


Lovász Meets Weisfeiler and Leman


Holger Dell

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

 <https://orcid.org/0000-0001-8955-0786>


Martin Grohe

RWTH Aachen University, Aachen, Germany
grohe@informatik.rwth-aachen.de

 <https://orcid.org/0000-0002-0292-9142>

Gaurav Rattan

RWTH Aachen University, Aachen, Germany
grohe@informatik.rwth-aachen.de

 <https://orcid.org/0000-0002-5095-860X>

Abstract

In this paper, we relate a beautiful theory by Lovász with a popular heuristic algorithm for the graph isomorphism problem, namely the color refinement algorithm and its k -dimensional generalization known as the Weisfeiler-Leman algorithm. We prove that two graphs G and H are indistinguishable by the color refinement algorithm if and only if, for all trees T , the number $\text{Hom}(T, G)$ of homomorphisms from T to G equals the corresponding number $\text{Hom}(T, H)$ for H .

There is a natural system of linear equations whose nonnegative integer solutions correspond to the isomorphisms between two graphs. The nonnegative real solutions to this system are called fractional isomorphisms, and two graphs are fractionally isomorphic if and only if the color refinement algorithm cannot distinguish them (Tinhofer 1986, 1991). We show that, if we drop the nonnegativity constraints, that is, if we look for arbitrary real solutions, then a solution to the linear system exists if and only if, for all t , the two graphs have the same number of length- t walks.

We lift the results for trees to an equivalence between numbers of homomorphisms from graphs of tree width k , the k -dimensional Weisfeiler-Leman algorithm, and the level- k Sherali-Adams relaxation of our linear program. We also obtain a partial result for graphs of bounded path width and solutions to our system where we drop the nonnegativity constraints.

A consequence of our results is a quasi-linear time algorithm to decide whether, for two given graphs G and H , there is a tree T with $\text{Hom}(T, G) \neq \text{Hom}(T, H)$.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis, Mathematics of computing → Graph theory

Keywords and phrases graph isomorphism, graph homomorphism numbers, tree width

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.40

Related Version A full version of the paper is available at <https://arxiv.org/abs/1802.08876>.

1 Introduction

An old result due to Lovász [16] states a graph G can be characterized by counting homomorphisms from all graphs F to G . That is, two graphs G and H are isomorphic if and only if, for all F , the number $\text{Hom}(F, G)$ of homomorphisms from F to G equals the number $\text{Hom}(F, H)$ of homomorphism from F to H . This simple result has far reaching consequences,



© Holger Dell, Martin Grohe, and Gaurav Rattan;
licensed under Creative Commons License CC-BY

45th International Colloquium on Automata, Languages, and Programming (ICALP 2018).
Editors: Ioannis Chatzigiannakis, Christos Kaklamani, Dániel Marx, and Donald Sannella;
Article No. 40; pp. 40:1–40:14



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



because mapping graphs G to their *homomorphism vectors* $\text{HOM}(G) := (\text{Hom}(F, G))_{F \text{ graph}}$ (or suitably scaled versions of these infinite vectors) allows us to apply tools from functional analysis in graph theory. This is the foundation of the beautiful theory of graph limits, developed by Lovász and others over the last 15 years (see [17]).

However, from a computational perspective, representing graphs by their homomorphism vectors has the disadvantage that the problem of computing the entries of these vectors is NP-complete. To avoid this difficulty, we may want to restrict the homomorphism vectors to entries from a class of graphs for which counting homomorphisms is tractable. That is, instead of considering the full homomorphism vector $\text{HOM}(G)$ we consider the vector $\text{HOM}_{\mathcal{F}}(G) := (\text{Hom}(F, G))_{F \in \mathcal{F}}$ for a class \mathcal{F} of graphs such that the problem of computing $\text{Hom}(F, G)$ for given graphs $F \in \mathcal{F}$ and G is in polynomial time. Arguably the most natural example of such a class \mathcal{F} is the class of all trees. More generally, computing $\text{Hom}(F, G)$ for given graphs $F \in \mathcal{F}$ and G is in polynomial time for all classes \mathcal{F} of bounded tree width, and under a natural assumption from parameterized complexity theory, it is not in polynomial time for any class \mathcal{F} of unbounded tree width [10]. This immediately raises the question what the vector $\text{HOM}_{\mathcal{F}}(G)$, for a class \mathcal{F} of bounded tree width, tells us about the graph G .

A first nice example (Proposition 9) is that the vector $\text{HOM}_{\mathcal{C}}(G)$ for the class \mathcal{C} of all cycles characterizes the spectrum of a graph, that is, for graphs G, H we have $\text{HOM}_{\mathcal{C}}(G) = \text{HOM}_{\mathcal{C}}(H)$ if and only if the adjacency matrices of G and H have the same eigenvalues with the same multiplicities. This equivalence is a basic observation in spectral graph theory (see [23, Lemma 1]). Before we state deeper results along these lines, let us describe a different (though related) motivation for this research.

Determining the similarity between two graphs is an important problem with many applications, mainly in machine learning, where it is known as “graph matching” (e.g. [9]). But how can the similarity between graphs be measured? An obvious idea is to use the *edit distance*, which simply counts how many edges and vertices have to be deleted from or added to one graph to obtain the other. However, two graphs that have a small edit distance can nevertheless be structurally quite dissimilar (e.g. [17, Section 1.5.1]). The edit distance is also very hard to compute as it is closely related to the notoriously difficult quadratic assignment problem (e.g. [3, 19]).

Homomorphism vectors offer an alternative, more structurally oriented approach to measuring graph similarity. After suitably scaling the vectors, we can compare them using standard vector norms. This idea is reminiscent of the “graph kernels” used in machine learning (e.g. [24]). Like the homomorphism vectors, many graph kernels are based on the idea of counting certain patterns in graphs, such as paths, walks, cycles or subtrees, and in fact any inner product on the homomorphism vectors yields a graph kernel.

A slightly different type of graph kernel is the so-called Weisfeiler-Leman (subtree) kernel [20]. This kernel is derived from the *color refinement* algorithm (a.k.a. the *1-dimensional Weisfeiler-Leman algorithm*), which is a simple and efficient heuristic to test whether two graphs are isomorphic (e.g. [11]). The algorithm computes a coloring of the vertices of a graph based on the iterated degree sequences, we give the details in Section 3. To use it as an isomorphism test, we compare the color patterns of two graphs. If they are different, we say that color refinement *distinguishes* the graphs. If the color patterns of the two graphs turn out to be the same, the graphs may still be non-isomorphic, but the algorithm fails to detect this.

Whether color refinement is able to distinguish two graphs G and H has a very nice linear-algebraic characterization due to Tinhofer [21, 22]. Let V and W be the vertex sets and let $A \in \{0, 1\}^{V \times V}$ and $B \in \{0, 1\}^{W \times W}$ be the adjacency matrices of G and H , respectively.

Now consider the system $F_{\text{iso}}(G, H)$ of linear equations:

$$F_{\text{iso}}(G, H) : \begin{cases} AX = XB & \text{(F1)} \\ X\mathbf{1}_W = \mathbf{1}_V & \text{(F2)} \\ \mathbf{1}_V^T X = \mathbf{1}_W^T & \text{(F3)} \end{cases}$$

In these equations, X denotes a $(V \times W)$ -matrix of variables and $\mathbf{1}_U$ denotes the all-1 vector over the index set U . Equations (F2) and (F3) simply state that all row and column sums of X are supposed to be 1. Thus the nonnegative integer solutions to $F_{\text{iso}}(G, H)$ are permutation matrices, which due to (F1) describe isomorphisms between G and H . The nonnegative real solutions to $F_{\text{iso}}(G, H)$, which in fact are always rational, are called *fractional isomorphisms* between G and H . Tinhofer proved that two graphs are fractionally isomorphic if and only if color refinement does not distinguish them.

For every $k \geq 2$, color refinement has a generalization, known as the *k-dimensional Weisfeiler-Leman algorithm (k-WL)*, which colors not the vertices of the given graph but k -tuples of vertices. Atserias and Maneva [4] (also see [18]) generalized Tinhofer's theorem by establishing a close correspondence between k -WL and the level- k Sherali-Adams relaxation of $F_{\text{iso}}(G, H)$.

Our results

How expressive are homomorphism vectors $\text{HOM}_{\mathcal{F}}(G)$ for restricted graph classes \mathcal{F} ? We consider the class \mathcal{T} of trees first, where the answer is surprisingly clean.


► **Theorem 1.** *For all graphs G and H , the following are equivalent:*

- i $\text{HOM}_{\mathcal{T}}(G) = \text{HOM}_{\mathcal{T}}(H)$.
- ii *Color refinement does not distinguish G and H .*
- iii *G and H are fractionally isomorphic, that is, the system $F_{\text{iso}}(G, H)$ of linear equations has a nonnegative real solution.*

As mentioned before, the equivalence between ii and iii is due to Tinhofer [21, 22]. An unexpected consequence of our theorem is that we can decide in time $O((n+m) \log n)$ whether $\text{HOM}_{\mathcal{T}}(G) = \text{HOM}_{\mathcal{T}}(H)$ holds for two given graphs G and H with n vertices and m edges. (If two graphs have a different number of vertices or edges, then their homomorphism counts already differ on the 1-vertex or 2-vertex trees.) This is remarkable, because every known algorithm for computing the entry $\text{Hom}(T, G)$ of the vector $\text{HOM}_{\mathcal{T}}(G)$ requires quadratic time when T and G are given as input.

It is a consequence of the proof of Theorem 1 that, in order to characterize an n -vertex graph G up to fractional isomorphisms, it suffices to restrict the homomorphism vector $\text{HOM}_{\mathcal{T}}(G)$ to trees of height at most $n - 1$. What happens if we restrict the structure of trees even further? In particular, let us restrict the homomorphism vector to its path entries, that is, consider $\text{HOM}_{\mathcal{P}}(G)$ for the class \mathcal{P} of all paths. Figure 1 shows an example of two graphs G and H with $\text{HOM}_{\mathcal{P}}(G) = \text{HOM}_{\mathcal{P}}(H)$ and $\text{HOM}_{\mathcal{T}}(G) \neq \text{HOM}_{\mathcal{T}}(H)$.

Despite their weaker distinguishing capabilities, the vectors $\text{HOM}_{\mathcal{P}}(G)$ are quite interesting. They are related to graph kernels based on counting walks, and they have a clean algebraic description: it is easy to see that $\text{Hom}(P_k, G)$, the number of homomorphisms from the path P_k of length k to G , is equal to the number of length- k walks in G , which in turn is equal to $\mathbf{1}^T A^k \mathbf{1}$, where A is the adjacency matrix of G and $\mathbf{1}$ is the all-1 vector of appropriate length.



■ **Figure 1** Two fractionally non-isomorphic graphs with the same path homomorphism counts.

► **Theorem 2.** *For all graphs G and H , the following are equivalent:*

- i $\text{HOM}_{\mathcal{P}}(G) = \text{HOM}_{\mathcal{P}}(H)$.
- ii The system $F_{\text{iso}}(G, H)$ of linear equations has a real solution.

While the proof of Theorem 1 is mainly graph-theoretic—we establish the equivalence between the assertions i and ii by expressing the “colors” of color refinement in terms of specific tree homomorphisms—the proof of Theorem 2 is purely algebraic. We use spectral techniques, but with a twist, because neither does the spectrum of a graph G determine the vector $\text{HOM}_{\mathcal{P}}(G)$ nor does the vector determine the spectrum. This is in contrast with $\text{HOM}_{\mathcal{C}}(G)$ for the class \mathcal{C} of all cycles, which, as we already mentioned, distinguishes two graphs if and only if they have the same spectrum.

Let us now turn to homomorphism vectors $\text{HOM}_{\mathcal{T}_k}(G)$ for the class \mathcal{T}_k of all graphs of tree width at most k . We will relate these to k -WL, the k -dimensional generalization of color refinement. We also obtain a corresponding system of linear equations. Let G and H be graphs with vertex sets V and W , respectively. Instead of variables X_{vw} for vertex pairs $(v, w) \in V \times W$, as in the system $F_{\text{iso}}(G, H)$, the new system has variables X_{π} for $\pi \subseteq V \times W$ of size $|\pi| \leq k$. We call $\pi = \{(v_1, w_1), \dots, (v_{\ell}, w_{\ell})\} \subseteq V \times W$ a *partial bijection* if $v_i = v_j \iff w_i = w_j$ holds for all i, j , and we call it a *partial isomorphism* if in addition $v_i v_j \in E(G) \iff w_i w_j \in E(H)$ holds for all i, j . Now consider the following system $L_{\text{iso}}^k(G, H)$ of linear equations:

$$L_{\text{iso}}^k(G, H) : \begin{cases} \sum_{v \in V} X_{\pi \cup \{(v, w)\}} = X_{\pi} & \text{for all } \pi \subseteq V \times W \text{ of size } |\pi| \leq k - 1 \text{ and all } w \in W & \text{(L1)} \\ \sum_{w \in W} X_{\pi \cup \{(v, w)\}} = X_{\pi} & \text{for all } \pi \subseteq V \times W \text{ of size } |\pi| \leq k - 1 \text{ and all } v \in V & \text{(L2)} \\ X_{\pi} = 0 & \text{for all } \pi \subseteq V \times W \text{ of size } |\pi| \leq k \text{ such that } \pi \text{ is not a partial isomorphism from } G \text{ to } H & \text{(L3)} \\ X_{\emptyset} = 1 & & \text{(L4)} \end{cases}$$

This system is closely related to the Sherali-Adams relaxations of $F_{\text{iso}}(G, H)$: Every solution for the level- k Sherali-Adams relaxation of $F_{\text{iso}}(G, H)$ yields a solution to $L_{\text{iso}}^k(G, H)$, and every solution to $L_{\text{iso}}^k(G, H)$ yields a solution to the level $k - 1$ Sherali-Adams relaxation of $F_{\text{iso}}(G, H)$ [4, 12]. Our result is this:

► **Theorem 3.** *For all $k \geq 1$ and for all graphs G and H , the following are equivalent:*

- i $\text{HOM}_{\mathcal{T}_k}(G) = \text{HOM}_{\mathcal{T}_k}(H)$.
- ii k -WL does not distinguish G and H .
- iii $L_{\text{iso}}^{k+1}(G, H)$ has a nonnegative real solution.

The equivalence between ii and iii is implicit in previous work [14, 4, 12]. The system $L_{\text{iso}}^k(G, H)$ has another nice interpretation related to the proof complexity of graph isomorphism: it is shown in [7] that $L_{\text{iso}}^k(G, H)$ has a real solution if and only if a natural system of polynomial equations encoding the isomorphisms between G and H has a degree- k solution in the Hilbert Nullstellensatz proof system [6, 8]. In view of Theorem 2, it is tempting to conjecture that the solvability of $L_{\text{iso}}^{k+1}(G, H)$ characterizes the expressiveness of the homomorphism vectors $\text{HOM}_{\mathcal{P}_k}(G)$ for the class \mathcal{P}_k of all graphs of path width k . Unfortunately, we only prove one direction of this conjecture.

► **Theorem 4.** *Let k be an integer with $k \geq 2$ and let G, H be graphs. If $L_{\text{iso}}^{k+1}(G, H)$ has a real solution, then $\text{HOM}_{\mathcal{P}_k}(G) = \text{HOM}_{\mathcal{P}_k}(H)$.*

Combining this theorem with a recent result from [13] separating the nonnegative from arbitrary real solutions of our systems of equations, we obtain the following corollary.

► **Corollary 5.** *For every k , there are graphs G and H with $\text{HOM}_{\mathcal{P}_k}(G) = \text{HOM}_{\mathcal{P}_k}(H)$ and $\text{HOM}_{\mathcal{T}_2}(G) \neq \text{HOM}_{\mathcal{T}_2}(H)$.*

2 Preliminaries

Basics. Graphs in this paper are simple, undirected, and finite (even though our results transfer to directed graphs and even to weighted graphs). For a graph G , we write $V(G)$ for its vertex set and $E(G)$ for its edge set. For $v \in V(G)$, the set of neighbors of v are denoted with $N_G(v)$. For $S \subseteq V(G)$, we denote with $G[S]$ the subgraph of G induced by the vertices of S . A *rooted graph* is a graph G together with a designated root vertex $r(G) \in V(G)$. We write multisets using the notation $\{\{1, 1, 6, 2\}\}$.

Matrices. An LU -decomposition of a matrix A consists of a lower triangular matrix L and an upper triangular matrix U such that $A = LU$ holds. Every finite matrix A over \mathbf{R} has an LU -decomposition. We also use *infinite matrices* over \mathbf{R} , which are functions $A : I \times J \rightarrow \mathbf{R}$ where I and J are locally finite posets and countable. The matrix product AB is defined in the natural way via $(AB)_{ij} = \sum_k A_{ik}B_{kj}$ if all of these inner products are finite sums, and otherwise we leave it undefined. An $n \times n$ real symmetric matrix has real eigenvalues and a corresponding set of orthogonal eigenspaces. The spectral decomposition of a real symmetric matrix M is of the form $M = \lambda_1 P_1 + \dots + \lambda_l P_l$ where $\lambda_1, \dots, \lambda_l$ are the eigenvalues of M with corresponding eigenspaces W_1, \dots, W_l . Moreover, each P_j is the projection matrix corresponding to the projection onto the eigenspace W_j . Usually, P_j is expressed as $P_j = UU^T$ for a matrix U whose columns form an orthonormal basis of W_j .

Homomorphism numbers. Recall that a mapping $h : V(F) \rightarrow V(G)$ is a homomorphism if $h(e) \in E(G)$ holds for all $e \in E(F)$ and that $\text{Hom}(F, G)$ is the number of homomorphisms from F to G . Let $\text{Surj}(F, G)$ be the number of homomorphisms from F to G that are surjective on both the vertices and edges of G . Let $\text{Inj}(F, G)$ be the number of injective homomorphisms from F to G . Let $\text{Sub}(F, G) = \text{Inj}(F, G)/\text{Aut}(F)$, where $\text{Aut}(F)$ is the number of automorphisms of F . Observe that $\text{Sub}(F, G)$ is the number of subgraphs of G that are isomorphic to F . Where convenient, we view the objects Hom , Surj , and Inj as infinite matrices; the matrix indices are all unlabeled graphs, sorted by their size. However, we only use one representative of each isomorphism class, called the *isomorphism type* of the graphs in the class, as an index in the matrix. Then Surj is lower triangular and Inj is upper triangular, so $\text{Hom} = \text{Surj} \cdot \text{Sub}$ is an LU -decomposition of Hom . Finally, $\text{Ind}(F, G)$ is

the number of times F occurs as an induced subgraph in G . Similarly to the homomorphism vectors $\text{HOM}_{\mathcal{F}}(G)$ we define vectors $\text{INJ}_{\mathcal{F}}(G)$ and $\text{IND}_{\mathcal{F}}(G)$. Finally, let G, H be rooted graphs. A homomorphism from G to H is a graph homomorphism that maps the root of G to the root of H . Moreover, two rooted graphs are isomorphic if there is an isomorphism mapping the root to the root.

3 Homomorphisms from trees

3.1 Color refinement and tree unfolding

Color refinement iteratively colors the vertices of a graph in a sequence of *refinement rounds*. Initially, all vertices get the same color. In each refinement round, any two vertices v and w that still have the same color get different colors if there is some color c such that v and w have a different number of neighbors of color c ; otherwise they keep the same color. We stop the refinement process if the vertex partition that is induced by the colors does not change anymore, that is, all pairs of vertices that have the same color before the refinement round still have the same color after the round. More formally, we define the sequence $C_0^G, C_1^G, C_2^G, \dots$ of colorings as follows. We let $C_0^G(v) = 1$ for all $v \in V(G)$, and for $i \geq 0$ we let $C_{i+1}^G(v) = \{\{ C_i^G(u) : u \in N_G(v) \}\}$. We say that color refinement *distinguishes* two graphs G and H if there is an $i \geq 0$ with

$$\{\{ C_i^G(v) : v \in V(G) \}\} \neq \{\{ C_i^H(v) : v \in V(H) \}\}. \quad (1)$$

We argue now that the color refinement algorithm implicitly constructs a tree at v obtained by simultaneously taking all possible walks starting at v (and not remembering nodes visited in the past). For a rooted tree T with root r , a graph G , and a vertex $v \in V(G)$, we say that T is *a tree at v* if there is a homomorphism f from T to G such that $f(r) = v$ and, for all non-leaves $t \in V(T)$, the function f induces a bijection between the set of children of t in T and the set of neighbors of $f(t)$ in G . In other words, f is a homomorphism from T to G that is *locally bijective*. If T is an infinite tree at v and does not have any leaves, then T is uniquely determined up to isomorphisms, and we call this *the infinite tree at v* (or the *tree unfolding of G at v*), denoted with $T(G, v)$. For an infinite rooted tree T , let $T_{\leq d}$ be the finite rooted subtree of T where all leaves are at depth exactly d . For all finite trees T of depth d , define $\text{Cr}(T, G) \in \{0, \dots, |V(G)|\}$ to be the number of vertices $v \in V(G)$ for which T is isomorphic to $T(G, v)_{\leq d}$. Note that this number is zero if not all leaves of T are at the same depth d or if some node of T has more than $n - 1$ children. The CR-vector of G is the vector $\text{CR}(G) = (\text{Cr}(T, G))_{T \in \mathcal{T}_r}$, where \mathcal{T}_r denotes the family of all rooted trees. The following connection between the color refinement algorithm and the CR-vector is known.

► **Lemma 6** (Angluin [2], also see Krebs and Verbitsky [15, Lemma 2.5]). *For all graphs G and H , color refinement distinguishes G and H if and only if $\text{CR}(G) \neq \text{CR}(H)$ holds.*

3.2 Proof of Theorem 1

Throughout this section, we work with rooted trees. For a rooted tree T and an (unrooted) graph G , we simply let $\text{Hom}(T, G)$ be the number of homomorphisms of the plain tree underlying T to G , ignoring the root.

Let T and T' be rooted trees. A homomorphism h from T to T' is *depth-preserving* if, for all vertices $v \in V(T)$, the depth of v in T is equal to the depth of $h(v)$ in T' . Moreover, a homomorphism h from T to T' is *depth-surjective* if the image of T under h contains vertices

at every depth present in T' . We define $\overrightarrow{\text{Hom}}(T, T')$ as the number of homomorphisms from T to T' that are both depth-preserving and depth-surjective. Note that $\overrightarrow{\text{Hom}}(T, T') = 0$ holds if and only if T and T' have different depths.

► **Lemma 7.** *Let T be a rooted tree and let G be a graph. We have*

$$\text{Hom}(T, G) = \sum_{T'} \overrightarrow{\text{Hom}}(T, T') \cdot \text{Cr}(T', G), \quad (2)$$

where the sum is over all unlabeled rooted trees T' . In other words, the matrix identity $\text{Hom} = \overrightarrow{\text{Hom}} \cdot \text{Cr}$ holds.

Proof. Let d be the depth of T and let r be the root of T . Every T' with $\overrightarrow{\text{Hom}}(T, T') \neq 0$ has depth d too and there are at most n non-isomorphic rooted trees T' of depth d with $\text{Cr}(T', G) \neq 0$. Thus the sum in (2) has only finitely many non-zero terms and is well-defined.

For a rooted tree T' and a vertex $v \in V(G)$, let $H(T', v)$ be the set of all homomorphisms h from T to G such that $h(r) = v$ holds and the tree unfolding $T(G, v)_{\leq d}$ is isomorphic to T' . Let $H(T') = \bigcup_{v \in V(G)} H(T', v)$ and observe $|H(T', v)| = \overrightarrow{\text{Hom}}(T, T')$. Since $\text{Cr}(T', G)$ is the number of $v \in V(G)$ with $T(G, v)_{\leq d} \cong T'$, we thus have $|H(T')| = \overrightarrow{\text{Hom}}(T, T') \cdot \text{Cr}(T', G)$. Since each homomorphism from T to G is contained in exactly one set $H(T')$, we obtain the desired equality (2). ◀

For rooted trees T and T' , let $\overrightarrow{\text{Surj}}(T, T')$ be the number of depth-preserving and surjective homomorphisms from T to T' . In particular, not only do these homomorphisms have to be depth-surjective, but they should hit every vertex of T' . For rooted trees T and T' of the same depth, let $\overrightarrow{\text{Sub}}(T, T')$ be the number of subgraphs of T' that are isomorphic to T (under an isomorphism that maps the root to the root); if T and T' have different depths, we set $\overrightarrow{\text{Sub}}(T, T') = 0$.

► **Lemma 8.** $\overrightarrow{\text{Hom}} = \overrightarrow{\text{Surj}} \cdot \overrightarrow{\text{Sub}}$ is an LU-decomposition of $\overrightarrow{\text{Hom}}$, and $\overrightarrow{\text{Surj}}$ and $\overrightarrow{\text{Sub}}$ are invertible.

As is the case for finite matrices, the inverse of a lower (upper) triangular matrix is lower (upper) triangular. As the matrix $\overrightarrow{\text{Surj}}$ is lower triangular and the matrix $\overrightarrow{\text{Sub}}$ is upper triangular, their inverses are as well. We are ready to prove our first main theorem.

Proof of Theorem 1. We only need to prove the equivalence between assertions i and ii. For every graph G , let $\text{HOM}_r(G) := (\text{Hom}(T, G))_{T \in \mathcal{T}_r}$. By our convention that for a rooted tree T and an unrooted graph G we let $\text{Hom}(T, G)$ be the number of homomorphisms of the plain tree underlying T to G , for all G and H we have $\text{HOM}_r(G) = \text{HOM}_r(H) \iff \text{HOM}(G) = \text{HOM}(H)$. By Lemma 6, it suffices to prove for all graph G, H that

$$\text{CR}(G) = \text{CR}(H) \iff \text{HOM}_r(G) = \text{HOM}_r(H). \quad (3)$$

We view the vectors $\text{HOM}_r(G)$ and $\text{CR}(G)$ as infinite column vectors. By Lemma 7, we have

$$\text{HOM}_r(G) = \overrightarrow{\text{Hom}} \cdot \text{CR}(G) \text{ and } \text{HOM}_r(H) = \overrightarrow{\text{Hom}} \cdot \text{CR}(H). \quad (4)$$

The forward direction of (3) now follows immediately.

It remains to prove the backward direction. Since $\overrightarrow{\text{Hom}} = \overrightarrow{\text{Surj}} \cdot \overrightarrow{\text{Sub}}$ holds by Lemma 8 for two invertible matrices $\overrightarrow{\text{Surj}}$ and $\overrightarrow{\text{Sub}}$, we can first left-multiply with $\overrightarrow{\text{Surj}}^{-1}$ to obtain the equivalent identities

$$\overrightarrow{\text{Surj}}^{-1} \cdot \text{HOM}_r(G) = \overrightarrow{\text{Sub}} \cdot \text{CR}(G) \text{ and } \overrightarrow{\text{Surj}}^{-1} \cdot \text{HOM}_r(H) = \overrightarrow{\text{Sub}} \cdot \text{CR}(H). \quad (5)$$

Now suppose $\text{HOM}_r(G) = \text{HOM}_r(H)$ holds, and set $\mathbf{v} = \text{HOM}_r(G)$. Then $\overrightarrow{\text{Surj}}^{-1} \cdot \mathbf{v}$ is well-defined, because $\overrightarrow{\text{Surj}}$ and its inverse are lower triangular. Thus we obtain $\overrightarrow{\text{Sub}} \cdot \text{CR}(G) = \overrightarrow{\text{Sub}} \cdot \text{CR}(H)$ and set $\mathbf{w} = \text{CR}(G)$. Unfortunately, $\overrightarrow{\text{Sub}}^{-1} \cdot \mathbf{w}$ may be undefined, since $\overrightarrow{\text{Sub}}^{-1}$ is upper triangular. While we can still use a matrix inverse, the argument becomes a bit subtle. The crucial observation is that $\text{Cr}(T', G)$ is non-zero for at most n different trees T' , and all such trees have maximum degree at most $n - 1$. Thus we do not need to look at *all* trees but only those with maximum degree n . Let $\tilde{\mathcal{T}}$ be the set of all unlabeled rooted trees of maximum degree at most n . Let $\text{CR}' = \text{CR}|_{\tilde{\mathcal{T}}}$, let $\mathbf{w}' = \mathbf{w}|_{\tilde{\mathcal{T}}}$, and let $\overrightarrow{\text{Sub}}' = \overrightarrow{\text{Sub}}|_{\tilde{\mathcal{T}} \times \tilde{\mathcal{T}}}$. Then we still have the following for all $T \in \tilde{\mathcal{T}}$ and G :

$$\mathbf{w}'_T = \sum_{T' \in \tilde{\mathcal{T}}} \overrightarrow{\text{Sub}}'(T, T') \cdot \text{Cr}(T', G). \tag{6}$$

The new matrix $\overrightarrow{\text{Sub}}'$ is a principal minor of $\overrightarrow{\text{Sub}}$ and thus remains invertible. Moreover, $\overrightarrow{\text{Sub}}'^{-1} \cdot \mathbf{w}'$ is well-defined, since

$$\sum_{T' \in \tilde{\mathcal{T}}} \overrightarrow{\text{Sub}}'^{-1}(T, T') \cdot \mathbf{w}'_{T'} \tag{7}$$

is a finite sum for each T : The number of (unlabeled) trees $T' \in \tilde{\mathcal{T}}$ that have the same depth d as T is bounded by a function in n and d . Thus $\overrightarrow{\text{Sub}}'^{-1} \cdot \mathbf{w}' = \text{CR}'(G)$. By a similar argument, we obtain $\overrightarrow{\text{Sub}}'^{-1} \cdot \mathbf{w}' = \text{CR}'(H)$. This implies $\text{CR}'(G) = \text{CR}'(H)$ and thus $\text{CR}(G) = \text{CR}(H)$. ◀

4 Homomorphisms from cycles and paths

While the arguments we saw in the proof of Theorem 1 are mainly graph-theoretic, the proof of Theorem 2 uses spectral techniques. To introduce the techniques, we first prove a simple, known result already mentioned in the introduction. We call two square matrices *co-spectral* if they have the same eigenvalues with the same multiplicities, and we call two graphs *co-spectral* if their adjacency matrices are co-spectral.

► **Proposition 9** (e.g. [23, Lemma 1]). *Let \mathcal{C} be the class of all cycles (including the degenerate cycle of length 0, which is just a single vertex). For all graphs G and H , we have $\text{HOM}_{\mathcal{C}}(G) = \text{HOM}_{\mathcal{C}}(H)$ if and only if G and H are co-spectral.*

For the proof, we review a few simple facts from linear algebra. The trace $\text{tr}(A)$ of a square matrix $A \in \mathbb{R}^{n \times n}$ is the sum of the diagonal entries. If the eigenvalues of A are $\lambda_1, \dots, \lambda_n$, then $\text{tr}(A) = \sum_{i=1}^n \lambda_i$. Moreover, for each $\ell \geq 0$ the eigenvalues of the matrix A^ℓ are $\lambda_1^\ell, \dots, \lambda_n^\ell$, and thus $\text{tr}(A^\ell) = \sum_{i=1}^n \lambda_i^\ell$. The following technical lemma encapsulates the fact that the information $\text{tr}(A^\ell)$ for all $\ell \in \mathbb{N}$ suffices to reconstruct the spectrum of A with multiplicities. We use the same lemma to prove Theorem 2, but for Proposition 9 a less general version would suffice.

► **Lemma 10.** *Let $X, Y \subseteq \mathbf{R}$ be two finite sets and let $c \in \mathbf{R}_{\neq 0}^X$ and $d \in \mathbf{R}_{\neq 0}^Y$ be two vectors. If the equation*

$$\sum_{x \in X} c_x x^\ell = \sum_{y \in Y} d_y y^\ell \tag{8}$$

holds for all $\ell \in \mathbf{N}$, then $X = Y$ and $c = d$.

Proof. We prove the claim by induction on $k := |X| + |Y|$. For $k = 0$, the claim is trivially true since both sums in (8) are equal to zero by convention.

Let $\hat{x} = \arg \max\{|x| : x \in X \cup Y\}$ and let $\hat{x} \in X$ without loss of generality. If $\hat{x} = 0$, then $X = \{0\}$ and we claim that $Y = \{0\}$ holds. Clearly (8) for $\ell = 0$ yields $0 \neq c_0 = \sum_{y \in Y} d_y$. In particular, $Y \neq \emptyset$ holds. Since $\hat{x} = 0$ is the maximum of $X \cup Y$ in absolute value, we have $Y = \{0\}$ and thus also $c = d$.

Now suppose that $\hat{x} \neq 0$ holds. We consider the sequences $(a_\ell)_{\ell \in \mathbf{N}}$ and $(b_\ell)_{\ell \in \mathbf{N}}$ with

$$a_\ell = \frac{1}{\hat{x}^\ell} \cdot \sum_{x \in X} c_x x^\ell \quad \text{and} \quad b_\ell = \frac{1}{\hat{x}^\ell} \cdot \sum_{y \in Y} d_y y^\ell. \quad (9)$$

Note that $a_\ell = b_\ell$ holds for all $\ell \in \mathbf{N}$ by assumption. Observe the following simple facts:

- 1) If $-\hat{x} \notin X$, then $\lim_{\ell \rightarrow \infty} a_\ell = c_{\hat{x}}$.
- 2) If $-\hat{x} \in X$, then $\lim_{\ell \rightarrow \infty} a_{2\ell} = c_{\hat{x}} + c_{-\hat{x}}$ and $\lim_{\ell \rightarrow \infty} a_{2\ell+1} = c_{\hat{x}} - c_{-\hat{x}}$.

As well as the following exhaustive case distinction for Y :

- a) If $\hat{x}, -\hat{x} \notin Y$, then $\lim_{\ell \rightarrow \infty} b_\ell = 0$.
- b) If $\hat{x} \in Y$ and $-\hat{x} \notin Y$, then $\lim_{\ell \rightarrow \infty} b_\ell = d_{\hat{x}}$.
- c) If $\hat{x} \notin Y$ and $-\hat{x} \in Y$, then $\lim_{\ell \rightarrow \infty} b_{2\ell} = d_{-\hat{x}}$ and $\lim_{\ell \rightarrow \infty} b_{2\ell+1} = -d_{-\hat{x}}$.
- d) If $\hat{x}, -\hat{x} \in Y$, then $\lim_{\ell \rightarrow \infty} b_{2\ell} = d_{\hat{x}} + d_{-\hat{x}}$ and $\lim_{\ell \rightarrow \infty} b_{2\ell+1} = d_{\hat{x}} - d_{-\hat{x}}$.

If $-\hat{x} \notin X$ holds, we see from 1) that a_ℓ converges to the non-zero value $c_{\hat{x}}$. Since the two sequences are equal, the sequence b_ℓ also converges to a non-zero value. The only case for Y where this happens is b), and we get $\hat{x} \in Y$, $-\hat{x} \notin Y$, and $c_{\hat{x}} = d_{\hat{x}}$. On the other hand, if $-\hat{x} \in X$, we see from 2) that a_ℓ does not converge, but the even and odd subsequences do. The only cases for Y where this happens for b_ℓ too are c) and d). We cannot be in case c), since the two accumulation points of b_ℓ just differ in their sign, while the two accumulation points of a_ℓ do not have the same absolute value. Thus we must be in case d) and obtain $x, \hat{x} \in Y$ as well as

$$c_{\hat{x}} + c_{-\hat{x}} = d_{\hat{x}} + d_{-\hat{x}} \quad \text{and} \quad c_{\hat{x}} - c_{-\hat{x}} = d_{\hat{x}} - d_{-\hat{x}}.$$

This linear system has full rank and implies $c_{\hat{x}} = d_{\hat{x}}$ and $c_{-\hat{x}} = d_{-\hat{x}}$.

Either way, we can remove $\{\hat{x}\}$ or $\{\hat{x}, -\hat{x}\}$ from both X and Y and apply the induction hypothesis on the resulting instance X', Y', c', d' . Then $(X, c) = (Y, d)$ follows as claimed. \blacktriangleleft

Proof of Proposition 9. For all $\ell \geq 0$, the number of homomorphisms from the cycle C_ℓ of length ℓ to a graph G with adjacency matrix A is equal to the number of closed length- ℓ walks in G , which in turn is equal to the trace of A^ℓ . Thus for graphs G, H with adjacency matrices A, B , we have $\text{HOM}_{\mathcal{C}}(G) = \text{HOM}_{\mathcal{C}}(H)$ if and only if $\text{tr}(A^\ell) = \text{tr}(B^\ell)$ holds for all $\ell \geq 0$.

If A and B have the same spectrum $\lambda_1, \dots, \lambda_n$, then $\text{tr}(A^\ell) = \lambda_1^\ell + \dots + \lambda_n^\ell = \text{tr}(B^\ell)$ holds for all $\ell \in \mathbf{N}$. For the reverse direction, suppose $\text{tr}(A^\ell) = \text{tr}(B^\ell)$ for all $\ell \in \mathbf{N}$. Let $X \subseteq \mathbf{R}$ be the set of eigenvalues of A and for each $\lambda \in X$, let $c_\lambda \in \{1, \dots, n\}$ be the multiplicity of the eigenvalue λ . Let $Y \subseteq \mathbf{R}$ and d_λ for $\lambda \in Y$ be the corresponding eigenvalues and multiplicities for B . Then for all $\ell \in \mathbf{N}$, we have

$$\sum_{\lambda \in X} c_\lambda \lambda^\ell = \text{tr}(A^\ell) = \text{tr}(B^\ell) = \sum_{\lambda \in Y} d_\lambda \lambda^\ell.$$

By Lemma 10, this implies $(X, c) = (Y, d)$, that is, the spectra of A and B are identical. \blacktriangleleft

In the following example, we show that the vectors $\text{HOM}_{\mathcal{C}}$ for the class \mathcal{C} of cycles and $\text{HOM}_{\mathcal{T}}$ for the class \mathcal{T} of trees are incomparable in their expressiveness.



■ **Figure 2** Two co-spectral graphs

► **Example 11.** The graphs G and H shown in Figure 2 are co-spectral and thus $\text{HOM}_{\mathcal{C}}(G) = \text{HOM}_{\mathcal{C}}(H)$, but it is easy to see that $\text{HOM}_{\mathcal{P}}(G) \neq \text{HOM}_{\mathcal{P}}(H)$ for the class \mathcal{P} of all paths.

Let G' be a cycle of length 6 and H' the disjoint union of two triangles. Then obviously, $\text{HOM}_{\mathcal{C}}(G') \neq \text{HOM}_{\mathcal{C}}(H')$. However, color refinement does not distinguish G' and H' and thus $\text{HOM}_{\mathcal{T}}(G') = \text{HOM}_{\mathcal{T}}(H')$.

Let us now turn to the proof of Theorem 2.

Proof of Theorem 2. Let A and B be the adjacency matrices of G and H , respectively. Since A is a symmetric and real matrix, its eigenvalues are real and the corresponding eigenspaces are orthogonal and span \mathbf{R}^n . Let $\mathbf{1}$ be the n -dimensional all-1 vector, and let $X = \{\lambda_1, \dots, \lambda_k\}$ be the set of all eigenvalues of A whose corresponding eigenspaces are not orthogonal to $\mathbf{1}$. We call these eigenvalues the *useful* eigenvalues of A and without loss of generality assume $\lambda_1 > \dots > \lambda_k$. The n -dimensional all-1 vector $\mathbf{1}$ can be expressed as a direct sum of eigenvectors of A corresponding to useful eigenvalues. In particular, there is a unique decomposition $\mathbf{1} = \sum_{i=1}^k u_i$ such that each u_i is a non-zero eigenvector in the eigenspace of λ_i . Moreover, the vectors u_1, \dots, u_k are orthogonal. For the matrix B , we analogously define its set of useful eigenvalues $Y = \{\mu_1, \dots, \mu_{k'}\}$ and the direct sum $\mathbf{1} = \sum_{i=1}^{k'} v_i$.

We prove the equivalence of the following three assertions (of which the first and third appear in the statement of Theorem 2).

1. $\text{HOM}_{\mathcal{P}}(G) = \text{HOM}_{\mathcal{P}}(H)$.
2. A and B have the same set of useful eigenvalues $\lambda_1, \dots, \lambda_k$ and $\|u_i\| = \|v_i\|$ holds for all $i \in \{1, \dots, k\}$. Here, $\|\cdot\|$ denotes the Euclidean norm with $\|x\|^2 = \sum_j x_j^2$.
3. The system $F_{\text{iso}}(G, H)$ of linear equations has a real solution.

Note that in 2, we do not require that the useful eigenvalues occur with the same multiplicities in A and B . We show the implications $(1 \Rightarrow 2)$, $(2 \Rightarrow 3)$, and $(3 \Rightarrow 1)$.

$(1 \Rightarrow 2)$: Suppose that $\text{Hom}(P_\ell, G) = \text{Hom}(P_\ell, H)$ holds for all paths P_ℓ . Equivalently, this can be stated in terms of the adjacency matrices A and B : for all $\ell \in \mathbf{N}$, we have $\mathbf{1}^T A^\ell \mathbf{1} = \mathbf{1}^T B^\ell \mathbf{1}$. We claim that A and B have the same useful eigenvalues, and that the projections of $\mathbf{1}$ onto the corresponding eigenspaces have the same lengths.

Note that $A^\ell \mathbf{1} = \sum_{i=1}^k \lambda_i^\ell u_i$ holds. Thus we have

$$\mathbf{1}^T A^\ell \mathbf{1} = \left(\sum_{i=1}^k u_i^T \right) \left(\sum_{i=1}^k \lambda_i^\ell u_i \right) = \sum_{i=1}^k \|u_i\|^2 \cdot \lambda_i^\ell. \quad (10)$$

The term $\mathbf{1}^T B^\ell \mathbf{1}$ can be expanded analogously, which together yields

$$\sum_{i=1}^k \|u_i\|^2 \cdot \lambda_i^\ell = \sum_{i=1}^{k'} \|v_i\|^2 \cdot \mu_i^\ell \quad \text{for all } \ell \in \mathbf{N}. \quad (11)$$

Since all coefficients $c_{\lambda_i} = \|u_i\|^2$ and $d_{\mu_i} = \|v_i\|^2$ are non-zero, we are in the situation of Lemma 10. We obtain $k = k'$ and, for all $i \in \{1, \dots, k\}$, we obtain $\lambda_i = \mu_i$ and $\|u_i\| = \|v_i\|$. This is exactly the claim that we want to show.

(2 \Rightarrow 3): We claim that the $(n \times n)$ -matrix X defined via

$$X = \sum_{i=1}^k \frac{1}{\|u_i\|^2} \cdot u_i v_i^T \tag{12}$$

satisfies the F_{iso} equations $AX = XB$ and $X\mathbf{1} = \mathbf{1} = X^T\mathbf{1}$. Indeed, we have

$$AX = \sum_{i=1}^k \frac{1}{\|u_i\|^2} \cdot A u_i v_i^T = \sum_{i=1}^k \frac{\lambda_i}{\|u_i\|^2} \cdot u_i v_i^T = \sum_{i=1}^k \frac{1}{\|u_i\|^2} \cdot u_i v_i^T B^T = X B^T = X B, \tag{13}$$

This follows, since $A u_i = \lambda_i u_i$, $B v_i = \lambda_i v_i$, and B is symmetric. Moreover, we have

$$X\mathbf{1} = \sum_{i=1}^k \frac{1}{\|u_i\|^2} \cdot A u_i v_i^T \mathbf{1} = \sum_{i=1}^k \frac{1}{\|u_i\|^2} \cdot u_i v_i^T \sum_{j=1}^k v_j = \sum_{i=1}^k \frac{1}{\|u_i\|^2} \cdot u_i \cdot v_i^T v_i = \mathbf{1}. \tag{14}$$

This holds by definition of u_i and v_i and from $v_i^T v_i = \|v_i\|^2 = \|u_i\|^2$. The claim $X^T\mathbf{1} = \mathbf{1}$ follows analogously.

(3 \Rightarrow 1): Suppose there is a matrix X with $X^T\mathbf{1} = X\mathbf{1} = \mathbf{1}$ and $AX = XB$. We obtain $A^\ell X = X B^\ell$ by induction for all $\ell \in \mathbf{N}_{>0}$. For $\ell = 0$, this also holds since $A^0 = I_n$ by convention. As a result, we have $\mathbf{1}^T A^\ell \mathbf{1} = \mathbf{1}^T A^\ell X \mathbf{1} = \mathbf{1}^T X B^\ell \mathbf{1} = \mathbf{1}^T B^\ell \mathbf{1}$ for all $\ell \in \mathbf{N}$. Since these scalars count the length- ℓ walks in G and H , respectively, we obtain $\text{Hom}(P_\ell, G) = \text{Hom}(P_\ell, H)$ for all paths P_ℓ as claimed. \blacktriangleleft

5 Homomorphisms from bounded tree width and path width

We briefly outline the main ideas of the proofs of Theorems 3 and 4; the technical details are deferred to the full version of this paper. In Theorem 3, the equivalence between ii and iii is essentially known, so we focus on the equivalence between i and ii. The proof is similar to the proof of Theorem 1 in Section 3.

Let us fix $k \geq 2$. The idea of the k -WL algorithm is to iteratively color k -tuples of vertices. Initially, each k -tuple (v_1, \dots, v_k) is colored by its *atomic type*, that is, the isomorphism type of the labeled graph $G[\{v_1, \dots, v_k\}]$. Then in the refinement step, to define the new color of a k -tuple \bar{v} we look at the current color of all k -tuples that can be reached from \bar{v} by adding one vertex and then removing one vertex.

Similar to the tree unfolding of a graph G at a vertex v , we define the *Weisfeiler-Leman tree unfolding* at a k -tuple \bar{v} of vertices. These objects have some resemblance to the pebbling comonad, which was defined by Abramsky, Dawar, and Wang [1] in the language of category theory. The WL-tree unfolding describes the color of \bar{v} computed by k -WL; formally it may be viewed as a pair (T, F) consisting of a graph F together with a “rooted” tree decomposition (potentially infinite, but again we cut it off at some finite depth). Similar to the numbers $\text{Cr}(T, G)$ and the vector $\text{CR}(G)$, we now have numbers $\text{WL}((T, F), G)$ and a vector $\text{WL}(G)$ such that $\text{WL}(G) = \text{WL}(H)$ holds if and only if k -WL does not distinguish G and H . Then we define a linear transformation Φ with $\text{HOM}_{\mathcal{T}_k}(G) = \Phi \text{WL}(G)$. The existence of this linear transformation directly yields the implication ii \Rightarrow i of Theorem 3. To prove the converse, we show that the transformation Φ is invertible by giving a suitable LU -decomposition of full rank. This completes our sketch of the proof of Theorem 3.

The proof of Theorem 4 requires a different argument, because now we have to use a solution (X_π) of the system $L_{\text{iso}}^{k+1}(G, H)$ to prove that the path width k homomorphism vectors $\text{HOM}_{\mathcal{P}_k}(G)$ and $\text{HOM}_{\mathcal{P}_k}(H)$ are equal. The key idea is to express entries of a suitable variant of $\text{HOM}_{\mathcal{P}_k}(G)$ as a linear combinations of entries of the corresponding vector for H using the values X_π as coefficients.

6 Conclusions

We have studied the homomorphism vectors $\text{HOM}_{\mathcal{F}}(G)$ for various graph classes \mathcal{F} , focusing on classes \mathcal{F} where it is tractable to compute the entries $\text{Hom}(F, G)$ of the vector. Our main interest was in the “expressiveness” of these vectors, that is, in the question what $\text{HOM}_{\mathcal{F}}(G)$ tells us about the graph G . For the classes \mathcal{C} of cycles, \mathcal{T} of trees, \mathcal{T}_k of graphs of tree width at most k , and \mathcal{P} of paths, we have obtained surprisingly clean answers to this question, relating the homomorphism vectors to various other well studied formalisms that on the surface have nothing to do with homomorphism counts.

Some interesting questions remain open. The most obvious is whether the converse of Theorem 4 holds, that is, whether for two graphs G, H with $\text{HOM}_{\mathcal{P}_k}(G) = \text{HOM}_{\mathcal{P}_k}(H)$, the system $L_{\text{iso}}^{k+1}(G, H)$ has a real solution (and hence the Nullstellensatz propositional proof system has no degree- $(k+1)$ refutation of G and H being isomorphic).

Another related open problem in spectral graph theory is to characterize graphs which are identified by their spectrum, up to isomorphism. In our framework, Proposition 9 ensures that we can equivalently ask for the following characterization: for which graphs G does the vector $\text{HOM}_{\mathcal{C}}(G)$ determine the entire homomorphism vector $\text{HOM}(G)$?

Despite the computational intractability, it is also interesting to study the vectors $\text{HOM}_{\mathcal{F}}(G)$ for classes \mathcal{F} of unbounded tree width. Are there natural classes \mathcal{F} (except of course the class of all graphs) for which the vectors $\text{HOM}_{\mathcal{F}}(G)$ characterize G up to isomorphism? For example, what about classes of bounded degree or the class of planar graphs? And what is the complexity of deciding whether $\text{HOM}_{\mathcal{F}}(G) = \text{HOM}_{\mathcal{F}}(H)$ holds when G and H are given as input? Our results imply that this problem is in polynomial time for the classes \mathcal{T} , \mathcal{T}_k , and \mathcal{P} . For the class of all graphs, it is in quasi-polynomial time by Babai’s quasi-polynomial isomorphism test [5]. Yet it seems plausible that there are classes \mathcal{F} (even natural classes decidable in polynomial time) for which the problem is co-NP-hard.

Maybe the most interesting direction for further research is to study the graph similarity measures induced by homomorphism vectors. A simple way of defining an inner product on the homomorphism vectors is by letting

$$\langle \text{HOM}_{\mathcal{F}}(G), \text{HOM}_{\mathcal{F}}(H) \rangle := \sum_{\substack{k \geq 1 \\ \mathcal{F}_k \neq \emptyset}} \frac{1}{k^k |\mathcal{F}_k|} \sum_{F \in \mathcal{F}_k} \text{Hom}(F, G) \text{Hom}(F, H),$$

where \mathcal{F}_k denotes the class of all graph $F \in \mathcal{F}$ with k vertices. The mapping $(G, H) \mapsto \langle \text{HOM}_{\mathcal{F}}(G), \text{HOM}_{\mathcal{F}}(H) \rangle$ is what is known as a *graph kernel* in machine learning. It induces a (pseudo)metric $d_{\mathcal{F}}$ on the class of graphs. It is an interesting question how it relates to other graph similarity measures, for example, the metric induced by the Weisfeiler-Leman graph kernel. Our Theorem 1 implies that the metric $d_{\mathcal{T}}$ for the class \mathcal{T} of trees and the metric induced by the Weisfeiler-Leman graph kernel have the same graphs of distance zero.

References

- 1 Samson Abramsky, Anuj Dawar, and Pengming Wang. The pebbling comonad in finite model theory. In *Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 1–12. IEEE Computer Society, 2017. doi:10.1109/LICS.2017.8005129.
- 2 Dana Angluin. Local and global properties in networks of processors (extended abstract). In Raymond E. Miller, Seymour Ginsburg, Walter A. Burkhard, and Richard J. Lipton, editors, *Proceedings of the 12th Annual ACM Symposium on Theory of Computing*, pages 82–93. ACM, 1980. doi:10.1145/800141.804655.
- 3 Vikraman Arvind, Johannes Köbler, Sebastian Kuhnert, and Yadu Vasudev. Approximate graph isomorphism. In *Proceedings of the 37th International Symposium on Mathematical Foundations of Computer Science*, volume 7464 of *Lecture Notes in Computer Science*, pages 100–111. Springer, 2012. doi:10.1007/978-3-642-32589-2_12.
- 4 Albert Atserias and Elitza N. Maneva. Sherali–Adams relaxations and indistinguishability in counting logics. *SIAM Journal on Computing*, 42(1):112–137, 2013. doi:10.1137/120867834.
- 5 László Babai. Graph isomorphism in quasipolynomial time [extended abstract]. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 684–697. ACM, 2016. doi:10.1145/2897518.2897542.
- 6 Paul Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, and Pavel Pudlák. Lower bounds on Hilbert’s Nullstellensatz and propositional proofs. *Proceedings of the London Mathematical Society*, s3-73(1):1–26, 1996. doi:10.1112/plms/s3-73.1.1.
- 7 Christoph Berkholz and Martin Grohe. Limitations of algebraic approaches to graph isomorphism testing. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming, Part I*, volume 9134 of *Lecture Notes in Computer Science*, pages 155–166. Springer Verlag, 2015. doi:10.1007/978-3-662-47672-7_13.
- 8 Samuel R. Buss. Lower bounds on Nullstellensatz proofs via designs. In *Proof Complexity and Feasible Arithmetics*, pages 59–71. American Mathematical Society, 1998.
- 9 Donatello Conte, Pasquale Foggia, Carlo Sansone, and Mario Vento. Thirty years of graph matching in pattern recognition. *IJPRAI*, 18(3):265–298, 2004. doi:10.1142/S0218001404003228.
- 10 Víctor Dalmau and Peter Jonsson. The complexity of counting homomorphisms seen from the other side. *Theor. Comput. Sci.*, 329(1-3):315–323, 2004. doi:10.1016/j.tcs.2004.08.008.
- 11 Martin Grohe, Kristian Kersting, Martin Mladenov, and Pascal Schweitzer. Color refinement and its applications. To appear. URL: <https://lii.rwth-aachen.de/images/Mitarbeiter/pub/grohe/cr.pdf>.
- 12 Martin Grohe and Martin Otto. Pebble games and linear equations. *Journal of Symbolic Logic*, 80(3):797–844, 2015. doi:10.1017/jsl.2015.28.
- 13 Martin Grohe and Wied Pakusa. The descriptive complexity of solving linear equation systems and its applications. In *Proceedings of the 32nd ACM-IEEE Symposium on Logic in Computer Science*, 2017. doi:10.1109/LICS.2017.8005081.
- 14 Neil Immerman and Eric Lander. Describing graphs: A first-order approach to graph canonization. In *Complexity Theory Retrospective: In Honor of Juris Hartmanis on the Occasion of His Sixtieth Birthday, July 5, 1988*, pages 59–81. Springer New York, 1990. doi:10.1007/978-1-4612-4478-3_5.
- 15 Andreas Krebs and Oleg Verbitsky. Universal covers, color refinement, and two-variable counting logic: Lower bounds for the depth. In *Proceedings of the 30th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 689–700. IEEE Computer Society, 2015. doi:10.1109/LICS.2015.69.
- 16 László Lovász. Operations with structures. *Acta Mathematica Hungarica*, 18:321–328, 1967.
- 17 László Lovász. *Large Networks and Graph Limits*. American Mathematical Society, 2012.

- 18 Peter N. Malkin. Sherali–Adams relaxations of graph isomorphism polytopes. *Discrete Optimization*, 12:73–97, 2014. doi:10.1016/j.disopt.2014.01.004.
- 19 Viswanath Nagarajan and Maxim Sviridenko. On the maximum quadratic assignment problem. *Math. Oper. Res.*, 34(4):859–868, 2009. doi:10.1287/moor.1090.0418.
- 20 Nino Shervashidze, Pascal Schweitzer, Erik Jan van Leeuwen, Kurt Mehlhorn, and Karsten M. Borgwardt. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research*, 12:2539–2561, 2011. URL: <https://dl.acm.org/citation.cfm?id=2078187>.
- 21 Gottfried Tinhofer. Graph isomorphism and theorems of Birkhoff type. *Computing*, 36(4):285–300, 1986. doi:10.1007/BF02240204.
- 22 Gottfried Tinhofer. A note on compact graphs. *Discrete Applied Mathematics*, 30(2-3):253–264, 1991. doi:10.1016/0166-218X(91)90049-3.
- 23 Edwin R. Van Dam and Willem H. Haemers. Which graphs are determined by their spectrum? *Linear Algebra and its applications*, 373:241–272, 2003. doi:10.1016/S0024-3795(03)00483-X.
- 24 S. V. N. Vishwanathan, Nicol N. Schraudolph, Risi Kondor, and Karsten M. Borgwardt. Graph kernels. *Journal of Machine Learning Research*, 11:1201–1242, 2010. URL: <https://portal.acm.org/citation.cfm?id=1859891>.