Order-Invariance in the Two-Variable Fragment of First-Order Logic

Julien Grange

Order-Invariance in the Two-Variable Fragment of First-Order Logic

Julien Grange ✉
LACL, Université Paris-Est Créteil, France

Abstract

We study the expressive power of the two-variable fragment of order-invariant first-order logic. This logic departs from first-order logic in two ways: first, formulas are only allowed to quantify over two variables. Second, formulas can use an additional binary relation, which is interpreted in the structures under scrutiny as a linear order, provided that the truth value of a sentence over a finite structure never depends on which linear order is chosen on its domain.

We prove that on classes of structures of bounded degree, any property expressible in this logic is definable in first-order logic. We then show that the situation remains the same when we add counting quantifiers to this logic.

2012 ACM Subject Classification Theory of computation → Finite Model Theory

Keywords and phrases Finite model theory, Two-variable logic, Order-invariance

Digital Object Identifier 10.4230/LIPIcs.CSL.2023.23

1 Introduction

The restriction of first-order logic to two variables (FO^2) holds an important place among the fragments of first-order logic (FO), since it is the maximal fragment, with respect to the number of variables, for which the finite and general satisfiability problems are decidable [14] – see [5] for a complete survey on these issues. They become undecidable as soon as we allow formulas to make use of three variables, as three variables are enough to encode grids, and thus runs of Turing machines. Tame as it is in this regard, FO^2 is still fairly expressive (in particular, it embeds modal logic). It is thus natural to investigate its order-invariant extension $<$-inv FO^2.

In the order-invariant extension $<$-inv $\mathcal{L}$ of a logic $\mathcal{L}$, one can make use in the $\mathcal{L}$-formulas of a linear order on the vertices of the structures at hand, provided that the validity of said formulas in a given finite structure does not depend on the choice of a particular order. Such a notion is very natural and useful both in database theory (where it corresponds to the requirement for a query to be independent of the order in which the data is stored on disk) and in descriptive complexity (where structures are needed to be ordered for a logic to capture complexity classes such as PTIME [10] and PSPACE [17]).

If $\mathcal{L} = \mathrm{FO}$, we get $<$-inv FO, whose syntax is not recursively enumerable. It has however been proven by Harwath and Zeume [19] that on the other hand, $<$-inv FO^2 has a recursive syntax. The natural follow-up to this result is to study the expressive power of $<$-inv FO^2. This shall be our endeavor in this article.

Perhaps surprisingly, it has been shown by Gurevich (see Section 5.2 of [13]) that such an order can indeed bring additional expressive power to FO, even when restricted in this way: there exist properties which are not definable in FO, but which can be expressed as soon as the use of a linear order is authorized, even in an invariant fashion. In symbols: $<$-inv FO $\not\subseteq$ FO. It is not hard to observe that $<$-inv $\mathrm{FO}^2 \not\subseteq \mathrm{FO}^2$ (for instance, one can state in $<$-inv FO^2 that a set has at least three elements, which is not possible in FO^2). It is however not clear whether $<$-inv $\mathrm{FO}^2 \subseteq \mathrm{FO}$, or whether even when restricted to two variables, the addition of an order allows one to express properties beyond the scope of FO.[a] This question is asked in [19]. Is this paper, we prove that on any class of bounded degree, the inclusion $<$-inv $\mathrm{FO}^2 \subseteq \mathrm{FO}$ indeed holds - this is Theorem 1, whose proof is the object of Sections 4 and 5. We then explain in Section 6 how to extend this result from $<$-inv FO^2 to $<$-inv C^2, where C^2 is the extension of FO^2 with counting quantifiers. Precisely, we show with Theorem 10 that $<$-inv $\mathrm{C}^2 \subseteq \mathrm{FO}$ when the degree is bounded. Let us already mention that both of these inclusions are strict. This matter is further discussed in the conclusion.

Related work

The question of the order-invariance of an FO^2-sentence has been shown to be decidable in [19]. The expressive power of $<$-inv FO^2 is only mentioned there as a follow up question.

We borrow the dichotomy between rare and frequent neighborhood types when the degree is bounded from [7]. Beyond that, the philosophies of the constructions differ widely: in [7], successor relations are constructed in a very regular way, in order to create as few neighborhood types as possible in the structures with successor. On the other hand, we make sure to realize all the possible types in our construction.

One line of research investigates the expressive power of $<$-inv FO. Let us mention [3] and [8], which prove that $<$-inv FO has the same expressive power as FO respectively on trees and on hollow trees. The present paper focuses on a weaker logic, but in a broader setting. Furthermore, while the techniques used in these papers involve the construction of several intermediate orders, making only localized changes at each step (in the fashion of [9], in which it is proved that $<$-inv FO retains the locality of FO), we equip in one go each of our structures with a linear order.

Although not directly related to our inquiries, [18] is also concerned with the expressive power of FO^2 on ordered structures. This paper establishes a strict hierarchy, based on the quantifier rank and quantifier alternation, among properties definable in FO^2 on words.

2 Preliminaries

We use the standard definition of first-order logic $\mathrm{FO}(\Sigma)$ with equality (written FO when Σ is clear from the context) on a finite signature Σ composed of relation and constant symbols. By FO^2 we denote the fragment of FO in which the only two variables are x and y.

Structures are denoted by a name that starts with (or consists of) a calligraphic upper-case letter, while their universes are denoted by the same name starting with a standard upper-case letter instead of the calligraphic one; for instance, Ex is the universe of the structure $\mathcal{E}x$. Throughout this paper, we consider only finite structures.

A sentence $\varphi \in \mathrm{FO}^2(\Sigma \cup \{<\})$, where $<$ is a binary relation symbol not belonging to Σ, is said to be **order-invariant** if for every finite Σ-structure $\mathcal{A}$, and every pair of strict linear orders $<_0$ and $<_1$ on A, $(\mathcal{A}, <_0) \models \varphi$ iff $(\mathcal{A}, <_1) \models \varphi$. It is then convenient to omit the interpretation for the symbol $<$, and to write $\mathcal{A} \models \varphi$ iff $(\mathcal{A}, <_0) \models \varphi$ for any (or, equivalently, every) linear order $<_0$. The set of order-invariant sentences using two variables is denoted $<$-inv FO^2.

[a] We have recently become aware [2] of an upcoming result leading us to believe that, in general, $<$-inv FO^2 can capture properties that are not FO-definable. More precisely, this result states the existence of two classes of structures (of unbounded degree) $\mathcal{C}' \subseteq \mathcal{C}$ such that $\mathcal{C}'$ is not definable in FO, but can be defined in $\mathcal{C}$ by an FO^2-sentence using a linear order, which is order-invariant on all structures of $\mathcal{C}$. Although this does not exactly imply $<$-inv $\mathrm{FO}^2 \not\subseteq \mathrm{FO}$ according to our definition of invariance (since said sentence is not invariant on all finite structures), it leads us to believe that $<$-inv FO^2 can express properties beyond FO's reach.

Let $\mathcal{L}, \mathcal{L}'$ be two logics defined over the same signature, and $\mathcal{C}$ be a class of finite structures on this signature. We say that a property $\mathcal{P} \subseteq \mathcal{C}$ is **definable** (or expressible) in $\mathcal{L}$ if there exists an $\mathcal{L}$-sentence φ such that $\mathcal{P} = \{\mathcal{A} \in \mathcal{C} : \mathcal{A} \models \varphi\}$. We say that $\mathcal{L} \subseteq \mathcal{L}'$ on $\mathcal{C}$ if every property on $\mathcal{C}$ definable in $\mathcal{L}$ is also definable in $\mathcal{L}'$.

It is quite clear that $\mathrm{FO}^2 \subseteq <\text{-inv FO}^2$: any sentence which does not make use of the order is indeed order-invariant. Furthermore, this inclusion is strict. For instance, over the empty signature, the property of having at least three elements is not definable in FO^2 (this can easily be seen with the tools presented in Section 5.1), but can be expressed in $<$-inv FO^2, for instance via the formula $\exists x\ \exists y\ (x < y\ \wedge\ (\exists x\ y < x))$.

The **quantifier rank** of a formula is the maximal number of quantifiers in a branch of its syntactic tree. Given two Σ-structures $\mathcal{A}_0$ and $\mathcal{A}_1$, and $\mathcal{L}$ being one of $\mathrm{FO}, \mathrm{FO}^2$ and $<$-inv FO^2, we write $\mathcal{A}_0 \equiv_k^{\mathcal{L}} \mathcal{A}_1$ if $\mathcal{A}_0$ and $\mathcal{A}_1$ satisfy the same $\mathcal{L}$-sentences of quantifier rank at most k. In this case, we say that $\mathcal{A}_0$ and $\mathcal{A}_1$ are **$\mathcal{L}$-similar at depth k**.

We write $\mathcal{A}_0 \simeq \mathcal{A}_1$ if $\mathcal{A}_0$ and $\mathcal{A}_1$ are isomorphic.

Atomic types

Let a be an element of a structure $\mathcal{A}$. The **atomic type** $\mathrm{tp}^0_{\mathcal{A}}(a)$ of a in $\mathcal{A}$ is the set of atomic formulas φ with at most one free variable x such that $\mathcal{A}, x \mapsto a \models \varphi$.

We define similarly the atomic type $\mathrm{tp}^0_{\mathcal{A}}(a,b)$ of a pair (a,b) of elements of $\mathcal{A}$ as the set of atomic formulas φ with free variables in $\{x,y\}$ such that $\mathcal{A}, x \mapsto a, y \mapsto b \models \varphi$.

Given a linearly ordered Σ-structure $(\mathcal{A}, <)$, $\mathrm{tp}^0_{(\mathcal{A},<)}(a,b)$ can be divided into $\mathrm{tp}^0_{<}(a,b)$ and $\mathrm{tp}^0_{\mathcal{A}}(a,b)$, where $\mathrm{tp}^0_{<}(a,b)$ is one of $\{x < y\}, \{x > y\}$ and $\{x = y\}$.

Gaifman graphs

The **Gaifman graph** $\mathcal{G}_{\mathcal{A}}$ of a structure $\mathcal{A}$ is defined as (A, E) where $(a,b) \in E$ iff a and b appear in the same tuple of a relation of $\mathcal{A}$. Notice that if a graph is seen as a structure on the signature consisting of a single binary relation symbol, its Gaifman graph is none other than the unoriented version of itself.

By $\mathrm{dist}_{\mathcal{A}}(a,b)$, we denote the distance between a and b in $\mathcal{G}_{\mathcal{A}}$. For $B \subseteq A$, we note $N_{\mathcal{A}}(B)$ the set of elements at distance exactly 1 from B in $\mathcal{G}_{\mathcal{A}}$. In particular, $B \cap N_{\mathcal{A}}(B) = \emptyset$.

The **degree** of $\mathcal{A}$ is the maximal degree of its Gaifman graph, and a class $\mathcal{C}$ of Σ-structures is said to have **bounded degree** if there exists some $d \in \mathbb{N}$ such that the degree of every $\mathcal{A} \in \mathcal{C}$ is at most d.

3 Main result

We are now able to state the main result of this article. Remember that $<$-inv FO^2 allows us to express properties that are beyond the scope of plain FO^2. We give an upper bound to its expressive power, when the degree is bounded:

▶ **Theorem 1.** *Let $\mathcal{C}$ be a class of structures of bounded degree.*
Then $<$-inv $FO^2 \subseteq FO$ on $\mathcal{C}$.

For the remainder of this paper, we fix a signature Σ, an integer d and a class $\mathcal{C}$ of Σ-structures of degree at most d.

Let us now show the skeleton of our proof. The technical part of the proof will be the focus of Sections 4 and 5. Our general strategy is to show the existence of a function $f : \mathbb{N} \to \mathbb{N}$ such that every formula $\varphi \in <\text{-inv FO}^2$ of quantifier rank k is equivalent on $\mathcal{C}$ (i.e. satisfied by the same structures of $\mathcal{C}$) to an FO-formula ψ of quantifier rank at most $f(k)$.

To prove this, it is enough to show that for any two structures $\mathcal{A}_0, \mathcal{A}_1 \in \mathcal{C}$ such that $\mathcal{A}_0 \equiv^{\mathrm{FO}}_{f(k)} \mathcal{A}_1$, we have $\mathcal{A}_0 \equiv^{<\text{-inv FO}^2}_k \mathcal{A}_1$. Indeed, the class of structures satisfying a formula $\varphi \in <$-inv FO^2 of quantifier rank k is a union of equivalence classes for the equivalence relation $\equiv^{<\text{-inv FO}^2}_k$, whose intersection with $\mathcal{C}$ is in turn the intersection of $\mathcal{C}$ with a union of equivalence classes for $\equiv^{\mathrm{FO}}_{f(k)}$. It is folklore (see, e.g., [13]) that the equivalence relation $\equiv^{\mathrm{FO}}_{f(k)}$ has finite index, and that each of its equivalence classes is definable by an FO-sentence of quantifier rank $f(k)$. Then ψ is just the finite disjunction of these FO-sentences.

In order to show that $\mathcal{A}_0 \equiv^{<\text{-inv FO}^2}_k \mathcal{A}_1$, we will construct in Section 4 two particular orders $<_0, <_1$ on these respective structures, and we will prove in Section 5 that

$$(\mathcal{A}_0, <_0) \equiv^{\mathrm{FO}^2}_k (\mathcal{A}_1, <_1)\,. \tag{1}$$

This concludes the proof, since any sentence $\theta \in <$-inv FO^2 with quantifier rank at most k holds in $\mathcal{A}_0$ iff it holds in $(\mathcal{A}_0, <_0)$ (by definition of order-invariance), iff it holds in $(\mathcal{A}_1, <_1)$ (by (1)), iff it holds in $\mathcal{A}_1$.

4 Constructing linear orders on $\mathcal{A}_0$ and $\mathcal{A}_1$

Recall from Section 3 that our goal is to find a function f such that, given two structures $\mathcal{A}_0, \mathcal{A}_1$ in $\mathcal{C}$ such that

$$\mathcal{A}_0 \equiv^{\mathrm{FO}}_{f(k)} \mathcal{A}_1\,, \tag{2}$$

we are able to construct two linear orders $<_0, <_1$ such that $(\mathcal{A}_0, <_0) \equiv^{\mathrm{FO}^2}_k (\mathcal{A}_1, <_1)$.

In this section, we define f and we detail the construction of such orders. The proof of $<$-inv FO-similarity between $(\mathcal{A}_0, <_0)$ and $(\mathcal{A}_1, <_1)$ will be the focus of Section 5.

Let us now explain how we define f. For that, we need to introduce the notion of neighborhood and neighborhood type. These notions are defined in Section 4.1. We then explain in Section 4.2 how to divide neighborhood types into rare ones and frequent ones. Finally, the details of the construction are given in Section 4.3.

4.1 Neighborhoods

Let us now define the notion of neighborhood of an element in a structure.

Let c be a new constant symbol, and let $\mathcal{A} \in \mathcal{C}$. For $k \in \mathbb{N}$ and $a \in A$, the (pointed) ***k*-neighborhood** $\mathcal{N}^k_{\mathcal{A}}(a)$ of a in $\mathcal{A}$ is the $(\Sigma \cup \{c\})$-structure whose restriction to the vocabulary Σ is the substructure of $\mathcal{A}$ induced by the set $N^k_{\mathcal{A}}(a) = \{b \in A : \mathrm{dist}_{\mathcal{A}}(a,b) \leq k\}$, and where c is interpreted as a. In other words, it consists of all the elements at distance at most k from a in $\mathcal{A}$, together with the relations they share in $\mathcal{A}$; the center a being marked by the constant c. We sometimes refer to $N^k_{\mathcal{A}}(a)$ as the k-neighborhood of a in $\mathcal{A}$ as well, but the context will always make clear whether we refer to the whole substructure or only its domain. The ***k*-neighborhood type** $\tau = \text{neigh-tp}^k_{\mathcal{A}}(a)$ of a in $\mathcal{A}$ is the isomorphism class of its k-neighborhood. We say that τ is a k-neighborhood type over Σ, and that a is an **occurrence** of τ. We denote by $|\mathcal{A}|_\tau$ the number of occurrences of τ in $\mathcal{A}$, and we write $[\![\mathcal{A}_0]\!]_k =^t [\![\mathcal{A}_1]\!]_k$ to mean that for every k-neighborhood type τ, $|\mathcal{A}_0|_\tau$ and $|\mathcal{A}_1|_\tau$ are either equal, or both larger than t.

Let $\textsc{NeighType}^d_k$ denote the set of k-neighborhood types over Σ occurring in structures of degree at most d. Note that $\textsc{NeighType}^d_k$ is a finite set.

The interest of this notion resides in the fact that when the degree is bounded, FO is exactly able to count the number of occurrences of neighborhood types up to some threshold [4]. We will only use one direction of this characterization, namely:

▶ **Proposition 2.** *For all integers k and t, there exists some $\hat{f}(k,t) \in \mathbb{N}$ (which also depends on the bound d on the degree of structures in $\mathcal{C}$) such that for all structures $\mathcal{A}_0, \mathcal{A}_1 \in \mathcal{C}$,*

$$\mathcal{A}_0 \equiv^{FO}_{\hat{f}(k,t)} \mathcal{A}_1 \quad \to \quad [\![\mathcal{A}_0]\!]_k =^t [\![\mathcal{A}_1]\!]_k \,.$$

We now exhibit a function $\Theta : \mathbb{N} \to \mathbb{N}$ such that, if $[\![\mathcal{A}_0]\!]_k =^{\Theta(k)} [\![\mathcal{A}_1]\!]_k$, then one can construct $<_0, <_1$ satisfying (1). Proposition 2 then ensures that $f : k \mapsto \hat{f}(k, \Theta(k))$ fits the bill. Let us now explain how the function Θ is chosen.

4.2 Frequency of a neighborhood type

Let us denote $|\text{NeighType}^d_k|$ as N.

Recall that every $\mathcal{A} \in \mathcal{C}$ has degree at most d. What this means is that if we consider the set $\text{Freq}[\mathcal{A}]_k$ of k-neighborhood types that have enough occurrences in $\mathcal{A}$ (where "enough" will be given a precise meaning later on), each type in $\text{Freq}[\mathcal{A}]_k$ must have many occurrences that are scattered across $\mathcal{A}$. Not only that, but we can also make sure that such occurrences are far from all the occurrences of every k-neighborhood type not in $\text{Freq}[\mathcal{A}]_k$, which by definition have few occurrences in $\mathcal{A}$. Since the degree is bounded, N is bounded too, which prevents our distinction (which will be formalized later on) between rare neighborhood types and frequent neighborhood types from being circular.

Such a dichotomy is introduced and detailed in [7]; we simply adapt this construction to our needs. In the remainder of this section, we describe this construction at a high level, and leave the technical details (such as the exact bounds) to the reader.

The proof of the following lemma (in the vein of [1]) is straightforward, and relies on the degree boundedness hypothesis. Intuitively, Lemma 3 states that when the degree is bounded, it is not possible for all the elements of large sets to be concentrated in one corner of the structure, thus making it possible to pick elements in each set that are scattered across the structure.

▶ **Lemma 3.** *Given three integers m, δ, s, there exists a threshold $g(m,\delta,s) \in \mathbb{N}$ such that for all $\mathcal{A} \in \mathcal{C}$, all $B \subseteq A$ of size at most s, and all subsets $C_1, \cdots, C_n \subseteq A$ (with $n \leq N$) of size at least $g(m,\delta,s)$, it is possible to find elements $c^1_j, \cdots, c^m_j \in C_j$ for all $j \in \{1, \cdots, n\}$, such that for all $j, j' \in \{1, \cdots, n\}$ and $i, i' \in \{1, \cdots, m\}$, $dist_{\mathcal{A}}(c^i_j, B) > \delta$ and $dist_{\mathcal{A}}(c^i_j, c^{i'}_{j'}) > \delta$ if $(j,i) \neq (j',i')$.*

Note that the N is this lemma could be replaced by any constant.

Our goal is, given a structure $\mathcal{A} \in \mathcal{C}$, to partition the k-neighborhood types into two classes: the frequent types, and the rare types. The property we wish to ensure is that there exist in $\mathcal{A}$ some number m (which will be made precise later on, but only depends on k) of occurrences of each one of the frequent k-neighborhood types which are both

- at distance greater than δ (which, as for m, is a function of k and will be fixed in the following) from one another, and
- at distance greater than δ from every occurrence of a rare k-neighborhood type.

To establish this property, we would like to use Lemma 3, with s being the total number of occurrences of all the rare k-neighborhood types, and $C_1, \cdots, C_n$ being the sets of occurrences of the n distinct frequent k-neighborhood types.

The number N of different k-neighborhood types of degree at most d is bounded by a function of k (as well as Σ and d, which are fixed). Hence, we can proceed according to the following (terminating) algorithm to make the distinction between frequent and rare types:

1. First, let us mark every k-neighborhood type as frequent.
2. Among the types which are currently marked as frequent, let τ be one with the smallest number of occurrences in $\mathcal{A}$.
3. If $|\mathcal{A}|_\tau$ is at least $g(m, \delta, s)$ (g being the function from Lemma 3) where s is the total number of occurrences of all the k-neighborhood types which are currently marked as rare, then we are done and the marking frequent/rare is final. Otherwise, mark τ as rare, and go back to step 2 if there remains at least one frequent k-neighborhood type.

Notice that we can go at most N times through step 2, where N depends only on k. Furthermore, each time we add a type to the set of rare k-neighborhood types, we have the guarantee that this type has few occurrences (namely, less than $g(m, \delta, s)$, where s can be bounded by a function of k).

It is thus apparent that the threshold t such that a k-neighborhood type τ is frequent in $\mathcal{A}$ iff $|\mathcal{A}|_\tau \geq t$ can be bounded by some T depending only on k - importantly, T is the same for all structures of $\mathcal{C}$.

Let us now make the above more formal. For $t \in \mathbb{N}$ and $\mathcal{A} \in \mathcal{C}$, let $\textsc{Freq}[\mathcal{A}]_k^{\geq t} \subseteq \textsc{NeighType}_k^d$ denote the set of k-neighborhood types which have at least t occurrences in $\mathcal{A}$. By applying the procedure presented above, we derive the following lemma:

▶ **Lemma 4.** *Let $k, m, \delta \in \mathbb{N}$. There exists $T \in \mathbb{N}$ such that for every $\mathcal{A} \in \mathcal{C}$, there exists some $t \leq T$ such that*

$$t \geq g(m, \delta, \sum_{\tau \notin \textsc{Freq}[\mathcal{A}]_k^{\geq t}} |\mathcal{A}|_\tau)\,.$$

Let $\textsc{Freq}[\mathcal{A}]_k := \textsc{Freq}[\mathcal{A}]_k^{\geq t}$ for the smallest threshold t given in Lemma 4. Some k-neighborhood type $\tau \in \textsc{NeighType}_k^d$ is said to be **frequent** in $\mathcal{A} \in \mathcal{C}$ if it belongs to $\textsc{Freq}[\mathcal{A}]_k$; that is, if $|\mathcal{A}|_\tau \geq t$. Otherwise, τ is said to be **rare**. With the definition of g in mind, Lemma 4 can then be reformulated as follows: in every structure $\mathcal{A} \in \mathcal{C}$, one can find m occurrences of each frequent k-neighborhood type which are at distance greater than δ from one another and from the set of occurrences of every rare k-neighborhood type.

All that remains is for us to give a value (depending only on k) to the integers m and δ: let $M := \max\{|\tau| : \tau \in \textsc{NeighType}_k^d\}$ (M indeed exists, and is a function of k - recall that the signature Σ and the degree d are assumed to be fixed). Let us consider

$$m := 2 \cdot (k+1) \cdot M! \qquad \text{and} \qquad \delta := 4k\,. \tag{3}$$

We then define $\Theta(k)$ as the integer T provided by Lemma 4 for these values of m and δ. The threshold $\Theta(k)$ indeed only depends on k. Finally, notice that if $[\![\mathcal{A}_0]\!]_k =^{\Theta(k)} [\![\mathcal{A}_1]\!]_k$, then $\textsc{Freq}[\mathcal{A}_0]_k = \textsc{Freq}[\mathcal{A}_1]_k$.

As discussed in Section 4.1, there exists a function f such that $\mathcal{A}_0 \equiv^{\mathrm{FO}}_{f(k)} \mathcal{A}_1$ entails $[\![\mathcal{A}_0]\!]_k =^{\Theta(k)} [\![\mathcal{A}_1]\!]_k$. We also make sure that $f(k) \geq \Theta(k) \cdot N + 1$ for every k.

Let us now consider $\mathcal{A}_0, \mathcal{A}_1 \in \mathcal{C}$ such that $\mathcal{A}_0 \equiv^{\mathrm{FO}}_{f(k)} \mathcal{A}_1$ for such an f. If $\textsc{Freq}[\mathcal{A}_0]_k = \emptyset$, then $|\mathcal{A}_0| \leq \Theta(k) \cdot N$. This guarantees that $\mathcal{A}_0 \simeq \mathcal{A}_1$, and in particular that $\mathcal{A}_0 \equiv_k^{<\text{-inv FO}^2} \mathcal{A}_1$. From now on, we suppose that there is at least one frequent k-neighborhood type.

The construction of two linear orders $<_0$ and $<_1$ satisfying $(\mathcal{A}_0, <_0) \equiv_k^{\mathrm{FO}^2} (\mathcal{A}_1, <_1)$ is the object of Section 4.3.

4.3 Construction of $<_0$ and $<_1$

This section is dedicated to the definition of two linear orders $<_0, <_1$ on $\mathcal{A}_0, \mathcal{A}_1 \in \mathcal{C}$. We then prove in Section 5 that $(\mathcal{A}_0, <_0)$ and $(\mathcal{A}_1, <_1)$ are FO^2-similar at depth k.

Recall that by hypothesis, $\mathcal{A}_0$ and $\mathcal{A}_1$ are FO-similar at depth $f(k)$, which entails that they have the same number of occurrences of each $\tau \in \text{NEIGHTYPE}_k^d$ up to a threshold $\Theta(k)$.

To construct our two linear orders, we need to define the notion of k-environment: given $\mathcal{A} \in \mathcal{C}$, a linear order $<$ on A, $k \in \mathbb{N}$ and an element $a \in A$, we define the ***k*-environment** $\mathcal{E}\text{nv}^k_{(\mathcal{A},<)}(a)$ **of a in $(\mathcal{A},<)$** as the restriction of $(\mathcal{A},<)$ to the k-neighborhood of a in $\mathcal{A}$, where a is the interpretation of the constant symbol c. Note that the order is not taken into account when determining the domain of the substructure (it would otherwise be A, given that any two distinct elements are adjacent for $<$). The ***k*-environment type** $\text{env-tp}^k_{(\mathcal{A},<)}(a)$ is the isomorphism class of $\mathcal{E}\text{nv}^k_{(\mathcal{A},<)}(a)$. In other words, $\text{env-tp}^k_{(\mathcal{A},<)}(a)$ contains the information of $\mathcal{N}^k_{\mathcal{A}}(a)$ together with the order of its elements in $(\mathcal{A},<)$.

Given $\tau \in \text{NEIGHTYPE}_k^d$, we define $\text{ENV}(\tau)$ as the set of k-environment types whose underlying k-neighborhood type is τ.

For $i \in \{0,1\}$, we aim to partition A_i into $2(2k+1)+2$ segments:

$$A_i = X_i \cup \bigcup_{j=0}^{2k} (L_i^j \cup R_i^j) \cup M_i\,.$$

Once we have set a linear order on each segment, the linear order $<_i$ on A_i will result from the concatenation of the orders on the segments as follows:

$$(A_i, <_i) := X_i \cdot L_i^0 \cdot L_i^1 \cdots L_i^{2k} \cdot M_i \cdot R_i^{2k} \cdots R_i^1 \cdot R_i^0\,.$$

Each segment L_i^j, for $j \in \{0,\cdots,2k\}$ is itself decomposed into two segments $NL_i^j \cdot UL_i^j$. The UL_i^j for $j \in \{k+1,\cdots,2k\}$ will be empty; they are defined solely in order to keep the notations uniform. The 'N' stands for "neighbor" and the 'U' for "universal", for reasons that will soon become apparent. Symmetrically, each R_i^j is decomposed into $UR_i^j \cdot NR_i^j$, with empty UR_i^j as soon as $j \geq k+1$.

For $i \in \{0,1\}$ and $r \in \{0,\cdots,2k\}$, we define S_i^r as

$$S_i^r := X_i \cup \bigcup_{j=0}^{r} (L_i^j \cup R_i^j)\,.$$

Let us now explain how the segments are constructed in $\mathcal{A}_0$; see Figure 1 for an illustration.

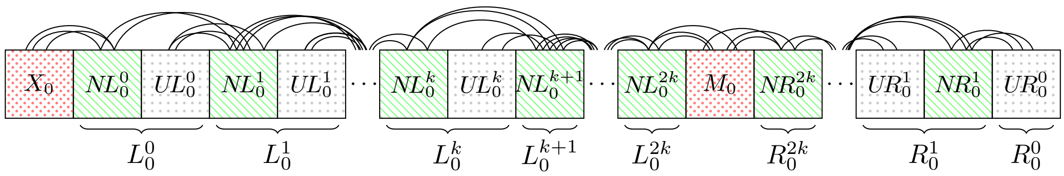

Figure 1 The black curvy edges represent the edges between elements belonging to different segments. Edges between elements of the same segment are not represented here. The order $<_0$ grows from the left to the right.

For every $\tau \in \text{FREQ}[\mathcal{A}_0]_k$, let $\tau_1,\cdots,\tau_{|\text{ENV}(\tau)|}$ be an enumeration of $\text{ENV}(\tau)$. Recall that we defined M in Section 4.2 as $\max\{|\tau| : \tau \in \text{NEIGHTYPE}_k^d\}$. Thus, we have $|\text{ENV}(\tau)| \leq M!$ for every $\tau \in \text{NEIGHTYPE}_k^d$.

In particular, by definition of frequency, and by choice of m and δ in (3), Lemma 4 ensures that we are able to pick, for every $\tau \in \text{FREQ}[\mathcal{A}_0]_k$, every $l \in \{1,\cdots,|\text{ENV}(\tau)|\}$ and every $j \in \{0,\cdots,k\}$, two elements $a[\tau_l]_L^j$ and $a[\tau_l]_R^j$ which have τ as k-neighborhood type in $\mathcal{A}_0$, such that all the $a[\tau_l]_*^j$ are at distance at least $4k+1$ from each other and from any occurrence of a rare k-neighborhood type in $\mathcal{A}_0$.

Construction of X_0 and NL_0^0

We start with the set X_0, which contains all the occurrences of rare k-neighborhood types, together with their k-neighborhoods.
Formally, the domain of X_0 is $\bigcup_{a\in A_0:\ \text{neigh-tp}^k_{\mathcal{A}_0}(a)\notin \textsc{Freq}[\mathcal{A}_0]_k} N^k_{\mathcal{A}_0}(a)\,.$
We set $NL_0^0 := N_{\mathcal{A}_0}(X_0)$ (the set of neighbors of elements of X_0), and define the order $<_0$ on X_0 and on NL_0^0 in an arbitrary way.

Construction of UL_0^j

If $k < j \leq 2k$, then we set $UL_0^j := \emptyset$. Otherwise, for $j \in \{0, \cdots, k\}$, once we have constructed $L_0^0, \cdots, L_0^{j-1}$ and NL_0^j, we construct UL_0^j as follows.
The elements of UL_0^j are $\bigcup_{\tau\in\textsc{Freq}[\mathcal{A}_0]_k} \bigcup_{l=1}^{|\textsc{Env}(\tau)|} N^k_{\mathcal{A}_0}(a[\tau_l]^j_L)\,.$
Note that UL_0^j does not intersect the previously constructed segments, by choice of the $a[\tau_l]^j_L$ and of $\delta = 4k$ in (3). Furthermore, the $N^k_{\mathcal{A}_0}(a[\tau_l]^j_L)$ are pairwise disjoint, hence we can fix $<_0$ freely and independently on each of them. Unsurprisingly, we order each $N^k_{\mathcal{A}_0}(a[\tau_l]^j_L)$ so that $\text{env-tp}^k_{(\mathcal{A}_0,<_0)}(a[\tau_l]^j_L) = \tau_l$. This is possible because for every $\tau \in \textsc{Freq}[\mathcal{A}_0]_k$ and each l, $\text{neigh-tp}^k_{\mathcal{A}_0}(a[\tau_l]^j_L) = \tau$ by choice of $a[\tau_l]^j_L$.
Once each $N^k_{\mathcal{A}_0}(a[\tau_l]^j_L)$ is ordered according to τ_l, the linear order $<_0$ on UL_0^j can be completed in an arbitrary way. Note that every possible k-environment type extending a frequent k-neighborhood type in $\mathcal{A}_0$ occurs in each UL_0^j. The UL_0^j are *universal* in that sense.

Construction of NL_0^j

Now, let us see how the NL_0^j are constructed. For $j \in \{1, \cdots, 2k\}$, suppose that we have constructed $L_0^0, \cdots, L_0^{j-1}$. The domain of NL_0^j consists of all the neighbors (in $\mathcal{A}_0$) of the elements of L_0^{j-1} not already belonging to the construction so far. Formally, $N_{\mathcal{A}_0}(L_0^{j-1}) \setminus \big(X_0 \cup \bigcup_{m=0}^{j-2} L_0^m\big)$.
The order $<_0$ on NL_0^j is chosen arbitrarily.

Construction of R_0^j

We construct similarly the R_0^j, for $j \in \{0, \cdots, 2k\}$, starting with $NR_0^0 := \emptyset$, then UR_0^0 which contains each $a[\tau_l]^0_R$ together with its k-neighborhood in $\mathcal{A}_0$ ordered according to τ_l, then $NR_0^1 := N_{\mathcal{A}_0}(R_0^0)$, then UR_0^1, etc.
Note that the $a[\tau_l]^j_R$ have been chosen so that they are far enough in $\mathcal{A}_0$ from all the segments that have been constructed so far, allowing us once more to order their k-neighborhood in $\mathcal{A}_0$ as we see fit.

Construction of M_0

M_0 contains all the elements of A_0 besides those already belonging to S_0^{2k}. The order $<_0$ chosen on M_0 is arbitrary.

Transfer on $\mathcal{A}_1$

Suppose that we have constructed S_0^{2k}. We can make sure, retrospectively, that the index $f(k)$ in (2) is large enough so that there exists a set $S \subseteq A_1$ so that $\mathcal{A}_0|_{S_0^{2k}\cup N_{\mathcal{A}_0}(S_0^{2k})} \simeq \mathcal{A}_1|_S$ (this is ensured as long as $f(k) \geq |S_0^{2k} \cup N_{\mathcal{A}_0}(S_0^{2k})| + 1$, which can be bounded by a function of k, independent of $\mathcal{A}_0$ and $\mathcal{A}_1$).

Let $\varphi_0 : \mathcal{A}_0|_{S_0^{2k}} \to \mathcal{A}_1|_{S'}$ be the restriction to S_0^{2k} of said isomorphism, and let φ_1 be its converse. By construction, the k-neighborhood of every $a \in S_0^k$ is included in S_0^{2k}; hence every such a has the same k-neighborhood type in $\mathcal{A}_0$ as has $\varphi_0(a)$ in $\mathcal{A}_1$.
We transfer alongside φ_0 all the segments, with their order, from $(\mathcal{A}_0, <_0)$ to $\mathcal{A}_1$, thus defining $X_1, NL_1^j, UL_1^j, \cdots$ as the respective images by φ_0 of $X_0, NL_0^j, UL_0^j, \cdots$, and define M_1 as the counterpart to M_0. Note that the properties concerning neighborhood are transferred; e.g. all the neighbors of an element in L_1^j, $1 \leq j < 2k$, belong to $L_1^{j-1} \cup L_1^j \cup L_1^{j+1}$.

By construction, we get the following lemma:

▶ **Lemma 5.** *For each $a \in S_0^k$, we have env-tp$^k_{(\mathcal{A}_0,<_0)}(a) =$ env-tp$^k_{(\mathcal{A}_1,<_1)}(\varphi_0(a))$.*

Lemma 5 has two immediate consequences:

- The set X_1 contains the occurrences in $\mathcal{A}_1$ of all the rare k-neighborhood types (just forget about the order on the k-environments, and remember that $\mathcal{A}_0$ and $\mathcal{A}_1$ have the same number of occurrences of each rare k-neighborhood type).
- All the universal segments UL_1^j and UR_1^j, for $0 \leq j \leq k$, contain at least one occurrence of each environment in $\textsc{Env}(\tau)$, for each $\tau \in \textsc{Freq}[\mathcal{A}_0]_k$.

Our construction also guarantees the following result:

▶ **Lemma 6.** *For each $a, b \in S_0^k$, we have tp$^0_{(\mathcal{A}_0,<_0)}(a,b) =$ tp$^0_{(\mathcal{A}_1,<_1)}(\varphi_0(a), \varphi_0(b))$.*

In particular, for $a = b \in S_0^k$, we have $\mathrm{tp}^0_{(\mathcal{A}_0,<_0)}(a) = \mathrm{tp}^0_{(\mathcal{A}_1,<_1)}(\varphi_0(a))$.

5 Proof of the FO^2-similarity of $(\mathcal{A}_0, <_0)$ and $(\mathcal{A}_1, <_1)$

In this section, we aim to show the following result:

▶ **Proposition 7.** *We have that $(\mathcal{A}_0, <_0) \equiv_k^{FO^2} (\mathcal{A}_1, <_1)$.*

5.1 The two-pebble Ehrenfeucht-Fraïssé game

To establish Proposition 7, we use Ehrenfeucht-Fraïssé games with two pebbles. These games have been introduced by Immerman and Kozen [11]. Let us adapt their definition to our context.

The **k-round two-pebble Ehrenfeucht-Fraïssé game on $(\mathcal{A}_0, <_0)$ and $(\mathcal{A}_1, <_1)$** is played by two players: the spoiler and the duplicator. The spoiler tries to expose differences between the two structures, while the duplicator tries to establish their indistinguishability.

There are two pebbles associated with each structure: p_0^x and p_0^y on $(\mathcal{A}_0, <_0)$, and p_1^x and p_1^y on $(\mathcal{A}_1, <_1)$. Formally, these pebbles can be seen as the interpretations in each structure of two new constant symbols, but it will be convenient to see them as moving pieces.

At the start of the game, the duplicator places p_0^x and p_0^y on elements of $(\mathcal{A}_0, <_0)$, and p_1^x and p_1^y on elements of $(\mathcal{A}_1, <_1)$. The spoiler wins if the duplicator is unable to ensure that $\mathrm{tp}^0_{(\mathcal{A}_0,<_0)}(p_0^x, p_0^y) = \mathrm{tp}^0_{(\mathcal{A}_1,<_1)}(p_1^x, p_1^y)$. Otherwise, the proper game starts. Note that in the usual definition of the starting position, the pebbles are not on the board; however, it will be convenient to have them placed in order to uniformize our invariant. This change is not profound and does not affect the properties of the game.

For each of the k rounds, the spoiler starts by choosing a structure and a pebble in this structure, and places this pebble on a element of the chosen structure. In turn, the duplicator must place the corresponding pebble in the other structure on an element of that structure. The spoiler wins at once if $\mathrm{tp}^0_{(\mathcal{A}_0,<_0)}(p_0^x, p_0^y) \neq \mathrm{tp}^0_{(\mathcal{A}_1,<_1)}(p_1^x, p_1^y)$. Otherwise, another round is played. If the spoiler has not won after k rounds, then the duplicator wins.

The main interest of these games is that they capture the expressive power of FO^2 [11]. We will only need the fact that these games are correct:

▶ **Theorem 8.** *If the duplicator has a winning strategy in the k-round two-pebble Ehrenfeucht-Fraïssé game on $(\mathcal{A}_0,<_0)$ and $(\mathcal{A}_1,<_1)$, then $(\mathcal{A}_0,<_0) \equiv_k^{FO^2} (\mathcal{A}_1,<_1)$.*

Thus, in order to prove Proposition 7, we show that the duplicator wins the k-round two-pebble Ehrenfeucht-Fraïssé game on $(\mathcal{A}_0,<_0)$ and $(\mathcal{A}_1,<_1)$. For that, let us show by a decreasing induction on $r = k, \cdots, 0$ that the duplicator can ensure, after $k-r$ rounds, that the three following properties (described below) hold:

$$\forall i \in \{0,1\}, \forall \alpha \in \{x,y\},\ p_i^\alpha \in S_i^r \ \to\ p_{1-i}^\alpha = \varphi_i(p_i^\alpha) \tag{S_r}$$

$$\forall \alpha \in \{x,y\},\ \text{env-tp}^r_{(\mathcal{A}_0,<_0)}(p_0^\alpha) = \text{env-tp}^r_{(\mathcal{A}_1,<_1)}(p_1^\alpha) \tag{E_r}$$

$$\text{tp}^0_{(\mathcal{A}_0,<_0)}(p_0^x,p_0^y) = \text{tp}^0_{(\mathcal{A}_1,<_1)}(p_1^x,p_1^y) \tag{T_r}$$

The first property, (S_r), guarantees that if a pebble is close (in a sense that depends on the number of rounds left in the game) to one of the $<_i$-minimal or $<_i$-maximal elements, the corresponding pebble in the other structure is located at the same position with respect to this $<_i$-extremal element.
As for (E_r), it states that two corresponding pebbles are always placed on elements sharing the same r-environment type. Once again, the satefy distance decreases with each round that goes.
Finally, (T_r) controls that both pebbles have the same relative position (both with respect to the order and the original vocabulary) in the two ordered structures. In particular, the duplicator wins the game if (T_r) is satisfied at the begining of the game, and after each of the k rounds of the game.

5.2 Base case: proofs of (S_k), (E_k) and (T_k)

We start by proving (S_k), (E_k) and (T_k).
At the start of the game, the duplicator places both p_0^x and p_0^y on the $<_0$-minimal element of $(\mathcal{A}_0,<_0)$, and both p_1^x and p_1^y on the $<_1$-minimal element of $(\mathcal{A}_1,<_1)$. In particular,

$$p_1^x = p_1^y = \varphi_0(p_0^x) = \varphi_0(p_0^y)\,.$$

This ensures that (S_k) holds, while (E_k) and (T_k) respectively follow from Lemma 5 and Lemma 6.

5.3 Strategy for the duplicator

We now describe the duplicator's strategy to ensure that (S_r), (E_r) and (T_r) hold no matter how the spoiler plays.

Suppose that we have (S_{r+1}), (E_{r+1}) and (T_{r+1}) for some $0 \le r < k$, after $k-r-1$ rounds of the game. Without loss of generality, we may assume that, in the $(k-r)$-th round of the Ehrenfeucht-Fraïssé game between $(\mathcal{A}_0,<_0)$ and $(\mathcal{A}_1,<_1)$, the spoiler moves p_0^x in $(\mathcal{A}_0,<_0)$. Let us first explain informally the general idea behind the duplicator's strategy.

1. If the spoiler plays around the endpoints (by which we mean the elements that are $<_i$-minimal and maximal), the duplicator has no choice but to play a tit-for-tat strategy, i.e. to respond to the placement of p_i^α near the endpoints by moving p_{1-i}^α on $\varphi_i(p_i^\alpha)$.

If the duplicator does not respond this way, then the spoiler will be able to expose the difference between $(\mathcal{A}_0, <_0)$ and $(\mathcal{A}_1, <_1)$ in the subsequent moves, by forcing the duplicator to play closer and closer to the endpoint, which will prove to be impossible at some point.

On top of that, the occurrences of rare neighborhood types are located in $(\mathcal{A}_i, <_i)$ near the $<_i$-minimal element. If the duplicator does not play according to φ_0 in this area, it will be easy enough for the spoiler to win the game.

The reason we introduced the segments NL_i^j, UL_i^j, NR_i^j and UR_i^j is precisely to bound the area in which the duplicator must implement the tit-for-tat strategy. Indeed, as soon as a pebble is placed in M_i, there is no way for the spoiler to join the endpoints in less than k moves while forcing the duplicator's hand.

The case where the spoiler plays near the endpoints corresponds to Case (I) below, and is detailed in Section 5.4.

2. Next, suppose that the spoiler places a pebble, say p_0^x, next (in $\mathcal{A}_0$) to p_0^y, i.e. such that $p_0^x \in N^1_{\mathcal{A}_0}(p_0^y)$. The duplicator must place p_1^x on an element whose relative position to p_1^y is the same as the relative position of p_0^x with respect to p_0^y. Note that once this is done, the spoiler can change variable, and place p_0^y (or p_1^y, if they decide to play in $(\mathcal{A}_1, <_1)$) in $N^1_{\mathcal{A}_0}(p_0^x)$, thus forcing the duplicator to play near p_1^x. In order to prevent the spoiler from being able, in k such moves, to expose the difference between $(\mathcal{A}_0, <_0)$ and $(\mathcal{A}_1, <_1)$, the duplicator must make sure, with r rounds left, that p_0^x and p_1^x (as well as p_0^y and p_1^y) share the same r-environment in $(\mathcal{A}_0, <_0)$ and $(\mathcal{A}_1, <_1)$. This will guarantee that the duplicator can play along if the spoiler decides to take r moves adjacent (in $\mathcal{A}_i$) to one another.

   The case where the spoiler places a pebble next (in the structure without ordering) to the other pebble is our Case (II), and is treated in Section 5.5.

3. Suppose now that the spoiler's move does not fall under the previous templates. Let us assume that the spoiler plays in $(\mathcal{A}_0, <_0)$, and moves p_0^x to the left of p_0^y (i.e. such that $(\mathcal{A}_0, <_0) \models p_0^x < p_0^y$).

   In order to play according to the remarks from Cases 1 and 2, the duplicator must place p_1^x on an element which shares the same r-environment with p_0^x (where r is the number of rounds left in the game), which is not near the endpoints.

   It must be the case that the k-neighborhood type of p_0^x in $\mathcal{A}_0$ is frequent, since it is not near the endpoints of $(\mathcal{A}_0, <_0)$, hence not in X_0. By construction, every universal segment UL_1^j, for $0 \leq j \leq k$, contains elements of each k-environment type extending any frequent k-neighborhood type. In particular, it contains an element having the same r-environment as p_0^x. The duplicator will place p_1^x on such an element in the leftmost segment UL_1^j which is not considered to be near the endpoints (this notion depends on the number r of rounds left in the game). This is detailed in Cases (III) and (V) (for the symmetrical case where p_0^x is placed to the right of p_0^y) below.

   However, we have to consider a subcase, where p_1^y is itself in the leftmost segment L_1^j which is not near the endpoints. Indeed, in this case, placing p_1^x as discussed may result in p_1^x being to the right of p_1^y, or being in $N^1_{\mathcal{A}_1}(p_1^y)$; either of which being game-losing to the duplicator. However, since p_1^y was considered to be near the endpoints in the previous round of the game, we know that the duplicator played a tit-for-tat strategy at that point, which allows us to replicate the placement of p_0^x according to φ_0. This subcase, as well as the equivalent subcase where the spoiler places p_0^x to the right of p_0^y, are formalized in Cases (IV) and (VI) below.

We are now ready to describe formally the strategy implemented by the duplicator:

(I) If $p_0^x \in S_0^r$, then the duplicator responds by placing p_1^x on $\varphi_0(p_0^x)$.
This corresponds to the tit-for-tat strategy implemented when the spoiler plays near the endpoints, as discussed in Case 1.

(II) Else, if $p_0^x \notin S_0^r$, and $p_0^x \in N^1_{\mathcal{A}_0}(p_0^y)$, then (E_{r+1}) ensures that there exists an isomorphism $\psi : \mathcal{E}\text{nv}^{r+1}_{(\mathcal{A}_0,<_0)}(p_0^y) \to \mathcal{E}\text{nv}^{r+1}_{(\mathcal{A}_1,<_1)}(p_1^y)$. The duplicator responds by placing p_1^x on $\psi(p_0^x)$.
This makes formal the duplicator's response to a move next to the other pebble, as discussed in Case 2 above.

(III) Else suppose that $(\mathcal{A}_0,<_0) \models p_0^x < p_0^y$ and $p_0^y \notin L_0^{r+1}$. Note that $\tau := \text{neigh-tp}^k_{\mathcal{A}_0}(p_0^x) \in \textsc{Freq}[\mathcal{A}_0]_k$, since $p_0^x \notin X_0$. Let $\tau_l := \text{env-tp}^k_{(\mathcal{A}_0,<_0)}(p_0^x)$.
The duplicator responds by placing p_1^x on $\varphi_0(a[\tau_l]_L^{r+1})$.

(IV) Else, if $(\mathcal{A}_0,<_0) \models p_0^x < p_0^y$ and $p_0^y \in L_0^{r+1}$, then the duplicator moves p_1^x on $\varphi_0(p_0^x)$ (by (S_{r+1}), p_0^x indeed belongs to the domain of φ_0).

(V) Else, suppose that $(\mathcal{A}_0,<_0) \models p_0^y < p_0^x$ and $p_0^y \notin R_0^{r+1}$. This case is symmetric to Case (III).
Similarly, the duplicator opts to play p_1^x on $\varphi_0(a[\tau_l]_R^{r+1})$, where $\tau_l := \text{env-tp}^k_{(\mathcal{A}_0,<_0)}(p_0^x)$.

(VI) If we are in none of the cases above, it means that the spoiler has placed p_0^x to the right of p_0^y, and that $p_0^y \in R_0^{r+1}$. This case is symmetric to Case (IV).
Once again, the duplicator places p_1^x on $\varphi_0(p_0^x)$.

It remains to show that this strategy satisfies our invariants: under the inductive assumption that (S_{r+1}), (E_{r+1}) and (T_{r+1}) hold, for some $0 \leq r < k$, we need to show that this strategy ensures that (S_r), (E_r) and (T_r) hold.

We treat each case in its own section: Section 5.4 is devoted to Case (I) while Section 5.5 covers Case (II). Both Cases (III) and (IV) are treated in Section 5.6. Cases (V) and (VI), being their exact symmetric counterparts, are left to the reader.

► Note 9. Note that some properties need no verification. Since p_0^y and p_1^y are left untouched by the players, (S_{r+1}) ensures that half of (S_r) automatically holds, namely that

$$\forall i \in \{0,1\}, \quad p_i^y \in S_i^r \quad \to \quad p_{1-i}^y = \varphi_i(p_i^y)\,.$$

Similarly, the part of (E_r) concerning p_0^y and p_1^y follows from (E_{r+1}):

$$\text{env-tp}^r_{(\mathcal{A}_0,<_0)}(p_0^y) = \text{env-tp}^r_{(\mathcal{A}_1,<_1)}(p_1^y)\,.$$

Lastly, notice that once we have shown that (E_r) holds, it follows that

$$\begin{cases} \text{tp}^0_{\mathcal{A}_0}(p_0^x) = \text{tp}^0_{\mathcal{A}_1}(p_1^x) \\ \text{tp}^0_{\mathcal{A}_0}(p_0^y) = \text{tp}^0_{\mathcal{A}_1}(p_1^y) \end{cases}$$

5.4 When the spoiler plays near the endpoints: Case (I)

In this section, we treat the case where the spoiler places p_0^x near the $<_0$-minimal or $<_0$-maximal element of $(\mathcal{A}_0,<_0)$. Obviously, what "near" means depends on the number of rounds left in the game; the more rounds remain, the more the duplicator must be cautious regarding the possibility for the spoiler to reach an endpoint and potentially expose a difference between $(\mathcal{A}_0,<_0)$ and $(\mathcal{A}_1,<_1)$.

As we have stated in Case (I), with r rounds left, we consider a move on p_0^x by the spoiler to be near the endpoints if it is made in S_0^r. In that case, the duplicator responds along the tit-for-tat strategy, namely by placing p_1^x on $\varphi_0(p_0^x)$.

Let us now prove that this strategy guarantees that (S_r), (E_r) and (T_r) hold. Recall from Note 9 that part of the task is already taken care of.

Proof of (S_r) in Case (I)

We have to show that $\forall i \in \{0,1\},\ p_i^x \in S_i^r \ \rightarrow\ p_{1-i}^x = \varphi_i(p_i^x)$. This follows directly from the duplicator's strategy, since $p_1^x = \varphi_0(p_0^x)$ (thus $p_0^x = \varphi_1(p_1^x)$).

Proof of (E_r) in Case (I)

We need to prove that $\text{env-tp}^r_{(\mathcal{A}_0,<_0)}(p_0^x) = \text{env-tp}^r_{(\mathcal{A}_1,<_1)}(p_1^x)$, which is a consequence of Lemma 5 given that $p_1^x = \varphi_0(p_0^x)$ and $r < k$.

Proof of (T_r) in Case (I)

First, suppose that $p_0^y \in S_0^{r+1}$. By (S_{r+1}), we know that $p_1^y = \varphi_0(p_0^y)$. Thus, Lemma 6 allows us to conclude that $\text{tp}^0_{(\mathcal{A}_0,<_0)}(p_0^x, p_0^y) = \text{tp}^0_{(\mathcal{A}_1,<_1)}(p_1^x, p_1^y)$.

Otherwise, $p_0^y \notin S_0^{r+1}$ and (S_{r+1}) entails that $p_1^y \notin S_1^{r+1}$.
We have two points to establish:

$$\text{tp}^0_{\mathcal{A}_0}(p_0^x, p_0^y) = \text{tp}^0_{\mathcal{A}_1}(p_1^x, p_1^y) \tag{4}$$
$$\text{tp}^0_{<_0}(p_0^x, p_0^y) = \text{tp}^0_{<_1}(p_1^x, p_1^y) \tag{5}$$

Notice that

$$\begin{cases} \text{tp}^0_{\mathcal{A}_0}(p_0^x, p_0^y) = \text{tp}^0_{\mathcal{A}_0}(p_0^x) \cup \text{tp}^0_{\mathcal{A}_0}(p_0^y) \\ \text{tp}^0_{\mathcal{A}_1}(p_1^x, p_1^y) = \text{tp}^0_{\mathcal{A}_1}(p_1^x) \cup \text{tp}^0_{\mathcal{A}_1}(p_1^y) \end{cases}$$

This is because, by construction, the neighbors in $\mathcal{A}_i$ of an element of S_i^r all belong to S_i^{r+1}.
Equation (4) follows from this remark and Note 9.
As for Equation (5), either

$$p_0^x \in X_0 \cup \bigcup_{0 \le j \le r} L_0^j \quad \text{and} \quad p_1^x \in X_1 \cup \bigcup_{0 \le j \le r} L_1^j\,,$$

in which case $\text{tp}^0_{<_0}(p_0^x, p_0^y) = \{x < y\} = \text{tp}^0_{<_1}(p_1^x, p_1^y)$, or

$$p_0^x \in \bigcup_{0 \le j \le r} R_0^j \quad \text{and} \quad p_1^x \in \bigcup_{0 \le j \le r} R_1^j\,,$$

in which case $\text{tp}^0_{<_0}(p_0^x, p_0^y) = \{x > y\} = \text{tp}^0_{<_1}(p_1^x, p_1^y)$.

5.5 When the spoiler plays next to the other pebble: Case (II)

Suppose now that the spoiler places p_0^x next to the other pebble in $\mathcal{A}_0$ (i.e. $p_0^x \in N^1_{\mathcal{A}_0}(p_0^y)$), but not in S_0^r (for that move would fall under the jurisdiction of Case (I)). In that case, the duplicator must place p_1^x so that the relative position of p_1^x and p_1^y is the same as that of p_0^x and p_0^y.

For that, we can use (E_{r+1}), which guarantees that $\text{env-tp}^{r+1}_{(\mathcal{A}_0,<_0)}(p_0^y) = \text{env-tp}^{r+1}_{(\mathcal{A}_1,<_1)}(p_1^y)$.
Thus there exists an isomorphism ψ between $\mathcal{E}\text{nv}^{r+1}_{(\mathcal{A}_0,<_0)}(p_0^y)$ and $\mathcal{E}\text{nv}^{r+1}_{(\mathcal{A}_1,<_1)}(p_1^y)$. Note that this isomorphism is unique, by virtue of $<_0$ and $<_1$ being linear orders.
The duplicator's response is to place p_1^x on $\psi(p_0^x)$. Let us now prove that this strategy is correct with respect to our invariants (S_r), (E_r) and (T_r).

Proof of (S_r) in Case (II)

Because the spoiler's move does not fall under Case (I), we know that $p_0^x \notin S_0^r$.
Let us now show that p_1^x is not near the endpoints either: suppose that $p_1^x \in S_1^r$. By construction, since p_1^x and p_1^y are neighbors in $\mathcal{A}_1$, this entails that $p_1^y \in S_1^{r+1}$. But then, we know by (S_{r+1}) that $p_0^y = \varphi_1(p_1^y)$; and because ψ is the unique isomorphism between $\mathcal{E}\mathrm{nv}^{r+1}_{(\mathcal{A}_0,<_0)}(p_0^y)$ and $\mathcal{E}\mathrm{nv}^{r+1}_{(\mathcal{A}_1,<_1)}(p_1^y)$, ψ is equal to the restriction $\widetilde{\varphi_0}$ of φ_0:

$$\widetilde{\varphi_0} \;:\; \mathcal{E}\mathrm{nv}^{r+1}_{(\mathcal{A}_0,<_0)}(p_0^y) \;\to\; \mathcal{E}\mathrm{nv}^{r+1}_{(\mathcal{A}_1,<_1)}(p_1^y)\,.$$

Thus $p_0^x = \psi^{-1}(p_1^x) = \widetilde{\varphi_0}^{-1}(p_1^x) = \varphi_1(p_1^x)$, and by definition of the segments on $(\mathcal{A}_1,<_1)$, which are just a transposition of the segments of $(\mathcal{A}_0,<_0)$ via φ_0, $p_1^x \in S_1^r$ then entails that $p_0^x \in S_0^r$, which is absurd.
Since we neither have $p_0^x \in S_0^r$ nor $p_1^x \in S_1^r$, (S_r) holds - recall from Note 9 that the part concerning p_0^y and p_1^y is always satisfied.

Proof of (E_r) in Case (II)

Recall that the duplicator placed p_1^x on the image of p_0^x by the isomorphism

$$\psi : \mathcal{E}\mathrm{nv}^{r+1}_{(\mathcal{A}_0,<_0)}(p_0^y) \to \mathcal{E}\mathrm{nv}^{r+1}_{(\mathcal{A}_1,<_1)}(p_1^y)\,.$$

It is easy to check that the restriction $\widetilde{\psi}$ of ψ: $\widetilde{\psi} \;:\; \mathcal{E}\mathrm{nv}^{r}_{(\mathcal{A}_0,<_0)}(p_0^x) \;\to\; \mathcal{E}\mathrm{nv}^{r}_{(\mathcal{A}_1,<_1)}(p_1^x)$ is well defined, and is indeed an isomorphism.
This ensures that $\text{env-tp}^r_{(\mathcal{A}_0,<_0)}(p_0^x) = \text{env-tp}^r_{(\mathcal{A}_1,<_1)}(p_1^x)$, thus completing the proof of (E_r).

Proof of (T_r) in Case (II)

This follows immediately from the fact that the isomorphism ψ maps p_0^x to p_1^x and p_0^y to p_1^y: all the atomic facts about these elements are preserved.

5.6 When the spoiler plays to the left: Cases (III) and (IV)

We now treat our last case, which covers both Cases (III) and (IV), i.e. the instances where the spoiler places p_0^x to the left of p_0^y (formally: such that $(\mathcal{A}_0,<_0) \models p_0^x < p_0^y$), which do not already fall in Cases (I) and (II).
Note that the scenario in which the spoiler plays to the right of the other pebble is the exact symmetric of this one (since the X_i play no role in this case, left and right can be interchanged harmlessly).

The idea here is very simple: since the spoiler has placed p_0^x to the left of p_0^y, but neither in S_0^r nor in $N^1_{\mathcal{A}_0}(p_0^y)$, the duplicator responds by placing p_1^x on an element of UL_1^{r+1} (the leftmost universal segment not in S_1^r) sharing the same k-environment. This is possible by construction of the universal segments: if $\tau_l := \text{env-tp}^k_{(\mathcal{A}_0,<_0)}(p_0^x)$ (which must extend a frequent k-neighborhood type, since $p_0^x \notin X_0$), then $\varphi_0(a[\tau_l]_L^{r+1})$ satisfies the requirements.

There is one caveat to this strategy. If p_1^y is itself in L_1^{r+1}, two problems may arise: first, it is possible for p_1^x and p_1^y to be in the wrong order (i.e. such that $(\mathcal{A}_1,<_1) \models p_1^x > p_1^y$). Second, it may be the case that p_1^x and p_1^y are neighbors in $\mathcal{A}_1$, which, together with the fact that p_0^x and p_0^y are orthogonal in $\mathcal{A}_0$ (i.e. $\mathrm{tp}^0_{\mathcal{A}_0}(p_0^x,p_0^y) = \mathrm{tp}^0_{\mathcal{A}_0}(p_0^x) \cup \mathrm{tp}^0_{\mathcal{A}_0}(p_0^y)$), would break (T_r).
This is why the duplicator's strategy depends on whether $p_1^y \in L_1^{r+1}$:

- if this is not the case, then the duplicator places p_1^x on $\varphi_0(a[\tau_l]_L^{r+1})$. This corresponds to Case (III).
- if $p_1^y \in L_1^{r+1}$, then (S_{r+1}) guarantees that $p_0^y \in L_0^{r+1}$. Hence p_0^x, which is located to the left of p_0^y, is in the domain of φ_0: the duplicator moves p_1^x to $\varphi_0(p_0^x)$. This situation corresponds to Case (IV).

Let us prove that (S_r), (E_r) and (T_r) hold in both of these instances.

Proof of (S_r) in Case (III)

Since the spoiler's move does not fall under Case (I), we have that $p_0^x \notin S_0^r$.
By construction, $a[\tau_l]_L^{r+1} \in L_0^{r+1}$, thus $\varphi_0(a[\tau_l]_L^{r+1}) \in L_1^{r+1}$, and $p_1^x \notin S_1^r$.

Proof of (E_r) in Case (III)

It follows from $\text{env-tp}^k_{(\mathcal{A}_0,<_0)}(a[\tau_l]_L^{r+1}) = \tau_l$ together with Lemma 5 that

$$\text{env-tp}^k_{(\mathcal{A}_0,<_0)}(p_0^x) = \text{env-tp}^k_{(\mathcal{A}_1,<_1)}(p_1^x)\,.$$

A fortiori, $\text{env-tp}^r_{(\mathcal{A}_0,<_0)}(p_0^x) = \text{env-tp}^r_{(\mathcal{A}_1,<_1)}(p_1^x)$.

Proof of (T_r) in Case (III)

Because the spoiler's move does not fall under Case (II), $p_0^x \notin N^1_{\mathcal{A}_0}(p_0^y)$. In other words,

$$\text{tp}^0_{\mathcal{A}_0}(p_0^x, p_0^y) = \text{tp}^0_{\mathcal{A}_0}(p_0^x) \cup \text{tp}^0_{\mathcal{A}_0}(p_0^y)\,.$$

Recall the construction of UL_0^{r+1}: the whole k-neighborhood of $a[\tau_l]_L^{r+1}$ was included in this segment. In particular, $N^1_{\mathcal{A}_1}(p_1^x) = N^1_{\mathcal{A}_1}(\varphi_0(a[\tau_l]_L^{r+1})) \subseteq UL_1^{r+1}$. By assumption, $p_1^y \notin L_1^{r+1}$, which entails that $\text{tp}^0_{\mathcal{A}_1}(p_1^x, p_1^y) = \text{tp}^0_{\mathcal{A}_1}(p_1^x) \cup \text{tp}^0_{\mathcal{A}_1}(p_1^y)$.
It then follows from the last observation of Note 9 that $\text{tp}^0_{\mathcal{A}_0}(p_0^x, p_0^y) = \text{tp}^0_{\mathcal{A}_1}(p_1^x, p_1^y)$.

Let us now prove that $\text{tp}^0_{<_1}(p_1^x, p_1^y) = \{x < y\}$.
We claim that $p_1^y \notin X_1 \cup \bigcup_{0\le j\le r+1} L_1^j$. Suppose otherwise: (S_{r+1}) would entail that $p_0^y \in X_0 \cup \bigcup_{0\le j\le r+1} L_0^j$ which, together with the hypothesis $p_0^y \notin L_0^{r+1}$ and $p_0^x < p_0^y$, would result in p_0^x being in S_0^r, which is absurd.
Thus, $\text{tp}^0_{<_1}(p_1^x, p_1^y) = \{x < y\} = \text{tp}^0_{<_0}(p_0^x, p_0^y)$, which concludes the proof of (T_r).

Proof of (S_r), (E_r) and (T_r) in Case (IV)

Let us now move to the case where $p_1^y \in L_1^{r+1}$. Recall that under this assumption, $p_0^y = \varphi_1(p_1^y) \in L_0^{r+1}$ and since $p_0^x < p_0^y$ and $p_0^x \notin S_0^r$, we have that $p_0^x \in L_0^{r+1}$.
The duplicator places the pebble p_1^x on $\varphi_0(p_0^x)$; in particular, $p_1^x \in L_1^{r+1}$.

The proof of (S_r) follows from the simple observation that $p_0^x \notin S_0^r$ and $p_1^x \notin S_1^r$.
As for (E_r) and (T_r), they follow readily from Lemma 5 and 6 and the fact that $p_1^x = \varphi_0(p_0^x)$ and $p_1^y = \varphi_0(p_0^y)$.

6 Counting quantifiers

We now consider the natural extension C^2 of FO^2, where one is allowed to use counting quantifiers of the form $\exists^{\geq i}x$ and $\exists^{\geq i}y$, for $i \in \mathbb{N}$. Such a quantifier, as expected, expresses the existence of at least i elements satisfying the formula which follows it. This logic C^2 has been extensively studied. On an expressiveness standpoint, C^2 stricly extends FO^2

(which cannot count up to three), and contrary to the latter, C^2 does not enjoy the small model property (meaning that contrary to FO^2, there exist satisfiable C^2-sentences which do not have small - or even finite - models). However, the satisfiability problem for C^2 is still decidable [6, 15, 16]. To the best of our knowledge, it is not known whether $<$-inv C^2 has a decidable syntax. Let us now explain how the proof of Theorem 1 can be adapted to show the following stronger version:

▶ **Theorem 10.** *Let $\mathcal{C}$ be a class of structures of bounded degree. Then $<$-inv $C^2 \subseteq FO$ on $\mathcal{C}$.*

Proof. The proof is very similar as to that of Theorem 1. The difference is that we now need to show, at the end of the construction, that the structures $(\mathcal{A}_0, <_0)$ and $(\mathcal{A}_1, <_1)$ are not only FO^2-similar, but C^2-similar. More precisely, we show that for every $k \in \mathbb{N}$, there exists some $f(k) \in \mathbb{N}$ such that if $\mathcal{A}_0 \equiv^{FO}_{f(k)} \mathcal{A}_1$, then it is possible to construct two linear orders $<_0$ and $<_1$ such that $(\mathcal{A}_0, <_0)$ and $(\mathcal{A}_1, <_1)$ agree on all C^2-sentences of quantifier rank at most k, and with counting indexes at most k, which we denote $(\mathcal{A}_0, <_0) \equiv^{C^2}_{k,k} (\mathcal{A}_1, <_1)$. This is enough to complete the proof, as these classes of C^2-sentences cover all the C^2-definable properties.

In order to prove that $(\mathcal{A}_0, <_0) \equiv^{C^2}_{k,k} (\mathcal{A}_1, <_1)$, we need an Ehrenfeucht-Fraïssé-game capturing $\equiv^{C^2}_{k,k}$. It is not hard to derive such a game from the Ehrenfeucht-Fraïssé-game for C^2 [12]. This game only differs from the two-pebble Ehrenfeucht-Fraïssé-game in that in each round, once the spoiler has chosen a structure (say $(\mathcal{A}_0, <_0)$) and a pebble to move (say p^x_0), the spoiler picks not only one element of that structure, but a set P_0 of up to k elements. Then the duplicator must respond with a set P_1 of same cardinality in $(\mathcal{A}_1, <_1)$. The spoiler then places p^x_1 on any element of P_1, to which the duplicator responds by placing p^x_0 on some element of P_0. As usual, the spoiler wins after this round if $\mathrm{tp}^0_{(\mathcal{A}_0,<_0)}(p^x_0, p^y_0) \neq \mathrm{tp}^0_{(\mathcal{A}_1,<_1)}(p^x_1, p^y_1)$. Otherwise, the game goes on until k rounds are played.

It is not hard to establish that this game indeed captures $\equiv^{C^2}_{k,k}$, in the sense that $(\mathcal{A}_0, <_0) \equiv^{C^2}_{k,k} (\mathcal{A}_1, <_1)$ if and only if the duplicator has a winning strategy for k rounds of this game. The restriction on the cardinal of the set chosen by the spoiler (which is at most k) indeed corresponds to the fact that the counting indexes of the formulas are at most k. As for the number of rounds (namely, k), it corresponds as usual to the quantifier rank. This can be easily derived from a proof of Theorem 5.3 in [12], and is left to the reader.

Let us now explain how to modify the construction of $<_0$ and $<_1$ presented in Section 4 in order for the duplicator to maintain similarity for k-round in such a game. The only difference lies in the choice of the universal elements. Recall that in the previous construction, we chose, for each k-environment type τ_l extending a frequent k-neighborhood type and each segment UL^j_0, an element $a[\tau_l]^j_L$ whose k-environment type in $(\mathcal{A}_0, <_0)$ is destined to be τ_l (and similarly for UR^j_0 and $a[\tau_l]^j_R$).

In the new construction, we pick k such elements, instead of just one. Just as previously, all these elements must be far enough from one another in the Gaifman graph of $\mathcal{A}_0$. Once again, this condition can be met by virtue of the k-neighborhood type τ underlying τ_l being frequent, and thus having many occurrences scattered across $\mathcal{A}_0$ (remember that we have a bound on the degree of $\mathcal{A}_0$, thus all the occurrences of τ cannot be concentrated). We only need to multiply the value of m by k in (3).

When the spoiler picks a set of elements of size at most k in one of the structures (say P_0 in $(\mathcal{A}_0, <_0)$), the duplicator responds by selecting, for each one of the elements of P_0, an element in $(\mathcal{A}_1, <_1)$ along the strategy for the FO^2-game explained in Section 5.3. All that remains to be shown is that it is possible for the duplicator to answer each element of P_0 with a different element in $(\mathcal{A}_1, <_1)$.

Note that if the duplicator follows the strategy from Section 5.3, they will never answer two moves by the spoiler falling under different cases among Cases (I)-(VI) with the same element. Thus we can treat separately each one of these cases; and for each case, we show that if the spoiler chooses up to k elements in $(\mathcal{A}_0, <_0)$ falling under this case in P_0, then the duplicator can find the same number of elements in $(\mathcal{A}_1, <_1)$, following the aforementionned strategy.

- For Case (I), this is straightforward, since the strategy is based on the isomorphism between the borders of the linear orders. The same goes for Cases (II), (IV) and (VI), as the strategy in these cases also relies on an isomorphism argument.
- Suppose now that $p_0^y \notin L_0^{r+1}$, and assume that the spoiler chooses several elements to the left of p_0^y, but outside of S_0^r and not adjacent to p_0^y. This corresponds to Case (III). Recall that our new construction guarantees, for each k-environment type extending a frequent k-neighborhood type, the existence in L_1^{r+1} of k elements having this environment. This lets us choose, in L_1^{r+1}, a distinct answer for each element in the set selected by the spoiler, sharing the same k-environment type. Case (V) is obviously symmetric.

This concludes the proof of Theorem 10. ◀

7 Conclusion

We have established that, when the degree is bounded, properties definable in the order-invariant extension of the two-variable fragment of first-order logic with counting are definable in first-order logic. From there, there seem to be three axes in which one can try to complete the picture.

The first natural question is whether this inclusion of expressiveness still holds when we release the hypothesis on the degree. As stated in the introduction, an upcoming result [2] exhibits two classes of structures (of unbounded degree) $\mathcal{C}' \subseteq \mathcal{C}$ such that (i) $\mathcal{C}'$ is definable in $\mathcal{C}$ by an FO^2-sentence φ using an order, such that φ is order-invariant on all structures of $\mathcal{C}$ (ii) $\mathcal{C}'$ is not FO-definable, and (iii) $\mathcal{C}$ is FO^3-definable but not FO^2-definable. Note that according to the standard definition, φ is not a sentence of $<$-inv FO^2, because its order-invariance does not hold on all finite structures. However, in view of this construction, it stands to reason to believe that in general, $<$-inv $\mathrm{FO}^2 \not\subseteq$ FO. Formally proving this non-inclusion would be an interesting goal.

Second, it is quite clear that the inclusion $<$-inv $\mathrm{C}^2 \subseteq$ FO when the degree is bounded is a severe over-approximation of the expressive power of $<$-inv C^2. For instance, it is not hard to prove that $<$-inv C^2 cannot define the class of triangle-free graphs: no sentence from $<$-inv C^2 can make a distinction between a large enough cycle and the disjoint union of a large enough cycle together with a triangle. This can be seen, following the general strategy detailed in Section 3, by constructing two carefully chosen linear orders on these graphs. Figure 2 illustrates the construction of such orders. Notice that there are only three kinds of elements: those whose two neighbors are on their left, those for which they are on their right, and those which have one neighbor on each side. By making sure to always respond to a move by the spoiler with an element of the same kind (and, of course, by implementing a tit-for-tat strategy near the endpoints), the duplicator can easily win the Ehrenfeucht-Fraïssé-game capturing $\equiv_{k,k}^{\mathrm{C}^2}$, provided that the cycles are long enough with respect to k.

It would be interesting to find upper bounds for $<$-inv FO^2 and $<$-inv C^2 tighter than FO; that is, tighter than the fragment $\exists^*\forall^*\exists^*$FO, to which FO collaspes when the degree is bounded [4] (since already in this fragment, one can count the number of occurrences of neighborhood types up to some threshold). Let us briefly explain why we fall short of giving

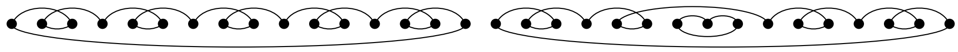

Figure 2 Illustration of two linear orders (growing from left to right) on a cycle (left figure) and the disjoint union of a cycle and a triangle (right figure), which are indistinguishable by any C^2-sentence of small enough quantifier rank and maximal counting index (where "small" is understood with respect to the length of the cycles).

such a bound: in such an attempt, the initial assumption about the similarity between $\mathcal{A}_0$ and $\mathcal{A}_1$ would be weaker than FO-similarity, and it would not be possible to base our work on neighborhoods. In this context, the starting hypothesis on the two structures would lack the rigidity which seems necessary to construct linear orders preserving FO^2-similarity or C^2-similarity. Establishing such a tighter bound thus seems to call for new techniques.

Last, we conjecture that the inclusion still holds when we lift the restriction on the number of variables; namely that $<$-inv FO $=$ FO when the degree is bounded. This would generalize the equality Succ-inv FO $=$ FO when the degree is bounded, obtained in [7]. Our hope is that a construction inspired by this one, albeit significantly refined, and in particular by the alternation of universal and neighbors segments, could possibly lead to establish such a result. We leave these three questions, as well as the issue of the syntactic decidability of $<$-inv C^2, for further research.

References

1. Albert Atserias, Anuj Dawar, and Martin Grohe. Preservation under extensions on well-behaved finite structures. *SIAM Journal on Computing*, 2008.
2. Bartosz Bednarczyk. Personal communication, July 2022.
3. Michael Benedikt and Luc Segoufin. Towards a characterization of order-invariant queries over tame graphs. *J. Symb. Log.*, 2009.
4. Ronald Fagin, Larry J. Stockmeyer, and Moshe Y. Vardi. On monadic NP vs. monadic co-NP. *Inf. Comput.*, 1995.
5. Erich Grädel and Martin Otto. On logics with two variables. *Theor. Comput. Sci.*, 1999.
6. Erich Grädel, Martin Otto, and Eric Rosen. Two-variable logic with counting is decidable. In *LICS*, 1997.
7. Julien Grange. Successor-invariant first-order logic on classes of bounded degree. *Log. Methods Comput. Sci.*, 2021.
8. Julien Grange and Luc Segoufin. Order-Invariant First-Order Logic over Hollow Trees. In *Computer Science Logic, CSL*, 2020.
9. Martin Grohe and Thomas Schwentick. Locality of order-invariant first-order formulas. *ACM Trans. Comput. Log.*, 2000. doi:10.1145/343369.343386.
10. Neil Immerman. Relational queries computable in polynomial time. *Information and Control*, 1986.
11. Neil Immerman and Dexter Kozen. Definability with bounded number of bound variables. *Inf. Comput.*, 1989.
12. Neil Immerman and Eric Lander. Describing graphs: A first-order approach to graph canonization. In *Complexity theory retrospective*. Springer, 1990.
13. 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.
14. Michael Mortimer. On languages with two variables. *Math. Log. Q.*, 1975.
15. Ian Pratt-Hartmann. Complexity of the guarded two-variable fragment with counting quantifiers. *J. Log. Comput.*, 2007.
16. Ian Pratt-Hartmann. The two-variable fragment with counting revisited. In *WoLLIC*, 2010.

17 Moshe Y. Vardi. The complexity of relational query languages. In *Proceedings of the 14th Annual ACM Symposium on Theory of Computing*, 1982.

18 Philipp Weis and Neil Immerman. Structure theorem and strict alternation hierarchy for FOˆ2 on words. *Log. Methods Comput. Sci.*, 2009.

19 Thomas Zeume and Frederik Harwath. Order-invariance of two-variable logic is decidable. In *LICS*, 2016.