A BRIEF EXPLORATION OF A HOMOLOGICAL THEORY OF FUNCTIONS

CHASE MEADORS

1. Introduction

The goal is this article is to explore the main ideas of a paper by Yang [5] that develops a so-called "homological theory of functions". The functions spoken of here refer to finite functions and especially finite Boolean functions that are often studied in the Computer Science field of learning theory. Yang uses heavy machinery from the fields of Combinatorial Commutative Algebra to develop a way of speaking about classes of functions using the language of homological algebra.

In the following, we will first cover some background theory from Combinatorial Commutative Algebra and Stanley-Reisner Theory, then discuss some foundational theory and basic core results of [5].

2. Stanley-Reisner Theory

Stanley-Reisner theory is a central topic in Combinatorial Commutative Algebra. At its core, it concerns a duality between (combinatorial or abstract) simplicial complexes and squarefree monomial ideals. Definitions of these notions are to come, but first we fix a notational convention: Throughout this section and most of the article, k will be a field that will not play any major role in our theory, and $R = k[\dots]$ will be the polynomial ring in finitely many indeterminates. Exactly what variables are used for these indeterminates will change according to the context or example in play, but R will always refer to the polynomial ring over the current variable set.

Definition 2.1. A *squarefree monomial ideal* $I \subset k[x_1, \dots, x_n]$ is a ring ideal whose (unique) minimal generators are monomials containing no variables raised to a power higher than 1.

Definition 2.2. An (abstract) *simplicial complex* Δ on $X = \{x_1, \dots, x_n\}$ is a family of subsets of X (called *faces*) closed under subset. That is, if $A \subset B \in \Delta$ then $A \in \Delta$. Δ can be uniquely specificed by its maximal faces called *facets*.

Remark 2.3. We have glossed over a subtle point that monomial ideals have a unique set of minimal monomial generators; this is a basic fact from Combinatorial Commutative Algebra;

c.f. Lemma 1.2 of [2]. An abstract simplicial complex can be thought of as a combinatorial way of completely capturing the data of a topological simplicial complex.

Remark 2.4. Just as one writes an ideal as ⟨minimal generators⟩, we will write a simplicial complex as {facets}, where we really mean the "closure under subset".

In fact these two notions are dual to each other via the *Stanley-Resiner* correspondence. One may suspect as much after the following observations: An ideal is closed under "growing", in that one can obtain larger monomials in the ideal by taking multiples of the generators, while a simplicial complex is, by definition, closed under "shrinking", since the inclusion of a face necessitates the inclusion of all its subsets. We will repeatedly make use of the observation that square-free monomials in an ideal and faces in a simplicial complex are both essentially subsets of the ambient variable set, and we will not make a notational distinction.

Definition 2.5. Given a squarefree monomial ideal $I \subset R$, its *Stanley-Reisner complex* Δ_I has faces given by the monomials *not* in I. That is, $\Delta_I = \{m : m \notin I\}$. Its facets are are the maximal such monomials not in I.

Definition 2.6. Given a simplicial complex Δ, its *Stanley-Reisner ideal* I_Δ is generated by the non-faces of Δ, that is $I_\Delta = \langle m : m \notin \Delta \rangle$. Its minimal generators are the minimal such non-faces.

It is straightforward to verify that performing one operation after the other results in the original object.

Simplicial complexes posess another type of duality called Alexander Duality:

Definition 2.7. Given a simplicial complex Δ, its Alexander Dual $\Delta^\vee$ has faces equal to the complements of non-faces of Δ. That is, $\Delta^\vee = \{m^c : m \notin \Delta\}$. Its facets are the complements of the minimal non-faces.

Again this operation is self-inverse. Using Stanley-Reisner duality, we can define the Alexander Dual *of an ideal*:

Definition 2.8. Given a squarefree monomial ideal $I \subset R$, its *Alexander dual* is the Stanley Reisner dual of the Alexander dual of the Stanley-Reisner dual of I.

A picture is worth far more than the preceding definition. Figure 1 explains how these two notions are related; the bottom arrow is *defined* by following the other three arrows around the square.

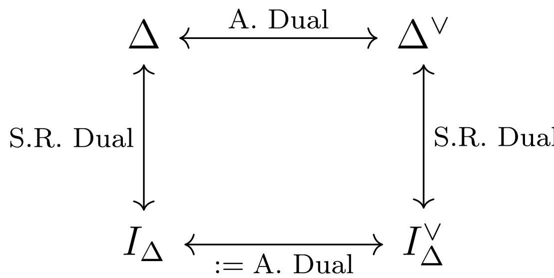

FIGURE 1. Stanley-Resiner duality and Alexander Duality

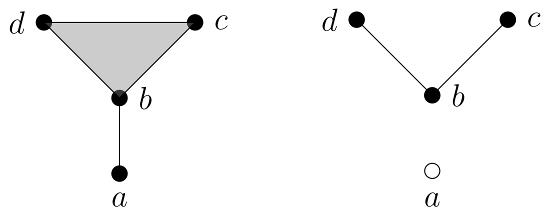

FIGURE 2. Δ and $\Delta^\vee$ in example 2.9

Example 2.9. Let the ambient variable set be $\{a, b, c, d\}$ and let $\Delta = \{ab, bcd\}$ (again we use monomial notation to refer to facets as well). Then $I_\Delta = \langle ac, ad\rangle \subset k[a, b, c]$. $\Delta^\vee = \{bc, bd\}$ and $I^\vee = \langle a, cd\rangle$. Figure 2 show Δ and $\Delta^\vee$.

Given the complex Δ, the ideal $I_{\Delta^\vee}$ is often important. It turns out that the minimal generators of this ideal are simply the complements of the facets of Δ:

Proposition 1. Given a complex Δ, $I_{\Delta^\vee} = \langle m^c : m \in \Delta\rangle$. Its minimal generators are the complements of the facets of Δ.

The interested reader is directed to consult [4] for a more thorough survey of Stanley-Reisner theory.

3. Free resolutions

Resolutions are constructions on modules studied heavily in commutative algebra and homological algebra. We will specialize some of the language here to our particular situation. Fix the ambient polynomial ring R. Note that R/I is a (non-free) R-module. A *free resolution* of R/I (in place of which we may just say a free resolution of I) is a way of measuring how far away this module is from being free. This is accomplished by fitting R/I into a long exact sequence of free R-modules.

Definition 3.1. A free resolution of I (i.e., a free resolution of R/I) is a long exact sequence

$$R^{\beta_\ell} \xrightarrow{d_\ell} \dots \xrightarrow{d_3} R^{\beta_2} \xrightarrow{d_2} R^{\beta_1} \xrightarrow{d_1} R \to R/I \to 0$$

where each integer β_i is the rank of the free R-module at that position.

It is a result of the Hilbert Syzygy Theorem that every finitely generated ideal in a polynomial ring has a finite minimal free resolution (here minimal means that the Betti numbers / ranks are minimal among all possible resolutions). Hence the following notions are well-defined:

Definition 3.2. Given a minimal free resolution of I, the integer β_i is called the i'th *Betti number* of I denoted $\beta_i(I)$. The integer ℓ is the *projective dimension* of I, denoted $\text{pd}(I)$.

This definition as given does not capture all the possible data it can about the resolution. In fact, there is a finer notion of Betti numbers called *multigraded* Betti numbers. Each Betti number is partitioned as:

$$\beta_i(I) = \sum_{m \in R} \beta_{i,m}(I)$$

where m ranges over all monomials in R (of course, $\beta_{i,m}$ is non-zero for finitely many m). $\beta_{i,m}$ is called the i'th Betti number of multidegree m. Another way to think of this is that we assign a multidegree (monomial) to the basis elements of each free R-module in the resolution. Though this can be made sense of formally by using the notion of $\mathbb{N}^n$-graded rings, modules, and resolutions (as is done in chapter 4 of [3]), we offer another definition of $\beta_{i,m}$. We assign a multidegree (monomial) to each basis element of the free modules in the resolution. Let $B_i = \left\{e_1^i, \ldots, e_{\beta_i}^i\right\}$ be a basis for R^{β_i} and denote by $\deg(e_j^i)$ its multidegree. The i'th Betti number in degree m becomes the count of the basis elements with that assigned degree:

$$\beta_{i,m}(I) = |\{e \in B_i : \deg(e) = m\}|$$

A way of defining this multidegree is as follows: Refer to the first free module R in the resolution as the *base module*. First set $\deg(e_j^1) = d_1(e_j^1)$ (d_1 sends each basis element to a monomial generator of I). For $n \geq 2$, consider the image of e_j^n in the base module by applying $d_n \circ \cdots \circ d_1$. Since the composition of at least two consecutive differentials is zero, this image is a polynomial which is equal to the zero polynomial. Thus, its monomial terms must all be identical, and this monomial will be $\deg(e_j^n)$. This will become more clear in an example.

Example 3.3 (c.f. [3], Example 4.4). Let $R = k[x,y]$ and $I = \langle x^3, xy, y^5 \rangle$. The following is a minimal free resolution of I:

$$R^2 \xrightarrow{\begin{pmatrix} -x^2 & -y^4 \\ y & 0 \\ 0 & x \end{pmatrix}} R^3 \xrightarrow{\begin{pmatrix} x\,y & x^3 & y^5 \end{pmatrix}} R \to R/I \to 0$$

as computed by Macaulay 2 [7] with code in code sample A.1. Hence the projective dimension of I is 2. Let the first free module R^3 have basis $\{f_1, f_2, f_3\}$ and the second free module R^2 have basis $\{e_1, e_2\}$. We have

$$\beta_1 = 3; \qquad \beta_{1,xy} = 1 \qquad \beta_{1,x^3} = 1 \qquad \beta_{1,y^5} = 1$$

Now we have

$$\begin{aligned} d_1(d_2(e_1)) &= d_1(-x^2 f_1 + y f_2) \\ &= -x^2 d_1(f_1) + y d_1(f_2) \\ &= -x^2(xy) + y(x^3) \\ &= -x^3 y + x^3 y \end{aligned}$$

hence $\deg(e_1) = x^3y$. Similarly,

$$(d_2 \circ d_1)(e_2) = -y^4(xy) + x(y^5) = -xy^5 + xy^5$$

and $\deg(e_2) = xy^5$. So,

$$\beta_2 = 2; \qquad \beta_{2,x^3y} = 1 \qquad \beta_{2,xy^5} = 1$$

And we have determined the multigraded Betti numbers of I.

Next, we discuss a common technique of finding free resolutions:

4. Cellular resolutions

Sometimes a resolution of I is provided by the reduced *cellular free complex* supported on some CW-complex. Recall that a CW-complex X comes with its *cellular chain complex* with groups

$$C_n = H_n(X_n, X_{n-1}) \cong \mathbb{Z}\{e_\sigma : \sigma \text{ an n-cell}\}$$

i.e., free $\mathbb{Z}$-modules over the n-cells of X, with boundary maps given by

$$\partial_n(e_\sigma) = \sum_{\tau \text{ (n-1)-cell}} [\sigma : \tau] e_\tau$$

where $[\sigma : \tau]$ is the degree of the attaching map. For all complexes we will consider (a subvariety called polytopal cell complexes), we may think of $[\sigma : \tau]$ as a sign coming from orientation. We can augment such a complex with a labeling obtained from I, which allows us to form an augmented chain complex called the *cellular free complex.*

Specifically, we label the vertices of X by the monomial generators of I. Each cell then inherits a label equal to the least common multiple of its boundaries. The *cellular free*

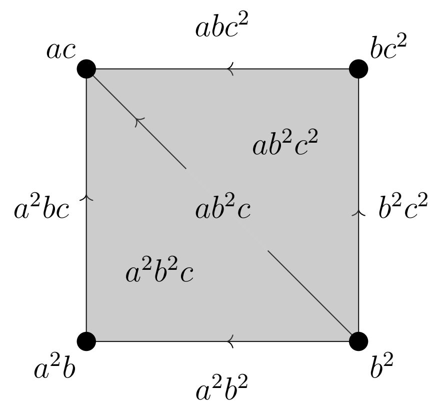

FIGURE 3. A labeled cell complex resolving I in Example 4.1

complex supported on X then consists of free R-modules:

$$C_n \cong R\{e_\sigma : \sigma \text{ an n-cell}\}$$

with boundary maps that incorporate the labels:

$$\partial_n(e_\sigma) = \sum_{\tau \text{ (n-1)-cell}} [\sigma : \tau] \frac{\ell(\sigma)}{\ell(\tau)} e_\tau$$

If the *reduced* cellular free complex has homology only in degree zero, then it is in fact a resolution of R/I (the image of its last augmentation map is the generated by the labels of the points!).

Example 4.1. Let $I = (a^2b, ac, bc^2, b^2) \subset k[a,b,c]$. A labeled cell complex that provides a cellular resolution is depicted in Figure 3. Its reduced cellular free complex is, as an explicit computation shows:

$$R^2 \xrightarrow{\begin{pmatrix} b & 0 \\ 0 & -b \\ 0 & -a \\ c & 0 \\ -a & c \end{pmatrix}} R^5 \xrightarrow{\begin{pmatrix} -c & 0 & 0 & b & 0 \\ a\,b & b\,c & 0 & 0 & b^2 \\ 0 & -a & b & 0 & 0 \\ 0 & 0 & -c^2 & -a^2 & -a\,c \end{pmatrix}} R^4 \xrightarrow{\begin{pmatrix} a^2b & a\,c & b\,c^2 & b^2 \end{pmatrix}} R^1 \to 0$$

One can further compute using Macaulay2 [7] with code in code sample A.2 that that homology of this chain complex is zero every except at the last free module R^1, where it is equal to R/I. Hence, by changing the end to $R \to R/I \to 0$ we obtain precisely a free resolution of I.

Cellular resolutions are desirable because one can infer a wealth of information about the Betti numbers of I just from the geometry and (plain) homology of the (plain) cell complex X. Denote by $X_{\preceq m}$ the subcomplex of X consisting of cells whose label divides m, and $X_{\prec m}$ similarly with labels strictly dividing m. We have the following collection of facts:

Proposition 2 (c.f. [2], Proposition 4.5)**.** The cellular free complex supported on X is a cellular resolution of R/I if and only if the (plain) complex $X_{\preceq m}$ is acyclic over k (has trivial reduced homology with coefficients in k) for all monomials m.

Proposition 3 (c.f. [5], Proposition 2.16)**.** A cellular resolution is minimal if, for each inclusion of cells $\sigma \subset \sigma'$, σ and σ' have distinct labels.

This is immediate from the following, noting that if there is such an inclusion of cells with the same label, the division of labels in the boundary map will result in a non-zero scalar entry in the boundary matrix:

Proposition 4 (c.f. [2], Definition 1.24)**.** A free resolution of I is minimal if and only if the only scalar entries in its boundary maps are 0.

Remark 4.2. This is taken as a definition of minimal in [2]. The fact that it is equivalent to the ranks of the modules being minimal is discussed therein.

Proposition 5. If the cellular resolution by X is minimal, then the i'th Betti number in degree m of I is the number of cells in X with label m.

Proof. Each n-cell σ has degree $\deg(e^n_\sigma)$ equal to its label $\ell(\sigma)$. The case of 0-cells is clear. Inductively, assume the the basis for the R-module generated by the $(n-1)$-dimensional faces e^{n-1}_σ have degree equal to their label $\ell(\sigma)$. The terms of $d_n(e^n_\sigma)$ are of the form $m_\tau e^{n-1}_\tau$, where $m_\tau \ell(\tau) = \ell(\sigma)$. But the degree of e^n_σ can be computed from any term of $d_n(e^n_\sigma)$ as $m \deg(e^{n-1}_\tau) = m\ell(\tau) = \ell(\sigma)$. Since the resolution is minimal and there are no duplicate labels, the claim holds. □

Since the Betti numbers of I are encoded in its minimal free resolution, one would expect that a non-minimal cellular resolution does not provide any useful information. However, this is not the case:

Proposition 6 (c.f. [2], Theorem 4.7)**.** If X is a labeled cell complex giving a cellular resolution of I, the Betti numbers of I are given by $\beta_{i,m}(I) = \dim_k \widetilde{H}_{i-1}(X_{\prec m}; k)$.

Thus even non-minimal cellular resolutions are useful for computation!

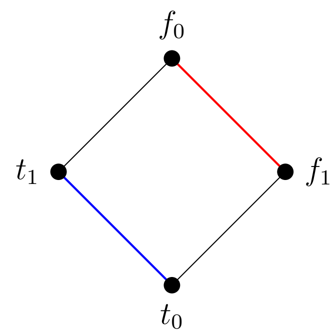

FIGURE 4. $\lozenge_{\mathtt{K}}$ for the complete class $\mathtt{K}$

5. The canonical ideal of a function class

We are now ready to discuss the main paper [5].

Remark 5.1. We will use somewhat different notation in the paper; this notation fails to capture some of the fully general theory of the paper, but will serve to "ground" the discussion for the subset of ideas we will cover.

The object of study in the paper is the following:

Definition 5.2. A *binary function class* or *concept class* is a family of finite functions contained in $[2]^{[n]}$

Here we use the convention that $[n] = \{0, 1, \ldots, n-1\}$. In other words, these are classes of finite functions taking n inputs with binary output.

The development starts by building a simplicial complex from the data of a function class:

Definition 5.3 (c.f. [5] Section 2.2)**.** For a concept class $\mathtt{C}$, the *canonical suboplex* $\lozenge_{\mathtt{C}}$ of $\mathtt{C}$ is a simplicial complex built as follows:

- $2n$ vertices given by t_i, f_i for each $i \in [n]$
- For each $g \in \mathtt{C}$, a facet $\{v_i(g) : i \in [n]\}$, where $v_i(g)$ is the *valuation of g at i*, equal to t_i if $g(i) = 1$ and f_i otherwise.

Remark 5.4. We may think of the facets as representing the "graphs" of functions; facets are then "glued together" along faces representing a partial function which is a common restriction of the two facets.

Example 5.5. Let $\mathtt{K}$ be the complete function class $[2]^{[2]}$. $\lozenge_{\mathtt{K}}$ is shown in Figure 4. The lower left blue line is the facet coming from the constant true function. The upper right red facet comes from the constant false function. The other two facets come from the other two functions in the class.

Immediately, Stanley-Resiner theory lets us define:

Definition 5.6 (c.f. [5], Definition 2.35)**.** The *Stanley-Resiner ideal* of a class $\mathtt{C}$, denoted $I_{\mathtt{C}}$, is the Stanley-Reisner ideal of $\Diamond_{\mathtt{C}}$.

Definition 5.7 (c.f. [5], Definition 2.35)**.** The *canonical ideal* of a class $\mathtt{C}$, denoted $I^{\star}_{\mathtt{C}}$, is the Alexander dual of $I_{\mathtt{C}}$ (following [5], we are using $\star$ for Alexander dual instead of $\vee$).

The philosophy of Stanley-Reisner theory is that thinking about any of $\Diamond_{\mathtt{C}}$, $I_{\mathtt{C}}$ or $I^{\star}_{\mathtt{C}}$ is "combinatorially equivalent", and we should exploit the one with the best description.

Remark 5.8. We will speak both of partial functions and total functions. Any function spoken of unqualified should be assumed to be total, as with the members of some concept class. We will emphasize whenever a function under consideration may be partial.

Denote by $\neg f$ the function taking on the opposite truth values as f. We will also need:

Definition 5.9. The *graph monomial* $\mathrm{gr}(p)$ of a (possibly partial) function p is given by $\prod_{i\in p} v_i(p)$, where $i \in p$ ranges over the i where p is defined.

Now we describe $I^{\star}_{\mathtt{C}}$ completely:

Theorem 5.10 (c.f. [5], following Definition 2.35)**.** $I^{\star}_{\mathtt{C}} = \langle \mathrm{gr}(\neg f) : f \in \mathtt{C}\rangle$ *(and these generators are minimal)*

Proof. This is a simple application of proposition 1. Using our new notation, the facets of $\Diamond_{\mathtt{C}}$ are $\mathrm{gr}(f)$ for $f \in \mathtt{C}$. Since f is total, the complement of such a facet in the ambient variable set $\{f_1, \ldots, f_n, t_1, \ldots, t_n\}$ is exactly $\mathrm{gr}(\neg f)$. □

Describing the Stanley-Resiner ideal $I_{\mathtt{C}}$ is a bit more complicated.

Definition 5.11 (c.f. [5], Proposition 2.38)**.** A partial function p is an *extenture* of the class $\mathtt{C}$ if p has *no* extension to a total function in $\mathtt{C}$, but every further restriction of p does.

Definition 5.12. We say that an input $i \in [n]$ is *trivial* for the class $\mathtt{C}$ if either all functions in $\mathtt{C}$ are true at i, or all functions in $\mathtt{C}$ are false at i. Else we call it *non-trivial.*

Theorem 5.13 (c.f. [5], Proposition 2.38)**.**

$$I_{\mathtt{C}} = \langle \mathrm{gr}(p) : p \text{ an extenture of } \mathtt{C},\ f_i t_i : i \text{ is non-trivial for } \mathtt{C}\rangle$$

(and these generators are minimal)

Proof. We must show that this set is exactly the set of minimal non-faces of $\lozenge_{\mathtt{C}}$. First consider the monomial $\mathrm{gr}(p)$ for an extenture p. It's not hard to see that the definition of extenture says that this is a minimal non-face of $\lozenge_{\mathtt{C}}$; p is not a face since p is not the restriction of any total $f \in \mathtt{C}$. However, any monomial strictly dividing $\mathrm{gr}(p)$ is the graph of a restriction of p, which does extend to some function in $\mathtt{C}$, so the graph of that restriction is a face. Second, consider $f_i t_i$ for i non-trivial. This cannot be a face since every face in $\lozenge_{\mathtt{C}}$ is the graph of a restriction of an $f \in \mathtt{C}$ (and as such the faces cannot contain f_i and t_i for the same i). But of course f_i and t_i are both faces since i was non-trivial and there are functions in $\mathtt{C}$ realizing both.

Conversely, suppose m is a minimal non-face of $\lozenge_{\mathtt{C}}$. Either m is the graph of a partial function p or not. If it is, then p must not be the restriction of some $f \in \mathtt{C}$ since it is not a face. But deleting any variable from m is also the graph of a restriction of p, and minimality says this must be a face, and so p must be the restriction of a function in $\mathtt{C}$; thus p is an extenture and m falls into the first set of generators. In the case that m is not the graph of a partial function, we must have $t_i, f_i \in m$ for some i. Deleting either t_i or f_i from m must produce a face by minimality. Since m/t_i is a face, it is the graph of a restriction of a function in $\mathtt{C}$, and likewise for m/f_i Since deleting one variable leaves the other, we can conclude immediately that i is non-trivial since we have restriction of two functions in $\mathtt{C}$ that are true at i and false at i, respectively. Finally, m cannot include t_j or f_j for $j \neq i$, if it did, deleting either one would produce a face, but the new monomial would still contain both f_i and t_i, and could not be the graph of a partial function. □

Example 5.14. For our complete class $\mathtt{K}$, the canonical ideal is:

$$I^{\star}_{\mathtt{K}} = \langle f_0 f_1, f_0 t_1, t_0 f_1, t_0 t_1 \rangle$$

where, for example, $f_0 f_1 = \mathrm{gr}(\neg f)$ for the constant-true function f. The Stanley-Reisner ideal is:

$$I_{\mathtt{K}} = \langle f_0 t_0, f_1 t_1 \rangle$$

since there are no extentures in the complete class, the generators are only of the second kind, and both 0 and 1 are certainly non-trivial for the complete class.

6. Cellular Resolutions of a function class

The theory developed so far now allows us to assign homological invariants to the class $\mathtt{C}$ through the canonical ideal $I^{\star}_{\mathtt{C}}$:

Definition 6.1 (c.f. [5], following Lemma 2.50)**.** The *multigraded Betti numbers* of $\mathtt{C}$ are $\beta_{i,m}(\mathtt{C}) = \beta_{i,m}(I^{\star}_{\mathtt{C}})$

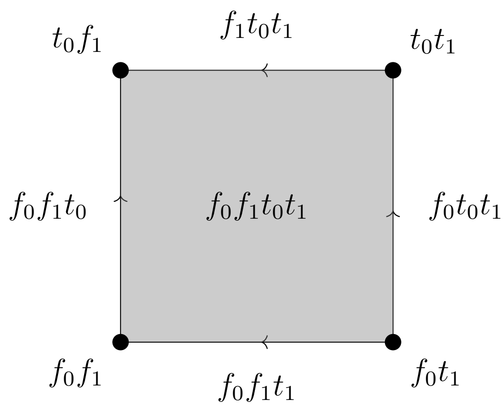

FIGURE 5. A cellular resolution of the complete class `K`

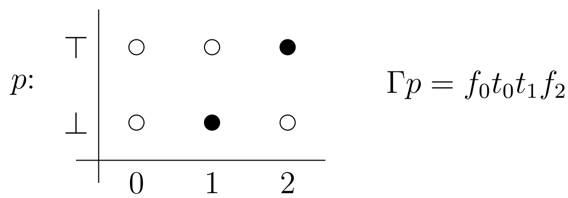

FIGURE 6. The definition of Γp

Definition 6.2 (c.f. [5], Definition 3.6)**.** The *homological dimension* of `C` is the projective dimension of $I^\star_{\mathtt{C}}$, $\dim_h(\mathtt{C}) = \mathrm{pd}(I^\star_{\mathtt{C}})$.

As both of these are defined in terms of the canonical ideal, resolving the canonical ideal becomes the main objective and the main tool used in [5] is finding cellular resolutions of this ideal. A cellular resolution for our running example `K` is picture in Figure 5, which we will generalize after discussing another way to think about labelled complexes in this context:

Definition 6.3. For a partial function p, denote by Γp the complement of the graph of p in $[n] \times [2]$.

Figure 6 shows an example of a partial function p and Γp. If we have a cellular resolution of $I^\star_{\mathtt{C}}$, then its vertices are labeled by $\mathrm{gr}(\neg f) = \Gamma f$ for a total function f.

Lemma 6.4. *If a face in a celluar resolution has boundaries labeled by* $\Gamma p_1, \ldots, \Gamma p_k$ *for possibly partial* p_i*, then its label is* $\Gamma(\bigcap_{i=1}^k p_i)$.

Proof. We wish to show the given expression is the least common multiple of the given boundary labels. Let $q = \bigcap_{i=1}^k p_i$ and $m = \Gamma q$. To see Γp_i divides m, note that if p_i is defined at j, so is q_i and they have the same value, hence whichever of t_j or f_j is in Γp_i is also in m. If m' were another monomial divisble by all the Γp_i, then it is also divisible by $\Gamma q = m$, so m is minimal. □

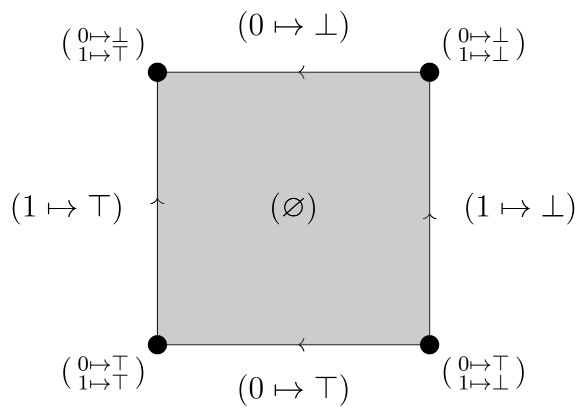

FIGURE 7. A cellular resolution of K, labeled by partial functions

This suggests that we can think of the faces in these cellular resolutions as being labeled by partial functions instead of monomials. Figure 7 shows our cellular resolution of K labeled by partial functions, where $(\varnothing)$ is the empty partial function. This example illustrates a kind of duality at play, where now larger faces are labeled by smaller objects. Labels dividing a given monomial $m = \Gamma p$ translates to partial functions extending p. Thus when labeling with partial functions, the subcomplex $X_{\preceq m}$ translates to what we will write as $X_{\supseteq p}$. If we want to resolve the canonical ideal of a class with such a partial-function-labeled complex, we must start by simply labeling the vertices by the functions in the class. We use this new method to generalize Figure 7:

Theorem 6.5 (c.f. [5], Theorem 2.46)**.** *The complete class* $[2]^{[n]}$ *is resolved by the* n*-cube.*

Proof. Use the standard CW-complex structure X on the cube with one n-cell. The 2^n vertices will be labeled with f for every $f : [n] \to [2]$; label them in such a way that adjacent vertices differ at exactly one input. In particular let vertex $(x_1, \ldots, x_n)$ be labeled by the function whose value at i is x_i (treating f as taking values in $\{0, 1\}$). We will use propositions 2 and 3. Consider a partial function p and $X_{\supseteq p}$. In our labeling scheme then, this subcomplex contains vertices $\{(x_1, \ldots, x_n) : x_i = p(i) \text{ for } i \in p\}$ and all faces contained among these vertices. Thus $X_{\supseteq p}$ is a a lower-dimensional cube, and such a cube is contractible and thus acyclic. So this is indeed a cellular resolution of the canonical ideal, and it is minimal since clearly every face has distinct partial function label. $\square$

[5] goes on in section 2.3 to construct cellular resolutions (and thus multigraded Betti numbers) for several function classes, and the sequel paper [6] constructs many more.

7. Applications

[5] discusses applications of this theory in chapter 3, of which we will highlight a few. The following is an important invariant of concept classes studied in learning theory:

Definition 7.1. For a function class $\mathtt{C} \subseteq [2]^{[n]}$, we way that a subset $U \subseteq [n]$ is *shattered* by $\mathtt{C}$ if the restriction of $\mathtt{C}$ to U is the complete function class on U. The collection of subsets of $[n]$ that are shattered by C form an abstract simplicial complex called the *shatter complex* $\mathcal{SH}_{\mathtt{C}}$. The *VC-dimension* of $\mathtt{C}$ is the largest size of a shattered subset $\dim_{VC}(\mathtt{C}) = \max\{|U| : U \text{ shattered by } \mathtt{C}\} = \dim \mathcal{SH}_{\mathtt{C}} + 1$.

The VC-dimension is important due to the following theorem from learning theory:

Theorem 7.2. *Given a function f that is known to be in a function class $\mathtt{C}$, the amount of sampling necessary to learn f (determine it with high probability) is proportional to $\dim_{VC}(\mathtt{C})$.*

A main result of the paper is that homological dimension dominates VC-dimension:

Theorem 7.3 (c.f. [5], Theorem 3.11, Corollary 3.34)**.** $\dim_{VC}(\mathtt{C}) \leq \dim_h(\mathtt{C})$*, with equality if $I_{\mathtt{C}}$ is Cohen-Macaulay.*

The Cohen-Macaulay condition is an algebraic condition on an ideal. There are several definitions, but we present the following:

Definition 7.4 (c.f. [2], Definition 5.52)**.** An ideal I is Cohen-Macaulay if the projective dimension of R/I coincides with its codimension: $\mathrm{pd}(I) = \mathrm{codim}(I)$.

For squarefree monomial ideals, the codimension has the following interpretation:

Proposition 7. For a squarefree monomial ideal I, $\mathrm{codim}(I)$ is the the smallest size of a set of variables such that every minimal generator of I contains a variable in that set.

So this condition is quite straightforward to check for function classes. As such, checking that $I_{\mathtt{C}}$ (which was completely described) is Cohen-Macaulay is a useful criterion for showing $\dim_{VC}(\mathtt{C}) = \dim_h(\mathtt{C})$.

Remark 7.5. One should notice that we are requiring $I_{\mathtt{C}}$ to be Cohen-Macaulay instead of the canonical ideal $I^{\star}_{\mathtt{C}}$, which we have mainly studied. In fact, requiring $I^{\star}_{\mathtt{C}}$ to be Cohen-Macaulay essentially trivializes the theory, and $\dim_{VC}(\mathtt{C}) = \dim_h(\mathtt{C}) = 1$ in this case ([5], Corollary 3.26). The Cohen-Macaulayness of I_C guarantees equality non-trivial function classes.

Another application is the separation of function classes. The problem is, given function classes $\mathtt{C} \subseteq \mathtt{D}$, show that they are actually different; i.e. $\mathtt{C} \subset \mathtt{D}$. If both classes are separately well-understood, we may possibly construct cellular resolutions for both. Then we may exploit the following

Proposition 8 (c.f. [5], Section 2.5, Proposition 2.89)**.** If X provides a cellular resolution of $\mathtt{C}$, and Y provides a cellular resolution of $\mathtt{D}$, then $X \star Y$ is a CW-complex with a labeling obtained from that of X and Y that provides a resolution of $\mathtt{C} \cup \mathtt{D}$.

Here $X \star Y$ is the *join* of two topological spaces generalizing the cone and the suspension; [1] page 9 provides a definition and discussion of the CW-complex structure.

If C and D really are the same class, then of course their Betti numbers will be the same. So to exhibit their separation, it suffices to exhibit a difference in some Betti number of C and D. But $\mathtt{D} = \mathtt{C} \cup \mathtt{D}$, so we can study the Betti numbers of D through the join of cellular resolutions of each individual class. Essentially, we are throwing the class D in with C and witnessing the homology change. This method is exploited in [5] to recover and extend the following well-known result in learning theory:

Theorem 7.6 (c.f. [5], Theorem 3.40, Corollary 3.41)**.** *The parity function is not computable by a polynomial threshold function unless its degree is maximal.*

Appendix A. Code for computed examples

Macaulay 2 is a language and collection of software packages for computational commutative algebra (`http://www2.macaulay2.com/Macaulay2/`). The following can be run in the online Macaulay2 intepreter:

Code A.1. Computes the resolution in example 3.3:

```
R = ZZ/101[x, y]
I = monomialIdeal(x^3, x*y, y^5)
-- compute a minimal free resolution of the R-module R/I:
chain = res (R^1/I)
-- display the full chain complex with boundary maps:
chain.dd
```

Code A.2. Computes the cellular resolution in example 4.1:

```
R = ZZ/101[a,b,c]
-- construct the maps of the cellular free complex manually:
d1 = matrix({{a^2*b, a*c, b*c^2, b^2}})
d2 = map(source d1, , {{-c, 0, 0, b, 0}, {a*b, b*c, 0, 0, b^2}, {0, -a, b,
    ↪ 0, 0}, {0, 0, -c^2, -a^2, -a*c}})
```

```
d3 = map(source d2, , {{b, 0}, {0, -b}, {0, -a}, {c, 0}, {-a, c}})
-- construct a chain complex from the maps:
C = new ChainComplex; C.ring = R;
C#0 = target d1; C#1 = source d1; C#2 = source d2; C#3 = source d3;
C.dd#1 = d1; C.dd#2 = d2; C.dd#3 = d3
-- print the chain complex with boundary maps
C.dd
-- compute the minimal presentation of the homology of C:
prune homology C
```

References

[1] Allen Hatcher. *Algebraic topology*. Cambridge ; New York: Cambridge University Press, 2002. ISBN: 9780521791601.

[2] Ezra Miller and Bernd Sturmfels. *Combinatorial Commutative Algebra.* English. Dordrecht: Springer-Verlag New York Inc., 2006. ISBN: 9780387271033.

[3] Irena Peeva. *Graded Syzygies*. en. Algebra and Applications. London: Springer-Verlag, 2011. ISBN: 9780857291769. DOI: 10.1007/978-0-85729-177-6.

[4] Christopher A. Francisco, Jeffrey Mermin, and Jay Schweig. "A Survey of Stanley–Reisner Theory". en. In: *Connections Between Algebra, Combinatorics, and Geometry*. Ed. by Susan M. Cooper and Sean Sather-Wagstaff. Springer Proceedings in Mathematics & Statistics. New York, NY: Springer, 2014, pp. 209–234. ISBN: 9781493906260. DOI: 10.1007/978-1-4939-0626-0_5.

[5] Greg Yang. "A Homological Theory of Functions". In: (Apr. 2017). arXiv: 1701.02302.

[6] Justin Chen et al. "Free resolutions of function classes via order complexes". In: (Sept. 2019). arXiv: 1909.02159 version: 1.

[7] Daniel R. Grayson and Michael E. Stillman. *Macaulay2, a software system for research in algebraic geometry*. Available at http://www.math.uiuc.edu/Macaulay2/.