Explainable AI (xAI) for Anatomic Pathology

Akif B. Tosun, PhD, Filippo Pullara, PhD,* Michael J. Becich, MD, PhD,*†*
D. Lansing Taylor, PhD,‡§ Jeffrey L. Fine, MD,*‖¶*
and S. Chakra Chennubhotla, PhD§*

Abstract: Pathologists are adopting whole slide images (WSIs) for diagnosis, thanks to recent FDA approval of WSI systems as class II medical devices. In response to new market forces and recent technology advances outside of pathology, a new field of computational pathology has emerged that applies artificial intelligence (AI) and machine learning algorithms to WSIs. Computational pathology has great potential for augmenting pathologists' accuracy and efficiency, but there are important concerns regarding trust of AI due to the opaque, black-box nature of most AI algorithms. In addition, there is a lack of consensus on how pathologists should incorporate computational pathology systems into their workflow. To address these concerns, building computational pathology systems with explainable AI (xAI) mechanisms is a powerful and transparent alternative to black-box AI models. xAI can reveal underlying causes for its decisions; this is intended to promote safety and reliability of AI for critical tasks such as pathology diagnosis. This article outlines xAI enabled applications in anatomic pathology workflow that improves efficiency and accuracy of the practice. In addition, we describe HistoMapr-Breast, an initial xAI enabled software application for breast core biopsies. HistoMapr-Breast automatically previews breast core WSIs and recognizes the regions of interest to rapidly present the key diagnostic areas in an interactive and explainable manner. We anticipate xAI will ultimately serve pathologists as an interactive computational guide for computer-assisted primary diagnosis.

Key Words: computational pathology, digital pathology, explainable artificial intelligence, machine learning, computer-assisted diagnosis

(*Adv Anat Pathol* 2020;27:241–250)

A growing number of pathologists have transitioned to viewing digital images of patient slides on computer monitors with the recent growth of digital pathology whole slide image (WSI) platforms.[1,2] Following a protracted course, the FDA recently approved 2 WSI systems as class II medical devices; Philips' IntelliSite digital imaging system[3] and Leica Biosystems' Aperio AT2 DX digital imaging system.[4] The interest in technologies that accomplish simple image analysis tasks such as detecting nuclei, labeling epithelial and stromal cells, cell counting, and measuring automatically led to the development of computational pathology mostly in the form of black-box machine learning (ML).[5–7] Computational pathology has great potential for augmenting the accuracy and efficiency of a pathology practice. However, many of these applications do not address the critical pain points in pathology practices (eg, diagnosing difficult cases involving atypia in breast lesions, improving efficiency of slide reading) and are mostly used for research only purposes with no clear plans for clinical integration. This has led to innovative new concepts such as computer-assisted diagnosis for pathologists (pCAD),[8,9] that propose methods for integration of ML tools into pathology workflow.[10–18]

Deep learning, in the form of convolutional neural networks, has been popular in early computational pathology efforts. While powerful in isolated and lower level applications, such as mitosis counting or malignancy and cancer detection,[19–21] deep learning has not yet yielded validated, comprehensive, and high level systems that critically address pathologists' workflow needs. There also exists a concern of trust in the application of artificial intelligence (AI)/ML to pathology because of the present black-box nature of AI/ML algorithms [22] (Fig. 1). There is also a lack of consensus on how pathologists should supervise or work with computational pathology systems. The explainable AI (xAI) concept has emerged as a powerful and transparent alternative to black-box AI.[23] Although xAI term recently used for explaining black-box AI tools, it is actually a broader concept in ML, covering AI algorithms that generate human-understandable statements for its conclusions. The main goal with xAI for pathology is to provide clear justifications to the user for the automated recommendations made in the diagnostic workflow (Fig. 1). This will promote safety, reliability, and accountability in addressing issues with *bias*, *transparency*, *safety*, and *causality*.[24]

The *causality* question seeks a mechanistic understanding of the problem from a data-driven learned model. For example, in the case of benign breast lesions, how does the spatial organization of a duct change from normal to atypia to in situ carcinoma to invasive cancer? *Safety* implies knowing the error bounds on unseen data, for example, can computational pathology systems ever achieve the same accuracy as the human experts or perform even better? *Bias* arises when cases such as atypical ductal hyperplasia (ADH) in breast biopsies do not occur frequently in the sample population, leading to unbalanced training datasets.[25] *Transparency* requires decisions affecting patients and pathologists should be clearly explained. xAI can justify its results to users, providing them with all the information needed to make good decisions based on xAI recommendations.

In pathology, patient *safety* is paramount and is the result of a complex interaction between pathologist, other

From the *SpIntellx Inc.; Departments of †Biomedical Informatics; §Computational and Systems Biology; ‖Pathology, University of Pittsburgh School of Medicine; ‡Drug Discovery Institute, University of Pittsburgh; and ¶UPMC Magee-Womens Hospital, Pittsburgh, PA.

This study was partially supported by NSF SBIR Phase I Award #1843825.

J.L.F., M.J.B., D.L.T., and S.C.C. are co-founders with equity in SpIntellx Inc. The remaining authors have no funding or conflicts of interest to disclose.

Reprints: S. Chakra Chennubhotla, PhD, SpIntellx Inc., 2425 Sidney St., Pittsburgh, PA 15203 (e-mail: chakra@spintellx.com).

All figures can be viewed online in color at www.anatomicpathology.com.

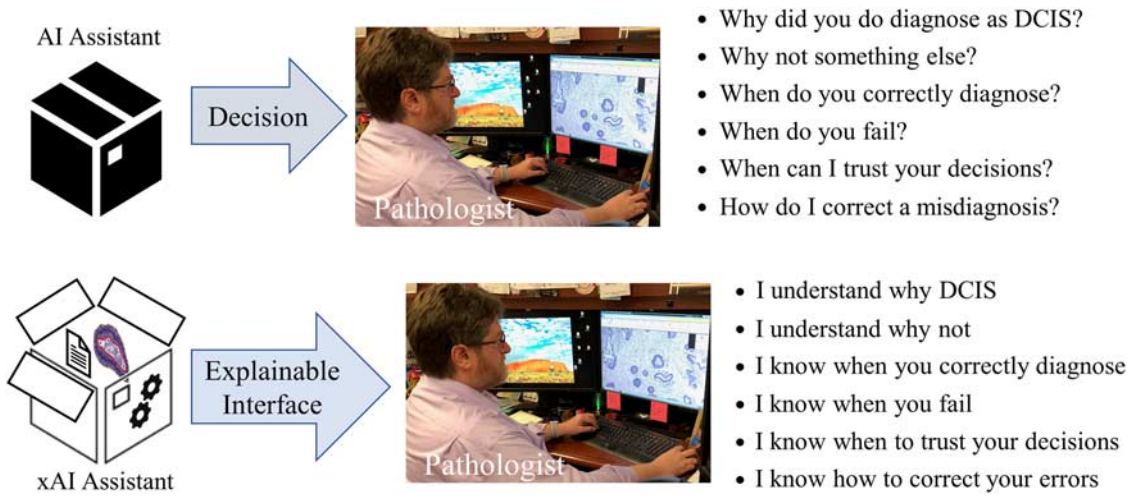

FIGURE 1. Machine learning/artificial intelligence (AI) models can be opaque, nonintuitive, and difficult for users to understand. The concept of explainable AI system must be integrated in pathology workflow. DCIS indicates ductal carcinoma in situ. Please see this image in color online.

physicians and laboratory personnel, and computer systems including computational pathology applications. Patient *safety* can be improved by xAI not only by reducing undetected *bias* or by providing *transparency* to the pathologist, it is also related to the pathologist being able to monitor xAI's functionality in real-time, on individual patient samples. Many ML efforts simply chase engineering statistics (eg, area under the curve) or trying to explain the black-box AI decisions by showing a group of pixels or an image region on a heat map that led the black-box to make a call without gaining new insight into the pathology itself. xAI on the other hand, allows researchers/pathologists to understand potential new disease mechanisms that can lead to meaningful diagnostic or therapeutic advances. These attributes make xAI a critical feature for computational pathology.

Pathology is considered a gold standard of medical diagnosis; therefore, pathologists have been conservative about making large practice changes. Interviews with >30 pathologists revealed how xAI might facilitate adoption of computational pathology by not replacing the pathologist. The goal is to allow the pathologist to focus on the most important decisions that only they can make.

Looking into the future of AI in medicine, the National Artificial Intelligence Research and Development Strategic Plan from White House (2019) (https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf) and European Commission's Report on The Ethics Guidelines for Trustworthy Artificial Intelligence (2019) (https://ec.europa.eu/futurium/en/ai-alliance-consultation) strongly recommend incorporating xAI capabilities for regulatory approvals. We believe that this paper presents the first reported applications of xAI in pathology and our examples will be on breast biopsies to assist anatomic pathologists.

Now there is tremendous interest in computational pathology despite any potential concerns, as seen at major pathology organizations (https://digitalpathologyassociation.org/, https://www.uscap.org/). Traditional manual pathology diagnosis, either with glass microscope slides or with manual WSIs, is inefficient and error prone. For example, one group reported low diagnostic concordance between pathologists when diagnosing difficult lesions in breast biopsies such as ADH (52% concordance); this report also noted that different kinds of pathologists (ie, breast pathologists vs. general pathologists) were subject to different levels of performance.[26]

Regarding efficiency, manual WSI viewing may not provide adequate efficiency, versus traditional glass slide microscopy, to justify its implementation for primary diagnosis given the cost and workflow changes of WSI including time delay for slide scanning. This is based on unpublished data but is also borne out by the sluggish adoption of digital pathology for primary diagnosis in recent years despite a regulatory opening in 2017.[27] According to our simulations of a pCAD model,[8] computational pathology could be 56% more efficient than traditional microscopy for breast core biopsies. Efficiency gains appeared to come from several factors, including earlier discovery of diagnostic regions of interest (ROIs); decreased uncertainty due to a triage effect; ability to review less diagnostic ROIs in an expedited manner after major diagnostic decisions were made.[13]

APPLICATIONS OF EXPLAINABLE ARTIFICIAL INTELLIGENCE IN ANATOMIC PATHOLOGY

We envision several applications of xAI in the anatomic pathology workflow.

Rapid Diagnosis Based on Triaged Region of Interest

The idea is to effectively change pathologists' view of a case from one or more WSIs into a guided series of triaged, diagnostically relevant ROIs. The xAI system previews an entire specimen's WSI data set to discover relevant structures and features. For breast core biopsies, it finds ducts, vascular structures, and stromal features that are then associated with ROIs. The ROIs can then be classified and/or measured with an xAI library of diagnostic labels. The combination of diagnostic labeling and ROI quantitation is then used to triage the ROIs taking its diagnostic impact into account. In an interactive work session, the pathologist reviews the entire case, ROI by ROI, in a triaged and optimized manner. This approach is highly efficient because the xAI interface presents the most clinically impactful ROIs to the pathologist first; this guidance enables the pathologist to focus on the hardest decisions first. If necessary, the

Copyright © 2020 Wolters Kluwer Health, Inc. All rights reserved.

system also keeps track of ROIs that may need further workup with additional stains, or ROIs that may require consultation with another pathologist. In initial setup, the pathologist is shown all ROIs to ensure nothing is missed; this could change in the future as with imaged Pap tests[28] that are screened-negative (ie, transition to human review of <100% of ROIs).

Pathologist-friendly Explainable Artificial Intelligence Interface

xAI analysis yields multiple morphologic features (with quantitation of those features) that can be described in pathologist-friendly language, thereby linking important features with diagnostic labels that can be used to explain the diagnoses in subsequent clinical sessions, providing pathologists with all the information that they need to make the best decisions possible. Decision support can show the pathologist examples of similar ROIs from other cases, from libraries of known diagnoses, or from didactic materials. The key is that the pathologist has good confidence based on having sufficient information to make difficult diagnostic decisions.

Explainable Artificial Intelligence Confidence Visualization in Real-time

A software "Why?" button can provide complete transparency by presenting additional information in real-time that explains the system's labels, fleshes out the relevant differential diagnosis, and acknowledges the strength of the analysis via a "confidence score." The system can display the results of its ROI analysis on individual patient specimens to the pathologists in real-time. This is unlikely if the AI is based on a black-box model. Recent explanation efforts that displays the image responsible for the classification label are not considered an "explanation" in pathologists' terms.

Intelligent Case Triaging

Confidence score based (content based) triage of patient biopsies; cases above a certain difficulty threshold can be distributed to subspecialty expert pathologists, whereas other cases can be sent to front-line general pathologists. xAI estimations of case difficulty and image volume can also be used to distribute pathology cases to a group of pathologists more evenly for better utilization of professional resources in larger practices. As with other applications, an xAI "Why?" button means that the system can justify its case distribution and triage decisions if necessary.

Rapid Ground-truth Collection to Facilitate Explainable Artificial Intelligence

The xAI system can expedite ground truth data labeling for training the ML. Using a preprocessing step to automatically find ROIs, pathologists can then rapidly view the ROIs and provide streamlined diagnosis labels. This is a novel approach that can provide hundreds or more pathologist-labeled ROIs per hour. This permits training by clinically busy and expert pathologists.

Postdiagnostic Quality Assurance

After cases are finalized, or signed out, they can be rereviewed by an xAI based system for concordance with the pathology diagnosis. The goal is to (a) reduce risks by detecting significant discrepancies (eg, benign vs. malignant), where "Why?" button provides transparency and shows why a case was potentially discrepant, (b) generate evidence to regulators of ongoing and effective QA activities, (c) improve clinician and patient confidence in diagnoses, especially difficult diagnoses such as breast atypias, (d) automate second review of pathologists' work by other pathologists (eg, QA reviews, standard second-opinion situations, etc.), and (e) monitor the performance of xAI system as part of a pathology practice's QA framework.

Explainable Artificial Intelligence for Educational Purposes

xAI powered content-based ROI comparison with known good examples of a diagnosis allows for trainee decision support. Further, the provision of supporting morphologic features helps provide mastery to trainees and to practicing pathologists. Also, the ability of an xAI system to provide a relevant differential diagnosis with acknowledgment of borderline or ambiguous ROIs is a powerful educational opportunity.

Explainable Artificial Intelligence for Content-based Image Retrieval Activities

Content-based image searching has other applications, such as for scholarly activities, for tumor board conferences, or other presentations. The xAI derived statistics of ROIs or of entire cases are a fingerprint of sorts that can also be used for other content-based purposes. The statistics include features present, quantitation of those features, and the previously mentioned confidence scores.

Explainable Artificial Intelligence for Clinical Trials

After a case has been finalized using an xAI tool, the diagnostic ROI information can be re-used for revisiting the case rapidly, such as for central review. In addition, it is possible to use xAI features to flag specimens for possible inclusion in clinical trials based on diagnostic labels and/or morphologic features, automatically. For example, a pathologist uses the tool to diagnose a patient with intermediate grade ductal carcinoma in situ (DCIS). The laboratory is collaborating with a surgeon on a DCIS clinical trial, and both the diagnosis (DCIS) and specific morphologic features (apocrine morphology) are desired in the study. xAI can flag the case and when the study pathologist reviews the case for suitability, it can display the diagnostic ROIs with DCIS immediately. This not only casts a reliable net for study inclusion, but it also greatly streamlines re-review of the case for determining whether the patient should be in the study.

Explainable Artificial Intelligence for Standardized Terminology

By supplying suggested diagnostic labels, the xAI system can encourage pathologists to report results in a more uniform fashion using standardized language, which is a significant quality improvement opportunity. As the pathologist and the xAI interactively work through the case together, the xAI system can use the pathologist's decisions to automatically construct a pathology result report in real-time, using standardized language. This standardizes practice and improves quality of a practice's reports by making them easier to understand due to standardization of practice. When supported by the electronic health record, this could also result in structured-data reporting for next-generation consumption of anatomic pathology results by clinical colleagues, patients, billing personnel, etc.

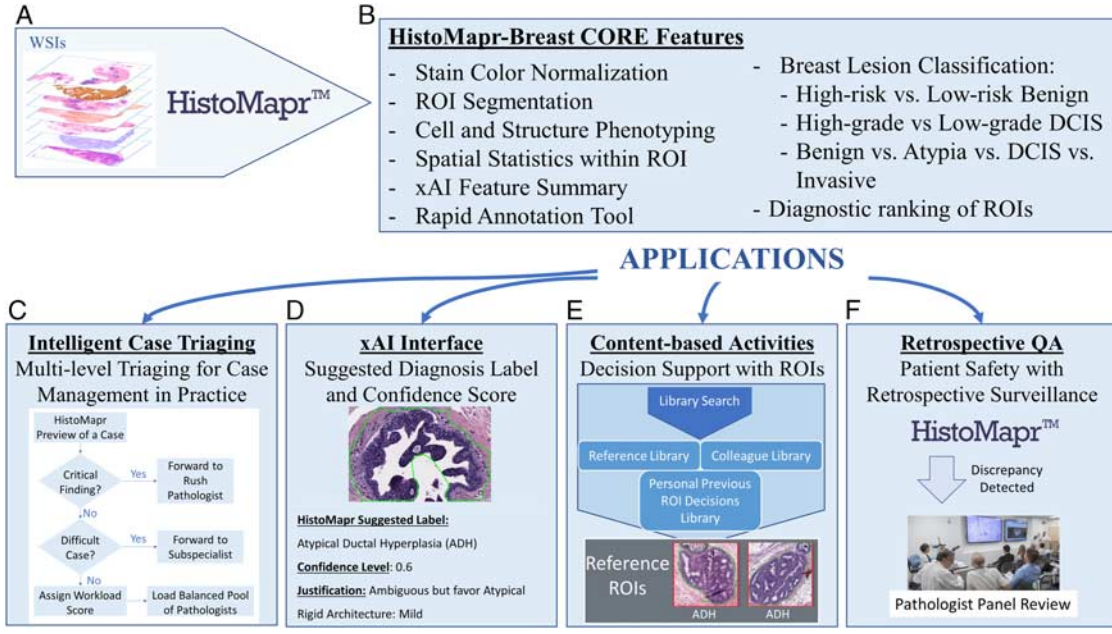

FIGURE 2. A, HistoMapr, explainable artificial intelligence (xAI) based software platform, analyzes whole slide images (WSIs) to discover diagnostic regions of interest (ROIs), (B) generating descriptive information required for xAI features. This in turn permits subsequent activities such as: (C) triaging cases based on their difficulties for subspecialists, (D) explanation and real-time awareness of HistoMapr's performance, (E) content-based image retrieval facilitating several activities such as decision support, education or clinical trial screening, and (F) quality assurance activities including discrepancy detection and improved reporting by standardized terminology. QA indicates quality assurance. Please see this image in color online.

HISTOMAPR-BREAST: AN EXPLAINABLE ARTIFICIAL INTELLIGENCE SYSTEM FOR BREAST LESIONS

To address the unmet needs of efficiency and accuracy in pathology diagnosis, and to accelerate the adoption of computational pathology in pathology practice we are developing an xAI platform called HistoMapr-Breast (Fig. 2). We believe that this is the first reported application of xAI in pathology and our example is on breast biopsies to assist pathologists.

HistoMapr as a Computational Pyramid

HistoMapr is designed for transmitted light applications where tissue samples are stained with Hematoxylin and Eosin (H&E), immunohistochemistry (IHC) labels, and/or other stains (eg, special stains, chromogenic in situ hybridization, enzyme metallography, etc.). HistoMapr xAI ecosystem reflects the natural hierarchy and the spatial organization found in human breast tissue and other organs (Fig. 3). Lowest level, simple image analyses, such as nuclear segmentation, size/shape measurements, or mitosis counts, can be integrated with our pointwise mutual information maps[29,30] (second level) to identify and classify tissue structures such as ducts in breast tissue. Diagnostic ROIs are based upon the tissue structures' spatial relationships, and the ROIs are labeled with diagnostic information, then triaged based on diagnostic significance. In this manner, the xAI HistoMapr guides emerge; WSIs are rerepresented as a guided review of triaged ROIs in the context of the pathologist's diagnostic tasks. Including relationships between distinct cellular populations of the tissue microenvironment (eg, tumor, stromal and immune cells), this enables precision medicine approaches to be incorporated into routine diagnostic and prognostic activities.

As an example, breast core biopsies can be difficult for pathologists to diagnose concordantly.[26,31,32] Therefore, we created a HistoMapr-Breast prototype as proof of concept (Fig. 4). Briefly, HistoMapr-Breast analyzes an entire breast core biopsy specimen, consisting of one or more WSIs, using both basic image analyses and pointwise mutual information maps to locate ROIs that contain diagnostic breast tissue structures such as ducts, blood vessels, or stromal infiltrates.[10,14] HistoMapr-Breast then analyzes the ROIs to capture diagnostically relevant features and quantitate them if present; this analysis is used to label the ROIs with diagnostic terms such as "ADH" or "invasive carcinoma"; this is based on analysis of the feature patterns and strengths.

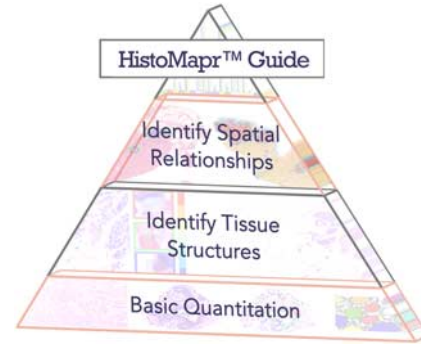

FIGURE 3. HistoMapr explainable artificial intelligence capabilities: hierarchical integration from cellular quantification to computational guides. Please see this image in color online.

Copyright © 2020 Wolters Kluwer Health, Inc. All rights reserved.

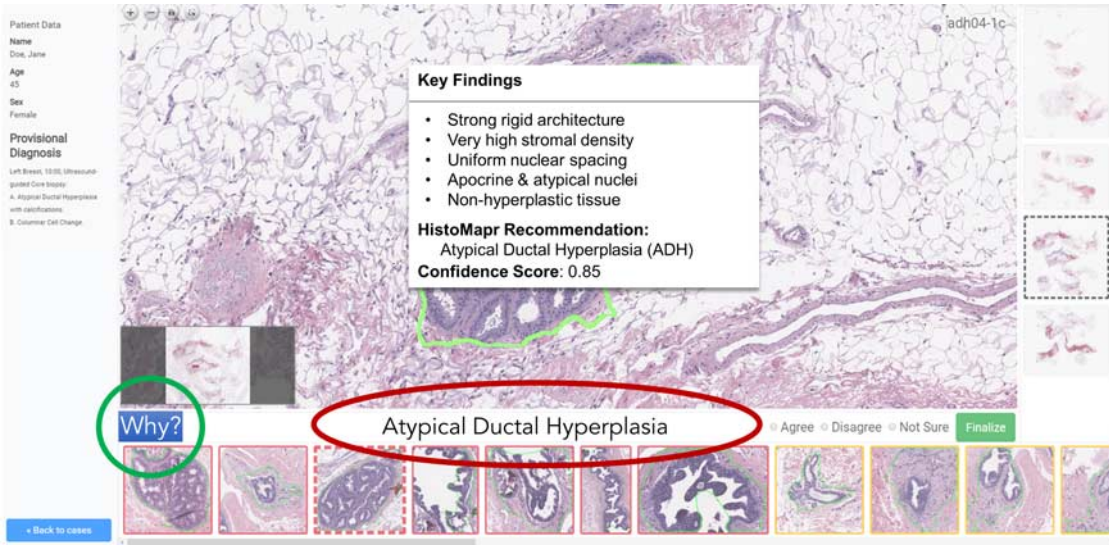

FIGURE 4. HistoMapr-Breast explainable artificial intelligence interface with a "Why?" button. Left panel shows patient information and provisional diagnosis and the right panel has thumbnail images of the patient slides. Regions of interest (ROIs) are automatically detected and presented in the bottom panel. ROIs are triaged based on diagnostic significance from left to right. In this example, HistoMapr analyzed the slide and recommended the diagnosis of atypical ductal hyperplasia for ROI in question, which is a challenging call. Pathologist can hit the "Why?" button to let HistoMapr display Key Findings that led to this recommendation. Please see this image in color online.

How HistoMapr-Breast EnablesExplainable Artificial Intelligence Applications in Anatomic Pathology

Rapid Diagnosis Based on Triaged Region of Interest

For example, a prototype version of HistoMapr-Breast used 18 features to classify ROIs as atypical or not atypical[10,14] (Fig. 5). In this analysis, HistoMapr finds which features are present in an ROI, and it quantitates each feature that it finds. For breast core biopsies, our system triages ROIs in a clinical spectrum from invasive cancer at one end, to low-risk benign at the other (Fig. 4). Using a pathologist-centric interactive interface, HistoMapr displays the ROIs in the

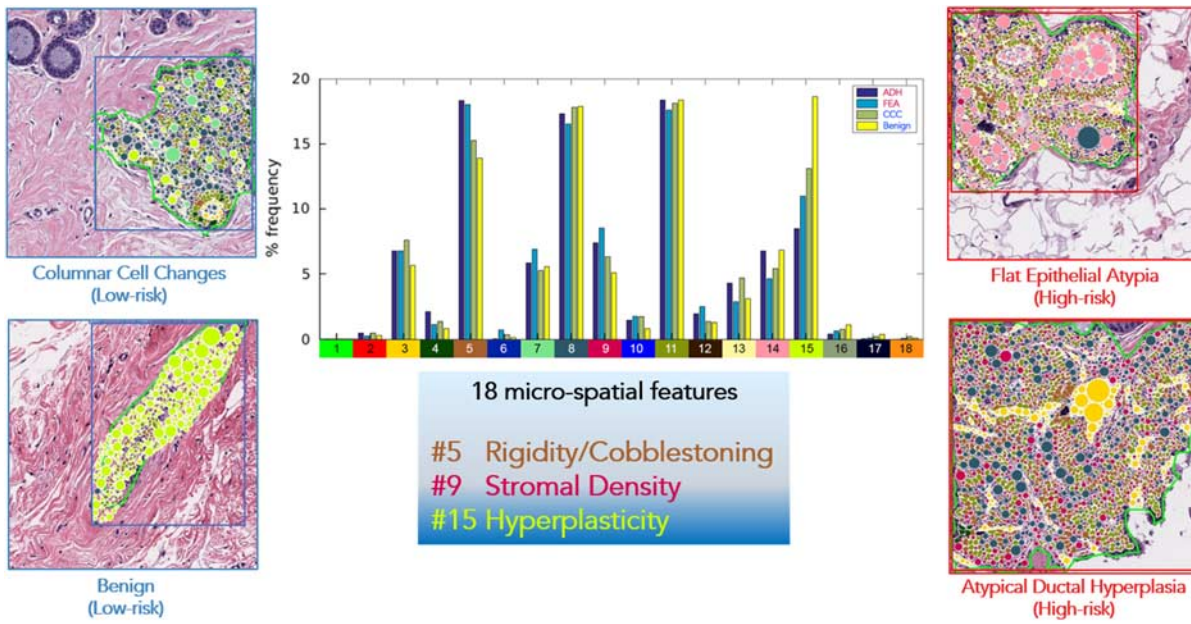

FIGURE 5. xAI: Micro-Spatial discriminant features in breast regions of interest (ROIs). Sample ductal ROIs representing (A) atypical ductal hyperplasia (ADH), (B) flat epithelial atypia (FEA), (C) columnar cell change (CCC), and (D) normal duct. Visualization of architectural patterns discovered in sample ROIs are overlaid on original images. Patterns are derived from a combination of cytologic and architectural features and visualized by color coded objects (see *x*-axis of histogram in center). Note the overexpression of pattern #5 in ADH, #7 in FEA, and #15 in normal ducts. This observation is further supported by the histogram in center, where we measure relative proportions of architectural patterns separately in each one of the categories: ADH, FEA, CCC, and normal. These architectural patterns are further labeled (eg, rigidity/cobblestoning) by expert pathologists based on the visual appearance of the underlying tissue architecture.

TABLE 1. Summary of Time Data for Glass Slide Reviews (Measured Times) and for pCAD[8] (Combination Of Simulated And Measured Times)

	Mean	SD
Time (glass slides, s)	221.6	141.3
Time (pCAD, s)	98.0	50.6
Time saved (s)	123.7	
Reduction (%)	55.8%	$P = 0.000266$

pCAD indicates computer-assisted diagnosis for pathologists.

triage order, so that the pathologist sees the most malignant or most atypical areas first, if present (Fig. 4). Critically, the pathologist is always fully in control and may take manual control of the WSI viewer software at any time if they need to review all or part of the WSIs manually. Early study of the HistoMapr-Breast prototype performed well and showed 83% f-measure concordance for ADH (N = 300 WSIs and ~2000 ROIs).[10] We particularly addressed the challenging problem of classifying high-risk benign breast lesions. This study focused on ADH, which is a known diagnostic dilemma and is also a diagnosis that confers significant cancer risk.[33] Flat epithelial atypia, columnar cell change, and several benign entities were also included in the study. Subsequent study of that computational pathology pipeline output yielded pathologist-friendly descriptions of some of the important diagnostic features that HistoMapr identified.

The results highlight the value of using HistoMapr-Breast to guide the pathologist in diagnosing atypical breast lesions. A key contribution of this study is in demonstrating how to encode morphometric properties of nuclear atypia (cytologic) and combine them with the spatial distribution of the nuclei in relationship to stroma and lumen (architectural). Although there are several studies on cancer detection in breast tissue images, our pipeline was the first of its kind in detecting high-risk benign breast lesions from WSIs.

Previously,[18] we also addressed the classification of breast lesions into benign usual ductal hyperplasia (UDH) versus malignant DCIS, which requires careful examination of diagnostic criteria, including cellular features, tissue architecture, and spatial extent. We developed a prototype HistoMapr algorithm using a publicly available data set that contained 167 cases (100 DCIS and 67 UDH). In a 4-way classification of UDH versus low, moderate or high-grade DCIS, HistoMapr achieved 88% area under the curve, which represented a 25% improvement (from 70%) over other previously published methods; C-path and DCIS CAD system.

We also timed pathologists diagnosing breast core biopsies using audio recording, which permitted us to know how long pathologists spent on each field of view under the microscope. Using models of HistoMapr style workflow we used this data to simulate the time it would take to diagnose those breast core biopsies with an AI driven pathology assistant. Comparing the simulated and actual biopsy review times, we observed that the HistoMapr-Breast driven diagnosis required 56% less time (Table 1).[13] This is a very significant and novel result that supports the hypothesis that a computer-augmented pathology diagnosis assistant is potentially much more efficient than current manual diagnostic practice.

This improvement is achieved because the pathologist was able to review the most diagnostically relevant ROI first (Fig. 6). HistoMapr was able to detect this ROI and provided it to the pathologist with a recommended diagnostic label, its own confidence score, and supporting explanations in pathology terms. After reviewing this ROI, HistoMapr presented the next relevant ROIs in a nonlinear manner (eg, next ROI may come from another slide within the case) since all slides belonging to the case are analyzed together by HistoMapr and ROIs triaged accordingly. This approach helps prevent problems that can affect pathologist's efficiency (eg, mandatory linear searching of all slides, chance of missing regions, diagnostically critical region is in the last slide).

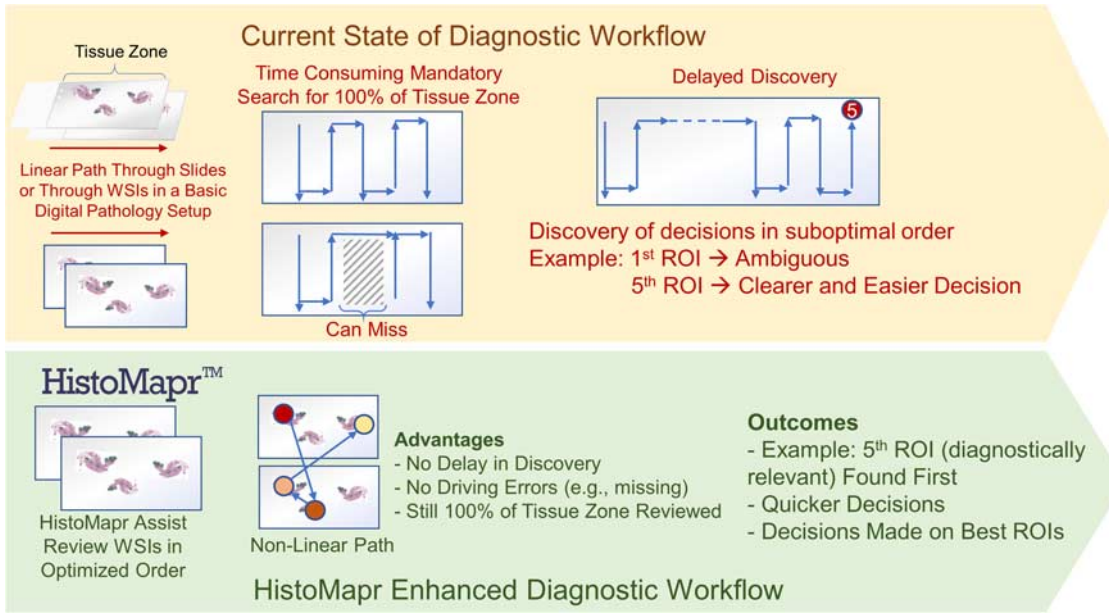

FIGURE 6. Triaged regions of interest (ROI) viewing helps pathologists sign-out more efficiently and more reliably. Please see this image in color online.

Pathologist Friendly Explainable Artificial Intelligence Interface

Using the output of our prototype computational pathology pipeline output, ROIs were reviewed in the context of the 18 features, especially features #5, #9 and #15 which were differentially seen in atypical versus nonatypical ducts (Fig. 5). Feature 5 appeared to correlate with architectural rigidity and cobblestoning of cells; feature 9 appeared to represent stromal density immediately surrounding ducts; and feature 15 seemed to correlate with hyperplasticity of the duct. Review of these ROIs suggested several opportunities for creation of software tools specifically for xAI (eg, an AI feedback overlay, automatic presentation of feature examples, etc.). HistoMapr-Breast xAI manifests as a "Why?" button that provides one or more panels of supplementary information (Fig. 4). This way HistoMapr has way more explanatory power than just showing a heatmap of pixels from a black-box AI model. The pathologist thus has complete situational awareness with assistance of HistoMapr and can make the very best diagnostic decisions. HistoMapr also facilitates the pathologist's work by managing diagnostic information and tracking the pathologist's agreement or disagreement with the provided diagnostic labels (see Agree, Disagree, Maybe button in the interface, Fig. 4); the pathologist may also indicate uncertainty and HistoMapr collects this information for possible additional stain work-up or consultation.

Explainable Artificial Intelligence Confidence Visualization in Real-time

HistoMapr can also indicate its confidence in the label using a "confidence score" that incorporates the features and feature quantities (Fig. 4). The labeled ROIs are then triaged based upon both the diagnostic labels and the confidence scores of those labels. For example, ROIs may be sorted from benign to malignant, or if cancer is not present, then from benign to atypical. Within a diagnostic category, ROIs can be triaged based on the confidence score. These steps occur before the pathologist begins viewing the case, possibly overnight or during weekend off hours.

A HistoMapr-Breast prototype uses a "Why?" button to achieve this. As previously mentioned, HistoMapr finds then quantitates features in the ROIs. When the pathologist presses the "Why?" button, they see a visualization of the ROI analysis by HistoMapr, in pathologist-friendly language (eg, strong rigid architecture, highly monomorphic nuclear patterns, etc.).

Pathologists can interpret the xAI output by accessing an interface panel or page through clicking a "Why?" or "Explain" button.

(a) Present histologic features in the ROI with quantitation (Fig. 2D).
(b) A confidence score analysis of the features, which transparently provides HistoMapr's estimate of the strength of its labeling and of the difficulty of the ROI (Fig. 2D).
(c) Examples of similar ROIs from other cases that can serve as a reference guide, for decision support (Fig. 2E).
(d) A cartoon representation of the features in question, with a control (eg, a slider) that allows the pathologist to view the continuum of that feature from low to high.
(e) HistoMapr xAI system presents its differential diagnosis and displays pros and cons of various diagnoses under consideration; if ambiguous HistoMapr can suggest further work-up with stains, or expedited consultation with another pathologist electronically.

For pathologist guide applications, it is possible for HistoMapr to provide data that supported its diagnostic label including a measure of the data's strength. This is a powerful communication from HistoMapr to the pathologist, for it permits the pathologist to understand why HistoMapr labeled the ROI as it did and whether HistoMapr considers the ROI to be difficult or ambiguous. This permits the pathologist to have all of the necessary information for making a diagnosis, and it permits the pathologist to fully examine HistoMapr's performance in real time.

Intelligent Case Triaging

The previously mentioned confidence scores are automatically generated by HistoMapr during WSI previewing. Aggregated confidence score data can then be used to estimate the difficulty of a case, and also combined with the number of ROIs to estimate the amount of worktime required to view the case. This permits work triage based upon case attributes; difficult cases might be assigned to an expert subspecialist pathologist rather than a generalist pathologist (Fig. 2C). Pathology work could also be distributed evenly to a pool of pathologists, thereby improving efficiency of pathologist time utilization. A "Why?" button would show how or why HistoMapr made its triage decision for a case, including confidence scores, and potential for ambiguity in the diagnosis.

Rapid Ground-truth Collection to Facilitate Explainable Artificial Intelligence

In automating ROI discovery, HistoMapr is much more efficient than traditional ground-truth labeling and effectively addresses what was a bottleneck in ML. HistoMapr has a pathologist-centric graphical user interface (GUI) for efficient annotation of segmented ROIs, which employs a GUI in Java environment to be easily used in any operation system, since the software will be used in pathologists' workstation or personal computer. The software is a plug-and-play, can be installed in an encrypted USB drive together with ROI images to be labeled. The GUI design for ground truth annotation is easy to learn and efficient to use. Pathologists do not need to hand- draw or type their inputs, instead they are shown a series of ROIs and asked to hit buttons from 1 to 5, each corresponding to labels "benign," "high-risk," "DCIS," "invasive carcinoma," and "other/don't know." The ROIs are shown with a green segmentation boundary to collect feedback from pathologists about the quality of our segmentation results. The option "other/don't know" is used by pathologists in case they cannot decide on the label of the ROI or there is a problem with the segmentation boundary. The pathologists can be asked to complete sets of 250 ROIs at a time, with an option of saving their progress in case the pathologist is interrupted during a session. HistoMapr has a database to store, manage and retrieve the ground truth annotations. In our initial experiments, pathologists were able to label around 1000 ROIs in an hour after getting used to the software.

Ground truth data labeling is necessary for ML training but has historically been a bottleneck.[34] Poorly implemented labeling tools can also squander scarce pathologist time. HistoMapr effectively addresses this with both automated ROI discovery and with the pathologist-friendly interface. As discussed above, HistoMapr can analyze one or more WSIs and extract ROIs. Using an efficient interface, pathologists can rapidly apply diagnostic labels to the ROIs. In our experience pathologists can label at least 1000 ROIs an hour. The ROI labels can then be used to train HistoMapr; corrected labels

derived from diagnostic work can also be used to improve HistoMapr based on real-world pathology diagnoses.

Postdiagnostic Quality Assurance

As noted above in the xAI discussion of causality; archived HistoMapr information could be used to facilitate subsequent revisiting of a prior biopsy case. This includes comparison of older biopsies with newer materials or centralized review for clinical trials; pathologists can immediately view the diagnostic ROIs without need to manually search through the WSIs. Previously recorded computational features may also play a role in research where diagnostic entities are reevaluated for risk or diagnostic criteria.

HistoMapr can review previously finalized cases for retrospective QA. The xAI features can be used to help identify potential pathology result discrepancies or errors. This important QA activity can reduce risk, improve clinician confidence and help pathologists monitor their diagnostic work in a timely manner. As with other HistoMapr applications, a "Why?" button would provide details that supported the possibility of a QA issue, such as strong likelihood of malignancy in a biopsy diagnosed as benign, or perhaps a very weak confidence score for the diagnosis that was rendered.

Example HistoMapr-Breast Session

Overnight, HistoMapr has previewed and then triaged a laboratory's breast core biopsies using xAI confidence scores. As a result, the cases are evenly distributed to multiple pathologists, and difficult cases have been assigned to breast pathologists (vs. generalists). The breast pathologist first looks at a worklist of specimens that need to be diagnosed; these are further triaged by HistoMapr based on a variety of factors (eg, many pathologists prefer to look at most-difficult cases early in a day to ensure adequate time for additional staining, consultation, or clinician communication). The pathologist selects the first case, and this brings up HistoMapr-Breast (Fig. 4). After reading the pertinent case information on the sidebar, the pathologist is ready to view the WSIs. Instead of manually and linearly proceeding through the WSIs, the pathologist begins to view the triaged ROIs.

HistoMapr has triaged the ROIs, and therefore the pathologist can see the most diagnostic areas immediately. In this case, the first ROI is labeled "ADH" and the pathologist feels confident that this does represent ADH. Clicking on the "Why?" button brings up a graphic displaying the relevant diagnostic features and the HistoMapr confidence score (Fig. 4); the pathologist was already confident, but viewing this additional information helps the pathologist make this critical diagnostic decision very early in the review process. This ROI is located near the edge of the last WSI in the case; without HistoMapr, the pathologist would have had to review almost all of WSIs before finding this ROI manually. The pathologist confirms the HistoMapr ADH label and then moves on to the next ROIs. As it happens, there are five more potential ADH ROIs. Because of the HistoMapr triage, each next ADH ROI has a lower confidence score and is less diagnostic to the pathologist, but this does not matter as the pathologist has already decided on the diagnosis in the first ROI. The pathologist therefore goes through these ROIs rapidly.

If the pathologist were manually reviewing the WSIs, they would have to spend time with each of these less diagnostic areas before eventually reaching the most-diagnostic ADH area. This is not only inefficient, if the pathologist feels uncertain about the diagnosis then that will also lead to slower manual WSI review as the diagnostic decision is harder until that last ADH area is encountered. In HistoMapr, next are benign ROIs with additional diagnoses such as columnar cell change, fibrocystic changes, and UDH. Throughout, HistoMapr-Breast tracks the pathologist's decisions about the diagnoses for later inclusion in the pathology results report. By the time the pathologist has finished viewing all ROIs, they have also finished viewing all of the specimen WSIs. Although the pathologist had the ability to manually navigate and view the WSIs using standard controls, it was not necessary this time.

After ROI review, the pathologist is ready to finalize the case and proceeds to the HistoMapr finalization screen. Here, the pathologist can review all of the ROI labels; although the pathologist was confident about the first ADH ROI, they were less confident in some of the later ROIs. This screen is an opportunity to review ROIs where the pathologist was uncertain or perhaps disagreed with HistoMapr. In addition, the pathologist can see the preliminary pathology diagnosis result that is based on the ROI workflow. This is an additional time saver, as the pathologist did not need to dictate or even choose what words to use, because HistoMapr-Breast has used standardized diagnostic terms for the report. Satisfied, the pathologist finalizes the result which is then transmitted to the electronic health record. In this instance, a radiologist reviews the report, finds it concordant with the radiographic impression, and then recommends a surgical consult for the patient. The use of standardized language is a quality improvement that facilitates this downstream workflow and also helps the surgeon to present a detailed risk assessment to the patient based on all of the available information.

DISCUSSION AND CONCLUSIONS

Enthusiasm for computational pathology has reignited the long smoldering field of digital pathology, albeit tempered by caution regarding AI. Currently, almost all AI in computational pathology is based on deep learning technology, and it is difficult to probe how deep learning methods work.[25] However, xAI software tools could assist pathologists in a transparent fashion that guides them to make the best diagnoses possible. Explainability is a crucial feature for early computational pathology applications, because it will be necessary for pioneer pathologists to begin building trust even with highly validated intelligent software guides (Fig. 1).[35,36]

In addition to facilitating pathology acceptance, xAI also has the potential to smooth or ameliorate potential regulatory issues for early computational pathology systems. For example, the on-demand ability to supervise or see what an xAI system is doing allows pathologists to be firmly in control—this should make it much more likely that pathologists will be able to detect whether a computational system is functioning as intended. This can help satisfy regulatory bodies because an xAI system augments pathologists and works with them but does not replace them; this is a critical point that can be misunderstood or that can be misrepresented.

A second regulatory issue regards static versus continually improving computational pathology systems. With prior-generation image analysis systems, algorithms would be cleared or approved by FDA or other groups in a locked-down state. In our experience this led to ossification of such systems, which cannot easily be improved or modified, as evidenced by the stagnation of simple biomarker image analysis systems that rely on equivalency studies based on

old original clearances. However, xAI offers an opportunity to provide continual improvement of algorithms that is well supervised and is based upon feedback from QA data and/or pathologist decision making. There has already been one FDA drug approval that was partially based upon nontrial, clinical patient data; xAI system data may present a similar opportunity.[37] Within carefully structured frameworks, xAI might be a key factor that provides regulatory agencies confidence that computational pathology systems can continually improve.

Finally, there is a growing movement or consensus that concerns ML and human rights. A proposed regulation before the European Union would prohibit "automatic processing" unless people are safeguarded (https://ec.europa.eu/futurium/en/ai-alliance-consultation). People now have a "right to an explanation" concerning algorithm-created decisions that are based on personal information. Future laws may further restrict AI use in professional practices, which represents a huge challenge to industry. Radiologists are already beginning to recognize the need for standards, for physician education, and for AI-specific guidelines in their practice, as exemplified in a draft document issued by The Royal Australian and New Zealand College of Radiologists.[38] Of the 8 proposed principles, the first 3 regard safety, avoidance of bias, and transparency; this is a clear demand for xAI and it is very likely that similar demands will be made of computational pathology systems as well.

From a health care system perspective, computational pathology is also desirable. Such systems can facilitate quickly getting the right case to the right pathologist (ie, generalist vs. subspecialist expert), which can reduce wasteful, unnecessary work-ups and consultations for some difficult cases. Intelligent xAI tools could also help pathologists perform at a higher level, thereby improving patients' access to higher quality diagnoses. Improvements in diagnosis could also lead to more efficient and accurate patient management by clinicians. For example, over-diagnosed breast core biopsies can lead to unnecessary surgical procedures, subsequent high-risk breast cancer screening, additional future biopsies related to extra screening, and lower quality of life for patients who endure the consequences. There is also a heavy potential consequence for under-diagnosed biopsies as well, due to delayed diagnosis and/or missed opportunity for higher-risk screening.

There is also a not-so-distant possibility of being able to learn more from pathology specimens than is currently possible using manual light microscopy. Computational pathology may ultimately enable highly sophisticated analytics of morphologic features, hyperplexed biomarker patterns, and regional genomics; leading to much more powerful diagnosis, prediction, and prognostication.

Computational pathology is poised to revolutionize digital pathology by delivering meaningful automation of anatomic pathology by augmenting pathologists without replacing them. There are important concerns about such automation that revolve around trusting and validating computational pathology tools. Using xAI, HistoMapr-Breast is a type of software platform that can move forward in a transparent manner, by permitting pathologists to understand how it is working and why it makes decisions (Fig. 4). It is anticipated that such an approach will also ease regulation of computational pathology software because transparency facilitates supervision and real-time monitoring in order to ensure safety. xAI guidance ensures that pathologists have all the necessary information to make the best diagnoses they can make. We have implemented this with HistoMapr, a transparent xAI system that can work for pathologists in a trustworthy fashion using its explanation interface.

Pathology is often considered to be the gold standard for patient diagnosis, and therefore pathologists have been conservative in adopting digital pathology. We regard xAI as a critical feature that differentiates computational tools like HistoMapr from other ML software. It is critical to emphasize that the purpose of the xAI is not to make a diagnosis independent from the pathologist, but to assist the pathologist in being more accurate and efficient. This approach is intended to maximize pathologists' ability to serve their patients and clinical colleagues.

REFERENCES

1. Volynskaya Z, Evans AJ, Asa SL. Clinical applications of whole-slide imaging in anatomic pathology. *Adv Anat Pathol*. 2017;24:215–221.
2. Cornish TC, Swapp RE, Kaplan KJ. Whole-slide imaging: routine pathologic diagnosis. *Adv Anat Pathol*. 2012;19:152–159.
3. Philips IntelliSite Pathology Solution (PIPS). US Food and Drug Administration. 2017. Available at: www.accessdata.fda.gov/cdrh_docs/pdf16/DEN160056.pdf. Accessed March 22, 2020.
4. Aperio AT2 DX System US Food and Drug Administration. 2019. Available at: www.accessdata.fda.gov/cdrh_docs/pdf19/K190332.pdf. Accessed March 22, 2020.
5. Louis DN, Feldman M, Carter AB, et al. Computational pathology: a path ahead. *Arch Pathol Lab Med*. 2016;140:41–50.
6. Fuchs TJ, Buhmann JM. Computational pathology: challenges and promises for tissue analysis. *Comp Med Imag Graphics*. 2011;35:515–530.
7. Leong F, Joel WM, Leong ASY. Digital imaging applications in anatomic pathology. *Adv Anat Pathol*. 2003;10:88–95.
8. Fine JL. 21(st) century workflow: a proposal. *J Pathol Inform*. 2014;5:44.
9. Farahani N, Liu Z, Jutt D, et al. Pathologists' computer-assisted diagnosis: a mock-up of a prototype information system to facilitate automation of pathology sign-out. *Arch Pathol Lab Med*. 2017;141:1413–1420.
10. Tosun AB, Nguyen L, Ong N, et al. Histological Detection of High-Risk Benign Breast Lesions from Whole Slide Images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Cham: Springer International Publishing; 2017:144–152.
11. Huang H, Tosun AB, Guo J, et al. Cancer diagnosis by nuclear morphometry using spatial information. *Pattern Recognit Lett*. 2014;42:115–121.
12. Tosun AB, Yergiyev O, Kolouri S, et al. Detection of malignant mesothelioma using nuclear structure of mesothelial cells in effusion cytology specimens. *Cytometry A*. 2015;87:326–333.
13. Jones T, Nguyen L, Tosun AB, et al. Computational Pathology versus Manual Microscopy: Comparison Based on Workflow Simulations of Breast Core Biopsies. In: Laboratory Investigation. Nature Publishing Group; 97;2017:398a.
14. Nguyen L, Tosun AB, Fine JL, et al. Spatial statistics for segmenting histological structures in H&E stained tissue images. *IEEE Trans Med Imaging*. 2017;36:1522–1532.
15. Tosun AB, Gunduz-Demir C. Graph run-length matrices for histopathological image segmentation. *IEEE Trans Med Imaging*. 2011;30:721–732.
16. Tosun AB, Kandemir M, Sokmensuer C, et al. Object-oriented texture analysis for the unsupervised segmentation of biopsy images for cancer detection. *Pattern Recognit*. 2009;42:1104–1112.
17. Ozolek JA, Tosun AB, Wang W, et al. Accurate diagnosis of thyroid follicular lesions from nuclear morphology using supervised learning. *Med Image Anal*. 2014;18:772–780.

18. Nguyen L, Tosun AB, Fine JL, et al. Architectural patterns for differential diagnosis of proliferative breast lesions from histopathological images. in 2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017); 2017:152–155.
19. Bejnordi BE, Veta M, Van Diest PJ, et al. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *JAMA*. 2017;318:2199–2210.
20. Liu Y, Gadepalli K, Norouzi M, et al. Detecting cancer metastases on gigapixel pathology images. arXiv preprint arXiv:1703.02442; 2017.
21. Bejnordi BE, Mullooly M, Pfeiffer RM, et al. Using deep convolutional neural networks to identify and classify tumor-associated stroma in diagnostic breast biopsies. *Mod Pathol*. 2018;31:1502–1512.
22. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat Mach Intell*. 2019;1:206–215.
23. Gunning D. Explainable artificial intelligence (xai). Defense Advanced Research Projects Agency (DARPA), nd Web; 2017.
24. Hoffman RR, Mueller ST, Klein G, et al. Metrics for Explainable AI: Challenges and Prospects. arXiv preprint arXiv:1812.04608; 2018.
25. Samek W, Wiegand T, Müller K-R. Explainable artificial intelligence: understanding, visualizing and interpreting deep learning models. arXiv preprint arXiv:1708.08296; 2017.
26. Elmore JG, Longton GM, Carney PA, et al. Diagnostic concordance among pathologists interpreting breast biopsy specimens. *JAMA*. 2015;313:1122–1132.
27. Montalto MC. An industry perspective: an update on the adoption of whole slide imaging. *J Pathol Inform*. 2016;7:18.
28. Thrall MJ. Automated screening of Papanicolaou tests: a review of the literature. *Diagn Cytopathol*. 2019;47:20–27.
29. Spagnolo DM, Gyanchandani R, Al-Kofahi Y, et al. Pointwise mutual information quantifies intratumor heterogeneity in tissue sections labeled with multiple fluorescent biomarkers. *J Pathol Inform*. 2016;7:47.
30. Chennubhotla SC, Taylor DL. Systems and methods for finding regions of interest in hematoxylin and eosin (h&e) stained tissue images and quantifying intratumor cellular spatial heterogeneity in multiplexed/hyperplexed fluorescence tissue images. US Patent Application 15/577,838, July 19, 2018. 2016. Available at: https://patents.google.com/patent/WO2016201186A1/en. Accessed May 25, 2020.
31. Rosai J. Borderline epithelial lesions of the breast. *Am J Surg Pathol*. 1991;15:209–221.
32. Dundar MM, Badve S, Bilgin G, et al. Computerized classification of intraductal breast lesions using histopathological images. *IEEE Trans Biomed Eng*. 2011;58:1977–1984.
33. Murray M. Pathologic high risk lesions, diagnosis and management. *Clin Obstet Gynecol*. 2016;59:727–732.
34. Campanella G, Hanna MG, Geneslaw L, et al. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nat Med*. 2019;25:1301–1309.
35. Gunning D, Chaudhri VK, Clark PE, et al. Project halo update—progress toward digital aristotle. *AI Magazine*. 2010;31:33–58.
36. Nguyen A, Yosinski J, Clune J. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2015:427–436.
37. Goldberg P. How FDA, Pfizer, and Flatiron Health did it: approval of Ibrance for men affords a glance at use of real world data. The Cancer Letter (cancerletter.com). April 19, 2019. Available at: https://cancerletter.com/articles/20190419_1. Accessed February 6, 2020.
38. The Royal Australian and New Zeland College of Radiologists. Ethical Principles for AI in Medicine. 2019. Available at: https://www.ranzcr.com/college/document-library/ethical-principles-for-ai-in-medicine. Accessed May 25, 2020.