

Homomorphisms Are a Good Basis for Counting Small Subgraphs*

Radu Curticapean[†]
Institute for Computer Science and
Control, Hungarian Academy of
Sciences (MTA SZTAKI)
Budapest, Hungary
radu.curticapean@gmail.com

Holger Dell
Saarland University and
Cluster of Excellence (MMCI)
Saarbrücken, Germany
hdell@mmci.uni-saarland.de

Dániel Marx[‡]
Institute for Computer Science and
Control, Hungarian Academy of
Sciences (MTA SZTAKI)
Budapest, Hungary
dmarx@cs.bme.hu

ABSTRACT

We introduce *graph motif parameters*, a class of graph parameters that depend only on the frequencies of constant-size induced subgraphs. Classical works by Lovász show that many interesting quantities have this form, including, for fixed graphs H , the number of H -copies (induced or not) in an input graph G , and the number of homomorphisms from H to G .

We use the framework of graph motif parameters to obtain faster algorithms for counting subgraph copies of fixed graphs H in host graphs G . More precisely, for graphs H on k edges, we show how to count subgraph copies of H in time $k^{O(k)} \cdot n^{0.174k+o(k)}$ by a surprisingly simple algorithm. This improves upon previously known running times, such as $O(n^{0.91k+c})$ time for k -edge matchings or $O(n^{0.46k+c})$ time for k -cycles.

Furthermore, we prove a general complexity dichotomy for evaluating graph motif parameters: Given a class \mathcal{C} of such parameters, we consider the problem of evaluating $f \in \mathcal{C}$ on input graphs G , parameterized by the number of induced subgraphs that f depends upon. For every recursively enumerable class \mathcal{C} , we prove the above problem to be either FPT or $\#W[1]$ -hard, with an explicit dichotomy criterion. This allows us to recover known dichotomies for counting subgraphs, induced subgraphs, and homomorphisms in a uniform and simplified way, together with improved lower bounds.

Finally, we extend graph motif parameters to colored subgraphs and prove a complexity trichotomy: For vertex-colored graphs H and G , where H is from a fixed class \mathcal{H} , we want to count color-preserving H -copies in G . We show that this problem is either polynomial-time solvable or FPT or $\#W[1]$ -hard, and that the FPT cases indeed need FPT time under reasonable assumptions.

*Part of this work was done while the authors were visiting the Simons Institute for the Theory of Computing, and at the Dagstuhl Seminar 17041 – “Randomization in Parameterized Complexity”.

[†]Dr. Curticapean was supported by ERC grant PARAMTIGHT (No. 280152).

[‡]Dr. Marx was supported by ERC grant PARAMTIGHT (No. 280152).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC’17, Montreal, Canada

© 2017 ACM. 978-1-4503-4528-6/17/06...\$15.00

DOI: 10.1145/3055399.3055502

CCS CONCEPTS

• **Mathematics of computing** → **Graph algorithms**; • **Theory of computation** → **Problems, reductions and completeness**; **Fixed parameter tractability**;

KEYWORDS

counting subgraphs, homomorphisms, fixed-parameter tractability, exponential time hypothesis

ACM Reference format:

Radu Curticapean, Holger Dell, and Dániel Marx. 2017. Homomorphisms Are a Good Basis for Counting Small Subgraphs. In *Proceedings of 49th Annual ACM SIGACT Symposium on the Theory of Computing, Montreal, Canada, June 2017 (STOC’17)*, 14 pages. DOI: 10.1145/3055399.3055502

1 INTRODUCTION

Deciding the existence of subgraph patterns H in input graphs G constitutes the classical *subgraph isomorphism problem* [13, 49], which generalizes NP-complete problems like the Hamiltonian cycle problem or the clique problem. In some applications however, it is not sufficient to merely know whether H occurs in G , but instead one wishes to determine the *number* of such occurrences. This is clearly at least as hard as deciding their existence, but it can be much harder: The existence of perfect matchings can be tested in polynomial time, but the counting version is $\#P$ -hard [50].

Subgraph counting problems have applications in areas like statistical physics, probabilistic inference, and network analysis. In particular, in network analysis, such problems arise in the context of discovering *network motifs*. These are small patterns that occur more often in a network than would be expected if the network was random. Through network motifs, the problem of counting subgraphs has found applications in the study of gene transcription networks, neural networks, and social networks [42], and there is a large body of work dedicated to the algorithmic discovery of network motifs [2, 10, 24, 32, 33, 45–47, 51].

Inspired by these applications, we study the algorithmic problem of counting occurrences of *small patterns* H in large host graphs G . The abstract notion of a “pattern occurrence” may be formalized in various ways, which may result in vastly different problems: To state only some examples, we may be interested in counting subgraph copies of a graph H , or induced subgraph copies of H , or homomorphisms from H to G , and we can also consider settings

where both pattern H and host graph G are colored and we wish to count subgraphs of G that are color-preserving isomorphic to H .

It may seem daunting at first to try to deal with all different types of pattern occurrences. Fortunately, Lovász [38, 39] defined a framework that allows us to express virtually all kinds of pattern types in a unified way. As it turns out, graph parameters such as the number of subgraph copies of H (induced or not) in a host graph G , or the number of graph homomorphisms from H to G are actually just “linear combinations” of each other in a well-defined sense. We build on this and define a general framework of so-called *graph motif parameters* to capture counting linear combinations of small patterns, into which (induced) subgraph or homomorphism numbers embed naturally as special cases.

In the remainder of the introduction, we first discuss algorithmic and complexity-theoretic aspects of counting (induced) subgraphs and homomorphisms in §1.1–§1.3 and state the results we derive for these special cases. In §1.4, we then give an introduction into the general framework of graph motif parameters, our interpretation of Lovász’s unified framework, which also provides the main techniques for our proofs. Finally, in §1.5 we give an exposition of our results for vertex-colored subgraphs.

1.1 Counting Small Subgraphs

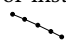
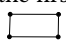
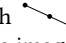
For any fixed k -vertex pattern graph H , we can count all subgraph copies of H in an n -vertex host graph G using brute-force for a running time of $O(n^k)$. While this running time is polynomial for any fixed H , it quickly becomes infeasible as k grows. Fortunately enough, non-trivial improvements on the exponent are known, albeit only for specific classes of patterns:

For instance, we can count triangles in the same time $O(n^\omega)$ that it takes to multiply two $(n \times n)$ -matrices [27]. It is known that $\omega < 2.373$ holds [23, 52]. This approach can be generalized from triangles to k -cliques with $k \in \mathbb{N}$ [43], for a running time of $n^{\omega k/3+O(1)}$. Fast matrix multiplication is also used to improve on exhaustive search for counting cycles of length at most seven [3] and various other problems [21, 34].

Secondly, for k -edge paths or generally any pattern of bounded pathwidth, a “meet in the middle” approach yields $n^{k/2+O(1)}$ time algorithms [4, 35]. For a while, this approach appeared to be a barrier for faster algorithms, until Björklund et al. [5] gave an algorithm for counting k -paths, matchings on k vertices, and other k -vertex patterns of bounded pathwidth in time $n^{0.455k+O(1)}$.

Finally, if $\text{vc}(H)$ is the *vertex-cover number* of H , that is, the size of its smallest vertex-cover, then we can count H -copies in time $n^{\text{vc}(H)+O(1)}$ [54] (also cf. [15, 36]). Essentially, one can exhaustively iterate over the image of the minimum vertex-cover in G , which gives rise to $n^{\text{vc}(G)}$ choices; the rest of H can then be embedded by dynamic programming. Note that $\text{vc}(H)$ may be constant even for large graphs H , e.g., if H is a star.

In this paper, we unify some of the algorithms above and generalize them to arbitrary subgraph patterns; in many cases our algorithms are faster. For two graphs H and G , let $\#\text{Sub}(H \rightarrow G)$ be the number of subgraphs of G that are isomorphic to H . Our main algorithmic result states that $\#\text{Sub}(H \rightarrow G)$ can be determined in time $O(n^{t+1})$, where t is the maximum treewidth (a very popular measure of tree-likeness) among the homomorphic images of H .

For our purposes, a homomorphic image of H is any simple graph that can be obtained from H by possibly merging non-adjacent vertices. For instance, identifying the first and the last vertex in the 4-path  yields the 4-cycle , and further identifying two non-adjacent vertices in the 4-cycle yields the 2-path . We define the *spasm* of H as the set of all homomorphic images of H , that is, as the set of “all possible non-edge contractions” of H . As an example, for the 4-path, we have

$$\text{Spasm}\left(\text{---}\right) = \left\{ \begin{array}{l} \text{---}, \text{---}, \text{---}, \square, \\ \text{---}, \text{---}, \text{---}, \text{---} \end{array} \right\}. \tag{1}$$

Our main algorithmic result can then be stated as follows:

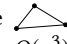
THEOREM 1.1. *Given a k -edge graph H and an n -vertex graph G , we can compute $\#\text{Sub}(H \rightarrow G)$ in time $k^{O(k)} \cdot n^{t+1}$, where t is the maximum treewidth in the spasm of H .*

As an example, for the 4-path, the largest treewidth among the graphs in the spasm is 2, and so Theorem 1.1 yields a running time of $O(n^3)$ for counting 4-paths. In fact, even the 6-path has only graphs with treewidth at most 2 in its spasm, so the same cubic running time applies.

Theorem 1.1 generalizes the vertex-cover based algorithm [54] mentioned before: Merging vertices can never increase the size of the smallest vertex-cover, and so the maximum treewidth in the spasm of H is bounded by $\text{vc}(H)$, since the vertex-cover number is an upper bound for the treewidth of a graph. Note that, while contracting edges cannot increase the treewidth of a graph, it is apparent from (1) that contracting non-edges might. In fact, if H is the k -edge matching, *all* k -edge graphs can be obtained by contracting non-edges, including expander graphs with treewidth $\Omega(k)$. However, Scott and Sorkin [48, Corollary 21] proved that every graph with at most k edges has treewidth at most $0.174 \cdot k + o(k)$. This bound enables the following immediate corollary to Theorem 1.1.

COROLLARY 1.2. *Given a k -edge graph H and an n -vertex graph G , we can compute $\#\text{Sub}(H \rightarrow G)$ in time $k^{O(k)} \cdot n^{0.174 \cdot k + o(k)}$.*

Our exponent is obviously smaller than the previously known $0.455 \cdot |V(H)| = 0.91 \cdot k$ for the k -edge matching and $0.455 \cdot k$ for the k -path, but somewhat surprisingly, our algorithm is also significantly simpler than its predecessors.

When H is the triangle , the algorithm from Theorem 1.1 matches the running time $O(n^3)$ of the exhaustive search method. To achieve a smaller constant in the exponent, we have to use matrix multiplication, since faster triangle detection is equivalent to faster Boolean matrix multiplication [53]. Indeed, we are able to generalize the $O(n^\omega)$ -time algorithm for counting triangles to arbitrary graphs H whose spasm has treewidth at most 2.

THEOREM 1.3. *If all graphs in the spasm of H have treewidth at most two, then we can compute $\#\text{Sub}(H \rightarrow G)$ in time $f(H) \cdot |V(G)|^\omega$, where $f(H)$ is a function that only depends on H .*

This algorithm applies, for example, to paths of length at most 6, thus providing an alternative and simplified way to obtain the corresponding results of [3].

We now turn to hardness results for counting subgraphs. Here, the vertex-cover number of the pattern H plays a special role:

When it is bounded by a fixed constant $b \in \mathbb{N}$, we have an $n^{b+O(1)}$ time algorithm even when the size of the pattern is otherwise unbounded. However when it is unbounded, e.g., for k -paths, the best known running times are of the form $n^{\epsilon k}$ for some $\epsilon \in (0, 1)$. Given these modest improvements for counting subgraph patterns of unbounded vertex-cover number, it is tempting to conjecture that “the exponent cannot remain constant” for such patterns. A result by a subset of the authors [15] shows that this conjecture is indeed true—for an appropriate formalization of the respective computational problem and under appropriate complexity-theoretic assumptions.

The *counting exponential time hypothesis* (#ETH) by Impagliazzo and Paturi [26], adapted to the counting setting [18], states that there is no $\exp(o(n)) \cdot \text{poly}(m)$ -time algorithm to count all satisfying assignments of a given 3-CNF formula with n variables and m clauses. More convenient for us is the following consequence of #ETH: There is no $f(k) \cdot n^{o(k)}$ time algorithm to count all cliques of size exactly k in a given n -vertex graph [9] — thus, the k -clique problem is a hard special case of the subgraph counting problem and clearly k -cliques have large vertex-cover number. Of course, this worst-case hardness of the most general subgraph counting problem does not directly help us understand the complexity of particular cases, such as counting k -matchings or k -paths.

We can model our interest in special cases by restricting the pattern graphs H to be from a fixed class \mathcal{H} of graphs: The computational problem $\#\text{Sub}(\mathcal{H})$ is to compute the number $\#\text{Sub}(H \rightarrow G)$ of H -copies in G when given two graphs $H \in \mathcal{H}$ and G as input. To prove that #ETH implies that no fixed-parameter tractable algorithm can exist for $\#\text{Sub}(\mathcal{H})$, one ultimately establishes a *parameterized reduction* from the k -clique problem to the H -subgraph counting problem, which has the important property that $\text{vc}(H)$ is bounded by $g(k)$, a function only depending on k .

The parameterized reduction in [15] was very complex with various special cases and a Ramsey argument that made g a very large function. While it was sufficient to conditionally rule out $f(k) \cdot n^c$ time algorithms for $\#\text{Sub}(\mathcal{H})$ for any constant c and graph class \mathcal{H} of unbounded vertex-cover number, it left open the possibility of, for example, $n^{\sqrt{\text{vc}(H)}}$ time algorithms. Running times of the form $n^{o(k/\log k)}$ could be ruled out under #ETH only for certain special cases, such as counting k -matchings or k -paths. In this paper, we obtain the stronger hardness result for all hard families \mathcal{H} . For technical reasons, we assume that \mathcal{H} can be recursively enumerated; without this assumption, we would however obtain a similar result under a non-uniform version of #ETH.

THEOREM 1.4. *Let \mathcal{H} be a recursively enumerable class of graphs of unbounded vertex-cover number. If #ETH holds, then $\#\text{Sub}(\mathcal{H})$ cannot be computed in time $f(H) \cdot n^{o(\text{vc}(H)/\log \text{vc}(H))}$.*

The log-factor here is related to an open problem in parameterized complexity, namely whether you can “beat treewidth” [40], i.e., whether there is an algorithm to find H as a subgraph in time $f(H) \cdot n^{o(\text{tw}(H))}$, or whether such an algorithm is ruled out by ETH for every graph class \mathcal{H} of unbounded treewidth. Indeed, replacing $\text{vc}(H)$ with $\text{tw}(H)$ in Theorem 1.4 essentially yields the hardness result in [40, Corollary 6.3], and since $\text{tw}(H) \leq \text{vc}(H)$, our theorem can be seen as a strengthening of this result in the counting world. In fact, our hardness proof is based on this weaker version.

Instead of relying on #ETH, we can also consider $\#\text{Sub}(\mathcal{H})$ from the viewpoint of fixed-parameter tractability. In this framework, the problem is parameterized by $|V(H)|$. A problem is $\#\text{W}[1]$ -complete if it is equivalent under parameterized reductions to the problem of counting k -cliques, where the allowed reductions are Turing reductions that run in time $f(k) \cdot \text{poly}(n)$. As mentioned before, it is known that #ETH implies $\text{FPT} \neq \#\text{W}[1]$. In this setting, Theorem 1.4 takes on the following form:

THEOREM 1.5 ([15]). *Let \mathcal{H} be a recursively enumerable class of graphs. If \mathcal{H} has bounded vertex-cover number, then $\#\text{Sub}(\mathcal{H})$ is polynomial-time computable. Otherwise, it is $\#\text{W}[1]$ -complete when parameterized by $|V(H)|$.*

The original proof of Theorem 1.5 relied on the $\#\text{W}[1]$ -completeness of counting k -matchings, which was nontrivial on its own [6, 14]. Then an extensive graph-theoretic analysis was used to find “ k -matching gadgets” in any graph class \mathcal{H} of unbounded vertex-cover number. These gadgets enable a parameterized reduction from counting k -matchings to $\#\text{Sub}(\mathcal{H})$, but they are also responsible for the uncontrollable blowup in the parameter that lead to highly non-tight results under #ETH. We obtain a much simpler proof of Theorem 1.5 that does not assign a special role to k -matchings.

Interestingly, Theorem 1.5 implies that there is no problem of the form $\#\text{Sub}(\mathcal{H})$ that is “truly” FPT—every such problem that is not $\#\text{W}[1]$ -complete is in fact already polynomial-time solvable. Thus, if we assume the widely-believed claim that $\text{FPT} \neq \#\text{W}[1]$ holds, then Theorem 1.5 exhaustively classifies the polynomial-time solvable problems $\#\text{Sub}(\mathcal{H})$. Indeed, such a sweeping dichotomy would not have been possible by merely assuming that $\text{P} \neq \text{P}^{\text{P}}$, since there exist artificial classes \mathcal{H} with $\#\text{P}$ -intermediate [12] $\#\text{Sub}(\mathcal{H})$.

We remark that a decision version of Theorem 1.5 is an open problem. It is known only for graph classes that are hereditary, that is, closed under induced subgraphs [28], where the dichotomy criterion is different. Moreover, certain non-hereditary cases, such as the $\text{W}[1]$ -completeness of deciding the existence of a bipartite clique or a grid, have been resolved only recently [11, 37].

1.2 Counting Small Homomorphisms

Similar classifications as Theorem 1.5 for counting subgraphs were previously known for counting homomorphisms [17] from a given class \mathcal{H} . Recall that a homomorphism from a pattern H to a host G is a mapping $f : V(H) \rightarrow V(G)$ such that $uv \in E(H)$ implies $f(u)f(v) \in E(G)$; we write $\#\text{Hom}(H \rightarrow G)$ for the number of such homomorphisms. In the context of pattern counting problems, we can interpret homomorphisms as a relaxation of the subgraph notion: If a homomorphism f from H to G is injective, then it constitutes a subgraph embedding from H to G . However, since we generally do *not* require injectivity in homomorphisms, these may well map into other *homomorphic images*. (Recall that the spasm of H contains exactly the loop-free homomorphic images of H .)

From an algorithmic viewpoint, not requiring injectivity makes counting patterns easier: One can now use separators to divide H into subpatterns and compose mappings of different subpatterns to a global one for H via dynamic programming. In this process, only the locations of separators under the subpattern mapping need to be memorized, while non-separator vertices of subpatterns may

be forgotten. As an example, note that counting k -paths is $\#W[1]$ -complete by Theorem 1.5, but counting homomorphisms from a k -path to a host graph G is polynomial-time solvable. Indeed, the latter problem amounts to counting k -walks in G , which can be achieved easily by taking the k -th power of the adjacency matrix of G , a process that can be interpreted as a dynamic programming algorithm: Given a table with the number of ℓ -walks from s to u for each u , we can compute a table with the number of $(\ell + 1)$ -walks from s to v for all v . That is, we only need to store the last vertex seen in a walk. This idea can be generalized easily to graphs of bounded treewidth using a straightforward dynamic programming approach on the tree decomposition of H .

PROPOSITION 1.6 (DÍAZ ET AL. [19]). *There is a deterministic $\exp(O(k)) + \text{poly}(k) \cdot n^{\text{tw}(H)+1}$ time algorithm to compute the number of homomorphisms from a given graph H to a given graph G , where $k = |V(H)|$, $n = |V(G)|$, and $\text{tw}(H)$ denotes the treewidth of H .*

This proposition is the basis of our main algorithmic result for subgraphs (Theorem 1.1) and other graph motif parameters. For graphs H of treewidth at most two, we can speed up the dynamic programming algorithm by using fast matrix multiplication:

THEOREM 1.7. *If H has treewidth at most 2, we can compute $\#\text{Hom}(H \rightarrow G)$ in time $\text{poly}(|V(H)|) \cdot |V(G)|^\omega$.*

Apart from being more well-behaved than subgraphs in terms of algorithmic tractability, homomorphisms also allow for simpler hardness proofs: Several constructions are significantly easier to analyze for homomorphisms than for subgraphs, as we will see in our proofs. This might explain why a dichotomy for counting homomorphisms from a fixed class \mathcal{H} was obtained an entire decade before its counterpart for subgraphs; it establishes the treewidth of \mathcal{H} as the tractability criterion for the problems $\#\text{Hom}(\mathcal{H})$.

THEOREM 1.8 (DALMAU AND JONSSON [17]). *Let \mathcal{H} be a recursively enumerable class of graphs. If \mathcal{H} has bounded treewidth, then the problem $\#\text{Hom}(\mathcal{H})$ of counting homomorphisms from graphs in \mathcal{H} into host graphs is polynomial-time solvable. Otherwise, it is $\#W[1]$ -complete when parameterized by $|V(H)|$.*

We also need the following hardness under $\#\text{ETH}$, the proof of which is a simple corollary of [40, Corollary 6.2 and 6.3].

PROPOSITION 1.9. *Let \mathcal{H} be a recursively enumerable class of graphs of unbounded treewidth. If $\#\text{ETH}$ holds, then there is no $f(H) \cdot |V(G)|^{o(\text{tw}(H)/\log \text{tw}(H))}$ time algorithm to compute $\#\text{Hom}(\mathcal{H})$ for given graphs $H \in \mathcal{H}$ and G .*

Finally, we remark that the decision version of Theorem 1.8, that is, the dichotomy theorem for deciding the existence of a homomorphism from $H \in \mathcal{H}$ to G is known and has a different criterion [25]: The decision problem is polynomial-time computable even when only the homomorphic cores of all graphs in \mathcal{H} have bounded treewidth, and it is $W[1]$ -hard otherwise.

1.3 Counting Small Induced Subgraphs

Let us also address the problem of counting small induced subgraphs from a class \mathcal{H} . This is a natural and well-studied variant of counting subgraph copies [12, 29–31, 34, 41], and for several applications it represents a more appropriate notion of “pattern

occurrence”. From the perspective of dichotomy results however, it is less intricate than subgraphs or homomorphisms: Counting induced subgraphs is known to be $\#W[1]$ -hard for any infinite pattern class \mathcal{H} , and even the corresponding decision version is $W[1]$ -hard.

THEOREM 1.10 ([12]). *Let \mathcal{H} be a recursively enumerable class of graphs. If \mathcal{H} is finite, then the problem $\#\text{Ind}(\mathcal{H})$ of counting induced subgraphs from \mathcal{H} is polynomial-time solvable. Otherwise, it is $\#W[1]$ -complete when parameterized by $|V(H)|$.*

Jerrum and Meeks [29–31, 41] introduced the following generalization of the problems $\#\text{Ind}(\mathcal{H})$ to fixed graph properties Φ : Given a graph G and $k \in \mathbb{N}$, the task is to compute the number of induced k -vertex subgraphs that have property Φ . Let us call this problem $\#\text{IndProp}(\Phi)$. They identified some classes of properties Φ that render this problem $\#W[1]$ -hard. Using our machinery, we get a full dichotomy theorem for this class of problems.

THEOREM 1.11 (SIMPLE VERSION). *If Φ is a decidable graph property, then $\#\text{IndProp}(\Phi)$ is fixed-parameter tractable or $\#W[1]$ -hard when parameterized by $|V(H)|$.*

1.4 A Unified View: Graph Motif Parameters

We now discuss our proof techniques on a high level. From a conceptual perspective, our most important contribution lies in finding a framework for understanding the parameterized complexity of subgraphs, induced subgraphs, and homomorphisms in a uniform context. Note that we quite literally *find* this framework: That is, we do not develop it ourselves, but we rather adapt works by Lovász et al. dating back to the 1960s [7, 38]. The most important observation is the following:

Many counting problems are actually linear combinations of homomorphisms in disguise!

That is, there are elementary transformations to express, say, linear combinations of subgraphs as linear combinations of homomorphisms, and vice versa.

The algorithms for subgraphs (Theorem 1.1 and Corollary 1.2) are based on a reduction from subgraph counting to homomorphism counting, so we want to find relations between the number of subgraphs and the number of homomorphisms. To get things started, note that injective homomorphisms from H to G , also called *embeddings*, correspond to a subgraph F of G that is isomorphic to H , and in fact, the number $\#\text{Emb}(H \rightarrow G)$ of embeddings is equal to the number $\#\text{Sub}(H \rightarrow G)$ of subgraphs times $\#\text{Aut}(H)$, the number of automorphisms of H .

Homomorphisms cannot map two adjacent vertices of H to the same vertex of G , assuming that G does not have any loops. For instance, every homomorphism from \triangle to G must be injective, and therefore the number of triangles in G is equal to the number of such homomorphisms, up to a factor of 6: the number of automorphisms of the triangle. Formally, we have $\#\text{Sub}(\triangle \rightarrow G) = \frac{1}{6} \cdot \#\text{Hom}(\triangle \rightarrow G)$ for every graph G that does not have loops.

Cases where homomorphisms from H to G are not automatically injective are more interesting. Clearly, the set of all homomorphisms contains the injective ones, which suggests we should simply count all homomorphisms and then subtract the ones that are not injective. Any non-injective homomorphism h from H to G has the property that there are at least two (non-adjacent) vertices

that it maps to the same vertex; in other words, its image $h(H)$ is isomorphic to some member of $\text{Spasm}(H)$ other than H itself. For example, $\#\text{Hom}(\text{---} \rightarrow G) - \#\text{Hom}(\text{---} \rightarrow G)$ is the number of injective homomorphisms from --- to G since the only way for such a homomorphism to be non-injective is that it merges the two degree-1 vertices.

In general, the number $\#\text{Emb}(H \rightarrow G)$ of injective homomorphisms is $\#\text{Hom}(H \rightarrow G) - \sum_{F \in \text{Spasm}(H) \setminus \{H\}} \#\text{Emb}(F \rightarrow G)$. Since each such F is strictly smaller than H , this fact yields a recursive procedure to compute $\#\text{Emb}(H \rightarrow G)$. However, there is a better way: We can use Möbius inversion over the partition lattice to obtain a closed formula. We already mentioned that the spasm of H can be obtained by consolidating non-adjacent vertices of H in all possible ways. This means that we consider partitions ρ of $V(H)$ in which each block is an independent set, and then form the *quotient graph* H/ρ obtained from H by merging each block of ρ into a single vertex. To express the injective homomorphisms from H to G (and hence the number of H -subgraphs) as a linear combination of homomorphisms, we consider all possible types in which a homomorphism h from H to G can fail to be injective. More precisely, we define this type ρ_h of h to be the partition of $V(H)$, where each block is the set of vertices of H that map to the same vertex of G under h . The homomorphism h is injective if and only if ρ_h is the finest partition, i.e., the partition where each block has size one.

The homomorphisms from H/ρ to G are precisely those homomorphisms from H to G that fail to be injective “at least as badly as ρ ”, that is, those homomorphisms f whose type ρ_f is a coarsening of ρ . As remarked by Lovász et al. [7, 38], one can then use Möbius inversion, a generalization of the inclusion–exclusion principle, to turn this observation into the “inverse” identity

$$\#\text{Sub}(H \rightarrow G) = \sum_{\rho} \mu_{\rho} \cdot \#\text{Hom}(H/\rho \rightarrow G), \text{ where} \quad (2)$$

$$\mu_{\rho} = \frac{(-1)^{|V(H)|-|V(H/\rho)|} \cdot \prod_{B \in \rho} (|B| - 1)!}{\#\text{Aut}(H)}. \quad (3)$$

The sum in (2) ranges over all partitions ρ of $V(H)$. Hence, the number of H -subgraphs in G is equal to a linear combination of the numbers of homomorphisms from graphs H/ρ to G . Each H/ρ is isomorphic to a graph in $\text{Spasm}(H)$, and so by collecting terms for isomorphic graphs, (2) represents the number of H -subgraphs as a linear combination of homomorphism numbers from graphs F in the spasm of H ; see Figure 1 for an example.

The identity (2) can be viewed as a basis transformation in a certain vector space of graph parameters, and we formalize this perspective in §3. A similar identity turns out to hold for counting induced subgraphs as well, so all three graph parameter types can be written as finite linear combinations of each other. This motivates the notion of a *graph motif parameter*, which is any graph parameter f that is a finite linear combination of induced subgraph numbers. That is, there are coefficients $\alpha_1, \dots, \alpha_t \in \mathbb{Q}$ and graphs H_1, \dots, H_t such that, for all graphs G , we have

$$f(G) = \sum_{i=1}^t \alpha_i \cdot \#\text{IndSub}(H_i \rightarrow G). \quad (4)$$

We study the problem of computing graph motif parameters f . For our results in parameterized complexity, we parameterize this

$$\begin{aligned} \text{Sub}(\text{---} \rightarrow \star) &= \\ & \frac{1}{2} \text{Hom}(\text{---} \rightarrow \star) \\ & - \text{Hom}(\text{---} \rightarrow \star) \\ & - \frac{1}{2} \text{Hom}(\text{---} \rightarrow \star) \\ & + \frac{1}{2} \text{Hom}(\text{---} \rightarrow \star) \\ & - \text{Hom}(\text{---} \rightarrow \star). \end{aligned} \quad \begin{aligned} & - \text{Hom}(\text{---} \rightarrow \star) \\ & - \frac{1}{2} \text{Hom}(\text{---} \rightarrow \star) \\ & + \frac{1}{2} \text{Hom}(\text{---} \rightarrow \star) \end{aligned}$$

Figure 1: An example for (2), where H is the path with four edges. The number of subgraphs is represented as a linear combination of homomorphisms from graphs $F \in \text{Spasm}(H)$. Each such F has treewidth at most two, so we can compute the homomorphism numbers in time $O(n^3)$ via Proposition 1.6. Computing the linear combination on the right side yields an $O(n^3)$ -time algorithm to count 4-paths, and in fact this is the algorithm in Theorem 1.1.

problem by the description length k of $\alpha_1, \dots, \alpha_t$ and H_1, \dots, H_t . Due to the basis transformation between induced subgraphs, subgraphs, and homomorphisms, writing $\#\text{Hom}(H_i \rightarrow G)$ instead of $\#\text{IndSub}(H_i \rightarrow G)$ in (4) yields an equivalent class of problems – switching bases only leads to a factor $g(k)$ overhead in the running time for some computable function g , which we can mostly neglect for our purposes.

Our main result is that the complexity of computing any graph parameter f is exactly governed by the maximum complexity of counting the homomorphisms occurring in its representation over the homomorphism basis. More precisely, let $\alpha_1, \dots, \alpha_t \in \mathbb{Q}$ and H_1, \dots, H_t be graphs such that $f(G) = \sum_i \alpha_i \cdot \#\text{Hom}(H_i \rightarrow G)$ holds for all graphs G . Our algorithmic results are based on the following observation: If each $\#\text{Hom}(H_i \rightarrow G)$ can be computed in time $O(n^c)$ for $n = |V(G)|$ and some constant $c \geq 0$, then the linear combination $f(G)$ can be computed in time $O(n^c)$ for the same constant c . However, we show that the reverse direction also holds: If f can be computed in time $O(n^c)$ for some $c \geq 0$, then each $\#\text{Hom}(H_i \rightarrow G)$ with $\alpha_i \neq 0$ can be computed in time $O(n^c)$ for the same constant c . The reduction that establishes this fine-grained equivalence gives rise to our results under #ETH and our new #W[1]-hardness proof.

Note that such an equivalence is *not* true for linear combinations of embedding numbers, as can be seen from the following example.

Example 1.12. Consider the following linear combination:

$$\begin{aligned} & \text{Emb}(\text{---} \rightarrow \star) + \text{Emb}(\text{---} \rightarrow \star) \\ & + \text{Emb}(\text{---} \rightarrow \star) + 2 \cdot \text{Emb}(\text{---} \rightarrow \star) \\ & + 2 \cdot \text{Emb}(\text{---} \rightarrow \star) + 3 \cdot \text{Emb}(\text{---} \rightarrow \star) \\ & + 4 \cdot \text{Emb}(\text{---} \rightarrow \star) + \text{Emb}(\text{---} \rightarrow \star). \end{aligned}$$

When this linear combination of embeddings is transformed into the homomorphism basis via (2), most terms cancel, and it turns out that it is equal to $\text{Hom}(\text{---} \rightarrow \star)$, that is, it counts the number of walks of length 4. Counting 4-walks can be done in time $O(n^2)$ via Proposition 1.6, but counting, for example, triangles is not known to be possible faster than $O(n^\omega)$. More generally, counting walks of

length k is in $O(n^2)$ -time, but counting *paths* of length k is $\#\text{W}[1]$ -hard and not in time $g(k) \cdot n^{o(k/\log k)}$ under $\#\text{ETH}$.

Thus, even a linear combination of subgraph numbers that looks complex at first and contains as summands embedding numbers that are fairly hard to compute can actually be quite a bit easier due to cancellation effects that occur when rewriting it as linear combination of homomorphism numbers.

As in the case of subgraphs in §1.1, we consider classes \mathcal{A} of linear combinations to get more expressive hardness results. That is, each element of \mathcal{A} is a pattern-coefficient list $(\alpha_1, H_1), \dots, (\alpha_t, H_t)$ as above. The graph motif problem $\#\text{Ind}(\mathcal{A})$ is then given a pattern-coefficient list from \mathcal{A} and a graph G , and is supposed to compute the linear combination (4). We have the following result for the complexity of computing graph motif parameters.

THEOREM 1.13 (INTUITIVE VERSION). *Let \mathcal{A} be a recursively enumerable class of pattern-coefficient lists. If the linear combinations (4) re-expressed as linear combinations of homomorphisms contain non-zero coefficients only for graphs of treewidth at most t , the problem $\#\text{Ind}(\mathcal{A})$ can be computed in time $f(\alpha) \cdot n^{t+1}$. Otherwise, the problem is $\#\text{W}[1]$ -hard parameterized by $|\alpha|$ and does not have $f(\alpha) \cdot n^{o(t/\log t)}$ time algorithms under $\#\text{ETH}$.*

With respect to fixed-parameter tractability vs. $\#\text{W}[1]$ -hardness, this theorem fully classifies the problems $\#\text{Ind}(\mathcal{A})$ for fixed classes of linear combinations \mathcal{A} . Of course we have similar (equivalent) formulations for $\#\text{Sub}(\mathcal{A})$ and $\#\text{Hom}(\mathcal{A})$, and thus we generalize the dichotomy theorems for subgraphs (Theorem 1.5), homomorphisms (Theorem 1.8), and induced subgraphs (Theorem 1.10). The dichotomy criterion is somewhat indirect; it addresses \mathcal{A} only through its representation as a linear combination of homomorphism numbers. However, we do not believe that there is a more 'native' dichotomy criterion on \mathcal{A} since seemingly complicated linear combinations can turn out to be easy – Example 1.12 gives an indication of this phenomenon; perturbing the coefficients just a tiny bit can turn a computationally easy linear combination into one that is hard.

Nevertheless, we can exhibit some interesting sufficient conditions. For example, for the problem $\#\text{Sub}(\mathcal{A})$, if all linear combinations of \mathcal{A} in fact feature exactly one pattern (so we are in the situation of Theorem 1.5), then the linear combination re-expressed over homomorphisms uses graphs of unbounded treewidth if and only if the patterns in \mathcal{A} have bounded vertex-cover number. We can hence recover Theorem 1.5 from Theorem 1.13.

1.5 Counting Vertex-colored Subgraphs

The techniques introduced above are sufficiently robust to handle generalizations to, e.g., the setting of vertex-colored subgraphs, where the vertices of H and G have colors and we count only subgraphs of G with isomorphisms to H that respect colors. A dichotomy for the special case of *colorful* patterns, where every vertex of the pattern H has a different color, follows from earlier results by observing that embeddings and homomorphisms are the same for colorful patterns. For colorful patterns, bounded treewidth is the tractability criterion.

THEOREM 1.14 ([15, 17, 41]). *Let \mathcal{H} be a recursively enumerable class of colorful vertex-colored graphs. If \mathcal{H} has bounded treewidth,*

then the problem $\#\text{Sub}(\mathcal{H})$ of counting colorful subgraphs from \mathcal{H} in vertex-colored host graphs is polynomial-time solvable. Otherwise, it is $\#\text{W}[1]$ -complete when parameterized by $|V(H)|$.

Theorems 1.5 and 1.14 characterize the two extreme cases of counting colored subgraphs: the uncolored and the fully colorful cases. But there is an entire spectrum of colored problems in between these two extremes. What happens when we consider vertex-colored graphs with some colors appearing on more than one vertex? As we gradually move from colorful to uncolored graphs, where exactly is the point when a jump in complexity occurs?

Answering such questions can be nontrivial even for simple patterns such as paths and matchings and can depend very much on how the colors appear on the pattern. Fortunately, by a basis change to (vertex-colored) homomorphisms via (2), we can answer such questions as easily as in the uncolored setting. The only technical change required is that we should consider only partitions ρ that respect the coloring of H , that is, the vertices of H that end up in the same block should have the same color. With these modifications, we can derive the following corollary from Theorem 1.13

THEOREM 1.15. *Let \mathcal{H} be a recursively enumerable class of vertex-colored patterns (or linear combinations thereof) and let \mathcal{A}_{hom} be the class of linear combinations of homomorphisms derived from \mathcal{H} by the identity (2) as discussed above. If there is a finite bound on the treewidth of graphs in \mathcal{A}_{hom} , then $\#\text{Sub}(\mathcal{H})$ is FPT. Otherwise, the problem is $\#\text{W}[1]$ -hard when parameterized by $|V(H)|$.*

Theorem 1.15 raises a number of questions. First, being a corollary of Theorem 1.13, the tractability criterion is quite indirect, whereas we may want to have a more direct structural understanding of the FPT cases. Secondly, Theorem 1.15 does not tell us whether the FPT cases are actually polynomial-time solvable or not. It is quite remarkable that in Theorems 1.5 and 1.14, all FPT cases are actually polynomial-time solvable, leaving no room for "true" FPT cases that are not polynomial-time solvable. It turns out however that, if we consider vertex-colored patterns in their full generality, then such pattern classes actually *do* appear. A prime example of this phenomenon is the case of *half-colorful matchings*, which are vertex-colored k -matchings such that one endpoint of each edge e_i for $i \in \{1, \dots, k\}$ is colored with 0, while the other is colored with i . Since counting perfect matchings in bipartite graphs is $\#\text{P}$ -hard, a trivial argument shows that counting half-colorful matchings is also $\#\text{P}$ -hard: if a bipartite graph with $n + n$ vertices is colored such that one part has color 0 and each vertex of the other part has a distinct color from 1 to n , then the number of half-colorful matchings of size n is exactly the number of perfect matchings. On the other hand, it is not difficult to show that counting the number of half-colorful matchings of size k in a graph colored with colors 0, 1, \dots , k is fixed-parameter tractable. It is essentially a dynamic programming exercise: for any subgraph $H' \subseteq H$ of the half-colorful matching and for any integer i , we want to compute the number of subgraphs of G isomorphic to H' that are allowed to use only the first i vertices of color class 0.

We give a complete classification of the polynomial-time solvable cases of counting colored patterns from a class \mathcal{H} . For classes of patterns (but not linear combinations thereof), we refine the FPT cases of Theorem 1.15 into two classes: those that are polynomial-time

solvable, and those that are not polynomial-time solvable under the complexity assumption Nonuniform Counting Exponential Time Hypothesis. This shows that the existence of half-colorful matchings is the canonical reason why certain classes of patterns require dynamic programming and therefore the full power given by the definition of FPT: those cases are polynomial where the size of the largest half-colorful matching appearing as a subgraph is at most logarithmic in the size of the pattern.

1.6 Organization of the Paper

In §2, we provide basics on parameterized complexity and the graph-theoretical notions used in this paper. We formalize *graph motif parameters* in §3, and in §3.1 we show how to switch between different useful representations of graph motif parameters. In §3.2, we then address computational aspects of graph motif parameters. These results are first put to use in §4, where we count subgraph patterns by reduction to homomorphisms. In §5, we prove hardness results for linear combinations of subgraphs and induced subgraphs under #ETH and $\text{FPT} \neq \#W[1]$. Due to space constraints, the colored subgraph problems are deferred to the full version.

2 PRELIMINARIES

For a proposition P , we use the *Iverson bracket* $[P] \in \{0, 1\}$ to indicate whether P is satisfied. For a potentially infinite matrix M , a *principal submatrix* M_S is a submatrix of M where the selected row and column index sets are the same set S .

2.1 Parameterized Complexity Theory

We refer to the textbooks [16, 22, 44] for background on parameterized complexity theory. Briefly, a *parameterized counting problem* is a function $\Pi : \{0, 1\}^* \rightarrow \mathbb{N}$ that is endowed with a *parameterization* $\kappa : \{0, 1\}^* \rightarrow \mathbb{N}$; it is *fixed-parameter tractable* (FPT) if there is a computable function $f : \mathbb{N} \rightarrow \mathbb{N}$ and an algorithm to compute $\Pi(x)$ in time $f(k) \cdot \text{poly}(n)$, where $n = |x|$ and $k = \kappa(x)$.

A *parameterized Turing reduction* is a Turing reduction from a problem (Π, κ) to a problem (Π', κ') such that the reduction runs in $f(\kappa(x)) \cdot \text{poly}(|x|)$ time on instances $\Pi(x)$ and each oracle query $\Pi'(y)$ satisfies $\kappa'(y) \leq g(k)$. Here, both f and g are computable functions. A parameterized problem is *#W[1]-hard* if there is a parameterized Turing reduction from the problem of counting the k -cliques in a given graph; since it is believed that the latter does not have an FPT-algorithm, #W[1]-hardness is a strong indicator that a problem is not FPT.

The exponential time hypothesis (ETH) by Impagliazzo and Paturi [26] asserts that satisfiability of 3-CNF formulas cannot be decided substantially faster than by trying all possible assignments. The counting version of this hypothesis [18] states that there is a constant $c > 0$ such that no deterministic algorithm can compute #3-SAT in time $\exp(c \cdot n)$, where n is the number of variables. Chen et al. [9] proved that ETH implies the hypothesis that there is no $f(k) \cdot n^{o(k)}$ -time algorithm to decide whether an n -vertex graph G contains a k -clique. Their reduction is parsimonious, so #ETH rules out $f(k) \cdot n^{o(k)}$ -time algorithms for counting k -cliques.

2.2 Graphs, Subgraphs, and Homomorphisms

Let \mathcal{G} be the set of all labeled, finite, undirected, and simple graphs; in particular, these graphs contain neither loops nor parallel edges. That is, there is a suitable fixed and countably infinite universe U , and \mathcal{G} contains all finite graphs G with vertex set $V(G) \subseteq U$ and edge set $E(G) \subseteq \binom{V(G)}{2}$.

Subgraphs. If G is a graph, a *subgraph* F of G is a graph with $V(F) \subseteq V(G)$ and $E(F) \subseteq E(G)$, and F is an *induced subgraph* of G if it is a subgraph with the additional property that, for all $uv \notin E(F)$ we have $uv \notin E(G)$. The set of subgraphs of G that are isomorphic to H is denoted with $\text{Sub}(H \rightarrow G)$, and the set of induced subgraphs of G that are isomorphic to H is denoted with $\text{IndSub}(H \rightarrow G)$.

Homomorphisms and related notions. If H and G are graphs, a *homomorphism* from H to G is a function $f : V(H) \rightarrow V(G)$ such that edges map to edges under f . That is, for all $\{u, v\} \in E(H)$, we have $\{f(u), f(v)\} \in E(G)$. The set of all homomorphisms from H to G is denoted with $\text{Hom}(H \rightarrow G)$. *Embeddings* are injective homomorphisms, and we denote the corresponding set with $\text{Emb}(H \rightarrow G)$. *Strong embeddings* are embeddings with the additional property that non-edges map to non-edges, that is, $\{f(u), f(v)\} \notin E(G)$ holds for all $\{u, v\} \notin E(H)$. We denote the set of strong embeddings with $\text{StrEmb}(H \rightarrow G)$. A homomorphism $f \in \text{Hom}(H \rightarrow G)$ is *surjective* if it hits all vertices and edges of G , that is, $f(V(H)) = V(G)$ and $f(E(H)) = E(G)$ hold. An *isomorphism* from H to G is a strong embedding from H to G that is also surjective, and it is an *automorphism* if additionally $H = G$ holds. We write $H \simeq G$ if H and G are isomorphic.

Colored graphs. We also use vertex-colored graphs G , where each vertex has a color from a finite set C of colors via a function $f : V(G) \rightarrow C$. We note that such a coloring is not necessarily proper, in the sense that any two adjacent vertices need to receive distinct colors. Each set $V_i(G) := f^{-1}(i)$ for $i \in C$ is a *color class* of G . A subgraph H of G is called (vertex-) *colorful* if $V(H)$ intersects each color class in exactly one vertex.

Treewidth. A *tree decomposition* of a graph G is a pair (T, β) , where T is a tree and β is a mapping from $V(T)$ to $2^{V(G)}$ such that, for all vertices $v \in V(G)$, the set $\{t \in V(T) : v \in \beta(t)\}$ is nonempty and connected in T , and for all edges $e \in E(G)$, there is some node $t \in V(T)$ such that $e \subseteq \beta(t)$. The set $\beta(t)$ is the *bag at t* . The *width* of (T, β) is the integer $\max\{|\beta(t)| - 1 : t \in V(T)\}$, and the *treewidth* $\text{tw}(G)$ of a graph G is the minimum possible width of any tree decomposition of G .

3 THE SPACE OF GRAPH MOTIF PARAMETERS

We develop our interpretation of the general setup of Lovász [39]. To this end, it will be useful to consider unlabeled graphs: For concreteness, we say that a graph $H \in \mathcal{G}$ is *unlabeled* if it is *canonically labeled*, that is, if it is the lexicographically first graph that is isomorphic to H . Then the set \mathcal{G}^* of unlabeled graphs is the subset of \mathcal{G} that contains exactly the canonically labeled graphs. *Graph parameters* are functions $f : \mathcal{G} \rightarrow \mathbb{Q}$ that are invariant under isomorphisms, and we view them as functions $f : \mathcal{G}^* \rightarrow \mathbb{Q}$.

For all $H, G \in \mathcal{G}^*$, we define $\text{IndSub}(H, G)$ as the number of (labeled) induced subgraphs of G that are isomorphic to H . We can view this function as an infinite matrix with indices from $\mathcal{G}^* \times \mathcal{G}^*$ and entries from \mathbb{N} . The matrix indices are ordered according to some fixed total order on \mathcal{G}^* that respects the total size $|V(F)| + |E(F)|$ of the graphs $F \in \mathcal{G}^*$. Among graphs of the same total size, ties may be broken arbitrarily. If H and G are graphs such that H has larger total size than G , then H cannot be an induced subgraph of G , that is, $\text{IndSub}(H, G) = 0$. We conclude that IndSub is an upper triangular matrix.

We define graph motif parameters as graph parameters that can be expressed as finite linear combinations of induced subgraph numbers. To obtain a clean formulation in terms of linear algebra, we represent these linear combinations as infinite vectors $\alpha \in \mathbb{Q}^{\mathcal{G}^*}$ of finite support. Here, the *support* $\text{supp}(\alpha)$ of a vector α is the set of all graphs $F \in \mathcal{G}^*$ with $\alpha_F \neq 0$.

Definition 3.1. A graph parameter $f : \mathcal{G}^* \rightarrow \mathbb{Q}$ is a *graph motif parameter* if there is a vector $\alpha \in \mathbb{Q}^{\mathcal{G}^*}$ with finite support such that $f(G) = \sum_{F \in \mathcal{G}^*} \alpha_F \cdot \text{IndSub}(F, G)$ holds for all $G \in \mathcal{G}^*$.

If we interpret f and α as row vectors, this definition can also be phrased as requiring $f = \alpha \cdot \text{IndSub}$ for α of finite support. Here, for two vectors $\alpha, \beta \in \mathbb{Q}^{\mathcal{G}^*}$, we define the scalar product (α, β) as $\sum_{F \in \mathcal{G}^*} \alpha_F \cdot \beta_F$ if this sum is defined.¹ The definition of the matrix-vector and matrix-matrix products is then as usual.

The set of all graph motif parameters, endowed with the operations of scalar multiplication and pointwise addition, forms an infinite-dimensional vector space. More specifically, it is the finitely supported row-span of the matrix IndSub . We remark that even if we drop the condition of α being finitely supported, the scalar product $\alpha \cdot \text{IndSub}$ remains well-defined, since every column of IndSub has finite support (as a consequence of every graph G having only finitely many induced subgraphs H). In fact, it can be verified that every graph parameter f can be written as $\alpha \cdot \text{IndSub}$ for some α .

3.1 Relations between Graph Motif Parameters

One may wonder why we chose induced subgraph numbers for our definition of graph motif parameters and not, say, the numbers of subgraphs or homomorphisms. It turns out that all of these choices lead to the same vector space: Subgraph and homomorphism numbers are graph motif parameters themselves, and indeed they also span the space of graph motif parameters.

Since some properties of graph motif parameters, such as their computational complexity, turn out to be easier to understand over the homomorphism basis, we show explicitly how to perform basis transformations. To this end, we first present the basis transformation between subgraphs and induced subgraphs, then we proceed with the basis transformation between subgraphs and homomorphisms.

Subgraphs and induced subgraphs. For graphs $H, G \in \mathcal{G}^*$, we first show how to express $\text{Sub}(H, G)$ as a linear combination of numbers $\text{IndSub}(F, G)$. To this end, note that every subgraph copy of H in G is contained in some induced subgraph F of G on $|V(H)|$ vertices. This induced subgraph F is isomorphic to a supergraph of H , and

¹In this paper, all such scalar products degenerate into finite sums, since the support of at least one of the vectors will be finite.

we call these supergraphs *F extensions*. More precisely, an *extension* of H is a (labeled) supergraph X of H with $V(X) = V(H)$.

Note that H might have different extensions that are isomorphic. Thus, given a graph F , let $\text{Ext}(H, F)$ be the number of extensions X of H that are isomorphic to F ; equivalently, we have

$$\text{Ext}(H, F) = [|V(H)| = |V(F)|] \cdot \text{Sub}(H, F),$$

and we thus obtain

$$\text{Sub}(H, G) = \sum_{F \in \mathcal{G}^*} \text{Ext}(H, F) \cdot \text{IndSub}(F, G). \quad (5)$$

Every graph H admits only finitely many extensions, and so the function $\text{Sub}(H, \star)$ is a graph motif parameter for every fixed H . In matrix notation, the identity (5) takes on the concise form

$$\text{Sub} = \text{Ext} \cdot \text{IndSub}. \quad (6)$$

Since Sub , Ext , and IndSub are upper triangular matrices whose diagonal entries are all equal to 1, every finite principal submatrix of any of these triangular matrices is invertible; indeed the entire matrix Ext has an inverse with $\text{IndSub} = \text{Ext}^{-1} \cdot \text{Sub}$. This implies that Sub also *spans* the space of graph motif parameters: We can express every function $\text{IndSub}(H, \star)$ as a linear combination of functions $\text{Sub}(F, \star)$ with coefficients $\text{Ext}^{-1}(H, F)$.

We remark that the values of the coefficients $\text{Ext}^{-1}(H, F)$ are actually well understood: The identity (5) can be interpreted as a zeta transform over the subset lattice [7, eq. (13) and (14)], so we can perform Möbius inversion to prove that

$$\text{Ext}^{-1}(H, F) = (-1)^{|E(F)| - |E(H)|} \cdot \text{Ext}(H, F) \quad (7)$$

holds for all graphs H and F . We will use this identity later to check that $\text{Ext}^{-1}(H, F) \neq 0$ holds for specific pairs (H, F) of graphs.

Homomorphisms and subgraphs. We wish to express $\text{Hom}(H, G)$ as a finitely supported linear combination $\sum_F \alpha_F \text{Sub}(F, G)$ of subgraph numbers. For a homomorphism h from H to G , let I be the image of h , that is, the graph with vertex set $f(V(H))$ and edge set $f(E(H))$; we observe that h is a surjective homomorphism from H to I and I is a subgraph of G . That is, every homomorphism from H to G can be written as a surjective homomorphism into a subgraph F of G . Writing $\text{Surj}(H, F)$ for the number of surjective homomorphisms from H to F , we have

$$\text{Hom}(H, G) = \sum_{F \in \mathcal{G}^*} \text{Surj}(H, F) \cdot \text{Sub}(F, G). \quad (8)$$

Note that $\text{Surj}(H, F) = 0$ holds if H is smaller than F in total size. Thus, analogously to the case of subgraphs, for each fixed H , we have $\text{Surj}(H, F) \neq 0$ only for finitely many $F \in \mathcal{G}^*$. Therefore $\text{Hom}(H, \star)$ is indeed a graph motif parameter for every fixed H . In matrix notation, we have

$$\text{Hom} = \text{Surj} \cdot \text{Sub}, \quad (9)$$

where Surj is a lower triangular matrix. Moreover, the diagonal entries of Surj satisfy $\text{Surj}(F, F) = \text{Aut}(F) \neq 0$, and hence each finite principal submatrix is invertible. In fact the entire matrix has an inverse Surj^{-1} satisfying $\text{Sub} = \text{Surj}^{-1} \cdot \text{Hom}$, and so Hom spans the space of graph motif parameters as well.

The inverse of Surj can be understood in terms of a Möbius inversion on a partition lattice. To this end, let us first consider the support of the vector $\text{Surj}(H, \star)$, that is, the set of all unlabeled

graphs that are homomorphic images of H . This set will play an important role throughout this paper, and we call it the *spasm* of H :

$$\text{Spasm}(H) = \{F \in \mathcal{G}^* : \text{Surj}(H, F) > 0\}. \quad (10)$$

In a more graph-theoretical interpretation, the elements in the spasm of H can also be understood as the unlabeled representatives of all graphs that can be obtained from H by merging independent sets. We make this more formal in the following definition. For each $H \in \mathcal{G}$, let $\text{Part}(H)$ be the set of all partitions of $V(H)$, where a partition is a set of disjoint non-empty subsets $B \subseteq V(H)$ whose union equals $V(H)$.

Definition 3.2. For a graph $H \in \mathcal{G}$ and a partition $\rho \in \text{Part}(H)$, the *quotient* H/ρ is the graph obtained by identifying, for each block $B \in \rho$, the vertices in B to a single vertex. This process may create loops or parallel edges; we keep loops intact in H/ρ , and we turn parallel edges into simple edges.

We get that $F \in \text{Spasm}(H)$ for $F \in \mathcal{G}^*$ if and only if there is a partition $\rho \in \text{Part}(H)$ with $F \simeq H/\rho$. Note that $F \in \mathcal{G}^*$ does not have loops, since we restricted the graphs in \mathcal{G}^* to be simple. Consequently, $F \simeq H/\rho$ can only hold if all blocks of ρ are independent sets of H , that is, if ρ represents a proper coloring of H .

Every surjective homomorphism from H to F can be interpreted as a pair (ρ, π) where $H/\rho \simeq F$ and $\pi \in \text{Aut}(F)$. Hence we have

$$\text{Surj}(H, F) = \#\text{Aut}(F) \cdot \sum_{\rho \in \text{Part}(H)} [H/\rho \simeq F], \quad (11)$$

For two partitions $\rho, \rho' \in \text{Part}(H)$, we write $\rho \geq \rho'$ if ρ is coarser than ρ' , that is, if every block of ρ' is contained in a block of ρ . This partial order gives rise to the *partition lattice* $(\text{Part}(H), \geq)$ whose minimal element \perp is the finest partition, i.e., the partition whose blocks all have size one. Now (8) can be viewed as a zeta-transformation on the partition lattice: Let H and G be fixed graphs. Let $f(\rho) = \#\text{Emb}(H/\rho \rightarrow G)$. Then consider its *upwards zeta-transform* on the partition lattice, i.e., the function \hat{f} defined by

$$\hat{f}(\rho) = \sum_{\rho' \geq \rho} f(\rho').$$

We observe that $\text{Hom}(H, G) = \hat{f}(\perp)$. By Möbius inversion, we get (see also [7, eq. (15)]):

$$f(\rho) = \sum_{\rho' \geq \rho} (-1)^{|\rho| - |\rho'|} \cdot \left(\prod_{B \in \rho'} (\lambda(\rho, \rho', B) - 1)! \right) \cdot \hat{f}(\rho'),$$

where $\lambda(\rho, \rho', B)$ is the number of blocks $C \in \rho$ with $C \subseteq B$. We set $\rho = \perp$ and collect terms ρ' that lead to isomorphic graphs H/ρ' . Note that, for a given graph isomorphism type, all terms leading to this type are non-zero and have the same sign. We obtain

$$\text{Surj}^{-1}(H, F) = \frac{(-1)^{|V(H)| - |V(F)|}}{\#\text{Aut}(H)} \cdot \sum_{\substack{\rho \in \text{Part}(H) \\ H/\rho \simeq F}} \prod_{B \in \rho} (|B| - 1)!. \quad (12)$$

In particular, this yields $\text{Surj}^{-1}(H, F) \neq 0$ if and only if $\text{Surj}(H, F) \neq 0$; that is, $F \in \text{Spasm}(H)$ is equivalent to $\text{Surj}^{-1}(H, F) \neq 0$. This observation will be crucial in the proof of our hardness result.

While we established before that Hom spans the space of graph motif parameters, it is not immediately clear that the homomorphism numbers form a *basis*, that is, that the rows of Hom are




Figure 2: The matrix identity $\text{Hom}_S = \text{Surj}_S \cdot \text{Sub}_S$, where S is the spasm of the path

consisting of the four graphs

linearly independent. The following proposition on the invertibility of certain principal submatrices will be important for our hardness results, and it implies that the rows of Hom are linearly independent with respect to finite linear combinations.

LEMMA 3.3 (PROPOSITION 5.43 IN [39]). *Let $S \subseteq \mathcal{G}^*$ be a finite set of graphs that is closed under surjective homomorphisms, that is, we have $\text{Spasm}(H) \subseteq S$ for all $H \in S$. Then the principal submatrix Hom_S of Hom is invertible and satisfies $\text{Hom}_S = \text{Surj}_S \cdot \text{Sub}_S$.*

PROOF. Let $F, G \in S$ and consider the expansion of $\text{Hom}(H, G)$ from (8). Since S is closed under surjective homomorphisms, only terms with $F \in S$ contribute to the sum. Hence we have $\text{Hom}_S = \text{Surj}_S \cdot \text{Sub}_S$. Since Surj_S and Sub_S are triangular matrices with non-zero diagonal entries, they are both invertible and so is Hom_S . \square

See Figure 2 for an example of Lemma 3.3. Let us also note three simple and useful properties that $\text{Spasm}(H)$ inherits from H .

FACT 3.4. *For all graphs H , the following properties hold:*

- (1) *Every graph $F \in \text{Spasm}(H)$ has at most $|V(H)|$ vertices and at most $|E(H)|$ edges.*
- (2) *If H has a vertex-cover of size $b \in \mathbb{N}$, then every graph $F \in \text{Spasm}(H)$ has a vertex-cover of size b .*
- (3) *If H contains a matching with $k \in \mathbb{N}$ edges, then every graph with k edges (and no isolated vertices) can be found as a minor of some graph in $\text{Spasm}(H)$.*

PROOF. Only the third claim merits some explanation. If M_k is the (not necessarily induced) k -matching in H and F is the k -edge graph we want to find, then we determine an arbitrary surjective homomorphism $g : V(M_k) \rightarrow V(F)$, which hits every edge of F . We are allowed to contract edges (due to the minor operation) or consolidate non-edges of H (due to the quotient operation) to build F . To this end, we simply identify all vertices of $g^{-1}(i)$, for each $i \in V(F)$, and delete all vertices in $V(H) \setminus V(M_k)$. \square

Embeddings and strong embeddings. For completeness, we define matrices for embeddings and strong embeddings. For $H, G \in \mathcal{G}^*$, let $\text{Emb}(H, G)$, $\text{StrEmb}(H, G)$, and $\text{Iso}(H, G)$ be the number of embeddings, strong embeddings, and isomorphisms from H to G , respectively. Clearly Iso is a diagonal matrix with $\text{Iso}(F, F) = \text{Aut}(F)$. We have $\text{Emb} = \text{Iso} \cdot \text{Sub}$ and $\text{StrEmb} = \text{Iso} \cdot \text{IndSub}$.

3.2 The Complexity of Graph Motif Parameters

Several computational problems can be associated with graph motif parameters, but perhaps the most natural one is the *evaluation problem*: Given as input a graph motif parameter $f : \mathcal{G}^* \rightarrow \mathbb{Q}$ and a graph $G \in \mathcal{G}^*$, compute the value $f(G)$.

This problem requires a suitable representation of the input f , and while we could choose any basis to represent f , the homomorphism basis turns out to be particularly useful for algorithmic purposes. That is, in this subsection, we represent graph motif parameters f as vector-matrix products $f = \alpha \cdot \text{Hom}$ for finitely supported row vectors $\alpha \in \mathbb{Q}^{\mathcal{G}^*}$. The input is then the coefficient vector α , encoded as a list of pairs (F, α_F) for $F \in \text{supp}(\alpha)$. Let $|\alpha|$ be the description length of α , and let $\text{tw}(\alpha)$ be the maximum treewidth $\text{tw}(F)$ among all graphs $F \in \text{supp}(\alpha)$. The following lemma is immediate:

LEMMA 3.5 (ALGORITHM). *There is a deterministic algorithm that is given α and G to compute $(\alpha \cdot \text{Hom})(G)$ in time $g(\alpha) + \text{poly}(|\alpha|) \cdot |V(G)|^{\text{tw}(\alpha)+1}$ for some computable function g depending only on α .*

PROOF. For each $F \in \text{supp}(\alpha)$, run the algorithm from Proposition 1.6 to compute the number $\text{Hom}(F, G)$ in time $\exp(O(k)) + \text{poly}(k) \cdot n^{\text{tw}(F)+1}$ where $k = |V(F)|$ and $n = |V(G)|$. Then output the sum $\sum_F \alpha_F \text{Hom}(F, G)$. \square

We could choose other representations for f , such as coefficient vectors α with $f = \alpha \cdot \text{Sub}$ or $f = \alpha \cdot \text{IndSub}$. However, switching between these representations only adds an overhead of $g(\alpha)$ in the running time, as can be seen from §3.1.

It is clear that the generic evaluation problem for graph motif parameters is $\#\text{W}[1]$ -hard, since it subsumes counting k -cliques as a special case. The following reduction shows that evaluating f with $f = \alpha \cdot \text{Hom}$ is at least as hard as every individual homomorphism problem $\text{Hom}(F, \star)$ for $F \in \text{supp}(\alpha)$. That is, if a linear combination of homomorphisms contains a “hard” pattern graph, then the entire linear combination is “hard”.

LEMMA 3.6 (EXTRACTING SUMMANDS). *There is a deterministic Turing reduction that is given a finitely supported vector $\alpha \in \mathbb{Q}^{\mathcal{G}^*}$, a graph $F \in \text{supp}(\alpha)$, and a graph $G \in \mathcal{G}^*$ to compute the number $\text{Hom}(F, G)$ with an oracle for the function $(\alpha \cdot \text{Hom})(\star)$. The reduction runs in time $g(\alpha) \cdot \text{poly}(|V(G)|)$ for some computable function g , makes at most $g(\alpha)$ queries to $(\alpha \cdot \text{Hom})(\star)$, and each queried graph has at most $\max_{H \in \text{supp}(\alpha)} |V(H)| \cdot |V(G)|$ vertices.*

PROOF. On input (α, F, G) , the reduction only makes queries of the form $(\alpha \cdot \text{Hom})(G \times X)$ for graphs X , where $G \times X$ is the categorical product, that is, the graph with vertex set $V(G) \times V(X)$ such that (v, x) and (v', x') are adjacent in $G \times X$ if and only if $vv' \in E(G)$ and $xx' \in E(X)$. The following holds [39, (5.30)]:

$$\text{Hom}(F, G \times X) = \text{Hom}(F, G) \cdot \text{Hom}(F, X). \quad (13)$$

Using this identity for various X , we aim at setting up a linear equation system that can be solved uniquely for $\text{Hom}(F, G)$. We expand the sum $(\alpha \cdot \text{Hom})(G \times X)$ and apply (13) to obtain

$$\sum_H \alpha_H \cdot \text{Hom}(H, G) \cdot \text{Hom}(H, X) = (\alpha \cdot \text{Hom})(G \times X). \quad (14)$$

For each graph X , the reduction can compute the right side of this linear equation using the oracle, and it can determine the numbers α_H and $\text{Hom}(H, X)$ in some time $f(\alpha)$. It remains to choose a suitable set S of graphs X so that the resulting system of linear equations can be uniquely solved for $\text{Hom}(F, G)$.

Let $S = \bigcup_{H \in \text{supp}(\alpha)} \text{Spasm}(H)$ be the closure of $\text{supp}(\alpha)$ under spasms. By Lemma 3.3, the matrix Hom_S is invertible. Moreover, we have $\alpha \cdot \text{Hom} = \alpha_S \cdot \text{Hom}_S$. We rewrite (14) as $\text{Hom}_S \cdot x = b$ where $b \in \mathbb{Q}^S$ and $\text{Hom}_S \in \mathbb{Q}^{S \times S}$ represent the known quantities with $b_X = (\alpha \cdot \text{Hom})(G \times X)$ and $\text{Hom}_S(H, X) = \text{Hom}(H, X)$, and the vector $x \in \mathbb{Q}^S$ represents the indeterminates with $x_H = \alpha_H \cdot \text{Hom}(H, G)$. We have $x = (\text{Hom}_S)^{-1} \cdot b$, so we can solve uniquely for the indeterminates. In particular, we can compute $\text{Hom}(F, G) = x_F / \alpha_F$, since $\alpha_F \neq 0$ holds by assumption.

The set S and the matrices Hom_S and $(\text{Hom}_S)^{-1}$ can be computed in time $g(\alpha)$ for some computable function g . For the vector b , we need to compute the product graphs and query the oracle. The number of queries is $|S|$ and thus bounded by $g(\alpha)$. Overall, the reduction takes time $g(\alpha) \cdot \text{poly}(|V(G)|)$. \square

Analogously to the problems $\#\text{Sub}(\mathcal{H})$ in §1.1, we consider restricted classes of graph motif parameters (represented as linear combinations of homomorphism numbers) to obtain more expressive hardness results.

Definition 3.7. Let $\mathcal{A} \subseteq \mathbb{Q}^{\mathcal{G}^*}$ be a set of finitely supported vectors. We let $\#\text{Hom}(\mathcal{A})$ be the computational problem whose task is to compute $(\alpha \cdot \text{Hom})(G)$ on input $\alpha \in \mathcal{A}$ and $G \in \mathcal{G}^*$.

We apply Lemma 3.6 to establish the hard cases of $\#\text{Hom}(\mathcal{A})$.

LEMMA 3.8 (HARDNESS). *Let $\mathcal{A} \subseteq \mathbb{Q}^{\mathcal{G}^*}$ be a recursively enumerable class of finitely supported vectors. If \mathcal{A} contains vectors of arbitrarily large treewidth $\text{tw}(\alpha)$, then $\#\text{Hom}(\mathcal{A})$ is $\#\text{W}[1]$ -hard when parameterized by $|\alpha|$. Moreover, the problem does not have $g(\alpha) \cdot |V(G)|^{o(\text{tw}(\alpha)/\log \text{tw}(\alpha))}$ time algorithms if $\#\text{ETH}$ holds.*

PROOF. For each $\alpha \in \mathcal{A}$, we select a graph $F_\alpha \in \text{supp}(\alpha)$ of maximum treewidth. Then the set \mathcal{F} with $\mathcal{F} = \{F_\alpha : \alpha \in \mathcal{A}\}$ has unbounded treewidth. To prove the hardness, we provide a parameterized Turing reduction from $\#\text{Hom}(\mathcal{F})$ to $\#\text{Hom}(\mathcal{A})$. Since \mathcal{F} has unbounded treewidth, Theorem 1.8 implies that $\#\text{Hom}(\mathcal{F})$ is $\#\text{W}[1]$ -hard when parameterized by $|V(F)|$ and Proposition 1.9 implies it cannot be computed in time $g(F) \cdot |V(G)|^{o(\text{tw}(F)/\log \text{tw}(F))}$ for any computable g if $\#\text{ETH}$ holds.

Let (F, G) with $F \in \mathcal{F}$ be an input for the reduction, whose goal is to compute $\text{Hom}(F, G)$ with oracle access to $\#\text{Hom}(\mathcal{A})$. First the reduction computes some $\alpha \in \mathcal{A}$ with $F_\alpha = F$. This is possible, since \mathcal{A} is recursively enumerable. The reduction then applies the algorithm from Lemma 3.6 on input (α, F, G) . The algorithm runs in time $g(\alpha) \cdot \text{poly}(|V(G)|)$ for a computable function g and makes queries to the function $(\alpha \cdot \text{Hom})(\star)$; its output is the desired number $\text{Hom}(F, G)$.

For the running time of the reduction, note that finding α takes time $h(k)$ for some computable function h , where $k = |V(F)|$. The algorithm from Lemma 3.6 runs in some time $g(\alpha) \cdot \text{poly}(n) \leq h(k) \cdot \text{poly}(n)$. Hence we indeed obtain a parameterized reduction from $\#\text{Hom}(\mathcal{F})$ to $\#\text{Hom}(\mathcal{A})$, which proves that $\#\text{Hom}(\mathcal{A})$ is $\#\text{W}[1]$ -hard when parameterized by $|\alpha|$. Finally, we have $\text{tw}(\alpha) = \text{tw}(F)$, so

if $\#\text{Hom}(\mathcal{A})$ can be computed in time $f(\alpha) \cdot n^{o(\text{tw}(\alpha)/\log \text{tw}(\alpha))}$, then $\#\text{Hom}(\mathcal{F})$ can be solved in time $f(F) \cdot n^{o(\text{tw}(F)/\log \text{tw}(F))}$, which by Proposition 1.9 is impossible if $\#\text{ETH}$ holds. \square

From the perspective of fine-grained complexity, for every *fixed* graph motif parameter f , Lemma 3.5 yields an algorithm for evaluating f on n -vertex graphs in some time $O(n^c)$. Here, c is a constant that depends on the largest treewidth in the homomorphism representation of f . To express this connection more precisely, given a graph motif parameter f , we define the constant

$$C(f) = \inf\{c \in \mathbb{R} : f \text{ can be computed in time } O(n^c)\}. \quad (15)$$

In particular, we can consider this constant for the graph parameters $\text{Hom}(F, \star)$ for fixed graphs F . Then the constant $C(\text{Hom}(F, \star))$ is the smallest possible exponent required for computing $\text{Hom}(F, G)$ on n -vertex graphs G . The proof of Lemma 3.5 implies

$$C(f) \leq \max_{F \in \text{supp}(\alpha)} C(\text{Hom}(F, \star)),$$

where $\alpha \in \mathbb{Q}^{\mathcal{G}^*}$ is the representation of f over the homomorphism basis, that is, the vector α with $f = \alpha \cdot \text{Hom}$. Lemma 3.6 implies the corresponding lower bound, so in fact we have

$$C(f) = \max_{F \in \text{supp}(\alpha)} C(\text{Hom}(F, \star)), \quad (16)$$

Proposition 1.9 implies that $C(\text{Hom}(F, \star)) \leq \text{tw}(F) + 1$ holds. Under the assumption $\text{FPT} \neq \#\text{W}[1]$, Theorem 1.8 implies that $C(\text{Hom}(F, \star))$ cannot be bounded by a universal constant, and under the stronger assumption $\#\text{ETH}$, Proposition 1.9 implies that $C(\text{Hom}(F, \star))$ is not bounded by $o(\text{tw}(F)/\log \text{tw}(F))$, even when F is restricted to be from any fixed family \mathcal{F} of graphs of unbounded treewidth. The k -clique hypothesis is that the current fastest algorithm for k -clique is optimal [1], which can be formalized as $C(\text{Hom}(K_k, \star)) = \omega k/3$. These facts suggest that the representation of a graph motif parameter f in the homomorphism basis and your favorite complexity hypotheses are all that is needed to understand the complexity of f (concerning the exponent of n).

Remark 3.9. Lovász's framework is easily adapted to the setting of vertex-colored pattern and host graphs. Here, we only wish to count homomorphisms (or embeddings, or strong embeddings) from H into G that map vertices of H to vertices of G of the same colors. In the definition of the spasm of H , one is then only allowed to identify non-adjacent vertices of the same color, that is, the allowed partitions ρ are those where each block is a monochromatic independent set. The algorithm (Lemma 3.5) and hardness result (Lemma 3.8) can also be adapted to the setting of vertex-colored homomorphisms without modifications. We excluded these variants from the main text to simplify the presentation.

4 ALGORITHMS FOR COUNTING SUBGRAPHS

In this section, we obtain algorithms for counting subgraphs and embeddings by expressing these quantities over the homomorphism basis via (8) and running an algorithm for counting homomorphisms. More concretely, recall that

$$\text{Sub}(H, G) = \sum_F \text{Surj}^{-1}(H, F) \cdot \text{Hom}(F, G). \quad (17)$$

We know from (12) that $\text{Surj}^{-1}(H, F) \neq 0$ is equivalent to $F \in \text{Spasm}(H)$. Hence for fixed patterns H , if we can compute the homomorphism numbers $\text{Hom}(F, G)$ for all $F \in \text{Spasm}(H)$ in time $O(n^c)$ on n -vertex graphs G , then we can compute $\text{Sub}(H, G)$ in time $O(n^c)$. When using the basic treewidth-based algorithm from Proposition 1.6, this running time $O(n^c)$ is governed by the maximum treewidth among the graphs in $\text{Spasm}(H)$.

Theorem 1.1 (restated). There is a deterministic algorithm that is given a k -edge graph H and an n -vertex graph G to compute $\#\text{Sub}(H \rightarrow G)$ in time $k^{O(k)} n^{t+1}$, where t is the maximum treewidth among all graphs in $\text{Spasm}(H)$.

PROOF. We obtain $\#\text{Sub}(H \rightarrow G)$ by evaluating the right side of (17) in the straightforward manner: First compute the spasm of H and all coefficients $\text{Surj}^{-1}(H, F)$ for $F \in \text{Spasm}(H)$ by first computing $\text{Surj}_{\text{Spasm}(H)}$ via brute-force and then inverting this triangular matrix. Then, for each $F \in \text{Spasm}(H)$, compute $\#\text{Hom}(F \rightarrow G)$ via Proposition 1.6. Finally, we compute the sum on the right side of (17) to obtain $\#\text{Sub}(H \rightarrow G)$.

For the running time claim, note that the size of $\text{Spasm}(H)$ is bounded by the number of partitions of the set $V(H)$, which in turn can be bounded crudely by $k^{O(k)}$. Each term $\#\text{Hom}(F \rightarrow G)$ can be computed in time $\exp(O(k)) \cdot n^{t+1}$ by Proposition 1.6. \square

Theorem 1.1 is particularly useful for sparse pattern graphs H : We mentioned in Fact 3.4 that any k -edge graph H only contains graphs with at most k edges in its spasm. Using this fact, we can exploit known bounds on the treewidth of sparse graphs to bound the running time guaranteed by Theorem 1.1: If F has k edges, then $\text{tw}(F) \leq ck + o(k)$ is known to hold for fairly small constants $c < 1$. Furthermore, there are k -edge graphs F with $\text{tw}(F) = \Theta(k)$. Since the best known upper and lower bounds on the treewidth of k -edge graphs are not tight, it will be useful to dedicate a universal constant ξ to the linear coefficient in the treewidth bound.

Definition 4.1. For $k \in \mathbb{N}$, let $\text{tw}^*(k)$ be defined as the maximum treewidth among all graphs with k edges. We define the constant $\xi \in \mathbb{R}$ as the limit superior

$$\xi = \limsup_{k \rightarrow \infty} \frac{\text{tw}^*(k)}{k}.$$

Using the existence of good 3-regular expanders, Dvořák and Norin [20, Corollary 7] find a family of 3-regular graphs on k edges and $n = \frac{2}{3}k$ vertices with treewidth at least $\frac{1}{24}n - 1 = \frac{1}{36}k - 1$ for all large enough k . On the other hand, Scott and Sorkin [48, Corollary 21] prove that $\text{tw}^*(k) \leq \frac{13}{75}k + o(k)$ holds. Collecting these two results, we get the following bounds on ξ .

THEOREM 4.2 ([20, 48]). *We have $\frac{1}{36} \leq \xi \leq \frac{13}{75}$.*

Since $\xi \leq \frac{13}{75} < 0.174$, this immediately implies upper bounds on the running times obtained from Theorem 1.1.

Corollary 1.2 (restated). There is a deterministic algorithm that is given H and G to compute $\#\text{Sub}(H \rightarrow G)$. The algorithm runs in time $f(H) \cdot |V(G)|^{\xi \cdot k + o(k)}$, where $k = |E(H)|$ and $\xi < 0.174$ is the constant from Definition 4.1.

Instead of relying on Proposition 1.6 to count homomorphisms in Theorem 1.1, we can use more sophisticated methods where available. For instance, in the full version, we use fast matrix multiplication to count homomorphisms from patterns H of treewidth at most two, thus proving Theorem 1.3.

5 COMPLEXITY OF LINEAR COMBINATION PROBLEMS

5.1 Linear Combinations of Subgraphs

We prove the dichotomy stated in Theorems 1.4 and 1.5. Recall that due to the basis transformation (17), we can express linear combinations of subgraph numbers as equivalent linear combinations of homomorphism numbers. By Lemma 3.8, the most difficult homomorphism number in this linear combination governs the complexity of the problem. Since the hardness criterion for homomorphisms is treewidth, and expressing subgraph numbers in the homomorphism basis yields terms for all graphs in the spasm, we first make the following observation.

FACT 5.1. *Let H be a graph. If H has a maximum matching of size k , the maximum treewidth among all graphs in $\text{Spasm}(H)$ is $\Theta(k)$.*

PROOF. Let k be the size of a maximum matching of H . Then the vertex-cover number of H is at least k and at most $2k$. By the second item of Fact 3.4, the vertex-cover number and thus the treewidth of graphs in the spasm of H is then also at most $2k$. For the lower bound, we use the third item of Claim 3.4: Since H contains a matching of size k , every k -edge graph occurs as a minor of some graph in $\text{Spasm}(H)$. By Theorem 4.2, there exist k -edge graphs F with treewidth $\text{tw}^*(k) \geq \Omega(k)$. Moreover, taking minors does not increase the treewidth, so there is a graph in $\text{Spasm}(H)$ with treewidth at least $\Omega(k)$. \square

Let $\mathcal{A} \subseteq \mathbb{Q}^{\mathcal{G}^*}$ be a family of finitely supported vectors, and recall the matrices Sub and Surj from Section 3.1. We define the set $\mathcal{A} \cdot \text{Surj}^{-1}$ as the set of all vectors $\alpha \cdot \text{Surj}^{-1}$ for $\alpha \in \mathcal{A}$. That is, for each $\alpha \in \mathcal{A}$, the graph motif parameter $\alpha \cdot \text{Sub}$ appears in the set $\mathcal{A} \cdot \text{Surj}^{-1}$ via its representation in the homomorphism basis.

THEOREM 5.2. *Let $\mathcal{A} \subseteq \mathbb{Q}^{\mathcal{G}^*}$ be a recursively enumerable family of finitely supported vectors. If there is a constant $t \in \mathbb{N}$ such that all vectors $\beta \in \mathcal{A} \cdot \text{Surj}^{-1}$ have treewidth $\text{tw}(\beta) \leq t$, then the problem $\#\text{Sub}(\mathcal{A})$ to compute*

$$\sum_{H \in \mathcal{G}^*} \alpha_H \cdot \#\text{Sub}(H \rightarrow G),$$

on input $\alpha \in \mathcal{A}$ and an n -vertex graph G , admits an algorithm with running time $g(\alpha) \cdot n^{t+1}$. Otherwise, it is $\#\text{W}[1]$ -hard parameterized by the description length of α , and it cannot be computed in time $g(\alpha) \cdot n^{o(t/\log t)}$ for $t = \text{tw}(\alpha \cdot \text{Surj}^{-1})$ unless $\#\text{ETH}$ fails.

PROOF. Recall that $\text{Sub} = \text{Surj}^{-1} \text{Hom}$ holds by (9), so we have $\alpha \text{Sub} = \alpha \text{Surj}^{-1} \text{Hom}$ for all $\alpha \in \mathcal{A}$. This implies the algorithmic claim via Lemma 3.5 and the hardness claims via Lemma 3.8. \square

Since any family of subgraph patterns \mathcal{H} can be represented as a family \mathcal{A} of linear combinations in which each vector has support one, we obtain Theorem 1.4 and Theorem 1.5 as a special case. The quantitative lower bound regarding the vertex-cover number

follows from the relationship between the vertex-cover number and the largest treewidth in the spasm via Fact 5.1.

PROOF OF THEOREMS 1.4 AND 1.5. Let \mathcal{H} be a graph family of unbounded vertex-cover number. Its closure $\bigcup_{H \in \mathcal{H}} \text{Spasm}(H)$ has unbounded treewidth by Fact 5.1. We want to apply Theorem 5.2. Let \mathcal{A} be the set of all α^H where $\alpha_F^H = [H = F]$ holds for some graph $H \in \mathcal{H}$. That is, α^H contains H with coefficient 1, and no other graphs. Clearly $\#\text{Sub}(\mathcal{A})$ is equivalent to $\#\text{Sub}(\mathcal{H})$.

Now consider the set \mathcal{B} with $\mathcal{B} = \mathcal{A} \cdot \text{Surj}^{-1}$. We need to prove that the graph class $\bigcup_{\beta \in \mathcal{B}} \text{supp}(\beta)$ has unbounded treewidth. We do so by showing that this class is in fact equal to $\bigcup_{H \in \mathcal{H}} \text{Spasm}(H)$. By definition, for each $H \in \mathcal{H}$, the set \mathcal{B} contains a vector β^H with $\beta^H = \alpha^H \cdot \text{Surj}^{-1}$. Expanding this vector-matrix product, we get

$$\beta_F^H = \sum_{J \in \mathcal{G}^*} \alpha_J^H \cdot \text{Surj}^{-1}(J, F).$$

Since $\alpha_J^H = [H = J]$, we have $\beta_F^H = \text{Surj}^{-1}(H, F)$. By (12), the latter is non-zero if and only if $F \in \text{Spasm}(H)$, so the claim follows. \square

We present an example that does not directly follow from [15], but that does follow from Theorem 5.2. The following statement was proved recently using a more complicated method.

COROLLARY 5.3 ([8]). *Given k and G , counting all trees with k vertices in G is $\#\text{W}[1]$ -hard when parameterized by k .*

PROOF. The number of k -vertex trees can be seen as a linear combination of subgraph numbers; for each fixed k , we set $\alpha_F = 1$ for all unlabeled graphs $F \in \mathcal{G}^*$ such that F is a k -vertex tree, and $\alpha_F = 0$ otherwise. Let \mathcal{A} be the family of all such α over all $k \in \mathbb{N}$. Since the class of all trees has unbounded vertex cover number, Fact 5.1 shows that the union of spasms of k -vertex trees has unbounded treewidth as k grows.

Write \mathcal{T}_k for the set of all k -vertex trees. For each $k \in \mathbb{N}$, pick a graph $F_k \in \text{Spasm}(\mathcal{T}_k)$ such that the sequence F_1, F_2, \dots has unbounded treewidth. In order to apply Theorem 5.2, we need to prove that some vector $\beta \in \mathcal{A} \cdot \text{Surj}^{-1}$ indeed has F_k in its support: To this end, let $\alpha \in \mathcal{A}$ be the vector corresponding to all k -vertex trees and let $\beta = \alpha \cdot \text{Surj}^{-1}$. We claim that F_k is contained in the support of β . To see this, we expand the matrix-vector product:

$$\beta_{F_k} = (\alpha \cdot \text{Surj}^{-1})_{F_k} = \sum_{T \in \mathcal{T}_k} \text{Surj}^{-1}(T, F_k). \quad (18)$$

Recall from (12) that $\text{Surj}^{-1}(T, F_k) \neq 0$ if and only if $F_k \in \text{Spasm}(T)$. The same equation implies that the sign of $\text{Surj}^{-1}(T, F_k)$ is equal to $(-1)^{|V(T)| - |V(F_k)|}$ for all $T \in \mathcal{T}_k$ with $F_k \in \text{Spasm}(T)$. Since $|V(T)| = k$ holds for all $T \in \mathcal{T}_k$, the sign is in fact $(-1)^{k - |V(F_k)|}$. Therefore, all terms in the sum of (18) have the same sign and at least one term is non-zero; thus $\beta_{F_k} \neq 0$ holds as claimed. \square

The preceding corollary also holds for counting k -vertex forests, and can be generalized to families \mathcal{A} where each linear combination α contains only graphs with the same number of vertices, and where \mathcal{A} contains graphs of unbounded vertex-cover number in their support. In fact, the graphs in $\text{supp}(\alpha)$ do not even need to have the same number of vertices, but the same *parity* of number of vertices suffices. The proof is analogous to that of Corollary 5.3.

5.2 Linear Combinations of Induced Subgraphs

Chen, Thurley, and Weyer [12] consider the restriction of computing $\#\text{StrEmb}(H \rightarrow G)$, parameterized by $k = |V(H)|$, when the graphs H are chosen from some class \mathcal{H} . They prove that the problem is $\#\text{W}[1]$ -hard if \mathcal{H} is infinite, otherwise it is polynomial-time computable. Recall that counting strong embeddings is equivalent to counting induced subgraphs. In full analogy to Theorem 5.2, we can generalize their result to a classification of linear combinations of induced subgraph numbers. For the following statement, recall the matrices Ext and Surj from Section 3.1.

Theorem 1.13 (restated). Let $\mathcal{A} \subseteq \mathbb{Q}^{\mathcal{G}^*}$ be a recursively enumerable family of finitely supported vectors. If there is a constant $t \in \mathbb{N}$ such that all vectors $\beta \in \mathcal{A} \cdot \text{Ext}^{-1} \cdot \text{Surj}^{-1}$ have treewidth $\text{tw}(\beta) \leq t$, then the problem $\#\text{Ind}(\mathcal{A})$ to compute

$$\sum_{H \in \mathcal{G}^*} \alpha_H \cdot \#\text{IndSub}(H \rightarrow G),$$

on input $\alpha \in \mathcal{A}$ and an n -vertex graph G , admits an algorithm with running time $g(\alpha) \cdot n^{t+1}$. Otherwise, it is $\#\text{W}[1]$ -hard parameterized by the description length of α , and it cannot be solved in time $g(\alpha) \cdot n^{o(t/\log t)}$ for $t = \text{tw}(\alpha \cdot \text{Ext}^{-1} \cdot \text{Surj}^{-1})$ unless $\#\text{ETH}$ fails.

The proof is analogous to that of Theorem 5.2, using the basis change matrix $\text{Ext}^{-1} \cdot \text{Surj}^{-1}$ rather than Surj^{-1} .

On a related note, Jerrum and Meeks [29–31, 41] introduced the following generalization of counting induced subgraphs: Let Φ be some graph property. The problem $\#\text{IndProp}(\Phi)$ is, given a graph G and an integer k , to compute the number of induced k -vertex subgraphs that have property Φ . Let us write this number as $I_{\Phi, k}(G)$. Since $I_{\Phi, k}(G)$ can be expressed as a sum $\sum_H \text{IndSub}(H, G)$ over all k -vertex graphs H that satisfy Φ , Theorem 1.13 immediately implies a complexity dichotomy for these problems.

Corollary 1.11 (restated). Let Φ be any decidable graph property. The problem $\#\text{IndProp}(\Phi)$ is fixed-parameter tractable if all $I_{\Phi, k}$ can be represented as linear combinations of homomorphisms from graphs of bounded treewidth. Otherwise, the problem is $\#\text{W}[1]$ -hard when parameterized by k .

Finally, let us sketch how to recover the hardness result of Chen, Thurley, and Weyer [12] for counting induced subgraphs from a fixed class \mathcal{H} as a special case of Corollary 1.11: When representing $\#\text{IndSub}(H \rightarrow G)$ for a k -vertex graph H as a linear combination of subgraph numbers $\#\text{Sub}(H' \rightarrow G)$ via (5), this linear combination has a non-zero coefficient $\text{Ext}^{-1}(H, H')$ for the clique $H' = K_k$. Indeed, the k -clique extends every graph on k vertices, so we have $\text{Ext}(H, H') \neq 0$, which in turn implies $\text{Ext}^{-1}(H, H') \neq 0$ by (7). When further writing each term $\#\text{Sub}(H' \rightarrow G)$ as a linear combination of homomorphisms $\#\text{Hom}(H'' \rightarrow G)$ for $H'' \in \text{Spasm}(H')$, we get exactly one term for $H'' = H' = K_k$, since K_k only occurs in its own spasm, and we have $\text{Surj}^{-1}(H', H'') = 1/\#\text{Aut}(H')$. Hence the coefficient of $\#\text{Hom}(K_k \rightarrow G)$ in the representation of $\#\text{IndSub}(H \rightarrow G)$ is non-zero. Thus, if \mathcal{H} contains infinitely many graphs, we have unbounded cliques in the homomorphism representation, and the problem of counting induced subgraphs from \mathcal{H} is $\#\text{W}[1]$ -hard by Corollary 1.11.

6 OPEN PROBLEMS

We have defined the space of graph motif parameters and explored three useful bases thereof, namely, Hom , Sub , and IndSub . These bases capture well-studied classes of counting problems, and we could use basis changes to transfer results between these classes. Are there other computationally interesting bases? Moreover, are there other interesting subspaces of the space of all graph parameters other than the graph motif parameters?

ACKNOWLEDGMENTS

Thanks a lot to Édouard Bonnet for pointing out [48] and [20].

REFERENCES

- [1] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. If the current clique algorithms are optimal, so is Valiant's parser. In *Proceedings of the 56th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 98–117, 2015.
- [2] Noga Alon, Phuong Dao, Iman Hajirasouliha, Fereydoun Hormozdiari, and S Cenk Sahinalp. Biomolecular network motif counting and discovery by color coding. *Bioinformatics*, 24(13):i241–i249, 2008.
- [3] Noga Alon, Raphael Yuster, and Uri Zwick. Finding and counting given length cycles. *Algorithmica*, 17(3):209–223, 1997.
- [4] Andreas Björklund, Thore Husfeldt, Petteri Kaski, and Mikko Koivisto. Counting paths and packings in halves. In *Proceedings of the 17th Annual European Symposium on Algorithms (ESA)*, pages 578–586, 2009.
- [5] Andreas Björklund, Petteri Kaski, and Lukasz Kowalik. Counting thin subgraphs via packings faster than meet-in-the-middle time. In *Proceedings of the 25th Annual Symposium on Discrete Algorithms (SODA)*, pages 594–603, 2014.
- [6] Markus Bläser and Radu Curticapean. Weighted counting of k -matchings is $\#\text{W}[1]$ -hard. In *Proceedings of the 7th International Symposium on Parameterized and Exact Computation (IPEC)*, pages 171–181, 2012.
- [7] Christian Borgs, Jennifer Chayes, László Lovász, Vera T Sós, and Katalin Vesztegombi. Counting graph homomorphisms. In *Topics in discrete mathematics*, pages 315–371. Springer, 2006.
- [8] Cornelius Brand and Marc Roth. Parameterized counting of trees, forests and matroid bases. In *Proceedings of the 12th International Computer Science Symposium in Russia (CSR)*, 2017 (to appear).
- [9] Jianer Chen, Benny Chor, Mike Fellows, Xiuzhen Huang, David W. Juedes, Iyad A. Kanj, and Ge Xia. Tight lower bounds for certain parameterized NP-hard problems. *Information and Computation*, 201(2):216–231, 2005.
- [10] Jin Chen, Wynne Hsu, Mong Li Lee, and See-Kiong Ng. Nemofinder: Dissecting genome-wide protein-protein interactions with meso-scale network motifs. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 106–115. ACM, 2006.
- [11] Yijia Chen, Martin Grohe, and Bingkai Lin. The hardness of embedding grids and walls. *arXiv preprint arXiv:1703.06423*, 2017.
- [12] Yijia Chen, Marc Thurley, and Mark Weyer. Understanding the complexity of induced subgraph isomorphisms. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 587–596, 2008.
- [13] Stephen A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the 3rd Annual Symposium on Theory of Computing (STOC)*, pages 151–158, 1971.
- [14] Radu Curticapean. Counting matchings of size k is $\text{W}[1]$ -hard. In *Proceedings of the 40th International Conference on Automata, Languages, and Programming (ICALP)*, pages 352–363, 2013.
- [15] Radu Curticapean and Dániel Marx. Complexity of counting subgraphs: Only the boundedness of the vertex-cover number counts. In *Proceedings of the 55th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 130–139, 2014.
- [16] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshantov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- [17] Victor Dalmau and Peter Jonsson. The complexity of counting homomorphisms seen from the other side. *Theoretical Computer Science*, 329(1):315–323, 2004.
- [18] Holger Dell, Thore Husfeldt, Dániel Marx, Nina Taslaman, and Martin Wahlén. Exponential time complexity of the permanent and the Tutte polynomial. *ACM Transactions on Algorithms*, 10(4):21:1–21:32, 2014.
- [19] Josep Diaz, Maria J. Serna, and Dimitrios M. Thilikos. Counting H -colorings of partial k -trees. *Theoretical Computer Science*, 281(1-2):291–309, 2002.
- [20] Zdeněk Dvořák and Sergey Norin. Strongly sublinear separators and polynomial expansion. *SIAM Journal on Discrete Mathematics*, 30(2):1095–1101, 2016.

- [21] Friedrich Eisenbrand and Fabrizio Grandoni. On the complexity of fixed-parameter clique and dominating set. *Theoretical Computer Science*, 326(1):57–67, 2004.
- [22] Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer-Verlag, Berlin, 2006.
- [23] François Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation (ISSAC)*, pages 296–303, 2014.
- [24] Joshua A Grochow and Manolis Kellis. Network motif discovery using subgraph enumeration and symmetry-breaking. In *Annual International Conference on Research in Computational Molecular Biology*, pages 92–106. Springer, 2007.
- [25] Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *Journal of the ACM*, 54(1):1:1–1:24, 2007.
- [26] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k -SAT. *Journal of Computer and System Sciences*, 62(2):367–375, 2001.
- [27] Alon Itai and Michael Rodeh. Finding a minimum circuit in a graph. *SIAM Journal on Computing*, 7(4):413–423, 1978.
- [28] Bart M. P. Jansen and Dániel Marx. Characterizing the easy-to-find subgraphs from the viewpoint of polynomial-time algorithms, kernels, and turing kernels. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 616–629, 2015.
- [29] Mark Jerrum and Kitty Meeks. The parameterised complexity of counting even and odd induced subgraphs. *CoRR*, abs/1410.3375, 2014.
- [30] Mark Jerrum and Kitty Meeks. The parameterised complexity of counting connected subgraphs and graph motifs. *Journal of Computer and System Sciences*, 81(4):702–716, 2015.
- [31] Mark Jerrum and Kitty Meeks. Some hard families of parameterized counting problems. *ACM Transactions on Computation Theory*, 7(3):11, 2015.
- [32] Zahra Razaghi Moghadam Kashani, Hayedeh Ahrabian, Elahe Elahi, Abbas Nowzari-Dalini, Elnaz Saberi Ansari, Sahar Asadi, Shahin Mohammadi, Falk Schreiber, and Ali Masoudi-Nejad. Kavosh: A new algorithm for finding network motifs. *BMC bioinformatics*, 10(1):318, 2009.
- [33] Nadav Kashtan, Shalev Itzkovitz, Ron Milo, and Uri Alon. Efficient sampling algorithm for estimating subgraph concentrations and detecting network motifs. *Bioinformatics*, 20(11):1746–1758, 2004.
- [34] Ton Kloks, Dieter Kratsch, and Haiko Müller. Finding and counting small induced subgraphs efficiently. *Information Processing Letters*, 74(3):115–121, 2000.
- [35] Ioannis Koutis and Ryan Williams. LIMITS and applications of group algebras for parameterized problems. *ACM Transactions on Algorithms*, 12(3):31:1–31:18, 2016.
- [36] Mirosław Kowaluk, Andrzej Lingas, and Eva-Marta Lundell. Counting and detecting small subgraphs via equations. *SIAM Journal on Discrete Mathematics*, 27(2):892–909, 2013.
- [37] Bingkai Lin. The parameterized complexity of k -biclique. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 605–615, 2015.
- [38] László Lovász. Operations with structures. *Acta Mathematica Hungarica*, 18(3-4):321–328, 1967.
- [39] László Lovász. *Large networks and graph limits*, volume 60. American Mathematical Society Providence, 2012.
- [40] Dániel Marx. Can you beat treewidth? *Theory of Computing*, 6(1):85–112, 2010.
- [41] Kitty Meeks. The challenges of unbounded treewidth in parameterised subgraph counting problems. *Discrete Applied Mathematics*, 198:170–194, 2016.
- [42] Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. Network motifs: Simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.
- [43] Jaroslav Nešetřil and Svatopluk Poljak. On the complexity of the subgraph problem. *Commentationes Mathematicae Universitatis Carolinae*, 26(2):415–419, 1985.
- [44] Rolf Niedermeier. *Invitation to fixed-parameter algorithms*, volume 31 of *Oxford Lecture Series in Mathematics and its Applications*. Oxford University Press, Oxford, 2006.
- [45] Saeed Omid, Falk Schreiber, and Ali Masoudi-Nejad. MODA: An efficient algorithm for network motif discovery in biological networks. *Genes & genetic systems*, 84(5):385–395, 2009.
- [46] Benjamin Schiller, Sven Jager, Kay Hamacher, and Thorsten Strufe. Stream - A stream-based algorithm for counting motifs in dynamic graphs. In *Proceedings of the 2nd International Conference on Algorithms for Computational Biology (AlCoB)*, pages 53–67, 2015.
- [47] Falk Schreiber and Henning Schwöbbermeyer. Frequency concepts and pattern detection for the analysis of motifs in networks. In *Transactions on computational systems biology III*, pages 89–104. Springer, 2005.
- [48] Alexander D. Scott and Gregory B. Sorkin. Linear-programming design and analysis of fast algorithms for Max 2-CSP. *Discrete Optimization*, 4(3-4):260–287, 2007.
- [49] Julian R Ullmann. An algorithm for subgraph isomorphism. *Journal of the ACM*, 23(1):31–42, 1976.
- [50] Leslie G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8(2):189–201, 1979.
- [51] Sebastian Wernicke. Efficient detection of network motifs. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 3(4), 2006.
- [52] Virginia Vassilevska Williams. Multiplying matrices faster than Coppersmith-Winograd. In *Proceedings of the 44th Symposium on Theory of Computing Conference (STOC)*, pages 887–898, 2012.
- [53] Virginia Vassilevska Williams and Ryan Williams. Subcubic equivalences between path, matrix and triangle problems. In *Proceedings of the 51st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 645–654, 2010.
- [54] Virginia Vassilevska Williams and Ryan Williams. Finding, minimizing, and counting weighted subgraphs. *SIAM Journal on Computing*, 42(3):831–854, 2013.