



- (51) International Patent Classification:
G06K 9/00 (2006.01) *G06T 7/00* (2017.01)
- (21) International Application Number:
PCT/US2019/069164
- (22) International Filing Date:
31 December 2019 (31.12.2019)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/787,047 31 December 2018 (31.12.2018) US
- (71) Applicant: **TEMPUS LABS, INC.** [US/US]; 600 West Chicago, Suite 510, Chicago, IL 60654 (US).
- (72) Inventors: **YIP, Stephen**; 3402 S. Oakley Ave., Chicago, IL 60608 (US). **HO, Irvin**; 600 West Chicago Ave, Chicago, IL 60654 (US). **SHA, Lingdao**; 600 West Chicago Ave, Chicago, IL 60654 (US). **OSINSKI, Boleslaw**; 600 West Chicago Ave, Chicago, IL 60654 (US).
- (74) Agent: **STEPHENS, Paul, B.** et al.; Marshall, Gerstgein & Borun LLP, 233 S. Wacker Drive, 6300 Willis Tower, Chicago, IL 60606-6357 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,

(54) Title: ARTIFICIAL INTELLIGENCE SEGMENTATION OF TISSUE IMAGES

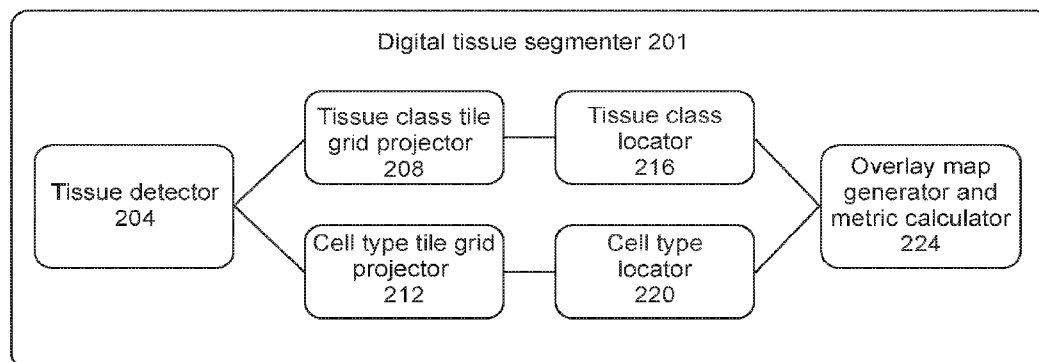


FIG. 2

(57) Abstract: Techniques for generating an overlay map on a digital medical image of a slide are provided, and include cell detection and tissue classification processes. Techniques include receiving a medical image, separating the image into tiles, and performing tile classifications and tissue classifications based on a multi-tile analysis. Techniques additionally include identifying cell objects in the image, separating the image into and displaying polygons identifying the cell objects and cell classifications. Generated displays may be overlays over the initial digital image.



MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

ARTIFICIAL INTELLIGENCE SEGMENTATION OF TISSUE IMAGES

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims benefit of priority to and claims under 35 U.S.C. §119(e)(1) the benefit of the filing date of U.S. provisional application serial number 62/787,047 filed December 31, 2018, the entire disclosure of which is incorporated herein by reference.

FIELD OF THE INVENTION

[0002] The present disclosure relates to techniques for the analysis of medical images and, more particularly, to techniques for analysis of histological slides other images of cancerous tissue.

BACKGROUND

[0003] The background description provided herein is for the purpose of generally presenting the context of the disclosure. Work of the presently named inventors, to the extent it is described in this background section, as well as aspects of the description that may not otherwise qualify as prior art at the time of filing, are neither expressly nor impliedly admitted as prior art against the present disclosure.

[0004] To guide a medical professional in diagnosis, prognosis and treatment assessment of a patient's cancer, it is common to extract and inspect tumor samples from the patient. Visual inspection can reveal growth patterns of the cancer cells in the tumor in relation to the healthy cells near them and the presence of immune cells within the tumor. Pathologists, members of a pathology team, other trained medical professionals, or other human analysts visually analyze thin slices of tumor tissue mounted on glass microscope slides to classify each region of the tissue as one of many tissue classes that are present in a tumor sample. This information aids the pathologist in determining characteristics of the cancer tumor in the patient, which can inform treatment decisions. A pathologist will often assign one or more numerical scores to a slide, based on a visual approximation. Numerical scores assigned during microscope slide analysis include tumor purity, which is the percentage of the tissue that is formed by tumor cells.

[0005] Characteristics of the tumor may include tumor grade, tumor purity, degree of invasiveness of the tumor, degree of immune infiltration into the tumor, cancer stage, and anatomic origin site of the tumor, which is especially important for diagnosing and treating a metastatic tumor. These details about the cancer can help a physician monitor the progression

of cancer within a patient and predict which anti-cancer treatments are likely to be successful in eliminating cancer cells from the patient's body.

[0006] Manually analyzing microscope slides is time consuming and requires a trained medical professional. Furthermore, because numerical scores are assigned by approximation, these scores are subjective, not quantitative.

[0007] High-resolution, digital images of microscope slides make it possible to use artificial intelligence to analyze the slides and classify the tissue components by tissue class.

[0008] A Convolutional Neural Network ("CNN") is a deep learning algorithm that analyzes digital images by assigning one class label to each input image. Slides, however, include more than one type of tissue, including the borders between neighboring tissue classes. There is a need to classify different regions as different tissue classes, in part to study the borders between neighboring tissue classes and the presence of immune cells among tumor cells. For a traditional CNN to assign multiple tissue classes to one slide image, the CNN would need to separately process each section of the image that needs a tissue class label assignment. Neighboring sections of the image overlap, so processing each section separately creates a high number of redundant calculations and is time consuming.

[0009] A Fully Convolutional Network (FCN) can analyze an image and assign classification labels to each pixel within the image, so a FCN is more useful for analyzing images that depict objects with more than one classification. A FCN generates an overlay map to show the location of each classified object in the original image. However, FCN deep learning algorithms that analyze digital slides would require training data sets of images with each pixel labeled as a tissue class, which requires too much annotation time and processing time to be practical. In digital images of slides, each edge of the image may contain more than 10,000-100,000 pixels. The full image may have at least $10,000^2$ - $100,000^2$ pixels, which forces long algorithm run times due to the intense computation required. The high number of pixels makes it infeasible to use traditional FCNs to segment digital images of slides.

SUMMARY OF THE INVENTION

[0010] In accordance with an example, a method for creating an overlay map on a digital image of a slide comprises: receiving the digital image; separating the digital image into a plurality of tiles; and identifying the majority class of tissue visible within each tile in the plurality of tiles, based on a multi-tile analysis.

[0011] In some examples, the method includes generating a digital overlay drawing of an outer edge of each cell in the image.

[0012] In some examples, the digital overlay drawing is prepared at the resolution level of an individual pixel.

[0013] In accordance with another example, a method for tissue classification of a digital image of a slide includes: receiving the digital image; generating a digital overlay drawing of a tissue region in the digital image; and displaying, in the digital overlay drawing, tiles over the tissue region and visually identifying predicted content of each tile, to generate a classification map classifying the digital image.

[0014] In some examples, the method further includes determining the predicted content for each tile using a classification model configured as a multi-resolution fully convolutional network, the multi-resolution fully convolution network configured to perform classification on digital images of different zoom levels.

[0015] In some examples, the method further includes using a classification model, for each tile, determining a classification from a plurality of classifications.

[0016] In some examples, the plurality of classifications include tumor epithelium, normal epithelium, immune, stroma, necrosis, blood, and fat.

[0017] In accordance with yet another example, a method for cell detection in a digital image of a slide includes: receiving the digital image; identifying cell objects in the digital image; generating a digital overlay drawing for the digital image; and displaying, in the digital overlay drawing, a polygon outlining each of the identified cell objects.

[0018] In some examples, the cell objects include lymphocyte cells and not lymphocyte cells.

[0019] In some examples, the cell objects comprise CD3, CD8, CD20, pancytokeratin, and smooth muscle actin.

[0020] In accordance with yet another example, a method for tissue classification of a digital image of a slide includes: receiving the digital image; segmenting the digital image into a plurality of tiles; determining a predicted class for each tile; identifying a plurality of cell objects in the digital image; determining a predicted class for each of the plurality of cell objects; and for

each of the plurality of tiles that corresponds to one of the plurality of cell objects, assigning the tile the predicted class of the corresponding cell object in place of the predicted class of the tile.

[0021] In some examples, the method further includes: storing, in a first file, for each tile, a tile position and the predicted class of the tile; and storing, in a second file, for each cell object, a polygon outlining the cell object and the predicted class of the cell object.

[0022] In some examples, the method further includes generating a digital overlay drawing for the digital image, where the digital overlay drawing is a cell mask displaying a polygon around each cell object.

[0023] In some examples, the method further includes generating a digital overlay drawing for the digital image, where the digital overlay drawing is a histology mask displaying the plurality of tiles and the predict class for each tile.

[0024] In some examples, the method further includes generating a digital overlay drawing for the digital image; and displaying, in the digital overlay drawing, the plurality of tiles and the predicted class for each tile that does not correspond to a cell object, and displaying, the plurality of cell objects and the predicted class of each cell object.

[0025] In some examples, the method further includes: generating a digital overlay drawing for the digital image, wherein the digital overlay drawing includes the digital image; and displaying the digital overlay drawing.

[0026] In some examples, the method further includes: generating a digital overlay drawing for the digital image, wherein the digital overlay drawing includes a generated version of the digital image; and displaying the digital overlay drawing.

[0027] In some examples, the method further includes: generating a plurality of digital overlay drawings for the digital image, where each digital overlay drawing corresponds to a different predicted class; and selectively displaying one of the plurality of digital overlay drawings.

[0028] In some examples, the method further includes generating a digital overlay drawing for the digital image, where the digital overlay drawing comprises percentages of predict classes corresponding to the digital image.

[0029] In some examples, the method further includes generating a digital overlay drawing for the digital image, where the digital overlay drawing comprises total counts of predict classes corresponding to the digital image.

[0030] In accordance with another example, a system includes a processor and a computer-readable memory is provided, the memory storing instructions that when executed by the processor cause the processor to perform the methods described herein.

[0031] In accordance with an example, a system includes a processor and a computer-readable memory storing instructions that when executed by the processor cause the processor to: receive the digital image; separate the digital image into a plurality of tiles; and identify the majority class of tissue visible within each tile in the plurality of tiles, based on a multi-tile analysis.

[0032] In accordance with an example, a system includes a processor and a computer-readable memory storing instructions that when executed by the processor cause the processor to: receive the digital image; generate a digital overlay drawing of a tissue region in the digital image; and display, in the digital overlay drawing, tiles over the tissue region and visually identify predicted content of each tile, to generate a classification map classifying the digital image.

[0033] In accordance with an example, a system includes a processor and a computer-readable memory storing instructions that when executed by the processor cause the processor to: receive the digital image; identify cell objects in the digital image; generate a digital overlay drawing for the digital image; and display, in the digital overlay drawing, a polygon outlining each of the identified cell objects.

[0034] In accordance with an example, a system includes a processor and a computer-readable memory storing instructions that when executed by the processor cause the processor to: receive the digital image; segment the digital image into a plurality of tiles; determine a predicted class for each tile; identify a plurality of cell objects in the digital image; determine a predicted class for each of the plurality of cell objects; and for each of the plurality of tiles that corresponds to one of the plurality of cell objects, assign the tile the predicted class of the corresponding cell object in place of the predicted class of the tile.

BRIEF DESCRIPTION OF THE DRAWINGS

[0035] The figures described below depict various aspects of the system and methods disclosed herein. It should be understood that each figure depicts an example of aspects of the present systems and methods.

[0036] FIG. 1 illustrates an example of a digital image of a histology slide, in accordance with an example.

[0037] FIG. 2 is an overview of a digital tissue segmenter, in accordance with an example.

[0038] FIG. 3A illustrates a tissue segmentation overlay map created by the digital tissue segmenter, in accordance with an example.

[0039] FIG. 3B illustrates a cell outer edge overlay map created by the digital tissue segmenter, in accordance with an example.

[0040] FIG. 3C illustrates a tumor probability overlay heatmap map created by the digital tissue segmenter, in accordance with an example.

[0041] FIG. 4 is a flowchart of a method for preparing digital images of histology slides for tissue segmentation and mapping analysis, in accordance with an example.

[0042] FIGS. 5A, 5B, and 5C each illustrate an exemplary grid overlay used for tissue segmentation on a digital image of a histology slide, in accordance with an example.

[0043] FIG. 6A illustrates the layers of an exemplary tissue segmentation algorithm, in accordance with an example.

[0044] FIG. 6B compares the layers of a known image classification algorithm with an exemplary tissue segmentation algorithm, in accordance with an example.

[0045] FIG. 6C is a visualization of the layers of an exemplary 3-dimensional input image matrix being convoluted by two exemplary 3-dimensional filter matrices, in accordance with an example.

[0046] FIG. 7 illustrates an exemplary grid overlay showing lymphocyte detections (little red dots) and the tile classifications, in accordance with an example.

[0047] FIG. 8 illustrates exemplary training set images for the cell type locator, in accordance with an example.

[0048] FIG. 9 illustrates example overlay maps identifying a programmed death ligand 1 (PD-L1) tumor positive score (TPS) using different immunohistochemical (IHC) stain images, in accordance with an example.

[0049] FIG. 10 a schematic illustration of an example computer processing system having a digital tissue image segmenter for performing techniques and methods herein, in accordance with an example.

DETAILED DESCRIPTION OF THE INVENTION

[0050] FIG. 1 illustrates an example of a digital image of a histology slide, also known as a pathology slide. In one example, the digital image is created by a scanner that visually captures a histology slide. In an alternative example, the digital image is created by a digital camera attached to a microscope. The histology slide may be made of two transparent glass layers with a slice of tissue affixed between the two layers of glass. The thin slice of tissue may be very thin, with a thickness, for instance, of approximately 5 microns. Tissue may be preserved in a fixative, including formaldehyde, formalin, and paraffin. The tissue contains a combination of many individual biological cells that are visible on the slide. The scanner may include a Philips digital pathology slide scanner, or any scanner known in the art that can create a digital image file.

[0051] In one example, the tissue slice contains stain that attaches to certain types of cells or cell parts within the tissue. The stain may include hematoxylin and eosin (H & E) stain and any immunohistochemical (IHC) stain. Hematoxylin is a stain that will bind to DNA and cause the nucleus of a cell to appear blue or purple. Eosin is a stain that will bind to proteins and cause all of the remaining parts of the cell, namely the cytoplasm interior, to appear red or pink. An IHC stain is comprised of an antibody coupled with a molecule that displays one of many colors. The antibody may be designed to bind to any surface shape to target a specific molecule such as a protein or a sugar. The IHC stain will result in a concentration of dye of the selected color near any copies of a specific target molecule present on the slide. Some commonly monitored proteins in tumor samples include programmed death ligand 1 (PD-L1), whose presence in a tumor region can indicate whether a tumor will respond to immunotherapy, and cluster of differentiation 3 (CD3), which is associated with T lymphocyte immune cells. The presence of

CD3 in a tumor region may be associated with tumor infiltrating lymphocytes which can indicate that the tumor will be susceptible to anti-cancer immunotherapy.

[0052] In some examples, the slide may also contain additional control slices of tissue that are not from the tumor sample, which serve as a positive and/or negative control for the staining process. Control tissue slices are more common on slides that have IHC staining.

[0053] FIG. 2 is an overview of a digital tissue image segmenter **201**. The digital tissue segmenter **201** may comprise a computational method and apparatus that receives a digital image of a slide, displaying a slice of a tumor sample, and creates a high-density, grid-based digital overlay map that identifies the majority class of tissue visible within each grid tile in the digital image. The digital tissue segmenter **201** may also generate a digital overlay drawing of the outer edge of each cell in the slide image, at the resolution level of an individual pixel.

[0054] In another example, the digital tissue segmenter **201** is a computational method and apparatus that receives a digital radiology image and efficiently creates a high-density, grid-based digital overlay map that identifies the majority class of tissue visible within each grid tile in the digital image. The radiology image may depict a tumor within a patient's body. The radiology image may be 3-dimensional (3-D) and the digital tissue segmenter **201** may receive 2-dimensional slices of the 3-dimensional image as an input image. Radiology images include but are not limited to X-rays, CT scans, MRI's, ultrasounds, and PET.

[0055] The digital tissue segmenter **201** shown at FIG. 2 includes a tissue detector **204** for detecting the areas of a digital image that have tissue, and storing data that includes the locations of the areas detected to have tissue. The tissue detector **204** transfers tissue area location data to a tissue class tile grid projector **208** and a cell tile grid projector **212**. The tissue class tile grid projector **208** receives the tissue area location data, as described in further detail below and with reference to FIGS. 5A and 5C. For each of several tissue class labels, the tissue class locator **216** calculates a percentage that represents the likelihood that the tissue class label accurately describes the image within each tile to determine where each tissue class is located in the digital image. For each tile, the total of all of the percentages calculated for all tissue class labels will sum to 1, which reflects 100%. In one example, the tissue class locator **216** assigns one tissue class label to each tile to determine where each tissue class is located in the digital image. The tissue class locator stores the calculated percentages and assigned tissue class labels associated with each tile.

[0056] Examples of tissue classes include but are not limited to tumor, stroma, normal, immune cluster, necrosis, hyperplasia/dysplasia, red blood cells, and tissue classes or cell types that are positive (contain a target molecule of an IHC stain) or negative for an IHC stain target molecule (do not contain that molecule). Examples also include tumor positive, tumor negative, lymphocyte positive, and lymphocyte negative. The grid-based digital overlay map or a separate digital overlay may also highlight individual immune cells, including lymphocytes, cytotoxic T cells, B cells, NK cells, macrophages, etc.

[0057] In one example, the digital tissue segmenter **201** includes a multi-tile algorithm that concurrently analyzes many tiles in an image, both individually and in conjunction with the portion of the image that surrounds each tile. The multi-tile algorithm may achieve a multiscale, multiresolution analysis that captures both the contents of the individual tile and the context of the portion of the image that surrounds the tile. The multi-tile algorithm is described further with reference to FIGS. 5A-5C and 6A-6B. Because the portions of the image that surround two neighboring tiles overlap, analyzing many tiles and their surroundings concurrently instead of separately analyzing each tile with its surroundings reduces computational redundancy and results in greater processing efficiency.

[0058] In one example, the digital tissue segmenter may store the analysis results in a 3-dimensional probability data array, which contains one 1-dimensional data vector for each analyzed tile. In one example, each data vector contains a list of percentages that sum to 100%, each indicating the probability that each grid tile contains one of the tissue classes analyzed. The position of each data vector in the orthogonal 2-dimensional plane of the data array, with respect to the other vectors, corresponds with the position of the tile associated with that data vector in the digital image, with respect to the other tiles.

[0059] The cell type tile grid projector receives the tissue area location data and projects a cell type tile grid onto the areas of an image with tissue, as described with further detail with respect to FIG. 8. The cell type locator may detect each biological cell in the digital image within each grid, prepare an outline on the outer edge of each cell, and classify each cell by cell type. The cell type locator stores data including the location of each cell and each pixel that contains a cell outer edge, and the cell type label assigned to each cell.

[0060] The overlay map generator and metric calculator may retrieve the stored 3-dimensional probability data array from the tissue class locator, and convert it into an overlay

map that displays the assigned tissue class label for each tile. The assigned tissue class for each tile may be displayed as a transparent color that is unique for each tissue class. In one example, the tissue class overlay map displays the probabilities for each grid tile for the tissue class selected by the user. The overlay map generator and metric calculator also retrieves the stored cell location and type data from the cell type locator, and calculates metrics related to the number of cells in the entire image or in the tiles assigned to a specific tissue class.

[0061] FIGS. 3A and 3B illustrate examples of a digital overlay created by the digital tissue segmenter **201**. FIG. 3A illustrates a tissue class overlay map created by the overlay map generator of the digital tissue segmenter **201**. FIG. 3B illustrates a cell outer edge overlay map created by the overlay map generator of the digital tissue segmenter **201**. The overlay map generator may display the digital overlays as transparent or opaque layers that cover the slide image, aligned such that the slide location shown in the overlay and the slide image are in the same location on the display. The overlay map may have varying degrees of transparency. The degree of transparency may be adjustable by the user. The overlay map generator may report the percentage of the labeled tiles that are associated with each tissue class label, ratios of the number of tiles classified under each tissue class, the total area of all grid tiles classified as a single tissue class, and ratios of the areas of tiles classified under each tissue class. FIG. 3C illustrates a tumor probability overlay heatmap map created by the digital tissue segmenter, in accordance with another example.

[0062] The digital tissue segmenter **201** may also report the total number of cells or the percentage of cells that are located in an area defined by either a user, the entire slide, a single grid tile, by all grid tiles classified under each tissue class, or cells that are classified as immune cells. The digital tissue segmenter **201** may also report the number of cells classified as immune cells that are located within areas classified as tumor or any other tissue class.

[0063] In one example, the digital tissue segmenter **201** is capable of calculating the percentage of cells that are colored by an IHC stain to highlight particular cells containing the molecule targeted by the stain. The percentage of cells may be specific to a tissue class region or a cell type (or cell object). For example, if the IHC stain targets programmed death ligand 1 (PD-L1) protein, the digital tissue segmenter **201** may determine the percentage of cancer cells in the tumor tissue class that contain PD-L1 protein. If the IHC stain targets cluster of differentiation 3 (CD3) protein, the digital tissue segmenter **201** may determine the percentage of lymphocytes or total cells that contain CD3. In other examples, the digital tissue segmenter

201 may determine the percentage of CD8, CD20, pancytokeratin, and/or smooth muscle actin in the digital image.

[0064] The map generator and metric calculator **224** may also create a digital overlay map, showing predicted IHC staining on a digital image of a slide that contains no IHC stain. In one example, the tissue class locator **216** can predict where IHC staining for a specific molecule would exist on a slide, or the percentage of cells that express a specific protein, based on input images that only contain H and E stain.

[0065] The digital overlays and reports generated by the digital tissue segmenter **201** can be used to assist medical professionals in more accurately estimating tumor purity, and in locating regions or diagnoses of interest, including invasive tumors having tumor cells that protrude into the non-tumor tissue region that surrounds the tumor. They can also assist medical professionals in prescribing treatments. For example, the number of lymphocytes in areas classified as tumor may predict whether immunotherapy will be successful in treating a patient's cancer.

[0066] The digital overlays and reports generated by the digital tissue segmenter **201** can also be used to determine whether the slide sample has enough high-quality tissue for successful genetic sequence analysis of the tissue. Genetic sequence analysis of the tissue on a slide is likely to be successful if the slide contains an amount of tissue and/or has a tumor purity value that exceeds a user-defined tissue amount and tumor purity thresholds. In one example, the digital tissue segmenter **201** may label a slide as accepted or rejected for sequence analysis, depending on the amount of tissue present on the slide and the tumor purity of the tissue on the slide. In some examples, the digital tissue segmenter **201** is configured to determine tumor-infiltrating lymphocytes (TILS). An example of a TILS process and engine is disclosed, for example, in U.S. Provisional Patent Application No. 62/889,521, titled "Determining Therapeutic Tumor-Infiltrating Lymphocytes (TILS) from Histopathology Slide Images," filed on August 20, 2019, which is incorporated herein by reference and in its entirety for all purposes.

[0067] The digital tissue segmenter **201** may also label a slide as uncertain, to recommend that it be manually reviewed by a trained analyst, who may be a member of a pathology team. In one example, if the amount of tissue present on the slide is approximately equal to the user-defined tissue amount threshold or within a user-defined range, the digital tissue segmenter **201**

may label the slide as uncertain. In one example, if the tumor purity of the tissue present on the slide is approximately equal to the user-defined tumor purity threshold or within a user-defined range, the digital tissue segmenter **201** may label the slide as uncertain.

[0068] In one example, the overlay map generator and metric calculator calculates the amount of tissue on a slide by measuring the total area covered by the tissue or by counting the number of cells on the slide. The number of cells on the slide may be determined by the number of cell nuclei visible on the slide. In one example, the digital tissue segmenter **201** calculates the proportion of tissue that is cancer cells by dividing the number of cell nuclei within grid areas that are labeled tumor by the total number of cell nuclei on the slide. The digital tissue segmenter **201** may exclude cell nuclei or outer edges of cells that are located in tumor areas but which belong to cells that are characterized as lymphocytes. The proportion of tissue that is cancer cells is known as the tumor purity of the sample. In some examples, the overlay map generator and metric calculator is configured to predict aligned tumor purity, for example, by calculating the fraction of tumor cells within a micro-dissected tissue area instead of within the whole tissue. These techniques herein may be configured into other tumor purity models based incorporated one or multiple definitions of tumor purity definition from pathologists.

[0069] In one example, the digital tissue segmenter **201** compares the tumor purity to the user-selected minimum tumor purity threshold and the number of cells in the digital image to a user-selected minimum cell threshold and approves the slide if both thresholds are exceeded. In one example, the user-selected minimum tumor purity threshold is 0.20, which is 20%.

[0070] In one example, the slide is given a composite tissue amount score that multiplies the total area covered by tissue detected on the slide by a first multiplier value, multiplies the number of cells counted on the slide by a second multiplier value, and sums the products of these multiplications.

[0071] The digital tissue segmenter **201** may calculate whether the grid areas that are labeled tumor are spatially consolidated or dispersed among non-tumor grid areas. If the digital tissue segmenter **201** determines that the tumor areas are spatially consolidated, the digital tissue segmenter **201** may produce a digital overlay of a recommended cutting boundary that separates the slide regions classified as tumor and the slide regions classified as non-tumor or within the areas classified as non-tumor, proximal to the areas classified as tumor. This recommended cutting boundary can be a guide to assist a technician in dissecting a slide to

isolate a maximum amount of tumor or non-tumor tissue from the slide, especially for genetic sequence analysis.

[0072] The digital tissue segmenter **201** may include clustering algorithms that calculate and report information about the spacing and density of type classified cells, tissue class classified tiles, or visually detectable features on the slide. For example, this may be accomplished by choosing a clustering algorithm and a set of objects to cluster, and the output can either be a qualitative description or a numerical value. An example clustering algorithm for cells would find spatial information of cells, such as, the center locations of tumor cells or lymphocytes, the size of each cell group, and the distance between the center of each cell group, where such spatial information is used for further outcome related analysis. The spacing information includes distribution patterns and heat maps for immune cells, tumor cells, or other cells. These patterns may include clustered, dispersed, dense, and non-existent. This information is useful to determine whether immune cells and tumor cells cluster together and what percentage of the cluster areas overlap, which may facilitate in predicting immune infiltration and patient response to immunotherapy.

[0073] The digital tissue segmenter **201** may also calculate and report average tumor cell roundness, average tumor cell perimeter length, and average tumor nuclei density.

[0074] The spacing information also includes mixture levels of tumor cells and immune cells. The clustering algorithms can calculate the probability that two adjacent cells on a given slide will be either two tumor cells, two immune cells, or one tumor cell and one immune cell.

[0075] The clustering algorithms can also measure the thickness of any stroma pattern located around an area classified as tumor. The thickness of this stroma surrounding the tumor region may be a predictor of a patient's response to treatment.

[0076] The digital tissue segmenter **201** may also calculate and report statistics including mean, standard deviation, sum, etc. for the following information in each grid tile of either a single slide image or aggregated from many slide images: red green blue (RGB) value, optical density, hue, saturation, grayscale, and stain deconvolution. Deconvolution includes the removal of the visual signal created by any individual stain or combination of stains, including hematoxylin, eosin, or IHC staining.

[0077] The digital tissue segmenter **201** may also incorporate known mathematical formulae from the fields of physics and image analysis to calculate visually detectable basic features for each grid tile. Visually detectable basic features, including lines, patterns of alternating brightness, and outlineable shapes, may be combined to create visually detectable complex features including cell size, cell roundness, cell shape, and staining patterns referred to as texture features.

[0078] The digital overlays, reports, statistics and estimates produced by the digital tissue segmenter **201** may be useful for predicting patient survival, PD-L1 status of a tumor or immune cluster, microsatellite instability, TILS%, and the origin of a tumor when the origin is unknown or the tumor is metastatic. Overlays of such status information can provide an illustrated estimate of how well a patient will respond to certain immunotherapies, in some examples. The overlay reports are not inherently biased towards any class of therapy, and may be used for indicating cell-based, vaccine, and cytokine therapies, for example. The visualized reports may be adjusted to better display tissue type, cell type, TILS%, etc. for the purpose of identifying certain immunotherapies. For example, the techniques herein can be adapted with requisite training data for specific cell- and tumor-type identification. For example, various models can be chosen to predict subtype of tumor (1 a model can predict DCIS, a subtype of tumor in breast cancer which is associated with higher risk of developing invasive breast cancer and 2 a model can predict squamous cell carcinoma vs adenocarcinoma in lung cancer). Another example is the model can be adjusted to predict subtypes of cells in lymphatic tissues, such as lymphocytes vs plasma cells.

[0079] The digital tissue segmenter **201** may calculate relative densities of each type of immune cell on an entire slide, in the areas designated as tumor or another tissue class. Immune tissue classes include lymphocytes, cytotoxic T cells, B cells, NK cells, macrophages, etc.

[0080] In one example, the act of scanning or otherwise digitally capturing a histology slide automatically triggers the digital tissue segmenter **201** to analyze the digital image of that histology slide.

[0081] In one example, the digital tissue segmenter **201** allows a user to edit a cell outer edge or a border between two tissue classes on a tissue class overlay map or a cell outer edge overlay map and saves the altered map as a new overlay.

[0082] FIG. 4 is a flowchart of a method for preparing digital images of histology slides for tissue classification and mapping analysis.

[0083] In one example, each digital image file received by the digital tissue segmenter **201** contains multiple versions of the same image content, and each version has a different resolution. The file stores these copies in stacked layers, arranged by resolution such that the highest resolution image containing the greatest number of bytes is the bottom layer. This is known as a pyramidal structure. In one example, the highest resolution image is the highest resolution achievable by the scanner or camera that created the digital image file.

[0084] In one example, each digital image file also contains metadata that indicates the resolution of each layer. The digital tissue segmenter **201** can detect the resolution of each layer in this metadata and compare it to user-selected resolution criteria to select a layer with optimal resolution for analysis. In one example, the optimal resolution is 1 pixel per micron (downsampled by 4).

[0085] In one example, the digital tissue segmenter **201** receives a Tagged Image File Format (TIFF) file with a bottom layer resolution of four pixels per micron. This resolution of 4 pixels per micron corresponds to the resolution achieved by a microscope objective lens with a magnification power of "40x". In one example, the area that may have tissue on the slide is up to 100,000 x 100,000 pixels in size.

[0086] In one example, the TIFF file has approximately 10 layers, and the resolution of each layer is half as high as the resolution of the layer below it. If the higher resolution layer had a resolution of four pixels per micron, the layer above it will have two pixels per micron. The area represented by one pixel in the upper layer will be the size of the area represented by four pixels in the lower layer, meaning that the length of each side of the area represented by one upper layer pixel will be twice the length of each side of the area represented by one lower layer pixel.

[0087] Each layer may be a 2x downsampling of the layer below it. Downsampling is a method by which a new version of an original image can be created with a lower resolution value than the original image. There are many methods known in the art for downsampling, including nearest-neighbor, bilinear, hermite, bell, Mitchell, bicubic, and Lanczos resampling.

[0088] In one example, 2x downsampling means that the red green blue (RGB) values from three of four pixels that are located in a square in the higher resolution layer are replaced by the RGB value from the fourth pixel to create a new, larger pixel in the layer above, which occupies the same space as the four averaged pixels.

[0089] In one example, the digital image file does not contain a layer or an image with the optimal resolution. In this case, the digital tissue segmenter **201** can receive an image from the file having a resolution that is higher than the optimal resolution and downsample the image at a ratio that achieves the optimal resolution.

[0090] In one example, the optimal resolution is 2 pixels per micron, or “20x” magnification, but the bottom layer of a TIFF file is 4 pixels per micron and each layer is downsampled 4x compared to the layer below it. In this case, the TIFF file has one layer at 40x and the next layer at 10x magnification, but does not have a layer at 20x magnification. In this example, the digital tissue segmenter **201** reads the metadata and compares the resolution of each layer to the optimal resolution and does not find a layer with the optimal resolution. Instead, the digital tissue segmenter **201** retrieves the 40x magnification layer, then downsamples the image in that layer at a 2x downsampling ratio to create an image with the optimal resolution of 20x magnification.

[0091] After the digital tissue segmenter **201** obtains an image with an optimal resolution, it transmits the image to the tissue detector **204**, which locates all parts of the image that depict tumor sample tissue and digitally eliminates debris, pen marks, and other non-tissue objects.

[0092] In one example, the tissue detector **204** differentiates between tissue and non-tissue regions of the image and uses gaussian blur removal to edit pixels with non-tissue objects. In one example, any control tissue on a slide that is not part of the tumor sample tissue can be detected and labeled as control tissue by the tissue detector or manually labeled by a human analyst as control tissue that should be excluded from the downstream tile grid projections.

[0093] Non-tissue objects include artifacts, markings, and debris in the image. Debris includes keratin, severely compressed or smashed tissue that cannot be visually analyzed, and any objects that were not collected with the sample.

[0094] In one example, a slide image contains marker ink or other writing that the tissue detector **204** detects and digitally deletes. Marker ink or other writing may be transparent over the tissue, meaning that the tissue on the slide may be visible through the ink. Because the ink

of each marking is one color, the ink causes a consistent shift in the RGB values of the pixels that contain stained tissue underneath the ink compared to pixels that contain stained tissue without ink.

[0095] In one example, the tissue detector **204** locates portions of the slide image that have ink by detecting portions that have RGB values that are different from the RGB values of the rest of the slide image, where the difference between the RGB values from the two portions is consistent. Then, the tissue detector may subtract the difference between the RGB values of the pixels in the ink portions and the pixels in the non-ink portions from the RGB values of the pixels in the ink portions to digitally delete the ink.

[0096] In one example, the tissue detector **204** eliminates pixels in the image that have low local variability. These pixels represent artifacts, markings, or blurred areas caused by the tissue slice being out of focus, an air bubble being trapped between the two glass layers of the slide, or pen marks on the slide.

[0097] In one example, the tissue detector **204** removes these pixels by converting the image to a grayscale image, passing the grayscale image through a gaussian blur filter that mathematically adjusts the original grayscale value of each pixel to a blurred grayscale value to create a blurred image. Other filters may be used to blur the image. Then, for each pixel, the tissue detector **204** subtracts the blurred grayscale value from the original grayscale value to create a difference grayscale value. In one example, if a difference grayscale value of a pixel is less than a user-defined threshold, it may indicate that the blur filter did not significantly alter the original grayscale value and the pixel in the original image was located in a blurred region. The difference grayscale values may be compared to a threshold to create a binary mask that indicates where the blurred regions are that may be designated as non-tissue regions. A mask may be a copy of an image, where the colors, RGB values, or other values in the pixels are adjusted to show the presence or absence of an object of a certain type to show the location of all objects of that type. For example, the binary mask may be generated by setting the binary value of each pixel to 0 if the pixel has a difference grayscale value less than a user-defined blur threshold and setting the binary value of each pixel to 1 if the pixel has a difference grayscale value higher than or equal to a user-defined blur threshold. The regions of the binary mask that have pixel binary values of 0 indicate blurred areas in the original image that may be designated as non-tissue.

[0098] The tissue detector **204** may also mute or remove extreme brightness or darkness in the image. In one example, the tissue detector **204** converts the input image to a grayscale image, and each pixel receives a numerical value according to how bright the pixel is. In one example, the grayscale values range from 0 to 255, where 0 represents black and 255 represents white. In pixels with a grayscale value above a brightness threshold value, the tissue detector will replace the grayscale value of those pixels with the brightness threshold value. For pixels with a grayscale value below a darkness threshold value, the tissue detector will replace the grayscale value of those pixels equal with the darkness threshold value. In one example, the brightness threshold value is approximately 210. In one example, the darkness threshold value is approximately 45. The tissue detector stores the image with the new grayscale values in a data file.

[0100] In one example, the tissue detector **204** analyzes the altered image for any artifacts, debris, or markings that remain after the first analysis. The tissue detector scans the image and categorizes any remaining groups of pixels with a certain color, size, or smoothness as non-tissue.

[0101] In one example, the slide has H and E staining and most tissue in the slide image will have a pink stain. In this example, the tissue detector **204** categorizes all objects without any pink or red hue, as determined by the RGB value of the pixels that represent the object, as non-tissue. The tissue detector **204** may interpret any color or the lack of any color in a pixel to indicate the presence or absence of tissue in that pixel.

[0102] In one example, the tissue detector **204** detects the contours of each object in the image in order to measure the size and smoothness of each object. Pixels that are very dark may be debris, and pixels that are very bright may be background, which are both non-tissue objects. Therefore, tissue detector **204** may detect the contours of each object by converting the image to grayscale, comparing the grayscale values of each pixel to a range of user-determined range of values that are not too bright or too dark, and determining whether the grayscale value is within the range to produce a binary image where each pixel is assigned one of two numerical values.

[0103] For example, to threshold an image, the tissue detector **204** may compare the grayscale values of each pixel to a user-defined range of values and replace each grayscale value outside of the user-defined range with the value 0 and each grayscale value within a user-

defined range with the value 1. Then, the tissue detector **204** draws all contours of all objects as the outer edge of each group of adjacent pixels having a value of 1. Closed contours indicate the presence of an object, and the tissue detector **204** measures the area within the contours of each object to measure the size of the object.

[0104] In one example, tissue objects on a slide are unlikely to make contact with the outer edges of the slide and the tissue detector **204** categorizes all objects that contact the edge of a slide as non-tissue.

[0105] In one example, after measuring the size of each object, the tissue detector **204** ranks the sizes of all objects and designates the largest value to be the size of the largest object. The tissue detector **204** divides the size of each object by the size of the largest object and compares the resulting size quotient to a user-defined size threshold value. If the size quotient for an object is smaller than the user-defined size threshold value, the tissue detector **204** designates that object as non-tissue. In one example, the user-defined size threshold value is 0.1.

[0106] Before measuring the size of each object, the tissue detector **204** may first downsample the input image to reduce the likelihood of designating portions of a tissue object as non-tissue. For example, a single tissue object may appear as a first tissue object portion surrounded by one or more additional tissue object portions having a smaller size. After thresholding, the additional tissue object portions may have a size quotient smaller than the user-defined size threshold value and may be erroneously designated as non-tissue. Downsampling before thresholding causes a small group of adjacent pixels having values of 1 surrounded by pixels having values of 0 in the original image to be included in a proximal, larger group of pixels having values of 1. The opposite may also be true, for small groups of adjacent pixels having values of 0 surrounded by pixels having values of 1 in the original image to be included in a proximal, larger group of pixels having values of 0.

[0107] In one example, the tissue detector **204** downsamples an image having 40x magnification by a ratio of 16x, so the magnification of the resulting downsampled image is 40/16x and each pixel in the downsampled image represents 16 pixels in the original image.

[0108] In one example, the tissue detector **204** detects the boundaries of each object on the slide as a cluster of pixels having binary or RGB values that do not equal zero, surrounded by pixels with RGB values equal to zero, indicating an object border. If the pixels forming the

boundaries lie on a relatively straight line, the tissue detector **204** classifies the object as non-tissue. For example, the tissue detector **204** outlines a shape with a closed polygon. If the number of vertices of the polygon is less than a user-defined minimum vertices threshold, the polygon is deemed to be a simple, inorganic shape that is too smooth, and marked as non-tissue.

[0109] FIGS. 5A, 5B, and 5C illustrate an exemplary grid overlay on a digital image of a histology slide for tissue classification. FIGS. 5A and 5B show examples of how a multi-field-of-view image tile input looks. FIG. 5A is the input image used to get prediction result of single central small tile. FIG. 5B is the input of image to get prediction results of a matrix of central small tiles. The grid overlays shown in FIGS. 5A and 5B, represent the tiling of the image that is performed by the PhiNet architecture, as described herein, e.g., at FIG 6A. In FIG. 5A, a single tile (teal) and its center (yellow) are shown. In FIG. 5B, shown is a grid of tiles, which is what is actually input into PhiNet architecture for inference. The PhiNet architecture processes each of the tiles in parallel, where each of the little yellow center squares receives a probability for each possible class (summing to 100% as described herein). FIG. 3A shows one possible output overlay map for a breast cancer model. There are as many overlay maps as there are models that may be displayed. For example, if the techniques include breast, colorectal, and lung models, then 3 overlay maps may be generated, and the techniques can also show the heat map (probability map) for a single class overlaid over the image.

[0110] After eliminating non-tissue pixels from the image, the tissue detector **204** stores data containing the locations of the remaining pixels representing tissue. In one example, the locations are stored as an image of the remaining pixels which could be a grayscale mask, a list of the locations of each vertex of the polygons that surround the remaining pixels, and/or other coordinates that indicate the locations of the tissue pixels.

[0111] The tissue class tile grid projector **208** receives the data containing the locations of the pixels representing tissue and overlays a tile grid over these pixels. In one example, the grid is composed of small square tiles that do not overlap.

[0112] The tissue class locator **216** includes a tissue classification algorithm (see FIGS. 6A through 6C) that assigns a tissue class label to the image represented in each small square tile. In one example, the digital tissue segmenter **201** may report the assigned tissue class label

associated with each small square tile by displaying a grid-based digital overlay map in which each tissue class is represented by a unique color (see FIG. 3A).

[0113] A smaller tile size may cause an increase in the amount of time required for the tissue class locator **216** to analyze the input image. Alternatively, a larger tile size may increase the likelihood that a tile will contain more than one tissue class and make it difficult to assign a single tissue class label to the tile. In this case, the tissue class locator **216** may calculate an equal probability for two or more tissue class labels being accurately assigned to a single small square tile instead of calculating that one of the tissue class labels has a higher probability of describing the image in the small square tile, compared to the other tissue class labels.

[0114] In one example, each side of each small square tile is approximately 32 microns long and approximately 5-10 cells fit in each small square tile. This small tile size allows the algorithm to create more spatially accurate borders when determining the boundary between two neighboring small square tile regions that depict two distinct tissue classes. In one example, each side of the small square tile can be as short as 1 micron.

[0115] In one example, the size of each tile is set by the user to contain a specific number of pixels. In this example, the resolution of the input image will determine the length of each side of the tile, as measured in microns. At different resolutions, the micron length of the tile side will vary and the number of cells in each tile may vary.

[0116] The algorithm recognizes various pixel data patterns in the portion of the digital image that is located within or near each small square tile and assigns a tissue class label to each small square tile based on those detected pixel data patterns. In one example, a medium square tile centered around a small square tile contains the area of a slide image that is close enough to the small square tile to contribute to the label assignment for that small square tile.

[0117] In one example, each side of a medium square tile is approximately 466 microns long, and each medium square tile contains approximately 225 (15x15) small square tiles. In one example, this medium tile size increases the likelihood that structural tissue features can fit within a single medium tile and provide context to the algorithm when labeling the central small square tile. Structural tissue features may include glands, ducts, vessels, immune clusters, etc.

[0118] In one example, this medium tile size is selected to counteract the shrinkage that occurs during convolution.

[0119] During convolution, an input image matrix is multiplied by a filter matrix to create a result matrix, and shrinkage refers to a case where the result matrix is smaller than the input image matrix. The dimensions of a filter matrix in a convolution layer affects the number of rows and columns lost to shrinkage. The total number of matrix entries that are lost to shrinkage by processing an image through a particular CNN can be calculated depending on the number of convolution layers in the CNN and the dimensions of the filter matrices in each convolution layer. (See FIGS. 6A through 6C)

[0120] In the example shown in FIG. 6B, the convolution layers in combination lose 217 total matrix rows or columns from the top, bottom, and two side edges of the matrix, so the medium square tile is set to equal the small square tile plus 217 pixels on each side of the small square tile.

[0121] In one example, two neighboring small square tiles share a side and are each at the center of a medium square tile. The two medium square tiles overlap. Of the 466×466 small pixels located in each medium square tile, the two medium square tiles will share all but 32×466 pixels. In one example, each convolution layer of the algorithm (see FIGS. 6A through 6B) analyzes both medium square areas simultaneously such that the algorithm produces two vectors of values (one for each of the two small square tiles).

[0122] The vector of values contains a probability value for each tissue class label, indicating the likelihood that the small square tile depicts that tissue class. The vectors of values will be arranged in a matrix, to form a 3-dimensional probability data array. The location of each vector in the 3-dimensional probability data array, relative to the other vectors, will correspond to the location of the associated small square tile, relative to the other small square tiles included in the algorithm analysis.

[0123] In the example, 434×434 (188,356) of the 466×466 (217,156) pixels in each medium square tile are common to both medium square tiles. By analyzing both medium square tiles simultaneously, the algorithm increases efficiency.

[0124] In one example, the algorithm can further increase efficiency by analyzing a large tile formed by multiple overlapping medium square tiles, each of which contains many small square tiles surrounding one center small square tile that receives a tissue class label. In this example, the algorithm still generates one data structure in the form of a 3-dimensional probability data array containing one vector of probabilities for each small square tile, wherein the location of the

vector within the 3-dimensional array corresponds to the location of the small tile within the large tile.

[0125] The tissue class locator **216** saves this 3-dimensional probability data array and the overlay map generator **224** converts the tissue class label probabilities for each small square tile into a tissue class overlay map. In one example, the overlay map generator **224** may compare the probabilities stored in each vector to determine the largest probability value associated with each small square tile. The tissue class label associated with that largest value may be assigned to that small square tile and only the assigned labels will be displayed in the tissue class overlay map.

[0126] In one example, matrices generated by each layer of the tissue class locator **216** for the large square tile are stored in graphics processing unit (GPU) memory. The capacity of the GPU memory and the amount of GPU memory required for each entry in the 3-dimensional probability data array may determine the maximum possible size of the large square tile. In one example, the GPU memory capacity is 250 MB and each entry in the matrices requires 4 bytes of GPU memory. This allows a large tile size of 4,530 pixels by 4,530 pixels, calculated as follows: $4 \text{ bytes/entry} * 4530 * 4530 * 3 \text{ entries for each large tile} = 246 \text{ (~250) MB of GPU memory required per large square tile}$. In another example, each entry in the matrices requires 8 bytes of GPU memory. In this example, a 16GB GPU can process 32 large tiles simultaneously, each large tile having dimensions of 4,530 pixels by 4,530 pixels, calculated as follows: $32 \text{ large tiles} * 8 \text{ bytes/entry} * 4530 * 4530 * 3 \text{ entries for each large tile} = 14.7 \text{ (~16) GB of GPU memory required}$.

[0127] In one example, each entry in the 3-dimensional probability data array is a single precision floating-point format (float32) data entry.

[0128] In one example, there are 16,384 (128^2) non-overlapping small square tiles that form a large square tile. Each small square tile is the center of a medium square tile having sides that are each approximately 466 pixels long. The small square tiles form a center region of a large square tile having sides that are each approximately 4,096 pixels long. The medium square tiles all overlap and create a border around all four sides of the center region that is approximately 217 pixels wide. Including the border, each large square tile has sides that are each approximately 4,530 pixels long. (See FIG. 5C)

[0129] In this example, this large square tile size allows simultaneous calculations that reduce the redundant computation percentage by 99%. This may be calculated as follows: first, select a pixel on the interior of a large square tile (any pixel at least 434 pixels from the edge of the large square tile). Construct a region that is the size of a medium square tile (466 pixels per edge) with this model pixel at the center. Then for any small square tile centered within this constructed region, the model pixel is contained within that small square tile's corresponding medium square tile. There are $(466/32)^2 = \sim 217$ such small square tiles within the large square tile. For pixels not on the interior of the large square tile, the number of small square tiles that satisfy this condition is smaller. The number decreases linearly as the distance between the selected small square tile and the edge of the large square tile decreases, then again as the distance between the selected small square tile and the corner decreases, where a small number of pixels ($\sim 0.005\%$) only contribute towards the classification of a single small square tile. Performing classification on a single large square tile means the computations for each pixel are only performed once, instead of once per small square tile. Thus, the redundancy is reduced by nearly 217-fold. In one example, redundancy is not completely eliminated because a slide may contain several large square tiles, each of which may overlap slightly with its neighbors.

[0130] An upper bound on the redundant calculation percentage can be established (slight deviation from this upper bound depends on the number of large square tiles needed to cover the tissue and the relative arrangement of these tiles). The redundancy percentage is $1 - 1/r$ where r is the redundancy ratio, and r can be calculated as $(T/N + 1)(\sqrt{N} \cdot E + 434)^2 / (\sqrt{T} \cdot E + 434)^2$; T is the total number of small square tiles on the slide, N is the number of small square tiles per large square tile, and E is the edge size of the small square tiles.

[0131] FIG. 6A illustrates the layers of an example tissue segmentation algorithm. FIG. 6B compares the layers of a known image classification algorithm with an example tissue segmentation algorithm.

[0132] In one example, the tissue class locator **216** includes a convolutional neural network (CNN), a deep learning algorithm that has been trained by a training set to recognize pixel data patterns.

[0133] In one example, the tissue class locator **216** includes a tile-resolution fully convolutional network (FCN) black box deep learning model based on a known CNN ResNet-18

image recognition model. In one example, the tile-resolution FCN is a multi-tile algorithm known as PhiNet, the layers of which are shown in FIG. 6A.

[0134] FIG. 6B illustrates the differences between the ResNet-18 algorithm on the left, and the tile-resolution FCN PhiNet shown in FIG. 6A and on the right half of FIG. 6B. Compared to the ResNet-18 algorithm, the tile-resolution FCN PhiNet included in the digital tissue segmenter **201** has additional layers of 1x1 convolution in a skip connection, downsampling by a factor of 8 in a skip connection, and a confidence map layer, and replaces an average pooling layer with a concatenation layer, and a fully connected (fc N) layer with a 1x1 convolution and softmax layer.

[0135] The added layers convert a classification task into a classification-segmentation task. This means that instead of receiving and classifying a whole image as one tissue class label, the added layers allow the tile-resolution FCN to classify each small tile in the user-defined grid as a tissue class.

[0136] These added and replacement layers convert a CNN to a tile-resolution FCN without requiring the upsampling performed in the later layers of traditional pixel-resolution FCNs. Upsampling is a method by which a new version of an original image can be created with a higher resolution value than the original image. Upsampling is a time-consuming, computation-intensive process.

[0137] There are many methods known in the art for upsampling, including nearest-neighbor, bilinear, hermite, bell, Mitchell, bicubic, and Lanczos resampling. In one example, 2x upsampling means that a pixel with red green blue (RGB) values will be split into four pixels, and the RGB values for the three new pixels may be selected to match the RGB values of the original pixel. In another example, the RGB values for the three new pixels may be selected as the average of the RGB values from the original pixel and the pixels that are adjacent to the neighboring pixel.

[0138] Because the RGB values of the new pixels may not accurately reflect the visible tissue in the original slide that was captured by the digital slide image, upsampling can introduce errors into the final image overlay map produced by the overlay map generator **224**.

[0139] Additionally, traditional pixel-resolution FCNs require training data wherein each pixel is labeled, which is not feasible to produce from digital histology slides because it is too time

consuming and because each individual pixel does not contain enough contextual information to identify a tissue class.

[0140] In one example, instead of labeling individual pixels, the tile-resolution FCN is programmed to analyze a large square tile made of small square tiles (see FIG. 5C), producing a 3D array of values that each represent the probability that one tissue class classification label matches the tissue class depicted in each small tile.

[0141] A convolution layer, known in the art, performs the multiplication of at least one input image matrix by at least one filter matrix. In the first convolution layer, the input image matrix has a value for every pixel in the large square tile input image, representing visual data in that pixel (for example, a value between 0 and 255 for each channel of RGB).

[0142] The filter matrix has dimensions selected by the user, and may contain weight values selected by the user or determined by backpropagation during CNN model training. In one example, in the first convolution layer, the filter matrix dimensions are 7x7 and there are 64 filters. The filter matrix may represent visual patterns that can distinguish one tissue class from another.

[0143] In an example where RGB values populate the input image matrix, the input image matrix and the filter matrices will be 3-dimensional. (See FIG. 6C)

[0144] Each filter matrix is multiplied by each input image matrix to produce a result matrix. All result matrices produced by the filters in one convolution layer may be stacked to create a 3-dimensional result matrix having dimensions such as rows, columns, and depth. The last dimension, depth, in the 3-D result matrix will have a depth equal to the number of filter matrices.

[0145] The result matrix from one convolution layer becomes the input image matrix for the next convolution layer.

[0146] A convolution layer title that includes “/n”, where n is a number, indicates that there is a downsampling (also known as pooling) of the result matrix produced by that layer. The n indicates the factor by which the downsampling occurs. Downsampling by a factor of 2 means that a downsampled result matrix with half as many rows and half as many columns as the original result matrix will be created by replacing a square of four values in the result matrix by

one of those values or a statistic calculated from those values. For example, the minimum, maximum, or average of the values may replace the original values.

[0147] The algorithm shown also adds skip connections (shown in FIG. 6A as black lines with arrows that connect blue convolution layers directly to the concatenation layer). The skip connection on the left includes downsampling by a factor of 8, and the skip connection on the right includes two convolution layers that multiply an input image matrix by filter matrices that each have dimensions of 1x1. Because of the 1x1 dimensions of the filter matrices in these layers, only an individual small square tile contributes to its corresponding probability vector in the result matrices created by the purple convolution layers. These result matrices represent a small focus of view.

[0148] In all of the other convolution layers, the larger dimensions of the filter matrices allow the pixels in each medium square tile, including the small square tile at the center of the medium square tile, to contribute to the probability vector in the result matrix that corresponds with that small square tile. These result matrices allow the contextual pixel data patterns surrounding the small square tile to influence the probability that each tissue class label applies to the small square tile. These result matrices represent a large focus of view.

[0149] The 1x1 convolution layers in the skip connection allow the algorithm to regard the pixel data patterns in the center small square tile as more or less important than pixel data patterns in the rest of the surrounding medium square tile. This is reflected by the weights that the trained model multiplies by the final result matrix from the skip connection layers (shown on the right side of FIG. 6A) compared to the weights that the trained model multiplies by the final result matrix from the medium tile convolution layers (shown in the center column of FIG. 6A) during the concatenation layer.

[0150] The downsampling skip connection shown on the left side of FIG. 6A creates a result matrix with a depth of 64. The 3x3 convolution layer having 512 filter matrices creates a result matrix with a depth of 512. The 1x1 convolution layer having 64 filter matrices creates a result matrix with a depth of 64. All three of these results matrices will have the same number of rows and the same number of columns. The concatenation layer concatenates these three results matrices to form a final result matrix with the same number of rows and the same number of columns as the three concatenated matrices, and a depth of $64+512+64$ (640). This final result matrix combines the large and small focus of view matrices.

[0151] The final result matrix may be flattened to 2 dimensions by multiplying a factor by every entry, and summing the products along each depth. Each factor may be selected by the user, or may be selected during model training by backpropagation. Flattening will not change the number of rows and columns of the final results matrix, but will change the depth to 1.

[0152] The 1x1 convolution layer receives the final result matrix and filters it with one or more filter matrices. The 1x1 convolution layer may include one filter matrix associated with each tissue class label in the trained algorithm. This convolution layer produces a 3-D result matrix that has a depth equal to the number of tissue class labels. Each depth corresponds to one filter matrix and along the depth of the result matrix there may be a probabilities vector for each small square tile. This 3-D result matrix is the 3-dimensional probability data array, and the 1x1 convolution layer stores this 3-D probability data array.

[0153] A softmax layer may create a 2-dimensional probability matrix from the 3-D probability data array by comparing every value in each probabilities vector and selecting the tissue class associated with the maximum value to assign that tissue class to the small square tile associated with that probabilities vector.

[0154] The stored 3-dimensional probability data array or the 2-D probability matrix may then be converted to a tissue class overlay map in the final confidence map layer in FIG. 6A, to efficiently assign a tissue class label to each tile.

[0155] In one example, to counteract shrinkage, input image matrices have added rows and columns on all four outer edges of the matrices, wherein each value entry in the added rows and columns is a zero. These rows and columns are referred to as padding. In this case, the training data input matrices will have the same number of added rows and columns with value entries equal to zero. A difference in the number of padding rows or columns in the training data input matrices would result in values in the filter matrices that do not cause the tissue class locator **216** to accurately label input images.

[0156] In the FCN shown in FIG. 6B, 217 total outer rows or columns on each side of the input image matrix will be lost to shrinkage before the skip connection, due to the gray and blue layers. Only the pixels located in the small square tiles will have a corresponding vector in the result matrices created by the green layers and beyond.

[0157] In one example, each medium square tile is not padded by adding rows and columns with value entries of zero around the input image matrix that corresponds to each medium square tile because the zeroes would replace image data values from neighboring medium square tiles that the tissue class locator **216** needs to analyze. In this case, the training data input matrices will not be padded either.

[0158] FIG. 6C is a visualization of each depth of an exemplary 3-dimensional input image matrix being convoluted by two exemplary 3-dimensional filter matrices.

[0159] In an example where an input image matrix contains RGB channels for each medium square tile, the input image matrix and filter matrices will be 3-dimensional. In one of the three dimensions, the input image matrix and each filter matrix will have three depths, one for red channel, one for green channel, and one for blue channel.

[0160] The red channel (first depth) of the input image matrix is multiplied by the corresponding first depth of the first filter matrix. The green channel (second depth) is multiplied in a similar fashion, and so on with the blue channel (third depth). Then, the red, green, and blue product matrices are summed to create a first depth of the 3-dimensional result matrix. This repeats for each filter matrix, to create an additional depth of the 3-dimensional result matrix that corresponds to each filter.

[0161] A variety of training sets may be used to train a CNN or FCN that is included in the tissue class locator **216**.

[0162] In one example, the training set may contain JPEG images of medium square tiles, each having a tissue class label assigned to its center small square tile, taken from at least 50 digital images of histology slides at a resolution of approximately 1 pixel per micron. In one example, a human analyst has outlined and labeled all relevant tissue classes or labeled each small square tile in each histology slide as non-tissue or as a specific type of cells. Classes of tissue may include tumor, stroma, normal, immune cluster, necrosis, hyperplasia/dysplasia, and red blood cells. In one example, each side of each center small square tile is approximately 32 pixels long.

[0163] In one example, the training set images are converted to input training image matrices and processed by the tissue class locator **216** to assign a tissue class label to each training set image. If the tissue class locator **216** does not accurately label the validation set of training

images to match the corresponding annotations added by a human analyst, the weights of each layer of the deep learning network may be adjusted automatically by stochastic gradient descent through backpropagation until the tissue class locator **216** accurately labels most of the validation set of training images.

[0164] In one example, the training data set has multiple classes where each class represents a tissue class. That training set will generate a unique model with specific hyperparameters (number of epochs, learning rate, etc.) that can recognize and classify the content in a digital slide image into different classes. Tissue classes may include tumor, stroma, immune cluster, normal epithelium, necrosis, hyperplasia/dysplasia, and red blood cells. In one example, the model can classify an unlimited number of tissue classes, provided each tissue class has a sufficient training set.

[0165] In one example, the training set images are converted into grayscale masks for annotation where different values (0-255) in the mask image represent different classes.

[0166] Each histology slide can exhibit large degrees of variation in visual features, including tumor appearance, so a training set may include digital slide images that are highly dissimilar to better train the model for the variety of slides that it may analyze.

[0167] A training set may also be specific to a cancer type. In this case, all of the histology slides that generated the digital images in a particular training set contain a tumor sample from the same cancer type. Cancer types may include breast, colorectal, lung, pancreatic, liver, stomach, skin, etc. Each training set may create a unique model specific to the cancer type. Each cancer type may also be split into cancer subtypes, known in the art.

[0168] In one example, a training set may be derived from histology slide pairs. A histology slide pair includes two histology slides that each have one slice of tissue, wherein the two slices of tissue were adjacent in the tumor sample. Therefore, the two slices of tissue are substantially similar. One of the slides in the pair is stained with H and E staining only, and the other slide in the pair is stained with IHC staining for a specific molecule target. The areas on the H and E stained slide that correspond to areas where IHC staining appears in the paired slide are annotated by a human analyst as containing a specific molecule target and the tissue class locator receives the annotated H and E slides as a training set.

[0169] In this case, the trained tissue class locator **216** receives digital images of H and E stained tissue to predict tiles that may contain a given molecule target and the overlay map generator **224** produces an overlay map showing which tiles contain a given molecule. In one example, the resolution of the overlay is at the level of an individual cell.

[0170] The overlay produced by a model trained by one or more training sets may be reviewed by a human analyst in order to annotate the digital slide image to add it to one of the training sets.

[0171] The pixel data patterns that the algorithm detects may represent visually detectable features. Some examples of those visually detectable features may include color, texture, cell size, shape, and spatial organization.

[0172] For example, color on a slide provides contextual information. An area on the slide that is purple has a higher density of cells and is more likely to be invasive tumor. Tumor also causes the surrounding stroma to become more fibrous in a desmoplastic reaction, which causes normally pink stroma to appear blue-grey. Color intensity also helps to identify individual cells of a certain type (for example, lymphocytes are uniformly very dark blue).

[0173] Texture refers to the distribution of stain within cells. Most tumor cells have a rough, heterogeneous appearance, with light pockets and dark nucleoli within their nuclei. A zoomed-out field of view with many tumor cells will have this rough appearance. Many non-tumor tissue classes each has distinguishing features. Furthermore, patterns of tissue classes that are present in a region can indicate the type of tissue or cell structures present in that region.

[0174] Additionally, cell size often indicates tissue class. If a cell is several times larger than normal cells elsewhere on the slide, the probability is high that it is a tumor cell.

[0175] The shape of individual cells, specifically how circular they are, can indicate what type of cell they are. Fibroblasts (stromal cells) are normally elongated and slim, while lymphocytes are very round. Tumor cells can be more irregularly shaped.

[0176] The organization of a group of cells can also indicate tissue class. Frequently, normal cells are organized in a structured and recognizable pattern, but tumor cells grow in denser, disorganized clusters. Each type and subtype of cancer can produce tumors with specific growth patterns, which include cell location relative to tissue features, the spacing of tumor cells relative to each other, formation of geometric elements, etc.

[0177] FIG. 7 illustrates an example of a first classification model for labeling an H&E Slide or IHC Slide using tile classifications and cell classifications. As shown, an exemplary model overlay may include red tiles for identifying clusters of lymphocytes, green tiles for identifying epithelium cells, and dark blue tiles for identifying tumor cells. The tile classifications are shown in the legend and generated using a PhiNet model architecture. The cell legend showing an example cell classification is shown, as generated used a UNet model architecture. In some examples, where a tile has a predicted tile classification and that tile corresponds to identified cell (i.e., cell object) with a classification, the cell classification may override, such that the segmenter **201** assigns the cell classification to that tile in place of the tile classification. In another example, an exemplary overlay may represent individual lymphocytes by coloring them red to distinguish between clusters of lymphocytes and individual lymphocytes infiltrating other tissues. In one example, individual lymphocytes detected within tumor tiles represent tumor-infiltrating lymphocytes which would be red dots within dark blue or green tiles.

[0178] In one example, the cell type locator **220** includes at least one cell outlining algorithm for detecting the outer edge of each cell and reporting the coordinates of every pixel that contains an outer edge of a cell to create a cell outer edge overlay map. (See FIG. 2B)

[0179] The cell type tile grid projector **212** may overlay a grid on the digital image of a histology slide. A smaller tile size may improve the training data set by increasing the number of tiles in the training data set, and may be easier for a human analyst to visually inspect before training the cell type locator **220**. The size of the tiles may be determined by the capacity of the memory.

[0180] The cell type locator **220** may recognize features that include changes in color and brightness between adjacent pixels, and the shape of areas with changes in color or brightness. The cell type locator **220** may also report the coordinates of the center of each detected cell and of all vertices of a polygon that encloses a single detected cell. The cell type locator **220** may include a radial symmetry filter or an edge detector such as a Laplacian edge filter or a Sobel edge filter. In one example, the cell type locator **220** is a UNet model, which is a three-class semantic segmentation FCN model.

[0181] In one example, the UNet is a three-class model, which means that the algorithm may classify each pixel as either background exterior to every cell, cell outer edge, or cell interior. This facilitates the counting of each individual cell, especially when two or more cells overlap

each other. In one example, tumor infiltrating lymphocytes will overlap tumor cells. In traditional two-class cell outlining models that only label whether a pixel contains a cell outer edge or not, each clump of two or more overlapping cells would be counted as one cell. Two-class UNet models with binary classification (cell vs background) are known in the art, but the three-class UNet model allows a type of classification that is not binary, which requires adaptation with the use of a different loss function. The classified results may be stored as matrix with different values to indicate cell exterior, cell outer edge, and cell interior. A training process is done by having an original image and a mask image together as input to a UNet model (e.g., the mask may be an image with pixel values at 0, 128 and 255, where 0 represent cell exterior, 128 represent cell outer edge, 255 represent cell interior).

[0182] In one example, to avoid the possibility that a cell that spans two tiles is counted twice, the algorithm may add a buffer around all four sides of each tile that is slightly wider than an average cell. The intention is to only count cells that appear in the center, non-buffered region for each tile. In this case, tiles will be placed so that the center, non-buffered region of neighboring tiles are adjacent and non-overlapping. Neighboring tiles will overlap in their respective buffer regions.

[0183] In one example, the cell segmentation algorithm includes two UNet models. One UNet model may be trained with images patches of mixed tissue classes, where a human analyst has highlighted a bounding polygon for each cell and classified each cell according to tissue class. In one example, training data includes digital slide images where every pixel has been labeled as either the interior of a cell, the outer edge of a cell, or the background which is exterior to every cell. In another example, the training data includes digital slide images where every pixel has been labeled with a yes or no to indicate whether it depicts the outer edge of a cell. The labels may be two-dimensional matrices with the same shapes as their corresponding image patches, where each matrix element takes on one of several values (the exact number of values is determined by the number of unique classes in the model).

[0184] Another UNet model may be trained with images of many cells of a single tissue class, or images of a diverse set of cells where cells of only one tissue class are outlined in a binary mask. In one example, the training set is labeled by associating a first value with all pixels showing a cell type of interest, a second value to all background pixels, and optionally a third value to pixels along the border of each cell (in the case of a three-class model instead of a two-class model). Visually, an image labeled in this way might appear as a black and white image

wherein all pixels showing a tissue class of interest would be white and all other pixels would be black, or vice versa. (See FIG. 8) For example, the images may have only labeled lymphocytes. This UNet model can recognize the outer edges of that particular cell type and assign a label to cells of that type in the digital image of the slide.

[0185] In one example, a non-UNet model may be used instead of a UNet model to label cells that are detected by a first UNet model. This non-UNet model could be another CNN that has a different architecture, or a non deep learning model like random forest or support vector machine.

[0186] In one example, the cell labels generated by the cell type locator **220** can be augmented by the tissue region labels generated by the tissue class locator **216**. If a cell is detected but not labeled a specific class by the cell type locator, it may be assigned a default label that is the label of the tissue region in which it is found.

[0187] In one example, the images in both training sets have a resolution of 2 pixels per micron.

[0188] The techniques herein, applying PhiNet architecture as described, may be used to predict other indicators, as well, such as microsatellite instability (MSI), consensus molecular subtypes (CMS), and other molecular biomarkers. Classification models may be trained to predict from H&E slides these biomarkers, using other additional information, such as other characteristics of the patient that may be included in the patient's electronic health record. Other biomarkers that can be predicted using trained PhiNet architecture include: tumor grade and/or stage, tumor origin (for TUO cases), histological subtypes (such as adenocarcinoma and squamous cell carcinoma), other biomarkers such as cell based, vaccine, or cytokine therapies).

[0189] Further, wherein exemplary embodiments disclosed herein are illustrated with H&E slides prepared by a pathologist, similar models may be generated for IHC slides as well. Exemplary IHC overlays may classify tiles according to a status of PD-L1 'presence' in the cell. In an example implementation, a digital tissue segmenter similar to that of segmenter **201** in FIG. 2 may be configured to as a digital PD-L1 tumor positive score (TPS) predictor. In an example, this PD-L1 TPS digital tissue segmenter includes a tissue detector, tissue class tile grid projector, and tissue class locator, like that of segmenter **201** in FIG. 2. The output of the tissue class locator 216 may be provided to a PD-L1 TPS metric calculator and overlay map generator (replacing the element **224** in the segmenter **201**) configured to count the number of

Tumor + tiles and divide by the total number of tumor tiles (i.e., the sum of Tumor + and Tumor - tiles), to determine a predicted PD-L1 TPS. Other quantities can also be calculated, such as the percent of stained lymphocytes or overall macrophage percentage. The PD-L1 tissue classifier in this example may be configured to predict any number of classes, including blood, epithelium, lymphocyte +, lymphocyte -, macrophage +, necrosis, stroma, tumor +, and tumor -, where "+" indicates positive PD-L1 expression and "-" indicates no PD-L1 expression. In this way, in some examples, a PD-L1 prediction process can be implemented by a digital tissue image segmenter implementing two main processes. First, the digital tissue image segmenter predicts different tissue classes overlaid onto an IHC image (i.e., a similar process as described in reference to H&E images described herein, but with different classes). Second, the digital tissue image segmenter obtains PD-L1 TPS by counting the number of predicted Tumor + tiles and dividing that by the total number of Tumor tiles (i.e. Tumor + added to Tumor -). An example overlay mapping process that may be performed by such a digital PD-L1 tumor positive score (TPS) predictor is shown in FIG. 9, showing different predicted tissue classifications. On the left are various IHC images and on the right are predicted TPS various of different tissue classifications overlaid on the IHC images.

[0190] FIG. 10 illustrates a system **1000** for performing the methods described herein may include a computing device **1002**, and more particularly may be implemented on one or more processing units **1004**, e.g., Central Processing Units (CPUs), and/or on one or more or Graphical Processing Units (GPUs) **1006**, including clusters of CPUs and/or GPUs. Features and functions described may be stored on and implemented from one or more non-transitory computer-readable media **1008** of the computing device. The computer-readable media **1008** may include, for example, an operating system **1010** and software modules, or "engines," that implement the methods described herein, including a digital tissue image segmenter **1012** implementing methods and processing illustrated and described herein. Example processes include tissue detection, tissue class tile grid projection, cell type tile grid projection, tissue class locator, cell type locator, overlay map generator and metric calculator processes. More generally, the computer-readable media **1008** may store batch normalization process instructions for the engines for implementing the techniques herein. The computing device **1002** may be a distributed computing system, such as an Amazon Web Services cloud computing solution. The computing device **1002** may be implemented on one network accessible processing device **1050** or distributed across multiple such devices **1050**, **1052**, **1054**, etc.

[0191] The computing device **1002** includes a network interface **1014** communicatively coupled to network **1016**, for communicating to and/or from a portable personal computer, smart phone, electronic document, tablet, and/or desktop personal computer, or other computing devices for communicating overlay maps, predicted tile classifications and locations, predicted cell classifications and locations, etc. Such information may also be stored in a database **1024**. The computing device **1002** further includes an I/O interface **1018** connected to devices, such as digital displays **1020** for displaying generator overlay maps, user input devices **1022**, etc. A dashboard generator **1026** may be used to generate GUI and/or other digital displays allowing a user to review and interact with and adjust generated overlay maps, prediction classifications, etc.

[0192] The functions of the engines may be implemented across distributed computing devices, etc. connected to one another through a communication link. In other examples, functionality of the system may be distributed across any number of devices, including the portable personal computer, smart phone, electronic document, tablet, and desktop personal computer devices shown. The computing device **1002** may be communicatively coupled to the network **1016** and another network. The networks **1016** may be public networks such as the Internet, a private network such as that of a research institution or a corporation, or any combination thereof. Networks can include, local area network (LAN), wide area network (WAN), cellular, satellite, or other network infrastructure, whether wireless or wired. The networks can utilize communications protocols, including packet-based and/or datagram-based protocols such as Internet protocol (IP), transmission control protocol (TCP), user datagram protocol (UDP), or other types of protocols. Moreover, the networks can include a number of devices that facilitate network communications and/or form a hardware basis for the networks, such as switches, routers, gateways, access points (such as a wireless access point as shown), firewalls, base stations, repeaters, backbone devices, etc.

[0193] The computer-readable media **1008** may include executable computer-readable code stored thereon for programming a computer (e.g., comprising a processor(s) and GPU(s)) to the techniques herein. Examples of such computer-readable storage media include a hard disk, a CD-ROM, digital versatile disks (DVDs), an optical storage device, a magnetic storage device, a ROM (Read Only Memory), a PROM (Programmable Read Only Memory), an EPROM (Erasable Programmable Read Only Memory), an EEPROM (Electrically Erasable Programmable Read Only Memory) and a Flash memory. More generally, the processing units of the computing device may represent a CPU-type processing unit, a GPU-type processing

unit, a field-programmable gate array (FPGA), another class of digital signal processor (DSP), or other hardware logic components that can be driven by a CPU.

[0194] Throughout this specification, plural instances may implement components, operations, or structures described as a single instance. Although individual operations of one or more methods are illustrated and described as separate operations, one or more of the individual operations may be performed concurrently, and nothing requires that the operations be performed in the order illustrated. Structures and functionality presented as separate components in example configurations may be implemented as a combined structure or component. Similarly, structures and functionality presented as a single component may be implemented as separate components or multiple components. These and other variations, modifications, additions, and improvements fall within the scope of the subject matter herein.

[0195] Additionally, certain embodiments are described herein as including logic or a number of routines, subroutines, applications, or instructions. These may constitute either software (e.g., code embodied on a machine-readable medium or in a transmission signal) or hardware. In hardware, the routines, etc., are tangible units capable of performing certain operations and may be configured or arranged in a certain manner. In example embodiments, one or more computer systems (e.g., a standalone, client or server computer system) or one or more hardware modules of a computer system (e.g., a processor or a group of processors) may be configured by software (e.g., an application or application portion) as a hardware module that operates to perform certain operations as described herein.

[0196] In various embodiments, a hardware module may be implemented mechanically or electronically. For example, a hardware module may comprise dedicated circuitry or logic that is permanently configured (e.g., as a special-purpose processor, such as a microcontroller, field programmable gate array (FPGA) or an application-specific integrated circuit (ASIC)) to perform certain operations. A hardware module may also comprise programmable logic or circuitry (e.g., as encompassed within a processor or other programmable processor) that is temporarily configured by software to perform certain operations. It will be appreciated that the decision to implement a hardware module mechanically, in dedicated and permanently configured circuitry, or in temporarily configured circuitry (e.g., configured by software) may be driven by cost and time considerations.

[0197] Accordingly, the term "hardware module" should be understood to encompass a tangible entity, be that an entity that is physically constructed, permanently configured (e.g., hardwired), or temporarily configured (e.g., programmed) to operate in a certain manner or to perform certain operations described herein. Considering embodiments in which hardware modules are temporarily configured (e.g., programmed), each of the hardware modules need not be configured or instantiated at any one instance in time. For example, where the hardware modules comprise a processor configured using software, the processor may be configured as respective different hardware modules at different times. Software may accordingly configure a processor, for example, to constitute a particular hardware module at one instance of time and to constitute a different hardware module at a different instance of time.

[0198] Hardware modules can provide information to, and receive information from, other hardware modules. Accordingly, the described hardware modules may be regarded as being communicatively coupled. Where multiple of such hardware modules exist contemporaneously, communications may be achieved through signal transmission (e.g., over appropriate circuits and buses) that connects the hardware modules. In embodiments in which multiple hardware modules are configured or instantiated at different times, communications between such hardware modules may be achieved, for example, through the storage and retrieval of information in memory structures to which the multiple hardware modules have access. For example, one hardware module may perform an operation and store the output of that operation in a memory device to which it is communicatively coupled. A further hardware module may then, at a later time, access the memory device to retrieve and process the stored output. Hardware modules may also initiate communications with input or output devices, and can operate on a resource (e.g., a collection of information).

[0199] The various operations of the example methods described herein can be performed, at least partially, by one or more processors that are temporarily configured (e.g., by software) or permanently configured to perform the relevant operations. Whether temporarily or permanently configured, such processors may constitute processor-implemented modules that operate to perform one or more operations or functions. The modules referred to herein may, in some example embodiments, comprise processor-implemented modules.

[0200] Similarly, the methods or routines described herein may be at least partially processor-implemented. For example, at least some of the operations of a method can be performed by one or more processors or processor-implemented hardware modules. The

performance of certain of the operations may be distributed among the one or more processors, not only residing within a single machine, but also deployed across a number of machines. In some example embodiments, the processor or processors may be located in a single location (e.g., within a home environment, an office environment or as a server farm), while in other embodiments the processors may be distributed across a number of locations.

[0201] The performance of certain of the operations may be distributed among the one or more processors, not only residing within a single machine, but also deployed across a number of machines. In some example embodiments, the one or more processors or processor-implemented modules may be located in a single geographic location (e.g., within a home environment, an office environment, or a server farm). In other example embodiments, the one or more processors or processor-implemented modules may be distributed across a number of geographic locations.

[0202] Unless specifically stated otherwise, discussions herein using words such as "processing," "computing," "calculating," "determining," "presenting," "displaying," or the like may refer to actions or processes of a machine (e.g., a computer) that manipulates or transforms data represented as physical (e.g., electronic, magnetic, or optical) quantities within one or more memories (e.g., volatile memory, non-volatile memory, or a combination thereof), registers, or other machine components that receive, store, transmit, or display information.

[0203] As used herein any reference to "one embodiment" or "an embodiment" means that a particular element, feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment. The appearances of the phrase "in one embodiment" in various places in the specification are not necessarily all referring to the same embodiment.

[0204] Some embodiments may be described using the expression "coupled" and "connected" along with their derivatives. For example, some embodiments may be described using the term "coupled" to indicate that two or more elements are in direct physical or electrical contact. The term "coupled," however, may also mean that two or more elements are not in direct contact with each other, but yet still co-operate or interact with each other. The embodiments are not limited in this context.

[0205] As used herein, the terms "comprises," "comprising," "includes," "including," "has," "having" or any other variation thereof, are intended to cover a non-exclusive inclusion. For

example, a process, method, article, or apparatus that comprises a list of elements is not necessarily limited to only those elements but may include other elements not expressly listed or inherent to such process, method, article, or apparatus. Further, unless expressly stated to the contrary, "or" refers to an inclusive or and not to an exclusive or. For example, a condition A or B is satisfied by any one of the following: A is true (or present) and B is false (or not present), A is false (or not present) and B is true (or present), and both A and B are true (or present).

[0206] In addition, use of the "a" or "an" are employed to describe elements and components of the embodiments herein. This is done merely for convenience and to give a general sense of the description. This description, and the claims that follow, should be read to include one or at least one and the singular also includes the plural unless it is obvious that it is meant otherwise.

[0207] This detailed description is to be construed as an example only and does not describe every possible embodiment, as describing every possible embodiment would be impractical, if not impossible. One could implement numerous alternative embodiments, using either current technology or technology developed after the filing date of this application.

Claims

What is claimed is:

1. A method for creating an overlay map on a digital image of a slide, the method comprising:

receiving the digital image;
separating the digital image into a plurality of tiles;
identifying the majority class of tissue visible within each tile in the plurality of tiles,
based on a multi-tile analysis.

2. The method of claim 1, further comprising:
generating a digital overlay drawing of an outer edge of each cell in the image.

3. The method of claim 2, wherein the digital overlay drawing is prepared at the resolution level of an individual pixel.

4. A method for tissue classification of a digital image of a slide, the method comprising:
receiving the digital image;
generating a digital overlay drawing of a tissue region in the digital image; and
displaying, in the digital overlay drawing, tiles over the tissue region and visually identifying predicted content of each tile, to generate a classification map classifying the digital image.

5. The method of claim 4, the method further comprising:
determining the predicted content for each tile using a classification model configured as a multi-resolution fully convolutional network, the multi-resolution fully convolution network configured to perform classification on digital images of different zoom levels.

6. The method of claim 4, the method further comprising:
using a classification model, for each tile, determining a classification from a plurality of classifications.

7. The method of claim 6, wherein the plurality of classifications include tumor epithelium, normal epithelium, immune, stroma, necrosis, blood, and fat.
8. A method for cell detection in a digital image of a slide, the method comprising:
 - receiving the digital image;
 - identifying cell objects in the digital image;
 - generating a digital overlay drawing for the digital image; and
 - displaying, in the digital overlay drawing, a polygon outlining each of the identified cell objects.
9. The method of claim 8, wherein the cell objects include lymphocyte cells and not lymphocyte cells.
10. The method of claim 8, wherein the cell objects comprise CD3, CD8, CD20, pancytokeratin, and smooth muscle actin.
11. A method for tissue classification of a digital image of a slide, the method comprising:
 - receiving the digital image;
 - segmenting the digital image into a plurality of tiles;
 - determining a predicted class for each tile;
 - identifying a plurality of cell objects in the digital image;
 - determining a predicted class for each of the plurality of cell objects; and
 - for each of the plurality of tiles that corresponds to one of the plurality of cell objects, assigning the tile the predicted class of the corresponding cell object in place of the predicted class of the tile.
12. The method of claim 11, further comprising:
 - storing, in a first file, for each tile, a tile position and the predicted class of the tile; and
 - storing, in a second file, for each cell object, a polygon outlining the cell object and the predicted class of the cell object.

13. The method of claim 11, further comprising:
generating a digital overlay drawing for the digital image, where the digital overlay drawing is a cell mask displaying a polygon around each cell object.
14. The method of claim 11, further comprising:
generating a digital overlay drawing for the digital image, where the digital overlay drawing is a histology mask displaying the plurality of tiles and the predict class for each tile.
15. The method of claim 11, further comprising:
generating a digital overlay drawing for the digital image; and
displaying, in the digital overlay drawing, the plurality of tiles and the predicted class for each tile that does not correspond to a cell object, and displaying, the plurality of cell objects and the predicted class of each cell object.
16. The method of claim 11, further comprising:
generating a digital overlay drawing for the digital image, wherein the digital overlay drawing includes the digital image; and
displaying the digital overlay drawing.
17. The method of claim 11, further comprising:
generating a digital overlay drawing for the digital image, wherein the digital overlay drawing includes a generated version of the digital image; and
displaying the digital overlay drawing.
18. The method of claim 11, further comprising:
generating a plurality of digital overlay drawings for the digital image, where each digital overlay drawing corresponds to a different predicted class; and
selectively displaying one of the plurality of digital overlay drawings.
19. The method of claim 11, further comprising:
generating a digital overlay drawing for the digital image, where the digital overlay drawing comprises percentages of predict classes corresponding to the digital image.
20. The method of claim 11, further comprising:

generating a digital overlay drawing for the digital image, where the digital overlay drawing comprises total counts of predict classes corresponding to the digital image.




FIG. 1




FIG. 2




FIG. 3A




FIG 3B




FIG. 3C

Receive digital image file →
Analyze metadata and retrieve image at known resolution →
(optional step) Downsample to create image with optimal resolution →
Identify and eliminate pixels that are too bright or too dark →
Identify and eliminate pixels with non-tissue objects

FIG 4




FIG. 5A






FIG. 5B



Prediction of center tile label based on cell and its surrounding context.

FIG. 5C

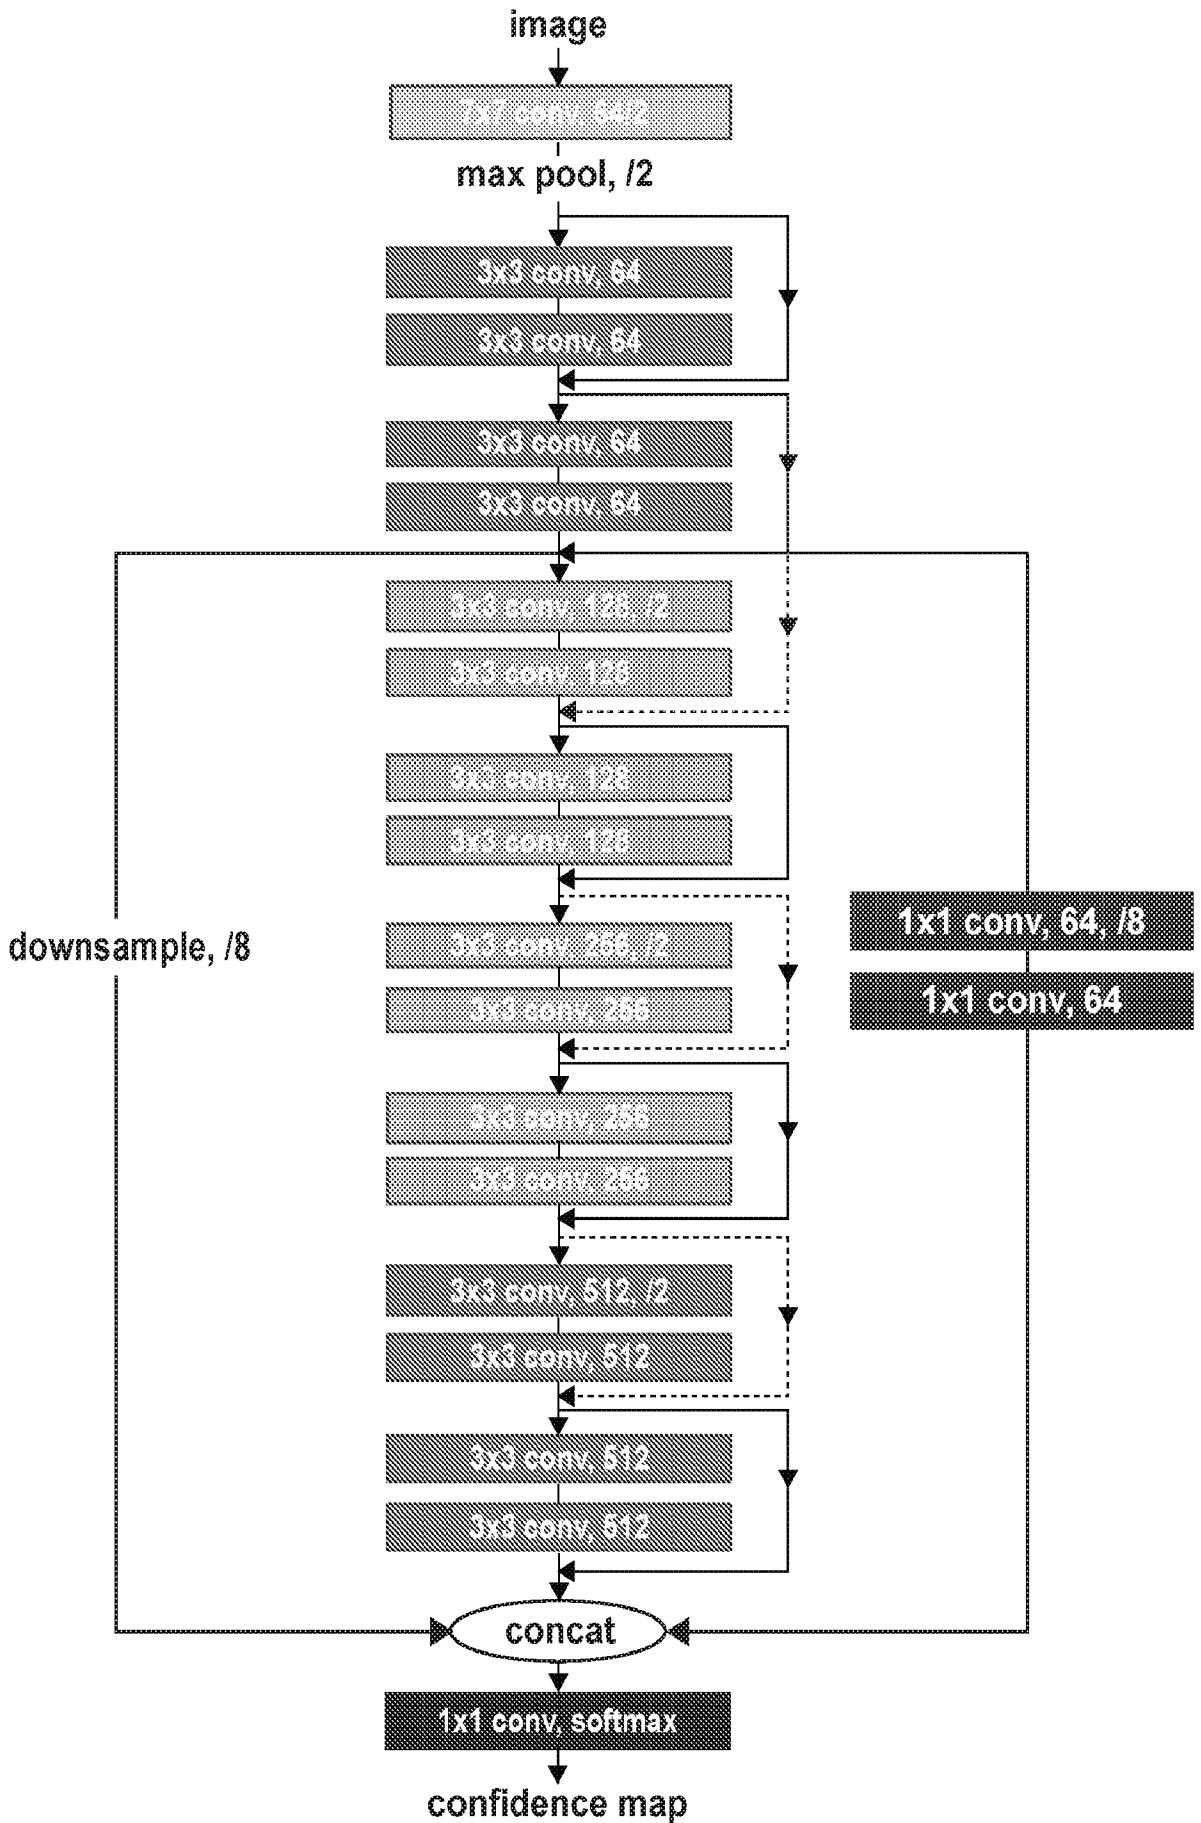


FIG. 6A

PhiNet

layer name	output size	sub layer
input layer	$(434+32N) \times (434+32N) \times 3$	
conv1	$(214+16N) \times (214+16N) \times 64$	
max pool	$(106+8N) \times (106+8N) \times 64$	
conv2_x	$(98+8N) \times (98+8N) \times 64$	Four 3x3 conv layers
conv3_x	$(42+4N) \times (42+4N) \times 128$	Four 3x3 conv layers
conv4_x	$(14+2N) \times (14+2N) \times 256$	Four 3x3 conv layers
conv5_x	$N \times N \times 512$	Four 3x3 conv layers
concat	$N \times N \times (512+64+64)$	
1x1 conv	$N \times N \times \#classes$	
softmax	$N \times N$	

ResNet-18

layer name	output size	sub layer
input layer	$224 \times 224 \times 3$	
conv1	$112 \times 112 \times 64$	
max pool	$56 \times 56 \times 64$	
conv2_x	$56 \times 56 \times 64$	Four 3x3 conv layers
conv3_x	$28 \times 28 \times 128$	Four 3x3 conv layers
conv4_x	$14 \times 14 \times 256$	Four 3x3 conv layers
conv5_x	$7 \times 7 \times 512$	Four 3x3 conv layers
avg. pool	$1 \times 1 \times 512$	
dense_layer	#classes	
softmax	#classes	

No padding

padding = 1 if filter_size = 3

Padding = 3 if filter_size = 7

FIG. 6B

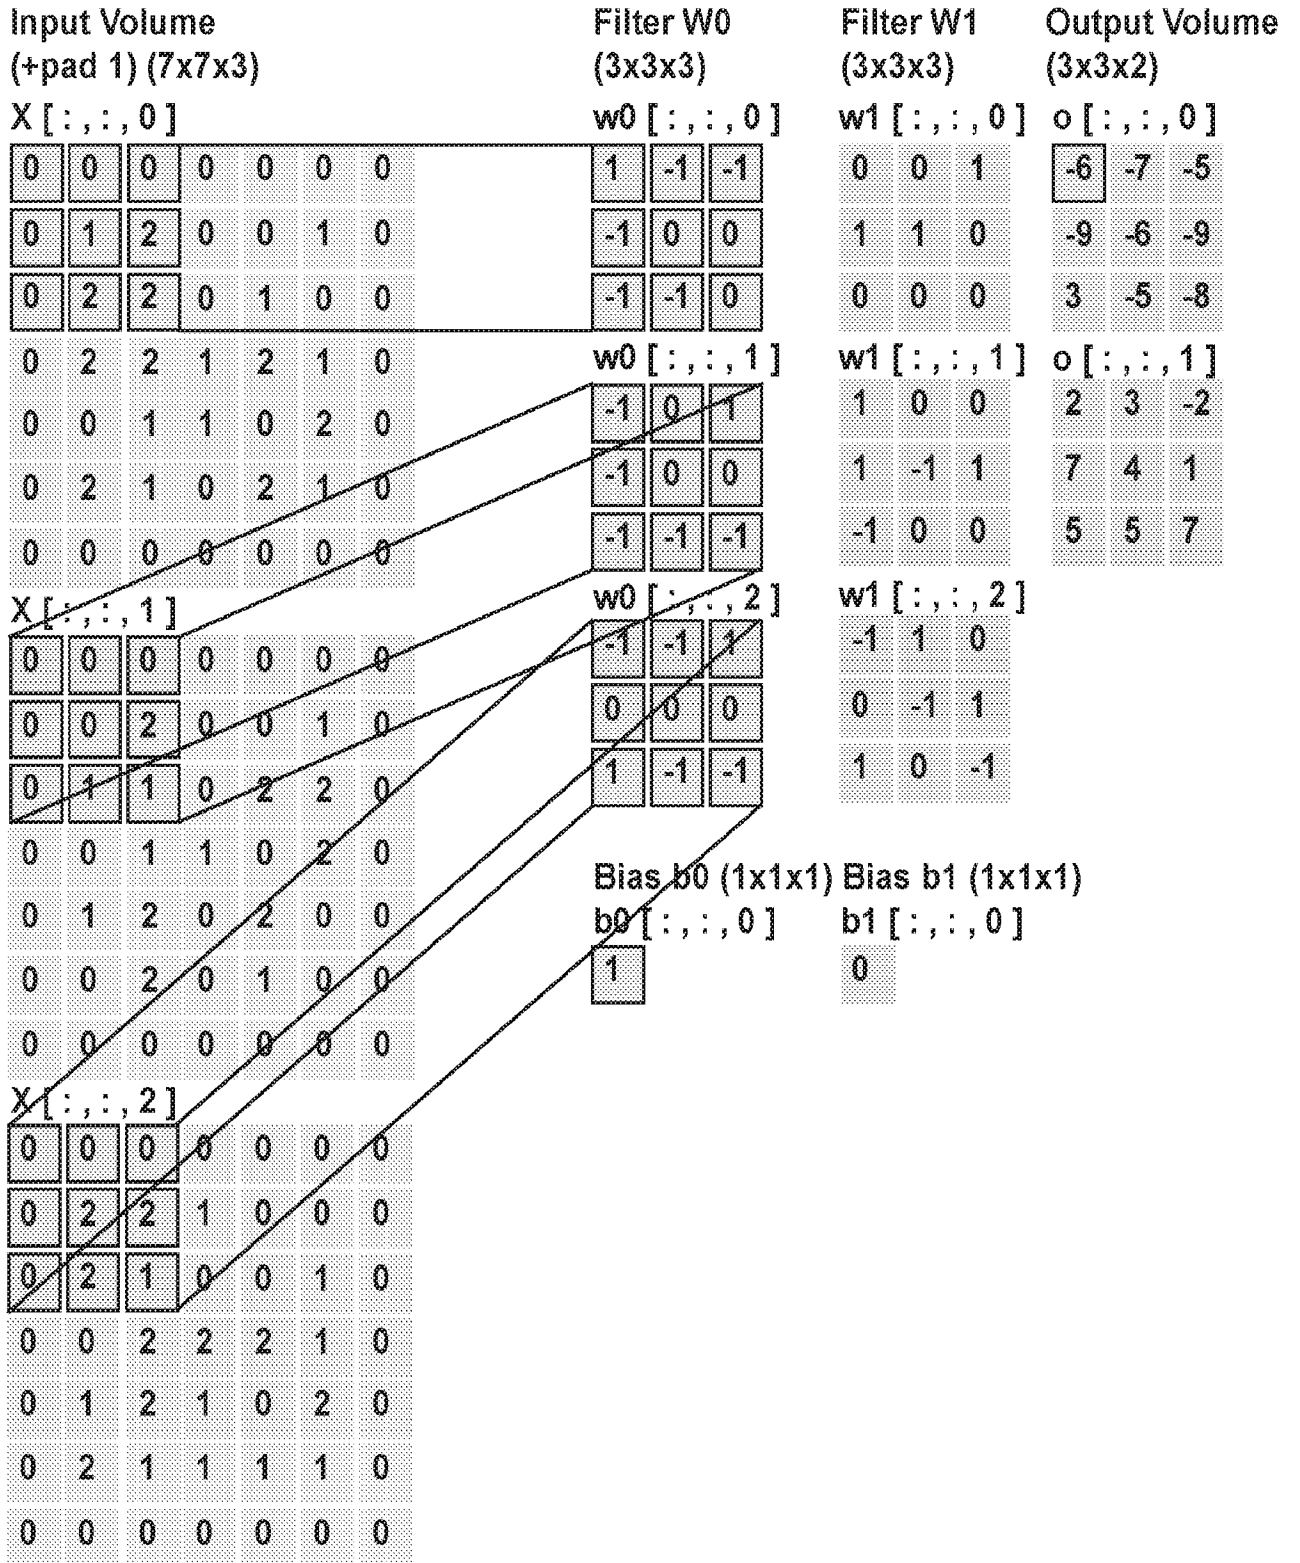
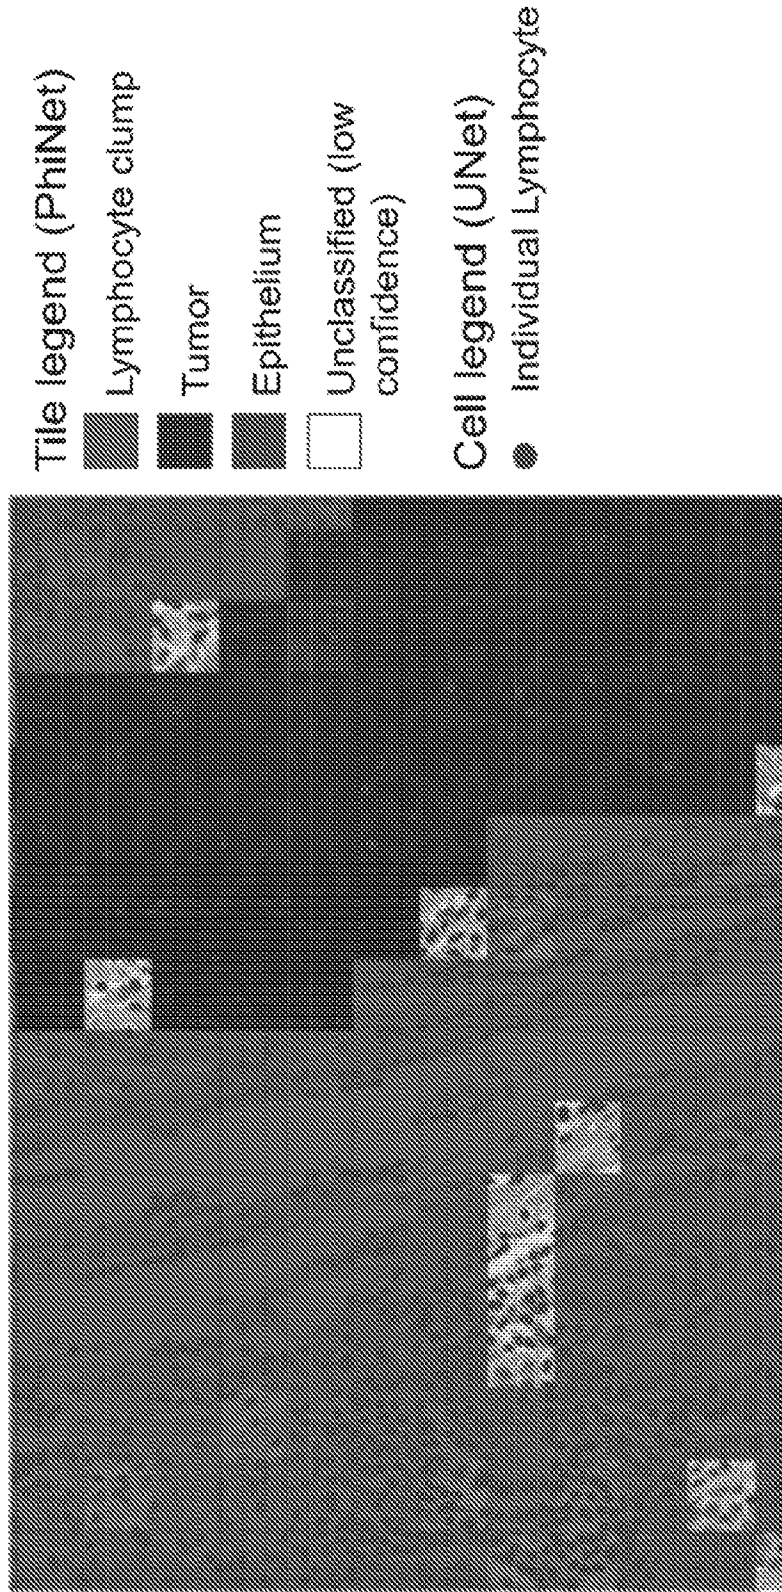


FIG. 6C

Union of tile classification (PhiNet) and
cell segmentation (UNet)



Detected lymphocytes (appearing as red dots) can be seen throughout tissue. The lymphocytes within the dark blue tumor tiles are tumor-infiltrating lymphocytes.

FIG. 7




FIG 8

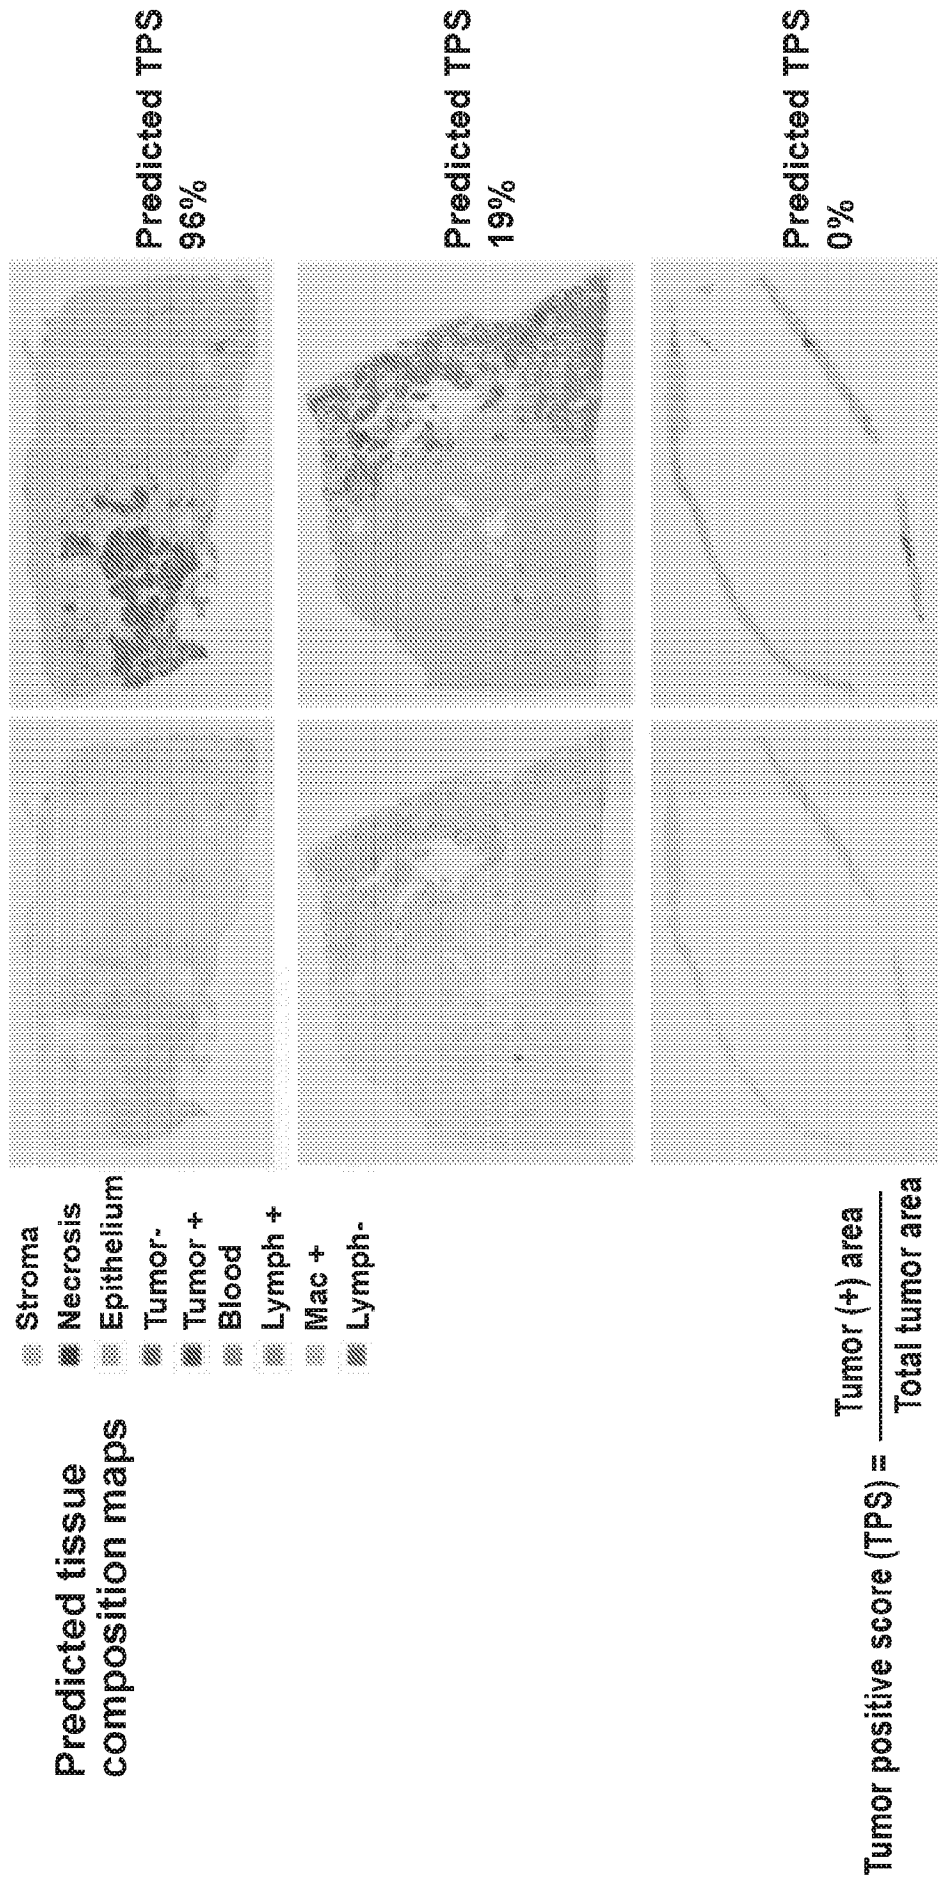


FIG. 9




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 19/69164

A. CLASSIFICATION OF SUBJECT MATTER
 IPC - G06K 9/00; G06T 7/00 (2020.01)
 CPC - G06T 7/0012; G06K 9/00127; G06K 9/0014; G06K 9/00147; G06T 7/0012; G06T 2207/10056;
 G06T 2207/20024

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 See Search History document

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
 See Search History document

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
 See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	US 2016/0063724 A1 (PATHXL LIMITED) 03 March 2016 (03.03.2016) Fig. 3, para [0025], [0026], [0028], [0029], [0033], [0034], [0036]-[0042], [0043], [0047], [0063], [0064], [0068], [0075], [0103], [0126]	1-4, 6-8, 11-14, 16 ----- 5, 9-10, 15, 17-20
Y	"Recurrent Multiresolution Convolutional Networks for VHR Image Classification" (Bergado et al.) 15 June 2018 (15.06.2018) <Downloaded from the Internet: https://arxiv.org/pdf/1806.05793.pdf > <Downloaded on: 15 March 2020> < DOI: 10.1109/TGRS.2018.2837357 >entire document, especially: pp 2 col left, para 2 pp 4 col left, para 1	5
Y	US 2014/0140607 A1 (Erjefalt) 22 May 2014 (22.05.2014) para [0078], [0113], [0190]	9-10
Y	US2018/0137394 A1 (KONINKLIJKE PHILIPS N. V.) 17 May 2018 (17.05.2018) para [0005], [0007]-[0010]	15
Y	US 2007/0216909 A1 (Everett et al.) 20 September 2007 (20.09.2007) entire document, especially: Fig. 3, para [0009], [0069], [0071]	17-18
Y	US 2014/0073922 (Mammone) 13 March 2014 (13.03.2014) entire document, especially: para [0051]	19-20

Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"D" document cited by the applicant in the international application	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"E" earlier application or patent but published on or after the international filing date	"&" document member of the same patent family
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 15 March 2020	Date of mailing of the international search report 09 APR 2020
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-8300	Authorized officer Lee Young Telephone No. PCT Helpdesk: 571-272-4300