



(51) International Patent Classification:

G06T 7/11 (2017.01) G06N 20/00 (2019.01)
G06T 3/40 (2006.01) G06N 3/08 (2006.01)
G06T 1/00 (2006.01)

(21) International Application Number:

PCT/US2020/051944

(22) International Filing Date:

22 September 2020 (22.09.2020)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/905,232 24 September 2019 (24.09.2019) US

(71) Applicant: **APPLIED MATERIALS, INC.** [US/US];
3050 Bowers Avenue, Santa Clara, California 95054 (US).

(72) Inventors: **JHA, Sumit Kumar**; Behind Vivek Memorial High School, New Nagar Bandha Gadi, Bariatu, Ranchi, Jharkhand 834009 (IN). **SISTA, Aditya**; H-906, Purva Fountain Square, Marathalli, Bangalore 560037 (IN). **RAGHUNATHAN, Ganesh Kumar Mohanur**; C-001, Sankalp Cherry Blossom, 19/4 Edgha Road, Varthur, Bangalore 560087 (IN). **KUMAR, Ubhay**; House No. 3, Ward No. 4, Bhuswar, Bibhutipur, Samastipur, Bihar 848220

(IN). **SAPRE, Kedar**; 5189 Paseo Olivos, San Jose, California 95130 (US).

(74) Agent: **GOREN, David J.** et al.; Fish & Richardson P.C., P.O. Box 1022, Minneapolis, Minnesota 55440-1022 (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, IT, 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, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, 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, 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).

(54) Title: INTERACTIVE TRAINING OF A MACHINE LEARNING MODEL FOR TISSUE SEGMENTATION

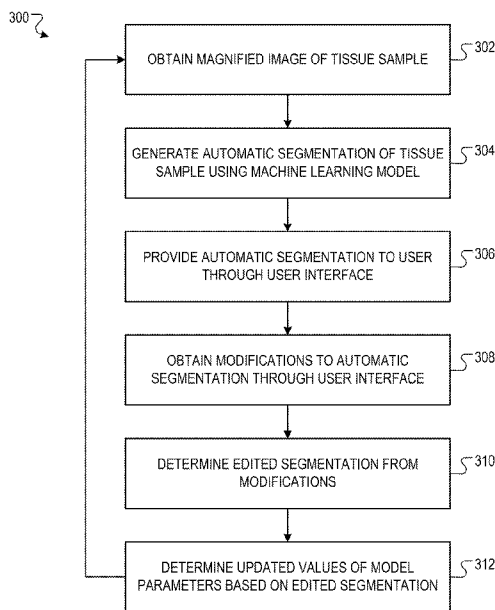


FIG. 3

(57) Abstract: Methods, systems, and apparatus, including computer programs encoded on a computer storage medium, for training a machine learning model to segment magnified images of tissue samples. In one aspect, a method comprises obtaining a magnified image of a tissue sample; processing an input comprising: the image, features derived from the image, or both, in accordance with current values of model parameters of a machine learning model to generate an automatic segmentation of the image into a plurality of tissue classes; providing, to a user through a user interface, an indication of: (i) the image, and (ii) the automatic segmentation of the image; determining an edited segmentation of the image, comprising applying modifications specified by the user to the automatic segmentation of the image; and determining updated values of the model parameters of the machine learning model based the edited segmentation of the image.



Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

INTERACTIVE TRAINING OF A MACHINE LEARNING MODEL FOR TISSUE SEGMENTATION

BACKGROUND

[0001] This specification relates to processing magnified images of tissue samples using machine learning models.

[0002] Machine learning models receive an input and generate an output, e.g., a predicted output, based on the received input. Some machine learning models are parametric models and generate the output based on the received input and on values of the parameters of the model.

[0003] Some machine learning models are deep models that employ multiple layers of models to generate an output for a received input. For example, a deep neural network is a deep machine learning model that includes an output layer and one or more hidden layers that each apply a non-linear transformation to a received input to generate an output.

SUMMARY

[0004] This specification describes a system implemented as computer programs on one or more computers in one or more locations for segmenting magnified images of tissue samples into respective tissue classes.

[0005] According to a first aspect there is provided a method performed by one or more data processing apparatus for training a machine learning model to segment magnified images of tissue samples. The method includes repeatedly performing operations including obtaining a magnified image of a tissue sample. An input including: the image, features derived from the image, or both, is processed in accordance with current values of model parameters of a machine learning model to generate an automatic segmentation of the image into a set of tissue classes, where the automatic segmentation specifies a respective tissue class corresponding to each pixel of the image. An indication of: (i) the image, and (ii) the automatic segmentation of the image, is provided to a user through a user interface. An input specifying one or more modifications to the automatic segmentation of the image is obtained from the user through the user interface. An edited segmentation of the image is determined, including applying the modifications specified by the user to the automatic segmentation of the image. Updated values of the model parameters of the machine learning model are determined based the edited segmentation of the image.

[0006] In some implementations, the set of tissue classes include cancerous tissue and non-cancerous tissue.

[0007] In some implementations, the image is a whole slide image (WSI) captured using a microscope.

[0008] In some implementations, the machine learning model includes a neural network model.

[0009] In some implementations, the neural network model is a convolutional neural network model including one or more convolutional neural network layers.

[0010] In some implementations, providing, to the user through the user interface, an indication of: (i) the image, and (ii) the automatic segmentation of the image includes providing, to the user through the user interface, a visualization that depicts the automatic segmentation overlaid on the image.

[0011] In some implementations, the visualization that depicts the automatic segmentation overlaid on the image indicates the respective tissue class corresponding to each pixel of the image.

[0012] In some implementations, each modification to the automatic segmentation of the image indicates, for one or more pixels of the image, a change to the respective tissue class specified for the pixel by the automatic segmentation. Applying the modifications specified by the user to the automatic segmentation of the image includes changing the tissue classes corresponding to one or more pixels of the image in accordance with the modifications specified by the user.

[0013] In some implementations, determining updated values of the model parameters based on the edited segmentation of the image includes: determining gradients of an objective function that characterizes a similarity between: (i) the automatic segmentation of the image, and (ii) the edited segmentation of the image; and adjusting the values of the model parameters using the gradients.

[0014] In some implementations, the method further includes determining an expertise score that characterizes a predicted skill of the user in reviewing and editing tissue segmentations.

[0015] According to another aspect, there is provided a system including: one or more computers; and one or more storage devices communicatively coupled to the one or more computers, where the one or more storage devices store instructions that, when executed by the one or more computers, cause the one or more computers to perform the operations of the method of the first aspect.

[0016] According to a third aspect, there are provided one or more non-transitory computer storage media storing instructions that when executed by one or more computers cause the one or more computers to perform the operations of the method of the first aspect.

[0017] Particular embodiments of the subject matter described in this specification can be implemented so as to realize one or more of the following advantages. The segmentation system described in this specification enables a user (e.g., a pathologist) to work in tandem with a machine learning model to segment magnified images of tissue samples into (target) tissue classes in a manner that is both time-efficient and highly accurate. This specification describes techniques for computing an “expertise” score for a user that characterizes the predicted skill of the user in manually reviewing and editing segmentations (i.e., of the target tissue classes). The expertise scores can be used to improve the performance of the segmentation system. For example, the expertise scores can be used to improve the quality of the training data used to train the segmentation system, e.g., by determining whether to include a segmentation generated by a user in the training data based on the expertise score of the user.

[0018] The details of one or more embodiments of the subject matter of this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0019] FIG. 1 shows an example segmentation system.

[0020] FIG. 2 is an illustration of an example segmentation of a magnified image of a tissue sample.

[0021] FIG. 3 is a flow diagram of an example process for repeatedly training a machine learning model to segment magnified images of tissue samples.

[0022] FIG. 4 is a flow diagram of an example process for determining an expertise score that characterizes the predicted skill of a user in reviewing and editing segmentations of target tissue classes in magnified images of tissue samples.

[0023] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0024] This specification describes a segmentation system for segmenting magnified images of tissue samples (e.g., that are generated using a microscope, e.g., an optical microscope) into respective tissue classes. More specifically, the segmentation system can process a magnified image of a tissue sample to identify a respective (target) tissue class corresponding to each pixel of the image. The (target) tissue class of a pixel in the image characterizes the type of tissue in the portion of the tissue sample corresponding to the pixel.

[0025] As used throughout this document, a “microscope” can refer to any system that can generate magnified images of a sample, e.g., using a 1-D array of photodetectors, or using a 2-D array of charge-coupled devices (CCDs).

[0026] The segmentation system can be configured to segment images into any appropriate set of tissue classes. In one example, the segmentation system may segment images into cancerous tissue and non-cancerous tissue. In another example, the segmentation system may segment images into: healthy tissue, cancerous tissue, and necrotic tissue. In another example, the segmentation system may segment images into: muscle tissue, nervous tissue, connective tissue, epithelial tissue, and “other” tissue. The segmentation system can be used in any of a variety of settings, e.g., to segment magnified images of tissue samples that are obtained from patients through biopsy procedures. The tissue samples can be samples of any appropriate sort of tissue, e.g., prostate tissue, breast tissue, liver tissue, or kidney tissue. The segmentations generated by the segmentation system can be used for any of a variety of purposes, e.g., to characterize the presence or extent of disease (e.g., cancer).

[0027] Manually segmenting a single magnified image of a tissue sample may be a challenging task that consumes hours of time, e.g., as a result of the high-dimensionality of the image, which can have on the order of 10^{10} pixels. On the other hand, a machine learning model can be trained to automatically segment magnified images of tissue samples in considerably less time (e.g., in seconds or minutes, e.g., 10-30 minutes). However, it may be difficult to train a machine learning model to achieve a level of accuracy that would be considered acceptable for certain practical applications, e.g., identifying cancerous tissue in biopsy samples. In particular, the microscopic appearance of tissue can be highly complex and variable due to factors that are both intrinsic to the tissue (e.g., the type and stage of the disease present in tissue) and extrinsic to the tissue (e.g., how the microscope is calibrated and the procedure used to stain the tissue). This makes it hard to aggregate a set of labeled training data (i.e., for training a machine learning model) that is sufficiently large to capture the full scope of possible variations in the microscopic appearance of tissue.

[0028] The segmentation system described in this specification enables a user (e.g., a pathologist) to work in tandem with a machine learning model to segment tissue samples in a manner that is both time-efficient and highly accurate. To segment an image, the machine learning model first generates an automatic segmentation of the image which is subsequently provided to the user through a user interface that enables the user to review and manually edit the automatic segmentation as necessary. The “edited” segmentation is provided by the segmentation system as an output, and is also used to update the parameter values of the

machine learning model (e.g., immediately or at a subsequent time point) to cause it to generate segmentations that more closely match those of the user.

[0029] In this manner, rather than being trained once on a static and limited set of training data (as in some conventional systems), the machine learning model continually learns and adapts its parameter values based on the feedback being provided by the user through the edited segmentations. Moreover, rather than being required to segment an image from scratch, the user can start from the automatic segmentation generated by the machine learning model, and may be required to make fewer corrections to the automatic segmentations over time as the machine learning model continually improves.

[0030] The term “tissue” here refers to a group of cells of similar structure and function, as opposed to individual cells. The color, texturing, and similar image properties of tissues are significantly different from those of individual cells, so image processing techniques applicable to cell classification often are not applicable to segmenting images of tissue samples and classifying those segments.

[0031] These features and other features are described in more detail below.

[0032] FIG. 1 shows an example segmentation system 100. The segmentation system 100 is an example of a system implemented as computer programs on one or more computers in one or more locations in which the systems, components, and techniques described below are implemented.

[0033] The segmentation system 100 is configured to process a magnified image 102 of a tissue sample to generate a segmentation 104 of the image 102 into respective tissue classes, e.g., cancerous and non-cancerous tissue classes.

[0034] The image 102 may be, e.g., a whole slide image (WSI) of a tissue sample mounted on a microscope slide, where the WSI is generated using an optical microscope and captured using a digital camera. The image 102 can be represented in any of a variety of ways, e.g., as a two-dimensional (2-D) array of pixels, where each pixel is associated with a vector of numerical values characterizing the appearance of the pixel, e.g., a 3-D vector defining the red-green-blue (RGB) color of the pixel. The array of pixels representing the image 102 may have a dimensionality on the order of, e.g., $10^5 \times 10^5$ pixels, and may occupy several gigabytes (GB) of memory. The system 100 may receive the image 102 in any of a variety of ways, e.g., as an upload from a user of the system using a user interface made available by the system 100.

[0035] The machine learning model 106 is configured to process the image 102, features derived from the image 102, or both, in accordance with current values of a set of model parameters 108 to generate an automatic segmentation 110 of the image 102 that specifies a

respective tissue class corresponding to each pixel of the image 102. The machine learning model 106 may be, e.g., a neural network model, a random forest model, a support vector machine model, or a linear model. In one example, the machine learning model may be a convolutional neural network having an input layer that receives the image 102, a set of convolutional layers that process the image to generate alternative representations of the image at progressively higher levels of abstraction, and a soft-max output layer. In another example, the machine learning model may be a random forest model that is configured to process a respective feature representation of each pixel of the image 102 to generate an output that specifies a tissue class for the pixel. In this example, a feature representation of a pixel refers to an ordered collection of numerical values (e.g., a vector of numerical values) that characterizes the appearance of the pixel. The feature representation may be generated using, e.g., histogram of oriented gradient (HOG) features, speeded up robust features (SURF), or scale-invariant feature transform (SIFT) features.

[0036] The model parameters 108 are a collection of numerical values that are learned during training of the machine learning model 106 and which specify the operations performed by the machine learning model 106 to generate an automatic segmentation 110 of the image 102. For example, if the machine learning model 106 is a neural network, the model parameters 108 may specify the weight values of each layer of the neural network, e.g., the weight values of the convolutional filters of each convolutional layer of the neural network. (The weight values for a given layer of the neural network may refer to the values associated with the connections between neurons of the given layer and neurons in the preceding layer of the neural network). As another example, if the machine learning model 106 is a random forest model, the model parameters 108 may specify the parameter values of the respective splitting function used at each node of each decision tree of the random forest. As another example, if the machine learning model 106 is a linear model, the model parameters 108 may specify the coefficients of the linear model.

[0037] The system 100 displays the image 102 and the automatic segmentation 110 of the image on a display device of a user interface 112. For example, the system 100 may display a visualization that depicts the automatic segmentation 110 overlaid onto the image 102, as illustrated with reference to FIG. 2. The user interface 112 may have any appropriate sort of display device, e.g., a liquid-crystal display (LCD).

[0038] The user interface 112 enables a user of the system (e.g., a pathologist) to view the image 102 and the automatic segmentation 110, and to edit the automatic segmentation 110 as necessary by specifying one or more modifications to the automatic segmentation 110.

Modifying the automatic segmentation 110 refers to changing the tissue class specified by the automatic segmentation 110 to a different tissue class for one or more pixels of the image 102. Generally, the user may edit the automatic segmentation 110 to correct any errors in the automatic segmentation 110. For example, the user interface 112 may enable the user to “deselect” a region of the image that is specified by the automatic segmentation as having a certain tissue class (e.g., cancerous tissue) by re-labeling the region as having a default tissue class (e.g., non-cancerous tissue). As another example, the user interface 112 may enable the user to “select” a region of the image and label the region as having a particular tissue class (e.g., cancerous tissue). As another example, the user interface 112 may enable the user to change the region of the image labelled as having a particular tissue class. The change in a region can be performed, e.g., by dragging corners of a polygon surrounding the region.

[0039] The user may interact with the user interface 112 to edit the automatic segmentation 110 in any of a variety of ways, e.g., using a computer mouse, a touch screen, or both. For example, to select a region of the image and label the region as having a tissue class, the user may use a cursor to draw a closed loop around the region of the image, and then select the desired tissue class from a drop down menu. The user may indicate that editing of the automatic segmentation is complete by providing an appropriate input to the user interface (e.g., clicking a “Finish” button), at which point the edited segmentation 114 (i.e., that has been reviewed and potentially modified by the user) is provided as an output. For example, the output segmentation 104 may be stored in a medical records data store in association with a patient identifier.

[0040] In addition to providing the edited segmentation 114 as an output, the system 100 may also use the edited segmentation 114 to generate a training example that specifies: (i) the image, and (ii) the edited segmentation 114, and store the training example in a set of training data 116. Generally, the training data 116 stores multiple training examples (i.e., that each specify a respective image and an edited segmentation), and may be continually augmented over time as users generate edited segmentations of new images. The system 100 uses a training engine 118 to repeatedly train the machine learning model 106 on the training data 116 by updating the model parameters 108 to encourage the machine learning model 106 to generate automatic segmentations that match the edited segmentations specified by the training data 116.

[0041] The training engine 118 may train the machine learning model 106 on the training data 116 whenever a training criterion is satisfied. For example, the training engine 118 may train the machine learning model 106 each time a predefined number of new training examples are added to the training data 116. As another example, the training engine 118 may train the

machine learning model 106 each time the machine learning model 106 generates an automatic segmentation 110 that differs substantially from the corresponding edited segmentation 114 that is specified by the user. In this example, the training engine 118 may use the substantial difference between the automatic segmentation 110 and the edited segmentation 114 as a cue that the machine learning model 106 failed to correctly segment an image and should be trained to avoid repeating the errors. The training engine 118 may determine that two segmentations are substantially different if a similarity measure between the segmentations (e.g., a Jaccard index similarity measure) does not satisfy a predefined threshold.

[0042] The manner in which the training engine 118 trains the machine learning model 106 on the training data 116 depends on the form of the machine learning model 106. In some cases, the training engine 118 may train the machine learning model 106 by determining an adjustment to the current values of the model parameters 108. In other cases, the training engine 118 may start by initializing the model parameters 108 to default values each time the machine learning model 106 is trained, e.g., values that are sampled from a predefined probability distribution, e.g., a standard Normal distribution.

[0043] Take, as an example, a case where the machine learning model 106 is a neural network model, and the training engine 118 trains the neural network model using one or more iterations of stochastic gradient descent. In this example, at each iteration, the training engine 118 selects a “batch” (set) of training examples from the training data 116, e.g., by randomly selecting a predefined number of training examples. The training engine 118 processes the image 102 from each selected training example using the machine learning model 106 in accordance with the current values of the model parameters 108, to generate a corresponding automatic segmentation. The training engine 118 determines gradients of an objective function with respect to the model parameters 108, where the objective function measures a similarity between: (i) the automatic segmentations generated by the machine learning model 106, and (ii) the edited segmentations specified by the training examples. The training engine 118 then uses the gradients of the objective function to adjust the current values of the model parameters 108 of the machine learning model 106. The objective function may be, e.g., a pixel-wise cross-entropy objective function, the training engine 118 may determine the gradients using backpropagation techniques, and the training engine 118 may adjust the current values of the model parameters 108 using any appropriate gradient descent technique, e.g., Adam or RMSprop.

[0044] Optionally, the training engine 118 may preferentially train the machine learning model 106 on training examples that were generated more recently, i.e., rather than treating each

training example equally. For example, the training engine 118 may train the machine learning model 106 on training examples that are sampled from the training data 116, where training examples that were generated more recently have a higher likelihood of being sampled than older training examples. Preferentially training the machine learning model 106 on training examples that were generated more recently can enable the machine learning model 106 to focus on learning from newer training examples while maintaining the insights gained from older training examples.

[0045] Generally, the system 100 trains the machine learning model 106 to generate automatic segmentations 110 that match edited segmentations 114 specified by users of the system 100, e.g., pathologists. However, certain users may be more skilled than others in reviewing and editing automatic segmentations generated by the machine learning model 106 for accuracy. For example, a more experienced pathologist may achieve a higher accuracy in reviewing and editing segmentations of complex and ambiguous tissue samples than a more junior pathologist. In some implementations, each user of the system 100 may be associated with an “expertise” score that characterizes the predicted skill of the user in reviewing and editing segmentations. In these implementations, the machine learning model 106 may be trained using only edited segmentations that are generated by users with a sufficiently high expertise score, e.g., an expertise score that satisfies a predetermined threshold. An example process for determining an expertise score for a user is described in more detail with reference to FIG. 4.

[0046] Determining whether to train the machine learning model 106 on an edited segmentation based on the expertise score of the user that generated the segmentation can improve the performance of the machine learning model 106 by improving the quality of the training data. Optionally, users of the system 100 may be compensated (e.g., financially or otherwise) for providing segmentations that are used to train the machine learning model 106. In one example, the amount of compensation provided to a user may depend on the expertise score of the user, and users with higher expertise scores may receive more compensation than users with lower expertise scores.

[0047] Optionally, the system 100 may be a distributed system where various components of the system are implemented remotely from one another and communicate over a data communication network, e.g., the Internet. For example, the user interface 112 (including the display device) may be implemented in a clinical environment (e.g., a hospital), while the machine learning model 106 and the training engine 118 may be implemented in a remote data center.

[0048] Optionally, a user of the system 100 may be provided the option of disabling the machine learning model 106. If this option is selected, the user can load images 102 and manually segment them without use of the machine learning model 106.

[0049] FIG. 2 is an illustration of a magnified image 200 of a tissue sample, where the regions 202-A-E (and the portion of the image outside of the regions 202-A-E) correspond to respective tissue classes.

[0050] FIG. 3 is a flow diagram of an example process 300 for repeatedly training a machine learning model to segment magnified images of tissue samples. For convenience, the process 300 will be described as being performed by a system of one or more computers located in one or more locations. For example, a segmentation system, e.g., the segmentation system 100 of FIG. 1, appropriately programmed in accordance with this specification, can perform the process 300.

[0051] The system obtains a magnified image of a tissue sample (302). For example, the image may be a magnified whole slide image of a biopsy sample from a patient that is generated using a microscope.

[0052] The system processes an input including: (i) the image, (ii) features derived from the image, or (iii) both, in accordance with current values of the model parameters of the machine learning model to generate an automatic segmentation of the image into a set of (target) tissue classes (304). The automatic segmentation specifies a respective tissue class corresponding to each pixel of the image. The tissue classes may include cancerous tissue and non-cancerous tissue. The machine learning model may be a neural network model, e.g., a convolutional neural network model with one or more convolutional layers.

[0053] The system provides an indication of: (i) the image, and (ii) the automatic segmentation of the image, to the user through a user interface (306). For example, the system may provide a visualization that depicts the automatic segmentation overlaid on the image through a display device of the user interface. The visualization of the automatic segmentation overlaid on the image may indicate the predicted tissue type of each of the regions delineated by the automatic segmentation. For example, the visualization may indicate the predicted tissue type of a region by colorizing the region based on the tissue type, e.g., cancerous tissue is colored red, while non-cancerous tissue is colored green.

[0054] The system obtains an input specifying one or more modifications to the automatic segmentation of the image from the user through the user interface (308). Each modification to the automatic segmentation may indicate, for one or more pixels of the image, a change to the respective tissue class specified for the pixel by the automatic segmentation.

[0055] The system determines an edited segmentation of the image (310). For example, the system may determine the edited segmentation of the image by applying the modifications specified by the user through the user interface to the automatic segmentation of the image.

[0056] The system determines updated values of the model parameters of the machine learning model based on the edited segmentation of the image (312). For example, the system may determine gradients of an objective function that characterizes a similarity between: (i) the automatic segmentation of the image, and (ii) the edited segmentation of the image, and then adjust the values of the model parameters using the gradients. In some cases, the system may determine updated values of the model parameters of the machine learning model only in response to determining that a training criterion is satisfied, e.g., that a predefined number of new edited segmentations have been generated since the last time the model parameters were updated. After determining updated values of the model parameters, the system may return to step 302. If the training criterion is not satisfied, the system may return to step 302 without training the machine learning model.

[0057] FIG. 4 is a flow diagram of an example process 400 for determining an expertise score that characterizes the predicted skill of a user in reviewing and editing segmentations of magnified images of tissue samples. For convenience, the process 400 will be described as being performed by a system of one or more computers located in one or more locations. For example, a segmentation system, e.g., the segmentation system 100 of FIG. 1, appropriately programmed in accordance with this specification, can perform the process 400.

[0058] The system obtains one or more tissue segmentations that were generated by the user (402). Each tissue segmentation corresponds to a magnified image of a tissue sample and specifies a respective tissue class for each pixel of the image. In some implementations, the user may have performed the segmentations from scratch, e.g., without the benefit of starting from automatic segmentations generated by a machine learning model.

[0059] The system obtains one or more features characterizing the medical experience of the user, e.g., in the field of pathology (404). For example, the system may obtain features characterizing one or more of: the number of years of experience of the user in the field of pathology, the number of academic publications of the user in the field of pathology, the number of citations of the academic publications of the user in the field of pathology, the academic performance of the user (e.g., in medical school), and the position currently held by the user (e.g., attending physician).

[0060] The system determines the expertise score for the user based on: (i) the tissue segmentations generated by the user, and (ii) the features characterizing the medical experience

of the user (406). For example, the system may determine the expertise score as a function (e.g., a linear function) of: (i) a similarity measure between the segmentations generated by the user and corresponding “gold standard” segmentations of the same images, and (ii) the features characterizing the medical experience of the user. A gold standard segmentation of an image may be a segmentation that is generated by a user (e.g., a pathologist) that is recognized as having a high level of expertise in performing tissue segmentations. A similarity measure between two segmentations of an image can be evaluated using, e.g., a Jaccard index. The expertise score for a user may be represented as a numerical value, e.g., in the range [0,1].

[0061] The system provides the expertise score for the user (408). For example, the system may provide the expertise score for the user for use in determining whether segmentations generated by the user should be included in training data used to train a machine learning model to perform automatic tissue sample segmentations. In this example, segmentations generated by a user may be included in the training data only if, e.g., the expertise score for the user satisfies a threshold. In another example, the system may provide the expertise score for the user for use in determining how the user should be compensated (e.g., financially or otherwise) for providing tissue sample segmentations, e.g., where having a higher expertise score may result in higher compensation.

[0062] This specification uses the term “configured” in connection with systems and computer program components. For a system of one or more computers to be configured to perform particular operations or actions means that the system has installed on it software, firmware, hardware, or a combination of them that in operation cause the system to perform the operations or actions. For one or more computer programs to be configured to perform particular operations or actions means that the one or more programs include instructions that, when executed by data processing apparatus, cause the apparatus to perform the operations or actions.

[0063] Embodiments of the subject matter and the functional operations described in this specification can be implemented in digital electronic circuitry, in tangibly-embodied computer software or firmware, in computer hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as one or more computer programs, i.e., one or more modules of computer program instructions encoded on a tangible non-transitory storage medium for execution by, or to control the operation of, data processing apparatus. The computer storage medium can be a machine-readable storage device, a machine-readable storage substrate, a random or serial access memory device, or a combination of one or more of them. Alternatively or in addition, the

program instructions can be encoded on an artificially-generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal, that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus.

[0064] The term “data processing apparatus” refers to data processing hardware and encompasses all kinds of apparatus, devices, and machines for processing data, including by way of example a programmable processor (e.g., central processing unit (CPU), graphics processing unit (GPU)), a computer, or multiple processors or computers. The apparatus can also be, or further include, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit). The apparatus can optionally include, in addition to hardware, code that creates an execution environment for computer programs, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

[0065] A computer program, which may also be referred to or described as a program, software, a software application, an app, a module, a software module, a script, or code, can be written in any form of programming language, including compiled or interpreted languages, or declarative or procedural languages; and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A program may, but need not, correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data, e.g., one or more scripts stored in a markup language document, in a single file dedicated to the program in question, or in multiple coordinated files, e.g., files that store one or more modules, sub-programs, or portions of code. A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a data communication network.

[0066] In this specification the term “engine” is used broadly to refer to a software-based system, subsystem, or process that is programmed to perform one or more specific functions. Generally, an engine will be implemented as one or more software modules or components, installed on one or more computers in one or more locations. In some cases, one or more computers will be dedicated to a particular engine; in other cases, multiple engines can be installed and running on the same computer or computers.

[0067] The processes and logic flows described in this specification can be performed by one or more programmable computers executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can

also be performed by special purpose logic circuitry, e.g., an FPGA or an ASIC, or by a combination of special purpose logic circuitry and one or more programmed computers.

[0068] Computers suitable for the execution of a computer program can be based on general or special purpose microprocessors or both, or any other kind of central processing unit. Generally, a central processing unit will receive instructions and data from a read-only memory or a random access memory or both. The essential elements of a computer are a central processing unit for performing or executing instructions and one or more memory devices for storing instructions and data. The central processing unit and the memory can be supplemented by, or incorporated in, special purpose logic circuitry. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto-optical disks, or optical disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device, e.g., a universal serial bus (USB) flash drive, to name just a few.

[0069] Computer-readable media suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks.

[0070] To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for example, by sending web pages to a web browser on a user's device in response to requests received from the web browser. Also, a computer can interact with a user by sending text messages or other forms of message to a personal device, e.g., a smartphone that is running a messaging application, and receiving responsive messages from the user in return.

[0071] Data processing apparatus for implementing machine learning models can also include, for example, special-purpose hardware accelerator units for processing common and compute-intensive parts of machine learning training or production, i.e., inference, workloads.

[0072] Machine learning models can be implemented and deployed using a machine learning framework, e.g., a TensorFlow framework, a Microsoft Cognitive Toolkit framework, an Apache Singa framework, or an Apache MXNet framework.

[0073] Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back-end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front-end component, e.g., a client computer having a graphical user interface, a web browser, or an app through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back-end, middleware, or front-end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network (LAN) and a wide area network (WAN), e.g., the Internet.

[0074] The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other. In some embodiments, a server transmits data, e.g., an HTML page, to a user device, e.g., for purposes of displaying data to and receiving user input from a user interacting with the device, which acts as a client. Data generated at the user device, e.g., a result of the user interaction, can be received at the server from the device.

[0075] While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any invention or on the scope of what may be claimed, but rather as descriptions of features that may be specific to particular embodiments of particular inventions. Certain features that are described in this specification in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially be claimed as such, one or more features from a claimed combination can in some

cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

[0076] Similarly, while operations are depicted in the drawings and recited in the claims in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system modules and components in the embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

[0077] Particular embodiments of the subject matter have been described. Other embodiments are within the scope of the following claims. For example, the actions recited in the claims can be performed in a different order and still achieve desirable results. As one example, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In some cases, multitasking and parallel processing may be advantageous.

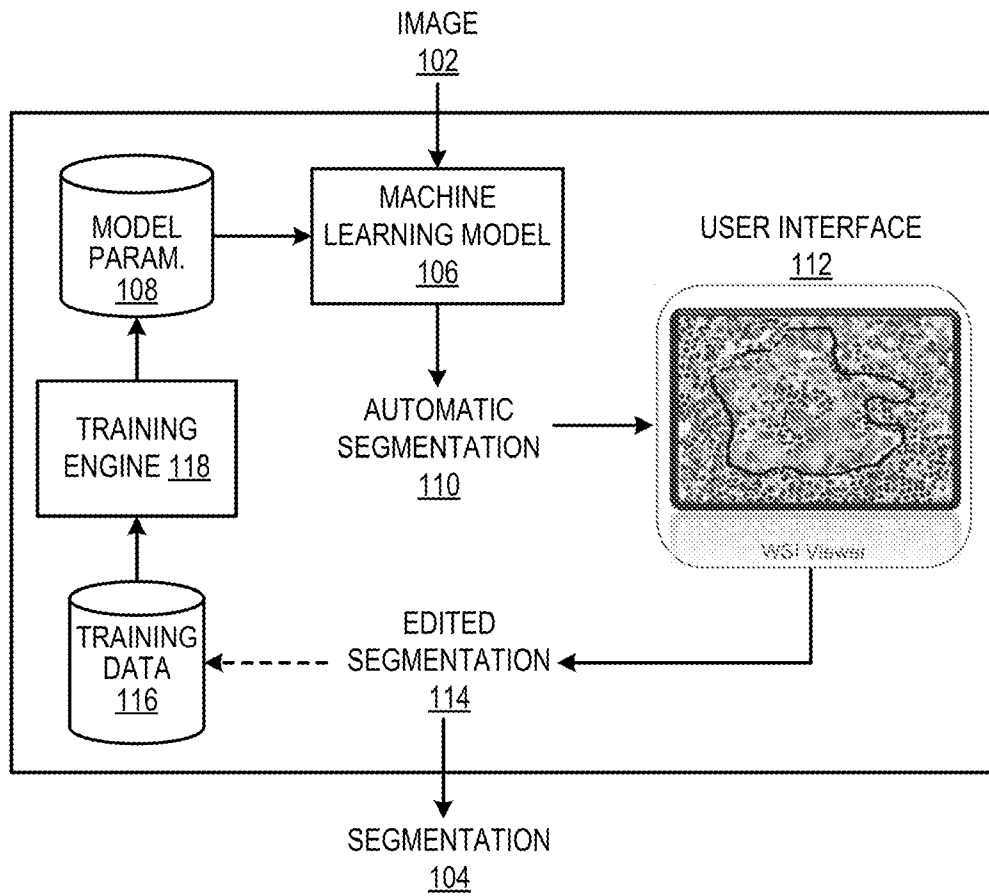
CLAIMS

What is claimed is:

1. A method performed by one or more data processing apparatus for training a machine learning model to segment magnified images of tissue samples, wherein the method comprises repeatedly performing operations including:
 - obtaining a magnified image of a tissue sample;
 - processing an input comprising: the image, features derived from the image, or both, in accordance with current values of model parameters of a machine learning model to generate an automatic segmentation of the image into a plurality of tissue classes, wherein the automatic segmentation specifies a respective tissue class corresponding to each pixel of the image;
 - providing, to a user through a user interface, an indication of: (i) the image, and (ii) the automatic segmentation of the image;
 - obtaining, from the user through the user interface, an input specifying one or more modifications to the automatic segmentation of the image;
 - determining an edited segmentation of the image, comprising applying the modifications specified by the user to the automatic segmentation of the image; and
 - determining updated values of the model parameters of the machine learning model based the edited segmentation of the image.
2. The method of claim 1, wherein the plurality of tissue classes comprise cancerous tissue and non-cancerous tissue.
3. The method of claim 1, wherein the image is a whole slide image (WSI) captured using a microscope.
4. The method of claim 1, wherein the machine learning model comprises a neural network model.
5. The method of claim 4, wherein the neural network model is a convolutional neural network model comprising one or more convolutional neural network layers.

6. The method of claim 4, wherein the model parameters may specify weight values of each layer of the neural network model.
7. The method of claim 1, wherein the machine learning model comprises a random forest model.
8. The method of claim 7, wherein the model parameters specify splitting functions used at nodes of decision trees of the random forest model.
9. The method of claim 1, wherein providing, to the user through the user interface, an indication of: (i) the image, and (ii) the automatic segmentation of the image comprises:
 - providing, to the user through the user interface, a visualization that depicts the automatic segmentation overlaid on the image.
10. The method of claim 9, wherein the visualization that depicts the automatic segmentation overlaid on the image indicates the respective tissue class corresponding to each pixel of the image.
11. The method of claim 1, wherein:
 - each modification to the automatic segmentation of the image indicates, for one or more pixels of the image, a change to the respective tissue class specified for the pixel by the automatic segmentation; and
 - applying the modifications specified by the user to the automatic segmentation of the image comprises changing the tissue classes corresponding to one or more pixels of the image in accordance with the modifications specified by the user.
12. The method of claim 1, wherein determining updated values of the model parameters based on the edited segmentation of the image comprises:
 - determining gradients of an objective function that characterizes a similarity between: (i) the automatic segmentation of the image, and (ii) the edited segmentation of the image; and
 - adjusting the values of the model parameters using the gradients.

13. The method of claim 1, further comprising determining an expertise score that characterizes a predicted skill of the user in reviewing and editing tissue segmentations.
14. A system comprising:
 - one or more computers; and
 - one or more storage devices communicatively coupled to the one or more computers, wherein the one or more storage devices store instructions that, when executed by the one or more computers, cause the one or more computers to perform the operations of the respective method of any one of claims 1-10.
15. One or more non-transitory computer storage media storing instructions that when executed by one or more computers cause the one or more computers to perform the operations of the respective method of any one of claims 1-10.



SEGMENTATION SYSTEM 100

FIG. 1

IMAGE 200




FIG. 2

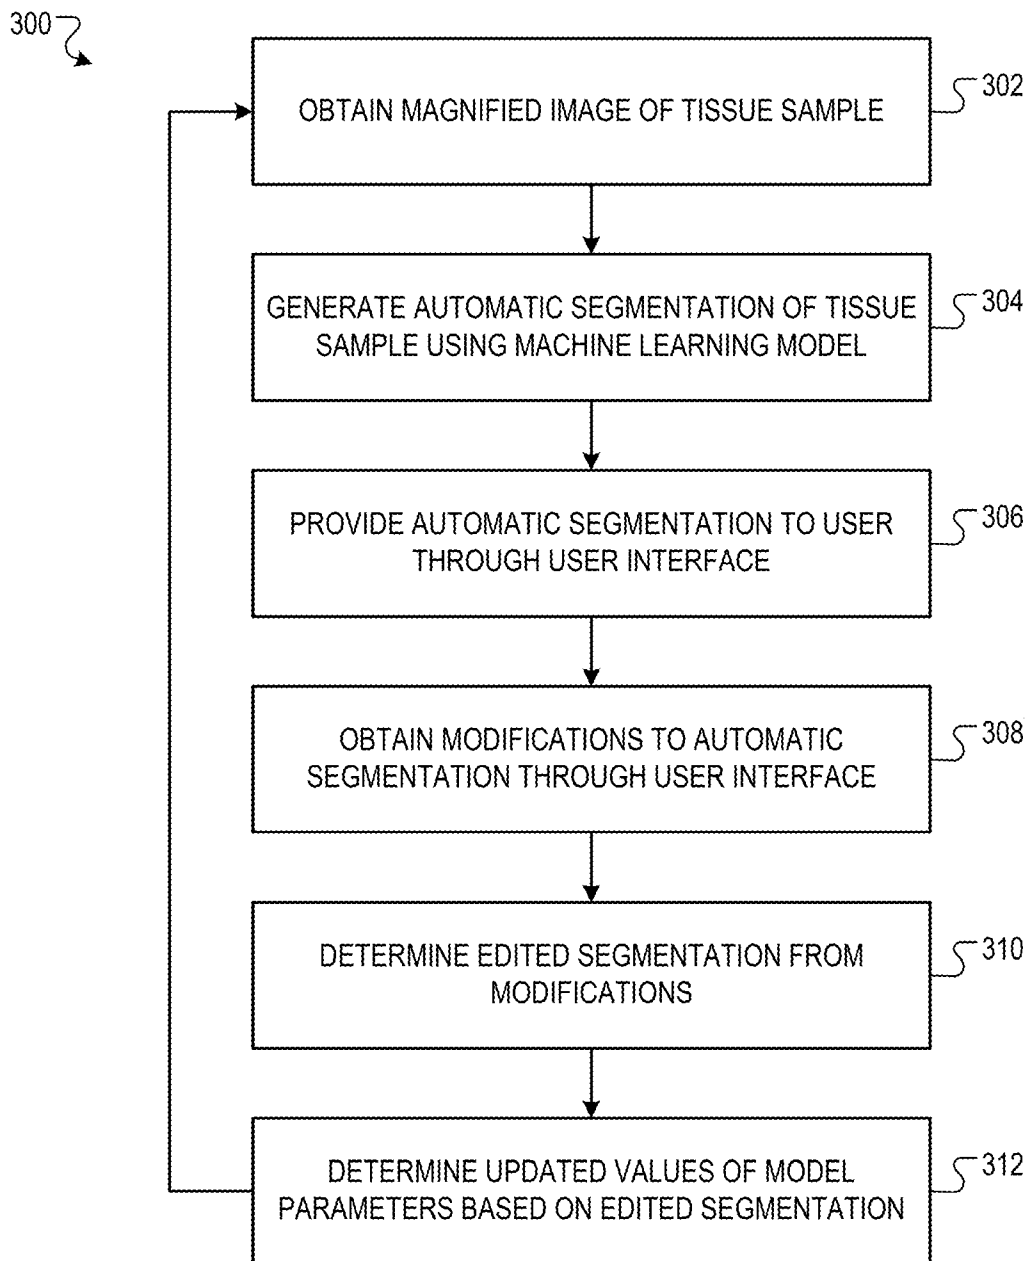


FIG. 3

4 of 4

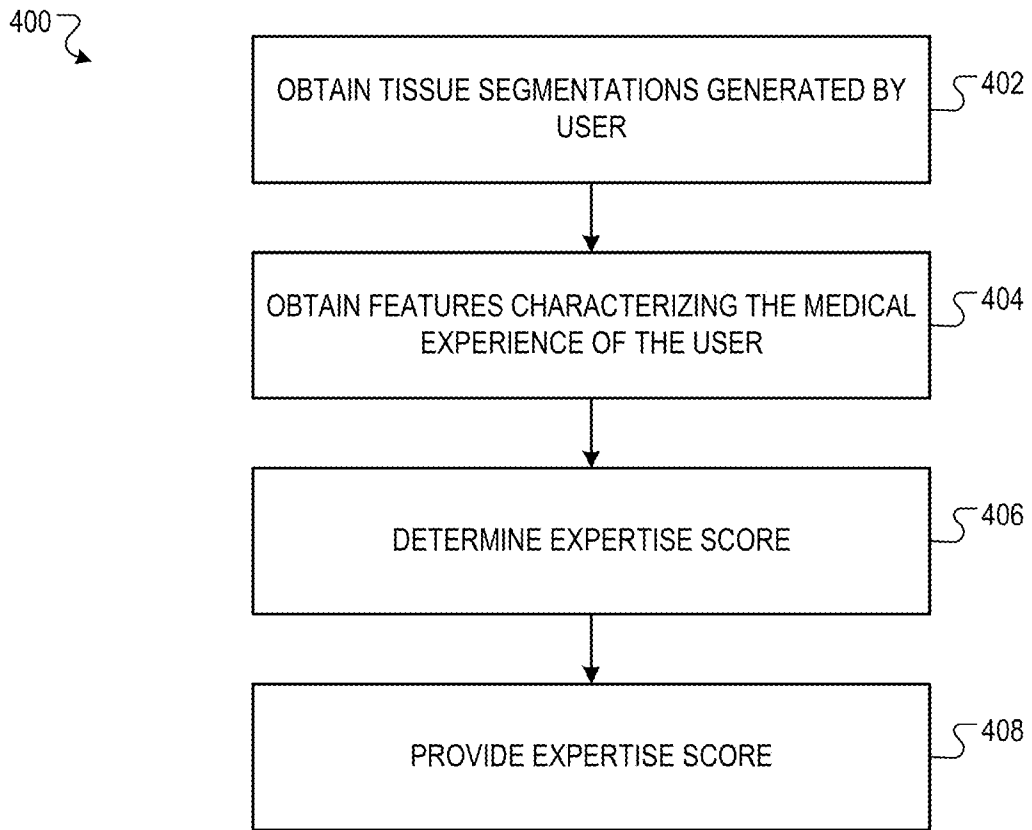


FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2020/051944

A. CLASSIFICATION OF SUBJECT MATTER

G06T 7/11(2017.01)i; G06T 3/40(2006.01)i; G06T 1/00(2006.01)i; G06N 20/00(2019.01)i; G06N 3/08(2006.01)i

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)

G06T 7/11(2017.01); A61B 6/00(2006.01); G06K 9/00(2006.01); G06K 9/62(2006.01); G06N 3/02(2006.01); G06N 99/00(2010.01); G06T 15/08(2011.01); G06T 7/00(2006.01); G06T 7/33(2017.01); G09G 5/00(2006.01)

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

Korean utility models and applications for utility models
Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: image, segmentation, machine learning, interface, update

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2010-0322489 A1 (HAMID REZA TIZHOOSH et al.) 23 December 2010 (2010-12-23) paragraphs [0058]-[0067], [0135]; claims 1-3; and figures 1-2, 15-16	1-15
Y	US 2014-0358825 A1 (CLOUDVU, INC.) 04 December 2014 (2014-12-04) paragraph [0119]; claim 1; and figure 3	1-15
A	US 2019-0188849 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION) 20 June 2019 (2019-06-20) claims 1-18	1-15
A	US 2019-0209116 A1 (PROGENICS PHARMACEUTICALS, INC. et al.) 11 July 2019 (2019-07-11) paragraph [0308]; and figure 34	1-15
A	WO 2019-084697 A1 (UNIVERSITY HEALTH NETWORK) 09 May 2019 (2019-05-09) paragraphs [0171]-[0177]; and figure 4	1-15

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

* Special categories of cited documents:

“A” document defining the general state of the art which is not considered to be of particular relevance

“D” document cited by the applicant in the international application

“E” earlier application or patent but published on or after the international filing date

“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

“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

“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

“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

“&” document member of the same patent family

Date of the actual completion of the international search

29 December 2020

Date of mailing of the international search report

30 December 2020

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon
35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

YANG JEONG ROK

Telephone No. +82-42-481-5709

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.

PCT/US2020/051944

Patent document cited in search report			Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
US	2010-0322489	A1	23 December 2010	US	8391603	B2	05 March 2013
				WO	2010-145007	A1	23 December 2010
US	2014-0358825	A1	04 December 2014	EP	3005056	A1	13 April 2016
				EP	3005138	A1	13 April 2016
				US	2014-0358828	A1	04 December 2014
				US	9218574	B2	22 December 2015
				WO	2014-194000	A1	04 December 2014
				WO	2014-194002	A1	04 December 2014
US	2019-0188849	A1	20 June 2019	US	10438351	B2	08 October 2019
US	2019-0209116	A1	11 July 2019	AU	2019-205766	A1	11 July 2019
				CN	111602174	A	28 August 2020
				EP	3738097	A2	18 November 2020
				WO	2019-136349	A2	11 July 2019
				WO	2019-136349	A3	15 August 2019
WO	2019-084697	A1	09 May 2019	CA	3081643	A1	09 May 2019
				US	2020-0272864	A1	27 August 2020