

Recognition of Graphs with Triangle-Free Treewidth-Two Thin Models is NP-Complete

1 September 2026

Abstract

We study recognition of graphs that admit a thick model whose hidden thin graph is triangle-free and has treewidth at most two. The model consists of a surjection onto the thin graph, with clique fibers and a one-way edge implication; in particular, thin edges need not be witnessed by cross-fiber edges of the input graph. We prove that the problem is NP-complete by a polynomial Karp reduction from 2-colorability of 3-uniform hypergraphs. The reduction uses two large cliques joined by a matching and a three-vertex path for each hyperedge. Triangle-freeness forces the two cliques either to be unsplit or to split into aligned pairs, while each hyperedge path rules out the unsplit case and monochromatic triples. A valid coloring yields a quotient assembled from a base 4-cycle and local 4-cycles or $K_{2,3}$ graphs, giving an explicit width-two decomposition. More generally, for each fixed $k \geq 2$ the same construction proves NP-completeness of $\text{thick}(\mathcal{W}_k \cap \mathcal{B})$ and $\text{thick}(\mathcal{W}_k \cap \mathcal{T})$; the fixed-width boundary is polynomial for $k \leq 1$ and NP-complete for every $k \geq 2$.

1 Introduction

Let G be a finite simple graph. In a thick model, vertices of G are grouped into cliques and the groups are mapped to vertices of another, hidden graph H . The recognition question asks whether such a grouping exists when H is required to belong to a specified graph class. This perspective is useful precisely because H and the grouping are not supplied with the input and may be highly non-unique.

We settle the problem posed as P2113: whether

$$G \in \text{thick}(\mathcal{W}_2 \cap \mathcal{T})$$

can be recognized in polynomial time, where $\mathcal{T}$ is the class of triangle-free graphs and $\mathcal{W}_k$ is the class of graphs of treewidth at most k . The question was listed as open in Section 7.2 of Dyer and Müller[6]. Our main result is the following.

Theorem 1.1 (Main theorem). *Recognition of $\text{thick}(\mathcal{W}_2 \cap \mathcal{T})$ is NP-complete under polynomial-time many-one (Karp) reductions, for the usual explicit adjacency-list (equivalently, adjacency-matrix) encoding of finite simple graphs.*

The proof has three features that are important for the exact model. First, we replace an arbitrary thin witness by the support quotient of its fibers; this is legitimate because triangle-freeness and bounded treewidth are preserved when edges are deleted. Second, the soundness argument uses only triangle-freeness, so it remains valid for several larger target classes. Third, the

positive quotient is not merely width two: it is bipartite and is formed by gluing short pieces along one interface edge. Consequently the same reduction yields a fixed-treewidth dichotomy.

The paper is organized as follows. Section 2 fixes the model and proves the support-quotient normal form. Sections 3–5 give the reduction and its two directions. Section 7 records the extensions and the $k = 0, 1, 2$ boundary. We close with limitations and reproducibility notes.

2 Model and support quotients

All graphs in this paper are finite, undirected, and simple.

Definition 2.1 (Thick model). Let G and H be graphs. A thick model of G with thin graph H is a surjection $\psi : V(G) \rightarrow V(H)$ such that

- (i) each fiber $\psi^{-1}(h)$ induces a clique in G ; and
- (ii) for every $uv \in E(G)$, either $\psi(u) = \psi(v)$ or $\psi(u)\psi(v) \in E(H)$.

The second condition is one-way: an edge of H need not be realized by a cross-fiber edge of G .

For a graph class $\mathcal{C}$, write $\text{thick}(\mathcal{C})$ for the class of graphs having a thick model with $H \in \mathcal{C}$. We use $\mathcal{T}$ for triangle-free graphs, $\mathcal{B}$ for bipartite graphs, and $\mathcal{W}_k = \{H : \text{tw}(H) \leq k\}$. We use the standard elimination characterization of treewidth: $\text{tw}(H) \leq k$ if and only if some ordering of $V(H)$, after fill edges are added among later neighbors at each step, eliminates every vertex with at most k later neighbors.

Definition 2.2 (Support quotient). Let Π be a partition of $V(G)$ into nonempty blocks, each of which is a clique in G . Its support quotient Q_Π has one vertex for each block, and two quotient vertices are adjacent exactly when some edge of G has one endpoint in each corresponding block.

Proposition 2.3 (Normal form). *Let $\mathcal{C}$ be closed under deleting edges. Then*

$$G \in \text{thick}(\mathcal{C}) \iff \exists \Pi [\Pi \text{ is a clique partition of } V(G) \text{ and } Q_\Pi \in \mathcal{C}].$$

Proof. Suppose (H, ψ) is a thick model and let Π be its fibers. Since every cross-fiber edge of G maps to an edge of H , Q_Π is a spanning subgraph of H . Closure under edge deletion gives $Q_\Pi \in \mathcal{C}$. Conversely, if $Q_\Pi \in \mathcal{C}$, use Q_Π itself as the thin graph and map each input vertex to its block. The clique and edge conditions are exactly the definitions of Π and Q_Π . $\square$

For P2113, Proposition 2.3 applies because both triangle-freeness and treewidth at most two are preserved by edge deletion. It also explains why no completeness assumption is made about a pair of fibers: only the existence of a cross-fiber edge is recorded in Q_Π .

3 Source problem and construction

We reduce from 3-Set Splitting, also called 2-colorability of a 3-uniform hypergraph or PROPERTY-B:

Given $U = \{1, \dots, n\}$ and a list F of 3-element subsets of U , decide whether there is a coloring $c : U \rightarrow \{0, 1\}$ such that no member of F is monochromatic.

The problem is NP-complete. This is the standard positive NAE-3SAT constraint-satisfaction problem (also called 3-Set Splitting or Property B); its NP-completeness follows from Schaefer's dichotomy[2]. The unrestricted Set Splitting problem is also catalogued by Garey–Johnson[5]. For an explicit treatment of the corresponding 3-Set Splitting optimization problem, see Kann–Lagergren–Panconesi[4]; related hardness results for 3-uniform hypergraph coloring appear in Dinur–Regev–Smyth[3]. Lovász's early work gives useful historical combinatorial background[1]; it is not needed for the precise complexity attribution here.

We use an explicit encoding listing the vertices and triples, with three distinct elements in every triple. Delete isolated elements first; this preserves colorability because colors extend arbitrarily over isolated vertices. If F is empty, add a fresh triple on three fresh elements; both instances are always splittable. Otherwise $n \leq 3|F|$, so the quadratic clique part of the construction below is polynomial in the source encoding. Thus we may assume $n \geq 3$ and F is nonempty.

Let F have m triples. Construct G_F as follows:

- $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_n\}$ are cliques;
- add the matching edges $a_i b_i$ for $1 \leq i \leq n$;
- for each $e = (i, j, k) \in F$, add vertices x_e, y_e, z_e and edges $x_e y_e, y_e z_e, b_i x_e, b_j y_e$, and $b_k z_e$;
- add no other edges.

The graph has $2n + 3m$ vertices and

$$2\binom{n}{2} + n + 5m = n(n-1) + n + 5m$$

edges.

4 The base partition lemma

The following lemma is the rigidity point of the reduction.

Lemma 4.1 (Base partition). *Let Π be a clique partition of G_F whose support quotient is triangle-free. Restrict the blocks to $A \cup B$ and delete empty intersections. Exactly one of the following holds:*

- (1) *all of A is one block and all of B is one block (the unsplit case); or*
- (2) *A is split into two blocks A_0, A_1 , B is split into two blocks B_0, B_1 , and, after swapping labels, a_i and b_i lie in corresponding blocks for every i .*

No block contains both an a_i and a b_i .

Proof. Any two distinct blocks meeting A are adjacent in the quotient because A is a clique, so triangle-freeness permits at most two such blocks. The same holds for B .

Suppose a block X contains a_i and b_i . For $t \neq i$, let P_t and R_t contain a_t and b_t . Neither is X . If $P_t \neq R_t$, the three blocks X, P_t, R_t form a quotient triangle using, respectively, $a_i a_t, a_t b_t$, and $b_t b_i$. Hence $P_t = R_t$. For two distinct $t, s \neq i$, these mixed blocks are distinct (otherwise one block contains the nonadjacent vertices a_t, b_s), and together with X again form a triangle through the two large cliques. Since $n \geq 3$, such t, s exist; there are no mixed blocks.

Let p and q be the numbers of blocks meeting A and B . They are each one or two. The cases $(p, q) = (1, 2)$ and $(2, 1)$ are impossible: the single block on one side is adjacent by matching edges to both blocks on the other side, and those two blocks are adjacent inside their clique. Thus $(p, q) = (1, 1)$ or $(2, 2)$. In the latter case, an A block cannot contain indices whose matched B vertices lie in both B_0 and B_1 , since those three quotient vertices would form a triangle. Each A block therefore maps wholly to one B block. The map is onto and both sides have two blocks, so it is a bijection; relabeling gives the alignment. $\square$

When restricting a partition to $A \cup B$, its quotient is a subgraph of the original quotient, so it remains triangle-free. This is the only fact about the auxiliary vertices needed in Lemma 4.1.

5 Soundness and completeness

5.1 Soundness

Assume G_F has a model with a triangle-free thin graph of treewidth at most two. By Proposition 2.3, take a clique partition Π whose support quotient is triangle-free, and apply Lemma 4.1.

The unsplit case is impossible. Let D be the block containing all of B , and fix $e = (i, j, k)$. Each of x_e, y_e, z_e is nonadjacent to at least two of the three distinct vertices b_i, b_j, b_k in D , so none lies in D ; each has a quotient edge to D . If the blocks containing x_e and y_e were distinct, their edge $x_e y_e$ together with D would make a triangle. They are therefore equal. The same argument for $y_e z_e$ puts all three vertices in one block, contradicting the nonedge $x_e z_e$ inside a clique block.

We are in the aligned case. Color i by $c(i) = t$ when $a_i \in A_t$ (equivalently $b_i \in B_t$). If a triple $e = (i, j, k)$ were monochromatic of color t , the block B_t would contain b_i, b_j, b_k . None of the three auxiliary vertices can join B_t , but each has a quotient edge to it. The two path edges then force the blocks of x_e, y_e, z_e to coincide, again contradicting $x_e z_e \notin E(G_F)$. Thus c is a valid Property-B coloring.

Notice that auxiliary vertices may join other blocks whenever the edge condition permits it; the argument only uses the relevant B block and does not assume a canonical partition.

5.2 Completeness

Conversely, let c be a valid Property-B coloring and set

$$A_t = \{a_i : c(i) = t\}, \quad B_t = \{b_i : c(i) = t\} \quad (t = 0, 1).$$

Both color classes are nonempty because F is nonempty and every triple is nonmonochromatic. These four sets are clique blocks. For each triple, use the following additional blocks:

color pattern	auxiliary blocks	quotient gadget
$c(i) = c(j) = t \neq c(k)$	$\{x_e, y_e\}, \{z_e\}$	C_4
$c(j) = c(k) = t \neq c(i)$	$\{y_e, z_e\}, \{x_e\}$	C_4
$c(i) = c(k) = t \neq c(j)$	$\{x_e\}, \{y_e\}, \{z_e\}$	$K_{2,3}$

The base quotient on A_0, A_1, B_0, B_1 is the cycle

$$A_0 - A_1 - B_1 - B_0 - A_0.$$

For either C_4 row, writing $U = B_t$ and $V = B_{1-t}$, the local quotient is $U - P - Q - V - U$. For the third row it has edges

$$U - V, U - X, U - Z, V - Y, X - Y, Y - Z,$$

which is $K_{2,3}$ with bipartition $\{U, Y\}$ and $\{V, X, Z\}$. Distinct gadgets share only the interface blocks U, V and have no edges between their auxiliary vertices. A triangle using two gadgets is therefore impossible; a triangle using at most one auxiliary vertex would require that vertex to be adjacent to both interfaces, which never occurs. The quotient is consequently triangle-free.

Here is an explicit width-two tree decomposition. Start with bags

$$S_0 = \{A_0, A_1, B_0\}, \quad S_1 = \{A_1, B_0, B_1\}.$$

Join S_0 and S_1 by a tree edge. For a C_4 gadget attach, as a path to S_1 , the bags

$$\{U, V, Q\}, \quad \{U, Q, P\}.$$

For a $K_{2,3}$ gadget attach to S_1 the star with center

$$\{U, V, Y\}$$

and leaves $\{U, Y, X\}$ and $\{U, Y, Z\}$. Since S_1 contains both B_0 and B_1 , all bags containing either interface remain connected. Every edge is covered and every vertex has a connected bag subtree; all bags have size at most three. Hence the quotient has treewidth at most two. By Proposition 2.3, this is a valid P2113 model.

6 Membership and the main conclusion

Let $N = |V(G)|$. A certificate labels each input vertex by an element of $[N]$ (the fiber labels) and supplies a permutation of the nonempty quotient vertices as an elimination ordering. Unused labels are ignored. The verifier checks that every nonempty label class is a clique, constructs its support quotient, and checks that no triangle occurs. It then checks the proposed width-two ordering by maintaining the fill graph, rejecting any step with more than two later neighbors and adding fill edges among those neighbors. The quotient has at most N vertices and the whole procedure is polynomial (a direct adjacency-matrix implementation is $O(N^3)$). Thus P2113 is in NP. Sections 3–5 give a polynomial Karp reduction from an NP-complete problem, proving Theorem 1.1.

7 Consequences

Corollary 7.1 (Larger target classes). *Recognition of $\text{thick}(\mathcal{T})$ is NP-complete. For every fixed integer $k \geq 2$, recognition of $\text{thick}(\mathcal{W}_k \cap \mathcal{T})$ and of $\text{thick}(\mathcal{W}_k \cap \mathcal{B})$ is NP-complete.*

Proof. The classes $\mathcal{T}$, $\mathcal{B}$, and $\mathcal{W}_k$ are closed under edge deletion, so Proposition 2.3 gives NP certificates (add a width- k elimination ordering for $\mathcal{W}_k$ and a two-coloring for $\mathcal{B}$). The quotient constructed from a Property-B coloring is bipartite and has treewidth two: the base and every local piece are bipartite, and all pieces are glued along the same interface edge B_0B_1 . It is therefore a valid thin graph for all these classes when $k \geq 2$. Conversely, every thin graph in $\mathcal{T}$ or $\mathcal{B}$ is triangle-free, so the soundness proof recovers a Property-B coloring. $\square$

Corollary 7.2 (Fixed-treewidth dichotomy). *For each $\mathcal{C} \in \{\mathcal{T}, \mathcal{B}\}$ and fixed k , recognition of $\text{thick}(\mathcal{W}_k \cap \mathcal{C})$ is polynomial-time solvable for $k \leq 1$ and NP-complete for every $k \geq 2$.*

Proof. $\mathcal{W}_0$ consists of edgeless graphs. A thick model over an edgeless thin graph exists exactly when every connected component of G is a clique, that is, when G is a cluster graph; this is directly recognizable in polynomial time. Since $\mathcal{W}_1$ is the class of forests and every forest is bipartite (and

triangle-free), both $\mathcal{W}_1 \cap \mathcal{T}$ and $\mathcal{W}_1 \cap \mathcal{B}$ are the forest class. The case $k = 1$ is therefore the thick-forest recognition problem, which has a polynomial-time algorithm by Dyer and Müller[6]. The cases $k \geq 2$ are the preceding corollary. $\square$

Remark 7.3 (Extra structure of the witnesses). *The explicit positive quotients are bipartite and have the dominating edge B_0B_1 : every vertex outside its endpoints is adjacent to at least one of them. Moreover, they have no induced cycle longer than four, so the witnesses are chordal-bipartite. Indeed, every piece (including the base) has terminals B_0, B_1 and only 4-cycles; a simple cycle meeting two pieces uses two internally disjoint terminal paths and therefore has the extra edge B_0B_1 as a chord. These properties are observations about the particular witnesses only; they are not used in the normal-form argument, since chordal-bipartiteness and the existence of a dominating edge are not preserved by arbitrary edge deletion.*

8 Scope and limitations

The NP-completeness result concerns unrestricted input graphs and does not rule out polynomial subclasses or useful parameterized variants. The reduction leaves open the complexity of restricted input classes and finer parameterized questions.

The reduction is intentionally economical: it uses only cliques, a matching, and constant-size path gadgets. No numerical optimization, random sampling, or large simulation is part of the proof. The remaining algorithmic question is therefore not the existence of a model, which is settled here, but whether additional input restrictions or parameterizations admit practical algorithms despite the unrestricted barrier.

9 Reproducibility

The repository contains small, deterministic regression checks. With NetworkX 3.6.1 installed, run:

```
python3 -B experiments/verify_reduction.py
python3 -B experiments/verify_small.py
python3 -B experiments/verify_small.py --atlas
```

The first script checks the base-partition lemma on small instances and the explicit quotient decompositions. The second checks structural families; its optional atlas mode exhausts all nonempty unlabeled graphs through seven vertices. These computations are sanity checks only and are not used as a substitute for the proofs.

Acknowledgements

The formulation follows the terminology of Dyer and Müller [6].

References

- [1] L. Lovász, *Coverings and colorings of hypergraphs*, Proceedings of the 4th Southeastern Conference on Combinatorics, Graph Theory and Computing, Utilitas Mathematica, 1973, pp. 3–12.
- [2] T. J. Schaefer, *The complexity of satisfiability problems*, Proceedings of the 10th ACM Symposium on Theory of Computing (STOC), 1978, pp. 216–226. DOI: 10.1145/800133.804350.

- [3] I. Dinur, O. Regev, and C. Smyth, *The hardness of 3-uniform hypergraph coloring*, *Combinatorica* 25(5) (2005), 519–535. DOI: 10.1007/s00493-005-0032-4.
- [4] V. Kann, J. Lagergren, and A. Panconesi, *Approximability of maximum splitting of k -sets and some other APX-complete problems*, *Information Processing Letters* 58 (1996), 105–110. DOI: 10.1016/0020-0190(96)00046-4.
- [5] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, 1979.
- [6] M. Dyer and H. Müller, *Thick forests*, *Discrete Applied Mathematics* 384 (2026), 177–200. See also arXiv:2309.01482 and DOI 10.1016/j.dam.2025.12.042.