Counting Small Induced Subgraphs with Hereditary Properties*

Jacob Focke
CISPA Helmholtz Center for Information Security
Saarbrücken, Germany

Marc Roth
Department of Computer Science, University of Oxford
Oxford, UK

ABSTRACT

We study the computational complexity of the problem #IndSub(Φ) of counting k-vertex induced subgraphs of a graph G that satisfy a graph property Φ. Our main result establishes an exhaustive and explicit classification for all hereditary properties, including tight conditional lower bounds under the Exponential Time Hypothesis (ETH):

- If a hereditary property Φ is true for all graphs, or if it is true only for finitely many graphs, then #IndSub(Φ) is solvable in polynomial time.
- Otherwise, #IndSub(Φ) is #W[1]-complete when parameterised by k, and, assuming ETH, it cannot be solved in time $f(k) \cdot |G|^{o(k)}$ for any function f.

This classification features a wide range of properties for which the corresponding detection problem (as classified by Khot and Raman [TCS 02]) is tractable but counting is hard. Moreover, even for properties which are already intractable in their decision version, our results yield significantly stronger lower bounds for the counting problem.

As additional result, we also present an exhaustive and explicit parameterised complexity classification for all properties that are invariant under homomorphic equivalence.

By covering one of the most natural and general notions of closure, namely, closure under vertex-deletion (hereditary), we generalise some of the earlier results on this problem. For instance, our results fully subsume and strengthen the existing classification of #IndSub(Φ) for monotone (subgraph-closed) properties due to Roth, Schmitt, and Wellnitz [FOCS 20].

A full version of our paper, containing all proofs, is available at https://arxiv.org/abs/2111.02277.

CCS CONCEPTS

• **Theory of computation → Problems, reductions and completeness**; • **Mathematics of computing → Combinatorics**; **Graph theory**.

*Research of the first author is part of a project that is supported by the European Research Council (ERC) consolidator grant No. 725978 SYSTEMATICGRAPH.

STOC '22, June 20–24, 2022, Rome, Italy

ACM ISBN 978-1-4503-9264-8/22/06…$15.00
https://doi.org/10.1145/3519935.3520008

KEYWORDS

Counting complexity, parameterized complexity, induced subgraphs, hereditary properties, fine-grained complexity, graph homomorphisms

ACM Reference Format:
Jacob Focke and Marc Roth. 2022. Counting Small Induced Subgraphs with Hereditary Properties. In *Proceedings of the 54rd Annual ACM SIGACT Symposium on Theory of Computing (STOC '22), June 20–24, 2022, Rome, Italy.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3519935.3520008

1 EXTENDED ABSTRACT

Detection and counting of patterns in networks belong to the most well-studied problems in theoretical computer science and have applications in a diverse set of disciplines such as database theory [17], statistical physics [22, 23, 33], and computational biology [16, 32]. In this work, we focus on counting *small* patterns in *large* networks. Among others, this task is motivated by the computation of so-called significance profiles of network motifs which play a central role in the analysis of complex networks [1, 26, 27, 31].

More formally, we consider the counting problem #IndSub(Φ) as introduced by Jerrum and Meeks [19].[1] Here, a graph property Φ is a function from the class of graphs to $\{0, 1\}$ that is invariant under graph isomorphisms. If $\Phi(H) = 1$ for a graph H, then H *satisfies* the property Φ. For any fixed graph property Φ, the problem #IndSub(Φ) asks, on input a graph G and a positive integer k, to compute the number of k-vertex induced subgraphs H in G that satisfy Φ. Observe that, for proper choices of Φ, the problem #IndSub(Φ) encodes a variety of well-studied counting problems such as counting of k-cliques, k-independent sets, induced k-cycles, and, to name a more intricate example, k-graphlets, that is, connected k-vertex induced subgraphs.

In recent years, the problem #IndSub(Φ) received significant attention [7, 11, 19–21, 25, 29, 30]. All of the previous works had the common goal of classifying the *parameterised* complexity of #IndSub(Φ) for a wide range of properties Φ. More precisely, the task is to identify those properties Φ for which the problem becomes *fixed-parameter tractable (FPT)*, i.e., solvable in time $f(k)\cdot|G|^{O(1)}$ for some computable function f. Note that a parameterised analysis of #IndSub(Φ) captures well the intuition that the size of the pattern k is significantly smaller than the size of the graph G, that is, we only aim for a running time which is polynomial in $|G|$ but may be super-polynomial in k.

Ideally, a complete classification of #IndSub(Φ) identifies not only the properties Φ for which the problem becomes FPT, but also establishes a hardness result for all remaining properties. A remarkable result due to Curticapean, Dell and Marx [7] shows that such a complete classification is possible: They prove that for every property Φ, the problem #IndSub(Φ) is either fixed-parameter

[1] In [19], #IndSub(Φ) is called #InducedUnlabelledSubgraphWithProperty(Φ).

tractable or complete for the parameterised class #W[1].[2] Unfortunately, their classification is implicit in the sense that, for most graph properties Φ, it is not clear how to pinpoint the complexity of #IndSub(Φ). More precisely, even for simple and natural properties such as $\Phi(H) = 1$ iff H is bipartite, or $\Phi(H) = 1$ iff H is acyclic, the complexity of #IndSub(Φ) is not easily deducible from the aforementioned classification. The corresponding hardness proofs turned out to be a non-trivial task [29]. Subsequent work focused on finding *explicit* criteria for tractability and hardness [11, 29, 30]. More details on the classification from [7] are given in Section 1.2.

The state of the art suggests that the only properties Φ for which #IndSub(Φ) is FPT are very restricted in the sense that they become "eventually trivial". More formally, we say that a property Φ is *meagre* if there exists a positive integer B, such that for each $k \geq B$ the property Φ is either constant false or constant true on the set of all k-vertex graphs. For example, the property Φ of having an even number of vertices is meagre, and it is easy to see that #IndSub(Φ) is trivial to solve: On input G and k, output 0 if k is odd, and output $\binom{|V(G)|}{k}$ if k is even. It is well-known that an analogue of the previous algorithm exists for every meagre property; this is made formal in the full version [15]. Conversely, as stated in [30], we conjecture that all non-meagre properties yield hardness:

Conjecture 1. Let Φ be a computable[3] graph property. If Φ is meagre then #IndSub(Φ) is fixed-parameter tractable. Otherwise, #IndSub(Φ) is #W[1]-complete.

Despite significant effort, we are nowhere close to a resolution of Conjecture 1 in its full generality. However, progress has been made for properties that satisfy certain closure criteria. For example, a result of Jerrum and Meeks [20] implies that Conjecture 1 is true for *minor-closed* graph properties. After several partial results, Conjecture 1 has recently also been established for the more general class of *monotone* properties, that is, properties that are closed under the removal of vertices *and* edges[4] [30].

There are two natural generalisations of the class of monotone properties:

(1) Properties that are closed under vertex-deletion, called *hereditary* properties (these properties are closed under taking induced subgraphs).
(2) Properties that are closed under edge-deletion, called *edge-monotone* properties.

To make this concrete, let us give some simple examples that are covered by the different notions of closure. Say Φ corresponds to the property of being "planar". Then Φ is closed under both vertex- and edge-deletion as well as under edge-contraction. Therefore, Φ is minor-closed and hence the complexity of #IndSub(Φ) is covered by [20]. The property of being "bipartite" is also closed under both vertex- and edge-deletion. However, edge-contractions can lead to non-bipartite graphs. So this is an example of a property that is monotone but not minor-closed, and the corresponding hardness result is from [30]. An example for a hereditary property that is not monotone is the property "claw-free", which refers to the absence of an induced claw. This property is closed under vertex-deletion, but not under edge-deletion. Conversely, the property of being "disconnected" is closed under edge-deletion but not under vertex-deletion. Hence, it is an example for an edge-monotone property that is not monotone (which implies that it is also not hereditary).

So far, there are only partial results on resolving Conjecture 1 for the hereditary and edge-monotone cases; see [30, Section 6] for hereditary properties defined by a single forbidden induced subgraph (which includes the property "claw-free"), and [25, 29] for some results on edge-monotone properties (including the property "disconnected"). Many natural properties are covered by these partial results but a full classification has remained elusive. In this work, we obtain a full classification for the first case, i.e., for all hereditary properties; our results are presented in Section 1.1. At this point let us give an example of a hereditary property that, as far as we are aware, has not been covered by previous work. Consider Φ with $\Phi(H) = 1$ iff H is "hole-free", which means that H does not have an induced cycle of length at least 5 (also known as a *hole*). First note, that Φ is closed under vertex-deletion, but not under edge-deletion. It is therefore hereditary but not monotone (or even minor-closed). Since Φ is characterised by multiple forbidden induced subgraphs it is not covered by [30, Theorem 4]. As Φ does not distinguish bicliques from independent sets it is not subject to [30, Theorem 2]. Finally, Φ does not have low Hamming-weight f-vectors, which is another criterion introduced in [30]. (For this fact, it is relevant that triangles are not forbidden by Φ.) Similar hereditary properties that have not been covered by previous work are "(odd-hole)-free", "(anti-hole)-free", etc. In Section 1.4, we give an example of an unresolved edge-monotone property.

1.1 Our Results

In addition to confirming Conjecture 1 for hereditary properties, we also establish a tight conditional lower bound under the Exponential Time Hypothesis; it turns out that a hereditary property is meagre if and only if either Φ is true for all graphs, or it is true only for finitely many graphs (an easy proof of this fact can be found in the full version [15]).

Theorem 2. *Let Φ be a computable hereditary graph property. If Φ is meagre, then #IndSub(Φ) is solvable in polynomial time. Otherwise #IndSub(Φ) is #W[1]-complete and, assuming the Exponential Time Hypothesis (ETH), cannot be solved in time $f(k) \cdot |G|^{o(k)}$ for any function f.*

Observe that our conditional lower bound under ETH rules out any significant improvement over the brute-force algorithm for #IndSub(Φ), which iterates over every k-vertex subset of $V(G)$ and counts those that induce a subgraph satisfying Φ. The running time of this algorithm is clearly bounded by $f(k) \cdot |V(G)|^{k+O(1)} \leq f(k) \cdot |G|^{O(k)}$ for some computable function f. Note further, that a stronger lower bound ruling out algorithms running in time $f(k) \cdot |V(G)|^{k-\varepsilon}$ for any $\varepsilon > 0$ is not possible: For the (hereditary) property Φ of being a complete graph, the problem #IndSub(Φ) is the problem of counting k-cliques, which can be solved in time

[2] The class #W[1] can be considered as a parameterised counting equivalent of NP; a formal definition is stated in the full version [15].
[3] We restrict ourselves to computable properties to avoid dealing with non-uniform fixed-parameter tractability.
[4] To avoid confusion, we remark that in some literature (e.g. [19, 25]) the term "monotone" is used for properties that are closed under the deletion (or addition) of edges only. The latter will be called *edge-monotone* in this work.

$|V(G)|^{\frac{\omega k}{3}+O(1)}$, where $\omega < 3$ is the matrix multiplication exponent [28].

To compare our results on exact counting with the complexity of decision and approximate counting we partition the class of all hereditary properties as follows; we write I_ℓ and K_ℓ for the independent set and the complete graph of size ℓ, respectively.

(1) Suppose there are positive integers s and t such that Φ is false on K_s and I_t. By Ramsey's Theorem and the fact that Φ is closed under taking induced subgraphs, Φ must then be false on all but finitely many graphs. The problem #IndSub(Φ) is thus solvable in polynomial time, and so are its decision and approximate counting versions.

(2) If, for all positive integers ℓ, the property Φ is true on K_ℓ and I_ℓ then Φ is not meagre, unless it is constant true. Khot and Raman [24] proved that deciding the existence of a k-vertex induced subgraph that satisfies Φ is fixed-parameter tractable. Furthermore, Meeks [25] established the existence of a "fixed-parameter tractable approximation scheme" (FPTRAS) for the counting problem, which can be considered the parameterised notion of an FPRAS (see [2] for a discussion). In sharp contrast, from Theorem 2 it follows that exact counting is intractable (unless Φ is trivially true, in which case #IndSub(Φ) is trivial).

(3) Otherwise, the decision version was shown to be W[1]-hard by Khot and Raman [24]. However, their reduction only yields an implicit ETH-based conditional lower bound of the form $f(k)\cdot|G|^{o(k^c)}$, where $0 < c < 1$ is a constant depending on the set of forbidden induced subgraphs of Φ.[5] While it is unsurprising that Theorem 2 yields #W[1]-hardness of exact counting (since the decision version is W[1]-hard), it is worth to point out that our conditional lower bound significantly improves upon the hardness of decision.

In summary, together with Khot and Raman [24], and Meeks [25], we fully complete the complexity landscape for detection, approximate counting and exact counting induced subgraphs with hereditary properties. In particular, we identify a significant variety of properties for which decision and approximate counting is easy, but exact counting is hard. A more concise overview is given in Table 1.

We note that our classification for hereditary properties subsumes and strengthens the classification of monotone properties due to Roth, Schmitt and Wellnitz [30][6], see Section 1.3.

In the course of establishing Theorem 2, we prove a much stronger technical intractability theorem which is stated in the full version [15]. We have not yet explored the full extent of its applicability and we believe that it will be useful in future work (see Section 1.4). For now, let us present one particular additional consequence: We say that a graph property Φ is *invariant under homomorphic equivalence* if $\Phi(H_1) = \Phi(H_2)$ whenever H_1 and H_2 are homomorphically equivalent, i.e., there are homomorphisms from H_1 to H_2 and from H_2 to H_1. Examples of properties invariant under homomorphic equivalence include

- $\Phi(H) = 1$ if and only if H has odd girth d.
- $\Phi(H) = 1$ if and only if H has clique number d.
- $\Phi(H) = 1$ if and only if H has chromatic number d.

Here, d can be any fixed positive integer. We note that none of the previous works on #IndSub(Φ) reveals its complexity for the previous three properties. We change that in the current work:

Theorem 3. *Let Φ be a computable graph property that is invariant under homomorphic equivalence. If Φ is meagre, then #IndSub(Φ) is solvable in polynomial time. Otherwise, #IndSub(Φ) is #W[1]-complete and, assuming ETH, cannot be solved in time $f(k)\cdot|G|^{o(k)}$ for any function f.*

As a consequence, for each $d \geq 1$ (and d odd in the case of odd girth), all three of the previous examples yield intractability of #IndSub(Φ).

1.2 Technical Overview

Similarly as in previous work [7, 11, 29, 30] we rely on the framework of graph motif parameters and the homomorphism basis as introduced by Curticapean, Dell and Marx [7]: Using the term $\#\mathsf{IndSub}(\Phi, k \to G)$ for the number of k-vertex induced subgraphs of G that satisfy Φ, it is known that there is a unique function $a_{\Phi,k}$ with finite support and independent from G, such that

$$\#\mathsf{IndSub}(\Phi, k \to G) = \sum_H a_{\Phi,k}(H)\cdot\#\mathsf{Hom}(H \to G)\,, \qquad (1)$$

where the sum is over all (isomorphism types of) graphs, and $\#\mathsf{Hom}(H \to G)$ denotes the number of graph homomorphisms from H to G. Let us emphasise that the sum is finite, since $a_{\Phi,k}$ has finite support, that is, $a_{\Phi,k}(H) \neq 0$ only for finitely many H. The *complexity monotonicity principle*, which was independently discovered by Curticapean, Dell and Marx [7] and by Chen and Mengel [3], states that computing a finite linear combination of homomorphism counts as in (1) is precisely as hard as computing its hardest term $\#\mathsf{Hom}(H \to G)$ with a non-zero coefficient $a_{\Phi,k}(H) \neq 0$. Since the complexity of counting homomorphisms from H to G is well-understood — the problem is feasible if and only if the treewidth[7] of H is small [9] — the complexity monotonicity principle shifted the study of the complexity of #IndSub(Φ) and related subgraph counting problems to the purely combinatorial problem of determining the treewidth of the graphs H with a non-zero coefficient $a_{\Phi,k}(H) \neq 0$. More formally, we can define a function t_Φ which maps a positive integer k to the maximum treewidth of a graph H with $a_{\Phi,k}(H) \neq 0$. We then obtain the following *implicit* classification:

Theorem 4 (Corollary 1.11 in [7]). *The problem #IndSub(Φ) is fixed-parameter tractable if t_Φ is bounded by a constant, and #W[1]-complete otherwise.*

For tight(er) lower bounds under ETH, it is additionally necessary that $t_\Phi(k) \in \Omega(k)$.

With Theorem 4 as a powerful tool at hand, recent work focused on establishing an *explicit* criterion for tractability of #IndSub(Φ). More concretely, we note that Conjecture 1 can be resolved if it is

[5] That is, a conditional lower bound that applies to all Φ is of the form $f(k)\cdot|G|^{o(g(k))}$ where g is asymptotically smaller than every proper rational function, e.g., $g(k) = \log(k)$.

[6] However, the classification of properties depending only on the number of edges in [30] is not subsumed.

[7] Intuitively, treewidth is a parameter that measures how tree-like a graph is. In this work, we will rely on treewidth purely in a black-box manner, and thus we refer the reader to Chapter 7 in [8] for a comprehensive treatment of treewidth.

Table 1: Finding and counting k-vertex induced subgraphs that satisfy a hereditary property Φ. The conditional lower bounds for exact counting assume the Exponential Time Hypothesis, and the absence of an FPTRAS for approximate counting is conditioned on the assumption that $\mathrm{W}[1]$ does not coincide with FPT under randomised parameterised reductions.

† **For the property** True **of being constant true, all versions of the problem become trivial.**
‡ **By Ramsey's Theorem, and since Φ is hereditary, the condition $\exists\, s,t : \Phi(K_s) = \Phi(I_t) = 0$ implies that Φ is false for all graphs with at least $R(s,t)$ vertices. As s and t are constants, all versions of the problem can be trivially solved in time $n^{O(R(s,t))} = n^{O(1)}$.**

Condition on $\Phi \neq$ True†	**Decision** Khot & Raman [24]	**Approx. Counting** Meeks [25]	**Exact Counting** This work (Theorem 2)
$\exists\, s,t : \Phi(K_s) = \Phi(I_t) = 0$‡	P	P	P
$\forall\, \ell : \Phi(K_\ell) = \Phi(I_\ell) = 1$	FPT	FPTRAS	$\#\mathrm{W}[1]$-hard, not in $f(k) \cdot \lvert G\rvert^{o(k)}$
Otherwise	$\mathrm{W}[1]$-hard	no FPTRAS	$\#\mathrm{W}[1]$-hard, not in $f(k) \cdot \lvert G\rvert^{o(k)}$

proved that t_Φ is bounded if and only if Φ is meagre. Unfortunately, it turned out that the analysis of t_Φ and thus the analysis of the coefficients $a_{\Phi,k}(H)$ in (1) is a very challenging task in its own right. The reason for the latter is that the coefficients $a_{\Phi,k}(H)$ often encode algebraic and even topological invariants.[8] Despite the latter difficulty, Theorem 4 was successfully used in previous works to resolve Conjecture 1 for some restricted classes of properties, which we will present in more detail in Section 1.4.

In the current work, we side-step the problem of analysing the coefficients in (1) for hereditary properties by considering a bipartite version of $\#\textsc{IndSub}(\Phi)$ as an intermediate step. For the definition of the intermediate problem, we need to consider bipartite graphs G with *fixed* bipartitions $V(G) = U \dot\cup V$ (note that a bipartite graph might have multiple bipartitions). Formally, we will write $\mathbf{G} = (U, V, E)$ to emphasise fixing the left- and right-hand side vertices U and V, respectively. Furthermore, two bipartite graphs $\mathbf{G}_1 = (U_1, V_1, E_2)$ and $\mathbf{G}_2 = (U_2, V_2, E_2)$ are said to be *consistently isomorphic* if there exists an isomorphism that maps U_1 to U_2 and V_1 to V_2, respectively. A *bipartite property* Ψ is then defined to be a function from bipartite graphs to $\{0, 1\}$ such that $\Psi(\mathbf{G}_1) = \Psi(\mathbf{G}_2)$ whenever $\mathbf{G}_1$ and $\mathbf{G}_2$ are consistently isomorphic.

Given a bipartite property Ψ, the problem $\#\textsc{BipIndSub}(\Psi)$ asks, on input a bipartite graph $\mathbf{G}$ (with fixed bipartition!) and a positive integer k, to compute the number of k-vertex induced subgraphs of $\mathbf{G}$ that satisfy Ψ; here, the bipartition of an induced subgraph of $\mathbf{G}$ is induced by the bipartition of $\mathbf{G}$. We stress that $\#\textsc{BipIndSub}(\Psi)$ is *not* the same as the restriction of $\#\textsc{IndSub}(\Phi)$ to bipartite input graphs (without fixed bipartition). For example, $\#\textsc{BipIndSub}(\Psi)$ allows us to express counting of $2k$-vertex induced subgraphs of $\mathbf{G}$ that have k vertices on the left-hand side and k vertices on the right-hand side, or, more interestingly, counting k-vertex induced subgraphs of $\mathbf{G}$ such that there is a vertex on the left-hand side that is adjacent to all vertices on the right-hand side. Both of those examples are not expressible by just restricting $\#\textsc{IndSub}(\Phi)$ to bipartite inputs.

Our proof of Theorem 2 can then be split into two essentially independent parts: First, we establish the following criterion for the intractability of $\#\textsc{BipIndSub}(\Psi)$. To this end, $\mathbf{I}_{k,k}$ denotes an independent set of size $2k$, with a fixed bipartition that contains k vertices on the left-hand side and k vertices on the right-hand side; and $\mathbf{B}_{k,k}$ denotes the complete bipartite graph with k vertices on the left-hand side and k vertices on the right-hand side. Furthermore, we call a set of integers $\mathcal{K}$ *dense* if there exists a constant c such that for every positive integer m, there is a $k \in \mathcal{K}$ with $m \le k \le c \cdot m$.

Theorem 5. *Let Ψ be a computable bipartite property. Let $\mathcal{K}$ be the set of primes k for which Ψ distinguishes $\mathbf{I}_{k,k}$ and $\mathbf{B}_{k,k}$, i.e., $\Psi(\mathbf{I}_{k,k}) \neq \Psi(\mathbf{B}_{k,k})$. If $\mathcal{K}$ is infinite then $\#\textsc{BipIndSub}(\Psi)$ is $\#\mathrm{W}[1]$-hard. Moreover, if $\mathcal{K}$ is dense then $\#\textsc{BipIndSub}(\Psi)$ cannot be solved in time $f(k) \cdot |G|^{o(k)}$ for any function f, assuming the ETH.*

In the second step, we show that for a wide range of properties Φ, including all hereditary properties, we can associate with Φ a bipartite property Ψ_Φ such that $\#\textsc{BipIndSub}(\Psi_\Phi)$ reduces to $\#\textsc{IndSub}(\Phi)$ with respect to parameterised reductions. Additionally, this reduction will be tight in the sense that all conditional lower bounds transfer. Finally, we show that, whenever Φ is not meagre, the bipartite property Ψ_Φ will satisfy the strong hardness condition in Theorem 5, yielding not only $\#\mathrm{W}[1]$-hardness, but also the conditional lower bound under ETH. [9]

In what follows, we will describe both steps in more detail separately.

1.2.1 Classification of Bipartite Properties. The main motivation of our consideration of $\#\textsc{BipIndSub}(\Psi)$ as an intermediate step is the "algebraic approach to hardness" as introduced in [11], which we will describe subsequently.

For a properly defined vertex-coloured version of $\#\textsc{IndSub}(\Phi)$, restricted to bipartite input graphs (but without fixed bipartition), a

[8] As a concrete example, it was shown in [29] that for edge-monotone Φ, the coefficient $a_{\Phi,k}(K_k)$ of the complete graph is equal to the so-called reduced Euler characteristic of the simplicial graph complex associated with Φ. As a consequence, it was established that $a_{\Phi,k}(K_k) \neq 0$ is a sufficient criterion for the property Φ to be *evasive* on k-vertex graphs. Therefore, a proof that $a_{\Phi,k}(K_k)$ does not vanish whenever Φ is non-trivial on k-vertex graphs would resolve Karp's famous Evasiveness-Conjecture. We refer the reader to [29] for a detailed treatment of the connection between the coefficients $a_{\Phi,k}(K_k)$ and the evasiveness of Φ.

[9] For readers familiar with the so-called bipartite double-cover $G \times K_2$, we wish to stress that the latter is *not* used in our reduction, even though this approach may seem tempting at first glance. Unfortunately, due to technical reasons which are out of the scope of this extended abstract, we were not able to obtain an easier proof via the bipartite double-cover.

transformation as linear combination of vertex-coloured homomorphism counts similar as in Equation (1) is known to hold. It was furthermore established that

(I) the complexity monotonicity principle (Theorem 4) remains true in the vertex-coloured setting, and
(II) the coefficient $a_{\Phi,k}(H)$ for H being a complete bipartite graph can be analysed much easier in the vertex-coloured case.

The reason for the simplified analysis of the coefficient in (II) ultimately relied on the fact that the complete bipartite graph is edge-transitive and that it can have a prime-power number of edges; we describe this in more detail when we apply the algebraic approach to the setting of fixed bipartitions further below.

In combination, (I) and (II) were shown to yield a classification similar to Theorem 5 but without considering fixed bipartitions. Unfortunately, our reduction from the bipartite to the non-bipartite case crucially depends on such fixed bipartitions. Therefore, we adapt the algebraic approach to #BipIndSub(Ψ) as follows. Writing $\#\mathsf{BipIndSub}(\Psi, k \to \mathbf{G})$ for the number of (bipartite) k-vertex induced subgraphs of $\mathbf{G}$ that satisfy Ψ, we establish a similar transformation as in Equation (1). We show that there exists a function $a_{\Psi,k}$ of finite support and independent of $\mathbf{G}$ such that[10]

$$\#\mathsf{BipIndSub}(\Phi, k \to \mathbf{G}) = \sum_{H} a_{\Psi,k}(H) \cdot \#\mathsf{Hom}(H \to G)\,, \qquad (2)$$

where the sum is again over all graphs (without fixed bipartitions) and G is the underlying graph of $\mathbf{G}$. Additionally, if $k = 2\ell$, we show that

$$a_{\Psi,k}(B_{\ell,\ell}) = \sum_{A \subseteq E(B_{\ell,\ell})} \Psi(\mathbf{B}_{\ell,\ell}[A]) \cdot (-1)^{\ell^2 - |A|}\,, \qquad (3)$$

where $B_{\ell,\ell}$ is the complete bipartite graph, i.e., the ℓ-by-ℓ biclique, $\mathbf{B}_{\ell,\ell}$ is the ℓ-by-ℓ biclique with fixed bipartition, and $\mathbf{B}_{\ell,\ell}[A]$ is obtained from $\mathbf{B}_{\ell,\ell}$ by removing all edges in $E(B_{\ell,\ell}) \setminus A$. The goal is to show that $a_{\Psi,k}(B_{\ell,\ell})$ in (3) is non-zero whenever ℓ is a prime and $\Psi(\mathbf{I}_{\ell,\ell}) \neq \Psi(\mathbf{B}_{\ell,\ell})$. Since the treewidth of $B_{\ell,\ell}$ is linear in ℓ, we can rely on a vertex-coloured version of complexity monotonicity such as in (I) to prove that $a_{\Psi,k}(B_{\ell,\ell}) \neq 0$ is sufficient for the classification of #BipIndSub(Ψ) (Theorem 5).

The subtle difference between (3) and the analysis in [11], which prevents us from using the main result of [11] in a black-box manner, is that the edge-subgraphs of $\mathbf{B}_{\ell,\ell}$ keep their fixed bipartition, and only the subgraphs for which the bipartite property Ψ holds contribute to the sum. More precisely, the main result of [11] is achieved by considering the canonical action of the automorphism group $\mathsf{Aut}(B_{\ell,\ell})$ on the set of edge-subsets and observing that the term $\Psi(B_{\ell,\ell}[A]) \cdot (-1)^{\ell^2-|A|}$ is invariant under this action if no bipartition is fixed and if Ψ is a graph property rather than a bipartite property. However, in our case, Ψ is a bipartite property that respects the bipartition. Thus, for an automorphism π of $B_{\ell,\ell}$ which maps vertices from the left-hand side to the right-hand side and vice versa, there might be an edge-subset A such that

$$\Psi(\mathbf{B}_{\ell,\ell}[A]) \neq \Psi(\mathbf{B}_{\ell,\ell}[\pi(A)])\,.$$

[10] In fact, for technical reasons, we establish Equation (2) in a vertex-coloured setting, which is, however, shown to be interreducible with the uncoloured setting. The formal treatment can be found in the full version [15].

Fortunately, we can easily solve this problem by restricting to automorphisms that are *consistent*, i.e., which map the left-hand side to the left-hand side and the right-hand side to the right-hand side. Writing $\mathsf{Aut}(\mathbf{B}_{\ell,\ell})$ for the set of consistent automorphisms, it is easy to see that $\mathsf{Aut}(\mathbf{B}_{\ell,\ell})$ still acts transitively on the edges of the complete bipartite graph $\mathbf{B}_{\ell,\ell}$, that is, for each pair of edges e and f of $\mathbf{B}_{\ell,\ell}$, there exists $\pi \in \mathsf{Aut}(\mathbf{B}_{\ell,\ell})$ such that $\pi(e) = f$. With that observation at hand, we can apply the algebraic approach similarly as in [11]; for now we provide a concise outline and refer the reader to the full version [15] for the detailed presentation.

To establish that $a_{\Psi,k}(B_{\ell,\ell})$ does not vanish under the previous constraints, we will first observe that $\#\mathsf{Aut}(\mathbf{B}_{\ell,\ell})$ is divisible by ℓ. Thus, we can show that there exists an ℓ-Sylow subgroup Γ of $\mathsf{Aut}(\mathbf{B}_{\ell,\ell})$ whose action on the edges of $\mathbf{B}_{\ell,\ell}$ is still transitive. Extending this action to edge-subsets of $\mathbf{B}_{\ell,\ell}$, we observe that for each pair A_1 and A_2 in the same orbit, we have that

$$\Psi(\mathbf{B}_{\ell,\ell}[A_1]) = \Psi(\mathbf{B}_{\ell,\ell}[A_2])\,.$$

Since the size of each orbit must divide the order of the group, which is the prime ℓ, we can take Equation (3) modulo ℓ, and observe that only the fixed points survive, that is $A = \emptyset$ and $A = E(B_{\ell,\ell})$. In other words, we obtain

$$a_{\Psi,k}(B_{\ell,\ell}) = \Psi(\mathbf{B}_{\ell,\ell}) + \Psi(\mathbf{I}_{\ell,\ell}) \cdot (-1)^{\ell^2} \mod \ell\,, \qquad (4)$$

which is non-zero whenever ℓ is a prime and $\Psi(\mathbf{B}_{\ell,\ell}) \neq \Psi(\mathbf{I}_{\ell,\ell})$. As outlined previously, this will suffice for establishing the classification of #BipIndSub(Ψ) (Theorem 5).

1.2.2 *Reducing from Bipartite Properties to Graph Properties using False Twins.*

Let Φ be a hereditary graph property that is not meagre. In order to confirm Conjecture 1 for such Φ we will relate Φ to a bipartite property Ψ_Φ, which satisfies the requirements of the hardness result from Theorem 5. Intuitively, this means that Ψ_Φ should distinguish, for certain k, the independent set $\mathbf{I}_{k,k}$ from the biclique $\mathbf{B}_{k,k}$. Additionally, we have to establish that #BipIndSub(Ψ_Φ) reduces to #IndSub(Φ).

How could we define such a property Ψ_Φ? Since Φ is hereditary it can be classified by a (possibly infinite) set of (inclusion-minimal) forbidden induced subgraphs $\Pi(\Phi)$. Since Φ is not meagre $\Pi(\Phi)$ contains at least one element, say H. Consider the following initial construction; an illustration is provided in Figure 1. Suppose that H contains an edge $e = \{u_1, u_2\}$. Then replacing this edge with a complete bipartite graph $\mathbf{B}_{k,k}$ yields a graph $H_{\mathbf{B}_{k,k}}$ that contains H as induced subgraph, which means that $\Phi(H_{\mathbf{B}_{k,k}}) = 0$. Now suppose further that replacing e with an independent set $\mathbf{I}_{k,k}$ yields a graph $H_{\mathbf{I}_{k,k}}$ for which $\Phi(H_{\mathbf{I}_{k,k}}) = 1$. Then the process of replacing the edge e with some bipartite graph $\mathbf{G}$ — note that the choice of the bipartition matters — and evaluating Φ for the resulting graph $H_{\mathbf{G}}$ defines a bipartite property that distinguishes $\mathbf{I}_{k,k}$ from $\mathbf{B}_{k,k}$.

However, for this approach to work in general, it is essential that $\Phi(H_{\mathbf{I}_{k,k}}) = 1$, i.e., that $H_{\mathbf{I}_{k,k}}$ does not contain any forbidden induced subgraph. This suggests some kind of "minimal" choice of the graph $H \in \Pi(\Phi)$ in the general case. Consider the graphs in Figure 2 as forbidden induced subgraphs. The graph on the left has fewer vertices and edges. However, when replacing any of its edges by $\mathbf{I}_{2,2}$, we obtain the graph on the right. So it does not suffice to look at the number of vertices or edges alone.

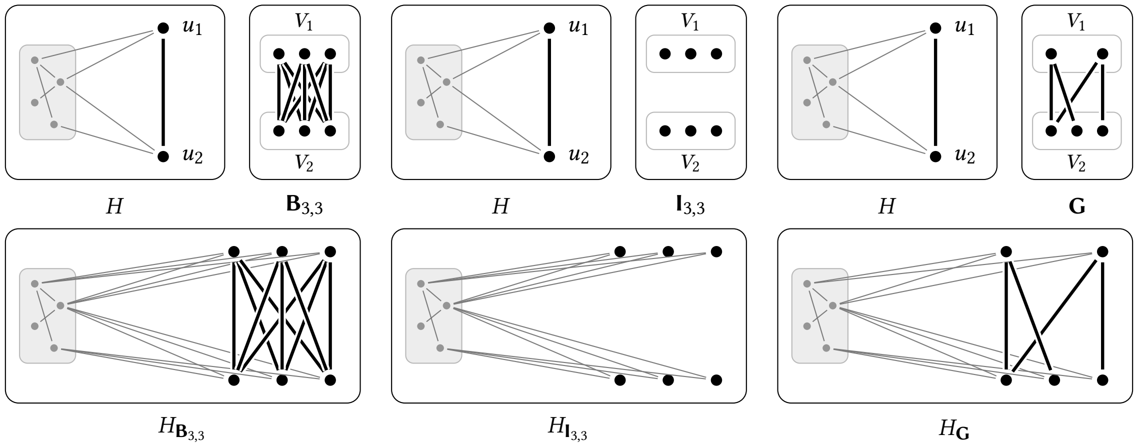

Figure 1: Replacing an edge e of H by the 3-by-3 biclique, the 3-by-3 independent set, and the bipartite graph G.

For the choice of H it turns out to be helpful to consider sets of vertices in H that have identical neighbourhood (so-called *false twins*). The false twin relation partitions the vertices of H into *blocks*. By $H\!\downarrow$ we denote the corresponding quotient graph, in which each block is replaced by a single vertex. We refer to it as the *twin-free quotient*.[11] Note that the twin-free quotient of the left-hand graph in Figure 2 is K_4 itself, whereas the twin-free quotient of the right-hand graph is K_4 minus an edge.

For our refined construction that will ultimately lead to the definition of the bipartite property Ψ_Φ, we will choose a graph H from $\Pi(\Phi)$ for which $H\!\downarrow$ has minimal number of edges. For an edge $\{u_1, u_2\}$ in H let B_1 and B_2 be the blocks containing u_1 and u_2, respectively. Given a bipartite graph $\mathbf{G} = (V_1, V_2, E)$ we define a graph $F_\mathbf{G}$, where we now replace not only the edge $\{u_1, u_2\}$ but the complete bipartite graph induced by B_1 and B_2, and insert in its stead the graph $\mathbf{G}$, where V_1 replaces B_1, and V_2 replaces B_2, see Figure 3 for an example. As in the initial construction, for sufficiently large k, H is an induced subgraph of $F_{\mathbf{B}_{k,k}}$, which implies $\Phi(F_{\mathbf{B}_{k,k}}) = 0$. However, one can also show that for every induced subgraph F' of $F_{\mathbf{I}_{k,k}}$, it holds that $|E(F'\!\downarrow)| \leq |E(F_{\mathbf{I}_{k,k}}\!\downarrow)| < |E(H\!\downarrow)|$, and so, by our choice of H, F' is not in $\Pi(\Phi)$, i.e., it is not a forbidden induced subgraph of Φ. Consequently, $\Phi(F_{\mathbf{I}_{k,k}}) = 1$ as intended. This way we establish that the bipartite property Ψ_Φ with $\Psi_\Phi(\mathbf{G}) := \Phi(F_\mathbf{G})$ distinguishes, for sufficiently large k, the independent set $\mathbf{I}_{k,k}$ from the biclique $\mathbf{B}_{k,k}$ and thereby satisfies the requirements of Theorem 5. This shows that #BipIndSub(Ψ_Φ) is #W[1]-hard with the corresponding conditional lower bound.

It is worth to mention that Conjecture 1 was previously confirmed for hereditary properties that only have a *single forbidden induced subgraph*, i.e., for the case $|\Pi(\Phi)| = 1$ [30]. The corresponding proof uses the idea of replacing a single edge that we described in the initial construction. It then boils down to a reduction from counting independent sets. In our work, we significantly generalise the gadget construction, and by using #BipIndSub(Ψ_Φ), where Ψ_Φ depends on Φ, we also broaden the class of problems we reduce from. Note that the problem of counting independent sets is a very special case of a property that distinguishes independent sets from bicliques.

We continue by giving an overview of the tight parameterised reduction from #BipIndSub(Ψ_Φ) to #IndSub(Φ). Given a bipartite graph $\mathbf{G} = (V_1, V_2, E)$ together with a positive integer k as input to #BipIndSub(Ψ_Φ) let $F_\mathbf{G}$ be as defined previously. Let $R = V(F_\mathbf{G}) \setminus (V_1 \cup V_2)$, so R contains the vertices of H that are outside of B_1 and B_2 (which were replaced by V_1 and V_2 in the construction of $F_\mathbf{G}$). Let $k' = k + |R|$. One can show that for the sought-for number of k-vertex induced subgraphs of $\mathbf{G}$ that satisfy Ψ_Φ we have

$$\#\mathsf{BipIndSub}(\Psi_\Phi, k \to \mathbf{G}) = \#\{S \in \mathsf{IndSub}(\Phi, k' \to F_\mathbf{G}) \mid R \subseteq S\}.$$

Then, from the standard inclusion-exclusion principle, it follows that

$$\#\{S \in \mathsf{IndSub}(\Phi, k' \to F_\mathbf{G}) \mid R \subseteq S\} = \sum_{J \subseteq R} (-1)^{|J|} \cdot \mathsf{IndSub}(\Phi, k' \to F_\mathbf{G} \setminus J),$$

where $F_\mathbf{G} \setminus J$ is the graph obtained from $F_\mathbf{G}$ by deleting the vertices in J. Thus, an algorithm that makes $2^{|R|} \in O(1)$ oracle calls, each of the form $(F_\mathbf{G} \setminus J, k')$, can compute the value $\#\mathsf{BipIndSub}(\Psi_\Phi, k \to \mathbf{G})$, which gives the sought-for reduction, i.e., the connection between the bipartite property Ψ_Φ and the original graph property Φ.

[11] The twin-free quotient was implicitly used in [30] as well, although in a much less general reduction.

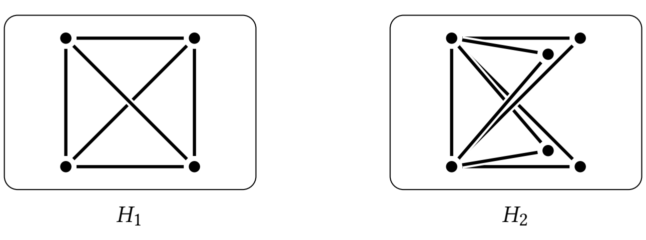

Figure 2: $H_1 = K_4$ and H_2 which can be obtained from K_4 by replacing an edge by $\mathbf{I}_{2,2}$.

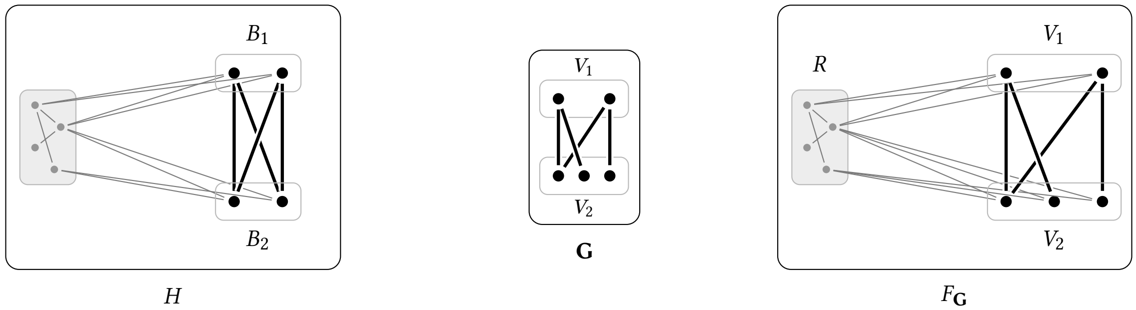

Figure 3: Replacing a pair of adjacent blocks B_1 and B_2 by a bipartite graph G.

There is an additional ingredient which so far we have swept under the rug. In the definition of $F_{\mathbf{G}}$ we replaced two adjacent blocks in the chosen graph H by the graph **G**. It is important that the blocks B_1 and B_2 share an edge. (Since they are blocks, this means that there is a complete set of edges between them.) However, it is possible that Φ contains an independent set as forbidden induced subgraph. It would follow that the graph H in $\Pi(\Phi)$ with edge-minimal $H\!\downarrow$ is an independent set, which would spoil the construction. In this case it helps to consider a closely related property. For each graph G, let $\overline{G}$ be the complement of G. Then we define $\overline{\Phi}$ with $\overline{\Phi}(G) = \Phi(\overline{G})$. It is known that $\overline{\Phi}$ is hereditary if and only if Φ is hereditary. Furthermore, $\#\textsc{IndSub}(\Phi)$ and $\#\textsc{IndSub}(\overline{\Phi})$ are known to be tightly interreducible by parameterised reductions [30]. So, for all our purposes, we are free to work with either one of Φ or $\overline{\Phi}$. By a simple application of Ramsey's theorem, we show that every hereditary property Φ, for which both Φ and $\overline{\Phi}$ have an independent set as forbidden induced subgraph, has to be meagre. Conversely, if Φ is not meagre then at least one of Φ or $\overline{\Phi}$ is a suitable candidate for our construction.

The approach of utilising bipartite properties and implanting bipartite graphs into some fixed graph H that depends on the graph property Φ is not only applicable to hereditary properties. With a slightly different construction we also prove Theorem 3, the classification for properties that are invariant under homomorphic equivalence. In this case it suffices to implant a graph **G** into an edge e in a graph H for which $\Phi(H) \neq \Phi(H - e)$ holds. We omit further details but it is worth to point out that we actually classify a less natural but even more general class of properties. A graph property Φ is *twin-invariant* if, for each pair of graphs H_1 and H_2 that have isomorphic twin-free quotients[12], we have $\Phi(H_1) = \Phi(H_2)$. This criterion covers previously unclassified properties such as "disconnected or bipartite" or "disconnected or triangle-free" but more importantly it is not hard to see that every property that is invariant under homomorphic equivalence is also twin-invariant. Thus, Theorem 3 is a direct consequence of the following more general result.

[12]In the definition we can even get away with only considering graphs whose twin-free quotient contains at least two vertices. This is a technicality which makes the class of covered properties more general and ensures that this result covers, for instance, also the property of being (dis)connected, which was of interest in some of the earlier works [19, 29].

Theorem 6. *Let Φ be a computable twin-invariant graph property. If Φ is meagre then $\#\textsc{IndSub}(\Phi)$ is solvable in polynomial time. Otherwise, $\#\textsc{IndSub}(\Phi)$ is $\#\mathrm{W}[1]$-complete and, assuming ETH, cannot be solved in time $f(k) \cdot |G|^{o(k)}$ for any function f.*

1.3 Further Related Work

The decision version of $\#\textsc{IndSub}(\Phi)$ for hereditary properties Φ was studied and fully classified by Khot and Raman [24] — their results are summarised in Table 1 — and also by Eppstein, Gupta and Havvaei [12] who additionally restricted the problem to hereditary classes of input graphs. Furthemore, if for each k, the property Φ is true for at most one k-vertex graph, the work of Chen, Thurley and Weyer [6] establishes hardness for both, decision and counting, whenever Φ is not meagre.[13]

The complexity of computing an ε-approximation of $\#\textsc{IndSub}(\Phi)$ was investigated by Jerrum and Meeks in a sequence of papers [19–21, 25], and in case of hereditary properties, it was ultimately resolved by Meeks [25] (see Table 1). For more general classes of properties, there are only partial results, and to the best of our knowledge the complexity of approximating $\#\textsc{IndSub}(\Phi)$ is still open for edge-monotone properties. However, there are strong recent meta-theorems such as the k-Hypergraph framework due to Dell, Lapinskas, and Meeks [10] that yield efficient approximation algorithms for $\#\textsc{IndSub}(\Phi)$ by reduction to vertex-coloured decision problems.

Most results on $\#\textsc{IndSub}(\Phi)$ are concerned with hardness of exact counting; we list them chronologically.

- In [19], Jerrum and Meeks proved that $\#\textsc{IndSub}(\Phi)$ always belongs to $\#\mathrm{W}[1]$, given that Φ is computable. Additionally, $\#\mathrm{W}[1]$-hardness was established for the property Φ of being connected.
- In [20], $\#\mathrm{W}[1]$-hardness was proved by the same authors for all properties with low edge-densities. This covers, for example, all non-trivial minor-closed properties.
- In the survey paper of Meeks [25] a $\#\mathrm{W}[1]$-hardness result was established for properties that are closed under the addition of edges, and whose edge-minimal elements have unbounded treewidth.

[13]We remark that our notion of meagre coincides with their notion of meagre in the special case where Φ is true for at most one k-vertex graph for each k, which applies, e.g., to the properties of being a path, a cycle, or a matching.

- In [21], Jerrum and Meeks proved #W[1]-hardness for the property of having an even, or an odd, number of edges.
- In the breakthrough paper of Curticapean, Dell and Marx [7], the principle of complexity monotonicity was introduced, and it was shown that #IndSub(Φ) is always either fixed-parameter tractable or #W[1]-hard, given that Φ is computable. (See Theorem 4.)
- In [29], Roth and Schmitt established #W[1]-hardness and a tight lower bound under ETH for edge-monotone properties that are non-meagre and false on odd cycles.
- In [11], Dörfler, Roth, Schmitt, and Wellnitz introduced the "algebraic approach to hardness" and established #W[1]-hardness for properties that distinguish independent sets from so-called wreath graphs. As a concrete example, their result applies to monotone properties that are non-trivial on bipartite graphs, in which case a tight lower bound under ETH is also achieved.
- Finally, in [30], Roth, Schmitt and Wellnitz established #W[1]-hardness and an *almost* tight conditional lower bound of the form $f(k) \cdot n^{o(k/\sqrt{\log k})}$ for non-meagre monotone properties. More generally, they proved the result for any property with so-called f-vectors of small Hamming weight; we refer the reader to [30, Sections 3 and 4] for a detailed exposition but remark that hereditary properties do not, in general, have f-vectors with small enough Hamming weight for the result in the current paper to be covered by their meta theorem. Additionally, Roth, Schmitt and Wellnitz proved Conjecture 1 for the restricted case of hereditary properties that are defined by a single forbidden induced subgraph.

None of the previous results comes close to resolving Conjecture 1 for all hereditary properties. In particular, since every monotone property is also hereditary, we not only subsume the classification for monotone properties from [30], but we also improve the conditional lower bound from almost tight to tight.

1.4 Open Problems

We conclude our presentation with two open problems and suggestions for further work.

First, the most important open question is whether Conjecture 1 is indeed true for all computable properties. With the case of hereditary properties now being resolved, the other central remaining family of properties with a natural closure condition is the class of all edge-monotone properties. Between the results from [25] and [29], large classes of edge-monotone properties are already covered, but there remains a significant gap towards a complete understanding. For example, none of the existing partial results resolves the complexity of #IndSub(Φ) for the following (slightly artificial) edge-monotone property:

$\Phi(H) = 1$ if and only if H is bipartite or has no apex.[14]

For this reason, we suggest to tackle Conjecture 1 for the case of edge-monotone properties as a concrete next step. Our hope is that the technical framework that we introduce in the work at hand will help to close the remaining gap for edge-monotone properties as well.

[14] An apex is a vertex adjacent to all other vertices.

However, we wish to point out that, in this case, strengthening the intractability part of Conjecture 1 by additionally aiming for tight conditional lower bounds might require hardness assumptions stronger than ETH. The reason for the latter is the existence of artificial "non-dense" edge-monotone properties, such as

$$\Phi(H) = \begin{cases} 1 & \exists \ell : H = I_{a(\ell,\ell)} \\ 0 & \text{otherwise}, \end{cases}$$

where a is the Ackermann function. Observe that Φ is closed under the removal of edges (but not under the removal of vertices). It is easy to establish #W[1]-hardness of #IndSub(Φ) by reducing from the parameterised problem of counting independent sets using standard methods. However, the parameter explodes drastically in this reduction due to the fact that we can only reduce to instances of #IndSub(Φ) in which k is in the image of the Ackermann function. This prevents us from coming even close to a tight ETH-based lower bound. One way to circumvent this problem is to restrict ourselves to graph properties that are *dense* in the sense that the set of ℓ for which Φ is non-trivial on ℓ-vertex graphs is a dense enough subset of the natural numbers; this approach was formalised and used in previous work [11, 29, 30].

Finally, given that this work provides even more evidence for the intractability of #IndSub(Φ), we stress that a relaxation of the problem is unavoidable if efficient algorithms are sought. The obvious and most promising candidate for such a relaxation is to only aim for an approximation of the solution.

As summarised in Table 1, Meeks [25] explicitly and exhaustively identified those hereditary properties Φ for which #IndSub(Φ) admits an FPTRAS. In particular, our main result shows that Meeks' result cannot be strengthened to yield fixed-parameter tractability of exact counting, unless ETH fails.

However, for general (not necessarily hereditary) properties much less is known about the complexity of the approximate counting variant of #IndSub(Φ); again, we refer the reader to the survey of Meeks for a comprehensive overview [25].

ACKNOWLEDGEMENTS

We are very grateful to Johannes Schmitt and Philip Wellnitz for helpful comments on early and recent drafts of this work.

REFERENCES

[1] Noga Alon, Phuong Dao, Iman Hajirasouliha, Fereydoun Hormozdiari, and S. Cenk Sahinalp. 2008. Biomolecular network motif counting and discovery by color coding. *Bioinformatics* 24, 13 (07 2008), i241–i249. https://doi.org/10.1093/bioinformatics/btn163

[2] Vikraman Arvind and Venkatesh Raman. 2002. Approximation Algorithms for Some Parameterized Counting Problems. In *Algorithms and Computation, 13th International Symposium, ISAAC 2002 Vancouver, BC, Canada, November 21-23, 2002, Proceedings.* 453–464. https://doi.org/10.1007/3-540-36136-7_40

[3] Hubie Chen and Stefan Mengel. 2016. Counting Answers to Existential Positive Queries: A Complexity Classification. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2016, San Francisco, CA, USA, June 26 - July 01, 2016.* 315–326. https://doi.org/10.1145/2902251.2902279

[4] Jianer Chen, Benny Chor, Mike Fellows, Xiuzhen Huang, David W. Juedes, Iyad A. Kanj, and Ge Xia. 2005. Tight lower bounds for certain parameterized NP-hard problems. *Inf. Comput.* 201, 2 (2005), 216–231. https://doi.org/10.1016/j.ic.2005.05.001

[5] Jianer Chen, Xiuzhen Huang, Iyad A. Kanj, and Ge Xia. 2006. Strong computational lower bounds via parameterized complexity. *J. Comput. Syst. Sci.* 72, 8 (2006), 1346–1367. https://doi.org/10.1016/j.jcss.2006.04.007

[6] Yijia Chen, Marc Thurley, and Mark Weyer. 2008. Understanding the Complexity of Induced Subgraph Isomorphisms. In *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part I: Tack A: Algorithms, Automata, Complexity, and Games*. 587–596. https://doi.org/10.1007/978-3-540-70575-8_48

[7] Radu Curticapean, Holger Dell, and Dániel Marx. 2017. Homomorphisms are a good basis for counting small subgraphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*. 210–223. https://doi.org/10.1145/3055399.3055502

[8] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. 2015. *Parameterized Algorithms*. Springer. https://doi.org/10.1007/978-3-319-21275-3

[9] Víctor Dalmau and Peter Jonsson. 2004. The complexity of counting homomorphisms seen from the other side. *Theor. Comput. Sci.* 329, 1-3 (2004), 315–323. https://doi.org/10.1016/j.tcs.2004.08.008

[10] Holger Dell, John Lapinskas, and Kitty Meeks. 2020. Approximately counting and sampling small witnesses using a colourful decision oracle. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, Shuchi Chawla (Ed.). SIAM, 2201–2211. https://doi.org/10.1137/1.9781611975994.135

[11] Julian Dörfler, Marc Roth, Johannes Schmitt, and Philip Wellnitz. 2022. Counting Induced Subgraphs: An Algebraic Approach to #W[1]-Hardness. *Algorithmica* 84, 2 (2022), 379–404. https://doi.org/10.1007/s00453-021-00894-9

[12] David Eppstein, Siddharth Gupta, and Elham Havvaei. 2021. Parameterized Complexity of Finding Subgraphs with Hereditary Properties on Hereditary Graph Classes. In *Fundamentals of Computation Theory - 23rd International Symposium, FCT 2021, Athens, Greece, September 12-15, 2021, Proceedings (Lecture Notes in Computer Science, Vol. 12867)*, Evripidis Bampis and Aris Pagourtzis (Eds.). Springer, 217–229. https://doi.org/10.1007/978-3-030-86593-1_15

[13] Jörg Flum and Martin Grohe. 2003. Describing parameterized complexity classes. *Inf. Comput.* 187, 2 (2003), 291–319. https://doi.org/10.1016/S0890-5401(03)00161-5

[14] Jörg Flum and Martin Grohe. 2006. *Parameterized Complexity Theory*. Springer. https://doi.org/10.1007/3-540-29953-X

[15] Jacob Focke and Marc Roth. 2021. Counting Small Induced Subgraphs with Hereditary Properties. *CoRR* abs/2111.02277 (2021). arXiv:2111.02277

[16] Joshua A. Grochow and Manolis Kellis. 2007. Network Motif Discovery Using Subgraph Enumeration and Symmetry-Breaking. In *Research in Computational Molecular Biology*, Terry Speed and Haiyan Huang (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 92–106.

[17] Martin Grohe, Thomas Schwentick, and Luc Segoufin. 2001. When is the evaluation of conjunctive queries tractable?. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing, July 6-8, 2001, Heraklion, Crete, Greece*. 657–666. https://doi.org/10.1145/380752.380867

[18] Russell Impagliazzo and Ramamohan Paturi. 2001. On the Complexity of k-SAT. *J. Comput. Syst. Sci.* 62, 2 (2001), 367–375. https://doi.org/10.1006/jcss.2000.1727

[19] Mark Jerrum and Kitty Meeks. 2015. The parameterised complexity of counting connected subgraphs and graph motifs. *J. Comput. Syst. Sci.* 81, 4 (2015), 702–716. https://doi.org/10.1016/j.jcss.2014.11.015

[20] Mark Jerrum and Kitty Meeks. 2015. Some Hard Families of Parameterized Counting Problems. *TOCT* 7, 3 (2015), 11:1–11:18. https://doi.org/10.1145/2786017

[21] Mark Jerrum and Kitty Meeks. 2017. The parameterised complexity of counting even and odd induced subgraphs. *Combinatorica* 37, 5 (2017), 965–990. https://doi.org/10.1007/s00493-016-3338-5

[22] Pieter W. Kasteleyn. 1961. The statistics of dimers on a lattice: I. The number of dimer arrangements on a quadratic lattice. *Physica* 27, 12 (1961), 1209–1225. https://doi.org/10.1016/0031-8914(61)90063-5

[23] Pieter W. Kasteleyn. 1963. Dimer Statistics and Phase Transitions. *J. Math. Phys.* 4, 2 (1963), 287–293. https://doi.org/10.1063/1.1703953

[24] Subhash Khot and Venkatesh Raman. 2002. Parameterized complexity of finding subgraphs with hereditary properties. *Theor. Comput. Sci.* 289, 2 (2002), 997–1008. https://doi.org/10.1016/S0304-3975(01)00414-5

[25] Kitty Meeks. 2016. The challenges of unbounded treewidth in parameterised subgraph counting problems. *Discrete Applied Mathematics* 198 (2016), 170–194. https://doi.org/10.1016/j.dam.2015.06.019

[26] Ron Milo, Shalev Itzkovitz, Nadav Kashtan, Reuven Levitt, Shai Shen-Orr, Inbal Ayzenshtat, Michal Sheffer, and Uri Alon. 2004. Superfamilies of Evolved and Designed Networks. *Science* 303, 5663 (2004), 1538–1542. https://doi.org/10.1126/science.1089167

[27] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. 2002. Network Motifs: Simple Building Blocks of Complex Networks. *Science* 298, 5594 (2002), 824–827. https://doi.org/10.1126/science.298.5594.824

[28] Jaroslav Nešetřil and Svatopluk Poljak. 1985. On the complexity of the subgraph problem. *Commentationes Mathematicae Universitatis Carolinae* 26, 2 (1985), 415–419.

[29] Marc Roth and Johannes Schmitt. 2020. Counting Induced Subgraphs: A Topological Approach to #W[1]-hardness. *Algorithmica* 82, 8 (2020), 2267–2291. https://doi.org/10.1007/s00453-020-00676-9

[30] Marc Roth, Johannes Schmitt, and Philip Wellnitz. 2020. Counting Small Induced Subgraphs Satisfying Monotone Properties. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*. IEEE, 1356–1367. https://doi.org/10.1109/FOCS46700.2020.00128

[31] Benjamin Schiller, Sven Jager, Kay Hamacher, and Thorsten Strufe. 2015. StreaM - A Stream-Based Algorithm for Counting Motifs in Dynamic Graphs. In *Algorithms for Computational Biology*, Adrian-Horia Dediu, Francisco Hernández-Quiroz, Carlos Martín-Vide, and David A. Rosenblueth (Eds.). Springer International Publishing, Cham, 53–67.

[32] Falk Schreiber and Henning Schwöbbermeyer. 2005. Frequency Concepts and Pattern Detection for the Analysis of Motifs in Networks. In *Transactions on Computational Systems Biology III*, Corrado Priami, Emanuela Merelli, Pablo Gonzalez, and Andrea Omicini (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 89–104.

[33] Harold N. V. Temperley and Michael E. Fisher. 1961. Dimer problem in statistical mechanics-an exact result. *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics* 6, 68 (1961), 1061–1063. https://doi.org/10.1080/14786436108243366