

RESEARCH ARTICLE

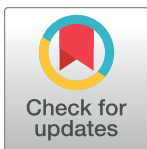
Determining clinically relevant features in cytometry data using persistent homology

Soham Mukherjee¹, Darren Wethington^{2,4}, Tamal K. Dey^{1*}, Jayajit Das^{2,3,4,5*}

1 Computer Science Department, Purdue University, West Lafayette, Indiana, United States of America, **2** Biomedical Sciences Graduate Program, College of Medicine, The Ohio State University, Columbus, Ohio, United States of America, **3** Biophysics Program, The Ohio State University, Columbus, Ohio, United States of America, **4** Battelle Center for Mathematical Medicine, Abigail Wexner Research Institute, Nationwide Children's Hospital, Columbus, Ohio, United States of America, **5** Departments of Pediatrics, Biomedical Informatics, and, Pelotonia Institute of Immuno-Oncology, College of Medicine, The Ohio State University, Columbus, Ohio, United States of America

☞ These authors contributed equally to this work.

* jayajit@gmail.com (JD); tamaldehy@purdue.edu (TKD)



OPEN ACCESS

Citation: Mukherjee S, Wethington D, Dey TK, Das J (2022) Determining clinically relevant features in cytometry data using persistent homology. PLoS Comput Biol 18(3): e1009931. <https://doi.org/10.1371/journal.pcbi.1009931>

Editor: James R. Faeder, University of Pittsburgh, UNITED STATES

Received: July 19, 2021

Accepted: February 16, 2022

Published: March 21, 2022

Copyright: © 2022 Mukherjee et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Our analysis focuses on the deposited data available at <https://premium.cytobank.org/cytobank/experiments/308357> for non-naïve CD8+ T cells collected at the time of admission (and not any later blood draws, such as at 7 days after admission). This data was originally published by Mathew et al. Please note that a Cytobank account is currently required for data access. Our current code is available at <https://github.com/soham0209/TopoCytometry> and will be updated for ease-of-use and performance enhancements.

Abstract

Cytometry experiments yield high-dimensional point cloud data that is difficult to interpret manually. Boolean gating techniques coupled with comparisons of relative abundances of cellular subsets is the current standard for cytometry data analysis. However, this approach is unable to capture more subtle topological features hidden in data, especially if those features are further masked by data transforms or significant batch effects or donor-to-donor variations in clinical data. We present that persistent homology, a mathematical structure that summarizes the topological features, can distinguish different sources of data, such as from groups of healthy donors or patients, effectively. Analysis of publicly available cytometry data describing non-naïve CD8+ T cells in COVID-19 patients and healthy controls shows that systematic structural differences exist between single cell protein expressions in COVID-19 patients and healthy controls. We identify proteins of interest by a decision-tree based classifier, sample points randomly and compute persistence diagrams from these sampled points. The resulting persistence diagrams identify regions in cytometry datasets of varying density and identify protruded structures such as 'elbows'. We compute Wasserstein distances between these persistence diagrams for random pairs of healthy controls and COVID-19 patients and find that systematic structural differences exist between COVID-19 patients and healthy controls in the expression data for T-bet, Eomes, and Ki-67. Further analysis shows that expression of T-bet and Eomes are significantly downregulated in COVID-19 patient non-naïve CD8+ T cells compared to healthy controls. This counter-intuitive finding may indicate that canonical effector CD8+ T cells are less prevalent in COVID-19 patients than healthy controls. This method is applicable to any cytometry dataset for discovering novel insights through *topological data analysis* which may be difficult to ascertain otherwise with a standard gating strategy or existing bioinformatic tools.

Funding: JD was supported by National Institutes of Health (NIH; <https://www.nih.gov>) under grant number R01-AI 143740 and R01-AI 146581, and by Nationwide Children's Hospital (Nationwide Children's; <https://www.nationwidechildrens.org>). TKD was supported by National Science Foundation (NSF; <https://www.nsf.gov>) under the grant number CCF 1839252 and CCF 2049010. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Author summary

Identifying differences between cytometry data seen as a point cloud can be complicated by random variations in data collection and data sources. We apply *persistent homology* used in *topological data analysis* to describe the shape and structure of the data representing immune cells in healthy donors and COVID-19 patients. By looking at how the shape and structure differ between healthy donors and COVID-19 patients, we are able to definitively conclude how these groups differ despite random variations in the data. Furthermore, these results are novel in their ability to capture shape and structure of cytometry data, something not described by other analyses.

1 Introduction

Cytometry data contain information about the abundance of proteins in single cells and are widely used to determine mechanisms and biomarkers that underlie infectious diseases and cancer. Recent advances in flow and mass cytometry techniques enable measurement of abundances of over 40 proteins in a single cell [1, 2]. Thus, in the space spanned by protein abundance values measured in cytometry experiments, a cytometry dataset is represented by a point cloud composed of thousands of points where each point corresponds to a single cell. Abundances of proteins or chemically modified forms (e.g., phosphorylated forms) of proteins in single immune cells change due to infection of the host by pathogens (e.g., a virus) or due to the presence of tumors which usually result in changes in the 'shape' of point cloud data measured in cytometry experiments [3–5]. Cytometry data analysis techniques commonly rely on Boolean gating and calculation of relative proportions of resulting populations as a method to compare datasets across control/healthy and experimental/diseased conditions. In recent years, state-of-the-art analyses based on sophisticated machine learning algorithms capable of mitigating batch effects, ad hoc gating assumptions, and donor-donor variability have been developed [6, 7]. However, these methods are not designed to quantitatively characterize shape features (e.g., connected clusters, cycles) in high dimensional cytometry datasets that can contain valuable information regarding unique co-dependencies of specific proteins in diseased individuals compared to healthy subjects.

Topological Data Analysis (TDA) aims to capture the underlying shape of a given dataset by describing its topological properties. Unlike geometry, topological features (e.g., the hole in a doughnut) are invariant under continuous deformation such as rotation, bending, twisting but not tearing and gluing. One of the tools by which TDA describes topological features latent in data is persistent homology [8, 9]. For example, for a point cloud data, persistent homology captures the birth and death of topological features (e.g., 'holes') in a dataset after building a scaffold called a simplicial complex out of the input points. This exercise provides details regarding topological features that 'persist' over a range of scale and thus contain information regarding the shape topology at different length scales (see S1 Fig for details). Persistent homology has been applied to characterize shapes and shape-function relationships in a wide variety of biological systems including skin pattern formation in zebra fish [10], protein structure, and pattern of neuronal firing in mouse hippocampus [11]. TDA has additionally previously been applied to identify immune parameters associated with transplant complications for patients undergoing allogenic stem cell transplant using populations of immune cell types assayed via mass cytometry [12]. However, this work did not use persistent homology or expression levels of proteins in their analysis, leaving the shape of cytometry data uncharacterized. Another work focuses on the use of TDA as a data reduction method for single-cell RNA

sequencing data [13], but again do not attempt to characterize how topologies derived from point clouds differ among disparate data sources such as healthy and diseased individuals.

The challenges of directly applying current persistence methodologies to cytometry data to characterize distinguishing features between healthy and diseased states are the following: 1. Features that separate healthy from diseased state can pertain to the change in density of points in a region in point cloud data—therefore, the information of local density should be incorporated in persistent homology methods, in particular in the filtration step that brings in sequentially the simplices connecting the points. In commonly used Rips filtration [14] the density of points is not included. 2. There can be shape changes giving a different length scale in the point cloud data, such as formation of an elbow, in a diseased condition. 3. There can be systematic differences between healthy and diseased states across batch effects and donor-donor variations. Topological features should capture these global differences being oblivious to the local variations caused by measurement noise.

We address the above challenges by developing an appropriate filtration function to compute persistence and applying the method to characterize distinguishing features of non-naïve CD8+ T cells between healthy and SARS-CoV-2 infected patients.

2 Results

2.1 Persistence framework for SARS-CoV-2 infection

Topological signatures given by persistence are stable, global, scale invariant and show resilience to local perturbations [15]. It is this property of persistent homology that motivates us to use TDA in distinguishing clinically relevant features in flow cytometry data in COVID-19 patients.

Persistent homology. Persistent homology builds on an algebraic structure called homology groups graded by its dimension i and denoted by H_i . Intuitively, they describe the shape of the data by ‘connectivity’ at different levels. For example, H_0 describes the number of connected components, H_1 describes the number of holes, and, H_2 describes the number of enclosed voids apparently present in the shape that the dataset represents. Three and higher dimensional homology groups capture analogous higher (≥ 3) dimensional features. A point cloud data (henceforth abbreviated as PCD) itself does not have much of a ‘connected structure’. So, a scaffold called a *simplicial complex* is built on top of it. This simplicial complex, in general, is made out of simplices of various dimensions such as vertices, edges, triangles, tetrahedra, and other higher dimensional analogues. Given a growing sequence of such complexes called *filtrations*, a persistence algorithm tracks information regarding the homology groups across this sequence. In our case, these complexes can be restricted only to vertices and edges. With the restriction that both vertices of an edge appear before the edge, we get a nested sequence of graphs

$$G_0 \subset G_1 \subset G_2 \subset \dots G_n$$

as the filtration. Fig 1 shows such a *filtration*.

Persistence diagram. Appearance (‘birth’) and disappearance (‘death’) of topological features, that is, cycles whose classes constitute the homology groups, can be captured by persistence algorithms [8, 16]. These ‘birth’ and ‘death’ events are represented as points in the so-called *persistence diagram*. If a topological feature is born at filtration step b and dies at step d , we represent this by *persistence pair* (b, d) with persistence $d - b$. The pair (b, d) becomes a point in the persistence diagram with the ‘birth’ as x-axis and ‘death’ as y-axis. This 2D plot summarizes topological features latent in the data. In the example-filtration of Fig 1, a new component gets ‘born’ when a vertex v_i appears in the filtration for the first time. When an

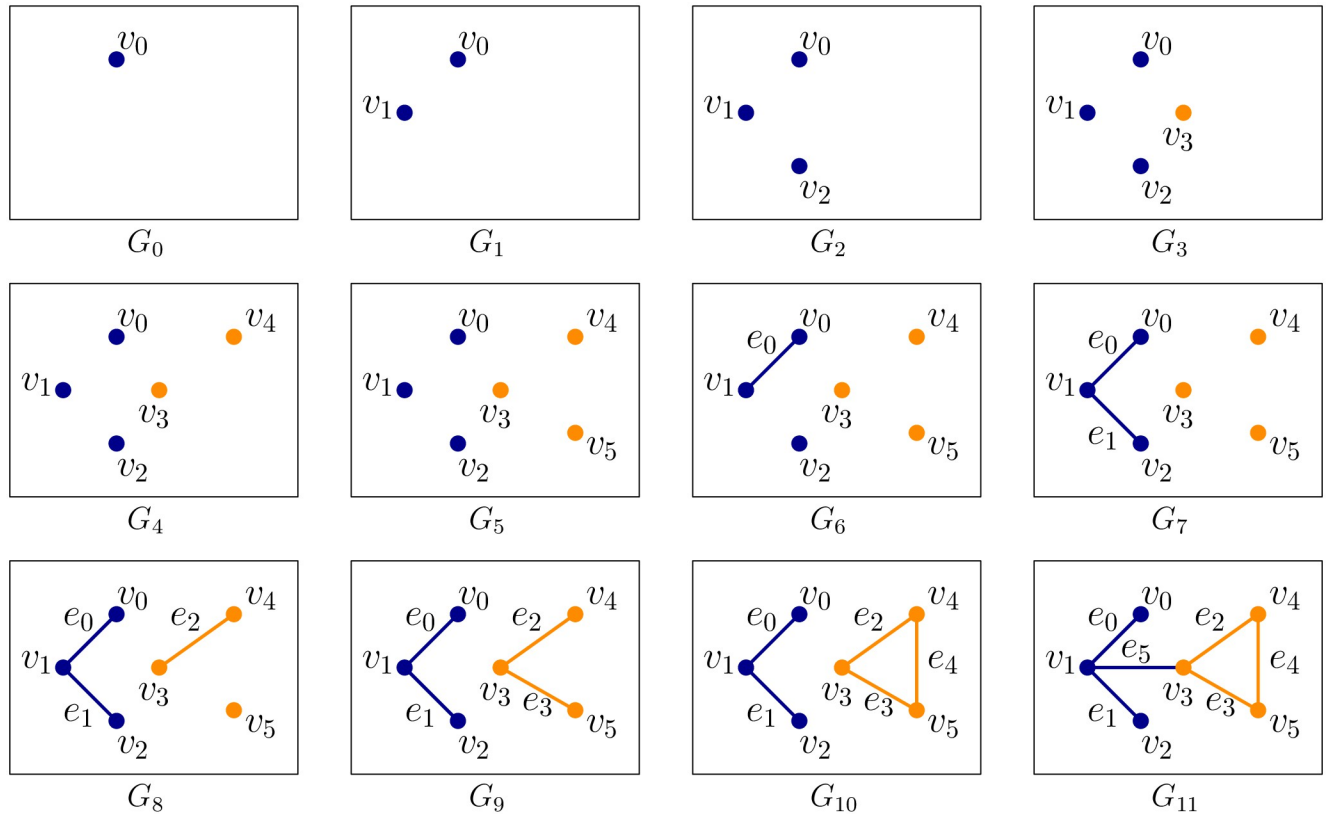


Fig 1. An example of filtration for a graph. The nested sequence of graphs $G_0 \subset G_1 \subset \dots \subset G_{11}$ forms a filtration of the final graph G_{11} . Each vertex v_i creates a new component in the nested sequence, and edges e_0, e_1, e_2, e_5 merge two components whereas e_4 creates a cycle (yellow).

<https://doi.org/10.1371/journal.pcbi.1009931.g001>

edge is introduced, one of the two things can happen—either two components are joined, or a cycle is created. In the first case, a ‘death’ happens for 0-th homology group H_0 , and in the second case, a ‘birth’ happens for the 1-st homology group H_1 . For example, when e_0 comes in the filtration (G_6), it merges two components created by v_0 and v_1 . By convention, we choose to kill the component that got created later in the filtration and thus we let the component created by v_1 die. We obtain a persistence pairing (1, 6) since edge e_0 at filtration step 6 kills the component created by v_1 at step 1. Similarly, we obtain pairs (2, 7), (4, 8), (5, 9), and (3, 11). These points, tracking the ‘birth’ and ‘death’ of components, produce the persistence diagram for the 0-th homology group H_0 and hence we refer to it as the H_0 -persistence diagram. Note that the edge e_4 creates a cycle (yellow) that never dies. In such cases, *i.e.* when a topological feature never dies, we pair it with ∞ . For the edge e_4 , we obtain a persistence pair (10, ∞). But, this feature concerns the 1-st homology group H_1 and thus it becomes a point in the persistence diagram for H_1 which we refer to as H_1 -persistence diagram. One way to leverage the above framework for studying a function is to assign function values to vertices and edges and construct a filtration by ordering them according to these assigned values. For such cases the persistence pairs take the form (b, d) where b is the value at which a feature is born and d is the value at which it dies. The function values that induce the filtration (Fig 2) are chosen to capture two features of the input PCD—(i) the density variations, and (ii) the *anisotropy* of the features, that is, how elongated it is in a certain direction, henceforth termed as *length scale* ‘of the feature’ or collectively ‘of the data’. In particular, length scales refer to the prominence of protrusions such as ‘elbows’ in COVID-19 data.

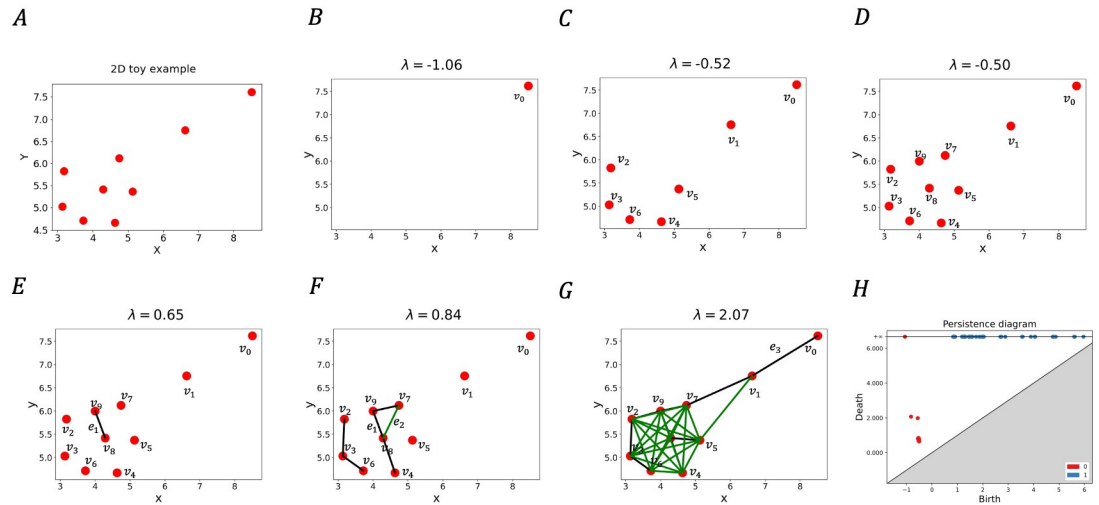


Fig 2. Illustration of persistence for a 2D point cloud data (PCD). (A), (H) shows a 2D PCD example and its computed persistence diagram. (B)-(G) shows important changes in topological feature as λ increases from $-\infty$ to ∞ . (B) At $\lambda = -1.06$ an isolated point, v_0 appears first. Note that each isolated vertex creates a new component. (C) At $\lambda = -0.52$ points in the denser region appears in the filtration, introducing more components. The indices of the vertices denote the order in which they appear in the filtration. (D) At $\lambda = -0.50$, all vertices appear in the filtration. Note that, the way we have chosen the filtration function f , vertices appear before the edges since $f_v(v)$ is always negative. (E) At $\lambda = 0.65$, the first edge e_1 appears merging two components. By persistence algorithm [8], we pair the edge e_1 with v_0 , since v_0 appears later in the filtration. Corresponding to this, we get a persistence pair $(f_v(v_0), f_e(e_1)) = (-0.50, 0.65)$. (F) At $\lambda = 0.84$, the green edge e_2 appears and creates a cycle. Since there is no 2-simplex(triangle) present, the cycle is never destroyed. In the persistence diagram we have this pair as $(f_e(e_3), \infty) = (0.84, \infty)$. (G) At $\lambda = 2.07$, the long edge e_3 appears joining v_0 and v_1 , yielding a persistence pair $(-1.06, 2.07)$.

<https://doi.org/10.1371/journal.pcbi.1009931.g002>

Below we briefly describe how we adapt the above persistence framework for analyzing point cloud data (PCD) representing CD8+ T cells in SARS-CoV-2 infection. Details regarding the approach are provided in Section 4 and [S2 Appendix](#) and [S1 Fig](#).

Computing persistent homology for cytometry datasets. Our datasets consist of cytometry data for non-naïve CD8+ T cells. Given protein expressions (real values) for d proteins in such a single cell, we can represent it as a d -dimensional point in \mathbb{R}^d . Considering a population of single cells, we get a point cloud (PCD) in \mathbb{R}^d . Now, we study the shape of this PCD using the persistence framework that we describe above. We compute persistence diagrams for the PCDs generated with protein expressions from different individuals and compare them. It turns out that, for computational purposes, we need a limit on the dimension d for PCD which means we need to choose carefully the proteins that differentiate effectively the subjects of our interest, namely the healthy individuals, COVID-19 patients, and recovered patients. We typically choose 3 (sometimes 2) protein expressions to generate the PCD and call it a PCD in the P_1, P_2, P_3 space if it is generated by proteins P_1, P_2 , and P_3 respectively.

Flow cytometry data for non-naïve CD8+ T cells in Mathew et al. [3] show generation of CD8+ T cells with larger abundances of the proteins CD38 and HLA-DR (CD38+HLA-DR+ cells) for some COVID-19 patients, forming an ‘elbow’ in the two dimensional PCD with CD38 and HLA-DR protein expressions (see [S2 Fig](#)). Moreover, there is an increase in the local density of the points (or single CD8+ T cells) in the ‘elbow’ region. This suggests that, to study the PCD generated by the protein expressions by persistence framework, we need to choose a filtration that is able to capture such geometric shapes and variations in the local density.

We briefly describe our choice of filtration by considering the example of a point cloud $P \subset \mathbb{R}^2$ shown in [Fig 2](#). Mathematical and computational details regarding the filtration are

provided in the Section 4. We build a filtration according to assigned values to the vertices and edges of a graph connecting the input points. For a vertex p which is a point in the input PCD P , we denote this value $f_v(p)$ (given by Eq 1 in Section 4). Similarly, we denote the assigned value to an edge e as $f_e(e)$ (given by Eq 2 in Section 4); see Fig 2. The values satisfy the conditions that $f_v(p) < 0$ and $f_e(e) \geq 0$; implications of this specific choice will become clear in the next paragraph. It is noteworthy to mention that $f_v(p)$ is the negative of distance-to-measure originally defined in [17] and later used in [18] for the PCD case and captures the density distribution of points, whereas $f_e(e)$ captures the inter-point distances between the points in the given point cloud.

The persistence algorithm processes each vertex and edge in the order of their appearance in the filtration. We execute it using a threshold value λ from $-\infty$ to ∞ and generate the persistence diagram accordingly. Intuitively, as λ is increased from $-\infty$ to ∞ , vertices p for which $f_v(p) \leq \lambda$ and edges for which $f_e(e) \leq \lambda$ appear in the filtration for a particular value of λ (see Fig 2). Since $f_v(p) < 0$ and $f_e(e) \geq 0$, all the vertices first appear as λ is increased from $-\infty$ to 0, and then edges start appearing as λ becomes positive. The birth-death events for H_0 and H_1 constituting the persistence diagram (Fig 2H) contain information about the density and length scales present in the point-cloud. For example, the points showing birth and death events for the H_0 -persistence diagram are more densely organized for the single cell protein expression data from the healthy donor than the SARS-CoV-2 infected patient in the HLA-DR—CD38 plane shown in S3 Fig. The denser organization of the birth-death events in the persistence diagram indicates a more homogeneous distribution of CD38 and HLA-DR proteins in the CD8+ T cells in healthy donors compared to that in infected patients. Most of the CD8+ T cells in healthy controls have low amounts of CD38 and HLA-DR abundances and few contain larger values of these proteins, indicating a greater degree of homogeneity. The birth-death events for H_1 in the persistence diagram (S3 Fig) in general contain information about the length scales of cyclic structures in the point cloud. It also can capture protrusions like ‘elbows’ that we have in COVID-19 data. Our filtration allows only birth (and not death) of 1-cycles and therefore, a λ value corresponding to the birth of a 1-cycle captures the length scale of the newly born cycle and hence an ‘elbow’. Our analysis of the PCDs in S3(D) and S3(H) Fig indeed shows that λ values for the birth of cycles for the COVID-19 patient is much larger compared to that for the healthy individual indicating the presence of larger length scales in the PCD which is consistent with the presence of an ‘elbow’ shape in the PCD for the patient.

2.2 Application of persistence to healthy and patient data

Our aim is to find out systematic differences in topological features extracted from cytometry data for healthy individuals and COVID-19 patients. Ideally one would like to compute persistence diagrams for all 25 proteins that were measured in single CD8+ T cells, however, this task encounters two major problems. First, as we mentioned before taking the full 25 dimensional PCD introduces the *curse of dimensionality* [19] making it computationally infeasible to produce the persistence diagrams. The second one is more subtle. In order to measure how the density of data differs from a healthy to infected person in a quantitative way, we need to ensure that the number of points in each PCD, to be analyzed by persistent homology, is the same. Cytometry data usually contain different numbers of single cells in datasets obtained from different donors or replicates. To address the *curse of dimensionality* we use a classifier (XGBoost [29]) that distinguishes single CD8+ T cells in healthy donors from those in COVID-19 patients and we choose the top r (taken to be 3) features (proteins) that are deemed important by the classifier while classifying the data points (cells). This reduces the dimension of the data from 25 to a much smaller value denoted r .

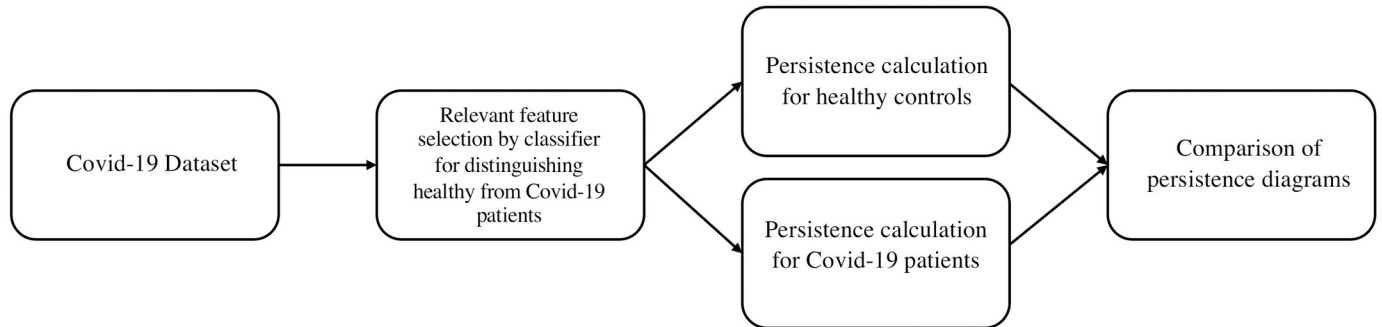


Fig 3. Flowchart of computation pipeline. The pipeline includes three main stages, namely, (i) relevant feature selection, (ii) persistence computation, and (iii) comparison of persistence diagrams.

<https://doi.org/10.1371/journal.pcbi.1009931.g003>

To address the second issue, we perform uniform random sampling on every r -dimensional dataset and take equal number of samples from it. We then use the filtration defined in Eqs 1 and 2 to construct persistence diagrams for each dataset. To quantify the structural differences in the datasets as captured by the corresponding persistence diagrams, we compute the Wasserstein distance [20] between persistence diagrams from randomly selected pairs of either two healthy donors (H×H) or a healthy donor and an infected patient (H×P) and compute distributions of the Wasserstein distances for a large number of (H×H) and (H×P) pairs. The comparison of these distributions via Kolmogorov-Smirnov (KS) tests provides information regarding the systematic differences in shape features in the CD8+ T cell cytometry data across healthy individuals and COVID-19 patients. The computational pipeline is summarized in (Fig 3). Below we describe results from the application of our computational pipeline to the CD8+ T cell cytometry data in Mathew et al. [3]

A few protein expressions in CD8+ T cells separate healthy donors from COVID-19 patients. We use XGBoost [29], a decision tree based classifier, to rank order proteins for their ability to distinguish CD8+ T cell point cloud data between healthy individuals and COVID-19 patients. The average accuracy of the classifier is about 92%. The classifier returns a feature score for each protein that characterizes its importance relative to other proteins in distinguishing cells from healthy individuals and COVID-19 patients. Intuitively, feature score is an indicator of the importance of a particular feature while classifying the data. By ranking the proteins by their feature scores, we can reduce our further analysis to only a subset of the most important proteins. Our analysis (Fig 4) shows that the top three most important proteins to the XGBoost classifier are proteins T-bet, Eomes, and Ki-67. T-bet induces gene expressions leading to an increase in cytotoxic functions of CD8+ T cells. CD8+ T cells with increased cytotoxic functions are known as ‘effector’ CD8+ T cells and these cells show higher T-bet abundances. Conversely, Eomes induces gene expressions that contribute towards increased life span and re-activation potential of CD8+ T cells to specific antigens [21]. These long-lived T cells are known as ‘memory’ T cells which show increased expressions of Eomes. Memory T cells provide key protection against re-exposure to the same infection. Ki-67 is a marker for actively proliferating cells [22]. Mathew et al. [3] identified Ki-67 as one marker that is upregulated (increased) in some COVID-19 patients. These three proteins are most likely to distinguish CD8+ T cells in healthy donors from those in patients. Further details regarding the application of the classifier are provided in the Materials and Methods section (Section 4).

Persistence diagrams distinguish structural features in CD8+ T cell data occurring in healthy individuals and COVID-19 patients across batch effects and donor-donor variations. We select the proteins T-bet, Eomes, and Ki-67 as relevant markers and compute the

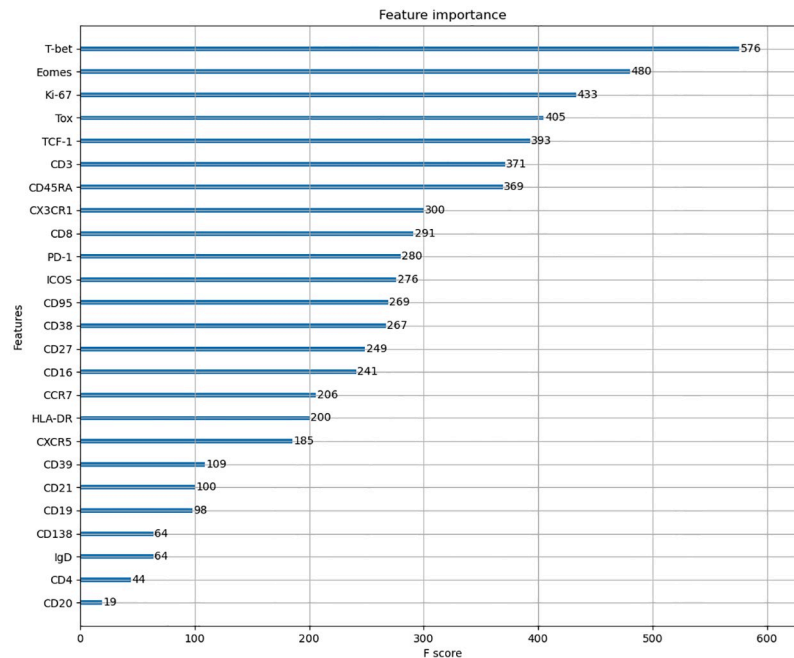


Fig 4. Rank ordering of proteins using a decision tree based classifier. Shows rank ordering of proteins by descending values of feature importance generated by the classifier XBoost.

<https://doi.org/10.1371/journal.pcbi.1009931.g004>

persistence diagrams of the PCD given by them for each individual belonging to groups of healthy donors, COVID-19 patients, and recovered patients. The persistence diagrams vary from individual to individual in each group and between groups which could arise due to batch effects in samples and/or donor-to-donor variations. To determine if there are systematic differences in persistence diagrams for individuals across the three groups (healthy, patient, and recovered), we compute Wasserstein distance between persistence diagrams for 3 categories of pairings: 1) two healthy donors (H×H), 2) one healthy donor and one patient (H×P), and 3) one healthy donor and one recovered individual (H×R). We compute distances for 100 randomly chosen pairs of individuals for each category of pairings. Wasserstein distances of the persistence diagrams for 0-th and 1-st homology groups H_0 and H_1 respectively are higher when comparing H×P pairs than when comparing H×H pairs (Fig 5). A 2-sided KS test showed that the Wasserstein distances for H×P and H×H belong to different probability distribution functions ($p \ll 0.01$); see Fig 5 and also the description of this test in Section 4. This indicates that systematic geometric differences in the flow cytometry PCD with T-bet, Eomes, and Ki-67 between individuals with and without COVID-19 are not attributable to batch effects or donor-to-donor variations alone. Increasing the number of randomly chosen pairs to 200 did not change this conclusion as Fig 5 and S4 Fig illustrate. The difference between H×H and H×R distributions of distances in the T-bet, Eomes, and Ki-67 space are less prominent (S5 Fig). We further test if such systematic differences are present for proteins that are at the bottom of the list in Fig 4 and find that the distributions of Wasserstein distances for corresponding persistence diagrams overlap between the H×H and H×P pairs (S6 Fig). This suggests that systematic differences in the geometry of the PCD occur only for specific sets of proteins. Details regarding computation of persistence diagrams and Wasserstein distances are given in Section 4.

Next, we select a comparison pair that generates a large Wasserstein distance between H_1 -persistence diagrams to further investigate what structural differences exist between the

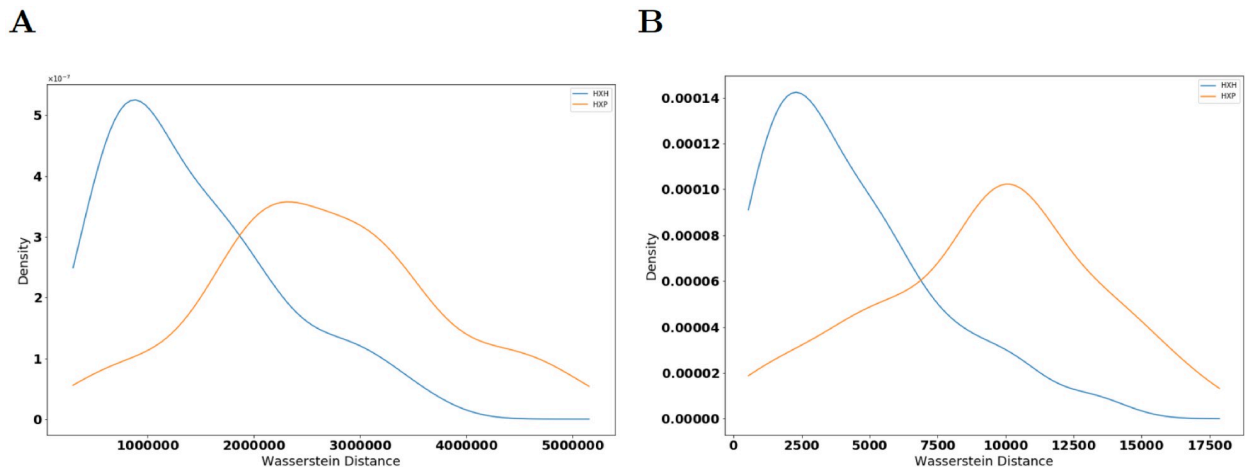


Fig 5. Distributions of Wasserstein distances between persistence diagrams. (A) Shows distributions of Wasserstein distance between H_0 -persistence diagrams for H×H (blue line) and H×P (orange line) pairs ($p = 8.77 \times 10^{-15}$, $QFD = 0.173$). (B) Shows distributions of Wasserstein distance between H_1 -persistence diagrams for H×H (blue line) and H×P (orange line) for the same pairs in (A) ($p = 3.04 \times 10^{-14}$, $QFD = 0.219$). Persistence diagrams are calculated from point clouds in the T-bet, Eomes, and Ki-67 axes. p-values are calculated from a 2-sided KS test.

<https://doi.org/10.1371/journal.pcbi.1009931.g005>

datasets. We choose one pair of a healthy control and patient that generated a Wasserstein distance of 4.0×10^6 units in their H_0 -persistence diagrams and 1.1×10^4 units in H_1 -persistence diagrams. These two individual PCDs and their resulting persistence diagrams are shown in Fig 6.

A readily apparent difference between the resulting persistence diagrams is given by the lower birth times in H_1 of the COVID-19 patient compared to the healthy control (Fig 6E and 6F). This result indicates that the length scale of the data is smaller in the COVID-19 patient, which can be visually confirmed in the scatter plots of the data (Fig 6A and 6B). Specifically, the single cell abundances of T-bet and Eomes in CD8+ T cells are clustered significantly tighter around the origin for the COVID-19 patients than for the healthy controls. Similar manual inspection of other H×P pairs that generate large Wasserstein distances between their persistence diagrams confirms that this trend is not limited to this pair alone.

Additionally, the points in the H_0 -persistence diagram are spread out more widely for the healthy control than the COVID-19 patient (Fig 6C and 6D). A wider distribution of births and deaths in the 0-th homology H_0 implies that there are regions of disparate densities. This suggests that the densities in the protein expressions of T-bet and Eomes are more homogeneous in the PCD in the COVID-19 patient than in the healthy control.

The structural change in the PCD for CD8+ T cells in the T-bet/Eomes plane that occurs during COVID-19 infection implies that T-bet and Eomes expression should be downregulated (decreased) in non-naïve CD8+ T cells. This result is consistent with analysis of clusters of CD8+ T cells by Mathew et al. [3] via a software package FlowSOM [27] that shows that clusters high in T-bet and/or Eomes are downregulated in COVID-19 patients.

The relevance of the above proteins in distinguishing healthy controls from patients is further demonstrated by the statistically significant differences (p-values $\ll 10^{-8}$) in the mean T-bet, Eomes, and Ki-67 abundances in the CD8+ T cells between the groups (S10 Fig). However, the distributions of the mean abundances for the above proteins also showed regions of overlap between healthy and patient populations (S10 Fig) indicating existence of H × P pairs with much smaller differences in the mean values between them than the population averaged mean values of these proteins. Our method specifically identifies differences in topological

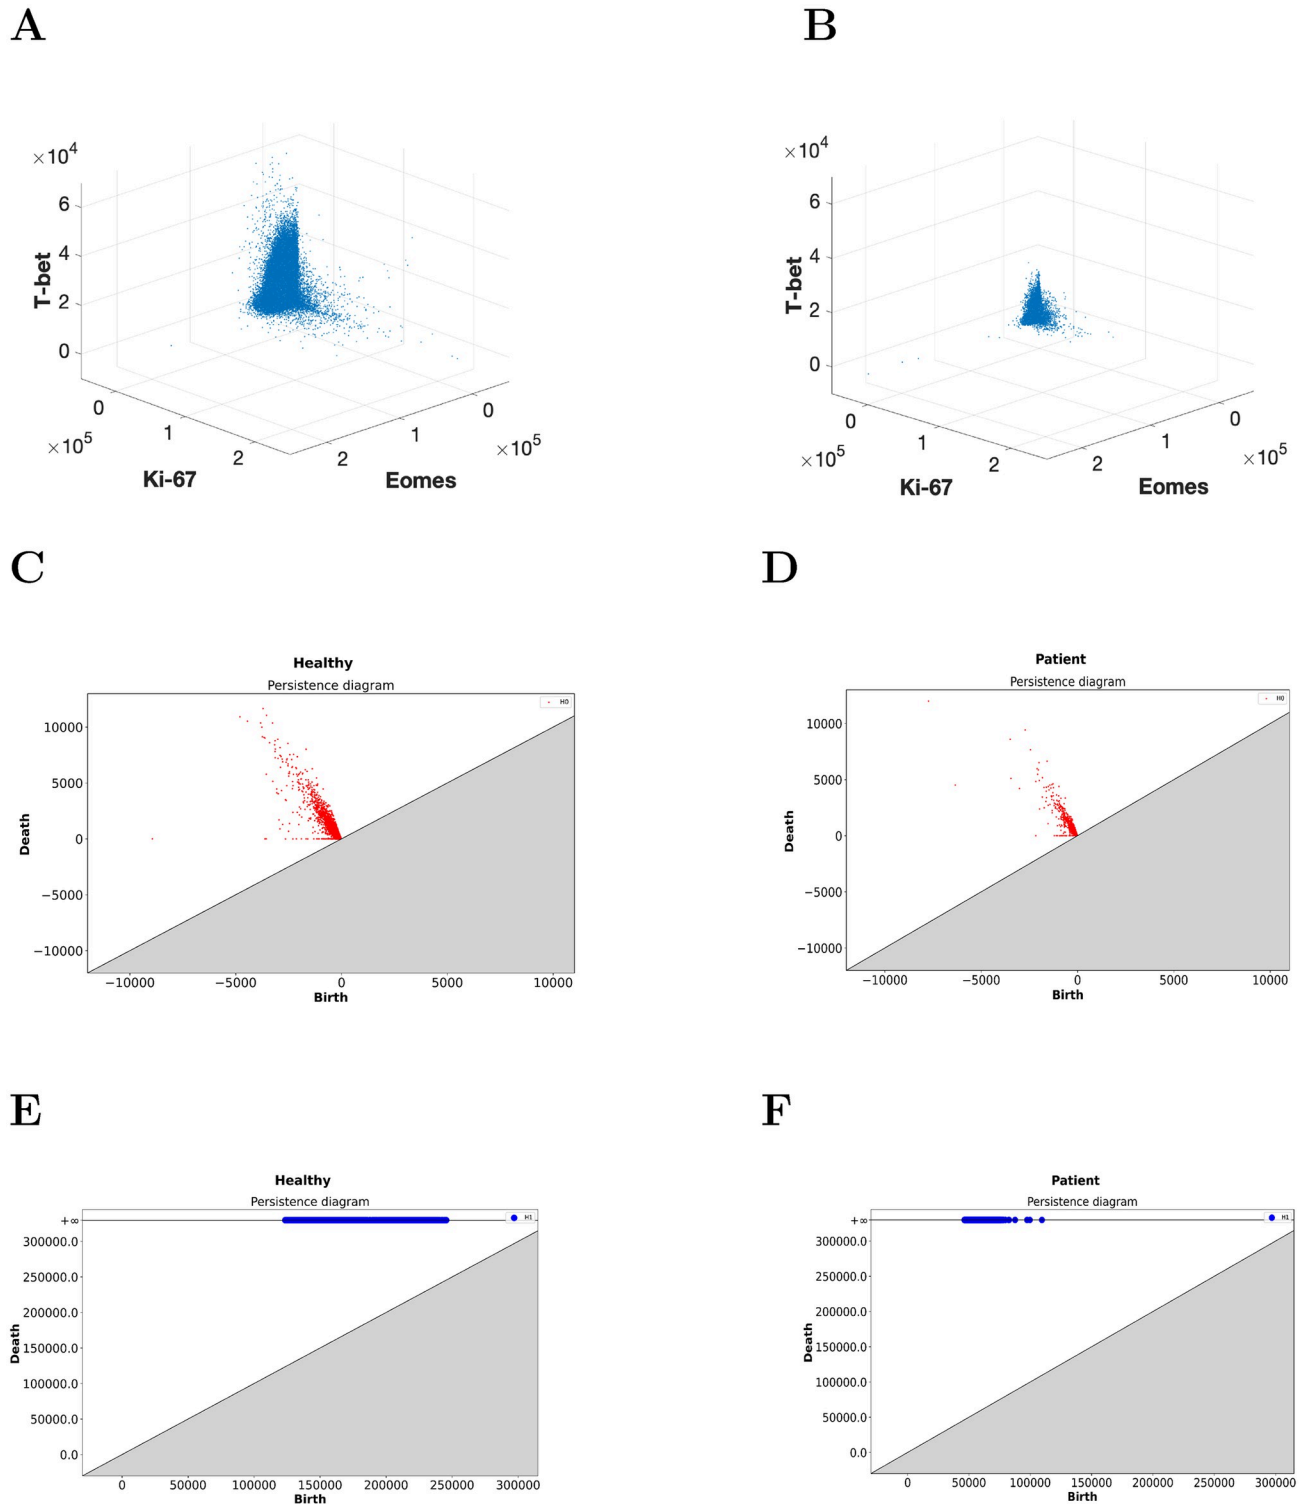


Fig 6. Differences in shape features in the 3D point cloud for CD8+ T cells in an HxP pair. CD8+ T cell point cloud for proteins Eomes, Ki-67, and T-bet for (A) a healthy control and (B) a COVID-19 patient. (C) Shows H_0 -persistence diagram for the healthy control in (A). (D) Shows the H_0 -persistence diagram for the COVID-19 patient in (B). (E) H_1 -persistence diagram for the healthy control in (A). (F) H_1 -persistence diagram for the COVID-19 patient in (B).

<https://doi.org/10.1371/journal.pcbi.1009931.g006>

features in the shape of the PCD between a $H \times P$ pair which can be present despite small differences in the mean values of specific proteins (e.g., T-bet). We further investigated this point by analyzing correlations between the Wasserstein distance between the persistence diagrams of $H \times P$ pairs with the difference in the mean protein abundances (S11(A) and S11(B) Fig) which showed moderate correlations (≤ 0.5). However, there are several instances in which the Wasserstein distance captures differences in the shape of the PCD via persistent homology even when the difference in mean protein abundance (e.g., mean T-bet abundance) is small (S11(C)–S11(F) Fig). This is due to changes in the shape of the point-cloud that are not easily captured by summary statistics such as the mean. Therefore, our analysis shows that healthy and COVID-19 pairs can be better separated by our persistent homology analysis than by summary statistics measures in such cases. The downregulation of T-bet and Eomes in response to viral infections is not well documented, as CD8+ T cells commonly differentiate into phenotypes with high T-bet, high Eomes, or both in response to infections [21, 23].

We next explore the application of our approach to other datasets. We apply our method to the single cell cytometry dataset in Mathew et al. [3] for B cells obtained from healthy donors and COVID-19 patients. The B cells are major orchestrators of the humoral component of adaptive immunity against infections. We compute persistence diagrams for the proteins CXCR5, PD-1, and TCF-1, identified by XGBoost as the three most important proteins for classifying healthy donors and patients. A chemokine receptor, CXCR5, is responsible for B cell trafficking and is found to be downregulated in B cells in COVID-19 patients [3]. Both PD-1, a checkpoint inhibitory receptor [24], and TCF-1, a transcription factor important for T cell differentiation and effector functions [25] are increased in B cells in infected individuals (S12 Fig). Immunosuppressive effects of high PD-1 expression in B cells have been reported earlier [24]. We find that Wasserstein distances of the persistence diagrams for both homology groups (H_0 and H_1) are significantly different (Fig 7) between the healthy donors and the patient population for B cells. This demonstrates that our approach is able to distinguish healthy individuals from patient populations using PCDs of other immune cells. Furthermore, we find that the margin of separation, quantified by the QF-distance (QFD), in these Wasserstein distances is smaller with B cells than the CD8+ T cells. This implies that the structure of PCDs for CD8+ T cells differ more between healthy controls and patients than that for the B cells. These findings may point to previously uncharacterized impact of PD-1 and TCF-1 on B cell function or phenotype in SARS-CoV-2 infection.

Distributions of mean PD-1 expression on B cells in COVID-19 patients is not largely different than that of healthy controls (S12(B) Fig). Therefore, we further analyzed how the difference between mean PD-1 values and the differences in PCDs quantified by the Wasserstein distances are related (S13(A)–S13(D) Fig). We found that unlike T-bet for CD8+ T cells, mean PD-1 expression differences in B cells are not tightly correlated with Wasserstein distances in healthy control-patient pairs (S13(A)–S13(B) Fig). We visualized the PCD in the space of TCF-1, PD-1, and CXCR5 (S13(E) and S13(F) Fig) to gain further insights regarding the differences in the shapes of the PCD in healthy control-patient pairs with similar mean PD-1 expressions. The differences in the shape of the PCDs for these pairs can be largely attributed to the higher expressions of TCF-1 in healthy controls (S13(E) and S13(F) Fig).

Comparison with existing methods. To determine how our selection of filtration compares with an existing method, we compare how our results might change if we use Rips filtration. In Rips filtration, simplices appear when all of their edges appear in the filtration. The edges appear in non-decreasing order of their lengths. We use the entire dataset as the PCD and generate persistence diagrams subsequently, using Rips filtration [8, 14, 26] rather than the filtration we use in our approach. Note that in the standard Rips filtration, all vertices appear at the same instant whereas in our case the vertices are ordered by Eq 1. We then

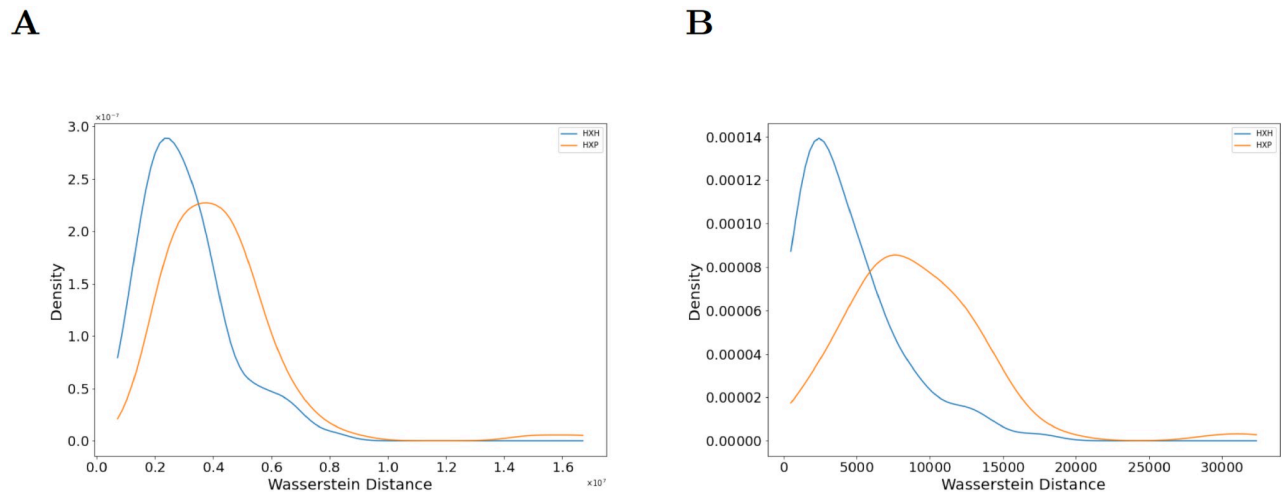


Fig 7. Distributions of Wasserstein distances between persistence diagrams for B cells. (A) Distributions of Wasserstein distance between H_0 -persistence diagrams for H×H (blue line) and H×P (orange line) pairs ($p = 6.28 \times 10^{-5}$, QFD = 0.0300). (B) Distributions of Wasserstein distance between H_1 -persistence diagrams for H×H (blue line) and H×P (orange line) for the same pairs in (A) ($p = 1.07 \times 10^{-12}$, QFD = 0.0946). Persistence diagrams are calculated from point clouds in the CXCR5, PD-1, and TCF-1 axes. p-values are calculated from a 2-sided KS test.

<https://doi.org/10.1371/journal.pcbi.1009931.g007>

compute Wasserstein distances as done previously. We find that Rips filtration is also able to distinguish persistence diagrams of healthy controls and patients, but the margin of separation is much lower, as evidenced by a higher p-value and lower QFD than our choice of filtration offers (S14 Fig). This indicates that our method, which is designed to identify protrusions such as “elbows” in the data, characterizes greater differences in CD8+ T cell protein expression structures than existing methods such as Rips filtration.

We then compare our TDA approach with an existing algorithm FlowSOM [27], widely used for visualizing, clustering, and analyzing PCDs from cytometry experiments. FlowSOM uses a self-organizing map algorithm for generating single cell subsets with unique marker protein expressions. FlowSOM is capable of clustering similar cells together and offers a robust way to determine which cellular subsets are differentially expressed between data sources [28]. We run a FlowSOM analysis and clustering on the CD8+ T cell data and determine that 6 of the 15 clusters are differentially expressed between healthy controls and COVID-19 patients (S15 Fig).

We select one FlowSOM cluster which is differentially expressed (Cluster #3) and one that is not differentially expressed (Cluster #1) between healthy donors and patients for downstream analysis. Visual inspection of these clusters in the 3-dimensional space of T-bet, Eomes, and Ki-67 shows that PCD structure may be different between healthy controls and patients in Cluster #1 (S15(C) Fig). This is because proteins can co-vary in different ways in healthy controls and patients, affecting topological features hidden in the PCD. We then perform our persistent homology analysis to determine if the structure of PCDs for proteins Eomes, Ki-67, and T-bet in subsets of single cells associated with these FlowSOM clusters differ between healthy controls and patients. We find distributions of Wasserstein distances between the persistence diagrams obtained for the above FlowSOM clusters are significantly different (Fig 8 and S16 Fig).

Performing FlowSOM on B cells reveals 12 of the 15 clusters are differentially expressed between healthy controls and COVID-19 patients (S17(A) and S17(B) Fig). The three clusters (Cluster #1, #2, and #14) that are not differentially expressed are all high in PD-1. Visualization

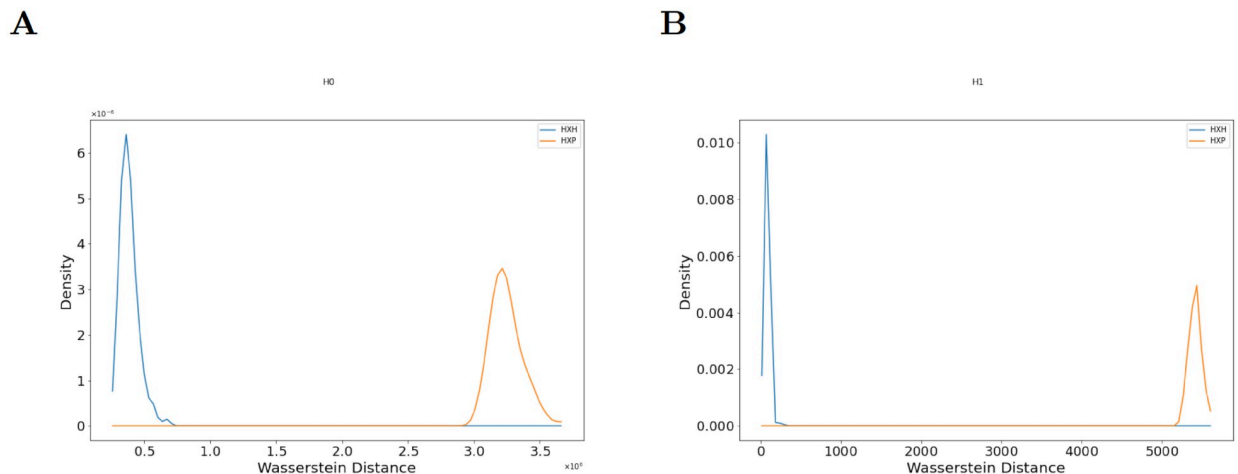


Fig 8. Distributions of Wasserstein distances between persistence diagrams from a FlowSOM cluster (Cluster #1) that is not differentially expressed between healthy controls and COVID-19 patients. (A) Distributions of Wasserstein distance between H_0 -persistence diagrams for H×H (blue line) and H×P (orange line) pairs ($p = 2.21 \times 10^{-59}$, QFD = 1.638) computed for the PCD for Eomes, Ki-67, and T-bet for CD8+ T cells. (B) Distributions of Wasserstein distance between H_1 -persistence diagrams for H×H (blue line) and H×P (orange line) for the same pairs in (A) ($p = 2.21 \times 10^{-59}$, QFD = 1.903).

<https://doi.org/10.1371/journal.pcbi.1009931.g008>

of the PCDs for Cluster #2 and Cluster #4 shows that regions high in TCF-1 [S17\(C\)](#) and [S17\(D\)](#) Fig) distinguish the PCDs for the pairs. Thus, our method is able to identify topological features hidden in the PCD that separate healthy controls from COVID-19 patients in FlowSOM clusters, including clusters which are not differentially expressed between these groups.

3 Discussions and conclusions

We develop a persistent homology based approach to determine topological features hidden in point cloud data representing single cell protein abundances measured in cytometry data. In particular, we characterize the number of connected components or H_0 , and the number of holes or H_1 in our persistence calculations, and show that our approach is able to determine systematic shape differences in the cytometry data for CD8+ T cells obtained from healthy individuals and COVID-19 patients. Therefore, the approach is able to successfully determine systematic shape differences that exist in the presence of batch effect noise and donor-donor variations in the cytometry data. Furthermore, our approach does not use data transformations (e.g., arc-sinh transformation) or any ad-hoc subtype gating to determine these systematic differences, thus we expect persistent homology based approaches will be especially useful in identifying high-dimensional structural trends hidden in cytometry data.

We determine structural changes in T-bet and Eomes abundances in single CD8+ T cells in COVID-19 patients that can be summarized as downregulation. This result is non-intuitive as previous findings show that T-bet and Eomes protein abundances are highest in effector CD8+ T cells, which are induced in response to acute infections, suggesting T-bet and Eomes expressions should be upregulated [21, 23]. The clinical implications of this result are unclear. Mathew et al. [3] describe an immunophenotype in which Eomes+, T-bet+, CD8+ T cells are more abundant in COVID-19 patients who respond poorly to Remdesivir and NSAIDs, have high levels of IL-6, and have fewer eosinophils. Our analysis identifies that this immunophenotype (i.e., Eomes+, T-bet+, CD8+ T cells) is systematically less prevalent in COVID-19 patients than in healthy controls. The ability of our approach to identify shape features in single immune cell PCD without any ‘supervision’ (e.g., specific gating) of the cytometry data shows

that it can potentially determine more complicated immunologically relevant shape features. Furthermore, our approach inferred finer geometric structures from PCDs in B cells for proteins CRCX5, PD-1, and TCF-1 which helped distinguish COVID-19 patients from healthy individuals. These proteins are associated with cell migration, immunosuppression, and effector functions in lymphocytes and can potentially provide further insights into B cell response in COVID-19.

We compare our approach with an existing algorithm FlowSOM which is widely used for analyzing and visualizing multidimensional cytometry data. Our comparison reveals PCDs for subtypes of CD8+ T cells in FlowSOM clusters that do not differentiate healthy controls and COVID-19 patients contain topological features separating the above groups. Therefore, detecting topological features hidden in the PCDs can provide important biological insights regarding response of the lymphocytes in COVID-19.

Our approach integrates cellular comparisons with dataset comparisons. First, the classifier pools all data and determines which proteins are significant in discriminating whether cells come from healthy controls or COVID-19 patients. In this way, the classifier identifies a way to compare cellular phenotypes across experimental groups. Next, the computation of Wasserstein distances for persistence diagrams compares individuals against each other, integrating cellular phenotypes with donor information (e.g., healthy and COVID-19 patients). Thus, this approach allows us to automatically identify individuals that are associated with distinguishing structural features in the point cloud data.

Currently, the limitations are mostly due to the curse of dimensionality that increases the computational complexity. Since we are computing pairwise distances between datapoints to obtain the persistence diagram (S2 Appendix), computation time increases as the dimension of data increases. Computational cluster resources that we use currently complete all computations in about 20 minutes. This is comparable to other data science applications using large datasets, but this approach can be a barrier to those without access or experience with computational clusters. Additionally, it is unclear how additional dimensions impact the statistical properties of the data and interpretability of the results. To expand into many (i.e. 25) dimensions, computational interpretation and validation tools are necessary.

4 Materials and methods

Relevant feature selection by the XGBoost classifier

Let $D = \{c_1, c_2, \dots, c_m\}$ be the collection of m cytometry datasets. Each dataset, c_i , can be viewed as a $\mathcal{M}_{n \times p}$ matrix where n is the number of datapoints (cells) and p is the number of proteins with which each c_i is generated. We denote the collection of cytometry datasets of healthy individuals as $C_{\mathcal{H}} \subset D$ and similarly the set of individuals infected with SARS-CoV-2 as $C_p \subset D$. We proceed to label the data in the following manner: If $c_i \in C_{\mathcal{H}}$ then we assign the label + 1 to each of the n datapoints, similarly we assign -1 if $c_i \in C_p$. Essentially, we now have a binary classification problem where our labeled dataset is $D' = \bigcup D = c_1 \cup c_2 \cup \dots \cup c_j$, with labels defined as above. We solve this binary classification problem with XGBoost [29], a gradient boosted decision tree based classifier, and as a byproduct we get feature scores that correspond directly to each feature's importance in the classification. The higher the score for a protein, the more important it is for the classifier's decision. After our classifier orders the proteins by their scores, we take first r proteins to construct the point-cloud on which persistence diagrams are computed. We set $r = 3$ for all our analysis reported here. We used data from 56 healthy individuals and 108 COVID-19 patients for our feature selection.

The XGBoost classifier was implemented using the open-source python XGBoost package [29]. The model was then trained and validated with K -fold cross-validation, with $K = 10$. The

average accuracy of the classifier was $92.14 \pm 0.04\%$. The protein scores are shown in Fig 4.

Random subsampling of the datapoints

Each PCD can be thought as a set of indexed points. These indices were first shuffled randomly and then 20,000 indices and hence respective points were sampled uniformly from this shuffled set. The samples drawn from each PCD were further analyzed using persistent homology. We discarded datasets that had less than 20,000 data points. Among 55 healthy individuals only 1 had less than 20,000 data points. For the patient data, the number of such datasets was 34.

Details of persistent homology computation

As mentioned before (Section 2.1), computation of persistence diagrams needs a *filtration*. We set the filtration induced by the function $f = \{f_v, f_e\}$ where $f_v(p)$ computes an ‘average’ Euclidean distance between the vertex p and its k neighbors according to Eq (1) and $f_e(e)$ computes the length of the edge e according to Eq (2).

$$f_v(p) = -\frac{1}{k} \sqrt{\sum_i^k \|p - q_i\|^2}, p \in P, \text{ and } q_i \in k\text{-Nearest Neighbors of } p. \quad (1)$$

The term $\|p - q_i\|$ in the above equation is the Euclidean distance between the vertices p and q_i . The function value $f_e(e)$ for an edge $e = (p, q)$ is given by the Euclidean distance between p and q . For the experiments, the number of nearest neighbors is fixed to $k = 40$.

$$f_e(e) = \|p - q\|, \forall p, q \in P \text{ and } p \neq q \quad (2)$$

We begin by sampling every cytometry PCD c_i and take $n (= 20,000)$ samples. We do this to make c_i uniform *w.r.t.* number of data points (single CD8+ T cells). We compute a complete weighted graph $G(V, E)$ with vertices in the sampled data. This complete graph G is a key-step that enables us to compute the *persistence diagram*, $Dgm(c_i)$ of the dataset c_i , *w.r.t.* the filtration function f . We show the algorithm (Algorithm B in S2 Appendix) that executes this step in detail in the supplementary material. Notice that the graph G is weighted in the sense that each vertex $v \in V$ and edge $e \in E$ carries a weight of $f_v(v)$ and $f_e(e)$ respectively. Observe that $f : V \cup E \rightarrow \mathbb{R}$ constitutes a valid filtration of G .

We compute persistence diagrams for each $c_i \in D$ according to Algorithm C in S2 Appendix. The next step involves comparing the persistence diagrams. We do this by computing the Wasserstein distance between persistence diagrams and plotting their distributions. We take two persistence diagrams of randomly selected healthy individuals and compute the Wasserstein distance between them with the help of Gudhi [20, 30] and scikit-learn Python library [31]. Similarly, we compute Wasserstein distance between persistence diagrams of a healthy and an infected individual (both are randomly drawn from the collection). We plot the resulting distances. We do this for 108 pairs to obtain two distributions. Note that, results described in Section 2.1 still holds for 200 pairs (S4 Fig). Intuitively, a large Wasserstein distance between two persistence diagrams implies the datasets on which they were constructed are structurally very different while a small distance implies they are structurally similar.

Statistical testing of difference in Wasserstein distance distributions

2-sided Kolmogorov–Smirnov (KS) tests were performed on Wasserstein distances for pairs of individuals to determine if they arise from the same or different probability distribution functions [32]. MATLAB’s subroutine *kstest2* was used to determine p-values, where the null hypothesis is that the Wasserstein distances from H×H comparisons and the experimental

condition (either $H \times P$ or $H \times R$) arise from the same non-parametric distribution, and the alternative hypothesis is that they come from different distributions. A p -value of ≤ 0.05 (or ≥ 0.05) indicates the support for the alternate hypothesis (i.e., data occur from different distributions) is statistically significant (or not significant).

Computing quadratic form (QF) distance

To measure the dissimilarity between a pair of Wasserstein distance distributions, quadratic form (QF) distance was computed as proposed by Bernas et al. in [33] (originally introduced in [34]). The QF distance was calculated using the formula

$$D^2(\mathbf{h}, \mathbf{f}) = (\mathbf{h} - \mathbf{f})^T \mathbf{A}_j^i (\mathbf{h} - \mathbf{f}) = \sum_{i=1}^n \sum_{j=1}^n a_{ij} (h_i - f_i)(h_j - f_j) \quad (3)$$

where \mathbf{f} and \mathbf{h} are two vectors that list counts corresponding to two histogram bin counts. The quantities \mathbf{f} and \mathbf{h} can be normalised so that $\sum_i f_i = \sum_i h_i = 1$ when indexed by i . In our case, $\mathbf{A}_j^i = [a_{ij}]$ and defined as $a_{ij} = 1 - \sqrt{\frac{(i-j)^2}{d_{max}^2}}$ with d_{max} being maximum distance between bins.

FlowSOM analysis

FlowSOM was performed on the entire non-naïve CD8+ T cell dataset, and separately, the non-naïve B cell dataset. Data was scaled with the transform $\text{asinh}(x/150)$ before analysis. See Code Availability for FlowSOM source code. Comparisons of relative cluster abundances between healthy controls and COVID-19 patients were performed with a Wilcoxon rank sum test. Subsequent persistent homology computation was performed on the selected clusters by sampling 20,000 cells from either the aggregated healthy control data or aggregated COVID-19 patient data. This sampling was repeated to form “synthetic” individual healthy control or COVID-19 patient data. We cannot sample data from each individual, as was done in the prior computations, because many individuals display too few cells in the selected clusters to reliably sample 20,000 cells. In the CD8+ T cell analysis, clusters #1 and #3 were chosen for persistent homology calculations because they contain the most cells, and thus are most likely to have cells in each sample and be unaffected by random sampling.

Supporting information

S1 Appendix. Some omitted details.

(PDF)

S2 Appendix. Algorithms.

(PDF)

S1 Fig. Intuition of persistence. (A) A set of points P sampled from a curve. (B) An Euclidean ball of radius ϵ is grown around each point in P . (C) As ϵ increases the smaller hole gets filled up. (D) The larger hole still ‘persists’ even though the smaller hole gets filled. Figures are adopted from [26, Fig 4.2].

(TIF)

S2 Fig. Scatter plots demonstrating change in structure of cytometry data. Transformed scatter plot for HLA-DR/CD38 axes for CD8+ T cell PCD in a singular (A) healthy donor and (B) COVID-19 patient. This plot demonstrates the ‘elbow’ found by the authors in [3]. The x-axis is $\text{asinh}(\text{HLA-DR}/200)$ and the y-axis is $\text{asinh}(\text{CD38}/500)$.

(TIF)

S3 Fig. Persistence calculations and comparisons for HLA-DR/CD38 axes for CD8+ T cell PCDs shown in Mathew et. al. [3]. (A) Point cloud for individual healthy control in HLA-DR/CD38 expression levels. (B) Complete persistence diagram for the healthy control shown in (A). Boxes indicate zoomed regions for figures (C) and (D); (C) Zoomed in region from (B) of H_0 persistence diagram. Box shows area of low density compared to patient persistence diagram; (D) Zoomed in region from (B) of H_1 persistence diagram; (E)-(H) Same as (A)-(D), but for an individual COVID-19 patient. The box in (G) is more densely populated than the identical box in (C).

(TIF)

S4 Fig. Distributions of Wasserstein distances between persistence diagrams calculated from 200 pairs of individuals. Distributions of Wasserstein distances between (A) H_0 -persistence diagrams ($p = 6.75 \times 10^{-20}$, QFD = 0.190) and (B) H_1 -persistence diagrams ($p = 4.74 \times 10^{-24}$, QFD = 0.220) for CD8+ T cells. Distances between pairs of healthy controls (H \times H) and pairs of a healthy control and a COVID-19 patient (H \times P) are overlaid. Persistence diagrams are calculated from point clouds in the T-bet, Eomes, and Ki-67 axes. This figure plots distributions of 200 randomly selected pairs, while Fig 5 plots distributions of 100 randomly selected pairs.

(TIF)

S5 Fig. Distributions of Wasserstein distances between persistence diagrams for healthy controls and recovered individuals calculated using 3 most important proteins for XGBoost classification of CD8+ T cells from healthy or infected individuals. Distributions of Wasserstein distances between (A) H_0 -persistence diagrams ($p = 0.131$, QFD = 0.005) and (B) H_1 -persistence diagrams ($p = 0.344$, QFD = 0.001) for CD8+ T cells. Distances between pairs of healthy controls (H \times H) and pairs of a healthy control and a individual that recovered from COVID-19 (H \times R) are overlaid. Persistence diagrams are calculated from point clouds in the T-bet, Eomes and Ki-67 axes. p-values are calculated from a 2-sided KS test.

(TIF)

S6 Fig. Distributions of Wasserstein distances between persistence diagrams calculated using least important proteins for XGBoost classification of CD8+ T cells. Distributions of Wasserstein distances between (A) H_0 -persistence diagrams ($p = 2.75 \times 10^{-8}$, QFD = 0.051) and (B) H_1 -persistence diagrams ($p = 0.111$, QFD = 0.022) for CD8+ T cells. Distances between pairs of healthy controls (H \times H) and pairs of a healthy control and a COVID-19 patient (H \times P) are overlaid. Persistence diagrams are calculated from point clouds in the IgD, CD4 and CD20 axes. p-values are calculated from a 2-sided KS test.

(TIF)

S7 Fig. Distributions of Wasserstein distances between persistence diagrams calculated using two most important proteins for XGBoost classification of CD8+ T cells from healthy or infected individuals. Distributions of Wasserstein distances between (A) H_0 -persistence diagrams ($p = 6.31 \times 10^{-19}$, QFD = 0.261) and (B) H_1 -persistence diagrams ($p = 6.31 \times 10^{-19}$, QFD = 0.276) for CD8+ T cells. Distances between pairs of healthy controls (H \times H) and pairs of a healthy control and a COVID-19 patient (H \times P) are overlaid. Persistence diagrams are calculated from point clouds in the T-bet and Eomes axes. p-values are calculated from a 2-sided KS test.

(TIF)

S8 Fig. Distributions of Wasserstein distances between persistence diagrams calculated using four most important proteins for XGBoost classification of CD8+ T cells from

healthy or infected individuals. Distributions of Wasserstein distances between (A) H_0 -persistence diagrams ($p = 3.35 \times 10^{-13}$, QFD = 0.267) and (B) H_1 -persistence diagrams ($p = 3.04 \times 10^{-14}$, QFD = 0.265) for CD8+ T cells. Distances between pairs of healthy controls (H × H) and pairs of a healthy control and a COVID-19 patient (H × P) are overlaid. Persistence diagrams are calculated from point clouds in the T-bet, Eomes, Tox and TCF-1 axes. p-values are calculated from a 2-sided KS test.

(TIF)

S9 Fig. Distributions of Wasserstein distances between persistence diagrams calculated using most important proteins for XGBoost classification of recovered CD8+ T cells. Distributions of Wasserstein distances between (A) H_0 -persistence diagrams ($p = 0.908$, QFD = 0.002) and (B) H_1 -persistence diagrams ($p = 0.994$, QFD = 0.001) for CD8+ T cell. Distances between pairs of healthy controls (H × H) and pairs of a healthy control and an individual that recovered from COVID-19 (H × R) are overlaid. Persistence diagrams are calculated from point clouds in the CD45RA, Eomes and TCF-1 axes. These 3 proteins are the best distinguishing features for the XGBoost classifier to distinguish cells from healthy controls from those from recovered individuals. p-values are calculated from a 2-sided KS test.

(TIF)

S10 Fig. Probability distribution function (pdf) of mean abundances of proteins identified by XGBoost to be important to patient classification in CD8+ T cells. PDFs for (A) T-bet, (B) Eomes, and (C) Ki-67 for CD8+ T cell PCDs in each healthy control and COVID-19 patient shown using violin plots. The thickness of the “violin” denotes the value of the pdf. Data distributions are calculated from the mean protein abundances across all non-naïve CD8+ T cells for each individual. Red dots represent the mean of the data, and red lines represent the standard deviation.

(TIF)

S11 Fig. Demonstration of persistent homology’s ability to capture more information than change in magnitude of single protein measurements in CD8+ T cells. (A-B) Scatter plot showing the relationship between Wasserstein distance between H_0 and H_1 persistence diagrams and the difference in mean T-bet abundance for CD8+ T cells for random pairs of healthy controls and COVID-19 patients. Red circles highlight a pair of individuals that generated a large Wasserstein distance despite a small difference in mean T-bet expressions, which are further analyzed in (C-F). (C-D) Histograms showing the T-bet expression of non-naïve CD8+ T cells for the healthy control (blue) and COVID-19 patient (red) from the points circled above in (A-B). (E-F) Scatter plots showing the T-bet, Eomes, and Ki-67 abundances for the healthy control (blue) and COVID-19 patient (red) from the points circled in (A-B).

(TIF)

S12 Fig. Probability distribution function (pdf) of mean abundances of proteins identified by XGBoost to be important to patient classification in B cells. PDFs for (A) CXCR5, (B) PD-1, and (C) TCF-1 in B cells of each healthy control and COVID-19 patient shown using violin plots. The thickness of the “violin” denotes the value of the pdf. Data distributions are calculated from the mean protein abundances across all B cells for each individual. Red dots represent the mean of the data, and red lines represent the standard deviation.

(TIF)

S13 Fig. Demonstration of persistent homology’s ability to capture more information than change in magnitude of single protein measurements in B cells. (A-B) Scatter plot showing the relationship between Wasserstein distance between H_0 and H_1 persistence diagrams and

the difference in mean PD-1 abundance for B cells for random pairs of healthy controls and COVID-19 patients. Red circles highlight a pair of individuals that generated a large Wasserstein distance despite a small difference in mean PD-1 expressions, which are further analyzed in (C-F). (C-D) Histograms showing the PD-1 expression of non-naïve B cells for the healthy control (blue) and COVID-19 patient (red) from the points circled above in (A-B). (E-F) Scatter plots showing the CXCR5, PD-1, and TCF-1 abundances for the healthy control (blue) and COVID-19 patient (red) from the points circled in (A-B). Protein expression axes in (C)-(F) are scaled to $\text{asinh}(x/150)$, where x is the expression of the given protein.

(TIF)

S14 Fig. Distributions of Wasserstein distances between persistence diagrams calculated using Rips filtration. (A) Shows distributions of Wasserstein distance between H_0 -persistence diagrams for H×H (blue line) and H×P (orange line) pairs ($p = 1.20 \times 10^{-4}$, QFD = 0.0575) for CD8+ T cells. (B) Shows distributions of Wasserstein distance between H_1 -persistence diagrams for H×H (blue line) and H×P (orange line) for the same pairs in (A) ($p = 3.73 \times 10^{-3}$, QFD = 0.0343).

(TIF)

S15 Fig. Results of FlowSOM analysis for CD8+ T cells. (A) t-SNE projection of protein expression data for CD8+ T cell PCD in Mathew et al. Each point represents a cell with 25 protein expressions. Colors represent 15 clusters identified by FlowSOM. Cluster #1 and Cluster #3 are selected for further topological analysis. (B) Heatmap showing scaled MFI for T-bet, Eomes, and Ki-67 for each cluster. Each entry in the first three columns is the MFI scaled by the average MFI of the column. The fourth column shows p-values determining differential expression of the cluster between healthy controls and COVID-19 patients. Note that Cluster #1 has $p > 0.05$ and Cluster #3 has $p < 0.05$. (C-D) Scatter plots showing the T-bet, Eomes, and Ki-67 abundances for all healthy controls (blue) and COVID-19 patients (red) from the cells in Cluster #1 (C) and Cluster #3 (D). Black circles in (C) indicate regions of single cell protein expressions which contribute to the differences in the PCD structure for the FlowSOM clusters. Axes are scaled to $\text{asinh}(x/150)$, where x is the expression of the given protein.

(TIF)

S16 Fig. Distributions of Wasserstein distances between persistence diagrams from a FlowSOM cluster (Cluster #3) that is differentially expressed between healthy controls and COVID-19 patients. (A) Shows distributions of Wasserstein distance between H_0 -persistence diagrams for H×H (blue line) and H×P (orange line) pairs ($p = 2.20 \times 10^{-59}$, QFD = 1.604).

(B) Shows distributions of Wasserstein distance between H_1 -persistence diagrams for H×H (blue line) and H×P (orange line) for the same pairs in (A) ($p = 2.21 \times 10^{-59}$, QFD = 1.573).

(TIF)

S17 Fig. Results of FlowSOM analysis for B cells. (A) t-SNE projection of protein expression data for B cell PCD in Mathew et al. Each point represents a cell with 25 protein expressions. Colors represent 15 clusters identified by FlowSOM. Cluster #2 and Cluster #4 are selected for further topological analysis. (B) Heatmap showing scaled MFI for CXCR5, PD-1, and TCF-1 for each cluster. Each entry in the first three columns is the MFI scaled by the average MFI of the column. The fourth column shows p-values determining differential expression of the cluster between healthy controls and COVID-19 patients. Note that Cluster #2 has $p > 0.05$ and Cluster #4 has $p < 0.05$. (C-D) Scatter plots showing the CXCR5, PD-1, and TCF-1 abundances for all healthy controls (blue) and COVID-19 patients (red) from the cells in Cluster #2 (C) and Cluster #4 (D). Axes are scaled to $\text{asinh}(x/150)$, where x is the expression of the given

protein.
(TIF)

Acknowledgments

The authors would like to thank John Wherry and members of his lab for depositing their data and helping us to understand it.

Author Contributions

Conceptualization: Tamal K. Dey, Jayajit Das.

Data curation: Darren Wethington.

Formal analysis: Soham Mukherjee, Darren Wethington.

Funding acquisition: Tamal K. Dey, Jayajit Das.

Methodology: Soham Mukherjee, Darren Wethington, Tamal K. Dey, Jayajit Das.

Project administration: Tamal K. Dey, Jayajit Das.

Resources: Tamal K. Dey, Jayajit Das.

Software: Soham Mukherjee, Darren Wethington.

Supervision: Tamal K. Dey, Jayajit Das.

Visualization: Soham Mukherjee, Darren Wethington.

Writing – original draft: Soham Mukherjee, Darren Wethington, Tamal K. Dey, Jayajit Das.

Writing – review & editing: Soham Mukherjee, Darren Wethington, Tamal K. Dey, Jayajit Das.

References

1. Spitzer M, Nolan G. Mass Cytometry: Single Cells, Many Features. *Cell*. 2016; 165(4):780–791. <https://doi.org/10.1016/j.cell.2016.04.019> PMID: 27153492
2. Simoni Y, Chng MHY, Li S, Fehlings M, Newell EW. Mass cytometry: a powerful tool for dissecting the immune landscape. *Current Opinion in Immunology*. 2018; 51:187–196. <https://doi.org/10.1016/j.coi.2018.03.023> PMID: 29655022
3. Mathew D, Giles JR, Baxter AE, Oldridge DA, Greenplate AR, Wu JE, et al. Deep immune profiling of COVID-19 patients reveals distinct immunotypes with therapeutic implications. *Science*. 2020; 369(6508). <https://doi.org/10.1126/science.abc8511> PMID: 32669297
4. Strauss-Albee DM, Fukuyama J, Liang EC, Yao Y, Jarrell JA, Drake AL, et al. Human NK cell repertoire diversity reflects immune experience and correlates with viral susceptibility. *Science Translational Medicine*. 2015; 7(297):297ra115–297ra115. <https://doi.org/10.1126/scitranslmed.aac5722> PMID: 26203083
5. Wargo JA, Reddy SM, Reuben A, Sharma P. Monitoring immune responses in the tumor microenvironment. *Current Opinion in Immunology*. 2016; 41:23–31. <https://doi.org/10.1016/j.coi.2016.05.006> PMID: 27240055
6. Azad A, Rajwa B, Pothan A. Immunophenotype Discovery, Hierarchical Organization, and Template-Based Classification of Flow Cytometry Samples. *Frontiers in Oncology*. 2016; 6:188. <https://doi.org/10.3389/fonc.2016.00188> PMID: 27630823
7. del Barrio E, Inouzhe H, Loubes JM, Matrán C, Mayo-Íscar A. optimalFlow: optimal transport approach to flow cytometry gating and population matching. *BMC Bioinformatics*. 2020; 21(1):479. <https://doi.org/10.1186/s12859-020-03795-w> PMID: 33109072
8. Edelsbrunner H, Harer J. *Computational topology: an introduction*. American Mathematical Soc.; 2010.
9. Zomorodian A. *Topological data analysis. Advances in applied and computational topology*. 2012; 70:1–39. <https://doi.org/10.1090/psapm/070/587>

10. McGuirol MR, Volkening A, Sandstede B. Topological data analysis of zebrafish patterns. *Proceedings of the National Academy of Sciences*. 2020; 117(10):5113–5124. <https://doi.org/10.1073/pnas.1917763117> PMID: 32098851
11. Komendantov AO, Venkadesh S, Rees CL, Wheeler DW, Hamilton DJ, Ascoli GA. Quantitative firing pattern phenotyping of hippocampal neuron types. *Scientific Reports*. 2019; 9(1):1–17. <https://doi.org/10.1038/s41598-019-52611-w> PMID: 31784578
12. Lakshmikanth T, Olin A, Chen Y, Mikes J, Fredlund E, Remberger M, et al. Mass cytometry and topological data analysis reveal immune parameters associated with complications after allogeneic stem cell transplantation. *Cell reports*. 2017; 20(9):2238–2250. <https://doi.org/10.1016/j.celrep.2017.08.021> PMID: 28854371
13. Rizvi AH, Camara PG, Kandror EK, Roberts TJ, Schieren I, Maniatis T, et al. Single-cell topological RNA-seq analysis reveals insights into cellular differentiation and development. *Nature Biotechnology*. 2017; 35(6):551–560. <https://doi.org/10.1038/nbt.3854> PMID: 28459448
14. Buchet M, Chazal F, Oudot SY, Sheehy DR. Efficient and robust persistent homology for measures. *Computational Geometry*. 2016; 58:70–96. <https://doi.org/10.1016/j.comgeo.2016.07.001>
15. Cohen-Steiner D, Edelsbrunner H, Harer J. Stability of Persistence Diagrams. In: *Proceedings of the Twenty-First Annual Symposium on Computational Geometry*. SCG'05. New York, NY, USA: Association for Computing Machinery; 2005. p. 263–271. Available from: <https://doi.org/10.1145/1064092.1064133>.
16. Carlsson G, Zomorodian A, Collins A, Guibas LJ. Persistence barcodes for shapes. *International Journal of Shape Modeling*. 2005; 11(02):149–187. <https://doi.org/10.1142/S0218654305000761>
17. Chazal F, Cohen-Steiner D, Mérigot Q. Geometric Inference for Probability Measures. *Foundations of Computational Mathematics*. 2011; 11(6):733–751. <https://doi.org/10.1007/s12028-011-9098-0>
18. Buchet M, Dey TK, Wang J, Wang Y. Decluster and resample: Towards parameter free denoising. In: *33rd International Symposium on Computational Geometry*, SoCG 2017. Schloss Dagstuhl, Leibniz-Zentrum für Informatik GmbH; 2017. p. 231–2316.
19. Bellman R. Dynamic Programming. *Science*. 1966; 153(3731):34–37. <https://doi.org/10.1126/science.153.3731.34> PMID: 17730601
22. Kerber M, Morozov D, Nigmatov A. Geometry Helps to Compare Persistence Diagrams. *ACM J Exp Algorithmics*. 2017; 22. <https://doi.org/10.1145/3064175>
21. Takemoto N, Intlekofer AM, Northrup JT, Wherry EJ, Reiner SL. Cutting Edge: IL-12 inversely regulates T-bet and eomesodermin expression during pathogen-induced CD8+ T cell differentiation. *The Journal of Immunology*. 2006; 177(11):7515–7519. <https://doi.org/10.4049/jimmunol.177.11.7515> PMID: 17114419
22. Scholzen T, Gerdes J. The Ki-67 protein: From the known and the unknown. *Journal of Cellular Physiology*. 2000; 182(3):311–322. [https://doi.org/10.1002/\(SICI\)1097-4652\(200003\)182:3%3C311::AID-JCP1%3E3.0.CO;2-9](https://doi.org/10.1002/(SICI)1097-4652(200003)182:3%3C311::AID-JCP1%3E3.0.CO;2-9) PMID: 10653597
23. Knox JJ, Cosma GL, Betts MR, McLane LM. Characterization of T-Bet and Eomes in Peripheral Human Immune Cells. *Frontiers in Immunology*. 2014; 5:217. <https://doi.org/10.3389/fimmu.2014.00217> PMID: 24860576
24. Thibault ML, Mamessier E, Gertner-Dardenne J, Pastor S, Just-Landi S, Xerri L, et al. PD-1 is a novel regulator of human B-cell activation. *International immunology*. 2013; 25(2):129–137. <https://doi.org/10.1093/intimm/dxs098> PMID: 23087177
25. Wang Y, Hu J, Li Y, Xiao M, Wang H, Tian Q, et al. The Transcription Factor TCF1 Preserves the Effector Function of Exhausted CD8 T Cells During Chronic Viral Infection. *Frontiers in Immunology*. 2019; 10:169. <https://doi.org/10.3389/fimmu.2019.00169> PMID: 30814995
26. Dey TK, Wang Y. *Computational Topology for Data Analysis*. Cambridge University Press; 2022. Available from: <https://books.google.com/books?id=PWtYAAAQBAJ>.
27. Van Gassen S, Callebaut B, Van Helden MJ, Lambrecht BN, Demeester P, Dhaene T, et al. FlowSOM: Using self-organizing maps for visualization and interpretation of cytometry data. *Cytometry Part A*. 2015; 87(7):636–645. <https://doi.org/10.1002/cyto.a.22625> PMID: 25573116
28. Quintelier K, Couckuyt A, Emmaneel A, Aerts J, Saey Y, Van Gassen S. Analyzing high-dimensional cytometry data using FlowSOM. *Nature Protocols*. 2021; p. 1–27. PMID: 34172973
29. Chen T, Guestrin C. XGBoost: A Scalable Tree Boosting System. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD'16. New York, NY, USA: Association for Computing Machinery; 2016. p. 785–794. Available from: <https://doi.org/10.1145/2939672.2939785>.
30. The GUDHI Project. *GUDHI User and Reference Manual*. 3.4.1 ed. GUDHI Editorial Board; 2021. Available from: <https://gudhi.inria.fr/doc/3.4.1/>.

31. Buitinck L, Louppe G, Blondel M, Pedregosa F, Mueller A, Grisel O, et al. API design for machine learning software: experiences from the scikit-learn project. In: ECML PKDD Workshop: Languages for Data Mining and Machine Learning; 2013. p. 108–122.
32. Panchenko D. Statistics for Applications: 18.650; 2006. Available from: <https://ocw.mit.edu>.
33. Bernas T, Asem EK, Robinson JP, Rajwa B. Quadratic form: a robust metric for quantitative comparison of flow cytometric histograms. *Cytometry Part A: the journal of the International Society for Analytical Cytology*. 2008; 73(8):715–726. <https://doi.org/10.1002/cyto.a.20586> PMID: 18561196
34. Hafner J, Sawhney HS, Equitz W, Flickner M, Niblack W. Efficient color histogram indexing for quadratic form distance functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 1995; 17(7):729–736. <https://doi.org/10.1109/34.391417>