Journal of Symbolic Computation 59 (2013) 81–112

Contents lists available at SciVerse ScienceDirect

Journal of Symbolic Computation

www.elsevier.com/locate/jsc

The Graph Isomorphism Problem and approximate categories

Harm Derksen[1]

University of Michigan, Department of Mathematics, 530 Church Street, Ann Arbor, MI 48109-1043, United States

ARTICLE INFO

Article history:
Received 6 May 2013
Accepted 1 June 2013
Available online 1 July 2013

Keywords:
Graph Isomorphism Problem
Weisfeiler–Lehman algorithm
Complexity Theory
Categories

ABSTRACT

It is unknown whether two graphs can be tested for isomorphism in polynomial time. A classical approach to the Graph Isomorphism Problem is the d-dimensional Weisfeiler–Lehman algorithm. The d-dimensional WL-algorithm can distinguish many pairs of graphs, but the pairs of non-isomorphic graphs constructed by Cai, Fürer and Immerman it cannot distinguish. If d is fixed, then the WL-algorithm runs in polynomial time. We will formulate the Graph Isomorphism Problem as an Orbit Problem: Given a representation V of an algebraic group G and two elements $v_1, v_2 \in V$, decide whether v_1 and v_2 lie in the same G-orbit. Then we attack the Orbit Problem by constructing certain approximate categories $\mathcal{C}_d$, $d \in \mathbb{N} = \{1, 2, 3, \ldots\}$ whose objects include the elements of V. We show that v_1 and v_2 are not in the same orbit by showing that they are not isomorphic in the category $\mathcal{C}_d$ for some $d \in \mathbb{N}$. For every d this gives us an algorithm for isomorphism testing. We will show that the WL-algorithms reduce to our algorithms, but that our algorithms cannot be reduced to the WL-algorithms. Unlike the Weisfeiler–Lehman algorithm, our algorithm can distinguish the Cai–Fürer–Immerman graphs in polynomial time.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction and main results

1.1. The Graph Isomorphism Problem

Suppose that Γ_1 and Γ_2 are two graphs on n vertices. The *Graph Isomorphism Problem* asks whether they are isomorphic or not. In Computational Complexity Theory, the Graph Isomorphism Problem

E-mail address: hderksen@umich.edu.

[1] The author is supported by the NSF, grant DMS 0901298.

0747-7171/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jsc.2013.06.002

plays an important role, because it lies in the complexity class **NP**, but it is not known whether it lies in **P** or **NP-complete**. See Köbler et al. (1993) for more details. Based on Valiant's algebraic version of the **P** versus **NP** problem (Valiant, 1979), Mulmuley and Sohoni reformulated Valiant's **P** versus **NP** problem into a question about orbits of algebraic groups in Mulmuley and Sohoni (2001, 2008). In this paper, we will study the Graph Isomorphism Problem in terms of orbits of algebraic groups, but our approach is not closely related to the work of Mulmuley and Sohoni.

For special families of graphs there are polynomial time algorithms for the Graph Isomorphism Problem. Polynomial time algorithms were found for trees (Edmonds' algorithm, see Busacker and Saaty, 1965, p. 196), planar graphs (Hopcraft and Tarjan, 1973; Hopcroft and Wong, 1974) and more generally for graphs of bounded genus (Filotti and Mayer, 1980; Miller, 1980), for graphs with bounded degree (Luks, 1982), for graphs with bounded eigenvalue multiplicity (Babai et al., 1982), and for graphs with bounded color class size (Luks, 1986).

A general approach to the Graph Isomorphism Problem was developed by Weisfeiler and Lehman in the 1960s. The d-dimensional Weisfeiler–Lehman algorithm $\mathbf{WL}_d$ systematically colors e-tuples of vertices ($e \leqslant d$) until a stable coloring is obtained (see Weisfeiler and Lehman, 1968; Weisfeiler, 1976). The d-dimensional WL-algorithm terminates with a proof that the two graphs are not isomorphic, or it terminates with an inconclusive result. If $d \geqslant n$, then the d-dimensional Weisfeiler–Lehman algorithm will distinguish all non-isomorphic graphs with n vertices. For fixed d, the Weisfeiler–Lehman algorithm runs in polynomial time. The higher-dimensional Weisfeiler–Lehman algorithm can distinguish graphs in many families of graphs. However, Cai, Fürer and Immerman showed in Cai et al. (1992) that for every d, there exists a pair of non-isomorphic graphs with degree 3 and $O(d)$ vertices which cannot be distinguished by the d-dimensional Weisfeiler–Lehman algorithm. The set of Weisfeiler–Lehman algorithms $\mathbf{WL} = \{\mathbf{WL}_d\}_{d\in\mathbb{N}}$ is an example of what we will call a *family of GI-algorithms*:

Definition 1.1. A family of GI-algorithms is a collection of algorithms $\mathbf{A} = \{\mathbf{A}_d\}_{d\in\mathbb{N}}$ such that

(1) The input of $\mathbf{A}_d$ consists of two graphs with the same number of vertices. The value of the output is either "`non-isomorphic`" or "`inconclusive`". If the output is "`non-isomorphic`" then the graphs are not isomorphic and we say that $\mathbf{A}_d$ distinguishes the two graphs.
(2) If the graphs are not isomorphic, then $\mathbf{A}_d$ distinguishes them for some d.
(3) For fixed d, $\mathbf{A}_d$ runs in polynomial time.

Besides the Weisfeiler–Lehman algorithm, there are other families of polynomial time algorithms for the Graph Isomorphism Problem. In order to compare various algorithms, we make the following definition (see also Evdokimov et al., 1999, §6):

Definition 1.2. For two families of GI-algorithms $\mathbf{A} = \{\mathbf{A}_d\}_{d\in\mathbb{N}}$ and $\mathbf{B} = \{\mathbf{B}_d\}_{d\in\mathbb{N}}$ we say that **A** is reducible to **B** if there exists a function $f : \mathbb{N} \to \mathbb{N}$ such that for every d and every pair of graphs which $\mathbf{A}_d$ distinguishes, the pair can be distinguished by $\mathbf{B}_{f(d)}$. We say that **A** and **B** are equivalent if **A** is reducible to **B** and **B** is reducible to **A**.

The Weisfeiler–Lehman algorithm is combinatorial in nature. There are also more algebraic approaches to the Graph Isomorphism Problem. The 2-dimensional Weisfeiler–Lehman algorithm can be formulated in terms of *cellular algebras* (see Weisfeiler, 1976).[2] These algebras were introduced by Weisfeiler and Lehman, and independently by D. Higman under the name *coherent algebras* (see Higman, 1987; Friedland, 1989). In Evdokimov et al. (1999), Evdokimov, Karpinski and Ponomarenko introduced the d-closure of a cellular algebra. One may view d-closed cellular algebras as higher-dimensional analogs of the cellular algebras. The algorithm based on this d-closure will be denoted by $\mathbf{CA}_d$. In Evdokimov et al. (1999) it was shown that the algorithm $\mathbf{CA}_d$ distinguishes any two graphs which can be distinguished by $\mathbf{WL}_d$. In Evdokimov and Ponomarenko (1999, Theorem 1.4)

[2] These cellular algebras should not be confused with a different, seemingly unrelated notion of cellular algebras introduced in Graham and Lehrer (1996).

it was shown that $\mathbf{WL}_{3d}$ can distinguish any two graphs which can be distinguished by $\mathbf{CA}_d$. So the approach with cellular algebras $\mathbf{CA} = \{\mathbf{CA}_d\}_{d\in\mathbb{N}}$ is equivalent to the Weisfeiler–Lehman algorithm $\mathbf{WL} = \{\mathbf{WL}_d\}_{d\in\mathbb{N}}$.

In this paper we will define a family of GI-algorithms $\mathbf{AC}_d = \{\mathbf{AC}_d\}_{d\geqslant 0}$ using approximate categories. We will show that **WL** reduces to **AC** and that **AC** does *not* reduce to **WL**.

1.2. Finite variable logic

Pairs of non-isomorphic graphs that can be distinguished with the Weisfeiler–Lehman algorithm can be characterized in terms of finite variable logic.

A graph is a pair $\Gamma = \langle Y, E\rangle$ where Y is a finite set, and $E \subseteq Y \times Y$ is a symmetric relation. We assume that there are no loops or multiple edges. We will write xEy if $(x, y) \in E$. We also consider graphs with colored vertices. A graph with m colors is a tuple $\langle Y, E, Y_1, Y_2, \dots, Y_m\rangle$ where $\langle Y, E\rangle$ is a graph, and Y is the disjoint union of $Y_1, \dots, Y_m$. Two colored graphs $\Gamma = \langle Y, E, Y_1, \dots, Y_m\rangle$ and $\Gamma' = \langle Y', E', Y'_1, \dots, Y'_m\rangle$ are isomorphic if there is a bijection $\phi : Y \to Y'$ such that $x \in Y_i \Leftrightarrow \phi(x) \in Y'_i$ for all $x \in Y$ and all i, and $xEy \Leftrightarrow \phi(x)E'\phi(y)$ for all $x, y \in Y$.

We will view a graph with m colors as a structure with 1 binary relation and m unitary relations. To such a structure one can associate a first order language $\mathscr{L}$. If φ is a closed formula in $\mathscr{L}$, then we will write $\Gamma \models \varphi$ if the formula φ is true for Γ.

Let $\mathscr{L}_d$ be the d-variable first order language. Formulas in $\mathscr{L}_d$ involve at most d variables, although variables may be re-used. For example

$$\varphi(x_1, x_2) = \exists x_3 \bigl(\exists x_2\,(x_1Ex_2 \wedge x_2Ex_3) \wedge x_3Ex_2\bigr)$$

is a formula in $\mathscr{L}_3$ which expresses that there exists a path of length 3 from x_1 to x_2. Note that in this formula, we re-use the variable x_2.

A more expressive language is $\mathscr{C}_d$, the d-variable first order language with counting. In this language we allow quantifiers such as $\exists_m$. A formula $\exists_m x\,\varphi(x)$ is true if there are exactly m elements $x \in Y$ for which $\varphi(x)$ is true.

Definition 1.3. We say that the language $\mathscr{C}_d$ distinguishes the colored graphs Γ and Γ' if there exists a closed formula φ in $\mathscr{C}_d$ such that $\Gamma \models \varphi$ and $\Gamma' \models \neg\varphi$.

Theorem 1.4. *(See §5 of Cai et al., 1992.) The language $\mathscr{C}_d$ distinguishes the colored graphs Γ and Γ' if and only if the $(d-1)$-dimensional Weisfeiler–Lehman algorithm distinguishes the two graphs.*

1.3. Approximate categories

Suppose that Y is a set of n vertices, and let $U = kY \cong k^n$ be the vector space with basis Y. The symmetric group $\Sigma_n = \Sigma(Y)$ acts on the vector space U by permuting the basis elements. We can identify $U \otimes U$ with the space $\mathrm{Mat}_{n,n}(k)$ of $n \times n$ matrices. Define

$$V = U \otimes U \oplus \underbrace{U \oplus \cdots \oplus U}_{m} \oplus k.$$

To an m-colored graph $\Gamma = \langle Y, E, Y_1, \dots, Y_m\rangle$ we can associate an element

$$A_\Gamma = \left(\sum_{(y_1,y_2)\in E} y_1 \otimes y_2, \sum_{y\in Y_1} y, \dots, \sum_{y\in Y_m} y, 1\right) \in V.$$

If $\Gamma' = \langle Y, E', Y'_1, \dots, Y'_m\rangle$ is another m-colored graph, then Γ and Γ' are isomorphic if and only if A_Γ and $A_{\Gamma'}$ lie in the same Σ_n-orbit.

In this paper we construct a category $\mathcal{C}_d = \mathcal{C}_d(V)$ for every integer $d \geqslant 2$. Objects in this categories are elements of $\mathrm{Aff}(V)$, the set of all affine subspaces of V. So an object is either the empty set or of the form $v + Z$ where $v \in V$ and $Z \subseteq V$ is a subspace. We will identify $v \in V$ with the affine subspace $\{v\}$. So we can view elements of V as objects in the category. If A and B are isomorphic in $\mathcal{C}_d$, then we write $A \cong_d B$.

Theorem 1.5. *Suppose that k is a field of characteristic* 0 *or characteristic p with $p > n$, and that Γ_1, Γ_2 are m-colored graphs on n vertices. We have the following implications*:

$$\begin{array}{rc}
\text{(i)} & \textit{the language } \mathscr{C}_d \textit{ distinguishes } \Gamma_1 \textit{ and } \Gamma_2 \\
 & \Updownarrow \\
\text{(ii)} & \textit{the } (d-1)\textit{-dimensional Weisfeiler–Lehman algorithm distinguishes } \Gamma_1 \textit{ and } \Gamma_2 \\
 & \Downarrow \\
\text{(iii)} & A_{\Gamma_1} \ncong_{2d} A_{\Gamma_2} \\
 & \Downarrow \\
\text{(iv)} & A_{\Gamma_1} \textit{ and } A_{\Gamma_2} \textit{ do not lie in the same } \Sigma_n\textit{-orbit} \\
 & \Updownarrow \\
\text{(v)} & \Gamma_1 \ncong \Gamma_2.
\end{array}$$

Theorem 1.6. *Suppose that $k = \mathbb{F}_p$ where $p = p(n)$ is a prime for all n and $\log p(n)$ grows at most polynomially. Then one can check whether two objects in $\mathcal{C}_d$ are isomorphic in polynomial time.*

The proofs of Theorems 1.5 and 1.6 are in Section 4.2.

Theorem 1.7. *If $k = \mathbb{F}_2$ and Γ_1, Γ_2 are non-isomorphic* 2*-colored graphs constructed following the Cai–Fürer–Immerman method, then A_{Γ_1} and A_{Γ_2} are not isomorphic in $\mathcal{C}_3$.*

The proof of Theorem 1.7 is in Section 5.

Algorithm 1.8 *(The family of GI-algorithms $\{\mathbf{AC}_d\}_{d \in \mathbb{N}}$).* To check whether two graphs Γ_1 and Γ_2 on n vertices are isomorphic, we can determine whether A_{Γ_1} and A_{Γ_2} are isomorphic in $\mathcal{C}_d$, where we work over the field $\mathbb{F}_p$ and p runs over all primes $\leqslant 2n$.

Note that there exists a prime between n and $2n$ by Bertrand's postulate (see Ramanujan, 1919). By Theorem 1.6, **WL** reduces to **AC**, because $\mathbf{AC}_{2d+2}$ distinguishes any two graphs that are distinguished by $\mathbf{WL}_d$. On the other hand, for every d, Cai, Fürer and Immerman constructed pairs of graphs that cannot be distinguished by $\mathbf{WL}_d$. By Theorem 1.7, these graphs are distinguished by $\mathbf{AC}_3$. This shows that **AC** cannot be reduced to **WL**. If there exists a fixed d such that $\mathbf{AC}_d$ distinguishes all pairs of non-isomorphic graphs, then the Graph Isomorphism Problem can be solved in polynomial time. It is not known to the author whether such an integer d exists. If it does not exist, then for every positive integer d there exists a pair of non-isomorphic graphs that cannot be distinguished by $\mathbf{AC}_d$. But constructing such families of non-isomorphic graphs seems to be a very hard problem.

1.4. Orbit problems

The approach in this paper to the Graph Isomorphism Problem naturally generalizes to isomorphism problems in which the symmetry groups are algebraic groups. Fix a field k, and let $\bar{k}$ be its algebraic closure. Suppose that G is an algebraic group defined over k, and V is a representation of G (over k). Let $G(\bar{k})$ be the set of $\bar{k}$-rational points of G.

Orbit Problem. Given $v_1, v_2 \in V$, determine whether v_1 and v_2 lie in the same $G(\bar{k})$-orbit.

In Section 1.3 we formulated the Graph Isomorphism Problem as an orbit problem, where $G = \Sigma_n$. In Section 4 we will also translate the Module Isomorphism Problem to an orbit problem. In this paper we attack the orbit problem as follows. Suppose that V is a representation of G. Let $\mathrm{Aff}(V)$ denote the set of affine subspaces of V. The group G acts on $\mathrm{Aff}(V)$. We also may view $\mathrm{Aff}(V)$ as a subset of $\mathrm{Aff}(V \otimes_k \bar{k})$ by identifying $v + Z \in \mathrm{Aff}(V)$ with $v \otimes 1 + Z \otimes \bar{k} \in \mathrm{Aff}(V \otimes_k \bar{k})$ for every $v \in V$ and every subspace $Z \subseteq V$. For every d we construct a category $\mathcal{C}_d(V)$. For $X_1, X_2 \in \mathrm{Aff}(V)$, we will write $X_1 \cong_d X_2$ if X_1 and X_2 are isomorphic in $\mathcal{C}_d(V)$. The categories $\mathcal{C}_d(V)$ have the following properties:

(1) The set of objects of $\mathcal{C}_d(V)$ is $\mathrm{Aff}(V)$. In particular, elements of V are objects in $\mathcal{C}_d(V)$.
(2) $\mathcal{C}_d(V)$ is a k-category, i.e., for every two objects X_1, X_2 the set $\mathrm{Hom}_d(X_1, X_2)$ is a k-vector space, and if X_3 is another object, then the composition map

$$\mathrm{Hom}_d(X_1, X_2) \times \mathrm{Hom}_d(X_2, X_3) \to \mathrm{Hom}_d(X_1, X_3)$$

is bilinear.
(3) For $X_1, X_2 \in \mathrm{Aff}(V)$ we have $X_1 \cong_{d+1} X_2 \Rightarrow X_1 \cong_d X_2$.
(4) Two affine subspaces $X_1, X_2 \in \mathrm{Aff}(V)$ lie in the same $G(\bar{k})$-orbit if and only if $X_1 \cong_d X_2$ for all d.

An *equivariant* $f : V \to V'$ is a polynomial map between two representations which is G-equivariant. An equivariant $f : V \to V'$ for which V' is an irreducible representation is called a *covariant*. If k is algebraically closed, then f being equivariant means that $f(g \cdot v) = g \cdot f(v)$ for all $v \in V$ and all $g \in G$. In the case where k is not algebraically closed, equivariance is defined in Definition 2.16. We say that an equivariant $f : V \to V'$ distinguishes two elements $v_1, v_2 \in V$ if either $f(v_1) = 0$ and $f(v_2) \neq 0$, or, $f(v_1) \neq 0$ and $f(v_2) = 0$. It is well known that if $v_1, v_2 \in V$ are not in the same $G(\bar{k})$-orbit, then they can be distinguished by some equivariant.

For a representation V, $\mathscr{E}(V)$ denotes the class of all equivariants $f : V \to V'$, where V' is any representation. For every positive integer d, we will define in Section 3.2 a subset $\mathscr{E}_d(V) \subseteq \mathscr{E}(V)$. Elements of $\mathscr{E}_d(V)$ are called *d-constructible equivariants*. We have

$$\mathscr{E}_1(V) \subseteq \mathscr{E}_2(V) \subseteq \mathscr{E}_3(V) \subseteq \cdots$$

and $\bigcup_{d=1}^{\infty} \mathscr{E}_d(V) = \mathscr{E}(V)$.

For a representation V and a positive integer d we will define in Section 3.4 a class $\mathcal{F}_d(V)$ of functors $\mathcal{F} : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ where V' is some representation. Elements of $\mathcal{F}_d(V)$ are called *d-constructible functors*. We have

$$\mathscr{F}_1(V) \subseteq \mathscr{F}_2(V) \subseteq \mathscr{F}_3(V) \subseteq \cdots.$$

The constructible functors are more general than the constructible equivariants in the following sense: If $f : V \to V'$ is a d-constructible equivariant, then there exists a constructible functor $\mathcal{F} : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ such that $\mathcal{F}(\{v\}) = \{f(v)\}$ for all $v \in V$ (see Lemma 3.20). We will say that a functor $\mathcal{F} : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ distinguishes $X_1, X_2 \in \mathrm{Aff}(V)$ if $\dim \mathcal{F}(X_1) \neq \dim \mathcal{F}(X_2)$. Here, we use the convention $\dim(\emptyset) = -\infty$.

Theorem 1.9. *Suppose that $X_1, X_2 \in \mathrm{Aff}(V)$. We have the following implications*:

(i) *the d-constructible functors distinguish X_1 and X_2*

$$\Downarrow$$

(ii) $$X_1 \not\cong_d X_2$$

$$\Downarrow$$

(iii) *X_1 and X_2 lie in distinct $G(\bar{k})$-orbits.*

Theorem 1.10. *Suppose that $v_1, v_2 \in V$. We have the following implications*:

$$\begin{array}{ll} \text{(i)} & \textit{the } d\textit{-constructible equivariants distinguish } v_1 \textit{ and } v_2 \\ & \Downarrow \\ \text{(ii)} & \textit{the } d\textit{-constructible functors distinguish } v_1 \textit{ and } v_2 \\ & \Downarrow \\ \text{(iii)} & v_1 \not\cong_d v_2 \\ & \Downarrow \\ \text{(iv)} & v_1 \textit{ and } v_2 \textit{ lie in distinct } G(\bar{k})\textit{-orbits.} \end{array}$$

Theorems 1.9 and 1.10 are proven in Section 4.2.

The following proposition follows from Proposition 4.8 and shows that the implication (i) ⇒ (ii) in Theorem 1.10 cannot be reversed.

Proposition 1.11. *For every d there exist examples, where v_1 and v_2 can be distinguished by* 3*-constructible functors, but not by d-constructible equivariants.*

2. The construction of the approximate categories

2.1. Truncated ideals

To define the approximate categories, we will need the notion of a truncated ideal. Suppose that k is a field, and R is a finitely generated commutative k-algebra with a filtration

$$R_0 \subseteq R_1 \subseteq R_2 \subseteq \cdots$$

such that $R_0 = k$, R_i is a finite-dimensional vector space for all i and $R_iR_j \subseteq R_{i+j}$ for all i, j.

Definition 2.1. If $S \subseteq R_d$ then we define

$$(S)_d = \sum_{e=0}^{d} (S \cap R_e)R_{d-e}.$$

Definition 2.2. A subset $S \subseteq R_d$ is called a *d-truncated ideal* if $(S)_d = S$.

We have a chain

$$(S)_d \subseteq ((S)_d)_d \subseteq (((S)_d)_d)_d \subseteq \cdots.$$

Since R_d is finite dimensional, this chain stabilizes to a subspace of R_d which we will denote by $((S))_d$. It is clear that $((S))_d$ is the smallest d-truncated ideal containing S. We will call it the *d-truncated ideal generated by S*.

Example 2.3. Consider the polynomial ring $R := k[x, y]$ in two variables with the usual grading, where R_d is the space of polynomials of degree at most d. We have

$$y - x^2 = -x\bigl(x - y^2\bigr) - y(xy - 1) \in \bigl(x - y^2, xy - 1\bigr),$$

but $y - x^2 \notin ((x - y^2, xy - 1))_2$.

Remark 2.4. Much of the Gröbner basis theory generalizes to truncated ideals. Suppose that $R = k[x_1, \dots, x_n]$ is the polynomial ring. In a polynomial ring we can choose a monomial ordering which is compatible with the grading: if one monomial has higher degree than another monomial, then it is larger in the monomial ordering. Then, a subset $\mathscr{G}$ of a d-truncated ideal J is a truncated Gröbner basis if the ideal generated by the leading monomials of elements of J is the same as the ideal generated by the leading monomials of elements of $\mathscr{G}$. There is also an analog of Buchberger's algorithm. Starting with a set of generators of J, one obtains a truncated Gröbner basis by reducing S-polynomials whose total degree is $\leqslant d$. Since R_d is a finite-dimensional vector space, computations with truncated ideals can be done by just using linear algebra. However, using truncated Gröbner bases exploits the ring structure and may speed up the computations. For complexity bounds, the linear algebra approach will be good enough, so we will not explore the truncated Gröbner bases in detail here.

Proposition 2.5. *Suppose that $R_e = R_1^e$ for all $e \geqslant 1$, i.e., R_e is spanned by all products $f_1 f_2 \cdots f_e$ with $f_1, \dots, f_e \in R_1$. Then, there exists a constant $C(d)$ (depending on d, R and the filtration) such that $(\!(S)\!)_e = (S) \cap R_e$ for all $e \geqslant C(d)$ and all $S \subseteq R_d$.*

Proof. Define a ring homomorphism

$$\gamma : k[x_1, \dots, x_n] \to R$$

such that $\gamma(x_1), \dots, \gamma(x_n)$ span R_1. Suppose that $h_1, \dots, h_r \in k[x_1, \dots, x_n]$ generate the kernel of γ, and let l be the maximum of the degrees of $h_1, \dots, h_r$. Assume that $S \subseteq R_d$ is a subspace spanned by $f_1, \dots, f_s$. For all i, choose $\widetilde{f_i} \in k[x_1, \dots, x_n]$ of degree $\leqslant d$ with $\gamma(\widetilde{f_i}) = f_i$. Let J be the ideal generated by the set $\mathcal{G} = \{\widetilde{f}_1, \dots, \widetilde{f}_s, h_1, \dots, h_r\}$. Then we have $\gamma(J) = (S)$. We can apply Buchberger's algorithm to the generator set $\mathcal{G}$ to obtain a Gröbner basis for the ideal J. It was shown in Weispfenning (1988) that there exists a universal bound $C(d)$ (depending only on d, n, and l) such that all polynomials in the reduced Gröbner basis, and all polynomials appearing in intermediate steps of Buchberger's algorithm have degree $\leqslant C(d)$. Following the Buchberger algorithm, it is easy to see that $\gamma(u) \in (\!(S)\!)_{C(d)}$ for all elements u in the Gröbner basis of J. Suppose that $f \in (S) \cap R_e$ and $e \geqslant C(d)$. We can lift f to an element $\widetilde{f}$ such that $\deg(\widetilde{f}) \leqslant e$ and $\gamma(\widetilde{f}) = f$. We can write $\widetilde{f} = \sum_{i=1}^{t} a_i u_i$ where $u_1, \dots, u_t$ are elements of the Gröbner basis, and $\deg(a_i u_i) \leqslant e$ for all i. From this it follows that

$$f = \gamma(\widetilde{f}) = \sum_{i=1}^{t} \gamma(a_i)\gamma(u_i) \in \sum R_{e-j}\big((\!(S)\!)_e \cap R_j\big) = \big((\!(S)\!)_e\big)_e = (\!(S)\!)_e. \qquad \square$$

2.2. The category $\mathcal{C}_d$

For sets S and T we denote the set of functions $S \to T$ by T^S. Consider the group algebra $k\Sigma_n$, and let $R = k^{\Sigma_n}$ be the commutative ring of all k-valued functions on Σ_n. The spaces $k\Sigma_n$ and R are dual to each other, and carry the usual Hopf algebra structures. We define a ring homomorphism

$$\Delta : R \to k^{\Sigma_n \times \Sigma_n}$$

as follows. If $f \in R = k^{\Sigma_n}$, then $\Delta(f)$ is the function on $\Sigma_n \times \Sigma_n$ given by

$$\Delta(f)(g_1, g_2) = f(g_1 g_2).$$

There exists a linear isomorphism between $R \otimes R$ and $k^{\Sigma_n \times \Sigma_n}$ where $f_1 \otimes f_2$ is mapped to the function

$$(g_1, g_2) \mapsto f_1(g_1) f_2(g_2).$$

We will identify $R \otimes R$ with $k^{\Sigma_n \times \Sigma_n}$. The linear map

$$\Delta : R \to R \otimes R$$

is the *co-multiplication* on R. The dual map

$$\Delta^\star : k\Sigma_n \otimes k\Sigma_n \to k\Sigma_n$$

is just the multiplication

$$k\Sigma_n \times k\Sigma_n \to k\Sigma_n$$

in the group algebra $k\Sigma_n$.

We will now define a filtration on R. We can identify Σ_n with the set of $n \times n$ permutation matrices. Every k-valued function on Σ_n can be extended to a polynomial function on $\mathrm{Mat}_{n,n}(k)$. Let $x_{i,j}$ with $1 \leqslant i, j \leqslant n$ be the coordinate functions on $\mathrm{Mat}_{n,n}(k)$. Define $R_d \subseteq R = k^{\Sigma_n}$ as the subspace of all k-valued functions that are given by a polynomial in $x_{1,1}, x_{1,2}, \ldots, x_{n,n}$ of degree $\leqslant d$.

Lemma 2.6. *We have $R_n = R$.*

Proof. If $f \in R$, then there exists a polynomial

$$F(x_{1,1}, x_{1,2}, \ldots, x_{n,n}) \in k[x_{1,1}, x_{1,2}, \ldots, x_{n,n}]$$

whose restriction to Σ_n is f. Using the relations $x_{i,j}^2 = x_{i,j}$ on Σ_n, we may assume that all monomials in F are square-free. Using the relations $x_{i,j}x_{i,k} = 0 = x_{j,i}x_{k,i}$ for $j \neq k$ we can assume that these monomials are of the form $x_{i_1,j_1}x_{i_2,j_2}\cdots x_{i_r,j_r}$ where $i_1, \ldots, i_r$ are distinct, and $j_1, \ldots, j_r$ are distinct. With these assumptions, F has degree $\leqslant n$. This shows that $R_n = R$. □

We have a filtration

$$k = R_0 \subseteq R_1 \subseteq R_2 \subseteq \cdots \subseteq R_n = R.$$

Suppose that $f \in R_d$. So f is the restriction of a polynomial $F(x_{1,1}, x_{1,2}, \ldots, x_{n,n})$ of degree $\leqslant d$. Because matrix multiplication is bilinear, the function

$$\Delta(f) : (g_1, g_2) \mapsto f(g_1 g_2)$$

is the restriction of the polynomial

$$F\left(\sum_j x_{1,j}y_{j,1}, \sum_j x_{1,j}y_{j,2}, \ldots, \sum_j x_{1,n}y_{j,n}\right)$$

to $\Sigma_n \times \Sigma_n$. This polynomial has degree $\leqslant d$ in the x-variables and degree $\leqslant d$ in the y-variables. This shows that

$$\Delta(R_d) \subseteq R_d \otimes R_d.$$

Dualizing the inclusion $R_d \subseteq R$ yields a surjective map

$$\Xi_d : R^\star = k\Sigma_n \to R_d^\star$$

which is defined by

$$\Xi_d\left(\sum_{g\in\Sigma_n} c_g g\right)(f) = \sum_{g\in\Sigma_n} c_g f(g)$$

for $\sum_{g\in\Sigma_n} c_g g \in k\Sigma_n$ and $f \in R_d$. We have a commutative diagram

$$\begin{tikzcd}
R_d \arrow[r] \arrow[d, hook] & R_d \otimes R_d \arrow[d, hook] \\
R \arrow[r, "\Delta"'] & R \otimes R.
\end{tikzcd}$$

Dualizing gives a commutative diagram

$$\begin{tikzcd} k\Sigma_n \otimes k\Sigma_n \arrow[r, "\Delta^\star"] \arrow[d, two heads] & k\Sigma_n \arrow[d, two heads] \\ R_d^\star \otimes R_d^\star \arrow[r] & R_d^\star. \end{tikzcd}$$

This means that the multiplication in the group algebra $k\Sigma_n$ induces a multiplication on $R_d^\star$. So the quotient map $k\Sigma_n \to R_d^\star$ is an algebra homomorphism.

Let $U = k^n$ be the representation of Σ_n on which Σ_n acts by permuting the coordinates. Define

$$V = (U \otimes U) \oplus U^m \oplus k.$$

This representation has dimension $r = n^2 + mn + 1$ and is given by a group homomorphism

$$\rho : \Sigma_n \to \mathrm{GL}_r(k).$$

We can write ρ as a matrix

$$\rho = \begin{pmatrix} \rho_{1,1} & \cdots & \rho_{1,r} \\ \vdots & & \vdots \\ \rho_{r,1} & \cdots & \rho_{r,r} \end{pmatrix}$$

where $\rho_{i,j} \in k^{\Sigma_n} = R$ for all i, j. The complexity of this representation is 2, meaning that $\rho_{i,j} \in R_2$ for all i and j.

We are now ready to define the category $\mathcal{C}_d$. Suppose that X_1, X_2 are affine subspaces of V. Let $S(X_1, X_2) \subseteq R_2$ be the subspace generated by all functions $\Sigma_n \to k$ of the form

$$g \in \Sigma_n \mapsto f(g \cdot v)$$

where $v \in X_1$, $g \in \Sigma_n$ and f an affine function on V which vanishes on X_2. Define

$$I_d(X_1, X_2) = \big(\big(S(X_1, X_2)\big)\big)_d.$$

We define

$$\mathrm{Hom}_d(X_1, X_2) = \big(R_d / I_d(X_1, X_2)\big)^\star \subseteq R_d^\star.$$

We will show that the multiplication in $R_d^\star$ induces a bilinear map

$$\mathrm{Hom}_d(X_1, X_2) \times \mathrm{Hom}_d(X_2, X_3) \to \mathrm{Hom}_d(X_1, X_3),$$

which will be the composition map in the category. For the proof in Section 2.4 we consider the Hopf algebra structure of R. For $f \in R$ we define $\iota(f) \in R = k^{\Sigma_n}$ as the function defined by

$$g \mapsto f\big(g^{-1}\big).$$

The ring homomorphism

$$\iota : R \to R$$

is called the *antipode map*. It is dual to the linear map $k\Sigma_n \to k\Sigma_n$ that sends g to g^{-1} for all $g \in \Sigma_n$. For a permutation matrix g, its inverse is equal to the transpose. So the map $\Sigma_n \to \Sigma_n$ that sends $g \in \Sigma_n$ to its inverse is the restriction of a transpose map $\mathrm{Mat}_{n,n}(k) \to \mathrm{Mat}_{n,n}(k)$ which is linear. From this it follows that $\iota(R_d) \subseteq R_d$ for all d.

The structure of R as a commutative ring, together with the co-multiplication, antipode map and co-unit satisfies the axioms of a Hopf algebra.

For $v \in V$ we define $\mu(v) \in V^{\Sigma_n}$ as the function that maps $g \in \Sigma_n$ to $g \cdot v$. We can identify

$$V^{\Sigma_n} \cong V \otimes k^{\Sigma_n} = V \otimes R.$$

This yields a linear map

$$\mu : V \to V \otimes R$$

with the following property: If $v \in V$ and $\mu(v) = \sum_i v_i \otimes f_i$, then $g \cdot v = \sum_i f_i(g) v_i$ for all $v \in V$ and $g \in \Sigma_n$. The map μ determines the representation V (see Definition 2.15).

2.3. Calculating with polynomial functions on the symmetric group

We will identify Σ_n with the permutation matrices in $\mathrm{Mat}_{n,n}(k)$. Let $x_{i,j}$ with $1 \leqslant i, j \leqslant n$ be the coordinate functions on $\mathrm{Mat}_{n,n}(k)$. The dimensions of $k[x_{1,1}, \dots, x_{n,n}]_{\leqslant d}$ and R_d are bounded by $\binom{n^2+d}{d} = O(n^{2d})$.

If $\mathfrak{a}$ is the vanishing ideal of $\Sigma_n \subseteq \mathrm{Mat}_{n,n}(k)$ then R is isomorphic to the quotient ring $k[x_{1,1}, \dots, x_{n,n}]/\mathfrak{a}$. If $\mathfrak{a}_{\leqslant d} = k[x_{1,1}, \dots, x_{n,n}]_{\leqslant d} \cap \mathfrak{a}$, then R_d is isomorphic to the quotient $k[x_{1,1}, \dots, x_{n,n}]_{\leqslant d}/\mathfrak{a}_{\leqslant d}$. If we are given an explicit Gröbner basis for $\mathfrak{a}$ with respect to some total degree monomial ordering, then it is easy to determine $\mathfrak{a}_{\leqslant d}$ and R_d. However, the author does not know of any explicit formulas for the Gröbner basis of $\mathfrak{a}$ for large n.

In this section, we will discuss an efficient way of calculating in R_d and $R_d^\star$ without using Gröbner bases. By definition, R_d consists of all restrictions of polynomials of degree $\leqslant d$ to the symmetric group Σ_n. Unfortunately, the cardinality of Σ_n is $n!$, which grows super-polynomially. We will see that it suffices to restrict polynomials to a small subset of Σ_n.

Lemma 2.7. *Let $H = \langle (1\ 2), \dots, (2d+1\ 2d+2) \rangle \subseteq \Sigma_n$ be the subgroup isomorphic to $(\mathbb{Z}/2\mathbb{Z})^{d+1}$. If f is a polynomial of degree $\leqslant d$ that vanishes on $H \setminus \{e\}$, then it also vanishes on e.*

Proof. Let Z be the vector space of all restrictions $F|_H$ where F is a polynomial on $\mathrm{Mat}_{n,n}(k)$ of degree $\leqslant d$. Suppose that $f \in Z$. There exists a polynomial F of degree $\leqslant d$ in the variables

$$x_{2i+1,2i+1},\ x_{2i+1,2i+2},\ x_{2i+2,2i+1},\ x_{2i+2,2i+2} \quad (0 \leqslant i \leqslant d)$$

whose restriction to H is the same as f. By replacing $x_{2i+2,2i+2}$ by $x_{2i+1,2i+1}$ and replacing $x_{2i+1,2i+2}$ and $x_{2i+2,2i+1}$ by $1 - x_{2i+1,2i+1}$, we may assume that F is a polynomial of degree $\leqslant d$ in the variables $x_{2i+1,2i+1}$, $i = 0, 1, 2, \dots, d$. Moreover, because $x_{2i+1,2i+1}^2 = x_{2i+1,2i+1}$ on H, we can also assume that F is multilinear. So F lies in the span of all square-free monomials in the variables $x_{2i+1,2i+1}$, $0 \leqslant i \leqslant d$ of degree $\leqslant d$. There are $2^{d+1} - 1$ such monomials, so $\dim Z \leqslant 2^{d+1} - 1$.

Suppose that $f \in Z$ vanishes on $H \setminus \{e\}$. Assume that $f(e) \neq 0$. We will derive a contradiction. For every permutation $g \in H$, the function

$$g \cdot f : h \mapsto f(gh)$$

also lies in Z. For all $g, h \in H$ we have

$$(g \cdot f)\bigl(h^{-1}\bigr) = f\bigl(gh^{-1}\bigr) \neq 0 \quad \Leftrightarrow \quad g = h.$$

This shows that the functions

$$g \cdot f, \qquad g \in H$$

are linearly independent, so $\dim Z \geqslant 2^{d+1}$. Contradiction. We conclude that $f(e) = 0$. □

Lemma 2.8. *Let $G_{d,n} \subseteq \Sigma_n$ be the subset of all permutations with at least $n - d$ fixed points. If $f \in R_d$ that vanishes on $G_{3d,n}$, then f vanishes on Σ_n.*

Proof. Suppose that $f \in R_d$ vanishes on $G_{3d,n}$. Let g be a permutation with $n-m$ fixed points. We will show by induction on m that $f(g)=0$. If $m \leqslant 3d$ then $f(g)=0$ by assumption. Suppose that $m \geqslant 3d+1$. Let $T \subseteq \{1,2,\dots,n\}$ be the set of non-fixed points of g. The set T has at least $3d+1$ elements. We can find a sequence $a_1,a_2,\dots,a_{d+1} \in \{1,2,\dots,n\}$ such that $a_j \notin \{g^{-1}(a_i), a_i, g(a_i)\}$ for all $i<j$. Let H be the group generated by the 2-cycles $(a_i, g(a_i))$ for $i=1,2,\dots,d+1$. Note that these 2-cycles are pairwise disjoint by construction. For every permutation $h \in H \setminus \{e\}$, hg has $>n-m$ fixed points. By the induction hypothesis, f vanishes on hg for all $h \in H \setminus \{e\}$. But then f also vanishes on $g = eg$ by Lemma 2.7. □

Lemma 2.9. *The cardinality of $G_{d,n}$ is at most n^d.*

Proof. For every subset S of $\{1,2,\dots,n\}$ with $n-d$ elements, there exist $d!$ permutations that fix S. Since there are $\binom{n}{d}$ such subsets, there are at most

$$d!\binom{n}{d} = \frac{n!}{(n-d)!} \leqslant n^d$$

permutations which have $\geqslant n-d$ fixed points. □

Let

$$\Lambda : R_d \to k^{G_{3d,n}}$$

be the function that maps $f \in R_d \subseteq k\Sigma_n$ to the restriction of f to $G_{3d,n}$. By Lemma 2.8 this map is injective. Dualizing yields a surjective map

$$\Lambda^\star : kG_{3d,n} \to R_d^\star.$$

If $f_1, f_2 \in R_d^\star$ then we can compute their product $f_1 f_2$ as follows. First, we find $F_1, F_2 \in kG_{3d,n}$ such that $\Lambda^\star(F_i) = f_i$ for $i=1,2$. Within the group algebra $k\Sigma_n$, we compute the product F_1F_2. Note that $F_1F_2 \in kG_{6d,n}$. Because F_1, F_2, F_1F_2 are sparse, we can compute the product F_1F_2 in polynomial time. Finally, $f_1f_2 = \Xi_d(F_1F_2)$.

2.4. Co-algebras associated to algebraic groups

Suppose that G is a linear algebraic group over k. Let $R := k[G]$ be the coordinate ring of G. The identity element $e \in G$ is defined over the field k. The multiplication $G \times G \to G$ corresponds to a homomorphism of k-algebras

$$\Delta : R \to R \otimes R,$$

where $\otimes$ denotes the tensor product as k-vector spaces. The ring R is a Hopf algebra with co-multiplication Δ, and co-unit $\sigma_e : R \to k$, where σ_e is evaluation at $e \in G$. The inverse function $G \to G$ defined by $g \mapsto g^{-1}$ defines an antipode map $\iota : R \to R$.

Suppose for the moment that k is algebraically closed. The group G acts on itself by left multiplication and it acts on the right by right multiplication. These actions correspond to right and left action of G on R respectively. If $g \in G$, then g acts on R on the right as the automorphism

$$(\sigma_g \otimes \mathrm{id}) \circ \Delta : R \to R$$

where $\sigma_g : R \to k$ is evaluation at g. The element g acts on the left by the automorphism

$$(\mathrm{id} \otimes \sigma_g) \circ \Delta : R \to R.$$

A subspace $W \subseteq R$ is stable under the right action if

$$\Delta(W) \subseteq R \otimes W$$

and stable under the left action if

$$\Delta(W) \subseteq W \otimes R.$$

It is stable under both actions if

$$\Delta(W) \subseteq W \otimes W.$$

Let k again be an arbitrary field, and suppose that $W \subseteq R$ is a finite-dimensional subspace such that

(1) W contains k;
(2) W generates R;
(3) $\Delta(W) \subseteq W \otimes W$;
(4) $\iota(W) \subseteq W$.

We will call such a subspace W a *stable generating subspace*.

Remark 2.10. For every algebraic group G, there exists a stable generating subspace $W \subseteq R = k[G]$. For example, if k is algebraically closed then we can construct such a subspace W as follows. Choose a finite-dimensional subspace $S \subseteq k[G]$ which contains k and generates the ring $k[G]$. By replacing S by $S + \iota(S)$ we may assume that $\iota(S) = S$. The group $G \times G$ acts on G by left and right multiplication. Let W be the space spanned by all $(g_1, g_2) \cdot f$ with $(g_1, g_2) \in G \times G$ and $f \in S$. Then W is finite dimensional and satisfies the four properties.

We define a filtration on R by $R_0 := k$ and $R_d := W^d$ for $d > 0$. We have $\Delta(R_d) \subseteq R_d \otimes R_d$ and $\iota(R_d) \subseteq R_d$, so R_d is a co-associative co-algebra. We will call $R = \bigcup_{d=0}^{\infty} R_d$ the *stable filtration of* R with respect to the stable generating subspace W. The dual $R_d^\star$ of R_d is an associative algebra.

Example 2.11. Suppose that $G = \mathbb{G}_a = (k, +)$ is the additive group. Then $R = k[\mathbb{G}_a]$ is isomorphic to $k[t]$, the polynomial ring in one variable. The identity element is $e = 0 \in k$. So the co-unit is σ_0, which is defined by:

$$\sigma_0\bigl(f(t)\bigr) = f(0).$$

The co-multiplication

$$\Delta : k[t] \to k[t] \otimes k[t] \cong k[t, s]$$

is defined by

$$\Delta\bigl(f(t)\bigr) = f(t + s)$$

and the $\iota : k[t] \to k[t]$ is defined by

$$\iota\bigl(f(t)\bigr) = f(-t).$$

We can take $W = k \oplus k \cdot t \subseteq k[t]$. Then we have $k \subseteq W$, W generates $k[t]$, $\Delta(W) \subseteq W \otimes W$ and $\iota(W) \subseteq W$. Now $R_d = W^d \subseteq k[t]$ consists of all polynomials of degree $\leqslant d$. This is a natural filtration on the ring $k[t]$.

Example 2.12. Suppose that $G = \mathbb{G}_m = (k^\star, \cdot)$ is the multiplicative group. Then $R = k[\mathbb{G}_m]$ is isomorphic to the ring $k[t, t^{-1}]$ of Laurent polynomials. The identity element is $e = 1 \in \mathbb{G}_m$. So the co-unit is

$$\sigma_1 : k\bigl[t, t^{-1}\bigr] \to k$$

defined by

$$\sigma_1\big(f(t)\big) = f(1).$$

The co-multiplication

$$\Delta : k\big[t, t^{-1}\big] \to k\big[t, t^{-1}\big] \otimes k\big[t, t^{-1}\big] \cong k\big[t, s, t^{-1}, s^{-1}\big]$$

is defined by

$$\Delta\big(f(t)\big) = f(ts)$$

and $\iota : k[t, t^{-1}] \to k[t, t^{-1}]$ is defined by

$$\iota\big(f(t)\big) = f\big(t^{-1}\big).$$

Define $W \subseteq k[t, t^{-1}]$ by $W = kt^{-1} \oplus k \oplus kt$. Then we have $k \subseteq W$, W generates $k[t, t^{-1}]$, $\Delta(W) \subseteq W \otimes W$ and $\iota(W) \subseteq W$. The space $R_d = W^d$ is the space of all Laurent polynomials of the form

$$\sum_{i=-d}^{d} a_i t^i$$

where $a_{-d}, a_{1-d}, \ldots, a_d \in k$.

Example 2.13. We go back to the notation of Sections 1.3 and 2.2. Let Y be a set of n vertices, and $\Sigma_n = \Sigma(Y)$. We identify Σ_n with the $n \times n$ permutation matrices and $R = k^{\Sigma_n} = k[\Sigma_n]$ is the space of all k-valued functions on Σ_n. Let $W = R_1$ be the space of restrictions of all linear functions on $\mathrm{Mat}_{n,n}(k)$ to Σ_n. Then W is a stable generating subspace of R, and $R_d = W^d$ consists of the restrictions of polynomials of degree $\leqslant d$ on $\mathrm{Mat}_{n,n}(k)$ to the set of permutation matrices.

Lemma 2.14. *Suppose that $R = \bigcup_{d=0}^{\infty} R_d$ is a stable filtration of $R = k[G]$ and assume that A, B, C are subspaces of R_d with*

$$\Delta(A) \subseteq B \otimes R_d + R_d \otimes C.$$

Then we have

$$\Delta\big((A)_d\big) \subseteq (B)_d \otimes R_d + R_d \otimes (C)_d \tag{1}$$

and

$$\Delta\big(((A))_d\big) \subseteq ((B))_d \otimes R_d + R_d \otimes ((C))_d. \tag{2}$$

Proof. For any $e \leqslant d$, the space $R_e/(B \cap R_e) \otimes R_e/(C \cap R_e)$ is a subspace of $R_d/B \otimes R_d/C$. It follows that

$$\begin{aligned}\Delta(A \cap R_e) &\subseteq \Delta(A) \cap \Delta(R_e) \subseteq (B \otimes R_d + R_d \otimes C) \cap (R_e \otimes R_e)\\ &= (B \cap R_e) \otimes R_e + R_e \otimes (C \cap R_e),\end{aligned}$$

where the last equality follows from elementary properties of tensor products. Therefore we have

$$\begin{aligned}\Delta\big((A \cap R_e)R_{d-e}\big) &= \Delta(A \cap R_e)\Delta(R_{d-e}) \subseteq \big((B \cap R_e) \otimes R_e + R_e \otimes (C \cap R_e)\big)(R_{d-e} \otimes R_{d-e})\\ &\subseteq \big((B \cap R_e)R_{d-e}\big) \otimes R_d + R_d \otimes \big((C \cap R_e)R_{d-e}\big).\end{aligned}$$

This shows (1). Now (2) follows by iteration. □

2.5. The complexity of a representation

Let G be a linear algebraic group over k and fix a stable generating subspace W.

Definition 2.15. A rational representation of G is a finite-dimensional vector space V with a k-linear map

$$\mu : V \to V \otimes R$$

such that the diagram

$$\begin{tikzcd}
V \arrow[r, "\mu"] \arrow[d, "\mu"'] & V \otimes R \arrow[d, "\mu\otimes\mathrm{id}"] \\
V \otimes R \arrow[r, "\mathrm{id}\otimes\Delta"'] & V \otimes R \otimes R
\end{tikzcd}$$

commutes, and $(\mathrm{id} \otimes \sigma_e) \circ \mu = \mathrm{id}$.

If k is algebraically closed, then we define

$$g \cdot w = (\mathrm{id} \otimes \sigma_g) \circ \mu(w)$$

for all $g \in G$ and $w \in V$. We have the following commutative diagram

$$\begin{tikzcd}
V \arrow[r, "\mu"] \arrow[d, "\mu"'] & V \otimes R \arrow[r, "\mathrm{id}\otimes\sigma_h"] \arrow[d, "\mu\otimes\mathrm{id}"'] & V \arrow[d, "\mu"] \\
V \otimes R \arrow[r, "\mathrm{id}\otimes\Delta"] \arrow[dr, "\mathrm{id}\otimes\sigma_{gh}"'] & V \otimes R \otimes R \arrow[r, "\mathrm{id}\otimes\mathrm{id}\otimes\sigma_h"] \arrow[d, "\mathrm{id}\otimes\sigma_g\otimes\sigma_h"] & V \otimes R \arrow[dl, "\mathrm{id}\otimes\sigma_g"] \\
 & V. &
\end{tikzcd}$$

This shows that

$$(gh) \cdot v = (\mathrm{id} \otimes \sigma_{gh}) \circ \mu(v) = (\mathrm{id} \otimes \sigma_g) \circ \mu \circ (\mathrm{id} \otimes \sigma_h) \circ \mu(v) = g \cdot (h \cdot v).$$

Definition 2.16. Suppose that V and V' are finite-dimensional rational representations of G given by $\mu : V \to V \otimes R$ and $\mu' : V' \to V' \otimes R$. A linear map $f : V \to V'$ is called G-equivariant if the following diagram commutes:

$$\begin{tikzcd}
V \arrow[r, "\mu"] \arrow[d, "f"'] & V \otimes R \arrow[d, "f\otimes\mathrm{id}"] \\
V' \arrow[r, "\mu'"'] & V' \otimes R.
\end{tikzcd}$$

If k is algebraically closed, $v \in V$ and $g \in G$ then we have

$$f(g \cdot v) = f \circ (\mathrm{id} \otimes \sigma_g) \circ \mu(v) = (\mathrm{id} \otimes \sigma_g) \circ (f \otimes \mathrm{id}) \circ \mu = (\mathrm{id} \otimes \sigma_g) \circ \mu \circ f(v) = g \cdot f(v).$$

Since V and $\mu(V)$ are finite dimensional, there exists a nonnegative integers l such that

$$\mu(V) \subseteq R_l \otimes V.$$

Let $\ell_W(V)$ be the smallest possible value of l for which this is true. The number $\ell_W(V)$ depends on the choice of W, but we will often drop the subscript and just write $\ell(V)$ if W is fixed. We can think of $\ell(V)$ as a measure of the complexity of the representation V.

Lemma 2.17.

(1) $\ell(V \oplus V') = \max\{\ell(V), \ell(V')\}$;
(2) $\ell(V \otimes V') \leqslant \ell(V) + \ell(V')$;
(3) $\ell(V) = \ell(V^\star)$.

Proof.

(1) This is straightforward.
(2) The representation V and V' are given by $\mu : V \to V \otimes R$ and $\mu' : V' \to V' \otimes R$. We have $\mu(V) \subseteq V \otimes R_{\ell(V)}$ and $\mu'(V') \subseteq V' \otimes R_{\ell(V')}$. The representation $V \otimes V'$ is given by the composition μ'' defined by:

$$\begin{tikzcd} V \otimes V' \arrow[r, "\mu\otimes\mu'"] & V \otimes R \otimes V' \otimes R \arrow[r, "\cong"] & V \otimes V' \otimes R \otimes R \arrow[r, "\mathrm{id}\otimes\mathrm{id}\otimes m"] & V \otimes V' \otimes R, \end{tikzcd}$$

where $m : R \otimes R \to R$ is the usual multiplication given by $\sum_i a_i \otimes b_i \mapsto \sum_i a_i b_i$. We have

$$\mu \otimes \mu'\bigl(V \otimes V'\bigr) \subseteq V \otimes R_{\ell(V)} \otimes V' \otimes R_{\ell(V')},$$

so

$$\begin{aligned}\mu''\bigl(V \otimes V'\bigr) &\subseteq (\mathrm{id} \otimes \mathrm{id} \otimes m)\bigl(V \otimes V' \otimes R_{\ell(V)} \otimes R_{\ell(V')}\bigr) \subseteq V \otimes V' \otimes R_{\ell(V)} R_{\ell(V')} \\ &\subseteq V \otimes V' \otimes R_{\ell(V)+\ell(V')}.\end{aligned}$$

(3) Let $\mu^\star : V^\star \to V^\star \otimes R$ be the dual representation of $\mu : V \to V \otimes R$. We have

$$(f \otimes \iota) \circ \mu(v) = (v \otimes \mathrm{id}) \circ \mu^\star(f)$$

where $v \in V = V^{\star\star}$ and $f \in V^\star$. If $\ell(V) = d$, then $\mu(v) \in V \otimes R_d$, and

$$(v \otimes \mathrm{id}) \circ \mu^\star(f) = (f \otimes \iota)\bigl(\mu(v)\bigr) \in (f \otimes \iota)(V \otimes R_d) \subseteq \iota(R_d) \subseteq R_d.$$

It follows that $\mu^\star(f) \subseteq V^\star \otimes R_d$. This shows that $\ell(V^\star) \leqslant \ell(V) \leqslant d$. Similarly, we have $\ell(V) = \ell(V^{\star\star}) \leqslant \ell(V^\star)$. □

Example 2.18. Suppose that $G = \mathbb{G}_a$ is the additive group as in Example 2.11 and that k is a field of characteristic 0. Let V_d be the $(d+1)$-dimensional indecomposable representation of $\mathbb{G}_a$ ($d \geqslant 0$). We can choose a basis $x_0, x_1, \ldots, x_d$ of V_d such that the action $\mu_d : V_d \to V_d \otimes R$ is given by:

$$\mu_d(x_i) = x_i \otimes 1 + x_{i-1} \otimes t + x_{i-2} \otimes \frac{t^2}{2!} + \cdots + x_0 \otimes \frac{t^i}{i!}.$$

It follows that $\mu(V_d) \subseteq V_d \otimes R_d$ and $\mu_d(V_d) \nsubseteq V_d \otimes R_{d-1}$. We conclude that $\ell(V_d) = d$. If V is any representation, then V is of the form

$$V = V_{d_1} \oplus \cdots \oplus V_{d_r}$$

and

$$\ell(V) = \max\{d_1, \ldots, d_r\}.$$

Example 2.19. Suppose that $G = \mathbb{G}_m$ is the multiplicative group as in Example 2.12. For $d \in \mathbb{Z}$, let $V_d \cong k$ be the irreducible 1-dimensional representation of $\mathbb{G}_m$ defined by $\mu_d : V_d \to V_d \otimes R$, where μ_d is given by

$$\mu_d(1) = 1 \otimes t^d.$$

Then we clearly have $\ell(V_d) = |d|$. Since V_d, $d \in \mathbb{Z}$ are all irreducible representations, any representation V can be written as

$$V = V_{d_1} \oplus \cdots \oplus V_{d_r}.$$

Then we have

$$\ell(V) = \max\bigl\{|d_1|, |d_2|, \ldots, |d_r|\bigr\}.$$

2.6. Definition of the approximate categories

For a subspace $Z \subseteq V$ we define $Z^\perp = \{f \in V^\star \mid \forall v \in V\ f(v) = 0\}$.

Suppose that $X_1, X_2 \in \mathrm{Aff}(V)$. If k is algebraically closed, then the inclusion $g \cdot X_1 \subseteq X_2$ $(g \in G)$ yields a system of polynomial equations in $k[G]$. Namely, if X_1 and X_2 are nonempty then we can write $X_1 = v_1 + Z_1$ and $X_2 = v_2 + Z_2$. Let $V^\star$ be the dual of V and $Z_2^\perp$ be the space of all $f \in V^\star$ which vanish on Z_2. For every function $f \in Z_2^\perp \subseteq V^\star$ and every $w \in X_1$ we have the equation $f(g \cdot w) = f(v_2)$. In other words,

$$(f \otimes \mathrm{id}) \circ \mu(w) - f(v_2) \otimes 1 \in k[G]$$

vanishes on g. The latter equation makes sense, even if k is not algebraically closed.

Definition 2.20. Let $S(X_1, X_2)$ be the span of all

$$(f \otimes \mathrm{id}) \circ \mu(w) - f(v_2) \otimes 1 = (f \otimes \mathrm{id})\bigl(\mu(w) - v_2 \otimes 1\bigr)$$

with $f \in Z_2^\perp$ and $w \in X_1$. We define $S(\emptyset, X) = \{0\}$ for $X \in \mathrm{Aff}(V)$ and $S(X, \emptyset) = \{1\}$ if $X \in \mathrm{Aff}(V) \setminus \{\emptyset\}$.

Lemma 2.21. *If $X_1, X_2, X_3 \in \mathrm{Aff}(V)$, then we have*

$$\Delta\bigl(S(X_1, X_3)\bigr) \subseteq S(X_2, X_3) \otimes R_{\ell(V)} + R_{\ell(V)} \otimes S(X_1, X_2).$$

Proof. Suppose that $X_i = v_i + Z_i$ for $i = 1, 2, 3$. We have

$$\mu(v_2) - v_3 \otimes 1 \in V \otimes S(X_2, X_3) + Z_3 \otimes R_{\ell(V)},$$

and

$$\mu(Z_2) \subseteq V \otimes S(X_2, X_3) + Z_3 \otimes R_{\ell(V)}.$$

It follows that

$$\begin{aligned}
&(\mathrm{id} \otimes \Delta)\bigl(\mu(v_1) - v_3 \otimes 1\bigr) \\
&\quad = (\mathrm{id} \otimes \Delta) \circ \mu(v_1) - v_3 \otimes 1 \otimes 1 \\
&\quad = (\mu \otimes \mathrm{id}) \circ \mu(v_1) - v_3 \otimes 1 \otimes 1 = (\mu \otimes \mathrm{id}) \circ \bigl(\mu(v_1) - v_2 \otimes 1\bigr) + \bigl(\mu(v_2) - v_3 \otimes 1\bigr) \otimes 1 \\
&\quad \in (\mu \otimes \mathrm{id})\bigl(V \otimes S(X_1, X_2) + Z_2 \otimes R_{\ell(V)}\bigr) + V \otimes S(X_2, X_3) \otimes R_{\ell(V)} + Z_3 \otimes R_{\ell(V)} \otimes R_{\ell(V)} \\
&\quad \subseteq V \otimes R_{\ell(V)} \otimes S(X_1, X_2) + V \otimes S(X_2, X_3) \otimes R_{\ell(V)} + Z_3 \otimes R_{\ell(V)} \otimes R_{\ell(V)}.
\end{aligned}$$

If $f \in Z_3^\perp$, then we have

$$\begin{aligned}
&\Delta\bigl((f \otimes \mathrm{id})\bigl(\mu(v_1) - v_3 \otimes 1\bigr)\bigr) \\
&\quad = (f \otimes \mathrm{id} \otimes \mathrm{id})\bigl(\mathrm{id} \otimes \Delta\bigl(\mu(v_1) - v_3 \otimes 1\bigr)\bigr) \\
&\quad \subseteq (f \otimes \mathrm{id} \otimes \mathrm{id})\bigl(V \otimes R_{\ell(V)} \otimes S(X_1, X_2) + V \otimes S(X_2, X_3) \otimes R_{\ell(V)} + Z_3 \otimes R_{\ell(V)} \otimes R_{\ell(V)}\bigr) \\
&\quad \subseteq S(X_2, X_3) \otimes R_{\ell(V)} + R_{\ell(V)} \otimes S(X_1, X_2). \qquad \square
\end{aligned}$$

For $X_1, X_2 \in \mathrm{Aff}(V)$, define

$$I_d(X_1, X_2) = \big(\big(S(X_1, X_2)\big)\big)_d$$

for all $d \geqslant \ell(V)$. We also define

$$I_\infty(X_1, X_2) = \big(S(X_1, X_2)\big) \subseteq R.$$

For $d \geqslant \ell(V)$ we have

$$I_d(X_1, X_2) \subseteq I_{d+1}(X_1, X_2) \subseteq \cdots,$$

where $I_\infty(X_1, X_2) = \bigcup_{j \geqslant d} I_j(X_1, X_2)$. For $d \geqslant \ell(V)$ we have a natural linear map $\psi_d : R_d/I_d(X_1, X_2) \to R_{d+1}/I_{d+1}(X_1, X_2)$. This gives us a chain of linear maps

$$R_d/I_d(X_1, X_2) \to R_{d+1}/I_{d+1}(X_2, X_3) \to R_{d+2}/I_{d+2}(X_1, X_2) \to \cdots. \tag{3}$$

There also is a natural linear map $\gamma_d : R_d/I_d(X_1, X_2) \to R/I_\infty(X_1, X_2)$ for all d. We have $\gamma_{d+1} \circ \psi_d = \gamma_d$ for all d. By Proposition 2.5, there exists a constant $C = C(\ell(V))$ such that γ_d is injective for large $d \geqslant C$. This implies that ψ_d is injective for large d. This shows that $R/I_\infty(X_1, X_2)$ is the direct limit of the chain (3).

If k is algebraically closed $g, h \in G$ and $gX_1 \subseteq X_2$, $hX_2 \subseteq X_3$ then we have $(hg)X_1 \subseteq X_3$. The corollary below expresses this simple fact in terms of truncated Hopf algebras, and it also makes sense and is true if k is not algebraically closed.

Corollary 2.22. *For $X_1, X_2, X_3 \in \mathrm{Aff}(V)$ we have*

$$\Delta\big(I_d(X_1, X_3)\big) \subseteq I_d(X_2, X_3) \otimes R_d + R_d \otimes I_d(X_1, X_2).$$

Proof. This follows from Lemma 2.14 and Lemma 2.21. □

For $X_1, X_2 \in \mathrm{Aff}(V)$ and $d \geqslant \ell(V)$, define

$$\mathrm{Hom}_d(X_1, X_2) = \big(R_d/I_d(X_1, X_2)\big)^\star.$$

We also define

$$\mathrm{Hom}_\infty(X_1, X_2) = \big(R/I_\infty(X_1, X_2)\big)^\star.$$

It follows from the definitions and Hilbert's Nullstellensatz that

$$\mathrm{Hom}_\infty(X_1, X_2) \neq 0 \quad \Leftrightarrow \quad I_\infty(X_1, X_2) \neq R \quad \Leftrightarrow \quad \exists g \in G(\bar{k}), \quad g \cdot X_1 \subseteq X_2.$$

Now $\mathrm{Hom}_\infty(X_1, X_2)$ is the inverse limit of the diagram

$$\cdots \to \mathrm{Hom}_{d+2}(X_1, X_2) \to \mathrm{Hom}_{d+1}(X_1, X_2) \to \mathrm{Hom}_d(X_1, X_2).$$

The map $\gamma_d^\star : \mathrm{Hom}_\infty(X_1, X_2) \to \mathrm{Hom}_d(X_1, X_2)$ is onto for $d \geqslant C$ large enough, where $C = C(\ell(V))$ is a constant depending on $\ell(V)$.

Definition 2.23. For $d \geqslant \ell(V)$, the category $\mathcal{C}_d(V)$ is the category whose objects are elements of $\mathrm{Aff}(V)$ and in which for $X_1, X_2 \in \mathrm{Aff}(V)$, $\mathrm{Hom}_d(X_1, X_2)$ is the set of morphisms from X_1 to X_2.

Corollary 2.24. *There exists a constant C such that for all $X_1, X_2 \in \mathrm{Aff}(V)$ and $d \geqslant C$, we have*

$$\mathrm{Hom}_d(X_1, X_2) \neq 0 \Leftrightarrow \exists g \in G(\bar{k}), \quad g \cdot X_1 \subseteq X_2$$

and

$$\begin{gathered} X_1, X_2 \textit{ are isomorphic in } \mathcal{C}_d(V) \\ \Updownarrow \\ \operatorname{Hom}_d(X_1, X_2) \neq 0 \textit{ and } \operatorname{Hom}_d(X_2, X_1) \neq 0 \\ \Updownarrow \\ X_1, X_2 \textit{ are in the same } G(\bar{k})\textit{-orbit.} \end{gathered}$$

We can view $\operatorname{Hom}_d(X_1, X_2)$ as a subspace of $R_d^\star$. From Corollary 2.22 it follows that Δ induces a linear map

$$\overline{\Delta} : R_d/I_d(X_1, X_3) \to R_d/I_d(X_2, X_3) \otimes R_d/I_d(X_1, X_2).$$

Dualizing gives a linear map

$$\operatorname{Hom}_d(X_1, X_2) \otimes \operatorname{Hom}_d(X_2, X_3) \to \operatorname{Hom}_d(X_1, X_3)$$

which corresponds to a bilinear multiplication

$$\operatorname{Hom}_d(X_1, X_2) \times \operatorname{Hom}_d(X_2, X_3) \to \operatorname{Hom}_d(X_1, X_3).$$

This multiplication is associative, because $R_d^\star$ is associative as a Hopf algebra.

Example 2.25. Consider the group $G = \mathbb{G}_m$ as in Examples 2.12 and 2.19. Let $V = V_3 \oplus V_5$. We have $\ell(V) = 5$. Let $v_1 = (1, 1), v_2 = (2, 1) \in V$. We will compute $\operatorname{Hom}_5(v_1, v_2)$. The equation $t \cdot v_1 = v_2$ gives us the equations

$$t^3 - 2,\ t^5 - 1.$$

So $S(v_1, v_2)$ is spanned by these two polynomials. We have

$$\begin{aligned} 2t^2 - 1 &= \left(t^5 - 1\right) - t^2 \cdot \left(t^3 - 2\right) \in \left(S(v_1, v_2)\right)_5, \\ t - 4 &= 2\left(t^3 - 2\right) - t\left(2t^2 - 1\right) \in \left(\left(S(v_1, v_2)\right)_5\right)_5 \subseteq I_5(v_1, v_2), \\ 31 &= \left(2t^2 - 1\right) - 2(t + 4)(t - 4) \in I(v_1, v_2)_5. \end{aligned}$$

Let us assume that 31 is invertible in k. Then we have $1 \in I(v_1, v_2)_5$, so $I(v_1, v_2)_5 = R_5$ and $\operatorname{Hom}_5(v_1, v_2) = 0$.

Suppose that $X_1 = \{v_1\}$ and $X_2 = \{(x, x) \mid x \in k\} \subseteq V$. We will compute $\operatorname{Hom}_5(X_1, X_2)$. The subspace X_2 is defined by $x_2 - x_1 = 0$, and $t \cdot v_1 = (t^3, t^5)$, so $S(X_1, X_2)$ is spanned by the polynomial $t^5 - t^3$. We have $t^2 - 1 = t^{-3}(t^5 - t^3) \in (S(X_1, X_2))_5$. We have that

$$\left(S(X_1, X_2)\right)_5 = I(X_1, X_2)_5$$

is the space spanned by

$$t^3\left(t^2 - 1\right),\ t^2\left(t^2 - 1\right),\ \ldots,\ t^{-5}\left(t^2 - 1\right).$$

The space $R_5/I_5(X_1, X_2)$ is 2-dimensional and spanned by $1 + I_5(X_1, X_2), t + I_5(X_1, X_2)$. So $\operatorname{Hom}_5(X_1, X_2)$ is 2-dimensional as well.

Remark 2.26. Let $G = \mathbb{G}_a$ as in Examples 2.11 and 2.18. Suppose that V is a representation with $\ell(V) \leqslant d$, and $X_1, X_2 \in \operatorname{Aff}(V)$. Then $S(X_1, X_2)$ is spanned by polynomials of degree $\leqslant d$. From the Euclidean algorithm in $k[\mathbb{G}_a] \cong k[t]$ it follows that $(S(X_1, X_2)) \cap R_d = I_d(X_1, X_2)$. So we have

$$\operatorname{Hom}_d(X_1, X_2) \neq 0 \quad \Leftrightarrow \quad \exists t \in \bar{k}, \quad t \cdot X_1 \subseteq X_2.$$

In particular, X_1 and X_2 are in the same $G(\bar{k})$-orbit if and only if $X_1 \cong_d X_2$.

The same result holds for $G = \mathbb{G}_m$, because $k[\mathbb{G}_m] = k[t, t^{-1}]$ is also a Euclidean domain. For other groups G, it is possible that $X_1 \cong_d X_2$ for $X_1, X_2 \in \mathrm{Aff}(V)$ with $\ell(V) = d$, but X_1, X_2 are not in the same $G(\overline{k})$-orbit. Still, we know that if X_1, X_2 are not in the same $G(\overline{k})$-orbit, then $X_1 \ncong_e X_2$ for some $e \gg 0$.

3. Properties of the approximate categories

3.1. Some elementary properties

If Z is a representation with $\ell(Z) \leqslant d$, then the homomorphism $\mu : Z \to Z \otimes R$ restricts to

$$\mu : Z \to Z \otimes R_d.$$

If $f \in R_d^\star$, then $(\mathrm{id} \otimes f) \circ \mu$ is an endomorphism of Z. One can verify that Z has the structure of an $R_d^\star$-module, where the multiplication is defined by

$$f \cdot w := (\mathrm{id} \otimes f) \circ \mu(w).$$

For any $X_1, X_2 \in \mathrm{Aff}(V)$, $\mathrm{Hom}_d(X_1, X_2)$ is a subspace of $R_d^\star$, so $\mathrm{Hom}_d(X_1, X_2)$ acts on any representation Z with $\ell(Z) \leqslant d$.

Lemma 3.1. *Suppose that $X_1 = v_1 + Z_1$ and $X_2 = v_2 + Z_2$. If $f \in \mathrm{Hom}_d(X_1, X_2)$ then $f \cdot X_1 \subseteq f(1)v_2 + Z_2$.*

Proof.
For $w \in X_1$ and $h \in Z_2^\perp$ we have

$$(h \otimes \mathrm{id}) \circ \mu(w) = h(v_2) \otimes 1.$$

If we apply f, we have

$$h(f \cdot w) = h \circ (\mathrm{id} \otimes f) \circ \mu(w) = f \circ (h \otimes \mathrm{id}) \circ \mu(w) = f\bigl(h(v_2) \otimes 1\bigr) = f(1)h(v_2).$$

Therefore, $h(f \cdot w - f(1)v_2) = 0$, so we get that $f \cdot w \in f(1)v_2 + Z_2$. □

Corollary 3.2. *If $X_1 \cong_d X_2$ then $\dim X_1 = \dim X_2$.*

Lemma 3.3. *Suppose that $0 \in X_1$ and $0 \notin X_2$. Then $\mathrm{Hom}_d(X_1, X_2) = 0$.*

Proof. Write $X_2 = v_2 + Z_2$, and choose $f \in Z_2^\perp$ with $f(v_2) \neq 0$. Then

$$(f \otimes \mathrm{id}) \circ \mu(0) - f(v_2) \otimes 1 = -f(v_2) \otimes 1$$

is a nonzero multiple of $1 \in R_{\ell(V)}$. It follows that $I_d(X_1, X_2) = R_d$ and $\mathrm{Hom}_d(X_1, X_2) = 0$. □

3.2. Constructible equivariants

Definition 3.4. Suppose that d is a positive integer. We inductively define the notion of a d-constructible equivariant:

(1) if $f : V \to V'$ is G-equivariant and linear, and $\ell(V), \ell(V') \leqslant d$, then f is d-constructible;
(2) if $f_1, f_2 : V \to V'$ are d-constructible, and $\lambda_1, \lambda_2 \in k$, then $\lambda_1 f_1 + \lambda_2 f_2$ is d-constructible;
(3) if $\ell(V_1) + \ell(V_2) \leqslant d$, then the bilinear map $V_1 \oplus V_2 \to V_1 \otimes V_2$ defined by $(v_1, v_2) \mapsto v_1 \otimes v_2$ is d-constructible;
(4) if $f_1 : V_1 \to V_2$ and $f_2 : V_2 \to V_3$ are d-constructible, then the composition $f_2 \circ f_1$ is d-constructible.

We will denote the class of d-constructible equivariants with domain V by $\mathscr{E}_d(V)$.

Proposition 3.5. *Suppose that $f : V \to V'$ is a d-constructible equivariant with $f(v_1) = 0$ and $f(v_2) \neq 0$. Then we have* $\operatorname{Hom}_d(v_1, v_2) = 0$. *In particular, v_1 and v_2 are not isomorphic in $\mathcal{C}_d(V)$.*

The proof will be given after Lemma 3.20.

3.3. The class of $(2d)$-constructible equivariants $\mathscr{E}_{2d}$ is at least as powerful as $\mathscr{C}_d$, the d-variable order logic with counting

Consider again the setup of Sections 1.3 and 2.2. Let Y be a set with n elements. Consider the symmetric group $G = \Sigma(Y) \cong \Sigma_n$ and let $U \cong k^n$ be the vector space with basis Y. We will write $U^{\otimes m}$ for

$$\underbrace{U \otimes \cdots \otimes U}_{m}.$$

To a subset $S \subseteq Y^m$, we can associate an element $\operatorname{tensor}(S) \in U^{\otimes m}$ defined by

$$\operatorname{tensor}(S) = \sum_{(x_1,\dots,x_m)\in Y^m} x_1 \otimes \cdots \otimes x_m.$$

We can construct a bilinear multiplication $\star : U^{\otimes m} \oplus U^{\otimes m} \to U^{\otimes m}$ as follows. If $x_1, \dots, x_m, y_1, \dots, y_m \in Y$, then we define

$$(x_1 \otimes \cdots \otimes x_m) \star (y_1 \otimes \cdots \otimes y_m) = \begin{cases} x_1 \otimes \cdots \otimes x_m & \text{if } (x_1, \dots, x_m) = (y_1, \dots, y_m); \\ 0 & \text{otherwise.} \end{cases}$$

Define $\mathbf{1} = \sum_{x \in Y} x$. For every i, we define the linear map $\operatorname{pr}_i : U^{\otimes d} \to U^{\otimes d}$ by

$$\operatorname{pr}_i(x_1 \otimes \cdots \otimes x_d) = x_1 \otimes \cdots \otimes x_{i-1} \otimes \mathbf{1} \otimes x_{i+1} \otimes \cdots \otimes x_m$$

for $x_1, \dots, x_m \in Y$. Note that $\frac{1}{n}\operatorname{pr}_i$ is a projection.

For $m \leqslant d$, the equivariant maps $\star$ and pr_i defined above lie in $\mathscr{E}_{2d}(U^{\otimes d})$.

Suppose that $m_1, \dots, m_s$ are positive integers. Define

$$V = U^{\otimes m_1} \oplus \cdots \oplus U^{\otimes m_s} \oplus k$$

and let $\mathscr{E}_d = \mathscr{E}_d(V)$. Define $\mathbf{1}_d : V \to U^{\otimes d}$ by

$$\mathbf{1}_d(v_1, \dots, v_s, a) = a(\mathbf{1} \otimes \cdots \otimes \mathbf{1}).$$

Then $\mathbf{1}_d$ lies in $\mathscr{E}_{2d}(V)$.

If $Y_i \subseteq Y^{m_i}$ for $i = 1, 2, \dots, s$, then $\Gamma = \langle Y, Y_1, \dots, Y_s\rangle$ is a structure with s relational symbols. Let $\mathscr{L}_d = \mathscr{L}_d(m_1, \dots, m_s)$ be the d-variable first order language for this structure, and let $\mathscr{C}_d = \mathscr{C}_d(m_1, \dots, m_s)$ be the d-variable language with counting. For $\Gamma = \langle Y, Y_1, \dots, Y_s\rangle$, define

$$A_\Gamma := \big(\operatorname{tensor}(Y_1), \dots, \operatorname{tensor}(Y_s), 1\big) \in V.$$

Definition 3.6. Suppose that $\varphi(x_1, \dots, x_d)$ is a formula in $\mathscr{C}_d$, and

$$f : V \to U^{\otimes d}.$$

We say that f represents φ, if

$$f(A_\Gamma) = \sum_{\Gamma \models \varphi(x_1,\dots,x_d)} x_1 \otimes \cdots \otimes x_d$$

for all $Y_1, \dots, Y_s$.

Theorem 3.7. *Suppose that k is a field of characteristic 0 or $p > n$. Then every formula $\varphi(x_1, \dots, x_d)$ in $\mathscr{C}_d$ is represented by an equivariant $f \in \mathscr{E}_{2d}(V)$.*

Proof. For $y_1, \dots, y_{m_i} \in \{x_1, \dots, x_d\}$, the formula $Y_i(y_1, \dots, y_{m_i})$ is represented by an equivariant linear map

$$V \to U^{\otimes d}.$$

The formula $x_i = x_j$ is represented by an equivariant linear map.

Suppose that $\varphi_1(x_1, \dots, x_d)$ and $\varphi_2(x_1, \dots, x_d)$ are represented by the covariants $f_1, f_2 \in \mathscr{E}_{2d}(V)$ respectively. Then $f_1 \star f_2$ represents the formula $\varphi_1 \wedge \varphi_2$, and $f_1 \star f_2 \in \mathscr{E}_{2d}(V)$.

If $\varphi(x_1, \dots, x_d)$ is represented by $f \in \mathscr{E}_{2d}(V)$, then $\neg\varphi(x_1, \dots, x_d)$ is represented by $\mathbf{1}_d - f$.

Suppose that $q(t)$ is a polynomial in t. Define an equivariant $[q(t)] : U^{\otimes d} \to U^{\otimes d}$ by

$$\bigl[q(t)\bigr]\biggl(\sum_{x_1,\dots,x_d \in Y} a_{x_1,\dots,x_d} x_1 \otimes \dots \otimes x_d\biggr) = \sum_{x_1,\dots,x_d \in Y} q(a_{x_1,\dots,x_d}) x_1 \otimes \dots \otimes x_d.$$

If we write $q(t) = tu(t) + a$ then we have

$$\bigl[q(t)\bigr](v) = \bigl[u(t)\bigr](v) \star v + av.$$

It follows by induction on the degree of $q(t)$ that $[q(t)]$ lies in $\mathscr{E}_{2d}(U^{\otimes d})$ for all polynomials $q(t)$.

Suppose that $\varphi(x_1, \dots, x_d)$ is represented by an equivariant f. There exists a polynomial $q(t) \in k[t]$ with $q(b) = 1$ and $q(j) = 0$ for $j \in \{1, 2, \dots, n\} \setminus \{b\}$. The formula

$$\exists_b x_i \, \varphi(x_1, \dots, x_d)$$

is represented by the covariant $[q(t)] \circ \mathrm{pr}_i \circ f$ for $i = 1, 2, \dots, d$. □

Corollary 3.8. *Suppose that k is a field of characteristic 0 or $p > n$. Suppose that $\Gamma = \langle Y, Y_1, \dots, Y_s\rangle$, $\Gamma' = \langle Y, Y_1', \dots, Y_s'\rangle$ are two structures (for example graphs) and*

$$f(A_\Gamma) = 0 \Leftrightarrow f(A_{\Gamma'}) = 0$$

for every $f \in \mathscr{E}_{2d}(V)$. Then we have

$$\Gamma \models \varphi \Leftrightarrow \Gamma' \models \varphi'$$

for every closed formula φ in $\mathscr{C}_d$.

3.4. Constructible functors

For the following definition, the reader should bear in mind that $\mathrm{Hom}_d(X_1, X_2)$ is a subspace of $R_d^\star$ for every representation V with $\ell(V) \leqslant d$ and every two objects X_1, X_2 of $\mathcal{C}_d(V)$.

Definition 3.9. We will call a covariant functor $\mathcal{F} : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ *very faithful* if $\mathcal{F}(\phi) = \phi$ for every morphism. A contravariant functor $\mathcal{F} : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ is called *very faithful* if $\mathcal{F}(\phi) = \phi \circ \iota$ for every morphism ϕ.

Note that a very faithful functor $\mathcal{F} : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ is uniquely determined by how it acts on objects.

Lemma 3.10. *Suppose that $\ell(V), \ell(V') \leqslant d$. There exist very faithful covariant functors $\mathcal{F}, \mathcal{G}, \mathcal{H} : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ such that $\mathcal{F}(X) = \emptyset$, $\mathcal{G}(X) = \{0\}$, $\mathcal{H}(X) = V'$ for all $X \in \mathrm{Aff}(V)$. Also, for every $\lambda \in k$ there exists a very faithful functor $\mathcal{I} : \mathcal{C}_d(V) \to \mathcal{C}_d(k)$ such that $\mathcal{I}(X) = \{\lambda\} \in \mathrm{Aff}(k)$ for all $X \in \mathrm{Aff}(V)$.*

Proof. This is clear because

$$\mathrm{Hom}_d(\emptyset,\emptyset)=\mathrm{Hom}_d\big(\{0\},\{0\}\big)=\mathrm{Hom}_d\big(V',V'\big)=R_d^\star$$

and for $\{\lambda\}\in\mathrm{Aff}(k)$, we have $\mathrm{Hom}_d(\{\lambda\},\{\lambda\})=R_d^\star$ as well. □

Lemma 3.11. *Suppose that $\ell(V),\ell(V')\leqslant d$ and $f:V\to V'$ is G-equivariant and linear. Then there exists a very faithful covariant functor $\mathcal{F}:\mathcal{C}_d(V)\to\mathcal{C}_d(V')$ such that $\mathcal{F}(X)=f(X)$ for all $X\in\mathrm{Aff}(V)$.*

Proof. Suppose that $X_1=v_1+Z_1$, $X_2=v_2+Z_2$ lie in $\mathcal{C}_d(V)$. Because f is equivariant, we have a commutative diagram

$$\begin{tikzcd} V \arrow[r, "\mu"] \arrow[d, "f"'] & V\otimes R_d \arrow[d, "f\otimes\mathrm{id}"] \\ V' \arrow[r, "\mu"'] & V'\otimes R_d. \end{tikzcd}$$

We can write $f(X_1)=f(v_1)+f(Z_1)$ and $f(X_2)=f(v_2)+f(Z_2)$. The space $S(f(X_1),f(X_2))$ is spanned by elements of the form

$$(g\otimes\mathrm{id})\circ\mu\big(f(w)\big)-g\big(f(v_2)\big)\otimes 1$$

where $g\in f(Z_2)^\perp$ and $w\in X_1$. Define $h=g\circ f\in Z_2^\perp$. We have

$$\begin{aligned}(g\otimes\mathrm{id})\circ\mu\big(f(w)\big)-g\big(f(v_2)\big)\otimes 1&=(g\otimes\mathrm{id})\circ(f\otimes\mathrm{id})\circ\mu(w)-g\big(f(v_2)\big)\otimes 1\\&=(h\otimes\mathrm{id})\circ\mu(w)-h(v_2)\otimes 1\in S_d(X_1,X_2).\end{aligned}$$

This shows that $S(f(X_1),f(X_2))\subseteq S(X_1,X_2)$. Following the definitions, it is easy to see that this implies $\mathrm{Hom}_d(X_1,X_2)\subseteq\mathrm{Hom}_d(\phi(X_1),\phi(X_2))$. □

Lemma 3.12. *Suppose that $\ell(V),\ell(V'),\ell(V'')\leqslant d$, and $\mathcal{F}':\mathcal{C}_d(V)\to\mathcal{C}_d(V')$ and $\mathcal{F}'':\mathcal{C}_d(V)\to\mathcal{C}_d(V'')$ are very faithful covariant (resp. contravariant) functors. Then there exists a very faithful covariant (resp. contravariant) functor $\mathcal{F}:\mathcal{C}_d(V)\to\mathcal{C}_d(V'\oplus V'')$ such that $\mathcal{F}(X)=\mathcal{F}'(X)\oplus\mathcal{F}''(X)$ for all $X\in\mathrm{Aff}(V)$.*

Proof. Suppose that $X_1,X_2\in\mathrm{Aff}(V)$, and let $X_1'=\mathcal{F}'(X_1)$, $X_2'=\mathcal{F}'(X_2)$, $X_1''=\mathcal{F}''(X_1)$, $X_2''=\mathcal{F}''(X_2)$. It is straightforward to verify that

$$S\big(X_1'\oplus X_1'',X_2'\oplus X_2''\big)=S\big(X_1',X_2'\big)+S\big(X_1'',X_2''\big).$$

We have

$$S\big(X_1',X_2'\big)\subseteq\big(\!\big(S(X_1,X_2)\big)\!\big)_d.$$

Similarly, we have

$$S\big(X_1'',X_2''\big)\subseteq\big(\!\big(S(X_1,X_2)\big)\!\big)_d,$$

so

$$\big(\!\big(S\big(X_1'\oplus X_1'',X_2'\oplus X_2''\big)\big)\!\big)_d\subseteq\big(\!\big(S(X_1,X_2)\big)\!\big)_d.$$

This implies that

$$\mathrm{Hom}_d(X_1,X_2)\subseteq\mathrm{Hom}_d\big(X_1'\oplus X_1'',X_2'\oplus X_2''\big)=\mathrm{Hom}_d\big(\mathcal{F}(X_1),\mathcal{F}(X_2)\big).\qquad\square$$

Definition 3.13. Suppose that $X \subseteq V$, $X' \subseteq V'$ are affine subspace. We define $X \otimes X' \subseteq V \otimes V'$ as the affine subspace spanned by all $x \otimes x'$ with $x \in X$ and $x' \in X'$. Suppose we write $X = v + Z$ and $X' = v' + Z'$ where $v \in V$, $v' \in V'$ and $Z \subseteq X, Z' \subseteq X'$ are subspaces. Then we have

$$X' \otimes X' = v \otimes v' + Z \otimes Z' + kv \otimes Z' + Z \otimes kv'.$$

Lemma 3.14. *Suppose that $\ell(V), \ell(V') + \ell(V'') \leqslant d$, and $\mathcal{F}' : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ and $\mathcal{F}'' : \mathcal{C}_d(V) \to \mathcal{C}_d(V'')$ are very faithful functors. Assume that both are covariant (resp. contravariant). Then there exists a very faithful covariant (resp. contravariant) functor $\mathcal{F} : \mathcal{C}_d(V) \to \mathcal{C}_d(V' \otimes V'')$ such that $\mathcal{F}(X) = \mathcal{F}'(X) \otimes \mathcal{F}''(X)$ for all $X \in \mathrm{Aff}(V)$.*

Proof. Let $e = \ell(V')$. One can verify that

$$\begin{aligned} S\bigl(X_1' \otimes X_1'', X_2' \otimes X_2''\bigr) &\subseteq S\bigl(X_1', X_2'\bigr)R_{d-e} + R_e S\bigl(X_1'', X_2''\bigr) \\ &\subseteq \bigl(\bigl(S(X_1, X_2)\bigr)\bigr)_e R_{d-e} + R_e \bigl(\bigl(S(X_1, X_2)\bigr)\bigr)_{d-e} \subseteq \bigl(\bigl(S(X_1, X_2)\bigr)\bigr)_d. \end{aligned}$$

It follows that

$$\bigl(\bigl(S\bigl(X_1' \otimes X_1'', X_2' \otimes X_2''\bigr)\bigr)\bigr)_d \subseteq \bigl(\bigl(S(X_1, X_2)\bigr)\bigr)_d$$

and therefore

$$\mathrm{Hom}_d(X_1, X_2) \subseteq \mathrm{Hom}_d\bigl(X_1' \otimes X_1'', X_2' \otimes X_2''\bigr) = \mathrm{Hom}_d\bigl(\mathcal{F}(X_1), \mathcal{F}(X_2)\bigr). \qquad \square$$

Definition 3.15. If $X \subseteq V$ is an affine subspace, then we define

$$X^+ = \bigl\{ f \in V^\star \bigm| f(x) = 1 \text{ for all } x \in X \bigr\}.$$

If $0 \in X$, then $X^+ = \emptyset$. If $0 \notin X$, then $X^{++} = X$.

Lemma 3.16. *Suppose that $\ell(V) \leqslant d$. There exists a very faithful contravariant functor $\mathcal{D} : \mathcal{C}_d(V) \to \mathcal{C}_d(V^\star)$ such that*

$$\mathcal{D}(X) = X^+$$

for all $X \in \mathrm{Aff}(V)$.

Proof. Suppose that $X_1, X_2 \in \mathrm{Aff}(V)$. The action of G on V is given by

$$\mu : V \to V \otimes R_d.$$

The action of G on $V^\star$ is given by

$$\mu^\star : V^\star \to V^\star \otimes R_d$$

such that

$$(h \otimes \iota) \circ \mu(v) = (v \otimes \mathrm{id}) \circ \mu^\star(h)$$

for all $h \in V^\star$, $v \in V = V^{\star\star}$. Suppose that $X_1 = v_1 + Z_1, X_2 = v_2 + Z_2 \in \mathrm{Aff}(V)$. The case $0 \in X_2$ is clear, because then $\mathcal{D}(X_2) = \emptyset$ and $\mathrm{Hom}_d(\mathcal{D}(X_2), \mathcal{D}(X_1)) = R_d^\star$. The case where $0 \in X_2$ and $0 \notin X_1$ is also clear, because we have $\mathrm{Hom}_d(X_1, X_2) = 0$. So we may assume that $0 \notin X_1$ and $0 \notin X_2$. Choose $u_1, u_2 \in V^\star$ with $u_1(X_1) = \{1\}$, and $u_2(X_2) = \{1\}$. Then we have $\mathcal{D}(X_1) = u_1 + Y_1$ and $\mathcal{D}(X_2) = u_2 + Y_2$, where $Y_i = (kv_i + Z_i)^\perp$ for $i = 1, 2$. The space $S(\mathcal{D}(X_2), \mathcal{D}(X_1))$ is spanned by elements of the form

$$(f \otimes \mathrm{id}) \circ \mu^\star(w) - f(u_1) \otimes 1 = (w \otimes \iota) \circ \mu(f) - f(u_1) \otimes 1 \tag{4}$$

with $f \in Y_1^{\perp} = kv_1 + Z_1$, and $w \in u_2 + Y_2 \subseteq Z_2^{\perp}$. In fact we only need those f for which $f \in X_1 = v_1 + Z_1$. Then (4) is equal to

$$\begin{aligned}(f \otimes \mathrm{id}) \circ \mu^{\star}(w) - 1 = (w \otimes \iota) \circ \mu(f) - 1 = (w \otimes \iota) \circ \mu(f) - w(v_1)\\ = \iota\big((w \otimes \mathrm{id}) \circ \mu(f) - w(v_1)\big) \in \iota\big(S(X_1, X_2)\big).\end{aligned}$$

From this it follows that

$$\mathrm{Hom}_d(X_1, X_2) \subseteq \iota^{\star}\big(\mathrm{Hom}_d\big(\mathcal{D}(X_2), \mathcal{D}(X_1)\big)\big). \qquad \square$$

Definition 3.17. We inductively define the notion of a d-constructible functor.

(1) The constant functors $\mathcal{F}$, $\mathcal{G}$, $\mathcal{H}$, $\mathcal{I}$ in Lemma 3.10 and the duality functor $\mathcal{D}$ in Lemma 3.16 are d-constructible. The functor $\mathcal{F}$ associated to a G-equivariant linear map as in Lemma 3.11 is d-constructible.
(2) If $\mathcal{F}'$, $\mathcal{F}''$ are as in Lemma 3.12 and they are d-constructible, then the very faithful functor $\mathcal{F}$ defined by $\mathcal{F}(X) = \mathcal{F}'(X) \oplus \mathcal{F}''(X)$ is d-constructible.
(3) If $\mathcal{F}'$, $\mathcal{F}''$ are as in Lemma 3.14 and they are d-constructible, then the very faithful functor $\mathcal{F}$ defined by $\mathcal{F}(X) = \mathcal{F}'(X) \otimes \mathcal{F}''(X)$ is d-constructible.
(4) A composition of d-constructible functors is again d-constructible.

Corollary 3.18. *If $X_1, X_2 \in \mathrm{Aff}(V)$, $X_1 \cong_d X_2$ and $\mathcal{F} : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ is a d-constructible functor, then $\mathcal{F}(X_1) \cong_d \mathcal{F}(X_2)$. In particular, we have*

$$\dim \mathcal{F}(X_1) = \dim \mathcal{F}(X_2).$$

Lemma 3.19. *Suppose that $\mathcal{F}, \mathcal{F}' : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ are d-constructible functors, either both covariant or both contravariant. Then there exists a d-constructible functor $\mathcal{G} : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ defined by $\mathcal{G}(X) = \mathcal{F}(X) \cap \mathcal{F}'(X)$.*

Proof. If $0 \notin X_1$ and $0 \notin X_2$, then we have $(X_1^+ + X_2^+)^+ = X_1 \cap X_2$. In $V \oplus k$, we have

$$\big((X_1 \times \{1\})^+ + (X_2 \times \{1\})^+\big)^+ = X_1 \cap X_2 \times \{1\}.$$

So if $\mathcal{I} : \mathcal{C}_d(V) \to \mathcal{C}_d(V \oplus k)$ is just the inclusion, and $\mathcal{P} : \mathcal{C}_d(V \oplus k) \to \mathcal{C}_d(V)$ is just the projection, then we define $\mathcal{G}$ by

$$\mathcal{G}(X) = \mathcal{P} \circ \mathcal{D}\big(\mathcal{D} \circ \mathcal{I} \circ \mathcal{F}(X) + \mathcal{D} \circ \mathcal{I} \circ \mathcal{F}'(X)\big). \qquad \square$$

Lemma 3.20. *Suppose that $f : V \to V'$ is a d-constructible equivariant. Then there exists a d-constructible functor $\mathcal{F} : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ with $\mathcal{F}(\{v\}) = \{f(v)\}$ for all $v \in V$.*

Proof. This follows easily from the inductive definitions (Definitions 3.4 and 3.17). □

Proof of Proposition 3.5. Suppose that $f : V \to V'$ is a d-constructible equivariant with $f(v_1) = 0$ and $f(v_2) \neq 0$. By Lemma 3.20 there is a d-constructible functor $\mathcal{F} : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ with $\mathcal{F}(\{v\}) = \{f(v)\}$ for all $v \in V$. We have $\mathrm{Hom}_d(v_1, v_2) \subseteq \mathrm{Hom}_d(f(v_1), f(v_2)) = 0$ by Lemma 3.3. □

Lemma 3.21. *Suppose that $v_1, v_2 \in V$ and*

$$\dim \mathcal{F}(v_1) = \dim \mathcal{F}(v_2)$$

for all d-constructible functors $\mathcal{F}$. Then $f(v_1) = 0 \Leftrightarrow f(v_2) = 0$ for every d-constructible equivariant.

Proof. Suppose that $f : V \to V'$ is a d-constructible equivariant. There exists a d-constructible functor $\mathcal{F} : \mathcal{C}_d(V) \to \mathcal{C}_d(V')$ with $\mathcal{F}(\{w\}) = \{f(w)\}$ for every $w \in V$ by Lemma 3.20. Define a d-constructible functor $\mathcal{F}'$ with $\mathcal{F}'(X) = \mathcal{F}(X) \cap \{0\}$. Suppose that $w \in V$. If $f(w) = 0$, then $\mathcal{F}'(\{w\}) = \{0\}$ and $\dim \mathcal{F}'(\{w\}) = 0$. If $f(w) \neq 0$, then $\mathcal{F}'(\{w\}) = \emptyset$ and $\dim \mathcal{F}'(\{w\}) = -\infty$. □

4. The module isomorphism problem

4.1. Reformulation of the module isomorphism problem

Suppose that M and N are (left) n-modules of the free associative algebra $T = k\langle x_1, \dots, x_r\rangle$, and we would like to test whether M and N are isomorphic. We can choose a basis in M and identify M with k^n. The action of x_i is given by a matrix A_i. Similarly we can identify N with k^n. The action of x_i is given by a matrix B_i. An isomorphism is an invertible linear map $C : M \to N$ such that $CA_i = B_iC$ for all i. This is equivalent to $CA_iC^{-1} = B_i$ for all i. Let $V = \mathrm{Mat}_{n,n}(k)^r$. Then $\mathrm{GL}_n(k)$ acts on V by simultaneous conjugation. The following lemma follows from the discussion above:

Lemma 4.1. *The modules M and N are isomorphic if and only if $A = (A_1, \dots, A_r)$ and $B = (B_1, \dots, B_r)$ lie in the same $\mathrm{GL}_n(k)$-orbit.*

Proposition 4.2. *Let $\bar{k}$ be the algebraic closure of k. The modules $M \otimes_k \bar{k}$ and $N \otimes_k \bar{k}$ are isomorphic if and only if $M \otimes_k \bar{k}$ and $N \otimes_k \bar{k}$ are isomorphic as $T \otimes_k \bar{k}$-modules. In other words, A and B lie in the same $\mathrm{GL}_n(k)$-orbit if and only if they lie in the same $\mathrm{GL}_n(\bar{k})$-orbit.*

Proof. Suppose that $M \otimes_k \bar{k}$ and $N \otimes_k \bar{k}$ are isomorphic $T \otimes_k \bar{k}$-modules. Then there exists an invertible matrix $C \in \mathrm{GL}_n(\bar{k})$ with entries in $\bar{k}$ such that $CA_i = B_iC$ for all i. There exists a finite field extension L of k such that all entries of C lie in L. It follows from Kraft and Riedtmann (1986, §5, Lemma 1) that M and N are isomorphic T-modules if k is a finite field. Suppose that k is infinite. Choose a basis $h_1, h_2, \dots, h_r$ of L as a k-vector space. We can write

$$C = \sum_j h_j C_j$$

where C_j is an R-module homomorphism from M to N for all j. Let $C(s_1, s_2, \dots, s_r) = \sum_{j=1}^r s_j C_j$ where $s_1, \dots, s_r$ are indeterminates. Since $C(h_1, \dots, h_r)$ is invertible, we have $\det C(h_1, \dots, h_r) \neq 0$. So $\det C(s_1, \dots, s_r)$ is not the zero polynomial. Since k is infinite, we can choose $a_1, \dots, a_r \in k$ such that $\det C(a_1, \dots, a_r) \neq 0$. Then $C(a_1, \dots, a_r)$ is an isomorphism between M and N. □

Theorem 4.3. *(See Chistov et al., 1997; Brooksbank and Luks, 2008.) There exists an algorithm for determining whether two n-dimensional modules M and N are isomorphic which requires only a polynomial number (polynomial in n and r) of arithmetic operations in the field k.*

In the theorem, the field k is fixed. Even if k is not fixed, say $k = \mathbb{F}_q$ and $\log q$ grows polynomially, then the algorithm still runs in polynomial time.

4.2. The isomorphism problem in k-categories

A category $\mathcal{C}$ is a k-category, if $\mathrm{Hom}_{\mathcal{C}}(M, N)$ is a vector space for every two objects M and N, and the composition map

$$\mathrm{Hom}_{\mathcal{C}}(M, N) \times \mathrm{Hom}_{\mathcal{C}}(N, P) \to \mathrm{Hom}_{\mathcal{C}}(M, P)$$

is k-bilinear for all objects M, N, P. Assume we have any k-category $\mathcal{C}$ with finite-dimensional Hom-spaces. Suppose that M and N are two isomorphic objects in $\mathcal{C}$, and let $T = \mathrm{Hom}_{\mathcal{C}}(N, N)$. Then T and $\mathrm{Hom}_{\mathcal{C}}(M, N)$ are isomorphic as left T-modules.

Lemma 4.4. *Suppose that M and N are isomorphic, and $\psi : T \to \operatorname{Hom}_{\mathcal{C}}(M,N)$ is an isomorphism of T-modules. Then $\varphi = \psi(1)$ is an isomorphism between M and N.*

Proof. Since 1 generates T as a T-module, $\varphi = \psi(1)$ generates $\operatorname{Hom}_{\mathcal{C}}(M,N)$ as a T-module. Suppose that $\gamma : M \to N$ is an isomorphism. Since $\gamma \in T\varphi$, there exists $\tau \in \operatorname{Hom}_{\mathcal{C}}(N,N) = T$ such that $\gamma = \tau\varphi$. So φ has a left inverse. The map

$$\Phi : \operatorname{Hom}_{\mathcal{C}}(N,M) \to \operatorname{Hom}_{\mathcal{C}}(N,N)$$

defined by $\Phi(\lambda) = \varphi\lambda$ is injective because φ has a left inverse. Since $\dim \operatorname{Hom}_{\mathcal{C}}(N,M) = \dim \operatorname{Hom}_{\mathcal{C}}(N,N) < \infty$ we have that Φ is surjective. Therefore id_N lies in the image of Φ. This implies that φ has a right inverse as well. □

To test whether any two objects M, N are isomorphic, we can proceed as follows.

(1) First test whether T and $\operatorname{Hom}_{\mathcal{C}}(M,N)$ are isomorphic as T-modules. If they are not isomorphic, then M and N are not isomorphic. Otherwise let $\psi : T \to \operatorname{Hom}_{\mathcal{C}}(M,N)$ be an isomorphism of R-modules.

(2) Let $\varphi = \psi(1)$. Test whether φ is an isomorphism. This is easy, because testing whether φ has a left and a right inverse just boils down to a system of linear equations. Now M and N are isomorphic if and only if φ is an isomorphism.

Proof of Theorem 1.6. We can use this approach for the categories $\mathcal{C}_d$. Note that

$$\dim \operatorname{Hom}_d(X_1, X_2) \leqslant \dim R_d$$

for all d because $\operatorname{Hom}_d(X_1, X_2)$ is a subspace of $R_d^\star$. For fixed d, $\dim R_d$ is bounded above by a polynomial in n, the number of vertices. Following the discussion in Section 2.4 we can compute $\operatorname{Hom}_d(X_1, X_2)$ in polynomial time. We have reduced the isomorphism problem in $\mathcal{C}_d$ to the isomorphism problem of modules, and by Theorem 4.3 the isomorphism problem of modules can be solved in a polynomial number of arithmetic operations in the field k. □

Let $\bar{k}$ be the algebraic closure of k. We construct a new category $\mathcal{C} \otimes_k \bar{k}$, where the objects are the same as the objects of $\mathcal{C}$, but

$$\operatorname{Hom}_{\mathcal{C}\otimes_k \bar{k}}(M,N) = \operatorname{Hom}_{\mathcal{C}}(M,N) \otimes_k \bar{k}.$$

Proposition 4.5. *Suppose that M, N are objects in $\mathcal{C}$. If $M \otimes_k \bar{k}$, $N \otimes_k \bar{k}$ are isomorphic in $\mathcal{C} \otimes_k \bar{k}$, then they are isomorphic in $\mathcal{C}$.*

Proof. Suppose M and N are objects in $\mathcal{C}$ which are isomorphic in $\mathcal{C} \otimes_k \bar{k}$. Let $T = \operatorname{Hom}_{\mathcal{C}}(N,N)$. Then $T \otimes_k \bar{k}$ is isomorphic to $\operatorname{Hom}_{\mathcal{C}}(M,N) \otimes_k \bar{k}$ as a $T \otimes_k \bar{k}$-module. From Proposition 4.2 it follows that T and $\operatorname{Hom}_{\mathcal{C}}(M,N)$ are isomorphic as T-modules. Let $\psi : T \to \operatorname{Hom}_{\mathcal{C}}(M,N)$ be an isomorphism and define $\varphi = \psi(1)$. Then ψ extends to an isomorphism $\psi \otimes \mathrm{id} : T \otimes_k \bar{k} : T \otimes_k \bar{k} \to \operatorname{Hom}_{\mathcal{C}}(M,N) \otimes_k \bar{k}$ of $T \otimes_k \bar{k}$-modules and $\varphi \otimes 1 = \psi(1)$ is an isomorphism by Lemma 4.4. We can write

$$(\varphi \otimes 1)^{-1} = \sum_{i=1}^{l} \gamma_i \otimes a_i$$

where $a_1, a_2, \ldots, a_l \in \bar{k}$ are linearly independent over k and $a_1 = 1$. Then we have

$$\mathrm{id} = (\varphi \otimes 1) \circ (\varphi \otimes 1)^{-1} = \sum_{i=1}^{l} (\varphi\gamma_i) \otimes a_i.$$

It follows that $\varphi\gamma_i = \mathrm{id}$ for $i = 1$ and $\varphi\gamma_i = 0$ for $i > 1$. Therefore, φ has a right inverse. Similarly φ has a left inverse, so φ is an isomorphism. □

Proof of Theorem 1.9. The implication (i) ⇒ (ii) follows from Corollary 3.18. It is easy to verify that the category $\mathcal{C}_d(V \otimes_k \bar{k})$ (working over the field $\bar{k}$) is equal to $\mathcal{C}_d(V) \otimes_k \bar{k}$. Suppose that $X_1, X_2 \in \mathrm{Aff}(V)$ are in the same $G(\bar{k})$-orbit, say $g \cdot X_1 = X_2$ for some $g \in G(\bar{k})$. We may view g as an element of $R_d^* \otimes_k \bar{k}$ if we identify g with the function $R_d \otimes_k \bar{k} \to \bar{k}$ which is evaluation at g. Then $g \in \mathrm{Hom}_d(X_1, X_2) \otimes_k \bar{k}$, and $g^{-1} \in \mathrm{Hom}_d(X_2, X_1) \otimes_k \bar{k}$ is its inverse. This shows that X_1, X_2 are isomorphic in $\mathcal{C}_d(V) \otimes_k \bar{k}$. By Proposition 4.5, we have that X_1, X_2 are isomorphic in $\mathcal{C}_d(V)$. The implication (ii) ⇒ (iii) follows. □

Proof of Theorem 1.10. The implication (i) ⇒ (ii) follows from Lemma 3.21. The other implications follow from Theorem 1.9. □

Proof of Theorem 1.5. The equivalence (i) ⇔ (ii) is Theorem 1.4. The implication (ii) ⇒ (iii) follows from Corollary 3.8 and Theorem 1.10. The implication (iii) ⇒ (iv) follows from Theorem 1.10. The equivalence (iv) ⇔ (v) is clear. □

4.3. *The categories $\mathcal{C}_d(V)$ for the general linear group*

Let G be the group $\mathrm{GL}_n(k)$. Let $U = k^n$ be the standard n-dimensional representation. We can identify G with the variety

$$\bigl\{(C, D) \in \mathrm{Hom}\bigl(k^n, U\bigr) \times \mathrm{Hom}\bigl(U, k^n\bigr) \bigm| DC = I_n\bigr\} \subseteq U^n \times \bigl(U^\star\bigr)^n.$$

Let W be the subspace of $k[G]$ spanned by the constant functions, and the functions induced by linear functions on $U^n \times (U^\star)^n$. So W is isomorphic to $U^n \oplus (U^\star)^n \oplus k$ as a representation of G. We have $\ell(U) = \ell(U^\star) = 1$. This choice of W gives us now a filtration of $R = k[G]$. For an n-dimensional vector space V every weakly decreasing sequence $\lambda = (\lambda_1, \dots, \lambda_n) \in \mathbb{Z}^n$ corresponds to an irreducible representation $S^\lambda(V)$ of $\mathrm{GL}(V)$. If $\lambda_r > 0$ and $\lambda_{r+1} \leqslant 0$ for some r, then we have that $S^\lambda(V)$ is a subrepresentation of

$$U^{\otimes(\lambda_1+\cdots+\lambda_r)} \otimes \bigl(U^\star\bigr)^{\otimes(-\lambda_{r+1}-\cdots-\lambda_n)}.$$

It follows that $\ell(S^\lambda(V)) \leqslant \sum_{i=1}^n |\lambda_i|$, where $|\cdot|$ denotes the absolute value.

Define

$$V = \mathrm{Mat}_{n,n}(k)^r = \mathrm{End}(U)^r$$

where G acts on V by simultaneous conjugation. We have $\ell(V) = 2$.

The remainder of the section is dedicated to the proof of Proposition 1.11. Let $T = k\langle x_1, \dots, x_r\rangle$ be the free associative algebra with r generators, and M and N be an n-dimensional T-modules represented by $A = (A_1, \dots, A_r) \in V$ and $B = (B_1, \dots, B_r) \in V$ respectively.

Let T-mod be the category of finite-dimensional left T-modules.

Proposition 4.6. *There exists a functor $\mathcal{F} : \mathcal{C}_3(V) \to T$-mod such that for every n-dimensional T-module M that is represented by $A = (A_1, \dots, A_r)$ we have $\mathcal{F}(A) \cong M$.*

Proof. Let $(C, D) \in G$. We can write $C = (c_{i,j})$ and $D = (d_{i,j})$. Then $I_3(A, B)$ is the 3-truncated ideal generated by the entries of the matrices $CD - I$, $DC - I$ and $CA_iD - B_i$ for $i = 1, 2, \dots, r$. Then the entries of $CA_i - B_iC = (CA_iD - B_i)C - CA_i(DC - I)$ also lie in $I_3(A, B)$. The coordinate functions $C = (c_{i,j})$ define a linear map $\pi : R^\star \to \mathrm{Hom}_k(k^n, k^n)$. It follows that $\pi(\mathrm{Hom}_d(M, N)) \subseteq \mathrm{Hom}_R(M, N)$. So we define $\mathcal{F}(\phi) = \pi(\phi)$ for all $\phi \in \mathrm{Hom}_d(M, N)$. □

Corollary 4.7. *The elements $A = (A_1, \dots, A_r)$, $B = (B_1, \dots, B_r)$ lie in the same orbit if and only if A and B are isomorphic to $\mathcal{C}_3(V)$.*

Consider now the case were $r = 1$. As the following proposition shows, the size needed for a covariant to distinguish two orbits may be excessively large:

Proposition 4.8. *Let*

$$C = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \quad D = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} \in \mathrm{Mat}_{2,2}(\mathbb{C})$$

and define the block matrices

$$A = \begin{pmatrix} C & & & & \\ & C & & & \\ & & \ddots & & \\ & & & C & \\ & & & & C \end{pmatrix} \quad \textit{and} \quad B = \begin{pmatrix} C & & & & \\ & C & & & \\ & & \ddots & & \\ & & & C & \\ & & & & D \end{pmatrix}$$

in $\mathrm{Mat}_{2n,2n}(\mathbb{C})$*. The group* $\mathrm{GL}_{2n}(\mathbb{C})$ *acts on* $\mathrm{Mat}_{2n,2n}(\mathbb{C})$ *by conjugation. Then A, B do not lie in the same* $\mathrm{GL}_{2n}(\mathbb{C})$*-orbit. If $\varphi : V \to V'$ is a covariant which distinguishes the orbits of A and B respectively, then we have* $\dim(V') \geqslant 3^n$ *and* $\ell(V') \geqslant 2n$*.*

Proof. Define $V = \mathrm{End}(U)$. Then $V \cong \mathrm{Mat}_{2n,2n}(\mathbb{C})$, and $\mathrm{GL}(U) \cong \mathrm{GL}_{2n}(\mathbb{C})$. We can write $U = U_1 \oplus \cdots \oplus U_n$ where $U_i \cong \mathbb{C}^2$. We can view $\mathrm{End}(U_1) \oplus \cdots \oplus \mathrm{End}(U_n)$ as a subalgebra of $\mathrm{End}(U)$. Now $A, B \in \mathrm{End}(U_1) \oplus \cdots \oplus \mathrm{End}(U_n) \subseteq \mathrm{End}(U)$ are given by $A = (C, C, \dots, C)$ and $B = (C, C, \dots, C, D)$.

Suppose that $\varphi : V \to V'$ is a covariant, where V' is an irreducible representation of $\mathrm{GL}(V)$. If $\varphi(A) = 0$, then $\varphi(B) = 0$ because B lies in the orbit closure of A. Suppose that $\varphi(A) \neq 0$ and $\varphi(B) = 0$. As a representation of $\mathrm{GL}(U_1) \times \cdots \times \mathrm{GL}(U_n)$, V' may not be irreducible. Let $Z_1 \otimes \cdots \otimes Z_n$ be an irreducible summand of V' as a $\mathrm{GL}(U_1) \times \cdots \times \mathrm{GL}(U_n)$ representation, such that $p(\varphi(A)) \neq 0$, where p is the $\mathrm{GL}(U_1) \times \cdots \times \mathrm{GL}(U_n)$-equivariant projection $V' \to Z_1 \otimes \cdots \otimes Z_n$, and Z_i is an irreducible representation of $\mathrm{GL}(U_i)$ for all i. Let $q : \mathrm{End}(U_1) \oplus \cdots \oplus \mathrm{End}(U_n) \to Z_1 \otimes \cdots \otimes Z_n$ be the restriction of $p \circ \varphi$. We have $q(A) \neq 0$. Suppose that $\dim Z_i = 1$ for some i. Let $B' = (C, \dots, C, D, C, \dots, C)$. Then B and B' are in the same $\mathrm{GL}(V)$-orbit, so $\varphi(B') = 0$, and hence $q(B') = 0$. Now B' lies in the $\mathrm{SL}(U_i)$-closure of A. Since q is $\mathrm{SL}(U_i)$-invariant, we get $q(A) = q(B') = 0$. Contradiction. Hence $\dim Z_i \geqslant 2$. Since Z_i must be an irreducible representation of $\mathrm{PSL}(U_i)$, we even have $\dim Z_i \geqslant 3$. It follows that

$$\dim V' \geqslant (\dim Z_1) \cdots (\dim Z_n) \geqslant 3^n.$$

Let us write $V' = S^{\lambda}(U)$ for some $\lambda = (\lambda_1, \dots, \lambda_{2n}) \in \mathbb{Z}^{2n}$ and $Z_i = S^{\mu^{(i)}}(U_i)$, where $\mu^{(i)} = (\mu_1^{(i)}, \mu_2^{(i)})$.

From $\dim(Z_i) \geqslant 3$ it follows that $\mu_1^{(i)} - \mu_2^{(i)} \geqslant 2$ for all i. The representation of $\mathrm{GL}(U)$ restricts to $\mathrm{GL}(U_1) \times \cdots \times \mathrm{GL}(U_n)$ according to the Littlewood–Richardson rule. We have the following inequalities:

$$\lambda_1 \geqslant \sum_i \mu_1^{(i)}$$

and

$$\lambda_{2n} \leqslant \sum_i \mu_2^{(i)}.$$

Taking the difference gives us

$$|\lambda| = \sum_i |\lambda_i| \geqslant \lambda_1 - \lambda_{2n} \geqslant \sum_{i=1}^{n} \mu_1^{(i)} - \mu_2^{(i)} \geqslant 2n.$$

It follows that $\ell(V') \geqslant 2n$. □

Remark 4.9. Define $\varphi : \mathrm{End}(U) \to \mathrm{End}(\bigwedge^n U)$ by

$$\varphi(E) = E \wedge \cdots \wedge E.$$

Then $\varphi(A) \neq 0$ and $\varphi(B) = 0$. Note that $\mathrm{End}(\bigwedge^n U)$ is not irreducible. There exists an irreducible summand W of $\mathrm{End}(\bigwedge^n U)$ such that $p(\varphi(C)) \neq 0$, where $p : \mathrm{End}(\bigwedge^n U) \to W$ is the projection. If we set $q = p \circ \varphi$, then q is a covariant that distinguishes the orbits of A and B. Note that $\dim W \leqslant \dim \mathrm{End}(\bigwedge^n U) \leqslant 4^n$.

5. The Cai–Fürer–Immerman examples

Cai, Fürer and Immerman showed that for every positive integer d there exist non-isomorphic 2-colored graphs Γ and Γ' such that $\Gamma \sim_d \Gamma'$. To explain this result, we need to describe the construction of Cai, Fürer and Immerman which, given a graph Q, produces two non-isomorphic 2-colored graphs $\Gamma(Q)$ and $\Gamma'(Q)$ (see Cai et al., 1992, §6).

Suppose that $Q = \langle Y, R\rangle$ is a graph. Let $E = \{\{x, y\} \mid (x, y) \in R\}$ be the set of edges in the graph. For every vertex $x \in Y$, we define $E(x) = \{e \in E \mid x \in e\}$, the set of edges which are incident with x. We define a vertex set $Y(Q) = Y_1(Q) \cup Y_2(Q)$, where

$$Y_1(Q) = \bigl\{c_{x,S} \bigm| x \in S,\ S \subseteq E(x),\ |S| \text{ is even}\bigr\},$$

and

$$Y_2(Q) = \bigl\{a_{x,e} \bigm| x \in Y,\ e \in E(x)\bigr\} \cup \bigl\{b_{x,e} \bigm| x \in Y,\ e \in E(x)\bigr\}.$$

We define the edge set $E(Q)$ by

$$\begin{aligned} E(Q) = {} & \bigl\{\{a_{x,e}, c_{x,S}\} \bigm| x \in Y,\ e \in S\bigr\} \cup \bigl\{\{b_{x,e}, c_{x,Y}\} \bigm| x \in Y,\ e \notin S\bigr\} \\ & \cup \bigl\{\{a_{x,e}, a_{a,e}\} \bigm| x, y \in Y,\ e \in E(x) \cap E(y)\bigr\} \cup \bigl\{\{b_{x,e}, b_{a,e}\} \bigm| x, y \in Y,\ e \in E(x) \cap E(y)\bigr\}. \end{aligned}$$

We also define another edge set $E'(Q)$ as follows: We choose two special vertices $\widetilde{x}$, $\widetilde{y}$ such that $\widetilde{e} = \{\widetilde{x}, \widetilde{y}\} \in E$ is an edge. To obtain $E'(Q)$, remove $\{a_{\tilde{x},\tilde{e}}, a_{\tilde{y},\tilde{e}}\}$ and $\{b_{\tilde{x},\tilde{e}}, b_{\tilde{y},\tilde{e}}\}$ from $E(Q)$ and add $\{a_{\tilde{x},\tilde{e}}, b_{\tilde{y},\tilde{e}}\}$ and $\{a_{\tilde{y},\tilde{e}}, b_{\tilde{x},\tilde{e}}\}$. Let $R(Q)$ and $R'(Q)$ be the symmetric relations corresponding to the edge sets $E(Q)$ and $E(Q')$ respectively. We now have two 2-colored graphs: $\Gamma(Q) = (Y(Q), R(Q), Y_1(Q), Y_2(Q))$ and $\Gamma'(Q) = (Y(Q), R'(Q), Y_1(Q), Y_2(Q))$.

The following proposition follows from Lemma 6.2 of Cai et al. (1992). We will give a proof here, because a crucial lemma is based on this proof.

Proposition 5.1. *The graphs $\Gamma(Q)$ and $\Gamma'(Q)$ are not isomorphic.*

Proof. Let M be the adjacency matrix of $\Gamma(Q)$ with entries in the field $\mathbb{F}_2$. Since $Y(Q) = Y_1(Q) \cup Y_2(Q)$, M has the following block form:

$$M = \begin{pmatrix} A_{1,1} & A_{1,2} \\ A_{2,1} & A_{2,2} \end{pmatrix}$$

where $A_{1,1}$, $A_{2,2}$ are symmetric and $A_{1,2} = A_{2,1}^t$. Similarly, let

$$M' = \begin{pmatrix} A_{1,1} & A_{1,2} \\ A_{2,1} & A'_{2,2} \end{pmatrix}$$

be the adjacency matrix for $\Gamma'(Q)$.

Let

$$B = (\, A_{2,1} \quad A_{2,2} \,), \quad \text{and} \quad B' = \bigl(A_{2,1} \quad A'_{2,2}\bigr).$$

The proposition now follows from the lemma below. □

Lemma 5.2. *We have*

$$\operatorname{rank}(B) = 3|E| + |Y| - 2$$

and

$$\operatorname{rank}\bigl(B'\bigr) = 3|E| + |Y| - 1.$$

Proof. The image $\operatorname{im}(B)$ of B is equal to $\operatorname{im}(A_{2,1}) + \operatorname{im}(A_{2,2})$. The space $\operatorname{im}(A_{2,1})$ is spanned by all

$$\sum_{e \in S} a_{x,e} + \sum_{e \in E(x) \setminus S} b_{x,e}$$

with $x \in Y$, $S \subseteq E(x)$ with $|S|$ even, and $\operatorname{im}(A_{2,2})$ is spanned by all

$$a_{x,e} + a_{y,e}, \; b_{x,e} + b_{y,e}$$

with $x \in Y$ and $e \in E(x)$. It is clear that $\dim \operatorname{im}(A_{2,2}) = 2|E|$. For $e = \{x, y\} \in E$, define $a_e = a_{x,e} + \operatorname{im}(A_{2,2}) = a_{y,e} + \operatorname{im}(A_{2,2})$ and $b_e = b_{x,e} + \operatorname{im}(A_{2,2}) = b_{y,e} + \operatorname{im}(A_{2,2}) \in k^{4|E|}/\operatorname{im}(A_{2,2})$. Now $\operatorname{im}(B)/\operatorname{im}(A_{2,2})$ is spanned by all

$$\sum_{e \in S} a_e + \sum_{e \in E(x) \setminus S} b_e$$

where $x \in Y$ and $S \subseteq Y$ with $|S|$ even. Note that $a_e + b_e + a_f + b_f \in \operatorname{im}(B)/\operatorname{im}(A_{2,2})$ for all $x \in Y$, $e, f \in E(x)$. Since Q is connected, it follows that $a_e + b_e + a_f + b_f \in \operatorname{im}(B)/\operatorname{im}(A_{2,2})$ for all $e, f \in E$. Let $Z \subseteq \operatorname{im}(B)$ containing $\operatorname{im}(A_{2,2})$ such that $Z/\operatorname{im}(A_{2,2})$ is spanned by all $a_e + b_e + a_f + b_f$. The dimension of $Z/\operatorname{im}(A_{2,2})$ is $|E| - 1$. Now $\operatorname{im}(B)/Z$ is spanned by all elements of the form

$$\sum_{e \in E(x)} b_e + Z$$

with $x \in Y$. Since Q is connected, it follows that $\dim \operatorname{im}(B)/Z = |Y| - 1$. We conclude that

$$\operatorname{rank}(B) = \dim \operatorname{im}(B) = 2|E| + \bigl(|E| - 1\bigr) + \bigl(|Y| - 1\bigr) = 3|E| + |Y| - 2.$$

We can do a similar computation for $\operatorname{rank}(B')$. First of all $\dim \operatorname{im}(A'_{2,2}) = 2|E|$. Let $\widetilde{e} = \{\widetilde{x}, \widetilde{y}\}$ be the special edge. For $e = \{x, y\} \neq \widetilde{e}$, we define $a'_e = a_{x,e} + \operatorname{im}(A'_{2,2})$ and $b'_e = b_{x,e} + \operatorname{im}(A'_{2,2})$. For $\widetilde{e} = \{\widetilde{x}, \widetilde{y}\}$ we define $a'_{\widetilde{e}} = a_{\tilde{x},\tilde{e}} + \operatorname{im}(A'_{2,2}) = b_{\tilde{y},\tilde{e}} + \operatorname{im}(A'_{2,2})$ and $b'_{\widetilde{e}} = b_{\tilde{x},\tilde{e}} + \operatorname{im}(A'_{2,2}) = a_{\tilde{y},\tilde{e}} \in \operatorname{im}(A'_{2,2})$. Let $Z' \subseteq B$ be the space containing $\operatorname{im}(A'_{2,2})$ such that $Z'/\operatorname{im}(A'_{2,2})$ is spanned by all $a'_e + b'_e + a'_f + b'_f$ with $e, f \in E$. We have $\dim(Z'/\operatorname{im}(A'_{2,2})) = |E| - 1$. Finally, $\operatorname{im}(B')/\dim(Z')$ is spanned by all

$$\sum_{e \in E(x)} b'_e + Z'$$

with $x \in Y$ and $x \neq \widetilde{y}$, and

$$\Biggl(\sum_{e \in E(\widetilde{x}) \setminus \{\widetilde{e}\}} a'_e\Biggr) + b_{\widetilde{e}'}.$$

It is easy to see that $\dim(\operatorname{im}(B')/\dim(Z')) = |Y|$. So we obtain

$$\dim\bigl(\operatorname{im}\bigl(B'\bigr)\bigr) = 2|E| + \bigl(|E| - 1\bigr) + |Y| = 3|E| + |Y| - 1. \qquad \square$$

Definition 5.3. A separator of a graph $Q = (Y, R)$ is a subset $S \subset Y$ such that the induced subgraph on $Y \setminus S$ has no connected component with more than $|Y|/2$ vertices.

The following theorem is Theorem 6.4 in Cai et al. (1992).

Theorem 5.4. *Suppose that Q is a graph such that every separator of Q has at least $d+1$ vertices. Then $\Gamma(Q)$ and $\Gamma'(Q)$ cannot be distinguished by the d-variable logic with counting.*

There exists a family of graphs T_d with the following properties: T_d has $O(d)$ vertices, every vertex in T_d has degree 3, and every separator has at least $d+1$ vertices. Then $\Gamma(T_d)$ and $\Gamma'(T_d)$ have $O(d)$ vertices, and $\Gamma(T_d) \sim_d \Gamma(T'_d)$. Every vertex of $\Gamma(T_d)$ or $\Gamma'(T_d)$ has degree 3. This shows that for fixed d, the d-dimensional Weisfeiler–Lehman algorithm cannot distinguish all graphs of degree 3. However, it is possible to distinguish graphs of bounded degree in polynomial time. Such an algorithm was given in Luks (1982).

Proof of Theorem 1.7. Suppose that $k = \mathbb{F}_2$ and $Q = \langle Y, E\rangle$. We will show that $A_{\Gamma(Q)}$ and $A_{\Gamma'(Q)}$ can be separated by a 3-constructible functor. We have $A_\Gamma, A_{\Gamma_1} \in V = U \otimes U \oplus U \oplus U \oplus k$. Let $p_1, p_2 : V \to U$ be the two projections onto U, and $q : V \to U \otimes U \cong \operatorname{End}(U)$ be the projection onto $U \otimes U$. Then we have $p_2(A_{\Gamma(Q)}) = p_2(A_{\Gamma'(Q)}) = \sum_{x \in X_2(Q)} x$. Let $\delta : U \to U \otimes U$ be defined by $\delta(x) = x \otimes x$ for all $x \in X(Q)$. Then $\delta(p_2(A_{\Gamma(Q)})) = \delta(p_2(A_{\Gamma'(Q)})) \in U \otimes U \cong U \otimes U^\star \cong \operatorname{End}(U)$ is the projection of onto the span of $X_2(Q)$. The compositions $q(A_{\Gamma(Q)}) \circ \delta(p_2(A_{\Gamma(Q)}))$ and $q(A_{\Gamma'(Q)}) \circ \delta(p_2(A_{\Gamma'(Q)}))$ are given by the matrices B and B' in the proof of Proposition 5.1. Define the following 3-constructible functors: The functor

$$\mathcal{F}_1 : \mathcal{C}_3(V) \to \mathcal{C}_3\bigl(\operatorname{End}(U)\bigr)$$

is defined by the 3-constructible equivariant linear map $\delta \circ p_2$. The functor

$$\mathcal{F}_2 : \mathcal{C}_3\bigl(\operatorname{End}(U)\bigr) \to \mathcal{C}_3\bigl(\operatorname{End}(U) \otimes U\bigr)$$

is defined by

$$\mathcal{F}_2(Z) = Z \otimes U.$$

The functor

$$\mathcal{F}_3 : \mathcal{C}_3\bigl(\operatorname{End}(U) \otimes U\bigr) \to \mathcal{C}_3(U)$$

is defined by the equivariant $f \otimes v \mapsto f(v)$. Let $\mathcal{F}_4 : \mathcal{C}_3(V) \to \mathcal{C}_3(\operatorname{End}(U))$ be defined by the equivariant linear map q. Then $\mathcal{F}_3 \circ \mathcal{F}_2 \circ \mathcal{F}_1$ and $\mathcal{F}_4$ are 3-constructible, and

$$\mathcal{F}_4 \otimes (\mathcal{F}_3 \circ \mathcal{F}_2 \circ \mathcal{F}_1) : \mathcal{C}_3(V) \to \mathcal{C}_3\bigl(\operatorname{End}(U) \otimes U\bigr)$$

is constructible. Define a 3-constructible functor $\mathcal{G} : \mathcal{C}_3(V) \to \mathcal{C}_3(U)$ by

$$\mathcal{G} = \mathcal{F}_3 \circ \bigl(\mathcal{F}_4 \otimes (\mathcal{F}_3 \circ \mathcal{F}_2 \circ \mathcal{F}_1)\bigr).$$

Then we have $\mathcal{G}(A_{\Gamma(Q)}) = \operatorname{im} B$ and $\mathcal{G}(A_{\Gamma'(Q)}) = \operatorname{im} B'$. By Lemma 5.2, we have $\dim \mathcal{G}(A_{\Gamma(Q)}) \neq \dim \mathcal{G}(A_{\Gamma'(Q)})$, so $\mathcal{G}$ distinguishes $A_{\Gamma(Q)}$ and $A_{\Gamma'(Q)}$. □

6. Open problems

We finish with some open questions:

Problem 6.1. Does $\mathbf{AC}_d$ distinguish all pairs of non-isomorphic graphs for some d?

A positive answer to this problem implies that the Graph Isomorphism Problem lies in the complexity class **P**.

Suppose that Γ_1, Γ_2 are (colored) graphs constructed using the Cai–Fürer–Immerman method. We know that A_{Γ_1} and A_{Γ_2} are non-isomorphic in $\mathcal{C}_3(V)$, assuming we are working over the field $\mathbb{F}_2$ (see

Theorem 1.7). The proof heavily relies on the fact that we are working over the field $\mathbb{F}_2$. So a natural question to ask is:

Problem 6.2. Are A_{Γ_1} and A_{Γ_2} non-isomorphic in $\mathcal{C}_3(V)$, even if we are working over a field of characteristic other than 2?

If we work over a base field $k = \mathbb{Q}$, then the size of the rational numbers may grow exponentially if we do arithmetic operations such as multiplications and additions. So it is a priori not clear whether algorithms for testing isomorphism in $\mathcal{C}_d(V)$ run in polynomial time.

Problem 6.3. If we work over the base field $k = \mathbb{Q}$, can we test for isomorphism in $\mathcal{C}_d(V)$ in polynomial time?

One may expect that there is a probabilistic algorithm for testing isomorphism in $\mathcal{C}_d(V)$ by working over $\mathbb{F}_p$ for various random primes p for which $\log(p)$ is polynomial in the number of vertices.

References

Babai, L., Grigoryev, D.Yu., Mount, D.M., 1982. Isomorphism of graphs with bounded eigenvalue multiplicity. In: Proceedings of the 14th Annual ACM Symposium on Theory of Computing, pp. 310–324.

Brooksbank, P.A., Luks, E.M., 2008. Testing isomorphism of modules. J. Algebra 320 (11), 4020–4029.

Busacker, R.G., Saaty, T.L., 1965. Finite Graphs and Networks. McGraw–Hill, New York.

Cai, J.-Y., Fürer, M., Immerman, N., 1992. An optimal lower bound for the number of variables for graph identification. Combinatorica 12.

Chistov, A., Ivanyos, G., Karpinski, M., 1997. Polynomial-time algorithms for modules over finite dimensional algebras. In: Proceedings of the 1997 International Symposium on Symbolic and Algebraic Computation (ISSAC), pp. 68–74.

Evdokimov, S., Karpinski, M., Ponomarenko, I., 1999. On a new higher dimensional Weisfeiler–Lehman algorithm. J. Algebr. Comb. 10.

Evdokimov, S., Ponomarenko, I., 1999. On highly closed cellular algebras and highly closed isomorphisms. Electron. J. Comb. 6, #18.

Filotti, I.S., Mayer, J.N., 1980. A polynomial-time algorithm for determining the isomorphism of graphs of fixed genus. In: Proceedings of the 12th Annual ACM Symposium on Theory of Computing, pp. 236–243.

Friedland, S., 1989. Coherent algebras and the graph isomorphism problem. Discr. Appl. Math. 25, 73–98.

Graham, J.J., Lehrer, G.I., 1996. Cellular algebras. Invent. Math. 123 (1), 1–34.

Higman, D.G., 1987. Coherent algebras. Linear Algebra Appl. 93, 209–240.

Hopcraft, J.E., Tarjan, R.E., 1973. A $V \log V$ algorithm for isomorphism of triconnected planar graphs. J. Comput. Syst. Sci. 7, 323–331.

Hopcroft, J.E., Wong, J.K., 1974. Linear time algorithm for isomorphism of planar graphs. In: Proceedings of the Sixth Annual ACM Symposium on Theory of Computing. Assoc. Comput. Mach., pp. 172–184.

Köbler, J., Schöning, U., Torán, J., 1993. The Graph Isomorphism Problem: Its Structural Complexity. Progr. Theoret. Comput. Sci.. Birkhäuser, Boston.

Kraft, H., Riedtmann, Ch., 1986. Geometry of representations of quivers. In: Webb, P. (Ed.), Representations of Algebras: Proceedings of the Durham Symposium, 1985. In: Lond. Math. Soc. Lect. Note Ser., vol. 116. Cambridge University Press.

Luks, E.M., 1982. Isomorphism of graphs of bounded valence can be tested in polynomial time. J. Comput. Syst. Sci. 25 (1), 42–65.

Luks, E.M., 1986. Parallel algorithms for permutation groups and graph isomorphism. In: Proceedings of the 27th IEEE Symposium on Foundations of Computer Science, pp. 292–302.

Miller, G., 1980. Isomorphism testing for graphs of bounded genus. In: Proceedings of the 12th Annual ACM Symposium on Theory of Computing, pp. 225–235.

Mulmuley, K., Sohoni, M., 2001. Geometric Complexity Theory I: An approach to the P vs. NP and related problems. SIAM J. Comput. 31, 496–526.

Mulmuley, K., Sohoni, M., 2008. Geometric Complexity Theory II: Towards explicit obstructions for embeddings among class varieties. SIAM J. Comput. 38.

Ramanujan, S., 1919. A proof of Bertrand's postulate. J. Indian Math. Soc. 11, 181–182. Also in: Collected Papers of Srinivasa Ramanujan. AMS Chelsea Publishing, Providence, RI, 2000, pp. 208–209.

Valiant, L.G., 1979. The complexity of the permanent. Theoret. Comp. Sci. 8, 189–201.

Weisfeiler, B.Yu. (Ed.), 1976. On Construction and Identification of Graphs. Lect. Notes Math., vol. 558. Springer, Berlin–New York.

Weisfeiler, B.Yu., Lehman, A.A., 1968. A reduction of a graph to a canonical form and an algebra arising during this reduction. Naucho-Techn. Inf. Ser. 2 9, 12–16 (in Russian).

Weispfenning, V., 1988. Some bounds for the construction of Gröbner bases. In: Applicable Algebra, Error-Correcting Codes Combinatorics and Computer Algebra. Karlsruhe, 1986. In: Lect. Notes Comput. Sci., vol. 307. Springer, Berlin, pp. 195–201.