

Representation Theory of Combinatorial Categories

by

John D. Wiltshire-Gordon

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mathematics)
in The University of Michigan
2016

Doctoral Committee:

Associate Professor David E. Speyer, Chair
Professor William Fulton
Professor Jeffrey C. Lagarias
Associate Professor Yaoyun Shi
Assistant Professor Andrew Snowden

© John D. Wiltshire-Gordon 2016
All Rights Reserved

ACKNOWLEDGEMENTS

Thanks to my helpful, knowledgeable, and wise advisor, David Speyer. Many thanks as well to Thomas Church and Benson Farb for their inspirational paper [CF13] and their encouragement when I was an undergraduate at the University of Chicago and after. Thank you to my fellow graduate students at the University of Michigan for providing a stimulating mathematical environment. Thanks to the NSF for support through a graduate research fellowship, and the Rackham Graduate School for support through a predoctoral fellowship. It was a pleasure to discuss the ideas of this thesis in helpful conversations with Daniel Barter, E. Hunter Brooks, Thomas Church, Jordan Ellenberg, Pavel Etingof, Benson Farb, Nic Ford, Trevor Hyde, Gene Kopp, Visu Makam, Jeremy Miller, Andrew Putman, Victor Reiner, Julian Rosen, Steven Sam, Andrew Snowden, Benjamin Steinberg, Phil Tosteson, and Jennifer Wilson. Finally, my heartfelt thanks go to Hannah Fehlner-Peach, my parents, and my siblings for their continued encouragement and support.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	ii
LIST OF FIGURES	vi
LIST OF TABLES	vii
ABSTRACT	viii
 CHAPTER	
I. Introduction	1
1.0.1 This thesis in broad strokes	3
1.0.2 Outline	4
1.1 What is a representation?	5
1.1.1 Categories and their representations	6
1.1.2 Matrices and additive categories	7
1.1.3 Matrices as presentations	9
1.2 What are the basic examples?	10
1.2.1 A familiar example of a finite presentation over a category	10
1.2.2 Another beginning example	11
1.2.3 Configurations of distinct points in \mathbb{C}	13
1.2.4 Monomials with combinatorial indexing	14
1.3 How do we handle a presentation matrix classically...	16
1.3.1 When \mathcal{D} is a finite group?	16
1.3.2 When \mathcal{D} is a graded polynomial algebra?	18
1.4 Generalizing these classical techniques to representations of categories	20
1.4.1 The dimension of a category \mathcal{D}	20
1.4.2 Sam-Snowden Gröbner theory	21
1.4.3 Representation stability	21
II. Basic constructions	23

2.1	Categories, functors, and natural transformations	23
2.1.1	Basic definitions of category theory	23
2.1.2	R -linear categories	24
2.2	Representations and the category $\mathbf{Mod}_R^{\mathcal{D}}$	25
2.3	Matrices over a category with coefficients in a ring	27
2.3.1	The R -linearization of a category	28
2.4	Yoneda's lemma and basic projectives in the category $\mathbf{Mod}_R^{\mathcal{D}}$	28
2.4.1	Finitely generated representations of a category	30
2.5	Idempotents and projectives	31
2.6	Left and right Kan extensions	32
2.6.1	Construction of the functors $f_!$ and f_*	35
2.6.2	Two easy facts about the functors $f_!$, f^* , and f_*	37
III. Irreducible representations, injective hulls, and projective covers		39
3.1	Simple representations, finite length representations	39
3.1.1	The Krull-Schmidt decomposition	40
3.1.2	Length and the three functors $f_!$, f^* , and f_*	41
3.2	Projective covers and injective hulls	43
3.3	The intermediate extension functor $f_{!*}$	44
3.3.1	Filtered categories, associated graded categories, and the classification of irreducible representations	46
3.3.2	Computing the dimensions of the irreducible representation $i_{!*}p^*W$	48
IV. Imaginations, homological moduli, and categories of dimension zero		49
4.1	The object preorders \leq_d	49
4.1.1	Useful facts about the preorders \leq_d	51
4.2	Upper bounds for the preorder \leq_d	55
4.3	Dimension zero categories	59
4.3.1	Computing simple multiplicities from a presentation matrix	61
4.3.2	Bounded models	63
4.3.3	The Dold-Kan Correspondence	65
4.3.4	Further examples of categories of dimension zero	66
4.3.5	Schur projectives	67
4.4	\mathcal{D} -regular functions	69
V. Case study: Representations of the category of finite sets		73

5.1	Summary of results for the representation theory of the category of finite sets	73
5.2	First experiments	75
5.2.1	Computing in \mathcal{D}	75
5.2.2	Computing the preorders \leq_d	76
5.2.3	A first pass at the irreducible representations	79
5.3	More experiments suggested by the results of the first experiments	83
5.3.1	A Morita equivalent category separating columns and non-columns	83
5.3.2	Computing the preorders \leq_π for the Morita equivalent category \mathcal{C}	84
5.3.3	Explicit construction of the irreducibles	85
5.3.4	What we know after constructing the irreducibles	88
5.3.5	What's left?	89
5.4	Preparing to apply the Multiplicity Theorem 4.3.5	89
5.4.1	Experimentally computing the indecomposable injectives	89
5.4.2	Construction of the indecomposable injectives	92
5.5	The multiplicity theorem for presentations	94
5.5.1	Demonstration computation	95
5.5.2	A computer program	96
5.6	What can be said over \mathbb{Z}	97
	BIBLIOGRAPHY	98

LIST OF FIGURES

Figure

1.1	The Kronecker quiver	1
1.2	A map of representations and the two squares that must commute .	1
1.3	The poset category (\mathbb{N}, \leq)	2
1.4	The trivial representation of (\mathbb{N}, \leq)	2
1.5	Multiplying matrices over \mathcal{D} with coefficients in R	9
1.6	A valid 3-coloring of the five-cycle	12

LIST OF TABLES

Table

1.1	Examples of monomials for combinatorial categories	16
2.1	Useful properties of Kan extensions	34
5.1	Recording the inequality $[3] \leq_{[1]} [2]$	77
5.2	The preorder $\leq_{[1]}$	77
5.3	The preorder $\leq_{[2]}$	77
5.4	Dimensions of the irreducible representation with $\lambda = \square\square$	82
5.5	Dimensions of the irreducible representation with $\lambda = \begin{smallmatrix} \square \\ \square \end{smallmatrix}$	82
5.6	Dimensions of the irreducible representations with λ a partition of 3	82
5.7	Dimensions of the indecomposable injective corresponding to $\square\square$	91
5.8	Dimensions of the indecomposable injectives for partitions of 3	91

ABSTRACT

Representation Theory of Combinatorial Categories

by

John D. Wiltshire-Gordon

Chair: David Speyer

A representation V of a category \mathcal{D} is a functor $\mathcal{D} \rightarrow \mathbf{Mod}_R$; the representations of \mathcal{D} form an abelian category with natural transformations as morphisms. Say V is finitely generated if there exist finitely many vectors $v_i \in Vd_i$ so that any strict subrepresentation of V misses some v_i . If every finitely generated representation satisfies both ACC and DCC on subrepresentations, we say \mathcal{D} has dimension zero over R . The main theoretical result of this thesis is a practical recognition theorem for categories of dimension zero (Theorem 4.3.2). The main computational result is an algorithm for decomposing a finitely presented representation of a category of dimension zero into its multiset of irreducible composition factors (Theorem 4.3.5). Our main applications take \mathcal{D} to be the category of finite sets; we explain how the general results of this thesis suggest specific experiments that lead to structure theory and practical algorithms in this case.

CHAPTER I

Introduction

Just as the definition of a **group** axiomatizes the notion of **symmetry**, the definition of a **category** axiomatizes the notion of **transformation**. Accordingly, a **representation of a category** provides a rule by which abstract transformations may be converted to concrete linear transformations.

Definition 1.0.1. A representation V of a category \mathcal{D} is a functor $V: \mathcal{D} \rightarrow \text{Vect}_{\mathbb{Q}}$.

One might call the study of such representations the **representation theory of categories**, although this broad subject goes by many names.

Here is a very small category:



Figure 1.1: The Kronecker quiver

To form a representation, one places a vector space on each object (\bullet), and a linear map on each morphism (\rightarrow). A **map of representations** consists of a vertical map for each object so that these maps commute with corresponding horizontal maps.




Figure 1.2: A map of representations and the two squares that must commute

When the vertical maps are inclusions, V is called a **subrepresentation** of W . Any collection of vectors **generates** a subrepresentation, which is the smallest containing those vectors. A representation is said to be **finitely generated** if it admits a finite generating set. One imagines the vectors being pushed around by the arrows until they span every vector space.

The **length** $\ell(V)$ of a representation V is the number of steps in a maximal, strictly increasing chain of subrepresentations. If there is no maximal chain, $\ell(V) = \infty$. Say \mathcal{D} has **dimension zero** if every finitely generated representation has finite length.

Theorem 1.0.2. *Every category with finitely many morphisms has dimension zero.*

Proof. Suppose V has generators v_1, \dots, v_g and \mathcal{D} has morphisms f_1, \dots, f_m . Vectors of the form $v_i \cdot f_j$ span V , so $\ell(V) \leq g \cdot m < \infty$. \square

We call a category with finitely many morphisms a **finite category**, and admire the computability and explicitness of its representation theory.

As a first example of an infinite category, take the natural numbers as objects and say $\text{Hom}(n, m)$ has one or zero elements depending if $n \leq m$:




Figure 1.3: The poset category (\mathbb{N}, \leq)

A representation of this category consists of a sequence of vector spaces, one for each natural number, and a linear map from each to the next. In other words, representations of (\mathbb{N}, \leq) are \mathbb{N} -graded $\mathbb{Q}[T]$ -modules where multiplication by T pushes vectors along the path.

Unsurprisingly, (\mathbb{N}, \leq) does not have dimension zero.¹ The representation

$$\mathbb{Q} \xrightarrow{1} \mathbb{Q} \xrightarrow{1} \mathbb{Q} \xrightarrow{1} \mathbb{Q} \xrightarrow{1} \mathbb{Q} \xrightarrow{1} \mathbb{Q} \xrightarrow{1} \mathbb{Q} \xrightarrow{1} \mathbb{Q} \xrightarrow{1} \dots$$

Figure 1.4: The trivial representation of (\mathbb{N}, \leq)

¹It has dimension one, matching the Krull dimension of $\mathbb{Q}[T]$.

is generated by a single vector in the leftmost vector space, but has an evident infinite descending chain of subrepresentations.

But there is a more pressing issue. Whereas a finitely generated representation of a finite category is described by a finite collection of matrices, a naïve description of a representation of an infinite category requires an infinite amount of data. And yet, graded $\mathbb{Q}[T]$ -modules, for example, manage to be practical objects. The secret: one specifies a finitely generated $\mathbb{Q}[T]$ -module not by its graded pieces with their transition maps, but by a presentation matrix. Indeed, any finitely generated module is the cokernel of a finite matrix with entries in $\mathbb{Q}[T]$.

1.0.1 This thesis in broad strokes

- We explain how to write a presentation matrix for a representation of a category.
- We explain how to interpret presentation matrices, and find them in nature wearing thin disguises.
- We give a combinatorial condition on \mathcal{D} that detects if \mathcal{D} is dimension zero. There exist infinite categories of dimension zero! Our flagship example is the category of finite sets.
- We give an algorithm by which a finitely presented representation of a dimension zero category may be decomposed into its multiset of irreducible composition factors.

In sum, we strive to make infinite diagrams of vector spaces as explicit and computational as finite diagrams.

1.0.2 Outline

Chapter I explains how to build a matrix over \mathcal{D} with scalars in R , and how to use it to get a presentation for a representation of \mathcal{D} . We give numerous examples of such presentations, and indicate the general circumstances in which they appear. Several examples take \mathcal{D} to be the category of finite sets.

Chapter II gives basic definitions and constructions of category theory. This chapter may be skipped completely if the reader is already familiar with left and right Kan extensions.

Chapter III explains how to construct the irreducible representations of a category from the irreducible representations of the endomorphism algebras of its objects. The results of this chapter may not be new, but are perhaps not widely known or anywhere assembled. Theorem 3.3.4 gives the recipe for extending an irreducible representation of an endomorphism algebra to a representation of \mathcal{D} . We give a similar construction for extending projective covers and injective hulls. Theorem 3.3.7 gives a convenient classification/description of the irreducible representations. Theorem 3.3.10 tells how to compute the dimension of an irreducible representation evaluated at an object.

Chapter IV explains when every finitely presented representation is built up from finitely many irreducibles (in the sense of composition series, not direct sums). It gives the algorithm that takes a presentation matrix over \mathcal{D} with coefficients in a field and returns the multiplicities of irreducible representations present in a composition series. The results of this section are due to the author, with many appearing already in [WG15]. Besides the main theorems (Theorem 4.3.2 detecting categories of dimension zero and Theorem 4.3.5 computing multiplicities), we give a result explaining that categories of dimension zero exhibit a strong form of homological stability (Corollary 4.2.6) and are Morita equivalent to categories that are “obviously” dimension zero (Theorem 4.3.11). We also give several results on Hilbert series, including

a generalization that works over the integers (see Corollary 4.4.5).

Chapter V gives the complete story for the category of finite sets, relying freely on the constructions of Chapter II and the results of Chapters III and IV. The level of computational detail may be excessive for most readers; a summary of the results in this chapter may be found in Section 5.1. Theorem 5.1.6 for computing simple multiplicities from a presentation matrix is new. Most of the other results can be found in [Rai09] who relies heavily on [Put96]. Our account situates these results in an organized general framework.

1.1 What is a representation?

We have already suggested that a functor $\mathcal{D} \rightarrow \mathbf{Mod}_R$ is somehow a representation of \mathcal{D} . This definition is appealing and concise; we now explain where it comes from and what it means.

A classical representation of a group G converts group elements $g \in G$ to matrices $\varphi(g)$ so that

$$\begin{aligned}\varphi(1) &= 1 \\ \varphi(g) \circ \varphi(h) &= \varphi(g \circ h).\end{aligned}$$

Since the axioms of a group capture the notion of *symmetry*, a representation of a group converts abstract symmetries (group elements) to concrete linear symmetries (invertible square matrices).

But “symmetry” means “invertible self-transformation;” we should study the more fundamental notion. Our view:

A representation converts **abstract transformations** to
concrete linear transformations.

Dropping the assumption of invertibility reflects the irreversibility of many important processes; dropping the assumption of composability reflects that some processes cannot be performed one after the other.²

1.1.1 Categories and their representations

Just as the group axioms formalize the notion of symmetry, the category axioms formalize the notion of transformation. For example, linear transformations between vector spaces form the **category of vector spaces**. A category feels a bit different from a group. Whereas any two invertible $n \times n$ matrices compose happily, linear transformations do not always compose. Two transformations $\mathbb{Q}^m \rightarrow \mathbb{Q}^n$ and $\mathbb{Q}^p \rightarrow \mathbb{Q}^q$ are **composable** exactly when $n = p$. In other words, a transformation passes from a **source** to a **target**, and two transformations compose only when the target of the first matches the source of the second. The possible sources and targets form the collection of **objects** of the category, and the transformations themselves form the collection of **morphisms**. If \mathcal{D} is a category, and $x, y \in \text{Ob}(\mathcal{D})$ are objects, write $\text{Hom}_{\mathcal{D}}(x, y)$ for the set of morphisms passing from x to y . Each hom-set $\text{Hom}_{\mathcal{D}}(x, x)$ has a distinguished element 1_x representing the identity transformation.

A **representation V of a category \mathcal{D}** consists of a vector space Vx for every object $x \in \mathcal{D}$ and a linear map $Vf: Vx \rightarrow Vy$ for every morphism $f \in \text{Hom}_{\mathcal{D}}(x, y)$ so that

$$V1_x = 1_{Vx}$$

$$Vf \circ Vg = V(f \circ g)$$

for any pair f, g of composable morphisms. If \mathcal{D} has a single object (call it $*$) and every morphism has a compositional inverse, then \mathcal{D} “is” a group in the sense that

²For example, it’s impossible to reply-all to an email before it’s been sent. Or to fly a connecting itinerary DTW-LGA, LAX-ORD. Or to castle kingside, then castle queenside.

$\text{Hom}_{\mathcal{D}}(*, *)$ is a set of elements closed under an associative composition law for which all elements are invertible and which has an identity. Further, a representation of \mathcal{D} is a vector space $V*$ with an action of $\text{Hom}_{\mathcal{D}}(*, *)$, and we recover the usual notion of a representation.

A **map of representations** $\varphi: V \rightarrow W$ is a collection of linear maps $\varphi_d: Vd \rightarrow Wd$ indexed by objects $d \in \text{Ob}(\mathcal{D})$. These maps are required to intertwine the structure maps of V and W in the sense that $\varphi_y \circ (Vf) = (Wf) \circ \varphi_x$ for any \mathcal{D} -morphism $f: x \rightarrow y$. In other words, we require a map of representations to “commute with the action of \mathcal{D} .” In Chapter II we shall explain that taking kernels, images, and cokernels still makes sense for maps of representations.

It bears mentioning that the representation theory of a category is sometimes called the representation theory of a **quiver with relations**, particularly if \mathcal{D} is described by generating morphisms subject to certain identifications. Quiver theory usually tries to answer finer questions than we will address, and usually restricts itself to finite categories \mathcal{D} , which we will not.

There is also a connection to M. Auslander’s theory of **coherent functors** [Aus66]. In his original formulation, the category \mathcal{D} is assumed to be abelian, and a coherent functor is a cokernel of a map between representables. Subsequent researchers, especially Henning Krause, have continued the study using a broader notion that assumes only that \mathcal{D} is additive [Kra02b] [Kra02a] [Kra03]. In this case, a finitely presented representation of a category \mathcal{D} is the same as a coherent functor on its additivization \mathcal{D}^{\oplus} .

1.1.2 Matrices and additive categories

Representation theory inherits its power from the elegance and concreteness of linear algebra. We say, “linear transformations are concrete” because any such transformation may be written as a rectangle of numbers—a **matrix**—with composition

given by an explicit rule using multiplication and addition. Mathematicians are good at numbers, so it makes sense to use them to study transformations.

But the entries of a matrix need not be numbers. Indeed, in a certain view, the entries of a matrix are little linear self-maps of the base field. This view acknowledges that we often have **block matrices** whose entries are more complicated linear transformations. For example, a linear map from a direct sum to a direct sum naturally takes the form of a block matrix. Since any vector space is a direct sum of copies of the ground field, and every linear map from the ground field to itself is multiplication by some number, matrices of numbers do capture every linear map. But, as every student of linear algebra learns, this description can obscure a transformation's essence.

Let us allow the entries of a matrix to be transformations, which is to say, morphisms. For block matrix multiplication to make sense, we must be able to add any two morphisms which have matching source and target. The correct formalism is that of an **additive category**. Such a category comes equipped with a **direct sum**, and a map between direct sums is given by a block matrix. Composing two block matrices relies on both composition and addition of morphisms.

Any category \mathcal{D} gives rise in a natural way to an additive category \mathcal{D}^\oplus where we allow formal direct sums of objects and formal \mathbb{Z} -linear combinations of morphisms. We often use **R -linear categories** where the morphism coefficients come from some ring R in place of \mathbb{Z} . This construction appears, for example, in [Mit72].

For the purposes of demonstration, if $a, b, c, p, q, x, y \in \text{Ob}(\mathcal{D})$ are objects and

$$\begin{array}{lll}
 \omega: p \rightarrow a & \zeta: q \rightarrow b & \xi: q \rightarrow b \\
 \varphi: x \rightarrow p & \psi: x \rightarrow p & \chi: x \rightarrow q \\
 \eta: y \rightarrow q & \sigma: p \rightarrow c & \tau: p \rightarrow c
 \end{array}$$

are morphisms, we may form matrices and multiply:

$$\begin{pmatrix} & p & q \\ x & \begin{bmatrix} \varphi + 5\psi & -\chi \end{bmatrix} \\ y & \begin{bmatrix} 0 & 2\eta \end{bmatrix} \end{pmatrix} \cdot \begin{pmatrix} a & b & c \\ p & \begin{bmatrix} \omega & 0 \\ 0 & \zeta + \xi \end{bmatrix} \\ q & \begin{bmatrix} \omega & 0 \\ 0 & \zeta + \xi \end{bmatrix} \end{pmatrix} = \begin{pmatrix} & a & b & c \\ x & \begin{bmatrix} \omega \circ \varphi + 5\omega \circ \psi & -\zeta \circ \chi - \xi \circ \chi & \sigma \circ \varphi + 5\sigma \circ \psi - \tau \circ \varphi - 5\tau \circ \psi \end{bmatrix} \\ y & \begin{bmatrix} 0 & 2\zeta \circ \eta + 2\xi \circ \eta & 0 \end{bmatrix} \end{pmatrix}.$$

Figure 1.5: Multiplying matrices over \mathcal{D} with coefficients in R

The row and column labels allow us to read off sources and targets, so in the future there will be no need to announce every morphism before it appears in a matrix. The source and target data may be conveyed even more succinctly by saying that the first matrix lies in the **matrix space** $\text{Mat}_R^{\mathcal{D}}(x \oplus y \rightarrow p \oplus q)$ and the second in $\text{Mat}_R^{\mathcal{D}}(p \oplus q \rightarrow a \oplus b \oplus c)$. **Matrix multiplication** then gives a map

$$\text{Mat}_R^{\mathcal{D}}(x \oplus y \rightarrow p \oplus q) \otimes_R \text{Mat}_R^{\mathcal{D}}(p \oplus q \rightarrow a \oplus b \oplus c) \longrightarrow \text{Mat}_R^{\mathcal{D}}(x \oplus y \rightarrow a \oplus b \oplus c).$$

1.1.3 Matrices as presentations

One convenient way to specify an R -module is to give a presentation by generators and relations. This data takes the form of a rectangular matrix giving a map from the free module on the relations mapping to the free module on the generators; the cokernel of this map is the module in question. Computational algebra often performs constructions on modules by interacting with their presentation matrices.

Concretely, if $A \in \text{Mat}_R(p \times q)$ is a classical $p \times q$ matrix over R , we could define the R -module presented by A

$$M_A = \frac{\text{Mat}_R(p \times 1)}{A \cdot \text{Mat}_R(q \times 1)}.$$

In other words, M_A is the quotient of the R -module $\text{Mat}_R(p \times 1)$ of column vectors by the submodule spanned by vectors of the form $A \cdot B$ for some $B \in \text{Mat}_R(q \times 1)$. We think of B as performing column operations on A so that every vector in the column span of A shows up in the denominator for some appropriately-chosen B .

Definition 1.1.1. Every matrix $C \in \text{Mat}_R^{\mathcal{D}}(x^{\oplus} \rightarrow y^{\oplus})$ gives rise to a representation of \mathcal{D} called V_C . We say V_C is finitely presented with presentation matrix C . Here is the formula:

$$V_C d = \frac{\text{Mat}_R^{\mathcal{D}}(x^{\oplus} \longrightarrow d)}{C \cdot \text{Mat}_R^{\mathcal{D}}(y^{\oplus} \rightarrow d)}.$$

In other words, $V_C d$ is the quotient of the matrix space $\text{Mat}_R^{\mathcal{D}}(x^{\oplus} \rightarrow d)$ by the R -submodule of matrices that factor through the matrix C .

This construction gives a different R -module for every $d \in \text{Ob}(\mathcal{D})$, but these R -modules are related. If $f: d \rightarrow d'$ is a morphism in \mathcal{D} , we have a post-multiplication map $V_C f: V_C d \rightarrow V_C d'$ given by $M \mapsto M \cdot [f]$; this map respects the denominator of the quotient since any matrix of the form $C \cdot B$ maps to $C \cdot B \cdot [f]$, and so its image is still in the denominator taking $B' = B \cdot [f]$. It follows that V_C is a representation of \mathcal{D} .

Just as every matrix $A \in \text{Mat}_R$ gives a presentation for M_A , an R -module, so every matrix $C \in \text{Mat}_R^{\mathcal{D}}(x^{\oplus} \rightarrow y^{\oplus})$ gives a presentation for V_C , a representation of \mathcal{D} .

1.2 What are the basic examples?

1.2.1 A familiar example of a finite presentation over a category

Let \mathcal{D} be the category with objects $\text{Ob}(\mathcal{D}) = \mathbb{Z}$ and morphisms

$$\text{Hom}_{\mathcal{D}}(n, m) = \{ \text{monomials of degree } m - n \text{ in the variables } w, x, y, z \}$$

where composition is given by multiplication of monomials. Build the matrix

$$C = \begin{matrix} & 2 & 2 & 2 \\ 0 & \begin{bmatrix} xz - y^2 & yw - z^2 & xw - yz \end{bmatrix} \end{matrix}.$$

Then, for every integer $n \in \mathbb{Z}$, the vector space $V_C n$ consists of the vector space of degree- n homogeneous polynomials modulo those which can be written as a polynomial combination of the three entries of C . In other words, V_C is the graded $\mathbb{Q}[w, x, y, z]$ -module which is the quotient of $\mathbb{Q}[w, x, y, z]$ by the submodule generated by those three particular elements in degree 2. The idea that V_C gives rise to a vector space for every $n \in \mathbb{Z}$ is not exotic at all; indeed, we define the Hilbert series almost by reflex:

$$\varphi_C(t) = \sum_{n=-\infty}^{\infty} (\dim V_C n) \cdot t^n.$$

The program `Macaulay2` [GS] readily accepts the matrix C as input, computing

$$\varphi_C(t) = \frac{1 - 3t^2 + 2t^3}{(1 - t)^4} = 1 + 4t + 7t^2 + 10t^3 + 13t^4 + 16t^5 + 19t^6 + 22t^7 + 25t^8 + \dots$$

Be inspired! The definition of C was a single line, but determines an infinite sequence of interesting vector spaces. Then, by the magic of Gröbner bases, a computer calculates the entire sequence in a fraction of a second.

1.2.2 Another beginning example

Consider graph colorings of the five-cycle




Figure 1.6: A valid 3-coloring of the five-cycle

where no two adjacent nodes receive the same color. We may ask, “How many valid colorings will there be with n colors?” A full answer will be a sequence of numbers.

Since a “coloring” is just a certain kind of function from the nodes to the colors, we seek a category where morphisms are functions. Let \mathcal{D} be the category whose objects are the finite sets $[n] = \{1, \dots, n\}$ for $n \in \mathbb{N}$ and whose morphisms are functions written in one-line notation. (For example, the identity function on $[3]$ is written $\boxed{123}$ and the constant function $[4] \rightarrow [1]$ is written $\boxed{1111}$). Consider the matrix

$$C = \begin{matrix} & [4] & [4] & [4] & [4] & [4] \\ [5] & \boxed{11234} & \boxed{12234} & \boxed{12334} & \boxed{12344} & \boxed{12341} \end{matrix}.$$

For every $n \in \mathbb{N}$, the vector space $V_C[n]$ is the free vector space on the functions $[5] \rightarrow [n]$ modulo the subspace spanned by those functions that factor through one of the five entries. Labeling the nodes with the numbers 1 through 5 clockwise starting at the top node, we see that a function $[5] \rightarrow [n]$ is an invalid coloring if and only if it factors through one of the entries of C . In other words, the remaining basis vectors of $V_C[n]$ give the valid colorings, and $\dim V_C[n]$ is the number of valid colorings of the graph using n colors.

Suppose we add two entries to the matrix

$$C' = \begin{matrix} & [4] & [4] & [4] & [4] & [4] & [5] & [5] \\ [5] & \boxed{11234} & \boxed{12234} & \boxed{12334} & \boxed{12344} & \boxed{12341} & \boxed{12345} - \boxed{23451} & \boxed{12345} - \boxed{15432} \end{matrix}.$$

Now $\dim V_{C'}[n]$ counts valid colorings up to rotation and reflection. In §5.5.2 we use a program based on Theorem 4.3.5 to find the Hilbert series for this example.

1.2.3 Configurations of distinct points in \mathbb{C}

Let z_1, \dots, z_n be coordinates on \mathbb{C}^n , and let $X_n \subset \mathbb{C}^n$ be the open subset where the de Rahm forms

$$\omega_{ij} = \frac{dz_i - dz_j}{z_i - z_j}$$

are defined. The space X_n can be thought of as parametrizing ordered n -tuples of distinct points in \mathbb{C} , with the form ω_{ij} measuring the winding of the i^{th} point around the j^{th} . Certainly $\omega_{ij} = \omega_{ji}$, but Arnol'd [Arn69] found a further relation for every triple $i, j, k \in [n]$

$$\omega_{ij} \wedge \omega_{jk} + \omega_{jk} \wedge \omega_{ki} + \omega_{ki} \wedge \omega_{ij} = 0.$$

He subsequently proved that there are no other relations, and that the cohomology algebra takes the form

$$H_{\text{dR}}^*(X_n) = \frac{\bigwedge^* \{\omega_{ij}\}}{(\omega_{ij} - \omega_{ji} \quad \omega_{ij} \wedge \omega_{jk} + \omega_{jk} \wedge \omega_{ki} + \omega_{ki} \wedge \omega_{ij})}.$$

Where i, j, k are distinct elements of the set $[n]$.

Once again, let \mathcal{D} be the category whose objects are the finite sets $[n]$ for $n \in \mathbb{N}$ and whose morphisms are functions. Arnol'd's presentation easily leads to the following matrix C giving a presentation for the de Rahm cohomology $H^2(X_n)$ in the sense that

$V_C[n] \simeq H^2(X_n)$.

$$C = \begin{array}{cccccc} & [4] & [3] & [4] & & [3] \\ [4] & \left[\boxed{1234} - \boxed{2134} \right. & \boxed{1123} & \boxed{1234} + \boxed{3412} & \boxed{1223} + \boxed{2331} + \boxed{3112} & \left. \right] \end{array}.$$

To interpret a function $[4] \rightarrow [n]$ as a monomial in the algebra, just read subscripts from left to right. For example, the function $\boxed{3281}$ corresponds to the monomial $\omega_{32} \wedge \omega_{81}$. We return to this example in §5.5.1 using the multiplicity theorem to compute the Hilbert series of V_C .

1.2.4 Monomials with combinatorial indexing

Let $x_{S,T}$ be a variable for every pair of subsets $S, T \subseteq [n]$, and consider the vector space V_n of degree-three monomials $x_{A,B}x_{P,Q}x_{S,T}$ modulo the relations

$$x_{A,E}x_{Q,P}x_{S,T} + x_{A,B}x_{P,E}x_{T,S} + x_{B,A}x_{P,Q}x_{S,E} = 0$$

$$x_{A,B}x_{AUB,SUT}x_{S,T} = x_{A,AUS}x_{AUT,BUS}x_{BUT,T}$$

for any subsets $A, B, P, Q, S, T, E \subseteq [n]$. What is the dimension of this space as $[n]$ grows?

Even for $n = 3$, there are 8 subsets of $[n]$ and so $8^7 = 2097152$ possible substitutions for the first relation. (The second relation depends on four variables, so it has only 4096 substitutions). We have every reason to expect subtle linear dependencies among generators.

The good news is that this vector space has no actual relevance, to my knowledge. The bad news is that similarly impractical presentations do have a way of cropping up. We introduce this presentation only as an example of the sort of detective work that must go into figuring out what category is responsible.

In order to imitate the previous examples, we would like each monomial

$$x_{A,B}x_{P,Q}x_{S,T}$$

to behave as if it were a morphism in some category \mathcal{D} . Since we are trying to evaluate at the finite set $[n]$, we might take $[n]$ with $n \in \mathbb{N}$ to be the objects of \mathcal{D} . Each morphism monomial carries with it 6 subsets of $[n]$. This observation suggests that a morphism $[6] \rightarrow [n]$ should be determined by a 6-tuple of subsets of $[n]$. The sorts of constructions appearing in the relations give clues about composition. In this case, the composition law must use union in some way.

The answer: we take \mathcal{D} to be the category of finite sets with relations. Explicitly, the objects are natural numbers and a morphism $n \rightarrow m$ is an $n \times m$ binary matrix; composition is given by matrix multiplication where $+$ is replaced by \max . Define the matrix

$$C = [6] \left[\begin{array}{c} \boxed{\begin{matrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{matrix}} + \overset{[7]}{\boxed{\begin{matrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{matrix}}} + \boxed{\begin{matrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{matrix}} - \overset{[4]}{\boxed{\begin{matrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{matrix}}} - \boxed{\begin{matrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{matrix}} \right].$$

Indeed, we have $V_C[n] \simeq V[n]$, and once again the combinatorics of a presentation are perfectly captured by the composition law in some combinatorial category.

Typical monomial	Description	Governing category
$z_i z_j y_{iijj}$	Subscripts drawn from a finite set	Finite sets with functions
$r_{i \leq j \leq k}$	Subscripts drawn from a totally-ordered set	Totally-ordered sets with monotone functions
$\omega_{i \neq j}$	Subscripts drawn from a set without repetition	Finite sets with injections
$x_{S \cup T, U} \wedge x_{S, T \cup U}$	Subscripts unions of subsets of a set	Finite sets with relations

Table 1.1: Examples of monomials for combinatorial categories

Of course, every category gives rise to a corresponding style of monomial, and so the examples given must not come close to capturing the notion.

1.3 How do we handle a presentation matrix classically...

1.3.1 When \mathcal{D} is a finite group?

Upon arriving at the sand volleyball court, four friends wish to pick teams for doubles. In how many ways is this possible? (You are right if you think the answer is 3, but the following argument is likely different from the one you have in mind.)

Introduce the symbol $\begin{array}{|c|} \hline a \ b \\ \hline c \ d \\ \hline \end{array}$ indicating the position of the four players on the court. Sometimes two court positionings give the same teams:

$$\begin{array}{|c|} \hline a \ b \\ \hline c \ d \\ \hline \end{array} = \begin{array}{|c|} \hline b \ a \\ \hline c \ d \\ \hline \end{array} \quad \begin{array}{|c|} \hline a \ b \\ \hline c \ d \\ \hline \end{array} = \begin{array}{|c|} \hline a \ b \\ \hline d \ c \\ \hline \end{array} \quad \begin{array}{|c|} \hline a \ b \\ \hline c \ d \\ \hline \end{array} = \begin{array}{|c|} \hline c \ d \\ \hline a \ b \\ \hline \end{array}.$$

These basic relations may generate other more complex relations. Conversely, these relations may be redundant in some ways. Still, they perfectly capture the idea of “picking teams without caring about shuffling within a side or which team is on which side of the net.” Let us write a corresponding 1×3 matrix over the group algebra

$\mathbb{C}S_4$:

$$C = \begin{matrix} & [4] & & [4] & & [4] \\ [4] & \left[\boxed{1234} - \boxed{2134} \right] & - & \left[\boxed{1234} - \boxed{1243} \right] & & \left[\boxed{1234} - \boxed{3412} \right] \end{matrix}.$$

Squinting if necessary, we see that this matrix fully expresses the three relations written more pictorially above. What's more, C is a presentation matrix for the $\mathbb{C}S_4$ -module of solutions to our original problem (where S_4 acts by permuting the four friends). In other words, the dimension of $M = \text{coker} \left((\mathbb{C}S_4)^{\oplus 3} \xrightarrow{C} \mathbb{C}S_4 \right)$ as a \mathbb{C} -vector space answers the question.

Let's take no shortcuts, and determine the decomposition of M into irreducible representations. The group algebra $\mathbb{C}S_4$ is semisimple by Maschke's theorem, and decomposes explicitly using the following map $\varphi = (\varphi_1, \varphi_2, \varphi_3, \varphi_4, \varphi_5)$ of \mathbb{C} -algebras

$$\begin{aligned} \mathbb{C}S_4 &\longrightarrow M_1(\mathbb{C}) \times M_3(\mathbb{C}) \times M_2(\mathbb{C}) \times M_3(\mathbb{C}) \times M_1(\mathbb{C}) \\ \boxed{2134} &\longmapsto \begin{bmatrix} 1 \end{bmatrix}, \begin{bmatrix} 0 & -1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 1 \\ 0 & -1 \end{bmatrix}, \begin{bmatrix} -1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & -1 & 0 \end{bmatrix}, \begin{bmatrix} -1 \end{bmatrix} \\ \boxed{2341} &\longmapsto \begin{bmatrix} 1 \end{bmatrix}, \begin{bmatrix} -1 & -1 & 0 \\ 1 & 0 & -1 \\ -1 & 0 & 0 \end{bmatrix}, \begin{bmatrix} -1 & 0 \\ 1 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 1 & 1 \\ -1 & 0 & 0 \\ 0 & -1 & 0 \end{bmatrix}, \begin{bmatrix} 1 \end{bmatrix} \end{aligned}$$

Here we have specified the map of algebras on a generating set for the group S_4 . The actual matrices were computed by Sage [Dev14]. Applying each φ_i to the presentation

matrix C in turn,

$$\begin{aligned} \varphi_1(C) &= \left[0 \mid 0 \mid 0 \right] \\ \varphi_2(C) &= \left[\begin{array}{ccc|ccc} 1 & 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 2 \end{array} \mid \begin{array}{ccc} 1 & -1 & -1 \\ 0 & 2 & 0 \\ -1 & -1 & 1 \end{array} \right] \\ \varphi_3(C) &= \left[\begin{array}{cc|cc} 0 & -1 & 0 & -1 \\ 0 & 2 & 0 & 2 \end{array} \mid \begin{array}{cc} 0 & 0 \\ 0 & 0 \end{array} \right] \\ \varphi_4(C) &= \left[\begin{array}{ccc|ccc} 2 & 0 & 0 & 0 & -1 & -1 \\ 0 & 1 & 1 & 0 & 2 & 0 \\ 0 & 1 & 1 & 0 & 0 & 2 \end{array} \mid \begin{array}{ccc} 1 & 0 & -1 \\ 1 & 2 & 1 \\ -1 & 0 & 1 \end{array} \right] \\ \varphi_5(C) &= \left[2 \mid 2 \mid 0 \right]. \end{aligned}$$

The coranks of these matrices correspond to the multiplicities of irreducible representations in M . The matrices $\varphi_2(C)$, $\varphi_4(C)$, and $\varphi_5(C)$ have full rank, and $\varphi_1(C)$ and $\varphi_3(C)$ have corank 1. In other words, the computation gives that M is composed of two irreducibles: one 1-dimensional, one 2-dimensional. Adding these together, we see that there are exactly 3 ways to pick doubles teams among four players.

There is a similar algorithm for decomposing any finitely presented module over a finite dimensional algebra using the indecomposable injective modules in place of the φ_i . This is a standard fact of representation theory; see, for example, [EGH⁺11][Prop. 9.2.3]. The main computational result of this thesis Theorem 4.3.5 is a suitable generalization of this construction in the context of the representation theory of categories.

1.3.2 When \mathcal{D} is a graded polynomial algebra?

Nine lily pads float in a perfect square, each populated with some number of indistinguishable frogs. At any moment, a frog may leap over a frog on an adjacent pad (no diagonals allowed), landing on the pad one further in that direction (no jumping into the water allowed). We consider two frog arrangements the same if

one can be obtained from the other by a sequence of leaps. How many different frog arrangements are there with exactly n frogs?

Introduce a variable for each lily pad:

$$\begin{array}{ccc} a & b & c \\ d & e & f \\ g & h & i. \end{array}$$

A monomial in these variables has a simple interpretation; b^3f^5 , for example, indicates that three frogs are situated at matrix coordinate $(1, 2)$, and five frogs sit at $(2, 3)$.

Build the 1×6 matrix

$$M = \left[\begin{array}{cccccc} ab - bc & de - ef & gh - hi & ad - dg & be - eh & cf - fi \end{array} \right]$$

with entries in the polynomial algebra $R = \mathbb{C}[a, b, c, d, e, f, g, h, i]$. Each entry specifies a pair of frog-placements that are meant to be the same. Indeed, this matrix gives a presentation for the module of solutions to our problem—except this time the answer is a graded module, and we must think of R as a graded algebra and M as a homogeneous matrix. The program `Macaulay2` easily computes the Hilbert series for this module:

$$\frac{-1 - 4t - 4t^2 + 3t^3 + 4t^4}{(t - 1)^5} = 1 + 9t + 39t^2 + 112t^3 + 251t^4 + 481t^5 + \dots$$

where the $39t^2$ means, for example, that there are 39 arrangements of two frogs.

1.4 Generalizing these classical techniques to representations of categories

The previous two examples— \mathcal{D} a finite group and \mathcal{D} a polynomial algebra—both have extremely satisfactory classical solutions, but the solutions differ in certain respects. Column reduction of matrices is parallelizable, and proceeds at a brisk, polynomial clip. Gröbner reduction of an ideal, while undeniably fast for ideals of low degree without too many variables, slows considerably on larger problems. Even the best Gröbner algorithms run in highly non-polynomial time.

Since the representation theory of categories generalizes both the representation theory of finite groups and the module theory of polynomial rings, we have great freedom in our line of inquiry. Our chosen direction is to study categories for which the representation theory has a computational style similar to the story for finite groups (or Artinian algebras), but where the final answer is still a Hilbert series. More specifically, we are concerned with categories “of dimension zero,” as we now explain.

1.4.1 The dimension of a category \mathcal{D}

Let us call a category \mathcal{D} **dimension zero over a ring** R if every presentation matrix $C \in \text{Mat}_R^{\mathcal{D}}(x^{\oplus} \rightarrow y^{\oplus})$ gives rise to a representation V_C that has a **composition series**, which is to say, a finite filtration

$$0 = V_C^0 \subsetneq V_C^1 \subsetneq \cdots \subsetneq V_C^k = V_C$$

so that the successive quotients V_C^i/V_C^{i-1} are **irreducible representations**, meaning they have no subrepresentations. A representation possessing this sort of filtration is said to have **finite length**. In summary, \mathcal{D} has dimension zero if and only if every finitely presented representation V_C has finite length.

When considered as a one-object category, any finite group is dimension zero over any Artinian ring R , and similarly for any category with a finite number of morphisms. For this reason, the representation theory of a dimension zero category “feels like” linear algebra over a field.

There is a more general notion of dimension due to Gabriel and Rentschler that assigns a number (or perhaps an ordinal) to every category (see [MR01], for example). For example, a graded polynomial algebra in 7 variables over an Artinian ring will have dimension 7. The representation theory of a higher-dimensional category “feels like” module theory more than representation theory.

1.4.2 Sam-Snowden Gröbner theory

The most successful approach to higher dimensional categories can be found in the important paper of Steven Sam and Andrew Snowden [SS]. They explain what extra structure on a category leads to a theory of Gröbner bases. Introducing the notion of a **lingual category**, they give practical results on the dimensions of representations of such categories. Many natural categories are lingual: the category of finite total orders with injections, the opposite of the category of finite sets with surjections, the category of finite sets with injections. In theory, their methods should be able to take matrices over lingual categories and compute Hilbert series as rational functions or as implicit power series satisfying some polynomial.

1.4.3 Representation stability

Thomas Church and Benson Farb’s work [CF13] isolated and named the **representation stability** phenomenon, which crops up whenever a sequence of representations seems to stabilize even though they may be growing and their actions come from different groups (Hilbert series are the basic example of growth that still “stabilizes;” representation stability generalizes this idea to representations of symmetric

groups and similar cases). Subsequent papers clarified and expanded the notion, offering rich results in varied subjects: [Chu12], [CEF15], [CEFN14], [Put13], [Put15], [CEF14], [Wil12], [JR11]. Sam-Snowden simultaneously spearheaded an overlapping story deepening certain aspects considerably and providing a connection to twisted commutative algebras and equivariant ideals in infinitely many variables, as well as solutions to combinatorial problems: [Sno13], [SS16a], [NSS16], [SS15], [SS16b].

In [WG14], this author took up the subject of presentation matrices over the category of finite sets, proving a strong form of representation stability. Those results were put to use in [EWG15]. The present work addresses a point raised in [SS, 11.1] indicating interest in a practical characterization of dimension zero categories.

CHAPTER II

Basic constructions

We introduce the fundamental notions of the representation theory of categories. Two canonical references are [ML98], [Mit72].

2.1 Categories, functors, and natural transformations

2.1.1 Basic definitions of category theory

Definition 2.1.1. A category $\mathcal{D} = (\text{Ob}(\mathcal{D}), \text{Hom}_{\mathcal{D}}, \circ_{\mathcal{D}})$ consists of

- a collection of **objects** $\text{Ob}(\mathcal{D})$,
- for any two objects $x, y \in \text{Ob}(\mathcal{D})$, a set of **morphisms** $\text{Hom}_{\mathcal{D}}(x, y)$,
- for any object $x \in \text{Ob}(\mathcal{D})$, a morphism $1_x \in \text{Hom}_{\mathcal{D}}(x, x)$, and
- for any triple of objects $x, y, z \in \text{Ob}(\mathcal{D})$, a **composition law**

$$\begin{aligned} \text{Hom}_{\mathcal{D}}(x, y) \times \text{Hom}_{\mathcal{D}}(y, z) &\longrightarrow \text{Hom}_{\mathcal{D}}(x, z) \\ (f, g) &\longmapsto g \circ_{\mathcal{D}} f, \end{aligned}$$

satisfying the **identity condition** $1_y \circ_{\mathcal{D}} f = f = f \circ_{\mathcal{D}} 1_x$ and the **associativity condition** $(h \circ_{\mathcal{D}} g) \circ_{\mathcal{D}} f = h \circ_{\mathcal{D}} (g \circ_{\mathcal{D}} f)$ whenever both sides of this last equation make sense.

Remark 2.1.2. Later, we shall generalize the composition law in \mathcal{D} to a product for matrices over \mathcal{D} . In this context, one may consider a morphism to be a 1×1 matrix over \mathcal{D} . For two morphisms f and g , this product will be written $f \cdot g$, abbreviating the diagrammatic style $x \xrightarrow{f} y \xrightarrow{g} z$. For the time being, we shall compose morphisms using the \circ -infix style $g \circ f$.

Definition 2.1.3. Let $f : x \rightarrow y$ be a morphism in a category \mathcal{D} . Say f is **epic** if for every $z \in \text{Ob}(\mathcal{D})$, the **precomposition function**

$$\begin{aligned} \text{Hom}_{\mathcal{D}}(y, z) &\rightarrow \text{Hom}_{\mathcal{D}}(x, z) \\ \varphi &\mapsto \varphi \circ f \end{aligned}$$

is an injection. Say f is **monic** if for every $w \in \text{Ob}(\mathcal{D})$, the **postcomposition function** $\text{Hom}_{\mathcal{D}}(w, x) \rightarrow \text{Hom}_{\mathcal{D}}(w, y)$, $\psi \mapsto f \circ \psi$ is an injection.

Definition 2.1.4. Let $f : x \rightarrow y$ be a morphism in a category \mathcal{D} . Say f is **split epic** if there exists a morphism $g : y \rightarrow x$ so that $f \circ g = 1_y$. Say f is **split monic** if there exists a morphism $g : y \rightarrow x$ so that $g \circ f = 1_x$. Say f is an **isomorphism** if there exists an **inverse morphism** $g : y \rightarrow x$ so that $fg = 1_x$ and $gf = 1_y$. Two objects x and y are called **isomorphic** if there exists an isomorphism between them.

2.1.2 R -linear categories

In the event that the sets $\text{Hom}_{\mathcal{D}}(x, y)$ come with the structure of an R -module for some commutative ring R , and that the composition laws are bilinear over R , we shall call \mathcal{D} an **R -linear category**.

Example 2.1.5. Given R , a commutative ring, the collection of R -modules and the linear transformations between them form a category called $\mathbf{Mod}_R = (\text{Mod}_R, \text{Hom}_R, \circ)$. Given two R -modules $M, N \in \text{Mod}_R$, the set of morphisms $\text{Hom}_R(M, N)$ comes with an addition and an action of R , and so \mathbf{Mod}_R is an R -linear category.

Definition 2.1.6. Given categories \mathcal{C} and \mathcal{D} , a **functor** $f : \mathcal{C} \longrightarrow \mathcal{D}$ consists of

- an object $fc \in \mathcal{D}$ for every object $c \in \mathcal{C}$, and
- a morphism $f\varphi \in \text{Hom}_{\mathcal{D}}(fc, fc')$ for every $\varphi \in \text{Hom}_{\mathcal{C}}(c, c')$ with $c, c' \in \text{Ob}(\mathcal{C})$,

so that $f1_c = 1_{fc}$ for all objects $c \in \mathcal{C}$, and $f\varphi \circ_{\mathcal{D}} f\psi = \mathcal{F}(\varphi \circ_{\mathcal{C}} \psi)$ whenever both sides of this last equation make sense. If \mathcal{C} and \mathcal{D} are R -linear categories, then f is said to be an **R -linear functor** if the function

$$\begin{aligned} \text{Hom}_{\mathcal{C}}(c, c') &\longrightarrow \text{Hom}_{\mathcal{D}}(fc, fc') \\ f &\longmapsto ff \end{aligned}$$

is R -linear. A functor f is said to be **full** if this function is surjective, and **faithful** if injective. A functor f is said to be **essentially surjective on objects** if every object $d \in \mathcal{D}$ is isomorphic to an object of the form fc for some $c \in \mathcal{C}$. A functor is called an **equivalence of categories** if it is full, faithful, and essentially surjective on objects.

Definition 2.1.7. Given functors $f, g : \mathcal{C} \longrightarrow \mathcal{D}$, a **natural transformation** $\phi : f \longrightarrow g$ provides a **component morphism** $\phi_c \in \text{Hom}_{\mathcal{D}}(fc, gc)$ for every object $c \in \mathcal{C}$, so that, for every $\psi \in \text{Hom}_{\mathcal{C}}(c, c')$, $(g\psi) \circ_{\mathcal{D}} \phi_c = \phi_{c'} \circ_{\mathcal{D}} (f\psi)$.

2.2 Representations and the category $\text{Mod}_R^{\mathcal{D}}$

Definition 2.2.1. A **representation** of a category \mathcal{D} over a commutative ring R is a functor $V : \mathcal{D} \longrightarrow \text{Mod}_R$. If \mathcal{D} is an R -linear category, then we require that V be an R -linear functor.

Definition 2.2.2. If V, W are representations of \mathcal{D} over R , a **map of representations** $\phi : V \longrightarrow W$ is a collection of R -linear maps $\phi_d : Vd \longrightarrow Wd$ indexed by

objects $d \in \text{Ob}(\mathcal{D})$ so that for any pair of objects $d, d' \in \mathcal{D}$ and morphism $f : d \rightarrow d'$, the following square commutes in the the category of R -modules:

$$\begin{array}{ccc} Vd & \xrightarrow{\phi_d} & Wd \\ Vf \downarrow & & \downarrow Wf \\ Vd' & \xrightarrow{\phi_{d'}} & Wd'. \end{array}$$

In other words, it is a natural transformation from V to W . Evidently, the representations of \mathcal{D} over a commutative ring R form a category where natural transformations compose componentwise; write $\mathbf{Mod}_R^{\mathcal{D}} = (\text{Mod}_R^{\mathcal{D}}, \text{Hom}_R^{\mathcal{D}}, \circ)$ for the category whose objects are representations of \mathcal{D} over R and whose morphisms are maps of representations.

Observation 2.2.3. *Kernels and cokernels in $\mathbf{Mod}_R^{\mathcal{D}}$ may be computed objectwise in \mathbf{Mod}_R . More precisely, any map of representations $\phi : V \rightarrow W$ determines two more representations, called $\ker \phi$ and $\text{coker } \phi$, fitting into a commutative diagram*

$$\begin{array}{ccccccccc} 0 & \longrightarrow & (\ker \phi)d & \longrightarrow & Vd & \xrightarrow{\phi_d} & Wd & \longrightarrow & (\text{coker } \phi)d & \longrightarrow & 0 \\ \downarrow & & (\ker \phi)f \downarrow & & Vf \downarrow & & \downarrow Wf & & \downarrow (\text{coker } \phi)f & & \downarrow \\ 0 & \longrightarrow & (\ker \phi)d' & \longrightarrow & Vd' & \xrightarrow{\phi_{d'}} & Wd' & \longrightarrow & (\text{coker } \phi)d' & \longrightarrow & 0 \end{array}$$

where the rows are exact sequences of R -modules. These representations behave in perfect analogy with usual kernels and cokernels of maps between R -modules. We may similarly construct the image as the kernel of the map to the cokernel. A **subrepresentation of V** is one that appears as a kernel of a map out of V ; a **quotient representation of V** is one that appears as a cokernel of a map into V . Be assured that these constructions behave exactly as expected. See [Wei94, 1.6.4].

2.3 Matrices over a category with coefficients in a ring

Definition 2.3.1. For \mathcal{D} a category and objects $x_1, \dots, x_k; y_1, \dots, y_l \in \text{Ob}(\mathcal{D})$, define the R -module of matrices over \mathcal{D} ,

$$\text{Mat}_R^{\mathcal{D}}(x_1 \oplus \dots \oplus x_k \rightarrow y_1 \oplus \dots \oplus y_l) = \bigoplus_{i=1}^k \bigoplus_{j=1}^l R \cdot \text{Hom}_{\mathcal{D}}(x_i, y_j).$$

where each summand is a free module on a hom-set of \mathcal{D} . Matrices over a category enjoy a natural **matrix multiplication**

$$\begin{aligned} \text{Mat}_R^{\mathcal{D}}(\oplus_i x_i \rightarrow \oplus_j y_j) \otimes_R \text{Mat}_R^{\mathcal{D}}(\oplus_j y_j \rightarrow \oplus_k z_k) &\longrightarrow \text{Mat}_R^{\mathcal{D}}(\oplus_i x_i \rightarrow \oplus_k z_k) \\ (M \otimes N) &\longmapsto M \cdot N \end{aligned}$$

where the (i, k) -entry of $M \cdot N$ is given by the formula $\sum_j N_{jk} \circ M_{ij}$ extending the composition law in \mathcal{D} by linearity.

If x^{\oplus} and y^{\oplus} are formal sums of objects of \mathcal{D} , any functor $f : \mathcal{C} \rightarrow \mathcal{D}$ induces an R -linear map

$$\text{Mat}_R^{\mathcal{C}}(x^{\oplus} \rightarrow y^{\oplus}) \longrightarrow \text{Mat}_R^{\mathcal{D}}(fx^{\oplus} \rightarrow fy^{\oplus})$$

by entrywise application of f . Evidently, this construction commutes with matrix multiplication.

Definition 2.3.2 (Evaluating a representation on a matrix). Given $V \in \text{Mod}_R^{\mathcal{D}}$ and $M \in \text{Mat}_R^{\mathcal{D}}(\oplus_{i=1}^k x_i \rightarrow \oplus_{j=1}^l y_j)$, we form a $k \times l$ block matrix

$$V(M) : \bigoplus_{i=1}^k Vx_i \longrightarrow \bigoplus_{j=1}^l Vy_j$$

by entrywise application of V .

2.3.1 The R -linearization of a category

In the same way that a representation of a group is the same data as a representation of the group algebra, there is a natural way to convert a category \mathcal{D} to an R -linear category \mathcal{D}^\oplus . An object $d^\oplus \in \mathcal{D}^\oplus$ is a formal direct sum of objects of \mathcal{D} , and a morphism is a matrix over \mathcal{D} with coefficients in R . Matrix multiplication gives the composition law. Since our representations have a right action of \mathcal{D} , our convention is that matrix multiplication proceeds left-to-right along the direction of the morphisms of \mathcal{D} , in contrast to the \circ -infix notation that reverses the order of arrows. For this reason, we shall not reason about \mathcal{D}^\oplus as a category, preferring to think of matrices over \mathcal{D} as computational objects that may be added and multiplied.

2.4 Yoneda's lemma and basic projectives in the category

$\mathbf{Mod}_R^{\mathcal{D}}$

Definition 2.4.1 (Basic projectives). For any object $d \in \mathcal{D}$ and commutative ring R , define a representation $P^d \in \mathbf{Mod}_R^{\mathcal{D}}$, the **basic projective at d** , given on objects by the formula $P^d(x) = R \cdot \mathrm{Hom}_{\mathcal{D}}(d, x)$ and on morphisms $\varphi: x \rightarrow y$ by the formula

$$\begin{aligned} \mathrm{Hom}_{\mathcal{D}}(d, x) &\longrightarrow \mathrm{Hom}_{\mathcal{D}}(d, y) \\ f &\longmapsto \varphi \circ f \end{aligned}$$

extended by linearity to the free R -module $R \cdot \mathrm{Hom}_{\mathcal{D}}(d, d')$. The finite direct sum $P = P^{x_1} \oplus \cdots \oplus P^{x_k}$ admits the easy description

$$Py = \mathrm{Mat}_R^{\mathcal{D}}(x_1 \oplus \cdots \oplus x_k \rightarrow y);$$

in this description, a morphism $\varphi: y \rightarrow z$ acts by right matrix multiplication with the single-entry matrix $[\varphi] \in \mathrm{Mat}_R^{\mathcal{D}}(y \rightarrow z)$.

The next lemma shows that P^d behaves like a module “freely generated by the vector $1_d \in R \cdot \text{Hom}_{\mathcal{D}}(d, d) = P^d d$.”

Lemma 2.4.2 (Yoneda). *Let $V \in \text{Mod}_R^{\mathcal{D}}$ be a representation, and let $d \in \text{Ob}(\mathcal{D})$ be an object. There is an isomorphism of R -modules*

$$\text{Hom}_R^{\mathcal{D}}(P^d, V) \xrightarrow{\sim} Vd$$

given by $\phi \mapsto \phi_d(1_d)$.

Proof. The inverse map is given by the formula $v \mapsto (f \mapsto (Vf)(v))$. □

As a consequence, the representations P^d are projective, for the same reason that free modules are projective: a lift of the destination of the identity vector 1_d provides a lift of the map. The following basic fact grants us access to the representation theory of \mathcal{D} using matrices over \mathcal{D} .

Corollary 2.4.3. *By the Yoneda lemma, every matrix*

$$M \in \text{Mat}_R^{\mathcal{D}}(x_1 \oplus \cdots \oplus x_k \rightarrow y_1 \oplus \cdots \oplus y_l)$$

gives rise to a unique map

$$M^{\natural} : P^{y_1} \oplus \cdots \oplus P^{y_l} \longrightarrow P^{x_1} \oplus \cdots \oplus P^{x_k}$$

fitting into the commutative diagram

$$\begin{array}{ccc} P^{y_1}d \oplus \cdots \oplus P^{y_l}d & \xrightarrow{(M^{\natural})_d} & P^{x_1}d \oplus \cdots \oplus P^{x_k}d \\ \sim \downarrow & & \downarrow \sim \\ \text{Mat}_R^{\mathcal{D}}(y_1 \oplus \cdots \oplus y_l \rightarrow d) & \xrightarrow[M \cdot -]{} & \text{Mat}_R^{\mathcal{D}}(x_1 \oplus \cdots \oplus x_k \rightarrow d). \end{array}$$

The assignment takes matrix multiplication to composition: $(M \cdot N)^{\natural} = M^{\natural} \circ N^{\natural}$.

2.4.1 Finitely generated representations of a category

We provide many equivalent ways to think about finite generation.

Proposition 2.4.4. *The following are equivalent conditions on a representation $V \in \mathbf{Mod}_R^{\mathcal{D}}$:*

- *There exists a finite list of objects $d_i \in \mathbf{Ob}(\mathcal{D})$ and a surjection $\oplus_i P^{d_i} \rightarrow V$;*
- *There exist finitely many vectors $v_i \in Vd_i$ so that any subrepresentation $W \subseteq V$ with each $v_i \in Wd_i$ has $W = V$;*
- *There exist finitely many vectors $v_i \in Vd_i$ so that any strict subrepresentation $W \subset V$ misses some v_i ;*
- *There exists a row vector $v \in \oplus_i Vd_i$, $i \in \{1, \dots, k\}$, so that any other vector $w \in Vd$ has the form $v \cdot (VM)$ for some matrix $M \in \mathbf{Mat}_R^{\mathcal{D}}(\oplus_{i=1}^k d_i \rightarrow d)$;*
- *If $V = \sum_{\alpha} V_{\alpha}$ where each $V_{\alpha} \subseteq V$, then there exists a finite list $\alpha_1, \dots, \alpha_k$ so $V = \sum_{i=1}^k V_{\alpha_i}$.*

Definition 2.4.5. Say $V \in \mathbf{Mod}_R^{\mathcal{D}}$ is **finitely generated** if it satisfies any of the equivalent hypotheses of Proposition 2.4.4. Say it is **finitely presented** if there exists some matrix $M \in \mathbf{Mat}_R^{\mathcal{D}}(\oplus_i x_i \rightarrow \oplus_j y_j)$ so that $V = \text{coker } M^{\natural}$; we say M is a **presentation** for V .

Certainly any finitely presented representation is finitely generated using the basis vectors 1_{x_i} . In the case where these two notions coincide, we borrow the corresponding term from module theory.

Definition 2.4.6. The category \mathcal{D} is **Noetherian** over R if every finitely generated $V \in \mathbf{Mod}_R^{\mathcal{D}}$ is finitely presented, if and only if every subrepresentation of a finitely generated representation is finitely generated.

This definition, which may appear unassuming, captures an extremely subtle property of \mathcal{D} , but one which is absolutely crucial for computations. Succinctly: if \mathcal{D} is Noetherian, every matrix over \mathcal{D} has a nullspace given by another matrix over \mathcal{D} .

2.5 Idempotents and projectives

Definition 2.5.1. A square matrix $\pi \in \text{Mat}_R^{\mathcal{D}}(d^{\oplus} \rightarrow d^{\oplus})$ is **idempotent** if $\pi \cdot \pi = \pi$. It is an **indecomposable idempotent** if $\pi \neq 0$ and any decomposition $\pi = \pi' + \pi''$ as a sum of two idempotents $\pi', \pi'' \in \text{Mat}_R^{\mathcal{D}}(d^{\oplus} \rightarrow d^{\oplus})$ has either $\pi' = 0$ or $\pi'' = 0$.

Proposition 2.5.2. *Every finitely generated projective $P \in \mathbf{Mod}_R^{\mathcal{D}}$ is a direct summand of a finite direct sum of basic projectives. Equivalently, there exists an idempotent matrix π over \mathcal{D} with coefficients in R so that $P \simeq \text{im } \pi^{\natural}$. Conversely, if π is idempotent, then $\text{im } \pi^{\natural}$ is projective. For any $V \in \text{Mod}_R^{\mathcal{D}}$, there is a natural isomorphism*

$$\text{Hom}_R^{\mathcal{D}}(P, V) \simeq \text{Im } V(\pi),$$

where $V(\pi)$ denotes entrywise application of the representation V to the matrix π as in Definition 2.3.2.

Proof. Since P is finitely generated, it admits a surjection from some direct sum of basic projectives $P^{d_1} \oplus \dots \oplus P^{d_k}$ using the first condition of Proposition 2.4.4. By projectivity of P , this surjection splits. The composite is an endomorphism of $P^{d_1} \oplus \dots \oplus P^{d_k}$ which must come from an idempotent matrix by Corollary 2.4.3. Conversely, by Corollary 2.4.3, π^{\natural} is an idempotent endomorphism of a direct sum of basic projectives. Since a summand of a sum of projectives is projective, $\text{im } \pi^{\natural}$ is projective as well. The last isomorphism is a consequence of Yoneda's Lemma 2.4.2.

□

Definition 2.5.3. The **category of idempotents** $\Pi_R(\mathcal{D})$ of a category \mathcal{D} over a

commutative ring R is the R -linear category whose objects are idempotent matrices

$$\text{Ob}(\Pi_R(\mathcal{D})) = \{\pi \in \text{Mat}_R^{\mathcal{D}}(d^{\oplus} \rightarrow d^{\oplus}) \mid d^{\oplus} \in \mathcal{D}^{\oplus}, \pi \cdot \pi = \pi\}$$

and whose morphisms from $\pi_1 \in \text{Mat}_R^{\mathcal{D}}(d_1^{\oplus} \rightarrow d_1^{\oplus})$ to $\pi_2 \in \text{Mat}_R^{\mathcal{D}}(d_2^{\oplus} \rightarrow d_2^{\oplus})$ are given by

$$\text{Hom}_{\Pi(\mathcal{D})}(\pi_1, \pi_2) = \{\gamma \in \text{Mat}_R^{\mathcal{D}}(d_1^{\oplus} \rightarrow d_2^{\oplus}) \mid \pi_1 \cdot \gamma \cdot \pi_2 = \gamma\}.$$

Composition in $\Pi_R(\mathcal{D})$ is matrix multiplication.

The category of idempotents goes by many other names: the Cauchy completion, the Karoubi envelope, the idempotent completion. Its construction is standard in category theory and algebra. See [KS06, p. 66], for example.

Proposition 2.5.4. *The category of idempotents $\Pi_R(\mathcal{D})$ is equivalent to the full subcategory of $(\mathbf{Mod}_R^{\mathcal{D}})^{op}$ spanned by the finitely generated projectives.*

Proof. The equivalence is given by the construction $\pi \mapsto \pi^{\natural}$ together with Corollary 2.4.3. □

2.6 Left and right Kan extensions

Induction and coinduction play a fundamental role in classical representation theory. In module theory, these notions correspond to extension and coextension of scalars. The generalization of these familiar concepts in the setting of categories: left and right Kan extension.

Any functor $f : \mathcal{C} \rightarrow \mathcal{D}$ determines a **pullback functor**

$$f^* : \mathbf{Mod}_R^{\mathcal{D}} \rightarrow \mathbf{Mod}_R^{\mathcal{C}}.$$

Definition 2.6.1. For any representation $V \in \mathbf{Mod}_R^{\mathcal{D}}$, define the **pullback repre-**

sentation $f^*V \in \text{Mod}_R^{\mathcal{C}}$ on an object $c \in \mathcal{C}$ by the formula $(f^*V)c = V(fc)$, and on a morphism $\varphi : c \rightarrow c'$ by $(f^*V)\varphi = V(f\varphi)$.

The next result provides left and right adjoints for the pullback functor. The details about units and counits will be necessary for proofs, but not for casual reading; in any event, suitable background on adjoint functors may be found in [ML98, IV].

Proposition 2.6.2. *There exists a pair of functors $f_!, f_* : \text{Mod}_R^{\mathcal{C}} \rightarrow \text{Mod}_R^{\mathcal{D}}$ called the **left and right Kan extensions along f** with the properties that, for any $V \in \text{Mod}_R^{\mathcal{C}}$ and $W \in \text{Mod}_R^{\mathcal{D}}$,*

$$\text{Hom}_R^{\mathcal{D}}(V, f_*W) \simeq \text{Hom}_R^{\mathcal{C}}(f^*V, W)$$

$$\text{Hom}_R^{\mathcal{D}}(f_!V, W) \simeq \text{Hom}_R^{\mathcal{C}}(V, f^*W).$$

Moreover, these isomorphisms are natural in the variables V and W . Concisely, we have described two adjunctions: $f_! \dashv f^*$ and $f^* \dashv f_*$. The unit and counit for the first adjunction will be written $\eta : 1 \rightarrow f^*f_!$ and $\varepsilon : f_!f^* \rightarrow 1$. The unit and counit for the second adjunction will be written $\eta' : 1 \rightarrow f_*f^*$ and $\varepsilon' : f^*f_* \rightarrow 1$. They satisfy the **unit-counit equations** $1 = \varepsilon f_! \circ f^* \eta$, $1 = f^* \varepsilon \circ \eta f_!$ and $1 = \varepsilon' f^* \circ f_* \eta'$, $1 = f_* \varepsilon' \circ \eta f^*$.

Proof. The existence of left and right Kan extensions may be found in [ML98, X]. The unit-counit equations may be found in [ML98, IV]. \square

It may be helpful to think of $f_!V$ as the universal representation generated by V and f_*V as the universal representation cogenerated by V .

Readers familiar with the usual tensor-hom adjunction of module theory may think of a “counit” as some sort of “evaluation map” and a unit as some sort of “identity matrix.” Shortly, we will give concrete constructions for the left and right Kan extensions of a representation along a functor using a direct analog of the tensor-hom adjunction. We first list some appealing formal properties.

Proposition 2.6.3. *The functor f^* is exact, $f_!$ is right exact, and f_* is left exact. The functor $f_!$ takes projectives to projectives, and the functor f_* takes injectives to injectives. In the following table, an entry in the first column implies the other entries in its row.*

f	$f_!$	f^*	f_*	$f_!f^* \xrightarrow{\epsilon} 1$	$1 \xrightarrow{\eta} f^*f_!$	$f^*f_* \xrightarrow{\epsilon'} 1$	$1 \xrightarrow{\eta'} f_*f^*$
essentially surjective on objects	\implies	faithful		epic			monic
full + faithful	\implies	full + faithful	full + faithful		iso	iso	
full + ess. surj. on objects	\implies	full + faithful		iso			iso

Table 2.1: Useful properties of Kan extensions

Proof. Basic properties of adjoint functors relate the conditions appearing within a row; see [Wei94, 2.6]. The top row is obvious. The middle row may be found in [ML98, X.3.3], and has recently been used in a similar context by Gan and Li [GL]. The bottom row may be found in [Ves08, Proposition A.2]. \square

Proposition 2.6.4. *If M is a matrix over \mathcal{C} giving a presentation for V , then applying a functor $f: \mathcal{C} \rightarrow \mathcal{D}$ entrywise to the matrix M gives a presentation for $f_!V$:*

$$f_!V \simeq \operatorname{coker}(f(M))^{\natural}.$$

Proof. By right-exactness of $f_!$ it suffices to check that $f_!M^{\natural} = (fM)^{\natural}$. Let $c \in \operatorname{Ob}(\mathcal{C})$. Yoneda's Lemma 2.4.2 gives $P^c \simeq {}_iR$ where i denotes the inclusion of c with its

identity morphism into \mathcal{C} . Since $f \circ i = (f \circ i)_!$, we have that $f \circ i P^c \simeq P^{f^c}$. So

$$\begin{aligned} \mathrm{Hom}_R^{\mathcal{D}}(f_! P^c, f_! P^{c'}) &\simeq \mathrm{Hom}_R^{\mathcal{D}}(P^{f^c}, P^{f^{c'}}) \\ &\simeq R \cdot \mathrm{Hom}_{\mathcal{D}}(f^c, f^{c'}) \end{aligned}$$

again by Yoneda. Applying this computation to the entries of M gives the result. \square

2.6.1 Construction of the functors $f_!$ and f_*

If $f: R \rightarrow S$ is a ring homomorphism, then any S -module restricts to an R -module. This construction has a left adjoint called extension of scalars which may be thought of as a tensor product. For example, if M_R is a right R -module, then we form the tensor product

$$M_R \otimes_R R S_S$$

where R acts on S from the left using f . The usual tensor-hom adjunction reads

$$\mathrm{Hom}_S(M_R \otimes_R R S_S, N_S) \simeq \mathrm{Hom}_R(M_R, \mathrm{Hom}_S({}_R S_S, N_S))$$

where we recognize that $\mathrm{Hom}_S({}_R S_S, N_S)$ is a complicated way to refer to the restricted module $f^* N_S$. In other words, the tensor product on the left side of the adjunction provides an explicit model for a left adjoint to restriction, which we should call $f_!$:

$$f_! M_R \simeq M_R \otimes_R R S_S.$$

This discussion motivates the following definition. Further explanation and context may be found in [Rie14, §4.1].

Definition 2.6.5. If $V: \mathcal{C} \rightarrow \mathbf{Mod}_R$ is a right \mathcal{C} -module and $W: \mathcal{C}^{op} \rightarrow \mathbf{Mod}_R$ is a

left \mathcal{C} -module, the **functor tensor product** is defined

$$V \otimes_{\mathcal{C}} W = \frac{\bigoplus_{c \in \text{Ob}(\mathcal{C})} Vc \otimes_R Wc}{v\varphi \otimes w - v \otimes \varphi v}$$

where $\varphi: c \rightarrow c'$ ranges over all morphisms in the category \mathcal{C} and $v \in Vc$, $w \in Wc'$.

Remark 2.6.6. If the category \mathcal{C} has infinitely many objects, then this description of the functor tensor product cannot be used directly for computation. However, if V is finitely generated, then we may restrict this sum to the degrees of the generators of V . Indeed, any vector can be written in the span of the images of the generators under the action of \mathcal{C} , and these morphisms may be moved across the tensor sign $\otimes_{\mathcal{C}}$.

Proposition 2.6.7. *If $f: \mathcal{C} \rightarrow \mathcal{D}$ is a functor and $V \in \text{Mod}_R^{\mathcal{C}}$ a representation, the left Kan extension of V along f takes the form*

$$f_! Vd \simeq V \otimes_{\mathcal{C}} \text{Hom}_{\mathcal{D}}(f-, d),$$

where $\text{Hom}_{\mathcal{D}}(f-, -)$ carries a left action of \mathcal{C} in its first coordinate, and a right action of \mathcal{D} in its second coordinate.

Proof. We show that the functor tensor product satisfies a direct analog of the usual tensor-hom adjunction. Let M be a right \mathcal{C} -module, B be a \mathcal{C}, \mathcal{D} -bimodule, and N be a left \mathcal{D} -module. We have an isomorphism

$$\text{Hom}_R^{\mathcal{C}}(M \otimes_{\mathcal{C}} B, N) \xrightarrow{\Phi} \text{Hom}_R^{\mathcal{D}}(M, \text{Hom}_R^{\mathcal{C}}(B, N))$$

natural in the variables M and N . The map is defined by the formula $\Phi(\varphi)(m)(b) = \varphi(m \otimes b)$, and its inverse by the formula $\Phi^{-1}(\psi)(m \otimes b) = \psi(m)(b)$. With this adjunction in hand, it is enough to observe that f^* is represented by the bimodule $\text{Hom}_{\mathcal{D}}(f-, -)$; the formula then follows by uniqueness of adjoints. A more detailed

proof of this adjunction may be found in [AW92]. A direct proof of this construction of $f_!$ may be found in [Rie14, Example 4.1.5]. \square

We state the corresponding formula for right Kan extensions.

Corollary 2.6.8. *If R is a field and $V \in \text{Mod}_R^{\mathcal{C}}$ is a representation, then the right Kan extension of V along a functor $f: \mathcal{C} \rightarrow \mathcal{D}$ takes the form*

$$f_*Vd \simeq V \otimes_{\mathcal{C}} \text{Hom}_{\mathcal{D}}(d, f-)^{\vee},$$

where we have used a dual vector space to reverse the variance of both actions on the bimodule $\text{Hom}_{\mathcal{D}}(-, f-)$.

Proof. Apply the dual vector space functor to Proposition 2.6.7. \square

2.6.2 Two easy facts about the functors $f_!$, f^* , and f_*

The following two facts will be used in Proposition 3.1.11 to give circumstances under which $f_!$ and f_* preserve indecomposability, and later to help with the proof of Theorem 3.3.4. In what follows, let V be a representation of \mathcal{C} , and let $f: \mathcal{C} \rightarrow \mathcal{D}$.

Proposition 2.6.9. *Suppose $U \subseteq f_*V$. If $f^*U = 0$, then $U = 0$.*

Proof. By naturality of the unit η' we have a commuting square

$$\begin{array}{ccc} U & \xrightarrow{\subseteq} & f_*V \\ \eta'_U \downarrow & & \downarrow \eta'_{f_*V} \\ f_*f^*U & \longrightarrow & f_*f_*f^*V. \end{array}$$

The right vertical arrow is split monic by the unit-counit equations, so the clockwise composite map is injective, and the left vertical map must be injective as well. On the other hand, the lower left representation is zero since $f^*U = 0$. \square

Corollary 2.6.10. *Dually, if $f_!V \twoheadrightarrow Q$ with $f^*Q = 0$, then $Q = 0$.*

Proof. An indication of the dual argument: the composite $ff^*fV \rightarrow fV \rightarrow Q$ is surjective, and so by naturality of the counit ε , the representation Q is a quotient of $ff^*Q = 0$. □

CHAPTER III

Irreducible representations, injective hulls, and projective covers

3.1 Simple representations, finite length representations

Definition 3.1.1. A representation $V \in \text{Mod}_R^{\mathcal{D}}$ is **irreducible** (we also say **simple**) if it is nonzero and admits no nontrivial subrepresentation. A representation $V \in \text{Mod}_R^{\mathcal{D}}$ is **indecomposable** if it admits no nontrivial direct summand.

The notion of an irreducible representation of a category, being easy and natural, has been studied before. The definition appears already in [Ner91], for example.

We introduce notation that will be useful in several upcoming proofs.

Definition 3.1.2. Given $V \in \text{Mod}_R^{\mathcal{D}}$, write

$$\text{IrrQuot}(V) = \{V/W \mid W \subseteq V, \ell(V/W) = 1\} / \simeq$$

for the set of simple quotients of V up to isomorphism. Similarly, write

$$\text{IrrSubQuot}(V) = \{W'/W \mid W \subseteq W' \subseteq V, \ell(W'/W) = 1\} / \simeq$$

for the set of simple subquotients of V up to isomorphism.

Definition 3.1.3. A representation $V \in \text{Mod}_R^{\mathcal{D}}$ is said to have **finite length** if it admits a **composition series**

$$0 = V_0 \subsetneq V_1 \subsetneq V_2 \subsetneq \cdots \subsetneq V_\ell = V$$

with the property that each $S_i = V_i/V_{i-1}$ is simple. The S_1, S_2, \dots, S_ℓ are called the **composition factors** of V , and the number $\ell = \ell(V)$ is called the **length** of V . The Jordan-Hölder theorem guarantees that every composition series has the same composition factors up to reordering; see [KS06, p. 205] for details. We have $\text{IrrSubQuot}(V) = \{S_1, \dots, S_\ell\}$. If V does not have finite length, we set $\ell(V) = \infty$.

3.1.1 The Krull-Schmidt decomposition

We give the reassuring fact that every finitely generated projective can be written as a direct sum in an essentially unique way.

Theorem 3.1.4. *If R is Artinian and each $\text{Hom}_{\mathcal{D}}(x, y)$ is finite, every finitely generated projective representation $P \in \text{Mod}_R^{\mathcal{D}}$ decomposes as a direct sum of indecomposable projectives*

$$P \simeq P_1 \oplus \cdots \oplus P_k$$

where the summands are determined uniquely up to isomorphism and reordering.

Proof. Since any finitely generated projective is a summand of a finite sum of basic projectives by Proposition 2.5.2, the endomorphism ring of P is a corner ring of $\text{End}(d^\oplus)$ for some formal sum $d^\oplus \in \mathcal{D}^\oplus$. It follows that $\text{End}(P)$ is finitely generated as an R -module, and hence is itself Artinian. Since Artinian rings are semiperfect—their finitely generated modules have projective covers—we may apply the following result, which appears as [Kra14, Corollary 4.4].

Theorem 3.1.5. *An additive category is a Krull-Schmidt category if and only if it*

has split idempotents and the endomorphism ring of every object is semiperfect.

In Krause's terminology, this means that every finitely generated projective splits essentially uniquely as a sum of projectives with local endomorphism rings. In particular, the summands are indecomposable, and the result follows. \square

3.1.2 Length and the three functors f , f^* , and f_*

We must see how our existing toolkit of adjunctions interacts with the notion of length. Let $f : \mathcal{C} \rightarrow \mathcal{D}$ be a functor and $V \in \text{Mod}_R^{\mathcal{D}}$ a representation.

Proposition 3.1.6. *If f is essentially surjective on objects, then $\ell(V) \leq \ell(f^*V)$.*

Proof. Since f is essentially surjective on objects, we have $f^*V \simeq 0$ if and only if $V \simeq 0$. Any nontrivial inclusion $U \subsetneq U' \subseteq V$ has $U'/U \neq 0$, but since f^* is exact, $f^*U'/f^*U \simeq f^*(U'/U) \neq 0$ and so f^* takes nontrivial inclusions to nontrivial inclusions. \square

Proposition 3.1.7. *If f is full and essentially surjective on objects, then $\ell(f^*V) = \ell(V)$. In particular, V is simple if and only if f^*V is simple.*

Proof. By exactness, f^* takes subobjects of V to subobjects of f^*V . This assignment is injective as in the proof of Proposition 3.1.6. By [Ves08, Prop. A.2], any subrepresentation $U \subseteq f^*V$ has the form $U \simeq f^*W$ for some $W \in \text{Mod}_R^{\mathcal{D}}$, and so this assignment is surjective as well. \square

Proposition 3.1.8. *If f is full and faithful, then $\ell(f^*V) \leq \ell(V)$.*

Proof. Any nontrivial inclusion $U \subsetneq U'$ between subrepresentations of f^*V gives rise to an inclusion $f_*U \subseteq f_*U'$ by left exactness of f_* . To see that this inclusion remains nontrivial, note that applying f^* gives back the original inclusion $U \subsetneq U'$ because, using Proposition 2.6.3, the counit $\varepsilon' : f^*f_* \rightarrow 1$ is an isomorphism. \square

Proposition 3.1.9. *If f is full, then $\ell(f^*V) \leq \ell(V)$.*

Proof. Factor $f = i \circ g$ where g is full and essentially surjective on objects, and i is full and faithful, and apply Propositions 3.1.6 and 3.1.8. \square

The following fact—a simple extends to at most one simple—should be considered a basic fact of representation theory. Its first appearance seems to be in the master’s thesis of Yoshioka [Yos93]; our proof mirrors the English account in [Kos99, Theorem 2.8].

Proposition 3.1.10 (Yoshioka). *If $f: \mathcal{C} \rightarrow \mathcal{D}$ is full and faithful and $V, W \in \text{Mod}_R^{\mathcal{D}}$ are simples with $f^*V \simeq f^*W \neq 0$, then $V \simeq W$.*

Proof. Let $U \in \text{Mod}_R^{\mathcal{C}}$ be a representation isomorphic to both f^*V and f^*W . By additivity of the adjoint functors $f_{\dagger} \dashv f^*$:

$$\text{Hom}_R^{\mathcal{C}}(U, f^*V \oplus f^*W) \simeq \text{Hom}_R^{\mathcal{D}}(f_{\dagger}U, V \oplus W).$$

The left side of this isomorphism has an element $\varphi \oplus \psi$, the direct sum of two isomorphisms, corresponding on the other side to an element $\varphi^{\dagger} \oplus \psi^{\dagger}$. We claim that $\varphi^{\dagger} \oplus \psi^{\dagger}$, which is nonzero because $\varphi \oplus \psi \neq 0$, is also non-surjective, and so its image is simple since $\ell(V \oplus W) = 2$.

If $\varphi^{\dagger} \oplus \psi^{\dagger}$ were surjective, then its restriction $f^*\varphi^{\dagger} \oplus f^*\psi^{\dagger}$ would be as well, since f^* is exact. But f is full and faithful, so Proposition 2.6.3 implies that $f^*\varphi^{\dagger} \oplus f^*\psi^{\dagger}$ is surjective if and only if $\varphi \oplus \psi$ is surjective. And $\varphi \oplus \psi$ is not surjective since its source is simple and its target is a direct sum of two nonzero representations.

To conclude the argument, note that both maps φ^{\dagger} and ψ^{\dagger} induce nonzero maps from $\text{im}(\varphi^{\dagger} \oplus \psi^{\dagger})$ which must be isomorphisms since we have shown that both the source and target are simple. It follows that $V \simeq \text{im}(\varphi^{\dagger} \oplus \psi^{\dagger}) \simeq W$ as required. \square

We end this section with a method of building new indecomposables from old.

Proposition 3.1.11. *If f is full and faithful and V is indecomposable, then both $f!V$ and f_*V are indecomposable.*

Proof. Suppose $f_*V = A \oplus B$ decomposes as a direct sum. Applying the additive functor f^* to both sides, we see that $f^*f_*V = f^*A \oplus f^*B$. Since f is full and faithful, the counit provides an isomorphism $f^*f_*V \simeq V$. Since V is indecomposable, either $f^*A = 0$ or $f^*B = 0$. Applying Proposition 2.6.9, either $A = 0$ or $B = 0$. A dual argument proves that $f!V$ is indecomposable using Corollary 2.6.10. \square

3.2 Projective covers and injective hulls

The results of this section are standard. We refer the reader to [Kra14] for a detailed treatment.

Definition 3.2.1. A projective $P \in \text{Mod}_R^{\mathcal{D}}$ with a surjection $\varphi: P \rightarrow V$ is called a **projective cover** of $V \in \text{Mod}_R^{\mathcal{D}}$ if for all $X \in \text{Mod}_R^{\mathcal{D}}$ and $\psi \in \text{Hom}_R^{\mathcal{D}}(X, P)$, $\varphi \circ \psi$ is surjective if and only if ψ is surjective. Dually, an injective $I \in \text{Mod}_R^{\mathcal{D}}$ with an injection $\iota: V \hookrightarrow I$ is called an **injective hull** of V if for all $Y \in \text{Mod}_R^{\mathcal{D}}$ and $\kappa \in \text{Hom}_R^{\mathcal{D}}(I, Y)$, $\kappa \circ \iota$ is injective if and only if κ is injective.

Intuitively, a projective cover for a module is the “best approximating projective module.” A rephrasing of the definition is that any submodule $U \subseteq V$ with $U + \ker \varphi = V$ actually has $U = V$. In order to guarantee that every finitely generated representation has a projective cover, we usually assume that R is Artinian and that the hom-sets of \mathcal{D} are finite.

Proposition 3.2.2. *Suppose R is Artinian and the hom-sets of \mathcal{D} are finite. If V is simple, then its projective cover P is indecomposable and satisfies $\text{IrrQuot}(P) = \{V\}$. Conversely, every finitely generated indecomposable projective is a projective cover for some simple which is its unique simple quotient and appears as a composition factor of any nonzero quotient of V .*

Proof. This follows from [Kra14, Lemma 3.6]. □

Definition 3.2.3. A **cover-hull description of a simple** V is a composite map

$$\left(P \xrightarrow{p} V \xrightarrow{i} I \right)$$

from an indecomposable projective to an indecomposable injective so $V \simeq \text{im}(i \circ p)$.

3.3 The intermediate extension functor $f_{!*}$

Recall that a functor $f : \mathcal{C} \rightarrow \mathcal{D}$ induces two adjunctions: left and right Kan extension along f . In the event that f is full and faithful, we have seen that $f_!$ interacts well with indecomposable projectives and f_* interacts well with indecomposable injectives. To handle simples, we introduce a new functor. In what follows, write $\eta : 1 \rightarrow f^*f_!$ for the unit of the left Kan extension adjunction, and $\eta' : 1 \rightarrow f_*f^*$ for the unit of the right Kan extension adjunction.

Definition 3.3.1. If $f : \mathcal{C} \rightarrow \mathcal{D}$ is full and faithful, (and so the unit for left Kan extension has an inverse $\eta^{-1} : f^*f_! \rightarrow 1$), then define the the **intermediate extension functor** $f_{!*} : \text{Mod}_R^{\mathcal{C}} \rightarrow \text{Mod}_R^{\mathcal{D}}$ as the image of the composite

$$f_! \xrightarrow{\eta'f_!} f_*f^*f_! \xrightarrow{f_*\eta^{-1}} f_*.$$

Proposition 3.3.2. *If $f : \mathcal{C} \rightarrow \mathcal{D}$ is full and faithful, and $V \in \text{Mod}_R^{\mathcal{C}}$, then*

$$f^*f_{!*}V \simeq V.$$

Proof. This is an immediate consequence of Proposition 2.6.3. □

Remark 3.3.3. The intermediate extension functor appears in [BBD82].

The next theorem lets us upgrade cover-hull descriptions of simples for a full subcategory to cover-hull descriptions of simples for the ambient category. The most important application will be when the full subcategory has a single object. In this case, the Theorem 3.3.4 upgrades representation theory of an endomorphism algebra to representation theory of a category.

Theorem 3.3.4. *If $f: \mathcal{C} \rightarrow \mathcal{D}$ is full and faithful, and $(P \xrightarrow{p} V \xrightarrow{i} I)$ is a cover-hull description of a simple, then f_*V is simple with cover-hull description*

$$\left(f_!P \xrightarrow{p'} f_*V \xrightarrow{i'} f_*I \right),$$

satisfying $i' \circ p' = (f_*i) \circ (f_*\eta^{-1}) \circ (\eta'f) \circ (f_!p)$.

Proof. The representation $f_!P$ is projective since $f_!$ takes projectives to projectives. It is indecomposable by Proposition 3.1.11. Similarly, f_*I is an indecomposable injective.

Since f_*i is an injection (by left-exactness of f_*) and $f_!p$ is a surjection (similarly), we see that the image of the composite $(f_*i) \circ (f_*\eta^{-1}) \circ (\eta'f) \circ (f_!p)$ matches the image of the middle two maps, which is exactly the definition of f_*V . Letting p' and i' form an epi-mono factorization of this long composite, it remains only to show that f_*V is simple.

Let $U \subseteq f_*V$ with quotient $Q = f_*V/U$. We have a short exact sequence

$$0 \rightarrow U \rightarrow f_*V \rightarrow Q \rightarrow 0,$$

which, after applying the exact functor f^* , becomes

$$0 \rightarrow f^*U \rightarrow f^*f_*V \rightarrow f^*Q \rightarrow 0.$$

The middle representation is simple by Proposition 3.3.2, so one of the outer two representations must be zero. But $U \subseteq f_*V \subseteq f_*V$, and similarly Q is a quotient of

fV , so either Proposition 2.6.9 or its dual Corollary 2.6.10 applies, and f_*V is simple as required. \square

3.3.1 Filtered categories, associated graded categories, and the classification of irreducible representations

Definition 3.3.5. A category \mathcal{D} is called an \mathbb{N} -**filtered category** if it comes equipped with a **degree function** $\deg : \text{Ob}(\mathcal{D}) \rightarrow \mathbb{N}$. For any $k \in \mathbb{N}$, define the subcategory $\mathcal{D}_{\leq k} \subseteq \mathcal{D}$ which is full on the objects of degree at most k . Similarly, define the subcategory $\mathcal{D}_{< k}$ full on the objects of degree less than k , and \mathcal{D}_k full on the subobjects of degree equal to k .

Definition 3.3.6. If \mathcal{D} is a filtered category, the **associated graded** R -linear category $\mathcal{A} = \sqcup_k \mathcal{A}_k$ is the universal quotient category of \mathcal{D} sending to zero all maps that factor through lower degree. Concretely, each category \mathcal{A}_k is the quotient of \mathcal{D}_k by the two-sided ideal generated by maps that factor through an object of degree less than k . For each k , write $i_k : \mathcal{D}_k \hookrightarrow \mathcal{D}$ for the inclusion, and $p_k : \mathcal{D}_k \twoheadrightarrow \mathcal{A}_k$ for the projection.

The next theorem shows how a choice of \mathbb{N} -filtration for \mathcal{D} gives a canonical graded bijection between the simples of \mathcal{A} and the simples of \mathcal{D} . Since the associated graded category \mathcal{A} is often much more tractable, Theorem 3.3.7 helps classify irreducible representations of \mathcal{D} .

Theorem 3.3.7. *If \mathcal{D} is an \mathbb{N} -filtered category with associated graded $\mathcal{A} = \sqcup_k \mathcal{A}_k$, then every simple $V \in \text{Mod}_R^{\mathcal{D}}$ can be written $V = (i_k)_{!*} \circ (p_k)^* W$ for some $k \in \mathbb{N}$ and some unique-up-to-isomorphism simple $W \in \text{Mod}_R^{\mathcal{A}_k}$.*

Proof. If $W \in \text{Mod}_R^{\mathcal{A}_k}$ is simple, then $p_k W$ is also simple since p_k is full and essentially surjective on objects, using Corollary 3.1.7. By Proposition 3.3.4, $(i_k)_{!*}$ takes simples to simples, and so any V constructed in this way is simple.

Suppose that V is simple. Since $V \neq 0$, there must be some object $d \in \mathcal{D}$ with $Vd \neq 0$; let $k \in \mathbb{N}$ be the minimal degree of such an object. So $(i_k)^*V \neq 0$, and hence by Proposition 3.1.9, $(i_k)^*V$ is simple. Every morphism of \mathcal{D} that factors through an object of degree less than k acts by zero on $(i_k)^*V$ since k is minimal. It follows that the action of \mathcal{D}_k on $(i_k)^*V$ descends to the quotient category \mathcal{A}_k .

By Proposition 3.3.2, $i^* \circ (i_k)_{!*} \circ (p_k)^*W \simeq (p_k)^*W$. Since $(p_k)_{!}(p_k)^*W \simeq W$ by Proposition 2.6.3, we may recover W from $(i_k)_{!*} \circ (p_k)^*W$ by applying $(p_k)_{!} \circ (i^*)$. Theorem 3.1.10 (Yoshioka's theorem) says that two simples are isomorphic if and only if their restrictions to some full subcategory are nonzero and isomorphic. Since i is the inclusion of a full subcategory, and we have produced at least one irreducible representation with the correct restriction, this representation is unique up to isomorphism and the constructions provide a bijection. \square

3.3.1.1 A corollary in the representation theory of algebras

If $d_1, \dots, d_n \in \text{Ob}(\mathcal{D})$ is a sequence of objects with each d_{i-1} a retract of d_i (meaning that the identity map on d_{i-1} factors through d_i), write A_i for the quotient algebra $\text{End}(d_i)/(\text{morphisms factoring through } d_{i-1})$.

Corollary 3.3.8. *The number of irreducible representations of $\text{End}(d_n)$ is given by the sum*

$$\sum_{i=1}^n \#\{\text{irreducible representations of } A_i\}.$$

Example 3.3.9. The category of finite sets satisfies these hypotheses, setting $d_i = \{1, 2, \dots, i\}$. The algebras A_i are group algebras for the symmetric group. It follows that the irreducible representations of the **full transformation monoid** $\text{End}(d_n)$ are in bijection with the irreducible representations of the symmetric groups of degree $1, 2, \dots, n$. The result holds over any field. This result is well-known, since it follows from the general representation theory of semigroups; see, for example [CP61, Chapter 5].

3.3.2 Computing the dimensions of the irreducible representation i_*p^*W

We give the explicit method by which the matrices of an intermediate extension from an endomorphism algebra may be computed. Let $i: \text{End}(d) \rightarrow \mathcal{D}$ be the inclusion. We assume that R is a field.

Let $V = p^*W$ be an irreducible right $\text{End}(d)$ -module, and $\varphi: \text{End}(d) \rightarrow R$ a matrix entry of V . For any $x \in \mathcal{D}$, let H_x be the $\text{Hom}(d, x) \times \text{Hom}(x, d)$ matrix with (f, g) -entry $\varphi(g \circ f)$.

Theorem 3.3.10. *We have*

$$\dim_R (i_*p^*W)x = \text{rank}_R H_x.$$

Proof. In fact, the row span of H_x coincides with the image of the natural map $i_*p^*W \rightarrow i_*p^*W$. The proof is just an unraveling of the construction of the left and right Kan extensions already described in Proposition 2.6.7 and Corollary 2.6.8. \square

This sort of construction is called a “monomial representation” in the representation theory of monoids. See [Ste, §5.5] for this perspective.

CHAPTER IV

Imaginations, homological moduli, and categories of dimension zero

4.1 The object preorders \leq_d


We introduce the technical heart of the thesis: preorders on the objects of \mathcal{D} that control the subrepresentations of a finitely generated representation.

Definition 4.1.1. If $d, x, y \in \text{Ob}(\mathcal{D})$, say $x \leq_d y$ if for all representations V generated by vectors in Vd , every subrepresentation $U \subseteq V$ generated by vectors in Ux is also generated by vectors in Uy .

In a moment, we give several conditions equivalent to $x \leq_d y$, each suited to a different purpose; let us develop a bit more notation before stating the result as Proposition 4.1.3.

Let $S(x, y)$ denote the set of self-maps of x that factor through y . For each $s \in S(x, y)$, define a square 0-1-matrix M_s whose rows and columns are indexed by $\text{Hom}_{\mathcal{D}}(d, x)$, putting a 1 in position (f, g) whenever $s \circ f = g$. Equivalently, the entries of M_s

record the commutativity of the following diagram:



Definition 4.1.2. Given representations $P, Q \in \mathbf{Mod}_R^{\mathcal{D}}$, the **imagination of P in Q** is the set of submodules of Q that admit surjections from a finite direct sum of copies of P :

$$\text{Imag}(P \rightarrow Q) = \{\text{im } \psi \mid \text{for } k \in \mathbb{N} \text{ and } \psi: P^{\oplus k} \rightarrow Q\}.$$

Finally, recall that the basic projective P^d associated to $d \in \text{Ob}(\mathcal{D})$ is the representation $x \mapsto R \cdot \text{Hom}_{\mathcal{D}}(d, x)$ where a map $h: x \rightarrow y$ acts by $\varphi \mapsto h \circ \varphi$.

Proposition 4.1.3. *If $\text{Hom}_{\mathcal{D}}(d, x)$ is finite, the following are equivalent to the statement $x \leq_d y$:*

1. *The condition given in Definition 4.1.1;*
2. *The identity matrix is in the R -span of the matrices M_s ;*
3. *There exists an invertible matrix is in the R -span of the matrices M_s ;*
4. $\text{Imag}(P^x \rightarrow P^d) \subseteq \text{Imag}(P^y \rightarrow P^d)$;
5. *There are matrices $\alpha \in \text{Mat}_R^{\mathcal{D}}(x \rightarrow y^{\oplus k})$, $\beta \in \text{Mat}_R^{\mathcal{D}}(y^{\oplus k} \rightarrow x)$ so that for any matrix $f \in \text{Mat}_R^{\mathcal{D}}(d \rightarrow x)$, we have $f \cdot \alpha \cdot \beta = f$;*
6. *A condition to be given soon as Proposition 4.1.5;*
7. *A condition relying on further assumptions to be given later as Proposition 4.3.4.*

Remark 4.1.4. The second condition gives an algorithmic test to check if $x \leq_d y$. The third condition gives a probabilistic test (since a random linear combination will be invertible). From the fourth condition, we see that \leq_d is reflexive and transitive, and so defines a preorder.

4.1.1 Useful facts about the preorders \leq_d

To shorten proofs in this section, we introduce concise notation for useful submodules of the R -linearized hom-bimodules coming from \mathcal{D} . If $d_1, d_2, \dots, d_k \in \mathcal{D}$ are objects, write (d_1, d_2, \dots, d_k) for the R -submodule of $\text{Mat}_R^{\mathcal{D}}(d_1 \rightarrow d_k)$ spanned by all maps $d_1 \rightarrow d_k$ that can be written as a composite $d_1 \rightarrow d_2 \rightarrow \dots \rightarrow d_k$. For example, $(x, y) = \text{Mat}_R^{\mathcal{D}}(x \rightarrow y)$, and $(x, y, z) = \text{Mat}_R^{\mathcal{D}}(x \rightarrow y) \cdot \text{Mat}_R^{\mathcal{D}}(y \rightarrow z)$. As we already do for matrices, write \cdot for the R -bilinearized composition law, so $(x, y) \cdot (y, z) = (x, y, z)$, for example. If $N \subseteq (x, y)$ is a submodule, define left and right annihilator submodules of N

$$N_d = \{f \in N : f \cdot (y, d) = 0\}$$

$${}_d N = \{f \in N : (d, x) \cdot f = 0\}.$$

Proposition 4.1.5. *We have $x \leq_d y$ if and only if $(x, x) = {}_d(x, x) + (x, y, x)$.*

Proof of Propositions 4.1.3 and 4.1.5. We prove that the first six conditions of Proposition 4.1.3 are equivalent; Proposition 4.1.5 appears in this list as condition (6).

(6) \iff (2): The matrices M_s with $s \in S(x, y)$ give the action of $(x, y, x) \subseteq (x, x)$ on (d, x) by postcomposition. Their span is a (possibly unitless) R -algebra isomorphic to $(x, y, x)/({}_d(x, x) \cap (x, y, x))$, by construction. This algebra has a unit exactly when $1 \in {}_d(x, x) + (x, y, x)$, which is condition (2); we are done since this sum is a two-sided ideal of (x, x) containing 1.

(2) \iff (3): The Cayley-Hamilton theorem gives the identity matrix in the R -span of the powers of any invertible matrix.

(5) \iff (6): We have $1 = \delta + \omega$ with $\delta \in {}_d(x, x)$ and $\omega \in (x, y, x)$ if and only if $f \cdot \omega = f$ for every f . Any element of (x, y, x) expands as a finite linear combination of morphisms factoring through y ; this decomposition translates directly into a pair of matrices.

(5) \implies (4): Every map $\psi: (P^x)^{\oplus k} \rightarrow P^d$ satisfies $\psi = \psi \circ \alpha^{\natural} \circ \beta^{\natural}$ by Corollary 2.4.3, and so $\text{im}(\psi \circ \alpha^{\natural}) = \text{im} \psi$.

(4) \implies (6): For each $f \in \text{Mat}_R^D(d \rightarrow x)$ we get a map $f^{\natural}: P^x \rightarrow P^d$. Since P^y is at least as imaginative as P^x is in P^d , we have a surjection $\varphi_f: (P^y)^{\oplus k} \rightarrow \text{im}(f^{\natural})$. By the projectivity of P^x , we may lift the map f^{\natural} along this surjection using some map $\alpha_f: P^x \rightarrow (P^y)^{\oplus k}$. Similarly, the map φ_f lifts along the surjection $f^{\natural}: P^x \rightarrow P^d$ using some map $\beta_f: (P^y)^{\oplus k} \rightarrow P^x$. The composite map satisfies $f \circ \beta_f \circ \alpha_f = f$, and so for each vector $f \in (d, x)$, the ideal (x, y, x) contains an element ω_f with the property that $f \cdot \omega_f = f$. If the action of (x, y, x) on (d, x) is to be unitless, (d, x) must have a subquotient for which the action is identically zero. This is impossible since a generator f for the subquotient is fixed by ω_f . It follows that there is some element $\omega \in (x, y, x)$ that acts on (d, x) by the identity, and so $1 = \delta + \omega$ for some $\delta \in {}_d(x, x)$.

(1) \iff (4): A representation V being generated by k vectors in Vd is equivalent to the existence of a surjection $(P^d)^{\oplus k} \twoheadrightarrow V$. It follows that any $U \subseteq V$ is generated by finitely many vectors in Ux if and only if it is in the imagination of P^x .

Writing V as an increasing union of finitely generated subrepresentations gives the result. \square

Proposition 4.1.6. *We have $x \leq_x y$ exactly when x is a retract of y .*

Proof. If x is a retract of y , then $(x, y, x) \subseteq (x, x)$ is a two-sided ideal containing 1, and so $(x, y, x) = (x, x)$. Similarly, if $x \leq_x y$, then $(x, y, x) = (x, x)$ since ${}_x(x, x) = 0$. \square

Definition 4.1.7. Write $x \leq^d y$ whenever $x \leq_d y$ in the opposite category \mathcal{D}^{op} .

Proposition 4.1.8. *If $c \leq^x d$, then ${}_d(x, y) \subseteq {}_c(x, y)$.*

Proof. By Proposition 4.1.5, $(c, c)_x + (c, d, c) = (c, c)$. Let $f : x \rightarrow y$, and suppose $(d, x) \cdot f = 0$. Then

$$\begin{aligned}
(c, x) \cdot f &= (c, c) \cdot (c, x) \cdot f \\
&= [(c, c)_x + (c, d, c)] \cdot (c, x) \cdot f \\
&= (c, c)_x \cdot (c, x) \cdot f + (c, d, c) \cdot (c, x) \cdot f \\
&= 0 + (c, d, c, x) \cdot f \\
&= (c, d) \cdot (d, c, x) \cdot f \\
&\subseteq (c, d) \cdot (d, x) \cdot f \\
&\subseteq 0,
\end{aligned}$$

and so $f \in {}_c(x, y)$ as required. \square

Lemma 4.1.9. *If $x \leq_d y$ and $c \leq^x d$, then $x \leq_c y$.*

Proof. By Proposition 4.1.8, ${}_d(x, x) \subseteq {}_c(x, x)$, and by Proposition 4.1.3, ${}_d(x, x) + (x, y, x) = (x, x)$. So

$$(x, x) = {}_d(x, x) + (x, y, x) \subseteq {}_c(x, x) + (x, y, x) \subseteq (x, x)$$

and $x \leq_c y$ as required. \square

The following result of Andrew Gitlin [Git15] provides a useful combinatorial strategy for proving $x \leq_d y$.

Lemma 4.1.10 (Gitlin's trick). *If $\text{Hom}_{\mathcal{D}}(d, x)$ is finite and there exists a partial order \preceq on $\text{Hom}_{\mathcal{D}}(d, x)$ and a function $s : (d, x) \rightarrow (x, y, x)$ such that $s(f) \circ f = f$ and $h \preceq s(f) \circ h$ for all $f, h \in \text{Hom}_{\mathcal{D}}(d, x)$, then $x \leq_d y$ over any infinite field.*

Proposition 4.1.11. *If $x, y \in \text{Ob}(\mathcal{D})$ satisfy $x \leq_y y$ and p denotes the projection*

$$p: R \cdot \text{End}(x) \rightarrow R \cdot \text{End}(x)/(x, y, x),$$

then the induced left Kan extension $p_!$ is exact.

Proof. We use the statement $x \leq_y y$ to produce an idempotent π , and show that the functor $p_!$ coincides with taking the image of π , proving exactness. By Proposition 4.1.3, we have

$$(x, x) = {}_y(x, x) + (x, y, x)$$

By the second isomorphism theorem,

$$\frac{(x, x)}{(x, y, x)} \simeq \frac{{}_y(x, x)}{{}_y(x, x) \cap (x, y, x)},$$

and so there is an element $\delta \in {}_y(x, x)$ that acts as an identity for this second algebra. Setting $B = (x, y, x)$, choose some vector space $A \subseteq \text{End}(x)$ so that $\text{End}(x) = A \oplus B$. Since δ annihilates every map from y , we see that $B \cdot \delta = 0$, and so a block matrix for δ has the form $\begin{bmatrix} 1 & * \\ 0 & 0 \end{bmatrix}$, from which it follows that δ is idempotent. We claim that any right $\text{End}(x)$ -module V satisfies $p_!V \simeq \text{im}(V \xrightarrow{1-\delta} V)$. Indeed, we have an isomorphism $p_!V \simeq V/(V \cdot (x, y, x)) \simeq V/(V \cdot \begin{bmatrix} 0 & * \\ 0 & 0 \end{bmatrix}) \simeq \ker(V \xrightarrow{\delta} V)$ by direct multiplication of block matrices. □

Corollary 4.1.12. *If $x \leq_y y$ then p^* takes injectives to injectives.*

Proof. The functor p^* admits an exact left adjoint. □

4.2 Upper bounds for the preorder \leq_d

We begin with a technical lemma that compiles a finite list of upper bounds into a convenient pair of matrices. Using the lemma, we are able to prove Theorem 4.2.2 explaining the usefulness of upper bounds in the preorder \leq_d . Finally, in Definition 4.2.4, we introduce terminology for a complete system of joint upper bounds for the preorders \leq_d .

Lemma 4.2.1. *Suppose $V \in \text{Mod}_R^{\mathcal{D}}$ is finitely generated in degrees $d_1, \dots, d_k \in \text{Ob}(\mathcal{D})$, and suppose we have a system of upper bounds on some $x \in \text{Ob}(\mathcal{D})$:*

$$x \leq_{d_1} m_1 \quad x \leq_{d_2} m_2 \quad \cdots \quad x \leq_{d_k} m_k$$

for $m_1, \dots, m_k \in \text{Ob}(\mathcal{D})$. Then there exists some formal sum $m^\oplus \in \mathcal{D}^\oplus$ with summands drawn (possibly with repetition) from the list m_1, \dots, m_k , and a pair of matrices

$$\alpha \in \text{Mat}_R^{\mathcal{D}}(x \rightarrow m^\oplus) \quad \beta \in \text{Mat}_R^{\mathcal{D}}(m^\oplus \rightarrow x)$$

so that $V(\alpha \cdot \beta) = 1_{Vx}$.

Proof. Let $\text{Mat}_R^{\mathcal{D}}(x \rightarrow m_\bullet \rightarrow x)$ denote the R -submodule of $\text{Mat}_R^{\mathcal{D}}(x \rightarrow x)$ spanned by all matrices that factor through one of the m_i . Since any matrix $\omega \in \text{Mat}_R^{\mathcal{D}}(x \rightarrow m_\bullet \rightarrow x)$ mentions only finitely many monomials, it is of the form $\omega = \alpha \cdot \beta$ for some pair of matrices

$$\alpha \in \text{Mat}_R^{\mathcal{D}}(x \rightarrow m^\oplus) \quad \beta \in \text{Mat}_R^{\mathcal{D}}(m^\oplus \rightarrow x)$$

where m^\oplus has summands drawn from the list m_1, \dots, m_k . If we find such a matrix ω with the property that $V(\omega) = 1_{Vx}$, then we are done.

For each $i \in \{1, \dots, k\}$, the corresponding inequality $x \leq_{d_i} m_i$ provides an element $\omega_i \in \text{Mat}_R^{\mathcal{D}}(x \rightarrow m_i \rightarrow x)$ with the property that, for any matrix $\gamma \in \text{Mat}_R^{\mathcal{D}}(d_i \rightarrow x)$, $\gamma \cdot \omega_i = \gamma$. We claim

$$\omega = 1 - (1 - \omega_1) \cdot (1 - \omega_2) \cdot \dots \cdot (1 - \omega_k)$$

satisfies $V(\omega) = 1_{Vx}$, completing the proof. (Note that the leading 1's cancel and so $\omega \in \text{Mat}_R^{\mathcal{D}}(x \rightarrow m_{\bullet} \rightarrow x)$.)

By Proposition 2.4.4, every vector $u \in Vx$ is of the form $u = v \cdot V(\gamma)$ for some $v \in Vd^{\oplus}$, and $\gamma \in \text{Mat}_R^{\mathcal{D}}(d^{\oplus} \rightarrow x)$, where the formal sum d^{\oplus} has summands drawn (possibly with repetition) from the finite list d_1, \dots, d_k . To show that $V(\omega)$ fixes $u \in Vx$, it suffices to show that $\gamma \cdot \omega = \gamma$, since then

$$u \cdot V(\omega) = (v \cdot V(\gamma)) \cdot V(\omega) = v \cdot V(\gamma \cdot \omega) = v \cdot V(\gamma) = u.$$

We show that for all $\gamma \in \text{Mat}_R^{\mathcal{D}}(d^{\oplus} \rightarrow x)$, we have $\gamma \cdot \omega = \gamma$ for each $\gamma' \in \text{Mat}_R^{\mathcal{D}}(d_i \rightarrow x)$, a matrix entry appearing in γ . Compute

$$\begin{aligned} \gamma' \cdot \omega &= \gamma' \cdot (1 - (1 - \omega_1) \cdot (1 - \omega_2) \cdot \dots \cdot (1 - \omega_k)) \\ &= \gamma' - \gamma' \cdot (1 - \omega_1) \cdot (1 - \omega_2) \cdot \dots \cdot (1 - \omega_k) \\ &= \gamma' - [\gamma' \cdot (1 - \omega_1) \cdot \dots \cdot (1 - \omega_{i-1})] \cdot (1 - \omega_i) \cdot \dots \cdot (1 - \omega_k) \\ &= \gamma' \end{aligned}$$

since ω_i fixes the matrix $[\gamma' \cdot (1 - \omega_1) \cdot \dots \cdot (1 - \omega_{i-1})] \in \text{Mat}_R^{\mathcal{D}}(d_i \rightarrow x)$. It follows that ω fixes every entry of γ , and so $\gamma \cdot \omega = \gamma$ as required. \square

Theorem 4.2.2. *Suppose $W \in \text{Mod}_R^{\mathcal{D}}$ is generated by vectors in Wd_1, \dots, Wd_k with $d_i \in \text{Ob}(\mathcal{D})$. If $i: \mathcal{M} \subseteq \mathcal{D}$ is a full subcategory with the property that for all $x \in \text{Ob}(\mathcal{D})$*

there exist $m_1, \dots, m_k \in \text{Ob}(\mathcal{M})$ such that $x \leq_{d_i} m_i$, then i^* induces isomorphisms

$$\text{Hom}_R^{\mathcal{D}}(V, W) \xrightarrow{\sim} \text{Hom}_R^{\mathcal{M}}(i^*V, i^*W)$$

for all $V \in \text{Mod}_R^{\mathcal{D}}$.

Proof. The map $\text{Hom}_R^{\mathcal{D}}(V, W) \rightarrow \text{Hom}_R^{\mathcal{M}}(i^*V, i^*W)$ is easily seen to be surjective since i is full and faithful; indeed, Proposition 2.6.3 gives that $i_*\varphi \mapsto \varphi$ for any $\varphi \in \text{Hom}_R^{\mathcal{D}}(i^*V, i^*W)$.

To prove injectivity, we show that any map $\psi \in \text{Hom}_R^{\mathcal{D}}(V, W)$ can be recovered from its restriction $i^*\psi$. More specifically, for any $d \in \mathcal{D}$, we give a formula for the component ψ_d depending only on the components ψ_m with $m \in \mathcal{M}$. Applying Lemma 4.2.1 to the representation W , obtain a formal sum $m^\oplus \in \mathcal{M}^\oplus$ and two matrices

$$\alpha \in \text{Mat}_R^{\mathcal{D}}(d \rightarrow m^\oplus) \quad \beta \in \text{Mat}_R^{\mathcal{D}}(m^\oplus \rightarrow d)$$

so that $W(\alpha) \cdot W(\beta) = 1_{Wd}$. Since ψ is a map of representations, components of ψ interpose this factorization:

$$\begin{array}{ccccc} Vd & \xrightarrow{V(\alpha)} & Vm^\oplus & \xrightarrow{V(\beta)} & Vd \\ \psi_d \downarrow & & \oplus \psi_m \downarrow & & \downarrow \psi_d \\ Wd & \xrightarrow{W(\alpha)} & Wm^\oplus & \xrightarrow{W(\beta)} & Wd. \\ & & \text{W(1)} & & \end{array}$$

We deduce the formula $\psi_d = W(\beta) \circ W(\alpha) \circ \psi_d = W(\beta) \circ (\oplus \psi_m) \circ V(\alpha)$, and note that it references only components at objects of \mathcal{M} as required. \square

Corollary 4.2.3. *If $W \in \text{Mod}_R^{\mathcal{D}}$ and $i: \mathcal{M} \subseteq \mathcal{D}$ satisfy the hypotheses of Theorem 4.2.2, any subrepresentation $S \subseteq W$ is determined by its restriction $i^*S \subseteq i^*W$. In particular, $\ell(W) = \ell(i^*W)$.*

We make a definition collecting a system of joint upper bounds (as in the hypotheses of Lemma 4.2.1) into a single object that may or may not exist.

Definition 4.2.4. A **homological modulus** for a category \mathcal{D} over a ring R is a function

$$\mu : \text{Ob}(\mathcal{D}) \longrightarrow \{\text{finite subsets of } \text{Ob}(\mathcal{D})\}$$

such that for every $d, x \in \text{Ob}(\mathcal{D})$, there exists $y \in \mu(d)$ so that $x \leq_d y$. By convention, we extend μ to finite subsets of $\text{Ob}(\mathcal{D})$ by the formula

$$\mu(S) = \bigcup_{s \in S} \mu(s).$$

Here is an example of the sort of result that becomes easy to state with Definition 4.2.4 in hand.

Observation 4.2.5. *If \mathcal{D} admits a homological modulus μ over R , then any representation $V \in \text{Mod}_R^{\mathcal{D}}$ generated in degrees $D \subseteq \text{Ob}(\mathcal{D})$ has a projective resolution of the form*

$$\cdots \longrightarrow P_2 \longrightarrow P_1 \longrightarrow P_0 \longrightarrow V \longrightarrow 0$$

where each P_i is a direct sum of basic projectives P^d with $d \in \mu^i(D)$ (here, μ^i denotes the i -fold composition of μ).

Proof. Follows directly from Definitions 4.1.1 and 4.2.4 and Proposition 4.1.3. □

Corollary 4.2.6 (Homological stability from a homological modulus). *If a representation $V \in \text{Mod}_R^{\mathcal{D}}$ is generated in degrees $D \subseteq \text{Ob}(\mathcal{D})$, then $\text{Ext}^k(V, W) \simeq \text{Ext}^k(i^*V, i^*W)$, for i the inclusion of the full subcategory on the objects $\mu^{k+1}(D)$.*

Proof. Follows from Theorem 4.2.2 and the Yoneda Lemma 2.4.2. □

4.3 Dimension zero categories

Definition 4.3.1. Say a category \mathcal{D} is **dimension zero over a commutative ring** R if every finitely generated $V \in \text{Mod}_R^{\mathcal{D}}$ has finite length.

Theorem 4.3.2. *A category \mathcal{D} is dimension zero over a commutative ring R if and only if R is Artinian, each hom-set $\text{Hom}_{\mathcal{D}}(x, y)$ is finite, and \mathcal{D} admits a homological modulus over R .*

Proof. Let us assume that R is Artinian, that \mathcal{D} has finite hom-sets, and μ is a homological modulus for \mathcal{D} over R . Suppose $V \in \text{Mod}_R^{\mathcal{D}}$ is finitely generated in degrees d_1, \dots, d_k . Let $i: \mathcal{M} \subseteq \mathcal{D}$ be the full subcategory on the finite set of objects $\mu(\{d_1, \dots, d_k\})$. By Corollary 4.2.3, $\ell(V) = \ell(i^*V)$. Since \mathcal{D} has finite hom-sets and V is finitely generated, each Vd is a finitely generated R -module. Since R is Artinian, the finite sum $\bigoplus_{m \in \mathcal{M}} Vm$ is finite length. It follows that i^*V is finite length, and hence V as well.

In the other direction, suppose every finitely generated representation $V \in \text{Mod}_R^{\mathcal{D}}$ is finite length. We rely on the following result of Zel'manov (for an English account, see Okniński's book [Okn91, p. 172, Theorem 23]):

Theorem 4.3.3 ([Zel77]). *If R is a commutative ring and M is a monoid, then the monoid ring RM is Artinian exactly when R is Artinian and M is finite.*

We first show that every endomorphism monoid of \mathcal{D} is finite, and then that every hom-set is finite as well. If $i: \text{End}(d) \subseteq \mathcal{D}$ is the inclusion of the full subcategory on the single object $d \in \text{Ob}(\mathcal{D})$, then the length of $i^*P^d = R\text{End}(d)$ as a module over itself is bounded above by the length of P^d using Proposition 3.1.9. It follows that $R\text{End}(d)$ is Artinian, and so R is Artinian and each endomorphism monoid is finite by the theorem of Zel'manov. Similarly, $i^*P^{d'} = R \cdot \text{Hom}(d', d)$ is finite length as a right $R\text{End}(d)$ -module, and so $\text{Hom}(d', d)$ is finite as well.

To prove that \mathcal{D} admits a homological modulus over R , we need the following.

Proposition 4.3.4. *If $d, x, y \in \text{Ob}(\mathcal{D})$ with either P^d or P^x finite length, we have $x \leq_d y$ if and only if*

$$\text{IrrQuot}(P^x) \cap \text{IrrSubQuot}(P^d) \subseteq \text{IrrQuot}(P^y) \cap \text{IrrSubQuot}(P^d).$$

Assuming Proposition 4.3.4, we produce a homological modulus as follows. If a collection of simples appears as a subset of $\text{IrrQuot}(P^z)$ for some $z \in \text{Ob}(\mathcal{D})$, we call that collection **attainable**, and we say that d **witnesses attainability**.

Fix $d \in \text{Ob}(\mathcal{D})$. For every attainable subset $Q \subseteq \text{IrrSubQuot}(P^d)$, let y_Q witness attainability:

$$Q \subseteq \text{IrrQuot}(P^{y_Q}).$$

We claim that we may take $\mu(d) = \{y_Q \mid Q \text{ is an attainable subset of } \text{IrrSubQuot}(P^d)\}$. Specifically, given $x \in \text{Ob}(\mathcal{D})$, Proposition 4.3.4 gives that $x \leq_d y_Q$ with $Q = \text{IrrQuot}(P^x) \cap \text{IrrSubQuot}(P^d)$. \square

Proof of Proposition 4.3.4. Suppose $x \leq_d y$, and let

$$S \in \text{IrrQuot}(P^x) \cap \text{IrrSubQuot}(P^d).$$

Choose $A \subset B \subseteq P^d$ so that $S \simeq B/A$. Since S is a quotient of P^x and P^x is projective, B is also a quotient of P^x . So $B \in \text{Imag}(P^x \rightarrow P^d)$. By Proposition 4.1.3, $B \in \text{Imag}(P^y \rightarrow P^d)$ as well, and so S is also a simple quotient of P^y .

In the other direction, if $B \in \text{Imag}(P^x \rightarrow P^d)$, then B is finite length since at least one of P^d and P^x is finite length. Let $A \subsetneq B$ be a maximal subrepresentation. We have that the quotient $S = B/A$ is simple and that $\ell(A) + 1 = \ell(B)$. Assume by way of induction on the length that $A \in \text{Imag}(P^y \rightarrow P^d)$; we must show that

$B \in \text{Imag}(P^y \rightarrow P^d)$ as well. Since S a quotient of P^x and a subquotient of P^d ,

$$S \in \text{IrrQuot}(P^x) \cap \text{IrrSubQuot}(P^d),$$

so S is also a quotient of P^y since we have assumed that $\text{IrrQuot}(P^x) \cap \text{IrrSubQuot}(P^d) \subseteq \text{IrrQuot}(P^y) \cap \text{IrrSubQuot}(P^d)$. The outer representations of the short exact sequence

$$0 \longrightarrow A \longrightarrow B \longrightarrow S \longrightarrow 0$$

have been established as elements of $\text{Imag}(P^y \rightarrow P^d)$. Using the projectivity of P^y , we may lift any surjection $P^y \twoheadrightarrow S$ to B . This proves that $B \in \text{Imag}(P^y \rightarrow P^d)$ as required. \square

4.3.1 Computing simple multiplicities from a presentation matrix

If \mathcal{D} is dimension zero over R , then any presentation matrix

$$M \in \text{Mat}_R^{\mathcal{D}}(x^{\oplus} \rightarrow y^{\oplus})$$

gives rise to some finite length representation $V_M \in \text{Mod}_R^{\mathcal{D}}$. If we have already computed the indecomposable injective representations of \mathcal{D} , the following result gives the best way to compute simple multiplicities.

Theorem 4.3.5. *If R is an algebraically closed field, V_M has finite length, and $S \in \text{Mod}_R^{\mathcal{D}}$ is irreducible, then the multiplicity of S in a composition series for V_M is given by the corank of the matrix $I(M)$ where I is the injective hull of S . (Here $I(M)$ denotes the the block matrix obtained by applying I to the entries of M .)*

If \mathcal{D} is dimension zero with homological modulus μ over R , then V_M automatically has finite length. Writing $i: \mathcal{C} \rightarrow \mathcal{D}$ for the inclusion of the full subcategory on the finite set $\mu(X)$ where $X \subseteq \text{Ob}(\mathcal{D})$ is a finite set containing the row labels of M ,

every irreducible composition factor S of V_M satisfies $S \simeq i_* i^* S$. Since there are only finitely many irreducible representations S with $i^* S \neq 0$, only a finite number of coranks need be computed.

Proof. We have $\dim_R \text{Hom}_R^{\mathcal{D}}(S, I) = 1$ since this module is free of rank one over the division algebra $\text{End}(S)$, which must be one-dimensional since R is algebraically closed. (If R is not algebraically closed, we have a similar result, but we would have to divide the coranks by the dimension of $\text{End}(S)$.) Since $\text{Hom}_R^{\mathcal{D}}(-, I)$ is an exact functor, an easy induction on the length of V proves that $\dim_R \text{Hom}_R^{\mathcal{D}}(V, I)$ gives the multiplicity of the simple S in a composition series for V . The result follows by applying Proposition 2.5.2 to the opposite category \mathcal{D}^{op} wherein the linear dual of I is projective.

Corollary 4.2.3 gives that the restriction $i^* V_M$ has the same length as V_M . It follows that each composition factor of V_M remains nonzero under the exact functor i^* . By Proposition 3.1.8 $\ell(i^* S) \leq \ell(S) = 1$, and since $i^* S \neq 0$, $i^* S$ is irreducible. By Yoshioka's theorem (Proposition 3.1.10) S is the unique irreducible representation that restricts to $i^* S$. But Theorem 3.3.4 gives that $i^* i_* i^* S \simeq i^* S$ as well, and so $S \simeq i_* i^* S$ as required.

The computation becomes finite because the category \mathcal{C} has finitely many objects (by construction) and finite hom-sets since \mathcal{D} does (by Theorem 4.3.2), and so \mathcal{C} is a finite category with finitely many irreducible representations up to isomorphism. Each of these extends uniquely to an irreducible representation of \mathcal{D} by Yoshioka's theorem, and these are the only irreducibles that may appear in a composition series for V_M . □

Remark 4.3.6 (Column reduction of a matrix over a category). We return to the notion of column reduction advertized in the introduction. The idea is that the indecomposable injectives provide a faithful functor $f: \mathcal{D} \hookrightarrow \prod_{\lambda} \mathbf{Vect}_{\mathcal{C}}$ by the formula $f = \prod_{\lambda} I_{\lambda}$. Under f , every arrow of \mathcal{D} maps to an infinite list of matrices, almost all

of which are zero (this is equivalent to \mathcal{D} being dimension zero over R). Applying f to a matrix M over \mathcal{D} gives a λ -indexed system of linear maps over R . The purpose of Theorem 4.3.5 is to explain how traditional (but λ -parallel) Gaussian elimination on fM computes the K-class of V_M .

In light of the theorem, we have no need (or desire) to introduce a separate version of Gaussian elimination for every dimension zero category. Instead, we have a general method by which such problems may be converted to usual Gaussian elimination.

4.3.2 Bounded models

This section provides a result saying that every category of dimension zero has the same representation theory as a category that “obviously” has dimension zero because its basic projectives have finite length when considered as R -modules. The fundamental example is the equivalence of cosimplicial R -modules and cochain complexes of R -modules supported in nonnegative degree. We return to this example in §4.3.3.

Definition 4.3.7. Two categories \mathcal{C} and \mathcal{D} are called **Morita equivalent over a commutative ring R** if their categories of representations $\mathbf{Mod}_R^{\mathcal{D}}$ and $\mathbf{Mod}_R^{\mathcal{C}}$ are equivalent as R -linear categories.

Lemma 4.3.8. *If $\mathcal{P} \subseteq \Pi_R(\mathcal{D})$ is a full subcategory of the category of idempotents with the property that for all nonzero $V \in \mathbf{Mod}_R^{\mathcal{D}}$ there exists some $\pi \in \text{Ob}(\mathcal{P})$ so that $V(\pi) \neq 0$, then \mathcal{P} is Morita equivalent to \mathcal{D} .*

Proof. Given a representation $V \in \mathbf{Mod}_R^{\mathcal{D}}$, we have a representation $V': \mathcal{P} \rightarrow \mathbf{Mod}_R$ given by the formula $V'\pi = \text{im } V\pi$. Certainly this assignment is functorial, full, and faithful. It is essentially surjective on objects since it is exact and the basic projective representations of \mathcal{P} come from a collection of enough projectives in $\mathbf{Mod}_R^{\mathcal{D}}$. \square

Definition 4.3.9. A **Cauchy model** for a category \mathcal{D} over a ring R is a full subcategory of $\Pi_R(\mathcal{D})$ spanned by idempotents π_λ so that the representations $\text{im } \pi_\lambda^{\natural} \in \mathbf{Mod}_R^{\mathcal{D}}$ form a complete set of representatives for the indecomposable projectives up to isomorphism.

By the Krull-Schmidt Theorem 3.1.4, if \mathcal{D} is dimension zero over R , every basic projective splits uniquely as a direct sum of indecomposable projectives, so a Cauchy model is always Morita equivalent to \mathcal{D} by Lemma 4.3.8.

Definition 4.3.10. An R -linear category \mathcal{C} is called **bounded** if for every object $c \in \text{Ob}(\mathcal{C})$, the set of objects $\{c' \in \text{Ob}(\mathcal{C}) \mid \text{Hom}_{\mathcal{C}}(c, c') \neq 0\}$ is finite and every hom- R -module $\text{Hom}_{\mathcal{C}}(c, c')$ is finitely generated.

Theorem 4.3.11 (Bounded models). *A category \mathcal{D} is dimension zero over an Artinian ring R if and only if it is Morita equivalent to a bounded category \mathcal{C} over R . In this case, we may take \mathcal{C} to be a Cauchy model of \mathcal{D} over R .*

Question 4.3.12. If a category admits a homological modulus over \mathbb{Z} , is it Morita equivalent over \mathbb{Z} to a bounded \mathbb{Z} -linear category?

Remark 4.3.13. Theorem 4.3.11 tells us that all dimension zero categories are abstractly Morita equivalent to categories where the basic projectives are finite length as R -modules, but the result has computational use only when a bounded model may be computed, as we may for example in §4.3.3. Typically, however, using Theorem 4.3.2 to show that a category is dimension zero is much easier than finding a bounded model, and almost as useful in computations.

Proof of Theorem 4.3.11. Certainly if \mathcal{D} has a bounded model over an Artinian ring R , then it is dimension zero over R ; indeed, finitely generated representations have finite length even after forgetting the action of the bounded model.

In the other direction, suppose Λ is an indexing set for the simple representations of \mathcal{D} so that each simple representation is isomorphic to one of the form $S_\lambda \in \text{Mod}_R^{\mathcal{D}}$

for a unique $\lambda \in \Lambda$. By Proposition 3.2.2, $S_\lambda \in \mathbf{Mod}_R^{\mathcal{D}}$ has a projective cover P_λ , which can be written $P_\lambda = \text{im } \pi_\lambda^{\natural}$ for some idempotent matrix $\pi_\lambda \in \text{Mat}_R^{\mathcal{D}}(d_\lambda^{\oplus} \rightarrow d_\lambda^{\oplus})$ by Proposition 2.5.2. Let \mathcal{P} be the full subcategory of $\Pi_R(\mathcal{D})$ on the idempotents $\{\pi_\lambda\}_{\lambda \in \Lambda}$. We claim that \mathcal{P} is bounded and Morita equivalent to \mathcal{D} . Moreover, since any indecomposable projective is a projective cover for some simple (again, by Proposition 3.2.2), the category \mathcal{P} is a Cauchy model for \mathcal{D} .

Suppose that $\text{Hom}_{\Pi(\mathcal{D})}(\pi_\lambda, \pi_{\lambda'}) \neq 0$. By Proposition 2.5.4, there must be a nonzero map $\varphi : P_{\lambda'} \rightarrow P_\lambda$ since $P_\lambda = \text{im } \pi_\lambda^{\natural}$ and $P_{\lambda'} = \text{im } \pi_{\lambda'}^{\natural}$. Since $\text{im } \varphi$ is a quotient of $P_{\lambda'}$, it must have $S_{\lambda'}$ as a composition factor by Proposition 3.2.2, and so P_λ must have $S_{\lambda'}$ as a composition factor as well. Since P_λ has only finitely many composition factors, \mathcal{P} is bounded.

We prove Morita equivalence using the condition from in Proposition 4.3.8: given a nonzero representation $V \in \text{Mod}_R^{\mathcal{D}}$, we must produce λ so that $V(\pi_\lambda) \neq 0$. Pick some nonzero vector in V , and let $V_0 \subseteq V$ be the subrepresentation it generates inside V . Since V_0 is nonzero and finitely generated, it is finite length and has minimal subrepresentation S_λ . So P_λ admits a nonzero map to V_0 and $\text{Hom}_R^{\mathcal{D}}(P_\lambda, V_0) \subseteq \text{Hom}_R^{\mathcal{D}}(P_\lambda, V)$ are both nonzero. By Proposition 2.5.2, this last R -module equals $\text{im } V(\pi_\lambda)$, and the claim is proved. \square

4.3.3 The Dold-Kan Correspondence

We give the most famous example of a bounded model: the category of cochain complexes. Let Δ be the category whose objects are the finite sets $[1], [2], \dots$ and whose morphisms are functions that are weakly order preserving. A functor $V : \Delta \rightarrow \mathbf{Mod}_R$ is more commonly known as a **cosimplicial R -module**. We state the Dold-Kan correspondence for cosimplicial R -modules.

Theorem 4.3.14. *The category of cosimplicial R -modules is equivalent to the category of cochain complexes supported in nonnegative degree.*

A cochain complex is a representation of the **cochain category** \mathcal{Ch} , a bounded category whose objects are the natural numbers and where $\text{Hom}_{\mathcal{Ch}}(n, n) = R \cdot 1_n$, $\text{Hom}_{\mathcal{Ch}}(n, n+1) = R \cdot \partial$, and $\text{Hom}_{\mathcal{Ch}}(n, m) = 0$ for any other $m \in \mathbb{N}$. Evidently, the basic projectives for this bounded model are indecomposable, so the cochain category provides a Cauchy model for Δ . Standard proofs of the Dold-Kan Correspondence (for example, [Wei94, Chapter 8]) give the following stronger result.

Theorem 4.3.15. *The basic projective $P^{[n]}: \Delta \rightarrow \mathbf{Mod}_R$ decomposes as a direct sum of indecomposable projectives*

$$P^{[n]} \simeq \bigoplus_{k=1}^n (P^k)^{\oplus \binom{n-1}{k-1}}$$

where P^k is a basic projective in the cochain model.

4.3.4 Further examples of categories of dimension zero

Example 4.3.16 (A homological modulus for the category of finite sets). The paper [WG14] shows that, over the rationals, the category of finite sets has a homological modulus given by $\mu(0) = \{0, 1\}$, $\mu(d) = \{d+1\}$ for $d \neq 0$, which also gives the more precise result that one may take $\mu(\varepsilon_k) = \{\varepsilon_k, \varepsilon_{k+1}\}$ and $\mu(\theta_k) = \{\theta_k, \varepsilon_k\}$.

Example 4.3.17 (Noncommutative finite sets are dimension zero). In [EWG15], Pirashvili-Richter’s category of noncommutative finite sets [PR02] is shown to be dimension zero over any field, and the simples in characteristic zero are deduced from [WG14] (these simples appeared earlier in [Rai09], and their images under Schur-Weyl duality even earlier in [Rud74]).

Example 4.3.18 (The category of finite sets with relations is dimension zero). Andrew Gitlin [Git15] has produced a homological modulus over \mathbb{Q} for the category of finite sets with relations, wherein $\mu(d) = \{2^d\}$. This is the first example known to the author of a homological modulus that is not “linear in d .” The irreducible representations

over \mathbb{Q} have now been constructed by Serge Bouc and Jacques Thévenaz [BT15].

Example 4.3.19 (Homological modulus for finite dimensional vector spaces over a finite field). The category of finite dimensional vector spaces over a finite field of characteristic p admits a homological modulus over $\mathbb{Z}[\frac{1}{p}]$ by work of Kuhn [Kuh15], who relies on idempotents constructed much earlier by Kovács [Kov92]. Working in a skeleton where objects are natural numbers and morphisms are \mathbb{F}_q -matrices, we may take $\mu(n) = \{n\}$. Indeed, the two-sided ideal $(n+1, n, n+1) \subseteq (n+1, n+1)$ is generated by an explicit idempotent that acts as an identity for this ideal considered as a subalgebra, and so witnesses $n+1 \leq_n n$. Appropriate application of Lemma 4.1.9 finishes the proof.

4.3.5 Schur projectives

We give the definition and basic properties of Schur projectives. These representations were useful in [WG14], which used them to give an analogue of Hilbert’s syzygy theorem for the category of finite sets. They will not be used further in this thesis.

Suppose \mathcal{D} is an \mathbb{N} -filtered category with associated graded $\mathcal{A} = \sqcup_k \mathcal{A}_k$, inclusions of homogeneous subcategories $i_k : \mathcal{D}_k \subseteq \mathcal{D}$ and projections $p_k : \mathcal{D}_k \twoheadrightarrow \mathcal{A}_k$. Choose subcategories $j_k : \mathcal{B}_k \hookrightarrow \mathcal{D}_k$ so that the composites $p_k \circ j_k$ are full and essentially surjective on objects.

In this section, we assume that \mathcal{D} is dimension zero over R . Suppose the irreducible representations of \mathcal{A} are indexed by an \mathbb{N} -graded set Λ , and write $W_\lambda \in \text{Mod}_R^{\mathcal{A}_k}$ for the irreducible corresponding to $\lambda \in \Lambda$. Similarly, set $V_\lambda = i_{k*} p^* W_\lambda$. Every irreducible representation of \mathcal{D} is uniquely of this form by Theorem 3.3.7. Finally, using Theorem 3.3.4, fix a projective cover P_λ and injective hull I_λ for each simple V_λ . For each $\lambda \in \Lambda$, use Proposition 2.5.2 to choose a formal sum $d_\lambda^\oplus \in \mathcal{D}^\oplus$ and $\pi_\lambda \in \text{Mat}_R^{\mathcal{D}}(d_\lambda^\oplus \rightarrow d_\lambda^\oplus)$ with $P_\lambda \simeq \text{im } \pi_\lambda^\natural$.

Definition 4.3.20. Given $\lambda \in \Lambda$, a **Schur projective** $U_\lambda \in \text{Mod}_R^{\mathcal{D}}$ is defined

$$U_\lambda = (i_k \circ j_k)!T_\lambda$$

where $T_\lambda \in \text{Mod}_R^{\mathcal{B}^k}$ is a projective cover of the simple $(p_k \circ j_k)^*W_\lambda$.

Proposition 4.3.21. *The Schur projective U_λ decomposes upper-triangularly*

$$U_\lambda = P_\lambda \oplus \bigoplus_{\deg(\lambda') < \deg(\lambda)} P_{\lambda'}^{\oplus m_{\lambda'}}$$

as a direct sum of indecomposable projectives.

Proof. A decomposition into indecomposable projectives exists and is unique by the Krull-Schmidt Theorem 3.1.4; we must prove upper-triangularity. Let $\lambda, \lambda' \in \Lambda$, supposing (for now) $\deg(\lambda) = \deg(\lambda') = k$. Compute

$$\begin{aligned} \text{Hom}_R^{\mathcal{D}}(U_\lambda, V_{\lambda'}) &\simeq \text{Hom}_R^{\mathcal{D}}((i_k \circ j_k)!T_\lambda, V_{\lambda'}) \\ &\simeq \text{Hom}_R^{\mathcal{D}^k}((j_k)!T_\lambda, (i_k)^*V_{\lambda'}) \\ &\simeq \text{Hom}_R^{\mathcal{D}^k}((j_k)!T_\lambda, (i_k)^*(i_k)!_*(p_k)^*W_{\lambda'}) \\ &\simeq \text{Hom}_R^{\mathcal{D}^k}((j_k)!T_\lambda, (p_k)^*W_{\lambda'}) \\ &\simeq \text{Hom}_R^{\mathcal{B}^k}(T_\lambda, (p_k \circ j_k)^*W_{\lambda'}). \end{aligned}$$

If $\lambda = \lambda'$, then all nonzero maps give projective covers, and this last module is free of rank 1 when considered as a right D -module where D is the division ring $\text{End}(W_{\lambda'}) \simeq \text{End}(V_{\lambda'})$ acting by postcomposition. It follows that the projective cover of V_λ occurs as a summand of S_λ with multiplicity one. If $\lambda \neq \lambda'$, then this last module is zero since T_λ is the projective cover of $(p_k \circ j_k)^*W_\lambda$ and cannot map to any other simple. If, on the other hand, $\deg(\lambda) < \deg(\lambda')$; then $(i_{\deg(\lambda)})^*V_{\lambda'} = 0$ since $V_{\lambda'}d = 0$ for any $d \in \mathcal{D}$ with $\deg(d) < \deg(\lambda')$, and so the second line in the

computation vanishes as well. □

4.4 \mathcal{D} -regular functions

Every representation $V \in \mathbf{Mod}_R^{\mathcal{D}}$ determines a dimension function $d \mapsto \dim Vd$. In this section, we develop tools to explain the behavior of dimension functions, including a generalization to the case $R = \mathbb{Z}$.

Definition 4.4.1. A **basic \mathcal{D} -regular function over R** is a function $\mathrm{Ob}(\mathcal{D}) \rightarrow \mathbb{N}$ of the form

$$d \mapsto (\text{multiplicity of } P_\lambda \text{ in a direct sum decomposition of } P^d)$$

for some $\lambda \in \Lambda$, where we have assumed that this multiplicity makes sense (if R is Artinian and the hom-sets of \mathcal{D} are finite, we may rely on the Krull-Schmidt Theorem 3.1.4; other situations will need their own Krull-Schmidt theorems). A **\mathcal{D} -regular function over R** is any \mathbb{N} -linear combination of basic \mathcal{D} -regular functions. Similarly, a **signed \mathcal{D} -regular function** (defined to take values in \mathbb{Z} instead of \mathbb{N}) is any \mathbb{Z} -linear combination of basic \mathcal{D} -regular functions.

Example 4.4.2. By Theorem 4.3.15, the basic Δ -regular functions (here, Δ is the category of finite, nonempty total orders and monotone maps) over any Artinian ring are of the form $[n] \mapsto \binom{n-1}{k-1}$ for various k .

Theorem 4.4.3. *If $V \in \mathrm{Mod}_R^{\mathcal{D}}$ is a finitely generated representation of a dimension zero category \mathcal{D} over a commutative ring R (which is necessarily Artinian by Theorem 4.3.2), then there exists a finite list of indecomposable R -modules M_1, \dots, M_k and a corresponding list of basic \mathcal{D} -regular functions $\varphi_1, \dots, \varphi_k$ so that*

$$Vd \simeq M_1^{\oplus \varphi_1(d)} \oplus \dots \oplus M_k^{\oplus \varphi_k(d)}.$$

Proof. By Theorem 4.3.11, a Cauchy model \mathcal{C} for \mathcal{D} is bounded and has $\mathbf{Mod}_R^{\mathcal{D}} \simeq \mathbf{Mod}_R^{\mathcal{C}}$. Proposition 2.4.4 provides an intrinsic description of being finitely generated (any formula of the form $V = \sum_{\alpha} V_{\alpha}$ for some infinite collection of subrepresentations V_{α} has infinite redundancy), so V is still finitely generated when considered as a representation of \mathcal{C} . Since \mathcal{C} is bounded, there are only finitely many $c \in \mathcal{C}$ for which $Vc \neq 0$; write $\text{Supp}(V) = \{c \in \mathcal{C} \mid Vc \neq 0\}$ for this finite set of objects. For each $c \in \text{Supp}(V)$, the previous discussion gives $\ell(Vc) < \infty$ and so by the (classical) Krull Schmidt Theorem, $Vc \simeq M_{c,1} \oplus \cdots \oplus M_{c,k_c}$ where each $M_{c,i} \in \text{Mod}_R$ is indecomposable.

Let $d \in \mathcal{D}$. By the Krull-Schmidt Theorem 3.1.4, P^d decomposes as a direct sum of indecomposable projectives. In a Cauchy model, every indecomposable projective $P \in \text{Mod}_R^{\mathcal{D}}$, considered as an object of $\mathbf{Mod}_R^{\mathcal{C}}$, is isomorphic to a unique projective of the form $P^c \in \text{Mod}_R^{\mathcal{C}}$ for some $c \in \mathcal{C}$. So P^d corresponds under the equivalence $\mathbf{Mod}_R^{\mathcal{D}} \simeq \mathbf{Mod}_R^{\mathcal{C}}$ to a representation isomorphic to $P^{c_1} \oplus \cdots \oplus P^{c_l}$ for some objects $c_1, \dots, c_l \in \mathcal{C}$. Fixing $c \in \mathcal{C}$, we may ask for the number of appearances of c among the list c_1, \dots, c_l ; indeed, this number is exactly the multiplicity of P^c appearing as a direct summand of P^d . In other words, as $d \in \mathcal{D}$ varies, this multiplicity coincides with some basic \mathcal{D} -regular function $\varphi_c(d)$. To conclude the proof, compute

$$\begin{aligned}
Vd &\simeq \text{Hom}_R^{\mathcal{D}}(P^d, V) \\
&\simeq \text{Hom}_R^{\mathcal{C}}(P^{c_1} \oplus \cdots \oplus P^{c_l}, V) \\
&\simeq Vc_1 \oplus \cdots \oplus Vc_l \\
&\simeq \bigoplus_{c \in \text{Supp}(V)} (Vc)^{\oplus \varphi_c(d)} \\
&\simeq \bigoplus_{c \in \text{Supp}(V)} (M_{c,1})^{\oplus \varphi_c(d)} \oplus \cdots \oplus (M_{c,k_c})^{\oplus \varphi_c(d)}.
\end{aligned}$$

□

Corollary 4.4.4. *If $V \in \text{Mod}_R^{\mathcal{D}}$ satisfies the hypotheses of Theorem 4.4.3 and $N \in$*

Mod_R is a simple R -module, then the function

$$d \longmapsto (\text{multiplicity of } N \text{ in a composition series for } Vd)$$

is \mathcal{D} -regular. Similarly, $d \longmapsto \ell(Vd)$ is \mathcal{D} -regular. More generally, if $\chi: \text{Mod}_R \rightarrow \mathbb{Z}_{\geq 0}$ (resp. \mathbb{Z}) is a function that's additive on direct sums, then $d \longmapsto \chi(Vd)$ is \mathcal{D} -regular (resp. signed \mathcal{D} -regular).

Proof. Use additivity of χ on the direct sum decomposition in Theorem 4.4.3. \square

Corollary 4.4.5. *If R is a PID, each hom-set $\text{Hom}_{\mathcal{D}}(x, y)$ is finite, and \mathcal{D} admits a homological modulus over R , then for each finitely generated $V \in \text{Mod}_R^{\mathcal{D}}$ there exists a finite list of prime powers $q_1, \dots, q_k \in R$ and a corresponding list of functions $\varphi_1, \dots, \varphi_k$ along with an extra function ψ so that*

$$Vd \simeq R^{\oplus \psi(d)} \oplus R/(q_1)^{\oplus \varphi_1(d)} \oplus \dots \oplus R/(q_k)^{\oplus \varphi_k(d)},$$

where the function ψ is \mathcal{D} -regular over the field of fractions of R , and each function φ_i is signed \mathcal{D} -regular over $R/(q_i)$.

Proof. First, let us show that the list of prime powers appearing is finite. By Theorem 4.2.2, every map of representations $P^d \rightarrow V$ is determined by its restriction along $i: \mathcal{M} \rightarrow \mathcal{D}$, where \mathcal{M} is the full subcategory spanned by the objects μ (location of generators for V). By Yoneda's lemma, $\text{Hom}_R^{\mathcal{D}}(P^d, V) \simeq Vd$, and so $Vd \simeq \text{Hom}_R^{\mathcal{M}}(i^*P^d, i^*V) \subseteq \prod_{m \in \mathcal{M}} \text{Hom}_{\mathbb{Z}}(P^d m, Vm) \simeq \prod_{m \in \mathcal{M}} (Vm)^{\oplus |\text{Hom}_{\mathcal{D}}(d, m)|}$ is a submodule of a module built from sums of the modules Vm . It follows that the prime powers that appear must divide the order of their primary component inside the finitely generated module $\oplus_{m \in \mathcal{M}} Vm$.

Tensoring V with the field of fractions of R , we kill all non-free summands and so the dimension matches the free rank; it follows that ψ is \mathcal{D} -regular over this

field. Pick i so that q_i doesn't divide any of the other prime powers. Tensoring with $R/(q_i)$, we pick up copies of $R/(q_i)$ from free copies of R as well as legitimate copies of $R/(q_i)$. This overcount is still \mathcal{D} -regular over $R/(q_i)$ because it gives the multiplicities of $R/(q_i)$ in the finitely generated representation $R/(q_i) \otimes_R V/\text{torsion}$, and so φ_i is signed- \mathcal{D} -regular over the ring $R/(q_i)$. A similar inclusion-exclusion gives the result. □

CHAPTER V

Case study: Representations of the category of finite sets

The purpose of this chapter is two-fold: to exhibit the power of the general toolkit in a special case, and to advertise its usability by showing the ease with which we compute examples and formulate conjectures. The tools themselves were developed in Chapters III and IV, but we strive to make this chapter reasonably self-contained.

A future researcher might use this chapter as a schematic for attacking a new category suspected of being dimension zero.

5.1 Summary of results for the representation theory of the category of finite sets

Let \mathcal{D} denote the category whose objects are the finite sets $[n] = \{1, \dots, n\}$ for $n \in \mathbb{N}$ and whose morphisms are the functions between them. Since any other finite set is isomorphic to one of these, the representation theory of \mathcal{D} matches the representation theory of the category of finite sets.

The main computational result of this chapter (Theorem 5.1.6) is novel, giving a practical way to deal with representations of \mathcal{D} arising in nature. The construction of the indecomposable injective modules in Theorem 5.4.2 appears to be new as well.

Finally, Theorem 5.6.1 gives a new result about representations over \mathbb{Z} .

The other results were obtained by this author in the early years of his Ph.D., but in actuality are due to (or were already known to) Eric Rains [Rai09] who deduces them from classical work of Putcha [Put96] on the representation theory of the full transformation monoid. The proofs below are due to the author.

We begin with a concrete characterization of the finitely generated representations.

Theorem 5.1.1. *If $V: \mathcal{D} \rightarrow \mathbf{Vect}_{\mathbb{Q}}$ is a representation, then V is finitely generated if and only if the sequence $n \mapsto \dim_{\mathbb{Q}} V$ is bounded above by a polynomial in n .*

We give the classification of irreducible representations.

Theorem 5.1.2. *The irreducible representations of \mathcal{D} are in bijection with integer partitions. If λ is a partition, the matrices X_{λ} given in Theorem 5.3.3 provide an explicit model of the irreducible representation corresponding to λ . In particular, the dimension sequences satisfy*

$$\dim V_{\lambda}[n] = \begin{cases} \dim W_{\lambda} \cdot \binom{n}{k} & \text{for } \lambda \text{ not a column} \\ \binom{n-1}{k-1} & \text{for } \lambda \text{ a column, } n > 0 \end{cases}$$

where W_{λ} denotes an irreducible representation of $S_{|\lambda|}$.

We now describe the precise way in which every finitely generated representation is built up from these irreducibles.

Theorem 5.1.3. *Every finitely generated representation $V: \mathcal{D} \rightarrow \mathbf{Vect}_{\mathbb{Q}}$ admits a filtration*

$$0 = V_0 \subsetneq V_1 \subsetneq V_2 \subseteq \cdots \subsetneq V_k = V$$

so that each successive quotient V_i/V_{i-1} is irreducible. As usual, any two such filtrations gives the same multiset of irreducibles.

Since dimension functions add in short exact sequences, we obtain the following.

Corollary 5.1.4. *A representation V is finitely generated if and only if the sequence $n \mapsto \dim_{\mathbb{Q}} V$ coincides with a polynomial for all $n > 0$.*

The next result guarantees the existence of a presentation matrix for every finitely generated representation.

Theorem 5.1.5. *A representation $V: \mathcal{D} \rightarrow \mathbf{Vect}_{\mathbb{Q}}$ is finitely generated if and only if it admits a presentation matrix over \mathcal{D} with coefficients in \mathbb{Q} .*

Once we have a presentation matrix for V , the last result (and the only novel result in this summary) lets us compute its multiset of irreducible composition factors.

Theorem 5.1.6. *If M is a presentation matrix for a representation $V: \mathcal{D} \rightarrow \mathbf{Vect}_{\mathbb{Q}}$ the multiplicity in V of the irreducible representation associated to λ may be computed as the corank of the matrix obtained by applying Y_{λ} (see Theorem 5.4.2) to the entries of M . In particular, if the largest row label is $[r]$, then no partition of size greater than $r + 1$ appears with positive multiplicity, and any of size exactly $r + 1$ must be a column partition.*

The rest of the chapter is an extended, fully-worked example.

5.2 First experiments

5.2.1 Computing in \mathcal{D}

To prepare for the computations required in further steps, we must program a computer to perform composition in \mathcal{D} and find hom-sets. A few keystrokes should tell us that

$$\mathrm{Hom}_{\mathcal{D}}([2], [3]) = \{\boxed{11}, \boxed{12}, \boxed{13}, \boxed{21}, \boxed{22}, \boxed{23}, \boxed{31}, \boxed{32}, \boxed{33}\}$$

or that

$$\boxed{2231} \circ \boxed{32441} = \boxed{32112}.$$

It will also be convenient to allow linear combinations of morphisms in \mathcal{D} , where composition extends by linearity:

$$\left(\boxed{412} - 3\boxed{222}\right) \circ \left(\boxed{12} + \boxed{23}\right) = \boxed{12} - 6\boxed{22} + \boxed{41}.$$

5.2.2 Computing the preorders \leq_d

According to Definition 4.1.1, each object $d \in \mathcal{D}$ induces a preorder \leq_d on $\text{Ob}(\mathcal{D})$. Theorem 4.3.2 characterizes categories of dimension zero as those for which every preorder \leq_d has a finite “joint maximum,” which is to say, a finite collection of objects $\mu(d) \subseteq \text{Ob}(\mathcal{D})$ so that for any $x \in \mathcal{D}$ there exists $m \in \mu(d)$ so that $x \leq_d m$.

We use a computational description of $\leq_{[d]}$ suitable for experimentation (see Proposition 4.1.3). For any composite map $s: [n] \rightarrow [m] \rightarrow [n]$, we build the $\text{Hom}_{\mathcal{D}}([d], [n]) \times \text{Hom}_{\mathcal{D}}([d], [n])$ matrix M_s whose (f, g) entry is 1 if $s \circ f = g$ and 0 otherwise. We have $[n] \leq_{[d]} [m]$ exactly when the identity matrix is in the \mathbb{Q} -linear span of the M_s .

Let’s check if $[3] \leq_{[1]} [2]$. For each

$$s \in \{\boxed{111}, \boxed{112}, \boxed{113}, \boxed{121}, \boxed{122}, \boxed{131} \dots, \boxed{333}\}$$

in the set of self-maps of $[3]$ that factor through $[2]$, build the 1-0-matrix M_s :

$$\begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}, \dots, \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix}.$$

Since

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} - \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix},$$

the inequality holds. Record this result in a table for $[n] \leq_{[1]} [m]$.

$\leq_{[1]}$	[0]	[1]	[2]	[3]	[4]
[0]					
[1]					
[2]					
[3]			✓		
[4]					

Table 5.1: Recording the inequality $[3] \leq_{[1]} [2]$

A computer calculation fills out the table.

$\leq_{[1]}$	[0]	[1]	[2]	[3]	[4]
[0]	✓	✓	✓	✓	✓
[1]	✗	✓	✓	✓	✓
[2]	✗	✗	✓	✓	✓
[3]	✗	✗	✓	✓	✓
[4]	✗	✗	✓	✓	✓

Table 5.2: The preorder $\leq_{[1]}$

It seems as though [2] might be a maximum element for this preorder. A similar computation for $\leq_{[2]}$ gives the table

$\leq_{[2]}$	[0]	[1]	[2]	[3]	[4]
[0]	✓	✓	✓	✓	✓
[1]	✗	✓	✓	✓	✓
[2]	✗	✗	✓	✓	✓
[3]	✗	✗	✗	✓	✓
[4]	✗	✗	✗	✓	✓

Table 5.3: The preorder $\leq_{[2]}$

suggesting that [3] might be a maximum for the preorder $\leq_{[2]}$. Since it also seems that [1] is a maximum for the preorder $\leq_{[0]}$, we feel justified making the following guess.

Guess 5.2.1. Every object $[n] \in \mathcal{D}$ satisfies $[n] \leq_{[d]} [d + 1]$.

If true, this statement is an elementary fact of combinatorics and linear algebra, and should be amenable to direct proof. The general theory provides a simplifying fact, which is a corollary of Lemma 4.1.9:

Proposition 5.2.2. *If the identity morphism of c can be factored through d , then $x \leq_d y \implies x \leq_c y$.*

Since the identity function on $[n]$ factors through $[n+1]$ for $n \geq 1$, the preorders \leq_d only gain checkmarks as d gets smaller. Consequently, to prove Guess 5.2.1, it suffices to prove that for each $n \geq 0$, we have

$$[n+2] \leq_{[n]} [n+1].$$

Proof of Guess 5.2.1. Define the set

$$H = \{h \in \text{Hom}_{\mathcal{D}}([n+2], [n+2]) : i \leq h(i) \leq i+1 \text{ for all } i\}.$$

We claim that

$$\sum_{h \in H} \left((-1)^{\sum_{i=1}^{n+2} (h(i)-i)} \right) \cdot M_h = 0,$$

from which the result follows since the identity matrix shows up in this sum exactly once with coefficient 1. We must show that the (f, g) -entry of this matrix is zero for any $f, g \in \text{Hom}_{\mathcal{D}}([n], [n+2])$. To this end, choose some $k \notin \text{Im}(f)$, $k \neq n+2$, and define an involution $\tau: H \rightarrow H$ flip-flopping function values at k ; precisely, $\tau(h)(i) = h(i)$ when $i \neq k$ and $t(h)(k) = 2k+1-h(k)$. We see that $h \circ f = \tau(h) \circ f$, and so the (f, g) -entry, which depends only on the truth of the equation $h \circ f = g$, is fixed under the action of τ . On the other hand, τ changes exactly one term in the sum $\sum_{i=1}^{n+2} (h(i)-i)$, switching between even and odd; so τ also acts on the sum by negation. □

Using terminology to be introduced in Definition 4.2.4 and applying Theorem 4.3.2,

we have proved the following result:

Theorem 5.2.3. *The function $\mu([0]) = \{[0], [1]\}$, $\mu([d]) = \{[d + 1]\}$ is a homological modulus for \mathcal{D} over \mathbb{Z} , and therefore \mathcal{D} is dimension zero over any Artinian ring R .*

Let us list some concrete statements that follow from our work so far. Every finitely generated representation has finite length, meaning that every representation has a finite filtration whose successive quotients are simple. If we later get a grip on the simples, we will have gone a long way toward understanding a general representation. By Proposition 4.1.3, the homological modulus tells us that if U is generated by vectors in $U[k]$ (with $k > 0$, say), any subrepresentation $U' \subseteq U$ is generated by vectors in $U'[k + 1]$.

These theoretical consequences are complemented by a crucial computational fact.

Corollary 5.2.4. *Every finitely generated representation of \mathcal{D} over a commutative ring R can be described completely by a finite presentation matrix whose entries are R -linear combinations of morphisms in \mathcal{D} . The category of finitely generated representations is closed under taking kernels, images, and cokernels. Furthermore, these constructions may be carried out algorithmically from knowledge of the homological modulus.*

A representation of \mathcal{D} is an infinite amount of highly compatible data. To be able to write one succinctly as an explicit matrix seems extremely convenient. Corollary 5.2.4 guarantees that homological computation on finitely generated representations may be carried out with explicit presentation matrices!

5.2.3 A first pass at the irreducible representations

The first step in constructing the simples is to pick an \mathbb{N} -filtration of \mathcal{D} , which is to say, a degree function $\deg: \text{Ob}(\mathcal{D}) \rightarrow \mathbb{N}$. We write \mathcal{D}_k for the full subcategory on the

objects of degree $k \in \mathbb{N}$. Since it will be in our interest to pick the filtration in a natural way, we set $\deg([n]) = n$. Our \mathbb{N} -filtration induces an “associated graded” category $\mathcal{A} = \sqcup_k \mathcal{A}_k$, where each $\mathcal{A}_k = \mathcal{D}_k / (\text{morphisms factoring through objects with degree } < k)$.

Theorem 3.3.7 tells us that the irreducible representations of \mathcal{D} are in canonical bijection with the irreducible representations of \mathcal{A} ; it also give a computational recipe. Let Λ_k be an indexing set for the irreducible representations of \mathcal{A}_k , and set $\Lambda = \sqcup_k \Lambda_k$ an \mathbb{N} -graded indexing set for the irreducible representations of \mathcal{A} .

In our case, \mathcal{A}_k is defined to be the algebra of endomorphisms of $[k]$ modulo the ideal spanned by functions factoring through smaller sets. Since these functions are exactly the non-bijections, we have a canonical isomorphism $\mathcal{A}_k \simeq \mathbb{Q}S_k$ with the group algebra of the symmetric group. The irreducible representations of the symmetric group S_k are indexed by partitions of size k , so we may as well think of the elements $\lambda \in \Lambda$ as partitions. We have a result.

Theorem 5.2.5. *The irreducible representations of \mathcal{D} are in bijection with partitions.*

Given $\lambda \in \Lambda_k$, write W_λ for the irreducible representation of S_k corresponding to λ . We describe the general procedure by which an irreducible representation of \mathcal{D} may be computed from knowledge of the irreducibles of \mathcal{A} . Write V_λ for the irreducible representation of \mathcal{D} coming from W_λ . By Yoshioka’s Theorem (given as Proposition 3.1.10), the representation V_λ is the unique irreducible with the property that $V_\lambda[k] \simeq W_\lambda$.

As before, we proceed computationally in the hopes of making a conjecture. Fix a matrix coefficient of the irreducible representation W_λ , and precompose with the natural map $\text{Hom}_{\mathcal{D}}([k], [k]) \rightarrow \mathcal{A}_k$ to obtain $\psi: \text{Hom}_{\mathcal{D}}([k], [k]) \rightarrow \mathbb{Q}$. Theorem 3.3.10 tells us that $V_\lambda[n]$ may be computed as the row space of the $\text{Hom}_{\mathcal{D}}([k], [n]) \times \text{Hom}_{\mathcal{D}}([n], [k])$ whose (p, q) -entry is given by $\psi(q \circ p)$.

For example, let $\lambda = \square\square$, the partition of 2 with one part, and suppose we wish to compute $V_\lambda[3]$. In this case, the function ψ is given by $\psi(\boxed{ij}) = 1$ if $i \neq j$

and $\psi(\boxed{\mathbf{ii}}) = 0$ (for more general λ , we could use the coefficients of the Young symmetrizer). Build the multiplication table for $\text{Hom}_{\mathcal{D}}([2], [3]) \times \text{Hom}_{\mathcal{D}}([3], [2])$

	$\boxed{111}$	$\boxed{112}$	$\boxed{121}$	$\boxed{122}$	$\boxed{211}$	$\boxed{212}$	$\boxed{221}$	$\boxed{222}$
$\boxed{11}$	$\boxed{11}$	$\boxed{11}$	$\boxed{11}$	$\boxed{11}$	$\boxed{22}$	$\boxed{22}$	$\boxed{22}$	$\boxed{22}$
$\boxed{12}$	$\boxed{11}$	$\boxed{11}$	$\boxed{12}$	$\boxed{12}$	$\boxed{21}$	$\boxed{21}$	$\boxed{22}$	$\boxed{22}$
$\boxed{13}$	$\boxed{11}$	$\boxed{12}$	$\boxed{11}$	$\boxed{12}$	$\boxed{21}$	$\boxed{22}$	$\boxed{21}$	$\boxed{22}$
$\boxed{21}$	$\boxed{11}$	$\boxed{11}$	$\boxed{21}$	$\boxed{21}$	$\boxed{12}$	$\boxed{12}$	$\boxed{22}$	$\boxed{22}$
$\boxed{22}$	$\boxed{11}$	$\boxed{11}$	$\boxed{22}$	$\boxed{22}$	$\boxed{11}$	$\boxed{11}$	$\boxed{22}$	$\boxed{22}$
$\boxed{23}$	$\boxed{11}$	$\boxed{12}$	$\boxed{21}$	$\boxed{22}$	$\boxed{11}$	$\boxed{12}$	$\boxed{21}$	$\boxed{22}$
$\boxed{31}$	$\boxed{11}$	$\boxed{21}$	$\boxed{11}$	$\boxed{21}$	$\boxed{12}$	$\boxed{22}$	$\boxed{12}$	$\boxed{22}$
$\boxed{32}$	$\boxed{11}$	$\boxed{21}$	$\boxed{12}$	$\boxed{22}$	$\boxed{11}$	$\boxed{21}$	$\boxed{12}$	$\boxed{22}$
$\boxed{33}$	$\boxed{11}$	$\boxed{22}$	$\boxed{11}$	$\boxed{22}$	$\boxed{11}$	$\boxed{22}$	$\boxed{11}$	$\boxed{22}$

and apply ψ :

	$\boxed{111}$	$\boxed{112}$	$\boxed{121}$	$\boxed{122}$	$\boxed{211}$	$\boxed{212}$	$\boxed{221}$	$\boxed{222}$
$\boxed{11}$	0	0	0	0	0	0	0	0
$\boxed{12}$	0	0	1	1	1	1	0	0
$\boxed{13}$	0	1	0	1	1	0	1	0
$\boxed{21}$	0	0	1	1	1	1	0	0
$\boxed{22}$	0	0	0	0	0	0	0	0
$\boxed{23}$	0	1	1	0	0	1	1	0
$\boxed{31}$	0	1	0	1	1	0	1	0
$\boxed{32}$	0	1	1	0	0	1	1	0
$\boxed{33}$	0	0	0	0	0	0	0	0

Since this matrix has rank 3, we see that $\dim V_{\square}[3] = 3$. Performing similar computations, we obtain the following table

n	0	1	2	3	4	5
$\dim V_{\square\square}[n]$	0	0	1	3	6	10

Table 5.4: Dimensions of the irreducible representation with $\lambda = \square\square$

The resulting function seems to match $n \mapsto \binom{n}{2}$. Computing more data,

n	0	1	2	3	4	5
$\dim V_{\square}[n]$	0	0	1	2	3	4

Table 5.5: Dimensions of the irreducible representation with $\lambda = \begin{smallmatrix} \square \\ \square \end{smallmatrix}$

and we see that the partition $\begin{smallmatrix} \square \\ \square \end{smallmatrix}$ seems to give the function $n \mapsto n - 1$ for $n > 0$. With no immediate general pattern to conjecture, we compute the values for the partitions of 3:

$\begin{smallmatrix} \square & \square & \square \end{smallmatrix}$	0	1	2	3	4	5
$\begin{smallmatrix} \square & \square \\ \square \end{smallmatrix}$	0	0	0	1	4	10
$\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}$	0	0	0	2	8	20
$\begin{smallmatrix} \square & \square \\ \square & \square \\ \square \end{smallmatrix}$	0	0	0	1	3	6

Table 5.6: Dimensions of the irreducible representations with λ a partition of 3

Each sequence seems to be a multiple of a binomial sequence, but if the diagram is a column, the sequence is shifted by one.

Guess 5.2.6. If λ is a partition of k , the irreducible representation V_λ satisfies

$$\dim V_\lambda[n] = \begin{cases} \dim W_\lambda \cdot \binom{n}{k} & \text{for } \lambda \text{ not a column} \\ \binom{n-1}{k-1} & \text{for } \lambda \text{ a column, } n > 0. \end{cases}$$

5.3 More experiments suggested by the results of the first experiments

5.3.1 A Morita equivalent category separating columns and non-columns

According to Guess 5.2.6, the construction of V_λ is likely to depend on whether or not λ is a column. We introduce a technique for encoding this casework into a category. The idea is that there may be other categories with the same representation theory as \mathcal{D} . These categories can be thought of as “Morita equivalent” to \mathcal{D} by analogy to the case of two rings having equivalent categories of modules.

Lemma 4.3.8 provides a rich source of categories Morita equivalent to \mathcal{D} using linear combinations of morphisms that are idempotent under composition. With this construction in mind, define two idempotents in the algebra $\mathbb{C} \cdot \text{Hom}_{\mathcal{D}}([k], [k])$

$$\varepsilon_k = \frac{1}{k!} \sum_{\sigma \in S_k} \text{sign}(\sigma) \cdot \sigma$$

$$\theta_k = 1 - \varepsilon_k.$$

Since a column partition gives rise to the alternating representation, these idempotents neatly separate the cases of being a column or non-column.

Definition 5.3.1. Let \mathcal{C} be the category with three infinite families of objects $1_{[0]}, 1_{[1]}, 1_{[2]} \dots$, $\varepsilon_0, \varepsilon_1, \dots$, and $\theta_2, \theta_3, \dots$ where morphisms in \mathcal{C} are linear combinations of morphisms in \mathcal{D} according to the rule

$$\text{Hom}_{\mathcal{C}}(\pi_k, \pi'_l) = \{ \varphi \in \mathbb{C} \cdot \text{Hom}_{\mathcal{D}}([k], [l]) \mid \pi'_l \circ \varphi \circ \pi_k = \varphi \},$$

and π_k can stand for any of the idempotents $1_{[k]}$, ε_k , or θ_k , similarly for π'_l . Composition is inherited from composition in \mathcal{D} .

The category \mathcal{C} satisfies the hypotheses of Lemma 4.3.8 because it contains “enough

idempotents” in the sense that any nonzero representation U of \mathcal{D} must have $U\pi \neq 0$ for some $\pi \in \text{Ob}(\mathcal{C})$.

5.3.2 Computing the preorders \leq_π for the Morita equivalent category \mathcal{C}

Once again we seek information from the preorders \leq_c . Experimentation directly analogous to the earlier computations leads to the following guess, which we state as a proposition since we prove it forthwith.

Proposition 5.3.2. *For each $k \geq 1$, $\theta_{k+1} \leq_{1_{[k]}} 1_{[k]}$ in the category \mathcal{C} .*

Proof. Unraveling the definitions, we must find an element

$$\delta_k \in \mathbb{Q} \cdot \text{Hom}_{\mathcal{D}}([k+1], [k+1])$$

so that for any map $f \in \text{Hom}_{\mathcal{D}}([k], [k+1])$, $\delta_k \circ f = 0$ and $\delta_k \cong \theta_k$ modulo the ideal of maps factoring through $[k]$. Define the set

$$H = \{h \in \text{End}_{\mathcal{D}}([k+1]) : i \leq h(i) \leq i+1 \text{ if } i \leq k \text{ and } k \leq h(k+1) \leq k+1\}.$$

An argument similar to the proof of Guess 5.2.1 shows that

$$\sum_{h \in H} \left((-1)^{\sum_{i=1}^{k+1} (h(i)-i)} \right) \cdot M_h = 0,$$

and that this element is congruent (modulo maps factoring through $[k]$) to $1 + \tau$, the sum of the identity and the transposition flipping k and $k+1$. This element of $\mathbb{Q}S_k$ generates the two-sided ideal that misses ε since any representation sending $1 + \tau$ to zero must send τ to -1 (along with every other transposition, since these are conjugate in the group). In particular, the idempotent θ_k is in this ideal, and so we may pre- and post-multiply this sum by elements of $\mathbb{Q}S_k$ to obtain an element that is congruent to θ_k modulo maps that factor through $[k]$, as required. \square

Let us gain more fluency in the meaning of the preorders \leq_d . Earlier, we computed $\mu([k]) = \{[k+1]\}$ for any $k \geq 1$, telling us that if U is generated by vectors in $U[k]$, any subrepresentation $U' \subseteq U$ is generated by vectors in $U'[k+1]$. The result we have just obtained is more refined; it says that if U is generated by vectors in $U[k]$ and $U' \subseteq U$ is a subrepresentation generated by vectors in $U'[k+1]$ that are fixed by the action of θ_{k+1} , then U' is equally-well generated by vectors in $U'[k]$. The detailed grammar of the previous sentence may be a bit hard to parse—therein lies the value of the notation \leq_d .

Armed with Proposition 5.3.2, we return our attention to understanding the irreducible representations V_λ . We had conjectured their dimensions in Guess 5.2.6, which we are now ready to prove.

5.3.3 Explicit construction of the irreducibles

In the case of λ a non-column partition of k , the conjectured dimension function $n \mapsto \dim W_\lambda \cdot \binom{n}{k}$ suggests a possible form for the matrices of V_λ : block matrices with rows and columns coming from k -subsets of $[n]$ and entries coming from the classical irreducible S_k -representation W_λ . If λ is a column, we hope for matrices labeled by those k -subsets of $[n]$ that contain n .

To this end, pick some favored injections

$$\text{ChosenInj}([k], [n]) \subseteq \{\iota: [k] \hookrightarrow [n]\}$$

so that each k -subset of $[n]$ is realized exactly once as the image of some $\iota \in \text{ChosenInj}([k], [n])$. (For example, we could take $\text{ChosenInj}([k], [n])$ to be the monotone injections.) Similarly, pick

$$\text{ChosenColumnInj}([k], [n]) \subseteq \{\iota: [k] \hookrightarrow [n] \text{ such that } \iota(k) = n\}.$$

Theorem 5.3.3. *If λ is a non-column partition of k , then for each function $f: [n] \rightarrow [m]$, build the $\text{ChosenInj}([k], [m]) \times \text{ChosenInj}([k], [n])$ block matrix $X_\lambda(f)$ with (κ, ι) -entry*

$$X_\lambda(f)_{\kappa, \iota} = \begin{cases} W_\lambda \sigma & \text{if } f \circ \iota = \kappa \circ \sigma \text{ for some } \sigma \in S_k \\ 0 & \text{otherwise.} \end{cases}$$

If λ is a column of height k , build the $\text{ChosenColumnInj}([k], [m]) \times \text{ChosenColumnInj}([k], [n])$ matrix $X_\lambda(f)$ with (κ, ι) -entry

$$X_\lambda(f)_{\kappa, \iota} = \begin{cases} \text{sgn}(\sigma) & \text{if } g = \kappa \circ \sigma \text{ for some } \sigma \in S_k \\ 0 & \text{otherwise,} \end{cases}$$

where $g: [k] \rightarrow [m]$ denotes the function

$$g(p) = \begin{cases} m & \text{if } f(\iota(p)) \notin \text{im}(\kappa) \\ f(\iota(p)) & \text{otherwise.} \end{cases}$$

The assignment $f \mapsto X_\lambda(f)$ is an irreducible representation of \mathcal{D} isomorphic to V_λ .

Proof. If we argue that these are irreducible representations and give the right S_k -representation when evaluated on $[k]$, the required isomorphism will follow by Yoshioka's theorem (Proposition 3.1.10).

Let us begin with the case where λ is not a column. It is straightforward to check that $X_\lambda(1) = 1$. Given functions $f: [a] \rightarrow [b]$ and $g: [b] \rightarrow [c]$, we verify that $X_\lambda(g \circ f) = X_\lambda(g) \circ X_\lambda(f)$ entry by entry. For all $\alpha \in \text{ChosenInj}([k], [a])$ and $\gamma \in \text{ChosenInj}([k], [c])$, we must check that

$$X_\lambda(g \circ f)_{\gamma, \alpha} = \sum_{\beta \in \text{ChosenInj}([k], [b])} X_\lambda(g)_{\gamma, \beta} \circ X_\lambda(f)_{\beta, \alpha}.$$

Each summand on the right vanishes unless $g \circ \beta = \gamma \circ \sigma$ and $f \circ \alpha = \beta \circ \tau$ for some

uniquely-determined $\sigma, \tau \in S_k$. Since fixing γ and α fixes $\beta = f \circ \alpha \circ \tau^{-1}$, there is only one term in the sum. We have $g \circ f \circ \alpha = g \circ \beta \circ \tau = \gamma \circ \sigma \circ \tau$, and so the claim follows because $W_\lambda \sigma \circ W_\lambda \tau = W_\lambda(\sigma \circ \tau)$.

Still assuming λ is not a column, let us verify irreducibility. Certainly $X_\lambda[k] \simeq W_\lambda$, by construction. We also see that the representation X_λ is generated by any vector in $X_\lambda[k]$ since any vector with nonzero values in a single block inside $X_\lambda[l]$ is the image of the corresponding vector in $X_\lambda[k]$ under the corresponding chosen injection. By Proposition 5.3.2, any subrepresentation $X' \subseteq X_\lambda$ generated by vectors in $X'[k+1]$ that are fixed by θ_{k+1} must be generated by vectors in $X'[k]$. Since $X_\lambda[k]$ is irreducible, it suffices to show that every vector in $X_\lambda[k+1]$ is fixed by θ_k —equivalently, that every vector in $X_\lambda[k+1]$ is killed by ε_{k+1} .

Let $v \in X_\lambda[k]$ be a nonzero vector. Since X_λ is generated by v , the vector space $X_\lambda[k+1]$ is spanned by vectors of the form $(X_\lambda f)(v)$ as f ranges over the set $\text{Hom}_{\mathcal{D}}([k], [k+1])$. Since λ is not a column, $(V_\lambda \theta_k)(v) = v$. So $(X_\lambda \varepsilon_{k+1} \circ X_\lambda f)(v) = (X_\lambda \varepsilon_{k+1} \circ X_\lambda f \circ X_\lambda \theta_k)(v)$. In a moment we shall see that

$$\text{Hom}_{\mathcal{C}}(\theta_k, \varepsilon_{k+1}) = 0$$

and so ε_{k+1} acts by zero on a spanning set for $X_\lambda[k+1]$. We conclude the proof by showing that this hom-vector-space is zero. Let $\sum_{f \in \text{Hom}_{\mathcal{D}}([k], [k+1])} \alpha_f \cdot f$ be a linear combination of morphisms from \mathcal{D} that is fixed under precomposition by θ_k and postcomposition by ε_{k+1} . If f is non-injective, then $\alpha_f = 0$; indeed, otherwise the image of f misses two points, and so f is fixed under postcomposition with some transposition. Since all injections $[k] \hookrightarrow [k+1]$ are in the same free orbit under postcomposition by elements of the symmetric group S_{k+1} , every $\alpha_f = \pm \alpha_{f'}$ depending on the sign of the permutation $\sigma \in S_{k+1}$ for which $\sigma \circ f = f'$. In other words, $\dim \text{Hom}_{\mathcal{C}}(\theta_k, \varepsilon_{k+1}) \leq 1$. If $\iota: [k] \hookrightarrow [k+1]$ is some chosen injection, we must prove

that the vector $\partial = \sum_{\sigma \in \mathcal{S}_{k+1}} \text{sgn}(\sigma) \cdot (\sigma \circ \iota)$ is killed by precomposition with θ_k . It suffices to check that ∂ is fixed by precomposition with ε_k ; this is clear since precomposition by a transposition $(a \ b)$ matches postcomposition by the transposition $(\iota(a) \ \iota(b))$.

In the case where λ is a column of height k , we produce an irreducible representation D_k and claim that the matrices X_λ describe D_k in a natural basis. Define $D_k[n] \subseteq \text{Mat}_{\mathbb{Q}}^{\mathcal{D}}([k-1] \rightarrow [n])$ by the formula

$$D_k[n] = \varepsilon_{k-1} \cdot ([k-1] \subseteq [k]) \cdot \varepsilon_k \cdot \text{Mat}_{\mathbb{Q}}^{\mathcal{D}}([k], [n]),$$

where a function $f: [n] \rightarrow [m]$ acts by postcomposition. Each $\iota \in \text{ChosenColumnInj}([k], [n])$ corresponds to $\varepsilon_{k-1} \cdot ([k-1] \subseteq [k]) \cdot \varepsilon_k \cdot \iota$. A routine computation verifies that these vectors form a basis and transform according to the entries of X_λ .

To see that D_k is irreducible, note that D_k is a subrepresentation of $\text{Mat}_{\mathbb{Q}}^{\mathcal{D}}([k-1], -)$ which is generated by the matrix $1_{[k-1]} \in \text{Mat}_{\mathbb{Q}}^{\mathcal{D}}([k-1], [k-1])$ and so any subrepresentation $U \subseteq D_k$ is generated by vectors in $U[k]$ since $\mu([k-1]) = \{[k]\}$. But $D_k[k]$ is one-dimensional, and so $U = D_k$ or else $U = 0$. \square

5.3.4 What we know after constructing the irreducibles

Since Guess 5.2.6 holds, we have proved the following result.

Theorem 5.3.4. *If V is a finitely generated representation of \mathcal{D} , then the function $n \mapsto \dim V[n]$ coincides with a polynomial in n for $n > 0$.*

Proof. By Theorem 5.2.3, V has a finite composition series; by Theorem 5.3.3, its composition factors have the polynomiality property. \square

Surprisingly, we also get a converse—a numerical test for finite generation.

Theorem 5.3.5. *If V is a representation of \mathcal{D} so that the the function $n \mapsto \dim V[n]$ is bounded above by a polynomial in n , then V is finitely generated.*

Proof. We prove the formally stronger (but actually equivalent by Theorem 5.2.3) fact that V has finite length. Any increasing chain of subrepresentations of V gives a pointwise-nondecreasing sequence of nonnegative integer-valued polynomials. But there are only finitely many nonnegative integer-valued polynomials that never exceed a given nonnegative integer-valued polynomial, and so any chain of distinct subrepresentations must be finite. \square

Finally, we know a great deal about “representation stability” phenomena for representations of \mathcal{D} . If λ is not a column and $|\lambda| = k$, then our construction of the irreducibles gives that $X_\lambda[n] \simeq \text{Ind}_{S_k}^{S_n} W_\lambda$, which may be computed by Pieri’s rule. In other words, the story for non-column partitions perfectly matches the representation theory of $\text{FI}\#$. If λ is a column, then we obtain a sequence of hooks in the same way we would for FI (see [CEF15]).

5.3.5 What’s left?

Returning to the theme of Chapter 1, we may ask for an algorithm that takes a presentation matrix over \mathcal{D} and returns the composition factors of its representation and their multiplicities. This is enough information to read off the dimension polynomial or even the symmetric group characters.

In other words, we wish to apply the multiplicity Theorem 4.3.5.

5.4 Preparing to apply the Multiplicity Theorem 4.3.5

5.4.1 Experimentally computing the indecomposable injectives

Any minimal idempotent $\pi \in \mathbb{Q} \cdot \text{Hom}_{\mathcal{D}}(d, d)$ gives rise to an indecomposable injective representation using the formula

$$x \mapsto \{\varphi \in \mathbb{Q} \cdot \text{Hom}_{\mathcal{D}}(x, d) \mid \pi \circ \varphi = \varphi\}$$

where a function $f: x \rightarrow y$ acts by $\varphi \mapsto \varphi \circ f$. It is a general fact that every indecomposable injective has a unique irreducible subrepresentation, and moreover, that every indecomposable injective is determined up to isomorphism by this subrepresentation. In other words, any indecomposable injective is the injective hull of some irreducible representation. Write I_λ for the injective hull of V_λ .

Let's try the recipe with $\lambda = \square\square$. We start with the Young symmetrizer (some multiple of which is idempotent)

$$\text{YoungSymmetrizer}(\square\square) = \boxed{12} + \boxed{21} \in \mathbb{Q} \cdot \text{Hom}_{\mathcal{D}}([2], [2]).$$

Since the algebra $\mathbb{Q} \cdot \text{Hom}_{\mathcal{D}}([2], [2])$ is Artinian, any such idempotent in the quotient algebra $\mathbb{Q} \cdot S_2$ has a lift which is a minimal idempotent of $\mathbb{Q} \cdot \text{Hom}_{\mathcal{D}}([2], [2])$. Such a lift may be obtained computationally using the MeatAXE algorithm, for example, although this case requires no tools:

$$\pi = \frac{1}{2}(\boxed{12} + \boxed{21} - \boxed{11} - \boxed{22}).$$

In order to compute $\dim I_{\square\square}[3]$, for example, we must compute the span of the eight vectors

$$\{\pi \circ \boxed{111}, \pi \circ \boxed{112}, \pi \circ \boxed{121}, \pi \circ \boxed{122}, \pi \circ \boxed{211}, \pi \circ \boxed{212}, \pi \circ \boxed{221}, \pi \circ \boxed{222}\}$$

which, after expanding, take the form

$$\{0, -\boxed{111} + \boxed{112} + \boxed{221} - \boxed{222}, -\boxed{111} + \boxed{121} + \boxed{212} - \boxed{222}, \\ -\boxed{111} + \boxed{122} + \boxed{211} - \boxed{222}, -\boxed{111} + \boxed{122} + \boxed{211} - \boxed{222}, \\ -\boxed{111} + \boxed{121} + \boxed{212} - \boxed{222}, -\boxed{111} + \boxed{112} + \boxed{221} - \boxed{222}, 0\},$$

and after row-reducing,

$$\begin{aligned} & \{\boxed{111} - \boxed{122} - \boxed{211} + \boxed{222}, \\ & \boxed{112} - \boxed{122} - \boxed{211} + \boxed{221}, \\ & \boxed{121} - \boxed{122} - \boxed{211} + \boxed{212}\}. \end{aligned}$$

It follows that $\dim I_{\square\square}[3] = 3$. A few similar calculations give a table:

n	0	1	2	3	4	5
$\dim I_{\square\square}[n]$	0	0	1	3	7	15

Table 5.7: Dimensions of the indecomposable injective corresponding to $\square\square$

Probably the formula is $n \mapsto 2^{n-1} - 1$, at least for $n > 0$. We try the partitions of 3:

$\square\square\square$	0	1	2	3	4	5
$\square\square$	0	0	0	1	6	25
\square	0	0	0	2	12	50
\square	0	0	1	4	13	40

Table 5.8: Dimensions of the indecomposable injectives for partitions of 3

A bit of poking around in the OEIS [Slo] leads us to a conjecture.

Guess 5.4.1. If λ is a partition of k , the indecomposable injective I_λ satisfies

$$\dim I_\lambda[n] = \begin{cases} \dim W_\lambda \cdot \{^n_k\} & \text{for } \lambda \text{ not a column} \\ \{^n_k\} + \{^n_{k-1}\} & \text{for } \lambda \text{ a column,} \end{cases}$$

where $\{^n_k\}$ denotes a Stirling number of the second kind: the number of partitions of the set $[n]$ into k nonempty disjoint subsets.

5.4.2 Construction of the indecomposable injectives

As before, Guess 5.4.1 suggests a form for our matrices. We hope for block matrices where the rows and columns are indexed by k -element *quotients* of the set $\{1, \dots, n\}$ and the entries are still coming from W_λ . In the column case, we hope for rows and columns to be indexed by k - and $(k-1)$ -element quotients with entries coming from the sign representation of S_k .

Taking a cue from the construction of the irreducibles, pick some favored surjections $\text{ChosenSurj}([n], [k]) \subseteq \text{Hom}_{\mathcal{D}}([n], [k])$ so that any function $f: [n] \rightarrow [m]$ factors through exactly one element of $\text{ChosenSurj}([n], [\#(\text{im } f)])$. For example, we may choose the surjections $\zeta: [n] \rightarrow [k]$ with the property that $\min \zeta^{-1}(i)$ is an increasing sequence for $1 \leq i \leq k$.

Theorem 5.4.2. *If λ is a non-column partition of k , then for each function $f: [n] \rightarrow [m]$, build the $\text{ChosenSurj}([m], [k]) \times \text{ChosenSurj}([n], [k])$ block matrix $Y_\lambda(f)$ with (ζ, ξ) -entry*

$$Y_\lambda(f)_{\zeta, \xi} = \begin{cases} W_\lambda \sigma & \text{if } \zeta \circ f = \sigma \circ \xi \text{ for some } \sigma \in S_k \\ 0 & \text{otherwise.} \end{cases}$$

If λ is a column of height k , build the $(\text{ChosenSurj}([m], [k]) \sqcup \text{ChosenSurj}([m], [k-1])) \times (\text{ChosenSurj}([n], [k]) \sqcup \text{ChosenSurj}([n], [k-1]))$ matrix $Y_\lambda(f)$ with (ζ, ξ) -entry

$$Y_\lambda(f)_{\zeta, \xi} = \begin{cases} \text{sign}(\sigma) & \text{if } |\text{im}(\zeta)| = k \text{ and } \xi \circ f = \sigma \circ \zeta \text{ for some } \sigma \in S_k \\ \text{sign}(\tau) & \text{if } |\text{im}(\zeta)| < k \text{ and } \xi \circ f = \tau \circ \zeta \text{ for some } \tau \in S_{k-1} \\ \text{sign}(\sigma) & \text{if } |\text{im}(\zeta)| < k, \xi \circ f = \sigma \circ \iota \circ \zeta \text{ for some } \sigma \in S_k, \\ & \text{where } \iota: [k-1] \subseteq [k] \text{ denotes the natural inclusion} \\ 0 & \text{otherwise.} \end{cases}$$

Then the assignment $f \mapsto Y_\lambda(f)$ is a representation of \mathcal{D} isomorphic to the indecom-

possible injective I_λ which is the injective hull of the irreducible V_λ .

Proof. Let us prove that Y_λ is a representation when λ is not a column. It is straightforward to check that $Y_\lambda(1) = 1$. Given functions $f: [a] \rightarrow [b]$ and $g: [b] \rightarrow [c]$, let us verify that $Y_\lambda(g \circ f) = Y_\lambda(g) \circ Y_\lambda(f)$ entry-by-entry. For all $\alpha \in \text{ChosenSurj}([a], [k])$ and $\gamma \in \text{ChosenSurj}([c], [k])$, we must check that

$$Y_\lambda(g \circ f)_{\gamma, \alpha} = \sum_{\beta \in \text{ChosenSurj}([b], [k])} Y_\lambda(g)_{\gamma, \beta} \circ Y_\lambda(f)_{\beta, \alpha}.$$

The summand associated to β vanishes unless $\gamma \circ g = \sigma \circ \beta$ and $\beta \circ f = \tau \circ \alpha$ for some uniquely defined $\sigma, \tau \in S_k$. There can be at most one non-vanishing summand because $\beta = \sigma^{-1} \circ \gamma \circ g$ is determined. In this case, we have $\gamma \circ g \circ f = \sigma \circ \beta \circ f = \sigma \circ \tau \circ \alpha$, and so Y_λ is a representation since $W_\lambda(\sigma \circ \tau) = W_\lambda \sigma \circ W_\lambda \tau$.

We set about proving that I_λ is an indecomposable injective. By Proposition 5.3.2 and Corollary 4.1.12, inflation along the surjection $p_k: \mathcal{D}_k \rightarrow \mathcal{A}_k$ preserves injectivity of representations of \mathcal{A}_k on which θ_k acts by the identity. Since W_λ is always injective because $\mathcal{A}_k \simeq \mathbb{Q}S_k$ is semisimple, and θ_k acts by the identity whenever $\lambda \in \Lambda$ is not a column, we have an injective \mathcal{D}_k -representation $(p_k)^*W_\lambda$ by inflation for non-column λ . By Corollary 3.1.7, this inflated representation is still irreducible, in particular indecomposable. It follows by Theorem 3.3.4 that the right Kan extension $(i_k)_*(p_k)^*W_\lambda$ is an injective hull for the irreducible V_λ . The general formula for right Kan extension (Proposition 2.6.7)

$$((i_k)_*(p_k)^*W_\lambda)[n] = ((p_k)^*W_\lambda) \otimes_{\mathcal{D}_k} \text{Hom}_{\mathcal{D}}([n], [k])^*$$

simplifies in this case because maps $[k] \rightarrow [k]$ that factor through $[k-1]$ act by zero on $(p_k)^*W_\lambda$, which means that

$$(p_k)^*W_\lambda \otimes_{\mathcal{D}_k} \mathcal{A}_k \otimes_{\mathcal{A}_k} \mathcal{A}_k \simeq (p_k)^*W_\lambda,$$

as right \mathcal{D}_k -modules, and so

$$((i_k)_*(p_k)^*W_\lambda)[n] \simeq W_\lambda \otimes_{\mathcal{A}_k} \mathcal{A}_k \otimes_{\mathcal{D}_k} \text{Hom}_{\mathcal{D}}([n], [k])^*.$$

Since $\mathcal{A}_k \otimes_{\mathcal{D}_k} \text{Hom}_{\mathcal{D}}([n], [k])^* \simeq \text{Surj}([n] \rightarrow [k])^*$ is free as a $\mathbb{Q}S_k$ -module, we see that this tensor product may be computed by first picking orbit representatives for the free S_k -action on $\text{Surj}([n] \rightarrow [k])$. This step was accomplished earlier by picking $\text{ChosenSurj}([n], [k]) \subseteq \text{Surj}([n], [k])$. The rest of the construction mirrors the formation of this tensor product, and we are done.

In the case where λ is a column of height k , we produce an indecomposable injective representation J_k and claim that the matrices Y_λ describe J_k in a natural basis. Define $J_k[n] \subseteq \text{Mat}_{\mathbb{Q}}^{\mathcal{D}}([n] \rightarrow [k])$ the summand of the basic injective I_k given by the indecomposable idempotent ε_k

$$J_k[n] = \text{Mat}_{\mathbb{Q}}^{\mathcal{D}}([n] \rightarrow [k]) \cdot \varepsilon_k$$

where a function $f: [m] \rightarrow [n]$ acts by precomposition. Each $\zeta \in \text{ChosenSurj}([n], [k])$ corresponds to $\zeta \cdot \varepsilon_k$. A routine computation verifies that these vectors form a basis and transform according to the entries of Y_λ . \square

5.5 The multiplicity theorem for presentations

In this section, we switch to the complex numbers for consistency with Theorem 4.3.5. As it happens, all these computations will give the correct answer over \mathbb{Q} as well because the representation theory of S_k is defined over \mathbb{Q} .

5.5.1 Demonstration computation

Let $M \in \text{Mat}_{\mathbb{C}}^{\mathcal{D}}(x^{\oplus} \rightarrow y^{\oplus})$ be a presentation matrix with entries from \mathcal{D} and coefficients in \mathbb{C} . Recall that any such matrix gives rise to a representation $V_M \in \text{Mod}_{\mathbb{C}}^{\mathcal{D}}$ by the formula

$$V_M d = \frac{\text{Mat}_{\mathbb{C}}^{\mathcal{D}}(x^{\oplus} \longrightarrow d)}{M \cdot \text{Mat}_{\mathbb{C}}^{\mathcal{D}}(y^{\oplus} \rightarrow d)},$$

where a morphism $f \in \text{Hom}_{\mathcal{D}}(d, d')$ acts by postcomposition. The following result specializes Theorem 4.3.5 to the category of finite sets.

Theorem 5.5.1. *The multiplicity of the irreducible representation V_{λ} as a composition factor of the finitely presented representation V_M is given by the \mathbb{C} -corank of the block matrix $I_{\lambda}(M)$ obtained by applying the indecomposable injective I_{λ} to the presentation matrix M entrywise.*

No theorem could be easier to use. We demonstrate with the presentation matrix from §1.2.3

$$M = \begin{array}{cccc} & [4] & [3] & [4] & [3] \\ [4] & \left[\boxed{1234} - \boxed{2134} \quad \boxed{1123} \quad \boxed{1234} + \boxed{3412} \quad \boxed{1223} + \boxed{2331} + \boxed{3112} \right] \end{array},$$

an element of $\text{Mat}_{\mathbb{C}}^{\mathcal{D}}([4] \rightarrow [4] \oplus [3] \oplus [4] \oplus [3])$. Let us compute the multiplicity of the irreducible representation V_{\square} inside V_M . By the theorem, this multiplicity is given by the corank of the block matrix $I_{\square}(M)$. Using the explicit construction of I_{\square} from

Theorem 5.4.2 we obtain

$$\begin{array}{c}
 \boxed{1222} \quad \boxed{1122} \quad \boxed{2122} \quad \boxed{2221} \quad \boxed{1221} \quad \boxed{2212} \quad \boxed{1212} \quad \boxed{122} \quad \boxed{221} \quad \boxed{212} \quad \boxed{1222} \quad \boxed{1122} \quad \boxed{2122} \quad \boxed{2221} \quad \boxed{1221} \quad \boxed{2212} \quad \boxed{1212} \quad \boxed{122} \quad \boxed{221} \quad \boxed{212} \\
 \boxed{1222} \left[\begin{array}{cccccccc|cccc|cccc|cccc|cccc}
 1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & 1 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 -1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\
 0 & 0 & 0 & 0 & 1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 & 1 & 1 & 1 & 1 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & -1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 & 0 & 0 & 0
 \end{array} \right]
 \end{array}$$

a matrix of corank 0. It follows that V_M contains no copy of V_{\square} as a subquotient. A similar calculation carried out for all partitions computes the full K-class

$$[V_M] = [V_{\square}] + [V_{\square\square}].$$

This computation is actually finite because almost all I_λ vanish on the sets [3] and [4] by the dimension count. More generally, we could appeal to the homological modulus for an upper bound as explained in Theorem 4.3.5. By Theorem 3.3.10, we see that the dimension of V_M is given by the formula

$$\dim V_M[n] = 2 \cdot \binom{n}{3} + 3 \cdot \binom{n}{4}.$$

5.5.2 A computer program

The explicit construction for I_λ is easy to code. Let us run the computation on the example from §1.2.2 asking for the number of n -colorings of a five-cycle up to rotation and reflection.

```
In: coker([5], "[[aabcd, abbcd, abccd, abcd,
abcda, abcde-bcdea, abcde-aedcb]]", [4, 4, 4, 4, 4, 5, 5])
```

```

Out: This finitely presented representation decomposes into \\  

irreducible representations as follows:  

[coker] = 1*[2, 1] + 1*[3] + 1*[2, 1, 1] + 1*[2, 2] \\  

+ 2*[3, 1] + 1*[4] + 1*[1, 1, 1, 1, 1] \\  

+ 1*[2, 2, 1] + 1*[3, 2] + 1*[5] \\  

+ 1*[1, 1, 1, 1, 1, 1].  

In particular, the dimension is given \\  

as a function of n > 0:  

(1/10) * (n - 2) * (n - 1) * n * (n^2 - 2*n + 2).  

Total computation time: 5.51840806007 seconds

```

The source code for this program will be available on the author's website.

5.6 What can be said over \mathbb{Z}

Until now, we have not hesitated to assume that R to be Artinian, or even some convenient field like \mathbb{Q} or \mathbb{C} . Still, if we plan to apply these results to problems of algebraic topology, we must prepare. In Definition 4.4.1 we give the notion of a regular function on a category; in this case, the regular functions restrict to integer-valued, nonnegative polynomials on the sets $[1], [2], \dots$ (This fact is stronger than the fact that the irreducible representations of \mathcal{D} have polynomial dimension over any field). The following result is an immediate consequence of Corollary 4.4.5 where we know the form of the \mathcal{D} -regular functions by Example 4.4.2.

Theorem 5.6.1. *For every finitely generated representation $V: \mathcal{D} \rightarrow \mathbf{Mod}_{\mathbb{Z}}$ there exists a finite sequence of prime powers $q_1, \dots, q_k \in \mathbb{Z}$ with the property that*

$$V[n] \simeq \mathbb{Z}^{\oplus \varphi_0(n)} \oplus (\mathbb{Z}/q_1)^{\oplus \varphi_1(n)} \oplus \dots \oplus (\mathbb{Z}/q_k)^{\oplus \varphi_k(n)}$$

for certain integer-valued, nonnegative polynomials $\varphi_0, \varphi_1, \dots, \varphi_k$ in $n > 0$.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [Arn69] V.I. Arnol'd. The cohomology ring of the colored braid group. *Mathematical notes of the Academy of Sciences of the USSR*, 5(2):138–140, 1969. Available from: <http://dx.doi.org/10.1007/BF01098313>, doi:10.1007/BF01098313.
- [Aus66] Maurice Auslander. Coherent functors. In S. Eilenberg, D.K. Harrison, S. MacLane, and H. Rhl, editors, *Proceedings of the Conference on Categorical Algebra*, pages 189–231. Springer Berlin Heidelberg, 1966. Available from: http://dx.doi.org/10.1007/978-3-642-99902-4_8, doi:10.1007/978-3-642-99902-4_8.
- [AW92] William A. Adkins and Steven H. Weintraub. *Algebra*, volume 136 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1992. An approach via module theory.
- [BBD82] A. A. Beilinson, J. Bernstein, and P. Deligne. Faisceaux pervers. In *Analysis and topology on singular spaces, I (Luminy, 1981)*, volume 100 of *Astérisque*, pages 5–171. Soc. Math. France, Paris, 1982.
- [BT15] Serge Bouc and Jacques Thévenaz. The representation theory of finite sets and correspondences. 2015. [arXiv:1510.03034](https://arxiv.org/abs/1510.03034).
- [CEF14] Thomas Church, Jordan S. Ellenberg, and Benson Farb. Representation stability in cohomology and asymptotics for families of varieties over finite fields. In *Algebraic topology: applications and new directions*, volume 620 of *Contemp. Math.*, pages 1–54. Amer. Math. Soc., Providence, RI, 2014. doi:10.1090/conm/620/12395.
- [CEF15] Thomas Church, Jordan S. Ellenberg, and Benson Farb. FI-modules and stability for representations of symmetric groups. *Duke Math. J.*, 164(9):1833–1910, 2015.
- [CEFN14] Thomas Church, Jordan S. Ellenberg, Benson Farb, and Rohit Nagpal. FI-modules over Noetherian rings. *Geom. Topol.*, 18(5):2951–2984, 2014.
- [CF13] Thomas Church and Benson Farb. Representation theory and homological stability. *Adv. Math.*, 245:250–314, 2013. doi:10.1016/j.aim.2013.06.016.

- [Chu12] Thomas Church. Homological stability for configuration spaces of manifolds. *Invent. Math.*, 188(2):465–504, 2012. doi:10.1007/s00222-011-0353-4.
- [CP61] A. H. Clifford and G. B. Preston. *The algebraic theory of semigroups. Vol. I.* Mathematical Surveys, No. 7. American Mathematical Society, Providence, R.I., 1961.
- [Dev14] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 6.6)*, 2014. <http://www.sagemath.org>.
- [EGH⁺11] Pavel Etingof, Oleg Golberg, Sebastian Hensel, Tiankai Liu, Alex Schwendner, Dmitry Vaintrob, and Elena Yudovina. *Introduction to representation theory*, volume 59 of *Student Mathematical Library*. American Mathematical Society, Providence, RI, 2011. With historical interludes by Slava Gerovitch.
- [EWG15] Jordan S. Ellenberg and John D. Wiltshire-Gordon. Algebraic structures on cohomology of configuration spaces of manifolds with flows. 2015. arXiv:1508.02430.
- [Git15] Andrew Gitlin. New examples of dimension zero categories and a computer program for calculations in representation theory, 2015. Available from: https://www.lsa.umich.edu/UMICH/math/Home/Undergrad/Research%20and%20Career%20Opportunities/Research/Past%20REU%20Projects/Gitlin_Andrew%20REU%202015.pdf.
- [GL] Wee Liang Gan and Liping Li. On central stability. arXiv:1504.07675.
- [GS] Daniel R. Grayson and Michael E. Stillman. Macaulay2, a software system for research in algebraic geometry. Available at <http://www.math.uiuc.edu/Macaulay2/>.
- [JR11] Rita Jimenez Rolland. Representation stability for the cohomology of the moduli space M_g^n . *Algebr. Geom. Topol.*, 11(5):3011–3041, 2011. doi:10.2140/agt.2011.11.3011.
- [Kos99] Masashi Kosuda. The irreducible representations of the Hecke category. *J. Algebra*, 215(1):135–184, 1999. doi:10.1006/jabr.1998.7732.
- [Kov92] L. G. Kovács. Semigroup algebras of the full matrix semigroup over a finite field. *Proc. Amer. Math. Soc.*, 116(4):911–919, 1992.
- [Kra02a] Henning Krause. A Brown representability theorem via coherent functors. *Topology*, 41(4):853–861, 2002. doi:10.1016/S0040-9383(01)00010-6.
- [Kra02b] Henning Krause. Coherent functors in stable homotopy theory. *Fund. Math.*, 173(1):33–56, 2002. doi:10.4064/fm173-1-3.

- [Kra03] Henning Krause. Coherent functors and covariantly finite subcategories. *Algebr. Represent. Theory*, 6(5):475–499, 2003. doi:10.1023/B:ALGE.0000006492.02381.df.
- [Kra14] Henning Krause. Krull-schmidt categories and projective covers, 2014. arXiv:1410.2822.
- [KS06] Masaki Kashiwara and Pierre Schapira. *Categories and sheaves*, volume 332 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2006. doi:10.1007/3-540-27950-4.
- [Kuh15] Nicholas J. Kuhn. Generic representation theory of finite fields in non-describing characteristic. *Adv. Math.*, 272:598–610, 2015.
- [Mit72] Barry Mitchell. Rings with several objects. *Advances in Math.*, 8:1–161, 1972.
- [ML98] Saunders Mac Lane. *Categories for the working mathematician*, volume 5 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1998.
- [MR01] J. C. McConnell and J. C. Robson. *Noncommutative Noetherian rings*, volume 30 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, revised edition, 2001. With the cooperation of L. W. Small. doi:10.1090/gsm/030.
- [Ner91] Yu. A. Neretin. Extension of representations of classical groups to representations of categories. *Algebra i Analiz*, 3(1):176–202, 1991.
- [NSS16] Rohit Nagpal, Steven V. Sam, and Andrew Snowden. Noetherianity of some degree two twisted commutative algebras. *Selecta Math. (N.S.)*, 22(2):913–937, 2016. doi:10.1007/s00029-015-0205-y.
- [Okn91] Jan Okniński. *Semigroup algebras*, volume 138 of *Monographs and Textbooks in Pure and Applied Mathematics*. Marcel Dekker, Inc., New York, 1991.
- [PR02] T. Pirashvili and B. Richter. Hochschild and cyclic homology via functor homology. *K-Theory*, 25(1):39–49, 2002.
- [Put96] Mohan S. Putcha. Complex representations of finite monoids. *Proc. London Math. Soc. (3)*, 73(3):623–641, 1996. doi:10.1112/plms/s3-73.3.623.
- [Put13] Andrew Putman. The Torelli group and congruence subgroups of the mapping class group. In *Moduli spaces of Riemann surfaces*, volume 20 of *IAS/Park City Math. Ser.*, pages 169–196. Amer. Math. Soc., Providence, RI, 2013.

- [Wil12] Jennifer C. H. Wilson. Representation stability for the cohomology of the pure string motion groups. *Algebr. Geom. Topol.*, 12(2):909–931, 2012. doi:10.2140/agt.2012.12.909.
- [Yos93] Y. Yoshioka. Representation of the Jones category (in Japanese). Master's thesis, Osaka University, 1993.
- [Zel77] E. I. Zel'manov. Semigroup algebras with identities. *Sibirsk. Mat. Ž.*, 18(4):787–798, 956, 1977.