

THE RING OF ALGEBRAIC FUNCTIONS ON
PERSISTENCE BAR CODES

AARON ADCOCK, ERIK CARLSSON AND GUNNAR CARLSSON

(communicated by Ralph Cohen)

Abstract

Persistent homology is a rapidly developing field in the study of numerous kinds of data sets. It is a functor which assigns to geometric objects so-called *persistence bar codes*, which are finite collections of intervals. These bar codes can be used to infer topological aspects of the geometric object. The set of all persistence bar codes, suitably defined, is known to possess metrics that are quite useful both theoretically and in practice. In this paper, we explore the possibility of coordinatizing, in a suitable sense, this same set of persistence bar codes. We derive a set of coordinates using results about multi-symmetric functions, study the property of the corresponding ring of functions, and demonstrate in an example how they work.

1. Introduction

Persistent homology [4, 17] is a fundamental tool in the area of computational topology. It can be used to infer topological structure in data sets (see [1, 6]), but variations on the method can be applied to study aspects of the shape of point clouds which are not overtly topological [7, 11]. The methodology assigns to any finite metric space (such as are typically obtained in experimental data of various kinds) and non-negative integer k a *bar code*, by which we will mean a finite collection of intervals with endpoints on the real line. Note that we do *not* include infinite intervals, but the intervals are allowed to have length zero. The integer k specifies a dimension of a feature (zero-dimensional for a cluster, one-dimensional for a loop, etc.), and an interval represents a feature which is “born” at the value of a parameter (the persistence parameter) given by the left hand endpoint of the interval, and which “dies” at the value given by the right hand endpoint. These bar codes have been demonstrated to identify structure in spaces of image patches in [1] and [6], to distinguish between hand-drawn letters in [11], and to help uncover unusual events in viral evolution [8]. Because of the unusual structure of the invariant, i.e., as a collection of intervals rather than numerical quantities, the method currently requires substantial knowledge of topological methods. It would clearly be useful to assign and interpret various numerical quantities attached to bar codes, so that these outputs could be

Received October 29, 2014, revised June 24, 2015; published on May 31, 2016.

2010 Mathematics Subject Classification: 55N99, 62-07.

Key words and phrases: persistent homology, point cloud, metric space, data analysis.

Article available at <http://dx.doi.org/10.4310/HHA.2016.v18.n1.a21>

Copyright © 2016, International Press. Permission to copy for private use granted.

used as input to standard algorithms within machine learning, cluster analysis, and other methods. This is already being done in an ad hoc fashion in most applications of persistent homology, and it is the purpose of this paper to identify an algebra of such functions on the set of bar codes which is defined in a conceptually coherent way.

The main idea is the following. A bar code with exactly n intervals can be specified by a vector $(x_1, y_1, x_2, y_2, \dots, x_n, y_n)$, where x_i denotes the left endpoint of the i -th interval and y_i the right endpoint. However, this representation is many to one, in that the bar code structure does not retain the ordering on the intervals. We describe it as follows. For any set X , $Sp^n(X)$ is defined to be the orbit space of the action of the symmetric group on n letters on the product X^n given by permuting the coordinates, and is referred to as the n -fold symmetric product of X . It is clear that the set of bar codes with exactly n bars can be described as the subset of $Sp^n(\mathbb{R}^2)$ consisting of points of the form $\{(x_1, y_1), \dots, (x_n, y_n)\}$ where $x_i \leq y_i$ for all $1 \leq i \leq n$. On the other hand, the space $(\mathbb{R}^2)^n$ is the set of real points of the algebraic variety $(\mathbb{A}^2)^n$ [13]. It is the affine space of dimension $2n$, and the symmetric group action mentioned above is an algebraic action. It is then known (see [16]) that the orbit space inherits the structure of an algebraic variety, and the elements of its *affine coordinate ring* [13] are functions on the set of bar codes with exactly n intervals. This variety is denoted by $Sp^n(\mathbb{A}^2)$, and its real points will be denoted by $Sp^n(\mathbb{A}^2)(\mathbb{R})$.

Remark. This affine coordinate ring consists of functions which are defined on all of $Sp^n(\mathbb{R}^2)$, without taking into account the inequalities $x_i \leq y_i$. *Real algebraic geometry*, as developed, for example, in [3], studies rings equipped with the additional structure of a class of positive functions, and would presumably give additional useful information, but we will deal with the functions without this additional structure.

The affine coordinate rings described above are well known algebras referred to generically as rings of *multisymmetric polynomials* [12]. They can be quite complicated, since it turns out that any set of algebra generators for them will satisfy non-trivial relations or *syzygies*. It turns out, though, that there are inclusions of algebraic varieties

$$Sp^n(\mathbb{A}^2) \rightarrow Sp^{n+1}(\mathbb{A}^2), \quad (1.1)$$

which produce an inverse system of affine coordinate rings

$$\cdots \rightarrow A_{Sp^{n+1}(\mathbb{A}^2)}(\mathbb{R}) \rightarrow A_{Sp^n(\mathbb{A}^2)}(\mathbb{R}) \rightarrow \cdots,$$

whose inverse limit we will denote, by abuse of notation, by $A_{Sp^\infty(\mathbb{A}^2)}(\mathbb{R})$. This inverse limit is a complicated ring, but we define a natural subring of functions, consisting of the so-called *K-finite functions*, which is known to be freely generated on a set of minimal algebra generators [12].

The analysis of the system (1.1) above is not sufficient, though. This system identifies a point $((x_1, y_1), \dots, (x_n, y_n)) \in Sp^n(\mathbb{R}^2)$ with the point

$$((x_1, y_1), \dots, (x_n, y_n), (0, 0)) \in Sp^{n+1}(\mathbb{R}^2).$$

In other words, a set S of n intervals is identified with the set of $n+1$ intervals obtained by adjoining the interval of length zero whose two endpoints are zero. However, in the parametrization of the isomorphism classes of persistence vector spaces in [17] by bar codes, any interval of length zero is identified with the zero module.

So, we would like to determine the ring of all algebraic functions (i.e., the elements of $A_{Sp^\infty(\mathbb{A}^2)}(\mathbb{R})$) which have the property that they take the same value on any bar code as on the result of adjoining any interval of length zero to it. In this paper, we will identify this subring, describe its structure, and describe the algebra generators explicitly. Let \sim be the equivalence relation on $Sp^\infty(\mathbb{R}^2)$ generated by the relations $\{[x_1, y_1], \dots, [x_n, y_n], [x, x]\} \sim \{[x_1, y_1], \dots, [x_n, y_n]\}$, i.e., the relation which identifies two sets of intervals if one can be obtained from the other by adding and deleting intervals of length zero. The main result is as follows.

Theorem 1.1. *Let \mathfrak{B} denote the quotient of $Sp^\infty(\mathbb{R}^2)$ by the equivalence relation \sim defined above. There is an embedding of \mathfrak{B} into the points of an affine scheme*

$$\mathfrak{B}^{sch} = Spec(A_{\mathfrak{B}^{sch}}),$$

which is compatible with the evident maps $Sp^n(\mathbb{R}^2) \hookrightarrow Sp^n(\mathbb{A}^2)(\mathbb{R})$, in the sense that the diagrams

$$\begin{array}{ccc} Sp^n(\mathbb{R}^2) & \longrightarrow & \mathfrak{B} \\ \downarrow & & \downarrow \\ Sp^n(\mathbb{A}^2)(\mathbb{R}) & \longrightarrow & \mathfrak{B}^{sch}(\mathbb{R}) \end{array}$$

all commute.

\mathfrak{B}^{sch} admits an action by the algebraic group \mathbb{G}_m , which defines a natural subring

$$A_{\mathfrak{B}^{sch}}^{fin} \subseteq A_{\mathfrak{B}^{sch}}$$

consisting of all functions f for which all the translates of f span a finite dimensional real subspace of $A_{\mathfrak{B}^{sch}}$. We find that the ring $A_{\mathfrak{B}^{sch}}^{fin}$ is described by

$$A_{\mathfrak{B}^{sch}}^{fin} \cong \mathbb{R}[p_{a,b}; a \geq 0, b \geq 1],$$

where a and b are integers. For a bar code $\{[x_1, y_1], [x_2, y_2], \dots, [x_n, y_n]\}$, the function $p_{0,1}$ is given by $\sum_i (y_i - x_i)$, and $p_{1,1}$ is given by $\sum_i (y_i - x_i)(y_i + x_i)$.

As used in algebraic topology, homological invariants are useful guides in understanding the structure of a space. This method of applying homology was extended to point clouds in [6] and [8]. Another direction of applications, though, is in the study of data sets whose objects themselves support geometry, such as chemical compounds or images, where one can attach bar code signatures to the individual points in the data set. In this case, the data set produces new, derived data sets whose entries are bar codes. It is already known that the bar codes support several natural metrics, including the so-called *bottleneck distance*, which serve as a useful organizing principle for such databases. What we do is to produce useful coordinates on the space, and demonstrate their utility.

Acknowledgments

The authors wish to thank Allen Knutson for extremely helpful conversations.

2. Barcodes

We first discuss the set of bar codes, without any algebraic variety structures. For any set X , we let $Sp^n(X)$ denote the orbit set of the action of the symmetric group Σ_n on the set X^n , by permuting coordinates.

Definition 2.1. Let \mathfrak{J} denote the set of all ordered pairs (x, y) of real numbers, with $x \leq y$. For any non-negative integer n , a *bar code* of cardinality n is a finite multiset of elements in \mathfrak{J} , of cardinality n . The collection of all bar codes of cardinality n is then clearly identified with $Sp^n(\mathfrak{J})$.

The goal is to build a natural construction of the set of all bar codes of varying cardinality, which informally has the property that intervals of length zero are “ignored”. This is done by considering a relative version of the standard “infinite symmetric product” construction.

For a set X , form the infinite disjoint union

$$Sp(X) = \coprod_{n \geq 0} Sp^n(X),$$

where $Sp^n(X)$ is the n -fold symmetric product of X . Of course, $Sp(X)$ is the free commutative monoid on X . For a subset $Y \subseteq X$, we may also form $Sp(X, Y)$ by forming the quotient $Sp(X)/\sim_Y$, where \sim_Y is the equivalence relation on $Sp(X)$ given, for elements $w, w' \in Sp(X)$, by $w \sim_Y w'$ if and only if there are elements $u, u' \in Sp(Y) \subseteq Sp(X)$ so that $w + u = w' + u'$. This construction is clearly functorial for maps of pairs of sets.

Proposition 2.2. *$Sp(X, Y)$ is isomorphic to the free monoid on the set X/Y , i.e., the quotient of X by the equivalence relation in which all elements of Y are equivalent to each other, and all elements of $X - Y$ lie in one element equivalence classes. In other terms,*

$$Sp(X, Y) \cong Sp(X/Y, *),$$

where $*$ denotes the equivalence class Y . The construction $Sp(-, -)$ is functorial for maps of pairs of sets.

Proof. Clear. □

There are maps $\alpha_n : Sp^n(X, x) \rightarrow Sp^{n+1}(X, x)$ defined by

$$\alpha_n(\{x_1, x_2, \dots, x_n\}) = \{x_1, x_2, \dots, x_n, x\}.$$

Proposition 2.3. *$Sp(X, Y)$ is isomorphic to the colimit*

$$Sp^1(X/Y) \xrightarrow{\alpha_1} Sp^2(X/Y) \xrightarrow{\alpha_2} Sp^3(X/Y) \xrightarrow{\alpha_3} \dots$$

Proof. Clear. □

Definition 2.4. Let \mathfrak{J} be as above, and let $\Delta \subseteq \mathfrak{J}$ be the diagonal subset. Then we define the infinite bar code set \mathfrak{B} to be the set $Sp(\mathfrak{J}/\Delta, [\Delta])$, where $[\Delta]$ denotes the equivalence class of Δ .

Remark. This construction is exactly the quotient of the set of all bar codes by the equivalence relation which identifies two bar codes which differ only by the addition or deletion of intervals of length zero.

We note that the set \mathfrak{B} is equipped with two group actions. Let \mathbb{R}_+^* denote the group of positive real numbers under multiplication. Then \mathbb{R}_+^* acts on \mathcal{J} by $r \cdot (x, y) = (rx, ry)$. This action preserves the subset $\Delta \subseteq \mathcal{J}$, and it is easy to see that a group action with these two properties will induce an action on $Sp(\mathcal{J}/\Delta, [\Delta])$. We also let \mathbb{R}^+ denote the additive group of all real numbers. Then \mathbb{R}^+ also acts on \mathcal{J} by $r \cdot (x, y) = (x + r, y + r)$, and this action also preserves Δ and, therefore, induces an action on \mathfrak{B} . Taken together, we obtain an action of the group $\mathbb{R}_+^* \times \mathbb{R}^+$, where \mathbb{R}_+^* acts on \mathbb{R}^+ by multiplication.

3. Colimits of affine schemes

We consider first the category of affine schemes **Aff**, and for a given commutative ring A the category \mathbf{Aff}^A of affine schemes over $Spec(A)$, which we will also refer to as A -schemes. We also have the categories **Comm** of commutative rings, as well as the category \mathbf{Comm}^A of commutative algebras over a commutative ring A . Let Γ denote the global section functor, which assigns to a scheme its ring of global sections, which is inverse to $Spec$. We will denote $\Gamma(X)$ by A_X . It is well known that **Aff** and \mathbf{Aff}^A are equivalent (via the functor $Spec$) to the categories \mathbf{Comm}^{op} and $(\mathbf{Comm}^A)^{op}$, respectively. It follows easily that **Aff** and \mathbf{Aff}^A are complete and cocomplete. In particular, if a group G acts on an affine scheme or an A -scheme, the orbit scheme exists and corresponds under Γ to the ring of invariants of the induced group action on the ring of global sections.

We fix the base ring A to be \mathbb{R} throughout, unless otherwise indicated. For example, \mathbb{A} will denote the affine line over \mathbb{R} , but we permit ourselves other coefficient rings by including that ring as a subscript, so $\mathbb{A}_{\mathbb{C}}$ indicates the affine line over \mathbb{C} . We consider \mathbb{R} -schemes with group actions by a group G . For a given \mathbb{R} -scheme X , we can consider the set of \mathbb{R} -valued points of X , which we denote by $X(\mathbb{R})$. They are in bijective correspondence with the set of maximal ideals \mathfrak{m} in A_X whose residue class fields are isomorphic to \mathbb{R} via the structure map $\mathbb{R} \rightarrow A_X \rightarrow A_X/\mathfrak{m}$. We note that every element $\alpha \in A_X$ defines an \mathbb{R} -valued function $f_\alpha: X(\mathbb{R}) \rightarrow \mathbb{R}$ via the rule that $f_\alpha(x) = [\alpha]_x$, where $[\alpha]_x$ denotes the coset $\alpha + m_x$, where m_x is the maximal ideal corresponding to x . The quotient A_X/m_x is identified with \mathbb{R} via the composite isomorphism

$$\mathbb{R} \rightarrow A_X \rightarrow A_X/m_x.$$

The assignment

$$\alpha \rightarrow f_\alpha$$

defines a homomorphism of commutative rings

$$\theta_X: A_X \rightarrow F(X(\mathbb{R}), \mathbb{R}),$$

where $F(X, Y)$ denotes the set of set valued functions from X to Y , where the addition and multiplication on $F(X(\mathbb{R}), \mathbb{R})$ are given by pointwise multiplication.

Let $F: \mathbf{C} \rightarrow \mathbf{Aff}^{\mathbb{R}}$ be any functor, where \mathbf{C} is a small category. Then we have the colimit $\underset{\rightarrow}{\text{colim}} F$, which is itself an affine \mathbb{R} -scheme. There is a unique map of sets

$$\eta: \underset{\rightarrow}{\text{colim}} F(c)(\mathbb{R}) \rightarrow (\underset{\rightarrow}{\text{colim}} F)(\mathbb{R}) \tag{3.2}$$

so that for each object c of \mathbf{C} , the composite

$$F(c)(\mathbb{R}) \rightarrow \underset{\rightarrow}{\text{colim}} F(c)(\mathbb{R}) \xrightarrow{\eta} (\underset{\rightarrow}{\text{colim}} F(c))(\mathbb{R})$$

is equal to the map of sets of real valued points induced by the map of \mathbb{R} -schemes $F(c) \rightarrow \underset{\rightarrow}{\text{colim}} F(c)$.

Colimits of diagrams of affine schemes often produce schemes whose associated rings of global sections are quite complicated, and large. In this section, we show that additional structure on a scheme X , in the form of an action on X by the algebraic group \mathbb{G}_m , allows us to identify a subring of A_X which is much more manageable.

By definition, an action of an affine algebraic group G on an affine scheme X is a map of schemes $\alpha: G \times X \rightarrow X$ satisfying the analogues of the usual compatibilities of the action map with the multiplication map in the algebraic group. As usual, we will denote by \mathbb{G}_m the multiplicative group of units. An action of an algebraic group on an affine \mathbb{R} -scheme gives a group action of the (ordinary, not algebraic) group $G(\mathbb{R})$ of \mathbb{R} -valued points of G on the scheme X , and, therefore, an action of $G(\mathbb{R})$ on A_X . We will consider schemes over $\text{Spec}(\mathbb{R})$, with action of the scheme \mathbb{G}_m . Any such scheme X , therefore, admits an action by the group \mathbb{R}^* . For any $n \in \mathbb{Z}$, we will define an \mathbb{R} -submodule $A_X(n) \subseteq A_X$ by

$$A_X(n) = \{\alpha \in A_X \mid r \cdot \alpha = r^n \alpha\}.$$

Here $r \cdot \alpha$ is the result of the action of the element $r \in \mathbb{R}^*$ under the specified action, and $r\alpha$ denotes the multiplication by r using the \mathbb{R} -module structure on A_X . Since $A_X(m) \cdot A_X(n) \subseteq A_X(m+n)$, it is clear that the direct sum

$$A_X^{gr} = \bigoplus_n A_X(n) \subseteq A_X$$

is a graded \mathbb{R} -algebra. We let A_X^{fin} denote A_X^{gr} regarded as an algebra over \mathbb{R} , with the grading ignored. There is an evident inclusion $A_X^{fin} \hookrightarrow A_X$.

Proposition 3.1. *The constructions A_X^{gr} and A_X^{fin} are both functorial for maps of affine \mathbb{R} -schemes with \mathbb{G}_m -action. Let \mathbf{C} be a small category and F be a functor with values in the category of \mathbb{R} -schemes with \mathbb{G}_m -action, and let X_F denote its colimit. Then $A_{X_F}^{gr}$ is isomorphic to the inverse limit of the corresponding functor F^{op} from the category \mathbf{C}^{op} to the category of graded \mathbb{R} -algebras.*

Proof. Straightforward. □

Example 3.2. Equip \mathbb{A}^n with the \mathbb{G}_m action defined by

$$r \cdot (x_1, \dots, x_n, \dots) = (rx_1, \dots, r^n x_n)$$

and consider the directed system of affine \mathbb{R} -schemes with \mathbb{G}_m -action

$$\mathbb{A}^1 \hookrightarrow \mathbb{A}^2 \hookrightarrow \mathbb{A}^3 \hookrightarrow \dots, \tag{3.3}$$

where the inclusions $\mathbb{A}^n \rightarrow \mathbb{A}^{n+1}$ are given by

$$(x_1, \dots, x_n) \rightarrow (x_1, \dots, x_n, 0).$$

Let \mathbb{A}^∞ denote the colimit of this diagram in the category of affine \mathbb{R} -schemes with \mathbb{G}_m -action. Then we have that $A_{\mathbb{A}^\infty}$ consists of all formal power series in the infinite set of variables $\{\sigma_1, \sigma_2, \dots, \sigma_n, \dots\}$ with the property that the power series in the set $\{\sigma_1, \dots, \sigma_n\}$ obtained by setting all the variables $\{\sigma_j | j > n\}$ equal to zero is a finite sum, i.e., a polynomial. On the other hand, $A_{\mathbb{A}^\infty}^{gr}$ and $A_{\mathbb{A}^\infty}^{fin}$ are the graded and ungraded rings of polynomials in the set of variables $\{\sigma_i | i \geq 1\}$, respectively.

Definition 3.3. Let X be any affine \mathbb{R} -scheme. By $Sp^n(X)$, we will mean the orbit scheme

$$X^n / \Sigma_n.$$

The product is computed in the category of schemes over $Spec(\mathbb{R})$, so $A_{Sp^n(X)}$ is identified with the ring of invariants $(A_X^{\otimes n})^{\Sigma_n}$. If we fix an \mathbb{R} -valued point $*$ in X , there are maps of \mathbb{R} -schemes $i_n: Sp^n(X) \rightarrow Sp^{n+1}(X)$ defined as the composites

$$Sp^n(X) = Sp^n(X) \times_{Spec(\mathbb{R})} Spec(\mathbb{R}) \xrightarrow{id \times * } Sp^n(X) \times_{Spec(\mathbb{R})} X \rightarrow Sp^{n+1}(X), \quad (3.4)$$

where the right hand map is the projection $X^{n+1} / \Sigma_n \times \{e\} \rightarrow X^{n+1} / \Sigma_{n+1}$. We define $Sp^{sch}(X, *)$ to be the colimit of this system in the category of affine \mathbb{R} -schemes.

If X is equipped with an action by the group \mathbb{G}_m , then there is an associated action of $\mathbb{G}_m(\mathbb{R})$ on $Sp^n(X)$, defined by $r \cdot \{x_1, \dots, x_n\} = \{rx_1, \dots, rx_n\}$. If the point $x \in X(\mathbb{R})$ is fixed by this \mathbb{G}_m -action, then diagram defining $Sp(X, *)$ becomes a diagram of \mathbb{R} -schemes with \mathbb{G}_m -action, and, consequently, the colimit is endowed with a $\mathbb{G}_m(\mathbb{R})$ -action. In the special case where X is \mathbb{A}^1 , with the usual defining action of \mathbb{G}_m , and with $x = 0$, we find that $A_{Sp(\mathbb{A}^1, 0)}$ consists of the ring of power series in the elementary symmetric functions σ_i with the property that the sums involving any finite set of the σ_i 's are finite. On the other hand, $A_{Sp(\mathbb{A}^1, 0)}^{gr}$ and $A_{Sp(\mathbb{A}^1, 0)}^{fin}$ are the graded and ungraded polynomial rings in the elementary symmetric functions. Of course, this result is isomorphic to the result of Example 3.2 above.

4. The affine scheme \mathfrak{B}^{sch}

By analogy with the construction in Section 2, we now define \mathfrak{B}^{sch} to be $Sp^{sch}(\mathbb{A}^2 / \Delta, [\Delta])$, where Sp^{sch} is defined in Example 3.3, $\Delta \subseteq \mathbb{A}^2$ is the diagonal subscheme in \mathbb{A}^2 , and where \mathbb{A}^2 / Δ denotes the quotient in the category of affine \mathbb{R} -schemes. We note that the groups \mathbb{R}^* and \mathbb{R}^+ act on \mathbb{A}^2 via the formulae $r \cdot (x, y) = (rx, ry)$ and $r \cdot (x, y) = (x + r, y + r)$, respectively, and that both actions preserve the subscheme Δ . It follows that both groups act on \mathfrak{B}^{sch} , and the two actions assemble to an action of the semidirect product $\mathbb{R}^* \ltimes \mathbb{R}^+$. We first examine the ring $A_{\mathbb{A}^2 / \Delta}$. We first set $\xi = x + y$ and $\eta = x - y$, where x and y are the standard coordinates on \mathbb{A}^2 , so \mathbb{A}^2 is also isomorphic to $\mathbb{R}[\xi, \eta]$.

Proposition 4.1. *The ring $A_{\mathbb{A}^2 / \Delta}$ is isomorphic to the subring of $\mathbb{R}[\xi, \eta]$ spanned by the monomials $\xi^m \eta^n$ for which $n > 0$ whenever $m > 0$. In other words, it consists of*

the span of the set consisting of the trivial monomial 1 and all non-trivial monomials with positive η -exponent.

Proof. Since \mathbb{A}^2/Δ is described as the pushout of the diagram

$$\begin{array}{ccc} \Delta & \longrightarrow & * \\ \downarrow & & \\ \mathbb{A}^2 & & \end{array}$$

the ring $A_{\mathbb{A}^2/\Delta}$ is the pullback of the diagram of rings

$$\begin{array}{ccc} & \mathbb{R}[\xi, \eta] & \\ & \downarrow & \\ \mathbb{R} & \longrightarrow & \mathbb{R}[\eta], \end{array}$$

where the vertical map is the projection $\xi \rightarrow 0$. The result follows. □

We consider also the n -fold powers $(\mathbb{A}^2/\Delta)^n$. We describe the coordinate ring of $(\mathbb{A}^2)^n$ as $\mathbb{R}[\xi_1, \dots, \xi_n, \eta_1, \dots, \eta_n]$, where $\xi_i = x_i + y_i$ and $\eta_i = x_i - y_i$.

Corollary 4.2. *For any $n > 0$, the ring $A_{(\mathbb{A}^2/\Delta)^n}$ is isomorphic to the subring of $\mathbb{R}[\xi_1, \dots, \xi_n, \eta_1, \dots, \eta_n]$ spanned by all monomials with the property that if the exponent of ξ_i is > 0 , then the exponent of η_i is also > 0 .*

Proof. Easy verification using the tensor product description of the ring associated to the n -fold product. □

We note that the subscheme Δ is invariant under the $\mathbb{G}_m(\mathbb{R})$ -action, and, therefore, $A_{(\mathbb{A}^2/\Delta)^n}$ is equipped with a grading, compatible with its inclusion into the graded ring $\mathbb{R}[\xi_1, \dots, \xi_n, \eta_1, \dots, \eta_n]$.

From this point on, we will denote the ring $A_{(\mathbb{A}^2/\Delta)^n}$ by $D(n)$, and the ring $A_{(\mathbb{A}^2/\Delta)^n}^{gr}$ by $D(n)_*$. The symmetric group action permuting the ξ_i 's and η_i 's among themselves clearly preserves the subring $D(n)$. What we need is the ring $A_{Sp^n(\mathbb{A}^2/\Delta)}$, which is clearly identified with the ring of invariants $D(n)^{\Sigma_n}$.

The symmetric group action clearly preserves the basis of monomials within $D(n)$. Let $\{\mu_\alpha\}_{\alpha \in A}$ denote a set of orbit representatives of the Σ_n -action on the set of monomials defined in Corollary 4.2. Let σ_α denote the sum of all the elements in the orbit of μ_α .

Proposition 4.3. *We let $D(n)^{\Sigma_n}$ denote the subring of elements of $D(n)$ which are invariant under the action of Σ_n . Then the elements σ_α form a k -basis of $D(n)^{\Sigma_n}$.*

Proof. This result plainly holds for any algebra over a field of characteristic zero on which there is a G -action which preserves a basis of monomials. □

Recall the maps i_n in diagram (3.4) of Definition 3.3. They induce homomorphisms $i_n^* : D(n + 1) \rightarrow D(n)$ of rings of global sections, which are given by the composites

$$D(n + 1)^{\Sigma_{n+1}} \rightarrow D(n + 1)^{\Sigma_n \times \{e\}} \rightarrow D(n)^{\Sigma_n},$$

where the left hand map is an inclusion of rings and the right hand map is the projection which sends any monomial containing ξ_{n+1} or η_{n+1} to zero. These maps also respect the \mathbb{G}_m -action, and, therefore, give an inverse system of graded rings.

Proposition 4.4. $A_{\mathfrak{B}^{sch}}^{gr}$ is isomorphic to the inverse limit of the system of graded rings

$$\dots \longleftarrow D(n - 1)_*^{\Sigma_{n-1}} \xleftarrow{i_{n-1}^*} D(n)_*^{\Sigma_n} \xleftarrow{i_n^*} D(n + 1)_*^{\Sigma_{n+1}} \longleftarrow \dots \quad (4.5)$$

Proof. Immediate from Proposition 3.1. □

Recall the notation \mathcal{J} from Section 2. We now have the sequence of maps

$$\mathfrak{B} = Sp(\mathcal{J}/\Delta, [\Delta]) \hookrightarrow Sp(\mathbb{R}^2/\Delta, [\Delta]) \rightarrow Sp(\mathbb{A}^2(\mathbb{R})/\Delta(\mathbb{R}), [\Delta(\mathbb{R})]) \xrightarrow{\eta} Sp^{sch}(\mathbb{A}^2/\Delta, [\Delta])(\mathbb{R}).$$

The left hand inclusion is the inclusion of \mathcal{J} in \mathbb{R}^2 which “forgets” the constraint $x \leq y$. The middle arrow is induced by the evident map $\mathbb{R}^2/\Delta \rightarrow (\mathbb{A}^2/\Delta)(\mathbb{R})$. The rightmost arrow is the map η of diagram (3.2). We denote this composite by J . It is clear from the definitions that if we restrict the $\mathbb{R}^* \times \mathbb{R}^+$ action on $Sp^{sch}(\mathbb{A}^2/\Delta, [\Delta])(\mathbb{R})$ to $\Gamma = \mathbb{R}_+^* \times \mathbb{R}^+$, J is equivariant with respect to the Γ -action on \mathfrak{B} defined in Section 2.

We now have that pullback along J gives a homomorphism $J^* : A_{\mathfrak{B}^{sch}}^{fin} \rightarrow F(\mathfrak{B}, \mathbb{R})$. It remains to describe the structure of $A_{\mathfrak{B}^{sch}}^{fin}$, and to prove two simple properties of J^* .

5. The structure of the ring $A_{\mathfrak{B}^{sch}}^{fin}$

We now consider the ring $A_{\mathfrak{B}^{sch}}^{gr}$. This means that we must evaluate the inverse limit of the system of graded rings (4.5) above. We denote this inverse limit by \mathfrak{D}_* .

We next recall some of the notation and basic facts about multisymmetric polynomials, which can be found in Dalbec [12]. Let $R_{n,r}$ be the polynomial ring in nr variables,

$$R_{n,r} = k[x_{i,j}; 1 \leq i \leq n, 1 \leq j \leq r].$$

We let the symmetric group S_n act on $R_{n,r}$ via the formula $\sigma(x_{ij}) = x_{\sigma(i)j}$, and let

$$\Lambda_{n,r} = R_{n,r}^{S_n}$$

denote the ring of S_n invariants. There is an inverse system parallel to the one constructed above involving the rings $\Lambda_{n,r}$. We have evaluation maps

$$\pi_{n,m} : R_{n,r} \rightarrow R_{m,r}, \quad m \leq n$$

defined by setting $x_{ir} = 0$ if $i > m$. The map $\pi_{n,m}$ is S_m -equivariant, when $S_m \subseteq S_n$ is the subgroup of permutations of the first m elements of the set $\{1, \dots, m\}$. We

have the composites

$$\Lambda_{n,r} = R_{n,r}^{S_n} \hookrightarrow R_{n,r}^{S_m} \xrightarrow{\pi_{n,m}^{S_m}} R_{m,r}^{S_m} = \Lambda_{m,r},$$

which we denote by $\rho_{n,m}$. The inverse limit of the system

$$\cdots \xrightarrow{\rho_{n+1,n}} \Lambda_{n,r} \xrightarrow{\rho_{n,n-1}} \Lambda_{n-1,r} \xrightarrow{\rho_{n-1,n-2}} \cdots \xrightarrow{\rho_{2,1}} \Lambda_{1,r}$$

will be denoted by Λ_r , and referred to as the ring of r -multisymmetric functions. It has a grading

$$\Lambda_r = \bigoplus_k \Lambda_r^k$$

induced by the grading on $R_{n,r}$. There is an evident embedding $\mathfrak{D}_* \hookrightarrow \Lambda_2$. We will use this embedding to identify the structure of \mathfrak{D}_* .

The ring of multisymmetric functions has several interesting sets of generators. Given an array of nonnegative integers

$$\begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1r} \\ a_{21} & a_{22} & \cdots & a_{2r} \\ \vdots & \vdots & \ddots & \vdots \\ a_{k1} & a_{k2} & \cdots & a_{kr} \end{bmatrix},$$

with $k \leq n$ and $\mathbf{a}_i = (a_{i1}, \dots, a_{ir})$, we define the multisymmetric monomials by

$$m_{\mathbf{a}_1, \dots, \mathbf{a}_k} = \text{Sym}(x_{11}^{a_{11}} \cdots x_{kr}^{a_{kr}}) = \sum_{\sigma \in S_n} \prod_{i,j} x_{\sigma(i)j}^{a_{ij}} \in \Lambda_r.$$

Sym applied to a monomial yields the sum of all monomials which are in the orbit of the S_n -action.

They form a vector space basis of $\Lambda_{n,r}$, for any n . It is known that $\Lambda_{n,r}$ is generated as an algebra by the symmetrizations of monomials involving only $\{x_{11}, x_{12}, \dots, x_{1r}\}$. They are given by the formulae

$$p_{\mathbf{a}} = m_{\mathbf{a}} = \sum_i x_{i1}^{a_1} \cdots x_{ir}^{a_r}$$

and are called the multisymmetric *power sums*. While there are relations among the power sums in finitely many variables, they freely generate the inverse limit Λ_r , making it a polynomial algebra. See [12] for details.

Our interest is in the case $r = 2$. Let us set

$$\xi_i = x_{i1}, \quad \eta_i = x_{i2},$$

and as before

$$B_n = R_{n,2} = k[\xi_1, \eta_1, \dots, \xi_n, \eta_n].$$

The subalgebra $\mathfrak{D}_* \subseteq \Lambda_2$ now has the following characterization.

Theorem 5.1. *As a subalgebra of Λ_2 , \mathfrak{D}_* is freely generated by the set Δ of elements of the form $p_{a,b}$ where $b \geq 1$.*

Proof. We first consider the subalgebra $k[\Delta] \subseteq \mathfrak{D}_*$ generated by Δ . Because Δ is a subset of the free generating set of Λ_2 , it is clear that the composite

$$k[\Delta] \rightarrow \mathfrak{D}_* \hookrightarrow \Lambda_2$$

is injective and isomorphic onto a polynomial subalgebra, and, therefore, that $k[\Delta]$ is itself a polynomial subalgebra of \mathfrak{D}_* . In order to prove that $k[\Delta] = \mathfrak{D}_*$, we only need to count dimensions, and we formulate the counting in terms of the Hilbert series. Recall that for a graded k -vector space V_* , we have the Hilbert series

$$P(V_*) = \sum_i \dim_k(V_i)t^i.$$

The Hilbert series for a polynomial algebra on a single generator x of grading i is $(1 - t^i)^{-1}$. Moreover, if we are given two graded vector spaces V_* and W_* , then $P(V_* \otimes W_*) = P(V_*)P(W_*)$. Since there are n monomials of degree n in Δ , we find that the Hilbert series for $k[\Delta]$ is

$$P(k[\Delta]) = \prod_n (1 - t^n)^{-n}.$$

If we can show that the Hilbert series for \mathfrak{D}_* is equal to this series, the proof will be complete.

In Proposition 4.3, we found that a k -basis for \mathfrak{D}_* may be identified with a set of orbit representatives of the S_n -action on the set of all monomials which have the property that if the exponent of ξ_i is non-zero, then so is the exponent of η_i . Such a set of representatives is given by the set of monomials of the form

$$\xi_1^{a_1} \eta_1^{b_1} \cdots \xi_l^{a_l} \eta_l^{b_l}, \quad \varphi^{-1}(a_i, b_i) \geq \varphi^{-1}(a_{i+1}, b_{i+1}),$$

where $l \leq n$, and $\varphi: \mathbb{N}^+ \rightarrow \mathbb{N}^+ \times \mathbb{N}$ is the bijection

$$(\varphi_1, \varphi_2, \dots) = ((1, 0), (1, 1), (2, 0), (1, 2), (2, 1), (3, 0), (1, 3), \dots)$$

onto the set of possible nonzero exponents. The dimension of the k -graded component of $B_n^{S_n}$ is just the number of these monomials of degree k . Let us say that $(a, b) \leq (c, d)$ when $\varphi^{-1}(a, b) \leq \varphi^{-1}(c, d)$, and let $f(a, b, k)$ denote the number of sequences $(a_1, b_1, \dots, a_l, b_l)$ such that

$$(a, b) \geq (a_1, b_1) \geq \cdots \geq (a_l, b_l), \quad (a_i, b_i) \in \mathbb{N}^+ \times \mathbb{N}$$

and

$$\sum_{i=1}^l (a_i + b_i) = k,$$

with no restrictions on l . It is easy to check that it satisfies the recursion relation

$$f(a, b, k) = \sum_{(c,d) \leq (a,b)} f(c, d, k - c - d).$$

This corresponds to a rule for the generating function $f_{a,b}(t) = \sum_k f(a, b, k)t^k$,

$$f_{a,b}(t) = 1 + \sum_{(c,d) \leq (a,b)} t^{c+d} f_{c,d}(t) = 1 + \sum_{(c,d) < (a,b)} t^{c+d} f_{c,d}(t) + t^{a+b} f_{a,b}(t).$$

Solving for $f_{a,b}(t)$ gives

$$(1 - t^{a+b})f_{a,b}(t) = 1 + \sum_{(c,d) < (a,b)} f_{c,d}(t) = f_{a',b'}(t),$$

where (a', b') is the element immediately below (a, b) under φ . It is readily verified that the formula

$$g_{a,b}(t) = (1 - t^{a+b})^{-a} \prod_{1 \leq k \leq a+b-1} (1 - t^k)^{-k}$$

satisfies the same recursion relation, and, therefore, that $f_{a,b}(t) = g_{a,b}(t)$. Taking limits

$$\lim_{n \rightarrow \infty} P(B_n^{S_n}) = \lim_{(a,b) \rightarrow \infty} \sum_{k \geq 1} f(a, b, k)t^k = \prod_{k \geq 1} (1 - t^k)^{-k}$$

gives the result. □

6. Two properties of the ring of functions $A_{\mathfrak{B}^{sch}}^{fin}$

Proposition 6.1. *The map $J^* : A_{\mathfrak{B}^{sch}}^{fin} \rightarrow F(\mathfrak{B}, \mathbb{R})$ defined in Section 4 is injective.*

Proof. There is a commutative diagram of rings

$$\begin{array}{ccc} A_{\mathfrak{B}^{sch}}^{fin} & \xrightarrow{J^*} & F(\mathfrak{B}, \mathbb{R}) \\ \downarrow f & & \downarrow \\ A_{Sp(\mathbb{A}^2, \emptyset)}^{fin} & \xrightarrow{g} & F(Sp^n(\mathbb{R}^2), \mathbb{R}) \end{array}$$

induced by the map of schemes $Sp(\mathbb{A}^2, \emptyset) \rightarrow \mathfrak{B}^{sch}$, and corresponding sets of real points. We claim that the homomorphisms f and g are injective, which will give the result. The map f factors as

$$A_{\mathfrak{B}^{sch}}^{fin} \xrightarrow{f'} A_{Sp(\mathbb{A}^2, 0)}^{fin} \xrightarrow{f''} A_{Sp(\mathbb{A}^2, \emptyset)}^{fin}.$$

The map f' is injective because the map of inverse systems defining it is given by the inclusions $D(n)_* \rightarrow A_{Sp^n(\mathbb{R}^2)}^{gr}$. That the maps f'' and g are injective follows from the observation in [12] that the inverse limit defining \mathfrak{D}_* stabilizes in each graded dimension. □

Proposition 6.2. *The ring of functions $J^*(A_{\mathfrak{B}^{sch}}^{fin}) \subseteq F(\mathfrak{B}, \mathbb{R})$ separates points.*

Proof. We first remark that for any real variety V (i.e., the scheme attached to a reduced Noetherian \mathbb{R} -algebra), the affine coordinate ring $A(V)$ does separate points. The corresponding fact for varieties over \mathbb{C} is an immediate consequence of the Nullstellensatz, so if we have two distinct real points x and y of V , there is an element $f \in A(V)_{\mathbb{C}} = A(V) \otimes_{\mathbb{R}} \mathbb{C}$ with the property that $f(x) = 1$ and $f(y) = 0$. Complex conjugation acts on $A(V)_{\mathbb{C}}$ with $A(V)$ as its fixed point subring, and we may consider the

function $g = f\bar{f} \in A(V)$. From its definition it is clear that $g(x) = 1$ and $g(y) = 0$, since x and y are real points, which gives this result.

We next note that there is a map $e: \mathfrak{B}^{sch} \rightarrow Sp(\mathbb{A}^1, 0)$ defined by the induced map of the map of pairs $(\mathbb{A}^2, \Delta) \rightarrow (\mathbb{A}^1, 0)$ given by $(x, y) \rightarrow y - x$. We next claim that the ring of functions $A_{Sp(\mathbb{A}^1, 0)}^{fin}$ separates points in $Sp(\mathbb{R}, 0)$. Since points in $Sp(\mathbb{R}, 0)$ can be viewed as the quotient of the collection of finite submultisets of \mathbb{R} by the equivalence relation generated by the relation $\{x_1, \dots, x_n\} \sim \{x_1, \dots, x_n, 0\}$, if we are given two distinct elements $x, y \in Sp(\mathbb{R}, 0)$, we may assume that they lie within $Sp^n(\mathbb{R})$ for some n , and then obviously are distinct. Consequently, from the remarks above, there is an element f in $A_{Sp^n(\mathbb{A}^1)}$ which separates x and y . The restriction map $A_{Sp(\mathbb{A}^1, 0)}^{fin} \rightarrow A_{Sp^n(\mathbb{A}^1)}$ is surjective in this case, from which it readily follows that $A_{Sp(\mathbb{A}^1, 0)}^{fin}$ separates points as stated above. It now follows that any two points in \mathfrak{B} which are not separated by the ring of functions $A_{\mathfrak{B}^{sch}}^{fin}$ must have the same values under e . Since \mathfrak{B} can be described as the quotient of the collection of multisets in \mathcal{J} by the equivalence relation generated by the relation

$$\{[x_1, y_1], [x_2, y_2], \dots, [x_n, y_n]\} \sim \{[x_1, y_1], [x_2, y_2], \dots, [x_n, y_n], [x, x]\}$$

that any two points which are not separated by $A_{\mathfrak{B}^{sch}}^{fin}$ must be represented by families of intervals of the same cardinality, and without any intervals of length 0. Consequently, they can be represented as two real points ξ and η in the scheme $Sp^n(\mathbb{A}^2)$ which lie in the complement of the union of the images of all the hyperplanes $x_i = y_i$. There is, therefore, a function $f \in A_{Sp^n(\mathbb{A}^2)}$ which separates ξ and η . Let θ denote the product of all the linear terms η_i . Then θ is another element of $Sp^n(\mathbb{A}^2)$ which does not vanish on either ξ or η . Therefore, the function $\theta \cdot f$ also separates ξ and η , and because it vanishes on the entire union of the hyperplanes defined by $y_i - x_i$, it lies in the ring $D(n)_*$. Since by Proposition 3.1 $A_{\mathfrak{B}^{sch}}^{fin}$ is the inverse limit of the inverse system of the graded rings $D(n)_*$, and the maps in the inverse system are all surjections, there is an element of $A_{\mathfrak{B}^{sch}}^{fin}$ which restricts to $\theta \cdot f$, and, therefore, separates ξ and η . \square

Remark. The import of these two results is that they assert that our functions provide a family of coordinates that faithfully represent the space of barcodes, without redundancy. This is an important philosophical point, and we believe that it will be useful in applications as well.

7. Machine learning on \mathfrak{B} with examples

7.1. Digits example

To illustrate the classification potential of this technique, we apply it to the MNIST database [15], of handwritten digits. We emphasize that the aim is not to outperform existing machine learning algorithms for digit classification, but to present an example that demonstrates one way of combining this technique with existing machine learning techniques. While it is clear that pure topological classification cannot distinguish between the digits (there are three numbers that do not have any loops, three that always have loops, one that has two loops and three that have style-dependent loops),

we can use the power of persistent homology to sift out more information. We begin by showing the full analysis of a few digits and then give the empirical results of applying this technique to a subset of the MNIST database.

7.1.1. Topological methods

We begin by describing a particular graph construction given a digital image. We treat the pixels as vertices and add edges between adjacent pixels (including diagonals). We can now define a filtration on the vertices of the graph corresponding to the image pixels. A natural filtration could be constructed using the pixel intensities of the original image (see Figure 6, Section 7.2). Another filtration, used in [11], can be constructed by thresholding, to produce a binary image, and adding 1-pixels as we sweep across the image. This adds spatial information into what would otherwise be a purely topological measurement. Since the orientation of the digit matters (a 6 is the same as a 9 given a 180 degree rotation), we choose the latter approach and sweep across the rows and columns of each digit.

By taking into account spatial information, we get a rough view of the location of various topological features. For example, though a ‘9’ and ‘6’ both have one connected component and a single loop, the loop will appear at different locations in the top-down filtration for the ‘9’ and ‘6’. The digits and one of the resulting bar codes are shown in Figures 1 and 2. Using all four sweeps, and both the Betti 0 and Betti 1 bar codes, reveals additional differences between each of the digits. Here y_{max} denotes the maximum y -value occurring in the various barcodes constructed across all the images.

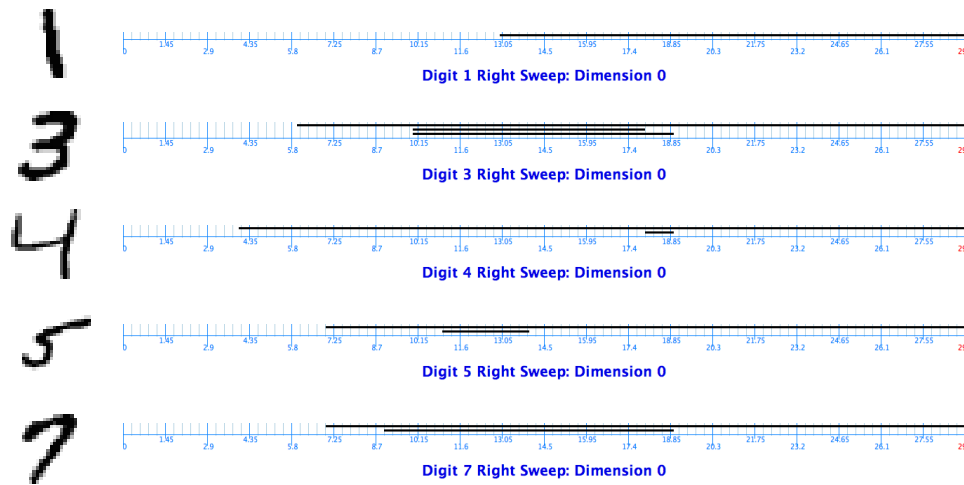


Figure 1: No loop digits with Betti 0 bar code, sweep to right

7.1.2. Feature selection

We can use the techniques described in this paper to coordinatize the bar code space \mathfrak{B} . In machine learning terminology, these coordinates are called *features*. This allows us to characterize the bar codes generated by each data point as a compact feature

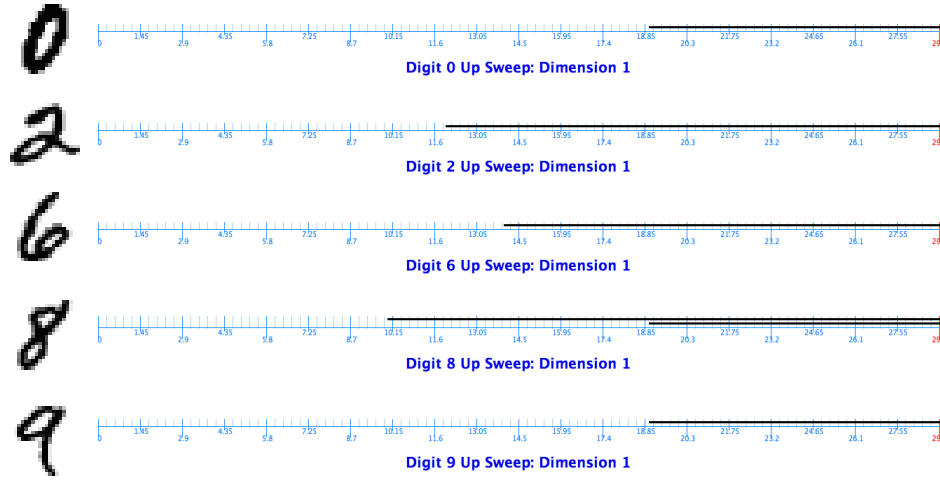


Figure 2: Loop digits with Betti 1 bar code, sweep to top

vector. This also gives us great flexibility in selecting features that work well with our data. We can then apply a standard machine learning algorithm, such as a support vector machine (SVM), to classify the data.

We selected a set of four features from the invariants discussed in this paper. Intuitively, the exponents in each polynomial will give the relative value of small bars or endpoints compared to large bars or endpoints. For example, if comparing two bars of length $\frac{b}{2}$ and b , the first bar will have more weight in an invariant linear polynomial than in an invariant quadratic polynomial. Indeed,


$$\begin{aligned} \left(\frac{b}{2}\right)^2 &= \frac{b^2}{4}, \\ \left(\frac{b}{2}\right)^3 &= \frac{b^3}{8}, \\ \left(\frac{b}{2}\right)^4 &= \frac{b^4}{16}, \\ &\vdots \end{aligned}$$

We selected four features:


$$\begin{aligned} &\sum_i x_i(y_i - x_i), \\ &\sum_i (y_{max} - y_i)(y_i - x_i), \\ &\sum_i x_i^2(y_i - x_i)^4, \\ &\sum_i (y_{max} - y_i)^2(y_i - x_i)^4, \end{aligned}$$

which when applied to the four sweeps, each with a 0-dimensional and 1-dimensional

bar code, gives a feature vector of total size 32 which we then arranged into a feature matrix. Intuitively speaking, the first two features take all of the bars, lengths and endpoints, into account. The second two features heavily favor the arrangement of longer bars. A visualization of a matrix of 10,000 digits using classical multidimensional scaling (MDS) is shown in Figure 3 and the spectrum of the matrix is shown in Figure 4.



(a) A 2D view of the data



(b) A 3D view of the data

Figure 3: Visualization of data using topological features




Figure 4: Normalized spectrum of topological feature matrix

As is typical when using a SVM, we scaled each coordinate such that the values were between 0 and 1. The SVM was implemented using software provided by [9].

7.1.3. Classification results

We applied these methods on a subset of 1000 digits from the MNIST database to tune parameters of the algorithm and test various kernels. For the radial basis function $e^{-\gamma|u-v|^2}$ (RBF, also known as the Gaussian kernel), we used $\gamma = 8$. For the polynomial kernel $(\gamma(u * v) + a)^d$, we used $d = 3$ with $\gamma = 2$ and $a = 2$. In both functions, u and v represent the calculated feature vectors. After this, we progressively increased the size of the subset to 10,000 handwritten digits.

The classification accuracy was measured by partitioning the data set into one hundred subsets and using cross-validation successively on each subset. The results are shown in Table 1.

Table 1: Classification accuracy of two SVM kernels

SVM	1000 Digits	5000 Digits	10000 Digits
Gaussian	87.70%	91.54%	92.04%
Polynomial	88.00%	91.62%	92.10%

With the polynomial kernel, an error of 7.9% is seen. As mentioned above, the

purpose of this test is not to outperform existing classification algorithms but to demonstrate one application of the topological features. In line with this, we examined some of the digits that the algorithm failed on. Figure 5 shows a few of the typical problem digits.

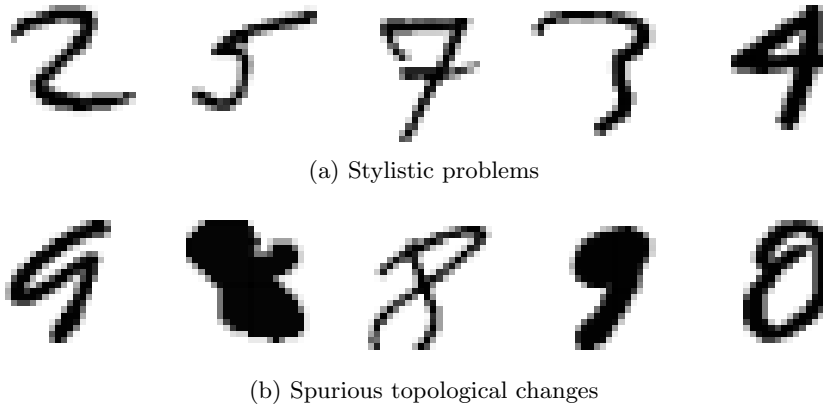


Figure 5: Common misclassifications

The most common confusion is between a ‘5’ and a ‘2’ written with no loop. Other confusions often occur between the shown style of ‘7’ and slanted ‘3’s and between a certain style of ‘4’ and a ‘9’. These confusions are not unexpected since these numbers are topologically the same. The extra spatial information added by the directional sweeps is sensitive to variations in the slant or style of handwriting and a visual inspection of these digits suggests why the algorithm has difficulty classifying these particular examples. Other common confusions occur when topological changes occurred to the digit, specifically when the writer adds or removes a loop.

7.2. Hepatic lesion classification

In this example, we apply topological features to classifying hepatic lesions. The dataset consists of computed tomography (CT) scans of 132 hepatic lesions that are outlined and annotated by radiologists. There are nine diagnoses represented in the data: cysts (45 lesions), metastases (45 lesions), hemangiomas (18 lesions), hepatocellular carcinomas (HCC, 11 lesions), focal nodules (5 lesions), abscesses (3 lesions), neuroendocrine neoplasms (NeN, 3 lesions), a single laceration and a single fat deposit. Additionally, there are no controls for the size of the lesion and the lesions vary from under 100 pixels to 10,000 pixels. Because of the unbalanced nature of the data, we focus on the subset of cysts, metastases, and hemangiomas.

Classification results using the bar code metric (matching metric) were first presented in [2], and we follow the same methods for processing and generating bar codes from the data. We will briefly describe the methods here. For a more detailed account, please read [2].

7.2.1. Topological methods

As mentioned above, a natural filtration for an image is to filter by the pixel intensity. An example of this filtration is given in Figure 6. The variation in pixel intensity allows

us to use a one-dimensional filtration on the pixel intensity, but as the results will show, the classification is improved when geometric information is added into the filtrations.

As there is no rotational orientation of the lesions, we cannot add in geometric information using the sweeps described in the previous section. Instead, we use the lesion border provided by the radiologist and assign each pixel its distance from the border. Then, by using two-dimensional homology, we achieve improved results, especially in the case of the hemangiomas which are characterized by large dense regions on the outer part of the lesion. Because two-dimensional filtrations are computationally intensive, we approximate the two-dimensional filtration with one-dimensional bar code ‘slices’ along the border filtration axis. We use 7 slices per lesion and both the Betti 0 and Betti 1 bar codes.

Note that we can look at each filtration from each direction and catch different features. The intensity filtration can add high intensity pixels first or low intensity pixels first. The boundary filtration can begin with pixels near the boundary first or pixels far from the boundary first. This yields 56 one-dimensional bar codes per lesion.

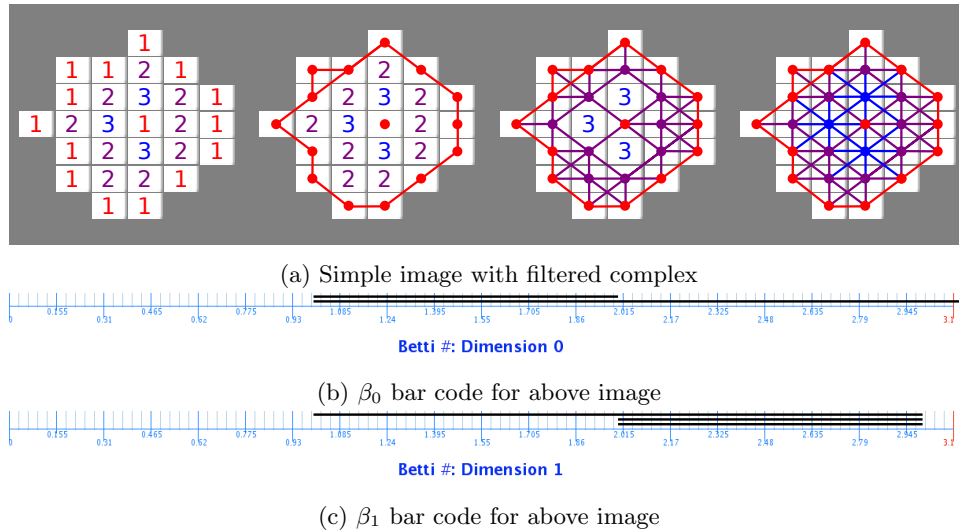


Figure 6: Constructing an increasing 1D-filtration on an image [2]

7.2.2. Feature selection

We use a slightly different set of four features as compared to the digits example. These features are shown below. The two sets of features that focus on long bars and features which take into account shorter bars is used here. In this application, this is analogous to filtering the bar code to remove the large number of smaller bars. Because of the variations in lesion size, we look at the average over each bar to try and eliminate the effects of large variations in lesion size:

$$\sum_i^n x_i(y_i - x_i)/n,$$

$$\begin{aligned} & \sum_i^n (y_{max} - y_i)(y_i - x_i)/n, \\ & \sum_i^n x_i^2(y_i - x_i)^4/n, \\ & \sum_i^n (y_{max} - y_i)^2(y_i - x_i)^4/n. \end{aligned}$$

Here y_{max} denotes the largest y -value occurring in the entire data set. As mentioned above, we have 56 bar codes per lesion. With four features, this yields a feature vector of 224 features for each lesion.

7.2.3. Classification results

We apply the SVM using only the Gaussian kernel and use an exponential parameter sweep to find optimal values of γ for each method. We use LOOCV (leave one out cross-validation) to calculate the classification accuracies. The results are shown below. Table 2 gives the results for 1D and 2D filtrations for several different datasets while Table 3 shows how well the algorithm performs on different lesion types for the different filtrations. Table 4 demonstrates the effect of size on classification.

Table 2: SVM classification accuracies for 1D and 2D filtrations

Filtration	Full	HcHeCM	HeCM	CM
1D (Intensity)	53.03%	59.66%	65.74%	75.56%
2D	67.42%	74.79 %	81.48%	86.67%

Table 3: HeCM % classification accuracy by lesion type

Filtration	% of HeCM	% of Heman.	% of Cysts	% of Metas.
1D	65.74%	33.33%	75.56%	68.89%
2D	81.48%	61.11%	86.67%	84.44%

Table 4: Classification by lesion size of HeCM

Lesion Size by Area	% Accu.	# of Heman.	# of Cysts	# of Metas.
All	81.48%	18	45	45
<10000 px	82.52%	18	42	43
<5000 px	84.78%	16	39	37
<2500 px	86.25%	14	32	34
<1250 px	88.514%	8	28	23

Using [2], we see that the topological features are comparable with using the matching metric to generate features. The results from the HeCM dataset for the two methods are shown below. They reflect the correct classification of a single lesion using the topological features, making the two methods virtually the same for this

subset of the data. Comparing with the other results in [2] shows that the two results are very close in most categories, with each slightly outperforming the other in certain subsets of the data.

Table 5: Classification methods

Filtration	bar code Features	Matching Metric
1D	65.74%	63.80%
2D	81.48%	80.56%

7.3. Discussion

These two examples demonstrate the classifying power of topological features when applied to real world datasets. This was done using off-the-shelf machine learning algorithms showing that these features can easily be combined with more traditional classification methods adding a set of additional classification features to the machine learning toolbox.

These examples also show the power of combining topology with geometry. In both datasets, this is an integral part of the classification procedure. The results in the hepatic lesion dataset provide an especially good example of the potential gains that can be achieved by combining both fields.

In summary, using algebraic geometry and invariant theory, we have identified a family of coordinates on the space of finite metric spaces, or sampled shapes. These coordinates can serve as a method for organizing the collection of all bar codes, and, therefore, any database whose members produce bar codes. Of course, we can also use various metrics on bar code space, such as the bottleneck or Wasserstein distances. It would be extremely interesting to analyze the relationship between these distances on bar code spaces with various more algebraic notions of distance on the bar code coordinates. It would also be very interesting to define and analyze analogous coordinates on spaces of multidimensional persistence modules, where they might give information which is currently not accessible due to the complexity of the algebraic descriptions of multidimensional persistence modules.

References

- [1] H. Adams and G. Carlsson, *On the non-linear statistics of range image patches*, SIAM J. Imaging Sci. 2 (1) (2009), 110–117.
- [2] A. Adcock, D. Rubin, and G. Carlsson, *Classification of hepatic lesions using the matching metric*, Computer Vision and Image Understanding 121 (2014), 36–42.
- [3] G. Brumfiel, *Partially Ordered Rings and Semi-Algebraic Geometry*, London Math. Soc. Lecture Note Ser., 37. Cambridge University Press, Cambridge, New York, 1979.
- [4] G. Carlsson, *Topology and data*, Bull. Amer. Math. Soc. (N.S.) 46 (2) (2009), 255–308.

- [5] G. Carlsson, *Topological pattern recognition for point cloud data*, Acta Numer. 23 (2014), 289–368.
- [6] G. Carlsson, T. Ishkhanov, V. de Silva, and A. Zomorodian, *On the local behavior of spaces of natural images*, Int. J. Comput. Vision 76 (1) (2008), 1–12.
- [7] G. Carlsson, A. Zomorodian, A. Collins, and L. Guibas, *Persistence bar codes for shapes*, Int. J. Shape Model. 11 (2005), 149–187.
- [8] J. Chan, G. Carlsson, and R. Rabadan, *Topology of viral evolution*, Proc. Nat. Acad. Sci. 110 (46) (2013), 18566–18571; published ahead of print October 29, 2013, doi:10.1073/pnas.1313480110.
- [9] Chih-Chung Chang and Chih-Jen Lin, *LIBSVM: a library for support vector machines*, ACM Trans. Intell. Syst. Technol. 2 (27) (2011), 1–27. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [10] F. Chazal, D. Cohen-Steiner, L. Guibas, F. Mémoli, and S. Oudot, *Gromov-Hausdorff stable signatures for shape using persistence*, Comput. Graph. Forum 28 (5) (2009), 1393–1403.
- [11] A. Collins, A. Zomorodian, G. Carlsson, and L. Guibas, *A bar code shape descriptor for curve point cloud data*, Comput. Graph. 28 (2004), 881–894.
- [12] J. Dalbec, *Multisymmetric functions*, Beiträge Algebra Geom. 40 (1) (1999), 27–51.
- [13] W. Fulton and R. Weiss, *Algebraic Curves: An Introduction to Algebraic Geometry*. Addison–Wesley, 1989.
- [14] J. Jantzen, *Representations of Algebraic Groups*, Second edition, Math. Surveys Monogr., 107. American Mathematical Society, Providence, RI, 2003.
- [15] Y. LeCun and C. Cortes, *The MNIST database*, Courant Institute NYU, Accessed 16 July 2012, <http://yann.lecun.com/exdb/mnist/>.
- [16] D. Mumford, J. Fogarty, and F. Kirwan, *Geometric Invariant Theory*. Springer Verlag, 2002.
- [17] A. Zomorodian and G. Carlsson, *Computing persistent homology*, Discrete Comput. Geom. 33 (2) (2005), 247–274.

Aaron Adcock aadcock@fb.com

Facebook, 8th Floor, 770 Broadway, New York, NY 10003, USA

Erik Carlsson erikcarlsson17@gmail.com

Center of Mathematical Sciences and Applications, Harvard University, Cambridge, MA 02138, USA

Gunnar Carlsson gunnar@math.stanford.edu

Department of Mathematics, Stanford University, Stanford, CA 94305, USA