Model Checking on Interpretations of Classes of Bounded Local Cliquewidth *

Édouard Bonnet† Jan Dreier‡ Jakub Gajarský§ Stephan Kreutzer¶
Nikolas Mählmann|| Pierre Simon** Szymon Toruńczyk††

February 25, 2022

Abstract

We present a fixed-parameter tractable algorithm for first-order model checking on interpretations of graph classes with bounded local cliquewidth. Notably, this includes interpretations of planar graphs, and more generally, of classes of bounded genus. To obtain this result we develop a new tool which works in a very general setting of dependent classes and which we believe can be an important ingredient in achieving similar results in the future.

1 Introduction

Algorithmic meta-theorems aim to explain the tractability of entire families of problems that can be specified in some logic. The prime example is Courcelle's theorem [6], stating that every problem expressible in monadic second-order logic (MSO) can be solved in linear time on every class of graphs with bounded treewidth. In this paper, we follow a

*This research was initiated at the Dagstuhl workshop *Sparsity in Algorithms, Combinatorics and Logic* (September 2021). We wish to thank the organizers and other participants. E.B. was supported by the ANR projects TWIN-WIDTH (ANR-21-CE48-0014) and Digraphs (ANR-19-CE48-0013). J.G. and S.T. were supported by the project BOBR that is funded from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreements No. 683080 and 948057, respectively). N.M. was supported by the German Research Foundation (DFG) with grant greement No. 444419611.

†Univ Lyon, CNRS, ENS de Lyon, UCBL 1, LIP UMR5668, France, edouard.bonnet@ens-lyon.fr
‡TU Wien, Austria, dreier@ac.tuwien.ac.at
§University of Warsaw, Poland, gajarsky@mimuw.edu.pl
¶TU Berlin, Germany, kreutzer@tu-berlin.de
||University of Bremen, Germany, maehlmann@uni-bremen.de
**University of Berkeley, USA, pierre.simon@berkeley.edu
††University of Warsaw, Poland, szymtor@mimuw.edu.pl

long line of research concerned with algorithmic meta-theorems for first-order logic (FO), on restricted classes of graphs. The central problem here is the first-order model checking problem, where one should decide whether a given FO sentence φ holds in a given graph G. A naive algorithm solves this problem in time $\mathcal{O}(|G|^{|\varphi|})$ whereas no algorithm can solve it in time $|G|^{o(|\varphi|)}$ in general, unless SAT admits a subexponential-time algorithm. The main goal of this line of research is to identify classes of graphs for which the problem is *fixed-parameter tractable* (FPT), i.e., solvable in time $f(\varphi) \cdot |G|^c$, for some constant c and computable function $f\colon \mathbb{N} \to \mathbb{N}$. Henceforth we call such classes *tractable*. Courcelle's theorem gives such an algorithm even for the more powerful logic MSO, on all classes of bounded treewidth.

The first result of this kind for FO, proven by Seese [30], states that FO model checking is FPT on every class of graphs with bounded maximum degree. This result is also the first application of the *locality method*, utilizing the locality of first-order logic, as formalized for example by Gaifman's locality theorem. Gaifman's theorem implies in particular that for two vertices u, v of a graph G (that are sufficiently far apart), whether or not u and v satisfy a fixed formula $\varphi(x, y)$ can be determined by looking only at neighborhoods of bounded radius around u and around v in G. The locality method was extended by Frick and Grohe [13] who showed that if there is an FPT algorithm for all classes $\mathcal{C}$ satisfying a certain property $\mathcal{P}$ (where the exponent in the run time of the algorithm is the same for all $\mathcal{C} \in \mathcal{P}$), then this immediately implies the existence of such an FPT algorithm for all classes $\mathcal{C}$ that *locally* have property $\mathcal{P}$. A class $\mathcal{C}$ has *locally* property $\mathcal{P}$ if for every fixed radius r, the class of all r-balls of graphs from $\mathcal{C}$ has property $\mathcal{P}$. For example, a class $\mathcal{C}$ has locally bounded treewidth if there is a function $f\colon \mathbb{N} \to \mathbb{N}$ such that for every $G \in \mathcal{C}$ and vertex $v \in V(G)$, the subgraph of G induced by the r-ball around v has treewidth at most $f(r)$. Such classes are also said to have *bounded local treewidth*. Planar graphs, graphs of bounded genus, and more generally, apex-minor-free graphs, have bounded local treewidth, so FO model checking is FPT on all those classes, by the observation of Frick and Grohe combined with the result of Courcelle. The locality method was subsequently combined with the graph minor theory of Robertson and Seymour, to capture all classes that exclude a minor [12], or more generally, classes that locally exclude a minor [9].

A new paradigm, based on sparsity theory developed by Nešetřil and Ossona de Mendez, has allowed to obtain further, more general tractability results. Dvořák, Král and Thomas showed that FO model checking is FPT for every class with *bounded expansion* [11]. And finally, Grohe, Kreutzer and Siebertz showed that the same holds for every *nowhere dense* graph class [20]. Those include all the classes mentioned above. See also Figure 1 for the relationship between these classes.

All graph classes we discussed so far are *monotone*, i.e., closed under removing vertices and edges. For sparse graph classes, monotonicity appears to be a reasonable assumption: after all, removing edges from a sparse graph should only make it even sparser.

For monotone graph classes, the aforementioned results are beautifully complemented by matching lower bounds. MSO_2 model checking is not FPT on monotone graph classes

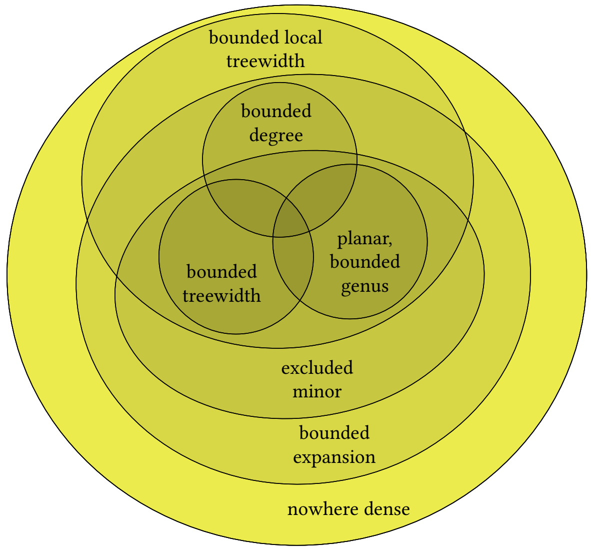

Figure 1: Inclusion diagram of selected monotone graph classes with FPT FO model checking.

whose treewidth is at least polylogarithmic with respect to the number of vertices [23, 18] and FO model checking is not FPT on monotone graph classes that are not nowhere dense [11, 22]. Thus, the aforementioned results yield a complete characterization of the monotone graph classes admitting FPT model checking of FO, and an almost complete characterization of the monotone graph classes admitting FPT model checking of MSO_2.

However, this is far from the complete picture, as this says nothing about the tractability for *dense* graph classes. Simple examples of graph classes that are not monotone, but admit efficient FO model checking are the class of complete graphs, or more generally, the class of edge complements of graphs from a fixed nowhere dense class. Those are not contained in any tractable monotone graph class, as every monotone graph class that contains cliques of unbounded size also contains all graphs. Thus, to make further progress, we need a paradigm shift towards considering non-monotone graph classes and width measures.

Dense Graph Classes. A graph class $\mathcal{C}$ is *hereditary* if $\mathcal{C}$ is closed under taking *induced* subgraphs, that is, under removing vertices. Since we do not assume closure under edge removal, hereditary classes are well suited to capture dense graph classes. After the question for monotone classes has been settled, the major next goal is to characterize hereditary graph classes for which FO model checking is FPT. This is again inspired by results for MSO model checking: The result by Courcelle, Makowsky, and Rotics [8], combined with the result of Oum and Seymour [28], shows that MSO model checking is FPT on classes of *bounded cliquewidth*. Cliquewidth is a generalization of the notion of treewidth to dense graphs. In particular, it is preserved by taking edge complements.

Applying again the locality argument to classes of bounded cliquewidth yields the following result, originating in the work of Frick and Grohe [13]. Say that a class $\mathcal{C}$ has *bounded local cliquewidth* if there is a function $f \colon \mathbb{N} \to \mathbb{N}$ such that for every number $r \in \mathbb{N}$, graph $G \in \mathcal{C}$, and vertex $v \in V(G)$, the subgraph of G induced by the r-ball around v has cliquewidth at most $f(r)$.

Theorem 1.1. *Let $\mathcal{C}$ be a class with bounded local cliquewidth. Then* FO *model checking is fixed-parameter tractable on $\mathcal{C}$.*

Currently, classes of bounded local cliquewidth are one of a few dense families for which FO model checking is known to be fixed-parameter tractable. However, there are many other graph classes that are conjectured to be tractable (for a more detailed discussion, see Figure 2 and Section 5). Those include, in particular, classes that can be obtained from tractable classes, using FO formulas, as we make precise now.

Interpretations. Let Σ and Γ be two signatures, where Γ is relational. A *simple interpretation* $\mathsf{I} \colon \Sigma \to \Gamma$ (here, *interpretation* for short) is specified by a *domain* formula $\delta(x)$, and one formula $\varphi_R(x_1, \ldots, x_k)$ for each relation symbol $R \in \Gamma$ of arity k, where all those formulas are in the signature Σ. For a given Σ-structure $\mathbb{A}$, the interpretation I outputs the structure $\mathsf{I}(\mathbb{A})$ whose domain is the set $\delta(\mathbb{A}) := \{a \in \mathbb{A} \mid \mathbb{A} \models \delta(a)\}$, and in which the interpretation of each relation $R \in \Gamma$ of arity k consists of those tuples $(a_1, \ldots, a_k) \in \delta(\mathbb{A})^k$ that satisfy $\mathbb{A} \models \varphi_R(a_1, \ldots, a_k)$. Usually we will be working with interpretations that map graphs with expanded signatures to uncolored, undirected graphs, having a single binary relation E. In this case, we write $\mathsf{I}_{\varphi,\delta}$ for the interpretation consisting of an irreflexive, symmetric formula $\varphi(x, y)$ interpreting the edge relation E and a domain formula $\delta(x)$. If $\delta(x)$ is equal to $x = x$, we will just write I_φ instead. For example, the interpretation I_φ with $\varphi(x, y) = \neg E(x, y)$ maps a given graph to its edge complement, and the interpretation I_ψ with with $\psi(x, y) = E(x, y) \vee \exists z.E(x, z) \wedge E(z, y)$ maps a given graph to its square.

The notion of an interpretation lifts to classes of structures, for which we denote with $\mathsf{I}(\mathcal{C}) := \{\mathsf{I}(G) \mid G \in \mathcal{C}\}$ the result of applying the interpretation I to the class $\mathcal{C}$. Say that a class of structures $\mathcal{C}$ *interprets* a class of structures $\mathcal{D}$, or that $\mathcal{D}$ *interprets in* $\mathcal{C}$, if there is an interpretation I such that $\mathcal{D} \subseteq \mathsf{I}(\mathcal{C})$. Note that this notion depends on the chosen underlying logic, which will be either FO or MSO in our discussion. We may write $\mathcal{L}$-*interpretation* for interpretation when the underlying logic is $\mathcal{L} \in \{\text{FO}, \text{MSO}\}$. This yields a transitive relation: if $\mathcal{C}$ interprets $\mathcal{D}$ and $\mathcal{D}$ interprets $\mathcal{E}$, then $\mathcal{C}$ interprets $\mathcal{E}$.

A class $\mathcal{C}$ of graphs has bounded cliquewidth if and only if the class of trees MSO-interprets $\mathcal{C}$ [3, Proposition 27]. In particular, bounded cliquewidth is preserved by MSO interpretations. Moreover, we may view interpretations as a tool to extend model checking results from sparse to dense graph classes. This invites the question, originally asked in [16], whether a similar statement holds for first-order logic.

Question 1.2. *[16] Let $\mathcal{C}$ be a class admitting an* FPT *algorithm for* FO *model checking, and $\mathcal{D}$ be a class that* FO*-interprets in $\mathcal{C}$. Does there exist an* FPT *algorithm for* FO *model checking on $\mathcal{D}$?*

The intuition underlying this question is that if a graph class $\mathcal{C}$ is sufficiently well-behaved, then a fixed formula $\varphi(x, y)$ should not be able to define complicated graphs in graphs from $\mathcal{C}$.

Thus, in particular, Question 1.2 suggests the existence of an FPT algorithm for FO model checking for any class $\mathcal{D}$ that interprets in some nowhere dense class $\mathcal{C}$. How could such an algorithm look like? To unravel this question, fix an interpretation I such that $\mathcal{D} \subseteq \mathsf{I}(\mathcal{C})$, where $\mathcal{C}$ is the class of k-colored graphs from $\mathcal{C}$. Given a graph $G \in \mathcal{D}$ and a first-order formula φ that we want to evaluate on G, a possible strategy is to try to "reverse the interpretation" and compute a graph $G' \in \mathcal{C}$ such that $\mathsf{I}(G') = G$. This process then yields a formula φ' such that $G \models \varphi$ if and only if $G' \models \varphi'$. Since G' comes from a nowhere dense class, one can then evaluate in FPT time whether $G' \models \varphi'$. However, reversing an interpretation seems to be a difficult task[1]. In this approach, we do not necessarily need to revert the interpretation I as described above – there may be some other nowhere dense class $\mathcal{C}'$ and interpretation I' that is easier to revert, such that $\mathcal{D} \subseteq \mathsf{I}'(\mathcal{C}')$.

So far, only for classes which interpret in bounded degree classes the method outlined above has been applied successfully [16]. For the more general classes interpreting in bounded expansion classes, an FPT FO model checking now boils down to efficiently computing so-called *low shrubdepth covers* [17], or Lacon or shrub decompositions [10].

Main result. In this paper, we extend the result of [16] significantly by proving that FO model checking is FPT for every class that interprets in a class with bounded local cliquewidth.

Theorem 1.3 (Main result)**.** *Let $\mathcal{C}$ be a graph class that interprets in a class of graphs with bounded local cliquewidth. Then* FO *model checking is fixed-parameter tractable on $\mathcal{C}$: there exists a function f and a constant c such that for every first-order sentence φ and graph $G \in \mathcal{C}$ one can decide in time $f(|\varphi|) \cdot n^c$ whether $G \models \varphi$.*

Thus, we make progress towards answering Question 1.2, by answering it positively in the case of interpretations of classes with bounded local cliquewidth. See Figure 2 for an overview on how our result relates to previous results. We remark that besides being much more general, our proof is also much simpler than the proof in [16]. As we explain in the proof outline below, our main lemma applies to much more general classes than just classes of bounded local cliquewidth – namely to all *NIP classes* – yielding a more general theorem than Theorem 1.3 (see Theorem 5.1). We proceed with a proof outline in Section 2, followed by the actual proofs and then an extended discussion in Section 5, comparing our results to existing results.

[1] For instance, it is NP-complete [25] to decide whether a given graph is a square of some graph.

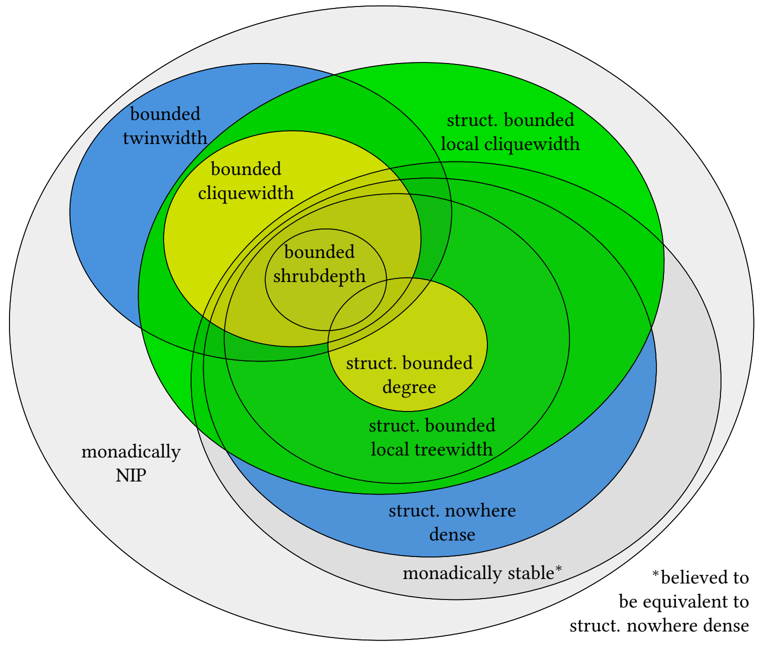

Figure 2: Inclusion diagram of selected transduction ideals, that is, properties of graph classes that are closed under transductions (that is, under interpretations of colorings of the graphs from the class). Yellow transduction ideals were previously known to admit an FPT FO model checking algorithm. Green transduction ideals admit an FPT FO model checking algorithm, as presented in this paper. Blue ideals admit an FPT FO model checking algorithm, assuming an appropriate decomposition is given as part of the input. Uncolored means unknown. The relevant notions are discussed in Section 5.

2 Proof outline

In this section, we sketch the proof of Theorem 1.3. This proof outline is not complete, and for simplicity of the description assumes interpretations in which the domain formula $\delta(x)$ holds for all x. For a complete proof see Sections 3 and 4.

We first describe a possible proof strategy for proving Theorem 1.3, outlined in [15], in order to isolate the main obstacle. The following lemma is an immediate consequence of Gaifman's locality theorem [14]. By $\mathrm{dist}^G(u,v)$ denote the distance between two vertices u and v in a graph G.

Lemma 2.1. *Let $\varphi(x,y)$ be an* FO *formula. Then there are numbers $r,t \in \mathbb{N}$ such that every graph G can be vertex-colored using t colors in such a way that for any two vertices $u,v \in V(G)$ with $\mathrm{dist}^G(u,v) > r$, whether or not $\varphi(u,v)$ holds depends only on the color of u and the color of v.*

Rephrasing, the conclusion of Lemma 2.1 says that there is a formula $\alpha(x, y)$, which is a Boolean combination of checks of the colors of x and y, and is such that the formula $\psi(x, y) := \varphi(x, y) \oplus \alpha(x, y)$ has *range* $\leqslant r$, that is, for every graph G and vertices $u, v \in V(G)$, if $\psi(u, v)$ holds then $\mathrm{dist}^G(u, v) \leqslant r$. Here, $\oplus$ denotes the exclusive or.

This has the following consequence, observed in [15, 27]. If G is a graph and $X, Y \subseteq V(G)$ are sets of vertices of G, then doing a *flip* between X and Y yields a new graph where the adjacency of all pairs $x \in X$ and $y \in Y$ is inverted: adjacent pairs become non-adjacent, and vice-versa.

Corollary 2.2 ([15, 27])**.** *For every formula $\varphi(x, y)$ there are $r, t \in \mathbb{N}$ and a formula $\psi(x, y)$ of range $\leqslant r$ such that for every graph G, the graph $\mathsf{I}_\varphi(G)$ can be obtained from the graph $\mathsf{I}_\psi(G)$ by performing flips between t pairs of sets.*

To see this, perform a flip for every pair of color classes C, D (as given by Lemma 2.1) such that $\varphi(u, v)$ holds for some $u \in C$ and $v \in D$ with $\mathrm{dist}^G(u, v) > r$. So the t in Corollary 2.2 is in fact at most the square of the t obtained from Lemma 2.1.

Now, suppose we are given a class $\mathcal{C}$ with bounded local cliquewidth and an interpretation I_φ, for some FO formula $\varphi(x, y)$, and want to solve the model checking problem on the class $\mathsf{I}_\varphi(\mathcal{C})$. In this problem, we are given as input a graph of the form $\mathsf{I}_\varphi(G)$, for some $G \in \mathcal{C}$ which is unknown, and a sentence α, and are to determine whether $\mathsf{I}_\varphi(G)$ satisfies α.

Let ψ be as in Corollary 2.2. As $\mathcal{C}$ has bounded local cliquewidth and $\psi(x, y)$ has range $\leqslant r$, it is not difficult to prove that $\mathsf{I}_\psi(\mathcal{C})$ is again a class with bounded local cliquewidth (this relies on the fact that classes with bounded cliquewidth are closed under FO-interpretations, and is proved in Lemma 4.6). Hence, FO model checking can efficiently be solved on the graph $\mathsf{I}_\psi(G)$ as given by Corollary 2.2. To model check the sentence α on $\mathsf{I}_\varphi(G)$ it is enough to model-check another sentence α' on the graph $\mathsf{I}_\psi(G)$ expanded with unary predicates marking the t pairs of sets that need to be flipped to obtain $\mathsf{I}_\psi(G)$ from $\mathsf{I}_\varphi(G)$. Here we use the fact that the same flips can be used to recover $\mathsf{I}_\varphi(G)$ from $\mathsf{I}_\psi(G)$, and the flipping process can be simulated by α'.

To summarize, to determine whether $\mathsf{I}_\varphi(G)$ satisfies α, it suffices to determine whether $\mathsf{I}_\psi(G)$ (with additional unary predicates) satisfies α', and this can be done efficiently since $\mathsf{I}_\psi(\mathcal{C})$ has bounded local cliquewidth. And moreover $\mathsf{I}_\psi(G)$ can be obtained from $\mathsf{I}_\varphi(G)$ by performing t flips between pairs of sets. The problem with this approach is: how to determine the t pairs of sets that need to be flipped in order to obtain $\mathsf{I}_\psi(G)$ from $\mathsf{I}_\varphi(G)$? Lemma 2.1 allows us to find those sets when G is given, but not when $\mathsf{I}_\varphi(G)$ is given.

Our main lemma overcomes this difficulty by proving a version of Corollary 2.2 in which the t flips can be efficiently computed, given $\mathsf{I}_\varphi(G)$. Before we can state it, we will need the following fundamental notions originating from learning theory.

VC-dimension and NIP classes. Say that a formula $\varphi(x, y)$ has *VC-dimension* at least N on a structure G if there exist elements v_i for $i = 1, \ldots, N$ and w_I for $I \subseteq \{1, \ldots, N\}$ such

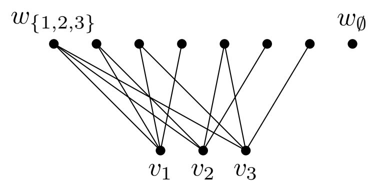

Figure 3: A graph for which the formula $\varphi(x, y) = E(x, y)$ has VC-dimension at least three.

that $\varphi(v_i, w_I)$ holds if and only if $i \in I$, for all $i = 1, \ldots, N$ and $I \subseteq \{1, \ldots, N\}$. See also Figure 3.

A class $\mathcal{C}$ of structures is *NIP* (or *dependent*) if for every first-order formula[2] $\varphi(x, y)$ there is some constant N such that the VC-dimension of φ on G is less than N for every $G \in \mathcal{C}$. Every class with bounded local cliquewidth is NIP [21]. There are many other known NIP classes $\mathcal{C}$, such as all nowhere dense classes, and more generally, all *monadically NIP classes* (see Section 5).

Main lemma. We are now ready to state our main technical lemma, in a form that parallels Lemma 2.1.

Lemma 2.3 (Main lemma)**.** *Let $\mathcal{C}$ be a class of graphs and let $\varphi(x, y)$ be an* FO *formula that has bounded VC-dimension on $\mathcal{C}$. Then there are numbers $s, r \in \mathbb{N}$ such that for every $G \in \mathcal{C}$ there is a set $S \subseteq V(G)$ of size at most s such that for any two vertices $u, v \in V(G)$ with $\mathrm{dist}^G(u, v) > r$, whether or not $\varphi(u, v)$ holds, depends only on $\varphi(u, S)$ and $\varphi(S, v)$.*

Here, $\varphi(u, S) := \{w \in S \mid G \models \varphi(u, w)\}$, and $\varphi(S, v)$ is defined symmetrically. In particular, if $\varphi(x, y) \leftrightarrow \varphi(y, x)$ holds, as is the case when considering formulas that define graphs, then we have that $\varphi(u, S) = \varphi(S, u)$. In what follows, we assume that $\varphi(x, y) \leftrightarrow \varphi(y, x)$ holds.

Given a set $S \subseteq V(G)$ define a coloring of $V(G)$ that colors a given $v \in V(G)$ with the set $\varphi(v, S) \subseteq S$. This coloring then uses at most 2^s colors, and is moreover definable in a straightforward way in the graph $\mathsf{I}_\varphi(G)$, by looking at the adjacencies between a given vertex and the vertices in S. The conclusion of the lemma says that for all u and v with $\mathrm{dist}^G(u, v) > r$, whether or not $\varphi(u, v)$ holds, depends only on the color of u and the color of v, that is, there is some binary relation $R \subseteq 2^S \times 2^S$ such that $G \models \varphi(u, v)$ if and only if the pair formed by the colors of u and v belongs to R. Hence, Lemma 2.3 can be seen as a variant of Lemma 2.1, where the coloring can moreover be efficiently computed, given the graph $\mathsf{I}_\varphi(G)$ and the set S.

[2] In the original definition [31], this condition is required for formulas $\varphi(\bar{x}, \bar{y})$, where $\bar{x}$ and $\bar{y}$ are tuples of variables. However, our proofs work with the weaker assumption.

Main algorithm. Using Lemma 2.3, we can now solve the model checking problem on $\mathsf{I}_\varphi(\mathcal{C})$, essentially in the way that was outlined above. More precisely, the algorithm works as follows. Given a graph $\mathsf{I}_\varphi(G) \in \mathsf{I}_\varphi(\mathcal{C})$ and an FO sentence α, in parallel for every set $S \subseteq V(G)$ with $|S| \leqslant s$, and every binary relation $R \subseteq 2^S \times 2^S$, do the following.

1. Compute the coloring of $V(G)$ as described above, using 2^S colors.
2. Compute the graph $\mathsf{I}_\psi(G)$ by performing flips between any pair of color classes such that belongs to R.
3. Check whether $\mathsf{I}_\psi(G)$ expanded with unary predicates marking the flipped sets, satisfies α', where α' is the formula that first recovers the graph $\mathsf{I}_\varphi(G)$ by undoing the flips, and then tests whether $\mathsf{I}_\varphi(G)$ satisfies α.

Whenever one of the parallel executions terminates, terminate with the same answer.

There is one technicality on which the proof of correctness of the above algorithm hinges. We do not know which of the parallel executions involves the "correct" set S and relation R resulting in a graph that belongs to a class of bounded local cliquewidth, but we know, by Lemma 2.3, that one of them does. So how do we know that we will receive a correct answer in the required running time?

First, we use the fact that interpretations with bounded-range formulas preserve classes with bounded local cliquewidth (Lemma 4.6). Second, we know that for every class $\mathcal{D}$ with bounded local cliquewidth there is a model checking algorithm that is guaranteed to be efficient on graphs from $\mathcal{D}$ only, but yields correct answers for all graphs (see Theorem 4.4). By applying this algorithm in parallel we are therefore guaranteed to efficiently get a correct answer. This completes the sketch of the proof of the main theorem, Theorem 1.3, using the main lemma. The details are presented in Section 4.

Proof of main lemma. We now outline the proof of the main lemma. See Section 3 for the complete argument. We use the following fundamental result based on the (p, q)-theorem [24] (see Theorem 3.4 below).

Theorem 2.4. *For every d there is a number k such that for every binary relation $E \subseteq A \times B$ of VC-dimension at most d, one of two cases holds:*

- *there is a set $A' \subseteq A$ with $|A'| \leqslant k$, such that for every $b \in B$ there is $a \in A'$ with $E(a, b)$, or*
- *there is a set $B' \subseteq B$ with $|B'| \leqslant k$, such that for every $a \in A$ there is $b \in B'$ with $\neg E(a, b)$.*

To prove Lemma 2.3, we proceed as follows. The starting point is again Lemma 2.1. Let r and t be given by that lemma. Fix a graph $G \in \mathcal{C}$ and its coloring as in Lemma 2.1. Assume, for the sake of simplicity, that every color class C is either *large*, that is contains 3 vertices with mutual distance larger than $2r$, or is *small*, that is, contains a *central* vertex

$c_0 \in C$ such that every vertex $v \in C$ is within distance at most $2r$ from c_0. This assumption is without much loss of generality, as every class that is neither large nor small can be partitioned into two new classes that are both small. We construct the set S as follows:

- for every large color class C, pick three elements which are mutually at distance larger than $2r$, and add them to S,
- for every small color class C, pick a central vertex $c_0 \in C$, and add it to S,
- for every pair A, B of color classes, let $S_{AB} \subseteq A \cup B$ be the result of applying Theorem 2.4 to the binary relation $E_\varphi \subseteq A \times B$ where $E_\varphi = \{(a,b) \in A \times B \mid G \models \varphi(a,b)\}$. Add S_{AB} to S.

This completes the construction of S. Note that $|S| \leqslant \mathcal{O}(t \cdot k^2)$, where k is given by Theorem 2.4. Correctness of the construction is verified for the radius $5r$. This amounts to proving that there are no vertices $u, v, u', v' \in V(G)$ such that:

- $\operatorname{dist}(u,v) > 5r$ and $\operatorname{dist}(u',v') > 5r$,
- $\varphi(u,S) = \varphi(u',S)$ and $\varphi(S,v) = \varphi(S,v')$,
- $\varphi(u,v)$ holds and $\neg\varphi(u',v')$ holds.

Assuming that such vertices exist, a contradiction is reached with the assumption that $\varphi(u,v)$ depends only on the color of u and the color of v whenever $\operatorname{dist}^G(u,v) > r$. This is done by performing a case analysis, depending on the sizes (large/small) of the color class $C(u')$ of u' and the color class $C(v)$ of v.

We showcase one of the four cases: when $C(u')$ and $C(v)$ are both small. As $\operatorname{dist}(u,v) > 5r$ and $C(v)$ is small, it follows that $\operatorname{dist}(u,w) > r$ for all $w \in C(v)$. Since $\varphi(u,v)$ holds, it follows that $\varphi(u,w)$ holds for all $w \in C(v)$. In particular, for $S(v) = S \cap C(v)$ we have $\varphi(u, S(v)) = S(v)$. As $\varphi(u,S) = \varphi(u',S)$ it follows that $\varphi(u', S(v)) = S(v)$ as well. By a symmetric argument, using the fact that $C(u')$ is small, we get that $\varphi(S(u'), v) = \emptyset$. This contradicts the construction of the set $S_{AB} \subseteq S$ for the pair $A = C(u')$ and $B = C(v)$.

The case when one of $C(u')$ and $C(v)$ is small and the other one is large uses similar arguments. The case when both classes are large is even more elementary, as it does not invoke the construction of the sets S_{AB}, and only relies on the existence of the three-element scattered sets in each of $C(u')$ and $C(v)$, that where selected to S.

This finishes the sketch of the proof of the main lemma, and hence also of the main theorem. Note that in Section 3 we state a slightly stronger version of Lemma 2.3, which is suited for treating interpretations in which the domain formula $\delta(x)$ is arbitrary. In Section 4, we prove Theorem 1.3.

3 Defining the relationship between far apart vertices

In this section we prove our main technical tool, Lemma 3.1. First we need some notation. For a formula $\varphi(x,y)$, elements u,v and a set S of elements of a structure G, write:

$$\varphi(u,S) := \{s \in S \mid G \models \varphi(u,s)\}$$
$$\varphi(S,v) := \{s \in S \mid G \models \varphi(s,v)\}.$$

Lemma 3.1 (Main lemma)**.** *Let $\mathcal{C}$ be a class of graphs and let $\varphi(x,y)$ be an* FO *formula that has bounded VC-dimension on $\mathcal{C}$. Then there are numbers $s,r \in \mathbb{N}$ such that for every graph $G \in \mathcal{C}$ and every $U \subseteq V(G)$ there is a set $S \subseteq U$ of size at most s such that for any two vertices $u,v \in U$ with* $\mathrm{dist}^G(u,v) > r$*, whether or not $\varphi(u,v)$ holds depends only on $\varphi(u,S)$ and $\varphi(S,v)$, where* depends only on *means that $G \models \varphi(u,v)$ iff $G \models \varphi(u',v')$ for any two pairs u,v and u',v' from U satisfying the following condition:*

$$\begin{aligned} \varphi(u,S) &= \varphi(u',S) \qquad \mathrm{dist}^G(u,v) > r, \\ \varphi(S,v) &= \varphi(S,v') \qquad \mathrm{dist}^G(u',v') > r. \end{aligned} \tag{$*$}$$

The following property will play a key role in the proof of the main results of this section.

Definition 3.2. *Let $E \subseteq A \times B$ be a binary relation. Say that E has a* duality of order k *if at least one of two cases holds:*

a) there is a set $A' \subseteq A$ of size at most k such that for every $b \in B$ there is some $a \in A'$ with $\neg E(a,b)$, or

b) there is a set $B' \subseteq B$ of size at most k such that for every $a \in A$ there is some $b \in B'$ with $E(a,b)$.

A *set system* $\mathcal{F}$ on a set X is a family $\mathcal{F}$ of subsets of X. The *VC-dimension* of $\mathcal{F}$ is the maximal size (or $+\infty$) of a subset $A \subseteq X$ such that $\{F \cap A \mid F \in \mathcal{F}\} = \mathcal{P}(A)$. For $m \in \mathbb{N}$ let $\pi_{\mathcal{F}}(m)$ denote the *shatter function* of $\mathcal{F}$, defined as

$$\pi_{\mathcal{F}}(m) := \max\Big\{\big|\{F \cap A : F \in \mathcal{F}\}\big| : A \subseteq X, |A| \leqslant m\Big\},$$

i.e., the maximum, over all sets $A \subseteq X$ with $|A| \leqslant m$, of the cardinality of $\{F \cap A \mid F \in \mathcal{F}\}$. It is well known that if $\mathcal{F}$ has VC-dimension d then $\pi_{\mathcal{F}}(m) = \mathcal{O}(m^d)$.

Define the VC-dimension of a binary relation $E \subseteq X \times Y$ as the VC-dimension of the set system $\{E(X,y) \mid y \in Y\}$ on Y.

The following is a special case of the (p,q)-*theorem*, stated below.

Theorem 3.3. *For every $d \in \mathbb{N}$ there is some $k \in \mathbb{N}$ such that the following holds. Let $E \subseteq A \times B$ have VC-dimension at most d, where A and B are finite. Then E has a duality of order k.*

This result follows from the proof of the conjecture of Hadwiger and Debrunner, see Matoušek [24, Theorem 4]. In the following formulation, which is dual to the formulation of Matoušek, the set system $\mathcal{F}$ is infinite.

Theorem 3.4 ([24]). *Let $\mathcal{F}$ be a set system on U with $\pi_{\mathcal{F}}(m) = o(m^k)$, for some integer k, and let $p \geqslant k$. Then there is a constant N such that the following holds for every finite set $V \subseteq U$: if for every $V' \subseteq V$ with $|V'| \leqslant p$ there is some $F \in \mathcal{F}$ containing V', then there is a family $\mathcal{F}' \subseteq \mathcal{F}$ with $|\mathcal{F}'| \leqslant N$ and $V \subseteq \bigcup \mathcal{F}'$.*

Proof of Theorem 3.3. Let $\mathcal{F}$ be the disjoint union of all finite set systems of VC-dimension at most d. Then $\mathcal{F}$ has VC-dimension at most d as well, and therefore $\pi_{\mathcal{F}}(m) = \mathcal{O}(m^d) = o(m^{d+1})$. Apply Theorem 3.4 to $p = d+1$, obtaining a number N with the following property: for every set system $\mathcal{G}$ on a finite set V of VC-dimension at most d, such that every p elements of V are contained in some element of $\mathcal{G}$, there is a set of at most N elements of $\mathcal{G}$ whose union contains V.

Let $E \subseteq A \times B$ have VC-dimension at most d, and let $\mathcal{G} = \{E(A,b) \mid b \in B\}$ be the corresponding set system on A.

Suppose there is a set $A' \subseteq A$ of size at most p such that for every $b \in B$ there is some $a \in A'$ with $\neg E(a,b)$. Then E has a duality of order $p = d+1$.

Otherwise, for every $A' \subseteq A$ of size at most p there is some $b \in B$ such that $E(a,b)$ holds for all $a \in A'$. This means that every subset of A of size at most p is contained in some element of $\mathcal{G}$. Hence, there is a subset $\mathcal{G}' \subseteq \mathcal{G}$ with $|\mathcal{G}'| \leqslant N$ such that $A = \bigcup \mathcal{G}'$. This means that there is a set $B' \subseteq B$ with $|B'| \leqslant N$ such that for every $a \in A$, $E(a,b)$ holds for some $b \in B'$. Then E has a duality of order N.

In either case, E has a duality of order $\max(d+1, N)$. $\square$

Fix a partition $\mathcal{P}$ of a set V. For an element $v \in V$, the *class* of v, denoted $C(v)$, is the unique $C \in \mathcal{P}$ containing v. In the context of the next theorem, a *pseudometric* is a symmetric function $f\colon V \times V \to \mathbf{R}^+ \cup \{+\infty\}$ satisfying the triangle inequality.

Theorem 3.5. *Fix $r, k, t \in \mathbb{N}$. Let V be a finite set equipped with:*

- *a binary relation $E \subseteq V \times V$ such that for all $A \subseteq V$ and $B \subseteq V$, $E \cap (A \times B)$ has a duality of order k,*
- *a pseudometric $\mathrm{dist}\colon V \times V \to \mathbf{R}_{\geqslant 0} \cup \{+\infty\}$,*
- *a partition $\mathcal{P}$ of V with $|\mathcal{P}| \leqslant t$,*

such that $E(u,v)$ depends only on $C(u)$ and $C(v)$ for all u, v with $\mathrm{dist}(u,v) > r$. Then there is a set $S \subseteq V$ of size $\mathcal{O}(kt^2)$ such that $E(u,v)$ depends only on $E(u,S)$ and $E(S,v)$ for all $u, v \in V$ with $\mathrm{dist}(u,v) > 5r$.

Proof. Say that a class $C \in \mathcal{P}$ is *large* if there are $s_1, s_2, s_3 \in C$ with mutual distance larger than $2r$. Say that a class $C \in \mathcal{P}$ is *small* if there is $c_0 \in C$ such that $\mathrm{dist}(c, c_0) \leqslant 2r$ for all $c \in C$. If a class $C \in \mathcal{P}$ is not large then there are $s_1, s_2 \in C$ such that $\mathrm{dist}(c, s_1) \leqslant 2r$ or $\mathrm{dist}(c, s_2) \leqslant 2r$ for all $c \in C$. For every class $C \in \mathcal{P}$ that is neither large nor small, pick arbitrarily any such s_1 and s_2 and let $C_1 = \{c \in C \mid c \neq s_2, \mathrm{dist}(c, s_1) \leqslant 2r\}$ and $C_2 := C - C_1$. Thus, by splitting the class $C \in \mathcal{P}$ into two classes C_1 and C_2, we arrive at the situation where both C_1 and C_2 are small. Hence, by at most doubling the number t of classes, we may assume that every class $C \in \mathcal{P}$ is either large or small.

Construction of S. We now construct the set S. For every ordered pair $(C, D) \in \mathcal{P}^2$ of classes let $S_{CD} \subseteq C \cup D$ be a duality of order k for $E \cap (C \times D)$, that is, $|S_{CD}| \leqslant k$ and one of two cases holds:

- for every $c \in C$ there is some $d \in S_{CD} \subseteq D$ with $E(c, d)$, or
- for every $d \in D$ there is some $c \in S_{CD} \subseteq C$ with $\neg E(c, d)$.

Such a set S_{CD} exists by the duality assumption of the lemma. Note that S_{CD} and S_{DC} are usually not the same and that we allow $C = D$ in the definition of S_{CD}. Let $S \subseteq V$ be the set containing the following elements:

- for every class $C \in \mathcal{P}$ that is large, any three elements $s_1, s_2, s_3 \in C$ with mutual distance larger than $2r$,
- a center c_0 of every small class C, so that $\mathrm{dist}(w, c_0) \leqslant 2r$ for all $w \in C$,
- all elements of S_{CD}, for every pair $(C, D) \in \mathcal{P}^2$.

Clearly, S has $\mathcal{O}(kt^2)$ elements.

Correctness. We show that S satisfies the condition in the lemma. Write $S(w)$ for $S \cap C(w)$, for $w \in V$. Towards a contradiction, suppose $u, v, u', v' \in V$ are such that:

1. $E(u, v)$ and $\neg E(u', v')$,
2. $\mathrm{dist}(u, v) > 5r$ and $\mathrm{dist}(u', v') > 5r$,
3. $E(u, S) = E(u', S)$,
4. $E(S, v) = E(S, v')$.

We show that this yields a contradiction.

For a pair of classes $C, D \in \mathcal{P}$, possibly with $C = D$, say that E *generically holds* between C and D if $E(c, d)$ holds for some $c \in C$ and $d \in D$ such that $\mathrm{dist}(c, d) > r$. Similarly define when $\neg E$ generically holds between C and D. Note that if E generically holds between C

and D then $E(c, d)$ holds for all $c \in C$ and $d \in D$ such that $\mathrm{dist}(c, d) > r$, by the assumption of the theorem. The same applies to $\neg E$.

By assumption, E generically holds between $C(u)$ and $C(v)$, whereas $\neg E$ generically holds between $C(u')$ and $C(v')$.

Claim 3.6. *The following hold:*

1. *If $C(u')$ is large, then $\neg E$ generically holds between $C(u')$ and $C(v)$.*
2. *If $C(v)$ is large, then E generically holds between $C(u')$ and $C(v)$.*

Proof. We prove the first item, as the other one follows by symmetry. The following situation is depicted in Figure 4.

Suppose $C(u')$ is large. Then there are three elements in $S(u')$ with mutual distance larger than $2r$. At most one of them can be at distance at most r from v'. So we have $s_1, s_2 \in S(u')$ with $\mathrm{dist}(s_i, v') > r$ for $i = 1, 2$. Then $\neg E(s_1, v')$ and $\neg E(s_2, v')$ holds since $\neg E$ generically holds between $C(u')$ and $C(v')$. Since $s_1, s_2 \in S$ and $E(S, v') = E(S, v)$, it follows that $\neg E(s_1, v)$ and $\neg E(s_2, v)$ hold as well. As above, v can be at distance at most r only from one of s_1 and s_2. It follows that $\neg E$ generically holds between $C(u')$ and $C(v)$. $\square$

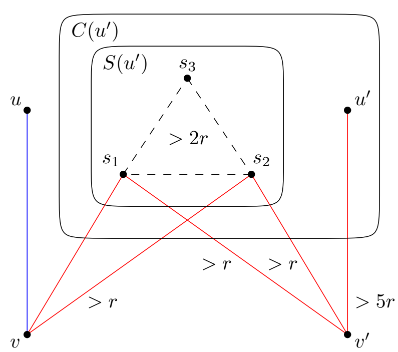

Figure 4: A visualization of Claim 3.6. Edges are annotated with the distances given by the pseudometric dist. A blue edge denotes an E connection. A red edge denotes a $\neg E$ connection. A dashed edge is used when only the distance is of relevance.

Consequently, $C(u')$ and $C(v)$ cannot both be large as it cannot be the case that simultaneously E and $\neg E$ generically hold between them.

We now show that we also arrive at a contradiction if both $C(u')$ and $C(v)$ are small. Later we will consider the case when one of them is small and the other one is large.

Claim 3.7. *The following hold:*

1. *If* $C(u')$ *is small, then* $\neg E(s, v)$ *holds for all* $s \in S(u')$*.*
2. *If* $C(v)$ *is small, then* $E(u', s)$ *holds for all* $s \in S(v)$*.*

Proof. Again we prove the first item, as the other one follows by symmetry. The following situation is depicted in Figure 5.

Fix $s \in S(u')$. Observe that $\mathrm{dist}(v', s) > r$. Indeed, suppose $\mathrm{dist}(v', s) \leqslant r$. As $C(u')$ is small, $\mathrm{dist}(s, u') \leqslant 4r$. Together this gives $\mathrm{dist}(v', u') \leqslant 5r$, a contradiction.

As $\neg E$ generically holds between $C(u')$ and $C(v')$, it follows that $\neg E(s, v')$ holds. Since $E(S, v) = E(S, v')$, we get that $\neg E(s, v)$ holds. $\square$

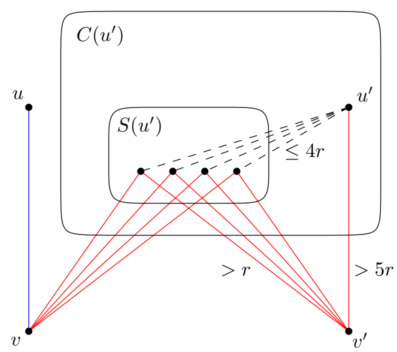

Figure 5: A visualization of Claim 3.7. The same notation as in Figure 4 is used.

Suppose both $C(u')$ and $C(v)$ are small. Then $E(u', s)$ holds for all $s \in S(v)$, and $\neg E(s, v)$ holds for all $s \in S(u')$, contradicting the construction of $S_{C(v)C(u')}$, as we have that either:

- for every $b \in C(u')$ there is some $a \in S(v)$ with $\neg E(a, b)$, a contradiction to u' being E-connected to every vertex from $S(v)$, or
- for every $a \in C(v)$ there is some $b \in S(u')$ with $E(a, b)$, a contradiction to v being E-connected to no vertex from $S(u')$.

So we are left with the case when exactly one of $C(u')$ and $C(v)$ is small. By symmetry, we may assume that $C(u')$ is small: otherwise $C(v)$ is small and, up to replacing $E(x, y)$ with $\neg E(y, x)$ and u, v, u', v' with v', u', v, u, we are in the same case.

So $C(u')$ is small and $C(v)$ is large. Then by Claim 3.6, E generically holds between $C(u')$ and $C(v)$. And by Claim 3.7, $\neg E(s, v)$ holds for all $s \in S(u')$.

Claim 3.8. *We have* $\mathrm{dist}(u', v) \leqslant 3r$*.*

Proof. Since $\neg E(s, v)$ holds for all $s \in S(u')$, in particular for the selected center $c_0 \in S(u')$ of the small class $C(u')$ we have that $\neg E(c_0, v)$ holds. Since E generically holds between $C(u')$ and $C(v)$, it must be the case that $\text{dist}(c_0, v) \leqslant r$. Together with $\text{dist}(u', c_0) \leqslant 2r$ this yields $\text{dist}(u', v) \leqslant 3r$. □

Since $\neg E(s, v)$ holds for all $s \in S(u')$, by construction of $S_{C(v)C(u')}$ there is some $s \in S(v)$ such that $\neg E(u', s)$ holds. Then also $\neg E(u, s)$ holds, as $E(u, S) = E(u', S)$. As E generically holds between $C(u')$ and $C(v)$, it follows that $\text{dist}(u', s) \leqslant r$. For a similar reason, $\text{dist}(u, s) \leqslant r$. Hence $\text{dist}(u, u') \leqslant 2r$. With Claim 3.8 this yields $\text{dist}(u, v) \leqslant 5r$, a contradiction. □

Lemma 3.1 now follows from Theorem 3.5.

Proof of Lemma 3.1. Let $\mathcal{C}$ and φ be as in the assumptions of the lemma and let d be the bound on the VC-dimension of φ on $\mathcal{C}$. Let $G \in C$ and $U \subseteq V(G)$. By Corollary 2.1 we know that there exist numbers r' and t such that $V(G)$ can be colored by at most t colors such that for all vertices u, v of G with $\text{dist}^G(u, v) > r'$, $\psi(u, v)$ depends only on the colors of u and v. Let $H = \mathsf{I}_\varphi(G)[U]$ and for each color class $C_i \subseteq V(G)$ with $i \in [t]$ let D_i be its restriction to the graph H, i.e. $D_i = C_i \cap U$.

We then have the following:

- For every $A \subseteq V(H)$ and $B \subseteq V(H)$ the VC-dimension of $(A \times B) \cap E(H)$ is bounded, and by Theorem 3.3 therefore $(A \times B) \cap E(H)$ has a duality of order k depending only on d.
- The function dist on $V(H) \times V(H)$ defined by setting $\text{dist}(u, v) := \text{dist}_G(u, v)$ for every $u, v \in V(H)$ is a pseudometric.
- $\mathcal{P} = \{D_i, \dots, D_t\}$ is a partition $V(H)$ into sets such $uv \in E(H)$ depends only on the classes of u and v in $\mathcal{P}$ for every $u, v \in V(H)$ with $\text{dist}(u, v) > r'$.

We can therefore apply Theorem 3.5 to obtain a subset S of $V(H)$ of size $\mathcal{O}(kt^2)$ such that $uv \in E(H)$ depends only on $E(u, S)$ and $E(S, v)$ for all $u, v \in V$ with $\text{dist}(u, v) > 5r'$.

Since $V(H) = U$ and $uv \in E(H)$ if and only if $G \models \varphi(u, v)$, this concludes the proof after setting $r := 5r'$. □

4 Model checking on interpretations of classes bounded local cliquewidth

In this section we prove the main result of the paper. Before we get started, we need to fix some notation.

4.1 Graph classes

We work with classes $\mathcal{C}$ of graphs that are possibly equipped with unary predicates, constants, and *flags*, that is, relation symbols of arity 0 (a flag f therefore evaluates to a Boolean $f_G \in \{true, false\}$, for each structure $G \in \mathcal{C}$). More precisely, each class $\mathcal{C}$ has a fixed finite signature Σ which contains the binary relation symbol E, and relation symbols of arity 0 or 1, and constant symbols. Moreover, E is interpreted as a symmetric, irreflexive relation in each $G \in \mathcal{C}$. By abuse of language, we call structures in $\mathcal{C}$ *graphs*. We will usually not mention the signature of a graph class explicitly, unless necessary. We say that a class $\mathcal{C}$ as above has *bounded local cliquewidth* if the class of underlying (usual) graphs has bounded local cliquewidth.

If Σ and Γ are two signatures with $\Sigma \subseteq \Gamma$, and G is a Σ-structure, then any Γ-structure G' obtained from G by interpreting the symbols from Γ not in Σ is called a Γ-*expansion* of G.

4.2 (Local) Cliquewidth

We assume familiarity with the notions of treewidth and of cliquewidth. We denote the cliquewidth of a graph G by $\mathrm{cw}(G)$. We will need the following results.

Theorem 4.1 ([7]). *Let $\mathcal{C}$ be a class of graphs which is interpretable in a graph class of bounded cliquewidth. Then $\mathcal{C}$ is of bounded cliquewidth.*

Theorem 4.2 ([7]). *There is a function $h\colon \mathbb{N} \to \mathbb{N}$, a constant c, and an algorithm that, given a (colored) graph G and a sentence $\varphi \in \mathrm{FO}$ decides whether $G \models \varphi$ in time $h(\mathrm{cw}(G) + |\varphi|) \cdot |G|^c$.*

We will also need the localized variant of cliquewidth. If G is a graph, $v \in V(G)$, and $r \geqslant 0$, then we denote by $N_r^G[v]$ the set of vertices in G of distance at most r from v. For $\bar{v} \in V(G)^l$ we define $N_r^G[\bar{v}] := \bigcup_{i=1}^{l} N_r^G[v_i]$.

Definition 4.3. *Let G be a graph. For $r \geqslant 0$ we define*

$$\mathrm{lcw}_r(G) := \max\{\mathrm{cw}\big(G[N_r^G[v]]\big) : v \in V(G)\}.$$

We say that a class $\mathcal{C}$ of graphs has bounded local cliquewidth *if there is a function $f : \mathbb{N} \to \mathbb{N}$ such that $\mathrm{lcw}_r(G) \leqslant f(r)$ for all $G \in \mathcal{C}$ and $r \geqslant 0$.*

Classes of graphs of bounded local cliquewidth include all classes of bounded local treewidth (defined analogously) such as classes of graphs of bounded degree, the class of planar graphs or more generally classes of graphs embedded on a surface of fixed genus. On the other hand, the class of apex graphs, i.e., graphs G which are planar after removal of a single vertex, does not have bounded local treewidth or cliquewidth. Another classical example of classes of graphs of bounded local cliquewidth are map graphs. The following result has its roots in the work of Frick and Grohe [13] (see also [9, 19]).

Theorem 4.4. *There is a function $h\colon \mathbb{N} \to \mathbb{N}$, a constant c, and an algorithm that, given a (colored) graph G and a sentence $\varphi \in \mathrm{FO}$ decides whether $G \models \varphi$ in time*

$$h(\mathrm{lcw}_{h(|\varphi|)}(G) + |\varphi|) \cdot |G|^c.$$

The theorem follows from the model checking algorithm for bounded cliquewidth as well as standard techniques, nicely presented in [19, Theorem 4.5]. Note that this algorithm yields correct answers on all classes of graphs, however it is only efficient on classes where the local cliquewidth is bounded.

We also need some notation related to Gaifman locality. A first-order formula $\varphi(x_1, \dots, x_l)$ is *r-local* if for every graph G and l-tuple $\bar{v} \in V(G)^l$, $G \models \varphi(\bar{v}) \Leftrightarrow G[N_r^G[\bar{v}]] \models \varphi(\bar{v})$.

Corollary 4.5. *For every formula $\psi(x, y)$ there exist numbers r and q such that the following holds: For every graph G there exists an r-local formula $\psi_G(x, y)$ of quantifier rank at most q such that for all vertices $u, v \in V(G)$ we have that $G \models \psi_G(u, v)$ if and only if $G \models \psi(u, v)$.*

We use Gaifman to prove that interpretations of bounded local cliquewidth have bounded local cliquewidth, too.

Lemma 4.6. *Let $\mathcal{C}$ be a class of graphs of bounded local cliquewidth and let $\mathsf{I}_{\varphi,\delta}$ be an interpretation such that there exists d such that $\varphi(x, y)$ has range at most d. Then $\mathsf{I}_{\varphi,\delta}(\mathcal{C})$ is a class of graphs of bounded local cliquewidth.*

Proof. Our task is to prove that for every $H \in \mathsf{I}_{\varphi,\delta}(\mathcal{C})$, vertex $v_0 \in V(H)$ and every r the graph $H[N_r^H[v_0]]$ has cliquewidth bounded in terms of r.

Let $G \in \mathcal{C}$ be such that $H = \mathsf{I}_{\varphi,\delta}(G)$. This means that $H = \mathsf{I}_\varphi(G)[U]$, where $U = \{u \in V(G) \mid G \models \delta(u)\}$. In particular, H is an induced subgraph of $\mathsf{I}_\varphi(G)$.

Let r' and $\varphi'(x, y)$ be the locality parameter and r'-local formula obtained by applying Gaifman's theorem in the form of Corollary 4.5 to $\varphi(x, y)$. Set $\ell := rd + r'$ and $G_0 = G[N_\ell^G[v_0]]$. We will consider the graph $\mathsf{I}_{\varphi'}(G_0)$ and show that

- $\mathsf{I}_{\varphi'}(G_0)$ has cliquewidth bounded by a function of r (here we consider d and r' to be fixed constants), and
- $H[N_r^H[v_0]]$ is an induced subgraph of $\mathsf{I}_{\varphi'}(G_0)$.

The lemma then follows because cliquewidth is preserved by taking induced subgraphs.

To show the first item we first note that G_0 has cliquewidth at most $g(\ell)$, where g is the local cliquewidth bounding function for $\mathcal{C}$. Let $\mathcal{C}_{g(\ell)}$ be the class of all graphs of cliquewidth at most $g(\ell)$. We then have that $\mathsf{I}_{\varphi'}(G_0) \in \mathsf{I}_{\varphi'}(\mathcal{C}_\ell)$, The claim then follows by Theorem 4.1.

It remains to show that $H[N_r^H[v_0]]$ is an induced subgraph of $\mathsf{I}_{\varphi'}(G_0)$, i.e., that $N_r^H[v_0] \subseteq V(G_0)$ and for all $u, v \in N_r^H[v_0]$ it holds that $uv \in E(H)$ if and only if $G_0 \models \varphi'(u, v)$. Let u, v be two vertices in $N_r^H[v_0]$. Since both u and v are at distance at most r from v_0 in H, by our assumption on the range of $\varphi(x,y)$ they are at distance at most dr from v_0 in G. Indeed, $ab \in E(H)$ is equivalent to $G \models \varphi(a,b)$, which implies, by assumption, $\mathrm{dist}^G(a,b) < d$. This means that $u,v \in N^G_{rd}[v_0]$ and so both u and v are in $V(G_0)$. Moreover, every vertex at distance at most r' from u or v in G is at distance at most $rd + r'$ from v in G, and so $N^G_{r'}[u] \cup N^G_{r'}[v] \subseteq N^G_\ell[v_0]$. Thus for the r'-local formula $\varphi'(x,y)$ it holds that $G_0 \models \varphi'(u,v)$ if and only if $G \models \varphi'(u,v)$, and from Corollary 4.5 we know that $G \models \varphi'(u,v)$ if and only if $G \models \varphi(u,v)$. We therefore get $G_0 \models \varphi'(u,v)$ if and only if $uv \in E(H)$, as desired. □

We will rely on the following theorem proved by Grohe and Turán [21, Lemma 22].

Theorem 4.7. *Let $\mathcal{C}$ be a class with bounded local cliquewidth. Then $\mathcal{C}$ is NIP.*

Lemma 4.6 and Theorem 4.7 both hold when $\mathcal{C}$ is a class of graphs equipped with unary predicates, constants, and flags.

4.3 Flips

For a graph G and $m \in \mathbb{N}$, an *m-flip* is an operation determined by a partition of $V(G)$ into sets $V_1, \dots, V_{m'}$ with $m' \leqslant m$ and a symmetric binary relation R on $[m]$. The resulting graph has the same vertex set as G and its edge relation is obtained from $E(G)$ by complementing the edges between any $x \in V_i$, $y \in V_j$ such that $(i,j) \in R$. Note that it can be $i = j$. We will call the output of an m-flip operation on a graph G also an *m-flip of G*. Also note that flips are reversible, that is if H is an m-flip of G, then G is an m-flip of H. Let $S \subseteq V(G)$. We say that an m-flip is *guarded* by S if each of the sets $V_1, \dots, V_{m'}$ is of the form $\{v \in V(G) \mid N_G(v) \cap S = A\}$ for some $A \subseteq S$. Note that in this case $m' \leqslant 2^{|S|}$.

Lemma 4.8. *Let $\mathcal{C}$ be a NIP class of graphs, $\mathsf{I}_{\varphi,\delta}$ be an interpretation, and $\mathcal{D} = \mathsf{I}_{\varphi,\delta}(\mathcal{C})$. There exist $s, r \in \mathbb{N}$, a signature Γ expanding the signature of graphs by constant symbols and relation symbols of arity 0, a formula $\psi(x,y)$ in the signature Γ that has range at most r, such that the following holds. For every $H \in \mathcal{D}$ there exists a graph $F(H)$ and a graph $\widehat{G}(H)$ in the signature Γ, such that:*

- *$F(H)$ is a 2^s-flip of H, guarded by a set $S \subseteq V(H)$ of size at most s,*
- *$\widehat{G}(H)$ is a Γ-expansion of some graph $G \in \mathcal{C}$,*
- *$F(H) = \mathsf{I}_{\psi,\delta}(\widehat{G}(H))$.*

Proof. Since $\mathcal{C}$ is NIP and therefore has bounded VC-dimension, we can apply Lemma 3.1 to $\mathcal{C}$ and φ to obtain numbers r and s with the properties claimed there.

Let $G \in C$ be such that $H = \mathsf{I}_{\varphi,\delta}(G)$. Let $U = \{v \in V(G) \mid G \models \delta(v)\}$. Note that U is the vertex set of H, meaning we have $H = \mathsf{I}_\varphi(G)[U]$. For any $u, v \in V(H)$ it therefore holds that $G \models \varphi(u,v)$ if and only if $uv \in E(H)$. Since $\varphi(x,y)$ is symmetric, for any $u \in U$ and $X \subseteq U$ it holds that $\varphi(u,X) = \varphi(X,u) = N_H(u) \cap X$. Lemma 3.1 then states the following: there is a set $S \subseteq V(H)$ of size at most s such that for all $u, v \in V(H)$ with $\mathrm{dist}^G(u,v) > r$, whether or not $uv \in E(H)$ holds depends only on $N_H(u) \cap S$ and $N_H(v) \cap S$.

We now describe the flip $F(H)$ of H from the statement of the lemma. Let $c_1, \ldots, c_s$ be an enumeration containing all the elements of S (possibly with repetitions, and possibly also containing some elements of $V(H) - S$). This exists, since $|S| \leqslant s$, and we may assume that H has at least one vertex, the other case being trivial. For every set $A \subseteq \{1, \ldots, s\}$, we define $V_A = \{v \in V(H) \mid N_H(v) \cap S = \{c_a \mid a \in A\}\}$. This determines a partition of $V(H)$ into exactly 2^s sets. Define a relation $R \subseteq 2^{[s]} \times 2^{[s]}$ as follows: $(A, B) \in R$ if and only if there are $u \in V_A$ and $v \in V_B$ such that $\mathrm{dist}^G(u,v) > r$ and $uv \in E(H)$. Let $F(H)$ be the 2^s-flip of H determined by R; it is guarded by S, and $|S| \leqslant s$. Thus, the first statement of the lemma is satisfied.

We next describe the graph $\widehat{G}(H)$ and the formula ψ. The graph $\widehat{G}(H)$ is obtained from G by

- Marking the elements $c_1, \ldots, c_s$ using s constant symbols, which we also denote $c_1, \ldots, c_s$.
- Encoding the relation R using $2^{[s]} \times 2^{[s]}$ flags $f_{A,B}$, for $A, B \subseteq \{1, \ldots, s\}$. Namely, $f_{A,B}$ is set to true in $\widehat{G}(H)$, for $A, B \subseteq \{1, \ldots, s\}$ if and only if the pair (A, B) belongs to the relation R.

Note that since G and $\widehat{G}(H)$ have the same vertices and edges, the distances between vertices are the same in G and $\widehat{G}$. Also note that since the signature of G is a subset of the signature of $\widehat{G}(H)$, the formulas $\delta(x)$ and $\varphi(x,y)$ can be evaluated in $\widehat{G}(H)$. To define the formula ψ, first define the following formulas. For $A \subseteq [s]$, let $P_A(x)$ be the formula

$$P_A(x) := \delta(x) \wedge \bigwedge_{i \in A} \varphi(x, c_i) \wedge \bigwedge_{i \in [s]-A} \neg\varphi(x, c_i),$$

expressing that x belongs to the part V_A. Further, let

$$\alpha_R(x,y) = \bigvee_{(A,B) \subseteq 2^{[s]} \times 2^{[s]}} f_{A,B} \wedge (P_A(x) \wedge P_B(y))$$

be the formula "encoding" the edges flipped according to R, so that $\widehat{G}(H) \models \alpha_R(u,v)$ holds for $u, v \in U = V(H)$ such if and only if $(A, B) \in R$, where V_A is the part containing u and V_B is the part containing v. In particular, by definition of R,

$$\widehat{G}(H) \models (\alpha_R(u,v) \leftrightarrow \varphi(u,v)) \qquad \text{for all } u, v \in U \text{ with } \mathrm{dist}^G(u,v) > r. \tag{1}$$

We now set $\psi(x,y) := (\varphi(x,y) \oplus \alpha_R(x,y)) \wedge (\mathrm{dist}(x,y) \leqslant r)$ where $\oplus$ denotes the *xor* operation, and $\mathrm{dist}(x,y) \leqslant r$ is the formula expressing the existence of a path of length at most r from x to y in the underlying graph (here is r is a fixed constant). By construction, ψ has range at most r. Moreover, by (1) we have

$$\widehat{G}(H) \models \psi(x,y) \leftrightarrow (\varphi(x,y) \oplus \alpha_R(x,y)). \tag{2}$$

It remains to argue that $F(H) = \mathsf{I}_{\psi,\delta}(\widehat{G})$. Clearly the vertex set of $\mathsf{I}_{\psi,\delta}(\widehat{G})$ is the same as $V(H)$, as both are equal to U. To see that they have the same edges, recall that $\alpha_R(x,y)$ encodes the flip which was used to obtain $F(H)$ from H, and the formula $\varphi(x,y) \oplus \alpha_R(x,y)$ can be viewed in the exactly the same way – it first introduces all edges of H via $\varphi(x,y)$ and then flips away exactly the edges $uv \in E(H)$ with $\mathrm{dist}^G(u,v) > r$. By (2), this proves $F(H) = \mathsf{I}_{\psi,\delta}(\widehat{G})$, as required. $\square$

4.4 Proof of the main theorem

At last, we can prove Theorem 1.3, which we restate here for convenience.

Theorem 1.3. *The first-order model checking problem is fixed-parameter tractable on any class $\mathcal{D}$ of graphs that is interpretable in a class $\mathcal{C}$ of graphs of bounded local cliquewidth.*

Proof. Let $\varphi(x,y)$ and $\delta(x)$ be the formulas defining an interpretation such that $\mathcal{D} \subseteq \mathsf{I}_{\varphi,\delta}(\mathcal{C})$. Since every class of graphs of bounded local cliquewidth is NIP by Theorem 4.7, we can apply Lemma 4.8 to $\mathcal{C}$ and $\mathsf{I}_{\varphi,\delta}$. This yields numbers r and s, a formula $\psi(x,y)$ of range at most r, and for every graph $H \in \mathcal{D}$ a graph $F(H)$ and an expansion $\widehat{G}(H)$ of some graph $G \in \mathcal{C}$, with the properties claimed there. We first establish a crucial property of graphs $F(H)$.

Claim 4.9. *The class $\mathcal{D}' := \{F(H) \mid H \in \mathcal{D}\}$ has bounded local cliquewidth.*

Proof. By Lemma 4.8, we have $\mathcal{D}' = \{\mathsf{I}_{\psi,\delta}(\widehat{G}(H)) \mid H \in \mathcal{D}\}$. As the class $\{\widehat{G}(H) \mid H \in \mathcal{D}\}$ consists of expansions of graphs from $\mathcal{C}$ by constant symbols and flags, it has has bounded local cliquewidth. As $\mathcal{D}'$ is the image of this class under the interpretation of $\mathsf{I}_{\psi,\delta}$, the claim follows from Lemma 4.6. ⌟

We now describe the model checking algorithm. Assume we want to determine whether $H \models \rho$ for some $H \in \mathcal{D}$ and $\rho \in \mathrm{FO}$. Let $\mathcal{S} = \{(S,R) \mid S \subseteq V(G), |S| \leqslant s, R \subseteq 2^S \times 2^S, R \text{ symmetric}\}$. Note that $|\mathcal{S}| = \mathcal{O}_s(|V(H)|^s)$ (meaning $|\mathcal{S}| \leqslant c_s \cdot |V(H)|^s$ for some constant c_s depending only on s). We will generate a collection of graphs $H_{S,R}$, $(S,R) \in \mathcal{S}$, colored with at most 2^s colors, together with sentences $\rho_{S,R}$ of length $\mathcal{O}_s(|\rho|)$ such that $H \models \rho \Leftrightarrow H_{S,R} \models \rho_{S,R}$, and for at least one "correct" choice $(S,R) \in \mathcal{S}$, $H_{S,R}$ is a vertex-coloring of $F(H)$, and therefore is a vertex-coloring of some graph in $\mathcal{D}'$. This then implies the theorem as follows.

Let $\widehat{\mathcal{D}}'$ be the class of all graphs from $\mathcal{D}'$ equipped with 2^s new unary predicates. Since at least one $H_{S,R}$ is a vertex-coloring of a graph in $\mathcal{D}'$, at least one $H_{S,R}$ is contained in $\widehat{\mathcal{D}}'$.

We now run the model checking algorithm from Theorem 4.4 for all $(S, R) \in \mathcal{S}$ in parallel to determine whether $H_{S,R} \models \rho_{S,R}$, and terminate the whole process once the first run stops. By Theorem 4.4, for any choice of $H_{S,R}$, $\rho_{S,R}$ this algorithm stops in time at most $h(\mathrm{lcw}_{h(|\rho_{S,R}|)}(H_{S,R}) + |\rho_{S,R}|) \cdot |H_{S,R}|^c$. For the "correct" choice $(S, R) \in \mathcal{S}$ (such that $H_{S,R}$ is a vertex-coloring of $F(H)$), $\mathrm{lcw}_w(H_{S,R}) = \mathrm{lcw}_w(F(H))$ for all $w \in \mathbb{N}$, and since $F(H)$ comes from a class of bounded local cliquewidth, $\mathrm{lcw}_w(F(H)) \leqslant f(w)$ for some function f depending only on $\mathcal{D}'$ (where $\mathcal{D}'$ depends only on φ, δ and $\mathcal{D}$). Moreover, $|\rho_{S,R}| \leqslant \mathcal{O}_s(|\rho|)$, so in total for this run the algorithm stops after at most $h(f(h(\mathcal{O}_s(|\rho|))) + |\rho_{S,R}|) \cdot |H|^c$ steps. Because there are $\mathcal{O}_s(|V(H)|^s)$ of runs executed in parallel, this bounds the run time of our algorithm by $\mathcal{O}_{s,|\rho|}(|H|^{c+s})$. Since s depends on $\mathcal{D}$ but not on ρ, we obtain FPT run time as desired.

We now describe the construction of the graphs $H_{S,R}$ and sentences $\rho_{S,R}$. For each set $S \subseteq V(H)$ of size at most s and symmetric binary relation R over the color set 2^S, define a colored graph $H_{S,R}$ obtained from H as follows. First, color each vertex $u \in V(H)$ by $\lambda_S(u) := N_H(u) \cap S$. Then, flip the adjacency (that is, an edge becomes a non-edge and vice versa) between every pair $u, v \in V(H)$ if and only if $(\lambda_S(u), \lambda_S(v)) \in R$. Since we go through all subsets S of $V(H)$ of size at most s and all possible flips guarded by S, for some choice of S and R, the graph $H_{S,R}$ is a coloring of $F(H)$.

To describe $\rho_{S,R}$, first let us consider the formula $\zeta_{S,R}$ which flips the edges of $H_{S,F}$ back to obtain H, given by

$$\zeta_{S,R}(x, y) := x \neq y \wedge \bigvee_{(\alpha,\beta) \in R} (c_\alpha(x) \wedge c_\beta(y)),$$

where $\alpha, \beta \in 2^S$ and the meaning of the predicate $c_\alpha(u)$ is that $\lambda_S(u) = \alpha$. Then $\rho_{S,R}$ is obtained from ρ by replacing each occurrence of $E(x, y)$ by $E(x, y) \oplus \zeta_{S,R}(x, y)$. We have $H \models \rho \Leftrightarrow H_{S,R} \models \rho_{S,R}$. $\square$

5 Discussion

We now discuss how our results fit into a broader picture. For the purpose of this discussion, it is slightly more convenient to replace interpretations with the more general *transductions*, which are defined below. *Simple non-copying transductions* (here, *transductions* for short) are defined similarly as interpretations. First, they may nondeterministically color the input graph G with a fixed number of colors, and afterwards they apply a fixed interpretation to the obtained colored graph, yielding an output graph. Thus, a transduction maps a single graph G to a set of possible output graphs, where the various possible outputs correspond to the various possible colorings. Say that a class $\mathcal{C}$ *transduces* a class $\mathcal{D}$, or that $\mathcal{D}$ *transduces in* $\mathcal{C}$, if there is a transduction T such that $\mathcal{D} \subseteq T(\mathcal{C})$. As in the case of interpretations, this defines a transitive relation. If the interpretation applied by the transduction T comes from a logic $\mathcal{L}$, we say that T is an $\mathcal{L}$-transduction.

In the previous section, replacing interpretations with transductions would not make a difference in most places. In particular, a class $\mathcal{C}$ is a transduction of a class with bounded local cliquewidth if and only if it is an interpretation of such a class, so our main result also holds for the more general notion. Such a replacement is not neutral in all contexts, however.

For a property $\mathcal{P}$ of graph classes, we say a class $\mathcal{D}$ of graphs has *structurally* $\mathcal{P}$, if $\mathcal{D}$ FO-transduces in some class $\mathcal{C}$ with property $\mathcal{P}$. So for example, a class $\mathcal{C}$ is structurally nowhere dense if it transduces (equivalently, interprets) in some nowhere dense class $\mathcal{D}$. Classes that transduce (equivalently, interpret) in a class with bounded local cliquewidth are exactly classes with structurally bounded local cliquewidth, and our main result concerns those classes.

A reformulation of Question 1.2, generalized to transductions instead of interpretations, therefore asks: are structurally tractable classes tractable? Let us evaluate the status of this question by listing classes $\mathcal{C}$ that are known to admit an FPT algorithm for FO model checking and discussing what can be said about transductions $\mathcal{D}$ thereof. See Figure 2 for an overview.

Particularly interesting cornerstones in this context are *transduction ideals*. We use this term to denote properties of hereditary graph classes, that are preserved by (first-order) transductions. By transitivity of the transduction relation, for every property $\mathcal{P}$ of graph classes, the property "structurally $\mathcal{P}$" forms a transduction ideal.

Question 1.2 suggests the existence of an FPT algorithm for FO model checking on structurally nowhere dense classes. Up to now, this has been only confirmed for classes of structurally bounded degree, and for classes of bounded shrubdepth (that is, transductions of classes of trees of bounded depth). Our main result in particular implies that the same holds for every class with structurally bounded local treewidth.

Besides classes of structurally bounded degree and classes of bounded shrubdepth, classes with structurally bounded local treewidth include structurally planar classes, classes with structurally bounded genus, and structurally apex-minor-free graph classes.

A next step would be to consider classes with structurally bounded expansion, which are strictly weaker than structurally nowhere dense classes. They do not include all classes with structurally bounded local treewidth, however (see Figure 1).

We note that the proof of tractability of nowhere dense classes [20] is based on an iterative application of locality arguments, combined with structural properties of nowhere dense classes. In particular, every r-ball in a graph from a nowhere dense class $\mathcal{C}$ belongs to a nowhere dense class that is simpler in some sense, as is formalized by the notion of splitter games [20]. Therefore, it is conceivable that an extension of our methods will allow to approach the problem of tractability of structurally nowhere dense classes.

Monadically stable classes. Structurally nowhere dense classes are further generalized by *monadically stable* classes. A class is monadically stable if it does not transduce[3] the class of all half-graphs. In other words, monadically stable classes form the largest transduction ideal that does not contain the class of half-graphs. Monadically stable classes were introduced by Baldwin and Shelah [2], and are a special case of *stable classes*, which are one of the central objects of interest in stability theory. Stability theory is now the main focus of model theory.

There are strong connections between stability theory and (structurally) sparse graph classes. Most notably, it was shown by Podewski and Ziegler [29] in the late 70's, long before the development of sparsity theory, that all nowhere dense classes (called *superflat* in their paper) are monadically stable.

The result of Podewski and Ziegler, connecting sparsity theory with stability theory, has been brought to the attention of the sparsity community by Adler and Adler [1], who observed that nowhere dense classes are the same as superflat classes. By the result of Podewski and Ziegler, nowhere dense classes, and therefore also structurally nowhere dense classes, are monadically stable. In the other direction, it is conjectured [26] that a graph class is monadically stable if and only if it is structurally nowhere dense.

Unstable classes. Classes with bounded cliquewidth are not necessarily monadically stable, as the class of all half-graphs has bounded cliquewidth and is not monadically stable by definition. On the other hand, the class of planar graphs is nowhere dense, and hence monadically stable, but has unbounded cliquewidth. Bounded cliquewidth is therefore incomparable to monadically stable (or nowhere dense) classes. Nevertheless, FO (and even MSO_1) model checking is FPT on these classes. Bounded cliquewidth also forms a transduction ideal (even for MSO transductions). Thus, if $\mathcal{C}$ has bounded cliquewidth and $\mathcal{D}$ is a transduction thereof, then $\mathcal{D}$ also has bounded cliquewidth.

Classes with bounded local cliquewidth are tractable (see Theorem 1.1), but do not form a transduction ideal, as they are not closed under edge-complementation. By our main result, all classes with structurally bounded local cliquewidth are also tractable, and those do form a transduction ideal.

Recently, Bonnet, Kim, Thomassé and Watrigant [5] introduced the notion of *twinwidth* and showed that classes of bounded twinwidth are preserved by FO transductions. So bounded twinwidth is a transduction ideal, and subsumes bounded cliquewidth, but is incomparable to structurally bounded local cliquewidth[4]. Moreover, FO model checking

[3]This is one of two places where the distinction between transductions and interpretations matters. For example, consider the class $\mathcal{C}$ of graphs that can be obtained from a clique by placing a vertex in the middle of every edge. Then $\mathcal{C}$ does not interpret the class of half-graphs, but $\mathcal{C}$ transduces the class of all graphs, since we can first color some subset of the middle vertices, and in this way encode any graph. The other place where the distinction matters is in the definition of monadically NIP classes, and the same example illustrates the issue.

[4]The class of cubic graphs has bounded local cliquewidth, but unbounded twinwidth [4]. On the other

is FPT on classes with bounded twinwidth, but only assuming an appropriate decomposition is given as additional input [5].

Let us stress that our algorithm captures all known transduction ideals for which the model checking problem is FPT (without an additional decomposition given as input).

(Monadically) NIP classes. The following notion, encompassing all the graph classes mentioned above, again originates in stability theory – despite its name, stability theory does not only concern stable classes. A class $\mathcal{C}$ is *monadically NIP*, or *monadically dependent*, if it does not transduce the class of all graphs. In other words, monadically NIP classes constitute the largest transduction ideal, apart from the one that contains all classes. All the aforementioned graph classes are monadically NIP: nowhere dense classes, classes of structurally bounded local cliquewidth, classes of bounded twinwidth, etc. It is conjectured [16, Conjecture 8.2] that FO model checking is FPT on all monadically NIP classes. Every monadically NIP class $\mathcal{C}$ is in particular NIP, that is, every formula $\varphi(x, y)$ has bounded VC-dimension on $\mathcal{C}$. Hence, Lemma 2.3 applies to all such classes.

With essentially the same proof as for Theorem 1.3, we can obtain the following, more general statement. We say the FO model checking problem is *conservatively FPT* on a class of structures $\mathcal{C}$ if there is an algorithm that, for every FO formula φ and structure $G \in \mathcal{C}$, decides whether $G \models \varphi$, and runs in time $f(\varphi) \cdot n^c$ for every $G \in \mathcal{C}$ with n elements. In contrast, an FPT model checking algorithm on $\mathcal{C}$ is not required to give correct answers for structures outside $\mathcal{C}$. All the FPT FO model checking algorithms we discussed so far are also conservatively FPT.

Let Σ be the signature of graphs. A formula $\varphi(x, y)$ has *bounded range* if it has range $\leqslant r$, for some $r \in \mathbb{N}$. An interpretation $\mathsf{I}\colon \Sigma \to \Gamma$, where Γ is a signature consisting of unary and binary relation symbols, has *bounded range* if for all binary symbols $R \in \Gamma$, the formula $\varphi_R(x, y)$ has bounded range. An interpretation *with parameters* $\mathsf{I}\colon \Sigma \to \Gamma$ is an interpretation $\mathsf{I}'\colon \Sigma' \to \Gamma$, where Σ' expands Σ with constant symbols and relation symbols of arity 0 (flags). For such an interpretation and class $\mathcal{C}$ of Σ-structures, write $\mathsf{I}(\mathcal{C})$ for the class of all structures of the form $\mathsf{I}'(\mathbb{A}')$, where $\mathbb{A}'$ is a Σ'-structure expanding some structure $\mathbb{A} \in \mathcal{C}$, by providing an interpretation of each constant symbol in Σ' and not in Σ. A class of *colored graphs* is a class of structures over a signature $\Gamma = \{E, U_1, \ldots, U_k\}$, where $U_1, \ldots, U_k$ are unary relation symbols, and E is interpreted as a binary symmetric, irreflexive relation.

Theorem 5.1. *Let $\mathcal{C}$ be an NIP class of graphs. The first-order model checking problem is conservatively FPT on graph classes that interpret in $\mathcal{C}$, if it is conservatively FPT on classes of colored graphs that interpret in $\mathcal{C}$ via a bounded-range interpretation with parameters.*

---

hand, consider the class $\mathcal{C}$ of graphs G such that each connected component of G is a grid with an added apex vertex. Then $\mathcal{C}$ has bounded twinwidth, but does not have structurally bounded local cliquewidth.

Bounded-range interpretations (with parameters) of classes with bounded local cliquewidth again have bounded local cliquewidth and therefore a conservative FPT model checking algorithm. Therefore, Theorem 1.3 merely describes part of a bigger picture painted by Theorem 5.1. We believe it will serve as a crucial tool towards answering Question 1.2 in other cases.

References

[1] Hans Adler and Isolde Adler. Interpreting nowhere dense graph classes as a classical notion of model theory. *Eur. J. Comb.*, 36:322–330, 2014.

[2] John Baldwin and Saharon Shelah. Second-order quantifiers and the complexity of theories. *Notre Dame Journal of Formal Logic*, 29, 07 1985.

[3] Achim Blumensath and Bruno Courcelle. Recognizability, hypergraph operations, and logical types. *Information and Computation*, 204(6):853–919, 2006.

[4] Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width II: small classes. In *Proc. of the ACM-SIAM Symposium on Discrete Algorithms (SODA 2021)*, pages 1977–1996, 2021.

[5] Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-Width I: Tractable FO Model Checking. *J. ACM*, 69(1):1–46, nov 2021.

[6] Bruno Courcelle. The monadic second-order logic of graphs I: Recognizable sets of finite graphs. *Inform. and Comput.*, 85:12–75, 1990.

[7] Bruno Courcelle and Joost Engelfriet. *Graph Structure and Monadic Second-Order Logic: a Language-Theoretic Approach*, volume 138 of *Encyclopedia of Mathematics and Its Applications*. Cambridge University Press, 2012.

[8] Bruno Courcelle, Johann A. Makowsky, and Udi Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000.

[9] Anuj Dawar, Martin Grohe, and Stephan Kreutzer. Locally excluding a minor. In *LICS'07*, pages 270–279. IEEE Computer Society, 2007.

[10] Jan Dreier. Lacon- and shrub-decompositions: A new characterization of first-order transductions of bounded expansion classes. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*, pages 1–13. IEEE, 2021.

[11] Zdenek Dvořák, Daniel Kráľ, and Robin Thomas. Deciding first-order properties for sparse graphs. In *FOCS'10*, pages 133–142. IEEE Computer Society, 2010.

[12] Jörg Flum and Martin Grohe. Fixed-parameter tractability, definability, and model-checking. *SIAM J. Comput.*, 31(1):113–145, 2001.

[13] Markus Frick and Martin Grohe. Deciding first-order properties of locally tree-decomposable structures. *J. ACM*, 48(6):1184–1206, 2001.

[14] Haim Gaifman. On local and non-local properties. In *Proceedings of the Herbrand Symposium*, volume 107 of *Stud. Logic Found. Math.*, pages 105 – 135. Elsevier, 1982.

[15] Jakub Gajarskỳ, Maximilian Gorsky, and Stephan Kreutzer. Differential games, locality and model checking for FO logic of graphs. *arXiv preprint arXiv:2007.11345*, 2020.

[16] Jakub Gajarský, Petr Hliněný, Jan Obdržálek, Daniel Lokshtanov, and M. S. Ramanujan. A new perspective on FO model checking of dense graph classes. *ACM Trans. Comput. Logic*, 21(4), July 2020.

[17] Jakub Gajarský, Stephan Kreutzer, Jaroslav Nesetril, Patrice Ossona de Mendez, Michal Pilipczuk, Sebastian Siebertz, and Szymon Torunczyk. First-order interpretations of bounded expansion classes. *ACM Trans. Comput. Log.*, 21(4):29:1–29:41, 2020.

[18] Robert Ganian, Petr Hliněný, Alexander Langer, Jan Obdržálek, Peter Rossmanith, and Somnath Sikdar. Lower bounds on the complexity of MSO1 model-checking. *Journal of Computer and System Sciences*, 80(1):180–194, 2014.

[19] Martin Grohe. Logic, graphs, and algorithms. *Electronic Colloquium on Computational Complexity (ECCC)*, 14(091):44, 2007.

[20] Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. *J. ACM*, 64(3):17:1–17:32, 2017.

[21] Martin Grohe and Gyorgy Turán. Learnability and definability in trees and similar structures. *Theory of Computing Systems*, 37:193–220, 2003.

[22] Stephan Kreutzer. Algorithmic meta-theorems. In *Finite and Algorithmic Model Theory*, volume 379 of *London Math. Soc. Lecture Note Ser.*, chapter 5, pages 177–270. Oxford University Press, 2011.

[23] Stephan Kreutzer and Siamak Tazari. Lower bounds for the complexity of monadic second-order logic. In *2010 25th Annual IEEE Symposium on Logic in Computer Science*, pages 189–198. IEEE, 2010.

[24] Jiří Matoušek. Bounded VC-dimension implies a fractional Helly theorem. *Discrete Comput. Geom.*, 31(2):251–255, 2004.

[25] Rajeev Motwani and Madhu Sudan. Computing roots of graphs is hard. *Discrete Appl. Math.*, 54(1):81–88, 1994.

[26] Jaroslav Nesetril, Patrice Ossona de Mendez, Michal Pilipczuk, Roman Rabinovich, and Sebastian Siebertz. Rankwidth meets stability. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 2014–2033. SIAM, 2021.

[27] Jaroslav Nešetřil, Patrice Ossona de Mendez, and Sebastian Siebertz. Structural properties of the first-order transduction quasiorder. *arXiv preprint arXiv:2010.02607*, 2021.

[28] Sang-il Oum and Paul D. Seymour. Approximating clique-width and branch-width. *J. Comb. Theory, Ser. B*, 96(4):514–528, 2006.

[29] Klaus-Peter Podewski and Martin Ziegler. Stable graphs. *Fundamenta Mathematicae*, 100(2):101–107, 1978.

[30] Detlef Seese. Linear time computable problems and first-order descriptions. *Math. Structures Comput. Sci.*, 6(6):505–526, 1996.

[31] Saharon Shelah. Stability, the f.c.p., and superstability; model theoretic properties of formulas in first order theory. *Annals of Mathematical Logic*, 3(3):271–362, 1971.