Kernelization Dichotomies for Hitting Minors Under Structural Parameterizations

Marin Bougeret
LIRMM, Université de Montpellier, CNRS, Montpellier, France

Eric Brandwein
Departamento de Computación, FCEyN, Universidad de Buenos Aires, Argentina

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

Abstract

For a finite collection of connected graphs $\mathcal{F}$, the $\mathcal{F}$-Minor Deletion problem consists in, given a graph G and an integer ℓ, deciding whether G contains a vertex set of size at most ℓ whose removal results in an $\mathcal{F}$-minor-free graph. We lift the existence of (approximate) polynomial kernels for $\mathcal{F}$-Minor Deletion by the solution size to (approximate) polynomial kernels parameterized by the vertex-deletion distance to graphs of bounded elimination distance to $\mathcal{F}$-minor-free graphs. This results in exact polynomial kernels for every family $\mathcal{F}$ that contains a planar graph, and an approximate polynomial kernel for Planar Vertex Deletion. Moreover, combining our result with a previous lower bound, we obtain the following infinite set of dichotomies, assuming $\mathsf{NP} \not\subseteq \mathsf{coNP}/\mathsf{poly}$: for any finite set $\mathcal{F}$ of biconnected graphs on at least three vertices containing a planar graph, and any minor-closed class of graphs $\mathcal{C}$, $\mathcal{F}$-Minor Deletion admits a polynomial kernel parameterized by the vertex-deletion distance to $\mathcal{C}$ if and only if $\mathcal{C}$ has bounded elimination distance to $\mathcal{F}$-minor-free graphs. For instance, this yields dichotomies for Cactus Vertex Deletion, Outerplanar Vertex Deletion, and Treewidth-t Vertex Deletion for every integer $t \geq 0$. Prior to our work, such dichotomies were only known for the particular cases of Vertex Cover and Feedback Vertex Set. Our approach builds on the techniques developed by Jansen and Pieterse [Theor. Comput. Sci. 2020] and also uses adaptations of some of the results by Jansen, de Kroon, and Włodarczyk [STOC 2021].

2012 ACM Subject Classification Theory of computation → Parameterized complexity and exact algorithms; Mathematics of computing → Combinatorics

Keywords and phrases hitting forbidden minors, parameterized complexity, polynomial kernel, structural parameterization, elimination distance, kernelization lower bound

Digital Object Identifier 10.4230/LIPIcs.STACS.2026.17

Related Version *Full Version*: `http://arxiv.org/abs/2512.13210`

Funding French project ELiT (ANR-20-CE48-0008-01).

1 Introduction

The field of *parameterized complexity* studies the computational complexity of problems when a parameter $k \in \mathbb{N}$ is given in addition to the input. One of the main objectives of the field is to find efficient preprocessing algorithms called *kernelization algorithms* (or *kernels*), which are polynomial-time algorithms that transform an instance of a parameterized problem into an equivalent instance whose size is bounded by a function of the parameter k. Of particular interest are *polynomial kernels*, which are kernels that produce instances of size bounded by a polynomial in k. See [13, 20, 23, 25, 46] for monographs on the area.

A very active direction within kernelization deals with so-called *structural parameters*. The idea is, for a given problem Π, to unveil the "smallest" parameter (usually related to the structure of the input graph) for which Π admits a polynomial kernel. Ideally, the holy grail

43rd International Symposium on Theoretical Aspects of Computer Science (STACS 2026).
Editors: Meena Mahajan, Florin Manea, Annabelle McIver, and Nguyễn Kim Thắng
Article No. 17; pp. 17:1–17:19
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

is to find a dichotomy describing which parameterizations allow for a polynomial kernel and which do not, subject to reasonable complexity assumptions. Not surprisingly, finding such dichotomies turns out to be quite hard, as we proceed to discuss.

The Vertex Cover problem, which consists in deciding whether a graph G contains a vertex set of size at most ℓ that intersects all edges, has usually served as a testbed for new techniques in parameterized complexity, and in particular in kernelization with structural parameters. Given that Vertex Cover is well-known to admit a polynomial kernel parameterized by the size of the desired solution [25], the challenge is to find parameters, potentially smaller than the size of a minimum vertex cover (which is called the *vertex cover number*), that still permit to obtain polynomial kernels. A very convenient and robust way of describing such structural parameters is by considering the *vertex-deletion distance* of the input graph G to a fixed graph class $\mathcal{C}$, defined as the minimum size of a vertex set X such that $G \setminus X \in \mathcal{C}$; such a set X is called a *modulator* to $\mathcal{C}$. Note that the vertex cover number corresponds to the vertex-deletion distance to the class of empty graphs.

Bodlaender and Jansen [35] proved a very influential result in this direction, namely a polynomial kernel for Vertex Cover parameterized by the *feedback vertex number* of the input graph, that is, the vertex-deletion distance to the class of forests. This result triggered a number of polynomial kernels for Vertex Cover parameterized by the vertex-deletion distance to other graph classes, such as graphs of maximum degree two [44], graphs of constant treedepth [8], pseudo-forests [26], or d-pseudo-forests [31]. It is worth noting that all the classes $\mathcal{C}$ mentioned so far are *minor-closed*, that is, if a graph is in $\mathcal{C}$, then any graph obtained from it by removing vertices or edges, or by contracting edges, is also in $\mathcal{C}$. Bougeret, Jansen, and Sau [5] culminated this line of research by proving the following dichotomy: assuming $\mathsf{NP} \not\subseteq \mathsf{coNP/poly}$ (which is the standard hypothesis in this area), Vertex Cover parameterized by the vertex-deletion distance to a minor-closed graph class $\mathcal{C}$ admits a polynomial kernel if and only if $\mathcal{C}$ has bounded *bridge-depth*. Here, bridge-depth is a newly introduced graph parameter that can be seen as a common generalization of feedback vertex number and tree-depth, in the sense that it is (functionally) smaller than both of them; see [5] for the precise definition.

It is worth mentioning that, using randomized algorithms with a small error probability, polynomial kernels for Vertex Cover are also known for several parameterizations by the vertex-deletion distance to graph classes that are not minor-closed, such as König graphs [42], bipartite graphs [42], and parameterizations based on the linear programming relaxation of Vertex Cover [32, 41]. However, for non-minor-closed graph classes, we are still far from a dichotomy. Thus, if one aims at obtaining similar dichotomies for generalizations of Vertex Cover, it is reasonable to stick to parameterizations defined as the vertex-deletion distance to a *minor-closed* graph class, and this is what we do in this article.

A very natural way of generalizing the Vertex Cover problem is by fixing a finite family of graphs $\mathcal{F}$ and considering the $\mathcal{F}$-Minor Deletion problem, defined as follows: given a graph G and an integer ℓ, the goal is to decide whether at most ℓ vertices can be removed from G so that the resulting graph does not contain any of the graphs in $\mathcal{F}$ as a minor. Note that Vertex Cover corresponds to the case $\mathcal{F} = \{K_2\}$. The $\mathcal{F}$-Minor Deletion problem has attracted great interest in the last years within the parameterized complexity community [3, 24, 27, 29, 38, 40, 49, 50], in particular in kernelization. Namely, when parameterizing by the solution size, Fomin et al. [24] showed that $\mathcal{F}$-Minor Deletion admits a *randomized* polynomial kernel whenever $\mathcal{F}$ contains at least one planar graph. It can be checked that the only randomized step in their kernel is a constant-factor approximation for the problem. In a subsequent work, Gupta et al. [29] provided a *deterministic* constant-factor approximation for $\mathcal{F}$-Minor Deletion, which together with the proof of Fomin et al. [24] yield a deterministic polynomial kernel for $\mathcal{F}$-Minor Deletion parameterized by the solution size when $\mathcal{F}$ contains a planar graph. For collections $\mathcal{F}$ containing only non-planar graphs, the existence of polynomial kernels for $\mathcal{F}$-Minor Deletion is one of the most notorious open problems in the field of kernelization [24, 25, 27, 34, 38, 51].

Probably, the most relevant open case is the case $\mathcal{F} = \{K_5, K_{3,3}\}$, commonly known as Planar Vertex Deletion, which is conjectured to admit a polynomial kernel [34]. On the positive side, recently Jansen and Włodarczyk [38] presented an *approximate* kernel for Planar Vertex Deletion by using intricate topological arguments.

Given the apparent hardness of finding polynomial kernels for the general $\mathcal{F}$-Minor Deletion problem parameterized by the solution size, it is natural to consider structural parameters that are not necessarily smaller than the solution size. In an article that is crucial to our work, Jansen and Pieterse [37] provided a polynomial kernel for $\mathcal{F}$-Minor Deletion, when $\mathcal{F}$ contains only connected graphs, parameterized by the vertex-deletion distance to a graph of constant treedepth. Note that this parameter may indeed be larger than the solution size.

So far, there is only one more particular case of $\mathcal{F}$-Minor Deletion, other than Vertex Cover, for which a kernelization dichotomy is known. This "outlier" is Feedback Vertex Set, corresponding to the case $\mathcal{F} = \{K_3\}$, and for which polynomial kernels by the solution size are well-known [25]. One may expect that, similarly to the dichotomy for Vertex Cover discussed above [5], the dichotomy for Feedback Vertex Set is also determined by (some variation of) bridge-depth. But somehow surprisingly, Dekker and Jansen [16] recently showed that, assuming $\mathsf{NP} \not\subseteq \mathsf{coNP/poly}$, Feedback Vertex Set parameterized by the vertex-deletion distance to a minor-closed graph class $\mathcal{C}$ admits a polynomial kernel if and only if $\mathcal{C}$ has bounded elimination distance to a forest. The *elimination distance* to a graph class $\mathcal{H}$ is a parameter introduced by Bulian and Dawar [9, 10] and defined as the minimum number of rounds needed to recursively delete one vertex from each connected component of the current graph until obtaining a graph that belongs to $\mathcal{H}$ (see Section 3 for the formal definition). Note that treedepth corresponds to the particular case where $\mathcal{H}$ is the class of empty graphs. Note also that, for any collection of graphs $\mathcal{F}$ and any graph G, a solution to $\mathcal{F}$-Minor Deletion in G is a modulator to the class of graphs of elimination distance zero to $\mathcal{F}$-minor-free graphs. Thus, a polynomial kernel for $\mathcal{F}$-Minor Deletion parameterized by the size of a solution is a weaker result than a polynomial kernel parameterized by the vertex-deletion distance to graphs of bounded elimination distance to $\mathcal{F}$-minor-free graphs.

Our results. In a nutshell, our contribution is to lift polynomial kernels for $\mathcal{F}$-Minor Deletion parameterized by the solution size (if they exist) to polynomial kernels parameterized by the vertex-deletion distance to graphs of bounded elimination distance to $\mathcal{F}$-minor-free graphs. This result also holds for approximate kernels by preserving the same approximation factor. We first provide a formal statement of our result and then discuss some of its consequences. We use the notation $\mathsf{ed}_{\mathcal{F}}(G)$ to denote the elimination distance of a graph G to the class of $\mathcal{F}$-minor-free graphs.

▶ **Theorem 1.** *For every fixed finite set $\mathcal{F}$ of connected graphs, every integer $\eta \geq 0$, and every positive constant α, if $\mathcal{F}$-Minor Deletion parameterized by the size of a given solution admits a polynomial (α-approximate) kernel, then $\mathcal{F}$-Minor Deletion parameterized by the size of a given modulator to graphs with $\mathsf{ed}_{\mathcal{F}} \leq \eta$ admits a polynomial (α-approximate) kernel.*

Recall that if $\mathcal{F}$ contains a planar graph, then $\mathcal{F}$-Minor Deletion is known to admit a polynomial kernel parameterized by the solution size [24, 29]. Thus, Theorem 1 implies polynomial kernels for $\mathcal{F}$-Minor Deletion parameterized by the size of a given modulator to graphs with bounded $\mathsf{ed}_{\mathcal{F}}$ whenever $\mathcal{F}$ contains at least one planar graph (in Subsection 4.2 we discuss that, in fact, the hypothesis that the modulator is given is not necessary). Prior to our work, this was only known for Vertex Cover ($\mathcal{F} = \{K_2\}$) [8] and Feedback Vertex Set ($\mathcal{F} = \{K_3\}$) [16]. Some relevant problems covered by our result are Cactus Vertex Deletion [2, 22, 53], Outerplanar Vertex Deletion ($\mathcal{F} = \{K_4, K_{2,3}\}$) [19], Pumpkin Hitting Set [39], d-Pseudoforest Deletion [47], or Treewidth-t Vertex Deletion [14, 52] for every integer $t \geq 0$, that is, the problem of finding a smallest modulator to graphs of treewidth at most t (note that the cases $t = 0$ and $t = 1$ correspond, respectively, to Vertex Cover and Feedback Vertex Set). For this latter problem, it is easy to verify that, for every $t \geq 0$, all minor obstructions to graphs of treewidth at most t are biconnected, and that at least one of them is planar. Other examples of problems encompassed by Theorem 1 are Pathwidth-t Vertex Deletion, Treedepth-t Vertex Deletion, and Branchwidth-t Vertex Deletion for every integer $t \geq 0$.

It turns out that Theorem 1 yields infinitely many kernelization dichotomies for $\mathcal{F}$-Minor Deletion. Indeed, Dekker and Jansen [16] proved that, assuming that $\mathsf{NP} \not\subseteq \mathsf{coNP/poly}$, for any finite collection $\mathcal{F}$ of biconnected graphs on at least three vertices containing at least one planar graph[1], the $\mathcal{F}$-Minor Deletion problem does not admit a polynomial kernel parameterized by the size of a given modulator to a graph of unbounded $\mathsf{ed}_{\mathcal{F}}$. Thus, this lower bound combined with Theorem 1 yields the following result.

▶ **Theorem 2.** *Let $\mathcal{C}$ be a minor-closed class of graphs and let $\mathcal{F}$ be a finite set of biconnected graphs on at least three vertices containing at least one planar graph. Assuming that $\mathsf{NP} \not\subseteq \mathsf{coNP/poly}$, $\mathcal{F}$-Minor Deletion admits a polynomial kernel in the size of a $\mathcal{C}$-modulator if and only if $\mathcal{C}$ has bounded elimination distance to the class of $\mathcal{F}$-minor-free graphs.*

Theorem 2 can be seen as a vast generalization of the dichotomy for Feedback Vertex Set by Dekker and Jansen [16], which was the only one known so far other than Vertex Cover [5]. Concrete examples of other problems covered by Theorem 2 are Cactus Vertex Deletion, Pumpkin Hitting Set, Outerplanar Vertex Deletion, Treewidth-t Vertex Deletion and Branchwidth-t Vertex Deletion for every $t \geq 0$, or C_p Hitting Set [28] for every $p \geq 3$ (that is, the problem of hitting all cycles of length at least p).

On the other hand, plugging the approximate kernel for Planar Vertex Deletion by Jansen and Włodarczyk [38] in Theorem 1 we get the following result, which is a significant strengthening of their kernel [38], corresponding to the case $\eta = 0$.

▶ **Theorem 3.** *For every integer $\eta \geq 0$, the* Planar Vertex Deletion *problem parameterized by the size of a given modulator to a graph of elimination distance to planar graphs at most η admits a polynomial α-approximate kernel, for some constant $\alpha > 1$.*

Donkers and Jansen [18] asked whether, for every collection $\mathcal{F}$, the $\mathcal{F}$-Minor Deletion problem admits a polynomial kernel when parameterized by the vertex-deletion distance to a linear forest, that is, a disjoint collection of paths. Theorem 1 provides a positive answer to

[1] As we discuss in Subsection 4.2, the statement of [16, Theorem 2] requires *all* the graphs in $\mathcal{F}$ to be planar, but the same proof goes through if only one of them is planar, as acknowledged by one of the authors [33].

their question for every collection $\mathcal{F}$ of connected graphs that are not paths containing a planar graph (indeed, in that case, no graph in $\mathcal{F}$ is a minor of a path, so a linear forest has elimination distance zero to the class of $\mathcal{F}$-minor-free graphs).

Finally, let us mention another interpretation of our results. Agrawal et al. [1] proved, among other results, that for every hereditary target graph class $\mathcal{C}$ satisfying some mild assumptions, parameterizing by the vertex-deletion distance to $\mathcal{C}$ and by the elimination distance to $\mathcal{C}$ are equivalent from the point of view of the existence of fixed-parameter tractable algorithms. Theorem 1 implies, in particular, that the same kind of equivalence holds with respect to the existence of polynomial (approximate) kernels in this "distance from triviality" setting, namely for problems defined by the exclusion of connected minors.

Discussion and further research. The most natural open problem that pops up from our work is to unveil the right kernelization dichotomies for the (connected) families $\mathcal{F}$ that are not covered by Theorem 2. For instance, the panorama is already hazy for $\mathcal{F} = \{P_3\}$, the path on three vertices, which is not biconnected. Does bridge-depth play a role in the dichotomy for this problem?

Another apparently challenging problem is to get rid of the connectivity assumption about the graphs in the collection $\mathcal{F}$ in Theorem 1. This is crucially exploited several times in our approach, which strongly builds on the one by Jansen and Pieterse for treedepth [37]. Note that the case where $\mathcal{F}$ contains disconnected graphs in indeed relevant, as for example it is well-known that the minor obstruction set of any surface of positive genus contains disconnected graphs [45].

Our main result (Theorem 1) fits within the active line of work of using *hybrid parameterizations*, which simultaneously capture the connectivity structure of the input instance (typically, via a width parameter) and properties of its optimal solutions; see [36] and the references therein for a complete account. One of the main takeaways of the breakthrough article by Jansen, de Kroon, and Włodarczyk [36] is that, for many natural vertex-deletion problems to a graph class $\mathcal{C}$, including $\mathcal{F}$-Minor Deletion, FPT algorithms parameterized by the solution size can be lifted to FPT algorithms parameterized by the elimination distance to $\mathcal{C}$. With this viewpoint, Theorem 1 can be seen as an analog to this lifting result of [36] with respect to the existence of polynomial kernels. In fact, the result of [36] even holds for a more general parameter recently defined by Eiben et al. [21] called $\mathcal{C}$-*treewidth*, which is a common generalization of treewidth and elimination distance. Thus, for the families $\mathcal{F}$ not covered by Theorem 2, we may hope to generalize Theorem 1 to the parameterization by the size of a modulator to graphs with bounded $\mathcal{H}$-treewidth, where $\mathcal{H}$ is the class of $\mathcal{F}$-minor-free graphs. However, this looks implausible, as Cygan et al. [14] showed that already the "simplest" case of Vertex Cover ($\mathcal{F} = \{K_2\}$) does not admit a polynomial kernel when parameterized by the size of a modulator to graphs with treewidth two, assuming $\mathsf{NP} \not\subseteq \mathsf{coNP/poly}$.

Another way we may try to go beyond Theorem 1 is by replacing the $\mathcal{F}$-Minor Deletion problem with its (induced) subgraph counterpart, where we want to delete vertices to make the graph $\mathcal{F}$-subgraph-free (resp. $\mathcal{F}$-induced-subgraph-free). However, Bougeret, Jansen, and Sau [6] recently showed that, for some choices of $\mathcal{F}$, these problems do not admit polynomial kernels when parameterized by the size of a modulator to graphs with bounded treedepth, unless $\mathsf{NP} \subseteq \mathsf{coNP/poly}$. As the elimination distance to any graph class is not greater than the treedepth, these problems also do not admit polynomial kernels when parameterized by the size of a modulator to graphs with bounded $\mathsf{ed}_{\mathcal{F}}$.

Organization. In Section 2 we present a summary of our techniques, which are based on proving two main ingredients, and a road map of the whole proof. In Section 3 we introduce the necessary definitions and notation, including the concepts of labeled graphs, and the extension of the minor relation to these graphs. In Section 4, assuming that the two main ingredients are proved, we provide the proof of Theorem 1, and its consequences mentioned above. Due to space limitations, the proof of the first and second ingredients can be found, respectively, in Sections 5 and 6 of the full version of this article available online. Whenever we make a citation to a result or to a section of the full version, we refer to this version available online: `http://arxiv.org/abs/2512.13210`.

2 Summary of our techniques

Our techniques are strongly based on those used by Jansen and Pieterse in [37], and in this section we explain the main ideas of this approach and which are our main technical contributions that allow us to obtain Theorem 1. We start in Subsection 2.1 by surveying which is the most common strategy used in the literature for kernelization with structural parameters, and we abstract it in terms of two main ingredients, which we explain in Subsection 2.2 and Subsection 2.3, respectively, for our particular setting.

Given the required amount of technical definitions and lemmas, the entire proof of these two ingredients has been moved to the full version available online, and the goal of this section is to provide some insight on the most important notions involved in the proofs, as well as highlighting our main technical novelties with respect to the proof in [37]. As it will become clear in this section, many of the ingredients that we need are either borrowed directly from [37], or follow from the corresponding results in [37] with very minor modifications. In the latter case, for the sake of completeness we provide (in the full version) both a full proof and a sketch of proof in which we just list which parts of the proof in [37] need to be changed.

We conclude this section with a road map of the whole proof (Figure 3), where one can see how our main technical contributions fit within the structure of the proof.

2.1 Typical approach for kernelization with structural parameters

Let us start by explaining how kernels usually work for hitting problems parameterized by the size of modulator to trivial classes $\mathcal{C}$. More precisely, we consider here $\mathcal{F}$-Minor Deletion problems where the input is (G, X, k), and we have to decide whether at most k vertices can be removed from G so that the resulting graph does not contain any of the graphs in $\mathcal{F}$ as a minor. The modulator $X \subseteq V(G)$ given in the input is such that $m(G \setminus X) \leq \eta$, for a fixed integer $\eta \geq 0$ and graph measure m, and the parameter is $|X|$. We assume that m is such that for disjoint graphs G_1, G_2, $m(G_1 \cup G_2) \leq \max(m(G_1), m(G_2))$, and for any connected graph G, there exists $v \in G$ such that $m(G - v) < m(G)$. Recall that this setup encapsulates several previous works, for example [7] where $\mathcal{F} = \{K_2\}$ and m is the treedepth, [15] where $\mathcal{F} = \{K_3\}$ and m is the elimination distance to a forest, and [37] where $\mathcal{F}$ is arbitrary and m is the treedepth.

Let $\mathcal{D}$ be the set of connected components of $G \setminus X$ and $n_D = |\mathcal{D}|$. All kernels in the mentioned articles follow the same two steps:

1. Remove some connected components of $\mathcal{D}$ until n_D becomes polynomial in $|X|$. This leads to an equivalent instance (G', X, k'). Let $\mathcal{D}'$ denote the set of connected components of $G' \setminus X$.

2. For each $C \in \mathcal{D}'$, find a vertex v_C such that $m(C - v_C) < m(C)$. Define $X' = X \cup \bigcup_{C \in \mathcal{D}'}$ and recurse on (G', X', k').

Informally, we moved vertices from $G' \setminus X$ to the modulator to get a slightly larger modulator X', but such that $m(G' \setminus X') < \eta$, implying that we can apply a recursive argument on the measure. Notice that when following this approach, the only challenge is to achieve item 1. As our setting for the $\mathcal{F}$-Minor Deletion problem and $m = \mathsf{ed}_{\mathcal{F}}$ fit into this framework, we also follow these two steps, and our only goal is to prove the following lemma, which formalizes item 1. This lemma corresponds exactly to [37, Lemma 6] (which provides a polynomial kernel for $\mathcal{F}$-Minor Deletion where $\mathsf{td}(G \setminus X) \leq \eta$, instead of $\mathsf{ed}_{\mathcal{F}}(G \setminus X) \leq \eta$ in our case), where td is replaced by $\mathsf{ed}_{\mathcal{F}}$.

▶ **Lemma 4** (Reduce Components Lemma – Generalized version of [37, Lemma 6]). *Let $\mathcal{F}$ be a finite set of connected graphs and let $\eta \geq 0$ be a constant. There is a polynomial-time algorithm that, given a graph G along with a modulator $X \subseteq V(G)$ such that $\mathsf{ed}_{\mathcal{F}}(G \setminus X) \leq \eta$, outputs an induced subgraph G' of G together with an integer Δ such that $\textsc{opt}_{\mathcal{F}}(G) = \textsc{opt}_{\mathcal{F}}(G') + \Delta$ and $G' \setminus X$ has at most $|X|^{O_{\mathcal{F},\eta}(1)}$ connected components. Moreover, a set Y' that hits all $\mathcal{F}$-minors in G' can be extended in polynomial time to a set Y of size $|Y'| + \Delta$ that hits all $\mathcal{F}$-minors in G.*

Assuming the above lemma, Theorem 1 follows immediately by induction on η (see Section 4 for the details), and thus in this overview we only focus on this lemma. Lemma 4 requires two ingredients: Lemma 5 and Lemma 6. These lemmas also correspond to the two ingredients required by [37, Lemma 6], where td is replaced by $\mathsf{ed}_{\mathcal{F}}$ here. Generalizing these two ingredients to $\mathsf{ed}_{\mathcal{F}}$ is the contribution of this paper, and we now aim at explaining the challenges and new ideas behind this generalization. To do so, and following the formalism of [37], we first need to introduce the notion of labeled minor (see Section 3 for formal definitions). For a set X, an *X-labeled graph* G is a graph where each vertex v is equipped with a set $\mathsf{Labels}_G(v) \subseteq X$ of labels. Given two X-labeled graphs G and H, we say that H is a *labeled minor* of G if H is a "classical minor" – ignoring the labels – of G, certified by a minor model ϕ, that additionally satisfies that for any $v \in V(H)$, $\mathsf{Labels}_H(v) \subseteq \bigcup_{u \in \phi(v)} \mathsf{Labels}_G(u)$.

Let us now discuss how labeled minors appear in the kernelization algorithm for $\mathcal{F}$-Minor Deletion (we also refer the reader to introduction of [37] for additional intuition on the role of labeled minors). Suppose that, given an instance (G, X, k) of $\mathcal{F}$-Minor Deletion, we want to remove a connected component C of $G \setminus X$ by defining $G' = G \setminus C$ and $k' = k - \textsc{opt}_{\mathcal{F}}(G[C])$ (where $\textsc{opt}_{\mathcal{F}}$ denotes the smallest solution size for the $\mathcal{F}$-Minor Deletion problem). To prove that (G', X, k') being a yes-instance implies that (G, X, k) is as well, a common approach is to consider a solution S' of (G', X, k'), and find $Y \subseteq C$ of size $\textsc{opt}_{\mathcal{F}}(G[C])$ such that $S = S' \cup Y$ is a solution of G. However, using an arbitrary local optimal solution Y of C (that "only" hits all $\mathcal{F}$ minors in $G[C]$) may not be enough, as there could be $\mathcal{F}$ minor models in G whose fragments intersect both C and $V(G) \setminus C$ (see Figure 1).

Typically, if we consider the F minor model and the labeled graph Q of Figure 1, to prevent this particular model of F, we need that there exists a local optimal solution $Y \subseteq C$ of size $\textsc{opt}_{\mathcal{F}}(G[C])$ which also hits in C the labeled minor Q, called *fragment*. In the real setting, a local optimal solution of $G[C]$ may be even asked to hit *a set* $\mathcal{Q}$ of labeled minors, corresponding to all possible fragments of an $\mathcal{F}$ minor model in G. Checking the existence of such special optimal solution in a connected component of $G \setminus X$ is precisely what we achieve in the first following ingredient.

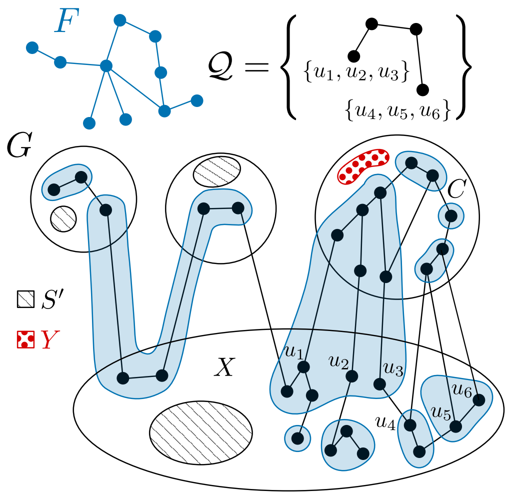

Figure 1 Example where adding Y, a local optimal solution to $\mathcal{F}$-Minor Deletion in $G[C]$, to S', an optimal solution in $G \setminus C$, misses an F-model. This particular model would have been hit if Y was also required to hit the graph in $\mathcal{Q}$ (as a labeled minor).

2.2 Ingredient 1: checking for the existence of special optimal $\mathcal{F}$-Minor Deletion solutions

In Section 5 of the full version we prove the following lemma.

▶ **Lemma 5** (Generalized version of [37, Lemma 5]). *Let $\mathcal{F}$ be a fixed set of connected (unlabeled) graphs, let $\eta \geq 0$ be a constant, and let X be a set. For any set $\mathcal{Q}$ of connected X-labeled graphs and X-labeled graph C with $\mathsf{ed}_{\mathcal{F}}(C) \leq \eta$, one can:*

1. *compute $\text{OPT}_{\mathcal{F}}(C)$ in $O_{\mathcal{F},\eta}(|V(C)|)$ time;*
2. *determine whether there is a solution $Y \in \text{OPTSOL}_{\mathcal{F}}(C)$ such that $C \setminus Y$ has no labeled $\mathcal{Q}$-minors, in time $f(\mathcal{F}, L, \sum_{H \in \mathcal{Q}} |V(H)|, \eta) \cdot |V(C)|^{O(1)}$ for some function f.*

Here, L is defined as the number of elements of X that appear in the labelset of at least one vertex in at least one graph of $\mathcal{Q}$.

The proof of Lemma 5 in [37] quickly follows from the fact that $\mathsf{td}(C) \leq \eta$ implies $\mathsf{tw}(C) \leq \eta$ (where td denotes the treedepth and tw the treewidth), and that the problem of finding labeled minors can be expressed as an MSOL formula. However, in our setting, a graph C with $\mathsf{ed}_{\mathcal{F}}(C) \leq \eta$ may have unbounded treewidth (because of the subgraphs induced by the leaves in the elimination distance decomposition), and we rather rely on the following reduction.

The problem we need to solve is what we later call the $\mathcal{F}$-Minor Deletion Hitting Labeled $\mathcal{Q}$ problem, where given a labeled graph G with $\mathsf{ed}_{\mathcal{F}}(G) \leq \eta$, and a set of labeled graphs $\mathcal{Q}$, one has to decide if there exists an optimal $\mathcal{F}$-Minor Deletion solution for G that also hits all $Q \in \mathcal{Q}$. To solve this problem (in FPT time parameterized by the total size of $\mathcal{F}$ and η), we first reduce (in Subsection 5.2 of the full version) to the unlabeled version called $\mathcal{F}$-Minor Deletion Hitting $\mathcal{Q}$. The idea of the reduction is the following. We define a gadget graph G_ℓ for each label ℓ in X, and glue to each vertex $v \in G$ all gadgets corresponding to labels of v (see the left part of Figure 2). We do the same for each graph in $\mathcal{Q}$. Moreover, we also add a last gadget G_ϵ that we glue to every vertex of G, every vertex of a graph in $\mathcal{Q}$, and every vertex of a graph in $\mathcal{F}$. Let G^+, $\mathcal{Q}^+$, and $\mathcal{F}^+$ denote, respectively, the obtained graphs. To guarantee that $(G, \mathcal{F}, \mathcal{Q}, k)$ is equivalent to $(G^+, \mathcal{F}^+, \mathcal{Q}^+, k)$, we need to control how models of $\mathcal{Q}^+$ and $\mathcal{F}^+$ live in G^+. For example, we want to avoid models of $\mathcal{F}^+$ or $\mathcal{Q}^+$ that invade partially a gadget G_ℓ, or models of $\mathcal{Q}^+$ where the part of the model corresponding to a gadget lives in $V(G)$ (see the right part of Figure 2). This is achieved through the notion of *nice gadgets*, which are informally (see Definition 20 of the full version) biconnected graphs that are pairwise incomparable with respect to the minor relation, and that are not minors of the host graph G. Notice that it is also crucial to guarantee that $\mathsf{ed}_{\mathcal{F}^+}(G^+) \leq \mathsf{ed}_{\mathcal{F}}(G)$, as this implies by the assumption on G that $\mathsf{ed}_{\mathcal{F}^+}(G^+) \leq \eta$, as we wish.

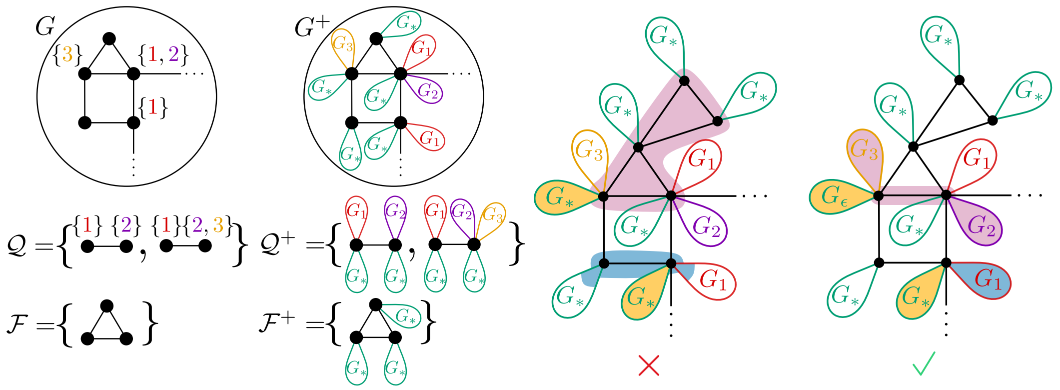

Figure 2 Left: How we compute $(G^+, \mathcal{F}^+, \mathcal{Q}^+)$ from $(G, \mathcal{F}, \mathcal{Q})$ to reduce to the unlabeled version. Right: Example of a "bad" minor model of the second graph in the set $\mathcal{Q}^+$ for two reasons. The first reason is that the pink branch set, which is a minor of G_2 glued with G_3, uses vertices of G instead of the attached G_ℓ. The second reason is that the blue branch set, which is a minor of G_1, uses partially vertices of some G_ℓ. The orange branch sets of this bad model behave as expected. A "good" minor model is depicted on the right.

It remains now to solve $\mathcal{F}$-Minor Deletion Hitting $\mathcal{Q}$ parameterized by the size of $\mathcal{F}$ and η, which we handle in Subsection 5.3 of the full version, by adapting the algorithm of Jansen, de Kroon, and Włodarczyk [36, Theorem 1.2] for $\mathcal{F}$-Minor Deletion parameterized by $\mathcal{H}_\mathcal{F}$-treewidth, where $\mathcal{H}_\mathcal{F}$ is the class of $\mathcal{F}$-minor-free graphs. This parameter generalizes both treewidth and $\mathsf{ed}_\mathcal{F}$. However, the simple strategy consisting in solving directly $(\mathcal{F} \cup \mathcal{Q})$-Minor Deletion using [36] as a black box is not possible, as the input graph G may have unbounded $\mathcal{H}_{\mathcal{F}\cup\mathcal{Q}}$-treewidth, even if its $\mathcal{H}_\mathcal{F}$-treewidth is at most η. Thus, we adapt their result to our case by proving that all key lemmas involved in the dynamic programming algorithm of [36] can be generalized to allow for the presence of $\mathcal{Q}$.

2.3 Ingredient 2: bounding the size of minimal blocking sets

Let us now turn to the second ingredient, which has a purely combinatorial flavor and is our main technical contribution. Recall that our goal is to prove Lemma 4, which removes some connected components of $G \setminus X$. To achieve this, we rely on the marking algorithm of [37]. We chose not to provide the details of this algorithm here (as we do not add any modification to it or to its proof), but only explain why this algorithm requires the second ingredient. To prevent scenarios like the one depicted in Figure 1, the algorithm has to check, before removing a connected component C, that for any set $\mathcal{Q}$ of labeled fragments, C is a yes-instance of $\mathcal{F}$-Minor Deletion Hitting Labeled $\mathcal{Q}$. Notice that for each $\mathcal{Q}$, Ingredient 1 allows us to perform this check in polynomial time, but the problem is that the list of all possible sets $\mathcal{Q}$ may have size exponential in $|X|$. Ingredient 2 (given by Lemma 6 below, which we prove in Section 6 of the full version) exactly fulfills this needs, by proving that if there exists a large set $\mathcal{Q}$ such that C is a no-instance of $\mathcal{F}$-Minor Deletion Hitting Labeled $\mathcal{Q}$, then there exists also a constant-sized subset $\mathcal{Q}^\star \subseteq \mathcal{Q}$ such that C remains a no-instance of $\mathcal{F}$-Minor Deletion Hitting $\mathcal{Q}^\star$. In that way, the marking algorithm will only enumerate (in polynomial time) such sets $\mathcal{Q}^\star$ of constant size.

▶ **Lemma 6** (Main Lemma – Generalized version of [37, Lemma 3]). *Let $\mathcal{F}$ be a finite set of (unlabeled) connected graphs, let X be a set of labels, let $\mathcal{Q}$ be a $(\min_{H\in\mathcal{F}}|V(H)|)$-saturated set of connected X-labeled graphs of at most $\max_{H\in\mathcal{F}}|E(H)|+1$ vertices each, and let C be an X-labeled graph. If all optimal solutions to $\mathcal{F}$-Minor Deletion on C leave a $\mathcal{Q}$-minor, then there is a subset $\mathcal{Q}^* \subseteq \mathcal{Q}$ whose size depends only on $(\mathcal{F}, \mathsf{ed}_{\mathcal{F}}(C))$, such that all optimal solutions leave a $\mathcal{Q}^*$-minor.*

We point out that there is an additional (very helpful) hypothesis required on $\mathcal{Q}$ in the above lemma that we did not discuss so far: the *saturated* property. Indeed, even if this hypothesis is unavoidable (see [37, Figure 11]), and may seem counter-intuitive (for example, the set $\mathcal{Q}$ in Figure 1 does not satisfy it), the fact that we are allowed to assume it comes from the details of the proof of the marking algorithm of [37], which we prefer to keep as a black box in this high-level summary.

We would also like to mention that such a set $\mathcal{Q}$ as in Lemma 6 that "affects" the behavior of all optimal solutions is often referred as a *blocking set* [5, 8, 30], and Lemma 6 can be rephrased as bounding the size of an inclusion-wise minimal blocking set (as invoking the lemma with an inclusion-wise minimal $\mathcal{Q}$ leads to $\mathcal{Q}^\star = \mathcal{Q}$, thus bounding $|\mathcal{Q}|$). One can also observe that the marking algorithm of [37] for $\mathcal{F}$-Minor Deletion corresponds to a generalized version of the marking algorithms used in [5, 8, 35], that is, for Vertex Cover parameterized by the feedback vertex number, the distance to constant treedepth, and the distance to constant bridge-depth, respectively.

Let us now discuss what differs between our proof of Lemma 6 and the proof of [37, Lemma 3]. The proof of Lemma 3 in [37] is inductive on the depth of the treedepth decomposition of C, and we also follow this approach. However, in our setting, C does not necessarily have bounded treedepth, but it does have bounded $\mathsf{ed}_{\mathcal{F}}$. Given the similarities between the two parameters, the only challenge is to add a base case for the leaves of the $\mathcal{F}$-elimination forest of C, which are $\mathcal{F}$-minor-free graphs instead of empty graphs. We consider this new base case, the $\mathcal{F}$-Minor-Free Base Case (Lemma 62 of the full version), to be our main technical contribution. The statement of this lemma is very technical and falls beyond the scope of this overview but, intuitively, it bounds the size of two objects that are relevant to the induction. The first one (denoted by $\mathcal{R}_N$ in the lemma) deals with the number of possible "remainders" of solutions that do *not* leave a $\mathcal{Q}$-minor, while the second one (denoted by $\mathcal{R}_{\mathcal{Q}}$ in the lemma) corresponds to the $\mathcal{Q}^\star$ discussed above. Let us now say a few words about the proof of the second item, since the first one is more technical and would require additional preliminaries.

To simplify the presentation, let us start by stating a simplified version of the setup of Lemma 62 of the full version. Consider a graph G and subgraphs G_A, G_C of G (we use these notation to match the notation of Lemma 62 of the full version, where G_B is assumed to be empty here) such that:

- $S := V(G_A) \cap V(G_C)$ is a separator in G, with $|S| \leq \eta$.
- $G'_A := G_A \setminus S$ is $\mathcal{F}$-minor-free.

Given a set $\mathcal{Q}$ of connected X-labeled graphs, the simplified goal of this second item is to define a set $\mathcal{Q}^\star \subseteq \mathcal{Q}$ whose size only depends on $\mathcal{F}$ and η, and such that for any optimal solution Y of the $\mathcal{F}$-Minor Deletion problem in G, if $G'_A \setminus Y$ leaves a $\mathcal{Q}$-minor, then it also leaves a $\mathcal{Q}^\star$-minor.

In order to bound the size of $\mathcal{Q}^\star$, our strategy is to identify a subset of labels $X' \subseteq X$ of size depending only on $\mathcal{F}$ and η, such that any such solution leaving a $\mathcal{Q}$-minor will also leave a $\mathcal{Q}$-minor that only uses labels from X'. Then, $\mathcal{Q}^*$ will be defined as all graphs of $\mathcal{Q}$ using only labels from X', implying immediately that $|Q^\star|$ only depends on $\mathcal{F}$ and η.

To identify such a restricted set X' of labels, we consider a set called $\mathsf{Breaker}$ corresponding to a minimum-size set in G'_A that hits all $\mathcal{Q}$-minors in G'_A. This setup can be seen as a generalization of the proof of [16, Lemma 4] by Dekker and Jansen, where in their setting (where there are no labels) $\mathcal{Q}$ corresponds to a set $T = \{(u_i, v_i)\}$ (encoding that solutions must hit all (u_i, v_i)-paths) and they consider a minimum-size set Z hitting all these paths.

Then, we use the crucial fact that any optimal solution Y of the $\mathcal{F}$-Minor Deletion problem in G is such that $|Y \cap V(G'_A)| \leq |S|$. Indeed, if Y used strictly more vertices in G'_A, then restructuring the solution by removing all vertices from G'_A and adding the whole of S instead would result in a smaller solution, as all graphs in $\mathcal{F}$ are connected; see Lemma 40 of the full version. We point out that this is the only place where the $\mathcal{F}$-minor-freeness of leaves is used, meaning in particular that we do not need to invoke any complex property on the structure of $\mathcal{F}$-minor-free graphs.

With this set $\mathsf{Breaker}$ and the property that $|Y \cap G'_A| \leq |S|$ at hand, we perform a marking scheme (that generalizes the marking scheme of [16, Lemma 4]) aiming at keeping only marked labels. This marking scheme marks some labels for each $v \in \mathsf{Breaker}$ (implying that we must ensure that $|\mathsf{Breaker}|$ is small), and uses the fact that, as $|Y \cap V(G'_A)| \leq |S|$, we only need to mark a small number of labels to ensure that one of the $\mathcal{Q}$-minors in $G'_A \setminus Y$ only uses marked labels.

Finally, let us just mention that our proof of the first item additionally introduces the definition of a *mandatory* vertex, which is a vertex that appears in every $\mathcal{F}$-Minor Deletion solution meeting some conditions. The set M of all such vertices is helpful in imposing a nice structure in the set of solutions that we need to consider (cf. Figure 10 of the full version).

In Figure 3 we provide a road map of the whole proof, by highlighting our main technical contributions.

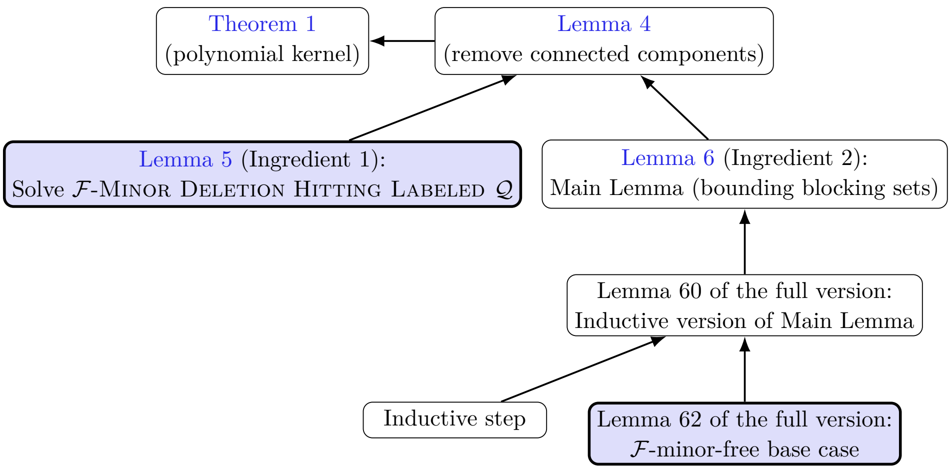

Figure 3 Road map of the proof, which has the same structure as the proof in [37]. Our two contributions correspond to the thicker blues boxes. Results in all other boxes can be obtained directly from the proofs of [37] by minor modifications, essentially by replacing $\mathtt{td}$ with $\mathtt{ed}_{\mathcal{F}}$.

3 Preliminaries

In this section we present some definitions used in this extended abstract. The additional definitions that are needed in the full version are deferred to the full version as well.

Parameterized complexity. A *parameterized problem* is a decision problem where each instance is a pair (I, k), with I the input and $k \in \mathbb{N}$ the parameter. A problem is *fixed-parameter tractable* (FPT) if it can be solved in time $f(k) \cdot |I|^{O(1)}$ for some computable function f.

A *kernelization algorithm* (or *kernel*) for a parameterized problem is a polynomial-time algorithm that transforms any instance (I, k) into an equivalent instance (I', k') (i.e., (I, k) is a yes-instance if and only if (I', k') is), such that $|I'| + k' \leq g(k)$ for some function g. It is known that a parameterized problem is FPT if and only if it admits a kernel [4]. If g is a polynomial, the kernel is called a *polynomial kernel*.

A related notion introduced by Lokshtanov et al. [43] is that of an *α-approximate kernel* for a parameterized optimization problem, for some constant $\alpha \geq 1$. The goal of a parameterized optimization problem is to find a solution of minimum *cost*. In the case of $\mathcal{F}$-Minor Deletion, the cost of a solution is its size. The formal definition of approximate kernelization can be found in [43]; below we give a simpler formulation that is enough for our purposes.

An α-approximate kernel is a polynomial-time algorithm that transforms any instance (I, k) of the parameterized optimization problem into an instance (I', k') such that:

- (I', k') has size at most $g(k)$ for some function g (as in standard kernelization), and
- any solution s' to (I', k') of cost β times the optimum for (I', k') can be transformed in time polynomial in $|I|$, k, $|I'|$, k', and s' into a solution s to (I, k) whose cost is at most $\alpha \cdot \beta$ times the optimum for (I, k).

If g is a polynomial, the kernel is called a *polynomial α-approximate kernel*.

Graphs. For graph notions not defined here, we refer the reader to standard textbooks such as [17]. For a set S, we use 2^S to denote the set of all subsets of S, and $\binom{S}{k}$ to denote the set of all subsets of S of size k. All graphs we consider are finite, undirected, and simple. A graph G consists of a vertex set $V(G)$ and edge set $E(G) \subseteq \binom{V(G)}{2}$. The open neighborhood of a vertex v is denoted $N_G(v)$. For a vertex set $S \subseteq V(G)$, its open neighborhood is $N_G(S) \coloneqq \bigcup_{v \in S} N_G(v) \setminus S$. For an edge $\{u, v\}$ in a graph G, *contracting* $\{u, v\}$ results in the graph G' obtained from G by removing u and v, and replacing them by a new vertex w with $N_{G'}(w) = N_G(\{u, v\})$. For a vertex set $S \subseteq V(G)$, we use $G \setminus S$ to denote the graph obtained from G by deleting all vertices in S and their incident edges. The subgraph of G induced by vertex set S is denoted $G[S]$. We use $\textsc{cc}(G)$ to denote the set of connected components of G. A graph is *biconnected* if it is connected and does not contain a *cut vertex*, i.e., a vertex whose removal increases the number of connected components of the graph. A *biconnected component* of a graph G is a maximal biconnected subgraph of G. A *grid graph of width w and height h* for two constants $w, h \geq 1$ is the graph G where $V(G)$ consists of the tuples (a, b) where $1 \leq a \leq w$ and $1 \leq b \leq h$, and two vertices (a_1, b_1) and (a_2, b_2) are adjacent if and only if $|a_1 - a_2| + |b_1 - b_2| \leq 1$.

For a set of graphs $\mathcal{F}$, we use $\|\mathcal{F}\|$ as a shorthand for $\max_{H \in \mathcal{F}} |V(H)|$. We say that a set $Y \subseteq V(G)$ *hits* all $\mathcal{F}$-minors in G if $G \setminus Y$ is $\mathcal{F}$-minor-free.

We denote the size of an optimal $\mathcal{F}$-MINOR DELETION solution on G by $\text{OPT}_{\mathcal{F}}(G)$, and the set of optimal solutions by $\text{OPTSOL}_{\mathcal{F}}(G)$. In our bounds, we use the notation $O_z(1)$ for a list of values z to denote a constant that only depends on z. For instance, $O_{\mathcal{F},\mathcal{Q}}(1)$ denotes a constant that only depends on the sets of graphs $\mathcal{F}$ and $\mathcal{Q}$.

▶ **Definition 7** (treedepth). *A* treedepth decomposition *of a connected graph G is a rooted tree T such that $V(T) = V(G)$, and for every edge $\{u, v\} \in E(G)$, vertex u is an ancestor of vertex v in T, or vice versa. A treedepth decomposition of a disconnected graph is just a disjoint union of treedepth decompositions for its connected components.*

The treedepth *of G, denoted by $\mathsf{td}(G)$, is the minimum depth (in number of vertices) of a treedepth decomposition of G.*

The following definition was introduced by Bulian and Dawar [9, 10]. In fact, the general definition is for a general target graph class $\mathcal{H}$, but since we only deal with $\mathcal{F}$-minor-free classes, we provide only the following more restricted definition.

▶ **Definition 8** (elimination distance). *Let $\mathcal{F}$ be a finite collection of graphs. An $\mathcal{F}$-*elimination forest *of a connected graph G is a pair (T, χ) of a rooted tree T and a function $\chi\colon V(T) \to 2^{V(G)}$, called the* bags of T, *such that:*

1. $\bigcup_{t \in V(T)} \chi(t) = V(G)$.
2. *For every two nodes $\{t, t'\} \subseteq V(T)$ we have that $\chi(t) \cap \chi(t') = \emptyset$.*
3. *Every internal node $t \in V(T)$ is such that $|\chi(t)| = 1$.*
4. *For every edge $\{u, v\} \in E(G)$ either:*
   - *there exists a leaf $t \in V(T)$ such that $\{u, v\} \subseteq \chi(t)$; or*
   - *there exist two nodes $\{t, t'\} \subseteq V(T)$ such that $u \in \chi(t)$, $v \in \chi(t')$, and t and t' are in an ancestor-descendant relationship in T.*
5. *For every leaf t of T, the induced subgraph $G[\chi(t)]$ is connected and $\mathcal{F}$-minor-free.*

An $\mathcal{F}$-elimination forest of a disconnected graph G is just a disjoint union of $\mathcal{F}$-elimination forests for its connected components.

The elimination distance *of G to the class of $\mathcal{F}$-minor-free graphs is the minimum depth (in number of edges) of an $\mathcal{F}$-elimination forest of G. We denote this distance by $\mathsf{ed}_{\mathcal{F}}(G)$.*

Note that the treedepth is exactly the elimination distance to the class of empty graphs, which is why the treedepth is measured in number of vertices, while the elimination distance is measured in number of edges.

▶ **Definition 9** (minor model). *A* minor model *of a graph H in a graph G is a mapping $\varphi\colon V(H) \to 2^{V(G)}$ assigning a* branch set *$\varphi(v) \subseteq V(G)$ to each vertex $v \in V(H)$, such that:*

- *$G[\varphi(v)]$ is non-empty and connected for all $v \in V(H)$,*
- *$\varphi(v) \cap \varphi(u) = \emptyset$ for all $u \neq v \in V(H)$, and*
- *if $\{u, v\} \in E(H)$, then there exist $u' \in \varphi(u)$ and $v' \in \varphi(v)$ such that $\{u', v'\} \in E(G)$.*

For a vertex set $S \subseteq V(H)$, we define $\varphi(S) := \bigcup_{v \in S} \varphi(v)$ to be the branch set of S. A minor model is minimal *if there is no minor model that results from removing a single vertex from a branch set $\varphi(v)$ for some $v \in V(H)$.*

Labeled graphs. We will be annotating the vertices of graphs with labels from a set X. This set X will be the modulator X that will be the parameter in our kernelization of $\mathcal{F}$-MINOR DELETION. The labels in each vertex will encode the adjacency of that vertex to the modulator X.

▶ **Definition 10** (labeled graph [37, Definition 2]). *Let X be a set. An X-*labeled graph *G is a graph G together with label function* $\mathsf{Labels}_G : V(G) \to 2^X$*, assigning a (potentially empty) subset of labels to each vertex in G.*

We will in fact be looking for minors in these labeled graphs which are connected to the modulator X in certain ways. This is captured by the following definition.

▶ **Definition 11** (labeled minor model [37, Definition 4]). *A* labeled minor model *of an X-labeled graph H in an X-labeled graph G is a mapping φ as in the definition of minor model (Definition 9), that additionally satisfies that for all $v \in V(H)$ and $\ell \in \mathsf{Labels}_H(v)$ there exists $v' \in \varphi(v)$ such that $\ell \in \mathsf{Labels}_G(v')$.*

If G contains a (labeled) minor model of H, then we say that G contains H as a (labeled) minor and denote this as $H \preceq_m G$. Observe that G contains H as a (labeled) minor if and only if H can be obtained from G by deleting edges and vertices (and potentially labels), and contracting edges (merging the labelsets of the corresponding vertices).

4 Proof assuming the two ingredients and consequences

In Subsection 4.1 we provide the proof of Theorem 1 assuming that we have at hand the two ingredients described in Section 2, and in Subsection 4.2 we discuss some of its consequences.

4.1 Proof of the polynomial kernel assuming the two ingredients

In this section we assume that Lemma 5 and Lemma 6 hold, and we prove the following lemma and Theorem 1. Unlike in all other results that we need to slightly modify from [37], for which we provide both a full proof and a sketch, we only provide a sketch of the proof of the following lemma, since it is almost exactly the same as the (very long) proof in [37].

▶ **Lemma 4** (Reduce Components Lemma – Adaptation of [37, Lemma 6]). *Let $\mathcal{F}$ be a finite set of connected graphs and let $\eta \geq 0$ be a constant. There is a polynomial-time algorithm that, given a graph G along with a modulator $X \subseteq V(G)$ such that $\mathsf{ed}_{\mathcal{F}}(G \setminus X) \leq \eta$, outputs an induced subgraph G' of G together with an integer Δ such that $\textsc{opt}_{\mathcal{F}}(G) = \textsc{opt}_{\mathcal{F}}(G') + \Delta$ and $G' \setminus X$ has at most $|X|^{O_{\mathcal{F},\eta}(1)}$ connected components. Moreover, a set Y' that hits all $\mathcal{F}$-minors in G' can be extended in polynomial time to a set Y of size $|Y'| + \Delta$ that hits all $\mathcal{F}$-minors in G.*

Sketch of proof. Modify the proof of Lemma 6 in [37] by replacing:
- treedepth with $\mathsf{ed}_{\mathcal{F}}$;
- usage of their Lemma 3 with the Main Lemma (Lemma 6); and
- usage of their Lemma 5 by our Lemma 5. Note that $\mathcal{H}$ in their proof consists of connected graphs, so the added requirement in Lemma 5 that $\mathcal{Q}$ is connected is not a problem. ◀

We are now ready to prove Theorem 1, but before that, two comments are in place. The first one is that the proof is almost the same as the proof of [37, Theorem 1], but we provide it here because we need to use a polynomial (α-approximate) kernel by the solution size in the case $\eta = 0$, and an extra argument to bound the approximation factor in the case of approximate kernels. The second one is that, for the sake of generality, we state the theorem for the case where the required modulator is *given* along with the input. In the first application discussed in Subsection 4.2 we show how to get rid of this hypothesis.

Proof of Theorem 1. Consider an input (G, X, k) to $\mathcal{F}$-Minor Deletion parameterized by the size of a given modulator X to graphs with $\mathsf{ed}_{\mathcal{F}} \leq \eta$, where k is the size of the sought solution. The proof is by induction on η.

($\eta = 0$) If $\mathsf{ed}_{\mathcal{F}}(G \setminus X) = 0$, set X is a modulator to an $\mathcal{F}$-minor-free graph. Thus, (G, X, k) is an instance of $\mathcal{F}$-Minor Deletion parameterized by the size of a given solution. We apply the polynomial (α-approximate) kernelization from the hypothesis to the input to obtain an instance (G', X', k') of $\mathcal{F}$-Minor Deletion parameterized by the size of a given solution, which serves as our kernel.

($\eta \geq 1$) We apply the Reduce Components Lemma (Lemma 4) on the input to obtain G' and Δ. We will augment the modulator X into a superset X' to ensure that $\mathsf{ed}_{\mathcal{F}}(G' \setminus X') < \eta$. To this end, we consider each connected component C of $G' \setminus X$. Since the Elimination Distance to $\mathcal{F}$-Minor-Free problem parameterized by the target width is fixed-parameter tractable [11], and η is a constant, we can decide if $\mathsf{ed}_{\mathcal{F}}(C)$ is smaller than η in $f(\eta) \cdot n^{O(1)}$ time. If it is, we do not need to add any vertex from C to X'. Otherwise, by the definition of the elimination distance there is a vertex x_C such that $\mathsf{ed}_{\mathcal{F}}(C \setminus \{x_C\}) < \mathsf{ed}_{\mathcal{F}}(C)$. We find such a vertex x_C by trying all options for x_C and computing the elimination distance to an $\mathcal{F}$-minor-free graph of the resulting graph, again in $f(\eta) \cdot n^{O(1)}$ time. We initialize X' as X. For each component C of $G' \setminus X$ with $\mathsf{ed}_{\mathcal{F}}(C) = \eta$, we add the corresponding vertex x_C to X'.

Since the Reduce Components Lemma (Lemma 4) guarantees that the number of connected components of $G' \setminus X$ is polynomial in $|X|$ for fixed $\mathcal{F}$ and η, the resulting modulator X' has size polynomial in $|X|$. Moreover, it guarantees that $\mathsf{ed}_{\mathcal{F}}(G' \setminus X') < \eta$. Hence we now have an instance $(G', X', k - \Delta)$ of $\mathcal{F}$-Minor Deletion parameterized by a modulator to $\mathsf{ed}_{\mathcal{F}} \leq \eta - 1$, with the same answer as (G, X, k). Using the inductive hypothesis, we apply the (α-approximate) kernel for the parameterization by a modulator to $\mathsf{ed}_{\mathcal{F}} \leq \eta - 1$, which outputs an instance $(G^\star, X^\star, k^\star)$. By induction, the size of $G^\star$ is bounded by some polynomial in $|X'|$, which is in turn bounded by a polynomial in $|X|$. Hence $G^\star$ has size $|X|^{O_{\mathcal{F},\eta}(1)}$ which (since $\mathcal{F}$ and η are constants) is bounded by $O(|X|^c)$ for some suitably chosen constant c, and we output $(G^\star, X^\star, k^\star)$ as the result of the kernelization.

If the kernelization algorithm from the hypothesis is exact, then $(G^\star, X^\star, k^\star)$ has the same answer as $(G', X', k - \Delta)$ and therefore as (G, X, k), which gives an exact kernelization.

Otherwise, following the definition of approximate kernelization due to Lokshtanov et al. in the case of a structural parameterization [43, Section 2.1], and since we parameterize by the size of a solution that we assume to be given in the input (and not by the size of sought solution), there exists a solution lifting algorithm $\mathcal{A}$ that transforms a solution $Y^\star$ for $(G^\star, X^\star, k^\star)$ into a solution Y' for $(G', X', k - \Delta)$ such that

$$\frac{|Y'|}{\text{OPT}_{\mathcal{F}}(G')} \leq \alpha \cdot \frac{|Y^\star|}{\text{OPT}_{\mathcal{F}}(G^\star)}.$$

Here, we follow the guidelines for problems with structural parameterizations given in [43], defining the value of a valid solution Y for an instance $(G^\star, X^\star, k^\star)$, where $X^\star$ is a valid modulator to graphs with bounded $\mathsf{ed}_{\mathcal{F}}$, as $|Y|$.

We construct our solution lifting algorithm by combining $\mathcal{A}$ with the algorithm described in the Reduce Components Lemma (Lemma 4), which outputs a solution Y of size $|Y'| + \Delta$ for G in polynomial time. As $\text{OPT}_{\mathcal{F}}(G') + \Delta = \text{OPT}_{\mathcal{F}}(G)$, we get that

$$\frac{|Y|}{\text{OPT}_{\mathcal{F}}(G)} = \frac{|Y'| + \Delta}{\text{OPT}_{\mathcal{F}}(G') + \Delta} \leq \frac{|Y'|}{\text{OPT}_{\mathcal{F}}(G')} \leq \alpha \cdot \frac{|Y^\star|}{\text{OPT}_{\mathcal{F}}(G^\star)}.$$

It follows that $(G^\star, X^\star, k^\star)$ is an α-approximate polynomial kernel for (G, X, k) as well. ◀

Lokshtanov et al. [43] additionally introduce the notion of *strict* approximate kernelization and α*-safe rules* to facilitate the usage of *reduction rules* commonly used in the literature. We could have used these notions for our proof as well, but as our kernelization does not use reduction rules that are applied exhaustively, we opted for the more direct approach above.

4.2 Consequences of Theorem 1

As said earlier, Fomin et al. [24] showed that $\mathcal{F}$-Minor Deletion parameterized by the size of a solution admits a randomized polynomial kernel whenever $\mathcal{F}$ contains at least one planar graph. The reliance on randomization of their algorithm lies in the use of a randomized constant-factor approximation algorithm for $\mathcal{F}$-Minor Deletion [24, Theorem 1]. Gupta et al. [29, Corollary 1.1] gave later a deterministic constant-factor approximation algorithm for the problem, which gives us the following corollary of Theorem 1.

▶ **Corollary 12.** *For every fixed finite set $\mathcal{F}$ of connected graphs containing at least one planar graph and every constant η, $\mathcal{F}$-Minor Deletion parameterized by the size of a modulator to graphs with $\mathsf{ed}_{\mathcal{F}} \leq \eta$ admits a polynomial kernel.*

Notice that we do not require that the modulator is given as part of the input in Corollary 12. This is because we can actually compute such a modulator (of size slightly larger) in polynomial time using the constant-factor approximation algorithm by Gupta et al. [29, Corollary 1.1] mentioned earlier. Indeed, when $\mathcal{F}$ contains at least one planar graph, the class of graphs with $\mathsf{ed}_{\mathcal{F}} \leq \eta$ is characterized by a finite set of forbidden minors that also contains a planar graph. This can be seen, for instance, by observing that if $\mathsf{ed}_{\mathcal{F}}(G) \leq \eta$, then $\mathsf{tw}(G) \leq t_\ell + \eta$, where t_ℓ is the maximum treewidth of a graph induced by the bag of a leaf in an $\mathcal{F}$-elimination forest of G. As the minor obstructions of graphs of bounded treewidth contain at least one planar graph [48], the claim follows. (We point out that a similar argument is used in [5, Section 3] to compute a modulator to bounded bridge-depth.)

On the other hand, Dekker and Jansen [16, Theorem 2] showed that, assuming $\mathsf{NP} \not\subseteq \mathsf{coNP/poly}$, $\mathcal{F}$-Minor Deletion does not admit a polynomial kernel when parameterized by the size of a modulator to graphs with unbounded $\mathsf{ed}_{\mathcal{F}}$ if $\mathcal{F}$ is a finite set of biconnected planar graphs on at least three vertices. The requirement that all graphs in $\mathcal{F}$ are planar exclusively comes from their Lemma 10. They define a structure called an $\mathcal{F}$*-necklace*, and state that every graph G that does not contain large $\mathcal{F}$-necklaces as minors has bounded treewidth. They prove this by assuming that G has large treewidth, and then using the Excluded Grid Theorem [12, 48] to show that G contains a large grid minor, which in turn contains every large enough planar graph as a minor. As $\mathcal{F}$-necklaces are planar if all graphs in $\mathcal{F}$ are planar, this shows that G contains a large $\mathcal{F}$-necklace as a minor, which is a contradiction. This argument still holds if not every graph in $\mathcal{F}$ is planar, as having just one planar graph guarantees that there exist planar $\mathcal{F}$-necklaces of arbitrarily large size [33]. Thus, G still has an $\mathcal{F}$-necklace as a minor in this case. Hence, their proof can be adapted to show the following theorem, which together with Corollary 12 yield Theorem 2.

▶ **Theorem 13** (cf. [16, Theorem 2]). *Let $\mathcal{C}$ be a minor-closed family of graphs and let $\mathcal{F}$ be a finite set of biconnected graphs on at least three vertices containing at least one planar graph. If $\mathcal{C}$ has unbounded elimination distance to an $\mathcal{F}$-minor-free graph, then $\mathcal{F}$-Minor Deletion does not admit a polynomial kernel in the size of a $\mathcal{C}$-modulator, unless $\mathsf{NP} \subseteq \mathsf{coNP/poly}$.*

Finally, combining Theorem 1 with the polynomial α-approximate kernel for Planar Vertex Deletion of Jansen and Włodarczyk [38] gives us Theorem 3 as stated in the introduction.

References

1 Akanksha Agrawal, Lawqueen Kanesh, Daniel Lokshtanov, Fahad Panolan, M. S. Ramanujan, Saket Saurabh, and Meirav Zehavi. Deleting, Eliminating and Decomposing to Hereditary Classes Are All FPT-Equivalent. In *Proc. of the 32st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1976–2004, 2022. doi:10.1137/1.9781611977073.79.

2 Yuuki Aoike, Tatsuya Gima, Tesshu Hanaka, Masashi Kiyomi, Yasuaki Kobayashi, Yusuke Kobayashi, Kazuhiro Kurita, and Yota Otachi. An improved deterministic parameterized algorithm for cactus vertex deletion. *Theory of Computing Systems*, 66(2):502–515, 2022. doi:10.1007/S00224-022-10076-X.

3 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, July 2023. doi:10.1137/21m140482x.

4 Hans L. Bodlaender. Kernelization: New upper and lower bound techniques. In *Proc. 4th IWPEC*, pages 17–37, 2009. doi:10.1007/978-3-642-11269-0_2.

5 Marin Bougeret, Bart M. P. Jansen, and Ignasi Sau. Bridge-depth characterizes which minor-closed structural parameterizations of vertex cover admit a polynomial kernel. *SIAM Journal on Discrete Mathematics*, 36(4):2737–2773, 2022. doi:10.1137/21m1400766.

6 Marin Bougeret, Bart M. P. Jansen, and Ignasi Sau. Kernelization dichotomies for hitting subgraphs under structural parameterizations. In *Proc. of the 51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*, volume 297 of *LIPIcs*, pages 33:1–33:20, 2024. doi:10.4230/LIPIcs.ICALP.2024.33.

7 Marin Bougeret and Ignasi Sau. How much does a treedepth modulator help to obtain polynomial kernels beyond sparse graphs? In Daniel Lokshtanov and Naomi Nishimura, editors, *Proc. 12th IPEC*, volume 89 of *LIPIcs*, pages 10:1–10:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.IPEC.2017.10.

8 Marin Bougeret and Ignasi Sau. How much does a treedepth modulator help to obtain polynomial kernels beyond sparse graphs? *Algorithmica*, 81(10):4043–4068, 2019. doi:10.1007/s00453-018-0468-8.

9 Jannis Bulian and Anuj Dawar. Graph isomorphism parameterized by elimination distance to bounded degree. *Algorithmica*, 75(2):363–382, 2016. doi:10.1007/s00453-015-0045-3.

10 Jannis Bulian and Anuj Dawar. Fixed-parameter tractable distances to sparse graph classes. *Algorithmica*, 79(1):139–158, 2017. doi:10.1007/s00453-016-0235-7.

11 Jannis Bulian and Anuj Dawar. Fixed-parameter tractable distances to sparse graph classes. *Algorithmica*, 79(1):139–158, 2017. doi:10.1007/S00453-016-0235-7.

12 Julia Chuzhoy and Zihan Tan. Towards tight(er) bounds for the excluded grid theorem. *Journal of Combinatorial Theory, Series B*, 146:219–265, January 2021. doi:10.1016/j.jctb.2020.09.010.

13 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.

14 Marek Cygan, Daniel Lokshtanov, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. On the hardness of losing width. *Theory of Computing Systems*, 54(1):73–82, 2014. doi:10.1007/s00224-013-9480-1.

15 David Dekker and Bart M. P. Jansen. Kernelization for Feedback Vertex Set via Elimination Distance to a Forest. In *Proc. of the 48th International Workshop Graph-Theoretic Concepts in Computer Science (WG)*, volume 13453 of *LNCS*, pages 158–172, 2022. doi:10.1007/978-3-031-15914-5_12.

16 David J. C. Dekker and Bart M. P. Jansen. Kernelization for feedback vertex set via elimination distance to a forest. *Discrete Applied Mathematics*, 346:192–214, 2024. doi:10.1016/j.dam.2023.12.016.

17 Reinhard Diestel. *Graph Theory*. Springer-Verlag, Heidelberg, 5th edition, 2016.

18 Huib Donkers and Bart M. P. Jansen. A Turing kernelization dichotomy for structural parameterizations of $\mathcal{F}$-Minor-Free Deletion. *Journal of Computer and System Sciences*, 119:164–182, 2021. doi:10.1016/J.JCSS.2021.02.005.

19 Huib Donkers, Bart M. P. Jansen, and Michal Wlodarczyk. Preprocessing for outerplanar vertex deletion: An elementary kernel of quartic size. *Algorithmica*, 84(11):3407–3458, 2022. doi:10.1007/S00453-022-00984-2.

20 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.

21 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.

22 Samuel Fiorini, Gwenaël Joret, and Ugo Pietropaoli. Hitting diamonds and growing cacti. In *Proc. of the 14th International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, volume 6080 of *LNCS*, pages 191–204, 2010. doi:10.1007/978-3-642-13036-6_15.

23 Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Springer-Verlag, 2006. doi:10.1007/3-540-29953-X.

24 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)*, 2012. doi:10.1109/FOCS.2012.62.

25 Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. *Kernelization: Theory of Parameterized Preprocessing*. Cambridge University Press, 2019. doi:10.1017/9781107415157.

26 Fedor V. Fomin and Torstein J. F. Strømme. Vertex cover structural parameterization revisited. In *Proc. of the 42nd International Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, volume 9941 of *LNCS*, pages 171–182, 2016. doi:10.1007/978-3-662-53536-3_15.

27 Archontia C. Giannopoulou, Bart M. P. Jansen, Daniel Lokshtanov, and Saket Saurabh. Uniform kernelization complexity of hitting forbidden minors. *ACM Transactions on Algorithms*, 13(3):35:1–35:35, 2017. doi:10.1145/3029051.

28 Alexander Göke, Dániel Marx, and Matthias Mnich. Hitting Long Directed Cycles Is Fixed-Parameter Tractable. In *Proc. of the 47th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 168 of *LIPIcs*, pages 59:1–59:18, 2020. doi:10.4230/LIPIcs.ICALP.2020.59.

29 Anupam Gupta, Euiwoong Lee, Jason Li, Pasin Manurangsi, and Michał Włodarczyk. *Losing Treewidth by Separating Subsets*, pages 1731–1749. SIAM, 2019. doi:10.1137/1.9781611975482.104.

30 Eva-Maria C. Hols, Stefan Disc.tsch, and Astrid Pieterse. Elimination distances, blocking sets, and kernels for vertex cover. *SIAM Journal on Discrete Mathematics*, 36(3):1955–1990, 2022. doi:10.1137/20m1335285.

31 Eva-Maria C. Hols and Stefan Kratsch. Smaller parameters for vertex cover kernelization. In *Proc. of the 12th International Symposium on Parameterized and Exact Computation (IPEC)*, volume 89 of *LIPIcs*, pages 20:1–20:12, 2017. doi:10.4230/LIPIcs.IPEC.2017.20.

32 Eva-Maria C. Hols, Stefan Kratsch, and Astrid Pieterse. Elimination distances, blocking sets, and kernels for vertex cover. In *Proc. 37nd STACS*, volume 154 of *LIPIcs*, pages 36:1–36:14, 2020. doi:10.4230/LIPIcs.STACS.2020.36.

33 Bart M. P. Jansen, 2025. Personal communication.

34 Bart M. P. Jansen. Quo vadis, kernelization? Invited survey, to appear in *Computer Science Review*, special issue dedicated to Michael R. Fellows on the occasion of his 70th birthday, 2025. URL: https://www.sciencedirect.com/special-issue/10540QH6GM2.

35 Bart M. P. Jansen and Hans L. Bodlaender. Vertex cover kernelization revisited - upper and lower bounds for a refined parameter. *Theory of Computing Systems*, 53(2):263–299, 2013. doi:10.1007/s00224-012-9393-4.

36 Bart M. P. Jansen, Jari J. H. de Kroon, and Michal Włodarczyk. Vertex deletion parameterized by elimination distance and even less. In *Proc. of the 53rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 1757–1769, 2021. `doi:10.1145/3406325.3451068`.

37 Bart M. P. Jansen and Astrid Pieterse. Polynomial kernels for hitting forbidden minors under structural parameterizations. *Theoretical Computer Science*, 841:124–166, 2020. `doi:10.1016/j.tcs.2020.07.009`.

38 Bart M. P. Jansen and Michal Włlodarczyk. Lossy Planarization: A Constant-Factor Approximate Kernelization for Planar Vertex Deletion. *SIAM Journal on Computing*, 54(1):1–91, 2025. `doi:10.1137/22M152058X`.

39 Gwenaël Joret, Christophe Paul, Ignasi Sau, Saket Saurabh, and Stéphan Thomassé. Hitting and harvesting pumpkins. *SIAM Journal on Discrete Mathematics*, 28(3):1363–1390, 2014. `doi:10.1137/120883736`.

40 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`.

41 Stefan Kratsch. A randomized polynomial kernelization for vertex cover with a smaller parameter. *SIAM Journal on Discrete Mathematics*, 32(3):1806–1839, 2018. `doi:10.1137/16M1104585`.

42 Stefan Kratsch and Magnus Wahlström. Representative sets and irrelevant vertices: New tools for kernelization. In *Proc. 53rd FOCS*, pages 450–459, 2012. `doi:10.1109/FOCS.2012.46`.

43 Daniel Lokshtanov, Fahad Panolan, M. S. Ramanujan, and Saket Saurabh. Lossy kernelization. In *Proce. of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 224–237. ACM, 2017. `doi:10.1145/3055399.3055456`.

44 Diptapriyo Majumdar, Venkatesh Raman, and Saket Saurabh. Polynomial kernels for vertex cover parameterized by small degree modulators. *Theory of Computing Systems*, 62(8):1910–1951, 2018. `doi:10.1007/s00224-018-9858-1`.

45 Bojan Mohar and Carsten Thomassen. *Graphs on Surfaces*. Johns Hopkins University Press, 2001.

46 Rolf Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006. `doi:10.1093/acprof:oso/9780198566076.001.0001`.

47 Geevarghese Philip, Ashutosh Rai, and Saket Saurabh. Generalized pseudoforest deletion: Algorithms and uniform kernel. *SIAM Journal on Discrete Mathematics*, 32(2):882–901, 2018. `doi:10.1137/16M1100794`.

48 Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *J. Comb. Theory, Ser. B*, 41(1):92–114, 1986. `doi:10.1016/0095-8956(86)90030-4`.

49 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`.

50 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`.

51 Saket Saurabh. Open problems from the workshop on kernelization (worker 2019, part 1). 2019. URL: `https://youtu.be/vCjG5zGjQr4?t=124`.

52 Jeroen L.G. Schols. Kernelization for Treewidth-2 Vertex Deletion. *CoRR*, abs/2203.10070, 2022. `doi:10.48550/arXiv.2203.10070`.

53 Dekel Tsur. Faster deterministic algorithm for cactus vertex deletion. *Information Processing Letters*, 179:106317, 2023. `doi:10.1016/J.IPL.2022.106317`.