfrak{A} \models \alpha(a)$. Similarly, for distinct elements $a, b \in A$, the notation $\mathrm{tp}^{\mathfrak{A}}[a, b]$ denotes the unique 2-type *realised* in $\mathfrak{A}$ by the pair $(a, b)$.

### 2.2 Scott Normal Form

A sentence $\varphi$ is in *Scott Normal Form* (Scott 1962) if it is in the shape of $\forall x, y.\, \psi_0(x, y) \wedge \bigwedge_{m=1}^{M} \forall x.\, \exists y.\, \psi_m(x, y)$, where $M$ denotes the number of Skolem conjuncts. For $m \in [0, M]$, the formula $\psi_m(x, y)$ is quantifier-free and constant-free, and uses only relation symbols of arity 1 and 2.

There is a standard polynomial-time procedure transforming any $\mathrm{FO}^2$ sentence into Scott Normal Form, preserving satisfiability over the same domains; see e.g., (Pratt-Hartmann 2023). It is readily verified that it is sound for the extensions of $\mathrm{FO}^2$ introduced in Subsection 1.1.

We standardise the special relation symbols in $\varphi$ by assuming they are the consecutive symbols $E_1, \ldots, E_K$ for some $K \in \mathbb{N}$. We do the same in the case of $\preceq_k$-symbols.

For convenience, we interpret additional symbol $E_0$ as the identity relation, and assume that $E_{K+1}$, the first symbol not used in $\varphi$, is interpreted as the universal relation. Note that this preserves nestedness: $E_0 \subseteq E_1$ and $E_K \subseteq E_{K+1}$.

## 3 Proof of Theorem 1

In this section, we establish our first main result.

**Theorem 1.** *Finitely satisfiable sentences of* $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$ *admit models of exponential size. The finite satisfiability problem for* $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$ *is* NEXPTIME*-complete.*

For this whole section, we fix a sentence $\varphi$ from the logic $\mathrm{FO}^2[<, \mathcal{EQ}^{\subseteq}]$. Assuming that $\varphi$ is finitely satisfiable, we aim to prove that it is actually satisfied by a model of exponential size. As discussed in Subsection 2.2, we may assume that $\varphi$ is in Scott Normal Form, i.e., in the shape of $\forall x, y.\, \psi_0(x,y) \wedge \bigwedge_{m=1}^{M} \forall x.\, \exists y.\, \psi_m(x,y)$, where $M$ denotes the number of Skolem conjuncts. By convention, $\varphi$ uses the first $K$ nested equivalence symbols $E_1, \ldots, E_K$; in addtion, we interpret $E_0$ and $E_{K+1}$ as the identity and universal relations, respectively. Let $\boldsymbol{\alpha}$ denote the set of all 1-types over the signature of $\varphi$.

The main ingredient of our proof is the following lemma, which allows us to replace a single equivalence class with its small counterpart, without touching the rest of the structure.

**Lemma 6.** *Let* $\mathfrak{A}$ *be a finite model of* $\varphi$*. Fix* $k \in [0, K]$*, and let* $\mathcal{C}$ *be an* $E_{k+1}$*-class in* $A/E_{k+1}^{\mathfrak{A}}$*. Then there exists a subset* $\mathcal{D} \subseteq \mathcal{C}$ *such that:*

1. $\mathcal{D}$ *is the union of at most* $12 \cdot M^3 \cdot |\boldsymbol{\alpha}|$ *many* $E_k$*-classes from* $\mathcal{C}/E_k^{\mathfrak{A}}$*.*
2. $\varphi$ *has a model* $\mathfrak{B}$ *over the domain* $B = (A \setminus \mathcal{C}) \cup \mathcal{D}$*.*

We will prove Lemma 6 in a moment. Let us first demonstrate how to derive Theorem 1 from it.

*Proof of Theorem 1.* Assume that $\varphi$ is finitely satisfiable, and consider a finite model $\mathfrak{A}$ of minimal size.

Because of this minimality, for each $k \in [0, K]$, every $E_{k+1}$-class of $\mathfrak{A}$ is partitioned by at most $12 \cdot M^3 \cdot |\boldsymbol{\alpha}|$ many $E_k$-classes. As otherwise, Lemma 6 would give us a model of strictly smaller size. Since each $E_0$-class of $\mathfrak{A}$ is of size 1, an immediate induction tells us that, for each $k \in [K+1]$, the $E_k$-classes of $\mathfrak{A}$ are of size at most $12^k \cdot M^{3k} \cdot |\boldsymbol{\alpha}|^k$.

Yet, $E_{K+1}^{\mathfrak{A}}$ having only one equivalence class, the size of $\mathfrak{A}$ is therefore bounded by $12^{K+1} \cdot M^{3K+3} \cdot |\boldsymbol{\alpha}|^{K+1}$. Clearly, $K$ and $M$ are both $\mathcal{O}(|\varphi|)$. As $\varphi$ does not use constants, the number of 1-types in $\boldsymbol{\alpha}$ is $2^{|\mathrm{Rels}|} = 2^{\mathcal{O}(|\varphi|)}$.

We conclude that the size of $\mathfrak{A}$ is indeed $2^{\mathcal{O}(|\varphi|^2)}$. □

We now prove Lemma 6.

Suppose that $\mathfrak{A}$ is a finite model of $\varphi$, $k$ is a number between 0 and $K$, and $\mathcal{C}$ is any $E_{k+1}$-class of $A/E_{k+1}^{\mathfrak{A}}$.

We assume w.l.o.g. that the domain $A$ is a finite subset of $\mathbb{N}$ and that the order $<^{\mathfrak{A}}$ of $\mathfrak{A}$ coincides with the usual order on $\mathbb{N}$, allowing us to write just $<$ instead of $<^{\mathfrak{A}}$.

We use the following terminology: if $m \in [M]$ and $a, b$ are elements of $A$, then we call $b$ a *witness* of $a$ whenever $\mathfrak{A} \models \psi_m(a,b)$. In addition, $b$ is an *internal* witness if $(a,b) \in E_k^{\mathfrak{A}}$; an *external* witness if $b \in A \setminus \mathcal{C}$ and $(a,b) \notin E_k^{\mathfrak{A}}$; and a $\mathcal{C}$*-witness* if $b \in \mathcal{C}$ and $(a,b) \notin E_k^{\mathfrak{A}}$.

We fix *witness functions* $\mathfrak{f}_m \colon A \to A$, i.e., such that for each $m \in [M]$ and every $a \in A$, $\mathfrak{A} \models \psi_m(a, \mathfrak{f}_m(a))$.

Before delving into the details, let us outline our proof strategy. We construct the set $\mathcal{D} \subseteq \mathcal{C}$, to replace the selected $E_{k+1}$-class $\mathcal{C}$, in three steps: First, we define a subset $W_{\mathrm{I}} \subseteq \mathcal{C}$ that will serve as witnesses for elements outside $\mathcal{C}$. Next, we set $W_{\mathrm{II}} := \bigcup_{m \in [M]} \mathfrak{f}_m[W_{\mathrm{I}}]$, which will provide witnesses for elements in $W_{\mathrm{I}}$. Similarly, we define $W_{\mathrm{III}} := \bigcup_{m \in [M]} \mathfrak{f}_m[W_{\mathrm{II}}]$, to serve as witnesses for elements in $W_{\mathrm{II}}$. It will be enough to define $\mathcal{D}$ as the $E_k$-closure of $(W_{\mathrm{I}} \cup W_{\mathrm{II}} \cup W_{\mathrm{III}}) \cap \mathcal{C}$. Finally, the structure $\mathfrak{B}$ will be defined as $\mathfrak{A} \restriction [(A \setminus \mathcal{C}) \cup \mathcal{D}]$ but with certain 2-types redefined.

In order to replace the original $\mathcal{C}$-witnesses by other elements, we use the following notion of *configurations*.

Let $\ell \in [M]$ and $a \in A$. Assume that $\langle b_1, \ldots, b_\ell \rangle$ and $\langle c_1, \ldots, c_\ell \rangle$ are tuples of pairwise distinct elements of $A$. We say that they realise the same $a$*-configuration* if the following conditions hold:

(i) For every $i \in [\ell]$, $\mathrm{tp}^{\mathfrak{A}}[b_i] = \mathrm{tp}^{\mathfrak{A}}[c_i]$.

(ii) For every $i \in [\ell]$, $a \bowtie b_i$ if and only if $a \bowtie c_i$, where $\bowtie$ is any of the symbols "$<$", "$=$" or "$>$".

(iii) For every $i \in [\ell]$ and every $j \in [K]$, $(a, b_i) \in E_j^{\mathfrak{A}}$ if and only if $(a, c_i) \in E_j^{\mathfrak{A}}$.

We define now the set $W_{\mathrm{I}} \subseteq \mathcal{C}$. For this, we introduce the notion of the $r$*-extremal* subset of a finite subset $S \subseteq \mathbb{N}$, of $\ell$ elements $a_1 < a_2 < \ldots < a_\ell$: if $\ell \leq 2 \cdot r$, then this subset, denoted by $\mathrm{extr}_r(S)$, is $S$ itself; if $\ell > 2 \cdot r$, then it is $\{a_1, \ldots, a_r, a_{\ell - r + 1}, \ldots, a_\ell\} \subset S$.

For each 1-type $\alpha$, we take the set $V_\alpha$ of the $M$ smallest and $M$ largest realisations of $\alpha$ from every $E_k$-class in $\mathcal{C}$:

$$V_\alpha := \bigcup_{\mathcal{E} \in \mathcal{C}/E_k^{\mathfrak{A}}} \mathrm{extr}_M(\{a \in \mathcal{E} \mid \mathrm{tp}^{\mathfrak{A}}[a] = \alpha\}).$$

We then define $W_{\mathrm{I}}$ as the set of $2M$ smallest and $2M$ largest elements of each $V_\alpha$:

$$W_{\mathrm{I}} := \bigcup_{\alpha \in \boldsymbol{\alpha}} \mathrm{extr}_{2 \cdot M}(V_\alpha).$$

In the following claim, we justify that any configuration of $\mathcal{C}$-witnesses can be realised within $W_{\mathrm{I}}$.

**Claim 7.** *Let* $\ell \in [M]$*, and let* $a \in A \setminus W_{\mathrm{I}}$*. Suppose that* $b_1, \ldots, b_\ell$ *are distinct elements of* $\mathcal{C} \setminus E_k^{\mathfrak{A}}[a]$*. Then there exist distinct elements* $c_1, \ldots, c_\ell$ *in* $W_{\mathrm{I}}$ *such that* $\langle b_1, \ldots, b_\ell \rangle$ *and* $\langle c_1, \ldots, c_\ell \rangle$ *realise the same* $a$*-configuration.*

*Proof.* We can partition the $b_i$'s into several subtuples by the 1-types that these elements realise. Considering these tuples independently, we obtain several tuples of the $c_i$'s to be glued together into a final tuple. Hence, w.l.o.g. assume that all the $b_i$'s realise the same 1-type $\alpha$, for some $\alpha \in \boldsymbol{\alpha}$.

Similarly, we can moreover assume that the $b_i$'s are related to $a$ in the same way, say $b_i < a$ holds for every $i \in [\ell]$.

Consider the set $S := \{c \in \mathcal{C} \mid c < a \text{ and } \mathrm{tp}^{\mathfrak{A}}[c] = \alpha\}$. Notice that $\{b_1, \ldots, b_\ell\} \subseteq S \setminus E_k^{\mathfrak{A}}[a]$. Moreover, any $\ell$-tuple of distinct element from $S \setminus E_k^{\mathfrak{A}}[a]$ realises the same $a$-configuration as the $b_i$'s. Hence, select $\ell$ minimal elements $c_1, \ldots, c_\ell$ from the set $S \setminus E_k^{\mathfrak{A}}[a]$. Because $\ell \leq M$, we have that $\{c_1, \ldots, c_\ell\} \subseteq V_\alpha$. Now, crucially among the $2 \cdot M$ minimal elements of the set $V_\alpha$, in addition to the elements $c_1, \ldots, c_\ell$, there are at most $M$ additional "*bad*" elements from the set $S \cap E_k^{\mathfrak{A}}[a]$. In consequence, we have that $\{c_1, \ldots, c_\ell\} \subseteq \mathrm{extr}_{2 \cdot M}(V_\alpha)$, which is a subset of $W_{\mathrm{I}}$. □

Naturally, elements of $W_{\mathrm{I}}$ require witnesses themselves, which is why, as announced above, we define $W_{\mathrm{II}} := \bigcup_{m\in[M]} \mathfrak{f}_m[W_{\mathrm{I}}]$. Similarly, elements of $W_{\mathrm{II}}$ need witnesses, so we define $W_{\mathrm{III}} := \bigcup_{m\in[M]} \mathfrak{f}_m[W_{\mathrm{II}}]$.

We do not introduce further sets $W_{\mathrm{IV}}, W_{\mathrm{V}}, \ldots$, as elements of $W_{\mathrm{III}} \setminus (W_{\mathrm{I}} \cup W_{\mathrm{II}})$ can be assigned $\mathcal{C}$-witnesses again in $W_{\mathrm{I}}$. This step resembles the construction of the exponential model property for FO$^2$ by Grädel, Kolaitis, and Vardi (Grädel, Kolaitis, and Vardi 1997), where three sets of elements provide witnesses for each other in a circular fashion. Our setting is slightly more intricate, but the general idea remains: $W_{\mathrm{II}}$ provides $\mathcal{C}$-witnesses for $W_{\mathrm{I}}$, $W_{\mathrm{III}}$ for $W_{\mathrm{II}}$, and, to close the cycle, we adjust the 2-types so that $W_{\mathrm{I}}$ provides $\mathcal{C}$-witnesses for $W_{\mathrm{III}} \setminus (W_{\mathrm{I}} \cup W_{\mathrm{II}})$. Moreover, $W_{\mathrm{I}}$ needs to provide $\mathcal{C}$-witnesses for elements outside $W_{\mathrm{I}} \cup W_{\mathrm{II}} \cup W_{\mathrm{III}}$. We ensure all of this in the subsequent steps of the proof.

We define the set $\mathcal{D}$ (from the statement of Lemma 6) as:

$$\mathcal{D} := \bigcup_{a\in(W_{\mathrm{I}}\cup W_{\mathrm{II}}\cup W_{\mathrm{III}})\cap\mathcal{C}} E_k^{\mathfrak{A}}[a].$$

This choice of $\mathcal{D}$ satisfies the first item of Lemma 6:

**Claim 8.** *The set $\mathcal{D}$ is a union of at most $12 \cdot M^3 \cdot |\boldsymbol{\alpha}|$ many $E_k$-classes of $\mathcal{C}/E_k^{\mathfrak{A}}$.*

*Proof.* We have that $|W_{\mathrm{I}}| \leq 4 \cdot M \cdot |\boldsymbol{\alpha}|$, $|W_{\mathrm{II}}| \leq M \cdot |W_{\mathrm{I}}|$, and $|W_{\mathrm{III}}| \leq M \cdot |W_{\mathrm{II}}|$. Therefore, each of the sets $W_{\mathrm{I}}$, $W_{\mathrm{II}}$ and $W_{\mathrm{III}}$ is of size at most $4 \cdot M^3 \cdot |\boldsymbol{\alpha}|$. We conclude, as $|\mathcal{D}| \leq |W_{\mathrm{I}} \cup W_{\mathrm{II}} \cup W_{\mathrm{III}}| \leq 12 \cdot M^3 \cdot |\boldsymbol{\alpha}|$. □

To conclude this proof, it remains to establish also the second item of Lemma 6: the sentence $\varphi$ has a model $\mathfrak{B}$ over the domain $B := (A \setminus \mathcal{C}) \cup \mathcal{D}$. In this structure, the 1-types, the linear order $<$, and the relations $E_k$ are induced from $\mathfrak{A}$, yet some 2-types will be redefined in a specific way.

In particular, this redefinition will ensure the elements of $B \setminus (W_{\mathrm{I}} \cup W_{\mathrm{II}})$ to have $\mathcal{C}$-witnesses in $W_{\mathrm{I}}$.

Let $N$ denote $|B \setminus (W_{\mathrm{I}} \cup W_{\mathrm{II}})|$. We fix an enumeration $a_1, \ldots, a_N$ of the set $B \setminus (W_{\mathrm{I}} \cup W_{\mathrm{II}})$. We inductively construct a sequence of structures $\mathfrak{B}_0, \mathfrak{B}_1, \ldots, \mathfrak{B}_N$: $\mathfrak{B}_0$ is the induced structure $\mathfrak{A}\restriction B$, and, for each $n \in [0, N-1]$, $\mathfrak{B}_{n+1}$ is obtained from $\mathfrak{B}_n$ by modifying some 2-types between $a_{n+1}$ and $W_{\mathrm{I}} \setminus E_k^{\mathfrak{A}}[a_{n+1}]$ in the way described below.

Let $W_{\mathcal{C}}(a_{n+1})$ denote the set of $\mathcal{C}$-witnesses of $a_{n+1}$ suggested by the witness functions $\mathfrak{f}_m$:

$$W_{\mathcal{C}}(a_{n+1}) := \{\mathfrak{f}_m(a_{n+1}) \mid m \in [M]\} \cap (\mathcal{C} \setminus E_k^{\mathfrak{A}}[a_{n+1}]).$$

We set $b_1, \ldots, b_\ell$ as the distinct elements of $W_{\mathcal{C}}(a_{n+1})$, with $\ell \leq M$. From Claim 7, there are distinct elements $c_1, \ldots, c_\ell$ from $W_{\mathrm{I}}$ realising the same $a_{n+1}$-configuration as $b_1, \ldots, b_\ell$. We define the structure $\mathfrak{B}_{n+1}$ almost as $\mathfrak{B}_n$, except that for each $c_j$ (which is in $W_{\mathrm{I}} \subseteq B$), the 2-type between $a_{n+1}$ and $c_j$ is defined as the 2-type between $a_{n+1}$ and $b_j$, i.e., $\mathrm{tp}^{\mathfrak{B}_{n+1}}[a_{n+1}, c_j] := \mathrm{tp}^{\mathfrak{B}_n}[a_{n+1}, b_j]$.

The following claim summarises the properties of $\mathfrak{B}_{n+1}$.

**Claim 9.** *Suppose that the numbers $n$ and $N$, the elements $a_1, \ldots, a_N$ of $B \setminus (W_{\mathrm{I}} \cup W_{\mathrm{II}})$, and the structure $\mathfrak{B}_{n+1}$ are all defined as above. Then:*

1. *$\mathfrak{B}_{n+1}$ realises only the 1-types and 2-types from $\mathfrak{A}$.*
2. *The linear order $<^{\mathfrak{B}_{n+1}}$ and the equivalence relations $E_j^{\mathfrak{B}_{n+1}}$, for every $j \in [K]$, are all inherited from $\mathfrak{A} \restriction B$.*
3. *For all distinct $a, b \in B$, $\mathrm{tp}^{\mathfrak{A}}[a, b] = \mathrm{tp}^{\mathfrak{B}_{n+1}}[a, b]$ unless $(a, b)$ or $(b, a)$ is in $(\{a_1, \ldots, a_{n+1}\} \times W_{\mathrm{I}}) \setminus E_k^{\mathfrak{A}}$.*
4. *For each $i \in [n+1]$ and each $b \in W_{\mathcal{C}}(a_i)$, there exists $c \in W_{\mathrm{I}}$ such that $\mathrm{tp}^{\mathfrak{A}}[a_i, b] = \mathrm{tp}^{\mathfrak{B}_{n+1}}[a_i, c]$.*

From the items of Claim 9, we conclude that the structure $\mathfrak{B}$, defined as the final structure $\mathfrak{B}_N$, is a model of $\varphi$ over the domain $(A \setminus \mathcal{C}) \cup \mathcal{D}$. This finishes the proof of the second item of Lemma 6.

## 4 How to show Theorems 2 and 4

Theorems 2 and 4 follow by similar arguments as Theorem 1. We sketch them here; details are in the full version.

*Theorem 2.* We proceed via a reduction to FO$^2[<, \mathcal{EQ}^{\subseteq}]$. Let $\mathfrak{A}_0$ be a finite FO$^2[\preceq^{\subseteq}]$-model of $\varphi$. We expand it to an FO$^2[<, \mathcal{EQ}^{\subseteq}]$-structure $\mathfrak{A}$ by: (i) setting $E_k^{\mathfrak{A}} := \preceq_k^{\mathfrak{A}_0} \cap (\preceq_k^{\mathfrak{A}_0})^{-1}$ for all $k$; and (ii) interpreting $<^{\mathfrak{A}}$ as a linear order derived from $\preceq_1{}^{\mathfrak{A}_0}$ by resolving ties within $E_1$-classes arbitrarily. Applying Lemma 6 to $\mathfrak{A}$ yields a model $\mathfrak{B}$ over the domain $B = (A \setminus \mathcal{C}) \cup \mathcal{D}$ with $\mathcal{D} \subseteq \mathcal{C}$. Item 2 of Claim 9 implies that $\preceq_k^{\mathfrak{B}} = \preceq_k^{\mathfrak{A}\restriction B}$ for all $k$. Since nested preorders are preserved under taking substructures, $\mathfrak{B}$ respects the semantics of $\preceq_k$-symbols, and thus it can be naturally transformed back into an FO$^2[\preceq^{\subseteq}]$-model of $\varphi$. Since the size bound analysis from the proof of Theorem 1 still applies, we get a small model and thereby establish Theorem 2.

*Theorem 4.* The argument proceeds analogously to that for Lemma 6 and Theorem 1, but now starting from a possibly infinite model $\mathfrak{A}$. In the absence of the linear order, we may choose arbitrary realisations of 1-types into $r$-extremal subsets (rather than the maximal and minimal ones). The rest of the reasoning carries over directly, with the only modification being that, due to the possible infinity of $\mathfrak{A}$, it is sometimes necessary to work with natural limit structures.

## 5 Proof of Theorem 3

In this section, we establish our second main result.

**Theorem 3.** *Finitely satisfiable sentences of* FO$^2[\preceq^{\subseteq}_{\mathrm{succ}}]$ *admit models of doubly exponential size. The finite satisfiability problem for* FO$^2[\preceq^{\subseteq}_{\mathrm{succ}}]$ *is* ExpSpace*-complete.*

For this section, we fix a sentence $\varphi$ of FO$^2[\preceq^{\subseteq}_{\mathrm{succ}}]$, and assume that it is in Scott Normal Form, i.e., in the shape of $\forall x, y.\, \psi_0(x, y) \wedge \bigwedge_{m=1}^{M} \forall x.\, \exists y.\, \psi_m(x, y)$, where $M$ denotes the number of Skolem conjuncts. As explained in Subsection 2.2, the problem of transforming any FO$^2$ sentence into Scott Normal Form, satisfiable over the same domains, is in polynomial-time. Hence, this assumption is without loss of generality. By convention, $\varphi$ uses the first $K$ special symbols $\preceq_k$ for $k \in [K]$, and the corresponding derived symbols $E_k$ and $\mathcal{S}_k$; moreover we interpret the symbols $E_0$ and $E_{K+1}$ as the identity and universal relations, respectively.

**Small model property.** We begin with the first part of Theorem 3, establishing that every finitely satisfiable sentence of

**Algorithm 1:** Finite Satisfiability for $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$

**Input:** A sentence $\varphi$ in Scott Normal Form
**Output:** Decides whether $\varphi$ is finitely satisfiable

**guess** the domain size $N \in \mathbb{N}$;
**for** $i \leftarrow 1, 2, \ldots, N$ **do**
  **guess** the number $r_i \in [K{+}1]$;
  **guess** the 1-type $\alpha_i \in \boldsymbol{\alpha}$;
  **if** $\alpha_i \not\models \psi_0(x,x)$ **then reject**;
  $W_i \leftarrow \{m \in [M] \mid \alpha_i \models \psi_m(x,x)\}$;
  **for** $j \in \{1, \ldots, i{-}1\}$ **do**
    **if** $q_j \leq r_i$ **then**
      $p_j \leftarrow q_j$;
      $q_j \leftarrow r_i$;
    **guess** the 2-type $\beta_{j,i} \in \boldsymbol{\beta}[\alpha_j, \alpha_i, p_j, q_j]$;
    **if** $\beta_{j,i} \not\models \psi_0(x,y) \wedge \psi_0(y,x)$ **then reject**;
    $W_j \leftarrow W_j \cup \{m \in [M] \mid \beta_{j,i} \models \psi_m(x,y)\}$;
    $W_i \leftarrow W_i \cup \{m \in [M] \mid \beta_{j,i} \models \psi_m(y,x)\}$;
  $p_i \leftarrow 1$; $q_i \leftarrow 1$;
**if** $W_i = [M]$ *for every* $i \in [N]$ **then accept**;
**else reject**;

$\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$ has a model of doubly exponential size. The key ingredient is the following lemma, analogous to Lemma 6; its proof is deferred to the full version of the paper.

**Lemma 10.** *Let $\mathfrak{A}$ be a finite model of $\varphi$. Fix $k \in [0, K]$, and let $\mathcal{C}$ be an $E_{k+1}$-class in $A/E^{\mathfrak{A}}_{k+1}$. Then there exists a subset $\mathcal{D} \subseteq \mathcal{C}$ such that:*

1. *$\mathcal{D}$ is the union of $2^{2^{\mathcal{O}(|\varphi|)}}$ many $E_k$-classes from $\mathcal{C}/E^{\mathfrak{A}}_k$.*
2. *$\varphi$ has a model $\mathfrak{B}$ with domain $B = (A \setminus \mathcal{C}) \cup \mathcal{D}$.*

The argument for deriving the small model bound from Lemma 10 is parallel to that of Theorem 1 from Lemma 6.

Let $\mathfrak{A}$ be a minimal-size finite model. By Lemma 10, for each $k \in [0, K]$, every $E_{k+1}$-class of $\mathfrak{A}$ is partitioned into at most $2^{2^{\mathcal{O}(|\varphi|)}}$ many $E_k$-classes. Since each $E_0$-class is a singleton and $E^{\mathfrak{A}}_{K+1}$ has a single equivalence class, it follows that $\mathfrak{A}$ is indeed of doubly exponential size:

$$|A| \leq \left(2^{2^{\mathcal{O}(|\varphi|)}}\right)^{K+1} = 2^{(K+1)\times 2^{\mathcal{O}(|\varphi|)}} = 2^{2^{\mathcal{O}(|\varphi|)}}.$$

**Satisfiability-checking procedure.** We now turn to the second part of Theorem 3, establishing an ExpSpace procedure deciding the finite satisfiability problem for $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$. A preliminary version of this procedure, operating in 2-ExpSpace, is given in Algorithm 1.

We call a 1-type or 2-type *proper* if it interprets all the special symbols $\preceq_k$, $E_k$, and $\mathcal{S}_k$ in a manner consistent with the semantics of $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$. In other words, proper types are exactly those that can occur in actual models of $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$. In particular, they respect the natural compatibility conditions between equivalences and preorders, e.g., if $E_k(x,y)$ holds, then $x \preceq_k y$ and $y \preceq_k x$ must both hold as well. For 2-types, we additionally require that they entail $x \preceq_k y$ for the relevant $k$; this condition guarantees compatibility with the fixed ordering alignment assumed in the algorithm. Denote by $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ the sets of proper, respectively, 1-types and 2-types over the signature of $\varphi$.

In the procedure, we use a table $\boldsymbol{\beta}[*,*,*,*]$. In this table, the entry $\boldsymbol{\beta}[\alpha_1, \alpha_2, p, q]$, for every $\alpha_1, \alpha_2 \in \boldsymbol{\alpha}$ and $1 \leq p \leq q \leq K{+}1$, stores a set of 2-types. We put a 2-type $\beta \in \boldsymbol{\beta}$ into the set $\boldsymbol{\beta}[\alpha_1, \alpha_2, p, q]$ if the following conditions hold:

(i) $\mathrm{tp}^{\beta}[x] = \alpha_1$ and $\mathrm{tp}^{\beta}[y] = \alpha_2$.

(ii) For each $k \in [K]$, $E_k(x,y) \in \beta$ iff $q \leq k$.

(iii) For each $k \in [K]$, $S_k(x,y) \in \beta$ iff $p \leq k < q$.

Algorithm 1 attempts to verify the existence of a model $\mathfrak{A}$ for $\varphi$ over the domain $\{1, \ldots, N\}$, for some $N \in \mathbb{N}$. We assume that the preorders $\preceq^{\mathfrak{A}}_k$ are aligned with the natural order $\leq$ on $\mathbb{N}$ (i.e., $i \prec_k j \rightarrow i < j$). In the $i$th iteration of the outer loop, the algorithm examines how the new element $i$ interacts with all previous elements $1, \ldots, i{-}1$, denoted by $j$ in the inner loop. The variable $q_j$ stores the smallest index $k$ such that $j$ and $i$ belong to the same $E_k$-class; and $K{+}1$ is a sentinel value. If $p_j < q_j$, then $p_j$ records the smallest index $k$ for which $j$ and $i$ lie in $\preceq_k$-consecutive $E_k$-classes; if instead $p_j = q_j$, then no such consecutive $E_k$-classes exist.

To illustrate better the idea of Algorithm 1, we visualise the structure in the following schematic, each horizontal brace groups together the elements of one $E_k$-class, and $\preceq_k$ linearly orders the $E_k$-classes from left to right:

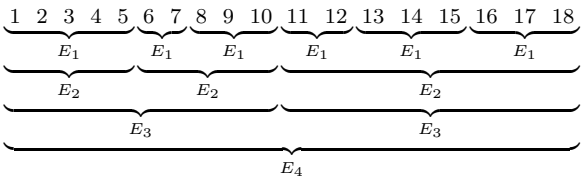


Consider now the 18th iteration of the outer loop, that is, $i = 18$. The values of the variables $p_j$ and $q_j$ for selected $j$ are: $p_5 = 3$, $q_5 = 4$; $p_9 = 2$, $q_9 = 4$; $p_{12} = 2$, $q_{12} = 2$.

Intuitively, these pairs $(p_j, q_j)$ capture the *relative position* of $j$ with respect to $i$ in the hierarchy of equivalence classes, allowing the algorithm to determine which 2-types between elements $j$ and $i$ are admissible.

The algorithm maintains, for each element $i \in [N]$, a set $W_i \subseteq [M]$ recording which Skolem conjuncts are already satisfied for $i$, and accepts precisely when every element has all its witnesses secured, i.e., when $W_i = [M]$ for all $i$.

We formalise the above explanations in the following claim, establishing completeness of the procedure.

**Claim 11.** *Suppose that $\varphi$ has a finite model $\mathfrak{A}$. Then there exists a sequence of guesses of Algorithm 1 such that, after $i$ iterations, where $i \in [N]$, the following invariants hold:*

- *For $j \in [i]$, $\mathrm{tp}^{\mathfrak{A}}[j] = \alpha_j$, and for $j < k \leq i$, $\mathrm{tp}^{\mathfrak{A}}[j,k] = \beta_{j,k}$.*
- *If $i > 1$, then, for $k \in [K]$, $(i{-}1, i) \in E^{\mathfrak{A}}_k$ iff $r_i \leq k$.*
- *For $j \in [i]$ and $k \in [K]$, $(j,i) \in E^{\mathfrak{A}}_k$ iff $q_j \leq k$.*
- *For $j \in [i]$ and $k \in [K]$, $(j,i) \in \mathcal{S}^{\mathfrak{A}}_k$ iff $p_j \leq k < q_j$.*
- *For $j \in [i]$, $W_j$ indicates witnesses of $j$ located in $\mathfrak{A} \restriction [i]$.*

Soundness follows as well. Given a transcription of an accepting run, we reconstruct a finite model of $\varphi$ as follows. Take the domain to be $\{1, \ldots, N\}$, assign to each element $i$

the 1-type $\alpha_i$, and for each ordered pair $(j, i)$ assign the 2-type $\beta_{j,i}$. Since each pair is processed exactly once, and all 2-types are determined consistently with both the preorder relations and their successors, no conflicts can arise. The resulting structure therefore satisfies $\varphi$.

**Corollary 12.** *Algorithm 1 is sound and complete.*

Algorithm 1 operates in space proportional to $N$. By the already established small model property, $N$ can be assumed doubly exponential in $|\varphi|$, yielding a 2-ExpSpace bound.

However, we note that the precise identities of the previously processed elements $j < i$ are irrelevant: all information needed about such an element $j$ is captured by the tuple $\langle \alpha_j, W_j, p_j, q_j \rangle$. Consequently, after the $i$th iteration, the algorithm's configuration can be represented compactly as a *counting function*, recording the number of such tuples:

$$F_i \colon \boldsymbol{\alpha} \times 2^{[M]} \times [K{+}1] \times [K{+}1] \to \mathbb{N}.$$

It is straightforward to adapt the procedure to work directly with this representation. Storing $F_i$ explicitly as a table requires $\mathcal{O}\big(|\boldsymbol{\alpha}| \cdot 2^M \cdot K^2 \cdot \log N\big)$ bits of space. Since $|\boldsymbol{\alpha}|$ is $2^{\mathcal{O}(|\varphi|)}$, $M$ and $K$ are $\mathcal{O}(|\varphi|)$, and $N$ is doubly exponential in $|\varphi|$, the overall space usage is singly exponential in $|\varphi|$.

**Corollary 13.** *Algorithm 1 can be implemented to operate in exponential space.*

**Lower bound.** The final part of the proof of Theorem 3 is to show the matching ExpSpace-hardness lower bound for the finite satisfiability problem for $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$.

Here we give a short proof sketch; full details are deferred to the full version. We reduce from the corridor tiling problem, which is ExpSpace-complete. An instance consists of a set of colours $\mathcal{C}$, horizontal and vertical adjacency constraints, colours for the top-left and bottom-right cells (initial conditions), and a number $n$ in unary. The decision question is whether there exists a number $m$ and a labelling with colours from $\mathcal{C}$ of a $2^n \times m$ grid that satisfies all constraints.

The reduction is realised by a conjunction of $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$ sentences that enforce the structure to encode such a grid. Introduce unary predicates $B_j$ for $0 \leq j < n$ together with unary predicates $P_c$ for each colour $c \in \mathcal{C}$. Interpretation: each domain element is a grid cell; the predicates $B_j$ encode the horizontal coordinate in binary (yielding width $2^n$), and each $E_1$-equivalence class represents a full row. Within a row, $P_c$ marks the colour of the cell at the horizontal position given by the bits. The successor relation $\mathcal{S}_1$ links consecutive $E_1$-classes, so rows form a vertical chain of length $m$. Expressing the horizontal and vertical adjacency constraints, as well as the initial conditions, in $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$ is routine.

**Corollary 14.** *The finite satisfiability problem for* $\mathrm{FO}^2[\preceq^{\subseteq}_{\mathrm{succ}}]$ *is* ExpSpace*-hard.*

## 6 Proof of Theorem 5

In this section, we establish our undecidability result.

**Theorem 5.** *The satisfiability and finite satisfiability problems are undecidable for the constant-free, equality-free, monadic fragment of* $\mathrm{FO}^2$ *extended with four special symbols $E_1, E_2, F_1, F_2$, interpreted as equivalence relations such that $E_2$ is coarser than $E_1$ and $F_2$ is coarser than $F_1$.*

To the end of this section, we prove Theorem 5.

The idea demonstrated here is similar to Pratt-Hartmann's proof of the undecidability of $\mathrm{FO}^2$ with counting and two equivalence relations (Pratt-Hartmann 2014). Hovewer, in the absence of the equality symbol, additional technical refinements are required.

The undecidability is proven via a reduction from the halting problem for *two-counter machines*. Such a machine $\mathbf{M}$ consists of a finite set $S = \{s_0, \ldots, s_F\}$ of states, one being initial (say $s_0$) and one being final (say $s_F$); two counters $c_1$ and $c_2$ holding non-negative integers; and a set of transitions $\delta \subseteq S \times Op \times S$ between states and operations over the counters; these operations are: increment a counter, decrement a counter (if it is non-zero), or test if a counter is zero. A *configuration* is a triple $\langle s, c_1, c_2 \rangle$ holding the state $s$ and the values of the two counters $c_1$ and $c_2$. A *run* is a (potentially infinite) sequence of configurations which starts with the initial state $s_0$ and respects the transition function. The machine *terminates* if it ever reaches the final state $s_F$. As two-counter machines can simulate Turing machines, the halting problem for two-counter machines is undecidable (more precisely, it is r.e.-complete).

We encode a run of some two-counter machine $\mathbf{M}$. Each configuration will be contained within an equivalence class of the intersection relation $E_2 \cap F_2$. We write $G_2(x, y)$ for the formula $E_2(x, y) \wedge F_2(x, y)$. We partition the domain with two unary predicates $d_E$ and $d_F$, and axiomatise that they are universal in each $G_2$-class:

$$\forall x.\, d_E(x) \leftrightarrow \neg d_F(x) \tag{1}$$

$$\forall x, y.\, G_2(x, y) \to \big(d_E(x) \leftrightarrow d_E(y)\big) \tag{2}$$

If a configuration has its elements satisfying $d_E$ (resp. $d_F$), we call it a *$d_E$-configuration* (resp. *$d_F$-configuration*). We specify that each $E_2$- and $F_2$-class has at most one configuration of each type:

$$\bigwedge_{\alpha \in \{E,F\}} \forall x, y.\, \big(E_2(x, y) \wedge d_\alpha(x) \wedge d_\alpha(y)\big) \to F_2(x, y) \tag{3}$$

$$\bigwedge_{\alpha \in \{E,F\}} \forall x, y.\, \big(F_2(x, y) \wedge d_\alpha(x) \wedge d_\alpha(y)\big) \to E_2(x, y) \tag{4}$$

We define now that each configuration is in precisely one state $s$ from the set of states $S$ of $\mathbf{M}$:

$$\forall x.\, \bigvee_{s \in S} \forall y.\, G_2(x, y) \to \Big(s(y) \wedge \bigwedge_{s' \in S\colon\, s' \neq s} \neg s'(y)\Big) \tag{5}$$

We order configurations with a definable binary relation $t$: among an $E_2$-class (resp. $F_2$-class), it links elements from the $d_E$-configuration to the successive $d_F$-one (resp. from the $d_F$-one to the successive $d_E$-one), see Figure 1:

$$\begin{aligned} t(x, y) := {} & \big(E_2(x, y) \wedge d_E(x) \wedge d_F(y)\big) \\ & \vee \big(F_2(x, y) \wedge d_F(x) \wedge d_E(y)\big) \end{aligned}$$

We have discussed so far how to axiomatise the succession of configurations, we explain now how to implement the counters with the help of the two thinner equivalence relations $E_1$ and $F_1$. We write $G_1(x, y)$ for the formula

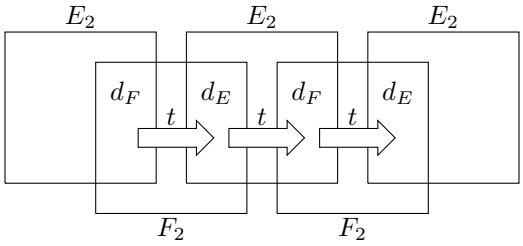


Figure 1: A succession of configurations.

$E_1(x,y) \vee F_1(x,y)$, defining the union of $E_1$ and $F_1$; a priori not an equivalence relation. We require first that inside any $G_2$-class, $E_1$ and $F_1$ define the same relation:

$$\forall x, y.\ \big(G_2(x,y) \wedge G_1(x,y)\big) \rightarrow \big(E_1(x,y) \wedge F_1(x,y)\big) \quad (6)$$

We introduce two unary predicates $c_1$ and $c_2$. Their intended meaning is as follows: If $\mathcal{E}$ is a $G_2$-class (representing a configuration), then, inside it, the value of the counter $c_i$ is the number of equivalence classes of the relation $E_1 \cap \mathcal{E}^2$ whose elements satisfy $c_i$. In particular, if no element of $\mathcal{E}$ satisfies $c_i$, then the value of the counter $c_i$ equals zero. We permit elements to satisfy only a state predicate, without being marked by $c_1$ or $c_2$. This ensures that configurations where both counters are zero can still be represented.

We require first that $c_1$ and $c_2$ are mutually exclusive:

$$\forall x.\ \neg c_1(x) \vee \neg c_2(x) \quad (7)$$

We also state that $c_1$ and $c_2$ propagate within $G_1$-connected components (and hence they do not interact via $E_1$ or $F_1$):

$$\bigwedge_{i=1}^{2} \forall x, y.\ \big(G_1(x,y) \wedge c_i(x)\big) \rightarrow c_i(y) \quad (8)$$

Then, the change of the counters is held as follows. In a $d_E$-configuration, $E_1$ may extend to the successive $d_F$-configuration and $F_1$ may extend to the preceding $d_F$-configuration. And, in a $d_F$-configuration, the situation is symmetric. This makes it possible to transfer the values of counters to the neighbouring configurations.

We define the formula $\varphi_{\exists\rightarrow}(x)$ stating that the element $x$ admits a $t$-successor linked via $G_1$:

$$\varphi_{\exists\rightarrow}(x) := \exists y.\ t(x,y) \wedge G_1(x,y).$$

Notice that if two elements of two distinct configurations are linked via $t$, then they are in particular linked via precisely one of $E_2$ and $F_2$, which does not leave any ambiguity whether they are linked via $E_1$ or $F_1$.

We also define the symmetric formula $\varphi_{\exists\leftarrow}(x)$ stating that $x$ admits a $t$-predecessor linked via $G_1$.

We define now, for each $i \in \{1,2\}$, a formula $\varphi_{\forall i\exists\rightarrow}(x)$ stating that every element $y$ of the $G_2$-class of $x$ which satisfies $c_i$ admits a $t$-successor via $G_1$:

$$\varphi_{\forall i\exists\rightarrow}(x) := \forall y.\ \big(G_2(x,y) \wedge c_i(y)\big) \rightarrow \varphi_{\exists\rightarrow}(y).$$

Because $E_1$ and $F_1$ are equivalence relations, we can see that if $\varphi_{\forall i\exists\rightarrow}(x)$ is true, then the configuration succeeding that of $x$ has a value of $c_i$ greater than or equal to the value of $c_i$ in the configuration of $x$. In order to make sure that it cannot increase by more than one, we add a condition stating that each two elements of the same configuration that satisfy $c_i$ and have no predecessor via $G_1$ are necessarily related via $E_1$, and hence at most one equivalence relation is added:

$$\bigwedge_{i=1}^{2} \forall x, y.\ \big(G_2(x,y) \wedge c_i(x) \wedge c_i(y) \wedge \neg\varphi_{\exists\leftarrow}(x) \wedge \neg\varphi_{\exists\leftarrow}(y)\big) \rightarrow E_1(x,y) \quad (9)$$

Symmetrically, we define $\varphi_{\forall i\exists\leftarrow}(x)$ stating that every element $y$ of the $G_2$-class of $x$ which satisfies $c_i$ admits a $t$-predecessor via $G_1$, and we also define the symmetric condition making sure that each counter decreases by at most one at every step:

$$\bigwedge_{i=1}^{2} \forall x, y.\ \big(G_2(x,y) \wedge c_i(x) \wedge c_i(y) \wedge \neg\varphi_{\exists\rightarrow}(x) \wedge \neg\varphi_{\exists\rightarrow}(y)\big) \rightarrow E_1(x,y) \quad (10)$$

We define below four auxiliary formulas $\varphi_{i,\text{zero}}(x)$, $\varphi_{i,\text{eq}}(x)$, $\varphi_{i,\text{incr}}(x)$ and $\varphi_{i,\text{decr}}(x)$.

First, the formula $\varphi_{i,\text{zero}}(x)$ states that no element of the $G_2$-class of $x$ satisfy $c_i$, and hence the counter $c_i$ is zero:

$$\varphi_{i,\text{zero}}(x) := \forall y.\ G_2(x,y) \rightarrow \neg c_i(y).$$

Then, the formula $\varphi_{i,\text{eq}}(x)$ states that the counter $c_i$ stays the same from the configuration of $x$ to the next one:

$$\varphi_{i,\text{eq}}(x) := \varphi_{\forall i\exists\rightarrow}(x) \wedge \forall y.\ t(x,y) \rightarrow \varphi_{\forall i\exists\leftarrow}(y).$$

Finally, the formulas $\varphi_{i,\text{incr}}(x)$ and $\varphi_{i,\text{decr}}(x)$ state that the counter $c_i$, respectively, increments and decrements from the configuration of $x$ to the successive one:

$$\varphi_{i,\text{incr}}(x) := \varphi_{\forall i\exists\rightarrow}(x) \wedge \forall y.\ t(x,y) \rightarrow \neg\varphi_{\forall i\exists\leftarrow}(y)$$
$$\varphi_{i,\text{decr}}(x) := \neg\varphi_{\forall i\exists\rightarrow}(x) \wedge \forall y.\ t(x,y) \rightarrow \varphi_{\forall i\exists\leftarrow}(y)$$

Having the above formulas, as well as the unary state predicates, we can easily axiomatise all the transitions of the two-counter machine $\mathbf{M}$ together with its initial and final configurations; we leave details to the reader.

Let $\Theta_{\mathbf{M}}$ be a sentence expressing the existence of a run of $\mathbf{M}$ that reaches the final state $s_F$. It is immediate to see that this sentence $\Theta_{\mathbf{M}}$ is finitely satisfiable if and only if $\mathbf{M}$ terminates. This allows us to conclude the undecidability of finite satisfiability. For general (infinite) satisfiability, we consider a sentence $\Theta'_{\mathbf{M}}$ expressing the existence of an infinite run of $\mathbf{M}$ (i.e. never reaching the final state $s_F$). Hence, the general satisfiability problem is also undecidable.

## 7 Conclusion

We analysed four extensions of $\text{FO}^2$ that enable decidable reasoning over nested equivalence relations: $\text{FO}^2[\mathcal{EQ}^{\subseteq}]$, $\text{FO}^2[<, \mathcal{EQ}^{\subseteq}]$, $\text{FO}^2[\preceq^{\subseteq}]$, and $\text{FO}^2[\preceq^{\subseteq}_{\text{succ}}]$, establishing the precise complexity of the finite satisfiability problems and bounds on the size of minimal finite models. Except $\text{FO}^2[\mathcal{EQ}^{\subseteq}]$ that enjoys the finite model property, understanding general (infinite) satisfiability remains open. Finally, we demonstrated that extending $\text{FO}^2$ with two independent families of nested equivalence relations leads to undecidability.

## Acknowledgments

The first and second authors were supported by the Polish National Science Center grant No. 2021/41/B/ST6/00996. The third author was supported by the ERC grant INFSYS, agreement No. 950398. We thank the first anonymous reviewer for his helpful feedback.

## References

Benaim, S.; Benedikt, M.; Charatonik, W.; Kieroński, E.; Lenhardt, R.; Mazowiecki, F.; and Worrell, J. 2016. Complexity of two-variable logic on finite trees. *ACM Trans. Comput. Log.* 17(4):32:1–32:38.

Björklund, H., and Bojańczyk, M. 2007. Shuffle expressions and words with nested data. In *Proceedings of the 32nd International Conference on Mathematical Foundations of Computer Science*, MFCS'07, 750–761.

Bojańczyk, M.; Muscholl, A.; Schwentick, T.; and Segoufin, L. 2009. Two-variable logic on data trees and xml reasoning. *J. ACM* 56(3).

Bojańczyk, M.; David, C.; Muscholl, A.; Schwentick, T.; and Segoufin, L. 2011. Two-variable logic on data words. *ACM Trans. Comput. Log.* 12(4):27.

Charatonik, W., and Witkowski, P. 2016. Two-variable logic with counting and trees. *ACM Trans. Comput. Log.* 17(4):31.

Danielski, D., and Kieroński, E. 2019. Finite satisfiability of unary negation fragment with transitivity. In *44th International Symposium on Mathematical Foundations of Computer Science, MFCS*, volume 138 of *LIPIcs*, 17:1–17:15.

Grädel, E.; Kolaitis, P.; and Vardi, M. Y. 1997. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic* 3(1):53–69.

Kahr, A.; Moore, E.; and Wang, H. 1962. Entscheidungsproblem reduced to the $\forall\exists\forall$ case. *Proc. Nat. Acad. Sci. U.S.A.* 48:365–377.

Kieroński, E., and Otto, M. 2012. Small substructures and decidability issues for first-order logic with two variables. *Journal of Symbolic Logic* 77:729–765.

Kieroński, E.; Michaliszyn, J.; Pratt-Hartmann, I.; and Tendera, L. 2014. Two-variable first-order logic with equivalence closure. *SIAM Journal of Computing* 43(3):1012–1063.

Kieroński, E. 2011. Decidability issues for two-variable logics with several linear orders. In *Computer Science Logic, 25th International Workshop / 20th Annual Conference of the EACSL, CSL 2011, Proceedings*, volume 12 of *LIPIcs*, 337–351.

Otto, M. 2001. Two-variable first-order logic over ordered domains. *Journal of Symbolic Logic* 66:685–702.

Pratt-Hartmann, I. 2014. Logics with counting and equivalence. In *Proceedings of the Joint Meeting of CSL and LICS*, CSL-LICS '14, 76:1–76:10. ACM.

Pratt-Hartmann, I. 2023. *Fragments of First-Order Logic*. Oxford Logic Guides. United Kingdom: Oxford University Press.

Schwentick, T., and Zeume, T. 2012. Two-variable logic with two order relations. *Logical Methods in Computer Science* 8(1).

Scott, D. 1962. A decision method for validity of sentences in two variables. *Journal Symbolic Logic* 27:477.

ten Cate, B., and Segoufin, L. 2013. Unary negation. *Logical Methods in Comp. Sc.* 9(3).

Zeume, T., and Harwath, F. 2016. Order-invariance of two-variable logic is decidable. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science*, LICS '16, 807–816.