Graph Isomorphism for Graph Classes Characterized by Two Forbidden Induced Subgraphs*

Stefan Kratsch[1] and Pascal Schweitzer[2]

[1] Utrecht University, the Netherlands
s.kratsch@uu.nl
[2] The Australian National University
pascal.schweitzer@anu.edu.au

Abstract. We study the complexity of the Graph Isomorphism problem on graph classes that are characterized by a finite number of forbidden induced subgraphs, focusing mostly on the case of two forbidden subgraphs. We show hardness results and develop techniques for the structural analysis of such graph classes, which applied to the case of two forbidden subgraphs give the following results: A dichotomy into isomorphism complete and polynomial-time solvable graph classes for all but finitely many cases, whenever neither of the forbidden graphs is a clique, a pan, or a complement of these graphs. Further reducing the remaining open cases we show that (with respect to graph isomorphism) forbidding a pan is equivalent to forbidding a clique of size three.

1 Introduction

Given two graphs G_1 and G_2, the Graph Isomorphism problem (GI) asks whether there exists a bijection from the vertices of G_1 to the vertices of G_2 that preserves adjacency. This paper studies the complexity of GI on graph classes that are characterized by a finite number of forbidden induced subgraphs, focusing mostly on the case of two forbidden subgraphs. For a set of graphs $\{H_1, \dots, H_k\}$ we let $(H_1, \dots, H_k)$-free denote the class of graphs G that do not contain any H_i as an induced subgraph.

As a first example, consider the class of graphs containing neither a clique K_s on s vertices, nor an independent set I_t on t vertices. Ramsey's Theorem [19] states that the number of vertices in such graphs is bounded by a function $f(s,t)$. Thus the classes (K_s, I_t)-free are finite and Graph Isomorphism is trivial on them. All other combinations of two forbidden subgraphs give graph classes of infinite size, since they contain infinitely many cliques or independent sets.

As a second example, consider the graphs containing no clique K_s on s vertices and no star $K_{1,t}$ (i.e., an independent set of size t with added universal vertex adjacent to every other vertex). On the one hand this class contains all

* In this version some proofs are omitted. For these the reader is referred to [12].

M.C. Golumbic et al. (Eds.): WG 2012, LNCS 7551, pp. 34–45, 2012.

graphs of maximum degree less than $\min\{s-1,t\}$, on the other hand, all graphs in $(K_s, K_{1,t})$-free have bounded degree: Indeed, if the degree of a vertex is sufficiently large, its neighborhood must contain a clique of size s or an independent set of size t by Ramsey's Theorem [19], leading to one of the two forbidden subgraphs. Thus, using Luks' algorithm [16] that solves Graph Isomorphism on graphs of bounded degree in polynomial time, isomorphism of $(K_s, K_{1,t})$-free graphs can also be decided in polynomial time.

To systematically study Graph Isomorphism on graph classes characterized by forbidden subgraphs, we ask: *Given a set of graphs* $\{H_1,\ldots,H_k\}$*, what is the complexity of Graph Isomorphism on the class of* $(H_1,\ldots,H_k)$-free *graphs?*

Related Work. The Graph Isomorphism problem is contained in the complexity class NP, since the adjacency preserving bijection (the isomorphism) can be checked in polynomial time. No polynomial-time algorithm is known and it is known that Graph Isomorphism is not NP-complete unless the polynomial hierarchy collapses [5]. More strongly, Graph Isomorphism is in the low hierarchy [21]. This has led to the definition of the complexity class of problems polynomially equivalent to Graph Isomorphism, the so-called GI-complete problems. There is a vast literature on the Graph Isomorphism Problem; for a general overview see [22] or [10], for results on its parameterized complexity see [13].

A question analogous to ours, asking about Graph Isomorphism on any class of $(H_1,\ldots,H_k)$-minor-free graphs, is answered completely by the fact that Graph Isomorphism is polynomially solvable on any non-trivial minor closed class [18]. Recently, the corresponding statement for topological minor free classes was also shown [8]. For the less restrictive family of *hereditary* classes, only closed under vertex deletion (i.e., classes $\mathcal{H}$-free for a possibly infinite set of graphs $\mathcal{H}$), both GI-complete and tractable cases are known: Graph Isomorphism is GI-complete on split graphs, comparability graphs, and strongly chordal graphs [23]. Graph Isomorphism is known to be polynomially solvable for circle graphs and circular-arc graphs [9], interval graphs [2,15], distance hereditary graphs [17], and graphs of bounded degree [16]. For various subclasses of these polynomially solvable cases results with finer complexity analysis are available, but of course the polynomial-time solvability for these subclasses follows already from polynomial-time solvability of the mentioned larger classes. Further results, in particular on GI-completeness, can be found in [4].

Concerning our question, for one forbidden subgraph, the answer, given by Colbourn and Colbourn, can be found in a paper by Booth and Colbourn [4]: If the forbidden induced subgraph H_1 is an induced subgraph of the path P_4 on four vertices, denoted by $H_1 \leq P_4$, then Graph Isomorphism is polynomial on H_1-free graphs, otherwise it is GI-complete.

Apart from the isomorphism problem, other studies aiming at dichotomy results for algorithmic problems on graph classes characterized by two forbidden subgraphs consider the chromatic number [11] and dominating sets [14].

Main Result. Let a graph be *basic* if it is an independent set, a clique, a $P_3 \dot\cup K_1$, or the complement of a $P_3 \dot\cup K_1$ (also called pan). If neither H_1 nor H_2 is basic,

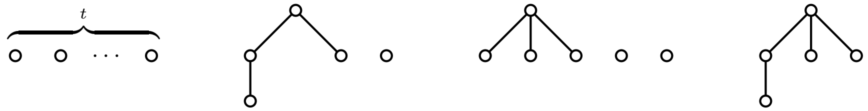

Fig. 1. The independent sets, the paths of length 3 with independent vertex, the claw with two independent vertices, and the fork, obtained by subdividing an edge in a claw

then we obtain a classification of (H_1, H_2)-free classes into polynomial and GI-complete cases, for all but a small finite number of classes. Theorem 1 justifies the terminology basic by showing that in our context forbidding a basic graph is equivalent to forbidding a complete graph. However, the case of forbidding a clique (alongside a second graph) appears to be structurally different and for a complete classification further new techniques are required.

Technical Contribution. Our main technical contribution lies in establishing tractability of Graph Isomorphism on four types of (H_1, H_2)-free classes (Theorem 4 in Section 4): A structural analysis of the classes enables reductions to the polynomially-solvable case of bounded color valence [1]. This reduction appears necessary since the polynomially-solvable classes of Theorem 4 encompass all classes of graphs of bounded degree, and for these Luks' group-theoretic approach [16] (implicit in [1]) is the only known polynomial-time technique. At the core of the proof of Theorem 4 lie individualization-refinement techniques and recursive structural analysis to allow for a reduction to the bounded color valence case.

However, to put these results in context and obtain the mentioned classification, we have to refine several known results for GI-completeness on bipartite, split, and line graphs (Section 3). In particular, we arrive at a set of four graph properties, which we call *split conditions*, such that Graph Isomorphism remains complete on any class (H_1, H_2)-free unless each property is true for at least one of the two forbidden subgraphs.

Based on this characterization we can state our results in more detail: If on the one hand neither of the two forbidden subgraphs H_1 and H_2 exhibits all four split conditions, then we have a dichotomy of GI on (H_1, H_2)-free classes into polynomial and GI-complete classes; the polynomial cases are due to Theorem 4 (see Section 4) as well as tractability on cographs (i.e., P_4-free graphs) [7], GI-completeness follows by using both known results as well as our strengthened reductions (see Section 3). Suppose on the other hand H_1 and H_2 are both not basic and H_1 simultaneously fulfills all four split conditions, then our hardness and tractability results resolve all but a finite number of cases (i.e., each case is one concrete class (H_1, H_2)-free), as showed in Theorem 6 (see Section 5). For these cases Figure 1 shows the relevant maximal graphs that adhere to all four split conditions.

2 Preliminaries

We write $H \leq G$ if the graph G contains a graph H as an induced subgraph. A graph G is H-free if $H \not\leq G$. It is $(H_1, \ldots, H_k)$-free, if it is H_i-free for all i. A graph class $\mathcal{C}$ is H-free (resp. $(H_1, \ldots, H_k)$-free) if this is true for all $G \in \mathcal{C}$. A graph class $\mathcal{C}$ is *hereditary* if it is closed under taking induced subgraphs. The class $(H_1, \ldots, H_k)$-free is the class of all $(H_1, \ldots, H_k)$-free graphs; each class $(H_1, \ldots, H_k)$-free is hereditary.

By I_t, K_t, P_t, and C_t we denote the independent set, the clique, the path, and the cycle on t vertices; $K_{1,t}$ is the claw with t leaves. By $H \mathbin{\dot\cup} H'$ we denote the disjoint union of H and H'; we use tK_2 for the disjoint union of t graphs K_2. By $\overline{G}$ we denote the (edge) complement of G. The graph $\overline{K_2 \mathbin{\dot\cup} I_2}$, i.e., the same as a K_4 minus one edge, is called diamond.

We recall that GI-completeness is inherited by superclasses while polynomial-solvability of Graph Isomorphism is inherited by subclasses. Also recall that Graph Isomorphism on a class $\mathcal{C}$ is exactly as hard as on $\overline{\mathcal{C}}$, the class of complements of graphs in $\mathcal{C}$. Note that any H-free graph is also H'-free if $H \leq H'$.

Proposition 1. *Let H_1, H_2 be graphs and let $\mathcal{C}$ be any hereditary graph class.*

1. $\overline{(H_1, H_2)\text{-free}} = (\overline{H_1}, \overline{H_2})\text{-free}$.
2. $(H_1, H_2)\text{-free} \subseteq (H_1', H_2')\text{-free}$ *for any H_1', H_2' with $H_1 \leq H_1'$ and $H_2 \leq H_2'$.*
3. $H_1, H_2 \notin \mathcal{C}$ *implies* $\mathcal{C} \subseteq (H_1, H_2)\text{-free}$.

Definition 1. *The* pan *is the graph $\overline{P_3 \mathbin{\dot\cup} K_1}$, i.e., a vertex and triangle joined by one edge. A graph is* basic*, if it is an independent set, a complete graph, the graph $P_3 \mathbin{\dot\cup} K_1$, or the pan.*

We now show that in the context of the isomorphism problem excluding a basic graph is equivalent to excluding a complete graph or an independent set.

Lemma 1. *Let G be a graph that contains P_4 as an induced subgraph.*

1. *If G is co-connected then it contains I_3 if and only if it contains $P_3 \mathbin{\dot\cup} K_1$.*
2. *If G is connected then it contains K_3 if and only if it contains a pan.*

Proof. By complementarity it suffices to prove Part 2. Fix a P_4 in the graph. Containment of a pan trivially implies containment of a triangle. For the converse, it can be easily verified that there is a pan, if some some triangle contains at least two vertices of the P_4. Else, if a triangle contains one vertex p of the P_4, we can add a vertex of the P_4 adjacent to p to the triangle, obtaining a pan. Else (i.e., if no triangle is incident with the P_4) consider the triangle closest to the P_4. Due to connectivity, there is a vertex that is adjacent to some vertex of the triangle and closer to the P_4. If this vertex is adjacent to exactly one vertex of the triangle, a pan arises. Otherwise we find a closer triangle, which contradicts our initial choice. □

Theorem 1. *Graph Isomorphism on a class $\mathcal{C}$ of K_3-free graphs is polynomial time equivalent to GI on the subclass of $\mathcal{C}$ that contains all pan-free graphs of $\mathcal{C}$.*

Proof. Since Graph Isomorphism for P_4-free graphs (so-called cographs) is solvable in polynomial time [7], the theorem follows from Lemma 1 and the fact that graph isomorphism can be solved by comparing connected components. □

3 Hardness Results

Our standard method to show GI-completeness for Graph Isomorphism on some graph class $\mathcal{H}$ works by reducing the isomorphism problem of a class $\mathcal{H}'$ for which Graph Isomorphism is known to be GI-complete to a subclass of $\mathcal{H}$. For this we require a mapping $\pi\colon\ \mathcal{H}' \to \pi(\mathcal{H}') \subseteq \mathcal{H}$ which is computable in polynomial time and for which the images of two graphs are isomorphic if and only if the two original graphs are. We call such a mapping π a *GI-reduction.* To show hardness for a class (H_1, H_2)-free it suffices to provide a GI-reduction π for which no graph $G \in \pi(\mathcal{H}')$ contains H_1 or H_2 as an induced subgraph, implying that $\pi(\mathcal{H}') \subseteq (H_1, H_2)$-free.

Our first reductions prove hardness results for bipartite graphs, split graphs, and line graphs. However, the (previously known) GI-completeness for these particular graph classes is not sufficient. We require hardness for more specific subclasses avoiding specific small graphs. Subsequently, using a more involved reduction, we show that isomorphism of $(P_4 \dot\cup K_1, K_4)$-free graphs is GI-complete.

3.1 Bipartite Graphs

A straightforward GI-reduction consists of subdividing each edge of a graph by a new vertex. Since the obtained graphs are bipartite, this proves that Graph Isomorphism remains GI-complete on bipartite graphs. This also implies that Graph Isomorphism remains GI-complete on (H_1, H_2)-free graphs unless one of the graphs is bipartite, since the class (H_1, H_2)-free contains all bipartite graphs if neither H_1 nor H_2 is bipartite. Let us observe however, that we can draw stronger conclusions namely that Graph Isomorphism remains GI-complete on connected bipartite graphs without induced cycles of length 4, for which the vertices in one of the partition classes have degree two. The following definition allows us to make a first structural observation for the graphs H_1, H_2:

Definition 2. *A* path-star *is a subdivision of the t-claw $K_{1,t}$, for some $t \in \mathbb{N}$.*

Lemma 2. *If neither H_1 nor H_2 is a disjoint union of path-stars, then Graph Isomorphism on the class (H_1, H_2)-free is GI-complete.*

Proof. If a graph is not a disjoint union of path-stars, then it either contains two vertices of degree at least 3 which are in the same connected component, or it contains a cycle. We use that two graphs G_1 and G_2 are isomorphic, if and only if the graphs obtained by subdividing each edge in G_1 and G_2 respectively are isomorphic. For any integer c there is an integer c' such that a graph that has been subdivided c' times neither contains a cycle of length at most c nor two vertices of degree at least three which are at a distance of at most c apart.

Thus, with a finite number of subdivision steps, we can reduce Graph Isomorphism on general graphs to isomorphism on (H_1, H_2)-free graphs. □

Using Part 1 of Proposition 1, we conclude that unless Graph Isomorphism is GI-complete on (H_1, H_2)-free, one of the graphs H_1, H_2 is a forest and one of the graphs is a co-forest.

Lemma 3. *A graph H and its complement $\overline{H}$ are forests, if and only if $H \leq P_4$.*

Proof. For the if part it suffices to observe that P_4 is a self-complementary forest. The only if part is true for graphs of size at most 4. A forest on $n \geq 5$ vertices has at most $n-1$ edges. Since $2(n-1) < \binom{n}{2}$ for $n \geq 5$ the statement follows. □

Graph Isomorphism for P_4-free graphs (cographs) is in P [7]. With the previous two lemmas one can conclude that it remains GI-complete on H-free graphs when H is not an induced subgraph of the P_4; this gives a simple dichotomy for the case of a single forbidden induced subgraph.

Theorem 2 (see [4]). *Let H be a graph. Graph Isomorphism on H-free graphs is in* P*, if $H \leq P_4$. GI on H-free graphs is GI-complete, if $H \not\leq P_4$.*

In the following, we focus on graph classes characterized by two forbidden induced subgraphs. Since isomorphism of P_4-free graphs is in P, we assume from now on that $H_1 \not\leq P_4$ and $H_2 \not\leq P_4$. Due to Lemmas 2 and 3 and Part 1 of Proposition 1 we may further assume that H_1 is a disjoint union of path-stars and H_2 is the complement of a disjoint union of path-stars.

Being forests, unions of path-stars are bipartite. Since bipartite graphs play a repeated role, we introduce some terminology: For a bipartite graph G, which has been partitioned into two classes , the *bipartite complement* is the graph obtained by replacing all edges that run between vertices from different partition classes by non-edges and vice versa. (Note that the bipartite complement for unpartitioned bipartite graphs is only well defined if the graph is connected.) A *crossing co-cycle* is a set of vertices that form a cycle in the bipartite complement.

Lemma 4. *Isomorphism of graphs that are (H_1, H_2)-free is GI-complete unless H_1 or H_2 can be partitioned as a bipartite graph without crossing co-cycle.*

Proof. Graph Isomorphism is GI-complete on connected graphs. By repeatedly subdividing a connected graph we produce a bipartite graph with an arbitrarily high girth. If at least three subdivisions have been performed on a non-trivial graph, its bipartite complement is connected. Thus taking the bipartite complement of such graphs is a GI-reduction (the bipartite complement of the bipartite complement is the original graph), and we obtain the lemma. □

Lemma 5. *For each $G \in \{3K_2, 2K_2 \dot\cup I_2, P_4 \dot\cup I_2\}$, Graph Isomorphism on the class on the class of (H_1, H_2)-free graphs is GI-complete unless one of the graphs H_1 and H_2 is bipartite and does not contain the graph G.*

Proof. Using Lemma 4 this follows, since none of the graphs $3K_2$, $2K_2 \dot\cup I_2$, and $P_4 \dot\cup I_2$ can be partitioned as bipartite graph without crossing co-cycle. □

3.2 Split Graphs

We now turn our attention to split graphs. A *split graph* is a graph whose vertices can be partitioned into an independent set and a clique. Recall that the split graphs are exactly the $(2K_2, C_4, C_5)$-free graphs. The reduction that subdivides each edge and connects all newly introduced vertices produces a split graph, and thus proves GI-completeness on that class. As in the previous section we are able to draw further conclusions about the obtained graphs.

Definition 3. *Let G be a split graph that has been partitioned into a clique K and an independent set I. We say that the partition is*

1. *of type 1, if the vertices in class K have at most 2 neighbors in class I.*
2. *of type 2, if the vertices in class K have at most 2 non-neighbors in class I.*
3. *of type 3, if the vertices in class I have at most 2 neighbors in class K.*
4. *of type 4, if the vertices in class I have at most 2 non-neighbors in class K.*

An unpartitioned graph G is said to fulfill split graph condition i (with $i \in \{1, 2, 3, 4\}$), if there is a split partition of the graph that is of type i.

Lemma 6. *For any $i \in \{1, 2, 3, 4\}$, Graph Isomorphism on the class of graphs that fulfill split graph condition i is GI-complete.*

For the proof, we refer the reader to [12]. Since the class of graphs which fulfill condition i is closed under taking induced subgraphs, the lemma implies that Graph Isomorphism on the class (H_1, H_2)-free is GI-complete unless for all $i \in \{1, 2, 3, 4\}$ one of the graphs H_1 or H_2 fulfills split graph condition i.

3.3 Line Graphs

The next graph class we consider is the class of line graphs. The line graph of a graph $G = (V, E)$ is the graph $L(G) = (E, E')$, in which two vertices are adjacent, if they represent two incident edges in the graph G. The line graph of a graph G encodes the isomorphism type of a graph G.

Lemma 7 ([24]). *Let G_1, G_2 be connected graphs such that neither is a triangle. Then G_1 and G_2 are isomorphic if and only if their line graphs are.*

The class of line graphs has a characterization by 9 forbidden subgraphs [3]. However, we reduce to a subclass characterized by three forbidden subgraphs.

Lemma 8. *Line graphs of graphs of girth at least 5 contain no $K_{1,3}$, no C_4, and no diamond.*

Proof. A claw $K_{1,3}$ in a line graph $L(G)$ would correspond to three edges in G that each share an endpoint with a fourth edge; then two of the three edges must share an endpoint (forcing an additional edge in $L(G)$). A C_4 in $L(G)$ corresponds to a C_4 in G. Finally, if three edges of a triangle-free graph G pairwise share an endpoint, then they all three share the same endpoint. A fourth edge can therefore not share an endpoint with exactly two of the edges without forming triangle in G, i.e., there can be no diamond in $L(G)$. □

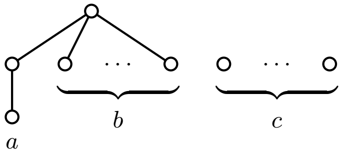

Fig. 2. A claw with a subdivision and added isolated vertices. Every split path-star is a subgraph of such a graph. They are denoted by $H(a, b, c)$ where a is the number of subdivided edges, $a + b$ the degree of the claw, and c the number of isolated vertices.

Since there is essentially a one-to-one correspondence between graphs and their line graphs, and Graph Isomorphism is GI-complete on triangle-free graphs (since $K_3 \not\leq P_4$), it is also GI-complete on line graphs of triangle-free graphs.

Lemma 9. *Graph Isomorphism is Graph Isomorphism-complete on the class of* $(diamond, claw, C_4)$-free *graphs.*

Proof. Since Graph Isomorphism is GI-complete on connected graphs of girth at least 5, and since graphs of girth at least 5 do not contain triangles, the lemma follows from Lemmas 7 and 8. □

3.4 A Reduction to $(P_4 \dot\cup K_1, K_4)$-Free Graphs

There is a reduction that reduces the class of all graphs to the $(P_4 \dot\cup K_1, K_4)$-free graphs. For an explicit description and correctness proof of the reduction, we refer the reader to [12]. Our reduction generalizes the reduction to bipartite graphs, making a replacement of the edges while additionally connecting some of the so-created independent sets. In this sense the reduction (as many other established reductions) is part of a larger scheme of GI-reductions, which use finitely many independent sets, cliques, and relationships between them to encode graphs. We obtain the following theorem.

Theorem 3. *Graph Isomorphism is GI-complete on* $(P_4 \dot\cup K_1, K_4)$-free *graphs.*

4 Structural Results and Polynomially-Solvable Cases

We have previously seen that Graph Isomorphism is GI-complete on graphs that are (H_1, H_2)-free unless each of the four split conditions is fulfilled by one of the two forbidden graphs (Lemma 6). This gives rise to two fundamental cases, namely whether or not one of the two graphs simultaneously fulfills all split conditions. In this section we address the case that neither graph fulfills all split conditions simultaneously. Amongst other conclusions this implies that both graphs must be split graphs or Graph Isomorphism will remain GI-complete. Recall also that one graph must be a path-star while the other must be the complement of a path-star or isomorphism of (H_1, H_2)-free graphs will be GI-complete (w.l.o.g. neither of the graphs is an induced subgraph of P_4 thus only H_i or $\overline{H}_i$ can be a forest or path-star). Using the results of the previous section we are able to fully characterize this case into classes with either polynomial or GI-complete isomorphism problems.

Without loss of generality we take H_1 to be a union of path-stars and H_2 to be the complement of a union of path-stars. We analyze first H_1; the graph H_2 must be a complement of the possible graphs we obtain. Since H_1 is split, i.e., $(2K_2, C_4, C_5)$-free, it is a $2K_2$-free path-star. Therefore it contains at most one non-trivial component and no induced path P_5 on five vertices. Thus if H_1 has a vertex v of degree three or larger, then at most one induced path of length two is emanating from v. Together these observations show that H_1 is an induced subgraph of the type of graph depicted in Figure 2.

We denote by $H(a, b, c)$ the graph that is depicted in Figure 2, with $a \in \{0, 1\}$, $b \in \mathbb{N}$, and $c \in \mathbb{N}$. We require that if $a = 1$ then $b > 0$, otherwise (if $a = 1$ and $b = 0$) we can reinterpret the graph with $b = 2$ and $a = 0$ (thus, $a = 1$ iff the graph contains a P_4). We also require $a + b \geq 1$ since the independent set and the clique fulfill all split conditions. Observe that any induced subgraph of some $H(a, b, c)$ is isomorphic to $H(a', b', c')$ for some values of a', b', c' (it suffices to consider the induced subgraphs of the claw with one subdivision).

We will argue that under these restrictions we may focus on the case that $a = a' = 0$, since Graph Isomorphism remains GI-complete otherwise.

Lemma 10. *Let $H_1 = H(a, b, c)$ and $H_2 = \overline{H(a', b', c')}$ such that neither graph fulfills all split conditions and such that $a + a' \geq 1$. Then GI remains GI-complete on (H_1, H_2)-free graphs.*

For the remaining discussion we may thus assume that $a = 0$ and $a' = 0$.

Theorem 4. *Isomorphism of $(H(0, b, c), \overline{H(0, b', c')})$-free graphs is in P when:*

1. *$b = 0$ or $b' = 0$ (i.e., one of the graphs is a clique or an independent set),*
2. *$c, c' \leq 1$ and $b, b' \geq 1$,*
3. *$c, c' \geq 2$ and $b, b' \in \{1, 2\}$,*
4. *$(c \geq 2,\ c' \leq 1,\ b \geq 1,\ b' \in \{1, 2\})$, or $(c' \geq 2,\ c \leq 1,\ b' \geq 1,\ b \in \{1, 2\})$.*

In all other cases it is GI-complete.

To prove the theorem we use vertex-colorings of the input graphs. (In the context of Graph Isomorphism the vertex colorings are not assumed to be proper). We say that a vertex-colored graph has *bounded color valences*, if there is a constant D, such that for every color class C every vertex v (possibly in C) has at most D neighbors or at most D non-neighbors in C. In a graph without $H(0, b, c)$ and $\overline{H(0, b', c')}$, bounded color valence within color classes implies bounded color valence overall. For a proof of this, we refer the reader to [12]. Bounding the color valence one can reduce the isomorphism problem to that of graphs of bounded degree.

Theorem 5 (Babai, Luks [1]). *Graph Isomorphism for colored graphs of bounded color valence is solvable in polynomial time.*

To prove Theorem 4 we distinguish cases according to the numbers c and c' in the forbidden subgraphs $H(0, b, c)$ and $\overline{H(0, b', c')}$.

Proof (General proof strategy for Theorem 4.). For the full proof of Theorem 4, we refer the reader to [12] and here, instead, provide a high level description of the general proof-strategy. When proving each of the four cases our strategy is as follows: The starting observation is that a colored $(H(0,b,c), \overline{H(0,b',c')})$-free graph, which has bounded degree or bounded co-degree within each color class, also has bounded color valence between different color classes. This enables the use of Theorem 5. Thus, we now intend to find a canonical (in particular isomorphism-respecting) way of coloring both input graphs, so that the color classes have bounded degree or bounded co-degree. We employ two methods to pick color classes, both of them ensure that the coloring preserves isomorphism. Either we choose the colors of the vertices by properties of the vertices that can be computed in polynomial time. Or we guess an ordered set of vertices of constant size, color the vertices in this set with singleton colors, and then color the remaining vertices according to their adjacencies to the vertices in the ordered set. Guessing a constant number k of vertices increases the running time by a factor of n^k, and can therefore be performed in polynomial time. The second coloring operation is typically referred to as individualization.

In Case 1 we individualize one vertex, and use induction to obtain a canonical coloring with the desired properties. In Case 2 we individualize one vertex, and use a combinatorial argument to show that this gives a canonical coloring with the desired properties. In Case 4, using Lemma 1, we reduce the problem to $(H(0,b,c), K_3)$-free graphs, and then apply induction on c to obtain the canonical coloring. Case 3 is the most interesting one (and rather involved). In this case, by individualizing a finite number of vertices, we can obtain a colored graph, in which each of the color classes is a cluster, or a co-cluster graph. (A cluster is a P_3-free graph or equivalently a disjoint union of cliques.) For our purpose this is not sufficient, as for example a cluster graph can have vertices that simultaneously have large degree and large co-degree. We call a cluster d-diverse if it contains at least d disjoint cliques of size at least d. A d-diverse co-cluster is the complement of a d-diverse cluster. We show that for large d a $(H(0,2,c), \overline{H(0,2,c')})$-free graph cannot contain a d-diverse cluster and a d-diverse co-cluster at the same time. With this (possibly taking complements) our situation simplifies to the case where there is one color class A that is a cluster, and all other color classes are of bounded degree or bounded co-degree. After splitting off a bounded number of cliques from A, we can show that for each of the remaining cliques there is only a bounded number of types by which the vertices are connected to the vertices outside the cluster. Using this we replace the cluster by a bounded number of representatives, one for each type, color-encoding the number of vertices of each type. This leaves a graph with bounded color valence and enables us to apply Theorem 5. □

5 The Remaining Cases

In the previous section we investigated the case when neither of the two forbidden induced subgraphs fulfills all split graph conditions. We now consider the case, where one of the two graphs fulfills all split graph conditions. W.l.o.g. we let H_1 be this graph and require that H_1 is a disjoint union of path-stars (otherwise we take complements); there are only few choices for H_1. For the proofs of Lemma 11 and Theorem 6, we refer the reader to [12].

Lemma 11. *If a graph G is a union of path-stars and fulfills all split graph conditions, then it is an induced subgraph of one of the following graphs (depicted in Figure 1): An independent set,* $P_4 \dot\cup K_1$, $K_{1,3} \dot\cup I_2$, *or the fork.*

Theorem 6. *Suppose* H_1 *is a nonbasic disjoint union of path-stars and fulfills all split graph conditions. If* H_2 *has more than 7 vertices, then an application of one of Lemmas 2, 3, 5, or 9, or Theorems 2, 3, or 4 determines that* $(H_1 \dot\cup H_2)$-free *is GI-complete or polynomial-time solvable. More strongly, this can be concluded unless* H_1 *is one of the graphs* $\{P_4 \dot\cup K_1, K_2 \dot\cup I_2, P_3 \dot\cup I_2\}$ *and* $\overline{H_2}$ *has at most 7 vertices and is a disjoint union of at most 3 paths.*

6 Conclusion

In order to initiate a systematic study of the Graph Isomorphism Problem on hereditary graph classes we considered graph classes characterized by two forbidden induced subgraphs. We presented an almost complete characterization of the case that neither of the two forbidden subgraphs is basic into GI-complete and polynomial cases, leaving only few pairs of forbidden subgraphs. Theorem 4 constitutes the main technical contribution towards this result. Together with the tractability of P_4-free graphs (Theorem 2, [7]) it establishes the polynomially solvable cases. On the other hand suppose H_1 and H_2 are non-basic and (H_1, H_2)-free is not a polynomial-time solvable case of Theorems 2 or 4. Then, Graph Isomorphism on the class of (H_1, H_2)-free graphs is GI-complete, unless for H_1 and H_2, or for $\overline{H_1}$ and $\overline{H_2}$, one of the graphs is in $\{P_4 \dot\cup K_1, K_2 \dot\cup I_2, P_3 \dot\cup I_2\}$, and the other graph has at most 7 vertices and is the complement of a union of at most 3 paths.

Several further cases, e.g., all cases involving the $\overline{P_6}$ or the $\overline{P_7}$, can be excluded by variants of the reduction used for Theorem 3. Of the remaining cases, in a preprint, Rao [20] resolves positively the case $(P_4 \dot\cup K_1, \overline{P_4 \dot\cup K_1})$-free and its subclasses; similar (modular) decomposition techniques appear to apply to other cases as well. Several of the remaining cases are classes of bounded clique-width [6], which could indicate their tractability.

For the case in which one of the forbidden graphs is basic, our reductions and our polynomial-time algorithms are still applicable and resolve a large portion of the cases. However, as mentioned in the introduction, complete resolution appears to require new techniques. Future steps for studying the hereditary graph classes include the resolution of the remaining cases and analysis of graph classes characterized by more than two forbidden subgraphs.

References

1. Babai, L., Luks, E.M.: Canonical labeling of graphs. In: STOC 1983, pp. 171–183 (1983)

2. Babel, L., Ponomarenko, I.N., Tinhofer, G.: The isomorphism problem for directed path graphs and for rooted directed path graphs. Journal of Algorithms 21(3), 542–564 (1996)
3. Beineke, L.W.: Characterizations of derived graphs. Journal of Combinatorial Theory, Series B 9(2), 129–135 (1970)
4. Booth, K.S., Colbourn, C.J.: Problems polynomially equivalent to graph isomorphism. Technical Report CS-77-04, Comp. Sci. Dep., Univ. Waterloo (1979)
5. Boppana, R.B., Hastad, J., Zachos, S.: Does co-NP have short interactive proofs? Information Processing Letters 25, 127–132 (1987)
6. Brandstädt, A., Dragan, F.F., Le, H., Mosca, R.: New graph classes of bounded clique-width. Theory of Computing Systems 38, 623–645 (2005)
7. Corneil, D.G., Lerchs, H., Stewart Burlingham, L.: Complement reducible graphs. Discrete Applied Mathematics 3(3), 163–174 (1981)
8. Grohe, M., Marx, D.: Structure theorem and isomorphism test for graphs with excluded topological subgraphs. In: STOC, pp. 173–192 (2012)
9. Hsu, W.L.: $O(m \cdot n)$ algorithms for the recognition and isomorphism problems on circular-arc graphs. SIAM Journal on Computing 24(3), 411–439 (1995)
10. Köbler, J., Schöning, U., Torán, J.: The graph isomorphism problem: its structural complexity. Birkhäuser Verlag, Basel (1993)
11. Král, D., Kratochvíl, J., Tuza, Z., Woeginger, G.J.: Complexity of Coloring Graphs without Forbidden Induced Subgraphs. In: Brandstädt, A., Van Le, B. (eds.) WG 2001. LNCS, vol. 2204, pp. 254–262. Springer, Heidelberg (2001)
12. Kratsch, S., Schweitzer, P.: Full version of paper. arXiv:1208.0142 [cs.DS]
13. Kratsch, S., Schweitzer, P.: Isomorphism for Graphs of Bounded Feedback Vertex Set Number. In: Kaplan, H. (ed.) SWAT 2010. LNCS, vol. 6139, pp. 81–92. Springer, Heidelberg (2010)
14. Lozin, V.V.: A decidability result for the dominating set problem. Theoretical Computer Science 411(44–46), 4023–4027 (2010)
15. Lueker, G.S., Booth, K.S.: A linear time algorithm for deciding interval graph isomorphism. Journal of the ACM 26(2), 183–195 (1979)
16. Luks, E.M.: Isomorphism of graphs of bounded valence can be tested in polynomial time. Journal of Computer and System Sciences 25(1), 42–65 (1982)
17. Nakano, S., Uehara, R., Uno, T.: A new approach to graph recognition and applications to distance-hereditary graphs. Journal of Computer Science and Technology 24, 517–533 (2009)
18. Ponomarenko, I.N.: The isomorphism problem for classes of graphs closed under contraction. Journal of Mathematical Sciences 55(2), 1621–1643 (1991)
19. Ramsey, F.P.: On a problem of formal logic. Proceedings of the London Mathematical Society s2-30(1), 264–286 (1930)
20. Rao, M.: Decomposition of (gem,co-gem)-free graphs (unpublished), `http://www.labri.fr/perso/rao/publi/decompgemcogem.ps`
21. Schöning, U.: Graph isomorphism is in the low hierarchy. Journal of Computer and System Sciences 37(3), 312–323 (1988)
22. Schweitzer, P.: Problems of unknown complexity: Graph isomorphism and Ramsey theoretic numbers. PhD thesis, Universität des Saarlandes, Germany (2009)
23. Uehara, R., Toda, S., Nagoya, T.: Graph isomorphism completeness for chordal bipartite graphs and strongly chordal graphs. Discrete Applied Mathematics 145(3), 479–482 (2005)
24. Whitney, H.: Congruent graphs and the connectivity of graphs. American Journal of Mathematics 54(1), 150–168 (1932)