

# Explainable AI for computational pathology identifies model limitations and tissue biomarkers

Jakub R. Kaczmarzyk<sup>1,2,3</sup>, Joel H. Saltz<sup>1</sup>, and Peter K. Koo<sup>2,\*</sup>

<sup>1</sup>Department of Biomedical Informatics, Stony Brook University, Stony Brook, NY, USA

<sup>2</sup>Simons Center for Quantitative Biology, Cold Spring Harbor Laboratory, Cold Spring Harbor, NY, USA

<sup>3</sup>Medical Scientist Training Program, Stony Brook University, Stony Brook, NY, USA

\*koo@cshl.edu

## ABSTRACT

Deep learning models have shown promise in histopathology image analysis, but their opaque decision-making process poses challenges in high-risk medical scenarios. Here we introduce HIPPO, an explainable AI method that interrogates attention-based multiple instance learning (ABMIL) models in computational pathology by generating counterfactual examples through tissue patch modifications in whole slide images. Applying HIPPO to ABMIL models trained to detect breast cancer metastasis reveals that they may overlook small tumors and can be misled by non-tumor tissue, while attention maps—widely used for interpretation—often highlight regions that do not directly influence predictions. By interpreting ABMIL models trained on a prognostic prediction task, HIPPO identified tissue areas with stronger prognostic effects than high-attention regions, which sometimes showed counterintuitive influences on risk scores. These findings demonstrate HIPPO's capacity for comprehensive model evaluation, bias detection, and quantitative hypothesis testing. HIPPO greatly expands the capabilities of explainable AI tools to assess the trustworthy and reliable development, deployment, and regulation of weakly-supervised models in computational pathology.

## Introduction

Digital pathology has emerged as a transformative force in medicine, ushering in an era where computational methods can augment and enhance the diagnostic and prognostic capabilities of pathologists. By digitizing whole slide images (WSIs) of tissue specimens, this field has opened up new avenues for applying advanced machine learning techniques to analyze complex histological patterns and features. The potential impact of computational pathology is far-reaching, promising to improve diagnostic accuracy, standardize interpretation, and uncover novel biomarkers that may inform personalized treatment strategies<sup>1–13</sup>.

Recently, attention-based multiple instance learning (ABMIL)<sup>14</sup> has emerged as a powerful approach to analyze WSIs for various pathological tasks, demonstrating performance that often rivals or surpasses that of expert pathologists<sup>15</sup>. ABMIL models treat each WSI as a collection of smaller image patches (instances) and use attention mechanisms to identify and focus on the most relevant regions for the task at hand. Importantly, multiple instance learning allows ABMIL models to learn from specimen-level labels, not requiring exhaustive pixel-level annotations, which are time-consuming and costly to obtain<sup>15</sup>. This feature makes ABMIL models particularly well-suited for tasks such as cancer detection<sup>16,17</sup>, diagnosis<sup>18–21</sup>, identification of primary cancer origin<sup>22</sup>, grading<sup>17,23,24</sup>, genomic aberration detection<sup>25–28</sup>, molecular phenotyping<sup>29–31</sup>, treatment response prediction<sup>32–34</sup>, and prognostication<sup>33,35–37</sup>.

However, the widespread adoption of ABMIL models in clinical settings is hindered by challenges in model interpretability and trustworthiness<sup>9,10,38,39</sup>. A key limitation lies in the heavy reliance of interpretations based on ABMIL's attention, which is often used as a proxy for understanding model behavior. While attention highlights regions of interest within a WSI, they do not necessarily reflect the direct influence of these regions on model predictions<sup>40,41</sup>. This disconnect between attention and model output can lead to misinterpretations of model behavior, potentially eroding trust in the model's decisions and limiting its clinical utility<sup>42–45</sup>. In addition, post hoc model explanations via attribution methods, such as LIME<sup>46</sup> and SHAP<sup>47</sup>, make restrictive additive or linear assumptions of individual pixels, which have been argued to not reflect a model's decision making process<sup>48</sup>.

To address these challenges, we introduce HIPPO (Histopathology Interventions of Patches for Predictive Outcomes), an explainable AI method designed to enhance trust in ABMIL models and provide deeper insights into their decision-making processes. HIPPO goes beyond traditional attention-based interpretations by quantitatively assessing the impact of specific tissue regions on model predictions. By simulating targeted interventions through the occlusion or inclusion of individual or groups of patches, HIPPO enables a more nuanced understanding of how different histological features influence ABMIL

model outputs.

We demonstrate the utility of HIPPO by applying it to two major tasks in computational pathology: metastasis detection and cancer prognosis prediction. In the context of metastasis detection, we evaluate five foundation models in pathology using the CAMELYON16 dataset<sup>16</sup>. Our analysis uncovers model-specific limitations and biases that would have remained hidden using attention mechanism alone. We reveal that some models rely heavily on extratumoral tissue for metastasis detection, while others are surprisingly insensitive to small tumor regions. These findings highlight the importance of rigorous model evaluation beyond standard performance metrics and underscore the potential of HIPPO in identifying when and why models might fail.

For cancer prognosis, we apply HIPPO to ABMIL models trained on breast cancer and cutaneous melanoma datasets from The Cancer Genome Atlas. Our results demonstrate that HIPPO can identify tissue regions more strongly associated with prognosis compared to those highlighted by attention. Strikingly, we find that high-attention regions can sometimes have counterintuitive effects on prognostic predictions, further emphasizing the limitations of relying solely on attention for model interpretation. By quantitatively assessing the impact of tumor-infiltrating lymphocytes (TILs) on model predictions, we confirm that the models have captured the known prognostic significance of TILs in both breast cancer and melanoma. This ability to link model behavior to established biological knowledge is crucial for building trust in AI-driven prognostic tools.

Beyond model interpretation, HIPPO opens up new possibilities for *in silico* experimentation in computational pathology. We showcase this potential by using HIPPO to simulate the effects of autologous TIL therapy in melanoma patients. By digitally replicating TIL-positive regions in specimens with poor prognosis, we demonstrate a proof-of-principle approach for identifying patients who potentially might benefit from this immunotherapy. This application illustrates how HIPPO can bridge the gap between computational predictions and clinically actionable insights.

As computational pathology continues to advance, the need for robust, interpretable, and trustworthy AI models becomes increasingly critical. HIPPO represents a significant step forward in this direction, offering a powerful tool for uncovering the strengths, limitations, and potential biases of ABMIL models in pathology. By providing a more comprehensive understanding of model behavior, HIPPO not only enhances the interpretability of existing models but also paves the way for developing more reliable and clinically relevant AI tools in pathology. As we demonstrate across multiple applications, from metastasis detection to prognostic modeling, HIPPO has the potential to accelerate the translation of computational pathology into clinical practice, ultimately improving patient care and outcomes.

## Results

### HIPPO: Histopathology Interventions of Patches for Predictive Outcomes

HIPPO is a specimen-level perturbation toolkit that explains weakly-supervised models in computational pathology (Fig. 1a). The fundamental goal of HIPPO is to explore counterfactual (i.e., “what if”) scenarios that are infeasible to realize in actual tissue samples. For instance, it would be impractical to directly manipulate the tumor microenvironment of a tissue specimen to understand its effect on a prognostic model. Instead, we can digitally modify a WSI that simulates this intervention. HIPPO enables *in silico* interventions through the occlusion or inclusion of single or multiple patches, utilizing the resulting ABMIL model predictions as counterfactual outcomes. HIPPO provides quantitative insights into how specific tissue alterations impact pathological assessments through the lens of the AI model. These assessments can include but are not limited to, patient prognosis, treatment response prediction, metastasis detection, inference of spatial transcriptomics, gene mutation detection, and microsatellite instability identification. Applying HIPPO to ABMIL models enables researchers, regulators, and clinicians to elucidate model behavior and assess the reliability of model outputs in high-risk clinical contexts.

Traditional approaches to digital interventions in medical imaging often require precise segmentation of objects for occlusion or inclusion<sup>49,50</sup>, as well as sophisticated inpainting techniques to maintain image integrity<sup>51–54</sup>. Alternatively, generative AI can generate counterfactual images<sup>55,56</sup>, but the quality of the generated images has not been thoroughly evaluated for histopathology. These manual or AI-assisted methods can introduce covariate shifts when imperfectly executed<sup>57</sup>, potentially leading to unreliable model predictions. The key insight for HIPPO is based on how data flows through ABMIL models. A WSI is treated as a bag of permutation-invariant patches, where the number and order of patches are allowed to vary<sup>14</sup>. Thus, an intervention can be achieved through two primary perturbation mechanisms: (1) removing specific patches, effectively excising tissue from the input specimen, or (2) including specific patches, simulating the addition of new tissue into the specimen. HIPPO leverages unique properties of multiple instance learning models to facilitate the generation of counterfactual images bypassing the complexities of direct image manipulation by creating hypothetical scenarios such as the introduction or removal of tumor patches or regions of tumor-infiltrating lymphocytes (TILs) from a patient’s specimen. Understanding when ABMIL models alter their predictions due to interventions provides quantitative insights into their decision making process, revealing important features and potential biases learned.

HIPPO offers hypothesis-driven and data-driven methods for intervention selection:

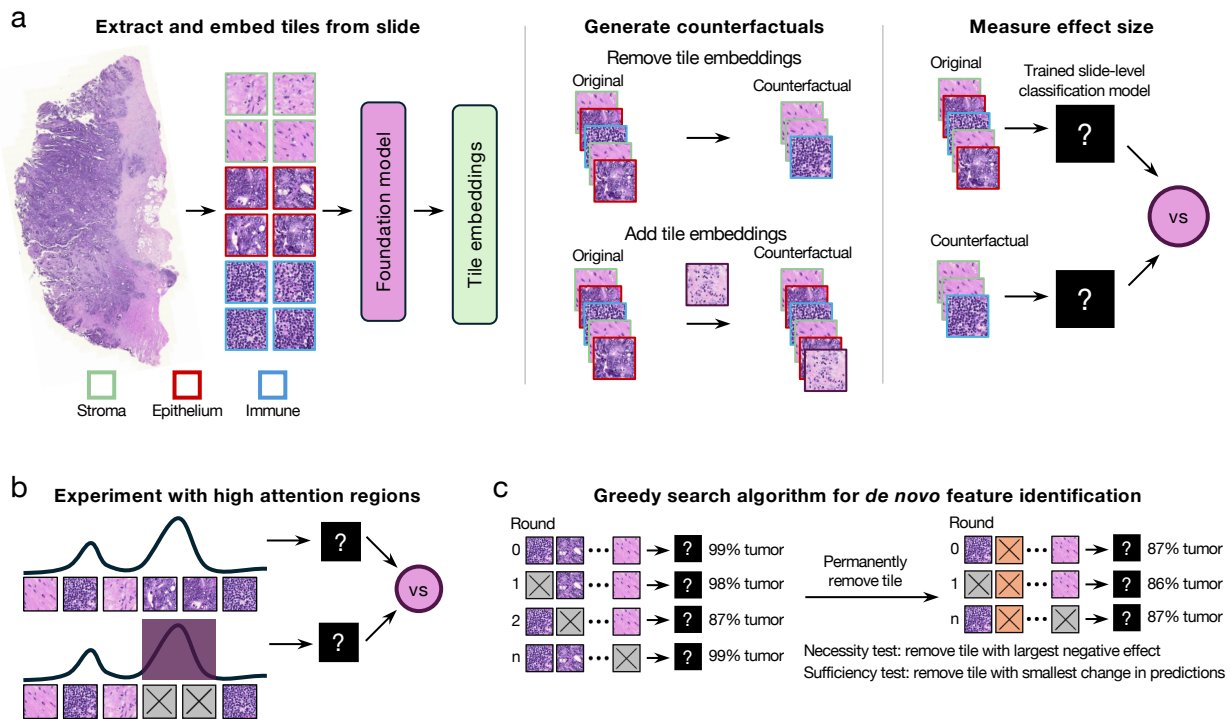
- *HIPPO-knowledge*: choosing a region based on prior knowledge or a well-defined hypothesis and quantifying its effect

by removing it from specimens or adding it to specimens without that region and measuring the change in model outputs (Fig. 1a).

- *HIPPO-attention*: quantifying the effect of high attention regions by removing the high attention regions and measuring the change in model outputs (Fig. 1b).
- *HIPPO-search-high-effect*: a greedy search algorithm to identify the regions that maximally drive a prediction. This can be used to identify regions necessary for a model’s output (Fig. 1c).
- *HIPPO-search-low-effect*: a greedy search algorithm to identify the regions that lead to the slightest change in model predictions. This can be used to identify the regions that do not change model predictions and, therefore, are not necessary for the output of a model. After removing unnecessary regions, the remaining regions can be considered sufficient for the model prediction (Fig. 1c).

HIPPO generates counterfactual outcomes through the lens of an AI model. Thus, if the AI model has limitations in fully capturing complex biological input-output relationships, HIPPO’s explanations will reflect these limitations. HIPPO provides unprecedented access to interrogating the biological factors underlying prediction, providing insight into an AI’s decision-making process.

With the advent of digital pathology foundation models, it is important to evaluate model robustness, generalizability, and potential biases and understand their limitations. Here, we showcase the breakthroughs made possible by HIPPO in rigorously evaluating models built on top of foundation models for breast metastasis detection and prognosis prediction tasks. We compare five foundation models in metastasis detection and identify model-specific limitations and biases. We also use HIPPO to study the effects of tissue components on prognostic models, demonstrating how HIPPO’s capabilities surpass attention in identifying



**Figure 1. HIPPO explainable AI toolkit.** HIPPO enables quantitative assessment of how specific tissue regions impact model predictions, enhancing interpretation and validation of AI models. **a**, Schematic of attention-based multiple instance learning. Whole slide images are divided into patches and embedded using a pretrained foundation model. ABMIL learns specimen-level labels from these bags of patches, assigning attention weights to each patch. Leveraging ABMIL’s invariance to patch order and count, we can create counterfactual specimens by adding or removing tissue regions within patches. Model outputs are then compared between original and counterfactual specimens to measure effects. **b**, HIPPO quantifies the effect of high-attention regions by removing them and measuring the resulting change in model outputs. **c**, HIPPO implements greedy search algorithms to identify necessary or sufficient tissue regions *de novo*.

low and high-risk drivers. We also measure the effect of TILs on breast cancer and melanoma patient prognosis and demonstrate digitally that autologous TILs improve predicted prognosis in a subset of melanoma patients, marking exciting progress in the field.

### **Do MIL models think tumor is necessary for breast cancer metastasis detection?**

Metastasis detection is a well-studied task, with well-defined features (i.e., tumor cells) that drive the label of whether or not a specimen contains metastasis. In a clinical setting, it is critical that metastases are identified; a false negative is unacceptable. Recent studies have shown that ABMIL models have strong performance in metastasis detection<sup>58</sup>. However, previous studies have also found that computer vision models can make the correct predictions for the wrong reasons, such as short-cut features or spurious correlations<sup>59,60</sup>. Thus, the degree to which AI models rely on the tumor regions remains to be seen, even for a relatively straightforward task like tumor detection. Understanding this is critical to elucidate the strengths and limitations of ABMIL models for metastasis detection, including potential biases.

To evaluate this, we trained several ABMIL models for breast metastasis detection using the CAMELYON16 dataset<sup>16</sup> (Fig. 2a). Several pathology foundation models have recently emerged, demonstrating near-human levels in metastasis detection. Here we consider five pathology foundation models (UNI<sup>58</sup>, REMEDIS<sup>61</sup>, Phikon<sup>62</sup>, CTransPath<sup>63</sup>, and RetCCL<sup>64</sup>). We trained five ABMIL models for each foundation model to distinguish whether or not a specimen contained metastasis. Similar to previously reported results<sup>58</sup>, UNI achieved a mean balanced accuracy of 0.982, REMEDIS 0.922, Phikon 0.907, CTransPath 0.858, and RetCCL 0.745. (Fig. 2b, Supplementary Table 1). For HIPPO explainability experiments, we used the best-performing model (out of 5 random initializations) on the test set for each foundation model. The best UNI model achieved balanced accuracy of 1.00, REMEDIS 0.949, Phikon 0.955, CTransPath 0.885, and RetCCL 0.769 (Supplementary Table 2).


In this dataset, expert pathologists finely annotated metastatic regions. This allows us to use HIPPO-knowledge to determine whether metastatic regions are necessary for detecting breast cancer metastasis. Specifically, for patients who were positive for metastasis, we removed the patches that intersected with the tumor annotations, effectively creating a version of the specimen that does not contain metastasis. We compared model predictions before and after the intervention. Specificity was calculated as the ratio of true negatives to all negative samples. In this set of counterfactuals, all specimens were negative, so the specificity represented the proportion of correct negative predictions by the models. Notably, the UNI-based model exhibited the lowest specificity (0.73) in these counterfactual examples despite achieving the highest balanced accuracy on the original test set (1.00). This discrepancy was particularly pronounced in counterfactual specimens that originally contained macrometastases (specificity 0.59), suggesting that the UNI-based ABMIL model uses tissue outside of the tumor region to drive positive metastasis predictions. The REMEDIS-based model exhibited a similar trend, with a specificity of 0.77 in counterfactuals derived from macrometastases. In contrast, the other models showed less dependence on extratumoral tissue (sensitivity of Phikon-based, 0.86; CTransPath-based, 0.92; RetCCL-based, 0.88), indicating that their predictions are primarily driven by tumor epithelial cells rather than other tissue components (Fig. 2c).

We hypothesized that the low specificity of the UNI-based ABMIL model may be attributed to metastasis-induced alterations in the surrounding tumor microenvironment. To investigate this, we used HIPPO-knowledge to remove increasingly larger regions surrounding the tumor annotation progressively and quantified the effect on metastasis detection. As the extent of peritumoral tissue removal increased, the UNI-based model was consistently more likely to predict the absence of metastasis. Specificity increased from 0.73 at dilation of 0  $\mu\text{m}$ , to 0.78 at 64  $\mu\text{m}$ , to 0.80 at 128  $\mu\text{m}$ , to 0.86 at 256  $\mu\text{m}$ , and to 0.88 at 1024  $\mu\text{m}$ . This was driven primarily by macrometastatic specimens, where specificity increased from 0.59 at dilation of 0  $\mu\text{m}$  to 0.68 at dilation of 64  $\mu\text{m}$ , to 0.73 at dilation of 128  $\mu\text{m}$ , to 0.82 at dilation of 256  $\mu\text{m}$ , to 0.86 at dilation of 1024  $\mu\text{m}$ . Notably, other ABMIL models remained largely unaffected by peritumoral tissue removal, highlighting a unique characteristic of the UNI-based model (Supplementary Fig. 1). In summary, HIPPO enabled the quantitative exploration of peritumoral tissue on metastasis detection.

### **Is tumor sufficient for breast cancer metastasis detection?**

While necessity assesses the importance of a feature or feature set, it does not inform whether the feature set is sufficient for model predictions. Metastasis detection models must be able to detect tumor regions no matter how small. Using HIPPO-knowledge, we tested the sufficiency of metastatic regions using two methods: removing all non-tumor patches and measuring model outputs and adding tumor regions to normal specimens and measuring model outputs.

First, we constructed counterfactual specimens (n=49) by removing all non-tumor tissue (i.e., removing patches that did not intersect with expert tumor annotations) and measuring model outputs. With only the tumor present, the true label for these images was “positive”, and the foundation models had the following sensitivity (true positive rate): UNI-based 0.98, REMEDIS-based 0.92, Phikon-based 0.98, CTransPath-based 0.96, RetCCL-based 0.82 (Fig. 2d). There is evidence to suggest that extratumoral tissue caused false negative predictions. Four of the five foundation models improved sensitivity when using only tumor tissue in micrometastases compared to the original positive samples, suggesting that extratumoral tissue drove false negative predictions. The sensitivity of CTransPath increased by 25%, Phikon by 4%, REMEDIS by 5%, and RetCCL by 100%.



**Figure 2. Understanding the role of tumor in detecting metastases.** **a**, Example WSI from the CAMELYON16 dataset containing a macrometastasis (specimen *test\_001*), with a  $128 \times 128 \mu\text{m}$  patch highlighted. **b**, Bar plot of balanced accuracy, sensitivity, and specificity on the CAMELYON16 test set ( $n=129$ , 80 negative, 49 positive) across five random initializations and five encoders, with mean values and 95% confidence intervals. The best-performing model for each encoder was used in subsequent experiments. **c-d**, Bar plots showing specificity when tumor-containing patches are removed (c) and sensitivity when only tumor tissue remains (d) in positive specimens ( $n=49$ , 22 macrometastases, 29 micrometastases), quantifying necessity and sufficiency of tumor regions for metastasis detection. **e**, Bar plot of sensitivity after adding metastases to negative specimens (3920 counterfactuals: 80 negative  $\times$  49 positive), further quantifying tumor sufficiency. **f**, Bar plot showing sensitivity of counterfactuals with a single  $128 \times 128 \mu\text{m}$  tumor patch in normal ( $n=80$ ) and metastatic ( $n=49$ ) specimens. **g**, Strip plot of model probabilities for tumor patches in specimen *test\_051* using the UNI-based ABMIL model, comparing original, tumor-removed, and single-tumor-patch ( $n=125$ ) conditions. **h**, Line plot relating tumor size to model sensitivity, with each point representing 3920 counterfactuals (80 negative  $\times$  49 positive) as tumor patches are added to negative specimens.

For UNI, however, using original WSIs resulted in a sensitivity of 1.0 on micrometastasis. However, when using only the tumor tissue, one false negative prediction suggested that the UNI-based model may use tissue outside of the metastatic region in its predictions. Critically, this demonstrated that the tumor was insufficient for a positive prediction in this specimen with the UNI-based model and that extratumoral tissue was solely driving the positive prediction. RetCCL had a true positive rate in macrometastases of 0.95 (21 predicted positive of 22 positive specimens). When using only tumor tissue, all macrometastases were detected successfully, demonstrating that tissue outside the metastatic region caused a false negative prediction.

We also evaluated whether tumor was sufficient for metastasis detection by embedding tumor regions in normal specimens. We embedded all patches intersecting with tumor annotations into normal specimens, resulting in 3,920 positive counterfactual examples (80 normal slides  $\times$  49 positive slides). Model outputs for these examples were recorded. The UNI-based model had



a sensitivity of 0.98, REMEDIS-based 0.86, Phikon-based 0.95, CTransPath-based 0.90, and RetCCL-based 0.63. Positive counterfactuals made with micrometastases were less likely to be detected by most models (UNI-based achieved sensitivity of 0.96, REMEDIS-based 0.75, Phikon-based 0.91, CTransPath-based 0.93, and RetCCL-based 0.40), suggesting that smaller tumors in the context of normal tissue are insufficient for positive metastasis detection (Fig. 2e).

The average treatment effect for each metastatic slide was calculated by averaging the model's probability of metastasis across all negative samples. This informs which positive slides can drive positive predictions across individuals. 100% of macrometastases (n=22) led to true positives in UNI-based, REMEDIS-based, Phikon-based, and CTransPath-based models. In the RetCCL-based model, 90% (n=20) of macrometastases had an average true positive effect. Micrometastases (n=27) were less likely to induce positive predictions on average, with 96% (n=26) positive in UNI, 93% (n=25) in Phikon, 81% (n=22) in CTransPath, 74% (n=20) in REMEDIS, and 37% (n=10) in RetCCL.

### Foundation models may miss small breast cancer metastases

To evaluate the sensitivity of ABMIL models to detect metastasis based on the size of the metastasis in a specimen, we analyzed the metastasis-positive specimens from the CAMELYON16 test set. Our methodology involved initially removing all tile embeddings that intersected with expert tumor annotations, effectively rendering the slide negative for metastases. A  $128 \times 128 \mu\text{m}$  region of tumor (shown in the right-hand side of Fig. 2a) was added to 80 normal specimens and 49 metastasis-removed positive specimens. When the single-patch tumor region was embedded in normal specimens, the REMEDIS-, Phikon-, and RetCCL-based ABMIL models detected 100% of counterfactuals as positive, highlighting their robustness to this small region of tumor. The UNI-based model, on the other hand, failed to detect 41% (n=33) of positive counterfactuals (n=80), and the CTransPath-based models failed to detect 35% (n=28) of positive counterfactuals. A similar trend was observed when the tumor region was embedded into the context of metastatic specimens (i.e., the positive specimen with metastasis removed). The REMEDIS-, Phikon-, and RetCCL-based models detected 100% of positive counterfactuals (n=49), whereas the UNI-based model missed 51% (n=25) and CTransPath-based missed 65% (n=32) of positive counterfactuals specimens (Fig. 2f). This result is surprising because the UNI-based model had perfect sensitivity in the original test set (Fig. 2b) as well as the highest sensitivity when larger tumors were embedded into normal tissue (Fig. 2e). This highlights that high classification performance on the held-out test set is insufficient to assess generalization to more nuanced downstream applications.

We also sought to quantify the sensitivity of models to each tumor patch in positive specimens, which can shed light on whether tumor patches carry different levels of informativeness for machine learning classifiers. To accomplish this, all tumor patches intersecting with expert tumor annotations were removed. Then, we reintroduced tiles fully within the expert tumor annotation, one at a time, to the tumor-removed specimen and evaluated the model outputs. These model outputs were compared to those when all tumor was removed. While some tumor patches could drive a positive prediction on their own, many could not (Fig. 2g for the UNI-based model, and other models and specimens are shown in Supplementary Figs. 2-6.).

To further quantify the effect of tumor size in metastasis detection, we added tumor patches into normal slides in a graded fashion and measured the sensitivity. All models exhibited a graded effect of tumor size, and UNI exhibited the highest sensitivity (Fig. 2h). Models tended to plateau in sensitivity at  $0.262144 \text{ mm}^2$  of tumor (16 patches) added. The RetCCL-based model showed the lowest sensitivity and the least sensitivity to smaller tumors.

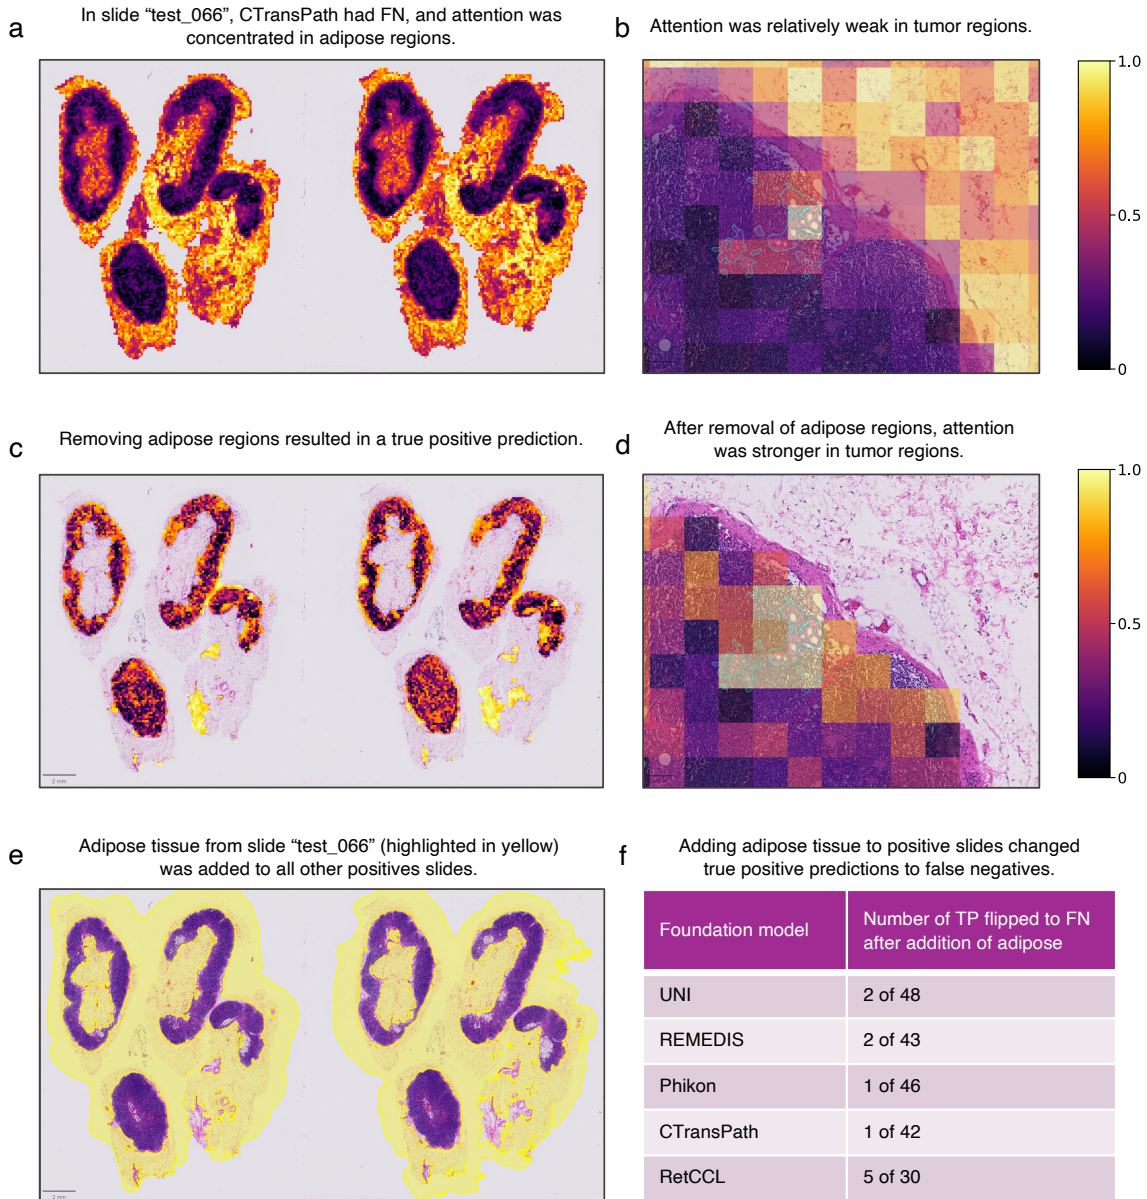
To identify the largest amount of tumor that would go undetected by a model, we also used a HIPPO search algorithm, *HIPPO-search-small-effect*. We found that in some cases, regions up to  $1.5 \text{ mm}^2$  could be added into a negative counterfactual while still maintaining a negative detection. Indeed, the tumor patches that were insufficient to drive large effect sizes were largely similar to the sufficient patches, though some insufficient patches contained adipose cells along with tumor epithelial cells (Supplementary Fig. 7). This shows that there exist regions within tumors that would go unseen by an ABMIL-based metastasis detection model. These biases should be explored further prior to the clinical use of metastasis detection models.

### Non-tumor tissue can cause false positive metastasis detections

Given the effect of peritumoral tissue on UNI-based model predictions, we also evaluated whether peritumoral tissue was sufficient for positive metastasis predictions. A halo of peritumoral tissue was extracted from metastasis-positive specimens (n=49) with a width of 64, 128, 256, or  $1024 \mu\text{m}$ , beginning at either the edge of the expert tumor annotation or  $256 \mu\text{m}$  outside of the tumor annotation. The patches intersecting with the peritumoral halos were added to normal specimens (n=80), resulting in 3920 counterfactual examples (80 normal  $\times$  49 positive specimens). Model predictions were averaged across normal specimens, resulting in the average treatment effect of the peritumoral region of each positive specimen. This was evaluated for the UNI-based and Phikon-based models.

Despite the lack of tumor cells, halos of peritumoral tissue were sufficient to drive a positive metastasis detection. Halos  $1024 \mu\text{m}$  in width beginning from the edge of the expert tumor annotation from 20% (n=10) of positive specimens caused positive predictions in the UNI-based model when embedded in normal specimens, whereas 10% (n=5) caused positive predictions in the Phikon-based model. When starting the  $1024 \mu\text{m}$  halo  $256 \mu\text{m}$  outside of the tumor annotation, there were 14% (n=7) positive predictions in the UNI-based model and 10% (n=5) for the Phikon-based model. Thinner halos of  $64 \mu\text{m}$

beginning 256  $\mu\text{m}$  outside of the tumor annotation also caused positive predictions: 10% (n=5) in the UNI-based model and 6% (n=3) in the Phikon-based model (Supplementary Fig. 8). These results demonstrate that the models learned an association between peritumoral tissue and the presence of metastasis despite the peritumoral tissue not containing any metastasis. This highlights an important bias that could not have been uncovered using attention alone. HIPPO enabled the quantitative assessment of peritumoral tissue on metastasis detection.



**Figure 3. Adipose tissue may cause models to miss breast cancer metastases.** The CTransPath-based ABMIL model misclassified specimen `test_066` as negative. **a-b**, Attention was strongest in adipose regions and relatively weak in the tumor. This led us to hypothesize that the adipose regions were causing the misclassification. **c-d**, The true positive prediction was rescued after adipose regions were removed. In addition, attention was concentrated in tumor regions (**d**). To evaluate how general the effect of this adipose tissue was on preventing true positive predictions, we added the embeddings from the adipose region (yellow highlighted region in **e**) to all positive slides except `test_066`. **f**, The addition of adipose tissue caused several true positive predictions to switch to false negatives, suggesting that the adipose tissue may cause false negatives.

### Adipose tissue can cause false negative metastasis detections

To investigate patient samples that consistently drove false negative predictions further, we visualized the attention maps of the ABMIL models to identify regions that it considers important. For the CTransPath-based model, we observed that attention was concentrated in adipose regions for specimen `test_066` (Figs. 3a and 3b).

Since attention maps only provide a qualitative visualization of regions in an image that the ABMIL models consider important, it is unclear to what extent adipose tissue directly affects model predictions. We address this with HIPPO-attention. The patches with adipose tissue were removed, and the effect of this perturbation on model outputs was quantified. We found that removing the adipose tissue rescued the true positive prediction in specimen `test_066`, suggesting that fat caused the false negative prediction (Figs. 3c and 3d). To test whether the adipose tissue from `test_066` would cause a misclassification in other models and specimens, we added the adipose regions from that specimen into the 48 other positive specimens. We then recorded whether the addition of adipose tissue caused false negative predictions. We found that true positives were flipped to false negatives in 2, 2, 1, 1, and 5 specimens for UNI, REMIDIS, Phikon, CTransPath, and RetCLL, respectively (Figs. 3e and 3f). This highlights how HIPPO can elucidate biases that cause misclassification. While attention alone could not inform that adipose was the cause of misclassification, it was useful to formulate a hypothesis that we can test with HIPPO-attention. This demonstrates how HIPPO can complement attention-based interpretability analysis to quantitatively test hypotheses of putatively important tissue regions.

### HIPPO identifies shortcut learning when attention struggles

Identifying spurious correlations in deep learning models for medical imaging is crucial to ensure reliable and clinically relevant results. To test HIPPO's ability to identify spurious correlations, we conducted an experiment where we deliberately introduced an artificial bias into the CAMELYON16 dataset (Figs. 4a and 4b). Specifically,  $768 \times 768 \mu\text{m}$  blue squares were added to all negative images. This mimics the plausible scenario in which a pathologist marks certain slides with a blue marker. However, in doing so, it introduces a strong spurious correlation with labels. We hypothesized that the models would learn that slides were negative if a blue region was present and that slides lacking this blue region are positive (as blue regions are easier to identify compared to more variable tumor regions).

An ABMIL model was trained on the modified training data using UNI embeddings. The model achieved a balanced accuracy of 1.0 on the test set, suggesting the spurious correlations created a trivial prediction task. By performing standard model interpretation using attention, we found that metastatic regions were considered highly important (Fig. 4c). However, removing these regions using HIPPO did not alter the model predictions, demonstrating that tumor regions were not important for model predictions despite a strong attention assignment. This highlights an important weakness of attention: the disconnect between attended regions and model predictions.

Knowing that the metastatic regions did not affect model outputs, we used the search algorithm *HIPPO-search-high-effect* to identify the regions that maximally drove positive tumor predictions in both models using one positive specimen, `test_002`. Given that the model trained with spurious correlations uses the lack of a blue square as a cue for positive specimens, we expected that no individual patches would drive the positive metastasis output and that tumor regions would not have a high effect on the prediction. Indeed, effect sizes were small and evenly distributed across the WSI (minimum  $2.1 \times 10^{-5}$ , maximum 0.02, mean  $9.4 \times 10^{-5}$ , and median  $5.5 \times 10^{-5}$ ), indicating that no single region contributed strongly to the model prediction (Fig. 4d). By contrast, applying this search algorithm to the model trained on the original CAMELYON16 dataset, we found that patch effect sizes were higher (minimum  $3.7 \times 10^{-8}$ , maximum 0.09, mean  $1.3 \times 10^{-4}$ , and median  $4.9 \times 10^{-8}$ ), and high effect patches were within expert tumor annotations (Fig. 4e). By tying interpretation analysis directly to predictions, HIPPO-based interpretations may provide more reliable explanations of model predictions.

Shortcut learning is an important bias that must be identified and addressed in deep learning on medical images. In this case, model performance and attention were insufficient to diagnose the shortcut learning. Observational analysis based on attention maps could easily mislead an observer to believe that tumor regions drive model predictions. Quantifying effect sizes of tumor regions using HIPPO addressed these limitations and diagnosed the shortcut learning.

### Refining the search for prognostic tissue biomarkers

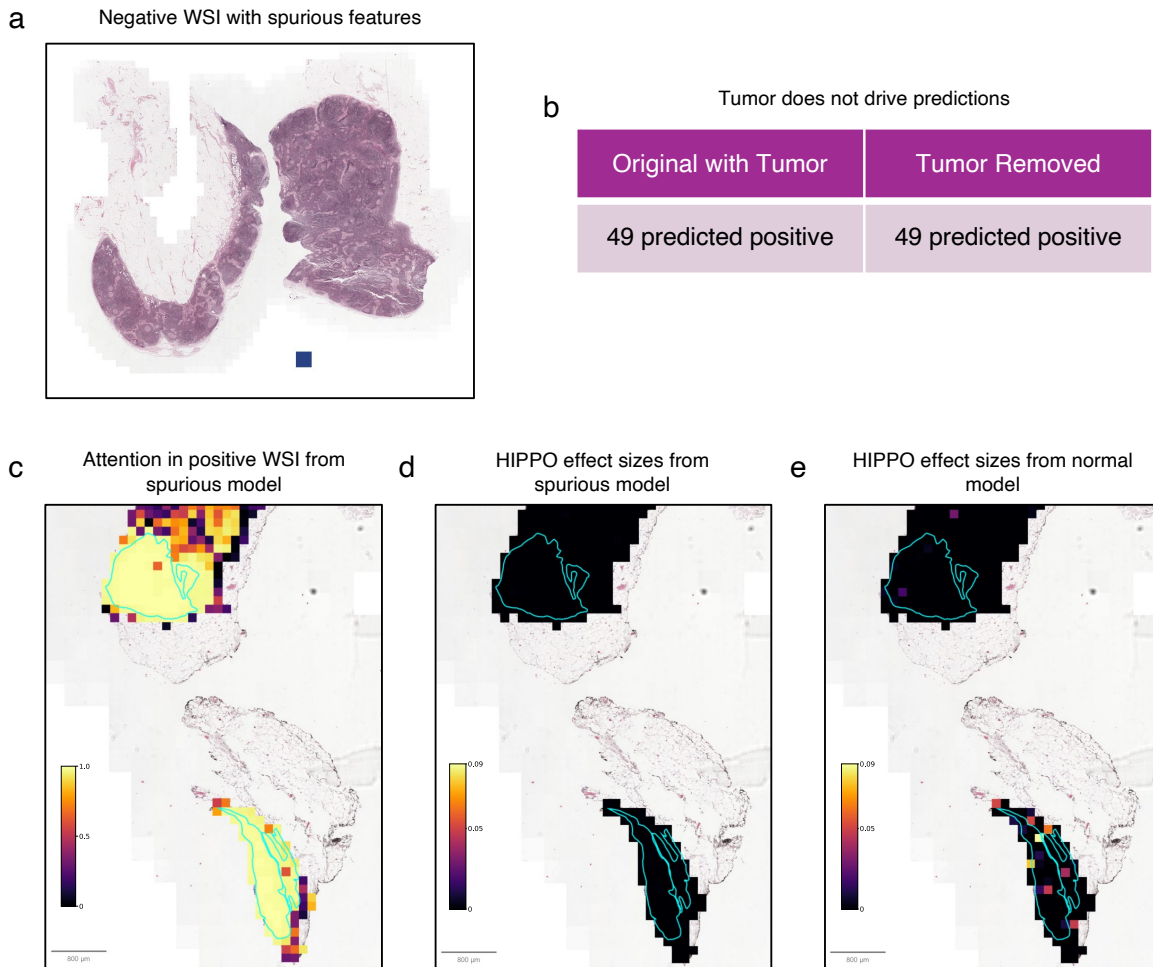
Having demonstrated HIPPO's effectiveness in metastasis detection, where the regions of interest are well-defined and were previously annotated by expert pathologists, we extended our investigation to the more complex domain of cancer prognosis. Unlike the clear delineation of tumor regions in metastasis detection, prognostic factors in WSIs are multifaceted and less clearly defined. We applied HIPPO to prognostic models that generate risk scores from WSIs, aiming to identify the tissue regions driving these predictions. Our experiments with HIPPO yielded two key insights. First, HIPPO's search algorithms demonstrated superior ability in identifying tissue patches that consistently and significantly influence risk predictions compared to conventional attention-based methods. While attention mechanisms yielded mixed effects — potentially identifying regions that counterintuitively drive lower risk in otherwise high-risk specimens — HIPPO provided a more consistent, reliable, and quantitative assessment of the regions that drive risk. Second, HIPPO's unique features enable *in silico* experiments to measure



the effects of targeted tissue interventions on prognostic outcomes through the lens of the ABMIL model. HIPPO’s potential to accelerate the discovery and validation of prognostic tissue biomarkers is an exciting development in cancer research, potentially bridging the gap between computational predictions and clinical actionability.

We trained prognostic ABMIL models using the PORPOISE framework<sup>36</sup>, a computational tool designed for predicting survival outcomes from histopathology images, to predict overall survival from WSIs in breast cancer (TCGA-BRCA) and cutaneous melanoma (TCGA-SKCM) (Supplementary Fig. 9). The same training and validation splits were used as in the original publication. Non-overlapping  $128 \times 128 \mu\text{m}$  patches from WSIs were embedded using the UNI model<sup>58</sup> (in the original PORPOISE publication, a truncated ResNet50<sup>65</sup> was used). Low and high risk were defined as the first and fourth quartiles of risk scores. High attention regions were defined as the top 1% of attended patches, and HIPPO search algorithms were also used to identify the top 1% of patches by effect size.

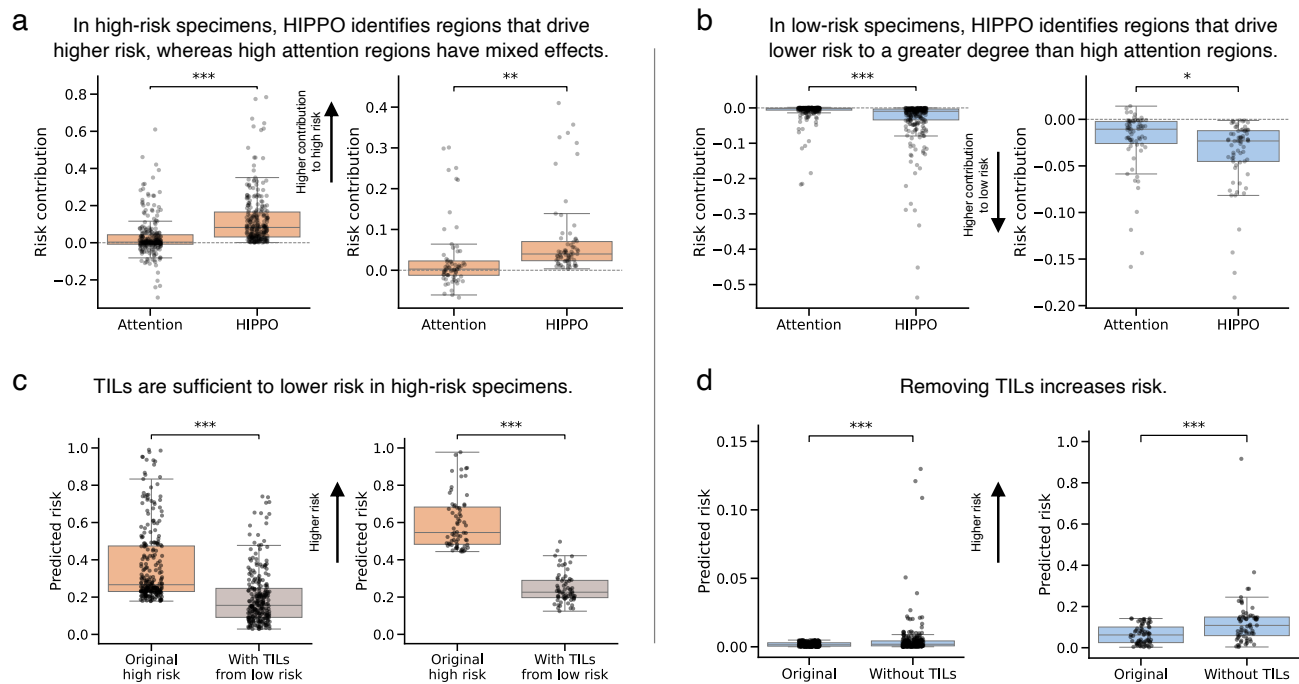
High attention regions drove counterintuitive effects in many specimens, while *HIPPO-search-low-effect* and *HIPPO-*



**Figure 4. HIPPO identifies shortcut learning when attention struggles.** **a**, Thumbnail of a negative specimen (normal\_009) with a  $768 \times 768 \mu\text{m}$  blue square added. A blue square was added to all negatives specimens ( $n=239$ ) in the CAMELYON16 dataset to promote shortcut learning. The UNI foundation model was used to embed the tissue and the blue squares. Positive samples were unaltered. **b**, All positive specimens were predicted as positive, and removal of tumor regions did not change model predictions. This suggested that the ABMIL models learned that if a blue patch is absent, the specimen is positive for metastasis. **c**, Attention heatmap for specimen test\_002, with expert tumor annotation in cyan. Despite tumor having no effect on model predictions, there was strong attention on tumor regions. **d**, Heatmap of patch effect sizes in specimen test\_002 using the ABMIL model trained on deliberate spurious specimens. Using “HIPPO-search-high-effect”, we searched for the patches with highest effect on model outputs. **e**, Heatmap of patch effect sizes in specimen test\_002 using the original ABMIL model, trained without deliberate spurious specimens.

*search-high-effect* identified more robust and consistent drivers of risk. High attention regions in high-risk cutaneous melanoma specimens (n=67) drove lower risk in 45% (n=30) of specimens. *HIPPO-search-high-effect*, on the other hand, identified regions that all drove higher risk and that more greatly contributed to high-risk predictions ( $t = 3.03$ ,  $p < 0.01$ , independent t-test). High attention in high-risk breast cancer specimens (n=256) drove lower risk in 40% (n=102) specimens. Again, HIPPO-search-high consistently identified regions that drove higher risk in the high-risk specimens ( $t = 8.83$ ,  $p < 0.0001$ , independent t-test) (Fig. 5a). High attention regions in low-risk SKCM specimens (n=67) drove higher risk in 10% (n=7). *HIPPO-search-low-effect* identified regions that all drove lower risk and more strongly contributed to lower risk predictions ( $t = -2.30$ ,  $p < 0.05$ , independent t-test). High attention regions in low-risk BRCA specimens (n=256) drove higher risk predictions in 8% (n=20) specimens. *HIPPO-search-low-effect* identified patches that consistently drove lower risk predictions ( $t = -5.43$ ,  $p < 0.0001$ , independent t-test) (Fig. 5b). This counterintuitive effect underscores that attention scores may not directly relate to model predictions. Thus, interpretations that solely rely on these features may be misguided. HIPPO search algorithms reliably identified the regions that drove risk predictions and may have value as a tool for prognostic biomarker search.

TILs are a well-known prognostic biomarker. We evaluated the necessity and sufficiency of TILs for low-risk predictions in BRCA and SKCM. To test sufficiency, we extracted TIL-positive patches from low-risk specimens and placed them in high-risk specimens. For each high-risk slide, we embedded the TILs from each low-risk slide, and we averaged the model predictions across the low-risk slides to compute the average treatment effect of TILs for each high-risk slide. In high-risk BRCA specimens (n=253, three specimens failed cell detection), the addition of TILs from low-risk specimens decreased the risk by 46% ( $t = 17.95$ ,  $p < 0.0001$ , paired t-test) from 0.37 (std. dev. 0.20) to 0.20 (std. dev. 0.15). In SKCM (n=67), the addition of TILs



**Figure 5. HIPPO outperforms attention in identifying prognostic tissue regions.** We studied prognostic ABMIL models in invasive breast carcinoma (BRCA) and cutaneous melanoma (SKCM) from The Cancer Genome Atlas. **a, b**, Box plots of the prognostic effects of patches selected using attention and HIPPO in high-risk (**a**) and low-risk (**b**) specimens. The y-axis depicts the risk contribution, which is calculated as the original predicted risk minus the predicted risk when using a specimen with high-attention or high-HIPPO patches removed. Positive values indicate contribution to higher risk, and negative values indicate contribution to lower risk (**b**). The x-axis is the method of patch selection (either the top 1% of attended patches or the top 1% of patches found using HIPPO-search-high-effect). **c**, Box plots showing the predicted risk scores before and after adding tumor-infiltrating lymphocytes (TILs) to high-risk BRCA (left, n=256) and SKCM (right, n=67) specimens. Orange boxes show the original risk scores, and gray boxes show risk scores after adding TILs from low-risk specimens and averaging across low-risk specimens. Lower risk scores indicate improved prognosis. **d**, Box plots showing the predicted risk scores before and after removing TILs from low-risk BRCA (left, n=256) and SKCM (right, n=67) specimens. Box plots show the first and third quartiles, the median (central line) and the range of data with outliers removed (whiskers), and significance is shown (\*:  $p < 0.05$ , \*\*:  $p < 0.01$ , \*\*\*:  $p < 0.001$ ). Sample sizes in high-risk (**a, c**) and low-risk (**b, d**) are n=256 for BRCA (left) and n=67 for SKCM (right).

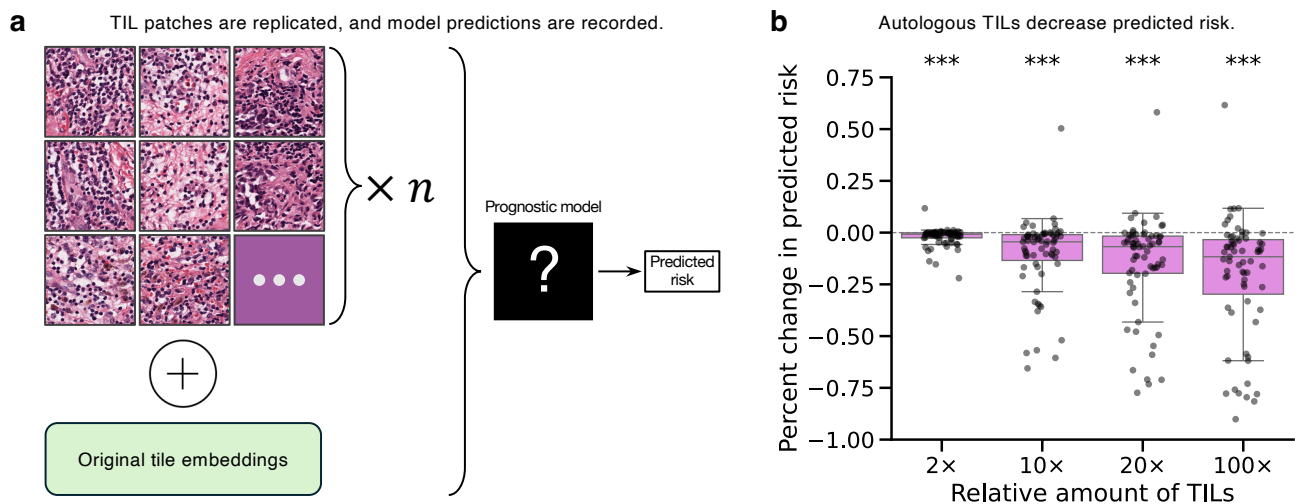
significantly decreased risk by 59% ( $t = -22.53, p < 0.0001$ , paired t-test) from 0.60 (std. dev. 0.14) to 0.25 (std. dev. 0.08) (Fig. 5c). To evaluate the necessity of TILs, we removed TIL-positive patches from low-risk specimens and measured the change in predictions. If TILs were necessary, then risk predictions would increase upon removal of TILs. In BRCA ( $n=254$ , two specimens failed cell detection), the removal of TILs significantly increased risk by 179% ( $t = 3.83, p < 0.001$ , paired t-test) from 0.002 (std. dev. 0.001) to 0.005 (std. dev. 0.014). In SKCM ( $n=67$ ), the removal of TILs increased risk by 98% ( $t = 4.27, p < 0.0001$ , paired t-test) from 0.064 (std. dev. 0.045) to 0.126 (std. dev. 0.123) (Fig. 5d). The removal of TILs did increase risk predictions, but the risk predictions did not reach the level of high-risk slides, suggesting that other features in the WSIs were also driving the low-risk predictions. HIPPO facilitated a quantitative evaluation of the role of TILs on prognosis, providing insights beyond those achievable through the attention mechanism of ABMIL.

### Generating hypotheses of which patients may benefit from autologous TIL therapy

Lifileucel is a promising immunotherapy for melanoma that involves isolating TILs from a patient’s tumor, replicating the TILs, and infusing them back into the patient<sup>1</sup>. In a phase II clinical trial, over 30% of patients responded to the therapy<sup>66</sup>. Identifying the patients that might respond to this therapy has the potential to improve patient outcomes and decrease costs (a single treatment may cost over \$500,000<sup>67</sup>). Therefore, we sought to explore whether we could emulate this with ABMIL and HIPPO. We conducted *in silico* experiments to measure the effect of autologous TILs on prognosis. We used the prognostic model for cutaneous melanoma described above, and we studied the high-risk specimens in TCGA-SKCM ( $n=67$  WSIs,  $n=54$  patients). Counterfactuals were designed to model the injection of autologous TILs. In each specimen, TIL-positive patches were replicated  $2\times$ ,  $10\times$ ,  $20\times$ , and  $100\times$  (Fig. 6a). TIL-positive patches were defined using the same heuristic as above (see Methods). The change in model predictions between original specimens and autologous counterfactuals was recorded to measure the effect of additional TILs on prognosis. Cohen’s  $d$  was also calculated to quantify effect sizes. Importantly, we do not claim to demonstrate the efficacy of autologous TIL therapy through HIPPO and TCGA-SKCM. Rather, we aim to show a proof-of-principle that HIPPO may be used for hypothesis generation.

Autologous TILs significantly lowered predicted risk in a dose-dependent manner. Risk decreased by  $-2.18\%$  ( $d = -0.50$ )

<sup>1</sup><https://www.fda.gov/news-events/press-announcements/fda-approves-first-cellular-therapy-treat-patients-unresectable-or-metastatic-melanoma>



**Figure 6. Autologous TILs improve predicted prognosis.** In high-risk slides of cutaneous melanoma (TCGA-SKCM,  $n=67$ ), TIL-positive patches were identified using a heuristic from<sup>36</sup>. High risk was defined as slides with the top 25% of predicted risk scores. **a**, The embeddings of TIL-positive regions were replicated and concatenated with the original embeddings (the ellipsis denotes that the displayed TIL patches are a representative sample of a larger set). Model predictions are then recorded for this counterfactual with additional autologous TILs. **b**, Box plot showing the difference in model predictions, relative to the original specimens. Differences are shown on the y-axis and were calculated as the predicted risks with autologous TILs minus the original predicted risk (negative values indicate that autologous TILs decreased predicted risk). The x-axis shows the amount of TILs relative to the original specimens. The sample size in each box is 67. Box plots show the first and third quartiles, the median (central line) and the range of data with outliers removed (whiskers), and significance is shown (\*\*\*:  $p < 0.001$ ).

at 2× dose ( $t = -4.06$ ,  $p < 0.001$ , paired t-test), -10.8 % ( $d = -0.56$ ) at 10× dose ( $t = -4.59$ ,  $p < 0.0001$ , paired t-test), -15.3 % ( $d = -0.62$ ) at 20× dose ( $t = -5.06$ ,  $p < 0.0001$ , paired t-test), and -20.8 % ( $d = -0.67$ ) at 100× dose ( $t = -5.49$ ,  $p < 0.0001$ , paired t-test) (Fig. 6b). Increasing the number of TILs by 100× decreased predicted risk scores by over half in 18 % of high-risk specimens. Together, we demonstrated a proof-of-principle in which we use HIPPO to identify patients who may benefit from autologous TIL therapy through improved predicted prognosis following the replication of their TILs.

## Discussion

In this study, we introduce HIPPO, an explainable AI method designed to enhance the interpretability and trustworthiness of ABMIL models in computational pathology. Our results demonstrate HIPPO’s ability to uncover hidden biases, quantify the impact of specific tissue regions on model predictions, and bridge the gap between computational outputs and clinically relevant insights. These findings may have significant implications for the development, regulation, and clinical application of AI in pathology.

One of the key strengths of HIPPO lies in its capacity to reveal model-specific limitations that are not apparent from performance metrics or attention mechanisms alone. In our evaluation of metastasis detection models, we uncovered surprising variations in how different foundation models process histological information. For instance, some models showed a strong reliance on peritumoral tissue, while others demonstrated unexpected insensitivity to small tumor regions. These findings underscore the importance of rigorous model evaluation beyond standard performance metrics and highlight potential pitfalls in clinical deployment.

The revelation that high-attention regions can sometimes have counterintuitive effects on prognostic predictions is particularly striking. This disconnect between attention and model output challenges the common practice of using attention maps as a primary means of model interpretation. Our results suggest that regulatory bodies and clinical teams should exercise caution when relying solely on attention-based explanations and should consider incorporating quantitative impact assessments, such as those provided by HIPPO, in their evaluation processes. For example, one may use HIPPO-knowledge to quantify the effect of high attention regions on model predictions.

HIPPO’s ability to verify that models have learned biologically relevant information, as demonstrated by our analysis of TILs, is crucial for building trust in AI-driven prognostic tools. This alignment between model behavior and established biological knowledge provides a foundation for explaining model decisions to clinicians and patients, potentially facilitating the integration of AI tools into clinical workflows. It is also possible to use HIPPO’s *de novo* search to identify sets of patches from which expert pathologists could interpret manually to identify new biomarkers, helping to understand disease progression.

The application of HIPPO to simulate the effects of autologous TIL therapy in melanoma patients showcases the potential for *in silico* experimentation in computational pathology. As ABMIL models improve, this approach could have far-reaching implications for personalized medicine, offering a computational method to predict treatment responses and guide therapy selection. However, it is important to note that these simulations are based on model predictions and would require extensive clinical validation before they can be considered for patient care.

Our findings suggest several key considerations for the future development and deployment of ABMIL models in clinical settings: (1) Model developers should incorporate robustness to tissue heterogeneity and small tumor regions as explicit design goals, potentially through targeted data augmentation using HIPPO-based counterfactuals; (2) Regulatory approval processes for AI tools in pathology may consider including comprehensive evaluations of model behavior across diverse tissue contexts, going beyond aggregate performance metrics; (3) The implementation of AI tools in clinical practice should be accompanied by clear explanations of model strengths and limitations, with HIPPO-like analyses providing quantitative assessments of model reliability for specific tissue types or patient subgroups; (4) Post-deployment monitoring of AI models should include ongoing analysis of model behavior in real-world settings, with HIPPO offering a means to detect potential shifts in model performance or the emergence of unexpected biases.

While our study demonstrates the potential of HIPPO, several limitations must be acknowledged. First, the counterfactual scenarios generated by HIPPO, while informative, may not always reflect biologically plausible tissue alterations. Future work should focus on refining these interventions to more closely mimic realistic tissue changes. Second, our analysis was limited to a specific set of foundation models and datasets. Broader evaluation across diverse pathology tasks and model architectures is needed to fully characterize the generalizability of our findings. In addition the interpretations offered by HIPPO are inherently bound by the underlying model’s capabilities and potential shortcomings in representing complex biological systems.

Looking ahead, several avenues for future research emerge from this work. The integration of HIPPO with multi-modal data, including genomic and clinical information, could provide even richer insights into model behavior and biological relevance. Additionally, exploring the use of HIPPO in guiding model refinement, such as targeted fine-tuning based on identified weaknesses, represents a promising direction for improving model robustness and clinical applicability.

In conclusion, HIPPO represents a major advance in the ability to interpret AI models in computational pathology. By providing a quantitative framework for assessing the impact of specific tissue regions on model predictions, HIPPO offers a



powerful tool for uncovering model limitations, verifying biological relevance, and biomarker discovery for various clinical applications. As the field of computational pathology continues to evolve, quantitative methods like HIPPO will be crucial in ensuring that AI tools are deployed responsibly and effectively in healthcare settings.

## Methods

### HIPPO toolkit

HIPPO (Histopathology Interventions of Patches for Predictive Outcomes) is an explainable AI toolkit for attention-based multiple instance learning models in computational pathology. It generates counterfactual examples by manipulating whole slide image patches to explain model behavior. In ABMIL, tissue from a WSI is divided into small tiles, which are embedded using a pre-trained model. These patch embeddings serve as input to the ABMIL model, which learns to map bags of patches to specimen-level labels. HIPPO is made possible by two key features of ABMIL: (1) models are invariant to patch order, and (2) models accommodate variable number of patches. Taking advantages of these features of ABMIL, HIPPO creates counterfactual examples by adding or removing patches, allowing for evaluation of hypothetical scenarios and quantification of the effects of tissue region on model predictions. This process can be hypothesis-driven (e.g., removing tumor regions to test their necessity in predicting breast metastases), based on model attention (e.g., removing high-attention regions), or identified by greedy search (described below). Patches can also be added to measure their impact, such as introducing tumor-infiltrating lymphocytes (TILs) to TIL-deficient specimens and measuring impact on prognostic models.

HIPPO employs a search algorithm to identify patches necessary or sufficient for predictions. This algorithm iteratively measures each patch's effect by removal and model re-evaluation. It can determine necessary patches (those causing the largest prediction drop when removed) or sufficient patches (those causing minimal change when removed). In multi-class prediction, this method can identify patches that drive the model towards specific outcomes. In regression models, the search algorithm can identify patches that either increase the output or decrease the output. The greedy search experiments in the present report were conducted on NVIDIA A100 GPUs.

### Deep neural network development

We used attention-based multiple instance learning (ABMIL) to learn specimen-level labels from whole slide images. For metastasis detection, we evaluated five different patch encoders: UNI<sup>58</sup>, REMEDIS<sup>61</sup>, CTransPath<sup>63</sup>, Phikon<sup>62</sup>, and RetCCL<sup>64</sup>. These encoders were used to embed non-overlapping  $128 \times 128 \mu\text{m}$  patches, with all encoders utilizing identical patches. We standardized hyperparameters across all ABMIL models, adapting from Chen et al.<sup>58</sup>. The architecture comprised a first hidden layer of 512 units and a second of 384 units, incorporating gated attention. During training, we applied a dropout rate of 0.25. The output layer performed binary classification, distinguishing between the presence and absence of metastasis. Models were trained using cross-entropy loss and the Adam optimizer with a learning rate of  $1 \times 10^{-4}$ , following a cosine learning rate scheduler. We used a batch size of 1 without gradient accumulation. Training continued for a maximum of 20 epochs, with the best model selected based on the highest ROC AUC on the validation set. To assess initialization variability, we trained five separate models with different random seeds for each patch encoder. For subsequent experiments, we selected the initialization yielding the highest balanced accuracy on the CAMELYON16 test set for each encoder. We visualized attention heatmaps using QuPath<sup>68</sup>. All models were implemented in PyTorch and trained on NVIDIA RTX 2080 Ti GPUs.

For prognostic models, we used the ABMIL models defined in<sup>36</sup>. The model was composed of a linear layer with 512 units, dropout with a rate of 0.25, and a second linear layer of 256 units. Gated attention was used. The model had four outputs, representing hazards at four points in time. Risk scores were calculated as in ref.<sup>36</sup> and were in range  $[0, 1]$ , where 0 indicates lowest probability of survival. Models were all implemented in PyTorch, and training was performed on NVIDIA RTX 2080 Ti GPUs.

## Datasets

### Breast cancer metastasis dataset

We used the CAMELYON16 dataset<sup>16</sup> to study breast cancer metastasis. This dataset consists of 399 images and has fine-grained tumor annotations made by expert pathologists. The training set was split into 90% training and 10% validation, stratified by the label of the specimen (i.e., normal or tumor). Training set consisted of 143 negative and 100 positive WSIs (52 macrometastases and 48 micrometastases). The validation set consisted of 16 negative and 11 positive WSIs (6 macrometastases and 5 micrometastases). We used the pre-defined test set, which consisted of 80 negative and 49 positive WSIs (22 macrometastases and 27 micrometastases). In the entire dataset, there were 160 metastasis-positive specimens. There was an average tumor area of  $12.26 \text{ mm}^2$  (std. dev.  $34.04 \text{ mm}^2$ ; minimum  $0.008 \text{ mm}^2$ ; and maximum  $276.09 \text{ mm}^2$ ). All 399 slides had pixel spacings between 0.226 and  $0.243 \frac{\mu\text{m}}{\text{px}}$  (MPP). The WSIs had  $10,250 \pm 6,672$  patches (mean  $\pm$  standard deviation), where each patch was  $128 \times 128 \mu\text{m}$ .

### **Prognostic datasets**

Prognostic models were trained and evaluated using the invasive breast carcinoma (BRCA) and cutaneous melanoma (SKCM) studies from The Cancer Genome Atlas. In TCGA BRCA, 1,022 WSIs from 956 patients were used (130 death events), and in TCGA SKCM, 268 slides from 230 patients were used (89 death events). Overall survival time and censoring was used and retrieved from the code repository<sup>2</sup> of ref.<sup>36</sup>. The training and validation splits for cross validation were accessed from the same code repository. The WSIs in TCGA BRCA had  $11,260 \pm 6,544$  patches (mean  $\pm$  standard deviation). The WSIs in TCGA SKCM had  $14,153 \pm 7,471$  patches.

### **Whole slide image processing**

Whole slide images were read using OpenSlide<sup>69</sup>, and a modified version of the CLAM toolkit<sup>70</sup> was used to segment tissue and calculate patch coordinates. Regions of tissue are identified to not spend computational resources on glass regions of the WSI. The image is converted to HSV color model (hue, saturation, value/brightness). The saturation channel is smoothed and thresholded to create a binary tissue image. Non-overlapping patches coordinates of  $128 \times 128 \mu\text{m}$  were calculated within the tissue regions. The CLAM toolkit<sup>70</sup> was modified to create patches at uniform physical sizes. The size of a patch in pixels can vary based on the spacing ( $\mu\text{m}/\text{px}$ , MPP), and the patch size in base pixels is calculated using,

$$\text{Patch size (px)} = \frac{\text{Patch size } (\mu\text{m})}{\text{WSI spacing } \left(\frac{\mu\text{m}}{\text{px}}\right)} \quad (1)$$

The  $128 \times 128 \mu\text{m}$  patches were then embedded using five pre-trained models (embedding dimensions in parentheses): UNI (1024)<sup>58</sup>, REMEDIS (4096)<sup>61</sup>, Phikon (768)<sup>62</sup>, CTransPath (768)<sup>63</sup>, and RetCCL (2048)<sup>64</sup>. When embedding a slide, patches were read directly from the WSI file. A batch of 64 patches was read from the WSI, and then that batch was processed by the model to compute embeddings. The embeddings of all patches in a WSI were concatenated into one array and saved to disk for reuse. These embeddings served as inputs to all ABMIL models in the present report.

### **HIPPO experiment details**

#### ***Testing the necessity of tumor regions***

The degree to which tumor regions influence ABMIL models for metastasis detection remains unclear. To test the necessity of tumor regions, all of the tumor in the 49 tumor-positive specimens was removed, and change in model outputs was recorded. The embeddings of all patches intersecting with expert tumor annotations were removed. See Supplementary Fig. 10 for the histogram of the number of patches removed from each WSI. Once all of the tumor patches were removed from the bag of embeddings, the specimen was called “negative” for metastasis. The modified bags of embeddings were run through the model, and outputs were recorded. The true negative rate (specificity) was calculated as the ratio of true negative detections to all negative samples. In this case, as all samples were negative, the true negative rate was the proportion of specimens called negative by the model. This was done for all patch embedding tested in the present report.

#### ***Testing the sufficiency of tumor regions***

The sufficiency of tumor regions for metastasis detection remains unclear. We evaluated this in two ways: by evaluating the use of only tumor tissue from positive specimens ( $n=49$ ), and by embedding metastatic patches from positive specimens ( $n=49$ ) into negative specimens ( $n=80$ ). In the first method, we removed all patches that did not intersect with the expert tumor annotations. This evaluated the hypothetical scenario that the specimen contained tumor and no other type of tissue. The labels of all specimens remained “positive”, and model outputs were recorded. Sensitivity was measured as the proportion of positive model predictions.

In the second method, we created counterfactual examples of metastasis-positive specimens ( $n=3920$ ) from normal specimens ( $n=80$ ). All combinations were evaluated: the patches that intersected expert tumor annotations from each positive slide were added to each negative slide, making a total of 3920 counterfactual examples ( $80 \text{ negative} \times 49 \text{ positive specimens}$ ). Each of these counterfactual examples was labeled “positive” because they contained tumor. These counterfactual examples were then run through the ABMIL models, and outputs were recorded. Sensitivity was measured as the proportion of positive model outputs.

#### ***Testing the effect of tumor size***

The extent to which tumor size affects specimen-level metastasis detection is incompletely understood. Conventionally, this analysis is limited to existing specimens. We explore a more rich set of tumor sizes using counterfactual examples. First, we evaluated the effect of a single  $128 \times 128 \mu\text{m}$  tumor region in normal and metastatic specimens. The tumor region was taken from specimen `test_001` at the coordinates (37878, 63530, 38444, 64096), indicating minimum X, minimum Y, maximum

<sup>2</sup><https://github.com/mahmoodlab/PORPOISE>

X, and maximum Y. For normal specimens, we added the embedding of this one patch into each of the 80 normal specimens and fed these bags of embeddings to the ABMIL model. Sensitivity was measured as the proportion of positive model predictions. We also evaluated this in the context of positive specimens. First, all tumor patches intersecting with expert tumor annotations were removed, and the single patch embedding was added to the bags of embeddings. 48 positive samples were used – the specimen that the patch came from was not included. Sensitivity was measured as the proportion of positive predictions.

In addition, the effect of each individual tumor patch was evaluated for metastasis detection. In the positive slides (n=49), all tumor patches intersecting expert tumor annotations were removed to render the slide negative for metastasis. Then, each tumor patch that was fully contained by the tumor annotations was added to the bag of embeddings one at a time, and model outputs were recorded. Model probabilities for tumor were recorded.

Last, the size of tumor was evaluated by sampling increasing numbers of tumor patches. First, all tumor patches intersecting with expert tumor annotations were removed. Then, tumor patches fully contained by the annotations were randomly sampled and added back to the bag of embeddings. This was evaluated over multiple numbers of sampled patches (i.e., 1, 2, 4, 8, 16, 32, 64). Sensitivity was evaluated as the proportion of positive predictions.

### ***Identifying the largest unseen tumor***

Motivated by the graded effect of tumor size on metastasis detection performance, we sought to identify the largest area of tumor that would still result in a negative prediction by the ABMIL models. To accomplish this, we used a HIPPO search algorithm. First, all patches that intersected the expert tumor annotation were removed, to render the specimen “negative” for metastasis. Then, tumor patches were added to the specimen one at a time, and model outputs were assessed. The tumor patch that resulted in the lowest model probability of tumor was kept in the bag, and the next round of the search was initiated. This was repeated until the model probability of tumor was greater than 0.5, which would trigger a positive prediction. The set of tumor patches that were in the bag prior to reaching a threshold of 0.5 were considered the largest area of tumor that could be present while maintaining a negative predictions.

### ***Testing the effect of adipose tissue on metastasis detection***

Upon inspection of attention maps for the CTransPath-based metastasis detection model, adipose regions had high attention in a false negative, leading us to hypothesize that adipose regions were driving the false negative in that specimen. Attention alone could not allow us to address this hypothesis, but HIPPO could. The adipose regions were annotated in QuPath. Patches that intersected with the adipose region were removed, while ensuring that no tumor patches were removed. To measure the effect of this adipose tissue in other specimens, patches intersecting with the adipose annotation were added to the other 48 metastasis-positive slides, and the number of changes from true positive to false negative were recorded.

### ***Diagnosing shortcut learning***

We sought to evaluate how HIPPO can uncover shortcut learning and how it compares to attention in this regard. To do this, we modified the normal specimens in the CAMELYON16 dataset to include a blue square (hexadecimal color code #284283). This is meant to mimic a plausible real world scenario in which a pathologist marked certain slides with a blue pen. In practice, we embedded one blue square of  $128 \times 128 \mu\text{m}$  using the UNI model<sup>58</sup> and replicated that embedding 36 times to create a  $768 \times 768 \mu\text{m}$  blue region. The embeddings of this blue region were concatenated with the patch embeddings of normal specimens. The specimens with metastasis were not modified. We reasoned that the ABMIL model would learn to distinguish normal from metastatic specimens by the presence of a blue region. To assess whether tumor regions were affecting model predictions in positive specimens, we removed all patches intersecting with tumor annotations in positive specimens and recorded model outputs. To visualize attention maps, we saved patch-wise attention weights in GeoJSON format and visualized the maps in QuPath<sup>68</sup>. We also used the search strategy *HIPPO-search-high-effect* to identify the regions with highest effect sizes *de novo*. We also did this using a UNI-based ABMIL model trained on the original, unaltered CAMELYON16 dataset, trained using the same hyperparameters and random seed.

### ***Identifying prognostic regions and comparing with attention***

We sought to compare the effectiveness of attention and HIPPO for identifying tissue regions related to predicted prognosis. TCGA BRCA and SKCM data were used in these experiments. For attention, regions assigned the top 1% of attention scores were selected. For HIPPO, the search strategy *HIPPO-search-high-effect* was used to identify the regions most contributing to high risk in high-risk specimens, and the search strategy *HIPPO-search-low-effect* was used to identify the regions most contributing to low risk in low-risk specimens. Low and high risk were defined as the first and fourth quartiles of predicted risk scores, respectively. The first 1% of patches identified by the HIPPO search algorithms were selected for evaluation. To quantify the effect of the selected regions on predicted prognosis, we calculated the difference between the predicted prognosis on the original specimens and the predicted prognosis on the specimens with the selected regions removed.

$$\text{Risk contribution of ROI} = \text{Risk using original WSI} - \text{Risk when ROI is removed} \quad (2)$$

Positive values indicated that the regions contributed to higher risk, and negative values indicated that the regions contributed to lower risk. Independent t-tests were used to assess significance of differences between attention and HIPPO.

### ***Effect of TILs on prognostic models***

In prognostic models, we measured the effects of tumor-infiltrating lymphocytes (TILs) on model behavior. The number of TILs was quantified using the same approach as Ref.<sup>36</sup>. Briefly, HoVer-Net<sup>71</sup> was used to outline and label the nuclei in TCGA BRCA and SKCM WSIs. The model labels nuclei as one of six categories: tumor epithelium, lymphocyte, stroma, necrosis, normal epithelium, and unknown. Each  $128 \times 128 \mu\text{m}$  was called TIL-positive if it contained more than 20 cells, more than 10 immune cells, and more than 5 tumor cells. In TCGA BRCA, HoVer-Net failed for 12 WSIs, some of which were missing pixel spacing information.

We measured the effect of TIL patches on predicted prognosis in TCGA BRCA AND SKCM by either removing TILs from low-risk specimens or adding TILs to high-risk specimens, where low-risk was defined as samples in the first quartile of predicted risk and high-risk were samples in the fourth quartile of predicted risk. The predicted prognoses were compared before and after the intervention. To evaluate the sufficiency of TILs for predicting low risk, we added TIL patches from low-risk specimens to high-risk specimens. Risk predictions of the model were recorded, and differences were tested using paired t-tests. To assess the necessity of TIL regions, we removed TIL-positive patches from low risk specimens and measured risk predictions. Differences were tested using paired t-tests.

### ***Evaluating autologous TILs***

Autologous TIL therapy is a promising immunotherapy. We explored how HIPPO could be used for hypothesis generation in the context of autologous TILs in high-risk SKCM specimens (n=67). We sought to assess the degree to which prognostic ABMIL models are effected by the number of TILs in a specimen. We do not claim to assess the efficacy of autologous TILs through HIPPO. The embeddings of TIL-positive regions were replicated 2×, 10×, 20×, or 100×, and the change in predicted risk was measured:

$$\text{Change in Risk} = \text{Risk with autologous TILs} - \text{Risk with original WSI} \quad (3)$$

Negative values indicated that the addition of TILs decreased risk. The change in risk from baseline was assessed using paired t-tests.

## **Data availability**

The CAMELYON16 dataset is available at <https://camelyon17.grand-challenge.org/Data/> under the CC0 license (public domain). The results shown here are in whole or part based upon data generated by the TCGA Research Network: <https://www.cancer.gov/tcga>. Clinical data and whole slide image files can be accessed at <https://portal.gdc.cancer.gov>. Training and validation splits for prognostic models were accessed at <https://github.com/mahmoodlab/PORPOISE>.

## **Code availability**

A Python package implementing HIPPO is available at <https://github.com/kaczmarj/HIPPO> and is licensed under the terms of the 3-Clause BSD License. HIPPO documentation is published at <https://github.com/kaczmarj/HIPPO> under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International copyright license (CC BY-NC-SA 4.0). Model weights and inference code are available at the following repositories: UNI (<https://huggingface.co/MahmoodLab/UNI>), REMEDIS (<https://github.com/google-research/medical-ai-research-foundations>), Phikon (<https://huggingface.co/owkin/phikon>), CTransPath (<https://github.com/Xiyue-Wang/TransPath>), and RetCCL (<https://github.com/Xiyue-Wang/RetCCL>). Model weights for the models trained for this report will be deposited to online repositories.

## **Acknowledgements**

This research was supported by National Science Foundation (NSF) grant IIS2212046, National Institutes of Health (NIH) grant UH3CA225012, and Stony Brook Profund 2022 seed funding. JRK was also supported by the Medical Scientist Training Program at Stony Brook University and NIH grant T32GM008444 (NIGMS). We would also like to acknowledge the Department of Biomedical Informatics at Stony Brook University and the Simons Center for Quantitative Biology at Cold Spring Harbor Laboratory.



## Author contributions

JRK and PKK conceived of the method and planned the experiments. JRK wrote the code and ran all experiments. JRK, PKK, and JHS interpreted the results of experiments. JHS and PKK supervised the project. JRK wrote the initial draft of the manuscript. All authors provided feedback on the manuscript and contributed to the final manuscript.

## Competing interests

The authors declare the following competing interests: J.H.S. is co-founder and chief executive officer of Chilean Wool, LLC. All other authors declare no competing interests.

## References

1. Van der Laak, J., Litjens, G. & Ciompi, F. Deep learning in histopathology: the path to the clinic. *Nat. medicine* **27**, 775–784 (2021).
2. Echle, A. *et al.* Deep learning in cancer pathology: a new generation of clinical biomarkers. *Br. journal cancer* **124**, 686–696 (2021).
3. Topol, E. J. High-performance medicine: the convergence of human and artificial intelligence. *Nat. medicine* **25**, 44–56 (2019).
4. Song, A. H. *et al.* Artificial intelligence for digital and computational pathology. *Nat. Rev. Bioeng.* **1**, 930–949 (2023).
5. Niazi, M. K. K., Parwani, A. V. & Gurcan, M. N. Digital pathology and artificial intelligence. *The lancet oncology* **20**, e253–e261 (2019).
6. Rakha, E. A. *et al.* Current and future applications of artificial intelligence in pathology: a clinical perspective. *J. clinical pathology* **74**, 409–414 (2021).
7. Cui, M. & Zhang, D. Y. Artificial intelligence and computational pathology. *Lab. Investig.* **101**, 412–422 (2021).
8. Morales, S., Engan, K. & Naranjo, V. Artificial intelligence in computational pathology—challenges and future directions. *Digit. Signal Process.* **119**, 103196 (2021).
9. Tizhoosh, H. R. & Pantanowitz, L. Artificial intelligence and digital pathology: challenges and opportunities. *J. pathology informatics* **9**, 38 (2018).
10. Tran, K. A. *et al.* Deep learning in cancer diagnosis, prognosis and treatment selection. *Genome Medicine* **13**, 1–17 (2021).
11. Cifci, D., Veldhuizen, G. P., Foersch, S. & Kather, J. N. Ai in computational pathology of cancer: improving diagnostic workflows and clinical outcomes? *Annu. Rev. Cancer Biol.* **7**, 57–71 (2023).
12. Shmatko, A., Ghaffari Laleh, N., Gerstung, M. & Kather, J. N. Artificial intelligence in histopathology: enhancing cancer research and clinical oncology. *Nat. cancer* **3**, 1026–1038 (2022).
13. Perez-Lopez, R., Ghaffari Laleh, N., Mahmood, F. & Kather, J. N. A guide to artificial intelligence for cancer researchers. *Nat. Rev. Cancer* 1–15 (2024).
14. Ilse, M., Tomczak, J. & Welling, M. Attention-based deep multiple instance learning. In *International conference on machine learning*, 2127–2136 (PMLR, 2018).
15. Gadermayr, M. & Tschuchnig, M. Multiple instance learning for digital pathology: A review of the state-of-the-art, limitations & future potential. *Comput. Med. Imaging Graph.* 102337 (2024).
16. Ehteshami Bejnordi, B. *et al.* Diagnostic Assessment of Deep Learning Algorithms for Detection of Lymph Node Metastases in Women With Breast Cancer. *JAMA* **318**, 2199–2210, DOI: [10.1001/jama.2017.14585](https://doi.org/10.1001/jama.2017.14585) (2017).
17. Bulten, W. *et al.* Artificial intelligence for diagnosis and gleason grading of prostate cancer: the panda challenge. *Nat. medicine* **28**, 154–163 (2022).
18. Qiu, S., Guo, Y., Zhu, C., Zhou, W. & Chen, H. Attention based multi-instance thyroid cytopathological diagnosis with multi-scale feature fusion. In *2020 25th International Conference on Pattern Recognition (ICPR)*, 3536–3541 (IEEE, 2021).
19. Butke, J. *et al.* End-to-end multiple instance learning for whole-slide cytopathology of urothelial carcinoma. In *MICCAI Workshop on Computational Pathology*, 57–68 (PMLR, 2021).
20. Del Amor, R. *et al.* Constrained multiple instance learning for ulcerative colitis prediction using histological images. *Comput. methods programs biomedicine* **224**, 107012 (2022).

21. Del Amor, R. *et al.* An attention-based weakly supervised framework for spitzoid melanocytic lesion diagnosis in whole slide images. *Artif. intelligence medicine* **121**, 102197 (2021).
22. Lu, M. Y. *et al.* Ai-based pathology predicts origins for cancers of unknown primary. *Nature* **594**, 106–110 (2021).
23. Su, Z. *et al.* Attention2majority: Weak multiple instance learning for regenerative kidney grading on whole slide images. *Med. Image Analysis* **79**, 102462, DOI: <https://doi.org/10.1016/j.media.2022.102462> (2022).
24. Li, J. *et al.* A multi-resolution model for histopathology image classification and localization with multiple instance learning. *Comput. Biol. Medicine* **131**, 104253, DOI: <https://doi.org/10.1016/j.combiomed.2021.104253> (2021).
25. Saldanha, O. L. *et al.* Self-supervised attention-based deep learning for pan-cancer mutation prediction from histopathology. *NPJ Precis. Oncol.* **7**, 35 (2023).
26. Dernbach, G. *et al.* Dissecting ai-based mutation prediction in lung adenocarcinoma: a comprehensive real-world study. *Eur. J. Cancer* 114292 (2024).
27. Zheng, Q. *et al.* Predicting tumor mutation burden and vhl mutation from renal cancer pathology slides with self-supervised deep learning. *Cancer Medicine* **13**, e70112 (2024).
28. Schirris, Y., Gavves, E., Nederlof, I., Horlings, H. M. & Teuwen, J. Deepsmile: Contrastive self-supervised pre-training benefits msi and hrd classification directly from h&e whole-slide images in colorectal and breast cancer. *Med. Image Analysis* **79**, 102464, DOI: <https://doi.org/10.1016/j.media.2022.102464> (2022).
29. Valieris, R. *et al.* Weakly-supervised deep learning models enable her2-low prediction from h & e stained slides. *Breast Cancer Res.* **26**, 124 (2024).
30. Boehm, K. M. *et al.* Multimodal histopathologic models stratify hormone receptor-positive early breast cancer. *bioRxiv* DOI: [10.1101/2024.02.23.581806](https://doi.org/10.1101/2024.02.23.581806) (2024). <https://www.biorxiv.org/content/early/2024/02/28/2024.02.23.581806.full.pdf>.
31. El Nahhas, O. S. *et al.* Regression-based deep-learning predicts molecular biomarkers from pathology slides. *nature communications* **15**, 1253 (2024).
32. Lipkova, J. *et al.* Deep learning-enabled assessment of cardiac allograft rejection from endomyocardial biopsies. *Nat. medicine* **28**, 575–582 (2022).
33. Jiang, Y. *et al.* Biology-guided deep learning predicts prognosis and cancer immunotherapy response. *Nat. Commun.* **14**, 5135 (2023).
34. Mallya, M., Mirabadi, A. K., Farahani, H. & Bashashati, A. Benchmarking histopathology foundation models for ovarian cancer bevacizumab treatment response prediction from whole slide images. *arXiv preprint arXiv:2407.20596* (2024).
35. Chen, R. J. *et al.* Multimodal co-attention transformer for survival prediction in gigapixel whole slide images. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4015–4025 (2021).
36. Chen, R. J. *et al.* Pan-cancer integrative histology-genomic analysis via multimodal deep learning. *Cancer Cell* **40**, 865–878, DOI: [10.1016/j.ccell.2022.07.004](https://doi.org/10.1016/j.ccell.2022.07.004) (2022).
37. Ammeling, J. *et al.* Attention-based multiple instance learning for survival prediction on lung cancer tissue microarrays. In *BVM Workshop*, 220–225 (Springer, 2023).
38. Plass, M. *et al.* Explainability and causability in digital pathology. *The J. Pathol. Clin. Res.* **9**, 251–260, DOI: <https://doi.org/10.1002/cjp2.322> (2023).
39. Castro, D. C., Walker, I. & Glocker, B. Causality matters in medical imaging. *Nat. Commun.* **11**, 3673 (2020).
40. Javed, S. A. *et al.* Additive mil: Intrinsically interpretable multiple instance learning for pathology. *Adv. Neural Inf. Process. Syst.* **35**, 20689–20702 (2022).
41. Nan, T. *et al.* Establishing truly causal relationship between whole slide image predictions and diagnostic evidence subregions in deep learning. *arXiv preprint arXiv:2407.17157* (2024).
42. Liu, X., Zhang, W. & Zhang, M.-L. Attention is not what you need: Revisiting multi-instance learning for whole slide image classification. *arXiv preprint arXiv:2408.09449* (2024).
43. Raff, E. & Holt, J. Reproducibility in multiple instance learning: a case for algorithmic unit tests. *Adv. Neural Inf. Process. Syst.* **36** (2024).
44. Lin, T., Yu, Z., Hu, H., Xu, Y. & Chen, C.-W. Interventional bag multi-instance learning on whole-slide pathological images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19830–19839 (2023).

45. Pocevičiūtė, M., Eilertsen, G. & Lundström, C. Survey of xai in digital pathology. *Artif. intelligence machine learning for digital pathology: state-of-the-art future challenges* 56–88 (2020).
46. Ribeiro, M. T., Singh, S. & Guestrin, C. "why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144 (2016).
47. Lundberg, S. M. & Lee, S.-I. A unified approach to interpreting model predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, 4768–4777 (Curran Associates Inc., Red Hook, NY, USA, 2017).
48. Lipton, Z. C. The mythos of model interpretability. *arXiv 1606.03490* (2016).
49. Zeiler, M. D. & Fergus, R. Visualizing and understanding convolutional networks. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part I 13*, 818–833 (Springer, 2014).
50. Lin, C., Covert, I. & Lee, S.-I. On the robustness of removal-based feature attributions. *Adv. Neural Inf. Process. Syst.* **36** (2024).
51. Bertalmio, M., Sapiro, G., Caselles, V. & Ballester, C. Image inpainting. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 417–424 (2000).
52. Tjoa, E. & Guan, C. A survey on explainable artificial intelligence (xai): Toward medical xai. *IEEE transactions on neural networks learning systems* **32**, 4793–4813 (2020).
53. Graziani, M., Lompech, T., Müller, H. & Andrearczyk, V. Evaluation and comparison of cnn visual explanations for histopathology. In *Proceedings of the AAAI Conference on Artificial Intelligence Workshops (XAI-AAAI-21), Virtual Event*, 8–9 (2021).
54. Afshar, P., Hashembeiki, S., Khani, P., Fatemizadeh, E. & Rohban, M. H. Ibo: Inpainting-based occlusion to enhance explainable artificial intelligence evaluation in histopathology. *arXiv preprint arXiv:2408.16395* (2024).
55. DeGrave, A. J., Cai, Z. R., Janizek, J. D., Daneshjou, R. & Lee, S.-I. Auditing the inference processes of medical-image classifiers by leveraging generative ai and the expertise of physicians. *Nat. Biomed. Eng.* 1–13 (2023).
56. Ktena, I. *et al.* Generative models improve fairness of medical classifiers under distribution shifts. *Nat. Medicine* 1–8 (2024).
57. Taori, R. *et al.* Measuring robustness to natural distribution shifts in image classification. *Adv. Neural Inf. Process. Syst.* **33**, 18583–18599 (2020).
58. Chen, R. J. *et al.* Towards a general-purpose foundation model for computational pathology. *Nat. Medicine* **30**, 850–862, DOI: [10.1038/s41591-024-02857-3](https://doi.org/10.1038/s41591-024-02857-3) (2024).
59. Geirhos, R. *et al.* Shortcut learning in deep neural networks. *Nat. Mach. Intell.* **2**, 665–673, DOI: <https://doi.org/10.1038/s42256-020-00257-z> (2020).
60. DeGrave, A. J., Janizek, J. D. & Lee, S.-I. Ai for radiographic covid-19 detection selects shortcuts over signal. *Nat. Mach. Intell.* **3**, 610–619 (2021).
61. Azizi, S. *et al.* Robust and efficient medical imaging with self-supervision. *arXiv preprint arXiv:2205.09723* (2022).
62. Filiot, A. *et al.* Scaling self-supervised learning for histopathology with masked image modeling. *medRxiv* DOI: [10.1101/2023.07.21.23292757](https://doi.org/10.1101/2023.07.21.23292757) (2023).
63. Wang, X. *et al.* Transformer-based unsupervised contrastive learning for histopathological image classification. *Med. Image Analysis* **81**, 102559, DOI: [10.1016/j.media.2022.102559](https://doi.org/10.1016/j.media.2022.102559) (2022).
64. Wang, X. *et al.* Retccl: Clustering-guided contrastive learning for whole-slide image retrieval. *Med. Image Analysis* **83**, 102645, DOI: [10.1016/j.media.2022.102645](https://doi.org/10.1016/j.media.2022.102645) (2023).
65. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778 (2016).
66. Sarnaik, A. A. *et al.* Lifileucel, a tumor-infiltrating lymphocyte therapy, in metastatic melanoma. *J. Clin. Oncol.* **39**, 2656–2666, DOI: [10.1200/JCO.21.00612](https://doi.org/10.1200/JCO.21.00612) (2021). PMID: 33979178.
67. Healey, N. Tumor infiltrating lymphocyte approval heralds new era for precision cancer immunotherapy. *Nat. Medicine* **30**, 1795–1796, DOI: [10.1038/d41591-024-00040-2](https://doi.org/10.1038/d41591-024-00040-2) (2024).

68. Bankhead, P. *et al.* Qupath: Open source software for digital pathology image analysis. *Sci. reports* **7**, 1–7, DOI: [10.1038/s41598-017-17204-5](https://doi.org/10.1038/s41598-017-17204-5) (2017).
69. Goode, A., Gilbert, B., Harkes, J., Jukic, D. & Satyanarayanan, M. Openslide: A vendor-neutral software foundation for digital pathology. *J. pathology informatics* **4**, 27 (2013).
70. Lu, M. Y. *et al.* Data-efficient and weakly supervised computational pathology on whole-slide images. *Nat. biomedical engineering* **5**, 555–570 (2021).
71. Graham, S. *et al.* Hover-net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images. *Med. image analysis* **58**, 101563 (2019).




Encoder	Balanced Accuracy	Sensitivity	Specificity	Precision	Weighted F1
UNI	0.982 (0.014)	0.976 (0.017)	0.988 (0.022)	0.980 (0.033)	0.983 (0.015)
REMEDIS	0.922 (0.031)	0.861 (0.062)	0.983 (0.014)	0.968 (0.024)	0.935 (0.025)
Phikon	0.907 (0.083)	0.845 (0.158)	0.970 (0.023)	0.943 (0.049)	0.920 (0.070)
CTransPath	0.858 (0.016)	0.784 (0.045)	0.933 (0.023)	0.879 (0.034)	0.874 (0.013)
RetCCL	0.745 (0.016)	0.567 (0.042)	0.922 (0.049)	0.827 (0.074)	0.777 (0.019)


Supplementary Table 1: Performance of metastasis detection models. Values shown are mean (standard deviation) across five random initializations. All models used the same hyperparameters and data splits. The training, validation, and test sets consisted of 243, 27, and 129 specimens, respectively. The test set consisted of 80 negative specimens, 22 samples with macrometastases, and 27 specimens with micrometastases.

Encoder	Balanced Accuracy	Sensitivity	Specificity	Precision	Weighted F1
UNI	1.000	1.000	1.000	1.000	1.000
REMEDIS	0.949	0.898	1.000	1.000	0.961
Phikon	0.955	0.959	0.950	0.922	0.954
CTransPath	0.885	0.857	0.912	0.857	0.891
RetCCL	0.769	0.612	0.925	0.833	0.799


Supplementary Table 2: Performance of the best metastasis detection models for each encoder. This table lists the performance metrics for the single best random initialization for each encoder. Performance was calculated on the CAMELYON16 test set. The models represented here were the ones used for downstream experiments in the present report.



Supplementary Figure 1: Peritumoral tissue may affect metastasis detection. The expert tumor annotations in the positive specimens (n=49) of the CAMELYON16 test set were systematically dilated, and then patches that intersected with these dilated regions were removed. This effectively removed tumor tissue and varying amount of peritumoral tissue, rendering the specimens negative for metastasis. Specificity of model outputs (true negative rate) was calculated. Dilation of tumor annotations did not change model outputs in REMEDIS-based or RetCCL-based models, suggesting that peritumoral tissue did not drive model outputs. The largest dilation (i.e., 1024  $\mu\text{m}$ ) increased specificity in Phikon-based and CTransPath-based models, suggesting that peritumoral was responsible to some degree for false positive predictions. The UNI-based model demonstrated a graded effect of dilation, suggesting that the tissue surrounding the tumor was driving positive model predictions. This effect was particularly strong in macrometastases (n=22). It appears that the UNI-based model relies on peritumoral tissue to some degree for positive predictions.




Supplementary Figure 2: Individual tumor patch effects on metastasis detection using the UNI-based model. Histograms show model probabilities of tumor, where the  $x$ -axis is model probability and  $y$ -axis is the number of examples in the bin. Each histogram represents a different positive specimen ( $n=49$ ) in the CAMELYON16 test set. First, all patches intersecting the expert tumor annotations were removed. Then, patches full contained within the annotation were added back into the specimen one at a time, and model predictions were recorded.




Supplementary Figure 3: Individual tumor patch effects on metastasis detection using the REMEDIS-based model. Histograms show model probabilities of tumor, where the  $x$ -axis is model probability and  $y$ -axis is the number of examples in the bin. Each histogram represents a different positive specimen ( $n=49$ ) in the CAMELYON16 test set. First, all patches intersecting the expert tumor annotations were removed. Then, patches full contained within the annotation were added back into the specimen one at a time, and model predictions were recorded.







Supplementary Figure 4: Individual tumor patch effects on metastasis detection using the Phikon-based model. Histograms show model probabilities of tumor, where the  $x$ -axis is model probability and  $y$ -axis is the number of examples in the bin. Each histogram represents a different positive specimen ( $n=49$ ) in the CAMELYON16 test set. First, all patches intersecting the expert tumor annotations were removed. Then, patches full contained within the annotation were added back into the specimen one at a time, and model predictions were recorded.




Supplementary Figure 5: Individual tumor patch effects on metastasis detection using the CTransPath-based model. Histograms show model probabilities of tumor, where the  $x$ -axis is model probability and  $y$ -axis is the number of examples in the bin. Each histogram represents a different positive specimen ( $n=49$ ) in the CAMELYON16 test set. First, all patches intersecting the expert tumor annotations were removed. Then, patches full contained within the annotation were added back into the specimen one at a time, and model predictions were recorded.



Supplementary Figure 6: Individual tumor patch effects on metastasis detection using the RetCCL-based model. Histograms show model probabilities of tumor, where the  $x$ -axis is model probability and  $y$ -axis is the number of examples in the bin. Each histogram represents a different positive specimen ( $n=49$ ) in the CAMELYON16 test set. First, all patches intersecting the expert tumor annotations were removed. Then, patches full contained within the annotation were added back into the specimen one at a time, and model predictions were recorded.

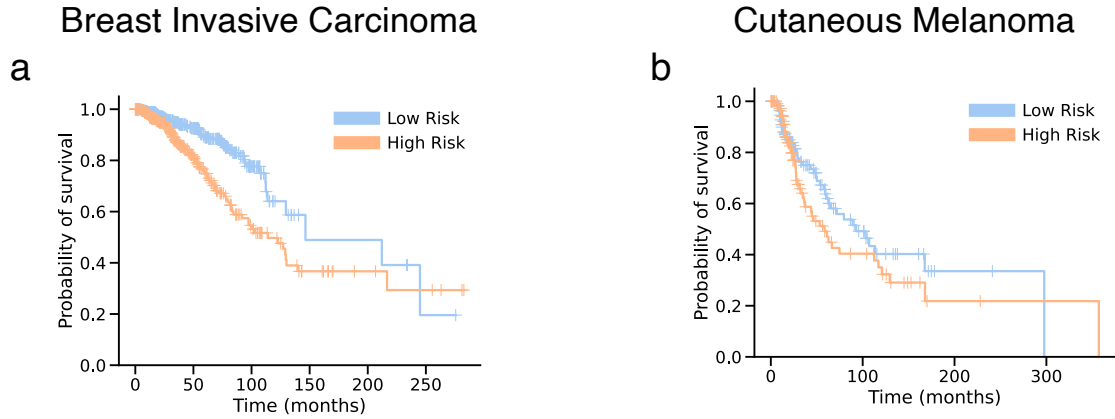


Supplementary Figure 7: Tumor patches have variable effects on metastasis detection, and some tumor regions go undetected entirely. To evaluate the effect of each individual tumor patch on metastasis detection, first all patches intersecting with expert tumor annotations was removed in the positive specimens ( $n=49$ ) of the CAMELYON16 test set. Then, patches that were fully contained in the tumor annotation were introduced into the specimen one at a time, and the model probability of metastasis was recorded. (a) shows a representative example of model probabilities of metastasis for each tumor patch, using the UNI-based model in specimen “test\_051”. The expert tumor annotation is outlined in cyan. A subset of patches was sufficient to drive a positive tumor prediction (model probability  $> 0.5$ ), but many tumor patches were insufficient to drive a positive prediction on their own. Some of these insufficient patches contained tumor epithelial cells along with adipose cells, but many did not. We also used a version of the search algorithm *HIPPO-search-low-effect* to identify the largest set of tumor patches that can be added to a negative counterfactual while still maintaining a negative prediction. First, all tumor patches intersecting the tumor boundary were removed. Then, we iteratively added tumor patches back into the specimen, and kept the tumor patch that drove the lowest probability of metastasis. This continued until the model probability was greater than 0.5. These “unseen” tumor regions are highlighted in yellow in (b) and (c), and the tumor region is highlighted in cyan. In (b) (specimen “test\_094”), we identified a  $0.95 \text{ mm}^2$  area of tumor that was undetected by the UNI-based model, and in (c), we identified a  $1.0 \text{ mm}^2$  area of tumor that was undetected by the UNI-based model.




Supplementary Figure 8: Non-tumor tissue is sufficient for positive detections in specimens without tumor. (a) To evaluate the sufficiency of non-tumor tissue from positive specimens to drive positive detections in negative specimens, halos of peritumoral tissue were selected. These halos did not intersect with the expert tumor annotations, and as such were considered to be entirely non-tumor. The patches intersecting with the halo but not intersecting with tumor annotations were added to normal specimens, resulting in 3,920 counterfactual examples (80 negative  $\times$  49 positive specimens). (b) The model's probability of metastasis was averaged across each negative specimen to evaluate the global effect of the peritumoral halo on model outputs. Four widths of halos were evaluated (i.e., 64, 128, 256, and 1024  $\mu\text{m}$ ), beginning at either the outer edge of the expert tumor annotation (left column) or 256  $\mu\text{m}$  outside of the annotation (right column). Two foundation models were evaluated: UNI (top row) and Phikon (bottom row). Multiple halos of non-tumor tissue were sufficient to drive false positive metastasis detection. In the UNI-based model, for example, a 64  $\mu\text{m}$  halo beginning at the tumor annotation border from 7 positive specimens was sufficient to drive false positives. In the Phikon-based model, a 1024  $\mu\text{m}$  halo was sufficient for false positive predictions from 5 positive specimens.





Supplementary Figure 9: **a**, **b**, Kaplan Meier plots for breast cancer (BRCA) (**a**) and cutaneous melanoma (SKCM) (**b**) in The Cancer Genome Atlas. Prognostic attention-based multiple instance learning models were trained to learn overall survival from whole slide images (WSIs), and risk scores were used to stratify patients. If a patient had multiple WSIs, the predicted prognoses were averaged across WSIs to arrive at a single predicted risk score per patient. Risk scores were then median split into low risk and high risk. BRCA overall survival had concordance index of 0.667 ( $p < 0.005$ , log-rank test), and for SKCM, concordance index was 0.557 ( $p > 0.05$ , log-rank test). Please note that for experiments in the main text, low and high risk were defined as the first and fourth quartiles of risk scores, respectively.



Supplementary Figure 10: Distribution of the number of metastasis-containing patches in the CAMELYON16 test set. The test consists of 49 specimens with metastasis (plotted here) and 80 specimens negative for tumor. The positive specimens contained an average of 1320 patches that intersected the tumor annotation, where each patch was  $128 \times 128 \mu\text{m}$ . This distribution was heavily right-skewed. The median number of tumor patches was 102, and the 25th and 75th percentiles were 18 and 616, respectively.