MIAAIM: Multi-omics image integration and tissue state mapping using topological data analysis and cobordism learning

Joshua M. Hess^{1,2}, Iulian Ilieș³, Denis Schapiro^{4,5,6,7}, John J. Iskra^{1,8}, Walid M. Abdelmoula^{9,10}, Michael S. Regan⁹, Georgios Theocharidis¹¹, Chin Lee Wu¹², Aristidis Veves¹¹, Nathalie Y.R. Agar^{9,13}, Ann E. Sluder¹, Mark C. Poznansky^{1*+}, Patrick M. Reeves^{1*+}, Ruxandra F. Sîrbulescu^{1*+}

¹Vaccine and Immunotherapy Center, Massachusetts General Hospital, Harvard Medical School, Boston, MA, USA.

²Tri-Institutional Training Program in Computational Biology and Medicine, Weill Cornell Medicine, New York, NY, USA

³Healthcare Systems Engineering Institute, Northeastern University, Boston, MA, USA.

⁴Laboratory of Systems Pharmacology, Harvard Medical School, Boston, MA, USA.

⁵Broad Institute of Harvard Medical School and Massachusetts Institute of Technology, Boston, MA, USA.

⁶Institute for Computational Biomedicine, Faculty of Medicine, Heidelberg University Hospital and Heidelberg University, Heidelberg, Germany

⁷Institute of Pathology, Heidelberg University Hospital, Heidelberg, Germany

⁸Department of Mathematics, Emory & Henry College, Emory, VA

⁹Department of Neurosurgery, Brigham and Women's Hospital, Harvard Medical School, Boston, MA, USA.

¹⁰Current address: Invicro, LLC, Boston, MA, USA

¹¹The Rongxiang Xu Center for Regenerative Therapeutics, Beth Israel Deaconess Medical Center, Harvard Medical School, Boston, MA, USA.

¹²Department of Pathology, Massachusetts General Hospital, Harvard Medical School, Boston, MA, USA.

¹³Department of Radiology, Brigham and Women's Hospital, Harvard Medical School, Boston, MA, USA.

*RFS, PMR and MCP contributed equally to this manuscript

+Corresponding authors

31 Corresponding authors:
 32 Ruxandra F. Sîrbulescu, Vaccine and Immunotherapy Center, 149 13th St, Charlestown, MA 02129;
 33 Telephone: 617-953-2316; email: rsirbulescu@mgh.harvard.edu
 34 Patrick M. Reeves, Vaccine and Immunotherapy Center, 149 13th St, Charlestown, MA 02129;
 35 Telephone: 617-726-3812; email: patrick_reeves@hms.harvard.edu
 36 Mark C. Poznansky, Vaccine and Immunotherapy Center, 149 13th St, Charlestown, MA 02129;
 37 Telephone: 617-724-6375; email: mpoznansky@mgh.harvard.edu

ABSTRACT

High-parameter tissue imaging enables detailed molecular analysis of single cells in their spatial environment. However, the comprehensive characterization and mapping of tissue states through multimodal imaging across different physiological and pathological conditions requires data integration across multiple imaging systems. Here, we introduce MIAAIM (Multi-omics Image Alignment and Analysis by Information Manifolds) a modular, reproducible computational framework for aligning data across bioimaging technologies, modeling continuities in tissue states, and translating multimodal measures across tissue types. We demonstrate MIAAIM's workflows across diverse imaging platforms, including histological stains, imaging mass cytometry, and mass spectrometry imaging, to link cellular phenotypic states with molecular microenvironments in clinical biopsies from multiple tissue types with high cellular complexity. MIAAIM provides a robust foundation for the development of computational methods to integrate multimodal, high-parameter tissue imaging data and enable downstream computational and statistical interrogation of tissue states.

INTRODUCTION

Increased availability of multiplexed imaging at cellular and subcellular resolution allows multiparameter molecular descriptions of organelle, cell, and tissue organization in a preserved spatial context¹⁻⁶. A variety of high-parameter imaging technologies have been developed, each capturing individual aspects of multi-faceted tissue representations – ranging from untargeted methods such as mass spectrometry imaging⁷ (MSI) and spatial transcriptomics^{8, 9}, which provide comprehensive molecular profiles, to targeted, multiplexed antibody-based labeling approaches¹⁻⁶. Compared to advances in data acquisition^{10, 11} and atlas mapping for single-cell multi-omics data in dissociated tissue^{12, 13}, analogous computational methods that preserve the spatial structure of tissues remain underdeveloped. Tissue cytoarchitecture and molecular composition are crucial for the comprehensive characterization of biological responses under different physiological conditions or disease states. Most importantly, multi-omics tissue portraits, including single-cell composition and spatial structure, could be used to infer equilibrium and transition states of tissue health and disease.

The integration of data across multiple imaging technologies can present significant challenges, including data set size and high dimensionality. In recent years, methods have been developed to align multimodal imaging data as it becomes increasingly relevant to analyze data acquired across multiple modalities^{14, 15}. To date, no approach to transform data from separate imaging modalities into the same

coordinate system that generalizes to the growing repertoire of high-parameter imaging systems has been developed. Dimensionality reduction is routinely used to model single-cell data, showing that generating functions of cellular phenotype and gene expression closely follow a nonlinear manifold structure. While manifold-based representations of single-cell systems enable cell state trajectory inference¹⁶⁻¹⁸, tissue states emerging from cellular behavior are often discretized. Clustering is typically used to partition cells into distinct populations^{19, 20} (i.e., cell types), and tissue states are defined based on overall abundances of cell types. We suggest that nonlinear mapping of a systems-level state space may better enable descriptions of continuous transitions from health to disease. For tissue atlases, such as the Human Tumor Atlas Network²¹ and the HuBMAP Consortium²², the transfer of information across tissue states may be used for downstream analysis, as is routinely done in dissociated tissue mapping. Comparable intact tissue mapping calls for a reproducible data integration framework that can be implemented at multiple institutions across computing architectures.

Here, we describe MIAAIM (Multi-omics Image Alignment and Analysis by Information Manifolds), a modular framework consisting of 4 new software packages (see **Methods, Key Resources Table**) with reproducible workflows to align, model, and transfer data across high-parameter imaging technologies. MIAAIM contains a manifold-based image registration approach, designed to compress and integrate data with diverse densities and spatial resolutions, that can be generalized to multiple high-parameter bioimaging systems. In addition, we introduce *cobordism learning* as a dimension reduction and domain transfer method that models phenotypic manifolds (i.e., tissue states) as the boundary of a higher-dimensional manifold, called a cobordism²³, that characterizes continuous transitions between tissue states. We demonstrate the applicability of MIAAIM using multimodal imaging data derived from a combination of mass spectrometry-based imaging technologies, including imaging mass cytometry (IMC) and mass spectrometry imaging (MSI), in addition to classical histology. Using this method, we investigate diverse cellular and molecular landscapes in tissue samples with high cellular complexity, including human tonsillar, diabetic foot ulcer (DFU) and prostate cancer biopsies.

RESULTS

MIAAIM workflow

MIAAIM is a sequential workflow aimed at providing comprehensive portraits of tissue states. It includes 4 processing stages: (i) image preprocessing and compression with the high-dimensional image preparation (HDIprep) workflow, (ii) image registration with the high-dimensional image registration

(HDIreg) workflow, (iii) tissue state modeling with cobordism approximation and projection (PatchMAP), and (iv) cross-modality information transfer with i-PatchMAP (**Fig. 1**). Image integration in MIAAIM begins with two or more assembled images (level 2 data)^{24, 25} or spatially resolved data sets (**Fig. 1**, assembled images). The size and standardized format of assembled images vary by technology. For example, cyclic fluorescence-based methods (e.g., CODEX³, CyCIF^{5, 6}) assemble BioFormats/OME-compatible²⁶ 20-60-plex full-tissue images after correcting uneven illumination (e.g., BaSiC²⁷) and stitching tiles (e.g., ASHLAR²⁸); other methods acquire 20 to 100-plex data within specified regions of interest (ROIs) (e.g. MIBI¹, IMC²) in the OME format. Additional methods quantify thousands of parameters at rasterized locations on full tissues or ROIs and are not stored in BioFormats/OME-compatible formats (e.g., the imzML format²⁹ often stores MSI data). MIAAIM accepts a variety of input data formats (**Fig. 1**), and its modular design will enable it to incorporate additional formats tailored to new imaging modalities.

Regardless of technology, assembled images contain high numbers of heterogeneously distributed parameters, which precludes comprehensive, manually-guided image alignment. In addition, high-dimensional imaging produces large feature spaces that challenge methods used in unsupervised settings. The HDIprep workflow in MIAAIM generates compressed images that preserve multiplex salient features to enable cross-technology statistical comparison while minimizing computational complexity (**Fig. 1**, HDIprep). For images acquired from histological staining, HDIprep provides parallelized smoothing and morphological operations that can be applied sequentially for preprocessing. Image registration with HDIreg produces transformations to combine modalities within the same spatial domain (**Fig. 1**, HDIreg). HDIreg uses Elastix³⁰, a parallelized image registration library to calculate transformations, and is optimized to transform large multichannel images with minimal memory use, while also supporting histological stains. HDIreg automates image resizing, padding, and trimming of borders prior to applying image transformations.

Aligned data are well-suited for established single-cell and spatial neighborhood analyses³¹ and they can be segmented to capture multimodal single-cell measures (level 3 and 4 data)^{24, 25}, such as average protein expression or spatial features of cells, or analyzed at the pixel level. A common goal in pathology is to use composite tissue portraits to map healthy-to-diseased transitions. To address this, similarities between systems-level tissue states can be visualized with the PatchMAP workflow (**Fig. 1**, PatchMAP). PatchMAP models tissue states as phenotypic manifolds that are stitched to form a higher-dimensional cobordism. The result is a nested model capturing nonlinear intra-system states and cross-system

continuities. The cobordism paradigm can be applied as a tissue-based atlas-mapping tool to transfer information across modalities with i-PatchMAP (**Fig. 1**, i-PatchMAP).

MIAAIM’s algorithms are nonparametric, using probability distributions supported by manifolds rather than training data models. MIAAIM is therefore technology-agnostic and generalizes to multiple imaging systems (**Supplementary Table 1**). Nonparametric image registration, however, is an iterative, parameter-tuning process rather than a “black-box” solution. This creates a substantial challenge for reproducible data integration across institutions and computing architectures. We therefore Dockerized containerized MIAAIM’s data integration workflows and developed a Nextflow³³ implementation to document human-in-the-loop processing and remove language-specific dependencies in accordance with the FAIR (findable, accessible, interoperable, and reusable) data stewardship principles³⁴. In the following sections, we describe algorithms unique to MIAAIM and their application to multimodal imaging data sets across technologies and diverse tissue specimens.

High-Dimensional Image Compression with HDIprep

To compress high-parameter images, HDIprep performs dimensionality reduction on pixels using Uniform Manifold Approximation and Projection (UMAP)³⁵ (**Fig. 2a**). We performed a stringent comparison of dimensionality reduction algorithms using imaging data sets of human DFU, prostate cancer, and tonsil tissue biopsies acquired using MSI, IMC, and hematoxylin and eosin (H&E). Based on dimensionality reduction benchmarks, UMAP consistently outperformed competing algorithms in its robustness to noise and ability to efficiently preserve data complexity while capturing morphological structure (**Supplementary Figs. 1, 2, 3**).

HDIprep retains data complexity with the fewest degrees of freedom necessary by detecting *steady-state manifold embeddings*. To identify steady-state dimensionalities, information captured by UMAP pixel embeddings is computed (**Methods, Definition 1**, cross-entropy) across a range of embedding dimensionalities, and the first dimension where the observed cross-entropy approaches the asymptote of an exponential regression fit is selected. Steady state embedding calculations scale quadratically with the number of pixels, HDIprep therefore embeds spectral landmarks in the pixel manifold representative of its global structure (**Supplementary Fig. 4**).

Pixel-level dimensionality reduction is computationally expensive for large images, i.e., at high resolution (e.g., $\leq 1 \mu\text{m}/\text{pixel}$). To reduce compression time while preserving quality, we developed a subsampling approach to embed a spatially representative subset of pixels prior to spectral landmark

selection and project out-of-sample pixels into embeddings (**Supplementary Figs. 5, 6, 7**). HDIprep combines all optimizations with a neural-network UMAP implementation³⁶ to scale to whole-tissue images. We demonstrate its efficacy on publicly available³⁷ 44-channel CyCIF images containing ~100 and ~256 million pixels (**Supplementary Fig. 8**). Thus, HDIprep presents an objective, pixel-level compression method applicable to multiple modalities (**Methods, Algorithm 1**).

High-Dimensional Image Registration (HDIreg)

MIAAIM connects the HDIprep and HDIreg workflows with a manifold alignment scheme parametrized by spatial transformations. To align pixel manifolds, we extended methodology to compute manifold α -entropy^{38,39} using entropic graphs to UMAP embeddings and adapted it to image registration using the entropic graph-based Rényi α -mutual information (α -MI)⁴⁰ (**Methods, HDIreg**). HDIreg selects for a transformation that maximizes image-to-image (manifold-to-manifold) α -MI (**Fig. 2b**). The entropic-graph based α -MI generalizes to arbitrary dimensionalities by considering distributions of k-nearest neighbor (KNN) graph lengths of compressed (i.e., embedded) pixels, rather than directly comparing the pixels themselves. Combining HDIprep compression with KNN α -MI extends intensity-based registration to images across technologies, without the need to match specific markers or analytes.

MIAAIM generates information on cellular phenotype, molecular ion distribution, and tissue state across scales

To demonstrate the capacity of MIAAIM to generate linked data between multimodal imaging with various scales and data densities, we applied the HDIprep and HDIreg workflows to MALDI-TOF MSI, H&E and IMC data from a DFU tissue biopsy containing a spectrum of tissue states, from the necrotic center of the ulcer to the healthy margin. Image acquisition covered 1.2 cm² for H&E and MSI data. Molecular imaging with MSI enabled untargeted mapping of lipids and small metabolites in the 400-1000 m/z range across the specimen at a resolution of 50 $\mu\text{m}/\text{pixel}$. Tissue morphology was captured with H&E at 0.2 $\mu\text{m}/\text{pixel}$, and 27-plex IMC data was acquired at 1 $\mu\text{m}/\text{pixel}$ resolution from 7 ROIs on an adjacent section.

Cross-modality alignment was performed in a global-to-local fashion (**Fig. 2c**). We utilized HDIprep compression for high-parameter data and HDIreg manifold alignment for registration of compressed images. Due to the destructive nature of IMC acquisition², we aligned full-tissue data from MSI, H&E (down-sampled to approximately 3.5 $\mu\text{m}/\text{pixel}$), and an IMC reference image first.

Deformations not captured at the full-tissue scale within each ROI were corrected using manual landmark guidance. Tissue deformations resulting from serial sectioning were accounted for with nonlinear transformations. Registrations were initialized by affine transformations for coarse alignment before nonlinear correction. Resolution differences were accounted for with a multiresolution smoothing scheme³⁰. Final alignment proceeded by composing modality and ROI-specific transformations.

After locating single-cells using image segmentation with Ilastik⁴¹, single-cell parameter quantification with MCMICRO²⁴, and antibody staining quality control, registered images yielded the following information for 7,114 cells: (i) average expression of 14 proteins (**Methods, Key Resources Table**) including markers for lymphocytes, macrophages, fibroblasts, keratinocytes, and endothelial cells, as well as extracellular matrix proteins, such as collagen and smooth muscle actin; (ii) morphological features, such as cell eccentricity, solidity, extent, and area, spatial positioning of each cell centroid; and (iii) the distribution of 9,753 m/z MSI peaks across the full tissue. Distances from each MSI pixel and IMC ROI to the center of the ulcer, identified by manual inspection of H&E, were also quantified. Through the integration of these modalities, MIAAIM provided cross-modal information that could not be gathered with a single imaging system, such as profiling of single-cell protein expression and microenvironmental molecular abundance.

Identification of molecular microenvironmental niches correlated with cell and disease states

We verified the existence of cross-modal associations from the integrated multimodal imaging data set by conducting a *microenvironmental correlation network analysis* (MCNA) on registered IMC and MSI data (**Fig. 3**). We performed community detection (i.e., clustering) on MSI analytes (m/z peaks) based on their correlations to single-cell protein measures and defined *microenvironmental correlation network modules* (MCNMs; different colors in **Fig. 3a**). Inspection of MCNMs with top correlations to protein levels identified with IMC revealed that sets of molecules, rather than individual peaks, were associated with cellular protein expression (**Fig. 3b**). MCNMs organized on an axis separating those with moderate positive correlations to cell markers indicative of inflammation and cell death (CD68, activated Caspase-3) and those with moderate positive correlations to markers of immune regulation (CD163, CD4, FoxP3) and vasculature (CD31). Some proteins, such as the myeloid marker CD14 and the cell proliferation marker Ki-67, were not strongly correlated with any of the detected m/z peaks across all cells.

To gain insights into the association of molecular distributions with tissue health, we plotted ion intensity distributions of MCNMs with respect to their proximity to the center of the ulcer (**Fig. 3c**). This

analysis revealed a change in molecular profiles about 6 mm from the ulcer center point, as the tissue state progressed from healthy to injured. We validated our observations and the performance of HDIreg to align micron-scale structures by visualizing the distribution of top correlated ions within cellular microenvironments (**Figs. 3d, 3e**).

An advantage of the MIAAIM approach is the capacity to identify molecular variations in one modality (here, MSI) that correlate with cell states identified using a different modality (here, IMC). We investigated whether m/z peaks differentially associate with cell proliferation (Ki-67 marker in IMC). We identified cell phenotypes through unsupervised clustering on IMC segmented cell-level expression patterns (**Fig. 3f**) and conducted a differential correlation analysis between phenotypes within the well-separated CD3+ cluster, which likely identifies infiltrating T cells at the wound site, and the CD3- cell populations (**Fig. 3g**). Interestingly, we found that correlations to Ki-67 expression shifted with near significance (2σ) for multiple m/z peaks (Fisher transformed, one-sided z-statistics; Bonferroni corrected P-values) between CD3- populations and the CD3+ population (**Fig. 3h**).

We then utilized spatial context preserved with MIAAIM and observed that ion intensities of m/z peaks positively correlated to Ki-67 in CD3+ cells increased with distance from the wound, while molecules with Ki-67 negative correlations specific to CD3+ cells showed the opposite trend (**Fig. 3i**). This distribution of Ki-67 correlates suggests that proliferation of CD3+ T cells occurs predominantly near the healthy margin of the DFU and confirms that molecular correlates of T-cell proliferation can be identified through this unbiased analysis. Collectively, these results provide insights to molecular microenvironments associated with different functional and metabolic states of cell subtypes, and how these microenvironments are distributed in the spatial context in a gradient from injured to healthy tissue.

Mapping tissue states via cobordism approximation and projection (PatchMAP)

To model transitions between tissue states, such as healthy or injured, we generalized manifold learning to higher-order instances by developing a new algorithm, called PatchMAP, that integrates mutual nearest-neighbor calculations with UMAP (**Fig. 4a** and **Methods, Algorithm 2**). We hypothesized that state spaces of system-level transitions could be represented as nonlinear manifolds. Accordingly, PatchMAP represents the disjoint union of phenotypic manifolds as the boundary of a nonlinear cobordism that maps the transition between phenotypic manifolds. Patches across boundary manifolds are connected by pairwise nearest-neighbor queries that represent cobordism geodesics and stitched using the t-norm to make their metrics compatible. Interpreting PatchMAP embeddings is analogous to existing

dimensionality reduction algorithms – similar data within or across boundary manifolds are located close to each other, while dissimilar data are farther apart. PatchMAP incorporates both boundary manifold topological structure and continuities across boundary manifolds to produce cobordisms.

Currently, no machine learning method to form cobordisms exists – methods closest to doing this are single-cell data integration algorithms^{13, 42, 43}. Therefore, to benchmark PatchMAP’s manifold stitching, we compared it to the data integration methods BBKNN⁴³, Seurat v3¹³, and Scanorama⁴² in a stitching simulation using the “digits” machine learning method development data set (**Fig. 4b**). We split data by label into boundary manifolds, then applied each method to re-stitch the full data set. In this task, perfect stitching would result in the separation of projected boundary manifolds, which we quantified with the silhouette coefficient (SC). For controlled visualization, UMAP was used after data integration to provide analogous embeddings to PatchMAP’s for all benchmark methods.

PatchMAP outperformed data integration methods at higher nearest-neighbor (NN) counts in the stitching simulation (**Fig. 4b**). All other methods incorrectly mixed boundary manifolds when there was no overlap, as expected given that lack of manifold connections violated their assumptions. By contrast, PatchMAP’s stitching uses a fuzzy set intersection³⁵, which prunes incorrectly connected data across manifolds while reinforcing correct connections. We also validated that PatchMAP preserves boundary manifold organization while embedding higher-order structures between similar boundary manifolds (**Supplementary Fig. 9**). At low NN values and when boundary manifolds are similar, PatchMAP resembles UMAP projections (**Supplementary Fig. 9**). At higher NN values, manifold annotations are strongly weighted, which results in less mixing and better manifold separation.

Information transfer across imaging technologies and tissues (i-PatchMAP)

We posited that the transfer of information across biological states should account for continuous transitions and be robust to manifold connection strength (including the lack thereof). The i-PatchMAP workflow therefore uses PatchMAP as a paired domain transfer and quality control visualization method to propagate information between different imaging methods (**Fig. 4a**, information transfer). To transfer information across data sets, i-PatchMAP first computes one-step Markov chain transition probabilities between boundary manifolds of “reference” and “query” data along cobordisms (**Fig. 4a**, transition probabilities), and then linearly interpolates measures from reference to query data (**Fig. 4a**, propagate information). Quality control of i-PatchMAP can be performed by visualizing connections between boundary manifolds in PatchMAP embeddings (**Fig. 4a**, visualize manifold connections).

To benchmark i-PatchMAP, we compared it to other nonparametric domain transfer tools, Seurat v3 and a modification of UMAP incorporating a transition probability-based interpolation similar to i-PatchMAP (UMAP+), first on data multimodal single-cell data from DFU ROIs (see **Fig. 3**) and then separately, using a publicly available cord blood mononuclear cell (CBMC) CITE-seq data set¹¹. UMAP+ utilized a directed NN graph from query to reference data for data interpolation, rather than PatchMAP's metric-compatible stitching, therefore acting as a control for PatchMAP. We tiled ROIs from to construct 23 evaluation instances and performed leave-one-out cross-validation using single-cell MSI profiles to predict IMC protein expression. We assessed accuracy using Spearman's correlation between the predicted spatial autocorrelation (Moran's I) and the true autocorrelation for each parameter to account for resolution differences between imaging modalities (**Fig. 4c**). i-PatchMAP outperformed tested methods on its ability to transfer IMC measures to query data based on MSI profiles (**Fig. 4d**). However, all methods demonstrated consistently poor performance for parameters with no original spatial autocorrelation within tiles (TGF- β , FoxP3, CD163). For the CITE-seq data set, we created 15 evaluation instances and used single-cell RNA profiles to predict antibody derived tag (ADT) abundance. We quantified performance using Pearson's correlation between true and predicted ADT values (**Fig. 4e**), finding that i-PatchMAP performed as well or slightly better than other tested methods for all parameters.

i-PatchMAP transfers multiplexed protein distributions across tissues based on molecular microenvironmental profiles

To assess whether i-PatchMAP can be used to transfer molecular signature information across imaging modalities and further, across different tissue samples, we used single-cell IMC/MSI protein measures (see **Fig. 3**) to extrapolate IMC information from discrete ROIs to the full DFU sample, as well as to distinct prostate tumor and tonsil specimens, based on MSI profiles. A PatchMAP embedding of single cells in DFU ROIs and individual pixels across tissues based on MSI parameters revealed that single-cell molecular microenvironments in the DFU ROIs provided a good representation of the overall DFU molecular profile (**Fig. 4f**). Accordingly, we used i-PatchMAP to transfer DFU single-cell protein measures from the ROIs to the full DFU tissue based on molecular similarities (**Fig. 4g**). i-PatchMAP predicted that the wound area of the DFU tissue would show high expression levels for CD68, a marker of pro-inflammatory macrophages, and activated Caspase-3, a marker of apoptotic cell death. By contrast, the healthy margin of the DFU biopsy was predicted to contain higher levels of CD4, indicating infiltrating T cells, and the cell proliferation marker Ki-67. Interestingly, the PatchMAP visualization revealed that

molecular microenvironments corresponding to specific single-cell measures in the DFU (e.g., CD4) were strongly connected with MSI pixels in the tonsil tissue sample (**Fig. 4h**). In the tonsil tissue, which is lymphocyte rich, i-PatchMAP predictions for CD4 matched well with follicular structure, and areas that lacked cell content were accurately predicted to not contain CD4. By contrast, there were no strong connections between the molecular profiles of the DFU biopsy and the prostate cancer sample, which showed poor immune cell infiltration by IMC examination (**Supplementary Fig. 10**). Thus, in the current datasets, strong inter-sample cellular and molecular connections are supported by the presence or absence of common immune cell populations.

DISCUSSION

We developed MIAAIM, a modular, reproducible computational framework to align, analyze, and transfer information across high-parameter bioimaging modalities. With MIAAIM we have demonstrated that highly detailed and mathematically grounded structures of high-dimensional bioimaging modalities can be linked and propagated across tissues, tissue states, and modalities. As part of MIAAIM, we have developed a manifold learning algorithm based on cobordism theory to translate multiplexed information across tissues. Our combined IMC-MSI approach provided insights into how molecular microenvironmental niches of individual cells may correlate with their states and localization within tissue sections. We showed MIAAIM's generalizability across diverse healthy and pathological human tissue types and imaging technologies, including those that generate large, whole-tissue datasets.

In its current version, MIAAIM does not incorporate deep learning, which would require large training sets across multiple imaging technologies for generalized application. To our knowledge, no learning models for multiplexed image registration exist, likely due to the large degree of variability in tissue organization across specimens. Similar limitations likely restrict generalized application of domain transfer neural networks for antibody-free labeling of growing repositories of publicly available images. Improvements in deep learning robustness may make its application routine in these settings – a key will be limiting model complexity through appropriate parametrization and regularization. MIAAIM currently provides a first framework to generate standardized training sets for the development of such methods. Its modular implementation will enable versatile workflow incorporation of new method developments such as those from the deep learning community.

With larger sample sizes, our multimodal imaging approach is well positioned to identify novel biomarkers to support diagnosis and prognosis in clinical cases in a disease-agnostic manner. PatchMAP

applied to single-cell data may provide a coordinate system on a manifold *between* single-cell network states (i.e., boundary manifolds) that could be used to infer a “single-cell network phenotypic tree” (i.e., cobordism). Computational views of biological systems have proven to be useful⁴⁴; cobordism theory in this context may indicate a deeper relevance of bridging pure mathematics with modern biology. Indeed, connections between cobordisms and vector spaces are the core of a topological field theory⁴⁵ – translating these concepts to biology may provide a formal, algebraic language to map transitions between cellular system states.

In conclusion, MIAAIM provides a conceptual and practical foundation for novel applications in high-dimensional bioimage integration, interrogating composite tissue states, and transferring data across growing numbers of high-parameter imaging modalities. MIAAIM is technology-agnostic and works with any data in BioFormats/OME-compatible, imzML, NIfTI, and HDF5 formats. We envision MIAAIM to have broad applicability, from single-cell proteomic and transcriptomic imaging to untargeted molecular imaging. The harmonization of these data is poised to provide more comprehensive tissue portraits across health and disease states, and transitions between them.

REFERENCES

1. Angelo, M. et al. Multiplexed ion beam imaging of human breast tumors. *Nature medicine* **20**, 436 (2014).
2. Giesen, C. et al. Highly multiplexed imaging of tumor tissues with subcellular resolution by mass cytometry. *Nature methods* **11**, 417-422 (2014).
3. Goltsev, Y. et al. Deep profiling of mouse splenic architecture with CODEX multiplexed imaging. *Cell* **174**, 968-981. e915 (2018).
4. Gut, G., Herrmann, M.D. & Pelkmans, L. Multiplexed protein maps link subcellular organization to cellular states. *Science* **361** (2018).
5. Lin, J.-R., Fallahi-Sichani, M. & Sorger, P.K. Highly multiplexed imaging of single cells using a high-throughput cyclic immunofluorescence method. *Nature communications* **6**, 1-7 (2015).
6. Lin, J.-R. et al. Highly multiplexed immunofluorescence imaging of human tissues and tumors using t-CyCIF and conventional optical microscopes. *Elife* **7** (2018).
7. McDonnell, L.A. & Heeren, R.M. Imaging mass spectrometry. *Mass spectrometry reviews* **26**, 606-643 (2007).
8. Vickovic, S. et al. High-definition spatial transcriptomics for in situ tissue profiling. *Nature methods* **16**, 987-990 (2019).
9. Rodriques, S.G. et al. Slide-seq: A scalable technology for measuring genome-wide expression at high spatial resolution. *Science* **363**, 1463-1467 (2019).
10. Peterson, V.M. et al. Multiplexed quantification of proteins and transcripts in single cells. *Nature biotechnology* **35**, 936-939 (2017).
11. Stoeckius, M. et al. Simultaneous epitope and transcriptome measurement in single cells. *Nature methods* **14**, 865-868 (2017).
12. Hao, Y. et al. Integrated analysis of multimodal single-cell data. *Cell*.
13. Stuart, T. et al. Comprehensive integration of single-cell data. *Cell* **177**, 1888-1902. e1821 (2019).
14. Abdelmoula, W.M. et al. Automatic 3D nonlinear registration of mass spectrometry imaging and magnetic resonance imaging data. *Analytical chemistry* **91**, 6206-6216 (2019).
15. Abdelmoula, W.M. et al. Automatic generic registration of mass spectrometry imaging data to histology using nonlinear stochastic embedding. *Analytical chemistry* **86**, 9204-9211 (2014).
16. Haghverdi, L., Büttner, M., Wolf, F.A., Buettner, F. & Theis, F.J. Diffusion pseudotime robustly reconstructs lineage branching. *Nature methods* **13**, 845 (2016).
17. Setty, M. et al. Characterization of cell fate probabilities in single-cell data with Palantir. *Nature biotechnology* **37**, 451-460 (2019).
18. Moon, K.R. et al. Visualizing structure and transitions in high-dimensional biological data. *Nature Biotechnology* **37**, 1482-1492 (2019).
19. Levine, J.H. et al. Data-driven phenotypic dissection of AML reveals progenitor-like cells that correlate with prognosis. *Cell* **162**, 184-197 (2015).
20. Samusik, N., Good, Z., Spitzer, M.H., Davis, K.L. & Nolan, G.P. Automated mapping of phenotype space with single-cell data. *Nature methods* **13**, 493-496 (2016).
21. Rozenblatt-Rosen, O. et al. The Human Tumor Atlas Network: charting tumor transitions across space and time at single-cell resolution. *Cell* **181**, 236-249 (2020).
22. Consortium, H. The human body at cellular resolution: the NIH Human Biomolecular Atlas Program. *Nature* **574**, 187 (2019).

- 407 23. Thom, R. Quelques propriétés globales des variétés différentiables. *Commentarii Mathematici*
408 *Helvetici* **28**, 17-86 (1954).
- 409 24. Schapiro, D. et al. MCMICRO: A scalable, modular image-processing pipeline for multiplexed
410 tissue imaging. *bioRxiv* (2021).
- 411 25. Schapiro, D. et al. MITI Minimum Information guidelines for highly multiplexed tissue images.
412 *arXiv preprint arXiv:2108.09499* (2021).
- 413 26. Linkert, M. et al. Metadata matters: access to image data in the real world. *Journal of Cell*
414 *Biology* **189**, 777-782 (2010).
- 415 27. Peng, T. et al. A BaSiC tool for background and shading correction of optical microscopy images.
416 *Nature communications* **8**, 1-7 (2017).
- 417 28. Muhlich, J.L., Chen, Y.-A., Russell, D.P.W. & Sorger, P.K. Stitching and registering highly
418 multiplexed whole slide images of tissues and tumors using ASHLAR software. *bioRxiv* (2021).
- 419 29. Schramm, T. et al. imzML—a common data format for the flexible exchange and processing of
420 mass spectrometry imaging data. *Journal of proteomics* **75**, 5106-5110 (2012).
- 421 30. Klein, S., Staring, M., Murphy, K., Viergever, M.A. & Pluim, J.P. Elastix: a toolbox for intensity-
422 based medical image registration. *IEEE transactions on medical imaging* **29**, 196-205 (2009).
- 423 31. Schapiro, D. et al. histoCAT: analysis of cell phenotypes and interactions in multiplex image
424 cytometry data. *Nature methods* **14**, 873 (2017).
- 425 32. Merkel, D. Docker: lightweight linux containers for consistent development and deployment.
426 *Linux journal* **2014**, 2 (2014).
- 427 33. Di Tommaso, P. et al. Nextflow enables reproducible computational workflows. *Nature*
428 *biotechnology* **35**, 316-319 (2017).
- 429 34. Wilkinson, M.D. et al. The FAIR Guiding Principles for scientific data management and
430 stewardship. *Scientific data* **3**, 1-9 (2016).
- 431 35. McInnes, L., Healy, J. & Melville, J. Umap: Uniform manifold approximation and projection for
432 dimension reduction. *arXiv preprint arXiv:1802.03426* (2018).
- 433 36. Sainburg, T., McInnes, L. & Gentner, T.Q. Parametric UMAP: learning embeddings with deep
434 neural networks for representation and semi-supervised learning. *arXiv preprint*
435 *arXiv:2009.12981* (2020).
- 436 37. Rashid, R. et al. Highly multiplexed immunofluorescence images and single-cell data of immune
437 markers in tonsil and lung cancer. *Scientific data* **6**, 1-10 (2019).
- 438 38. Costa, J.A. & Hero, A.O. in 2004 IEEE International Conference on Acoustics, Speech, and Signal
439 Processing, Vol. 3 iii-988 (IEEE, 2004).
- 440 39. Costa, J.A. & Hero, A.O. Geodesic entropic graphs for dimension and entropy estimation in
441 manifold learning. *IEEE Transactions on Signal Processing* **52**, 2210-2221 (2004).
- 442 40. Staring, M., Van Der Heide, U.A., Klein, S., Viergever, M.A. & Pluim, J.P. Registration of cervical
443 MRI using multifeature mutual information. *IEEE transactions on medical imaging* **28**, 1412-
444 1421 (2009).
- 445 41. Berg, S. et al. ilastik: Interactive machine learning for (bio) image analysis. *Nature Methods*, 1-7
446 (2019).
- 447 42. Hie, B., Bryson, B. & Berger, B. Efficient integration of heterogeneous single-cell transcriptomes
448 using Scanorama. *Nature biotechnology* **37**, 685-691 (2019).
- 449 43. Polański, K. et al. BBKNN: fast batch alignment of single cell transcriptomes. *Bioinformatics* **36**,
450 964-965 (2020).

- 451 44. Regev, A. & Shapiro, E. Cellular abstractions: Cells as computation. *Nature* **419**, 343-343 (2002).
- 452 45. Atiyah, M.F. Topological quantum field theory. *Publications Mathématiques de l'IHÉS* **68**, 175-
- 453 186 (1988).
- 454

METHODS

MIAAIM implementation. MIAAIM workflows are implemented in Python and connected via the Nextflow pipeline language to enable automated results caching and dynamic processing restarts after alteration of workflow parameters, and to streamline parallelized processing of multiple images. MIAAIM is also available as a Python package. Each data integration workflow is containerized to enable reproducible environments and eliminate any language-specific dependencies. MIAAIM's output interfaces with a number of existing image analysis software tools (see **Supplementary Note 1, Combining MIAAIM with existing bioimaging software**). MIAAIM therefore supplements existing tools rather than replaces them.

High-dimensional image compression and pre-processing (HDIprep). HDIprep is implemented by specifying sequential processing steps. Options include image compression for high-parameter data, and filtering and morphological operations for single-channel images. Processed images were exported as 32-bit NIfTI-1 images using the NiBabel library in Python. NIfTI-1 was chosen as the default file format for many of MIAAIM's operations due to its compatibility with Elastix, ImageJ for visualization, and its memory mapping capability in Python.

To compress high-parameter images, HDIprep identifies a steady-state embedding dimensionality for pixel-level data. Compression is initialized with optional, spatially-guided subsampling to reduce data set size. We then implement UMAP to construct a graph representing the data manifold and its underlying topological structure (*FuzzySimplicialSet*, **Algorithm 1**). UMAP aims to optimize an embedding of a high-dimensional fuzzy simplicial set (i.e., a weighted, undirected graph) so that the fuzzy set cross-entropy between the embedded simplicial set and the high-dimensional counterpart is minimized, where the fuzzy set cross-entropy is defined as³⁵:

Definition 1. Given a reference set A and membership functions $u: A \rightarrow [0,1]$, $v: A \rightarrow [0,1]$, the *fuzzy set cross-entropy* C of (A, u) and (A, v) is defined as:

$$C((A, u), (A, v)) \triangleq \sum_{a \in A} \left(u(a) \log \left(\frac{u(a)}{v(a)} \right) + (1 - u(a)) \log \left(\frac{1 - u(a)}{1 - v(a)} \right) \right)$$

The fuzzy set cross-entropy is a global measure of agreement between simplicial sets, aggregated across members of the reference set A (here, graph edges). Calculating its exact value scales quadratically

with the number of data points, restricting its use for large data sets. UMAP's current implementation does not, therefore, compute the exact cross entropy during its optimization of low-dimensional embeddings. Instead, it relies on probabilistic edge sampling and negative sampling to reduce runtimes for large data sets³⁵. Congruently, to identify steady-state embedding dimensionalities, we compute patches on the data manifold that are representative of its global structure, and we use these patches in the calculation of the exact cross-entropy after projecting them with UMAP over a range of dimensionalities. The result is a global estimate of the dimensionality required to accurately capture manifold complexity.

To identify globally representative patches on the data manifold, we subject the fuzzy simplicial set to a variant of spectral clustering. We iteratively project the spectral centroids into Euclidean spaces of increasing dimensionality using UMAP and calculate the fuzzy set cross-entropy in each case, then min-max normalize the obtained values. To identify the steady-state embedding dimensionality, we fit a least-squares exponential regression to normalized cross-entropy as a function of dimensionality, then simulate samples along the regression line to find the first dimensionality falling within the 95% confidence interval of the exponential asymptote. The subsampled data is embedded into the steady-state dimensionality, and out-of-sample pixels are projected into this embedding using the native nearest-neighbor based method in UMAP (*transform* () function). Finally, all pixels are mapped back to their original spatial coordinates to construct a compressed image with the number of channels equal to the steady-state embedding dimensionality. These steps are summarized in pseudo-code below:

Algorithm 1: Image Compression.

Input: Multichannel image (X), SVD dimensionality (b), k-means clusters (k), embedding dimensions (n)

Output: Compressed image (I)

function *Compress*

 {*Spatially Subsample Pixels*}

$p \leftarrow \subseteq \{a_{x,y} \mid a_{x,y} \in X\}$

 {*Compute Data Manifold*}

$g \leftarrow \text{FuzzySimplicialSet}(p)$

▷ UMAP

 {*Compute Spectral Centroids*}

$s \leftarrow \text{RandomizedSVD}(g, b)$

$c \leftarrow \text{MiniBatchKmeans}(s, k)$

 {*Compute Manifold Embedding Error*}

$g_c \leftarrow \text{FuzzySimplicialSet}(c)$

for $i = 1 \dots n$ **do**

```

519          $e_c^i \leftarrow \text{EmbedSimplicialSet}(g_c, i)$ 
520         {Compute Fuzzy Set Cross Entropy}
521          $E_i \leftarrow \text{FuzzySetCrossEntropy}(g_c, e_c^i)$ 
522          $E \leftarrow \frac{E - E_{\min}}{E_{\max} - E_{\min}}$  ▷ Min-max Normalization
523         {Compute Steady-state Dimensionality}
524          $\hat{y} \leftarrow (y_0 - c)e^{-kE} + c$  ▷ Exponential Regression
525         for  $i = 1 \dots n$  do
526             if  $\hat{y}_i$  in  $c \pm 1.96\sigma$  do ▷ Model Gaussian Residual Process
527                  $j \leftarrow i$ 
528                 break
529         {Compute Pixel Manifold Embedding}
530         for  $k_{x,y}$  in  $\{a_{x,y} \mid a_{x,y} \in X \wedge a_{x,y} \notin p\}$  do
531              $M \leftarrow \text{ProjectData}(k_{x,y}, e_c^j)$  ▷ Project data (UMAP transform())
532         {Reconstruct Image}
533         for  $k_{x,y} \in M$  do
534              $I(x, y) \leftarrow k_{x,y}$ 
535         return  $I$ 

```

Image data subsampling. Subsampling is performed at pixel level and is optional for image compression. Implemented options include uniformly spaced grids within the (x, y) plane, random coordinate selection, and random selection initialized with uniformly spaced grids (“pseudo-random”). HDIprep also supports specification of masks for sampling regions, which may be useful for extremely large data sets.

By default, images with fewer than 50,000 pixels are not subsampled, images with 50,000-100,000 pixels are subsampled using 55% pseudo-random sampling initialized with 2x2 pixel uniformly spaced grids, images with 100,000-150,000 pixels are subsampled using 15% pseudo-random sampling initialized with 3x3 pixel grids, and images with more than 150,000 pixels are subsampled with 3x3 pixel grids. These default values are based on empirical studies (**Supplementary Figs. 4, 5, and 6**).

No subsampling was used for presented MSI data. Subsampling rates used for presented IMC data were determined on a case-by-case basis from empirical studies and match those used in the spectral landmark sampling experiments. Subsampling with 10x10 pixel uniformly spaced grids was used for CyCIF data compression.

Fuzzy simplicial set generation. To construct a pixel-level data manifold, we represent each pixel as a d -dimensional vector, where d is the number of channels in the given high-parameter image (i.e., discarding

spatial information). We then implement the UMAP algorithm and extract the resulting fuzzy simplicial set representing the manifold structure of these d -dimensional points. For all presented results, we used the default UMAP parameters to generate this manifold: 15 nearest neighbors and the Euclidean metric.

Manifold landmark selection with spectral clustering. Spectral landmarks are identified using a variant of spectral clustering. We use randomized singular value decomposition (SVD) followed by mini-batch k-means to scale spectral clustering to large data sets, following the procedure introduced in the potential of heat diffusion for affinity-based transition embedding (PHATE) algorithm¹⁸. Given a symmetric adjacency matrix A representing pairwise similarities between nodes (here, pixels) originating from a d -dimensional space $\mathbb{R}^d$, we first compute the eigenvectors corresponding to the k largest eigenvalues of A . We then perform mini-batch k-means on the nodes of A using these k eigenvectors as features. Spectral landmarks are defined as the d -dimensional centroids of the resulting clusters.

By default, the input data is reduced to 100 components using randomized SVD and then split into 3,000 clusters using mini-batch k-means. These default parameter values are based on empirical studies (**Supplementary Fig. 7**). Due to steady-state embeddings of MSI and IMC data only being available after experimental tests, no landmark selection was used for processing or determining the optimal embedding dimensionality of these data sets. Instead, full or subsampled datasets were used. All other steady-state embeddings for image data was compressed using the above default parameters.

Steady-state UMAP embedding dimensionalities. By default, HDIprep embeds spectral landmarks into Euclidean spaces with 1-10 dimensions to identify steady-state embedding dimensionalities. Exponential regressions on the spectral landmark fuzzy set cross entropy are performed using built-in functions from the Scipy Python library. These default parameters were used for all presented data.

Histology image preprocessing. HDIprep processing options for hematoxylin and eosin (H&E) stained tissues and other low-channel histological stains include image filters (e.g., median), thresholding (e.g., manually set or automated), and successive morphological operations (e.g., thresholding, opening and closing). Presented H&E and toluidine-blue stained images were processed using median filters to remove salt-and-pepper noise, followed by Otsu thresholding to create a binary mask representing the foreground. Sequential morphological operations were then applied to the mask, including morphological opening to

remove small connected foreground components, morphological closing to fill small holes in foreground, and filling to close large holes in foreground.

Image compression with UMAP parametrized by a neuronal network. We implemented parametric UMAP³⁶ using the default parameters and neural architecture with a TensorFlow backend. The default architecture was comprised of a 3-layer 100-neuron fully connected neuronal network. Training was performed using gradient descent with a batch size of 1,000 edges and the Adam optimizer with a learning rate of 0.001.

High-dimensional image registration (HDIreg). HDIreg is a containerized workflow implementing the open-source Elastix software³⁰ in conjunction with custom-written Python modules to automate the image resizing, padding, and trimming often applied before registration. HDIreg incorporates several different registration parameters, cost functions, and deformation models, and additionally allows manual definition of point correspondences for difficult problems, as well as composition of transformations for fine-tuning (see **Supplemental Note 2, Notes on the HDIreg workflow’s expected performance**).

High-parameter images are registered using a manifold alignment scheme parametrized by spatial transformations, which aims to maximize image similarity. Formally, we view registration as the following optimization problem⁴⁰:

Given a fixed d -dimensional image $I_F: \mathbb{R}^d \rightarrow \mathbb{R}^2$ with domain Ω_F and a moving q -dimensional image $I_M: \mathbb{R}^q \rightarrow \mathbb{R}^2$ with domain Ω_m , we aim to optimize

$$\hat{\mu} = \arg \min_{\mu} S(\mathbf{T}_{\mu}; I_F, I_M, \Omega_F) \quad (1)$$

where $\mathbf{T}_{\mu}: \Omega_F \rightarrow \Omega_M$ is a smooth transformation defined by the vector of parameters $\mu \subset \mathbb{R}^M$, and S is a similarity measure maximized when $I_M \circ \mathbf{T}_{\mu}$ and I_F are aligned.

Differential geometry and manifold learning: MIAAIM’s manifold alignment scheme uses the entropic graph-based Rényi α -mutual information (α -MI)⁴⁰ as the similarity measure S in **Equation 1**, which extends to manifold representations of images (i.e., compressed images) embedded in Euclidean space with potentially differing dimensionalities. This measure is justified in the HDIreg manifold alignment scheme through the notion of *intrinsic manifold information* (i.e., entropy). In what follows, we introduce

basic differential geometric concepts that enable us to expand existing foundations of intrinsic manifold entropy estimation to the UMAP algorithm. We assume familiarity with the definition of a topological space (see Armstrong⁴⁶ for an introduction).

Definition 2: Let X and Y be topological spaces. A function $f: X \rightarrow Y$ is *continuous* if for each point $x \in X$ and each open neighborhood N of $f(x) \in Y$ the set $f^{-1}(N)$ is an open neighborhood of $x \in X$. A function $f: X \rightarrow Y$ is a *homeomorphism* if it is one-to-one, onto, continuous, and has a continuous inverse. When a homeomorphism exists between spaces X and Y , they are called *homeomorphic* spaces.

Definition 3. A *manifold* $\mathcal{M}$ of dimension n (i.e., an n -manifold) is a second-countable Hausdorff space, each point of which has an open neighborhood homeomorphic to n -dimensional Euclidean space, $\mathbb{R}^n$. For any open set $U \subset \mathcal{M}$ we can define a chart (φ, U) where $\varphi: U \rightarrow \tilde{U} \subset \mathbb{R}^n$ is a homeomorphism. We can say that (φ, U) acts as a local coordinate system for $\mathcal{M}$, and we can define a transition between two charts (φ, U) and (ω, V) as $\varphi \circ \omega^{-1}: \omega(U \cap V) \rightarrow \varphi(U \cap V)$ when $U \cap V$ is non-empty.

Definition 4. A *smooth manifold* is a manifold where there exists a smooth transition map between each chart of $\mathcal{M}$. A *Riemannian metric* g is a mapping that associates to each point $y \in \mathcal{M}$ an inner product $g_y(\cdot, \cdot)$ between vectors tangent to $\mathcal{M}$ at y . We denote tangent spaces of y as $T_y\mathcal{M}$. A *Riemannian manifold*, written $(\mathcal{M}, g)$, is a smooth manifold $\mathcal{M}$ together with a Riemannian metric g . Given a Riemannian manifold, the *Riemannian volume element* provides a means to integrate a function with respect to volume in local coordinates. Given $(\mathcal{M}, g)$, we can express the volume element ω in terms of the metric g and the local coordinates of a point $x = x_1, \dots, x_n$ as $\omega = g(x)dx_1 \wedge \dots \wedge dx_n$ where $g(x) > 0$ and $\wedge$ indicates the wedge product. The volume of $\mathcal{M}$ under this volume form is given by $Vol(\mathcal{M}) = \int_{\mathcal{M}} \omega$.

Definition 5. An *immersion* of a smooth n -manifold $\mathcal{M}$ into $\mathcal{N}$ is a differentiable mapping $\psi: \mathcal{M} \rightarrow \mathcal{N}$ such that $d\psi_p: T_p\mathcal{M} \rightarrow T_{\psi(p)}\mathcal{N}$ is injective for all points $p \in \mathcal{M}$. ψ is therefore an immersion if its derivative is everywhere injective.

Definition 6. An *embedding* between smooth manifolds $\mathcal{M}$ and $\mathcal{N}$ is a smooth function $f: \mathcal{M} \rightarrow \mathcal{N}$ such that f is an immersion and an embedding of topological spaces (i.e., is an injective homeomorphism).

Let $(\mathcal{M}, g)$ be a compact, n -dimensional Riemannian manifold immersed in an ambient $\mathbb{R}^d$, where $n \ll d$, and let $\mathcal{X}_j = \{X_1, \dots, X_j\}$ be a set of independent and identically distributed random vectors with values drawn from a distribution supported by $\mathcal{M}$. Define $\mathcal{U}_j = \{U_1, \dots, U_j\}$ to be open neighborhoods of the elements of $\mathcal{X}_j$. An aim of manifold learning is to approximate an embedding f such that a measure of distortion D is minimized between $\mathcal{U}_j$ and $f(\mathcal{U}_j) = \{f(U_1), \dots, f(U_j)\}$. The manifold learning problem can therefore be written as

$$\hat{f} = \arg \min_{f \in \mathcal{F}} D(\mathcal{U}_j, f(\mathcal{U}_j)) \quad (2)$$

where $\mathcal{F}$ represents the family of possible measurable functions taking $x \mapsto f(x)$. In machine learning settings, open neighborhoods $U_i \in \mathcal{U}_j$ about vectors $X_i \in \mathcal{X}_j$ are often defined to be KNN neighborhoods with geodesic distances (or probabilistic encodings thereof⁴⁷) within each open neighborhood approximated with a positive definite kernel⁴⁸, which allows the computation of inner products in a Riemannian framework (as compared with a pseudo-Riemannian framework which need not be positive definite). Measures of distortion vary by algorithm (see **Supplementary Note 3, HDIprep dimension reduction validation** for examples). Of interest in our exposition are the measures induced by the embedded geodesics via the volume elements of these coordinate patches. These provide the components needed to quantify the *intrinsic Rényi α -entropy* of embedded data manifolds.

Entropic graph estimators. Given Lebesgue density f and identically distributed random vectors $\mathbf{X}_1, \dots, \mathbf{X}_n$ with values in a compact subset of $\mathbb{R}^d$, the *extrinsic Rényi α -entropy* of f is given by:

$$H_{\alpha}^{\mathbb{R}^d}(f) = \frac{1}{1 - \alpha} \log \int_{\mathbb{R}^d} f^{\alpha}(x) dx \quad (3)$$

where $\alpha \in (0, 1)$.

Definition 7 (adapted from Costa and Hero³⁸). Let $\mathcal{X}_n = \{\mathbf{X}_1, \dots, \mathbf{X}_n\}$ be identically distributed random vectors with values in a compact subset of $\mathbb{R}^d$, the nearest neighbor of $\mathbf{X}_i \in \mathcal{X}_n$ under the Euclidean metric is given by:

$$\arg \max_{\mathbf{X} \in \mathcal{X}_n \setminus \{\mathbf{X}_i\}} \|\mathbf{X} - \mathbf{X}_i\|_2 \quad (4)$$

A k -nearest neighbor (KNN) graph puts an edge between each $\mathbf{X}_i \in \mathcal{X}_n$ and its k -nearest neighbors. Let $\Gamma_{k,i} = \Gamma_{k,i}(\mathcal{X}_n)$ be the set of k -nearest neighbors of $\mathbf{X}_i \in \mathcal{X}_n$. Then the total edge length of the KNN graph for $\mathcal{X}_n$ is given by:

$$L_{\gamma,k}(\mathcal{X}_n) = \sum_{i=1}^n \sum_{\mathbf{X} \in \Gamma_{k,i}} \|\mathbf{X} - \mathbf{X}_i\|_2^\gamma \quad (5)$$

where $\gamma > 0$ is a power-weighting constant.

In practice, the *extrinsic* Rényi α -entropy of f can be suitably approximated using a class of graphs known as continuous quasi-additive graphs, including k -nearest neighbor (KNN) Euclidean graphs⁴⁹, as their edge lengths asymptotically converge to the Rényi α -entropy of feature distributions as the number of feature vectors increases⁵⁰. This property leads to the convergence of KNN Euclidean edge lengths to the *extrinsic* Rényi α -entropy of a set of random vectors with values in a compact subset of $\mathbb{R}^d$ with $d \geq 2$ ³⁸. This is a direct corollary of the **Beardwood-Halton-Hammersley Theorem** outlined below.

Beardwood-Halton-Hammersley (BHH) Theorem^{38, 49}. Let $(\mathcal{M}, g)$ be a compact Riemannian m -manifold immersed in an ambient $\mathbb{R}^d$. Suppose $\mathcal{X}_n = \{\mathbf{X}_1, \dots, \mathbf{X}_n\}$ are identically distributed random vectors with values in a compact subset of $\mathbb{R}^d$ and Lebesgue density f . Assume $d \geq 2, 1 \leq \gamma < d$ and define $\alpha = \frac{(d-\gamma)}{d}$. Then with probability 1,

$$\lim_{n \rightarrow \infty} \frac{L_{\gamma,k}(\mathcal{X}_n)}{n^{\frac{d'-\gamma}{d'}}} = \beta_{d,\gamma,k} \int_{\mathcal{M}} f^\alpha(\mathbf{x}) d\mathbf{x} \quad (6)$$

The value that determines the right side of the limit in **Equation 6** is the extrinsic Rényi α -entropy given by Equation 4. When identically distributed random vectors are restricted to a compact smooth m -manifold, $\mathcal{M}$, in an ambient $\mathbb{R}^d$, the **BHH Theorem** generalizes to enable an estimation of *intrinsic* Rényi α -entropy $H_\alpha^\mathcal{M}(f)$ of the multivariate density f on $\mathcal{M}$ defined by

$$H_\alpha^\mathcal{M}(f) = \frac{1}{1-\alpha} \log \int_{\mathcal{M}} f^\alpha(\mathbf{y}) \mu_g(d\mathbf{y}) \quad (7)$$

by incorporating the measure μ_g naturally induced by the Riemannian metric via the Riemannian volume element³⁸. This is formalized by the following given by Costa and Hero³⁸:

Theorem 1 (Costa and Hero³⁸): Let $(\mathcal{M}, g)$ be a compact Riemannian m -manifold immersed in an ambient $\mathbb{R}^d$. Suppose $\mathcal{Y}_n = \{\mathbf{Y}_1, \dots, \mathbf{Y}_n\}$ are identically distributed random vectors of $\mathcal{M}$ with bounded density f relative to the differential volume element μ_g induced by the metric g . Assume $d \geq 2, 1 \leq \gamma < d$ and define $\alpha = \frac{(d-\gamma)}{d}$. Then with probability 1,

$$\lim_{n \rightarrow \infty} \frac{L_{\gamma,k}(\mathcal{Y}_n)}{n^{\frac{d'-\gamma}{d'}}} = \begin{cases} \infty, & d' < m \\ \beta_{m,\gamma,k} \int_{\mathcal{M}} f^\alpha(\mathbf{y}) \mu_g(d\mathbf{y}), & d' = m \\ 0, & d' > m \end{cases} \quad (8)$$

where $\beta_{m,\gamma,k}$ is a constant that is independent of f and $(\mathcal{M}, g)$. Similarly, the expectation $E[L_{\gamma,k}(\mathcal{Y}_n)]/n^\alpha$ converges to the same limit.

The quantity that determines the limit when $d' = m$ is the intrinsic Renyi alpha entropy of f given by **Equation 7**. **Theorem 1** has been used in conjunction with manifold learning algorithms Isomap and a variant C-Isomap to estimate the intrinsic dimensionality of embedded manifolds³⁹. In contrast to these results that use all pairwise geodesic approximations for each point in the data set to estimate the α -entropy, we aim to provide a similar formulation utilizing local information contained in data manifolds, following the results of our dimension reduction benchmark which shows that local information preserving algorithms are well-suited for the task of high-dimensional image data compression (**Supplementary Figs 1, 2, 3**). Indeed, the *information density* of volumes of continuous regions of model families (i.e., collections of output embedding spaces or input points) have been recognized in defining an information geometry of statistical manifold learning⁴⁷, motivating the title of our software.

Entropic graph estimators on local information of embedded manifolds: In what follows, we utilize two concepts to show that the intrinsic information of multivariate probability distributions supported by embedded manifolds in Euclidean space with the UMAP algorithm can be approximated using the **BHH Theorem**: (i.) the compactness of constructed manifolds and (ii.) the conservation of Riemannian volume elements. We address (i.) with a simple proof, and we provide a motivational example of conservation of volume elements using UMAP to address (ii.).

Definition 8. A topological space X is *compact* if every open cover $\mathcal{A}$ of X contains a finite sub collection that also covers X . By finite open cover, we mean that the elements of $A_i \subset \mathcal{A}$ with $i \in \{1, \dots, n\}$ are open sets in X , and that the union of the elements of $\mathcal{A}$ equals X : $\bigcup_{i=1}^n A_i = X$.

Proposition 1. Let $n > d$ and suppose that $\mathcal{M}$ is a compact manifold of dimension r with $r \leq d$ that is immersed in ambient $\mathbb{R}^n$. Then the image $f(\mathcal{M})$ of $\mathcal{M}$ under a projection $f: \mathcal{M} \rightarrow \mathbb{R}^d$ is compact.

Proof. Let $(\mathcal{M}, g)$ be a compact Riemannian manifold (e.g., a manifold constructed with UMAP) with metric g in an ambient $\mathbb{R}^n$ and f a projection from $\mathcal{M}$ to $\mathbb{R}^d$. Since f is a projection, it is continuous, which takes compact sets to compact sets. ■

Proposition 1 shows that a d -dimensional Euclidean projection of a compact Riemannian manifold take values in a compact subset of $\mathbb{R}^d$, a necessary condition in the **BHH Theorem**. The UMAP algorithm considers fuzzy simplicial sets, i.e., manifolds, constructed from finite extended pseudo-metric spaces (see finite fuzzy realization functor, Definition 7, given by McInnes and Healy³⁵). By finite, we mean that these extended pseudo-metric spaces are constructed from a finite collection of points. If one considers this finiteness condition, then the compactness of UMAP manifolds follows naturally from **Definition 8** – given an open cover on a manifold, one can find a finite subcover. Thus, UMAP projections are compact, following **Proposition 1**.

To extend the **BHH Theorem** to the calculation of intrinsic α -entropy of UMAP embeddings as in **Equation 7**, we must show that volume elements induced via embedding are well approximated. Note that these results apply to any dimension reduction algorithm that can provably preserve distances within open neighborhoods upon embedding a compact manifold in Euclidean space. In what follows, we do not provide a proof that UMAP preserves distances within open neighborhoods about points, although, this would be an ideal scenario. Rather, we assume that this ideal scenario exists, and we describe how to find the optimal dimensionality for projecting data to satisfy this assumption.

In contrast to global data preserving algorithms, such as Isomap that calculate all pairwise geodesic distances or approximations thereof with landmark based approaches, UMAP approximates geodesic distances in open neighborhoods local to each point (see **Lemma 2** below). Given Lebesgue density μ and identically distributed random vectors $\mathbf{Y}_1, \dots, \mathbf{Y}_n$ with values constrained to lie on a compact, uniformly distributed Riemannian manifold $\mathcal{M}$, geodesics between samples $\mathbf{Y}_i$ and $\mathbf{Y}_j$ drawn from μ are encoded probabilistically with UMAP and represent a scaled exponential distribution⁵¹:

$$P_{j|i} = \exp\left(\frac{-\|Y_i - Y_j\| - \rho_i}{\sigma_i}\right) \quad (9)$$

$$P_{ij} = P_{j|i} + P_{i|j} - P_{j|i}P_{i|j} \quad (10)$$

where ρ_i is the distance from vector Y_i to its nearest neighbor and σ_i is an adaptively chosen normalization factor. Using the terminology in **Equation 2**, the objective of embedding in UMAP is given by minimizing the fuzzy simplicial set cross-entropy (**Definition 1**), which represents distortion D . Formally, given probability distribution P_{ij} encoding geodesics between samples Y_i and Y_j and probability distribution Q_{ij} encoding distances between samples $f(Y_i)$ and $f(Y_j)$, we can represent the cross-entropy loss employed by UMAP as (adapted from Narayan et. al⁵¹):

$$CE(P\|Q) = -\sum_{ij} P_{ij} \log Q_{ij} + (1 - P_{ij}) \log(1 - Q_{ij}) \quad (11)$$

where Q_{ij} is the probability distribution formed from low dimensional positions of embedded vectors $f(Y_i)$ and $f(Y_j)$ by $Q_{ij}(a, b) = \frac{1}{(1 + a\|f(Y_i) - f(Y_j)\|_2^{2b})}$ with a, b user-defined parameters to control embedding spread.

Minimizing **Equation 11** is not, in general, a convex optimization problem. Optimization over the family $\mathcal{F}$ from **Equation 2** is restricted to a subset rather than the full family and thus represents, in the best case, a local optimum. We include a larger family of measurable functions in order to more accurately approach optimal embedding of geodesic distances within open neighborhoods of vectors through the identification steady-state embedding dimensionalities, as outlined in the HDIprep workflow in a “pseudo-global” optimization procedure.

Using the notation in **Equation 2**, the volume elements induced by the embedding $\hat{f}$ are similarly those that minimize distortion between the volume elements of $\mathcal{M}$ given for local coordinates $x_1, \dots, x_n$ by $\sigma = g(x)dx$ where $dx = dx_1 \wedge \dots \wedge dx_n$ and those of $\hat{f}(\mathcal{M})$ given for local coordinates $\hat{f}x_1, \dots, \hat{f}x_n$ by $\tau = \hat{f}g(x)d(\hat{f}x)$, assuming that the dimensionality n of local coordinates is preserved. Under a globally optimal solution to **Equation 7** (i.e., when $P_{ij} = Q_{ij}$), volume elements are preserved: $\int_{\mathcal{M}} \sigma = \int_{\hat{\mathcal{M}}} \tau$. The existence of volume preserving diffeomorphisms for compact manifolds in consideration are proven by Moser⁵².

In order to identify manifold embeddings that minimize distortion of volume elements, we viewed increases in the dimensionality of real-valued data in a natural way by modeling increases in dimensionality as exponential increases in potential positions of points (i.e., increasing copies of the real

line, $\mathbb{R}^n$). The relationship between radii of open neighborhoods, volume, and density of manifolds in a learning setting is formalized by Narayan et. al⁵¹ in an application where density is preserved with a given dimensionality for embeddings by altering open neighborhood radii; however, we can readily extend this in an adapted scenario of volume preservation under fixed radii to infer a dimensionality that satisfies the assumption of geodesic distance preservation. Consider the following example:

Let $\mathbf{Y}_i \in \mathcal{M}$ be a vector of manifold $\mathcal{M}$ immersed in an ambient $\mathbb{R}^d$, and assume that the k -nearest neighbors of $\mathbf{Y}_i$ are uniformly distributed in a ball B_{r_d} of radius r_d with proportional volume $V_d \propto r_d^d$. Assume that a map f takes the open neighborhood B_{r_d} of $\mathbf{Y}_i$ to an m -dimensional ball $B_{r_m} \subseteq \mathcal{N}$ of manifold $\mathcal{N}$ with radius r_m while preserving its structure, including the uniform distribution and the induced Riemannian volume element V_m . Following Narayan et al.⁵¹, we can use the proportion $V_d \propto r_m^m$ to infer a power law relationship $r_m \propto r_d^{d-m}$ between the local radii r_m in embedding spaces and the original radius of B_{r_d} .

To extend this example, suppose that r_m and r_d are fixed (we can assume that radii are fixed in the original space, and that those in the embedding space are controlled by the a, b parameters influencing Q_{ij} in **Equation 11**). Assume that the ambient metrics of the embedding space and the original space are the same, and that they give rise to geodesic distances within B_{r_m} and B_{r_d} using the native UMAP method (**Lemma 2** below). Since ambient metrics and radii are preserved and that $r_m \propto r_d^{d-m}$, this implies that the geodesic distances δ_m and δ_d between points in B_{r_m} and B_{r_d} also exhibit a power law relationship, $\delta_m \propto \delta_d^{d-m}$. If we further assume that $\delta_m = \delta_d^d$ (i.e., that geodesic distances between points in open neighborhoods are preserved, the ideal scenario) we can solve for the relationship between geodesics and dimensionality m . Specifically, suppose that $\delta_m = \delta_d^d$. Substituting δ_d^d for δ_m , we can see that

$$\begin{aligned}\delta_d^d &= \delta_d^{d-m} \\ \rightarrow \delta_d^d &= \frac{\delta_d^d}{\delta_d^m} \\ \rightarrow \delta_d^m &= \frac{\delta_d^d}{\delta_d^d} = 1\end{aligned}$$

which implies that the geodesics within open neighborhoods of points of the original space have an exponential relationship with dimensionality m .

Using the power-law relationship between volumes of open neighborhoods in UMAP manifolds and their embedded counterparts, we can attempt to identify the dimensionality m that such that geodesics

within open neighborhoods is preserved with exponential regression. Note that geodesic distance preservation is stronger than volume preservation; however, volume preservation is implied. As a result, steady-state manifold embeddings provide a Euclidean dimensionality to approximate manifold geodesics of vectors sampled from $\mathcal{M}$, and thus the volume elements of $\mathcal{M}$. KNN graph functionals calculated in the steady-state embedding space provide the necessary machinery to calculate the intrinsic α -entropy of embedded data manifolds in MIAAIM by applying the **BHH Theorem** using the induced measure across all coordinate patches as in **Theorem 1**.

Note, though, that the distances outside of open neighborhoods of points are not guaranteed to be accurately modelled in the embedding space with UMAP if one makes the assumptions introduced in our example. Therefore, applying **Theorem 1** in conjunction with UMAP by replacing KNN graph lengths with those obtained with length functionals of geodesic minimal spanning trees (GMST), another type of entropic graph, should not be expected to reproduce the intrinsic entropy originally reported by Costa and Hero³⁹. Our primary contribution here is combining the KNN intrinsic entropy estimator with a local information preserving dimension reduction algorithm. We wish to extend these results to the setting where two such manifolds are compared, as this serves the basis for our image registration application. An entropic graph-based estimator of α -MI in an image registration setting is described by the following⁴⁰:

Let $z(x_i) = [z_i(x_i), \dots, z_d(x_i)]$ be a d -dimensional vector encoding the features of point x_i . Let $Z_f(x) = \{z^f(x_1), \dots, z^f(x_N)\}$ be the feature set of a fixed image, $Z_m(\mathbf{T}_\mu(x)) = \{z^m(\mathbf{T}_\mu(x_1)), \dots, z^m(\mathbf{T}_\mu(x_N))\}$ be the feature set of a transformed moving image at points in $T_\mu(x)$, and $z_i^{fm} = [z^f(x_i), z^m(\mathbf{T}_\mu(x_i))]$ be the concatenation of the feature vectors of the fixed and transformed moving image at x_i . Then,

$$\widehat{\alpha\text{-MI}}(\mu; Z_f, Z_m, Z_{fm}) = \frac{1}{\alpha - 1} \log \frac{1}{N^\alpha} \sum_{i=1}^N \left(\frac{\Gamma_i^{fm}}{\sqrt{\Gamma_i^f \Gamma_i^m(\mu)}} \right)^{2\gamma} \quad (12)$$

is a graph-based estimator for α -MI, where $\gamma = d(1 - \alpha)$, $0 < \alpha < 1$, and three graphs

$$\Gamma_i^f = \sum_{i=1}^k \|z^f(x_i) - z^f(x_{ip})\| \quad (13)$$

$$\Gamma_i^m(\mu) = \sum_{i=1}^k \|z^m(\mathbf{T}_\mu(x_i)) - z^m(\mathbf{T}_\mu(x_{ip}))\| \quad (14)$$

$$\Gamma_i^{fm}(\mu) = \sum_{i=1}^k \|z^{fm}(x_i, T_\mu(x_i)) - z^{fm}(x_{ip}, T_\mu(x_{ip}))\| \quad (15)$$

represent the Euclidean graph functionals (lengths) of a feature vector z to its p^{th} nearest neighbor over k considered nearest neighbors.

The Rényi α -MI provides a quantitative measure of association between the intrinsic structure of multiple manifold embeddings constructed with the UMAP algorithm. The Rényi α -MI measure extends to feature spaces of arbitrary dimensionality, which MIAAIM utilizes in combination with its image compression method to quantify similarity between steady-state embeddings of image pixels in potentially differing dimensionalities.

Multimodal image registration studies. Since data acquisition removed tissue context at regions of interest on the IMC full-tissue reference image, we first aligned full tissue sections then used the coordinates of IMC regions to extract data from all modalities for fine-tuning. Custom Python scripts were used to propagate alignment across imaging modalities. Manual landmark correspondences were used where unsupervised alignment proved suboptimal. We accounted for alignment errors around IMC regions following full-tissue registration by padding regions with extra pixels prior to cropping. All registrations involving MSI or IMC data were conducted using KNN α -MI, as in **Equation 12**, with $\alpha = 0.99$ and 15 nearest neighbors. All registrations aligning low-channel slides (IMC reference toluidine blue image and H&E) were conducted using histogram-based MI after grayscale conversion for rapid processing.

For full-tissue images, a two-step registration process was implemented by first aligning images using an affine model for the vector of parameters μ (**Equation 1**) and subsequently aligning images with a nonlinear model parametrized by B-splines. Hierarchical Gaussian smoothing pyramids were used to account for resolution differences between image modalities, and stochastic gradient descent with random coordinate sampling was used for optimization. We additionally optimized final control point grid spacings for B-spline models and the number of hierarchical levels to include in pyramidal smoothing for each MSI data set to H&E alignment individually (**Supplementary Figs. 1, 2, 3**). A final control point spacing of 300 pixels for nonlinear B-spline registrations of MSI data to corresponding H&E data balanced correct alignment with unrealistic warping, which we identified visually and by inspecting the determinants of the spatial Jacobian matrices for values that deviated substantially from 1. H&E and IMC reference tissue registrations utilized a final grid spacing of 5 pixels. Similar optimizations for numbers of pyramidal levels were carried out for these data. All data that underwent image registration were

exported and stored as 32-bit NIfTI-1 images. IMC data was not transformed and was kept in 16-bit OME-TIF(F) format.

Cobordism Approximation and Projection (PatchMAP). PatchMAP is an algorithm that constructs a smooth manifold by representing lower dimensional manifolds as its boundary, and it projects the higher-order manifold in a lower dimensional space for visualization. Higher-order manifolds produced by PatchMAP can be understood as cobordisms, which are described by the following set of definitions:

Definition 9. Given a family of sets $\{S_i: i \in I\}$ and an index set I , their *disjoint union*, denoted $S = \coprod_{i \in I} S_i$ is a set together with injective functions $\phi_i: S_i \rightarrow S$ for each S_i . The disjoint union corresponds to the coproduct of sets.

Definition 10. Two closed n -manifolds $\mathcal{M}$ and $\mathcal{N}$ are *cobordant* if their disjoint union, denoted $\mathcal{M} \coprod \mathcal{N}$, is the boundary of an $n+1$ -manifold $\mathcal{W}$. We call the manifold $\mathcal{W}$ a *cobordism*. By boundary of a d -manifold, we mean the set of points on $\mathcal{M}$ which are homeomorphic to the upper half plane, $\mathbb{R}_+^d$. We denote the boundary of $\mathcal{W}$ as $\partial\mathcal{W}$.

PatchMAP addresses cobordism learning in a semi-supervised manner, where data is assumed to follow the structure of a nonlinear cobordism, and our task is to glue lower dimensional manifolds to the boundary of a higher-dimensional manifold to produce a cobordism. Here, we would like for coordinate transformations across the cobordism to have their own geometries independent of the metrics of boundary manifolds. In practice, this property allows us to explore data within boundary manifolds without being reliant on the specific geometry of the cobordism. Ultimately, cobordism geodesics are the base component of downstream applications such as the i-PatchMAP workflow. Furthermore, we would like for the cobordism to highlight boundary manifolds whose data points are highly similar. A natural way to satisfy both of these conditions is to use the fuzzy set-theoretic foundation of the UMAP algorithm.

The primary goal of PatchMAP then is to identify a smooth manifold whose boundary is the disjoint union of smooth manifolds of lower dimensionality, and which has a metric independent of each boundary manifolds' metric that we choose to represent. We address this with a two-step algorithm that uncouples the computation of boundary manifolds from the cobordism. First, boundary manifolds are computed by applying the UMAP algorithm to each set of data with a user-provided metric. Practically,

the result of this step are symmetric, weighted graphs that represent geodesics within each boundary manifold. Our task is to construct from a finite set of n boundary manifolds $\mathcal{F} = \{(\mathcal{M}_1, g_1) \dots (\mathcal{M}_n, g_n)\}$ a manifold $(\mathcal{M}_{\mathcal{F}}, g)$, with metric g such that the disjoint union $\coprod_{i \in I} \mathcal{M}_i = \partial \mathcal{M}_{\mathcal{F}}$. We wish to approximate the geodesics of $\mathcal{M}_{\mathcal{F}}$, and for each element of $p \in \mathcal{M}_{\mathcal{F}}$, the tangent space $T_p \mathcal{M}_{\mathcal{F}}$ with inner product g_p .

Lemma 2 (McInnes and Healy³⁵). *Let $(\mathcal{M}, g)$ be a Riemannian manifold in an ambient $\mathbb{R}^n$, and let $p \in \mathcal{M}$ be a point. If g is locally constant about p in an open neighborhood $U \subseteq \mathcal{M}$ such that g is a constant diagonal matrix in ambient coordinates, then in a ball $B \subseteq U$ centered at p with volume $\frac{\pi^{\frac{n}{2}}}{\Gamma(\frac{n}{2}+1)}$ with respect to g , the geodesic distance from p to any point in $q \in B$ is $\frac{1}{r} d_{\mathbb{R}^n}(p, q)$, where r is the radius of the ball in the ambient space $\mathbb{R}^n$ and $d_{\mathbb{R}^n}$ is the metric on the ambient space.*

If we assume that data points across boundary manifolds can be compared with a suitable metric, we can compute the geodesic distances between points on the projections of each boundary manifold under a user-provided ambient metric using **Lemma 2** (for a proof of **Lemma 2**, see Appendix A of McInnes and Healy³⁵) to their disjoint union. Given two boundary manifolds $\mathcal{M}_i$ and $\mathcal{M}_j$, we compute the pairwise geodesic distances between points $p_i, p_j \in \mathcal{M}_i \coprod \mathcal{M}_j$ where $p_i \in \mathcal{M}_i \rightarrow \mathcal{M}_i \coprod \mathcal{M}_j$ and $p_j \in \mathcal{M}_j \rightarrow \mathcal{M}_i \coprod \mathcal{M}_j$ to obtain the components needed to construct the metric on $\mathcal{M}_{ij}$ where the boundary $\partial \mathcal{M}_{ij} = \mathcal{M}_i \coprod \mathcal{M}_j$. By extension, concatenation of all pairwise distance calculations between boundary manifolds $\mathcal{M}_i$ and $\mathcal{M}_j$ where $i, j \in \{1, 2, \dots, n\}; i \neq j$ to their disjoint union $\coprod_{i \in I} \mathcal{M}_i$ provides the components to construct the geodesics on the full cobordism $\mathcal{M}_{\mathcal{F}}$ across all pairwise combinations of boundary manifolds, as desired. The result of using **Lemma 2** to approximate the projections of manifold geodesics, however, is that we have directed, incompatible views of the geodesics across the cobordism, to and from, boundary manifolds. That is, $\mathcal{M}_{ij} \neq \mathcal{M}_{ji}$. We can interpret these directed geodesics as being defined on *oriented* cobordisms. We aim to encode the directed geodesics in the oriented cobordisms and the boundary manifold geodesics with a single data representation.

Our goal then is to construct an *unoriented* cobordism, where the directed geodesics of oriented cobordisms are resolved into a single, symmetric matrix representation. To do this, we can translate each of the extended pseudo-metric spaces described above into a fuzzy simplicial set using the fuzzy singular set functor (see Definition 9 by McInnes and Healy³⁵), which captures both a topological representation

of the oriented cobordisms and the underlying metric information. The incompatibilities in oriented cobordism geodesics can be resolved with a norm of our choosing. A natural choice for the fuzzy set representation that we have chosen is the t-norm (also known as a *fuzzy intersection*). If we interpret the fuzzy simplicial set representations of oriented cobordisms in a probabilistic manner, then their intersection corresponds to the joint distribution of the oriented cobordism metric spaces, which highlights the directed cobordism geodesics that occur in both directions, to and from, with a high probability.

After symmetrizing the cobordisms between each pairwise boundary manifold, the final step is to integrate the boundary manifold geodesics with the symmetric cobordism geodesics obtained with the fuzzy set intersection. Additionally, we should ensure that cobordisms with shared domains (i.e., boundary manifolds) are identified and integrated with each other. We can do this by taking a *fuzzy set union* (probabilistic t-conorm) over the family of extended pseudo-metric spaces, as in the original UMAP implementation. The result is a cobordism that contains its own geometric structure that is captured in cobordism geodesics, in addition to individual boundary manifolds that contain their own geometries.

Optimizing a low dimensional representation of the cobordism can be achieved with a number of methods – we choose to optimize the embedding using the fuzzy set cross entropy (**Definition 1**), as in the original UMAP implementation for consistency. Note that since our algorithm produces a symmetric matrix, PatchMAP could be repeatedly applied to construct “nested” cobordisms of hierarchical dimensionality.

PatchMAP implementation. To construct cobordisms, PatchMAP first computes boundary manifolds by constructing fuzzy simplicial sets from each provided set of data, i.e., tissue state, by applying the UMAP algorithm (*FuzzySimplicialSet*, **Algorithm 2**). Then, pairwise, directed nearest neighbor (NN) queries between boundary manifolds are computed in the ambient space of the cobordism (*DirectedGeodesics*, **Algorithm 2**). Directed NN queries between boundary manifolds are weighted according to the native implementation in UMAP, the method of which we refer the reader to **Equations 5 and 6**. Resulting pairwise directed NN graphs between UMAP boundary manifolds are weighted, and they reflect Riemannian metrics that are not compatible. That is, they cannot be simply added or multiplied to integrate their weights. We therefore stitch the cobordism metric and make directed NN queries compatible by applying a fuzzy simplicial intersection, which results in a weighted, symmetric graph (*FuzzyIntersection*, **Algorithm 2**). The final cobordism produced by PatchMAP is obtained by taking a fuzzy union over the family of all fuzzy simplicial sets that represent cobordisms between boundary manifolds (*FuzzyUnion*,

Algorithm 2). To represent connections between boundary manifolds in PatchMAP cobordism projections, we implemented the hammer edge bundling algorithm in the Datashader Python library. Pseudo-code outlining the PatchMAP algorithm is shown in the following:

Algorithm 2: PatchMAP

Input: Data sets $(D_1, D_2 \dots D_n)$, boundary manifold ambient metric (g_f) , cobordism ambient metric (g)

Output: Cobordism $(\mathcal{W})$

function *Stitch*

 { *Compute boundary manifolds* }

for $i \dots n$ **do**

$\mathcal{M}_i \leftarrow \text{FuzzySimplicialSet}(D_i, g_f)$ ▷ UMAP

 { *Compute geodesic projections in cobordism* }

for $i \dots n$ **do**

$\mathcal{M}_{i \rightarrow j} \leftarrow \text{DirectedGeodesics}(D_j, g) : i \neq j$ ▷ Project data (UMAP transform ())

$\mathcal{M}_{j \rightarrow i} \leftarrow \text{DirectedGeodesics}(D_i, g) : i \neq j$

$\mathcal{M}_{ij} \leftarrow \text{FuzzyIntersection}(\mathcal{M}_{i \rightarrow j}, \mathcal{M}_{j \rightarrow i})$ ▷ t-norm

 { *Integrate Cobordism* }

$\mathcal{W} \leftarrow \text{FuzzyUnion}\{\mathcal{M}_i, \mathcal{M}_{ij}\}$ ▷ t-conorm

return $\mathcal{W}$

Domain/Information Transfer (i-PatchMAP). Let $\mathcal{M}_{rq}$ be a cobordism with boundary manifolds $\mathcal{M}_r$ and $\mathcal{M}_q$ obtained with PatchMAP that come a reference and query data set, respectively. Specifically, $\mathcal{M}_{rq}$ is a matrix where rows represent points in the reference boundary manifold, columns represent the nearest neighbors of reference manifold points in the query boundary manifold under a user defined metric, and the i, j^{th} entry represents the geodesics between points $p_i, p_j \in \mathcal{M}_r \amalg \mathcal{M}_q$ such that $p_i \in \mathcal{M}_r \rightarrow \mathcal{M}_r \amalg \mathcal{M}_q$ and $p_j \in \mathcal{M}_q \rightarrow \mathcal{M}_r \amalg \mathcal{M}_q$. We compute a new feature matrix of predictions P_q for the query data set by multiplying the feature matrix to be transferred, F , with the transpose of the weight matrix W_{rq} obtained through ℓ_1 normalization of $\mathcal{M}_{rq}$:

$$P_q = F W_{rq}^T \quad (16)$$

In this context, the matrix W_{rq} can be interpreted as a single-step transition matrix of a Markov chain between states p_i and p_j derived from geodesic distances on the cobordism $\mathcal{M}_{rq}$.

Biological Methods. All patient tissue samples were obtained with approval from the Institutional Review Boards (IRB) of Massachusetts General Hospital (protocol #2005P000774) and Beth Israel Deaconess Medical Center (protocol #2018P000581).

Generation of imaging mass cytometry data. Frozen tissues were sectioned serially at a thickness of 10 μm using a Microm HM550 cryostat (Thermo Scientific) and thaw-mounted onto SuperFrost™ Plus Gold charged microscopy slides (Fisher Scientific). After temperature equilibration to room temperature, tissue sections were fixed in 4% paraformaldehyde (Ted Pella) for 10 min, then rinsed 3 times with cytometry-grade phosphate-buffered saline (PBS) (Fluidigm). Unspecific binding sites were blocked using 5% bovine serum albumin (BSA) (Sigma Aldrich) in PBS including 0.3% Triton X-100 (Thermo Scientific) for 1 hour at room temperature. Metal conjugated primary antibodies (Fluidigm) at appropriately titrated concentrations were mixed in 0.5% BSA in DPBS and applied overnight at 4 °C in a humid chamber. Sections were then washed twice with PBS containing 0.1% Triton X-100 and counterstained with iridium (Ir) intercalator (Fluidigm) at 1:400 in PBS for 30 min at room temperature. Slides were rinsed in cytometry-grade water (Fluidigm) for 5 min and allowed to air dry. Data acquisition was performed using a Hyperion Imaging System (Fluidigm) and CyTOF Software (Fluidigm), in 33 channels, at a frequency of 200 pixels/second and with a spatial resolution of 1 μm . Images were visualized with MCD Viewer software (Fluidigm) before exporting the data as text files for further analysis. After imaging, slides were rapidly stained with 0.1% toluidine blue solution (Electron Microscopy Sciences) to reveal gross morphology. Slides were digitized at a resolution of approximately 2.75 μm /pixel using a digital camera.

Generation of mass spectrometry imaging data. Paired 10 μm thick sections from the same tissue blocks as used for imaging mass cytometry were thaw mounted onto Indium-Tin-Oxide (ITO) coated glass slides (Bruker Daltonics). Tissue sections were coated with 2,5-dihydroxybenzoic acid (40 mg/mL in 50:50 acetonitrile:water including 0.1% TFA) with an automated matrix applicator (TM-sprayer, HTX imaging). Mass spectrometry imaging of sections was performed using a rapifleX MALDI TissueTyper (Bruker Daltonics, Billerica, MA). Data acquisition was performed using FlexControl software (Bruker Daltonics, Version 4.0) with the following parameters: positive ion polarity, mass scan range (m/z) of 300–1000, 1.25 GHz digitizer, 50 μm spatial resolution, 100 shots per pixel, and 10 kHz laser frequency. Regions of interest for data acquisition were defined using FlexImaging software (Bruker Daltonics, version 5.0), and individual images were visualized using both FlexImaging and SCiLS Lab (Bruker Daltonics). After data

acquisition, sections were washed with PBS and subjected to standard hematoxylin and eosin histological staining followed by dehydration in graded alcohols and xylene. The stained tissue was digitized at a resolution of 0.5 $\mu\text{m}/\text{pixel}$ using an Aperio ScanScope XT brightfield scanner (Leica Biosystems).

Mass spectrometry imaging data preprocessing. Data were processed in SCiLS LAB 2018 using total ion count normalization on the mean spectra and peak centroiding with an interval width of $\pm 25\text{mDa}$. For all analyses, a peak range of m/z 400-1,000 was used after peak centroiding, which resulted in 9,753 m/z peaks. No peak-picking was performed for presented data unless explicitly stated. Data were exported from SCiLS Lab as imzML files for further analysis and processing.

Single-cell segmentation. To quantify parameters of single-cells in IMC and registered MSI data within the DFU dataset, we performed cell segmentation on IMC ROIs using the pixel classification module in Ilastik (version 1.3.2)⁴¹, which utilizes a random forest classifier for semantic segmentation. For each ROI, two 250 μm by 250 μm areas were cropped from IMC data and exported in the HDF5 format to use for supervised training. To ensure each cropped area was a representative training sample, a global threshold for each was created using Otsu thresholding on the Iridium (nuclear) stain with Scikit-image Python library. Cropped regions were required to contain greater than 30% of pixels over their respective threshold.

Training regions were annotated for “background”, “membrane”, “nuclei”, and “noise”. Random forest classification incorporated Gaussian smoothing features, edges features, including Laplacian of Gaussian features, Gaussian gradient magnitude features, and difference of Gaussian features, and texture features, including structure tensor eigenvalues and Hessian of Gaussian eigenvalues. The trained classifier was used to predict each pixels’ probability of assignment to the four classes in the full images, and predictions were exported as 16-bit TIFF stacks. To remove artifacts in cell staining, noise prediction channels were Gaussian blurred with a sigma of 2 and Otsu thresholding with a correction factor of 1.3 was applied, which created a binary mask separating foreground (high pixel probability to be noise) from background (low pixel probability to be noise). The noise mask was used to assign zero values in the other three probability channels from Ilastik (nuclei, membrane, background) to all pixels that were considered foreground in the noise channel. Noise-removed, three-channel probability images of nuclei, membrane, and background were used for single-cell segmentation in CellProfiler (version 3.1.8)⁵³.

Single-cell parameter quantification. Single-cell parameter quantification for IMC and MSI data were performed using an in-house modification of the quantification (MCQuant) module in the multiple-choice microscopy software (MCMICRO)²⁴ to accept NifFTI-1 files after cell segmentation. IMC single-cell measures were transformed using 99th percentile quantile normalization prior to downstream analysis.

Imaging mass cytometry cluster analysis. Cluster analysis was performed in Python using the Leiden community detection algorithm with the leidenalg Python package. UMAP's simplicial set (weighted, undirected graph) created with 15 nearest neighbors and Euclidean metric was used as input to community detection.

Microenvironmental correlation network analysis. To calculate associations across MSI and IMC modalities, we used Spearman's correlation coefficient in the Python Scipy library. *M/z* peaks from MSI data with no correlations to IMC data with Bonferroni corrected P-values above 0.001 were removed from the analysis. Correlation modules were formed with hierarchical Louvain community detection using the Scikit-network package. The resolution parameter used for community detection was chosen based on the elbow point of a graph plotting resolution vs. modularity of community detection results. UMAP's simplicial set, created with 5 nearest neighbors and the Euclidean metric, was used as input for community detection after inverse cosine transformation of Spearman's correlation coefficients to form metric distances⁵⁴. Visualization of MSI correlation module trends to IMC parameters were computed using exponential-weighted moving averages in the Pandas library in Python after standard scaling IMC and MSI single-cell data. MSI moving averages were additionally min-max scaled to a range of 0-1 for plotting purposes. Differential correlations of variables *u* from MSI data and *v* from IMC data between conditions *a* and *b* were quantified and ranked using the formula given by Hsu et al.⁵⁵:

$$D_{a \rightarrow b}^{u,v} = (\rho_b^{u,v} - \rho_a^{u,v}) \times \max(|\rho_a^{u,v}|, |\rho_b^{u,v}|) \quad (17)$$

where change in correlation coefficients for each pair *u, v* between conditions are weighted according to the maximal absolute correlation coefficient among both conditions. Significance of differential correlations $\rho_b^{u,v} - \rho_a^{u,v}$ were calculated using one-sided, Bonferroni corrected z-statistics after Fisher transformation.

Dimensionality reduction algorithm benchmarking. Methods used for benchmarking dimensionality reduction algorithms are outlined in **Supplementary Note 3, HDIprep dimension reduction validation.**

1082

1083 **Spatial subsampling benchmarking.** Default subsampling parameters in MIAAIM are based on
 1084 experiments across IMC data from DFU, tonsil, and prostate cancer tissues recording Procrustes
 1085 transformation sum of squares errors between subsampled UMAP embeddings with subsequent projection
 1086 of out-of-sample pixels and full UMAP embeddings using all pixels. Spatial subsampling benchmarking
 1087 was performed across a range of subsampling percentages.

1088

1089 **Spectral landmark benchmarking.** Subsampling percentages and dimensionalities to validate landmark-
 1090 based steady-state UMAP dimensionalities were determined on a case-by-case basis from empirical
 1091 studies and match those used for presented, registered data. Parameters were chosen due to the
 1092 computational burden of computing values for cross-entropy on large data. To compare landmark steady-
 1093 state dimensionality selections to subsampled data, we compared the shape of exponential regression fits
 1094 from both sets of data using sum of squared errors. Sum of squared errors were calculated across a range
 1095 of resulting landmarks.

1096

1097 **Cobordism boundary manifold stitching simulation.** Simulations were performed using the MNIST
 1098 digits dataset in the Python Scikit-learn library using the default parameters for BKNN, Seurat v3,
 1099 Scanorama, and PatchMAP across a range of nearest neighbor values. Data points were split into
 1100 according to their digit label and stitched together using each method. Integrated data from each tested
 1101 method excluding PatchMAP was then visualized with UMAP. Quality of boundary manifold stitching
 1102 for each algorithm was quantified using the silhouette coefficient in the UMAP embedding space,
 1103 implemented in Python with the Scikit-learn library. The silhouette coefficient is a measure of dispersion
 1104 for a partition of a dataset. A high value indicates that data from the same label/type are tightly grouped
 1105 together, whereas a lower value indicates that data from different types are grouped together. The
 1106 silhouette coefficient (SC) is the average silhouette score s computed across each data point in the dataset,
 1107 given by the following:

1108

$$s(i) = \begin{cases} 1 - \frac{a(i)}{b(i)}, & \text{if } a(i) < b(i) \\ 0, & \text{if } a(i) = b(i) \\ \frac{b(i)}{a(i)} - 1, & \text{if } a(i) > b(i) \end{cases} \quad (18)$$

where $a(i)$ is the average distance of data point i to all points with its label and $b(i)$ is the average distance of point i to all other data that do not have the same label.

CBMC CITE-seq data transfer. CBMC CITE-seq data were preprocessed according to the vignette provided by the Satija lab at https://satijalab.org/seurat/articles/multimodal_vignette.html. RNA profiles were log transformed and ADT abundances were normalized using centered log ratio transformation. RNA variable features were then identified, and the dimensionality of the RNA profiles of cells were reduced using principal components analysis. The first 30 principal components of single-cell RNA profiles were used to predict single-cell ADT abundances. The CBMC dataset was split randomly into 15 evaluation instances with 75% training and 25% testing data. Training data was used to predict test data measures. Prediction quality was quantified using Pearson's correlation coefficient between true and predicted ADT abundances. Correlations were calculated using the Python library, Scipy. Seurat was implemented using the TransferData function after finding transfer anchors (FindTransferAnchors function) using the default parameters. PatchMAP and UMAP+ were applied with 80 nearest neighbors and the Euclidean metric in the PCA space.

Spatially resolved image data transfer. To benchmark MSI to IMC information transfer, we performed leave-one-out cross validation with segmented single-cells from 23 image tiles (number of cells ranging from ~100 to ~500 each) from the DFU dataset. IMC ROIs were split into four evenly sized quadrants to create 24 tiles. One tile was removed due to lack of cell content. Data was transformed prior to information transfer using principal components analysis with 15 components using the Scikit-learn library. Seurat was implemented using the TransferData function after finding transfer anchors (FindTransferAnchors function) using the default parameters and 15 principal components. PatchMAP and UMAP+ were implemented using 80 nearest neighbors and the Euclidean metric in the PCA space. Information transfer quality was computed for each predicted IMC parameter by calculating Pearson's correlation in Python with the Scipy library between the Moran's autocorrelation index of ground truth data and predicted data. Moran's autocorrelation index (I) is given by the following¹³:

$$I = \frac{N}{W} \frac{\sum_i \sum_j w_{ij} (x_i - \bar{x})(x_j - \bar{x})}{\sum_i (x_i - \bar{x})^2} \quad (19)$$

where N is the number of spatial dimensions in the data (2 for our purposes), x is the abundance of a protein of interest, $\bar{x}$ is the mean abundance of protein x , w_{ij} is a spatial weight matrix, and W is the sum of all w_{ij} .

We acknowledge the Satija lab for utilizing spatial autocorrelation to quantify information transfer quality in spatial data¹³. We calculated Moran's I using the R package *ape*. Spatial weight matrices were calculated using the R package *stats*. To account for resolution differences between IMC and MSI imaging modalities, spatial weights for predicted IMC parameters were calculated using $w = \frac{1}{d^{50}}$ and weights for true IMC parameters were calculated using $w = \frac{1}{d}$ where d is the Euclidean distance between segmented cell centroids in the tiled image. Since there was no spatial deconvolution for spatial parameter predictions, IMC predictions were at a lower resolution than the true IMC data. Thus, a value of 50 was chosen to reduce the spatial weighting of predicted parameters and provide a fair comparison between true IMC data and predicted IMC data. Full tissue IMC parameter predictions for the DFU and tonsil biopsy tissue section were performed using the same parameters as the tiled ROIs from the benchmark study above.

Key Resources Table

Reagent or Resource	Source	Identifier
<i>Biological Samples</i>		
Diabetic Foot Ulcer IMC/MSI	Fresh operative tissue sample at BIDMC	N/A
Prostate Cancer MSI	Fresh operative tissue sample at MGH	N/A
Tonsil MSI	Fresh operative tissue sample at MGH	N/A
<i>Antibodies</i>		
anti-SMA 141 Pr	Fluidigm	Cat: 3141017D
anti-CD19 142 Nd	Fluidigm	Cat: 3142014D
anti-CD117 143 Nd	Fluidigm	Cat: 3143001B
anti-CD14 144 Nd	Fluidigm	Cat: 3144025D
anti-CD16 146 Nd	Fluidigm	Cat: 3146020D
anti-CD163 147 Sm	Fluidigm	Cat: 3147021D
anti-CD11b 149 Sm	Fluidigm	Cat: 3149028D
anti-CD86 150 Nd	Fluidigm	Cat: 3150020B
anti-CD31/PECAM-1 151 Eu	Fluidigm	Cat: 3151025D
anti-CD45 152 Sm	Fluidigm	Cat: 3152016D

anti-CD44 153 Eu	Fluidigm	Cat: 3153029D
anti-CD11c 154 Sm	Fluidigm	Cat: 3154025D
anti-FoxP3 155 Gd	Fluidigm	Cat: 3155016D
anti-CD4 156 Gd	Fluidigm	Cat: 3156033D
anti-CD73 158 Gd	Fluidigm	Cat: 3158031D
anti-CD68 159 Tb	Fluidigm	Cat: 3159035D
anti-CD20 161 Dy	Fluidigm	Cat: 3161029D
anti-Pan-Keratin 162 Dy	Fluidigm	Cat: 3162033D
anti-TGF- β 163 Dy	Fluidigm	Cat: 3163010B
anti-Arginase-1 164 Dy	Fluidigm	Cat: 3164027D
anti- β -Catenin 165 Ho	Fluidigm	Cat: 3165032D
anti-IL-10 166 Er	Fluidigm	Cat: 3166008B
anti-Ki-67 168 Er	Fluidigm	Cat: 3168022D
anti-Collagen type 1 169 Tm	Fluidigm	Cat: 3169023D
anti-CD3 170 Er	Fluidigm	Cat: 3170019D
anti-Caspase-3 172 Yb	Fluidigm	Cat: 3172027D
anti-Pan-Actin 175 Lu	Fluidigm	Cat: 3175032D
anti-Histone-3 176 Yb	Fluidigm	Cat: 3176023D

Critical Software/Algorithms

MIAAIM	This study	https://github.com/JoshuaHess12/miaaim and https://github.com/JoshuaHess12/miaaim-python
HDIprep application	This study	https://github.com/JoshuaHess12/hdi-prep
HDIreg application	This study	https://github.com/JoshuaHess12/hdi-reg
(i)-PatchMAP application	This study	https://github.com/JoshuaHess12/patchmap
HDIutils application	This study	https://github.com/JoshuaHess12/hdi-utils
Elastix	Klein and Staring, et al. ³⁰	https://github.com/SuperElastix/elastix
Docker	Merkel ³²	https://www.docker.com
Nextflow	Tommaso et al. ³³	https://www.nextflow.io
MCMICRO	Schapiro et al. ²⁴	https://mcmicro.org
SCiLS Lab (2018)	Bruker	https://scils.de
Ilastik	Berg et al. ⁴¹	https://www.ilastik.org
CellProfiler	McQuin et al. ⁵³	https://cellprofiler.org
Leiden Algorithm	Traag et al. ⁵⁶	https://leidenalg.readthedocs.io/en/stable/intro.html
Scipy	Virtanen et al. ⁵⁷	https://www.scipy.org/index.html
UMAP	McInnes and Healy ³⁵	https://umap-learn.readthedocs.io/en/latest/
Scikit-learn	Pedregosa et al. ⁵⁸	https://scikit-learn.org/stable/
Scikit-image	van der Walt et al. ⁵⁹	https://scikit-image.org
Scikit-network	Bonaldi et al. ⁶⁰	https://scikit-network.readthedocs.io/en/latest/
Seurat v3	Stuart et al. ¹³	https://satijalab.org/seurat/
BBKNN	Polański et al. ⁴³	https://github.com/Teichlab/bbknn

Scanorama	Hie et al. ⁴²	https://github.com/brianhie/scanorama
-----------	--------------------------	---

1152

- 1153 13. Stuart, T. et al. Comprehensive integration of single-cell data. *Cell* **177**, 1888-1902. e1821
1154 (2019).
- 1155 18. Moon, K.R. et al. Visualizing structure and transitions in high-dimensional biological data.
1156 *Nature Biotechnology* **37**, 1482-1492 (2019).
- 1157 24. Schapiro, D. et al. MCMICRO: A scalable, modular image-processing pipeline for multiplexed
1158 tissue imaging. *bioRxiv* (2021).
- 1159 30. Klein, S., Staring, M., Murphy, K., Viergever, M.A. & Pluim, J.P. Elastix: a toolbox for intensity-
1160 based medical image registration. *IEEE transactions on medical imaging* **29**, 196-205 (2009).
- 1161 32. Merkel, D. Docker: lightweight linux containers for consistent development and deployment.
1162 *Linux journal* **2014**, 2 (2014).
- 1163 33. Di Tommaso, P. et al. Nextflow enables reproducible computational workflows. *Nature*
1164 *biotechnology* **35**, 316-319 (2017).
- 1165 35. McInnes, L., Healy, J. & Melville, J. Umap: Uniform manifold approximation and projection for
1166 dimension reduction. *arXiv preprint arXiv:1802.03426* (2018).
- 1167 36. Sainburg, T., McInnes, L. & Gentner, T.Q. Parametric UMAP: learning embeddings with deep
1168 neural networks for representation and semi-supervised learning. *arXiv preprint*
1169 *arXiv:2009.12981* (2020).
- 1170 38. Costa, J.A. & Hero, A.O. in 2004 IEEE International Conference on Acoustics, Speech, and Signal
1171 Processing, Vol. 3 iii-988 (IEEE, 2004).
- 1172 39. Costa, J.A. & Hero, A.O. Geodesic entropic graphs for dimension and entropy estimation in
1173 manifold learning. *IEEE Transactions on Signal Processing* **52**, 2210-2221 (2004).
- 1174 40. Staring, M., Van Der Heide, U.A., Klein, S., Viergever, M.A. & Pluim, J.P. Registration of cervical
1175 MRI using multifeature mutual information. *IEEE transactions on medical imaging* **28**, 1412-
1176 1421 (2009).
- 1177 41. Berg, S. et al. ilastik: Interactive machine learning for (bio) image analysis. *Nature Methods*, 1-7
1178 (2019).
- 1179 42. Hie, B., Bryson, B. & Berger, B. Efficient integration of heterogeneous single-cell transcriptomes
1180 using Scanorama. *Nature biotechnology* **37**, 685-691 (2019).
- 1181 43. Polański, K. et al. BBKNN: fast batch alignment of single cell transcriptomes. *Bioinformatics* **36**,
1182 964-965 (2020).
- 1183 46. Armstrong, M.A. Basic topology. (Springer Science & Business Media, 2013).
- 1184 47. Sun, K. & Marchand-Maillet, S. in International Conference on Machine Learning 1-9 (PMLR,
1185 2014).
- 1186 48. Jayasumana, S., Hartley, R., Salzmann, M., Li, H. & Harandi, M. Kernel methods on Riemannian
1187 manifolds with Gaussian RBF kernels. *IEEE transactions on pattern analysis and machine*
1188 *intelligence* **37**, 2464-2477 (2015).
- 1189 49. Yukich, J.E. Probability theory of classical Euclidean optimization problems. (Springer, 2006).
- 1190 50. Hero, A.O., Ma, B., Michel, O.J. & Gorman, J. Applications of entropic spanning graphs. *IEEE*
1191 *signal processing magazine* **19**, 85-95 (2002).
- 1192 51. Narayan, A., Berger, B. & Cho, H. Assessing single-cell transcriptomic variability through density-
1193 preserving data visualization. *Nature Biotechnology*, 1-10 (2021).

- 1194 52. Moser, J. On the volume elements on a manifold. *Transactions of the American Mathematical*
1195 *Society* **120**, 286-294 (1965).
- 1196 53. McQuin, C. et al. CellProfiler 3.0: Next-generation image processing for biology. *PLoS biology*
1197 **16**, e2005970 (2018).
- 1198 54. Van Dongen, S. & Enright, A.J. Metric distances derived from cosine similarity and Pearson and
1199 Spearman correlations. *arXiv preprint arXiv:1208.3145* (2012).
- 1200 55. Hsu, C.-L., Juan, H.-F. & Huang, H.-C. Functional analysis and characterization of differential
1201 coexpression networks. *Scientific reports* **5**, 1-14 (2015).
- 1202 56. Traag, V.A., Waltman, L. & van Eck, N.J. From Louvain to Leiden: guaranteeing well-connected
1203 communities. *Scientific reports* **9**, 1-12 (2019).
- 1204 57. Virtanen, P. et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature*
1205 *methods* **17**, 261-272 (2020).
- 1206 58. Pedregosa, F. et al. Scikit-learn: Machine learning in Python. *the Journal of machine Learning*
1207 *research* **12**, 2825-2830 (2011).
- 1208 59. Van der Walt, S. et al. scikit-image: image processing in Python. *PeerJ* **2**, e453 (2014).
- 1209 60. Bonald, T., de Lara, N., Lutz, Q. & Charpentier, B. Scikit-network: Graph Analysis in Python.
1210 *Journal of Machine Learning Research* **21**, 1-6 (2020).
- 1211

Data availability statement

All newly generated data from this study are available at https://www.synapse.org/MIAAIM_manuscript.

Code availability statement

All software and code that produced the findings of the study, including all main and supplemental figures, are free and publicly available via GitHub at <https://github.com/JoshuaHess12/miaaim-development>. A user manual and training guide to get started using the MIAAIM software and associated workflows is available at <https://joshuahess12.github.io/miaaim-python>. Python modules that allow for individual applications of MIAAIM workflows can be found at <https://github.com/JoshuaHess12/hdi-prep>, <https://github.com/JoshuaHess12/hdi-reg>, and <https://github.com/JoshuaHess12/patchmap>. A Python package implementing MIAAIM is available at <https://github.com/JoshuaHess12/miaaim-python>.

Author contribution statements

J.M.H., M.C.P., P.M.R. and R.F.S. designed the study. J.M.H. conceived and designed software and algorithms. R.F.S., P.M.R., and M.C.P. supervised the work. R.F.S. and P.M.R. conceived and performed MSI, IMC, and H&E multimodal imaging. J.M.H. performed processing and analysis of MSI, IMC, H&E, and CyCIF images. J.M.H., R.F.S., I.I., M.C.P. and P.M.R. designed benchmark studies and interpreted analyses. J.M.H. and D.S. developed and implemented cell segmentation and quantification pipeline. J.M.H., J.J.I., and I.I. critically reviewed mathematics of the paper. M.S.R., G.T., C.L.W., A.V., N.Y.R.A. provided MSI, IMC, and H&E images. W.M.A. and A.E.S. provided critical input. J.M.H., I.I., M.C.P., P.M.R., and R.F.S. wrote the manuscript. J.M.H., I.I., M.C.P., P.M.R. and R.F.S. finalized the manuscript with input from all authors.

ACKNOWLEDGEMENTS

J.M.H, M.C.P., P.M.R. and R.F.S. were supported by the Vaccine and Immunotherapy Center Innovation and Education Funds at Massachusetts General Hospital. J.M.H was previously supported and PMR is also supported by a Congressionally Directed Medical Research Program grant W81XWH-20-1-0301. J.M.H. is supported by a National Science Foundation Graduate Research Fellowship under grant no. 1746886. D.S. was supported by the University of Zurich BioEntrepreneur-Fellowship (BIOEF-17-001), a Swiss National Science Foundation Early Postdoc Mobility fellowship (P2ZHP3_181475, BMBF 01ZZ2004) and was a Damon Runyon Fellow supported by the Damon Runyon Cancer Research

1243 Foundation (DRQ-03-20). N.Y.R.A was partially supported by NIH grants U54-CA210180 and P41-
 1244 EB028741. We thank Enterprise Research Infrastructure & Services at Partners HealthCare for their in-
 1245 depth support and provision of computing resources.

1246

1247 **OUTSIDE INTERESTS**

1248 D.S. is a scientific consultant to Roche Glycart AG; none of these relationships are directly related to the
 1249 topic of this study. R.F.S., J.M.H., P.M.R, and M.C.P. are authors on patent PCT/US21/48928 applied for
 1250 by the General Hospital Corporation covering the described methods.

FIGURE LEGENDS

Fig. 1 | Summary of MIAAIM workflow and key outputs

A schematic representation of the MIAAIM workflow to integrate level 2 multi-modal images derived from various imaging methodologies (cyan boxes) and tissue specimen types (red boxes) to enable tissue state mapping across modalities and tissues. Key inputs and outputs (yellow boxes) of each MIAAIM module (shaded boxes) are connected through MIAAIM's Nextflow implementation (solid arrows) or exploratory analysis modules in Python (dashed arrows). Algorithms unique to MIAAIM (orange boxes) are detailed in corresponding figures in the remainder of this study (black bolded text). Dataflows associated with high-parameter images (i.e., greater than 3 channels) follow the left-side arrows through the HDIprep (green) and HDIreg (navy blue) modules for pixel-level image compression and pixel manifold-based image registration, while low-parameter images follow the right-side arrows (HDIprep and HDIreg modules, white boxes). Aligned multi-modal data may be exported in various file formats (HDIreg, left-side yellow boxes), and MIAAIM's Nextflow implementation of the Docker containerized HDIprep and HDIreg modules enables level 2 data integration reproducibility (HDIreg, right-side yellow boxes). MIAAIM's nonparametric data integration allows for human-in-the-loop quality control of aligned multi-modal data (HDIreg, right-side white box). External analysis techniques that interface with MIAAIM's key outputs are included in the schematic (bottom right white boxes). Aligned multi-modal data or multi-modal single-cell data (i.e., from cell segmentation) across tissue types and tissue states can be modeled using the PatchMAP module (pink), and data can be transferred across tissue states and modalities using the i-PatchMAP module (light blue).

Fig. 2 | HDIprep compression and spatially constrained manifold alignment with HDIreg

a. HDIprep compression algorithm. **(i)** High-dimensional modality. **(ii)** Spatial subsampling. **(iii)** Data manifold represented as a fuzzy simplicial set (UMAP). Edge bundled connectivity of the manifold is shown on two axes of the resulting steady state embedding (*fractal-like structure may not reflect biologically relevant features). **(iv)** High-connectivity landmarks identified with spectral clustering. **(v)** Spectral landmarks are embedded into a range of dimensionalities and exponential regression identifies steady-state dimensionalities. Pixel locations are used to reconstruct compressed image. **b.** HDIreg manifold alignment. **(i)** Spatial transformation is optimized to align moving image to fixed image. KNN graph lengths between resampled points (yellow) are used to compute α -MI. Edge-length distribution panels show Shannon MI between distributions of intra-graph edge lengths at resampled locations before

and after alignment (α -MI converges to Shannon MI as $\alpha \rightarrow 1$). MI values show increase in information shared between images after alignment. KNN graph connections show correspondences across modalities. (ii) Optimized transformation aligns images. Shown are results of transformed H&E image (green) to IMC (red). c. Example alignment. (i) Full-tissue MSI-to-H&E registration produces T_0 . (ii) H&E is transformed to IMC full-tissue reference, producing T_1 . (iii) ROI coordinates extract underlying MSI and IMC data in IMC reference space. (iv) H&E ROI is transformed to correct in IMC domain, producing T_2 . Final alignment applies modality-specific transformations. Shown are results for an IMC ROI.

Fig. 3 | Microenvironmental correlation network analysis (MCNA) links protein expression with molecular distributions in the DFU niche

a. MCNA UMAP of m/z peaks grouped into modules. b. Exponential-weighted moving averages of normalized ion intensities for top five positive and negative correlates to proteins. Colors indicate module assignment. Heatmaps (right) indicate Spearman's rho. c. Exponential-weighted moving averages of normalized average ion intensity per modules ordered as distance from center of wound in DFU increases. d. Raw IMC nuclear (Ir) and CD3 staining in ROI (left) (scale bars = 80 μ m). Masks showing CD3 expression (middle-left). Aligned MSI showing one of top CD3 correlates (middle-right). Overlay of CD3 expression and a top molecular correlate (right). e. Same as d at different ROI. f. Unsupervised phenotyping. Shaded box indicates CD3+ population. Heatmap indicates normalized protein expression. g. MCNA UMAP colored to reflect ions' correlations to Ki-67 within CD3+ and CD3- populations. Colors indicate Spearman's rho and size of points indicates negative log transformed, Benjamini-Hochberg corrected P-values for correlations. h. Tornado plot showing top five CD3+ differential negative and positive correlates to Ki-67 compared to the CD3- cell populations. X-axis indicates CD3+ specific Ki-67 values. Color of each bar indicates change in correlation from CD3- to CD3+ populations. i. Boxplots showing ion intensity and of top differentially correlated ions (top, positive; bottom, negative) to CD3+ specific Ki-67 expression across ROIs on the DFU. Tissue maps of top differentially associated CD3+ Ki-67 correlates (top, positive; bottom, negative) with boxes (white) indicating ROIs on the tissue that contain CD3+ cells.

Fig. 4 | Cobordism projection with PatchMAP and domain transfer with i-PatchMAP

a. Schematic representing the PatchMAP algorithm's cobordism geodesic stitching between boundary manifolds (reference and query data) to form cobordism (represented as grey surface), information transfer

1313 across cobordism geodesics with the i-PatchMAP workflow (top right), and cobordism projection
 1314 visualization (bottom). In PatchMAP, cobordisms sharing common boundary manifolds are glued along
 1315 their shared domain, and geodesic distances in the cobordism are made compatible with boundary
 1316 manifold geodesics. **b.** Boundary manifold stitching simulation. PatchMAP projection (manually drawn
 1317 dashed lines indicate stitching) and UMAP projections of integrated data are shown at NN values that
 1318 maximized the SC for each method. **c.** MSI-to-IMC data transfer with i-PatchMAP. Line plots show
 1319 Spearman's rho between predicted and true spatial autocorrelation values. **d.** MSI-to-IMC data transfer
 1320 benchmark. **e.** CBMC multimodal CITE-seq data transfer benchmark. **f.** PatchMAP of DFU segmented
 1321 single-cells (blue), DFU full tissue (red), tonsil tissue (green), and prostate tissue (orange) pixels based on
 1322 MSI profile. Individual plots show IMC expression for DFU single-cells (right). **g.** MSI-to-IMC data
 1323 transfer from DFU single-cells to the DFU full tissue using i-PatchMAP. **h.** MSI-to-IMC data transfer
 1324 from DFU single-cells to the tonsil tissue.

a High-dimensional Image Data Compression

b Image Registration: Spatially Constrained Manifold Alignment

c Example Propagating Alignment Workflow (DFU)

Hess et al. MIAAIM. Figure 3

a Cobordism Learning and Domain Transfer with (i)-PatchMAP