Stitching and registering highly multiplexed whole slide images of tissues and tumors using ASHLAR software

Jeremy Muhlich, Yu-An Chen, Douglas Russell and Peter K Sorger

Human Tumor Atlas Network

Harvard Ludwig Cancer Center

Laboratory of Systems Pharmacology

Department of Systems Biology

Harvard Medical School, Boston, Massachusetts 02115, USA

† Peter Sorger
Warren Alpert 432
200 Longwood Avenue
Harvard Medical School,
Boston MA 02115
peter_sorger@hms.harvard.edu cc: sorger_admin@hms.harvard.edu
617-432-6901

Key Words: microscopy, computation, stitching, microenvironment, cancer, image processing, tissue imaging, histology

ORCID IDS
Jeremy Muhlich 0000-0002-0811-637X
Yu-An Chen 0000-0001-7228-4696
Douglas Russell 0000-0001-7446-2353
Peter K. Sorger 0000-0002-3364-1838

ABSRACT

Widespread use of highly multiplexed microscopy to study normal and diseased tissues at a single-cell level is complicated by underdevelopment of the necessary software. This is particularly true of high resolution whole-slide imaging (WSI), which involves gigapixel datasets of specimens as large as 5 cm^2. WSI is necessary for accurate spatial analysis and a diagnostic necessity. High resolution WSI requires collection of successive image tiles; multiplexing commonly involves successive data acquisition cycles, each with a subset of dyes, antibodies or oligonucleotides. We describe a new Python tool, ASHLAR (Alignment by Simultaneous Harmonization of Layer/Adjacency Registration), that coordinates stitching and registration and scales to 10^3 or more image tiles over many imaging cycles to generate accurate, high-plex image mosaics, the key type of data for downstream visualization and computational analysis. ASHLAR is more robust and accurate than existing methods and compatible with any scanner or microscope conforming to Open Microscopy Environment standards.

INTRODUCTION

Multiple approaches have been described for performing 20-60 plex subcellular resolution microscopy on normal and diseased tissues for research and diagnostic purposes. These methods make it possible to image differentiation markers, signaling proteins, cell cycle regulators, oncogenes, and drug targets in a preserved tissue context. The resulting data can be processed to precisely determine the molecular and physical relationships of cells to each other and to noncellular components of connective tissue, basement membranes and vessels. Multiplexed spatial profiling of proteins and mRNA is also the cornerstone of large scale atlas projects such as the Human Cell Atlas [1], NIH HuBMAP consortium [2], and NCI Human Tumor Atlas Network (HTAN) [3] and promises to impact clinical practice by improving how diseases are diagnosed and individual patients matched to optimal therapies.

Highly multiplexed imaging of proteins in tissues uses antibodies to detect specific antigens, building on 80 years of experience with immunohistochemistry in research and diagnostic settings [4].

Methods such as MxIF, CyCIF, CODEX, 4i, and mIHC use conventional fluorescence microscopes whereas MIBI and IMC vaporize specimens with ion beams or lasers followed by atomic mass spectrometry [5–9,9–11]. Approaches to imaging nucleic acids are based on hybridization [12,13] and sequencing [14]. Some methods require frozen samples, but those that make use of Formaldehyde Fixed Paraffin Embedded (FFPE) specimens of the type universally acquired for diagnostic purposes can tap into large archives of human biopsy and resection specimens and are ideally suited to analysis of genetically engineered mouse models.

Existing imaging methods differ in resolution, field of view and number of distinct antigens or genes that can be detected (the assay "plex"). Widely used immunofluorescence methods (e.g. MxIF, CyCIF, mIHC, CODEX and Immuno-SABER) are cyclic approaches in which high-plex data are generated by repeated acquisition of lower-plex images, each of which has 2-6 channels of information. Each channel represents an image acquired with excitation and emission filters matching one antibody or oligo-coupled fluorophore. This makes it possible to optimally exploit the optical properties of fluorescence microscopes while allowing 60 or more distinct antigens to be interrogated from a single specimen. Most instruments for multiplexed imaging can acquire data at subcellular resolution (commonly 0.3 to 1 μm laterally) to enable analysis of intracellular structures. At this resolution, imaging a whole slide requires a grid of many separate images (10^3 or more for large specimens), each from a different lateral (x, y) stage position. These "tiles" are each multi-wavelength megapixel-scale images. The precise number of wavelengths in each tile, the number of tiles and the number of imaging "cycles" (each of which involves acquisition of a full set of tiles), differs with the microscope and the multiplexing technology. However, it is universally true that all tiles from all cycles must be merged accurately into a single high-plex "mosaic" image. Such mosaic images represent the key Level 2 or 3 data type for all subsequent visualization and quantitative data analysis [15].

Even the best microscopes require computational alignment of tiles to form a mosaic, since mechanical tolerances and imperfect calibration introduce uncertainty into recorded tile positions. To

enable assembly of a mosaic, tiles are slightly overlapped so that each pair of adjacent tiles contains some of the same exact cells. Image features in these cells are then used as reference points for "stitching" adjacent tiles into a seamless mosaic. In cyclic imaging, all tiles from the second and subsequent cycles must also be aligned to the mosaic through "registration" of image features across entire tiles. DNA-stained nuclei serve as an excellent image feature for alignment since they stain well with a variety of fluorescent dyes, are present at suitable density in most tissue types, and have sharp edges with high contrast. Multiple tools exist for registering image stacks and stitching image tiles [16–18] and some are available in common image analysis packages such as Image J [19]. However, we have found that using existing tools to stitch and register cyclic whole-slide images is unsatisfactory with respect to speed, reliability and accuracy. Existing tools also struggle with very large images and generally require substantial format conversion and file renaming, a non-trivial task when confronted with 100 GB of data contained in 10^4 megapixel-scale image tiles (a large 10-cycle whole-slide image).

In this paper we report the development of a new open source Python package, ASHLAR (Alignment by Simultaneous Harmonization of Layer/Adjacency Registration), for coordinated stitching and registration of multiplexed, multi-tile images. The package offers both a command line file-oriented interface and a documented API for incorporation into other tools. ASHLAR can directly process any image format supported by the widely used Open Microscopy Environment (OME) BioFormats library [20] and it outputs standard OME-TIFF files. We describe ASHLAR's design and implementation and compare its performance to existing tools using high-plex CyCIF images. ASHLAR is available as a Docker or Singularity container and has been incorporated into MCMICRO [15], the Nextflow-based image processing pipeline developed by HTAN; as part of MCMICRO, ASHLAR has been tested with CyCIF, CODEX and mxIF image data.

RESULTS

Assembling cyclic multi-tile fluorescence images requires a unified approach to stitching and registration

ASHLAR operates in three broad phases to convert a multi-cycle multi-tile (Level 1) dataset into a cohesive (Level 2) mosaic image [15] (**Figure 1**): (i) tiles from the first imaging cycle are stitched (ii) tiles from the second and subsequent cycles are registered to tiles from the first cycle (iii) all tiles from all cycles are merged into a mosaic image. The output of stitching and registration is a list of new, corrected positions for all tiles in each cycle. Only in the final mosaic phase is the actual full-size many-channel mosaic image created. These mosaic images can be very large and contain information spanning length scales from <1 µm (subcellular structures) to cm in dimension (gross tissue morphology). In many cases the boundary of a tissue specimen is irregular and a significant fraction of the tiles in a rectangular data collection grid contain few if any cells, posing a challenge for stitching but also an opportunity to reduce data collection demands by creating irregular-shaped tile sets that closely follow the outlines of the tissue.

To generate an initial estimate of tile positions, ASHLAR uses data from image tiles (grids of pixels), recorded stage positions, and physical pixel dimensions. ASHLAR leverages the Open Microscopy Environment Bio-Formats library[20] to extract the necessary image data and metadata (stage position and pixel size) directly from native image files produced by the great majority of commercial microscopes, obviating the need for image format conversion and manual metadata extraction. For microscopes that do not support BioFormats, ASHLAR accepts a set of TIFF files using a configurable naming convention along with explicit specification of tile overlap and acquisition order. Subsequent stitching and registration involve aligning one image to another. In stitching, the small overlapping strips of adjacent tiles are aligned (**Figure 1B**), and in registration, full tiles that cover the same region of the sample but acquired in different cycles are aligned (**Figure 1C**). We performed stitching and registration only on the reference image channel (typically Hoechst-stained nuclei) and applied the resulting

positional corrections to all other channels recorded within that cycle. This is sufficient because the chromatic aberration exhibited by research-grade wide-field microscopes is not a major contributor to image inaccuracy at resolutions typically used for tissue imaging (10X to 40X magnification, 0.4 to 0.9 NA air objectives).

ASHLAR uses the phase correlation algorithm for image alignment during both stitching and registration phases. Phase correlation is a fast, parameter-free method that computes the image alignment having maximum cross-correlation, but it is only suitable for aligning images that are translated relative to each other in X and Y; it cannot align images that differ by rotation, scaling, skew, or non-affine transformations. This trade-off is acceptable for accurate stitching of multi-tile images on the five microscopes we have tested (see methods) since the rigidity of the sample and the construction of modern stages ensures that almost all of the discrepancy between recorded stage positions and true positions can be modeled by translation alone. Stage position errors encountered between multiple imaging cycles are also purely translational, as long as the slide is always placed at exactly the same angle on the stage; this can routinely be achieved with kinematic mounts that positively register the slide in a consistent position (this is a standard feature of contemporary slide-scanning microscopes).

One other limitation of classic phase correlation is that precision is limited to whole pixels. At pixel sizes around 1 μm or larger, this represents a substantial error relative to the size of a single cell. We overcame this limitation by using an improved phase correlation algorithm [21] that offers arbitrary sub-pixel precision with minimal extra computation. An alignment precision of 0.1 pixels produced a discernable improvement in final mosaic quality over whole-pixel alignment, with diminishing returns beyond that. Another enhancement to phase correlation implemented in ASHLAR is pre-filtering input images with the discrete Laplacian operator (or Laplacian of Gaussian operator – LoG – for noisy images) to eliminate auto-correlation. It has been understood for at least a century that computing cross-correlation can yield spurious results with signals that exhibit auto-correlation, but this fact is often overlooked in practice [22,23]. We are not aware of published or open-source image stitching tools using

phase correlation that also include a decorrelation step, but our work with ASHLAR shows that it substantially improves confidence in image tile alignments.

Image Stitching

The stitching procedure begins with the creation of a node-edge adjacency graph in which nodes represent tiles (**Figure 2** – step A1). Edges are added to the graph to connect overlapping tile pairs, which are initially identified by consulting recorded stage positions and other metadata. By reading recorded stage positions directly from BioFormats metadata, it is straightforward to support samples imaged with non-rectangular grids and irregular layouts. **Figure 3A** shows the overlap in adjacent tiles associated with one edge in the adjacency graph and reveals that one image is slightly translated relative to the next – this translation represents the stage positioning error we are trying to correct. When the overlap region contains many cell nuclei or other alignment features, phase correlation can accurately and confidently compute the correct translation between the images. Phase correlation will always return *some* value for translation of any two tiles, even when the overlap region is uninformative and contains only incidental signal or background noise in the registration channel; in these cases we rely on the recorded stage positions. Uninformative overlaps in tissue sections are most commonly encountered when nuclei are scant, such as in fat, connective tissue, or regions of necrosis, and in areas in which no tissue is present, such as along the edge of a specimen, between separate pieces of tissue, or between the circular cores that make up a tissue microarray (a regular grid of 50-200 0.3 to 1.2 mm pieces of tissue arrayed on a single microscope slide).

We use normalized cross correlation (NCC) to score how well the translation returned by phase correlation aligns images, but the threshold dividing an effective alignment from a spurious one varies by dataset. We estimate this threshold by the 99th percentile of NCC values computed from a permutation test that considers 1,000 randomly selected pairs of non-adjacent tiles (**Figure 2A** – step A2); this corresponds to the unadjusted one-sided empirical p-value threshold of 0.01. We use the

negative logarithm of NCC values, hereafter referred to as E_{NCC} (NCC error), and a user-provided translation limit parameter to determine whether to accept or discard an alignment of tile pairs (**Figure 3B**). The value of the translation limit is not particularly critical, as physical translation distances for spurious alignments tend to be fairly extreme (note that the Y axis in Figure 3B is on a log scale). When an alignment is discarded, we delete the corresponding edge from the adjacency graph (**Figure 2** - step A3, **Figure 3C**). To eliminate graph cycles (loops) from the adjacency graph, a minimum spanning tree is constructed with the E_{NCC} values as the edge weights (**Figure 2 -** step A4). Since the edge deletion process in step A3 potentially splits the graph into multiple disconnected pieces, we perform the spanning tree procedure independently for each piece. It is then straightforward to walk along the graph from tile to tile, adding up each pairwise alignment along the way to obtain final corrected positions (**Figure 3D)**. In images where step A3 splits the adjacency graph into multiple pieces, we adjust the relative positions of disconnected pieces by extrapolating from known-good tile positions using a linear regression model (step A5). Even though individual pairwise tile alignments correct primarily for local uncorrelated stage position error, the aggregation of all alignments provides information about systematic errors such as miscalibrated pixel size or improper camera-stage alignment. A linear model fit to all alignments across the image yields a single affine transformation that accounts for such errors. At the end of steps A1 to A5 (**Figure 2A),** optimized global positions have been determined for all tiles in the first cycle and stitching is complete.

Image Registration

The procedure for registering subsequent cycles against the first cycle uses many of the same tools as stitching, although the goal is aligning whole tiles across data acquisition cycles rather than aligning adjacent tiles edge-to-edge within a single cycle (**Figure 2B**). First we establish a correspondence between each tile in the target cycle (the one to be registered) and the nearest tile in the first cycle by comparing recorded stage positions (Step B1). Identifying these tile correspondences is

trivial when recorded stage positions are consistent from run to run and the geometry of the image acquisition grids is identical. However a significant shift in stage positions between cycles is often introduced by microscopes that lack a physical "homing" procedure to zero stage position encoders at start-up. Shifts also arise where the planned tile grid is significantly displaced or rearranged between runs. To account for this shift when comparing tile positions, we down-sample the data by 20x and assemble low-resolution "thumbnail" mosaic images for each cycle using the recorded stage positions. We then align the thumbnails using phase correlation with sub-pixel precision to obtain a coarse alignment between the first and target cycles. Working with low-resolution images in this step saves compute time and memory while providing sufficient precision to accurately recover inter-cycle tile correspondences. Next, each corresponding tile pair is cropped to mutually overlapping regions and aligned using phase correlation (step B3). The resulting alignments are then filtered using the user-specified translation limit. Note that we do not currently use a permutation test and E_{NCC} threshold in the registration phase (step B2) as the translation limit alone has been sufficient for all images processed to date. For alignments that pass the filter, the target-cycle tiles are positioned by adding the alignment translations to the corrected positions of the corresponding tiles from the first cycle (step B4). The remaining tiles (generally those with sparse or no tissue, or where the tissue was damaged significantly during inter-cycle sample handling) are positioned using the linear regression model constructed in the stitching phase (step B5). The registration steps described above are then applied to all other cycles, establishing corrected global positions for all remaining tiles.

Mosaic image generation

The result of the stitching and registration phases is a corrected global position for every image tile. To generate the final output image mosaic, we create an empty image large enough to encompass all corrected tile positions and copy each tile into it at the appropriate coordinates. Since each pairwise image registration is accurate to 0.1 pixels as described above, adding up several of these shifts to

determine the final coordinates for a given tile generally yields non-integer values. ASHLAR defaults to applying sub-pixel translations on the tile images to account for this, but some users may prefer to round the final positions to the nearest pixel instead. Where neighboring image tiles overlap in the mosaic, they are combined with linear blending or one of several other user-selectable blending functions. The final many-channel image is then written out as a standard OME-TIFF file containing a multi-resolution image pyramid to support efficient visualization.

Performance

We identified MIST[16] as the current state-of-the-art public domain tool for stitching and registering large, tiled, microscopy images. We used the evaluation framework described by Chalfoun at al. to compare the accuracy of stitching by ASHLAR and MIST using that manuscript's most challenging dataset, which consisted of colonies of GFP-labeled stem cells grown for 2 days and imaged with 10% tile overlap (**Figure 5A**). MIST and ASHLAR yielded the same false negative and false positive counts (FN=47, FP=4), the size error distributions were nearly the same and had indistinguishable medians (median S_{err} = 0.0474%), and the distance error distributions were also similar with a difference between medians that was not statistically significant (MIST median D_{err} = 10.8 pixels, ASHLAR median D_{err} = 11.5 pixels, Mood's median test p-value = 0.32). Runtime and RAM usage were comparable (see Methods for details). We conclude that MIST and ASHLAR have similar stitching accuracy when applied to a previously described test set involving cells grown *in vitro*.

To test ASHLAR and MIST on high-plex, whole-slide tissue data, a ~24 mm x 14 mm x 5 µm section of a human normal colon sample was subjected to 9-cycle t-CyCIF [10], generating a subcellular resolution 28-plex image. Acquisition involved a grid of 609 image tiles (29 x 21 tiles), each 832 x 702 µm (1280 x 1080 pixels; 0.65 µm per pixel) in size and overlapping by 20 µm (~31 pixels or 2-3%). Images were acquired with a 20X 0.75 NA objective using 2 x 2 pixel binning on an OME BioFormats-compatible RareCyte CyteFinder microscope (additional details of the optical system are provided in

Methods). Data acquisition yielded a total of 5,481 16-bit four-channel tiles comprising 61 GB of data. Cell nuclei were stained with Hoechst 33342 in every cycle, providing reference features for image alignment. **Figure 4A** shows the resulting image mosaic following registration and stitching with ASHLAR; the quality of the alignment is highlighted at four regions of tile-tile overlap in **Figure 4B.** These data are made freely available with this publication for use in other applications as part of the EMIT dataset.

To compare the registration accuracy of images stitched with MIST and ASHLAR we used both tools to stitch and register two cycles of the colon dataset, yielding mosaic images approximately 36,000 x 22,000 pixels in dimension. Optical flow was used to quantify and visualize local alignment errors in the Hoechst reference channel. Whereas ASHLAR performs stitching and registration in a coordinated process, it was necessary to independently register stitched MIST image mosaics (See Methods for details). We found that the local registration error, obtained from the magnitude of the flow vector field, had a median value of 3.1 pixels in cycle 2 of the MIST mosaic and 0.1 pixels for the corresponding ASHLAR mosaic (**Figure 5A, B**). At a magnification sufficient to see individual cells, the error generated by MIST was readily apparent at regions of tile overlap (**Figure 5C**). Visualizing the full vector field on top of the reference channel images (**Figure 5D**) showed that the field direction was often consistent across large regions but could change dramatically at tile boundaries. This most likely arises because small local stitching differences propagate across the mosaic in a manner that is uncorrelated between cycles, leading to inter-cycle shifts that cannot be corrected by any rigid adjustment of the entire mosaic. It is important to note that this is not a weakness of MIST *per se,* but rather a consequence of applying a tool designed for stitching to stitching and registration, a use case for which it was not designed. With the ASHLAR-generated mosaic, vector field visualizations confirmed a much lower level of registration error (**Figure 5E, F).** Close inspection the few regions with high error showed that they corresponded to parts of the tissue in which cells had become physically distorted or detached from the slide as a consequence of the cyclic staining procedure. Thus, remaining errors are not

a result of errors in registration and stitching but rather extrinsic processes that must be identified and accounted for by downstream error-checking procedures.

DISCUSSION

ASHLAR provides a robust and efficient way to generate large, multi-channel, mosaic images of tissues and other biological specimens by assembling individual megapixel image tiles collected at multiple wavelengths over multiple imaging cycles. The key innovation with respect to image quality is joint optimization of stitching and registration as opposed to stitching individual cycles independently and then attempting to register mosaic images against each other. Joint optimization becomes increasingly important as the size of the specimen increases. Use of ASHLAR coupled with tile-based image acquisition and cyclic data collection makes it possible to optimally balance the resolution, size and plex of a tissue image for reliable analysis of spatial features on a wide variety of scales. Although many recent highly multiplexed studies have relied on small fields of view and TMAs, whole-slide imaging is emerging not only as a diagnostic necessity [24] but also as a key requirement for basic research into the spatial organization of tissue and tumor microenvironments [25]. ASHLAR is optimized for these data acquisition requirements and is more rapid and accurate than existing open-source methods we have tested as well as commercial software available with many slide scanners. Moreover, once an optimized whole-slide image mosaic has been generated, it is often convenient to visualize or analyze limited regions of the image. ASHLAR therefore supports re-tiling using adjustable block sizes and overlaps while retaining subpixel registration and without losing any information along the original tile seams.

ASHLAR reads and writes files in the OME-TIFF standard and can process images from almost all commercial microscopes using the OME Bio-Formats library [20]. This greatly streamlines the stitching and registration process since little manual intervention is required. We have successfully used ASHLAR on several hundred specimens and TMAs collected using five different microscopes in three different laboratories which have implemented three different imaging methods (CyCIF, CODEX and

mxIF). The software has also been incorporated into the MCMICRO image processing pipeline [15] used by the HTAN consortium [3] for analysis of tumor and tissue images from nearly a dozen laboratories. ASHLAR is available under the permissive MIT open source software license, making it possible for commercial companies to modify and package it with their instruments. A particular opportunity exists for vendors to couple ASHLAR with microscope control software to create closed-loop image acquisition systems.

ASHLAR is effective not only with conventional rectangular image acquisition grids, but also with images involving multiple non-overlapping regions of interest and tiles that do not form a rectangular grid. The ability to manage irregular and disconnected specimens has emerged as a key requirement in the broader application of tissue imaging. By instructing a microscope to avoid imaging empty space lying outside of the margins of the tissue, imaging time and file size can be reduced, often by a factor of two or more (a significant advantage as datasets approach terabyte scale). We have successfully used ASHLAR to assemble images of tissue microarrays (TMAs), in which 0.3 to 1 mm diameter "cores" from multiple tissue specimens are positioned in a regular array on a glass slide, making it possible to analyze up to 100 specimens in parallel. This represents a potentially challenging stitching problem since much of the slide is devoid of sample, and individual cores are often divided among multiple fields. Core biopsies and fine needle aspirations are other samples in which collection of non-rectangular images is highly advantageous. Such biopsies are typically long and thin (0.3 x 10-20 mm) and rarely aligned along the axis of the slide, making it necessary to collect tiles on a diagonal. The ability of ASHLAR to reject spurious alignments using a permutation test followed by pruning of adjacency graphs make the algorithm robust to regions of the images such as these that contain little if any data in the registration channel.

Much of the recent discussion about multiplexed tissue imaging has focused on the importance of the number of channels (assay plex) [26], since the greater the number of channels the greater the number of proteins or genes that can be assayed and the more detailed the molecular insight. However,

two other considerations are at least as important: image resolution and field of view (speed also matters for high volume applications). In the case of fluorescent imaging of a single tile, resolution and field of view are functions of the optics, primarily the numerical aperture of the objective lens, the properties of the transfer optics and the number of camera pixels – which together specify pixel size [27]. However, for large whole slide images assembled from many tiles, the accuracy of image stitching and registration also becomes critical. ASHLAR directly addresses this requirement. In most applications, it is also combined with other software to optimize the quality of image mosaics. Prior to stitching and registering tiles using ASHLAR we perform illumination correction using BaSiC [28], which exploits low-rank and sparse decomposition to correct for uneven shading and background variation in microscope images. This is essential because the illumination of each tile is typically brightest at the center of the field (along the optical axis) and dimmer at the edges. Following illumination correction using BaSiC, stitching, and registration, mosaic images are ready for inspection, segmentation and pixel-level machine learning [29]. Many of these functions are available in the MCMICRO pipeline [15], which leverages Docker/Singularity containers [30,31] and the Nextflow language [32] to link ASHLAR with upstream and downstream image processing routines.

Limitations of the approach and future extensions

To date, ASHLAR has been used to stitch several hundred whole-slide images and we have found that the spanning tree approach used to combine individual pairwise tile alignments is broadly effective. However, one recurrent weakness observed with large specimens is that tiles at the margin of the tissue that are adjacent in physical space are often found to lie far apart in the adjacency graph. When corrected positions are determined, uncorrelated error accumulates as pairwise shifts are added up along the edges in the spanning tree. The resulting stitching error is rarely noticeable by eye in the resultant mosaic image, but we have identified this as an area for future improvement of the algorithm.

To achieve reasonable processing speed, ASHLAR makes some compromises with respect to the factors it accounts for during stitching and registration. For example, ASHLAR currently performs only rectilinear correction of tile position and assumes tiles have the same magnification. When images from different microscopes must be combined it is usually necessary to account for changes in camera angle due to rotation of individual cameras and their microscope stages relative to each other. Scaling is also frequently different across instruments, even when the same objective is used, due to differences in transfer optics and sensor configurations. We have never encountered the need to assemble an image from multiple microscopes (partly because many other batch effects are introduced) but corrections for image rotation and scaling can be performed through minor additions to the alignment procedure [33]; we will add these features to ASHLAR as the need arises, most likely as multiple imaging modalities are combined for collection of different types of data.

ACKNOWLEDGEMENTS

We thank Jerry Lin, Zoltan Maliga, Connor Jacobson, Artem Sokolov and other members of the Laboratory of Systems Pharmacology for their assistance. This work was funded by NCI Human Tumor Atlas Network grants U2C-CA233262 and U2C-CA233280 and by the Ludwig Cancer Center. The normal colon specimen used for this work was acquired from the Cooperative Human Tumor Network (https://www.chtn.org/). We thank the members of the HTAN Consortium for their support; a full list can be found at humantumoratlas.org.

AUTHOR CONTRIBUTIONS

Software development and testing: JM, YC and DR. Writing and editing manuscript: JM, YC and PKS. Supervision and funding: PKS

CONFLICT STATEMENT

Peter K. Sorger is a member of the SAB or BOD of Glencoe Software (which supports a commercial version of OME software), Applied Biomath and RareCyte Inc (which built the microscopes used in this study) and has equity in these companies; he is on the SAB of NanoString Inc. In the last five years the Sorger lab has received research funding from Novartis and Merck. Sorger declares that none of these relationships altered the conduct or reporting of this research.

CODE AVAILABILITY

Ashlar is available under the MIT open source license at https://github.com/labsyspharm/ashlar.

DATA AVAILABILITY

Data shown in **Figure 4** are available as part of the EMIT data set via the NCI-recognized repository for Human Tumor Atlas Network (HTAN) data at https://www.synapse.org/#!Synapse:syn22345748. Note to reviewers: access to this data requires a (free) Synapse account. An anonymous "reviewer-only" link can be provided prior to that by requesting it from the monitoring editor. A related colorectal cancer dataset processed using ASHLAR can be viewed on-line at https://www.cycif.org/crc1-intro as can a range of other ASHLAR-processed datasets at https://www.cycif.org/software/minerva.

METHODS

Software Implementation

ASHLAR is implemented in Python 3 and utilizes many numerical and image processing routines from the numpy, scipy, scikit-image, scikit-learn and networkx packages. The pyjnius Python-to-Java connector provides access to the Java BioFormats library for reading microscopy image files.

Performance Evaluation

The runtime and RAM usage for the stitching and registration phases of ASHLAR are each comparable to that of MIST stitching running on a single CPU core. On the first two cycles of the human colon dataset, MIST-FFTW required 250 seconds (125 seconds per cycle) and 2.5 GB RAM to compute corrected tile positions. ASHLAR required 186 seconds (72 seconds to stitch cycle 1 and 114 seconds to register cycle 2) and 2.1 GB RAM. MIST does execute more quickly when allowed to use multiple CPU cores or a GPU; we have not enabled parallel processing in ASHLAR but expect a similar increase in speed. All measurements were taken on a 3.3 GHz Intel Core i5-6600 CPU with 16 GB of RAM and a Samsung 960 EVO 500 GB SSD running Ubuntu Linux 20.04. Software versions were as follows: ASHLAR 1.13.0, MIST 2.0.5, Python 3.8.5, OpenJDK 14.0.2.

Large mosaic image registration

To register the stitched MIST mosaic images prior to computing the optical flow, we first downsampled them by 10X to obtain a manageable image size followed by alignment with subpixel-precision phase correlation (phase correlation on the full-size images required a computer with more RAM than was readily available to us). The computed shift distance was approximately 6 pixels. We performed the same registration process on the ASHLAR mosaics (which are already registered) as a check and verified that the shift distance was exactly 0 pixels.

Optical Flow Computation

Because we could not find any general-purpose optical flow implementations capable of processing gigapixel-scale images on a reasonable workstation computer, we implemented our own coarse approximation using a block-based approach which is memory-efficient and highly parallelizable. The two images to be compare are broken down into non-overlapping blocks of 200 x 200 pixels and the relative shift between each corresponding pair of blocks is computed using phase correlation. This

produces a 200X-downsampled vector flow field which can then be upsampled back to the original image dimensions. The results can be seen in **Figures 5D** and **F**.

Tissue Imaging

The normal colon specimen used for this work was acquired from the Cooperative Human Tumor Network (https://www.chtn.org/) and has been judged to be "not human subjects research." Tissue-based cyclic immunofluorescence[10] was used to generate images for performance comparison. The experimental protocol documented on protocols.io (https://dx.doi.org/10.17504/protocols.io.bjiukkew) was used and 9 cycles of imaging were performed. The antibodies are listed in **Table S1**. Imaging was performed on a RareCyte CyteFinder II HT Instrument with 20x 0.75 NA objective, with four excitation and emission filters having peak and full-width at half-maximum bandpass wavelengths (in nm) of: 395/25-438/26, 485/25-522/20, 555/20-590/20, and 651/10-692/40, respectively. 2 x 2 pixel binning was used during acquisition, yielding a 4-channel tile with the size of 1280 x 1080 pixels and a pixel size of 0.65 µm per pixel. To image the specimen of size X by Y mm, a grid of 609 (29 x 21) tiles was used and each tile overlapped by ~31 pixels (20 µm) in both directions. One OME-compliant RCPNL file that contains 16-bit imaging data from all tiles was generated from each cycle, the file size of each RCPNL was ~6.58 GB.

Microscopes Tested with ASHLAR

ASHLAR has been tested with the multiple slide-scanning microscopes operating in wide-filed and confocal modes using the objective lenses listed below.

Instrument	Type	Objective	Field of View	Nominal Resolution[1]
RareCyte CyteFinder	Slide scanner	10X/0.3NA	1.66 x 1.40 mm	1.06 µm
		20X/0.75NA	0.83 x 0.7 mm	0.42 µm
		40X/0.6NA	0.42 x 0.35 mm	0.53 µm

RareCyte Orion	Slide scanner	20X/0.75NA	0.66 x 0.66 mm	0.42 µm
		40X/0.95NA	0.33 x 0.33 mm	0.33 µm
GE IN Cell Analyzer 6000	Slide scanning mode	10X/0.45NA	1.3 x 1.3 mm	0.70 µm
		20X/0.75NA	0.66 x 0.66 mm	0.42 µm
		40X/0.95NA	0.33 x 0.33 mm	0.33 µm
		60X/0.95NA	0.22 x 0.22 mm	0.33 µm
GE IN Cell Analyzer 6000	Confocal mode	60X/0.95NA	0.22 x 0.22 mm	0.21 µm
Zeiss Axio Scan.Z1	Slide scanner	10X/0.45NA	1.3 x 1.3 mm	0.70 µm
		20X/0.8NA	0.66 x 0.66 mm	0.40 µm
Zeiss Axio Observer.Z1	Slide scanner	20X/0.8NA	0.83 x 0.66 mm	0.40 µm

[1]The nominal resolution (r) was determined using the formulae: (r) = 0.61λ/NA for widefield or (r) = 0.4λ/NA for confocal microscopy (λ = 520 nm). The actual resolution depends on optical properties, the thickness of the tissue section, and both the alignment and the quality of the optical components used.

FIGURE LEGENDS

Figure 1: Schematic of cyclic whole-slide data acquisition, stitching, and registration. (A) One cycle of whole-slide imaging (scanning) is achieved by moving the microscope stage along a controlled path and acquiring multichannel image tiles that overlap. Further cycles follow the same process after re-labeling the specimens with new antibodies or other detection reagents. To integrate information across a wide spatial context at high resolution, it is necessary to stitch neighboring image tiles within one cycle and also register tiles across different cycles. **(B)** Corners of four neighboring tiles (Hoechst channel) from one cycle, false-colored and positioned using the recorded microscope stage positions (upper panel) and the corrected stitched positions (lower panel). Arrows indicate two individual cells in the tile overlap regions before and after stitching. **(C)** Centers of three Hoechst-channel image tiles from different cycles false-colored and merged using recorded stage positions (upper panel) and post-ASHLAR registered positions (lower panel). Arrows indicate one cell before and after registration.

Figure 2: ASHLAR phases for aligning whole-slide scans. (A) Steps for stitching tiles within one cycle. **(B)** Steps for registering tiles across cycles. **(C)** Seamless mosaic generation enables whole-slide visualization and flexible re-tiling for downstream parallel processing. See text for details.

Figure 3: Visualizing stitching steps using a whole-slide scan from a colon specimen. All images and data in this figure derive from analysis of a large multi-tile image of human colon (see text for details). **(A)** Alignment of one pair of neighboring tiles from an image of human colon. Images of Hoechst 33342-stained nuclei in left and right tiles are pseudocolored in red and cyan respectively. The red and cyan images are overlaid before and after stitching to demonstrate the effect at the single-cell level. For context, in the remaining panels the location of this tile pair is denoted with a yellow X. **(B)** Alignment shift distance vs. E_{NCC} for all neighboring pairs, with E_{NCC} threshold and user-provided translation limit indicated. The null distribution generated by the permutation test (red) is overlaid on the E_{NCC} marginal distribution. **(C)** Adjacency graph with edges colored by E_{NCC} overlaid on the Hoechst image. Edges corresponding to discarded alignments (E_{NCC} or shift distance above the thresholds) are hidden. Hidden edges correlate with regions containing scant or no tissue. **(D)** Minimum spanning tree with edges colored by alignment shift distance.

Figure 4: ASHLAR mosaic results. All images and data in this figure derive from analysis of the multi-tile image of human colon shown in Figure 3. **(A)** Pseudocolor image showing 5 channels from a 28-plex (9-cycle) t-CyCIF image of a normal human colon section acquired using the antibodies described in **Supplementary Table 1**. Tiles, denoted by the gray grid, overlapped by ~31 pixels (20 µm) **Inset**: Hematoxylin and eosin (H&E) staining of an adjacent section of the same specimen. **(B)** Higher magnification view of the area surrounding a single tile showing 7 channels from 4 different cycles to highlight stitching and registration accuracy. Insets 1-4 depict regions of the tile overlap areas

at full resolution (note that the antibodies shown in panels **A** and **B** differ to make structures relevant to different spatial scales more apparent).

Figure 5: Performance comparison of ASHLAR and MIST software. All images and data in this figure derive from analysis of the multi-tile image of human colon shown in Figures 3-4. **(A)** Stitching error metric distributions (kernel density estimate) for MIST and ASHLAR computed according to the stitching evaluation framework of Chalfoun et al. [16]. Dotted lines indicate median values. **(B)** Local registration error distance distributions for MIST and ASHLAR mosaic images of two t-CyCIF cycles of a human colon section. Distances at the upper end in this plot as well as in panels C and E were clipped to a maximum value of 6, which is the 75th percentile of the MIST error values, to highlight the relevant data. **(C)** Full-resolution view of four regions from the MIST mosaics demonstrating local registration error in different directions. The Hoechst images of nuclei from cycles 1 and 2 are pseudocolored red and cyan, respectively, to visualize the effect of registration error at the single-cell level. **(D)** Heatmap of MIST local registration error direction (hue) and magnitude (intensity) at 200-pixel resolution overlaid on the Hoechst image (brighter colors indicate larger errors). Characteristic tile-sized scale of heatmap features suggests inconsistent stitching. Yellow X marks indicate locations highlighted in panel C. **(E)** The same regions as in panel C, but taken from ASHLAR mosaics. An identical pseudocoloring scheme is used; the red and cyan images, now more accurately registered, combine to appear nearly white. **(F)** Heatmap of Ashlar local registration error using the same intensity and hue scale as in panel D showing overall lower error and no apparent tile-scale features. Remaining small-scale errors represent damaged tissue that could not be registered.

LITERATURE CITED

1. Regev, A. *et al.* The Human Cell Atlas. *eLife* **6**, (2017).

2. HuBMAP Consortium. The human body at cellular resolution: the NIH Human Biomolecular Atlas Program. *Nature* **574**, 187–192 (2019).
3. 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).
4. Wick, M. R. Histochemistry as a tool in morphological analysis: a historical review. *Ann. Diagn. Pathol.* **16**, 71–78 (2012).
5. Angelo, M. *et al.* Multiplexed ion beam imaging of human breast tumors. *Nat. Med.* **20**, 436–442 (2014).
6. Gerdes, M. J. *et al.* Highly multiplexed single-cell analysis of formalin-fixed, paraffin-embedded cancer tissue. *Proc. Natl. Acad. Sci. U. S. A.* **110**, 11982–11987 (2013).
7. Giesen, C. *et al.* Highly multiplexed imaging of tumor tissues with subcellular resolution by mass cytometry. *Nat. Methods* **11**, 417–422 (2014).
8. Goltsev, Y. *et al.* Deep Profiling of Mouse Splenic Architecture with CODEX Multiplexed Imaging. *Cell* **174**, 968-981.e15 (2018).
9. Gut, G., Herrmann, M. D. & Pelkmans, L. Multiplexed protein maps link subcellular organization to cellular states. *Science* **361**, (2018).
10. Lin, J.-R. *et al.* Highly multiplexed immunofluorescence imaging of human tissues and tumors using t-CyCIF and conventional optical microscopes. *eLife* **7**, e31657 (2018).
11. Tsujikawa, T. *et al.* Quantitative Multiplex Immunohistochemistry Reveals Myeloid-Inflamed Tumor-Immune Complexity Associated with Poor Prognosis. *Cell Rep.* **19**, 203–217 (2017).
12. Chen, K. H., Boettiger, A. N., Moffitt, J. R., Wang, S. & Zhuang, X. RNA imaging. Spatially resolved, highly multiplexed RNA profiling in single cells. *Science* **348**, aaa6090 (2015).
13. Lee, J. H. *et al.* Highly multiplexed subcellular RNA sequencing in situ. *Science* **343**, 1360–1363 (2014).

14. Ståhl, P. L. *et al.* Visualization and analysis of gene expression in tissue sections by spatial transcriptomics. *Science* **353**, 78–82 (2016).

15. Schapiro, D. *et al.* MCMICRO: A scalable, modular image-processing pipeline for multiplexed tissue imaging. *bioRxiv* (2021) doi:10.1101/2021.03.15.435473.

16. Chalfoun, J. *et al.* MIST: Accurate and Scalable Microscopy Image Stitching Tool with Stage Modeling and Error Minimization. *Sci. Rep.* **7**, 4988 (2017).

17. Holtkamp, D. J. & Goshtasby, A. A. Precision Registration and Mosaicking of Multicamera Images. *IEEE Trans. Geosci. Remote Sens.* **47**, 3446–3455 (2009).

18. Hörl, D. *et al.* BigStitcher: reconstructing high-resolution image datasets of cleared and expanded samples. *Nat. Methods* **16**, 870–874 (2019).

19. Schneider, C. A., Rasband, W. S. & Eliceiri, K. W. NIH Image to ImageJ: 25 years of image analysis. *Nat. Methods* **9**, 671–675 (2012).

20. Li, S. *et al.* Metadata management for high content screening in OMERO. *Methods San Diego Calif* **96**, 27–32 (2016).

21. Guizar-Sicairos, M., Thurman, S. T. & Fienup, J. R. Efficient subpixel image registration algorithms. *Opt. Lett.* **33**, 156 (2008).

22. Dean, R. T. & Dunsmuir, W. T. M. Dangers and uses of cross-correlation in analyzing time series in perception, performance, movement, and neuroscience: The importance of constructing transfer function autoregressive models. *Behav. Res. Methods* **48**, 783–802 (2016).

23. Yule, G. U. Why do we Sometimes get Nonsense-Correlations between Time-Series?--A Study in Sampling and the Nature of Time-Series. *J. R. Stat. Soc.* **89**, 1 (1926).

24. Health, C. for D. and R. Technical Performance Assessment of Digital Pathology Whole Slide Imaging Devices. *U.S. Food and Drug Administration* http://www.fda.gov/regulatory-information/search-fda-guidance-documents/technical-performance-assessment-digital-pathology-whole-slide-imaging-devices (2019).

25. Lin, J.-R. *et al. Multiplexed 3D atlas of state transitions and immune interactions in colorectal cancer*. http://biorxiv.org/lookup/doi/10.1101/2021.03.31.437984 (2021) doi:10.1101/2021.03.31.437984.
26. Baharlou, H., Canete, N. P., Cunningham, A. L., Harman, A. N. & Patrick, E. Mass Cytometry Imaging for the Study of Human Diseases—Applications and Data Analysis Strategies. *Front. Immunol.* **10**, (2019).
27. Ghiran, I. C. Introduction to fluorescence microscopy. *Methods Mol. Biol. Clifton NJ* **689**, 93–136 (2011).
28. Peng, T. *et al.* A BaSiC tool for background and shading correction of optical microscopy images. *Nat. Commun.* **8**, (2017).
29. Komura, D. & Ishikawa, S. Machine Learning Methods for Histopathological Image Analysis. *Comput. Struct. Biotechnol. J.* **16**, 34–42 (2018).
30. Merkel, D. Docker: Lightweight Linux Containers for Consistent Development and Deployment. *Linux J* **2014**, (2014).
31. Kurtzer, G. M., Sochat, V. & Bauer, M. W. Singularity: Scientific containers for mobility of compute. *PLOS ONE* **12**, e0177459 (2017).
32. Di Tommaso, P. *et al.* Nextflow enables reproducible computational workflows. *Nat. Biotechnol.* **35**, 316–319 (2017).
33. Gonzalez, R. Improving Phase Correlation for Image Registration. in (University of Auckland, 2011).

Figure 1

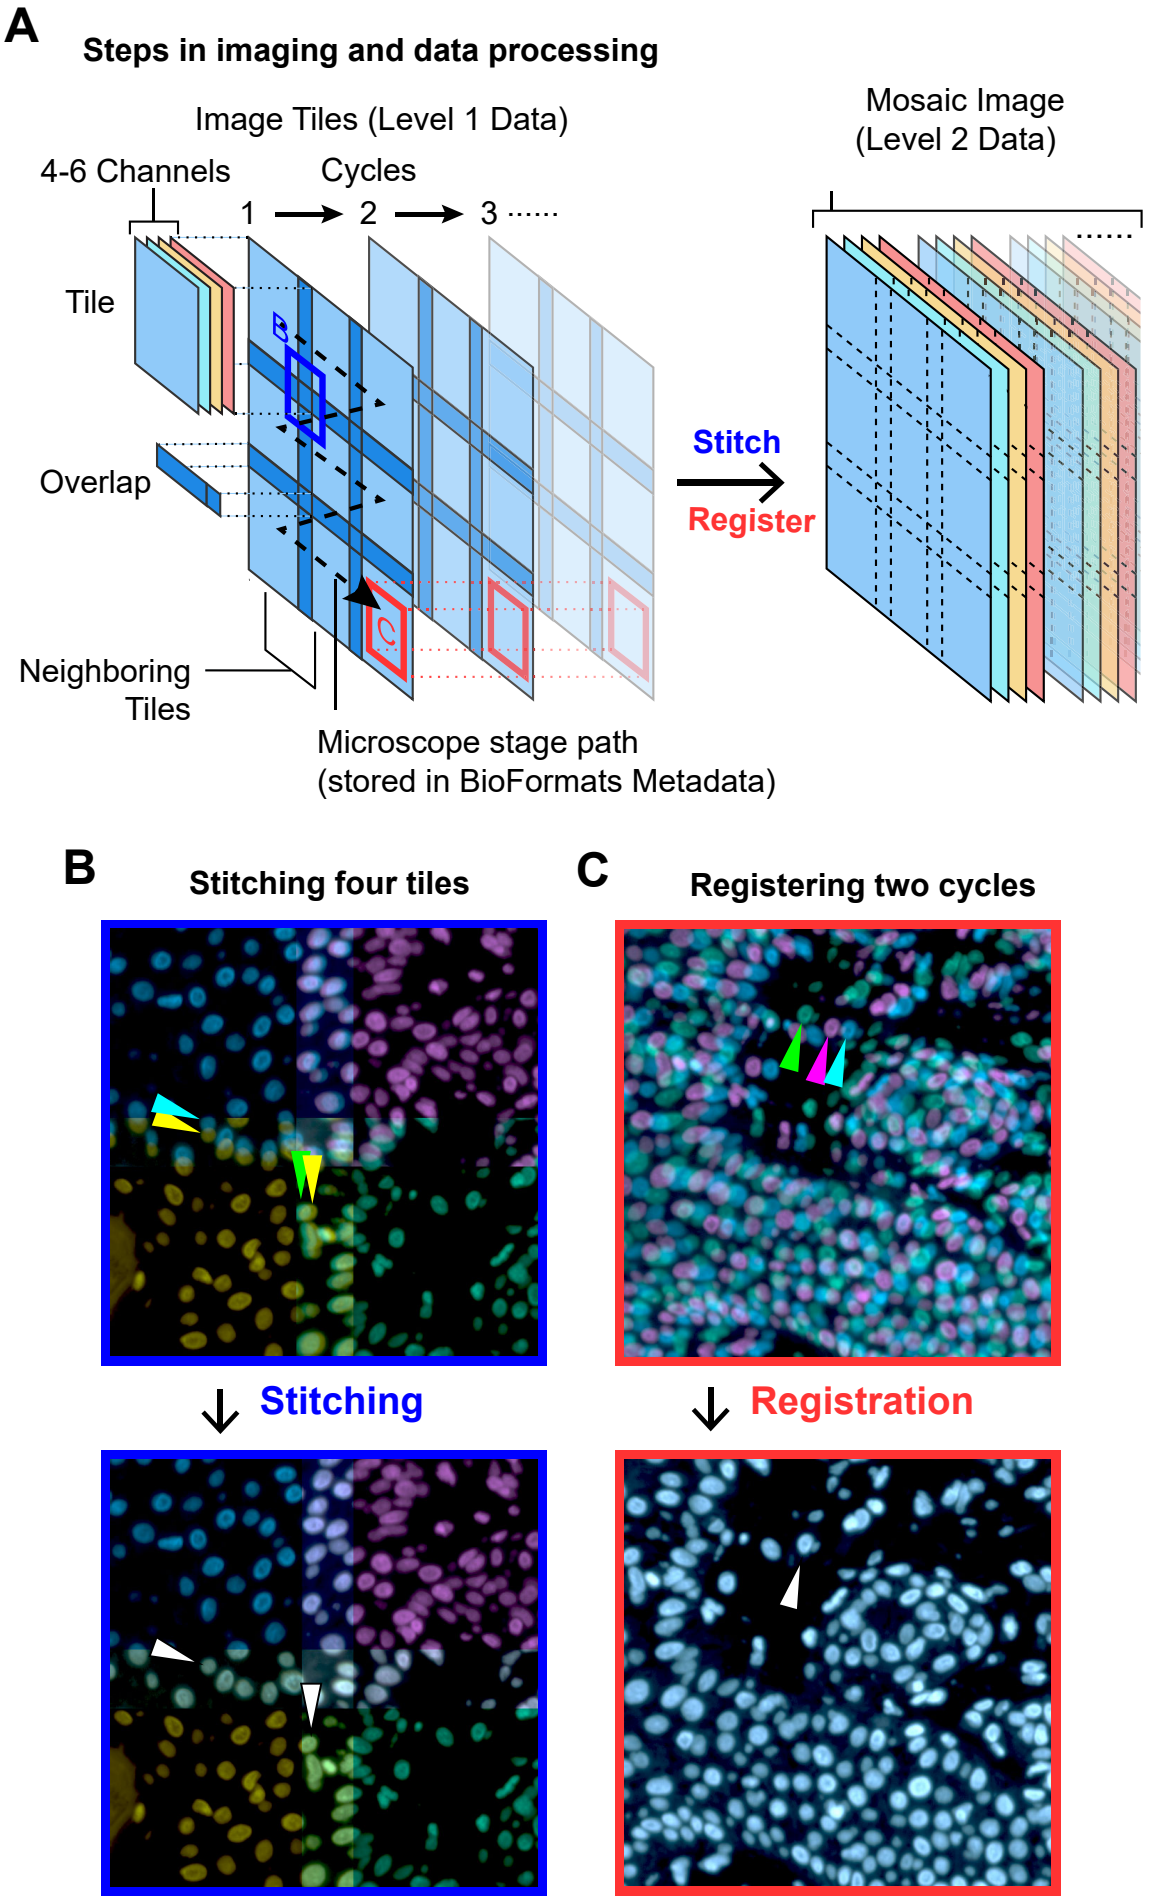
A
Steps in imaging and data processing
Image Tiles (Level 1 Data)
4-6 Channels
Cycles
1
2
3 .....
Tile
Overlap
Neighboring Tiles
Microscope stage path
(stored in BioFormats Metadata)
Stitch
Register
Mosaic Image
(Level 2 Data)
B
Stitching four tiles
Stitching
C
Registering two cycles
Registration

Figure 2

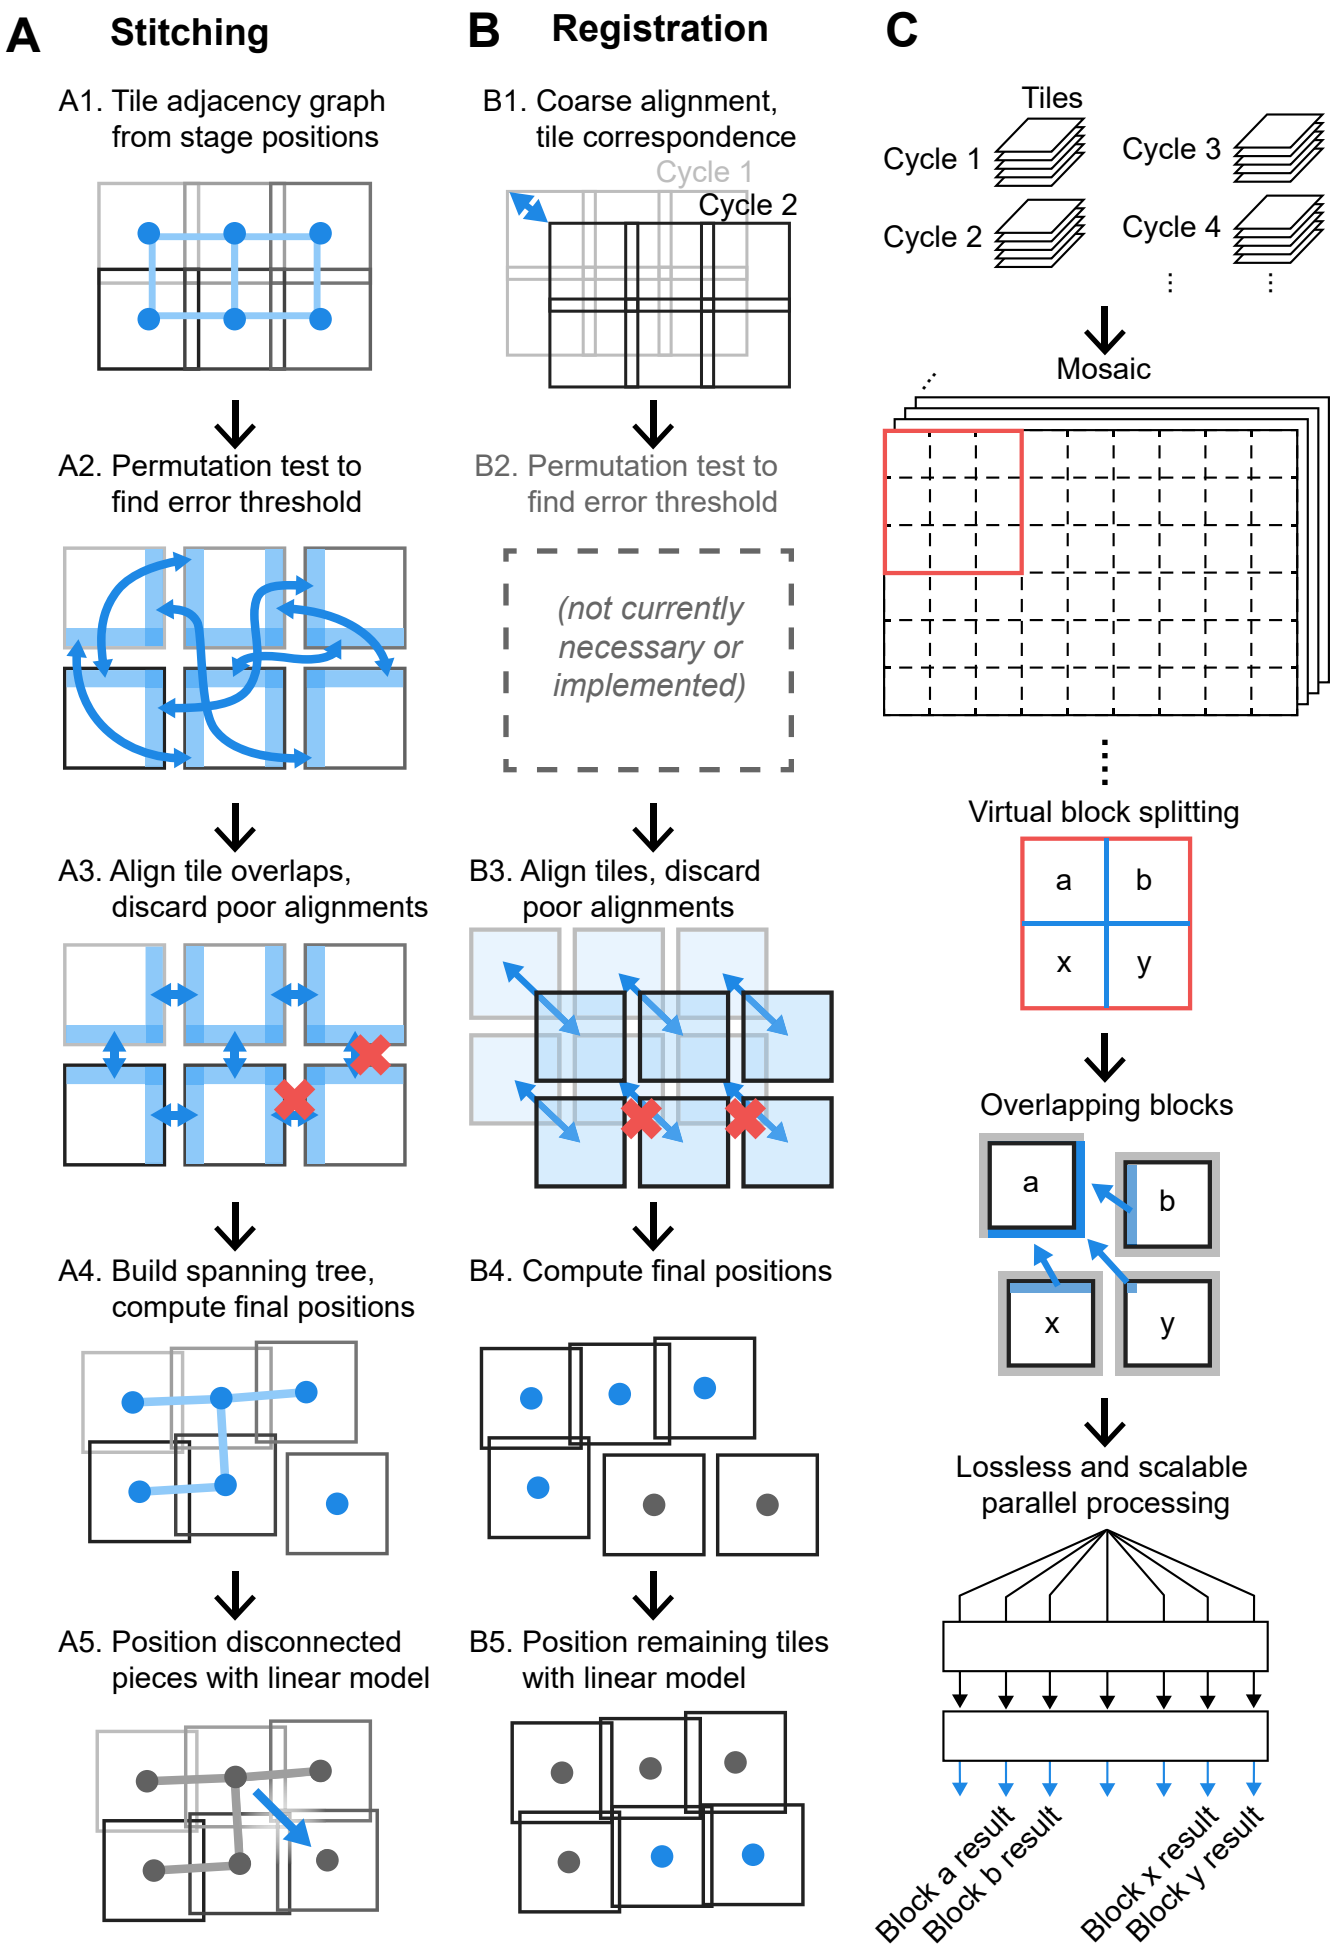
A
Stitching
A1. Tile adjacency graph from stage positions
A2. Permutation test to find error threshold
A3. Align tile overlaps, discard poor alignments
A4. Build spanning tree, compute final positions
A5. Position disconnected pieces with linear model
B
Registration
B1. Coarse alignment, tile correspondence
Cycle 1
Cycle 2
B2. Permutation test to find error threshold
(not currently necessary or implemented)
B3. Align tiles, discard poor alignments
B4. Compute final positions
B5. Position remaining tiles with linear model
C
Tiles
Cycle 1
Cycle 2
Cycle 3
Cycle 4
Mosaic
Virtual block splitting
a
b
x
y
Overlapping blocks
a
b
x
y
Lossless and scalable parallel processing
Block a result
Block b result
Block x result
Block y result

Figure 3

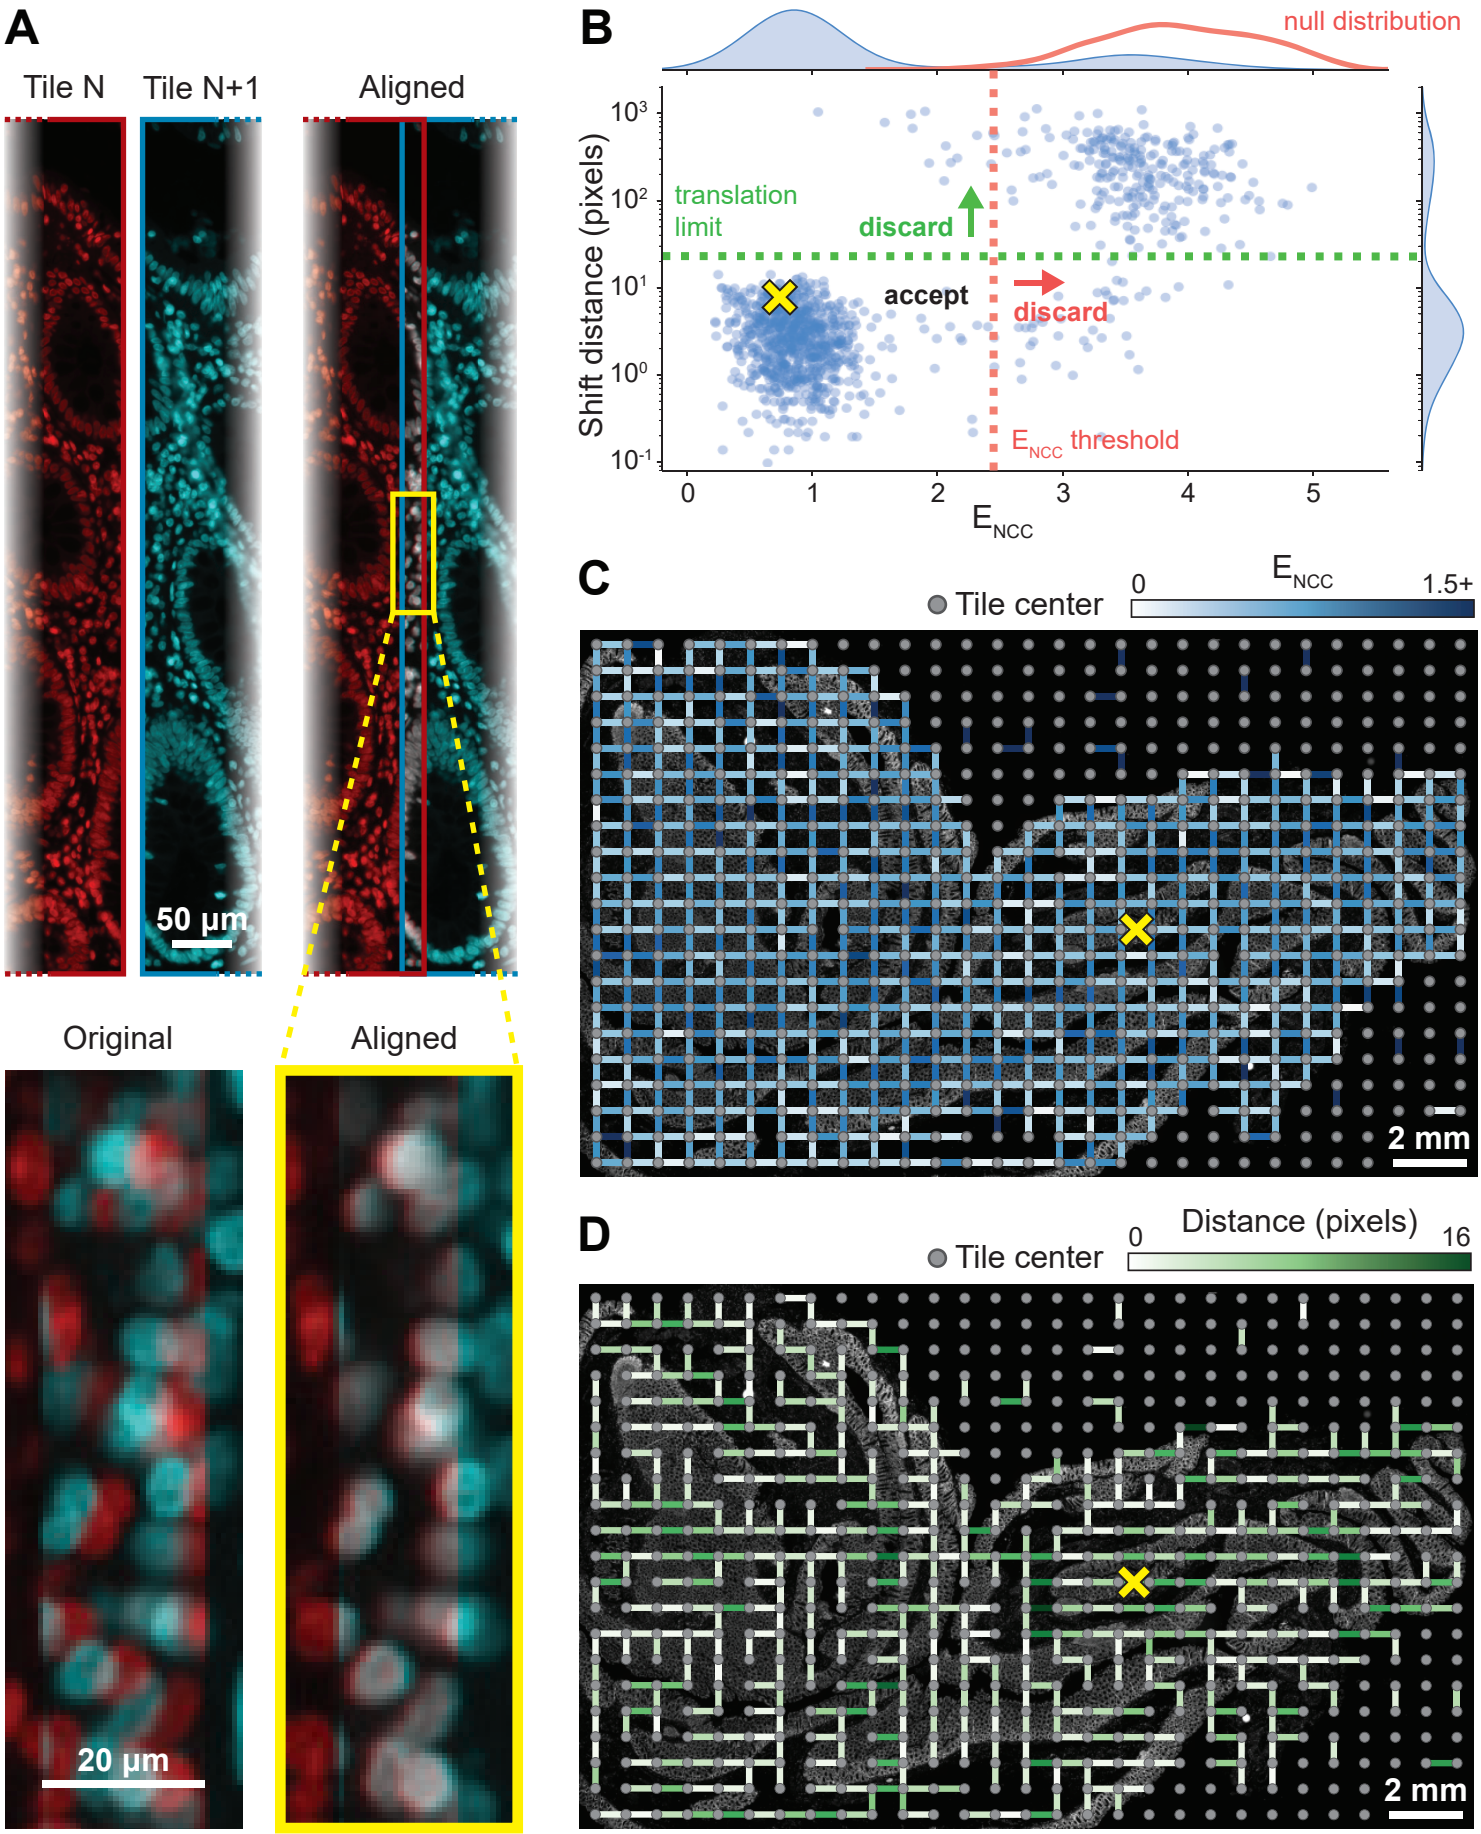
A
Tile N
Tile N+1
Aligned
50 µm
Original
Aligned
20 µm
B
null distribution
translation limit
discard
accept
discard
E_{NCC} threshold
Shift distance (pixels)
10^3
10^2
10^1
10^0
10^{-1}
0
1
2
3
4
5
E_{NCC}
C
Tile center
E_{NCC}
0
1.5+
2 mm
D
Tile center
Distance (pixels)
0
16
2 mm

Figure 4

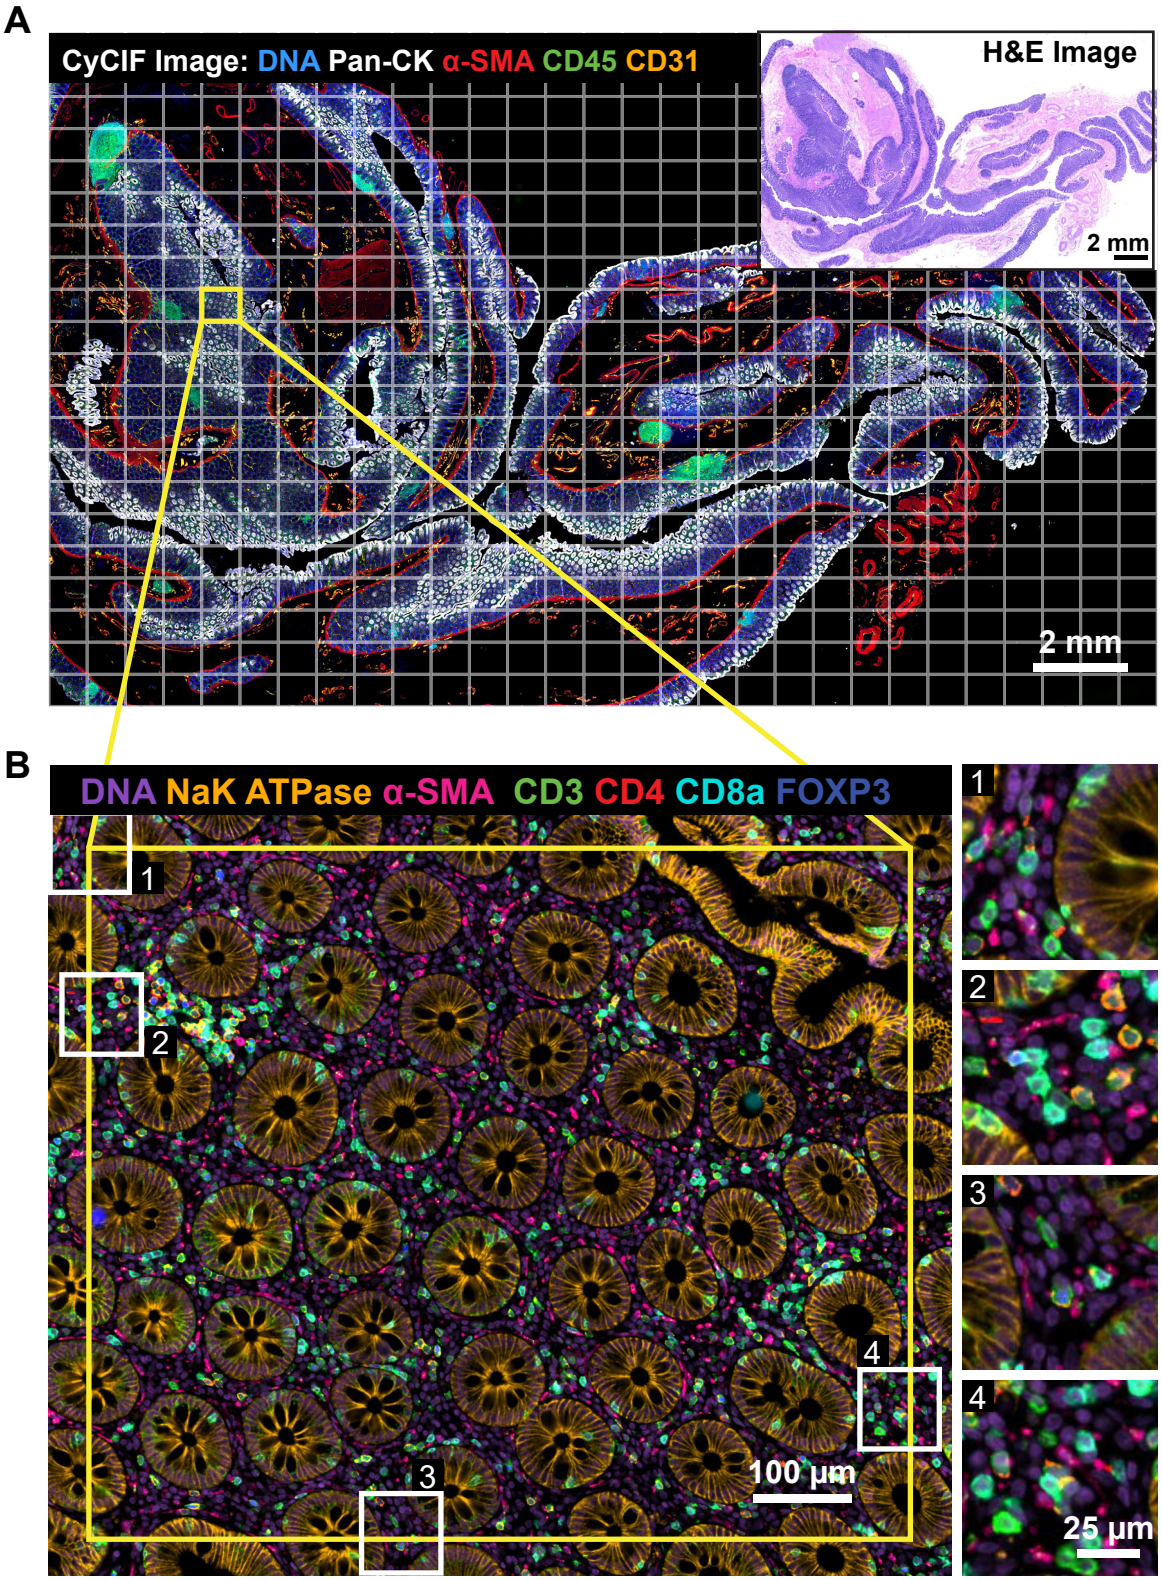
A
CyCIF Image: DNA Pan-CK α-SMA CD45 CD31
H&E Image
2 mm
2 mm
B
DNA NaK ATPase α-SMA CD3 CD4 CD8a FOXP3
1
2
3
4
100 µm
1
2
3
4
25 µm

Figure 5

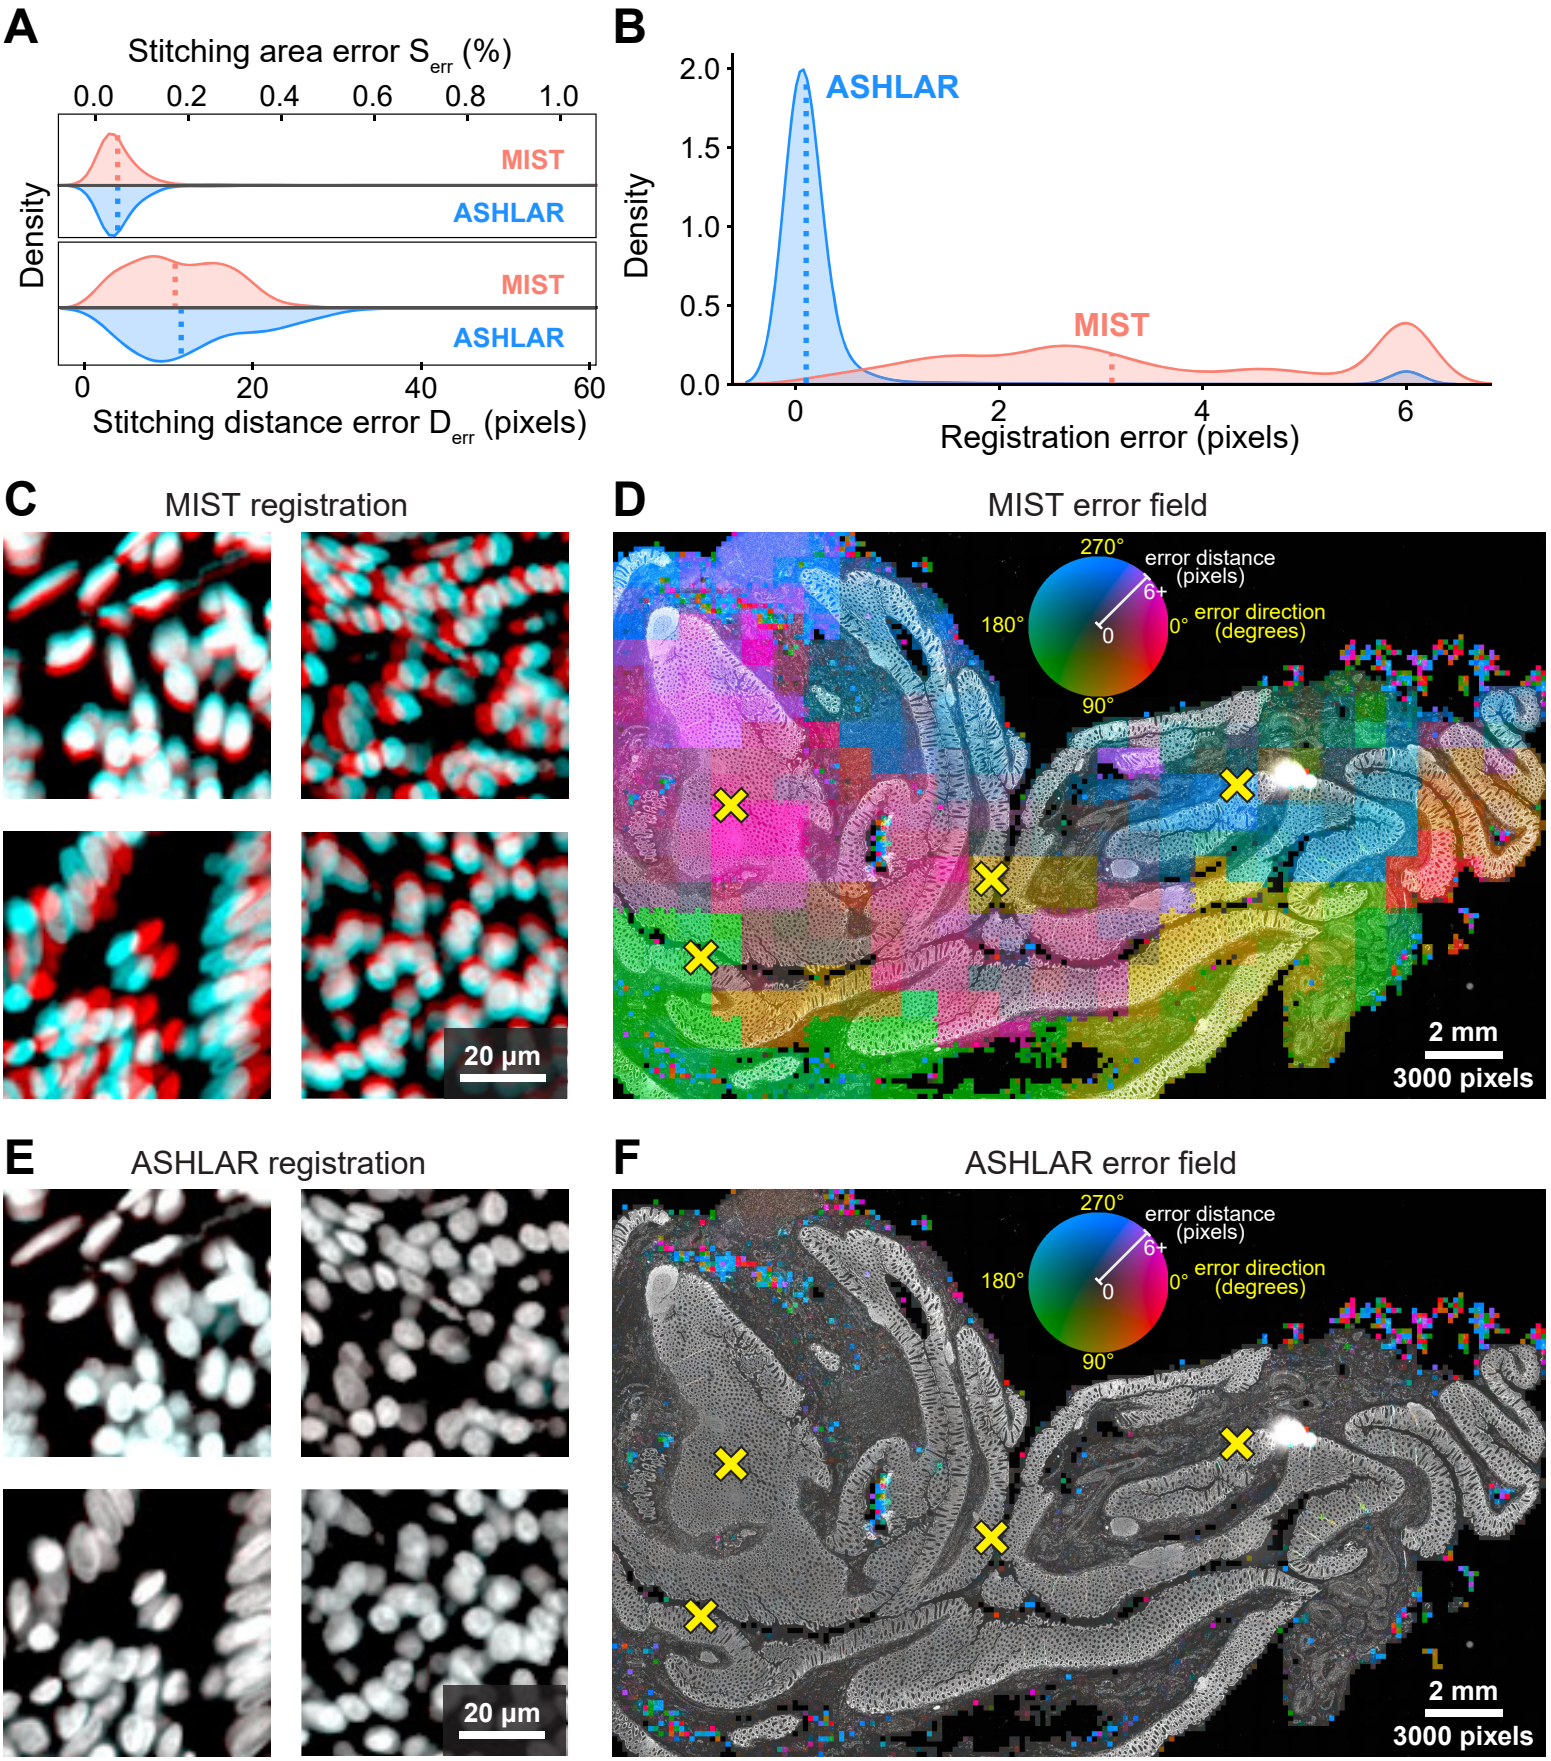
A
Stitching area error S_{err} (%)
0.0
0.2
0.4
0.6
0.8
1.0
MIST
ASHLAR
Density
MIST
ASHLAR
0
20
40
60
Stitching distance error D_{err} (pixels)
B
2.0
1.5
1.0
0.5
0.0
ASHLAR
MIST
Density
0
2
4
6
Registration error (pixels)
C
MIST registration
20 µm
D
MIST error field
270°
error distance
(pixels)
6+
180°
0
0°
error direction
(degrees)
90°
2 mm
3000 pixels
E
ASHLAR registration
20 µm
F
ASHLAR error field
270°
error distance
(pixels)
6+
180°
0
0°
error direction
(degrees)
90°
2 mm
3000 pixels