

The Complexity of Boolean-Weighted Graph Isomorphism

Francesco Cristiano*

Abstract

A Boolean-weighted graph is a finite graph whose edges carry DNF formulas over a common variable set. This paper studies two isomorphism problems on such graphs, distinguished by whether the per-edge condition requires the matched edge labels to be syntactically DNF-isomorphic (BWG-ISO) or syntactically DNF non-isomorphic (BWG-NI), each over a vertex bijection that preserves edges. The two sit at opposite ends of one phenomenon: imposing local syntactic DNF-isomorphism on every edge does not raise complexity above graph isomorphism, whereas replacing it by local non-isomorphism makes the problem NP-hard, the hardness being driven by the interaction of the per-edge non-isomorphism requirement with host automorphisms rather than by the complexity of the labels. For the isomorphism requirement, $\text{BWG-ISO} \equiv_m^P \text{GI}$: although syntactic DNF isomorphism is itself GI-complete [ACL12], attaching such a label to every edge leaves the global problem exactly at GI. For the non-isomorphism requirement, the per-edge test is a graph-non-isomorphism query, placing BWG-NI in NP^{GI} ; the paper shows it is GI-hard and GNI-hard and not coNP-hard unless $\text{PH} = \Sigma_2^P$, and proves it is NP-hard. The hardness already holds when labels are restricted to monotone DNFs, equivalently $\{0,1\}$ -matrices, so it stems from the interaction with host automorphisms, not from rich Boolean structure. This NP-hardness is unconditional; the hardest NP set known to $\leq_m^P$ -reduce to FI is GI [AT00]. BWG-NI is thereby a natural inhabitant of NP^{GI} : with the NP-hardness, NP^{GI} membership places it outside P^{GI} , coNP^{GI} , and Low_2 unless $\text{PH} = \Sigma_2^P$, making it a natural candidate to separate NP^{GI} from P^{GI} when the hierarchy is infinite. The second-level collapse that would be forced by its Σ_2^P -completeness is one level below the third-level collapse Agrawal and Thierauf obtain for the formula isomorphism problem FI, reflecting that BWG-NI, unlike the coNP-hard FI, lies in NP^{GI} . The counting problem $\#\text{BWG-NI}$ is $\#\text{P}$ -hard.

Keywords: Graph isomorphism, formula isomorphism, polynomial hierarchy, lowness, NP-hardness.

1 Introduction

Two isomorphism problems delimit the structural complexity of equivalence under variable renaming. Graph isomorphism GI lies in $\text{NP} \cap \text{coAM}$, is low for Σ_2^P [Sch88], and is not NP-complete unless the polynomial hierarchy collapses. Formula isomorphism FI [AT00] lies in Σ_2^P , is coNP-hard, and is not Σ_2^P -complete unless the hierarchy collapses. The two are connected by the factorisation $\text{FI} = \exists \lambda \cdot \text{EQ}$ of [Cri26], an existential over renamings applied to a Boolean-equivalence test, and by the result of [ACL12] that syntactic DNF isomorphism (DSFI) is GI-complete.

A Boolean-weighted graph (BWG) consists of a host graph whose edges carry DNF formulas over a fixed variable set. Fixing a host bijection, one may impose on each matched pair of edge labels either of two conditions: that they be DSFI-isomorphic or that they be

*fra.cristiano@gmail.com

DSFI non-isomorphic. The two conditions yield two problems, the isomorphism variant BWG-ISO and the non-isomorphism variant BWG-NI, of markedly different complexity. Attention is restricted throughout to general (unrestricted) host graphs.

Two features motivate the definition. First, a Boolean-weighted graph is a labelled graph whose labels are themselves instances of a hard comparison problem: comparing two DNF labels syntactically is as hard as comparing two graphs, DSFI being GI-complete [ACL12]. The label test and the host test are therefore of the same difficulty, so the labels are not a bounded annotation on an otherwise unchanged graph isomorphism question. Second, a per-edge comparison relation and its complement give the two conditions studied here. Requiring matched labels to be isomorphic constrains a host bijection to respect the labels; requiring them to be non-isomorphic constrains it to separate them. The two per-edge tests are complements of one another and are of matched difficulty, DSFI being GI-complete and its negation DSNI being GNI-complete, so a first estimate places BWG-ISO and BWG-NI at the same level. The estimate fails. BWG-ISO sits exactly at GI (Theorem 3.5) while BWG-NI is NP-hard (Theorem 5.1), and the source of the asymmetry is identifiable: per-edge isomorphism requirements compose into a single isomorphism query, whereas the reduction of Theorem 5.1 uses identical hosts, where a witnessing bijection is an automorphism, and deciding whether some automorphism of the host separates every matched label pair is already NP-hard.

For the isomorphism variant, Section 3 proves $\text{BWG-ISO} \equiv_m^P \text{GI}$: the per-edge isomorphism requirements can be encoded as a single coloured-graph isomorphism query, each labelled edge being replaced by a canonical anchor vertex carrying the label's gadget. The isomorphism variant therefore sits exactly at the GI degree, hence is low for Σ_2^P ; the per-edge labels, each of which is GI-complete in isolation [ACL12], are absorbed into a single GI instance without raising the complexity.

For the non-isomorphism variant, the picture is different. The per-edge condition is a graph non-isomorphism (GNI) test, which places BWG-NI in NP^{GI} (Proposition 4.1). Section 4 also shows BWG-NI is GI-hard and GNI-hard but not coNP-hard unless $\text{PH} = \Sigma_2^P$, and Section 5 proves it is NP-hard by a reduction from the fixed-point-free automorphism problem FPFAUT; the NP-hardness is carried by the host-automorphism search and survives severe restriction of the labels (Section 8). NP^{GI} membership together with the NP-hardness then gives (Sections 6–7) that BWG-NI is not Σ_2^P -complete and not $\leq_m^P$ -equivalent to FI unless the hierarchy collapses, and that it lies in none of P^{GI} , coNP^{GI} , Low_2 unless the hierarchy collapses, while the rigid-host restriction (BWG-NI on hosts with no nontrivial automorphism) is low for Σ_2^P , and GNI-complete (equivalently, $\leq_m^P$ -complete for the conjunctive closure of GNI) on rigid hosts drawn from the tractable classes of Section 7. The counting problem $\#\text{BWG-NI}$ is shown to be $\#\text{P}$ -hard.

Read together, the two variants localize the complexity: the existential search over host isomorphisms (trivial precisely when the host is rigid) is what lifts it to NP-hardness while the GI-complete labels keep BWG-NI inside NP^{GI} , so BWG-NI is a natural candidate inhabitant of NP^{GI} lying outside P^{GI} unless the hierarchy collapses. This membership also sharpens the comparison with FI: were BWG-NI Σ_2^P -complete the hierarchy would collapse to its second level, one step below the third-level collapse entailed by Σ_2^P -completeness of FI [AT00], so BWG-NI obstructs second-level completeness more strongly than the coNP-hard FI, a second-level collapse being the stronger consequence since it implies the third-level one but not conversely.

All conditional separations obtained here are of relativizing type and do not yield unconditional separations.

Contributions. The following table summarises the results; the third column gives the section in which each is established.

Result	Statement	Section
Isomorphism variant	BWG-ISO $\equiv_m^P$ GI	3
Membership	BWG-NI $\in$ NP ^{GI}	4
Hardness profile	GI-hard and GNI-hard; not coNP-hard unless PH = Σ_2^P	4
NP-hardness	NP-hard via FPFAUT	5
NP-hardness vs FI	BWG-NI NP-hard unconditionally; FI not known NP-hard, its known deterministic floor from NP only GI [AT00]	5, 6
Counting	#BWG-NI is #P-hard, in $\# \cdot P^{GI}$	5
Non-completeness	not Σ_2^P -complete and not $\equiv_m^P$ FI unless PH = Σ_2^P	6
General hosts	NP-hard $\Rightarrow \notin P^{GI}/\text{coNP}^{GI}/\text{Low}_2$ unless collapse; P^{GI} -membership $\Rightarrow$ NP $\subseteq$ SPP	7
Rigid hosts	rigid-host BWG-NI is low for Σ_2^P , GNI-complete on tractable rigid hosts (Theorem 7.2)	7
Label alphabet	monotone, two-symbol labels suffice for all of the above	8

The per-edge label-comparison relation is the syntactic formula-isomorphism relation of [ACL12], a GI-complete equivalence relation, so BWG-NI is a two-level isomorphism problem: an isomorphism test on structures whose components are themselves compared under isomorphism; comparing labels by plain string disequality instead yields unconditional NP-completeness (Remark 8.2). The host quantifier alone positions the problem across its range: with the per-edge comparison held fixed, the rigid-host restriction is GNI-complete (Theorem 7.2), while admitting hosts with nontrivial automorphisms raises the problem from the GNI degree to NP-hardness (Proposition 7.5).

Organisation. Section 2 fixes notation and records the anchoring facts. Section 3 establishes the isomorphism variant, BWG-ISO $\equiv_m^P$ GI. Sections 4–7 develop the non-isomorphism variant: membership and hardness profile (4), NP-hardness and the consequent counting hardness (5), non-completeness and non-equivalence with FI (6), and the closure, lowness, and host-automorphism analysis (7). Section 8 records the label-alphabet refinement and collects observations and a map of the resulting complexity landscape. Section 9 records scope, provenance, and open problems; Section 10 collects the references.

2 Preliminaries

Familiarity with the polynomial hierarchy is assumed. GI and GNI denote graph isomorphism and graph non-isomorphism; GI $\in$ NP, GNI $\in$ coNP $\cap$ AM [GMW91], and GI $\in$ Low₂, i.e. $\Sigma_2^{GI} = \Sigma_2^P$ [Sch88], where Low₂ = { A : $\Sigma_2^A = \Sigma_2^P$ } consists of the sets that give Σ_2^P no additional power when used as oracles. FI is the formula isomorphism problem; FI $\in$ Σ_2^P and is coNP-hard [AT00]. EQ is Boolean equivalence, which is coNP-complete. DSFI, for DNF syntactic formula isomorphism, is syntactic DNF isomorphism, which is GI-complete [ACL12]; its negation DSNI (DNF syntactic non-isomorphism) = $\neg$ DSFI is GNI-complete. NP^{GI} (resp. P^{GI}) is the class decided by an NP (resp. deterministic polynomial-time) machine with a GI oracle; both are closed under $\leq_m^P$, and NP^{GI} $\subseteq$ NP^{NP} = Σ_2^P . For a graph G, Aut(G) denotes its automorphism group, and G is rigid (asymmetric) if Aut(G) is trivial, that is, if its only automorphism is the identity.

Fix a variable set Z. A Boolean-weighted graph (BWG) is a pair A = (H_A, ℓ_A) in which H_A = (V_A, E_A) is a finite simple graph, the host graph of A, and $\ell_A : E_A \rightarrow \text{DNF}(Z)$ is the edge labelling, assigning to each edge a DNF formula over Z. The host graph carries the combinatorial structure on which candidate isomorphisms act; the label of an edge is the

additional per-edge datum that such an isomorphism must respect. The relation imposed on a matched pair of labels (DSFI in BWG-ISO, DSNI in BWG-NI) is the per-edge comparison. The unrestricted (general) case places no restriction on the host; a restricted host family (for instance, the rigid hosts of Section 7) constrains which graphs may serve as the host. For an edge $e = \{u, v\}$ and a bijection h , write $h(e) = \{h(u), h(v)\}$; write $n = |V_A|$, $m = |E_A|$, and $\langle A \rangle$ for the input size; for a BWG G we also write $V(G)$ and $E(G)$ for the vertex and edge sets of its host. Two DNF formulas φ, ψ are syntactically DNF-isomorphic when one is carried to the other by a renaming of the variables together with a reordering of the terms and of the literals within each term; DSFI denotes the associated decision problem, the set of pairs (φ, ψ) that are so related, so $(\varphi, \psi) \in \text{DSFI}$ records that φ and ψ are syntactically DNF-isomorphic. Throughout, a DNF formula is a finite multiset of terms, repeated terms being significant, and a term may contain a variable both negated and unnegated; the companion paper [Cri26] adopts the set-of-terms convention, and the two conventions coincide on formulas without repeated terms.

In the two definitions that follow, the clauses are tagged (S) for the structural condition on the hosts and (L^+) , (L^-) for the label condition in its positive and negative form; the two problems share (S) and differ only in the sign of the label clause.

Definition 2.1 (BWG-ISO). Given BWGs A, B with $|V_A| = |V_B|$, the pair (A, B) lies in BWG-ISO iff there is a bijection $h : V_A \rightarrow V_B$ that is (S) a graph isomorphism $H_A \rightarrow H_B$ and (L^+) satisfies $(\ell_A(e), \ell_B(h(e))) \in \text{DSFI}$ for every $e \in E_A$; equivalently $\exists h$ [h is a graph isomorphism between the host graphs H_A and $H_B \wedge \forall e : (\ell_A(e), \ell_B(h(e))) \in \text{DSFI}$].

Definition 2.2 (BWG-NI). Given BWGs G, H with $|V_G| = |V_H|$, the pair (G, H) lies in BWG-NI iff there is a bijection $f : V_G \rightarrow V_H$ that is (S) a graph isomorphism $H_G \rightarrow H_H$ and (L^-) satisfies $(\ell_G(e), \ell_H(f(e))) \notin \text{DSFI}$ for every $e \in E_G$; equivalently $\exists f$ [f is a graph isomorphism between the host graphs H_G and $H_H \wedge \forall e : (\ell_G(e), \ell_H(f(e))) \notin \text{DSFI}$].

Three facts anchor the development, stated as Facts 2.3–2.5 and referred to by number throughout.

Fact 2.3 (the syntactic comparison is GI-complete [ACL12]). *DSFI is GI-complete. Informally, gad turns each DNF formula into a plain finite graph, so that two formulas yield isomorphic graphs exactly when they are syntactically DNF-isomorphic. More formally, there is a polynomial-time map $\text{gad} : \text{DNF}(Z) \rightarrow \dot{G}$, where $\dot{G}$ is the class of all finite graphs, such that $\text{gad}(\varphi)$ has size polynomial in the length of φ , and, for all $\varphi, \psi \in \text{DNF}(Z)$, $\text{gad}(\varphi) \cong \text{gad}(\psi) \Leftrightarrow (\varphi, \psi) \in \text{DSFI}$. Concretely, gad is the composition of two reductions of [ACL12]. First, φ is written as its variable-by-term incidence matrix over $\{-1, 0, 1, 2\}$, with entry 0 if the variable is absent from the term, 1 if it occurs unnegated, -1 if it occurs negated, and 2 if it occurs both ways. Second, that matrix is turned into a graph with a vertex for each variable, each term, and each matrix entry, together with fixed cliques of distinct sizes that mark the variable side and the term side and that codify the four entry values; each entry-vertex is joined to its variable, to its term, and by a single edge to a designated vertex of the clique codifying its value (Figure 1). By [ACL12, Theorem 2] (Lemmas 4 and 5 there) two such graphs are isomorphic iff the incidence matrices agree up to row and column permutations, that is, iff the formulas are related by a variable renaming together with a term reordering: syntactic DNF isomorphism. [ACL12] state the construction for CNF; the DNF form used here is the closing observation there, the variable-by-term matrix with terms in place of clauses, where the DNF problem is named DSFI. Conversely, the $\text{GI} \leq_m^p \text{DSFI}$ direction is realised by a map dnf from graphs to monotone DNFs ([ACL12], Lemma 1 with Theorem 3; Corollary 3 there gives the monotone statement), the map used in Theorem*

7.2; in particular the monotone fragment (formulas with no negated variables) is already GI-complete [ACL12]. Hence DSNI is GNI-complete and $DSFI \equiv_m^p GI$.

Throughout, gad denotes the map of Fact 2.3, sending each DNF formula to the plain graph $\text{gad}(\varphi)$ [ACL12]; it supplies the per-label graphs compared in both variants. The construction that turns it into a global reduction object is developed in Section 3 and is referred to as the coloured incidence encoding.

Fact 2.4 (the semantic comparison [AT00]). *FI is coNP-hard and lies in Σ_2^P ; it is the semantic analogue of DSFI, asking for a renaming making two formulas logically equivalent. Its negation FNI = co-FI is NP-hard and lies in Π_2^P . Moreover $GI \leq_m^p FI$ [BRS98, Proposition 12], a reduction observed with R. Chang [BR93, Proposition 2] whose image $\bigvee_{(i,j) \in E} (v_i \wedge v_j)$ is itself a monotone DNF, so it lands in FI directly; independently, semantic and syntactic isomorphism coincide on canonical DNF inputs (minterm expansions), where $GI \equiv_m^p CFI$, canonical formula isomorphism, the restriction of FI to such inputs, and $GI \leq_m^p CFI$ [Cri26, Theorem 2], and the identity map on this restriction gives $CFI \leq_m^p FI$.*

Fact 2.5 (GI is low in the polynomial hierarchy [Sch88]). *$GI \in \text{Low}_2$, hence $P^{GI} \subseteq \text{Low}_2 \subseteq \Sigma_2^P \cap \Pi_2^P$. By contrast $P^{FI} \supseteq \text{coNP}$ (by Fact 2.4), so $P^{GI} \subseteq P^{FI}$, and the inclusion is strict unless PH collapses, since $\text{coNP} \subseteq P^{GI}$ would give $\text{coNP} \subseteq \text{Low}_2$, i.e. $PH = \Sigma_2^P$.*

Lemma 2.6 (label precomputation). *Let the edges of an instance carry labels from a set S with $|S|$ polynomial in the input. With $O(|S|^2)$ GI queries one can compute an index $\text{ind} : S \rightarrow \mathbb{N}$ such that two labels receive the same index iff they are syntactically DNF-isomorphic. After this precomputation every per-edge comparison is the integer test $\text{ind}(\ell(e)) \neq \text{ind}(\ell(e'))$ (for the isomorphism variant, $\text{ind}(\ell(e)) = \text{ind}(\ell(e'))$), which uses no oracle.*

Proof. There are at most $|S|$ distinct labels; compare every pair (φ, ψ) with one GI query on $(\text{gad}(\varphi), \text{gad}(\psi))$, correct by Fact 2.3, and take the equivalence classes of the resulting relation as index classes. The number of pairs is $O(|S|^2)$, and the resulting integers determine every per-edge comparison by equality with no further oracle access. ■

Definition 2.7 (precomputable-query fragment $\text{NPGI}_{\text{pre}}^{\text{GI}}$). An NP^{GI} computation has precomputable queries if every instance submitted to the oracle lies in a polynomial-size set determined by the input alone, independently of the nondeterministic guess; the oracle answers can then be tabulated in advance (Lemma 2.6) and the nondeterministic phase made oracle-free. Let $\text{NP}_{\text{pre}}^{GI}$ be the class decided by such a computation. $\text{BWG-NI} \in \text{NP}_{\text{pre}}^{GI}$: the only label graphs that occur are $\text{gad}(\ell(e))$ over the edges of the two inputs, so every comparison lies among the $O(m^2)$ input-fixed pairs, and a guessed f selects which of these are required to be non-isomorphic. The inclusion $\text{NP}_{\text{pre}}^{GI} \subseteq \text{NP}^{GI}$ is immediate; whether it is proper is open (tabulation is not obviously available in general, since the queries of an NP^{GI} computation may depend on the nondeterministic guess).

$\text{CGNI}^\wedge$ denotes the closure of GNI under polynomial conjunction $\bigwedge_i \text{GNI}(A_i, B_i)$; the class is defined precisely in Definition 7.1. Since GNI is AND-closed (Definition 7.1) it coincides with the class of sets $\leq_m^p$ -reducible to GNI, and in particular lies in the Boolean hierarchy over GI.

3 The isomorphism variant: BWG-ISO is equivalent to GI

This section treats the isomorphism variant and reduces it to a single graph-isomorphism query. The obstruction to a per-edge reduction is that gad carries no isomorphism-invariant

distinguished vertex (Fact 2.3): rooting at an arbitrary vertex of a label graph is not isomorphism-invariant. The remedy is a canonical per-edge anchor, external to the label graphs, onto which each label graph is coned.




Figure 1. The gadget $\text{gad}(\varphi)$ of Fact 2.3 for the non-monotone formula $\varphi = (z_1 \wedge z_2) \vee \neg z_1$, following [ACL12]. The formula is written as its variable-by-term matrix over $\{-1, 0, 1, 2\}$ (top right); the graph has one vertex per variable, per term, and per matrix entry, with marker cliques of sizes 1 and 2 separating terms from variables and value-cliques of sizes 3, 4, 5, 6 codifying the entries $-1, 0, 1, 2$. Each entry-vertex joins its variable, its term, and, by a single edge, a designated vertex of the clique of its value; by [ACL12, Theorem 2] graph isomorphism of these gadgets is exactly a variable renaming with a term reordering.

Construction 3.1 (coloured incidence encoding). The DNF label on each host edge e is encoded by a gadget Γ_e , the graph that the map gad of Fact 2.3 builds from that formula. More formally, for a BWG A (Figure 2) and each edge $e \in E_A$, write $\Gamma_e := \text{gad}(\ell_A(e))$. Define a vertex-coloured graph $\Phi(A) = (W_A, F_A, \kappa)$ over the colour set $\{h, m, g\}$. The vertex set W_A consists of the host vertices, one new edge-vertex m_e per host edge e , and the vertices of the gadget copies Γ_e , all pairwise disjoint; more formally, $W_A = V_A \uplus \{m_e : e \in E_A\} \uplus \bigsqcup_e V(\Gamma_e)$, where $\uplus$ and $\bigsqcup$ denote disjoint union and $V(\Gamma_e)$ is the vertex set of Γ_e . The colouring is $\kappa = h$ on V_A , $\kappa = m$ on the edge-vertices m_e , and $\kappa = g$ on each gadget copy (gad returns a plain graph, so a single gadget colour suffices). The edge set F_A has three types: (incidence) $u \sim m_e$ iff $u \in e$, so V_A together with the edge-vertices m_e forms the incidence graph of H_A ; (cone) $m_e \sim a$ for every $a \in V(\Gamma_e)$, i.e. each gadget




Figure 2. A Boolean-weighted graph A : host path $u_1-u_2-u_3$ with edge labels $\varphi_1 = (z_1 \wedge z_2) \vee (z_2 \wedge z_3)$ and $\varphi_2 = z_1 \wedge z_3$.

vertex is joined to the apex m_e (the cone over Γ_e with apex m_e); (gadget) the edges of each Γ_e ; Figure 3 shows the construction on a worked example. Write $\Phi(A) \cong_{\text{col}} \Phi(B)$ when a colour-preserving isomorphism $\Phi(A) \rightarrow \Phi(B)$ exists.




Figure 3. The reduction of Construction 3.1 on a small instance, with each edge's gadget drawn as a single g -coloured block $\text{gad}(\ell(e))$ whose internal structure is that of Figure 1. The pair (A, B) is isomorphic by a host reflection, so $\Phi(A) \cong_{\text{col}} \Phi(B)$; the pair (A, B') differs in its label multiset, so $\Phi(A) \not\cong \Phi(B')$. Labels $\alpha = z_1 \wedge \neg z_2$ and $\beta = \neg z_1 \wedge z_2 \wedge z_3$; the colour classes h, m, g separate host, incidence, and gadget.

Definition 3.2 (edge-anchor gadget). For $e \in E_A$, let $N_A(e)$ be $\Phi(A)$ with the single vertex m_e individualised, that is, assigned a fresh colour given to no other vertex. The anchor m_e is canonical: it is the unique m -vertex whose h -neighbourhood equals e , since H_A is simple and distinct edges have distinct endpoint sets.

Lemma 3.3 (gadget isolation and isomorphism transfer). In $\Phi(A)$: (i) the g -neighbourhood of m_e is exactly $V(\Gamma_e)$, and gadget vertices are adjacent only within their own gadget and to their apex m_e ; (ii) for any colour-preserving isomorphism $\varphi : \Phi(A) \rightarrow \Phi(B)$ with $\varphi(m_e) = m_{e'}$, the restriction $\varphi[V(\Gamma_e)]$ is a colour-preserving isomorphism $\Gamma_e \rightarrow \Gamma_{e'}$.

Proof. (i) is immediate from Construction 3.1: a vertex of Γ_e has neighbours only inside $V(\Gamma_e)$ (gadget edges) and $\{m_e\}$ (cone edges); distinct gadgets are mutually non-adjacent and touch no h -vertex. (ii) Since $\varphi(m_e) = m_{e'}$, φ carries the apex's g -neighbourhood $V(\Gamma_e)$ bijectively onto $V(\Gamma_{e'})$; by (i) the only edges among these vertices are gadget-internal and are preserved by φ , so the restriction is an isomorphism $\Gamma_e \rightarrow \Gamma_{e'}$. ■

Construction 3.4 (decolouring [KST93]). Let $\hat{\Phi}(A)$ be obtained from $\Phi(A)$ by attaching to every vertex a tag that encodes its colour, using the individualisation gadget of [KST93]: with the three colours indexed $i \in \{1, 2, 3\}$ and $N = \max(|V(\Phi(A))|, |V(\Phi(B))|)$ a common value for the pair (if this maximum is below 4 the hosts are edgeless and the instance is decided directly), a colour- i vertex v receives a fresh path of $2N + 3$ vertices joined at one end to v , together with a secondary path of i further vertices branching from the $(N + 2)$ -th vertex of that chain. These tags are pairwise non-isomorphic across colours and identifiable by the lengths of their secondary branches; each tag, rooted at its attachment vertex v , admits no nontrivial automorphism fixing v , since, measured from the branch vertex, the two arms of the chain have different lengths ($N + 2$ through v against $N + 1$ to the far leaf), breaking the reflection of the bare chain even when v is isolated, and none occurs elsewhere in $\hat{\Phi}(A)$, the chain length $2N + 3$ exceeding the vertex count of $\Phi(A)$ so that no such chain lies inside $\Phi(A)$ itself; hence no isomorphism of $\hat{\Phi}(A)$ can carry a tag vertex off its tag or across colours; hence $\hat{\Phi}(A) \cong \hat{\Phi}(B) \Leftrightarrow \Phi(A) \cong_{\text{col}} \Phi(B)$.

Figure 4 shows this discharge on one edge-vertex and its gadget.




Figure 4. The colour-discharge of Construction 3.4 on one edge-vertex and its gadget. Each colour class h, m, g is replaced by a private rigid tag (drawn schematically as a distinct pendant path per class; the actual tag is the branched chain of Construction 3.4), so colour-preserving isomorphism of Φ becomes plain isomorphism of $\hat{\Phi}$.

Theorem 3.5 (BWG-ISO $\equiv_{\text{pm}}$ GI). $\overset{p}{m}$ **GI**. The map ρ sending (A, B) with $|V_A| = |V_B|$ to $(\hat{\Phi}(A), \hat{\Phi}(B))$, and to a fixed non-isomorphic pair when $|V_A| \neq |V_B|$, is a polynomial-time many-one reduction with $\rho(A, B) \in \text{GI} \Leftrightarrow (A, B) \in \text{BWG-ISO}$. Together with the trivial reduction $\text{GI} \leq_m^p \text{BWG-ISO}$ (bare hosts carrying one fixed label on every edge, accepting iff the hosts are isomorphic), $\text{BWG-ISO} \equiv_m^p \text{GI}$.

Proof. By Construction 3.4 it suffices to show $\Phi(A) \cong_{\text{col}} \Phi(B) \Leftrightarrow (A, B) \in \text{BWG-ISO}$. ($\Leftarrow$) Let h witness BWG-ISO. For each e , $(\ell_A(e), \ell_B(h(e))) \in \text{DSFI}$ gives by Fact 2.3 an isomorphism $g_e : \Gamma_e \rightarrow \Gamma_{h(e)}$. Define φ by $\varphi[V_A = h, \varphi(m_e) = m_{h(e)}$, and $\varphi[V(\Gamma_e) = g_e$. It is colour-preserving and bijective on each colour block, and preserves all three edge types: incidence $u \sim m_e \mapsto h(u) \sim m_{h(e)}$ since $u \in e \Leftrightarrow h(u) \in h(e)$; cone $m_e \sim a \mapsto m_{h(e)} \sim g_e(a)$ since the apex joins all gadget vertices on both sides; gadget $a \sim b \mapsto g_e(a) \sim g_e(b)$ since

g_e is a gadget isomorphism. ($\Rightarrow$) Let $\varphi : \Phi(A) \rightarrow \Phi(B)$ be colour-preserving. Colours force $\varphi(V_A) = V_B$ and $\varphi(\{m_e\}) = \{m_{e'}\}$; set $h := \varphi \upharpoonright V_A$. Each m_e has h -neighbourhood exactly e , which in a simple host determines it, so $\varphi(m_e)$ is the m -vertex with h -neighbourhood $\{h(u), h(v)\}$; hence $\{h(u), h(v)\} \in E_B$ and $\varphi(m_e) = m_{h(e)}$. Applying the same to φ^{-1} , h is an incidence-preserving bijection carrying E_A onto E_B , i.e. an isomorphism $H_A \rightarrow H_B$. By Lemma 3.3(ii) each $\varphi \upharpoonright V(\Gamma_e)$ is a colour-preserving isomorphism $\Gamma_e \rightarrow \Gamma_{h(e)}$, whence $(\ell_A(e), \ell_B(h(e))) \in \text{DSFI}$ by Fact 2.3, so h witnesses BWG-ISO. For the running time, let $s = \max_e |V(\Gamma_e)| = \text{poly}(\max_e |\ell_A(e)|)$; then $|W_A| = O(n + ms)$ and $|F_A| = O(ms^2)$, both polynomial; the discharge of Construction 3.4 then attaches a tag of $O(N)$ vertices at every vertex, with $N = O(n + ms)$, so $|V(\hat{\Phi}(A))| = O(N^2)$, and every part is listed in time polynomial in these sizes, so ρ is polynomial-time. ■

Remark 3.6 (per-edge reading). For $e \in E_A$, $e' \in E_B$, $N_A(e) \cong_{\text{col}} N_B(e')$ iff there is a host isomorphism $h : H_A \rightarrow H_B$ with $h(e) = e'$ and $(\ell_A(f), \ell_B(h(f))) \in \text{DSFI}$ for every f , by the same argument with $m_e \mapsto m_{e'}$ forced by the individualisation. The matrix $[N_A(e) \cong_{\text{col}} N_B(e')]$ thus records which label-preserving host isomorphisms route through which edge, while, when $E_A \neq \emptyset$, its nonemptiness coincides with the single query $\Phi(A) \cong_{\text{col}} \Phi(B)$ of Theorem 3.5 (on edgeless hosts the matrix is empty while the query holds iff $|V_A| = |V_B|$).

The construction is specific to the isomorphism variant. It does not transfer to BWG-NI: a polynomial-time GI-reduction of BWG-NI would, by Theorem 7.6(a) below, collapse PH, because BWG-NI is NP-hard whereas GI is low. The non-iso condition cannot be folded into a single GI query in this way, since the displacing witness it requires is not preserved by coning further structure onto the labels.

4 Membership and the hardness profile

Proposition 4.1 (membership). $BWG\text{-}NI \in \text{NP}^{\text{GI}}$.

Proof. An NP machine with a GI oracle guesses one bijection f , verifies in polynomial time that it is a graph isomorphism of the bare hosts (the host graphs stripped of their labels), then, with gad the label-to-graph map of Fact 2.3, queries the oracle on the pairs $\text{gad}(\ell_G(e))$, $\text{gad}(\ell_H(f(e)))$ for every edge e of G ; by Fact 2.3 a query answers non-isomorphic exactly when $(\ell_G(e), \ell_H(f(e))) \notin \text{DSFI}$, which is condition (L^-) of Definition 2.2. It accepts iff all $|E|$ non-adaptive queries answer non-isomorphic. One existential guess followed by GI-oracle queries with a polynomial-time decision is an NP^{GI} computation. ■

Proposition 4.2 (hardness profile). $BWG\text{-}NI$ is GI-hard and GNI-hard, and is not coNP-hard unless $\text{PH} = \Sigma_2^P$.

Proof. GI-hardness: if $|V_G| \neq |V_H|$ output a fixed NO-instance, say two copies of a single-edge host bearing the same label; otherwise label every edge of G with a fixed DNF φ and every edge of H with ψ chosen so $(\varphi, \psi) \notin \text{DSFI}$; then (L^-) holds for any f , so BWG-NI accepts iff a host isomorphism exists (GI-hardness also follows from Theorem 5.1, GI being in NP; the direct reduction is kept for its independence from Section 5). GNI-hardness: a single-edge instance is $(\varphi, \psi) \notin \text{DSFI}$, i.e. DSNI, which is GNI-complete by Fact 2.3. Non-coNP-hardness: by Proposition 4.1, $BWG\text{-}NI \in \text{NP}^{\text{GI}}$, and $\text{coNP} \subseteq \text{NP}^{\text{GI}}$ already forces $\text{PH} = \Sigma_2^P$ (Lemma 6.1, proved in Section 6 independently of this section); so coNP-hardness of BWG-NI would collapse the hierarchy. The per-edge comparison alone is weaker: DSNI $\equiv$ GNI is coNP-hard iff GI is NP-complete. ■

5 NP-hardness via fixed-point-free automorphisms

Theorem 5.1 (NP-hardness). *BWG-NI is NP-hard, via a reduction from the fixed-point-free automorphism problem FPFAUT [Lub81].*

Reduction. The reduction maps a graph G to the BWG-NI instance $(R(G), R(G))$, with $R(G)$ built as follows. We may assume G has no isolated vertices: a single isolated vertex is fixed by every automorphism (map such G to a fixed NO-instance), and if the set I of isolated vertices has $|I| \geq 2$ then $G \in \text{FPFAUT}$ iff $G \setminus I \in \text{FPFAUT}$ (extend by any derangement of I), while $G = I$ lies in FPFAUT iff $|I| \geq 2$; this preprocessing is polynomial. First, subdivide every edge $\{u, v\}$ of G by a new vertex z_{uv} (the path $u-z_{uv}-v$), giving $S(G)$; the subdivision lets an automorphism interchange u and v by exchanging the two half-edges $\{u, z_{uv}\}$ and $\{z_{uv}, v\}$ rather than fixing the edge (Figure 5). Second, at every original vertex v attach a pinning gadget: a 3-star with centre c and three leaves, the attachment leaf l_1 being also joined to v and the other two, l_2 and l_3 , free (Figure 6). Finally, label every edge of $R(G)$ by a distinct DNF, the labels pairwise non-DSFI-isomorphic (Lemma 5.2), and take both hosts equal to this labelled $R(G)$ (Figure 7). ■




Figure 5. *The subdivision device. Left: an interchange of adjacent endpoints fixes the edge as a set. Right: after inserting z_{uv} the two half-edges are exchanged, so no edge is fixed.*




Figure 6. *The pinning gadget at an original vertex v : a 3-star with centre c and leaves l_1, l_2, l_3 , attached to v through l_1 . The gadget vertices c, l_1, l_2, l_3 have degree multiset $(1, 1, 2, 3)$ in $R(G)$; these degrees together with adjacency to a centre identify every gadget vertex. If an automorphism fixes v it fixes l_1 and c , hence the pinned edge $\{l_1, c\}$.*

Lemma 5.2 (availability of labels). *For every m there is a list $\varphi_1, \dots, \varphi_m$ of DNF formulas over a common variable set, each of size at most m , pairwise non-DSFI-isomorphic and computable in time polynomial in m .*

Proof. Over the variables $z_1, \dots, z_m$, let φ_i be the DNF consisting of the single monomial $z_1 \wedge \dots \wedge z_i$. A DSFI-isomorphism renames variables and carries one formula to the other




Figure 7. The construction $R(G)$ for $G = C_4$. The fixed-point-free rotation σ lifts to a fixed-edge-free automorphism of $R(C_4)$; with all edge labels distinct and pairwise non-DSFI-isomorphic, (L^-) reads off fixed-edge-freeness. Vertices of C_4 are labelled $v_0, \dots, v_3$ here and $1, \dots, 4$ in Example 5.5.

syntactically, hence preserves the multiset of monomial degrees; φ_i has one monomial, of degree i , so $(\varphi_i, \varphi_j) \in \text{DSFI}$ forces $i = j$. The φ_i are pairwise non-DSFI-isomorphic, each of size at most m , and the list is produced in time polynomial in m . ■

Lemma 5.3 (vertex classes and the induced automorphism). *Assume G has no isolated vertices and put $n = |V(G)|$. The free leaves (degree 1), the centres (degree 3 with two degree-1 neighbours), the attachment leaves (degree 2 adjacent to a centre), the originals (the non-centres adjacent to an attachment leaf), and the subdivision vertices (the remaining degree-2 vertices) partition $V(R(G))$, and every automorphism of $R(G)$ preserves these five classes. Each $f \in \text{Aut}(R(G))$ restricts on the originals to an automorphism σ of G with $f(z_{uv}) = z_{\sigma(u)\sigma(v)}$, and carries the gadget at v to the gadget at $\sigma(v)$; conversely every $\sigma \in \text{Aut}(G)$ is induced by exactly 2^n automorphisms of $R(G)$, differing by which free leaf maps to which in each of the n gadgets. In particular $|\text{Aut}(R(G))| = 2^n \cdot |\text{Aut}(G)|$.*

Class	Degree in $R(G)$	Identifying test
free leaves l_2, l_3	1	the degree-1 vertices (only neighbour a centre)
centres c	3	exactly two degree-1 neighbours (a degree-3 original has none)
attachment leaves l_1	2	adjacent to a centre (the other neighbour an original)
originals v	$\deg_G(v) + 1 \geq 2$	the non-centres adjacent to an attachment leaf
subdivision vertices z_{uv}	2	the remaining degree-2 vertices (both neighbours originals)

The vertex classification of Lemma 5.3. Degrees are taken in $R(G)$; each test uses the abstract graph alone, and every predicate is isomorphism-invariant.

Proof. In $R(G)$ a free leaf has degree 1 (its only neighbour is a centre); a centre has degree 3 (neighbours the attachment leaf and the two free leaves), two of them of degree 1; an attachment leaf has degree 2 (neighbours v and the centre); an original v has degree $\deg_G(v) + 1 \geq 2$ (its subdivision neighbours together with the attachment leaf); a subdivision vertex z_{uv} has degree 2 (neighbours u and v). The five predicates are therefore mutually exclusive and exhaustive: the degree-1 vertices are exactly the free leaves; among degree-3 vertices a centre is the one with two degree-1 neighbours, whereas an original of G -degree 2 also has degree 3 but all of its neighbours have degree ≥ 2 ; a degree-2 vertex is an attachment leaf if adjacent to a centre, an original of G -degree 1 if adjacent to an attachment leaf, and a subdivision vertex if its two neighbours are originals; these cases are distinct, since an attachment leaf has one original and one centre as neighbours, a G -degree-1 original has one attachment leaf and one subdivision, and a subdivision has two originals. Across all degrees the originals are exactly the non-centres adjacent to an attachment leaf (an attachment leaf's only neighbours being its original and its centre), and the remaining vertices are the subdivisions. Each predicate is invariant under isomorphism, so every $f \in \text{Aut}(R(G))$ preserves the classes. Hence f permutes the originals; as z_{uv} is the unique subdivision vertex adjacent to both u and v , the restriction σ of f to the originals satisfies $f(z_{uv}) = z_{\sigma(u)\sigma(v)}$ and sends edges of G to edges of G , so $\sigma \in \text{Aut}(G)$; and f carries the gadget at v to the gadget at $\sigma(v)$, taking its centre and attachment leaf to those at $\sigma(v)$ and its two free leaves onto the two free leaves at $\sigma(v)$. Conversely any $\sigma \in \text{Aut}(G)$ lifts: act by σ on the originals, by $z_{uv} \mapsto z_{\sigma(u)\sigma(v)}$ on the subdivisions, and by the induced map on each gadget, with a free choice in each gadget of the bijection between its two free leaves and those of the image gadget. These n choices are independent and the remainder of f is forced, so exactly 2^n automorphisms induce each σ ; summing over σ gives $|\text{Aut}(R(G))| = 2^n \cdot |\text{Aut}(G)|$. ■

Lemma 5.4 (fixed edges versus fixed points). *For $f \in \text{Aut}(R(G))$ with induced $\sigma \in \text{Aut}(G)$, f is fixed-edge-free iff σ is fixed-point-free. Consequently $R(G)$ has a fixed-edge-free automorphism iff $G \in \text{FPFAUT}$.*


Proof. If $\sigma(v) = v$ then f maps the gadget at v to itself and fixes its attachment leaf l_1 , the unique gadget vertex adjacent to the fixed v ; hence it also fixes the centre c , the unique neighbour of l_1 other than v , so it fixes the edge $\{l_1, c\}$ and is not fixed-edge-free. Conversely, let σ be fixed-point-free and f any automorphism inducing it; every edge of $R(G)$ is of one of three kinds. A subdivision edge $\{v, z_{uv}\}$ maps to $\{\sigma(v), z_{\sigma(u)\sigma(v)}\}$, which differs from it because $\sigma(v) \neq v$ and $\sigma(v)$ is an original while z_{uv} is a subdivision vertex; if σ interchanges the endpoints of $\{u, v\}$ the two half-edges are merely exchanged. An attachment edge $\{v, l_1\}$ maps to the attachment edge of the gadget at $\sigma(v)$ and moves since $\sigma(v) \neq v$. An internal gadget edge at v is carried into the gadget at $\sigma(v) \neq v$, a different gadget, so it moves, a free-leaf swap notwithstanding. Hence no edge is fixed and f is fixed-edge-free. Ranging over automorphisms and using Lemma 5.3 yields the final equivalence. ■

Proof of Theorem 5.1. The reduction outputs $(R(G), R(G))$ in polynomial time, its labels by Lemma 5.2. Since the labels are distinct and pairwise non-DSFI-isomorphic, for any host automorphism f and edge e the condition $(\ell(e), \ell(f(e))) \notin \text{DSFI}$ holds iff $f(e) \neq e$, so it holds at every edge iff f is fixed-edge-free. By Lemma 5.3 every host automorphism induces $\sigma \in \text{Aut}(G)$, and by Lemma 5.4, $R(G)$ has a fixed-edge-free automorphism iff $G \in \text{FPFAUT}$. Thus $(R(G), R(G)) \in \text{BWG-NI}$ iff $G \in \text{FPFAUT}$, and since FPFAUT is NP-complete [Lub81], BWG-NI is NP-hard. The reduction uses general fixed-point-free automorphisms, not only involutions; the subdivision-plus-pinning device is of the kind used by [HSV01] for fixed-edge-free involutory automorphisms.


Example 5.5 (the reduction on C_4 and on P_3). **Example 5.5 (the reduction on C_4 and on P_3).** Take $G = C_4$ with vertices 1, 2, 3, 4 in cyclic order (Figure 8). Then $R(C_4)$ has 24 vertices: the four originals, four subdivision vertices $z_{12}, z_{23}, z_{34}, z_{41}$, and four pinning gadgets, the i -th with centre c_i , attachment leaf l_1^i , and free leaves l_2^i, l_3^i . The degree multiset locates the classes: free leaves have degree 1, attachment leaves and subdivision vertices degree 2, centres degree 3; and so do the originals, since $\deg_G(v) = 2$ gives each original degree $\deg_G(v) + 1 = 3$. This degree-3 coincidence between originals and centres, the kind of collision that can hide a stray automorphism, is resolved by neighbour degrees exactly as in Lemma 5.3: a centre has two degree-1 neighbours (its free leaves), whereas an original has none, its three neighbours being two subdivision vertices and one attachment leaf, all of degree 2. The rotation $\sigma : i \mapsto i + 1 \pmod{4}$ is a fixed-point-free automorphism of C_4 . It lifts to $f \in \text{Aut}(R(C_4))$ acting by $i \mapsto i + 1$ on the originals, $z_{i,i+1} \mapsto z_{i+1,i+2}$ on the subdivisions, and carrying the gadget at i to the gadget at $i + 1$ (any of the 2^4 free-leaf matchings serves). No edge of $R(C_4)$ is fixed: a subdivision edge $\{i, z_{i,i+1}\} \mapsto \{i+1, z_{i+1,i+2}\}$, an attachment edge $\{i, l_1^i\} \mapsto \{i+1, l_1^{i+1}\}$, and an internal gadget edge at i is carried into the gadget at $i + 1 \neq i$, which is Lemma 5.4 made explicit. Labelling the edges of $R(C_4)$ by the pairwise non-DSFI-isomorphic monomials of Lemma 5.2, every label is displaced, so (L^-) holds at f and $(R(C_4), R(C_4)) \in \text{BWG-NI}$, certifying $C_4 \in \text{FPFAUT}$. Exhaustively, $|\text{Aut}(R(C_4))| = 2^4 \cdot |\text{Aut}(C_4)| = 16 \cdot 8 = 128$, of which $2^4 \cdot \#\text{FPFAUT}(C_4) = 16 \cdot 5 = 80$ are fixed-edge-free (Lemma 5.3, Remark 5.6). By contrast take $G = P_3$, the path 1–2–3. Its only nontrivial automorphism is the reflection $1 \leftrightarrow 3$, which fixes the centre 2; hence $P_3 \notin \text{FPFAUT}$. In $R(P_3)$ every automorphism fixes the original 2, therefore fixes its attachment leaf l_1^2 (the unique gadget vertex adjacent to 2) and then the centre c_2 (the unique neighbour of l_1^2 other than 2), and so fixes the pinned edge $\{l_1^2, c_2\}$: no

automorphism is fixed-edge-free, and $(R(P_3), R(P_3)) \notin \text{BWG-NI}$. The pinning gadget thus turns the fixed vertex 2 into a fixed edge, which is its role in Lemma 5.4.

(a) $G = C_4$: σ fixed-point-free $\Rightarrow$ every edge of $R(C_4)$ moves



(b) $G = P_3$: reflection fixes vertex 2 $\Rightarrow$ pinned edge $\{l_1^2, c_2\}$ fixed



● original ■ subdivision z_{uv} ● centre c ● attachment leaf l_1 ● free leaves l_2, l_3 — fixed edge

Figure 8. Example 5.5: the reduction on C_4 (a) and P_3 (b). (a) the rotation $\sigma : i \mapsto i + 1$ is fixed-point-free, and its lift to $R(C_4)$ displaces every edge, so $C_4 \in \text{FPFAUT}$. (b) the reflection of P_3 fixes vertex 2, so every lift fixes the pinned edge $\{l_1^2, c_2\}$ (bold) and $P_3 \notin \text{FPFAUT}$. Vertex classes are coloured as in Figures 5–6; the originals $v_0, \dots, v_3$ of Figure 7 are the vertices 1, $\dots$, 4 here.

Remark 5.6 (counting). By Lemma 5.3 the automorphisms of $R(G)$ inducing a given σ number 2^n , $n = |V(G)|$, and by Lemma 5.4 one of them is fixed-edge-free iff σ is fixed-point-free; hence $R(G)$ has exactly $2^n \cdot \#\text{FPFAUT}(G)$ fixed-edge-free automorphisms, and these are the accepting witnesses. Write $\#\text{BWG-NI}$ for the function mapping an instance (G, H) to the number of bijections f satisfying (S) and (L^-) . A reduction is parsimonious if it preserves the number of solutions and weakly parsimonious if the two counts are related by a polynomial-time computable factor; this reduction is weakly parsimonious. An asymmetric replacement gadget would make it exactly parsimonious, but the weak form suffices for Proposition 5.7.

NP-hardness is consistent with membership: $\text{NP} \subseteq \text{NP}^{\text{GI}}$, so an NP-hard problem may lie in NP^{GI} , unlike a coNP-hard one. The host-automorphism search carries the NP-hardness; the labels only encode the distinctness restriction. The construction uses unbounded-degree hosts; see Section 9.

Proposition 5.7 (counting hardness). $\#\text{BWG-NI}$ is $\#P$ -hard, and $\#\text{BWG-NI} \in \#P^{\text{GI}}$, where $\#P^{\text{GI}}$ is the counting operator applied to P^{GI} : the functions $x \mapsto |\{y : (x, y) \in B\}|$ over polynomially length-bounded witnesses y , for a fixed $B \in P^{\text{GI}}$. Counting reduces to decision for both neighbours of BWG-NI, $\#GI \in \text{FP}^{\text{GI}}$ [Mat79] and $\#FI \in \text{FP}^{\text{FI}}$ [AT00, Theorem 5.2], and for the syntactic comparison $\#DSFI \in \text{FP}^{\text{GI}}$ (DSFI-isomorphisms form a coset of the syntactic automorphism group, so Mathon’s pointwise-stabiliser counting [Mat79] applies with a DSFI, hence by Fact 2.3 a GI, oracle [ACL12]; variables are stabilised syntactically by appending each as a singleton term in a distinct multiplicity exceeding the term counts of the input formulas, which every syntactic isomorphism preserves); that

$\#BWG\text{-}NI \in \text{FP}^{\text{GI}}$ would instead force $PH = \Sigma_2^P$ makes this a three-way asymmetry, $BWG\text{-}NI$ breaking relative to its own natural oracle GI a counting-to-decision reduction that holds for both GI and FI relative to theirs.

Proof. By Remark 5.6, $\#BWG\text{-}NI(R(G), R(G)) = 2^n \cdot \#\text{FPFAUT}(G)$; dividing by the computable 2^n gives $\#\text{FPFAUT} \leq \#BWG\text{-}NI$ under polynomial-time counting reductions, so $\#BWG\text{-}NI$ is $\#P$ -hard, since $\#\text{FPFAUT}$ is $\#P$ -complete, by a parsimonious reduction from 3SAT [Lub81, Corollary 2]. The witness predicate uses the GI oracle, so $\#BWG\text{-}NI \in \#P^{\text{GI}}$. If $\#BWG\text{-}NI \in \text{FP}^{\text{GI}}$ then, since $\#BWG\text{-}NI$ is $\#P$ -hard, $\#P \subseteq \text{FP}^{\text{GI}}$, whence $PH \subseteq P^{\#P}$ [Tod91] $\subseteq P^{\text{GI}} \subseteq \text{Low}_2$ [Sch88], collapsing the hierarchy. ■

6 Non-completeness and non-equivalence with FI

Lemma 6.1 (after [Sch88]; cf. [AT00]). *If $\text{coNP} \subseteq \text{NP}^{\text{GI}}$ then $PH = \Sigma_2^P$.*

Proof. (i) $\Sigma_2^P \subseteq \text{NP}^{\text{GI}}$. Let $L \in \Sigma_2^P$, $L = \{x : \exists y (x, y) \in B\}$ with $B \in \text{coNP}$. By hypothesis $B \in \text{NP}^{\text{GI}}$, $B = \{(x, y) : \exists z M^{\text{GI}}(x, y, z) \text{ accepts}\}$. Then $L = \{x : \exists (y, z) M^{\text{GI}}(x, y, z) \text{ accepts}\}$, a single existential block over a polynomial-time GI -oracle predicate, hence NP^{GI} . As $GI \in \text{NP}$, $\text{NP}^{\text{GI}} \subseteq \text{NP}^{\text{NP}} = \Sigma_2^P$, so $\Sigma_2^P = \text{NP}^{\text{GI}}$. (ii) $\Sigma_3^P = \text{NP}^{\Sigma_2^P} = \text{NP}^{\text{NP}^{\text{GI}}} = \Sigma_2^{\text{GI}} = \Sigma_2^P$ by $GI \in \text{Low}_2$ [Sch88]. Hence $PH = \Sigma_2^P$. ■

Theorem 6.2 (non-completeness). *If $BWG\text{-}NI$ is Σ_2^P -complete then $PH = \Sigma_2^P$.*

Proof. If $BWG\text{-}NI$ is Σ_2^P -complete then $\Sigma_2^P \leq_m^P BWG\text{-}NI \in \text{NP}^{\text{GI}}$; as NP^{GI} is $\leq_m^P$ -closed (compose the reduction with the oracle machine), $\Sigma_2^P \subseteq \text{NP}^{\text{GI}}$, so $\text{coNP} \subseteq \text{NP}^{\text{GI}}$ and $PH = \Sigma_2^P$ by Lemma 6.1. ■

Corollary 6.3. *No coNP -hard problem lies in NP^{GI} unless $PH = \Sigma_2^P$. $BWG\text{-}NI$ is an inhabitant of NP^{GI} that is not coNP -hard unless $PH = \Sigma_2^P$ (Proposition 4.2).*

Proof. A coNP -hard $L \in \text{NP}^{\text{GI}}$ would place $\text{coNP} \subseteq \text{NP}^{\text{GI}}$ by the $\leq_m^P$ -closure of NP^{GI} , whence $PH = \Sigma_2^P$ by Lemma 6.1. ■

Theorem 6.4 (non-equivalence). *If $BWG\text{-}NI \equiv_m^P FI$ then $PH = \Sigma_2^P$. Already the single direction $FI \leq_m^P BWG\text{-}NI$ implies $\text{coNP} \subseteq \text{NP}^{\text{GI}}$, hence $PH = \Sigma_2^P$.*

Proof. FI is coNP -hard [AT00]. $FI \leq_m^P BWG\text{-}NI$ would make $BWG\text{-}NI$ coNP -hard; with $BWG\text{-}NI \in \text{NP}^{\text{GI}}$ and NP^{GI} $\leq_m^P$ -closed, $\text{coNP} \subseteq \text{NP}^{\text{GI}}$, so $PH = \Sigma_2^P$ by Lemma 6.1. ■

The mechanism is whether the per-edge comparison is an isomorphism test or its negation. The factorisation $FI = \exists \lambda \cdot EQ$ has a coNP -complete inner test, making FI coNP -hard; $BWG\text{-}NI$ never certifies a full coNP fact, its per-edge comparison being the negation of a GI -confined test ($\text{DSNI} \equiv \text{GNI}$), so the per-edge comparison alone is only GNI -hard. Section 7 makes this quantitative.

Remark 6.5 (unconditional versus randomized NP -hardness). Whether FI is NP -hard is a challenging open question [AT00]; the hardest NP problem known to reduce to it is GI [AT00] (the reduction itself is Chang's, Fact 2.4), so the established deterministic NP -hardness floor of FI reaches only the GI degree. $BWG\text{-}NI$ is NP -hard unconditionally (Theorem 5.1), so its floor is all of NP . FI is NP -hard under randomized reductions only by composition, from $\text{SAT} \leq_r \text{USAT}$ (unique satisfiability) [VV86], $\text{USAT} \leq_m^P \text{FC}$ ([BRS98, Proposition 10 with Corollary 11]; the map is restated in [AT00]), and $\text{FC} \equiv_m^P FI$ [BRS98, Theorem 6], FC being formula congruence (a variable permutation composed with negations

of variables); [AT00] state the randomized hardness for USAT, not as a theorem about FI. The remaining points of difference are coNP-hardness (FI yes, BWG-NI no) and NP^{GI} membership (BWG-NI yes, FI not unless the hierarchy collapses).

7 Degree, lowness, and the relativized separation

This section locates the source of the complexity: the single feature separating the NP-hard general problem from its tractable rigid restriction (hosts with no nontrivial automorphism) is the quantifier over host isomorphisms, the labels contributing only a polynomial-time distinctness condition. On rigid hosts the problem is captured by the conjunctive closure of GNI (Theorem 7.2) and is low for Σ_2^P (Theorem 7.4); Proposition 7.5 records that this quantifier alone accounts for each separation between the rigid and general cases. Throughout, BWG-NI_{rig} denotes the restriction of BWG-NI to instances in which both host graphs are rigid.

Definition 7.1 (conjunctive closure of GNI). $\text{CGNI}^\wedge$ is the class of L for which a polynomial-time map sends each x to graph pairs $(A_1, B_1), \dots, (A_m, B_m)$ with $x \in L \Leftrightarrow \bigwedge_i \text{GNI}(A_i, B_i)$. A problem is OR-closed (AND-closed) if the disjunction (conjunction) of any polynomially many of its instances reduces to a single instance. GNI is OR-closed (a disjunction of non-isomorphisms is one GNI instance, via disjoint unions whose components are individualised so that any isomorphism matches the i -th components), and it is also AND-closed, since an AND of non-isomorphisms is the complement of an OR of isomorphisms: GI has a polynomial-time computable or-function [KST93, Theorem 1.42] (originally [Cha89, LT92]), built from the ordered-pair encoding in the proof of [KST93, Theorem 1.40] together with the unordered-pair behaviour of disjoint unions of connected graphs [KST93, Fact 1.41], so by complementation GNI has a polynomial-time computable and-function. Hence $\text{CGNI}^\wedge$ coincides with the class of sets $\leq_m^p$ -reducible to GNI, and the $\text{CGNI}^\wedge$ -completeness below is GNI-completeness.

Theorem 7.2 (hardness for the closure; the rigid restriction). *Every $L \in \text{CGNI}^\wedge$ satisfies $L \leq_m^p \text{BWG-NI}$, that is, BWG-NI is $\text{CGNI}^\wedge$ -hard. On arbitrary rigid hosts $\text{BWG-NI}_{\text{rig}} \in P^{\text{GI}}$. On rigid hosts drawn from a polynomial-time-recognizable class on which an isomorphism is computable, and triviality of the automorphism group decidable, in polynomial time (for instance rigid trees), $\text{BWG-NI}_{\text{rig}}$ is moreover $\leq_m^p$ -complete for $\text{CGNI}^\wedge$, equivalently for GNI.*

Proof. Hardness. Given $\bigwedge_i \text{GNI}(A_i, B_i)$, $i = 1, \dots, m$, first pad with fixed true conjuncts (each a fixed non-isomorphic pair) so that $m \geq 3$. Let R be the tree obtained from the path $u_0 - u_1 - \dots - u_m$ by attaching at u_m two pendant paths (the markers), of lengths one and two. Then u_m is the unique vertex of degree three and its three branches have distinct lengths $m, 1, 2$, so R is rigid (a path graph, by contrast, retains the reversal automorphism); the unique isomorphism between two copies of R is the identity. Take both hosts equal to R ; label the i -th path edge $\{u_{i-1}, u_i\}$ with $\text{dnf}(A_i)$ in G and $\text{dnf}(B_i)$ in H , where dnf is the graphs-to-DNFs map of Fact 2.3 (so $A_i \cong B_i$ iff $(\text{dnf}(A_i), \text{dnf}(B_i)) \in \text{DSFI}$), and label every marker edge of G with one fixed DNF and every marker edge of H with a fixed non-DSFI-isomorphic partner. The unique host isomorphism is the identity, pairing edge i with edge i and marker edges with marker edges; the marker labels always disagree, so (L^-) holds iff every path edge disagrees, i.e. iff $\bigwedge_i \text{GNI}(A_i, B_i)$. The image lies in $\text{BWG-NI}_{\text{rig}}$ and is a rigid tree. Membership. The machine first rejects unless both hosts are rigid; rigidity is decidable in P^{GI} since GA, the problem of deciding whether a graph has a nontrivial automorphism, is Turing-reducible to GI [KST93, Example 1.10]. On arbitrary rigid hosts the unique isomorphism f , or the fact that none exists, is computable in FP^{GI} : search

reduces to decision for GI by individualisation [KST93], fixing f vertex by vertex in $O(n^2)$ oracle calls. A P^{GI} machine computes f and then decides the per-edge disequations $(\ell(e), \ell(f(e))) \notin \text{DSFI}$ with $|E|$ further GI queries (each DSFI test is a GI query by Fact 2.3), so $\text{BWG-NI}_{\text{rig}} \in P^{\text{GI}}$ with no tractability assumption. When the rigid hosts are moreover drawn from such a class, the map first checks in polynomial time that both hosts lie in the class and are rigid, outputting a fixed isomorphic pair (a false conjunct) otherwise; f is then computable without an oracle in polynomial time, and if none exists the map outputs the same false conjunct; otherwise output the gadget pairs $(\text{gad}(\ell(e)), \text{gad}(\ell(f(e))))$ over all edges e , so acceptance is $\bigwedge_e \text{GNI}(\text{gad}(\ell(e)), \text{gad}(\ell(f(e))))$ and $\text{BWG-NI}_{\text{rig}} \in \text{CGNI}^\wedge$. ■

Canonicity reaches the same degree by a different route: the complement of the canonical restriction of the semantic comparison is GNI-complete, $\text{CFNI} = \text{co-CFI} \equiv_m^p \text{GNI}$ [Cri26, Theorem 3]; rigidity removes the host quantifier, canonicity the representation ambiguity, and both land on the GNI degree.

Proposition 7.3 (the closure lies in the low hierarchy). $\text{GNI} \in \text{CGNI}^\wedge \subseteq P^{\text{GI}} \subseteq \text{Low}_2 \subseteq \Sigma_2^P$.

Proof. GNI is decided by one GI query; $\text{CGNI}^\wedge$ by m non-adaptive queries, so $\text{CGNI}^\wedge \subseteq P^{\text{GI}}$. Any $A \in P^{\text{GI}}$ satisfies $\Sigma_2^A \subseteq \Sigma_2^{\text{GI}} = \Sigma_2^P$ [Sch88], i.e. $A \in \text{Low}_2$. ■

The third inclusion is the operative one. $\text{CGNI}^\wedge \subseteq \text{Low}_2$, whereas $\text{coNP} \subseteq \text{Low}_2$ would force $\Sigma_3^P = \Sigma_2^P$ (since $\Sigma_2^{\text{coNP}} = \Sigma_3^P$). So unless the hierarchy collapses, the hardness floor of BWG-NI sits inside the low part of Σ_2^P , not in the coNP -complete part where the floor of FI sits.

Theorem 7.4 (rigid lowness). $\text{BWG-NI}_{\text{rig}}$ is low for Σ_2^P ; in particular it is not Σ_2^P -complete unless $\text{PH} = \Sigma_2^P$.

Proof. $\text{BWG-NI}_{\text{rig}} \in P^{\text{GI}} \subseteq \text{Low}_2$ for arbitrary rigid hosts (Theorem 7.2, Proposition 7.3); it is a decision problem, a language rather than a promise problem, since rigidity of the hosts is decidable in P^{GI} since GA is Turing-reducible to GI [KST93, Example 1.10], and ill-formed instances are rejected rather than left unpromised, so the Sch88 lowness applies to it as a set. A low set L satisfies $\Sigma_2^L = \Sigma_2^P$; were it Σ_2^P -complete, $\Sigma_3^P = \Sigma_2^P$. ■

Proposition 7.5 (the host-isomorphism quantifier accounts for the rigid/general distinction). BWG-NI and $\text{BWG-NI}_{\text{rig}}$ differ only in the quantifier over host isomorphisms. In BWG-NI the witness f ranges over all isomorphisms $G \rightarrow H$, a coset of $\text{Aut}(G)$, under the per-edge disequation predicate, which is not $\text{Aut}(G)$ -invariant; in $\text{BWG-NI}_{\text{rig}}$ the host is rigid, so this coset is a single map and the quantifier is trivial. This quantifier is the sole source of each complexity difference established above: (i) on rigid hosts drawn from a class as in Theorem 7.2, $\text{BWG-NI}_{\text{rig}}$ is $\leq_m^p$ -complete for $\text{CGNI}^\wedge$ (Theorem 7.2), whereas BWG-NI is $\text{CGNI}^\wedge$ -hard but does not lie in $\text{CGNI}^\wedge$ unless $\text{PH} = \Sigma_2^P$; (ii) $\text{BWG-NI}_{\text{rig}} \in P^{\text{GI}}$ (Theorem 7.2), whereas $\text{BWG-NI} \in \text{NP}^{\text{GI}}$ and $\text{BWG-NI} \notin P^{\text{GI}}$ unless $\text{PH} = \Sigma_2^P$ (Theorem 7.6(a), Corollary 7.7); (iii) $\text{BWG-NI}_{\text{rig}} \in \text{Low}_2$ (Theorem 7.4), whereas $\text{BWG-NI} \notin \text{Low}_2$ unless $\text{PH} = \Sigma_2^P$ (Theorem 7.6(c)); (iv) $\text{co-BWG-NI} \in \Pi_2^P$ natively, and $\text{BWG-NI} \in \text{coNP}^{\text{GI}}$ (equivalently $\text{co-BWG-NI} \in \text{NP}^{\text{GI}}$) implies $\text{PH} = \Sigma_2^P$ (Theorem 7.6(b)). When $\text{Aut}(G) = 1$ each statement reduces to its $\text{BWG-NI}_{\text{rig}}$ form.

Proof. Immediate: the two problems have identical definitions except that the host bijection is quantified over all isomorphisms $G \rightarrow H$ in BWG-NI and is fixed to the unique such isomorphism in $\text{BWG-NI}_{\text{rig}}$; (i)–(iv) are the cited results. ■

Theorem 7.6 (NP-hardness excludes P^{GI} , coNP^{GI} , and Low_2). *GI , coNP^{GI} , and Low_2). Each of the following implies $PH = \Sigma_2^P$: (a) $\text{BWG-NI} \in P^{GI}$; (b) $\text{BWG-NI} \in \text{coNP}^{GI}$; (c) $\text{BWG-NI} \in \text{Low}_2$. Moreover (a) also yields $NP \subseteq \text{SPP}$. Hence, unless $PH = \Sigma_2^P$, general BWG-NI lies in none of P^{GI} , coNP^{GI} , Low_2 .*

Proof. BWG-NI is NP-hard (Theorem 5.1). (a) $\text{BWG-NI} \in P^{GI}$ gives $NP \subseteq P^{GI} \subseteq \text{Low}_2$, so $\Sigma_2^{\text{SAT}} = \Sigma_2^P$; but $\Sigma_2^{\text{SAT}} = \Sigma_3^P$, so $\Sigma_3^P = \Sigma_2^P$. For the SPP consequence, $GI \in \text{SPP}$ [AK02] and SPP is self-low, $P^{\text{SPP}} = \text{SPP}$ [FFK94], so $NP \subseteq P^{GI} \subseteq P^{\text{SPP}} = \text{SPP}$. (b) $\text{BWG-NI} \in \text{coNP}^{GI}$ gives $\text{coNP} \subseteq \text{NP}^{GI}$, so $PH = \Sigma_2^P$ by Lemma 6.1. (c) NP-hardness gives $\Sigma_2^{\text{BWG-NI}} \supseteq \Sigma_2^{\text{SAT}} = \Sigma_3^P$; if $\text{BWG-NI} \in \text{Low}_2$ then $\Sigma_3^P = \Sigma_2^P$. ■

Corollary 7.7 (a witness to a relativized separation). *If PH is infinite then $P^{GI} \neq \text{NP}^{GI}$, and BWG-NI witnesses the separation: $P^{GI} = \text{NP}^{GI} \Rightarrow NP \subseteq P^{GI} \subseteq \text{Low}_2 \Rightarrow PH = \Sigma_2^P$, so under infinite PH , $\text{BWG-NI} \in \text{NP}^{GI} \setminus P^{GI}$.*

The collapse needs no witness problem: $NP \subseteq \text{NP}^{GI}$ always, so $\text{NP}^{GI} = P^{GI}$ gives $NP \subseteq P^{GI}$, and $P^{GI} \subseteq \text{Low}_2$ [Sch88] yields $PH = \Sigma_2^P$. The role of BWG-NI is to provide a natural problem in $\text{NP}^{GI} \setminus P^{GI}$: under $PH \neq \Sigma_2^P$ it lies there, with Theorem 7.6(a) adding $NP \subseteq \text{SPP}$ to the cost of the contrary. By the relativizing nature of the arguments this is conditional and does not yield an unconditional separation. A complementary barrier holds at the witness level: $\text{BWG-NI} \notin P^{GI}$ already implies $P \neq NP$, since $P = NP$ gives $PH = P$ and then $\text{BWG-NI} \in \text{NP}^{GI} \subseteq \Sigma_2^P = P \subseteq P^{GI}$; an unconditional proof of the exclusion is therefore at least as hard as separating P from NP .

Remark 7.8 (hardness floor). The hardness floor of BWG-NI is the join of two independent sources: a label part (GNI and $\text{CGNI}^\wedge$, inside $P^{GI} \subseteq \text{Low}_2$) and a search part (the NP-hardness, outside Low_2 unless collapse). Rigidity removes the search floor and leaves only the label floor. More generally, a host restriction lowers BWG-NI into P^{GI} when it removes the search part, that is, when the per-edge disequations no longer couple across the host-automorphism action; triviality of the host automorphism group is one sufficient condition for this decoupling but not evidently the only one, since the search part requires only that the disequation constraints separate, not that the automorphism coset reduce to a single point.

8 Label alphabet and the complexity landscape

The per-edge comparison of BWG-NI is matrix isomorphism over a bounded integer alphabet, which is GI-complete by Fact 2.3; the monotone fragment is already complete and binary matrices suffice [ACL12]. One refinement follows.

Corollary 8.1 (monotone, two-symbol labels suffice). *BWG-NI restricted to monotone-DNF edge labels, equivalently $\{0,1\}$ -matrix labels, remains NP-hard, GNI-hard, in NP^{GI} , and not Σ_2^P -complete unless $PH = \Sigma_2^P$.*

Proof. The single-edge case is the negation of monotone DSFI and is GNI-complete since the monotone fragment is GI-complete [ACL12]. Membership is unchanged. The reduction of Theorem 5.1 already uses monotone labels: the single-monomial DNFs of Lemma 5.2 are pairwise non-DSFI-isomorphic, monotone, and polynomial-time constructible, so its NP-hardness holds verbatim for monotone-DNF labels, equivalently $\{0,1\}$ -matrices. Non- Σ_2^P -completeness follows from NP^{GI} membership as in Theorem 6.2. ■

Collapse level. If BWG-NI is Σ_2^P -complete, Lemma 6.1(i) gives $\Sigma_2^P = \text{NP}^{\text{GI}}$ and (ii) gives $\text{PH} = \Sigma_2^P$; no deeper collapse is evident: $\Sigma_2^P = \text{NP}^{\text{GI}}$ is not known to imply $\text{NP} = \text{coNP}$. Thus the collapse level obtained is Σ_2^P , one level below the level forced by Σ_2^P -completeness of FI, which is $\text{PH} = \Sigma_3^P$ [AT00, Corollary 3.4]. Inside the formula-isomorphism family the same sharpening from the third level to the second is achieved by the canonical restriction [Cri26, Theorem 7]. The analogy is one of collapse depth; the second-level statements concern different completeness notions, NP-completeness of CFI there and Σ_2^P -completeness of BWG-NI here. Agrawal and Thierauf [AT00] show FI is not Σ_2^P -complete unless PH collapses, via a one-round interactive proof for the complementary problem FNI with an NP oracle, i.e. $\text{FNI} \in \text{AM}^{\text{NP}} = \text{BP} \cdot \Sigma_2^P$; this places the FI collapse at the third level, by the relativised Boppana–Håstad–Zachos argument [BHZ87]. The second-level collapse available here is a consequence of $\text{BWG-NI} \in \text{NP}^{\text{GI}}$: FI is coNP -hard (Fact 2.4) and so lies outside NP^{GI} unless $\text{PH} = \Sigma_2^P$ (Lemma 6.1), leaving the GI-oracle route unavailable to it.

Complement asymmetry. In the complement every bijection must either fail to be a host isomorphism or match some edge pair of syntactically isomorphic labels; formally $\text{co-BWG-NI} = \forall f [f \text{ a graph isomorphism} \Rightarrow \exists e (\ell_G(e), \ell_H(f(e))) \in \text{DSFI}]$. The inner comparison is positive and guessable ($\text{DSFI} \in \text{NP}$), so $\text{co-BWG-NI} \in \Pi_2^P$ natively. The GI oracle is used only on the BWG-NI side.

NP-completeness under GI $\in$ coNP. BWG-NI is in NP iff $\text{GNI} \in \text{NP}$, i.e. $\text{GI} \in \text{coNP}$: the forward direction is the GNI-hardness of Proposition 4.2 ($\text{GNI} \leq_m^P \text{BWG-NI}$, and NP is closed under $\leq_m^P$), and for the converse, if $\text{GNI} \in \text{NP}$ an NP machine guesses the host bijection f and, for each edge, a witness that $\ell(e)$ and $\ell(f(e))$ are non-DSFI-isomorphic, each such test being DSNI and hence GNI-complete. Under that hypothesis BWG-NI is NP-complete. The canonical restriction of the semantic comparison meets the same equivalence: $\text{CFI} \in \text{NP} \cap \text{coNP}$ iff $\text{GI} \in \text{coNP}$, equivalently $\text{NP}^{\text{GI}} = \text{NP}$ [Cri26, Theorem 8]. $\text{GI} \in \text{coNP}$ follows from $\text{AM} = \text{NP}$ (via $\text{GNI} \in \text{AM}$) and a fortiori from $\text{GI} \in \text{P}$. Each such hypothesis places $\text{BWG-NI} \in \text{NP}$, hence it is NP-complete; the label-disequality variant is NP-complete with no hypothesis (Remark 8.2). The unconditional statement is: NP-hard, in NP^{GI} , not coNP -hard or Σ_2^P -complete unless $\text{PH} = \Sigma_2^P$. Under $\text{GI} \in \text{P}$, $\text{NP}^{\text{GI}} = \text{NP}$ and Theorem 7.6 reads “ $\text{BWG-NI} \in \text{P} \Rightarrow \text{P} = \text{NP}$ ”.

Figure 9 records the landscape under the hypothesis $\text{PH} \neq \Sigma_2^P$. BWG-NI lies in NP^{GI} but outside P^{GI} , Low_2 , and coNP (Theorems 5.1, 6.2, 7.6), and outside NP when in addition $\text{GI} \notin \text{coNP}$ ($\text{BWG-NI} \in \text{NP} \Leftrightarrow \text{GI} \in \text{coNP}$); FI lies in Σ_2^P but outside NP^{GI} , its coNP -hardness being what excludes it (Theorem 6.4, Lemma 6.1); BWG-ISO sits with GI and GNI inside $P^{\text{GI}} \subseteq \text{Low}_2$ (Theorem 3.5). FACT, the decision form of integer factorisation (given n and k in binary, whether n has a nontrivial factor at most k), lies in $\text{NP} \cap \text{coNP}$ and appears for orientation: $\text{FACT} \leq_m^P \text{BWG-NI}$ by NP-hardness, while the converse reduction would give $\text{NP} = \text{coNP}$. Under $\text{GI} \in \text{coNP}$, BWG-NI drops into NP and is NP-complete.

Remark 8.2 (string comparator; unconditional NP-completeness). With labels compared by plain string disequality, a per-edge test in P, BWG-NI is NP-complete with no hypothesis: membership needs no oracle, since a machine guesses the host bijection and checks the per-edge string comparisons in polynomial time, and the reduction of Theorem 5.1 applies verbatim, its single-monomial labels being pairwise distinct as strings, so identical labels compare equal and distinct labels compare unequal under either comparator.

9 Scope, provenance, and open problems

Provenance and verification. Among the results above, the one resting on a construction rather than a composition with classical facts is the NP-hardness (Theorem 5.1); its




Figure 9. The complexity landscape under $PH \neq \Sigma_2^P$; the placement of $BWG-NI$ (star) outside NP additionally assumes $GI \notin coNP$, since $BWG-NI \in NP \Leftrightarrow GI \in coNP$ (Section 8). $BWG-NI$ lies in $NP^{GI} \setminus (P^{GI} \cup NP \cup coNP)$ (Theorem 7.6(a) and (b), the $coNP$ exclusion via $coNP \subseteq coNP^{GI}$); FI lies in $\Sigma_2^P \setminus NP^{GI}$; GI and GNI lie in $P^{GI} \subseteq Low_2$; $BWG-ISO$ coincides with the GI dot (Theorem 3.5). $FACT \in NP \cap coNP$ is shown for orientation. Positions not implied by the stated results are conjectural.

consequences (the exclusions from P^{GI} , coNP^{GI} , and Low_2 by Theorem 7.6, the rigid/general distinction, the witness to $P^{\text{GI}} \neq \text{NP}^{\text{GI}}$ by Corollary 7.7, and the counting hardness by Proposition 5.7) follow from it together with [Lub81, Sch88, AK02, FFK94, Tod91]. The isomorphism-variant reduction (Theorem 3.5) is a second construction. The membership (Proposition 4.1), Lemma 6.1, and Proposition 7.3 are compositions of standard arguments. Three remarks. (i) The NP-hardness rests on the reduction of Theorem 5.1, which is new relative to the literature. Its combinatorial core has been machine-checked exhaustively (see below). The reduction uses unbounded-degree hosts deliberately, since a bounded-degree reduction risks vacuity, FPFAUT being possibly tractable on bounded-degree graphs, consistent with [Lub81]’s hard instances being unbounded-degree. The exhaustive check: over all 155 isolate-free graphs on at most six vertices, the equivalence of Theorem 5.1 and the count identity of Remark 5.6 hold without exception, the Example 5.5 values reproducing exactly. (ii) The conjunctive closure $\text{CGNI}^\wedge$ coincides with the class of sets $\leq_m^p$ -reducible to GNI by [KST93, Theorem 1.42] (Definition 7.1), so the rigid-host completeness and lowness are GNI-level results and the conjunctive framing is retained only as notation; its relation to the bounded-query-over-GI literature [Wag90, AK00, LT92] is expository. (iii) Fact 2.3 is used under a fixed reading of the attachment clause of [ACL12], which says only that each entry-vertex is linked to the clique codifying its value: here each entry-vertex joins a designated vertex of that clique by a single edge; the same clause form governs the marker attachments, read here as joining each row-vertex to both vertices of the two-vertex marker clique, the one-vertex column marker being unambiguous ([ACL12] write clique of degree k for the clique with k vertices). Under this reading, graph isomorphism of the resulting gadgets and syntactic DNF isomorphism induce the same partition on all 45,759 DNFs with at most three variables and at most three term-occurrences, repeated terms included (checked exhaustively), so the concrete description in Fact 2.3 rests on [ACL12, Theorem 2] together with that verification. The attribution of the subdivision-plus-pinning device to [HSV01] is given as a comparison: that paper states its constructions for involutory fixed-edge-free automorphisms, whereas the present reduction uses general fixed-point-free automorphisms.

Tractable host restrictions. Rigidity places BWG-NI in P^{GI} by trivialising the host-automorphism search (Theorem 7.2). Any host restriction with the same effect is necessarily non-generic: if restricting the host to some class keeps BWG-NI in P^{GI} , then a many-one reduction of general BWG-NI to that restricted problem would place an NP-hard problem in $P^{\text{GI}} \subseteq \text{Low}_2$ (Theorem 5.1, [Sch88]) and collapse PH, so no tractable host restriction simulates the general problem under $\leq_m^p$ unless the hierarchy collapses. It is open which structural parameter of the host governs membership in P^{GI} (equivalently, which parameter controls whether the per-edge disequations decouple across the host-automorphism action) and how far beyond rigidity tractability extends.

Bounded-degree host-automorphism search. The NP-hardness reduction uses unbounded-degree hosts. Whether BWG-NI restricted to bounded-degree hosts is NP-hard depends on whether FPFAUT is NP-hard on bounded-degree graphs, a question about the host-automorphism search to which the group-theoretic machinery for bounded-valence isomorphism [Luk82] applies but which is not settled here.

Precomputable versus full oracle access. Whether $\text{NP}_{\text{pre}}^{\text{GI}} = \text{NP}^{\text{GI}}$ (whether host-automorphism structure can combine precomputed GI facts into genuinely new queries) is open. BWG-NI lies in $\text{NP}_{\text{pre}}^{\text{GI}}$. Two consequences follow. $\text{NP}_{\text{pre}}^{\text{GI}}$ is closed downward under $\leq_m^p$, the query set on x being the query set on $g(x)$ for a reduction g , so BWG-NI, lying in $\text{NP}_{\text{pre}}^{\text{GI}}$, is $\leq_m^p$ -complete for NP^{GI} only if $\text{NP}_{\text{pre}}^{\text{GI}} = \text{NP}^{\text{GI}}$. The question also trivialises under derandomisation: if $\text{AM} = \text{NP}$ then $\text{GNI} \in \text{NP}$, so $\text{GI} \in \text{coNP}$, $\text{NP}^{\text{GI}} = \text{NP} = \text{NP}_{\text{pre}}^{\text{GI}}$, and

BWG-NI is NP-complete by Section 8; hence $\text{NP}_{\text{pre}}^{\text{GI}} \neq \text{NP}^{\text{GI}}$ implies $\text{AM} \neq \text{NP}$.

References

- [ACL12] G. Ausiello, F. Cristiano, and L. Laura. Syntactic Isomorphism of CNF Boolean Formulas is Graph Isomorphism Complete. ECCC, Report No. TR12-122, 2012.
- [AK00] V. Arvind and J. Köbler. Graph isomorphism is low for $\text{ZPP}(\text{NP})$ and other lowness results. Proc. STACS 2000, LNCS 1770, Springer, 2000.
- [AK02] V. Arvind and P. P. Kurur. Graph Isomorphism is in SPP. Proc. FOCS 2002; Information and Computation 204(5):835–852, 2006.
- [AT00] M. Agrawal and T. Thierauf. The formula isomorphism problem. SIAM J. Comput. 30(3):990–1009, 2000.
- [BHZ87] R. Boppana, J. Håstad, and S. Zachos. Does co-NP have short interactive proofs? Inform. Process. Lett. 25:127–132, 1987.
- [BR93] B. Borchert and D. Ranjan. The Circuit Subfunction Relations are Σ_2^P -complete. Technical Report MPI-I-93-121, Max-Planck-Institut für Informatik, Saarbrücken, 1993.
- [BRS98] B. Borchert, D. Ranjan, and F. Stephan. On the computational complexity of some classical equivalence relations on Boolean functions. Theory Comput. Systems 31(6):679–693, 1998.
- [Cha89] R. Chang. On the structure of bounded queries to arbitrary NP sets. Proc. 4th Structure in Complexity Theory Conf., IEEE, 1989, 250–258.
- [Cri26] F. Cristiano. The complexity landscape of Boolean formula isomorphism: from graph isomorphism to the second level. Electronic Colloquium on Computational Complexity (ECCC), Report TR26-114, 2026.
- [FFK94] S. Fenner, L. Fortnow, and S. Kurtz. Gap-definable counting classes. J. Comput. System Sci. 48(1):116–148, 1994.
- [GMW91] O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. J. ACM 38(3):690–728, 1991.
- [HSV01] F. Harary, W. Slany, and O. Verbitsky. A Symmetric Strategy in Graph Avoidance Games. arXiv:cs/0110049; DBAI-TR-2001-42, TU Wien, 2001. Also in: More Games of No Chance, MSRI Publications 42, Cambridge University Press, 2002.
- [KST93] J. Köbler, U. Schöning, and J. Torán. The Graph Isomorphism Problem: Its Structural Complexity. Birkhäuser, Boston, 1993.
- [LT92] A. Lozano and J. Torán. On the nonuniform complexity of the graph isomorphism problem. Proc. 7th Structure in Complexity Theory, 1992, pp. 118–131.

- [Lub81] A. Lubiw. Some NP-complete problems similar to graph isomorphism. *SIAM J. Comput.* 10(1):11–21, 1981.
- [Luk82] E. M. Luks. Isomorphism of graphs of bounded valence can be tested in polynomial time. *J. Comput. System Sci.* 25(1):42–65, 1982.
- [Mat79] R. Mathon. A note on the graph isomorphism counting problem. *Inform. Process. Lett.* 8:131–132, 1979.
- [Sch88] U. Schöning. Graph isomorphism is in the low hierarchy. *J. Comput. System Sci.* 37:312–323, 1988.
- [Tod91] S. Toda. PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.* 20(5):865–877, 1991.
- [VV86] L. Valiant and V. Vazirani. NP is as easy as detecting unique solutions. *Theoret. Comput. Sci.* 47:85–93, 1986.
- [Wag90] K. W. Wagner. Bounded query classes. *SIAM J. Comput.* 19(5):833–846, 1990.