Graph Modification of Bounded Size to Minor-Closed Classes as Fast as Vertex Deletion

Laure Morelle
LIRMM, Université de Montpellier, CNRS, Montpellier, France

Ignasi Sau
LIRMM, Université de Montpellier, CNRS, Montpellier, France

Dimitrios M. Thilikos
LIRMM, Université de Montpellier, CNRS, Montpellier, France

Abstract

A *replacement action* is a function $\mathcal{L}$ that maps each graph H to a collection of graphs of size at most $|V(H)|$. Given a graph class $\mathcal{H}$, we consider a general family of graph modification problems, called $\mathcal{L}$-Replacement to $\mathcal{H}$, where the input is a graph G and the question is whether it is possible to replace some induced subgraph H_1 of G on at most k vertices by a graph H_2 in $\mathcal{L}(H_1)$ so that the resulting graph belongs to $\mathcal{H}$. $\mathcal{L}$-Replacement to $\mathcal{H}$ can simulate many graph modification problems including vertex deletion, edge deletion/addition/edition/contraction, vertex identification, subgraph complementation, independent set deletion, (induced) matching deletion/contraction, etc. We present two algorithms. The first one solves $\mathcal{L}$-Replacement to $\mathcal{H}$ in time $2^{\mathsf{poly}(k)} \cdot |V(G)|^2$ for every minor-closed graph class $\mathcal{H}$, where poly is a polynomial whose degree depends on $\mathcal{H}$, under a mild technical condition on $\mathcal{L}$. This generalizes the results of Morelle, Sau, Stamoulis, and Thilikos [ICALP 2020, ICALP 2023] for the particular case of Vertex Deletion to $\mathcal{H}$ within the same running time. Our second algorithm is an improvement of the first one when $\mathcal{H}$ is the class of graphs embeddable in a surface of Euler genus at most g and runs in time $2^{\mathcal{O}(k^9)} \cdot |V(G)|^2$, where the $\mathcal{O}(\cdot)$ notation depends on g. To the best of our knowledge, these are the first parameterized algorithms with a reasonable parametric dependence for such a general family of graph modification problems to minor-closed classes.

2012 ACM Subject Classification Theory of computation → Parameterized complexity and exact algorithms

Keywords and phrases Graph modification problems, Parameterized complexity, Graph minors, Flat Wall theorem, Irrelevant vertex technique, Algorithmic meta-theorem, Parametric dependence, Dynamic programming

Digital Object Identifier 10.4230/LIPIcs.ESA.2025.7

Related Version *Full Version*: `http://arxiv.org/abs/2504.16803`

Funding *Laure Morelle*: French project ELiT (ANR-20-CE48-0008-01)
Ignasi Sau: French project ELiT (ANR-20-CE48-0008-01)
Dimitrios M. Thilikos: ANR project GODASse ANR-24-CE48-4377, French-German Collaboration ANR/DFG Project UTMA (ANR-20-CE92-0027), and Franco-Norwegian project PHC AURORA 2024-25 (Projet n° 51260WL)

1 Introduction

A graph modification problem is typically determined by a target graph class $\mathcal{H}$ and a prescribed set of allowed local modifications $\mathcal{M}$, such as vertex/edge removal or edge addition/contraction or combinations of them, and the question is, given a graph G and an integer k, whether it is possible to transform G to a graph in $\mathcal{H}$ by applying at most k modification operations from $\mathcal{M}$. Graph modification problems are fundamental in algorithmic graph theory, as can be seen from the span of applications in domains as diverse as computational biology, computer vision, machine learning, networking, or sociology; see [19] and the references therein. Unfortunately, most of these problems are NP-complete [32, 49], and this justifies, among other approaches, to study them from the parameterized complexity viewpoint (see the monographs [10, 12, 15, 40] for an introduction to the field), where the number k of allowed modifications is taken as the parameter.

In the recent years, there has been a very active line of research about algorithmic meta-theorems for graph modification problems where the target class $\mathcal{H}$ is *minor-closed*, that is, closed under vertex deletion, edge deletion, and edge contraction. By Robertson and Seymour's seminal result [43], a minor-closed graph class $\mathcal{H}$ has a finite number of minor-obstructions, that is, graphs that are not in $\mathcal{H}$ but whose all proper minors are. Combined with a minor containment algorithm [30] (see also [26, 42]) running in almost-linear time, this implies that checking membership in a minor-closed graph class can be done in almost-linear time. For some modification problems where the target class $\mathcal{H}$ is minor-closed, such as vertex deletion, edge deletion, or vertex identification, the graphs G such that (G, k) is a yes-instance of the problem for a fixed k form a minor-closed graph class, which immediately implies an FPT-algorithm in almost-linear time for these problems. However, not all graph modification problems define a minor-closed graph class. For instance, edge contraction to planar graphs does not. Indeed, consider the graph $K_{3,4}^+$ obtained from $K_{3,4}$ by adding one edge e on the side with three vertices. Contracting e gives the planar graph $K_{2,4}$, but $K_{3,4}$ cannot be made planar by contracting one edge. Hence, some other algorithmic meta-theorems for graph modification problems to minor-closed graph classes were later introduced. Some of them are ad-hoc meta-theorems such as the one of Fomin, Golovach, and Thilikos [17] that gives quadratic FPT-algorithms for graph modification problems where $\mathcal{H}$ is the class of planar graphs and the modification is any combination of edge addition and edge deletion. Much more generally, there has been a recent line of research on model-checking on minor-closed graph classes [16, 47], which in particular implies quadratic FPT-algorithms for an extremely wide family of graph modification problems where the target class $\mathcal{H}$ is minor-closed. Unfortunately, all these algorithmic meta-theorems have a major drawback: the parametric dependence on the "amount of modification" is humongous; in fact, even a rough upper bound is not known.

On the other hand, another line of research has focused on optimizing the parametric dependence for some particular graph modification problems when the parameter is the solution size. When the target class $\mathcal{H}$ is minor-closed, such study usually does not go much beyond $\mathcal{H}$ being the class of forests and the class of union of paths [7, 9, 22, 23, 33, 39, 48]. To the best of our knowledge, only the case of vertex deletion has been studied in a a series of papers [18, 24, 25, 28, 29, 37, 38, 44, 45], focusing on optimizing the running time (both the dependence on k and n). In particular, when the minor-obstructions of $\mathcal{H}$ are connected and that one of them is planar, the currently fastest algorithm runs in time $2^{\mathcal{O}(k)} \cdot n \log^2 n$ [18], when $\mathcal{H}$ excludes a planar graph, in time $2^{\mathcal{O}(k)} \cdot n^2$ [28], when $\mathcal{H}$ is the class of planar graphs, in time $2^{\mathcal{O}(k \log k)} \cdot n$ [24], when $\mathcal{H}$ is the class of graphs embeddable in a surface of bounded genus, in time $2^{\mathcal{O}(k^2 \log k)} \cdot n^{\mathcal{O}(1)}$ [29], and when $\mathcal{H}$ is any minor-closed graph class, in time $2^{\mathsf{poly}(k)} \cdot n^2$ [38].

This article places itself in-between these two lines of research: we consider generic "meta-modification" operations (of course, much less generic than those of the currently most general algorithmic meta-theorem in [47], but still quite versatile), and we manage to achieve the same (very reasonable) parametric dependence as the currently best one for vertex deletion [38] when the target $\mathcal{H}$ is any minor-closed graph class. We hope that our work will trigger further research about *efficient* algorithmic meta-theorems for graph modification problems to minor-closed graph classes.

Our results. We define a graph modification problem, called $\mathcal{L}$-Replacement to $\mathcal{H}$ ($\mathcal{L}$-R-$\mathcal{H}$ for short), which, depending on the choice of the function $\mathcal{L}$, called a *replacement action*, can simulate vertex deletion, edge deletion, edge completion, edge edition, edge contraction, vertex identification, independent set deletion, matching deletion, matching contraction, star deletion, and subgraph complementation, to name a few (see section 3 for an exposition of some problems encompassed by our result). When $\mathcal{H}$ is minor-closed, we solve $\mathcal{L}$-R-$\mathcal{H}$ in time $2^{\mathsf{poly}(k)} \cdot n^2$ (Theorem 2), where the degree of the polynomial poly depends on the maximum size $s_{\mathcal{H}}$ of the minor-obstructions of $\mathcal{H}$. This is the same running time as the one achieved by the currently best algorithm for vertex deletion [38] (the degree of k in poly is the same as in [38] up to an extra additive constant of one that is absolutely negligible compared to the total degree that depends (wildly) on $s_{\mathcal{H}}$). For the other graph modification problems encompassed within $\mathcal{L}$-R-$\mathcal{H}$, to the authors' knowledge, the only minor-closed classes for which a good parametric dependence was previously known, if any, were the class of forests and the class of union of paths [9, 22, 33, 39, 48].

As it is usually the case concerning meta-theorems, the degree d of the polynomial poly in Theorem 2 is unfortunately huge. While we did not compute its exact value, we know that $d \geq 2^{2^{s_{\mathcal{H}}^{24}}}$. Nevertheless, d can improved for some specific target classes $\mathcal{H}$. The *Euler genus* of a surface Σ that is obtained from the sphere by adding h handles and c crosscaps is defined to be $c + 2h$. In particular, when $\mathcal{H}$ is the class of graphs embeddable in a surface of Euler genus at most g, we provide another algorithm solving $\mathcal{L}$-R-$\mathcal{H}$ in time $2^{\mathcal{O}(k^9)} \cdot n^2$ (Theorem 3), where the $\mathcal{O}(\cdot)$ notation depends on g. Note that, as opposed to Theorem 2, in Theorem 3 the contribution of the genus (that is, of the target graph class $\mathcal{H}$) does *not* affect the degree of the parameter k in the exponent.

Organization. In section 2 we give basic definitions and conventions, and we formally define the problem and state our results. In section 3 we provide a non-exhaustive list of problems generated by different instantiations of the replacement action $\mathcal{L}$, and hence encompassed by our results. In section 4 we present an overview of our techniques. Due to space limitations, all proofs have been deferred to the full version of this article. In section 5 we present some directions for further research.

2 Definition of the problem and formal statement of the results

In this section we formally define the $\mathcal{L}$-R-$\mathcal{H}$ problem and state our results, already informally discussed in the introduction. We use standard graph-theoretic notation, and complete preliminaries about graphs (including tree decompositions and minors) can be found in the full version. We define here the non-standard notions that are needed in order to state our results.

Minor-closed graph classes. A graph class $\mathcal{H}$ is *minor-closed* if, for each graph G and each minor H of G, the fact that $G \in \mathcal{H}$ implies that $H \in \mathcal{H}$. Given a collection of graphs $\mathcal{F}$, we denote by $\mathsf{exc}(\mathcal{F})$ the class of graphs that do not contain a graph in $\mathcal{F}$ as a minor. Obviously, $\mathsf{exc}(\mathcal{F})$ is minor-closed. A *(minor-)obstruction* of a graph class $\mathcal{H}$ is a graph F that is not in $\mathcal{H}$, but whose minors are all in $\mathcal{H}$. The set of all the obstructions of $\mathcal{H}$ is denoted by $\mathsf{obs}(\mathcal{H})$. By the seminal work of Robertson and Seymour [43], if $\mathcal{H}$ is a minor-closed graph class, then $\mathsf{obs}(\mathcal{H})$ is finite. Note that, if $\mathcal{F} = \mathsf{obs}(\mathcal{H})$, then $\mathsf{exc}(\mathcal{F}) = \mathcal{H}$. The *detail* of a graph G is $\max\{|V(G)|, |E(G)|\}$.

Ordered graphs. For the definitions of the next two paragraphs to be correct, we actually need to consider ordered graphs instead of graphs (see the "Graph modifications" paragraph). An *ordered graph* is a graph G equipped with a strict total order on $V(G)$, denoted by $<_G$. In other words, there exists an indexation $v_1, \dots, v_n$ of the vertices of $V(G)$ such that $v_1 <_G v_2 <_G \cdots <_G v_n$. A subgraph H of an ordered graph G naturally comes equipped with the strict order $<_H$ such that, for each distinct $u, v \in V(H)$, $u <_H v$ if and only if $u <_G v$.

Replacement actions. The *any-replacement action* is the function $\mathcal{M}$ that maps each ordered graph H_1 to the collection $\mathcal{M}(H_1)$ of all the pairs (H_2, ϕ), where H_2 is an ordered graph and $\phi : V(H_1) \to V(H_2) \cup \{\emptyset\}$ is a function such that:

- $|V(H_2)| \leq |V(H_1)|$,
- for each $v \in V(H_2)$, $\phi^{-1}(v) \neq \emptyset$, and
- $<_{H_2}$ is the strict total order such that, for each distinct $v_1, v_2 \in V(H_2)$, we have $v_1 <_{H_2} v_2$ if and only if $u_1 <_{H_1} u_2$ where, for $i \in [2]$, u_i is the smallest vertex (according to $<_G$) in $\phi^{-1}(v_i)$.

A *replacement action* (abbreviated as *R-action*) is any function $\mathcal{L}$ that maps an ordered graph (called a *pattern*) H_1 to a non-empty collection $\mathcal{L}(H_1) \subseteq \mathcal{M}(H_1)$ of its possible *pattern transformations*. See Figure 1 for an illustration. The vertices of H_1 mapped by ϕ to the empty set are said to be *deleted*, and two vertices of H_1 mapped by ϕ to the same vertex of H_2 are said to be *identified*. Given $S \subseteq V(H_1)$, we set $\phi^+(S) = \phi(S) \setminus \{\emptyset\}$. Note that, if $\phi(S) = \{\emptyset\}$, then $\phi^+(S) = \{\emptyset\} \setminus \{\emptyset\} = \emptyset$.

Graph modifications. Let $\mathcal{L}$ be an R-action, let G be an ordered graph, and $S \subseteq V(G)$. Let $(H_2, \phi) \in \mathcal{L}(G[S])$. We denote by $G^S_{(H_2,\phi)}$ the graph obtained from the disjoint union of $G - S$ and H_2 by adding an edge $u\phi(v)$ for each $u \in V(G) \setminus S$ and each $v \in \phi^{-1}(V(H_2))$ such that $uv \in E(G)$. We equip $G' := G^S_{(H_2,\phi)}$ with the strict total order $<_{G'}$ such that $v_1 <_{G'} v_2$ if and only if $u_1 <_G u_2$ where, for $i \in [2]$, $u_i := v_i$ if $v_i \in V(G) \setminus S$, and u_i is the smallest vertex in $\phi^{-1}(v_i)$ if $v_i \in V(H_2)$. We also set $\mathcal{L}_S(G) = \{G^S_{(H_2,\phi)} \mid (H_2, \phi) \in \mathcal{L}(G[S])\}$. See Figure 1 for an illustration.

Note that we consider ordered graphs merely so that the correspondence between the vertices in S and the vertices in $V(H_2)$ is well-defined. We actually omit the order from the statements, but it will be implicitly assumed that vertices have a label that allows us to keep track of them during the modification procedure.

Let $\mathcal{L}$ be an R-action and $\mathcal{H}$ be a graph class. We define the following problem.

$\mathcal{L}$-Replacement to $\mathcal{H}$ ($\mathcal{L}$-R-$\mathcal{H}$)

Input: A graph G and $k \in \mathbb{N}$.
Question: Is there a set $S \subseteq V(G)$ of size at most k such that $\mathcal{L}_S(G) \cap \mathcal{H} \neq \emptyset$?

Such a set S is called *solution* of $\mathcal{L}$-R-$\mathcal{H}$ for the instance (G, k).

We will use the following observation, which implies that a no-instance for Vertex Deletion to $\mathcal{H}$ is also a no-instance for $\mathcal{L}$-R-$\mathcal{H}$.

▶ **Observation 1.** *Let $\mathcal{H}$ be a hereditary graph class, let $\mathcal{L}$ be an R-action, let G be a graph, and let $S \subseteq V(G)$. If $\mathcal{L}_S(G) \cap \mathcal{H} \neq \emptyset$, then $G - S \in \mathcal{H}$.*

Proof. Indeed, suppose that there is $(H_2, \phi) \in \mathcal{L}(G[S])$ such that $G^S_{(H_2,\phi)} \in \mathcal{H}$. Then, because $\mathcal{H}$ is hereditary, $G^S_{(H_2,\phi)} - \phi^+(S) = G - S \in \mathcal{H}$. ◀

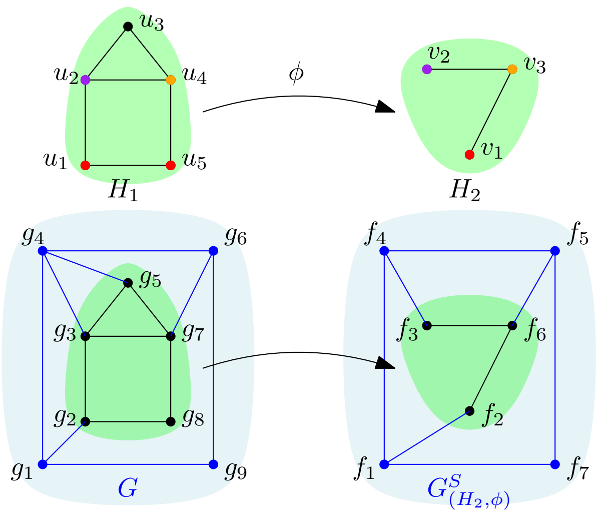

Figure 1 Example of $(H_2, \phi) \in \mathcal{L}(H_1)$ and of the graph modification $G^S_{(H,\phi)}$ where S is the set of black vertices of G. ϕ is represented by the colors, that is, $\phi(u_1) = \phi(u_5) = v_1$, $\phi(u_2) = \phi(v_2)$, $\phi(u_3) = \emptyset$, and $\phi(u_4) = v_3$. The order on the vertex sets of the depicted graphs is given by the corresponding labels.

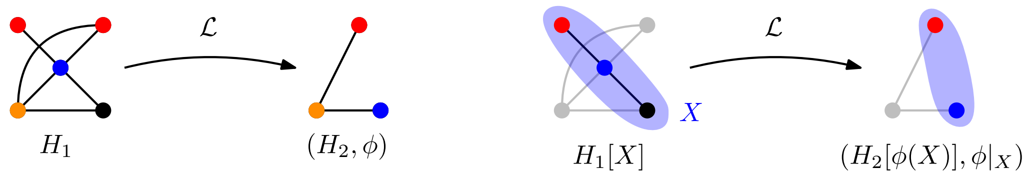

Figure 2 If $\mathcal{L}$ is hereditary, then a restriction of an allowed modification is also allowed.

Hereditary R-actions. An R-action is said to be *hereditary* if, for each ordered graph H_1, for each non-empty $X \subseteq V(H_1)$, and for each $(H_2, \phi) \in \mathcal{L}(H_1)$, we have $(H_2[\phi^+(X)], \phi|_X) \in \mathcal{L}(H_1[X])$. We say that $(H_2[\phi^+(X)], \phi|_X)$ is the *restriction* of (H_2, ϕ) to X. See Figure 2 for an illustration. Informally, an R-action is hereditary if, when a modification is allowed, then modifying "less" is allowed as well. For instance, if $\mathcal{L}$ allows us to delete *exactly* k vertices, then $\mathcal{L}$ also allows us to delete *at most* k vertices.

Some conventions. By convention, when there is no confusion, we set $n := |V(G)|$ and $m := |E(G)|$. In the rest of the paper, instead of considering a minor-closed graph class $\mathcal{H}$, we consider its obstruction set $\mathcal{F}$, and thus the minor-closed graph class $\mathsf{exc}(\mathcal{F})$. We define three constants depending on $\mathcal{F}$ that will be used throughout the paper whenever we consider such a collection $\mathcal{F}$. We define $a_{\mathcal{F}}$ as the minimum apex number of a graph in $\mathcal{F}$, we set $s_{\mathcal{F}} := \max\{|V(F)| \mid F \in \mathcal{F}\}$, and we define $\ell_{\mathcal{F}}$ to be the maximum detail of a graph in $\mathcal{F}$. Given a tuple $\mathbf{t} = (x_1, \ldots, x_\ell) \in \mathbb{N}^\ell$ and two functions $\chi, \psi : \mathbb{N} \to \mathbb{N}$, we write $\chi(n) = \mathcal{O}_{\mathbf{t}}(\psi(n))$ in order to denote that there exists a computable function $\phi : \mathbb{N}^\ell \to \mathbb{N}$ such that $\chi(n) = \mathcal{O}(\phi(\mathbf{t}) \cdot \psi(n))$. Notice that $s_{\mathcal{F}} \leq \ell_{\mathcal{F}} \leq s_{\mathcal{F}}(s_{\mathcal{F}} - 1)/2$, and thus $\mathcal{O}_{\ell_{\mathcal{F}}}(\cdot) = \mathcal{O}_{s_{\mathcal{F}}}(\cdot)$.

Our main result is the following.

▶ **Theorem 2.** *Let $\mathcal{F}$ be a finite collection of graphs and let $\mathcal{L}$ be a hereditary R-action. There is an algorithm that, given a graph G and $k \in \mathbb{N}$, runs in time $2^{\mathsf{poly}_{\mathcal{F}}(k)} \cdot n^2$ and either outputs a solution of $\mathcal{L}$-R-$\mathsf{exc}(\mathcal{F})$ for the instance (G, k) or reports a* no*-instance. Moreover,* $\mathsf{poly}_{\mathcal{F}}$ *is a polynomial whose degree depends on the maximum detail of a graph in $\mathcal{F}$.*

As mentioned in the introduction, the main result in [47] already implies that $\mathcal{L}$-R-$\mathcal{H}$ is solvable in time $f(k) \cdot n^2$ when $\mathcal{H}$ is minor-closed for some huge function f that is not even estimated. Our main contribution is an explicit and single-exponential dependence on k.

The degree of $\mathsf{poly}_{\mathcal{F}}(k)$ is quite big, but we can reduce it in some specific cases.

▶ **Theorem 3.** *Let $\mathcal{L}$ be a hereditary R-action and $\mathcal{H}$ be the class of graphs embeddable in a surface Σ of Euler genus at most g. There is an algorithm that, given a graph G and $k \in \mathbb{N}$, runs in time $2^{\mathcal{O}_g(k^9)} \cdot n^2$ and either outputs a solution of $\mathcal{L}$-R-$\mathcal{H}$ for the instance (G, k) or reports a* no*-instance.*

More generally, we study the annotated version of $\mathcal{L}$-R-$\mathcal{H}$. Let $\mathcal{L}$ be a hereditary R-action and $\mathcal{H}$ be a graph class. We define the following problem.

$\mathcal{L}$-Annotated Replacement to $\mathcal{H}$ ($\mathcal{L}$-AR-$\mathcal{H}$)

Input: A graph G, a set of annotated vertices $S' \subseteq V(G)$, $(H_2', \phi') \in \mathcal{L}(G[S'])$, and $k \in \mathbb{N}$.

Question: Is there a set $S \subseteq V(G)$ of size at most k and $(H_2, \phi) \in \mathcal{L}(G[S])$ such that (H_2', ϕ') is the restriction of (H_2, ϕ) to S' and $G^S_{(H_2,\phi)} \in \mathcal{H}$?

Obviously, we must have $S' \subseteq S$. Such a triple (S, H_2, ϕ) is called a *solution* of $\mathcal{L}$-AR-$\mathcal{H}$ for the instance (G, S', H_2', ϕ', k). An instance of $\mathcal{L}$-AR-$\mathcal{H}$ where $S' = \emptyset$ is an instance of $\mathcal{L}$-R-$\mathcal{H}$, so $\mathcal{L}$-AR-$\mathcal{H}$ generalizes $\mathcal{L}$-R-$\mathcal{H}$. Two instances $\mathcal{I}_1$ and $\mathcal{I}_2$ are *equivalent* instances of $\mathcal{L}$-AR-$\mathcal{H}$ if $\mathcal{I}_1$ is a yes-instance of $\mathcal{L}$-AR-$\mathcal{H}$ if and only if $\mathcal{I}_2$ is a yes-instance of $\mathcal{L}$-AR-$\mathcal{H}$.

In fact, we prove stronger statements of Theorem 2 and Theorem 3 that apply to their respective annotated versions.

3 Problems generated by different instantiations of $\mathcal{L}$

Many graph modification problems correspond to $\mathcal{L}$-R-$\mathcal{H}$ for a specific R-action $\mathcal{L}$ and a specific target graph class $\mathcal{H}$. We give a few examples below. Let $\mathcal{H}$ be a minor-closed graph class. For instance, $\mathcal{H}$ could be the class of edgeless graphs, of forests, of graphs whose connected components have size at most k, of planar graphs, or of graphs embeddable in a surface Σ. Note that we do not mention Edge Addition to $\mathcal{H}$ (nor Edge Edition to $\mathcal{H}$) here, because when $\mathcal{H}$ is a minor-closed graph class, adding edges is "unnecessary", in the sense that the edge deletion variant has the same expressive power, and we can solve it. Note also that $\mathcal{L}$-R-$\mathcal{H}$, and thus in particular all problems of this section, was already known to be solvable in FPT-time (when $\mathcal{H}$ is minor-closed) by the result of [47]. However, as mentioned before, the parametric dependence is huge and not even explicit in [47].

Given a set A, we denote the identity function mapping each $a \in A$ to itself by id_A.

Vertex Deletion to $\mathcal{H}$

Input: A graph G and $k \in \mathbb{N}$.

Question: Is there a set $S \subseteq V(G)$ of size at most k such that $G - S \in \mathcal{H}$?

Vertex Deletion to $\mathcal{H}$ reduces to $\mathcal{L}_{\mathsf{vDel}}$-R-$\mathcal{H}$, where $\mathcal{L}_{\mathsf{vDel}}$ is the function that maps any graph H_1 to the singleton containing the empty graph and the constant function $\phi : V(H_1) \to \{\emptyset\}$. Vertex Deletion to $\mathcal{H}$ is already known [38] to be solvable within the same running time as the one of Theorem 2. Hence, the result of Theorem 2 is not an improvement for this specific problem, but it shows that our result is tight compared to the currently best known result for Vertex Deletion to $\mathcal{H}$.

EDGE DELETION TO $\mathcal{H}$

Input: A graph G and $k \in \mathbb{N}$.
Question: Is there a set $F \subseteq E(G)$ of size at most k such that $G - F \in \mathcal{H}$?

(G, k) is a yes-instance of EDGE DELETION TO $\mathcal{H}$ if and only if $(G, 2k)$ is a yes-instance of $\mathcal{L}_{\mathsf{eDel},k}$-R-$\mathcal{H}$, where $\mathcal{L}_{\mathsf{eDel},k}$ is the function that maps each graph H_1 to the set of pairs $(H_1 - F, \mathsf{id}_{V(H_1)})$ over all $F \subseteq E(G)$ of size at most k. Algorithms with a nice parametric dependence are only known for specific target classes $\mathcal{H}$. Namely, when $\mathcal{H}$ is the class of forests, EDGE DELETION TO $\mathcal{H}$ corresponds to FEEDBACK EDGE SET, which can be solved in constant time given that the size of a minimum feedback edge set is $m - n + 1$ (assuming the graph is connected). When $\mathcal{H}$ is the class of graphs that are a union of paths, then there is a linear kernel for the problem [34], as well as a FPT algorithm with parametric dependence on k at most 2^k [48]. We refer the reader to the survey of [9], as well as [13], for other results with explicit dependence on k when $\mathcal{H}$ is not a minor-closed graph class.

Given a graph G and a set of edges $F \subseteq E(G)$, we denote by G/F the graph obtained from G after contracting the edges in F.

EDGE CONTRACTION TO $\mathcal{H}$

Input: A graph G and $k \in \mathbb{N}$.
Question: Is there a set $F \subseteq E(G)$ of size at most k such that $G/F \in \mathcal{H}$?

(G, k) is a yes-instance of EDGE CONTRACTION TO $\mathcal{H}$ if and only if $(G, 2k)$ is a yes-instance of $\mathcal{L}_{\mathsf{Con},k}$-R-$\mathcal{H}$, where $\mathcal{L}_{\mathsf{Con},k}$ is the function that maps each graph H_1 to the set of pairs $(H_1/F, \phi)$ over all $F \subseteq E(G)$ of size at most k, where ϕ maps $v \in V(H_1)$ to the corresponding vertex of H_1/F. An explicit parametric dependence was given in [22] when $\mathcal{H}$ is a class of paths (running time $2^{k+o(k)} + n^{\mathcal{O}(1)}$) or the class of trees (running time $4.98^k \cdot n^{\mathcal{O}(1)}$). Though these classes are not minor-closed, we can easily extend these results to the case when $\mathcal{H}$ is the class of unions of paths or the class of forests (up to a 2^k factor). FPT-algorithms with an explicit parametric dependence were also studied when $\mathcal{H}$ is a collection of generalization and restriction of trees [2,3], or when $\mathcal{H}$ is the class of cactus graphs [31]. We refer the reader to [21] for more results when the target class is not minor-closed.

VERTEX IDENTIFICATION TO $\mathcal{H}$

Input: A graph G and $k \in \mathbb{N}$.
Question: Is there a set $S \subseteq V(G)$ of size at most k and a partition $(X_1, \dots, X_p)$ of S such that the graph obtained after identifying the vertices in X_i to a single vertex x_i, for $i \in [p]$, belongs to $\mathcal{H}$?

VERTEX IDENTIFICATION TO $\mathcal{H}$ reduces to $\mathcal{L}_{\mathsf{Id}}$-R-$\mathcal{H}$, where $\mathcal{L}_{\mathsf{Id}}$ is the function that maps each graph H_1 to the set of pairs (H_2, ϕ), where H_2 can be obtained from H_1 after identifying each X_i of a partition $(X_1, \dots, X_p)$ of some set $S \subseteq V(H_1)$ to a single vertex x_i, and ϕ maps vertices of X_i to x_i and is the identity on $V(H_1) \setminus S$. VERTEX IDENTIFICATION TO $\mathcal{H}$ is known to admit a kernel of size $2k + 1$ when $\mathcal{H}$ is the class of forests [39]. To the authors' knowledge, this is the only known result for this problem.

INDEPENDENT SET DELETION TO $\mathcal{H}$

Input: A graph G and $k \in \mathbb{N}$.
Question: Is there an independent set $I \subseteq V(G)$ of size at most k such that $G - I \in \mathcal{H}$?

INDEPENDENT SET DELETION TO $\mathcal{H}$ reduces to $\mathcal{L}_{\mathsf{ISDel}}$-R-$\mathcal{H}$, where $\mathcal{L}_{\mathsf{ISDel}}$ is the function that maps any graph H_1 to the set of pairs $(H_1 - I, \phi)$ over all independent sets $I \subseteq V(H_1)$, where ϕ maps vertices of I to the empty set and is the identity on $V(H_1) \setminus I$.

When $\mathcal{H}$ is the class of forests, the problem is known to be solvable in time $3.62^k \cdot n^{\mathcal{O}(1)}$ [33]. Concerning other target classes that are not minor-closed, mainly bipartite graphs, let us mention [1, 6, 20].

To illustrate the versatility of $\mathcal{L}$-R-$\mathcal{H}$, let us present some other problems that can be defined by particular hereditary R-actions, though they do not seem to have been studied when parameterized by the solution size.

(INDUCED) MATCHING DELETION TO $\mathcal{H}$

Input: A graph G and $k \in \mathbb{N}$.
Question: Is there an (induced) matching $M \subseteq E(G)$ of size at most k such that $G - M \in \mathcal{H}$?

(G, k) is a yes-instance of (INDUCED) MATCHING DELETION TO $\mathcal{H}$ if and only if $(G, 2k)$ is a yes-instance of $\mathcal{L}_{\mathsf{mDel},k}$-R-$\mathcal{H}$, where $\mathcal{L}_{\mathsf{mDel},k}$ is defined similarly to $\mathcal{L}_{\mathsf{eDel},k}$ above, but for (induced) matchings. There are some results on MATCHING DELETION TO $\mathcal{H}$ when $k = n$ and $\mathcal{H}$ is the class of forests [36, 41] or bipartite graphs (see [35] for a small survey).

(INDUCED) MATCHING CONTRACTION TO $\mathcal{H}$

Input: A graph G and $k \in \mathbb{N}$.
Question: Is there an (induced) matching $M \subseteq E(G)$ of size at most k such that $G/M \in \mathcal{H}$?

(G, k) is a yes-instance of (INDUCED) MATCHING CONTRACTION TO $\mathcal{H}$ if and only if $(G, 2k)$ is a yes-instance of $\mathcal{L}_{\mathsf{mCon},k}$-R-$\mathcal{H}$, where $\mathcal{L}_{\mathsf{mCon},k}$ is defined similarly to $\mathcal{L}_{\mathsf{Con},k}$ above, but for (induced) matchings.

INDUCED STAR DELETION TO $\mathcal{H}$

Input: A graph G and $k \in \mathbb{N}$.
Question: Is there a set $F \subseteq E(G)$ inducing a star $K_{1,k'}$ with $k' \leq k$ such that $G - F \in \mathcal{H}$?

(G, k) is a yes-instance of STAR DELETION TO $\mathcal{H}$ if and only if $(G, k+1)$ is a yes-instance of $\mathcal{L}_{\mathsf{StarDel},k}$-R-$\mathcal{H}$, where $\mathcal{L}_{\mathsf{StarDel},k}$ is the function that maps any graph H_1 to the set of pairs $(H_1 - F, \mathsf{id}_{V(H_1)})$ over all sets $F \subseteq E(G)$ inducing a subgraph of $K_{1,k}$.

Given a graph G, the *complement* of G, denoted by $\overline{G}$, is graph with vertex set $V(G)$ and edge set the edges that do not belong to $E(G)$.

SUBGRAPH COMPLEMENTATION TO $\mathcal{H}$

Input: A graph G and $k \in \mathbb{N}$.
Question: Is there a set $S \subseteq V(G)$ of size at most k such that the graph obtained after replacing $G[S]$ with its complement $\overline{G[S]}$ belongs to $\mathcal{H}$?

SUBGRAPH COMPLEMENTATION TO $\mathcal{H}$ reduces to $\mathcal{L}_{\mathsf{Comp}}$-R-$\mathcal{H}$, where $\mathcal{L}_{\mathsf{Comp}}$ is the function that maps any graph H_1 to the singleton containing the pair $(\overline{H_1}, \mathsf{id}_{V(H_1)})$. The problem was recently studied when $k = n$ for various target classes; we refer the reader to [4].

4 Overview of our techniques

To handle several modification problems at once, we adapt the vocabulary of Fomin, Golovach, and Thilikos [17], who introduced the notion of replacement action. Intuitively (see section 2 for the formal definition), a *replacement action* is a function $\mathcal{L}$ that maps a graph H_1 to a collection $\mathcal{L}(H_1)$ of pairs (H_2, ϕ) where H_2 is a graph with at most $|V(H_1)|$ vertices and ϕ maps each vertex of H_1 to either a vertex of H_2 or the empty set. Mapping a vertex of H_1 to the empty set corresponds to a deletion, while mapping several vertices to the same vertex of H_2 corresponds to an identification. Replacement actions were originally defined in [17] to solve a collection of graph modification problems where only edges are modified and where the target class is the class of planar graphs. Compared to [17], however, the size of H_2 may here be smaller than the size of H_1, which happens when deleting or identifying vertices, while in [17] it is required that $|V(H_1)| = |V(H_2)|$. Let us fix a replacement action $\mathcal{L}$ and a target graph class $\mathcal{H}$. Recall that the $\mathcal{L}$-Replacement to $\mathcal{H}$ ($\mathcal{L}$-R-$\mathcal{H}$) problem asks, given a graph G and $k \in \mathbb{N}$, whether there is an induced subgraph H_1 of size at most k in G and a pair $(H_2, \phi) \in \mathcal{L}(H_1)$ such that H_1 can be replaced by H_2 such that the resulting graph G' belongs to $\mathcal{H}$ (for $u \in V(G) \setminus V(H_1)$ and $v \in V(H_2)$, $uv \in E(G')$ if and only if there is $v' \in \phi^{-1}(v)$ such that $uv' \in E(G)$). For our techniques to work (see the "irrelevant vertex technique" paragraph below for more precision), we require our function $\mathcal{L}$ to be *hereditary*, which essentially means if H_2 is in $\mathcal{L}(H_1)$, then for any induced subgraph H_1' of H_1, the corresponding induced subgraph of H_2 is in $\mathcal{L}(H_1')$ (cf. section 2 for the formal definition and Figure 2 for an illustration). For instance, this implies that we can ask whether it is possible to do *at most* k edge editions to get a graph in $\mathcal{H}$, but we cannot ask whether it is possible to do *exactly* k edge editions to get a graph in $\mathcal{H}$.

High-level description of our algorithms. The techniques that we employ for our first algorithm (that is, when $\mathcal{H}$ is any minor-closed graph class) are strongly inspired by those used by Morelle, Sau, Stamoulis, Thilikos [38] for the particular case of vertex deletion (see also [44]), namely Vertex Deletion to $\mathcal{H}$, achieving the same running time. Nevertheless, in order to deal with our "meta-modification" operations, we need several new technical insights compared to the approach of [44], which we proceed to sketch. In a nutshell, the algorithm of [38] employs a win/win strategy that proceeds as follows:

- If the treewidth of the input graph is small (as a function of the parameter k), then solve the problem via a dynamic programming approach.
- If the treewidth of the input graph is big, then either
  - (*irrelevant vertex*) find a vertex v such that (G, k) and $(G - v, k)$ are equivalent instances, or
  - (*branching case*) find a set $A \subseteq V(G)$ of small size such that there exists $v \in A$ such that (G, k) and $(G - v, k - 1)$ are equivalent instances,

  and recurse.

Hence, we require three ingredients: one to solve the problem parameterized by treewidth, one to find an irrelevant vertex, and one to find an "obligatory set" A, all with a "reasonable" parametric dependence on k. Then, we need to construct an algorithm so that one of these three cases always applies and such that the overall running time is still within the desired bound, which is one of the most convoluted parts of the proof. In what follows we provide further insights about these steps, by first saying a few words about flat walls.

The need for annotation. Let S' be the set of vertices recursively guessed to be modified in the branching step. An advantage when the modification consists in vertex deletion is that we can simply recurse on $(G - S', k - |S'|)$. For the more general case of $\mathcal{L}$-R-$\mathcal{H}$, we cannot simply delete S', as the considered modification may be different from vertex deletion. We need 1) to guess how $G[S']$ is modified, that is, to guess $(H_2', \phi') \in \mathcal{L}(G[S'])$ and 2) to remember S' and (H_2', ϕ') in order to check that we eventually find a set $S \supseteq S'$ and an allowed modification $(H_2, \phi) \in \mathcal{L}(G[S])$ whose restriction to S' is (H_2', ϕ') such that the modified graph is in $\mathcal{H}$. This is why we need to solve the *annotated version of the problem*, denoted by $\mathcal{L}$-AR-$\mathcal{H}$, where we add to the input a subset S' of vertices of G that are required to be part of H_1, as well as the modification (H_2', ϕ') made on S'.

Dynamic programming algorithm in the case of bounded treewidth. Note that we cannot just use Courcelle's theorem [8], since we require a nice parametric dependence on k. Hence, we need to design our own dynamic programming algorithm to solve $\mathcal{L}$-AR-$\mathcal{H}$ parameterized by the treewidth and k. Essentially, the idea is to guess, in each bag $\beta(t)$ of the decomposition, the set S_t of vertices that are modified as well as how they are modified, and to reduce the size of the graph G_t induced by the bag t and its children using the representative-based technique of [5]. This technique is essentially based on the property that, given a graph G in a minor-closed graph class $\mathcal{H}$ with a boundary B, there is a graph R of *bounded size* with same boundary B, called the *representative of* G, such that, for any graph H glued on B to get $G \oplus H$ and $R \oplus H$, $G \oplus H \in \mathcal{H}$ if and only if $R \oplus H \in \mathcal{H}$. G_t does not belong to $\mathcal{H}$, so we cannot find a representative of G_t, but we find instead a representative of the graph $G_t' \in \mathcal{H}$ modified from G_t according to the guessed modification on S_t and the previously guessed modification on the children of t. Given that we may need to identify together vertices that are far apart in the tree decomposition, we need to remember throughout the algorithm the vertices that are guessed to be part of the solution. The fact that we keep information about these at most k vertices explains the dependence on k of the dynamic programming algorithm. More precisely, we prove the following result.

▶ **Theorem 4.** *Let $\mathcal{F}$ be a finite collection of graphs and $\mathcal{L}$ be an R-action. There is an algorithm that, given $k \in \mathbb{N}$, a graph G of treewidth at most w, a set $S' \subseteq V(G)$ of size at most k, and $(H_2', \phi') \in \mathcal{L}(G[S'])$, in time $2^{\mathcal{O}_{\ell_{\mathcal{F}}}(k^2+(k+w)\log(k+w))} \cdot n$ either outputs a solution of $\mathcal{L}$-AR-$\mathsf{exc}(\mathcal{F})$ for the instance (G, S', H_2', ϕ', k), or reports a* no*-instance.*

The above result parameterized by treewidth and k may be of independent interest, given that it implies an algorithm with a good parametric dependence on the treewidth and k for a number of graph modification problems. Note that the question of whether $\mathcal{L}$-R-$\mathcal{H}$ is FPT parameterized by only treewidth is open. Even Courcelle's theorem only implies a running time of $f(\mathsf{tw}, k) \cdot n$, given that the size of the CMSO formula expressing yes-instances of $\mathcal{L}$-R-$\mathcal{H}$ depends on k. Note that, in [38], the bounded treewidth part consists just in a black-box application of the algorithm of Baste, Sau, and Thilikos [5].

Flat walls. An essential tool of our approach is the notion of *flat wall*, originating in the work of Robertson and Seymour [42]. Informally speaking, a flat wall is a structure made up of (non-necessarily planar) pieces, called *flaps*, that are glued together in a bidimensional grid-like way defining the so-called *bricks* of the wall. While such a structure may not be planar, it enjoys topological properties similar to those of planar graphs, in the sense that two paths that are not routed entirely inside a flap cannot "cross", except at a constant-sized vertex set A whose vertices are called *apices*. Hence, flat walls are only "locally non-planar", and after removing apices we can apply useful locality arguments, in the sense that two vertices that are in "distant" flaps should also be "distant" in the whole graph without the apices. In this article we apply some variants of one of the most celebrated results in the theory of Graph Minors by Robertson and Seymour [42, 43], known as the *Flat Wall theorem* (see also [27, 46] for recently proved variants), which informally states that graphs of large treewidth contain either a large clique minor or a large flat wall.

In order for our formal statements to be mathematically correct, we would need to introduce a number of notions originating in [46]. Unfortunately, in the attached full version several pages are required to provide all these technical notions. Due to space limitations, in this extended abstract we only provide intuitive descriptions of the main notions required to read the statement of the results, but we skip a number of technical terms (such as renditions, tilts, influence, regular flatness pairs, etc.) that are not the main focus of this sketch. All details can be found in the attached full version, and we refer the reader to [46] for a more detailed exposition of these definitions and the reasons for which they were introduced.

An r-*wall* is any graph W obtained from a so-called *elementary* r-wall $\bar{W}$ after subdividing edges: see Figure 3 for self-explanatory illustration of a 5-wall.

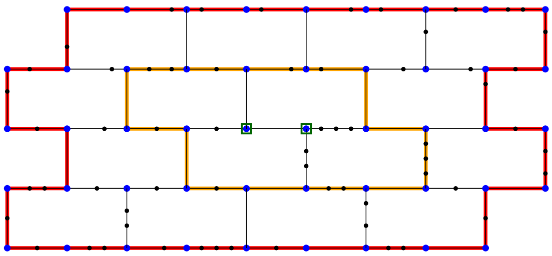

Figure 3 A 5-wall.

A flat wall is illustrated in Figure 4, where the flaps mentioned above correspond to the orange cells. The perimeter of a flat wall in a graph G separates $V(G)$ into two sets X and Y with Y containing the wall. The *compass* of a flat wall is $G[Y]$. For example, in Figure 4, X is the set of vertices in the green part, and Y the set of vertices in the orange part.

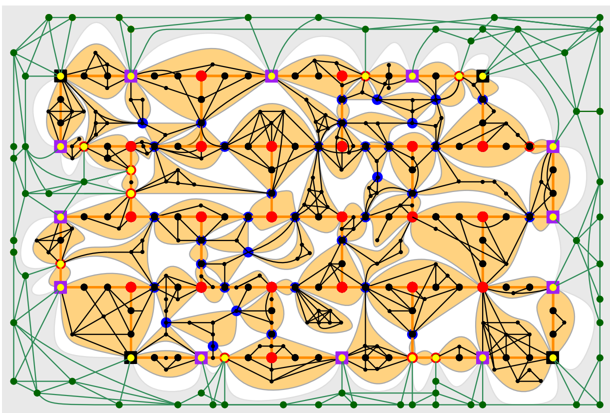

Figure 4 Illustration of flat wall, adapted from [5, Figure 4]). The edges of the subjacent wall are depicted in orange, defining the corresponding bricks.

In order to find an irrelevant vertex, we need to deal with *homogeneous* flat walls. Intuitively, homogeneous flat walls are flat walls that allow the routing of the same set of (topological) minors in the augmented flaps (i.e., the flaps together with the apex set) "cropped" by each one of their bricks. Such a homogeneous wall can be detected in a big enough flat wall and this "homogeneity" property implies that some central part of a big enough homogeneous wall can be declared irrelevant.

Another very useful notion is that of a *canonical partition* of a graph G with respect to some wall W of G. Informally, this refers to a partition of the vertex set of G into bags that follow the grid-like structure of W; see Figure 5. Essentially, the goal is to be able to contract each of these bags to obtain a grid that is a minor of W, and thus of G. In particular, we prove (see Theorem 5 below) that if G contains as a minor a grid Γ along with a set A whose vertices have sufficiently many neighbors in the grid, then some vertex in A is obligatory. We use canonical partitions here to easily find such a structure given a wall of G.

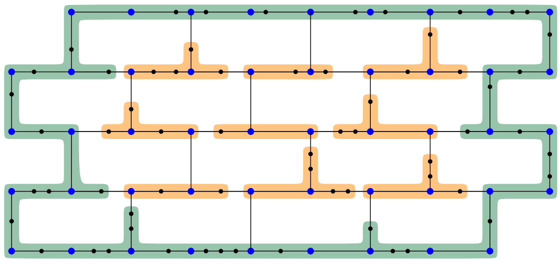

Figure 5 A 5-wall and its canonical partition $\mathcal{Q}$. The green bag is the external bag Q_{ext} and the orange bags are the internal bags of $\mathcal{Q}$. Contracting each internal bag of $\mathcal{Q}$ we obtain a (3×3)-grid.

Branching step. The branching case is not much different from what is done in [38] (originally from [45]): essentially, if there is a big enough wall W (cf. Figure 3) and a set A of vertices having many disjoint paths to W (cf. Figure 6), then some modification (H_A, ϕ_A) must happen in A and we can branch. Here, we however need to additionally prove that we must have $|\phi_A(A) \setminus \{\emptyset\}| < |A|$. We stress that it is important here to guess some modification in A that strictly decreases the size of A, so that, after applying this partial modification to G at the next step in the recursion, we will not find the exact same obligatory set A. Hence, in the algorithm with input (G, S', H_2', ϕ', k), at each step, either we find an irrelevant vertex and strictly decrease the size of G, or we branch and strictly increase the size of S'.

More precisely, the next result is the main technical ingredient in this part of the proof, essentially stating that a part of the solution S can be found in a set A of size $a_{\mathcal{F}}$ in which every vertex is adjacent to many vertices of a big enough wall. This is our "obligatory vertex" method. See Figure 6 for an illustration and Section 6 of the full version for the details.

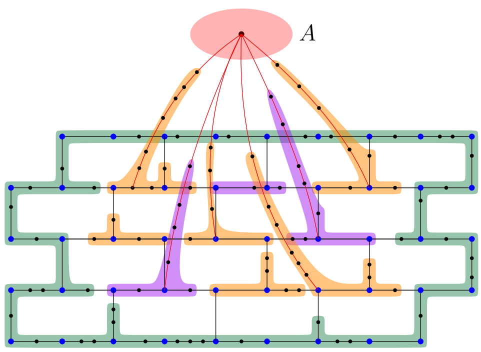

Figure 6 Illustration of Theorem 5.

▶ **Lemma 5.** *Let $\mathcal{F}$ be a finite collection of graphs and $\mathcal{L}$ be a hereditary R-action. There exist three functions $f_1, f_2, f_3 : \mathbb{N} \to \mathbb{N}$ such that the following holds. Let $k \in \mathbb{N}$. Let G be a graph, $S' \subseteq V(G)$ be a set of size at most k, and $(H_2', \phi') \in \mathcal{L}(G[S'])$. Suppose that $G' := G^{S'}_{(H_2',\phi')}$ contains a set $A \subseteq V(G')$ of size at least $a_{\mathcal{F}}$ and that there is a wall W in $G' - A$ of height $f_1(k)$. Suppose also that there is a W-canonical partition $\tilde{\mathcal{Q}}$ of $G' - A$ such that each vertex of A is adjacent to at least $f_2(k)$ many $f_3(k)$-internal bags of $\tilde{\mathcal{Q}}$. Then, for every solution (S, H_2, ϕ) of $\mathcal{L}$-AR-$\mathsf{exc}(\mathcal{F})$ for (G, S', H_2', ϕ'), it holds that $A' \neq \emptyset$, where $A' := (S \setminus S') \cap A$, and that $|\phi^+(A')| < |A'|$. Moreover $f_1(k) = \mathcal{O}_{s_{\mathcal{F}}}(k^2)$, $f_2(k) = \mathcal{O}_{s_{\mathcal{F}}}(k^3)$, and $f_3(k) = \mathcal{O}_{s_{\mathcal{F}}}(k^2)$.*

Finding an irrelevant vertex. As expected, we use the irrelevant vertex technique of Robertson and Seymour [42]. More specifically, we generalize the irrelevant vertex technique used in [38] (actually proved in [45]). This technique is based on the (intuitive but surprisingly hard to prove) fact that the central vertex of a homogeneous flat wall is always irrelevant. While our irrelevant vertex technique for $\mathcal{L}$-AR-$\mathcal{H}$ takes inspiration from [45], it is far more involved due to the annotation and the fact that we allow a wide variety of modifications. In particular, we need to redefine what it means to be homogeneous for a flat wall, to adapt it to our new setting. The previous definition was made to handle the case when we had to remove a small vertex set, called *apex set*, to find a flat wall, and more specifically to handle the fact that some vertices are possibly deleted from the apex set. Now, we also need to handle any other way the apex set may be modified, hence the new definition. The fact that we ask the replacement action $\mathcal{L}$ to be hereditary comes from the irrelevant vertex technique. Indeed, in order to prove that the central vertex v of a homogeneous flat wall W is irrelevant, we essentially prove that, for any solution (S, H_2, ϕ), we can delete a small part X of W containing v, and that the restriction of (S, H_2, ϕ) to $G - X$ is still a solution.

The following is the main technical result that we prove in this part, stating that an irrelevant vertex can be found in a big enough flat wall whose compass has bounded treewidth.

▶ **Theorem 6.** *Let $\mathcal{F}$ be a finite collection of graphs and $\mathcal{L}$ be a hereditary R-action. There exist a function $f_4 : \mathbb{N}^2 \to \mathbb{N}$, whose images are odd integers, and an algorithm with the following specifications:*

`Irrelevant-Vertex`$(G, S', H_2', \phi', k, A, a, W, \mathfrak{R}, t)$

Input: *Integers $k, a, t \in \mathbb{N}$, a graph G, a set $S' \subseteq V(G)$ of size at most k, $(H_2', \phi') \in \mathcal{L}(G[S'])$, a set $A \subseteq V(G')$ of size at most a, where $G' := G^{S'}_{(H_2,\phi)}$, and a regular flatness pair $(W, \mathfrak{R})$ of $G' - A$ of height at least $f_{??}(k, a)$ whose $\mathfrak{R}$-compass has treewidth at most t and does not intersect $\phi'(S')$.*

Output: *A vertex $v \in V(G) \setminus S'$ such that (G, S', H_2', ϕ', k) and $(G - v, S', H_2', \phi', k)$ are equivalent instances of $\mathcal{L}$-AR-$\mathsf{exc}(\mathcal{F})$. Moreover, $f_{??}(k, a) = \mathcal{O}_{a,\ell_{\mathcal{F}}}(k^c)$, where $c = \mathcal{O}_{a,\ell_{\mathcal{F}}}(1)$, and the algorithm runs in time $2^{\mathcal{O}_{a,\ell_{\mathcal{F}}}(k \log k + t \log t)} \cdot (n + m)$.*

We also prove the following result for the bounded genus case with a better dependence on k and a better running time. In this case, we do not ask for our flat wall to have bounded treewidth, but to have a planar embedding instead. Note that here, instead of a single vertex v, we might sometimes find an entire planar block of vertices V that is irrelevant.

▶ **Theorem 7.** *Let $\mathcal{L}$ be a hereditary R-action and $\mathcal{F}$ be the collection of obstructions of the graphs embeddable in a surface of genus at most g. There exist a function $f_5 : \mathbb{N} \to \mathbb{N}$, whose images are odd integers, and an algorithm with the following specifications:*

`Planar-Irrelevant-Vertex`$(G, S', H_2', \phi', k, W, \mathfrak{R})$

Input: *An integer $k \in \mathbb{N}$, a graph G, a set $S' \subseteq V(G)$ of size at most k, $(H_2', \phi') \in \mathcal{L}(G[S'])$, and a flatness pair $(W, \mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi))$ of $G^{S'}_{(H_2',\phi')}$ of height at least $f_{??}(k)$ whose $\mathfrak{R}$-compass does not intersect $\phi'(S')$ and is embeddable in a disk with $X \cap Y$ on the boundary.*

Output: *A non-empty set $Y \subseteq V(G) \setminus S'$ such that (G, S', H_2', ϕ', k) and $(G - Y, S', H_2', \phi', k)$ are equivalent instances of $\mathcal{L}$-AR-exc$(\mathcal{F})$.*

Moreover, $f_5(k) = \mathcal{O}(k)$ and the above algorithm runs in time $\mathcal{O}(n + m)$.

Piecing everything together. Finally, we combine the three ingredients discussed above to find an algorithm for $\mathcal{L}$-AR-$\mathcal{H}$. We essentially proceed as follows. Let (G, S', H_2', ϕ', k) be the instance we want to solve, and G' be obtained by doing the modification (H_2', ϕ') of S'. In the first steps, we either find that G has small treewidth, where we can use our dynamic programming algorithm to conclude, or that G' contains a wall W. Given W, we first try to find a flat wall W' inside, with all the necessary conditions to find an irrelevant vertex. If we manage to do so, we remove the irrelevant vertex and recurse. Otherwise, through a greedy procedure, we try to find an obligatory vertex set A with many disjoint paths to W in G'. If we find such a set, we branch and recurse. If not, we manage to argue that we must have a no-instance, and conclude.

The special case of bounded genus. Our second algorithm (Theorem 3), when $\mathcal{H}$ is a class of graphs embeddable in a surface of bounded Euler genus, uses two additional ideas to get an improved running time. The first one is that here, the obligatory set A is a singleton. Indeed, the size of A is the size of the minimum number of vertices one can remove from an obstruction of $\mathcal{H}$ to make it planar. It is well known that, when $\mathcal{H}$ is such a class, there is some integer t depending on the Euler genus such that $K_{3,t} \notin \mathcal{H}$, and thus, $|A| = 1$. In particular, this implies that we do not need to branch on A, but that we instead immediately find an obligatory vertex. The second idea is about homogeneous flat walls. In the running time $2^{\mathsf{poly}(k)} \cdot n^2$ of the first algorithm, the degree of poly essentially corresponds to the size of the required flat wall to find a big enough homogeneous flat wall, and hence an irrelevant vertex, inside of it. In the case where $\mathcal{H}$ is the class of graphs embeddable in a surface of Euler genus at most g, we prove that we can find a homogeneous flat wall inside a flat wall of smaller size, hence the improved running time. To do so, we prove that, after some preliminary processing, a flat wall that is furthermore embeddable in a disk with the perimeter on its boundary is already homogeneous. Hence, our second algorithm proceeds similarly to the first one, but if we find a flat wall W' in G', we divide W' into $k + 1$ disjoint smaller flat walls and check whether they belong to $\mathcal{H}$. By the pigeonhole principle, one of them, W_i, does not contain a modified vertex and must thus be in $\mathcal{H}$, otherwise we return a no-instance. Then, we argue, using a result from [11] to guarantee additional properties of the planar embedding that are needed for technical reasons, that we can find a smaller flat wall W_i' in W_i with a *planar embedding* (even if the genus of the target graph class is strictly positive). Hence, we find an irrelevant vertex in W_i' and conclude.

5 Conclusion

For a large family of graph modification problems involving a bounded number of vertices, if the target class $\mathcal{H}$ is minor-closed, we provided an algorithm solving the problem in time $2^{\mathsf{poly}(k)} \cdot n^2$. This is actually the same running time as the best known running time for Vertex Deletion to $\mathcal{H}$ [38]. For the other graph modification problems encompassed by our result, such as Edge Deletion to $\mathcal{H}$, Edge Contraction to $\mathcal{H}$, Vertex Identification to $\mathcal{H}$, or Independent Set Deletion to $\mathcal{H}$, the only minor-closed $\mathcal{H}$ for which an algorithm with an explicit parametric dependence in the solution size was known, to the authors' knowledge, were the class of forests and the class of union of paths. Other problems, such as Matching Deletion to $\mathcal{H}$, Matching Contraction to $\mathcal{H}$, Induced Star Deletion to $\mathcal{H}$, or Subgraph Complementation to $\mathcal{H}$, were not even considered yet from the parameterized complexity viewpoint, other than in the meta-theorem of [47].

The degree of $\mathsf{poly}(k)$ in the running time comes from the irrelevant vertex technique and is quite huge. In the bounded genus case, we reduce the running time to $2^{\mathcal{O}(k^9)} \cdot n^2$ thanks to some improvement on the irrelevant vertex technique. This does not match the parametric dependence in the running time of $2^{\mathcal{O}(k^2 \log k)} \cdot n^{\mathcal{O}(1)}$ for Vertex Deletion to $\mathcal{H}$ [29] for $\mathcal{H}$ of bounded genus, though we possibly have a better dependence on n. To the authors' knowledge, this is the first bounded genus result with an explicit parametric dependence in the solution size for the other graph modification problems encompassed by our result.

Improving more the parametric dependence in the general case would certainly require coming up with new techniques. On the other hand, given the recent results of [30] for minor containment, it is worth studying whether the quadratic dependence on n could be improved to an almost-linear dependence while maintaining a good dependence on k. Note that the approach of [30] heavily uses Courcelle's theorem [8], which would require to be translated to a plausibly very involved dynamic programming algorithm to keep a good parametric dependence on k.

Given that we require the replacement action $\mathcal{L}$ to be hereditary for our irrelevant vertex technique to work, we unfortunately restrict the graph modification problems that we solve. For instance, Planar Subgraph Isomorphism can be expressed as an $\mathcal{L}$-R-Planar problem for a specific $\mathcal{L}$, which is not hereditary. Hence, we do not encompass this problem in our general algorithm, while such an algorithm is provided in [17], where the constraint about $\mathcal{L}$ being hereditary is not required. While most of the "reasonable" modification problems correspond to a hereditary replacement action, it is worth investigating whether our result can be extended to non-hereditary replacement actions.

Here, we only consider modifications that affect a bounded number of vertices of the input graph. This is necessary as we want k to decrease by one each time we find an obligatory vertex (or, more precisely, as we want the size of the increasingly guessed partial solution to be bounded by k), so that the depth of the branching tree is bounded. Some relevant graph modification problems, however, such as Elimination Distance to $\mathcal{H}$ [38] or $\mathcal{H}$-Treewidth [14] (where we want to delete a vertex set X whose "torso" has bounded treedepth or treewidth, respectively, such that $G - X \in \mathcal{H}$), consider a modification that affects a set of vertices that may have *unbounded* size. In this case, the branching method does not seem applicable. However, the irrelevant vertex technique still works, and provided that we have a dynamic programming for graphs of bounded treewidth, an algorithm can still be designed in some cases, but with a worse parametric dependence on k. This is what is done, for instance, in [38] for Elimination Distance to $\mathcal{H}$. Therefore, we could consider extending the results of this paper to (some kinds of) modifications involving sets of vertices or edges of unbounded size.

References

1. Akanksha Agrawal, Pallavi Jain, Lawqueen Kanesh, Pranabendu Misra, and Saket Saurabh. Exploring the kernelization borders for hitting cycles. In *Proc. of the 13th International Symposium on Parameterized and Exact Computation (IPEC)*, volume 115 of *LIPIcs*, pages 14:1–14:14, 2018. doi:10.4230/LIPICS.IPEC.2018.14.
2. Akanksha Agrawal, Lawqueen Kanesh, Saket Saurabh, and Prafullkumar Tale. Paths to trees and cacti. *Theoretical Computer Science*, 860:98–116, 2021. doi:10.1016/J.TCS.2021.01.033.
3. Akanksha Agrawal, Saket Saurabh, and Prafullkumar Tale. On the parameterized complexity of contraction to generalization of trees. *Theory of Computing Systems*, 63(3):587–614, 2019. doi:10.1007/S00224-018-9892-Z.
4. Dhanyamol Antony, Sagartanu Pal, and R. B. Sandeep. Algorithms for subgraph complementation to some classes of graphs. *Information Processing Letters*, 188:106530, 2025. doi:10.1016/J.IPL.2024.106530.
5. Julien Baste, Ignasi Sau, and Dimitrios M. Thilikos. Hitting Minors on Bounded Treewidth Graphs. IV. An Optimal Algorithm. *SIAM Journal on Computing*, 52(4):865–912, 2023. doi:10.1137/21M140482X.
6. Marthe Bonamy, Konrad K. Dabrowski, Carl Feghali, Matthew Johnson, and Daniël Paulusma. Independent feedback vertex set for p_5-free graphs. *Algorithmica*, 81(4):1342–1369, 2019. doi:10.1007/S00453-018-0474-X.
7. Jianer Chen, Iyad A. Kanj, and Ge Xia. Improved upper bounds for vertex cover. *Theoretical Computer Science*, 411(40-42):3736–3756, 2010. doi:10.1016/j.tcs.2010.06.026.
8. Bruno Courcelle. The monadic second-order logic of graphs. I. recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, 1990. doi:10.1016/0890-5401(90)90043-H.
9. Christophe Crespelle, Pål Grønås Drange, Fedor V. Fomin, and Petr Golovach. A survey of parameterized algorithms and the complexity of edge modification. *Computer Science Review*, 48:100556, 2023. doi:10.1016/j.cosrev.2023.100556.
10. Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3.
11. Erik D. Demaine, MohammadTaghi Hajiaghayi, and Dimitrios M. Thilikos. The bidimensional theory of bounded-genus graphs. *SIAM Journal on Discrete Mathematics*, 20(2):357–371, 2006. doi:10.1137/040616929.
12. Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013. doi:10.1007/978-1-4471-5559-1.
13. Maël Dumas, Anthony Perez, Mathis Rocton, and Ioan Todinca. Polynomial kernels for edge modification problems towards block and strictly chordal graphs. *CoRR*, abs/2201.13140, 2022. arXiv:2201.13140.
14. Eduard Eiben, Robert Ganian, Thekla Hamm, and O-joung Kwon. Measuring what matters: A hybrid approach to dynamic programming with treewidth. *Journal of Computer and System Sciences*, 121:57–75, 2021. doi:10.1016/j.jcss.2021.04.005.
15. Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2006. doi:10.1007/3-540-29953-X.
16. Fedor V. Fomin, Petr A. Golovach, Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. Compound logics for modification problems. In *Proc. of the 50th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 261 of *LIPIcs*, pages 61:1–61:21, 2023. doi:10.4230/LIPICS.ICALP.2023.61.
17. Fedor V. Fomin, Petr A. Golovach, and Dimitrios M. Thilikos. Modification to planarity is fixed parameter tractable. In *Proc. of the 36th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 126 of *LIPIcs*, pages 28:1–28:17, 2019. doi:10.4230/LIPIcs.STACS.2019.28.

18 Fedor V. Fomin, Daniel Lokshtanov, Neeldhara Misra, and Saket Saurabh. Planar $\mathcal{F}$-Deletion: Approximation, Kernelization and Optimal FPT Algorithms. In *Proc. of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 470–479, 2012. doi:10.1109/FOCS.2012.62.

19 Fedor V. Fomin, Saket Saurabh, and Neeldhara Misra. Graph modification problems: A modern perspective. In *Proc. of the 9th International Workshop on Frontiers in Algorithmics (FAW)*, volume 9130 of *LNCS*, pages 3–6. Springer, 2015. doi:10.1007/978-3-319-19647-3_1.

20 Hanna Furmanczyk, Marek Kubale, and Stanislaw P. Radziszowski. On bipartization of cubic graphs by removal of an independent set. *Discrete Applied Mathematics*, 209:115–121, 2016. doi:10.1016/J.DAM.2015.10.036.

21 Petr A. Golovach, Pim van 't Hof, and Daniël Paulusma. Obtaining planarity by contracting few edges. *Theoretical Computer Science*, 476:38–46, 2013. doi:10.1016/J.TCS.2012.12.041.

22 Pinar Heggernes, Pim van 't Hof, Benjamin Lévêque, Daniel Lokshtanov, and Christophe Paul. Contracting graphs to paths and trees. *Algorithmica*, 68(1):109–132, 2014. doi:10.1007/S00453-012-9670-2.

23 Yoichi Iwata and Yusuke Kobayashi. Improved analysis of highest-degree branching for feedback vertex set. *Algorithmica*, 83(8):2503–2520, 2021. doi:10.1007/S00453-021-00815-W.

24 Bart M. P. Jansen, Daniel Lokshtanov, and Saket Saurabh. A near-optimal planarization algorithm. In *Proc. of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1802–1811, 2014. doi:10.1137/1.9781611973402.130.

25 Ken-ichi Kawarabayashi. Planarity allowing few error vertices in linear time. In *Proc. of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 639–648, 2009. doi:10.1109/FOCS.2009.45.

26 Ken-ichi Kawarabayashi, Yusuke Kobayashi, and Bruce A. Reed. The disjoint paths problem in quadratic time. *Journal of Combinatorial Theory, Series B*, 102(2):424–435, 2012. doi:10.1016/j.jctb.2011.07.004.

27 Ken-ichi Kawarabayashi, Robin Thomas, and Paul Wollan. A new proof of the flat wall theorem. *Journal of Combinatorial Theory, Series B*, 129:204–238, 2018. doi:10.1016/j.jctb.2017.09.006.

28 Eun Jung Kim, Alexander Langer, Christophe Paul, Felix Reidl, Peter Rossmanith, Ignasi Sau, and Somnath Sikdar. Linear kernels and single-exponential algorithms via protrusion decompositions. *ACM Transactions on Algorithms*, 12(2):21:1–21:41, 2016. doi:10.1145/2797140.

29 Tomasz Kociumaka and Marcin Pilipczuk. Deleting Vertices to Graphs of Bounded Genus. *Algorithmica*, 81(9):3655–3691, 2019. doi:10.1007/s00453-019-00592-7.

30 Tuukka Korhonen, Michal Pilipczuk, and Giannos Stamoulis. Minor containment and disjoint paths in almost-linear time. In *Proc. of the 65th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 53–61, 2024. doi:10.1109/FOCS61266.2024.00014.

31 R. Krithika, Pranabendu Misra, and Prafullkumar Tale. A single exponential-time FPT algorithm for cactus contraction. *Theoretical Computer Science*, 954:113803, 2023. doi:10.1016/J.TCS.2023.113803.

32 John M. Lewis and Mihalis Yannakakis. The node-deletion problem for hereditary properties is NP-complete. *Journal of Computer and System Sciences*, 20(2):219–230, 1980. doi:10.1016/0022-0000(80)90060-4.

33 Shaohua Li and Marcin Pilipczuk. An improved FPT algorithm for independent feedback vertex set. *Theory of Computing Systems*, 64(8):1317–1330, 2020. doi:10.1007/S00224-020-09973-W.

34 Wenjun Li, Qilong Feng, Jianer Chen, and Shuai Hu. Improved kernel results for some FPT problems based on simple observations. *Theoretical Computer Science*, 657:20–27, 2017. doi:10.1016/J.TCS.2016.06.012.

35 Carlos V. G. C. Lima, Dieter Rautenbach, Uéverton S. Souza, and Jayme Luiz Szwarcfiter. On the computational complexity of the bipartizing matching problem. *Annals of Operations Research*, 316(2):1235–1256, 2022. doi:10.1007/S10479-021-03966-9.

36 Carlos Vinícius G. C. Lima, Dieter Rautenbach, Uéverton S. Souza, and Jayme Luiz Szwarcfiter. Decycling with a matching. *Information Processing Letters*, 124:26–29, 2017. doi:10.1016/J.IPL.2017.04.003.

37 Dániel Marx and Ildikó Schlotter. Obtaining a planar graph by vertex deletion. *Algorithmica*, 62(3-4):807–822, 2012. doi:10.1007/s00453-010-9484-z.

38 Laure Morelle, Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. Faster parameterized algorithms for modification problems to minor-closed classes. *TheoretiCS*, 3, 2024. doi:10.46298/THEORETICS.24.19.

39 Laure Morelle, Ignasi Sau, and Dimitrios M. Thilikos. Vertex identification to a forest. *CoRR*, 2024. doi:10.48550/arXiv.2409.08883.

40 Rolf Niedermeier. *Invitation to fixed parameter algorithms*, volume 31. Oxford University Press, 2006. doi:10.1093/ACPROF:OSO/9780198566076.001.0001.

41 Fábio Protti and Uéverton S. Souza. Decycling a graph by the removal of a matching: new algorithmic and structural aspects in some classes of graphs. *Discrete Mathematics & Theoretical Computer Science*, 20(2), 2018. doi:10.23638/DMTCS-20-2-15.

42 Neil Robertson and Paul D. Seymour. Graph Minors. XIII. The Disjoint Paths Problem. *Journal of Combinatorial Theory, Series B*, 63(1):65–110, 1995. doi:10.1006/jctb.1995.1006.

43 Neil Robertson and Paul D. Seymour. Graph Minors. XX. Wagner's conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004. doi:10.1016/j.jctb.2004.08.001.

44 Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. k-apices of minor-closed graph classes. II. Parameterized Algorithms. *ACM Transactions on Algorithms*, 18(3):21:1–21:30, 2022. doi:10.1145/3519028.

45 Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. k-apices of minor-closed graph classes. I. Bounding the obstructions. *Journal of Combinatorial Theory, Series B*, 161:180–227, 2023. doi:10.1016/J.JCTB.2023.02.012.

46 Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. A more accurate view of the Flat Wall Theorem. *Journal of Graph Theory*, 107(2):263–297, 2024. doi:10.1002/jgt.23121.

47 Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. Parameterizing the quantification of CMSO: model checking on minor-closed graph classes. In *Proc. of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3728–3742, 2025. doi:10.1137/1.9781611978322.124.

48 Dekel Tsur. Faster deterministic algorithm for co-path set. *Information Processing Letters*, 180:106335, 2023. doi:10.1016/J.IPL.2022.106335.

49 Mihalis Yannakakis. Node- and Edge-Deletion NP-Complete Problems. In *Proc. of the 10th Annual ACM Symposium on Theory of Computing*, pages 253–264, 1978. doi:10.1145/800133.804355.