Lessons Learned from Designing an AI-Enabled Diagnosis Tool for Pathologists

HONGYAN GU, University of California, Los Angeles, USA

JINGBIN HUANG, University of California, Los Angeles, USA

LAUREN HUNG, Carnegie Mellon University, USA

XIANG 'ANTHONY' CHEN, University of California, Los Angeles, USA

10

Despite the promises of data-driven artificial intelligence (AI), little is known about how we can bridge the gulf between traditional physician-driven diagnosis and a plausible future of medicine automated by AI. Specifically, how can we involve AI usefully in physicians' diagnosis workflow given that most AI is still nascent and error-prone (e.g., in digital pathology)? To explore this question, we first propose a series of collaborative techniques to engage human pathologists with AI given AI's capabilities and limitations, based on which we prototype Impetus — a tool where an AI takes various degrees of initiatives to provide various forms of assistance to a pathologist in detecting tumors from histological slides. We summarize observations and lessons learned from a study with eight pathologists and discuss recommendations for future work on human-centered medical AI systems.

CCS Concepts: • **Human-centered computing** → **Empirical studies in HCI**; **Collaborative interaction**; **Collaborative and social computing systems and tools**.

Additional Key Words and Phrases: Digital pathology; Medical AI; Human-AI collaboration; Human-centered AI;

ACM Reference Format:

Hongyan Gu, Jingbin Huang, Lauren Hung, and Xiang 'Anthony' Chen. 2021. Lessons Learned from Designing an AI-Enabled Diagnosis Tool for Pathologists. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 10 (April 2021), 25 pages. <https://doi.org/10.1145/3449084>

1 INTRODUCTION

Recent advancements in machine learning techniques are rejuvenating the use of artificial intelligence (AI) in medicine that originally started over half a century ago. Enabled by data-driven statistical models, AI can already read X-Ray images [14, 32] and analyze histological slides [13, 44] with a performance on par with human experts.

Despite their promises of automating diagnosis, existing medical AI models tend to be 'imperfect' [11] — there remain inherent limitations in the models' performance and generalizability. For example, in digital pathology, scanned tissue slides are processed by AI to detect tumor cells. The problem is that such histological data (e.g., ovarian carcinoma) tends to have a high between-patient variance [35]; thus, a pre-trained model often struggles to generalize when deployed to a new set of patients. At present, it remains underexplored how to integrate such 'imperfect' AI usefully into physicians' existing workflow.

Authors' addresses: Hongyan Gu, ghy@ucla.edu, University of California, Los Angeles, USA; Jingbin Huang, jihuang97@ucla.edu, University of California, Los Angeles, USA; Lauren Hung, hsienhuh@andrew.cmu.edu, Carnegie Mellon University, USA; Xiang 'Anthony' Chen, xac@ucla.edu, University of California, Los Angeles, USA.

This work is licensed under a Creative Commons Attribution International 4.0 License.

2573-0142/2021/4-ART10. <https://doi.org/10.1145/3449084>

Researchers have long realized the limitation of using AI as a ‘Greek Oracle’. Miller and Masarie pointed out that a “mixed-initiative system” is mandatory whereby “physician-user and the consultant program should interact symbiotically” [43]. Some research focused on mimicking how doctors think [21, 54], such as using an attention-guided approach to extract local regions of interest on a thoracic X-ray image [25]; others developed explainable models [16, 69] or system designs [60, 62] that promote a doctor’s awareness of AI’s diagnosis process. Cai *et al.* developed a content-based image retrieval (CBIR) tool that allows a pathologist to search for similar cases by region, example, or concept [11]. Yang *et al.* conducted fieldwork to identify when and how AI can fit in the decision-making process of vascular assist device transplant [66, 67]. Despite such a growing body of work on mental models, explainability, CBIR tools, and field study, little has been done to answer the following question for medical AI: when AI is still nascent and error-prone, how can physicians still make use of such ‘imperfect’ AI in their existing workflow of diagnosis?

To ground the exploration of this question, we focus on medical imaging — the primary data sources in medicine [34]. Amongst various medical imaging techniques, histological data in digital pathology [28], in particular, presents some of the most difficult challenges for achieving AI-automated diagnosis, thus serving as an ideal arena to explore the interactional relationship between physicians and AI.

Focusing on digital pathology as a point of departure, we propose a series of physician-AI collaboration techniques, based on which we prototype Impetus — a tool where an AI aids a pathologist in histological slide tumor detection using multiple degrees of initiative. Trained on a limited-sized dataset, our AI model cannot fully automate the examination process; instead, Impetus harnesses AI to (i) guide pathologists’ attention to regions of major outliers, thus helping them prioritize the manual examination process; (ii) use agile labeling to train and adapt itself on-the-fly by learning from pathologists; and (iii) take initiatives appropriately for the level of performance confidence, from full automation, to pre-filling diagnosis, and to defaulting back to manual examination. We used the Impetus prototype as a medium to engage pathologists and observe how they perform diagnosis with AI involved in the process and elicit pathologists’ qualitative reactions and feedback on the aforementioned collaborative techniques. From work sessions with eight pathologists from a local medical center, we summarize lessons learned as follows.

Lesson #1 To explain AI’s guidance, suggestions and recommendations, the system should go beyond a one-size-fits-all concept and provide instance-specific details that allow a medical user to see evidence that leads to a recommendation.

Lesson #2 Medical diagnosis is seldom a one-shot task, thus AI’s recommendations need to continuously direct a medical user to filter and prioritize a large task space, taking into account new information extracted from a user’s up-to-date input.

Lesson #3 Medical tasks are often time-critical, thus the benefits of AI’s guidance, suggestions and recommendations need to be weighed by the amount of extra efforts incurred and the actionability of the provided information.

Lesson #4 To guide the examination process with prioritization, AI should help a medical user narrow in small regions of a large task space, as well as helping them filter out information within specific regions.

Lesson #5 It is possible for medical users to provide labels during their workflow with acceptable extra effort. However, the system should provide explicit feedback on how the model improves as a result, as a way to motivate and guide medical users' future inputs.

Lesson #6 Tasks treated equally by an AI might carry different weights to a medical user. Thus for medically high-staked tasks, AI should provide information to validate its confidence level.

Importantly, these lessons reveal what was unexpected as pathologists collaborated with AI using Impetus' techniques, which we further discuss as design recommendations for the future development of human-centered AI for medical imaging.

1.1 Contributions

Our contributions are as follows.

- The first suite of interaction techniques in medical diagnosis that instantiate mixed-initiative principles [29] for physicians to interact with AI with adaptive degree of initiatives based on AI's capabilities and limitations;
- A proof-of-concept system that embodies these techniques as an integrated diagnostic tool for pathologists to detect tumors from histological slides;
- A summary of observations and lessons learned from a study with eight pathologists that provides empirical evidence of employing mixed-initiative interaction in the medical imaging domain, thus informing future work on the design and development of human-centered AI systems.

2 RELATED WORK

Our review of literature starts from a general body of cross-disciplinary work on human-AI interaction, gradually drills down to the (medical) imaging domain, and finally summarizes the current status on digital pathology, which exemplifies the gap between traditional manual diagnosis and not-yet-available AI-enabled automation.

2.1 Human-AI Collaborative Work

Since J. C. R. Licklider's vision of 'man-machine symbiosis' [37], bringing human and AI to work together collaboratively has been a long-standing challenge across multiple fields.

Recent work has been employing human-AI interaction to utilize human input better and reduce manual effort in repetitive routines. Specifically, Amershi *et al.* propose a system that gives the user flexibility to provide better training examples in interactive concept learning [1]. The system also grants users control over the training process: users could decide to stop training to avoid overfitting. Chau *et al.* combine visualization, user interaction, and machine learning to guide users to explore correlations and understand the structure of large-scale network data [17]. Suh *et al.* show that classifier training with mixed-initiative teaching is advantageous over both computer-initiated and human-initiated counterparts. Specifically, mixed-initiative training could significantly reduce the labeling complexity across a broad spectrum of scenarios, from perfect, helpful teachers who always provide the most helpful teaching, to 'naïve' teachers who give unhelpful labels [57]. Felix *et al.* propose a topic modeling system that could find unknown labels for a group of documents: by integrating human-driven label refinement and machine-driven label recommendations, the system enables analysts to explore, discover and formulate labels [24].

Research has also shown that human-AI collaboration can enhance domain-specific tasks. For example, Forté enables designers to interactively express and refine sketch-based 2D design automated by topology optimization: specifically, a user can modify the optimization's result, which serves as input for the next iteration to reflect the user's intent [18]. Nguyen *et al.* combine human knowledge, automated information retrieval, and machine learning in a mixed-initiative approach to conduct fact-checking [45]. However, in the medical domain, Yang *et al.* point out that the involvement of the machine-initiative often fails to assist physicians in clinical practices [67]. In particular, "the disruptive, time-consuming" machine work would "conflict with the chaotic nature of clinical work", which undermines the mental need for machine support for medical users. Building on this finding, Yang *et al.* further adapted the concept of unremarkable computing, where doctors could seek machines' AI support on-demand to reduce AI's disruptive behavior [66].

To democratize the design of human-AI collaboration, Horvitz articulated a series of principles of mixed-initiative interaction via the example of an email-to-calendar scheduling agent [29]. Insights from Horvitz's work were renewed in a recent paper by Amershi *et al.*, which proposes and validates 18 guidelines for human-AI interaction, which includes, for instance, "support efficient correction", "make clear why the system did and what it did", "convey the consequences of user actions" [2]. However, it is "uniquely difficult" [65] to implement those guidelines. Due to the "uncertainty surrounding AI's capabilities" and "complexity in AI's output", it is hard for humans to control and trust AI systems [65]. Given such challenges in system design, human-AI collaboration guidelines in the medical domain focus on building human-centered systems that revolve around AI *as-is*. When AI makes mistakes, instead of modifying or correcting the model, current work emphasizes informing users of AI's pitfalls to ensure transparency [12]. To the best of our knowledge, there remains a lack of research on discussing how human-AI collaboration could provide new knowledge to AI in the medical domain.

Going beyond collaborating with AI *as-is*, our work is grounded in a proof-of-concept system situated in a specific application context, through which we found that the sheer amount of high-resolution medical imaging data posts unique challenges: (i) how to filter AI's analyses on such high-resolution data and communicate actionable information to advance a differential diagnosis that is by nature iterative and (ii) how to cost-effectively incorporate pathologists' input as new knowledge to AI without taxing them with labeling a large amount of data.

2.2 Data-Driven Digital Image Processing

Imaging provides an abundant source of clinical data in medicine [34]. Furthermore, data-driven AI has served as a powerful toolkit for processing digital images (e.g., the recent breakthrough in deep models for examining chest X-ray [32]). Remarkably, the increase of data availability in digital pathology [38, 47, 58, 61] has triggered a recent surge of data-driven techniques in a board range of applications, such as carcinoma detection [4, 7, 8], histological features detection [19, 59], and tumor grading [5, 22].

While some works mentioned above have reported expert-level performance of AI models, human involvement remains indispensable, primarily in the provision of training labels. However, labeling medical data is a non-trivial task, since it usually suffers from a high variance in tissue appearance [35], subjectivity in medical guidelines [12], and, most importantly, a rare availability of trained specialists [50]. Those barriers result in a low throughput in the medical image processing pipeline. Schaekermann *et al.* try to break this dilemma by training more-available general clinicians — their comprehension of difficult cases could be improved by being exposed to specialist adjudication feedback [50]. However, the validation study shows that labeling performance does not increase significantly, since training humans without comprehensive, informative teaching would incur confusion.

In lieu of training humans, another way to reduce labeling cognitive workload is by using human labels more efficiently. Specifically, multiple works [44, 56, 71] have employed the concept of active learning [53], where machines could learn from previous input iteratively and recommend the most uncertain cases to users for annotation. For example, Sommer *et al.* propose Ilastik for cell segmentation model training, which allows users to annotate by drawing strokes over cells [56]. However, the microscopic nature of stroke annotations demands much user effort to achieve good performance on the gigapixel whole-slide level. Nalysnik *et al.* implement HistomicsML to perform nucleus identification with a random forest classifier, which dynamically recommends the most uncertain patches for annotation in each training iteration [44]. Going beyond selecting the most uncertain samples, Zhu *et al.* add a diversity constraint to reduce recommended samples' over-concentration in a localized area [71]. However, due to the high variation of histological features, merely relying on limited human annotations without exposing the model to a comprehensive training set would cause bias. This would result in a dispersion of sensitive false-positives across the whole-slide image, and humans would lose trust in machines while being overwhelmed with false-positive information.

To address this 'low information input' issue without significantly increasing human burden, recent works [13, 30, 64] have applied multiple instance learning (MIL) [6, 70] approaches to digital pathology, where MIL learners could even learn from whole slide-level labels, thus dispensing with the need for users to label at the pixel-level. For example, Xu *et al.* employ MIL [64] to train a patch-level classifier to identify colon cancer based on slide-level annotations. Ilse *et al.* use an attention-based deep MIL approach to learn a Convolutional Neural Network (CNN) with image-level annotations. The learned CNN can highlight breast or colon cancer areas, with reported performance on par with other MIL approaches without sacrificing interpretability [30]. Campanella *et al.* combine MIL with Recurrent Neural Network (RNN) techniques and train with slide-level diagnosis for prostate cancer, basal cell carcinoma, and breast cancer metastases [13]. However, it usually requires a large amount of slide-level annotations for training (for example, [13] collected a database of >44,000 slides from >15,000 patients); otherwise, there is often a performance drop compared to using strongly-supervised labels on the same set of slides.

Different from previous work, we focus on training machines with reduced human burden, addressing the issue of adopting a generic, pre-trained model on a new set of data. Specifically, we apply a CNN feature extractor that has been previously trained on a larger dataset. The pre-trained CNN would generate a local embedding for the new data, which addresses the 'low information input' in traditional active learning techniques. Then, instead of training an MIL learner end-to-end, we implement an MIL learner that can learn jointly from the data embedding and user input labels. This training scheme harnesses the benefit of reducing human annotation labels with MIL technique, while avoiding having to retrain an MIL learner from scratch.

2.3 Interactive Tools for Digital Pathology

Digital pathology, similar to biology research, often deals with high-resolution, visually challenging images. Beyond involving data-driven models trained by domain experts, tools that allow pathologists to define, explore, and decide upon clinical or research problems are also needed.

ImageJ [52] and its distributions [20, 48, 51] are the primary scientific image analysis tools for digital pathology. They not only allow pathology researchers to perform image operations (e.g., cropping, measuring, editing), but support a wide range of plugins for immunohistochemistry (IHC) and fluorescence analysis [33]. Besides ImageJ, multiple interactive tools have been proposed to aid medical users to automate whole-slide image (WSI) analysis without coding, covering domains of phenotype analysis [15], segmentation [49], and IHC screening [42]. Recent research [62, 63, 66, 67] suggests that, besides reasoning with medical data, the design of a diagnosis tool

often needs to consider physicians' established workflow and other domain-specific behaviors. Some digital pathology tools further consider the social background of clinical use and support database constructions [26] and remote collaborations [41].

When it comes to diagnosis, Blois considers the relationship between physicians and computer systems as a 'funnel' as in Figure 2 [9]: from A to B, a physician needs to gather information, formulate hypotheses, perform tests and gradually arrive at a diagnosis. Many of the aforementioned tools tend to focus on and perform well at or near Point B, *e.g.*, performing a physician-defined measuring task. In comparison, we design Impetus to support physicians' examination at or near Point A as well, *e.g.*, suggesting where on a WSI to examine first.

Apart from this work, some recent interactive tools have also targeted supporting physicians in a border context. Specifically, Hegde *et al.* propose a deep-learning reverse image searching tool to help pathologists search for images with similar tissue type, histological feature, or disease state [27]. Given a specific image patch, the tool first calculates image embeddings through a deep-learning model. Then, it returns the images' nearest neighbors in the embedding space as a searching result. Building upon the searching algorithm, Cai *et al.* build a CBIR system that allows users to retrieve similar image patches from a database. Meanwhile, users could refine the searching results on-the-fly by regions, examples, or pre-defined histological concepts to deal with the searching algorithm's imperfection behavior [11]. Our work differs from [11, 27] in two ways: (i) instead of enabling users to *adjust* imperfect *searching* results while leaving the underlying model untouched, our tool seeks to address imperfect *carcinoma detection* algorithm by cost-effectively incorporating pathologists' input; (ii) due to the difference in tasks (retrieval vs. diagnosis), our work naturally goes beyond concentrating a localized region and discovers suspicious patterns with spatial structures at a whole-slide level, giving users a global view of the model behavior.

3 IMPETUS: AN AI-ENABLED TOOL FOR PATHOLOGISTS

Before we unfold our design process in the next section, we first give a background introduction on digital pathology and the motivation to involve AI. We then walk through Impetus's scenario to present a high-level overview of how the tool works with pathologists.

3.1 Background on Digital Pathology

Central to digital pathology are whole-slide images or WSI for short. WSIs are produced by high-speed slide scanners that digitalize the glass slide at very high resolution, resulting in gigapixel images [38]. Due to its large and high visual variance, a WSI cannot be directly fed into a model (such as a CNN) for classification. A WSI is usually divided into small patches, which are then classified by a CNN model (Figure 4(a)). These patch level predictions can then be assembled to create a tumor probability heatmap, from which a pathologist can derive a whole-slide level diagnosis.

In our study, we used a dataset containing Hematoxylin and Eosin (H&E) stained sentinel lymph node (SLN) sections of breast cancer patients [38]. The diagnosis of such specimens contains four main categories [3]:

- **Isolated tumor cells (ITC)** if the node contains a single tumor cell or cell deposits that are no larger than 0.2 mm or contain fewer than 200 cells;
- **Micro** if containing metastasis greater than 0.2 mm or more than 200 cells;
- **Macro** if containing metastasis greater than 2 mm;
- **Negative** if containing no tumor cells.

3.2 Promises & Challenges of AI for Digital Pathology

Digital pathology transforms traditional microscopic analysis of histological slides into high-resolution digital visualization [28]. Digital pathology allows pathologists to investigate fine-grained pathological information, transfer previously-learned knowledge to new tasks [28], and, most importantly, leverage the recent development of data-driven AI to augment their visual analytical tasks.

However, the main challenge for digital pathology is that, unlike other imaging modalities (*e.g.*, X-Ray, CT), histological data (*e.g.*, ovarian carcinoma) tends to have a high between-patient variance [35]; thus, a pre-trained model often struggles to generalize when deployed to a new set of patients. Such an uncertainty of performance creates a barrier that prevents AI from being adopted to assist diagnosis in digital pathology.

To overcome this barrier, one solution is to improve the machine learning model by training it on a sufficiently large amount of patient data using cost-effective labeling and learning schemes [13, 44]. However, such a ‘big data’ approach attempts to close the gap by (marginally) improving AI’s performance, while ignoring the opportunity to engage human physicians. As a result, efforts are often bound to repeat the ‘Greek Oracle’ pitfall pointed out by Miller and Masarie almost three decades ago [43]. The focus of our paper is to explore and study the oft-missed opportunity of combining physicians with an ‘imperfect’ AI: rather than awaiting AI to be fully automatable one day, how can we make use of its capability with limitations today?

Below we describe a scenario walkthrough of Impetus — a tool that explores how AI — without yet the ability to diagnose fully-automated — can still empower pathologists by becoming an integral part of their workflow.

3.3 Scenario Walkthrough

The user of Impetus, a pathologist, starts diagnosing a patient’s case by importing multiple Whole Slide Images (WSI) of the patient into Impetus.

First, the pathologist’s attention is drawn to the two boxes generated by the AI, which encompass regions of patches that visually appear to be ‘outliers’ from the majority of cells (Figure 1(a)), which suggests that these patches are likely to be tumor-positive. With these automatic recommendations, Impetus alleviates the pathologist’s burden of navigating a large, high-resolution image and having to go through a large number of areas that might or might not be as tumor-characteristic as the recommended regions.

Next, the pathologist performs diagnosis by marking each recommended region as either ‘tumor’ or ‘normal’, and continues to marquee-select and label a few more regions on the WSI (Figure 1(b)). As the pathologist makes such selections, their input is also collected by the back end AI and used as labels to adapt the model better to align itself with the pathologist’s domain knowledge.

Based on these diagnostic inputs and revisions from the pathologist, Impetus immediately adapts the underlying AI model accordingly. In contrast to conventional data labeling tasks, Impetus’ agile labeling is designed to be lightweight and can learn from pathologists’ input of coarsely marked regions without having to trace a precise contour of a tumor region. In this way, Impetus allows pathologists to agilely train an AI model as a natural and integral part of their existing workflow without incurring extra effort.

As the pathologist annotates more WSIs (which also trains the AI), they notice that some new slides are already marked as ‘diagnosed’ — AI takes the initiative to diagnose slides that it feels highly confident about. Thus the pathologist skips ahead to see other unlabeled slides, some of which, have pre-filled diagnosis dialogues (Figure 1(c)). In such cases, the pathologist examines the

Fig. 1. Key interactive features of Impetus: (a) as a pathologist loads a whole slide image, AI highlights areas of interest identified by outlier detection, shown as two yellow recommended boxes. (b) Agile labeling: a pathologist can drag and click to provide a label that can be employed to train the AI’s model. (c) Diagnosis dialogue, pre-filled with AI’s diagnosis, allows the pathologist to either confirm or disregard and proceed with manual diagnosis.

WSI to verify the AI’s hypothesis. In the rest of the WSIs, the AI almost becomes invisible (due to a lack of confidence), and the pathologist proceeds to finish the diagnostic tasks manually.

The above scenario demonstrates how an ‘imperfect’ AI can still benefit a pathologist without necessarily automating the user’s existing workflow: recommendation boxes suggestively prioritize pathologists’ manual searching process (Figure 1(a)), agile labeling adapts AI while minimizing the extra effort from the pathologists (Figure 1(b)), and as AI attempts to improve itself, it handles cases with different degrees of initiatives — from full automation to pre-filling plausible results to remaining complete ‘invisible’—based on its confidence (Figure 1(c)).

4 DESIGN AND IMPLEMENTATION

Below we first describe the design process then detail the specific interaction techniques and their implementation in Impetus.

4.1 Overview of the Design Process: Empirical & Theoretical Grounding

The design of Impetus is grounded in both empirical evidence and principles drawn from literature.

Fig. 2. A physician's differential diagnosis process is similar to a funnel, starting with a broad exploration of plausible conditions and gradually rule out less likely possibilities as more evidence (e.g., test results) is gathered until finally a single most probable conclusion can be drawn. Beyond mixed-initiatively automating certain diagnosis (near Point B), Impetus also supports exploration near Point A by enabling pathologists' initial exploration with recommended regions. Image modified based on Blois [9].

On the empirical side, we co-designed Impetus with our pathologist collaborator. Specifically, we learned that one major challenge for pathologists is efficiently and effectively navigating large, high-resolution WSIs. This suggests that AI, besides making diagnosis, can usefully serve to guide pathologists to navigate complex and high-resolution image space. We detail this design in §4.2.

On the theoretical side, Impetus goes beyond the singular objective of automation by offering a spectrum of AI-enabled assistance. As pointed out by Blois' seminal paper [9], a physician's differential diagnosis process is similar to a funnel, starting with a broad exploration of plausible conditions and gradually rule out less likely possibilities as more evidence (e.g., test results) is gathered until finally a single most probable conclusion can be drawn. According to Blois, AI has been canonically developed to optimize Point B, where a computer program can deterministically confirm whether a patient has a certain disease given all the evidence. As Blois foresaw, a recent development of AI starts to exhibit capabilities towards Point A, e.g., Stanford's CheXpert produces likelihoods of 10+ thoracic diseases based on a chest X-ray image [32]. Similarly, Impetus also aims at "reaching Point A" by enabling pathologists' initial exploration with recommended regions.

Overall, Impetus provides the first suite of interaction techniques in the medical imaging domain that instantiates mixed-initiative principles [29] for physicians to interact with AI with an adaptive degree of initiatives based on AI's capabilities and limitations. Specifically, we focus on the following principles in [29]:

- *Scoping precision of service to match uncertainty.* We first design a rule-based algorithm to identify three levels of uncertainty in AI's performance given a WSI (detailed in §4.4), based on which we then design the corresponding AI-initiated action appropriate for each level of uncertainty (Table 1).
- *Providing mechanisms for efficient agent-user collaboration to refine results.* For each AI-initiated action, we also design mechanisms to introduce physician-initiated actions aimed at confirming, refining, or even overriding AI's results (Table 1). Further, we extend this principle by leveraging physician-initiated input for 'machine teaching' [55], i.e., an agile labeling technique to dynamically adapt an AI by retraining it with examples of how a physician interpret a patient's histological data (detailed in §4.3).

Table 1. Spectrum of human and AI initiatives at different AI confidence levels.

AI Confidence	AI-Initiated Action	Physician-Initiated Action
High	Performing diagnosis automatically in the background; marking WSIs as diagnosed	Doing nothing and accepts AI's results; can re-open a WSI to overwrite AI's result
↑	Pre-filling the diagnosis box without directly labeling the WSI	Performing diagnosis with help from AI predictions; confirming or correcting the pre-filled dialogue
Low	Showing original WSI by default to prompt for manual diagnosis	Performing diagnosis with little input from AI

4.2 AI Guiding Pathologists' Attention to Regions of Major Outliers

In our communication with our pathologist collaborator, we learned that one major limitation of pathologists is the ability to efficiently and effectively navigate a large, high-resolution WSI. To address this limitation, we design AI to guide pathologists' attention to regions of major outliers that appear visually different from the rest of the WSI and are more likely to be tumors. Such guidance is manifested in two user interface elements:

(i) **Attention map** visualizes each patch's degree of outlying overlaid on the current WSI (Figure 3(a)); (ii) **Recommendation boxes** as a more explicit means to draw pathologists' attention to large clusters based on outlier detection results (Figure 3(a), yellow box) — these boxes are always visible, whether on the original WSI, on the attention, or on the prediction map (described below).

Implementation When the system is first loaded, a pre-trained InceptionResNetv2 model¹ [31] on PatchCamelyon dataset² extracts patch features (patch dimension= $96 \times 96 \times 3$, feature dimension= 1536×1) in WSIs (Figure 4(a)). Given the imbalance nature of tumor vs. normal tissues, in the first iteration, the system performs isolation forest (max_samples=256) [39] outlier detection based on extracted features (Figure 4(b)), and the detected outliers are highlighted in the attention map. In the following iterations, the attention map is a combination of outliers (from the initial detection) and high uncertainty patches (from specific models in each iteration)³. In order to obtain the recommendation boxes, the system uses a DBSCAN clustering algorithm (min_sample=10, epsilon=3) [23] to cluster WSI patches with attention value (Figure 4(c,k)). In order to reduce users' distraction, the recommendation boxes are selected as the two clusters that occupy the largest areas on the WSI in each iteration.

4.3 AI Using Agile Labeling to Train and Adapt Itself On-the-fly

In digital pathology, the main challenge for AI is that, unlike other imaging modalities (e.g., X-ray, CT), histological data (e.g., ovarian carcinoma) tends to have a high variance across slides of different patients (sometimes same patients as well) [35]. Thus a pre-trained model often struggles to generalize to new data. To address this limitation, Impetus enables pathologists to use agile labeling to train AI on the fly.

¹We trained this model with image augmentation preprocessing, Adadelata optimizer with initial learning rate 0.1, binary cross entropy loss, 100 iterations with early stopping on validation loss.

²<https://patchcamelyon.grand-challenge.org/>

³Uncertainty is calculated as $\text{Uncertainty} = 1 - |0.5 - \text{Probability}| \times 2$. The attention maps in the following iterations are calculated as the soft-OR of outliers and uncertainty: $\text{Attention} = \text{Uncertainty} \odot \text{Outlier}$.

Fig. 3. The two maps used by Impetus to provide guidance and communicate AI results. (a) Attention map, where outlier patches and high uncertainty patches are highlighted in red, while other patches are in blue. The yellow recommendation boxes are generated by clustering attention values. (b) Prediction map, where red shows a high probability of tumor, and white shows a low probability of tumor, as predicted by the AI. The green and red boxes are areas of “normal” and “tumor”, as labeled by the pathologist. Recommendation boxes generated by clustering attention values are also visible on this map.

Agile labeling allows a pathologist to directly label on recommendation boxes (Figure 1(a)), or to draw a bounding box of tumor-negative patches (Figure 1(b)), or a box containing a mix of negative and positive cells, which serve as labels to train an existing model further to incorporate pathologists’ domain knowledge. Importantly, such labeling technique is designed to be agilely achievable without incurring significant extra effort that interrupts the main diagnosis workflow.

Implementation Agile labeling does not specifically require users to provide the exact contour of tumor tissues in a WSI, as strongly-supervised learning does. Alternatively, a user can marquee-select a positive box over an area which contains *at least one* tumor patches, or a negative box on *all negative* regions. We implemented a weakly-supervised MIL [6, 70] to learn over such agile labels. To train the model, the system first initializes a *positive_set* and *negative_set*. For each box annotated by a user, Impetus first partitions the WSI areas into $96 \times 96 \times 3$ non-overlapping patches, and extracts the feature of each patch by the pre-trained CNN model from Section 4.2 (Figure 4(a)). Here, we denote each the extracted feature set as X_i and the box-level annotation from user as Y_i .⁴ For a negative box, all the patch features in the box can be included in *negative_set* (Figure 4(f)). For positive boxes, the system uses T-SNE [40] to represent the high-dimension features X_i with two-dimension embedding $\bar{X}_i$ ⁵ (Figure 4(g)). Then, K-Means clustering is used to split $\bar{X}_i$ into

⁴In the MIL setup, each box only has one box-level annotation Y_i .

⁵The embedding $\bar{X}_i$ is used for clustering for two reasons: (i) avoiding K-Means to process high-dimension data, which could prevent clustering performance degradation; (ii) better visualizing the high-dimensional embedding space.

Fig. 4. Overview of Impetus' AI backend. In the first iteration, Impetus first extracts the WSI to non-overlapping patches, followed by (a) feature extraction with a pre-trained CNN (InceptionResNetV2) model; (b) outlier detection by the isolation forest algorithm; (c) outlier clustering with the DBSCAN algorithm. Then, (d) an attention map with outlier clusters and recommendation boxes are generated. In the following iterations, the user first (e) annotates the recommendation boxes with agile labeling. Next, Impetus processes negative annotations by (f) adding negative box features to the negative set. For the positive annotation, Impetus uses (g) T-SNE to reduce the dimension of positive box features and applies (h) K-Means clustering to split them into two clusters. After that, Impetus (i) assigns the two clusters with labels by comparing them to the negative set and only adds features in the positive cluster to the positive set. Last, (j) a random forest classifier learns from the positive and negative set and predicts at a whole-slide level. The attention map and recommendation boxes are generated by (k) clustering from a combination of outliers and uncertain predictions. Procedures between (e - k) are repeated until the doctors are satisfied with the AI performance.

two clusters: $\bar{X}_i^{(1)}, \bar{X}_i^{(2)}$ (Figure 4(h)). After clustering, the algorithm compares the two clusters with negative samples from *negative_set* to pick the real positive cluster. After the positive cluster is recognized, all the instances in the positive cluster are included in the *positive_set* (Figure 4(i)). Finally, a random forest classifier (MIL-RF, 100 trees, max_depth=100) [10] is trained with the obtained *positive_set* and *negative_set*⁶ (Figure 4(j)), and the user can continuously provide more annotations until the trained classifier reaches a satisfactory level of performance.

4.4 AI Taking Initiatives Appropriately for the Level of Performance Confidence

Even with agile labeling, lightweight on-the-fly learning only has limited improvement compared to training extensively offline. Thus it is crucial to convey the level of AI's performance to the pathologists. In Impetus, AI takes initiatives appropriately for its performance confidence level, as manifested in the following two user interface elements: (i) **Prediction map** visualizes current AI's results overlaying the WSI, which serves to inform both the labeling and the usage of the current AI's model (Figure 3(b)). (ii) **Initiatives based on confidence**—the more uncertain the AI 'feels' about a WSI, the less initiative it takes, as shown in Table 1.

Implementation Impetus has a rule-based confidence-level classifier to sort slides into three categories: high-confidence, mid-confidence, and low-confidence. First, predictions of all the patches in the WSI are obtained. A patch has two characteristics: *is_positive* and *is_uncertain*. A patch

⁶The random forest algorithm is used since its "resistance to overfit" [44]. The notion of random forest has been applied to active learning [44], or multiple instance learning algorithms [36].

is positive if the MIL-RF classifier output > 0.5 , and is uncertain if MIL-RF classifier output $\in [0.25, 0.75]$. We empirically summarize the confidence-level decision rules⁷ as follows:

- If there are more than 200 positive patches AND the number of positive patches is greater than twice the number of uncertain patches, then the slide is predicted as high-confidence;
- Else if there are no outlier clusters, then the slide is predicted as low-confidence;
- Else if the number of uncertain patches is greater than 300, then the slide is predicted as low-confidence;
- Else if the number of positive patches is greater than 200, then the slide is predicted as high-confidence;
- For all other cases, the slide is predicted as mid-confidence.

5 WORK SESSIONS WITH PATHOLOGISTS

To validate our design of Impetus, we observed how pathologists used this tool to perform diagnosis on a clinical dataset [38]. Our goal is to study whether the AI in Impetus (*i*) can be compatibly integrated into pathologists' workflow and (*ii*) can provide added values to pathologists' diagnosis process.

Participants We recruited eight medical professionals from the pathology department in UCLA Health. The participants have experiences ranging from 1 to 43 years, including residents, fellows, and attending pathologists.

Data & apparatus We used the Camelyon 17 [38] dataset and selected 16 WSIs⁸ that were collected in the same medical center. Participants interacted with Impetus on a 15-inch laptop computer using a wired mouse. Impetus ran on a Microsoft Windows 10 Operating System using an Nvidia 960M GPU and 16GB RAM.

Design Our discussion with pathologists collaborators and an initial screening survey indicated that there was not a commonly-used digital pathology tool among the participants. To help pathologists calibrate their experience with Impetus, we introduced another tool – ASAP⁹, which represents a very basic manual tool for viewing and annotating digital pathology slides. Each pathologist interacted with both Impetus and ASAP, which were referred to as System A and System B, respectively, to avoid biasing the pathologists. The order of tools was counterbalanced across the eight pathologists. Twelve of the 16 slides were diagnosed using Impetus and the remaining four using ASAP: we chose to keep more slides for Impetus as it was the target of our study, whereas ASAP was just to calibrate pathologists' tool experience.

Tasks & Procedure After briefly introducing the background of computer-assisted diagnosis, we walked each pathologist through a tool and let them practice on a separate toy dataset also gathered from [38]. We then asked questions about how the participant understood different interactive components, whether the tool was easy to learn and use, and whether the tool was helpful to their diagnosis. Then, the primary task began, which was to diagnose the entire group of WSIs using the provided tool in each condition. A trial started with a participant clicking to open a WSI and finished when they selected a diagnosis and clicked the 'Confirm' button. After each condition, we further conducted a brief semi-structured interview for each participant to summarize their experience, feedback, and suggestions for the tool. Participants took a short break between the two conditions.

⁷... which can be easily modified as a configuration of our tool.

⁸Our pilot studies indicated that 16 is the number of WSIs that would allow us to finish the session in about an hour to most effectively use the pathologists' time.

⁹<https://computationalpathologygroup.github.io/ASAP/>

Fig. 5. We conducted work sessions with eight pathologists from a local medical center to observe how they used Impetus as part of their diagnosis process.

Analysis We employed an iterative open-coding method to analyze the qualitative data collected from the semi-structured interviews with pathologists. Two experimenters coded each participant's data within one day after the study. One experimenter performed the first pass of coding and updated a shared codebook, which was then reviewed by the other experimenter to resolve disagreements. The two experimenters alternated the roles of the first coder and reviewer. After all the participants' data were coded and consolidated, a third experimenter reviewed all the codes and transcripts and resolved disagreements through discussion with the previous two experimenters. Finally, we arrived at six high-level themes, which we summarize below as lessons learned.

6 FINDINGS, LESSONS LEARNED & DESIGN RECOMMENDATIONS

Based on the observations and data from the work sessions with pathologists, we present our findings below, which are summarized into six lessons. To maintain consistency, we organize these lessons using the same structure as §5.

6.1 AI Guiding Pathologists' Attention to Regions of Major Outliers

6.1.1 Recommendation boxes. (Figure 1a) were the most frequently used and discussed features during the study. We observed that in almost all the trials, pathologists started by zooming into the recommendation boxes and tried to provide annotations of the outlined region. Pathologists found it helpful to have such concrete start points in their examination.

... [recommendation boxes] narrow down the area of interest ... it helps (P7)

It was less effort because I was focusing only on the attention areas and not focusing on the other areas of the node so it was different from my usual way of looking at a slide. (P2)

However, pathologists did not always find the recommended regions matched their intuition, and they could not understand why certain regions were recommended.

... [the recommendation box] seems a little bit random. It's not necessarily areas that I would [look at] ... (P5)

The things it's focusing on does not correlate with at least what my brain thinks I am looking for. (P6)

A lack of transparency is not a new problem in recommender system research (e.g., [68]). When introducing Impetus, we did explain that recommendation boxes were based on a detection of visual outliers, and all pathologists acknowledged that they understood such a concept. Although such outliers were computed based on histological features (the PatchCamelyon dataset), they did not always agree with what pathologists intuited as ‘interesting’ regions worth examination. When such a mismatch occurred—i.e., an unexpected case of recommendation, pathologists could no longer reason about the recommendation boxes simply by referring to the abstract concept of ‘visual outliers’. At times, pathologists started to develop their own hypothesis of how AI was processing the WSI: “... it’s interesting that it’s picking area with fat as area of interest.” (P2)

Lesson #1 To explain AI’s guidance, suggestions and recommendations, the system should go beyond a one-size-fits-all concept and provide instance-specific details that allow a medical user to see evidence that leads to a recommendation.

Recommendation #1: an overview + instance-based explanation of AI’s suggestions. Currently, Impetus only provides an explanation of the suggested regions at the overview level: a textual description of the outlier detection method as part of the tutorial and visualization (i.e., attention map) that shows the degree of ‘outlying’ across the WSI. As an addition, we can further incorporate instance-based explanation, i.e., with information specific to a particular patient and a particular region on the patient’s slide. The idea is to allow pathologists to question why a specific region is recommended by clicking on the corresponding part of the slide, which prompts the system to show a comparison between the recommended region and a number of samples from non-recommended parts of the slide for the physician to contrast features in these regions extracted by AI. One important consideration is that such an additional explanation should be made available on-demand rather than shown by default, which could defeat the recommendation boxes’ purpose to accelerate the pathologists’ examination process.

We also found that pathologists wondered what they should do about the area outside of the recommendation boxes:

So I just look at the ones in the [recommendation] square? (P7)

Am I supposed to assume the rest of it is normal? I don’t have to go searching for the rest of the slides for [tumor]? (P2)

Pathologists understood the implication *in* the recommendation boxes, i.e., to prioritize certain regions of a WSI and to serve as a ‘shortcut’ in lieu of scanning the entire WSI. However, it was unclear what was the implication *outside* of the recommendation boxes. This is especially true when pathologists could not find signs of tumor in the recommended regions: the system did not continue to guide them on how to proceed with the rest of the WSI.

Lesson #2 Medical diagnosis is seldom a one-shot task, thus AI’s recommendations need to continuously direct a medical user to filter and prioritize a large task space, taking into account new information extracted from a user’s up-to-date input.

Recommendation #2: make AI-generated suggestions always available (and constantly evolving) throughout the process of a (manual) examination. For example, in Impetus, a straightforward design idea is to show recommendation boxes one after another. We believe this is especially helpful when the pathologist might be drawn to a local, zoomed-in region and neglect looking at the rest of the WSI. The always available recommendation boxes can serve as global anchors that inform pathologists of what might need to be examined elsewhere beyond the current

view. This is an example of a multi-shot diagnosis behavior where each shot is an attempt to find tumor cells in a selected region.

6.1.2 Attention map. (Figure 3a) visualizes outliers detected by the AI — the same information based on which the recommendation boxes were drawn. It was designed to complement recommendation boxes with a backdrop of detailed guidance. We expected pathologists to use the attention map similarly as the recommendation boxes, *i.e.*, to direct their attention to look for more outlying regions for examination. However, pathologists did not find attention map useful:

The attention map shows the same thing as the recommended box. The box is enough to direct my attention. (P2)

I don't really see the point of the attention map ... These two maps are redundant. (P4)

The main difference was that recommendation boxes cost less effort to process, while the attention map needed to be navigated (*i.e.*, panned and zoomed and interpreted (*i.e.*, mentally 'decoding' the color scheme). Further, recommendation boxes provided actionable information (*i.e.*, to look into this box first), while the attention map is action-neutral. Given that pathologists' overall goal is to eliminate the amount of area to study, they tended to prefer less extra effort and information with clearer actionability.

Lesson #3 Medical tasks are often time-critical, thus the benefits of AI's guidance, suggestions and recommendations need to be weighed by the amount of extra efforts incurred and the actionability of the provided information.

Recommendation #3: weigh the amount of extra efforts by co-designing a system with target medical users, as different physicians have different notions of time urgency. Emergency room doctors often deal with urgent cases by making decisions in a matter of seconds, and internists often perform examinations in 15-20 minutes per patient; oncologists or implant specialists might decide on a case via multiple meetings that span days. There is a sense of timeliness in all these scenarios, but the amount of time that can be budgeted differs from case to case. To address such differences, we further recommend modeling each interactive task in a medical AI system (*i.e.*, how long it might take for the user to perform each task) and providing a mechanism that allows physicians to 'filter out' interactive components that might take too much time (*e.g.*, the attention map in Impetus). Importantly, different levels of urgency should be modifiable (perhaps as a one-time setup) by physicians in different specialties.

6.2 AI Using Agile Labeling to Train and Adapt Itself On-the-Fly

Prediction map (Figure 3b) visualizes current AI's diagnosis of the WSI and was designed to help the pathologists assess the model's performance and decide where they could provide more labels.

However, pathologists used the prediction map differently than we expected. Pathologists would often zoom into recommendation boxes on the WSI, study the region for a few seconds, then switch to the prediction map for a few seconds, and switch back to WSI. They tended to use the prediction map as a tool to help them see if there is something 'interesting' in the current zoomed-in region. Sometimes pathologists used the prediction map for double-checking their developing diagnosis:

That was all negative, and I didn't get a strong heatmap signal, so it was confirmatory and somewhat helpful. (P6)

Interestingly, how pathologists used the prediction map seemed to complement the recommendation boxes: while recommendation boxes told pathologists which region is worth looking at (*i.e.*,

might contain tumors), prediction map confirmed pathologists' assumption when they thought a region was of little 'interest' (i.e., no signs of tumor).

Lesson #4 To guide the examination process with prioritization, AI should help a medical user narrow in small regions of a large task space, as well as helping them filter out information within specific regions.

Recommendation #4: use visualization to filter out information, i.e., leverage AI's results to reduce information load for the physicians. An example would be a spotlight effect that darkens parts of a WSI where AI detects little or no tumor cells. Based on our observation that pathologists used AI's results to confirm their examination of the original H&E WSI, such an overt visualization can help them filter out subsets of the WSI patches. Meanwhile, pathologists can also reveal a darkened region if they want to examine further AI's findings (e.g., when they disagree with AI, believing a darkened spot has signs of tumor).

The unexpected usage of the prediction map affected agile labeling, as we discuss below.

Agile labeling (Figure 1b) allows a pathologist to label on a recommendation box directly, or to marquee-select a region to coarsely annotate as normal or tumor. In the introduction phase, all pathologists reported having no problem understanding the idea of continuously labeling WSIs to improve the AI:

This is actually adding more work for me, but I would be willing to add labels knowing I would be improving the model (P4)

However, during the tasks, we noticed that almost all the labels were drawn only based on the recommendation boxes. Only one pathologist actively searched for other regions to draw and provide more labels. It seemed that recommendation boxes served as a prompt, and pathologists were unmotivated to label other regions if unprompted.

We believe one fundamental reason is a lack of feedback to inform pathologists how important their labels were to the model retraining. Without such feedback, it might have been unclear to pathologists whether they needed to provide labels at all, or how much labeling would be enough.

Do I need to add labels? (P6)

Should I have provided more labels? (P5)

We assume that once pathologists see how a prediction map contained inaccurate results, they would be motivated to provide more labels to improve the prediction. However, our observations show that pathologists were more likely to make a diagnosis directly by manual examination, instead of correcting AI's predictions as we expected. Falling back to manual examination seems a more cost-effective alternative to AI automation than than improving the AI iteratively.

Lesson #5 It is possible for medical users to provide labels during their workflow with acceptable extra effort. However, the system should provide explicit feedback on how the model improves as a result, as a way to motivate and guide medical users' future inputs.

Recommendation #5: when adapting the model on-the-fly, show a visualization that indicates the model's performance changes as the physician labels more data. There could be various designs of such information, from showing low-level technical details (e.g., the model's specificity vs. sensitivity), high-level visualization (e.g., charts that plot accuracy over WSIs read) and even actionable items (e.g., 'nudging' the user to label certain classes of data to balance the training set). There are two main factors to consider when evaluating a given design: (i) as we observed in our study, whether the design could inform the physician of the model's performance

improvement or degradation as they label more data, which can be measured quantitatively as the amount of performance gain divided by the amount of labeling work done; (ii) as we noted in Lesson #2, whether consuming the extra information incurs too much effort and slows down the agile labeling process, and whether there is actionability given the extra information about model performance changes.

6.3 AI Taking Initiatives Appropriately for the Level of Performance Confidence

As shown in Table 1, AI's level of initiative is mediated based on its level of confidence about the model's performance. For low-confidence cases, AI took no initiative, and all pathologists were mostly unaware of AI's presence, while they simply focused on performing the usual manual diagnosis. For high-confidence cases, as expected, pathologists quickly confirmed AI's proactive diagnosis of macro — the easiest type of tumor to detect by both pathologists and AI. However, when it comes to cases diagnosed as negative by the AI, pathologists tended to perform a manual diagnosis anyway:

On the ones that it said it's confident but didn't really tell you it's negative, I still felt like I had to look at those to confirm. I wasn't going to trust the system [to confirm] that it's negative (P2)

In pathology, in order to rule out tumors, pathologists have to thoroughly examine the entire WSI, whereas it only takes one positive case to diagnose the lymph node as positive. Thus there was a discrepancy of trust between macro vs. negative, despite that AI treats both equally as different labels of a slide image and categorizes both as high confidence.

Lesson #6 Tasks treated equally by an AI might carry different weights to a medical user. Thus for medically high-staked tasks, AI should provide information to validate its confidence level.

Recommendation #6: provide additional justification for a negative diagnosis of a high-staked disease. For example, when Impetus concludes a case as negative, the system can still display the top five regions wherein AI finds the most likely signs of tumor (albeit below a threshold of positivity). In this way, even if the result turned out to be a false negative, the physicians would be guided to examine regions where the actual tumor cells are likely to appear. Beyond such intrinsic details, it is also possible to retrieve extrinsic information, e.g., prevalence of the disease given the patient's population, or similar histological images for comparison. As suggested in [62], such extrinsic justification can complement the explanation of a model's intrinsic process, thus allowing physicians to understand AI's decision more comprehensively.

For the mid-confidence case, AI was designed to pre-fill the diagnosis dialog (but without any confirmative action) as a way to hint its prediction without signaling any conclusive decision. This design did not seem to have noticeable effects on the pathologists, which echoes Lesson #3 that information needs to present actionability in order to affect a medical user's workflow.

7 LIMITATIONS, DISCUSSION AND CONCLUSION

This paper explores how AI's capabilities with limitations can still benefit a traditional manual diagnosis process on histological data. We investigate this question through the design and study of Impetus, a tool where an AI takes multi-leveled initiatives to provide various forms of assistance to a pathologist performing tumor detection in whole-slide histological images. We conducted work sessions with eight medical professionals using Impetus and summarize our findings and lessons learned, based on which we provide design recommendations with concrete examples to inform future work on human-centered medical AI systems.

Due to the limited availability of pathologists, the number of participants ($N=8$) of this work was small, and all participants were from the same medical center. The performance of the multiple instance learning model was not evaluated because of the limited number of new labels obtained from pathologists in the work sessions, which we discussed in Lesson #5 on engaging pathologists to provide extra input.

Below, we discuss several other limitations encountered during the development of our system.

Detecting small lesion tissues in a WSI. We found it hard for the system to detect small lesion tissues in a WSI during the work session. However, from our interviews with medical professionals, it is more valuable to find those areas with AI, whereas large, macro tissues can be located quickly without assistance. Here we summarize the reasons why the machine learning algorithm fails to localize small lesion tissues. First, the system takes patches as input to extract the features and classifies them with trained MIL-RF. However, this approach can be problematic when detecting small-area lesion tissues, since the classification performance highly depends on color, and small metastasis do not change the color of patches significantly. Further, the machine learning model treats tissues in a WSI as separate patches without considering structural correlations among the tissues. Specifically, lymph node invasion starts from the perimeters of the node. Thus small lesion tissues are more likely to appear in those peripheral regions.

Integration into pathologists' workflow. Data encountered in a pathologist's day-to-day workflow are imbalanced by nature: metastasis areas often occupy small fractions of the entire WSI. As a result, there would be more negative annotations than positive ones. Furthermore, with the coarse labeling enabled by our MIL algorithm, only a subset of the patches in a positive annotation are truly positive patches. This imbalance in training data skews the model's predictions.

On the other hand, to use an AI system for diagnosis, the AI's performance must be validated. However, a trained-on-the-fly AI can not be practically validated, *i.e.*, it is not feasible to fully validate the model's performance after every iteration of labeling and training, even though the number of new training data is small. Such a dynamic system is hard to control to maintain a certain level of performance regardless of run-time user interaction.

We did not observe any automation bias (*i.e.*, physicians biasing their decision in favor of AI's output), primarily because (i) the current AI's performance is imperfect and pathologists often did not trust the model to automate a diagnosis; (ii) we designed the AI to reframe from automation in cases of a low confidence, thus prompting pathologists to take control and rely on their own expertise.

Combining prior knowledge. In the real diagnosis environment of a pathologist, extra patient information is necessary for diagnosing a glass slide, which often includes the patient's medical history and type of cancer as determined from previous examinations. From our interviews with pathologists, we learn that such information is crucial for diagnosis speed and accuracy, as it informs the pathologist on what to look for and where to find them. To better match the AI with a pathologist's mental model and provide better guidance and explanations, we should incorporate patient information into the diagnosis model. For example, the AI might use a different CNN to look for a specific type of tumor tissues given a particular cancer type.

Explicit vs. implicit feedback. So far, Impetus primarily relies on explicit feedback from physicians via agile labeling. Future work should leverage other information as implicit feedback, *e.g.*, what kinds of WSI areas a pathologist looks at first, or spends the most time examining, where and how much they zoom in. Inferring useful information for adapting the model presents new technical challenges for the machine learning community; for HCI and CSCW researchers, the challenge is making such inference transparent by informing pathologists how AI is learning from some of their implicit behavior.

ACKNOWLEDGMENTS

We thank the reviewers for their valuable feedback. We thank all the anonymous participants for their participation in our study. This work was funded in part by a Hellman Fellowship and the National Science Foundation under grant IIS-1850183.

REFERENCES

- [1] Saleema Amershi, James Fogarty, Ashish Kapoor, and Desney Tan. 2009. Overview based example selection in end user interactive concept learning. In *Proceedings of the 22nd annual ACM symposium on User interface software and technology*. ACM, 247–256.
- [2] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 3.
- [3] Sophia K Apple. 2016. Sentinel lymph node in breast cancer: review article from a pathologist's point of view. *Journal of pathology and translational medicine* 50, 2 (2016), 83.
- [4] Teresa Araújo, Guilherme Aresta, Eduardo Castro, José Rouco, Paulo Aguiar, Catarina Eloy, António Polónia, and Aurélio Campilho. 2017. Classification of breast cancer histology images using convolutional neural networks. *PloS one* 12, 6 (2017), e0177544.
- [5] Eirini Arvaniti, Kim S Fricker, Michael Moret, Niels Rupp, Thomas Hermanns, Christian Fankhauser, Norbert Wey, Peter J Wild, Jan H Rueschoff, and Manfred Claassen. 2018. Automated Gleason grading of prostate cancer tissue microarrays via deep learning. *Scientific reports* 8, 1 (2018), 1–11.
- [6] Boris Babenko. 2008. Multiple instance learning: algorithms and applications. *View Article PubMed/NCBI Google Scholar* (2008), 1–19.
- [7] Dalal Bardou, Kun Zhang, and Sayed Mohammad Ahmad. 2018. Classification of breast cancer based on histology images using convolutional neural networks. *IEEE Access* 6 (2018), 24680–24693.
- [8] Babak Ehteshami Bejnordi, Mitko Veta, Paul Johannes Van Diest, Bram Van Ginneken, Nico Karssemeijer, Geert Litjens, Jeroen AWM Van Der Laak, Meyke Hermesen, Quirine F Manson, Maschenka Balkenhol, et al. 2017. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *Jama* 318, 22 (2017), 2199–2210.
- [9] Marsden S. Blois. 1980. Clinical Judgment and Computers. *New England Journal of Medicine* 303, 4 (1980), 192–197. <https://doi.org/10.1056/NEJM198007243030405> arXiv:<https://doi.org/10.1056/NEJM198007243030405> PMID: 7383090.
- [10] Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.
- [11] Carrie J. Cai, Emily Reif, Narayan Hegde, Jason Hipp, Been Kim, Daniel Smilkov, Martin Wattenberg, Fernanda Viegas, Greg S. Corrado, Martin C. Stumpe, and Michael Terry. 2019. Human-Centered Tools for Coping with Imperfect Algorithms During Medical Decision-Making. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, Glasgow, Scotland Uk, 1–14. <https://doi.org/10.1145/3290605.3300234>
- [12] Carrie J Cai, Samantha Winter, David Steiner, Lauren Wilcox, and Michael Terry. 2019. "Hello AI": Uncovering the Onboarding Needs of Medical Practitioners for Human-AI Collaborative Decision-Making. *Proceedings of the ACM on Human-computer Interaction* 3, CSCW (2019), 1–24.
- [13] Gabriele Campanella, Matthew G Hanna, Luke Geneslaw, Allen Miraflor, Vitor Werneck Krauss Silva, Klaus J Busam, Edi Brogi, Victor E Reuter, David S Klimstra, and Thomas J Fuchs. 2019. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature medicine* 25, 8 (2019), 1301–1309.
- [14] R. Cao, A. M. Bajgirani, S. A. Mirak, S. Shakeri, X. Zhong, D. Enzmann, S. Raman, and K. Sung. 2019. Joint Prostate Cancer Detection and Gleason Score Prediction in mp-MRI via FocalNet. *IEEE Transactions on Medical Imaging* (2019), 1–1. <https://doi.org/10.1109/TMI.2019.2901928>
- [15] Anne E Carpenter, Thouis R Jones, Michael R Lamprecht, Colin Clarke, In Han Kang, Ola Friman, David A Guertin, Joo Han Chang, Robert A Lindquist, Jason Moffat, et al. 2006. CellProfiler: image analysis software for identifying and quantifying cell phenotypes. *Genome biology* 7, 10 (2006), R100.
- [16] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noémie Elhadad. 2015. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. <https://doi.org/10.1145/2783258.2788613>
- [17] Duen Horng Chau, Aniket Kittur, Jason I Hong, and Christos Faloutsos. 2011. Apollo: making sense of large network data by combining rich user interaction and machine learning. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 167–176.
- [18] Xiang 'Anthony' Chen, Ye Tao, Guanyun Wang, Runchang Kang, Tovi Grossman, Stelian Coros, and Scott E Hudson. 2018. Forte: User-Driven Generative Design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing*

- Systems*. ACM, 496.
- [19] Dan C Cireşan, Alessandro Giusti, Luca M Gambardella, and Jürgen Schmidhuber. 2013. Mitosis detection in breast cancer histology images with deep neural networks. In *International conference on medical image computing and computer-assisted intervention*. Springer, 411–418.
 - [20] Tony J Collins. 2007. ImageJ for microscopy. *Biotechniques* 43, S1 (2007), S25–S30.
 - [21] Trafton Drew, Karla Evans, Melissa L. H. Võ, Francine L. Jacobson, and Jeremy M. Wolfe. 2013. Informatics in Radiology: What Can You See in a Single Glance and How Might This Guide Visual Search in Medical Images? *RadioGraphics* 33, 1 (jan 2013), 263–274. <https://doi.org/10.1148/rg.331125023>
 - [22] Mehmet Günhan Ertoşun and Daniel L Rubin. 2015. Automated grading of gliomas using deep learning in digital pathology images: A modular approach with ensemble of convolutional neural networks. In *AMIA Annual Symposium Proceedings*, Vol. 2015. American Medical Informatics Association, 1899.
 - [23] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise.. In *Kdd*, Vol. 96. 226–231.
 - [24] Cristian Felix, Aritra Dasgupta, and Enrico Bertini. 2018. The exploratory labeling assistant: Mixed-initiative label curation with large document collections. In *The 31st Annual ACM Symposium on User Interface Software and Technology*. ACM, 153–164.
 - [25] Qingji Guan, Yaping Huang, Zhun Zhong, Zhedong Zheng, Liang Zheng, and Yi Yang. 2018. Diagnose like a radiologist: Attention guided convolutional neural network for thorax disease classification. *arXiv preprint arXiv:1801.09927* (2018).
 - [26] David A Gutman, Mohammed Khalilia, Sanghoon Lee, Michael Nalisnik, Zach Mullen, Jonathan Beezley, Deepak R Chittajallu, David Manthey, and Lee AD Cooper. 2017. The digital slide archive: A software platform for management, integration, and analysis of histology for cancer research. *Cancer research* 77, 21 (2017), e75–e78.
 - [27] Narayan Hegde, Jason D Hipp, Yun Liu, Michael Emmert-Buck, Emily Reif, Daniel Smilkov, Michael Terry, Carrie J Cai, Mahul B Amin, Craig H Mermel, et al. 2019. Similar image search for histopathology: SMILY. *NPJ digital medicine* 2, 1 (2019), 1–9.
 - [28] Andreas Holzinger, Bernd Malle, Peter Kieseberg, Peter M Roth, Heimo Müller, Robert Reihs, and Kurt Zatloukal. 2017. Towards the augmented pathologist: Challenges of explainable-ai in digital pathology. *arXiv preprint arXiv:1712.06657* (2017).
 - [29] Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *Conference on Human Factors in Computing Systems - Proceedings*. <https://doi.org/10.1145/302979.303030>
 - [30] Maximilian Ilse, Jakub M Tomczak, and Max Welling. 2018. Attention-based deep multiple instance learning. *arXiv preprint arXiv:1802.04712* (2018).
 - [31] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167* (2015).
 - [32] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Illcus, Chris Chute, Henrik Marklund, Behzad Haghighi, Robyn Ball, Katie Shpanskaya, Jayne Seekins, David A Mong, Safwan S Halabi, Jesse K Sandberg, Ricky Jones, David B Larson, Curtis P Langlotz, Bhavik N Patel, Matthew P Lungren, and Andrew Y Ng. 2019. CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison. (2019). [arXiv:1901.07031](http://arxiv.org/abs/1901.07031) [www.aaii.org/http://arxiv.org/abs/1901.07031](http://arxiv.org/abs/1901.07031)
 - [33] Ellen C Jensen. 2013. Quantitative analysis of histological staining and fluorescence using ImageJ. *The Anatomical Record* 296, 3 (2013), 378–381.
 - [34] Fei Jiang, Yong Jiang, Hui Zhi, Yi Dong, Hao Li, Sufeng Ma, Yilong Wang, Qiang Dong, Haipeng Shen, and Yongjun Wang. 2017. Artificial intelligence in healthcare: Past, present and future. , 230–243 pages. <https://doi.org/10.1136/svn-2017-000101>
 - [35] Martin Köbel, Steve E Kalloger, Patricia M Baker, Carol A Ewanowich, Jocelyne Arseneau, Viktor Zhrebickiy, Soran Abdulkarim, Samuel Leung, Máire A Duggan, Dan Fontaine, et al. 2010. Diagnosis of ovarian carcinoma cell type is highly reproducible: a transcanadian study. *The American journal of surgical pathology* 34, 7 (2010), 984–993.
 - [36] Christian Leistner, Amir Saffari, and Horst Bischof. 2010. MIForests: Multiple-instance learning with randomized trees. In *European Conference on Computer Vision*. Springer, 29–42.
 - [37] Joseph Carl Robnett Licklider. 1960. Man-computer symbiosis. *IRE transactions on human factors in electronics* 1 (1960), 4–11.
 - [38] Geert Litjens, Peter Bandi, Babak Ehteshami Bejnordi, Oscar Geessink, Maschenka Balkenhol, Peter Bult, Altuna Halilovic, Mayke Hermesen, Rob van de Loo, Rob Vogels, et al. 2018. 1399 H&E-stained sentinel lymph node sections of breast cancer patients: the CAMELYON dataset. *GigaScience* 7, 6 (2018), giy065.
 - [39] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. 2008. Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining*. IEEE, 413–422.
 - [40] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, Nov (2008), 2579–2605.

- [41] Raphaël Marée, Loïc Rollus, Benjamin Stévens, Renaud Hoyoux, Gilles Louppe, Rémy Vandaele, Jean-Michel Begon, Philipp Kainz, Pierre Geurts, and Louis Wehenkel. 2016. Collaborative analysis of multi-gigapixel imaging data using Cytomine. *Bioinformatics* 32, 9 (2016), 1395–1401.
- [42] Anne L Martel, Dan Hosseinzadeh, Caglar Senaras, Yu Zhou, Azadeh Yazdanpanah, Rushin Shojaii, Emily S Patterson, Anant Madabhushi, and Metin N Gurcan. 2017. An image analysis resource for cancer research: PIIP—pathology image informatics platform for visualization, analysis, and management. *Cancer research* 77, 21 (2017), e83–e86.
- [43] R. A. Miller and F. E. Masarie. 1990. The demise of the 'Greek Oracle' model for medical diagnostic systems.
- [44] Michael Nalisnik, Mohamed Amgad, Sanghoon Lee, Sameer H Halani, Jose Enrique Velazquez Vega, Daniel J Brat, David A Gutman, and Lee AD Cooper. 2017. Interactive phenotyping of large-scale histology imaging data with HistomicsML. *Scientific reports* 7, 1 (2017), 14588.
- [45] An T Nguyen, Aditya Kharosekar, Saumyaa Krishnan, Siddhesh Krishnan, Elizabeth Tate, Byron C Wallace, and Matthew Lease. 2018. Believe it or not: Designing a human-AI partnership for mixed-initiative fact-checking. In *The 31st Annual ACM Symposium on User Interface Software and Technology*. ACM, 189–199.
- [46] Nobuyuki Otsu. 1979. A threshold selection method from gray-level histograms. *IEEE transactions on systems, man, and cybernetics* 9, 1 (1979), 62–66.
- [47] Ludovic Roux, Daniel Racocanu, Nicolas Loménie, Maria Kulikova, Humayun Irshad, Jacques Klossa, Frédérique Capron, Catherine Genestie, Gilles Le Naour, and Metin N Gurcan. 2013. Mitosis detection in breast cancer histological images An ICPR 2012 contest. *Journal of pathology informatics* 4 (2013).
- [48] Curtis T Rueden, Johannes Schindelin, Mark C Hiner, Barry E DeZonia, Alison E Walter, Ellen T Arena, and Kevin W Eliceiri. 2017. ImageJ2: ImageJ for the next generation of scientific image data. *BMC bioinformatics* 18, 1 (2017), 529.
- [49] Joel Saltz, Ashish Sharma, Ganesh Iyer, Erich Bremer, Feiqiao Wang, Alina Jasniewski, Tammy DiPrima, Jonas S Almeida, Yi Gao, Tianhao Zhao, et al. 2017. A containerized software system for generation, management, and exploration of features from whole slide tissue images. *Cancer research* 77, 21 (2017), e79–e82.
- [50] Mike Schaeckermann, Carrie J Cai, Abigail E Huang, and Rory Sayres. 2020. Expert Discussions Improve Comprehension of Difficult Cases in Medical Image Assessment. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [51] Johannes Schindelin, Ignacio Arganda-Carreras, Erwin Frise, Verena Kaynig, Mark Longair, Tobias Pietzsch, Stephan Preibisch, Curtis Rueden, Stephan Saalfeld, Benjamin Schmid, et al. 2012. Fiji: an open-source platform for biological-image analysis. *Nature methods* 9, 7 (2012), 676.
- [52] Caroline A Schneider, Wayne S Rasband, and Kevin W Eliceiri. 2012. NIH Image to ImageJ: 25 years of image analysis. *Nature methods* 9, 7 (2012), 671.
- [53] Burr Settles. 2009. *Active learning literature survey*. Technical Report. University of Wisconsin-Madison Department of Computer Sciences.
- [54] Edward H. Shortliffe. 1993. The adolescence of AI in Medicine: Will the field come of age in the '90s? *Artificial Intelligence In Medicine* (1993). [https://doi.org/10.1016/0933-3657\(93\)90011-Q](https://doi.org/10.1016/0933-3657(93)90011-Q)
- [55] Patrice Y Simard, Saleema Amershi, David M Chickering, Alicia Edelman Pelton, Soroush Ghorashi, Christopher Meek, Gonzalo Ramos, Jina Suh, Johan Verwey, Mo Wang, et al. 2017. Machine teaching: A new paradigm for building machine learning systems. *arXiv preprint arXiv:1707.06742* (2017).
- [56] Christoph Sommer, Christoph Straehle, Ullrich Koethe, and Fred A Hamprecht. 2011. Ilastik: Interactive learning and segmentation toolkit. In *2011 IEEE international symposium on biomedical imaging: From nano to macro*. IEEE, 230–233.
- [57] Jina Suh, Xiaojin Zhu, and Saleema Amershi. 2016. The label complexity of mixed-initiative classifier training. In *International Conference on Machine Learning*. 2800–2809.
- [58] Mitko Veta, Yujing J Heng, Nikolas Stathonikos, Babak Ehteshami Bejnordi, Francisco Beca, Thomas Wollmann, Karl Rohr, Manan A Shah, Dayong Wang, Mikael Rousson, et al. 2019. Predicting breast tumor proliferation from whole-slide images: the TUPAC16 challenge. *Medical image analysis* 54 (2019), 111–121.
- [59] Mitko Veta, Paul J Van Diest, Stefan M Willems, Haibo Wang, Anant Madabhushi, Angel Cruz-Roa, Fabio Gonzalez, Anders BL Larsen, Jacob S Vestergaard, Anders B Dahl, et al. 2015. Assessment of algorithms for mitosis detection in breast cancer histopathology images. *Medical image analysis* 20, 1 (2015), 237–248.
- [60] Danding Wang, Qian Yang, Ashraf Abdul, and Brian Y Lim. 2019. Designing Theory-Driven User-Centric Explainable AI. (2019). <https://doi.org/10.1145/3290605.3300831>
- [61] John N Weinstein, Eric A Collisson, Gordon B Mills, Kenna R Mills Shaw, Brad A Ozenberger, Kyle Ellrott, Ilya Shmulevich, Chris Sander, Joshua M Stuart, Cancer Genome Atlas Research Network, et al. 2013. The cancer genome atlas pan-cancer analysis project. *Nature genetics* 45, 10 (2013), 1113.
- [62] Yao Xie, Melody Chen, David Kao, Ge Gao, and Xiang 'Anthony' Chen. 2020. CheXplain: Enabling Physicians to Explore and Understand Data-Driven, AI-Enabled Medical Imaging Analysis. In *To appear at the 2020 CHI Conference on Human Factors in Computing Systems*.

- [63] Yao Xie, Xiang ‘Anthony’ Chen, and Ge Gao. 2019. Outlining the Design Space of Explainable Intelligent Systems for Medical Diagnosis. In *ACM IUI*. <http://ceur-ws.org/Vol-2327/IUI19WS-ExSS2019-18.pdf>
- [64] Yan Xu, Jun-Yan Zhu, Eric Chang, and Zhuowen Tu. 2012. Multiple clustered instance learning for histopathology cancer image classification, segmentation and clustering. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 964–971.
- [65] Qian Yang, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. In *Proceedings of the 2020 chi conference on human factors in computing systems*. 1–13.
- [66] Qian Yang, Aaron Steinfeld, and John Zimmerman. 2019. Unremarkable AI: Fitting Intelligent Decision Support into Critical, Clinical Decision-Making Processes. *Conference on Human Factors in Computing Systems - Proceedings* (apr 2019). <https://doi.org/10.1145/3290605.3300468> arXiv:1904.09612
- [67] Qian Yang, John Zimmerman, Aaron Steinfeld, Lisa Carey, and James F. Antaki. 2016. Investigating the Heart Pump Implant Decision Process. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*. ACM Press, New York, New York, USA, 4477–4488. <https://doi.org/10.1145/2858036.2858373>
- [68] Yongfeng Zhang, Guokun Lai, Min Zhang, Yi Zhang, Yiqun Liu, and Shaoping Ma. 2014. Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. ACM, 83–92.
- [69] Zizhao Zhang, Yuanpu Xie, Fuyong Xing, Mason McGough, and Lin Yang. 2017. MDNet: A Semantically and Visually Interpretable Medical Image Diagnosis Network. (jul 2017). arXiv:1707.02485 <http://arxiv.org/abs/1707.02485>
- [70] Zhi-Hua Zhou. 2004. Multi-instance learning: A survey. *Department of Computer Science & Technology, Nanjing University, Tech. Rep* (2004).
- [71] Yan Zhu, Shaoting Zhang, Wei Liu, and Dimitris N Metaxas. 2014. Scalable histopathological image analysis via active learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 369–376.

A IMPLEMENTATION DETAILS OF IMPETUS SOFTWARE

We introduce the implementation of Impetus. Overall, the software was written in Python and the detailed implementation of Impetus' capabilities are shown as follows.

Attention map & recommendation boxes When the system is first loaded, Otsu method [46] is first utilized to separate foreground tissues and background. Then a pre-trained InceptionResNetv2 model [31] with PatchCamelyon dataset¹⁰ extracts patch features in WSIs. We trained this model with image augmentation preprocessing, Adadelta optimizer with initial learning rate 0.1, binary cross entropy loss, 100 iterations with early stopping on validation loss. The input dimension of the pre-trained network is $96 \times 96 \times 3$, and the extracted features have a dimension of 1536×1 . In the first iteration, the system performs isolation forest [39] outlier detection (max_samples=256) based on extracted features, and the detected outliers are highlighted in the attention map. In the following iterations, the attention map is a combination of outliers (from the initial detection) and high uncertainty patches (from specific models in each iteration). Uncertainty is calculated as $Uncertainty = 1 - |0.5 - Probability| \times 2$. The attention map is calculated as the soft-OR of outliers and uncertainty: $Attention = Uncertainty \odot Outlier$. In order to obtain the recommendation boxes, the system uses a DBSCAN clustering algorithm [23] (min_sample=10, epsilon=3) on the attention map to find clusters. In each iteration, Impetus only selects the two clusters which occupy the largest areas on the WSI as recommendation boxes to reduce users' distraction.

Algorithm 1 Impetus MIL Training

```

initialize positive_set = [ ], negative_set = [ ]
while Performance not satisfied do
    Annotate  $X_1, X_2, \dots, X_N$  boxes with  $Y_1, Y_2, \dots, Y_N$  labels
    for  $X_i$  do
        if  $Y_i == -1$  then
            negative_set.append( $X_i$ )
        else
            Embed  $X_i$  to  $\bar{X}_i$  with T-SNE
            Split  $\bar{X}_i$  into  $\bar{X}_i^{(1)}, \bar{X}_i^{(2)}$  with K-Means, map the split to original data  $X_i^{(1)}, X_i^{(2)}$ 
            Assign each instance in  $X_i^{(1)}$  with +1 labels, and  $X_i^{(2)}$  with -1 labels
            Train a random forest classifier with  $X_i^{(1)}$  and  $X_i^{(2)}$ 
            Predict negative box  $X_{neg}$ 
            Adjust the labels of  $X_i^{(1)}$  and  $X_i^{(2)}$ 
            Append positive set with positive instances
        end if
    end for
    Train a random forest classifier with positive_set and negative_set
end while

```

Agile labeling Impetus uses an agile labelling technique that does not specifically require users to provide the exact contour of tumor tissues in a WSI, as strongly-supervised learning does. Alternatively, a user can mark a positive box over an area which contains *at least one* tumor patches, or a negative box on *all negative* regions. We implemented a weakly-supervised multiple instance learning (MIL) to enable a traditional random forest algorithm learn over such agile labels. As shown in Algorithm 1, X_i is a feature set in within a box and only has one box-level label Y_i . For

¹⁰<https://patchcamelyon.grand-challenge.org/>

a negative box, all the instances in the box can be included in *negative_set*. In order to avoid the performance degradation of the following clustering algorithm, we first use manifold learning – T-SNE [40] – to represent the 1536 dimension feature X_i with two-dimension embedding $\bar{X}_i$. Here, it is assumed that the T-SNE would embed the positive box instances into one positive and one negative cluster, thus we use K-Means clustering to split $\bar{X}_i$ into two clusters: $\bar{X}_i^{(1)}$, $\bar{X}_i^{(2)}$. By using the split, the original 1536-dimension X_i can be partitioned into $X_i^{(1)}$ and $X_i^{(2)}$. After clustering, the algorithm compares the two clusters with negative samples from negative box X_{neg} to pick the real positive cluster. To achieve this goal, it first assigns each instance in $X_i^{(1)}$ with +1 labels, and that in $X_i^{(2)}$ with -1 labels, then trains a random forest classifier (100 trees, max_depth=100) with $X_i^{(1)}$ and $X_i^{(2)}$. The trained classifier is used to predict negative box instances X_{neg} previously provided by the user. Finally, the cluster which has the *opposite* prediction to the negative box is the positive cluster. After the positive cluster is recognized, all the instances in the positive cluster are included in the *positive_set*, and the instances in the rest cluster are aborted. Finally, a random forest classifier (MIL-RF) is trained with the obtained *positive_set* and *negative_set*. Here, the random forest algorithm is selected since its “resistance to overfit” [44]. What’s more, the notion of random forest has been applied to active learning [44], or multiple instance learning algorithms [36]. In the following iterations, the user can iteratively provide more annotations until the trained classifier reaches a satisfactory level of performance.

Confidence calculation Impetus has a rule-based confidence-level classifier to classify slides into three categories, namely high-confidence, mid-confidence and low-confidence. The system first obtains the predictions of all the patches in the WSI. A patch has two characteristics: *is_positive* and *is_uncertain*. A patch is positive if the MIL-RF classifier output > 0.5, and is uncertain if MIL-RF classifier output $\in [0.25, 0.75]$. We empirically summarize the confidence-level decision rules as follows:

- if there are more than 200 positive patches AND the number of positive patches is greater than twice the number of uncertain patches, then the slide is predicted as high-confidence;
- else if there are no outlier clusters, then the slide is predicted as low-confidence;
- else if the number of uncertain patches is greater than 300, then the slide is predicted as low-confidence;
- else if the number of positive patches is greater than 200, then the slide is predicted as high-confidence;
- for all other cases, the slide is predicted as mid-confidence.

Received June 2020; revised October 2020; accepted December 2020