

Implementing Trust in Non-Small Cell Lung Cancer Diagnosis with a Conformalized Uncertainty-Aware AI Framework in Whole-Slide Images

Xiaoge Zhang^{1,†,*}, Tao Wang^{1,†}, Chao Yan^{2,†}, Fedaa Najdawi³, Kai Zhou⁴, Yuan Ma⁵, Yiu-ming Cheung⁶, and Bradley A. Malin^{2,7,8,*}

¹Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Kowloon, Hong Kong

²Department of Biomedical Informatics, Vanderbilt University Medical Center, Nashville, TN, USA

³Department of Pathology, Microbiology and Immunology, Vanderbilt University Medical Center, Nashville, TN, USA

⁴Department of Computing, The Hong Kong Polytechnic University, Kowloon, Hong Kong

⁵Department of Mechanical Engineering and Research Institute for Intelligent Wearable Systems, The Hong Kong Polytechnic University, Kowloon, Hong Kong

⁶Department of Computer Science, Hong Kong Baptist University, Kowloon Tong, Kowloon, Hong Kong

⁷Department of Computer Science, Vanderbilt University, Nashville, TN, USA

⁸Department of Biostatistics, Vanderbilt University Medical Center, Nashville, TN, USA

[†]These authors contributed equally to this work.

*Corresponding author.

ABSTRACT

Ensuring trustworthiness is fundamental to the development of artificial intelligence (AI) that is considered societally responsible, particularly in cancer diagnostics, where a misdiagnosis can have dire consequences. Current digital pathology AI models lack systematic solutions to address trustworthiness concerns arising from model limitations and data discrepancies between model deployment and development environments. To address this issue, we developed TRUECAM, a framework designed to ensure both data and model trustworthiness in non-small cell lung cancer subtyping with whole-slide images. TRUECAM integrates 1) a spectral-normalized neural Gaussian process for identifying out-of-scope inputs and 2) an ambiguity-guided elimination of tiles to filter out highly ambiguous regions, addressing data trustworthiness, as well as 3) conformal prediction to ensure controlled error rates. We systematically evaluated the framework across multiple large-scale cancer datasets, leveraging both task-specific and foundation models, illustrate that an AI model wrapped with TRUECAM significantly outperforms models that lack such guidance, in terms of classification accuracy, robustness, interpretability, and data efficiency, while also achieving improvements in fairness. These findings highlight TRUECAM as a versatile wrapper framework for digital pathology AI models with diverse architectural designs, promoting their responsible and effective applications in real-world settings.

Main

In the field of digital pathology, artificial intelligence (AI) has shown significant potential for enhancing decision support across a wide spectrum of clinical tasks, from diagnosing and distinguishing between various types of cancers to detecting subtle histopathological changes that may predict disease progression [1–3]. However, the reliability of medical AI models can be significantly compromised by their inherent limitations, such as weak ability in measuring uncertainty and controlling error rates. These issues are further compounded by data discrepancies between model development and deployment environments, including variations in patient demographics, disease

NOTE: This preprint reports new research that has not been certified by peer review and should not be used to guide clinical practice.

characteristics, data acquisition techniques, staining protocols, and clinical practices. Addressing these challenges is essential to ensure model reliability, uphold patient safety, and establish trustworthiness in AI systems for medical applications [4, 5].

Rather than simply training AI models to produce the most likely outcome for patients as a point estimate—a practice that often lacks context and can lead to inaccurate interpretations—researchers have sought to quantify the confidence associated with AI generated results. This additional layer of insight can enable users to better interpret the reliability of a model’s prediction or classification, helping inform when to accept the results [6–8]. Multiple uncertainty quantification (UQ) approaches have been recently integrated into AI development to characterize confidence [9–12]. For instance, Olsson and colleagues [13] developed a conformalized deep convolutional neural network ensemble to generate all likely cancer diagnoses in biopsies and demonstrated the value of uncertainty estimates in flagging classifications with insufficient confidence. Additionally, Dolezal and colleagues [14] applied a Monte Carlo Dropout technique to quantify the uncertainty in classifications distinguishing between lung adenocarcinoma and squamous cell carcinoma, facilitating the recognition of high- versus low-confidence outputs to address domain shifts. And, most recently, Sun and colleagues [15] introduced TISSUE, a general framework for estimating uncertainty in spatial gene expression predictions that works by creating well-calibrated prediction intervals.

Despite advances in quantifying uncertainty, the current collection of approaches do not adequately align with several of the essential properties needed for safeguarding practical utilization of pathology AI, and medical AI more broadly [16]. First, an ideal UQ approach should consistently characterize a model’s confidence across varying levels of data risk. Beyond capturing uncertainty for in-domain (In-D) data, it should reliably identify out-of-domain (OOD) inputs, address potential distribution shifts between data used during model development and data encountered after deployment, and recognize regions within a slide that could adversely impact inference. At the present time, AI systems generally do not systematically address these challenges. Second, the claimed confidence interval (e.g., a set of likely cancer subtypes) or the error rate established for real-world inference must closely align with the observed outcomes. For example, if a model predicts with 95% confidence, then the true value should fall within the predicted range in 95% of the cases. Although existing methods, such as Platt [17–19] and temperature scaling [20, 21], seek to improve calibration, they neither ensure consistency between expected and observed confidence nor objectively represent the model’s confidence. Notably, these methods lack the ability to abstain from inputs that are difficult to infer accurately, a practice widely advocated for reliable medical AI with human in the loop [5]. Third, when estimating uncertainty, an effective UQ approach should enhance interpretability without adding substantial computational overhead beyond its deterministic counterpart (i.e., the original model without UQ). This is essential for medical imaging applications, where strict time constraints and resource limitations demand rapid and efficient processing without compromising reliability. Most UQ approaches rely on ensemble-based estimations, which require a large number of repeated inferences to be made using the same pathology AI models to estimate uncertainty for each patch (or tile) in a single slide. Given that these models typically contain millions of parameters, the resulting computational demand for generating uncertainty and interpretability can become prohibitively large, hindering their usability in practice [22].

In light of these issues, we set out to equip pathology AI with a formal, principled, and scalable UQ, as well as demonstrate how it empowers trustworthy subtyping of non-small-cell lung cancers (NSCLC) [23] using whole-slide images (WSI). We decompose the trustworthiness of medical AI into two constituent parts: 1) data trustworthiness, which ensures that the input data during deployment aligns with the model’s training scope and allows ambiguous slide regions to be excluded, and 2) model trustworthiness, which offers valid confidence intervals with a customizable level of coverage, ensuring the true label is covered in a specified proportion of classifications. To do so, we developed TRUECAM (Fig. 1a), a framework that provides **TR**ustworthiness-focused, **U**ncertainty-aware, **E**nd-to-end **C**ancer diagnosis with **M**odel-agnostic capabilities. TRUECAM consists of three components designed to simultaneously ensure data and model trustworthiness: 1) a spectral-normalized neural Gaussian process (SNGP) to establish informative data representation and uncertainty quantification, 2) an elimination of ambiguous tiles (EAT) mechanism for filtering out noisy patches from slides, and 3) conformal prediction (CP) to enable a statistically guaranteed error rate. We then conducted a wide range of experiments with several large cancer datasets, utilizing a widely adopted specialized model (i.e., Inception-v3 [24, 25]) and multiple

general-purpose foundation models (including UNI [26], CONCH [27], Prov-GigaPath [28], and TITAN [29]). Our results indicate that TRUECAM provides significant improvements in NSCLC subtyping accuracy and often enhances fairness in model performance across demographic groups without explicitly enforcing it into the model training process. Our analysis further show that TRUECAM is robust to OOD input and distribution shift, delivers more informative confidence intervals with statistical guarantee, enhances interpretability, and substantially reduces the computational inference burden. Together, these observations position TRUECAM as a general framework that can be integrated with medical AI systems of various sizes, architectures, purposes, and complexities to support their responsible and trustworthy application.

Results

An overview of TRUECAM

This study focuses on distinguishing between two main types of NSCLC—lung adenocarcinoma (LUAD) and lung squamous cell carcinoma (LUSC)—each of which is defined by distinct biological, morphological, molecular, and prognostic features that inform which therapeutic approach to undertake for a patient [23]. WSIs of lung tissue specimens were used for this specific task. In this paper, we use the term “prediction” interchangeably with “classification” where needed to align with the established terminology convention in conformal prediction.

We designed the TRUECAM framework (Fig. 1a) to wrap around deep learning models of various architectures, sizes, and purposes for imaging tasks (Fig. 1b). TRUECAM employs SNGP to perform distance-preserving transformations on input data and efficiently estimate data uncertainty, enhancing data trustworthiness by enabling OOD detection and data shift control before model inference. Furthermore, CP is applied on top of SNGP to calibrate model predictions (represented as a set of values, i.e., prediction set), providing a statistical guarantee on the model’s coverage of true labels. These components contribute complementary functions to assure the trustworthiness of AI for NSCLC subtyping. Specifically, SNGP establishes a valid foundation for CP to operate. This is important because CP generally requires that data in the deployment environment comes from the same underlying distribution as those in the development environment, known as data exchangeability—an assumption that might not hold true in the real-world setting, where OOD data and distribution shift are common. Meanwhile, CP calibrates the outputs of SNGP to ensure the statistical validity of model coverage. Notably, the distance-preserving data transformation by SNGP facilitates an effective identification of tiles that only add ambiguity to NSCLC subtyping. EAT can then be applied to enhance the supervisory signal—leading to a more reliable model—while significantly reducing the computational load during inference (Fig. 1c). The details of model design and implementation are described in the Methods.

In this study, we relied on two digital pathology datasets of NSCLC stained with hematoxylin and eosin (H&E) (Fig. 1d), both of which are large-scale, multi-institutional initiatives in the United States: 1) 941 WSIs from The Cancer Genome Atlas (TCGA), and 2) 1,306 WSIs from the Clinical Proteomics Tumor Analysis Consortium (CPTAC). Additionally, we constructed an OOD dataset from the same source as TCGA, which incorporates 631 WSIs of non-lung cancers. For simplicity, we refer to these three datasets as TCGA, CPTAC, and TCGA-OOD. The details for these datasets are provided in the Methods.

To assess the effectiveness of TRUECAM, we evaluated two types of AI models. We utilized Inception-v3, a widely used convolutional neural network architecture for image inference [24, 25], as a representative classifier for specialized medical imaging tasks within the TRUECAM framework. In this context, we used TCGA for model training and internal validation, while we used CPTAC as an independent external validation dataset to assess TRUECAM’s effectiveness in scenarios requiring model transferability. We further investigated four digital pathology foundation models that exemplify the recent advancements in large, general-purpose models in the field: 1) UNI [26], 2) CONCH [27], 3) Prov-GigaPath [28], and 4) TITAN [29]. Since foundation models are pretrained on diverse datasets, offering strong knowledge transferability as an image encoder, and downstream task-specific classifiers are typically lightweight and easy to train, our evaluation focused solely on the setting where each healthcare organization trains and evaluates its own site-specific model using its data. Accordingly, TCGA and CPTAC were utilized to mimic this setting, such that we

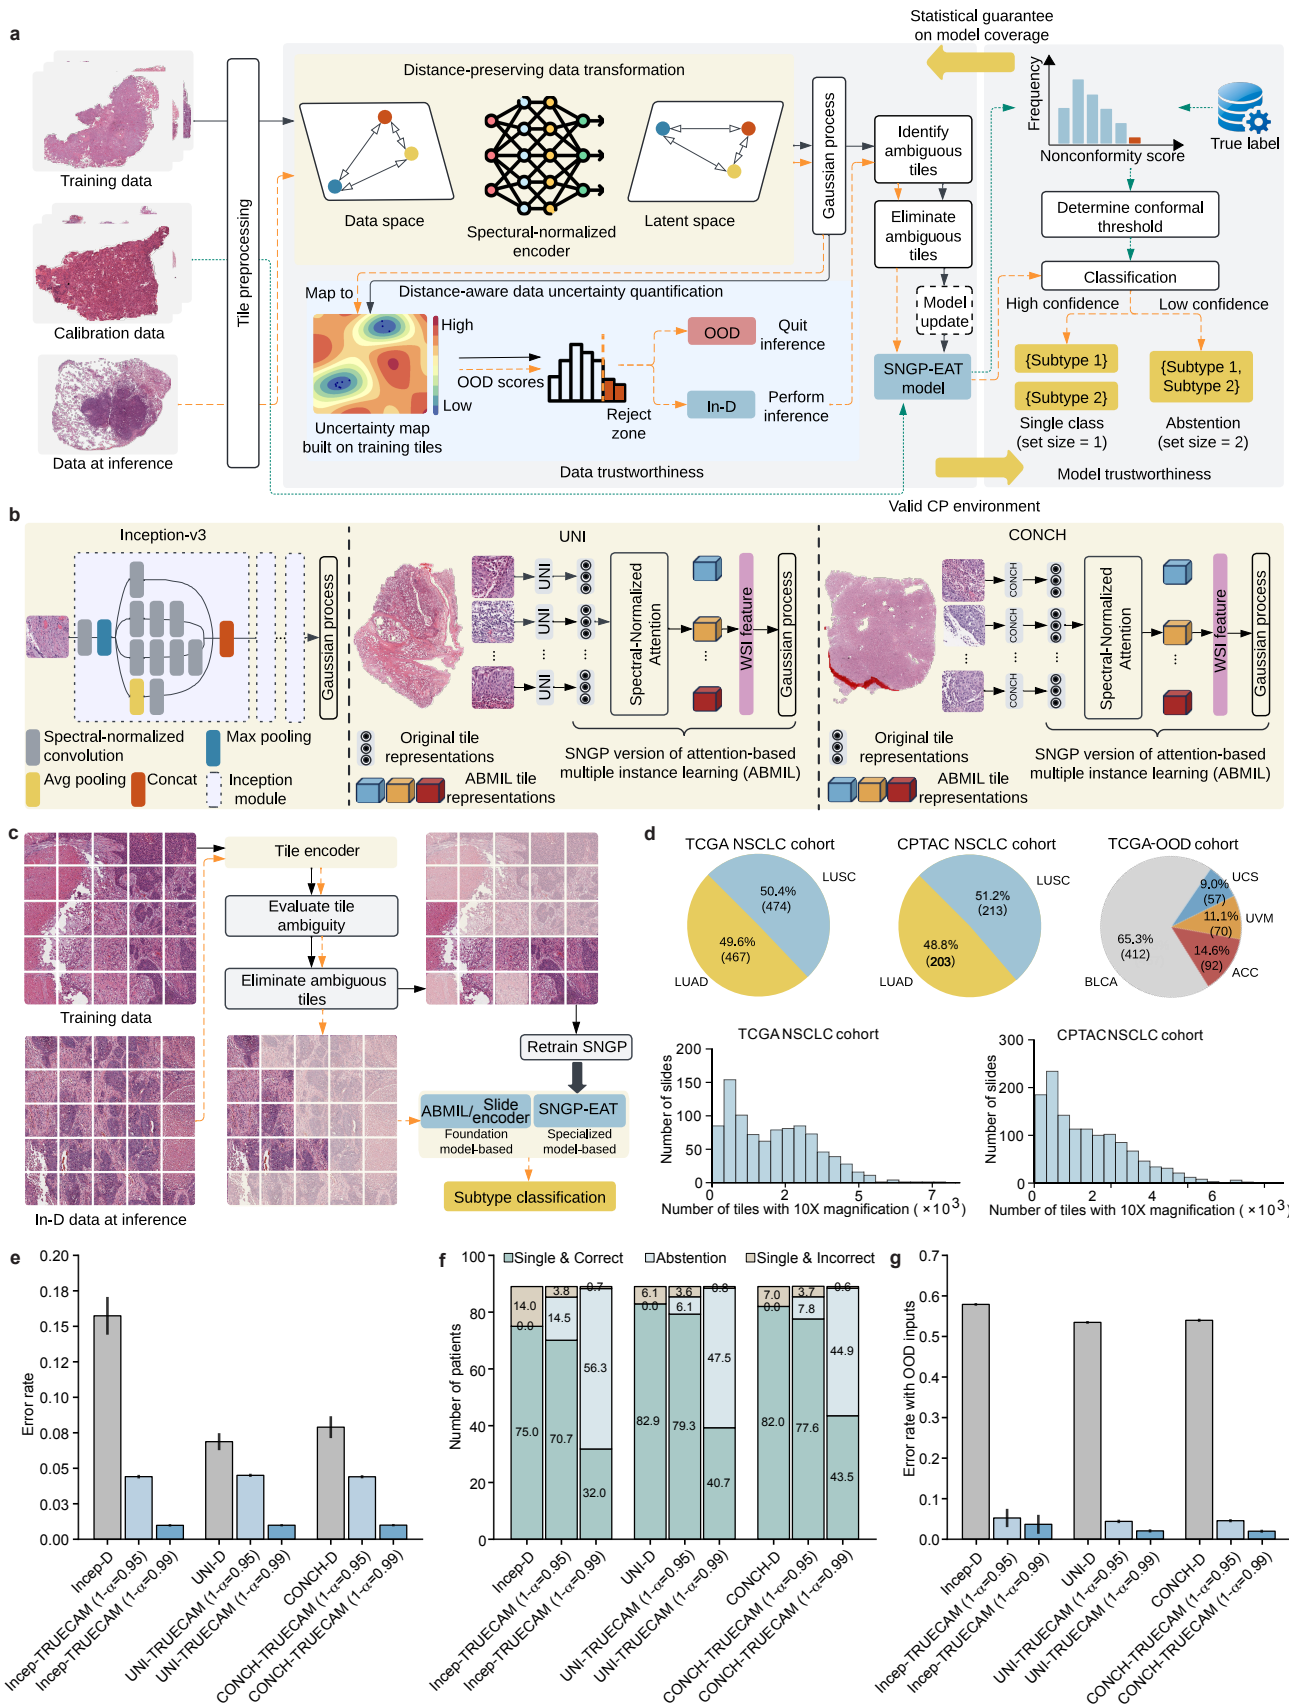


Figure 1. Overview of TRUECAM. TRUECAM is a versatile, model-agnostic digital pathology AI framework for reliable non-small-cell lung cancer (NSCLC) subtyping, achieving trustworthiness by substantially reducing errors, detecting OOD data and controlling distribution shifts pre-inference, identifying and eliminating ambiguous slide regions, and ensuring true label coverage with statistical guarantees via abstention on uncertain inputs. **a**, The architecture of TRUECAM designed to ensure both data and model trustworthiness. **b**, Customization illustration of TRUECAM for deep learning models of varying architecture, complexity, and purpose, including Inception-v3, UNI, and CONCH. **c**, Illustration of eliminating ambiguous tiles in TRUECAM for slide inference. **d**, Overview of the TCGA and CPTAC NSCLC datasets, as well as the out-of-domain dataset built from other cancer types within TCGA. **e**, TRUECAM significantly reduced NSCLC subtyping error rates across all model types (denoted using suffix “-TRUECAM”) compared to their original deterministic versions (denoted using suffix “-D”), adhering to the pre-specified true label coverage $1-\alpha$. **f**, Patient-level classification breakdown for models with and without TRUECAM. **g**, TRUECAM achieved significantly lower error rates in real-world NSCLC subtyping scenarios involving a 1:1 mix of in-domain and out-of-domain inputs. Results shown in **e-g** are based on the TCGA dataset. See Extended Data Fig. 1 for evaluations based on the CPTAC dataset and two other foundation models (Prov-GigaPath and TITAN). All mean values and 95% confidence intervals are based on 20 independently trained models, each with 500 conformal prediction evaluations. OOD, out-of-domain; In-D, in-domain; SNGP, spectral-normalized neural Gaussian process; EAT, elimination of ambiguous tiles; WSI, whole-slide image; LUAD, lung adenocarcinoma; LUSC, lung squamous cell carcinoma; BLCA, bladder urothelial carcinoma; USC, uterine carcinosarcoma; UVM, uveal melanoma; ACC, adrenocortical carcinoma; Incep, Inception-v3; D, Deterministic.

assumed each dataset represents a distinct healthcare organization, independently supporting the required model training and validation. Our evaluation primarily focused on UNI and CONCH, with additional results provided for Prov-GigaPath and TITAN to showcase the generalizability of TRUECAM’s advantages.

Our evaluation yielded several notable findings. First, TRUECAM significantly reduced NSCLC subtyping error rates (defined as the fraction of incorrect patient-level classifications) across all three models—Inception-v3, UNI, and CONCH—compared to their original versions, denoted as Deterministic (or D) (Fig. 1e). Defining α as the desired maximum error level (or significance level), applying TRUECAM to Inception-v3 (referred to as Incep-TRUECAM) with coverage of $1-\alpha=0.95$ and $1-\alpha=0.99$ resulted in error rate reductions by 72.0% and 93.8%, respectively. Notably, these empirical error rates closely complemented the desired coverage levels that reflect adaptable error tolerance for individual NSCLC subtyping services. Second, TRUECAM demonstrated an ability to identify and abstain from reasoning about challenging inputs, allowing such inputs to be deferred to pathologists. As shown in Fig. 1f, with coverage set at $1-\alpha=0.95$, Incep-TRUECAM abstained on approximately 15 patients, reducing the number of misclassified patients from 14 to fewer than 4. Third, the foundation model-based classifiers exhibited more promising results than Inception-v3, with inherently lower error rates and fewer abstentions to achieve the same desired coverage (Fig. 1e,f). Fourth, in simulated real-world settings where OOD slides may be unknowingly submitted for NSCLC subtyping, directly deploying the original versions of the considered models often incorrectly classified non-lung cancer slides as LUAD or LUSC. By contrast, TRUECAM reliably detected these OOD inputs prior to model inference, while maintaining the same error rates as the scenarios involving only In-D inputs (Fig. 1g). We observed similar findings in the evaluation of Prov-GigaPath and TITAN (Extended Data Fig. 1).

The following sections report on a systematic evaluation of TRUECAM’s performance in various dimensions and the effectiveness of its core modules, using Inception-v3 as the primary example. We then extend the analysis to foundation models UNI and CONCH.

Integration of distance-aware uncertainty estimation improves NSCLC subtyping performance

We trained three deep neural network models, all based on the Inception-v3 architecture, to discriminate between LUSC and LUAD: 1) a deterministic model without UQ, referred to as Deterministic, which represents the original Inception-v3, 2) a Monte Carlo Dropout-based model (referred to as MC Dropout) to enable UQ [14], and 3) a model that utilizes SNGP to quantify uncertainty (referred to as SNGP). We provide model details in the Methods. Following the design in [14], all models were trained at the tile level, with slide-level classifications obtained by averaging tile-level outputs across each WSI. For patients with multiple WSIs, the final patient-level diagnosis was derived by averaging the slide-level classifications across all WSIs for that patient.

We deactivated TRUECAM’s CP functionality and EAT to assess the effectiveness of SNGP. At the tile level (using tiles as subtyping units), SNGP consistently outperformed Deterministic on the TCGA dataset in both accuracy ($p < 0.001$) and the area under receiver operator curves (AUROC) ($p < 0.001$) (Fig. 2a). In addition, when evaluated on TCGA, SNGP significantly outperformed MC Dropout in terms of both accuracy ($p < 0.01$) and AUROC ($p < 0.001$). When evaluated on the external CPTAC dataset, SNGP

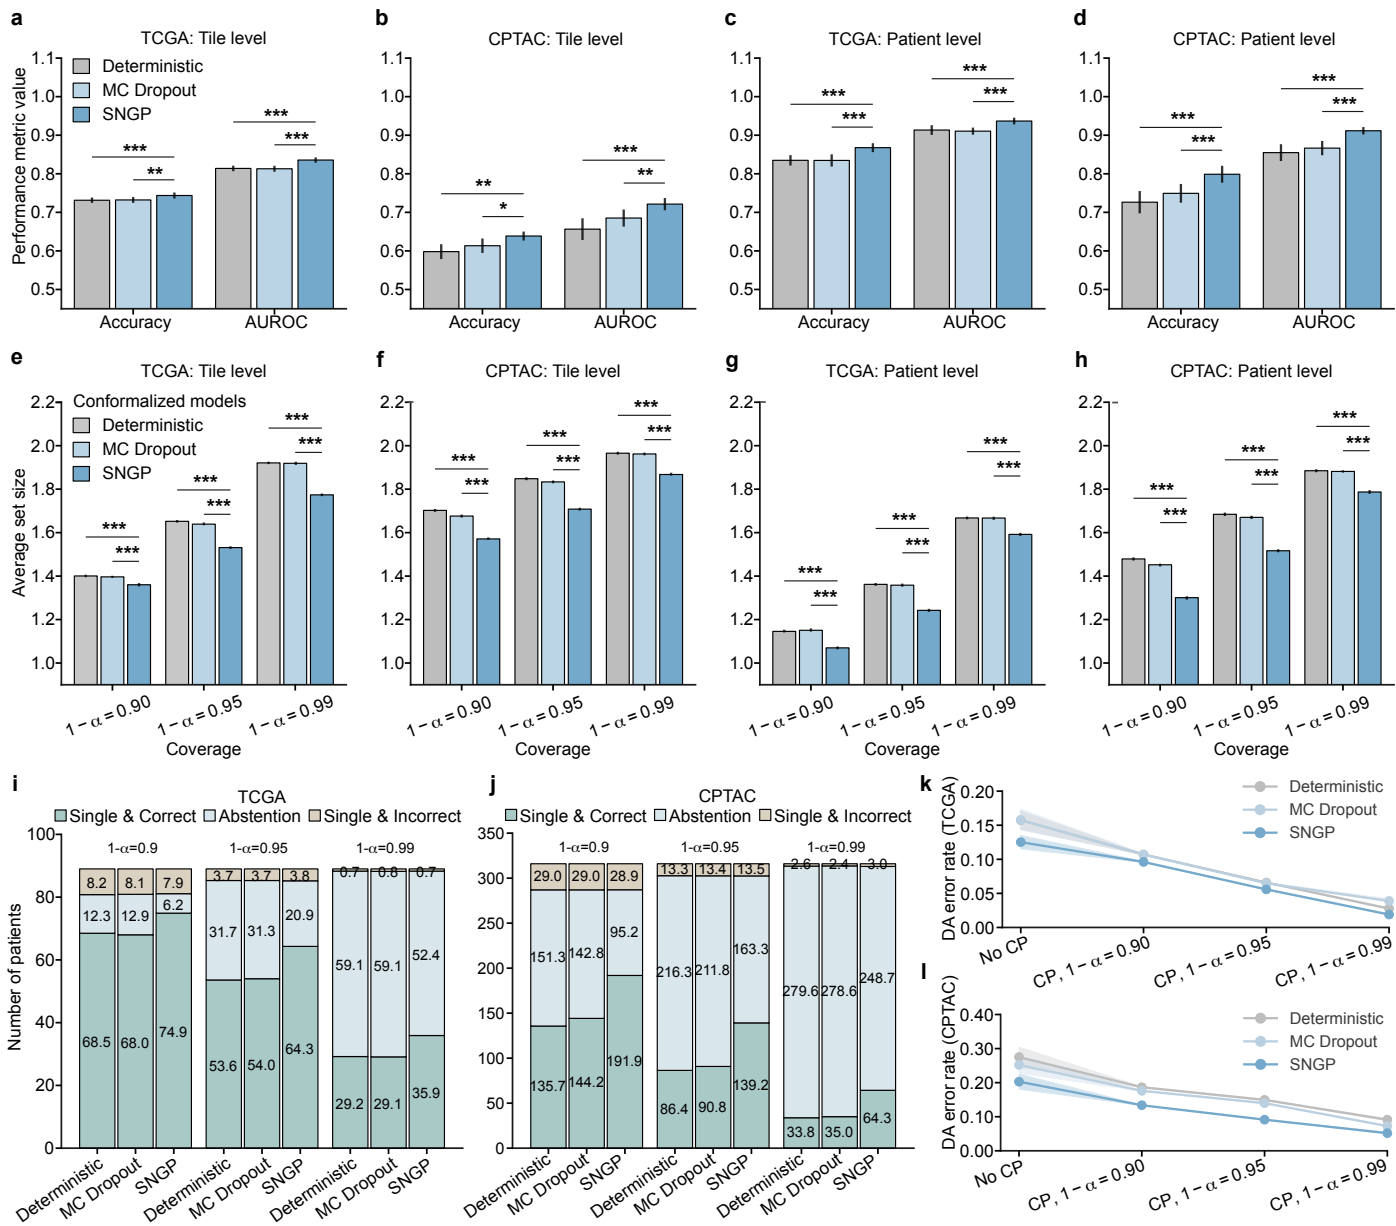


Figure 2. NSCLC subtyping performance of three Inception-v3-based deep neural network models and their conformalized counterparts. **a**, Tile-level performance with respect to the TCGA dataset in terms of classification accuracy and area under receiver operator curve (AUROC). **b**, Tile-level performance evaluated on the CPTAC dataset. **c**, Patient-level performance evaluated on the TCGA dataset. **d**, Patient-level performance evaluated on the CPTAC dataset. **e**, Tile-level prediction set size on TCGA for three distinct values of error level α . **f**, Tile-level prediction set size on CPTAC. **g**, Patient-level prediction set size on TCGA. **h**, Patient-level prediction set size on CPTAC. **i**, Patient-level classification breakdown for three conformalized models on the TCGA testing dataset. **j**, Patient-level classification breakdown on CPTAC. **k**, Patient-level DA error rate on TCGA before and after activating CP. **l**, Patient-level DA error rate on CPTAC before and after activating CP. Mean values and the corresponding 95% confidence intervals in **a-d** are derived from 20 independently trained models, evaluated using a combination of calibration and testing data. **e-l** are based on 20 independently trained models, each with 500 CP evaluations (See details in the Methods). One-sided Wilcoxon signed-rank test is utilized to calculate p values. *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. SNGP, spectral-normalized neural Gaussian process; CP, conformal prediction; DA, definitive-answer.

exhibited similar advantages over Deterministic and MC Dropout (Fig. 2b). At the patient level (using patients as subtyping units), the performance improvement enabled by SNGP was more pronounced. In this situation, Deterministic achieved an average accuracy of 0.843 and AUROC of 0.926 for TCGA, whereas SNGP achieved an accuracy of 0.875 and AUROC of 0.950 (Fig. 2c). Moreover, integrating SNGP into Inception-v3 led to a significant performance advantage over MC Dropout (accuracy: $p < 0.001$; AUROC: $p < 0.001$). Further, the performance gain enabled by SNGP generalized well to the external CPTAC dataset. Specifically, SNGP

outperformed MC Dropout by 4.9% ($p < 0.001$) and Deterministic by 7.2% ($p < 0.001$) in accuracy, and exceeded MC Dropout by 4.5% ($p < 0.001$) and Deterministic by 5.7% ($p < 0.001$) in AUROC (Fig. 2d).

Conformalized SNGP establishes statistical coverage guarantee while enhancing NSCLC subtyping efficiency

We enabled CP in TRUECAM and then assessed its effectiveness in calibrating SNGP-based NSCLC subtyping by generating a prediction set for each patient. A set size of 2 indicates that the model lacks sufficient confidence to definitively classify input data as either LUAD or LUSC, effectively resulting in abstention. Two broadly used measures in the literature, validity and efficiency [30–32] were evaluated. We consider a model to be valid at an error level α when the proportion of the prediction set containing the true subtype (i.e., model coverage) is at least $1 - \alpha$. Efficiency is determined by the size of the prediction sets—we deem a model to be more efficient when it produces smaller prediction sets, provided that validity is maintained. This is because more concise and informative outputs are generally preferred by both service providers and patients.

There are several findings we wish to highlight. First, all three of the models achieved the validity requirements at both the tile and patient levels across the TCGA and CPTAC testing datasets for three distinct error levels (Supplementary Tables 9, 10). Second, for each considered coverage $1 - \alpha$, conformalized SNGP consistently produced significantly smaller prediction sets than conformalized MC Dropout and Deterministic at both the tile and patient levels (Fig. 2e,g). Such an advantage generalized well to the external CPTAC dataset at both the tile and patient levels (Fig. 2f,h). More concretely, at $\alpha = 0.05$, conformalized SNGP produced 41,570 more tiles with a single NSCLC subtype output (set size of one) on average, which was a 31.4% increase over conformalized MC Dropout. In other words, when validity was ensured, conformalized SNGP exhibited a significantly higher efficiency in terms of the number of informative outputs in NSCLC subtyping than other models (Supplementary Tables 11,12). Third, as expected, lowering the value of α , which decreases the tolerance for model miscoverage, resulted in larger prediction sets across all models, datasets, and classification levels (Fig. 2e-h). In other words, more tiles and patients were classified as “unsure” (i.e., abstention) to reduce the likelihood of classifications containing only a single, but incorrect, subtype.

We further investigated the models’ patient-level classifications, which include the following categories 1) single and correct, 2) single but incorrect, and 3) uncertain between subtypes (abstention). For all considered error levels, conformalized MC Dropout achieved similar patient-level performance as Deterministic in terms of patient numbers under each category (Fig. 2i,j). By contrast, conformalized SNGP consistently produced more single and correct classifications in both of the TCGA and CPTAC testing datasets than the other two models, while incurring a similar number of single but incorrect classifications (Fig. 2i,j). Specifically, at $\alpha = 0.01$, conformalized SNGP classified approximately six more patients with a single and correct NSCLC subtype than conformalized MC Dropout, reflecting a 23.4% increase (Fig. 2i). This advantage was further amplified in the CPTAC testing dataset, where it classified around 30 more patients with single and correct subtype at $\alpha = 0.01$, corresponding to a 83.7% increase than conformalized MC Dropout (Fig. 2j). As an artifact of bounded error rates, the increase in correctly identified patients by conformalized SNGP was accompanied by a comparable reduction in classifications with abstention. These findings collectively highlight enhanced overall performance and greater efficiency of TRUECAM, as evidenced by the increased ratio of single and correct classifications, underscoring its superior practical utility.

Next, we examined the capability of TRUECAM to manage errors when committing to a single subtype prediction. We define the *definitive-answer error rate* (referred to as DA error rate) as the proportion of patients with single but incorrect subtype designation among all patients with single subtype classification. Before activating CP, TRUECAM demonstrated the lowest error rate for both TCGA and CPTAC (Fig. 2k,l). In other words, when the model outputs were not conformalized (i.e., classifications were binary, based solely on the highest probability between subtypes), SNGP outperformed both MC Dropout and Deterministic, a finding consistent with the results shown in Fig. 2c,d. When CP allows a model to abstain on uncertain inputs, all models achieved a reduction in DA error rate, with a steady decline as the significance level α decreases. Notably, at $\alpha = 0.01$, CP induced a DA error rate of 1.9% in TCGA, representing an 84.8% reduction compared to using SNGP alone. In other words, SNGP, without CP, made approximately one error for every eight patients, whereas with CP at $\alpha = 0.01$, SNGP made only one error for every 100 patients on average. These

findings, along with similar trends observed across other models, demonstrated that CP provided a significant advantage by universally reducing the DA error rate across various models. Additionally, for all considered coverage values $1 - \alpha$, conformalized SNGP exhibited the lowest DA error rates, indicating that TRUECAM delivered more accurate classifications than other models.

Eliminating ambiguous tiles (EAT) concurrently augments classification and CP performance in NSCLC subtyping

In clinical inference using a WSI, not all regions provide diagnostic value to pathologists (e.g., areas with normal, non-cancerous cells). However, the current practice in digital pathology typically involves utilizing all tiles—often numbering in the hundreds to tens of thousands from a single WSI (Fig. 1d)—paired with coarse-grained slide-level labels for training diagnostic models [33–37]. Such a weakly supervised learning process inherently introduces noise, as many normal tiles may be inaccurately associated with tumor labels [38]. This practice can lead to reliability issues and inefficiencies as models are forced to learn from irrelevant or non-informative regions (even with attention mechanisms in place). These, in turn, dilute the signal necessary for accurate and efficient NSCLC subtyping, which undermines data trustworthiness.

To filter out non-informative tiles for the NSCLC subtyping task, we applied k -means clustering to the tile representations of the TCGA training data extracted by SNGP. We calculated the Silhouette coefficient [39] and determined $k = 3$ as the optimal number of clusters. We observed that one cluster contained mostly tiles from WSIs with the true label of LUAD, another cluster contained mostly tiles with the true label of LUSC, and the final cluster was not dominated by either subtype (Fig. 3a). Using t-distributed stochastic neighbor embedding (t-SNE) visualization (Fig. 3b), we observed that the clustering pattern consistently aligned with the label distribution depicted in Fig. 3a.

Subsequently, we measured the ambiguity of all of the tiles to quantify model discernibility (formally defined in the Methods) and observed that the ambiguity scores were significantly higher in the cluster lacking clear label dominance (referred to as ambiguous cluster) than the other two clusters (Fig. 3c). This suggests that the tiles in the ambiguous cluster may only add ambiguity to NSCLC subtyping when aggregating tile-level classifications. Motivated by these observations, we hypothesized that eliminating ambiguous tiles through the EAT process would strengthen the supervisory signals for model training and lead to more accurate subtyping of NSCLC. To test this hypothesis, we compared the classification and CP performance of TRUECAM before and after activating EAT. To do so, we removed the training tiles from the ambiguous cluster (accounting for 66.7% of all training tiles) and used the remainder of the data to train a new classification model (Fig. 1c), which we refer to as SNGP-EAT. For new WSI inference, tiles from a new slide with their latent representations falling within the ambiguous cluster were excluded from the NSCLC subtyping task using SNGP-EAT.

In addition to comparing SNGP-EAT with SNGP, we established another baseline model, referred to as SNGP-RE, which follows the same procedure as SNGP-EAT but eliminates an equal number of tiles that are selected uniformly at random. SNGP-EAT achieved a 2.83% improvement in patient-level classification accuracy compared to SNGP ($p < 0.001$) on TCGA and an 8.05% increase ($p < 0.001$) on CPTAC (Fig. 3d, Supplementary Table 8). By contrast, SNGP-RE consistently performed worse than both SNGP and SNGP-EAT in terms of accuracy and AUROC (Extended Data Fig. 2a-d). Moreover, EAT improved CP efficiency by producing significantly smaller prediction set sizes (Extended Data Fig. 2e-h, Supplementary Tables 11, 12) and generated more single and correct classifications in both the TCGA and CPTAC cohorts (Fig. 3e,f), except for the TCGA setting with $\alpha = 0.01$, while yielding a comparable number of single but incorrect classifications as the models without EAT. Meanwhile, a comparison of DA error rates also favored SNGP-EAT (Extended Data Fig. 2i,j, Supplementary Tables 13,14).

TRUECAM achieves fairer NSCLC diagnosis compared to other methods

We investigated how TRUECAM affects fairness by comparing it to a set of baseline models. We first evaluated the classification performance gap, defined as the maximum difference in accuracy among patient subgroups based on race and sex (Supplementary Table 1). This was assessed in scenarios without CP. We then extended this evaluation to fairness in terms of average set sizes of CP across these subgroups. A large difference in set sizes between two subgroups suggests greater uncertainty in the model's predictions for one subgroup compared to the other.

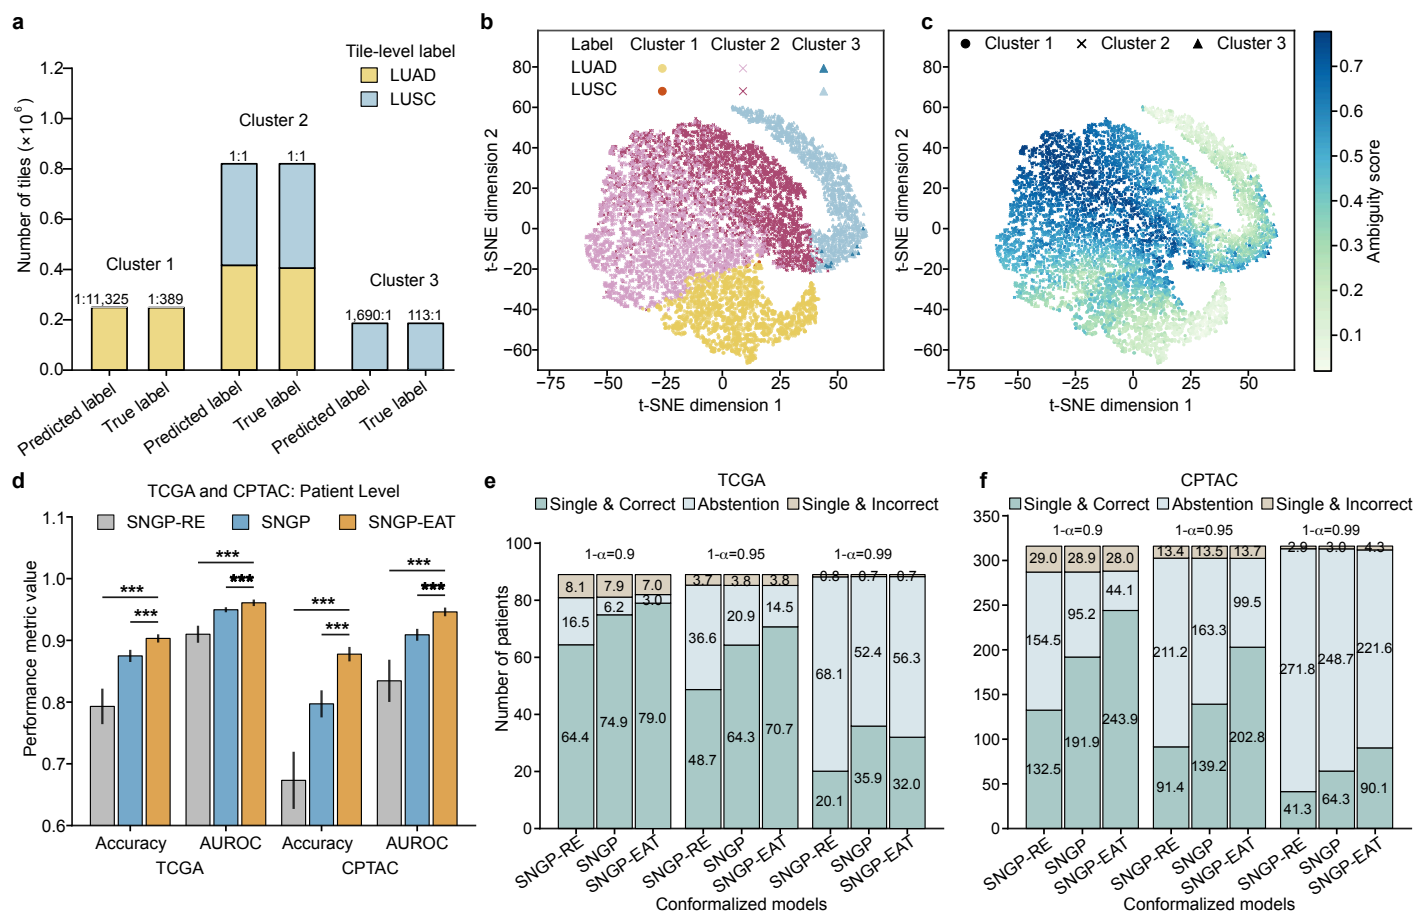


Figure 3. Evaluation of eliminating ambiguous tiles (EAT). **a**, The count and ratio of tiles predicted by SNGP as LUSC versus LUAD within each tile cluster identified using k -means on SNGP-based tile representations, along with the corresponding true labels derived from slide-level diagnosis. **b**, T-distributed stochastic neighbor embedding (t-SNE)-based illustration of tile clusters and true labels. **c**, Tile-wise ambiguity scores on top of the t-SNE visualization. For the sake of illustration, we randomly sampled 10,000 tiles to visualize the t-SNE distributions in **b** and **c**. **d**, Patient-level classification performance on the TCGA and CPTAC testing datasets. **e**, Patient-level classification breakdown of three conformed SNGP-based models on TCGA. **f**, Patient-level classification breakdown on CPTAC. Mean values and the corresponding 95% confidence intervals in **d-f** are based on 20 independently trained models, each with 500 CP evaluations (See details in the Methods). One-sided Wilcoxon signed-rank test is utilized to calculate p values. *** $p < 0.001$. LUAD, lung adenocarcinoma; LUSC, lung squamous cell carcinoma; SNGP, spectral-normalized neural Gaussian process; RE, random elimination; EAT, elimination of ambiguous tiles.

SNGP and SNGP-EAT consistently exhibited higher overall classification accuracy than Deterministic and MC Dropout, while also improving fairness in accuracy across racial and sexual subgroups when Deterministic's fairness gap was relatively high or preserving fairness when the gap was relatively low (Fig. 4a-d). Notably, SNGP-EAT achieved the highest overall accuracy on both of the internal and external datasets and exhibited the smallest fairness gaps in accuracy across racial subgroups. Specifically, SNGP-EAT led to a 38.1% and 78.3% reduction in fairness gaps across racial groups in TCGA and CPTAC, respectively, compared to Deterministic (Supplementary Tables 23-24). By contrast, randomly removing tiles through SNGP-RE did not result in a significant improvement in overall accuracy and fairness, again highlighting the benefit of EAT.

The advantages of conformed SNGP and SNGP-EAT were more pronounced for fairness in the average set sizes of CP (Fig. 4e-h). Remarkably, conformed SNGP-EAT (i.e., TRUECAM), consistently outperformed all other models by a large margin in terms of the average set size gap across sexual and racial subgroups. For example, at $1 - \alpha = 0.90$, compared to conformed Deterministic, conformed SNGP-EAT achieved a reduction of 42.2% in the set size gap across racial subgroups on TCGA and reduces such gap by 27.4% on CPTAC (Figs. 4f,h, Supplementary Tables 25-26). With respect to sex, conformed SNGP-EAT reduced the fairness gap for TCGA and CPTAC by 51.6% and 73.0%, respectively, approaching a near-perfect value (Fig. 4e,g,

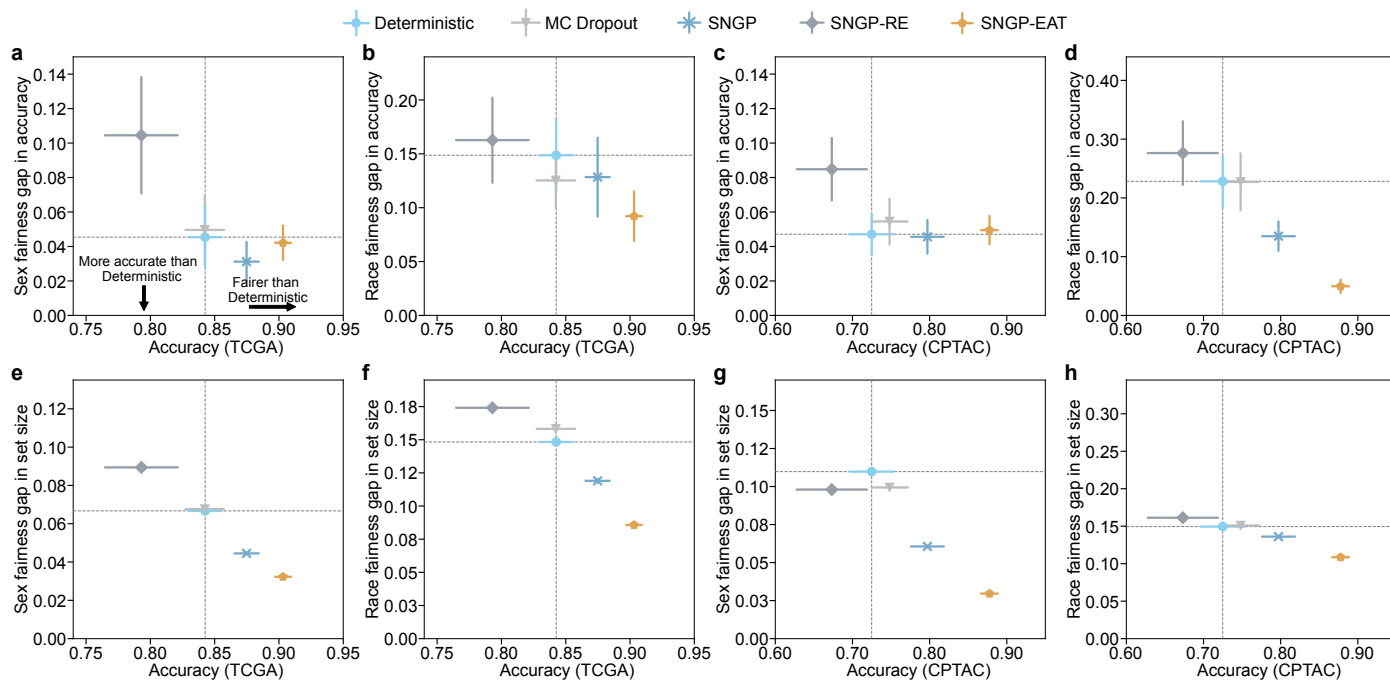


Figure 4. Fairness evaluation of TRUECAM based on metric value differences between the best- and worst-performing demographic subgroups against overall accuracy of NSCLC subtyping. **a-d**, Sex- and race-wise accuracy gap versus overall accuracy for discriminating between LUAD and LUSC on TCGA (n=189) and CPTAC (n=416) with CP deactivated. **e-h**, At $\alpha = 0.1$, sex- and race-wise set size gap versus overall accuracy on TCGA and CPTAC with CP activated. Evaluations on Deterministic or its conformalized version are marked as an overall baseline. Additional results are provided in Supplementary Tables 21-28.

Supplementary Tables 27-28). These improvements were accompanied by significant gains in the overall model accuracy.

TRUECAM enables effective OOD detection and distribution shift control

In pathology AI deployment, WSIs that differ from the model’s training data in aspects such as cohorts, tissue types, devices, and other factors are prevalent and difficult to detect. For models equipped with CP, the claimed model coverage may no longer hold valid under such condition. This potentially undermines reliability and poses a serious risk to patient health. We now show that TRUECAM is able to safeguard CP by detecting data that is outside of the model’s scope (i.e., OOD data).

We built the TCGA-OOD dataset that contains 698 WSIs from 631 patients with non-lung cancers (See Supplementary Table 1 and Methods for details), and then combined it with an In-D dataset, which integrates the TCGA calibration and testing datasets. We defined two distinct measures to quantify the degree to which an input WSI is out of the model’s scope: 1) an uncertainty-based OOD score, which aggregates tile-level uncertainty as determined by the model’s UQ module (e.g., Gaussian layer for SNGP and SNGP-EAT, stochastic dropout layer for MC Dropout), and 2) a probability-based OOD score, which aggregates tile-level classification probabilities, calculated as $1 - \frac{1}{N} \sum_{i=1}^N \max_k p(\hat{y}_k | \mathbf{x}_i)$, where $p(\hat{y}_k | \mathbf{x}_i)$ denotes the probability of tile i being classified as label \hat{y}_k . We observed that probability-based OOD scores from SNGP-EAT exhibited the strongest ability to distinguish between In-D and OOD WSIs, achieving the highest AUROC of 0.949, significantly outperforming Deterministic (0.876), MC Dropout (0.884), and SNGP (0.898) (Fig. 5a). At a 0.95 precision, SNGP-EAT achieved a true positive rate (TPR, i.e., recall) of 0.929, which is 46.0%, 41.5%, and 33.3% higher than Deterministic, MC Dropout, and SNGP, respectively (Fig. 5b). This observation also generalized to uncertainty-based OOD scores. Notably, for both Inception-v3-based SNGP models, probability-based OOD scores were more effective than uncertainty-based scores for identifying OOD inputs (Supplementary Table 15). As such, the probability-based OOD score was relied upon in the subsequent investigations related to Inception-v3-based models.

We further examined the effect of identifying and removing OOD inputs from inference enabled by TRUECAM on the empirical coverage of a model (i.e., model validity). When we randomly sampled an equal number of OOD patients to match the size of the In-D

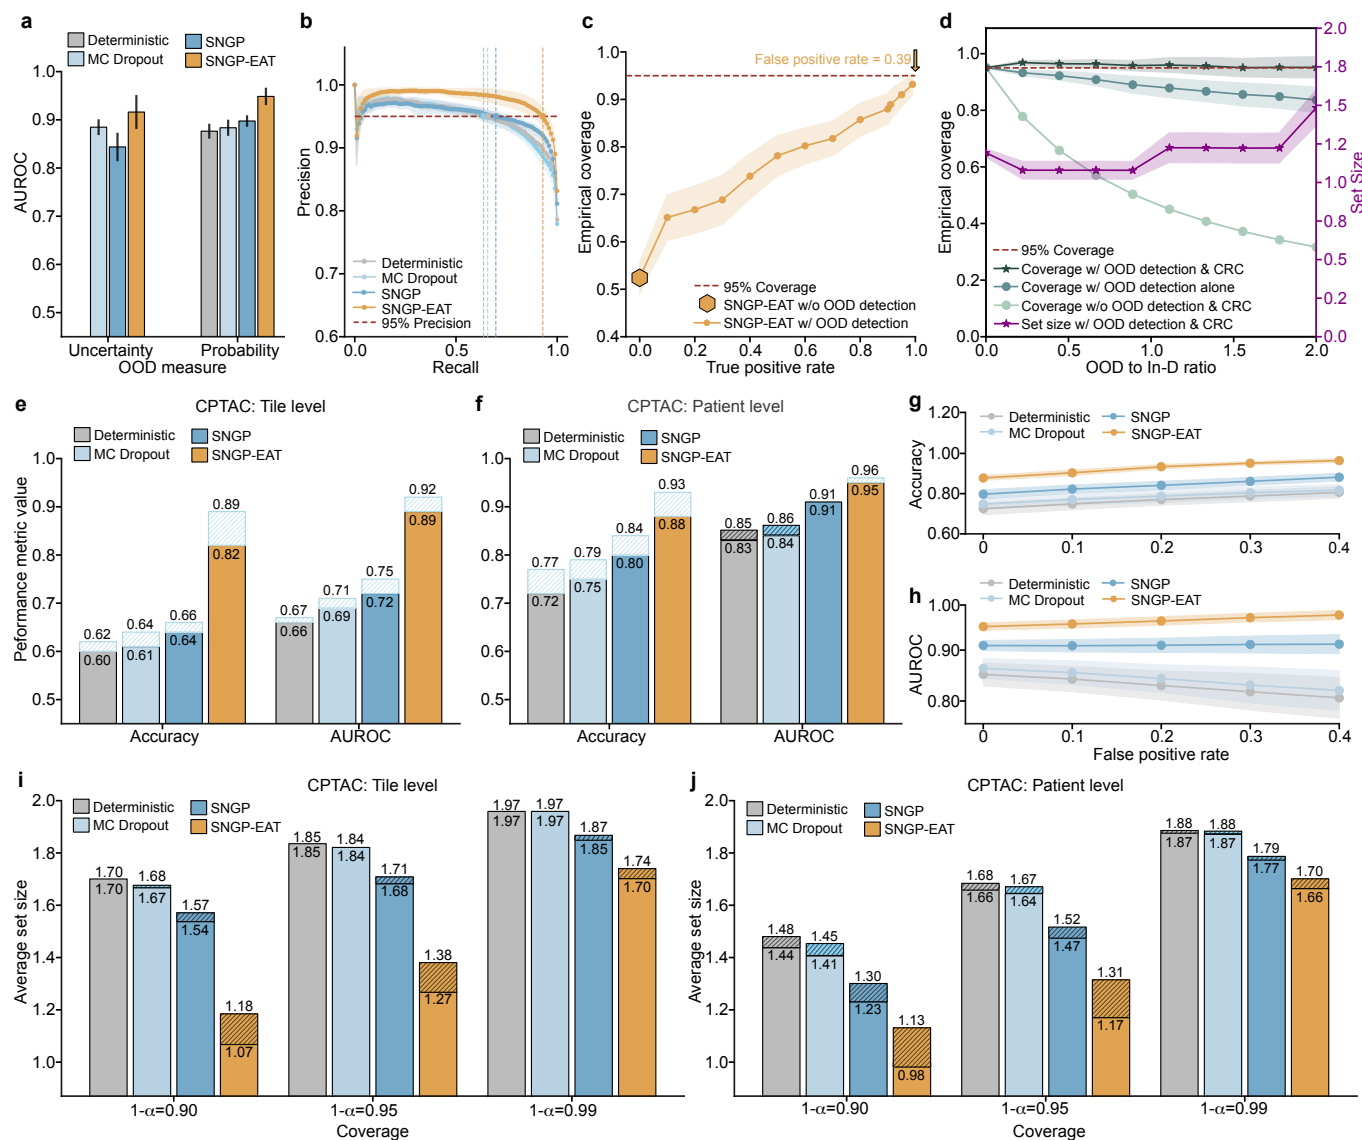


Figure 5. Evaluation of OOD detection enabled by TRUECAM. **a**, Performance of two distinct OOD scores in identifying slides outside a model's scope, using the top $\delta=200$ tiles in a slide with the lowest tile-level values for patient-level OOD score calculation. The bar for Deterministic is not shown as it does not measure uncertainty. See Supplementary Table 16 for sensitivity analyses for varying δ values. **b**, Precision-recall curves for OOD detection using the probability-based OOD score. **c**, Empirical model coverage with $\alpha = 0.05$ as a function of the true positive rate for OOD detection. **d**, Empirical coverage and prediction set size of CP for TRUECAM, compared to the scenario where conformal risk control and OOD detection were deactivated, shown as a function of the ratio of OOD to In-D data. **e**, Impact assessment of distribution shift control on the tile-level classification performance. Striped boxes in blue and black indicate performance increase and decrease, respectively, compared to the scenario without distribution shift control. **f**, Impact assessment of distribution shift control on the patient-level classification performance. **g,h**, Accuracy and area under the receiver operating characteristic (AUROC) for evaluation on CPTAC as a function of false positive rate in OOD detection. **i**, Impact assessment of distribution shift control on the tile-level average set size. Striped boxes in black indicate improvement in set size, compared to the scenario without distribution shift control. **j**, Impact assessment of distribution shift control on the patient-level average set size. SNGP, spectral-normalized neural Gaussian process; EAT, elimination of ambiguous tiles; CRC, conformal risk control.

dataset, the empirical coverage of SNGP-EAT approached the intended value of 0.95 ($\alpha=0.05$) as the true positive rate (TPR) of OOD detection neared 1 (Fig. 5c). This effectively restored data exchangeability. By contrast, without an OOD detection mechanism, the empirical coverage remained at 0.478, which fell significantly short of the target and signaled a failure in the model's reliability.

However, even with the best-performing model and OOD score, identifying 99.0% OOD inputs came at the cost of misclassifying 39% of In-D data as OOD (Fig. 5c), excluding them from model inference. Furthermore, while the OOD threshold was set to achieve TPR = 1 in testing, this level of coverage cannot be guaranteed in deployment, where unexpected OOD data may extend beyond the

threshold, thus compromising model coverage and trustworthiness. To address these challenges, TRUECAM pairs conformal risk control (CRC) with OOD detection, aiming to provide a more reliable model coverage even in the presence of undetected OOD data. CRC adjusts model outputs dynamically by accounting for the risks posed by OOD inputs that escape detection, thereby maintaining valid empirical coverage (see Methods for details).

We assessed TRUECAM's effectiveness in handling OOD data by simulating a realistic environment where the ratio of OOD to In-D data varied. Our findings indicated that combining OOD detection with CRC in TRUECAM resulted in a robust empirical coverage that closely matched the target of 0.95 ($\alpha=0.05$) across a broad range of OOD-to-In-D ratios (Fig. 5d). Notably, as the ratio increased, where more OOD data were misclassified as In-D, we observed a corresponding increase in the prediction set size to accommodate these risks (i.e., triggering more abstention) for maintaining a valid empirical coverage. By contrast, removing either CRC or OOD detection module significantly harmed the model's empirical coverage, with the decline becoming more severe as the proportion of OOD data increased.

In addition, we investigated if implementing an inspection procedure (as described in the OOD detection above) before applying trained models on a new population's data (CPTAC in our scenario) can address performance degradation (observed in Fig. 2a-d) due to distribution shift between datasets, thereby creating a more reliable model transfer environment. This mechanism, which we call distribution shift control (DSC), aims to ensure that deploying a trained model in different settings, particularly those without fine-tuning capabilities, is done with sufficient safeguards. We observed that excluding the CPTAC data identified by DSC at an OOD score threshold with FPR=0.2 enhanced NSCLC subtyping classification performance (Fig. 5e,f) and CP efficiency by reducing set sizes across all settings (Fig. 5i,j). A sensitivity analysis demonstrated that using an OOD score threshold corresponding to a higher FPR (and consequently higher TPR) consistently enhanced NSCLC subtyping accuracy (Fig. 5g), AUROC (Fig. 5h), average set size (Extended Data Fig. 3a), and DA error rate (Extended Data Fig. 3b). Again, SNGP-EAT surpassed all the other models by a large margin across the considered metrics (Supplementary Tables 17-20).

TRUECAM's benefits extend to digital pathology foundation models

Beyond specialized medical AI models like Inception-v3, which perform tile-level inference and aggregation, TRUECAM seamlessly integrates with general-purpose foundation models to make slide-level inference directly, without the need for tile-level result aggregation (Fig. 1b). Specifically, TRUECAM utilizes a foundation model's pretrained encoder to extract tile representations from each slide, assigns ambiguity values to discard confusing tiles, and processes remaining tiles either through ABMIL for tile-level models or a slide encoder for slide-level models. For evaluation, we systematically assessed TRUECAM's performance with two recently released foundation models for digital pathology: UNI [26], a visual model, and CONCH [27], a visual-language model. For these two models, TRUECAM enforces distance-preserving transformations within the widely-adopted attention-based multiple instance learning (ABMIL) architecture [40] by incorporating spectral normalization in fully-connected layers. The final layer is updated with a Gaussian process to enable UQ. To assess tile ambiguity for EAT within the context of foundation models, a tile-level classifier aligned with the target task remains essential. Thus, we integrated TRUECAM with an AutoML system, AutoGluon [41], to identify ambiguous tiles. Upon the completion of EAT, the trained ABMIL model was directly reused for inference without the need for retraining. For slide-level foundation models like Prov-GigaPath and TITAN, the adaptation of TRUECAM was slightly different. See the Methods for details.

We now show that TRUECAM established both data and model trustworthiness for NSCLC subtyping using digital pathology foundation models, addressing shortcomings of their original versions in deployment environments. We highlight the following key observations. First, when applied to foundation model-based ABMIL in the absence of CP, both UNI-TRUECAM and CONCH-TRUECAM, which removed 60.0% of tiles per slide based on the ambiguity score defined in the Methods, enhanced the patient-level NSCLC subtyping accuracy (Fig. 6a). Specifically, UNI-TRUECAM and CONCH-TRUECAM reduced the error rate by 10.04% and 8.48%, respectively, than their deterministic counterparts (Fig. 6i,j). This highlights the combined effectiveness of SNGP and EAT in overcoming model limitations that contribute to suboptimal performance. See Supplementary Table 29 for TRUECAM's impact on the

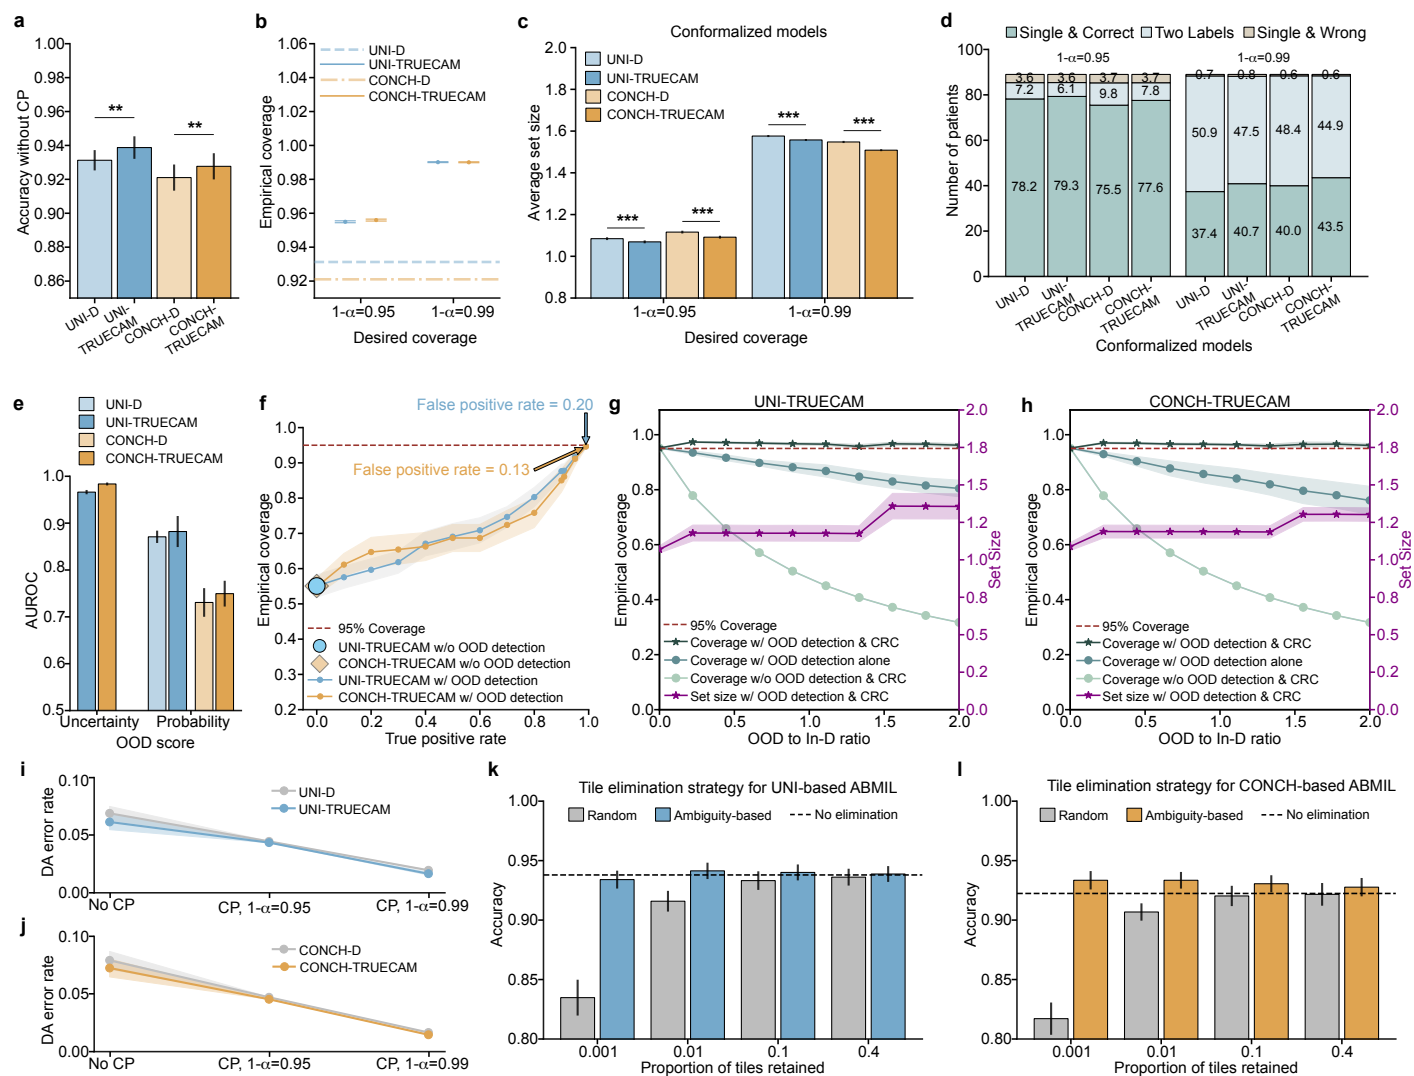


Figure 6. Assessment of TRUECAM using the foundation models UNI and CONCH. Results on TCGA ($n=189$) are reported. See Extended Data Fig. 4 for results on CPTAC, which produced similar observations. **a**, Classification accuracy of foundation model-based ABMIL with (denoted using suffix “-TRUECAM”) and without TRUECAM (i.e., the original version of foundation model-based ABMIL, denoted using suffix “-D”), evaluated without CP. **b**, Comparison of empirical coverage with and without TRUECAM. **c**, Comparison of average set size of CP with and without TRUECAM. **d**, Comparison of patient classification breakdown with and without TRUECAM. **e**, Performance of two distinct OOD scores in identifying slides outside a model’s scope. The uncertainty-based score is not applicable to the original versions of UNI and CONCH. **f**, Empirical coverage with $\alpha = 0.05$ as a function of the true positive rate for OOD detection. **g**, **h**, Empirical coverage and set size of CP for UNI-TRUECAM and CONCH-TRUECAM under conformal risk control with a false positive rate of 0.20 for UNI and 0.13 for CONCH, shown as a function of OOD-to-In-D ratio. **i**, **j**, Comparison of patient-level DA error rate with and without TRUECAM in varying contexts of CP. **k**, **l**, Comparison of classification accuracy for UNI-TRUECAM and CONCH-TRUECAM using TRUECAM’s ambiguity-based tile elimination strategy versus random tile elimination, shown as a function of the proportion of tiles retained per slide. One-sided Wilcoxon signed-rank test is utilized to calculate p values. ** $p < 0.01$, *** $p < 0.001$. ABMIL, attention-based multiple instance learning; CP, conformal prediction; SNGP, spectral-normalized neural Gaussian process; EAT, elimination of ambiguous tiles; DA, definitive-answer.

classification performance of TITAN and Prov-GigaPath. Second, with CP activated, the statistical guarantee of model coverage was established, as evidenced by the close alignment between desired and empirical coverage when applying TRUECAM (Fig. 6b). By comparison, the original foundation models paired with ABMIL for NSCLC subtyping achieved a coverage of only 0.931 (UNI-D) and 0.921 (CONCH-D) and generally lacked the flexibility to adjust to customized error levels. Third, when compared to the conformalized foundation model-based ABMIL (i.e., conformalized UNI-D and CONCH-D, without SNGP-EAT), TRUECAM displayed advantages by achieving smaller average set sizes (Fig. 6c) and producing a greater number of single and correct classifications (Fig. 6d). Fourth, TRUECAM enabled fairer or comparable classifications for both foundation model-based ABMIL setups compared to scenarios without TRUECAM in most cases (Extended Data Figs. 5 and 6).

In foundation model-based inference, TRUECAM identified and rejected OOD inputs with similar effectiveness as in the specialized model setting. Using the same OOD design and dataset as previously described, we observed that the uncertainty-based OOD score (AUROC of 0.967 for UNI-TRUECAM and 0.984 for CONCH-TRUECAM) consistently outperformed the probability-based score (AUROC of 0.882 for UNI-TRUECAM, and 0.750 for CONCH-TRUECAM) in distinguishing between In-D and OOD WSIs (Fig. 6e, Supplementary Table 40). Thus, the uncertainty-based OOD score was favored in the OOD evaluation. Notably, without OOD input identification, the empirical coverage of UNI-TRUECAM and CONCH-TRUECAM fell significantly below the desired coverage (Fig. 6f); however, TRUECAM's OOD detection capability assisted in achieving the desired coverage by adjusting the threshold for OOD score to approach a near perfect TPR. To manage undetected OOD inputs in environments with varying OOD-to-In-D ratios, TRUECAM integrates CRC with OOD detection to effectively address these inputs, achieving an empirical coverage that aligned closely with the desired level and maintaining consistency across a broad range of OOD-to-In-D ratios for both UNI-TRUECAM and CONCH-TURECAM (Fig. 6g,h). By contrast, neglecting either CRC or OOD detection invalidated empirical coverage.

We further evaluated the effectiveness of ambiguity-based tile elimination in TRUECAM's EAT by comparing it to eliminating tiles uniformly at random and then conducted a sensitivity analysis to assess how the proportion of retained tiles influences classification accuracy. We observed that ambiguity-based tile elimination outperformed random tile elimination across a wide range of tile retaining rates for both foundation model-based ABMIL models (Fig. 6k,l). It is remarkable that when ambiguity-based EAT was applied aggressively, even as low as a 0.1% tile retention rate, it did not compromise classification accuracy. By contrast, random tile elimination was much more sensitive to the tile retaining rate, leading to significant degradation in accuracy when performed at a comparable retention rate. This suggested that the ambiguity-based EAT was highly effective as it kept the most informative tiles in slides for NSCLC subtyping.

TRUECAM delivers efficient interpretation and inference

Compared to traditional approaches that rely solely on attention for interpretation, TRUECAM offers a diagnostic basis for model inference through two interpretability layers informed by its UQ module: 1) tile-level ambiguity scores, which assess each tile's potential to introduce confusion and help eliminate corresponding regions, and 2) a global attention map that shows the reliance of slide-level inference on each remaining tile after EAT.

To validate the diagnosis basis of TRUECAM across models, we randomly selected and visualize two WSIs from the TCGA testing cohort: one representing LUAD (Extended Data Fig. 7) and one for LUSC (Extended Data Fig. 8). A pathologist (F.N.), blinded to any model-derived information, received these WSIs to delineate tumor regions with significant diagnostic relevance for NSCLC subtyping. A random selection of tiles from various representative regions was also provided to the pathologist to assess their informativeness for distinguishing LUAD from LUSC. We observed that the regions with relatively low ambiguity scores closely aligned with the annotated areas by the pathologist, which facilitated their discrimination between LUAD and LUSC using established diagnostic criteria. Specifically, these low-ambiguity regions contained cancer epithelial areas, which were critical for distinguishing NSCLC subtypes. By contrast, regions with higher ambiguity scores lacked distinct morphological features of either NSCLC subtype. For example, tiles with blue borders from both LUAD and LUSC WSIs contained inflamed non-neoplastic lung parenchyma, areas of stroma, or necrosis, all of which were devoid of tumor-specific morphological hallmarks (Extended Data Fig. 7, 8). Some regions appeared as poorly preserved or fields with artifacts, further undermining their utility in subtype discrimination.

After applying EAT, the morphological features in regions with high attention weights became the primary basis for the model's classification. As confirmed by the pathologist, the high-attention regions highlighted in blue were generally consistent with the diagnostic criteria used to determine the NSCLC subtype. Tiles with red borders exemplified this alignment: LUAD-specific regions revealed glandular arrangements characteristic of adenocarcinoma, such as well-formed acinar patterns, while LUSC-specific regions displayed intercellular bridges, focal keratinization, and dense eosinophilic cytoplasm, all consistent with the morphological features used in pathologic classification. However, not all regions contributed equally to subtype discrimination. Regions lacking clear indicators for either subtype were recognized by TRUECAM with correspondingly low attention weights. For example, tiles with

purple borders included stroma, areas of inflammation, areas of necrosis, and poorly differentiated regions, none of which provided clear discriminative clues for LUAD vs LUSC.

We then quantified TRUECAM's interpretability compared to traditional models by introducing *attention efficiency*, which we defined as the fraction of high-attention tiles identified by the model that pathologists also consider highly relevant for subtyping decisions. This metric captures the precision of a model's attention in recognizing regions critical for accurate classification. TRUECAM consistently achieved an attention efficiency higher than traditional attention across all three model types (Extended Data Fig. 9a,b). Thus, the two-level interpretability offered by TRUECAM can not only aid pathologists in identifying task-relevant regions for more reliable conclusions, but also serve as a valuable verification source to build trust in AI applications.

Beyond enhanced model interpretability, TRUECAM exhibited advantages in inference speed over traditional and other baseline uncertainty quantification methods. Multiple key findings were observed. First, Inception-v3 equipped with SNGP (denoted as Incep-SNGP) achieved nearly the same tile inference speed as its deterministic counterpart (Extended Data Fig. 9c). Additionally, it was 5× faster than models using MC Dropout. This was primarily due to SNGP's single-forward pass mechanism. Second, the integration of SNGP introduced negligible computational burden on foundation models, partially attributed to their inherently slower encoding process. Third, while foundation models demonstrated significantly lower inference speeds compared to Inception-v3, they delivered superior classification performance (Supplementary Tables 8, 31). Fourth, EAT significantly boosted slide inference speed by excluding a significant proportion of ambiguous tiles, while simultaneously enhancing classification and CP performance (Extended Data Fig. 9d). This inference efficiency gain positioned TRUECAM closer to practical deployment, particularly in time-sensitive and resource-limited medical settings.

Discussions

Trustworthiness of pathology AI is often compromised by model limitations in delivering reliable performance and quantifying uncertainty with statistical implication. Moreover, data challenges like noisy patches and the complexity of model deployment environment, including OOD data and distribution shifts, further undermine trustworthiness of pathology AI. To address these issues, we developed TRUECAM, a framework designed to seamlessly integrate with models with various architectures, purposes, and complexities, to enable responsible and trustworthy AI-driven NSCLC subtyping in pathology.

TRUECAM holds significant implications for the practical application of pathology AI models in real-world clinical settings. The most immediate benefit of TRUECAM lies in its significant reduction in the likelihood of AI models producing incorrect results. However, its impact extends beyond improving diagnostic accuracy; it redefines how pathology AI can be responsibly integrated into clinical workflows of cancer diagnostics. To manage the risk of erroneous patient-level diagnoses and ensure statistically guaranteed error bounds, TRUECAM adaptively defers uncertain or OOD cases to expert pathologists, fostering a collaborative decision-making process. Importantly, the framework tackles a fundamental challenge in pathology AI: inference data may experience distribution shifts, whether due to temporal changes at the same site, inter-site variability, or demographic and medical practice-based differences, all of which can undermine model trustworthiness. By addressing these issues, TRUECAM ensures that AI acts as a reliable assistant rather than an inflexible tool, which complements clinical workflows while mitigating the risks of automation bias. Critically, TRUECAM's deferral mechanism is proactive, customizable, and strategic, abstaining less frequently than other models while achieving superior accuracy. This demonstrates the framework's ability to recognize and act on its own limitations, which confines AI usage within a trusted scope and reduces the cognitive load on pathologists, who can focus their expertise and effort on the most challenging cases [42–44]. By integrating reliability, adaptability, and collaboration, TRUECAM establishes itself as a transformative step toward making AI a trusted partner for clinical decision-making.

TRUECAM not only confines diagnostics to the trusted scope of pathology AI but also lays a solid foundation for responsibly expanding this scope through a built-in awareness of its limitations. In a human-AI collaborative framework, expert pathologists can assign definitive labels to slides flagged as OOD or those where the model abstains from making a prediction. These newly labeled data

points can then be incorporated into the uncertainty quantification process to enable iterative refinement of the model and its conformal predictor, progressively broadening the trusted scope of pathology AI in a controlled and reliable manner. Such a progressive approach outlines a responsible development paradigm, ensuring that pathology AI evolves in tandem with its real-world applications. As trustworthiness takes center stage in medical AI, the ability to systematically address foundational models' limitations in clinical settings becomes critical [45, 46]. By enabling pathology AI to adapt and expand their capabilities through iterative feedback and refinement, TRUECAM provides key insights that can guide the development of next-generation foundation models in pathology, paving the way for their scalable and trustworthy deployment in clinical workflows.

While SNGP and EAT are designed to address data trustworthiness and CP (with conformal risk control) to assure model trustworthiness, these components do not function independently but instead provide complementary benefits. SNGP-EAT establishes an environment that upholds data exchangeability for CP to function and ensures that the input data is within the scope of the trained pathology AI models. Meanwhile, CP guarantees that classifications based on SNGP-EAT are statistically valid in containing the ground truth labels with a pre-specified coverage guarantee. The absence of either component significantly compromises NSCLC subtyping reliability. For example, without SNGP-EAT-based OOD detection, the empirical coverage of CP for both specialized and foundation model-based models dropped substantially to 0.477 (compared to the desired coverage of 0.95), resulting in 52.3% incorrect patient subtyping even with CP in place (Fig. 5d). While retaining SNGP and EAT alone (without CP) improved classification performance compared to baseline models (Fig. 3), it lacked the ability to identify challenging inputs and abstain from providing a diagnosis, nor can it establish bounded error rates to ensure guaranteed validity. The interplay between these components creates a robust framework that not only enhances individual performance but also addresses the broader challenges of real-world variability in pathology workflows.

Aligned with the data valuation and data-centric AI paradigm emphasized by the machine learning community [42–44], EAT exemplifies the “less data, better performance” phenomenon and delivers transformative benefits for NSCLC subtyping in TRUECAM. While conceptually similar to the topK max-pooling operator used in MI-Zero [27, 47], which aggregates the most informative tile-level predictions, EAT takes this approach a step further. By filtering out ambiguous tiles and focusing on high-diagnostic-value regions, EAT creates a feedback loop that enhances model learning through reinforced supervisory signals from diagnostically valuable slide areas, leading to improved tile representations and amplifying both efficiency and accuracy.

EAT's advantages are manifold. First, it significantly reduces the computational burden for slide inference. For example, in the specialized model (Inception-v3), EAT decreased inference computational load by 66% (Fig. 3a), while in the foundation model-based setting, this efficiency gain reached up to 1000× without compromising accuracy (Fig. 6k,l). This efficiency lowers barriers to deploying advanced pathology AI models in resource-constrained settings, such as small clinics and rural hospitals, enabling broader access to AI-driven diagnostics. Second, EAT enhances interpretability by directing the model's focus to diagnostically relevant areas, closely aligning its attention with pathologists' annotations (Extended Data Fig. 7, 8). This not only reduces the cognitive load for pathologists when reviewing AI-generated results but also fosters greater trust in the model's outputs, since its decision-making process becomes more transparent and clinically meaningful. From a data annotation perspective, EAT offers a scalable and efficient mechanism for filtering out low-value tiles, substantially reducing the reliance on labor-intensive and costly pixel-level annotations. This streamlined process saves time and resources while addressing the need for domain-specific expertise, making EAT a practical tool for advancing future model development. Third, EAT has significant implications for promoting equity by addressing demographic fairness gaps in both classification performance and CP set size, as seen with the original specialized and foundation model-based models. By eliminating highly ambiguous tiles from each slide, EAT mitigates biases that could arise when traditional pathology AI approaches process all tiles indiscriminately, which may inadvertently reinforce harmful disparities in NSCLC subtyping inference. These disparities could result in unequal treatment of patients from different racial or sexual backgrounds [48]. Importantly, EAT achieved these fairness improvements without requiring explicit constraints during model training, which highlights its potential to improve the equity and clinical reliability of pathology AI systems. Fourth, EAT also has far-reaching implications for the development

of pathology AI foundation models, particularly in optimizing their pretraining process. The current standard self-supervised learning frameworks (e.g., DINOv2 [49]) for image encoder pretraining treat all tiles in a slide equally. However, it is well-established that only cancerous regions are critical for clinically relevant downstream tasks. Similarly, in visual-language foundation models like CONCH or TITAN, a large portion of tiles in each slide are irrelevant to the diagnostic reports and contribute little to aligning text and image modalities in the model's representation space. By applying EAT to exclude non-relevant tiles during pretraining, the process could become significantly more efficient by focusing on the meaningful cancerous areas. This would not only reduce the computational cost of pretraining but also enable the model to learn better representations, ultimately improving its performance across various downstream tasks.

For OOD detection, we observed that the uncertainty-based OOD score consistently outperformed the probability-based OOD score in foundation models (Fig. 6e, Extended Data Fig. 4e, Supplementary Table 40), but showed inferior performance compared to the probability-based OOD score in the specialized model, i.e., Inception-v3, (Fig. 5a). This phenomenon can be in part attributed to the representational power of the underlying model. Since the uncertainty estimated by SNGP is essentially determined by the distance between testing and training data, models with stronger feature representation capabilities—such as foundation models, which achieve higher classification accuracy than Inception-v3—provide more discriminative and semantically rich dimensions to distinguish between In-D and OOD data in the feature space. This also aligns with findings reported by existing research: distance-based methods can outperform probability-based methods in OOD detection tasks, particularly when the feature space is well-structured and the Gaussian assumption holds [50]. On the other hand, the performance difference in OOD detection could be due to the distinct mechanisms of probability-based and uncertainty-based OOD score for OOD detection [50]. In essence, for OOD detection, the probability-based score examines a data point's distance from the learned decision boundary in the feature space, while the uncertainty-based OOD score estimates the uncertainty of the data based on its distance from the training data. As foundation models learn better representations of training data than Inception-v3, the resultant representations enable fine-grained and informative estimates for the distance to the training data in the learned feature space.

There are several limitations that we wish to highlight as opportunities for furthering this line of research. First, this study focused on the binary NSCLC subtyping problem of discriminating LUAD vs LUSC. However, the generalizability of TRUECAM's advantages in addressing broader multi-class cancer subtyping tasks in digital pathology needs to be assessed in settings with a larger number of tumor histologies, pathological patterns, and dataset characteristics. Second, we evaluated TRUECAM using one specialized model and four cutting-edge pathology AI foundation models that lack conversational interaction capabilities. Yet with the emergence of interactive pathology AI foundation models [51–53], which lower the barrier to practical use by enabling conversational queries and reasoning regarding WSIs, it is important to investigate and assess how TRUECAM can be adapted to this new language-centered paradigm to further enhance model reliability. Third, while our findings suggested that distribution shift control enhanced the classification and CP performance on an external dataset (CPTAC) collected from a program distinct from the training dataset (TCGA) (Fig. 5e,f,i,j), the root cause of the distribution shift remains unidentified. Further research is warranted to systematically investigate these distribution shifts and pinpoint their underlying cause, which could guide the model refinement. Fourth, we did not evaluate the overall empirical performance of TRUECAM in a human-in-the-loop setting, particularly when abstention is triggered and pathologists are involved for independent inference. In the future, well-designed trials will be critical to evaluate the impact of TRUECAM-assisted clinical decision-making for NSCLC subtyping. In practice, pathologists may choose to accept or override the definitive answers provided by error-bounded TRUECAM, influenced by factors such as the pathological features of specific slides and their own experience and confidence. The effect of these decisions on patient outcomes remains unclear. Therefore, it is essential to evaluate and explore effective human-AI collaboration strategies before deploying TRUECAM in clinical settings.

Method

Dataset description

Dataset overview. Supplementary Table 1 summarizes the basic characteristics of the datasets used in this study. We utilized two publicly available digital pathology datasets of NSCLC stained with hematoxylin and eosin (H&E): 1) 941 WSIs from TCGA, comprising 467 LUAD and 474 LUSC cases, and 2) 1,306 WSIs from CPTAC, comprising 644 LUAD and 662 LUSC cases. In the TCGA dataset, each patient is represented by a single slide, whereas in the CPTAC dataset, each patient is associated with an average of 3.1 slides. Among the 416 patients in CPTAC, 203 are diagnosed with LUAD, while the remaining 213 are diagnosed with LUSC. To evaluate TRUECAM's ability in detecting OOD inputs, we constructed a TCGA-OOD dataset, consisting of 698 WSIs from 631 patients with non-lung, non-adenocarcinomas, and non-squamous tumors. TCGA-OOD includes slides from 92 patients with adrenocortical carcinoma (ACC), 70 patients with uveal melanoma (UVM), 57 patients with uterine carcinosarcoma (UCS), and 412 patients with bladder urothelial carcinoma (BLCA), representing all available WSIs from each considered cancer category in TCGA. Among these, except for UVM, one patient is associated with more than 1 tissue slides. We included BLCA, which is comprised of over 65% of the dataset, because its urothelial origin and histological spectrum—ranging from low-grade papillary lesions to high-grade invasive forms—can closely resemble poorly differentiated NSCLC. In particular, high-grade BLCA exhibits certain morphological features that make it very challenging to distinguish from NSCLC, such as solid nests or sheets of atypical cells. As a result, TCGA-OOD serves as a meaningful dataset for the OOD detection task.

Image preprocessing. Regardless of distinct model input requirements, all WSIs were segmented and split into appropriate tiles using a standard preprocessing pipeline implemented in Slideflow (version 2.1.0) [54]. Directly feeding large gigapixel WSIs into a neural network is computationally impractical, so image tiles of manageable sizes were extracted to align with the model's specifications. For example, foundation models like UNI, CONCH, and Prov-GigaPath accept image tiles of size 256×256 at 20X magnification; as we follow the Inception-v3 model developed in [14], it requires image tiles of 299×299 at 10X magnification; unlike UNI, CONCH, and Prov-GigaPath, the latest foundation model TITAN accepts image tiles with a size of 512×512 at 20X magnification. The preprocessing pipeline also includes several standard steps to create deep learning-ready datasets. Otsu's thresholding [55] was applied to differentiate the foreground (tissue) from the background (empty slide), followed by grayscale filtering to remove background tiles. All extracted tiles were then stain-normalized and saved in TFRecord format for model training, inference, and performance evaluation. Importantly, no pathologist-annotated regions of interest were used; instead, we relied solely on slide-level diagnoses and assigned a uniform weak label to all tiles extracted from a given WSI when needed for model training.

Dataset configuration. In the setup for the specialized model based on Inception-v3, we randomly partitioned the TCGA dataset, allocating 65% of the patients' data for training, 15% for validation, and 20% for calibration and testing. We repeated the random partitioning procedure 20 times, producing 20 independently trained models for evaluation. Within the 20% allocated for calibration and testing, data from 100 patients was used for calibration to establish CP (referred to as the TCGA calibration dataset), while data from 89 patients was set aside for testing (referred to as the TCGA testing dataset). To assess the variability of conformal prediction performance, the 20% split was randomized 500 times, generating distinct calibration and testing datasets for each model. As a result, we report model performance based on 20 models \times 500 evaluations per model. The CPTAC dataset was fully utilized for external validation to assess the effectiveness of TRUECAM in scenarios that need transferability, with data from 100 patients designated for calibration to establish conformal prediction (referred to as the external CPTAC calibration dataset) and the remaining 316 patients reserved for testing (referred to as the external CPTAC testing dataset). Similarly, the CPTAC dataset was randomly partitioned 500 times to evaluate the variability of conformal prediction performance.

In the scenario of using foundation models for NSCLC subtyping, which involves training a separate model using latent tile representations generated by foundation models, the TCGA and CPTAC datasets were utilized independently for the entire model training and evaluation process. Each dataset was partitioned into three subsets: 65% for training, 15% for validation, and 20% for calibration and testing. Consistent with the strategy used for Inception-v3 evaluation, this partitioning and model training process was

repeated 20 times, with each involving 500 random splits to establish and evaluate CP.

Spectral-normalized Neural Gaussian Process (SNGP)

We leveraged SNGP [50, 56] to provide a principled estimation of data uncertainty in deep learning models for NSCLC subtyping. SNGP introduces two important modifications to a regular neural network: 1) incorporating spectral normalization (SN) into the hidden layers and 2) replacing the conventional dense output layer with a Gaussian process (GP). Incorporating SN ensures that the transformation from the input \mathbf{x} in the original input space to the latent representation in the penultimate layer $h(\mathbf{x})$ preserves the relative distances. In other words, the distance $\|h(\mathbf{x}_1) - h(\mathbf{x}_2)\|_{\mathbb{H}}$ between representations of any two image data points $(\mathbf{x}_1, \mathbf{x}_2)$ in the latent space meaningfully reflects their distance $\|\mathbf{x}_1 - \mathbf{x}_2\|_{\mathbb{X}}$ in the original data space, relative to other data pairs. In practice, adding SN addresses a problem in training neural networks known as feature collapse [57]. This issue occurs when OOD data or data from different classes, which are geometrically distant from remaining data in the input space, are unexpectedly mapped to nearby points in the latent space, thus leading to unreasonable uncertainty estimations. Mathematically, the SN operation is equivalent to requiring the mapping function $h(\cdot)$ to satisfy the following bi-Lipschitz condition for any pair of inputs \mathbf{x}_1 and \mathbf{x}_2 [56, 58]:

$$L_1 \times \|\mathbf{x}_1 - \mathbf{x}_2\|_{\mathbb{X}} \leq \|h(\mathbf{x}_1) - h(\mathbf{x}_2)\|_{\mathbb{H}} \leq L_2 \times \|\mathbf{x}_1 - \mathbf{x}_2\|_{\mathbb{X}}, \quad (1)$$

where L_1 and L_2 are the positive lower and upper Lipschitz bounds ($0 < L_1 < 1 < L_2$) of feature extractor $h(\cdot)$, and $\|\cdot\|_{\mathbb{X}}$ and $\|\cdot\|_{\mathbb{H}}$ correspond to the distance metrics chosen for the original and latent space, respectively. The lower Lipschitz bound L_1 serves to preserve distances in the latent space and the upper Lipschitz bound L_2 helps enforce the smoothness and robustness of the neural network by limiting over-sensitive transformations to perturbations in the input space of \mathbf{x} .

In SNGP, the bi-Lipschitz constraint is enforced by applying spectral normalization to neural network weights $\{\mathbf{W}_l\}_{l=1}^{L-1}$ [59], where l represents the l -th hidden layer of the neural network. During each training step, the SN method first estimates the spectral norm $\hat{\lambda} \approx \|\mathbf{W}_l\|_2$ using the power iteration method [60], and then normalizes the neural network weights as:

$$\hat{\mathbf{W}}_l = \begin{cases} c\mathbf{W}_l/\hat{\lambda}, & \text{if } c < \hat{\lambda} \\ \mathbf{W}_l, & \text{otherwise} \end{cases}, \quad (2)$$

where c is a hyperparameter used to adjust the exact spectral norm upper bound on $\|\mathbf{W}_l\|_2$ (i.e., $\|\mathbf{W}_l\|_2 \leq c$).

SNGP then estimates data uncertainty by leveraging the preserved distances in the latent space, achieved by replacing the dense output layer with a GP. Consider a dataset consisting of N training points $\mathcal{D} = \{\mathbf{x}_i, y_i\}_{i=1}^N$, and let $h_i(\mathbf{x}_i)$ denote the representation of \mathbf{x}_i in the second-to-last layer of neural network. The output $\mathbf{f}_{N \times 1} = [f(h_1), f(h_2), \dots, f(h_N)]^T$ of GP follows a multivariate Gaussian distribution:

$$\mathbf{f}_{N \times 1} \sim \mathcal{N}(\mathbf{0}_{N \times 1}, \mathbf{K}_{N \times N}), \text{ where } \mathbf{K}_{i,j} = \exp\left(\frac{-\|h_i - h_j\|_2^2}{2}\right). \quad (3)$$

Conditional on the learned latent representation h_i , SNGP approximates GP using the Laplace approximation with the random Fourier feature (RFF) expansion [61]. Specifically, SNGP approximates the GP prior in Eq. (3) with a low-rank approximation to the kernel matrix $\mathbf{K} = \Phi\Phi^T$ using random features [62]:

$$\mathbf{f}_{N \times 1} \sim \mathcal{N}(\mathbf{0}_{N \times 1}, \Phi\Phi^T), \text{ where } \Phi_i = \sqrt{\frac{2}{D_L}} \cos(-\mathbf{W}_L h_i + \mathbf{b}_L), \quad (4)$$

where Φ_i is the last layer of the neural network with a dimension D_L , \mathbf{W}_L is a fixed weight matrix randomly generated from $\mathcal{N}(0, 1)$ and \mathbf{b}_L is a fixed bias vector randomly generated from a uniform distribution $\mathcal{U}(0, 2\pi)$. Since \mathbf{W}_L and \mathbf{b}_L are fixed, the RFF approximation to the GP prior for the k -th logit of the classification problem (in Eq. (3)) can be reformulated as:

$$g_k(h_i) = \sqrt{\frac{2}{D_L}} \cos(-\mathbf{W}_L h_i + \mathbf{b}_L)^T \times \beta_k, \quad k = 1, \dots, K, \quad (5)$$

where $\beta_k \sim \mathcal{N}(0, \mathbf{I})$, K denotes the number of classes in the classification problem, and $\boldsymbol{\beta} = \{\beta_k\}_{k=1}^K$ are the only learnable parameters in the output layer.

Overall, the trainable parameters consist of $\boldsymbol{\beta}$ along with the weights and biases $\boldsymbol{\alpha} = \{\mathbf{W}_l, \mathbf{b}_l\}_{l=1}^{L-1}$ in the first $L-1$ layers of the neural network. Notably, the RFF approximation to GP reduces the high-dimensional GP to a standard Bayesian linear model while also producing a closed-form posterior that is end-to-end trainable alongside the rest of the neural network using stochastic gradient descent. After inferring the MAP estimate of $\hat{\beta}_k$, for an unobserved point \mathbf{x}^* , the first $L-1$ layers of SNGP serve as a feature extractor to derive its latent representation $\Phi^* = \sqrt{\frac{2}{D_L}} \cos(-\mathbf{W}_L h(\mathbf{x}^*) + \mathbf{b}_L)$. Next, the mean and variance associated with the k -th logit in the classification problem are computed as:

$$\mu_k(\mathbf{x}^*) = [\Phi^*]^T \hat{\beta}_k, \quad \sigma_k(\mathbf{x}^*) = [\Phi^*]^T \Sigma_k \Phi^*. \quad (6)$$

Finally, the predictive distribution for \mathbf{x}^* is calculated as:

$$p(y^* | \mathbf{x}^*) = \int_{s \sim \mathcal{N}(\mu_k(\mathbf{x}^*), \sigma_k(\mathbf{x}^*))} \gamma(s), \quad (7)$$

where γ denotes the softmax activation function.

Conformal prediction

Conformal prediction, or CP, is a distribution-free approach for constructing prediction sets for arbitrary prediction or classification algorithm. It provides statistical guarantees of a bounded error rate for model outputs, rigorously aligned with the error level specified by developers or service providers under data exchangeability assumption [63]. Consider a classification problem comprising of K classes. Let $\mathcal{D} = \{\mathbf{x}_i, y_i\}_{i=1}^T$ be independently and identically distributed (i.i.d.) data points sampled from distribution $\mathcal{X} \times \mathcal{Y}$, where $\mathbf{x}_i \sim \mathcal{X}$, $y_i \sim \mathcal{Y}$, $y_i \in \{1, \dots, K\}$, and $\alpha \in [0, 1]$ denotes a specified error level (also referred to as significance level). Then the objective of CP is to construct a prediction set for each data point, denoted as

$$\Gamma: \mathcal{X} \rightarrow \{\text{subset of } \{1, 2, \dots, K\}\}, \quad (8)$$

such that for a new i.i.d. data point $(\mathbf{x}^*, y^*) \sim P(\mathcal{X}, \mathcal{Y})$, the probability that the prediction set fails to contain the truth label is α bounded. This is formally expressed as [64]:

$$1 - \alpha \leq \mathbb{P}(y^* \in \Gamma(\mathbf{x}^*) | \mathbf{x}^*) \leq 1 - \alpha + \frac{1}{T+1}, \quad (9)$$

where the probability is taken over the data $(\mathbf{x}_i, y_i)_{i=1}^T \cup (\mathbf{x}^*, y^*)$. By design, CP guarantees that the prediction set $\Gamma(\mathbf{x}^*)$ contains the true label y^* with a probability of at least $1 - \alpha$. The bounded error rate is achieved by generating a set of likely class labels for the inputs with insufficient certainty in a systematic way, in contrast to the single-value (in regression) or single-label (in classification) outputs produced by regular machine learning models.

Using two separate and disjoint datasets for model training and calibration, the development of an inductive conformal predictor consists of three key steps:

1. **Define a nonconformity score.** Conformity reflects how similar or conformal a new data point is to the training data [65] and is typically quantified using a nonconformity score. In classification tasks, the nonconformity score is oftentimes defined as $s(\mathbf{x}_i, y_i) = 1 - \hat{f}(\mathbf{x}_i)_{y_i}$, where $\hat{f}(\mathbf{x}_i)_{y_i}$ denotes the probability assigned by the trained model \hat{f} to the true class y_i for the input \mathbf{x}_i . Thus, the lower the model's confidence in correctly classifying \mathbf{x}_i as y_i , the higher the corresponding nonconformity score. We used this nonconformity score throughout the study.
2. **Determine the nonconformity threshold.** We then computed the nonconformity score $s(\mathbf{x}_i, y_i)$ for each data point \mathbf{x}_i ($i = T+1, \dots, T+R$) in the calibration set and obtain a set of nonconformity scores: $\{s(\mathbf{x}_{T+1}, y_{T+1}), s(\mathbf{x}_{T+2}, y_{T+2}), \dots, s(\mathbf{x}_{T+R}, y_{T+R})\}$. We ranked these scores in ascending order and determine the threshold \hat{q} as the $\frac{\lceil (R+1)(1-\alpha) \rceil}{R}$ quantile of these ordered scores. This threshold ensures that the resulting prediction sets have a coverage rate of at least $1 - \alpha$ on new i.i.d. datasets.

3. **Generate prediction sets.** For a new data point \mathbf{x}^* , the prediction set is formed by including all class labels $k \in \{1, 2, \dots, K\}$ whose corresponding scores $s(\mathbf{x}^*, k)$ are smaller than or equal to \hat{q} . Mathematically, we have:

$$\Gamma(\mathbf{x}^*) = \{k : s(\mathbf{x}^*, k) \leq \hat{q}\}, k = 1, \dots, K. \quad (10)$$

Unlike regular uncertainty quantification methods, CP converts the traditional heuristic concept of uncertainty, often represented by softmax outputs, into a statistically meaningful measure in the form of a prediction set. By following the steps described above, CP ensures to generate statistically valid prediction sets that contain the true label with a pre-specified coverage of $1 - \alpha$. This approach offers a standardized framework for constructing reliable confidence intervals that enables machine learning systems to deliver robust outputs while upholding responsible and trustworthy scientific inferences.

Specialized models based on Inception-v3

We trained three deep learning models based on the Inception-v3 architecture [24, 25]. The first is the original deterministic version, which does not include UQ. The other two models do incorporate UQ techniques: 1) MC Dropout [9] and 2) SNGP. Inception-v3 is a 48-layer deep convolutional neural network designed to capture spatial features at varying scales by employing parallel convolution operations with different filter sizes within the same layer. To address the vanishing gradient problem and provide additional supervision signals, Inception-v3 integrates an auxiliary classifier to propagate label information to earlier network layers. In addition, Inception-v3 utilizes bottleneck layers to reduce dimensionality and computational costs. These design features collectively enhance performance and resource efficiency, rendering Inception-v3 widely adopted for a diverse range of general computer vision and medical imaging applications. Supplementary Table 3 provides parameter and training details. The integration of MC dropout with Inception-v3 follows the design adopted by Dolezal et al. [14]. At the inference stage, uncertainty is estimated by generating an ensemble of predictions, where each image tile is passed through the dropout-enabled Inception-v3 five times.

We trained models from scratch using a 24GB NVIDIA RTX 4090 GPU with a batch size of 64, where image tiles are resized to 299×299 pixels. We adopted a learning rate scheduler with an initial learning rate of 0.0003 and a decay rate of 0.98 every 512 steps. Each model is trained for four epochs using the Adam optimizer with the 1st and 2nd moment exponential decay rate setting to 0.9 and 0.999, respectively. During model training, we monitored model performance on the validation dataset and save the best-performing model for evaluation on the testing data. Standard data augmentation techniques, including random horizontal and vertical flips, 90-degree rotations, JPEG compression with a 50% probability (quality range: 50-100), and Gaussian blur with a 10% probability (sigma: 0.5-2.0), are used to augment the training process of each model.

Foundation models

TRUECAM seamlessly integrates with digital pathology foundation models for trustworthy NSCLC subtyping. In this study, we explore integrating TRUECAM with two types of foundation models: tile-level (or patch-level) models, including UNI [26] and CONCH [27], and slide-level models, including Prov-GigaPath [28] and TITAN [29], both using the vision transformer-based architecture as image encoders. All of them are general-purpose models systematically pretrained on diverse sources of histopathology images using self-supervised strategies. Among them, CONCH, Prov-GigaPath, and TITAN are visual-language models designed to align image and text modalities within their representation space, whereas UNI is exclusively image-pretrained. These foundation models exhibit exceptional adaptability and knowledge transferability through image encoding, making them highly effective for a wide range of downstream tasks in histopathology, including image classification, region-of-interest retrieval, and cell type segmentation.

To adapt tile-level foundation models for NSCLC subtyping, we utilized the image encoders of UNI and CONCH to extract tile embedding (i.e., latent representations), concatenate the resulting representations for each slide, and apply the widely used attention-based multiple instance learning (ABMIL) for weakly supervised classification using slide-level diagnostic labels. The ABMIL model architecture and training configuration follow those established in UNI [26] and CONCH [27]. Specifically, the ABMIL architecture includes a fully-connected layer with a rectified linear unit (ReLU) activation function to transform tile-level features

into 512-dimensional embedding, followed by an intermediate layer with size 384. Finally, another fully-connected layer maps the attention-pooled slide-level representations to logits and class probabilities through softmax normalization. Dropout is applied at multiple stages: a rate of 0.1 for the input features and 0.25 after each intermediate layer. To incorporate SNGP, we added spectral normalization to all fully-connected layers and replace the final layer with a RFF-approximated Gaussian process. The ABMIL model is trained via a maximum of 20 epochs using the AdamW optimizer [66] with a cosine learning rate scheduler, an initial learning rate of 1×10^{-4} , and cross-entropy loss.

For slide-level foundation models, we fine-tuned the slide encoders of Prov-GigaPath and performed linear probing on the slide representations of TITAN for downstream classification tasks as suggested by [29] and [28]. For both models, we appended a Gaussian process layer to the slide-level representations without applying spectral normalization, as the representations extracted from Prov-GigaPath and TITAN inherently preserve distances effectively, due to their dedicated self-supervised training strategy.

Before the trained ABMIL models with SNGP and fine-tuned slide encoders with GP are utilized for NSCLC subtyping, EAT is applied to identify and remove ambiguous tiles from each slide. As the tile-level features extracted from foundation models are incompatible with the Inception-v3 architecture, we employed AutoGluon to train a machine learning model as a proxy for estimating ambiguity scores defined later. AutoGluon treats the extracted features as tabular data and searches the design space comprising multiple state-of-the-art models (e.g., LightGBM, NeuralNet, XGBoost, ExtraTrees) for NSCLC subtyping. It then automatically ensembles the weights of these base models to maximize the model performance on the validation data. Only the tiles with ambiguity score less than a given threshold are kept in the attention-pooled slide-level representations for NSCLC subtyping. Such a strategy eliminates 60.0% tiles for TCGA and CPTAC, a level comparable to that achieved in the specialized model setting.

Ambiguity score

Regardless of the underlying tile-level classification model, given a trained classification model M , we define ambiguity score of a tile \mathbf{x} as $s(\mathbf{x}, M) = 1 - |p(y(\mathbf{x}) = 0 | M) - p(y(\mathbf{x}) = 1 | M)|$, where $|p(y(\mathbf{x}) = 0 | M) - p(y(\mathbf{x}) = 1 | M)|$ measures how discernible tile \mathbf{x} is by the model M . Clearly, the higher the ambiguity score, the lower $|p(y(\mathbf{x}) = 0 | M) - p(y(\mathbf{x}) = 1 | M)|$ and the less informative tile \mathbf{x} to the model M in the NSCLC subtyping task. For Inception-v3, the model itself was used to calculate the ambiguity score. In contrast, for UNI, CONCH, and TITAN, their respective tile encoders were used to extract tile representations, and AutoGluon was trained from scratch to derive tile-level ambiguity scores. Prov-GigaPath utilized the AutoGluon model derived from CONCH to evaluate its transferability.

Uncertainty quantification

MC Dropout is employed as the primary baseline approach for uncertainty quantification [9]. By assuming that the prior of neural network weights follows a standard normal distribution, MC Dropout is mathematically equivalent to sampling from a variational posterior consisting of two independent Gaussians with fixed covariance. To implement MC Dropout in Inception-v3, we appended two fully-connected layers, each containing 1,024 neurons, to the end of the neural network and applied dropout with a rate of 0.1 after each fully-connected layer. During inference, each image tile was forward-passed through the dropout-enabled network five times. The standard deviation of the model's predictions across these five passes was calculated as the measure of tile-level uncertainty. The tile-level predictions and their corresponding uncertainty were then aggregated to produce slide- and patient-level predictions, along with the associated uncertainty estimates.

To derive data uncertainty in the context of SNGP, once the model is properly trained, the random Fourier features for \mathbf{x}^* can be calculated as $\Phi^* = \sqrt{\frac{2}{D_L}} \cos(-\mathbf{W}_L h_i(\mathbf{x}^*) + \mathbf{b}_L)$. With this D_L -dimensional latent representation of \mathbf{x}^* , the tile-level uncertainty estimates can be derived as $\sqrt{\tau [\Phi^*]^T (\Phi_{N \times D_L}^T \Phi_{N \times D_L} + \tau \mathbf{I})^{-1} \Phi^*}$, where τ denotes the ridge factor. SNGP requires only a single forward pass of an image tile through the neural network to extract the corresponding RFF features for uncertainty estimation. By contrast, MC Dropout requires multiple forward passes of the same image tile to quantify prediction uncertainty. This makes SNGP more computationally efficient during inference, which is particularly important for large-scale, safety-critical applications with

stringent real-time inference requirements like cancer diagnostics.

It is important to note that SNGP and CP differ in their approaches for uncertainty quantification. Beyond the regular classification probability of \mathbf{x}^* produced by a deep learning model, SNGP generates an estimated standard deviation $\sigma(\mathbf{x}^*)$ as the uncertainty measure for \mathbf{x}^* . By contrast, CP does not produce any explicit uncertainty measure. Instead, it operates within the class probability space established by the deep learning model and constructs a prediction set $\Gamma(\mathbf{x}^*)$ for \mathbf{x}^* by including labels with the corresponding prediction probabilities beyond a threshold. Thus, SNGP and CP fundamentally differ in how they quantify uncertainty.

For OOD detection and DSC, two scores were considered: 1) a probability-based OOD score, and 2) an uncertainty-based OOD score. The probability-based OOD score aggregates tile-level class probabilities and measures the degree of OOD as $1 - \frac{1}{N} \sum_{i=1}^N \max_k p(\hat{y}_k | \mathbf{x}_i)$, where $p(\hat{y}_k | \mathbf{x}_i)$ denotes the probability of tile i being classified as \hat{y}_k . Unlike the probability-based OOD score, the uncertainty-based OOD score quantifies tile-level prediction uncertainty using the standard deviation estimated by MC Dropout or SNGP. It then employs the first δ tiles with the lowest uncertainty for each patient to represent the patient-level prediction uncertainty. Clearly, the uncertainty-based OOD score offers additional insight beyond what deterministic models can provide, effectively highlighting the inherent decision limitations of deep learning models, as demonstrated by its superior performance in identifying OOD images in the foundation model-based approach.

Conformal risk control (CRC)

CRC extends regular CP to accommodate a broader range of loss functions that assess different aspects of prediction quality beyond miscoverage, such as accuracy, calibration, and other task-specific performance metrics [67]. More concretely, we define the prediction set for \mathbf{x} by model \hat{f} as $\Gamma_\rho(\mathbf{x}) = \{k : \hat{f}(\mathbf{x})_k \geq 1 - \rho\}$, where $k \in \{1, \dots, K\}$ denotes all possible class labels and ρ is a threshold that controls the conservativeness of the prediction set. Then, for any bounded loss function $l(\cdot)$ that is non-increasing as the prediction set $\Gamma_\rho(\mathbf{x})$ gets larger, CRC provides a statistical guarantee on the loss in the following form:

$$\mathbb{E} [l(\Gamma_\rho(\mathbf{x}), y)] \leq \alpha. \quad (11)$$

Compared to regular CP, CRC extends CP to control the expected value of any monotone loss function instead of only coverage. The goal of CRC is to determine the lowest possible $\hat{\rho}$ based on the calibration dataset, such that the risk control defined in Eq. (11) holds for unseen data. In the deployment of pathology AI where In-D and OOD data coexist, the observed coverage over the calibration data guides the identification of $\hat{\rho}$ via binary search. The threshold $\hat{\rho}$ can be iteratively adjusted until the target coverage $1 - \alpha$ is satisfied in the presence of OOD data. After the optimal $\hat{\rho}$ is identified, CRC then applies it to the unseen data such that their coverage is $1 - \alpha$ guaranteed.

Model fairness evaluation

For model fairness evaluation, we examined the performance gap between the best- and worst-performing subgroups based on sex and race [68]. Two performance dimensions were assessed: 1) accuracy of model predictions, and 2) size of CP prediction sets. Group-level average values of these metrics were calculated to represent the performance of each corresponding subgroup. Since the gap value is always non-negative, a smaller gap indicates a fairer model. Sex and race information is available in both the TCGA and CPTAC datasets. TCGA includes 569 male and 372 female patients, while CPTAC comprises 288 male and 127 female patients, with one patient's sex not reported. For race, categories with fewer than 20 patients are grouped as 'Others'. Specifically, in TCGA, patients are grouped into Others (n=98), White (n=678), and Not reported (n=165). In CPTAC, patients are grouped into Asian (n=123), Others (n=49), and White (n=244). Demographic information is summarized in Supplementary Table 1.

Acknowledgments

X.Z. and T.W. were partially supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. PolyU 25206422) and the National Natural Science Foundation of China (Grant No. 62406269). C.Y. was

supported by the U.S. National Institutes of Health (NIH) NLM Pathway to Independence Award 1K99LM014428-01A1.

Author contributions statement

X.Z., T.W., and C.Y. conceived the study and designed the experiments, X.Z. carried out data collection, T.W. wrote the codes, performed the experiments, and analyzed the experimental results. X.Z. and C.Y. wrote the original manuscript and summarized major experimental findings, Y.C. and B.M. reviewed and revised the manuscript, and F.N. annotated the two WSI slides and added descriptions on the role of different tiles for NSCLC subtyping. All authors contributed to manuscript preparation and approved the final version.

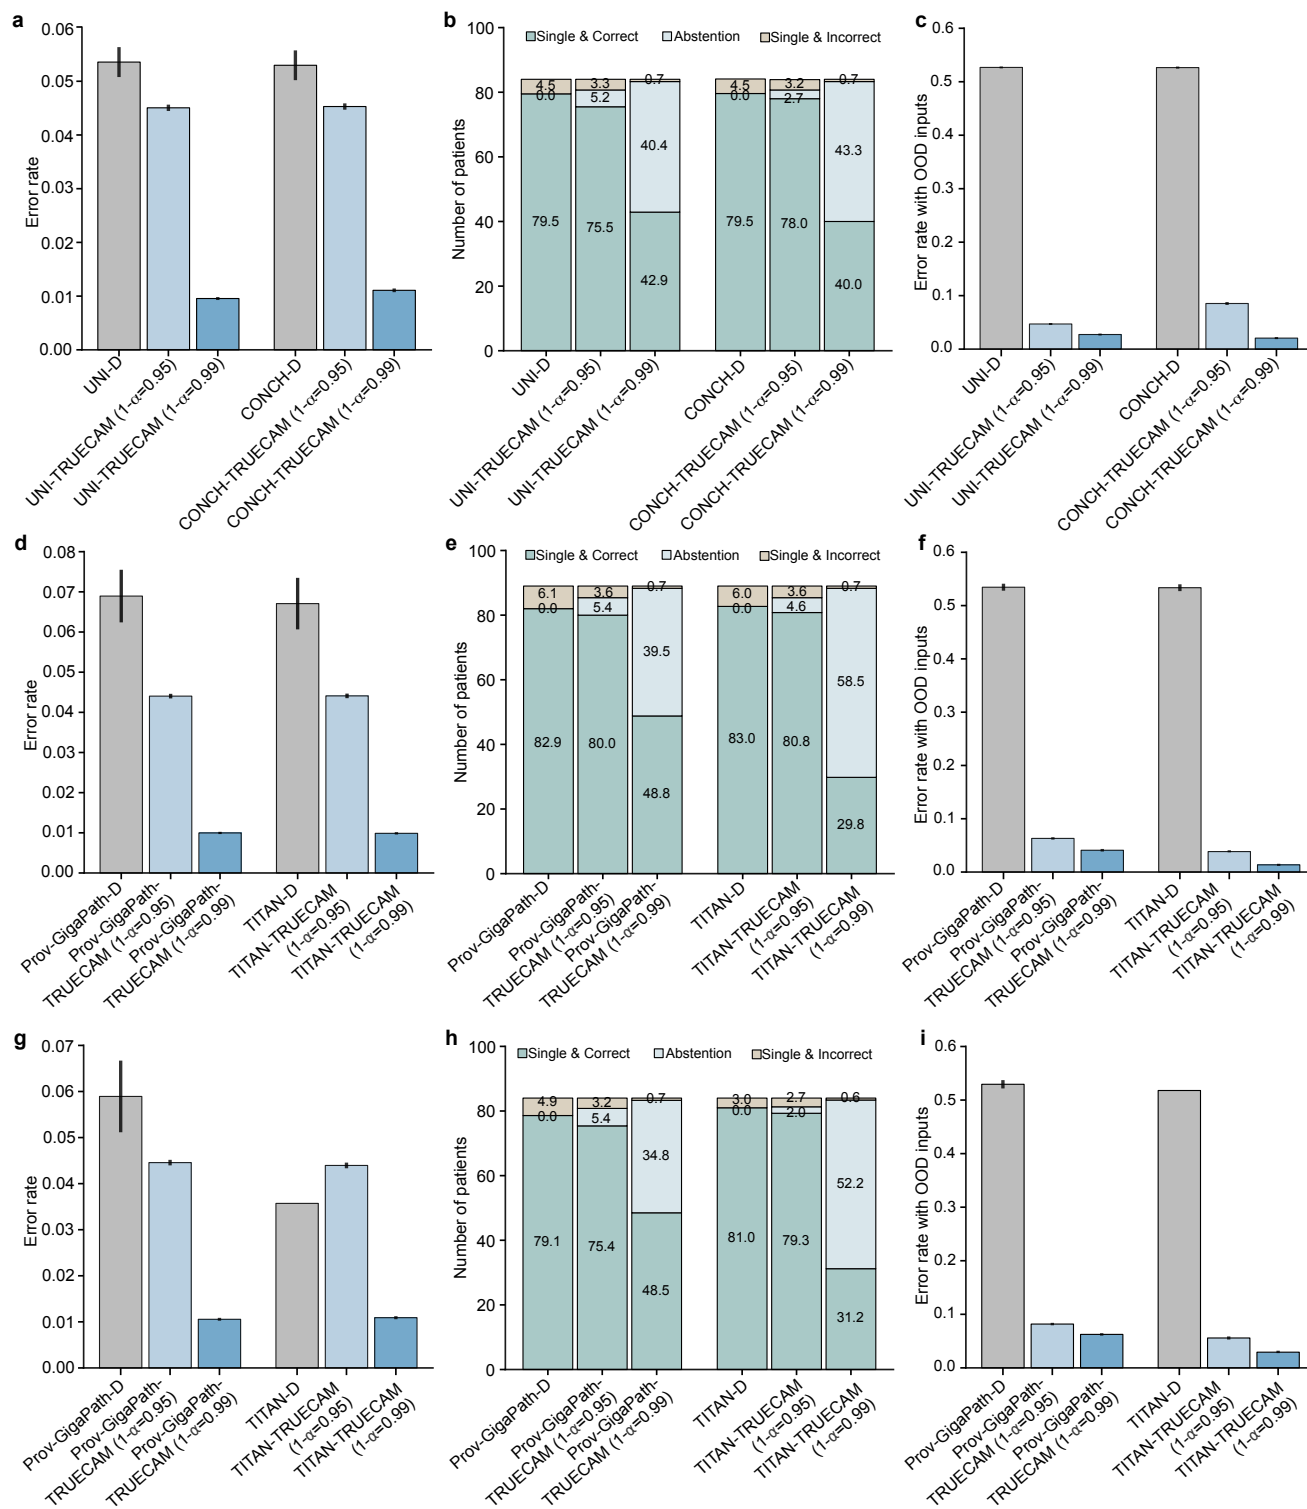
References

1. Bera, K., Schalper, K. A., Rimm, D. L., Velcheti, V. & Madabhushi, A. Artificial intelligence in digital pathology—new tools for diagnosis and precision oncology. *Nat. Rev. Clin. Oncol.* **16**, 703–715 (2019).
2. Rajpurkar, P., Chen, E., Banerjee, O. & Topol, E. J. AI in health and medicine. *Nat. Medicine* **28**, 31–38 (2022).
3. Carrillo-Perez, F. *et al.* Generation of synthetic whole-slide image tiles of tumours from rna-sequencing data via cascaded diffusion models. *Nat. Biomed. Eng.* 1–13 (2024).
4. Thirunavukarasu, A. J. *et al.* Large language models in medicine. *Nat. Medicine* **29**, 1930–1940 (2023).
5. Dvijotham, K. *et al.* Enhancing the reliability and accuracy of ai-enabled diagnosis via complementarity-driven deferral to clinicians. *Nat. Medicine* **29**, 1814–1820 (2023).
6. Begoli, E., Bhattacharya, T. & Kusnezov, D. The need for uncertainty quantification in machine-assisted medical decision making. *Nat. Mach. Intell.* **1**, 20–23 (2019).
7. Chua, M. *et al.* Tackling prediction uncertainty in machine learning for healthcare. *Nat. Biomed. Eng.* **7**, 711–718 (2023).
8. Banerji, C. R., Chakraborti, T., Harbron, C. & MacArthur, B. D. Clinical AI tools must convey predictive uncertainty for each individual patient. *Nat. Medicine* 1–3 (2023).
9. Gal, Y. & Ghahramani, Z. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, 1050–1059 (PMLR, 2016).
10. Abdar, M., Khosravi, A., Islam, S. M. S., Acharya, U. R. & Vasilakos, A. V. The need for quantification of uncertainty in artificial intelligence for clinical data analysis: increasing the level of trust in the decision-making process. *IEEE Syst. Man, Cybern. Mag.* **8**, 28–40 (2022).
11. Kompa, B., Snoek, J. & Beam, A. L. Second opinion needed: communicating uncertainty in medical machine learning. *NPJ Digit. Medicine* **4**, 4 (2021).
12. Luo, Y., Liu, Y. & Peng, J. Calibrated geometric deep learning improves kinase–drug binding predictions. *Nat. Mach. Intell.* 1–12 (2023).
13. Olsson, H. *et al.* Estimating diagnostic uncertainty in artificial intelligence assisted pathology using conformal prediction. *Nat. Commun.* **13**, 7761 (2022).
14. Dolezal, J. M. *et al.* Uncertainty-informed deep learning models enable high-confidence predictions for digital histopathology. *Nat. Commun.* **13**, 6572 (2022).

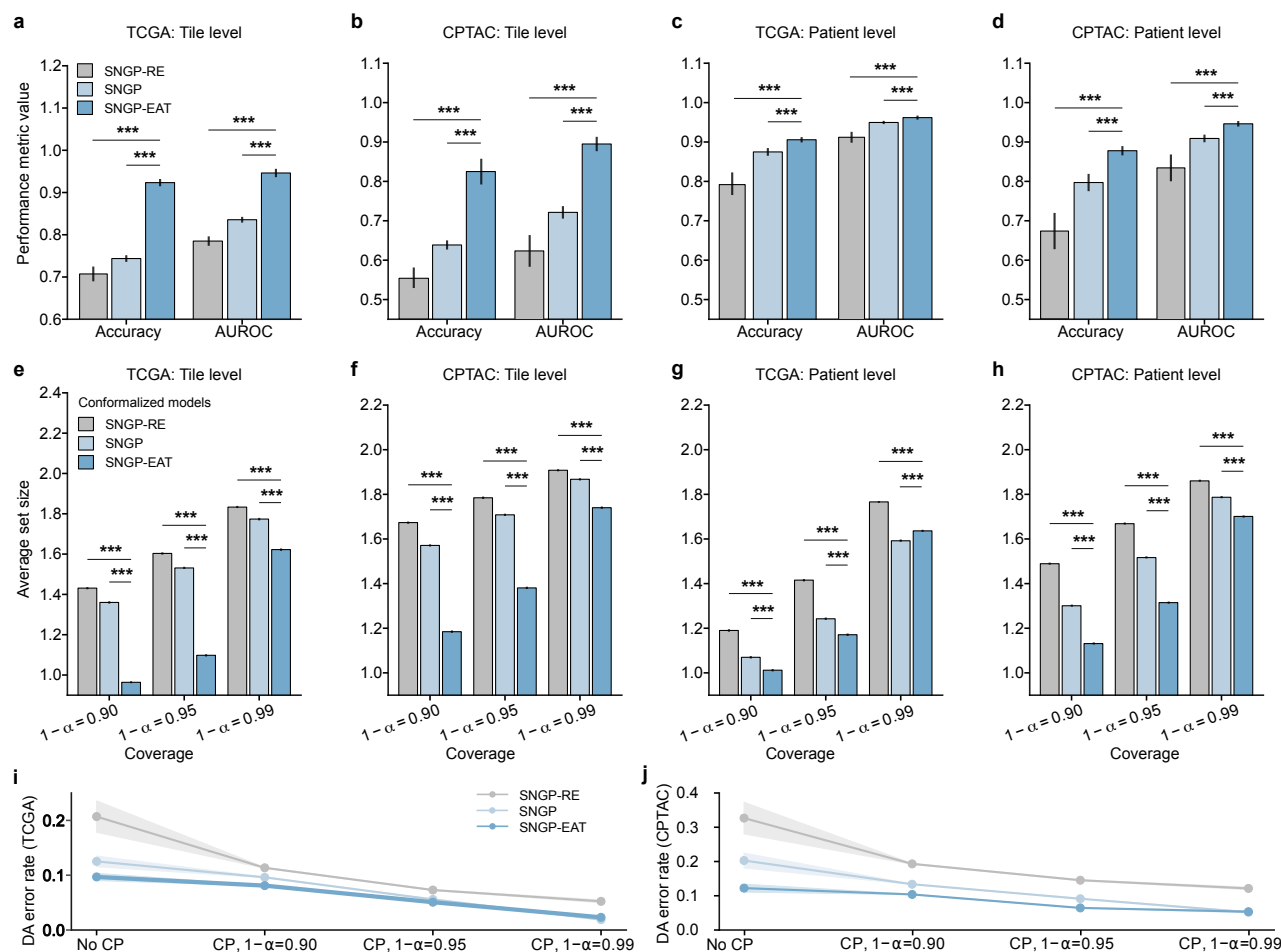
15. Sun, E. D., Ma, R., Navarro Negredo, P., Brunet, A. & Zou, J. TISSUE: uncertainty-calibrated prediction of single-cell spatial transcriptomics improves downstream analyses. *Nat. Methods* 1–11 (2024).
16. Tran, D. *et al.* Plex: Towards reliability using pretrained large model extensions. In *First Workshop on Pre-training: Perspectives, Pitfalls, and Paths Forward at ICML 2022* (2022).
17. Niculescu-Mizil, A. & Caruana, R. Predicting good probabilities with supervised learning. In *Proceedings of The 22nd International Conference on Machine Learning*, 625–632 (2005).
18. Kuleshov, V., Fenner, N. & Ermon, S. Accurate uncertainties for deep learning using calibrated regression. In *International Conference on Machine Learning*, 2796–2804 (PMLR, 2018).
19. Palmer, G. *et al.* Calibration after bootstrap for accurate uncertainty quantification in regression models. *NPJ Comput. Mater.* **8**, 115 (2022).
20. Guo, C., Pleiss, G., Sun, Y. & Weinberger, K. Q. On calibration of modern neural networks. In *International Conference on Machine Learning*, 1321–1330 (PMLR, 2017).
21. Ding, Z., Han, X., Liu, P. & Niethammer, M. Local temperature scaling for probability calibration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6889–6899 (2021).
22. Wilson, A. G. & Izmailov, P. Bayesian deep learning and a probabilistic perspective of generalization. *Adv. Neural Inf. Process. Syst.* **33**, 4697–4708 (2020).
23. Travis, W. D. *et al.* International association for the study of lung cancer/american thoracic society/european respiratory society international multidisciplinary classification of lung adenocarcinoma. *J. Thorac. Oncol.* **6**, 244–285 (2011).
24. Szegedy, C. *et al.* Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1–9 (2015).
25. Kers, J. *et al.* Deep learning-based classification of kidney transplant pathology: a retrospective, multicentre, proof-of-concept study. *The Lancet Digit. Heal.* **4**, e18–e26 (2022).
26. Chen, R. J. *et al.* Towards a general-purpose foundation model for computational pathology. *Nat. Medicine* 1–13 (2024).
27. Lu, M. Y. *et al.* A visual-language foundation model for computational pathology. *Nat. Medicine* **30**, 863–874 (2024).
28. Xu, H. *et al.* A whole-slide foundation model for digital pathology from real-world data. *Nature* 1–8 (2024).
29. Ding, T. *et al.* Multimodal whole slide foundation model for pathology. *arXiv preprint arXiv:2411.19666* (2024).
30. Shafer, G. & Vovk, V. A tutorial on conformal prediction. *J. Mach. Learn. Res.* **9** (2008).
31. Balasubramanian, V., Ho, S.-S. & Vovk, V. *Conformal prediction for reliable machine learning: theory, adaptations and applications* (Newnes, 2014).
32. Romano, Y., Patterson, E. & Candes, E. Conformalized quantile regression. *Adv. Neural Inf. Process. Syst.* **32** (2019).
33. Coudray, N. *et al.* Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning. *Nat. Medicine* **24**, 1559–1567 (2018).
34. Campanella, G. *et al.* Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nat. Medicine* **25**, 1301–1309 (2019).

35. Lu, M. Y. *et al.* Data-efficient and weakly supervised computational pathology on whole-slide images. *Nat. Biomed. Eng.* **5**, 555–570 (2021).
36. Chen, C.-L. *et al.* An annotation-free whole-slide training approach to pathological classification of lung cancer types using deep learning. *Nat. Commun.* **12**, 1193 (2021).
37. Claudio Quiros, A. *et al.* Mapping the landscape of histomorphological cancer phenotypes using self-supervised learning on unannotated pathology slides. *Nat. Commun.* **15**, 4596 (2024).
38. Jiang, R. *et al.* A transformer-based weakly supervised computational pathology method for clinical-grade diagnosis and molecular marker discovery of gliomas. *Nat. Mach. Intell.* **6**, 876–891 (2024).
39. Shahapure, K. R. & Nicholas, C. Cluster quality analysis using silhouette score. In *2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA)*, 747–748 (IEEE, 2020).
40. Ilse, M., Tomczak, J. & Welling, M. Attention-based deep multiple instance learning. In *International Conference on Machine Learning*, 2127–2136 (PMLR, 2018).
41. Erickson, N. *et al.* Autogluon-tabular: Robust and accurate automl for structured data. *arXiv preprint arXiv:2003.06505* (2020).
42. Cao, F., Chen, Q., Xing, Y. & Liang, J. Efficient classification by removing Bayesian confusing samples. *IEEE Transactions on Knowl. Data Eng.* **36**, 1084–1098 (2023).
43. Liang, W. *et al.* Advances, challenges and opportunities in creating data for trustworthy AI. *Nat. Mach. Intell.* **4**, 669–677 (2022).
44. Bernhardt, M. *et al.* Active label cleaning for improved dataset quality under resource constraints. *Nat. Commun.* **13**, 1161 (2022).
45. Hager, P. *et al.* Evaluation and mitigation of the limitations of large language models in clinical decision-making. *Nat. Medicine* 1–10 (2024).
46. Pfohl, S. R. *et al.* A toolbox for surfacing health equity harms and biases in large language models. *Nat. Medicine* 1–11 (2024).
47. Lu, M. Y. *et al.* Visual language pretrained multiple instance zero-shot transfer for histopathology images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19764–19775 (2023).
48. Chen, R. J. *et al.* Algorithmic fairness in artificial intelligence for medicine and healthcare. *Nat. Biomed. Eng.* **7**, 719–742 (2023).
49. Oquab, M. *et al.* DINOv2: Learning robust visual features without supervision. *Transactions on Mach. Learn. Res. J.* 1–31 (2024).
50. Liu, J. Z. *et al.* A simple approach to improve single-model deep uncertainty via distance-awareness. *J. Mach. Learn. Res.* **24**, 42–1 (2023).
51. Moor, M. *et al.* Foundation models for generalist medical artificial intelligence. *Nature* **616**, 259–265 (2023).
52. Huang, Z., Bianchi, F., Yuksekgonul, M., Montine, T. J. & Zou, J. A visual–language foundation model for pathology image analysis using medical twitter. *Nat. Medicine* **29**, 2307–2316 (2023).
53. Sun, Y. *et al.* Pathasst: A generative foundation ai assistant towards artificial general intelligence of pathology. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, 5034–5042 (2024).
54. Dolezal, J. M. *et al.* Slideflow: deep learning for digital histopathology with real-time whole-slide visualization. *BMC bioinformatics* **25**, 134 (2024).

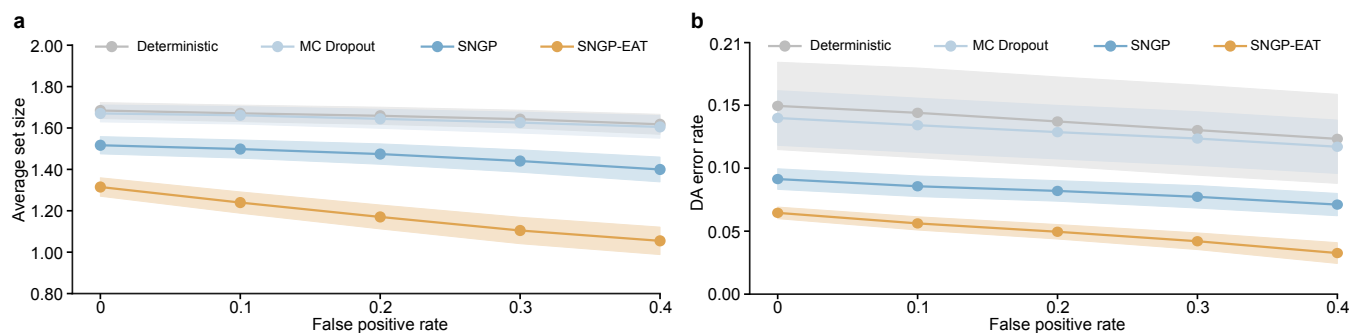
55. Otsu, N. *et al.* A threshold selection method from gray-level histograms. *Automatica* **11**, 23–27 (1975).
56. Liu, J. *et al.* Simple and principled uncertainty estimation with deterministic deep learning via distance awareness. *Adv. Neural Inf. Process. Syst.* **33**, 7498–7512 (2020).
57. Van Amersfoort, J., Smith, L., Teh, Y. W. & Gal, Y. Uncertainty estimation using a single deep deterministic neural network. In *International Conference on Machine Learning*, 9690–9700 (PMLR, 2020).
58. O’Searcoid, M. *Metric spaces* (Springer Science & Business Media, 2006).
59. Behrmann, J., Grathwohl, W., Chen, R. T., Duvenaud, D. & Jacobsen, J.-H. Invertible residual networks. In *International Conference on Machine Learning*, 573–582 (PMLR, 2019).
60. Gouk, H., Frank, E., Pfahringer, B. & Cree, M. J. Regularisation of neural networks by enforcing Lipschitz continuity. *Mach. Learn.* **110**, 393–416 (2021).
61. Williams, C. K. & Rasmussen, C. E. *Gaussian processes for machine learning*, vol. 2 (MIT press Cambridge, MA, 2006).
62. Rahimi, A. & Recht, B. Random features for large-scale kernel machines. *Adv. Neural Inf. Process. Syst.* **20** (2007).
63. Angelopoulos, A. N., Bates, S., Fannjiang, C., Jordan, M. I. & Zrnic, T. Prediction-powered inference. *Science* **382**, 669–674 (2023).
64. Lei, J., G’Sell, M., Rinaldo, A., Tibshirani, R. J. & Wasserman, L. Distribution-free predictive inference for regression. *J. Am. Stat. Assoc.* **113**, 1094–1111 (2018).
65. Angelopoulos, A. N. & Bates, S. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint arXiv:2107.07511* (2021).
66. Loshchilov, I. & Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations* (2019).
67. Angelopoulos, A. N., Bates, S., Fisch, A., Lei, L. & Schuster, T. Conformal risk control. *arXiv preprint arXiv:2208.02814* (2022).
68. Ktena, I. *et al.* Generative models improve fairness of medical classifiers under distribution shifts. *Nat. Medicine* 1–8 (2024).



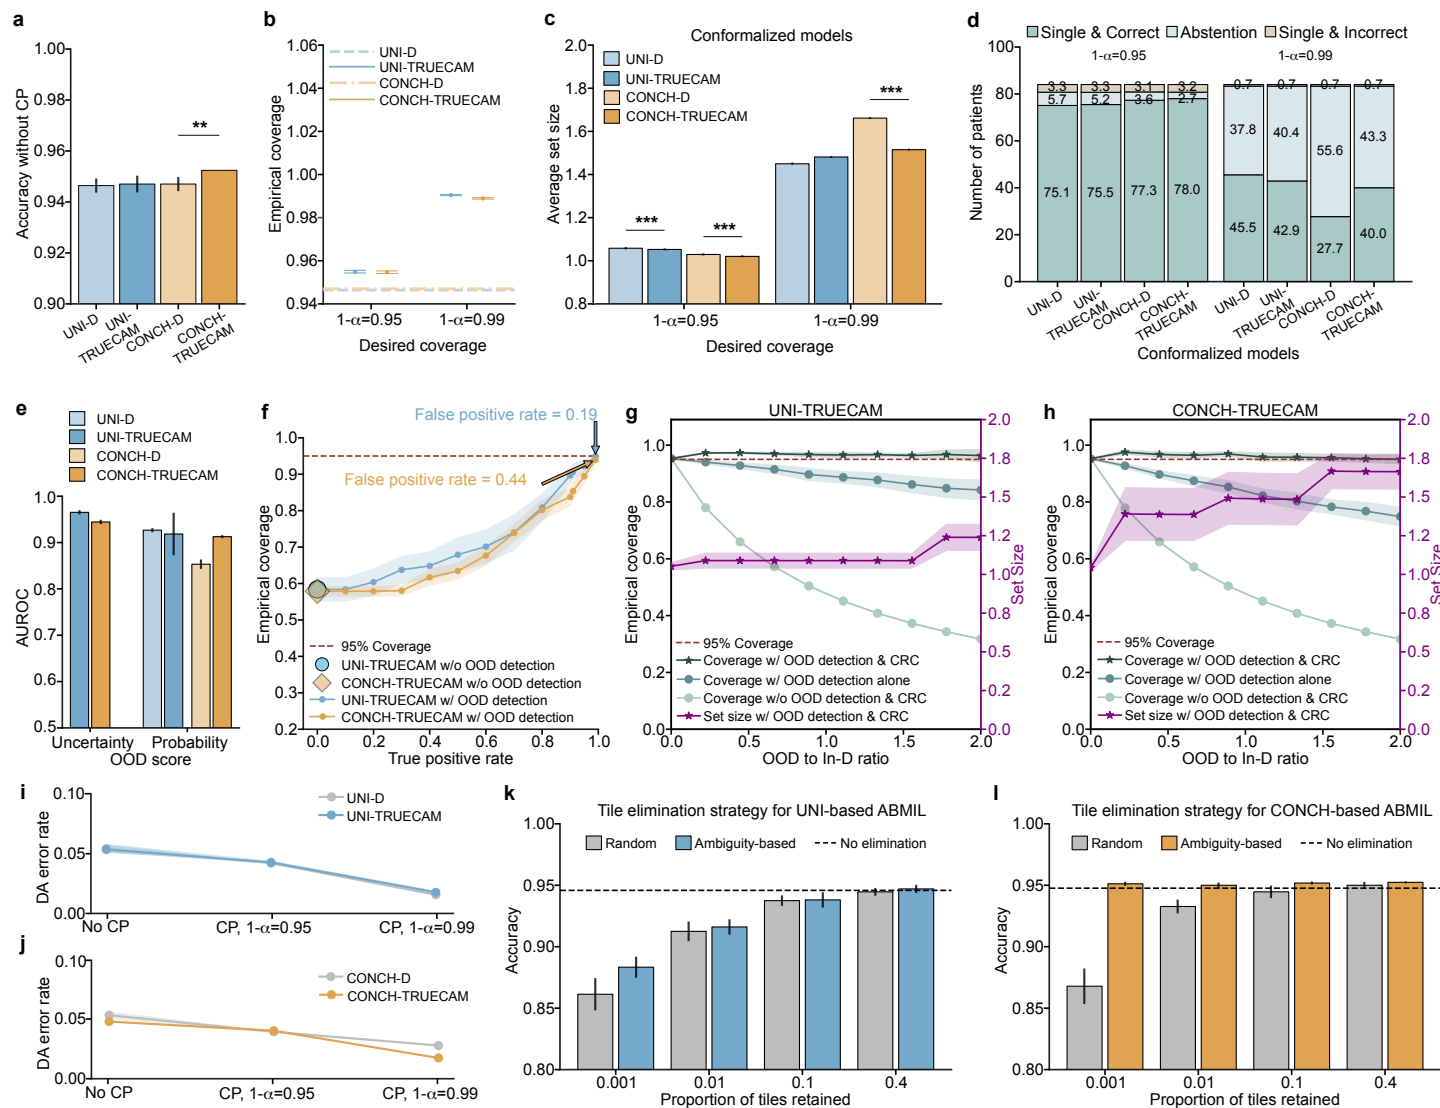
Extended Data Figure 1: Evaluation of TRUECAM’s performance across four foundation models: UNI, CONCH, Prov-GigaPath, and TITAN. **a, d, g,** TRUECAM significantly reduced NSCLC subtyping error rates across model types (denoted using suffix “-TRUECAM”) compared to their original deterministic versions (denoted using suffix “-D”), adhering to the pre-specified true label coverage $1-\alpha$. The only exception is TITAN-TRUECAM with $1-\alpha=0.95$, where TITAN-D achieved an error rate below 0.05. As a result, CP produced a zero prediction set size (no prediction) to maintain the desired coverage of 0.95. **b, e, h,** Patient-level classification breakdown for models with and without TRUECAM. **c, f, i,** TRUECAM achieved significantly lower error rates in real-world NSCLC subtyping scenarios involving a 1:1 mix of in-domain and out-of-domain inputs. Evaluations in **a-c** and **g-i** are based on the CPTAC testing data, whereas those in **d-f** are based on the TCGA testing data. Results shown in **a-c** are based on 20 independently trained models, each with 500 conformal prediction evaluations, and those in **d-i** are based on 20 independently trained models. OOD, out-of-domain; D, Deterministic.



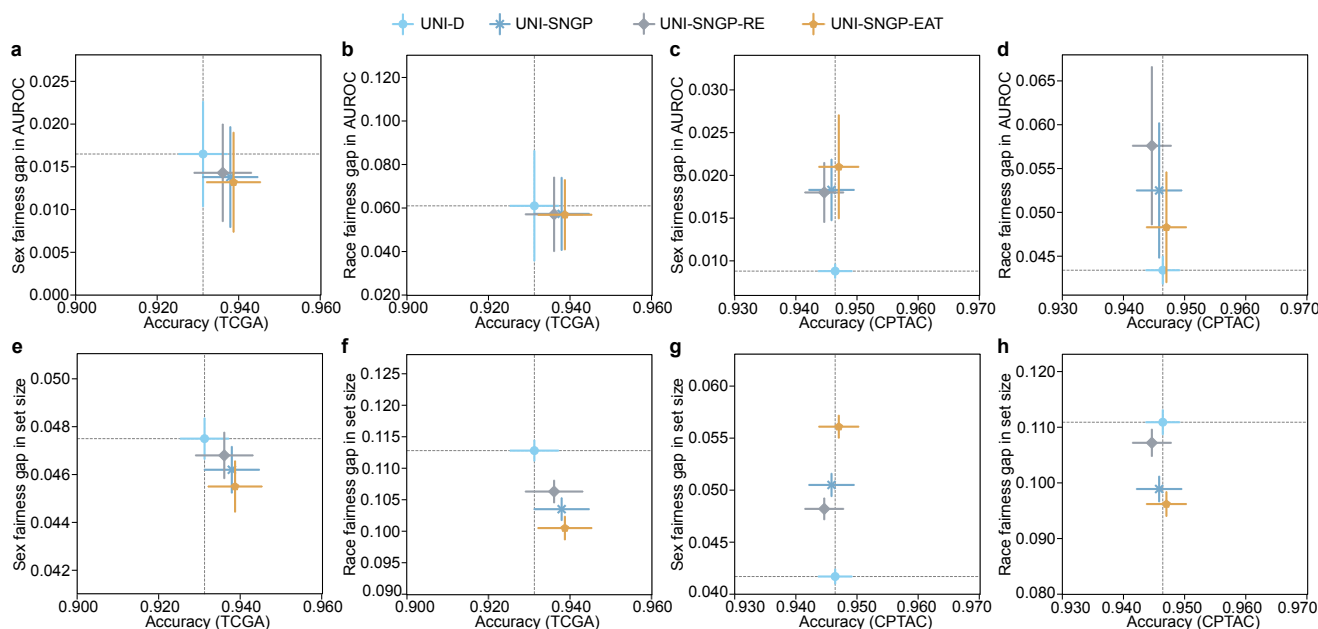
Extended Data Figure 2: Impact assessment of EAT on classification and CP performance. **a**, Tile-level classification performance on TCGA. **b**, Tile-level classification performance on CPTAC. **c**, Patient-level classification performance on TCGA. **d**, Patient-level classification performance on CPTAC. **e**, Tile-level prediction set size on TCGA for three distinct values of significance level α . **f**, Tile-level prediction set size on CPTAC. **g**, Patient-level prediction set size on TCGA. **h**, Patient-level prediction set size on CPTAC. **i**, Patient-level DA error rate on TCGA before and after activating CP. **j**, Patient-level DA error rate on CPTAC before and after activating CP. Mean values and the corresponding 95% confidence intervals in **a-j** are based on 20 independently trained models, each with 500 CP evaluations (See details in the Methods). One-sided Wilcoxon tests are utilized to calculate p values. *** $p < 0.001$. SNGP, spectral-normalized neural Gaussian process; RE, random elimination; EAT, elimination of ambiguous tiles; DA, definitive-answer.



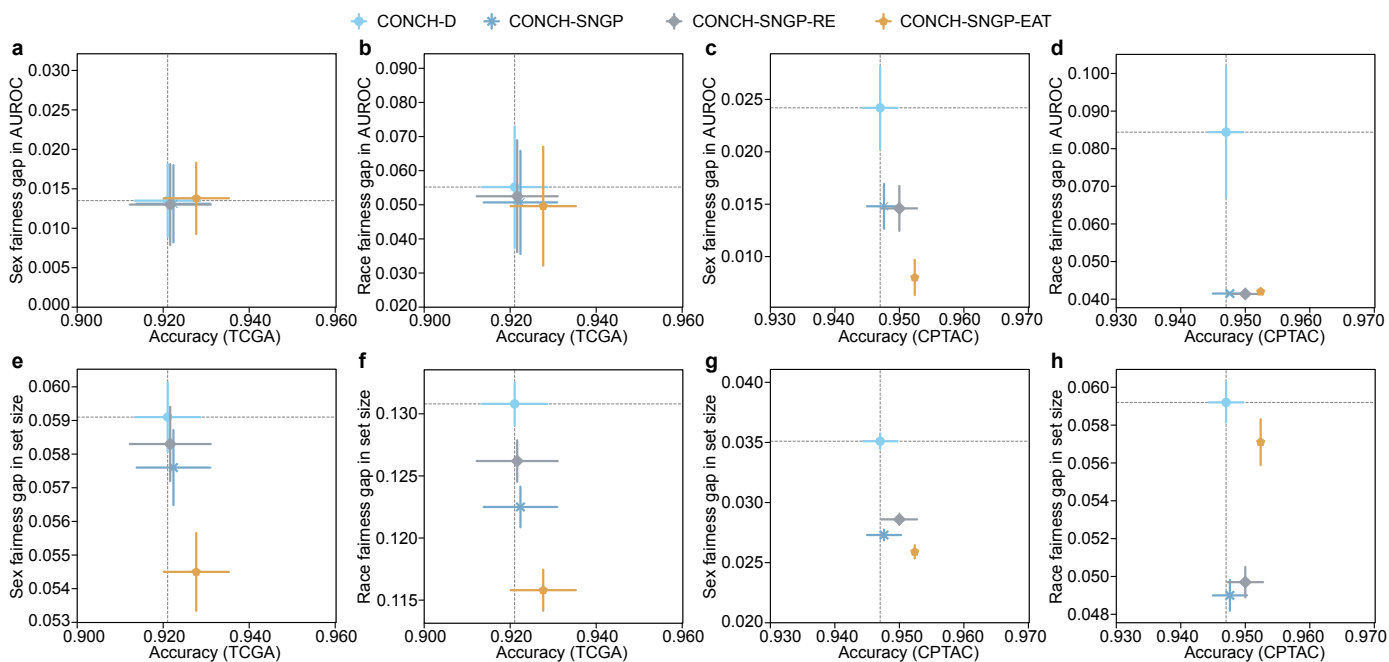
Extended Data Figure 3. Impact of OOD detection-based distribution shift control on the average set size of CP and the definitive-answer error rate in the CPTAC testing dataset. **a**, Average set size as a function of false positive rate for OOD detection after applying distribution shift control. **b**, Definitive-answer error rate as a function of false positive rate for OOD detection after applying distribution shift control. Mean values and the corresponding 95% confidence intervals are based on 20 independently trained models, each with 500 CP evaluations (See details in the Methods). SNGP, spectral-normalized neural Gaussian process; EAT, elimination of ambiguous tiles; DA, definitive-answer.



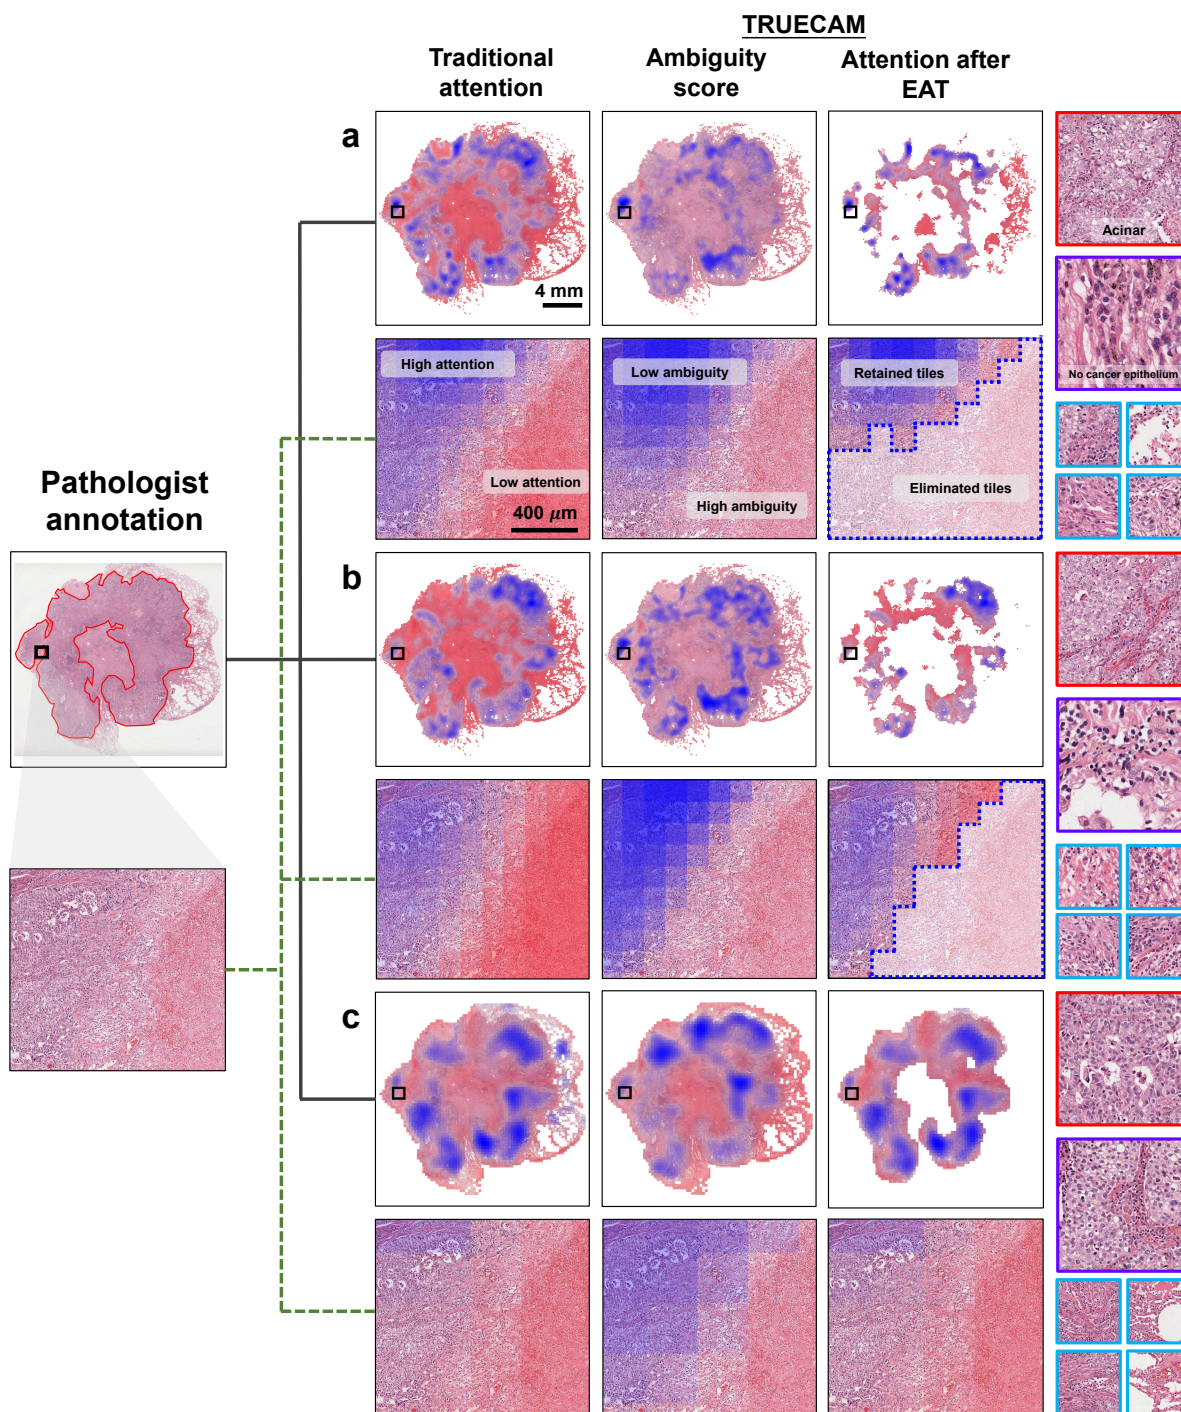
Extended Data Figure 4. Assessment of TRUEECAM using the foundation models UNI and CONCH. Results on CPTAC are reported. **a**, Classification accuracy of foundation model-based ABMIL with (denoted using suffix “-TRUEECAM”) and without TRUEECAM (i.e., the original version of foundation model-based ABMIL, denoted using suffix “-D”), evaluated without CP. **b**, Comparison of empirical coverage with and without TRUEECAM. **c**, Comparison of average set size of CP with and without TRUEECAM. **d**, Comparison of patient classification breakdown with and without TRUEECAM. **e**, Performance of two distinct OOD scores in identifying slides outside a model’s scope, which was established using the CPTAC training data. The uncertainty-based score is not applicable to the original versions of UNI and CONCH. **f**, Empirical coverage with $\alpha = 0.05$ as a function of the true positive rate for OOD detection. **g**, **h**, Empirical coverage and set size of CP for UNI-TRUEECAM and CONCH-TRUEECAM under conformal risk control with a false positive rate of 0.19 for UNI and 0.44 for CONCH, shown as a function of OOD-to-In-D ratio. **i**, **j**, Comparison of patient-level DA error rate with and without TRUEECAM in varying contexts of CP. **k**, **l**, Comparison of classification accuracy for UNI-TRUEECAM and CONCH-TRUEECAM using TRUEECAM’s ambiguity-based tile elimination strategy versus random tile elimination, shown as a function of the proportion of tiles retained per slide. One-sided Wilcoxon signed-rank test is utilized to calculate p values. ** $p < 0.01$, *** $p < 0.001$. ABMIL, attention-based multiple instance learning; CP, conformal prediction; SNGP, spectral-normalized neural Gaussian process; EAT, elimination of ambiguous tiles; DA, definitive-answer.



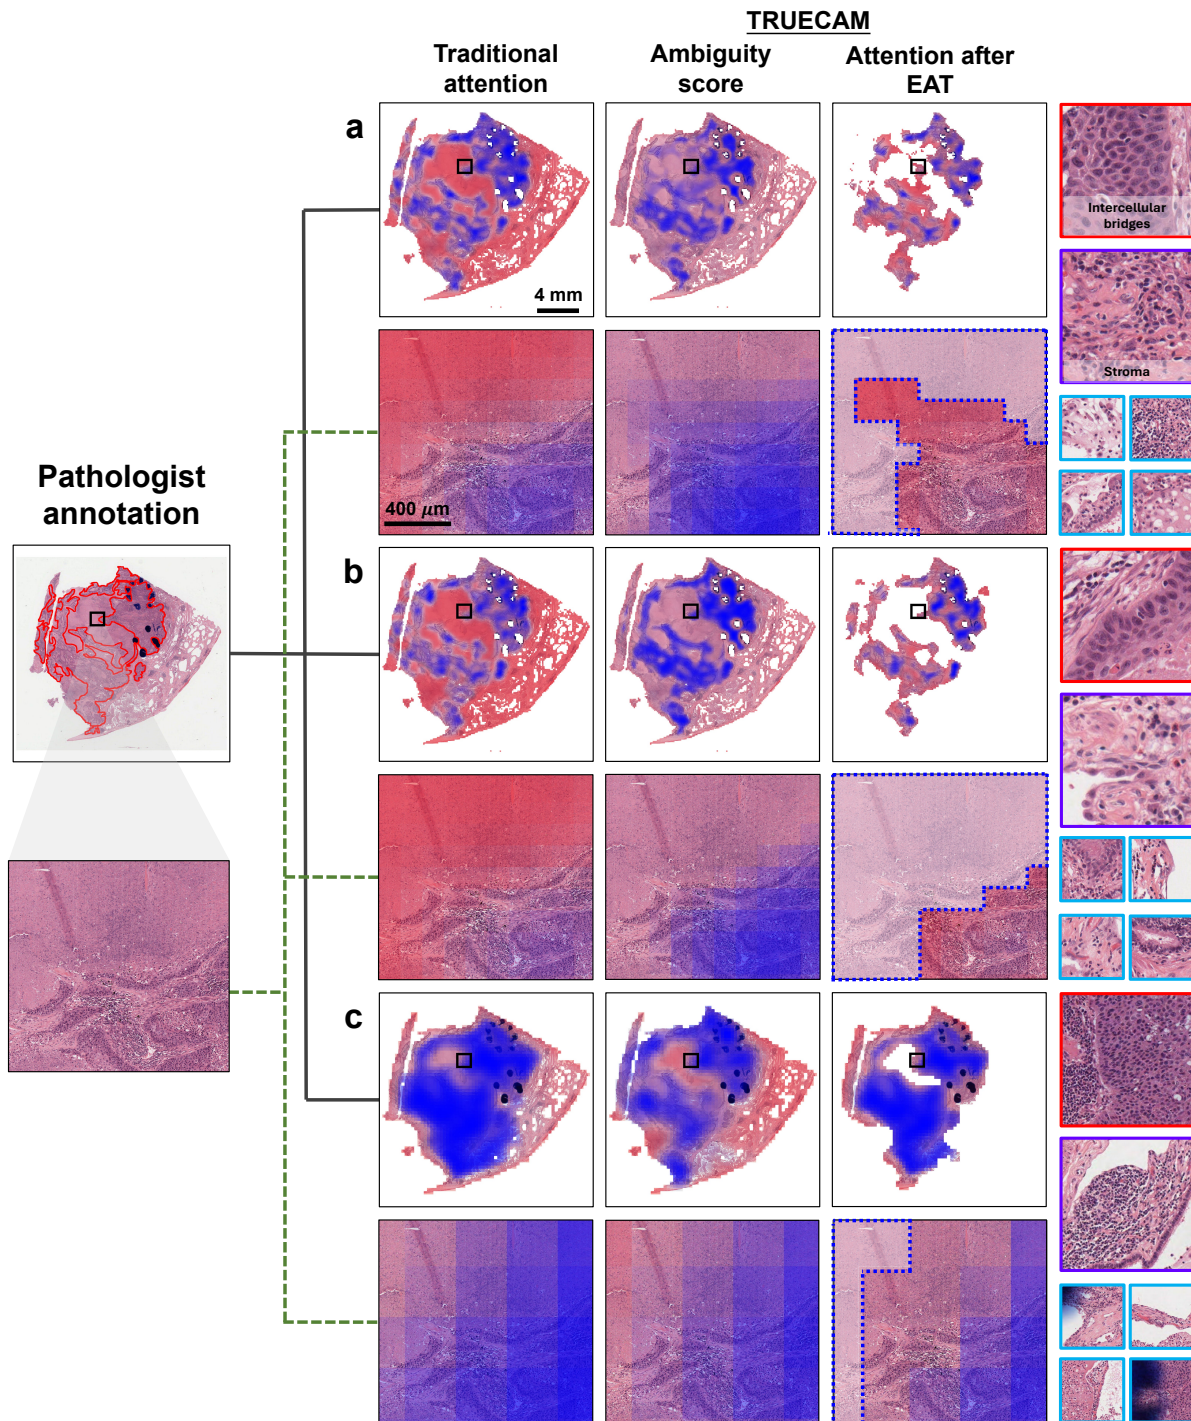
Extended Data Figure 5: Fairness evaluation of UNI-based ABMIL, measured by the metric value differences between the best- and worst-performing demographic subgroups against the overall accuracy of NSCLC subtyping. **a-d**, Sex- and race-wise AUROC gap versus overall accuracy for discriminating between LUAD and LUSC on TCGA (n=89) and CPTAC (n=84) with CP deactivated. **e-h**, At $\alpha = 0.05$, sex- and race-wise set size gap versus overall accuracy on TCGA and CPTAC with CP activated. Evaluations on UNI-D and its conformalized version are marked as an overall baseline.



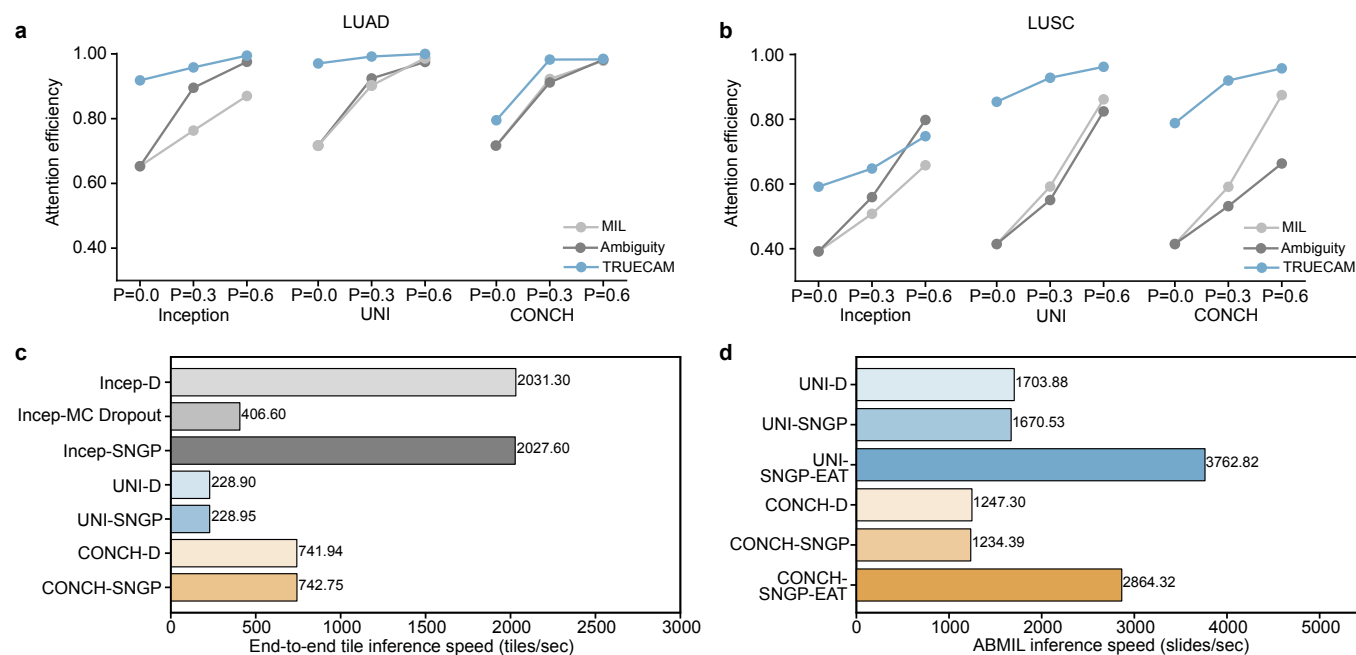
Extended Data Figure 6: Fairness evaluation of CONCH-based ABMIL, measured by the metric value differences between the best- and worst-performing demographic subgroups against the overall accuracy of NSCLC subtyping. **a-d**, Sex- and race-wise AUROC gap versus overall accuracy for discriminating between LUAD and LUSC on TCGA (n=89) and CPTAC (n=84) with CP deactivated. **e-h**, At $\alpha = 0.05$, sex- and race-wise set size gap versus overall accuracy on TCGA and CPTAC with CP activated. Evaluations on CONCH-D and its conformalized version are marked as an overall baseline.



Extended Data Figure 7: Visualization of a correctly classified LUAD slide randomly selected from the TCGA testing dataset, accompanied by a pathologist’s annotations and model-derived information supporting the inference. The whole-slide image (WSI) and a selected patch (with a zoomed-in view) are visualized in the context of **a**, UNI, **b**, CONCH, and **c**, Inception-v3. The red-highlighted region in the original WSI (first column) indicates the pathologist’s annotation used to guide NSCLC subtyping. Traditional interpretability (second column) was generated from the attention weights of the ABMIL module in foundation models or from the predicted tile-level subtype probabilities in Inception-v3. Tile-level ambiguity scores (third column) were derived from the UQ module or UQ-informed subtype probabilities in TRUECAM. After applying ambiguity-guided EAT, the attention map (fourth column) was calculated for inference using the SNGP-EAT version of each model with white regions indicating the eliminated tiles. Example tiles (all from the selected patch) visualized in the fifth column include: a randomly selected tile with low ambiguity and a high-attention weight (red border), a tile with low ambiguity and a low-attention weight (purple border), and four randomly selected tiles with high ambiguity (removed by EAT, blue border). The pathologist annotated these tiles as “good indication”, “cannot determine”, and “not relevant”, respectively, demonstrating strong alignment with the TRUECAM-derived interpretability.



Extended Data Figure 8: Visualization of a correctly classified LUSC slide randomly selected from the TCGA testing dataset, accompanied by a pathologist’s annotations and model-derived information supporting the inference. The whole-slide image (WSI) and a selected patch (with a zoomed-in view) are visualized in the context of **a**, UNI, **b**, CONCH, and **c**, Inception-v3. The red-highlighted region in the original WSI (first column) indicates the pathologist’s annotation used to guide NSCLC subtyping. Traditional interpretability (second column) was generated from the attention weights of the ABMIL module in foundation models or from the predicted tile-level subtype probabilities in Inception-v3. Tile-level ambiguity scores (third column) were derived from the UQ module or UQ-informed subtype probabilities in TRUECAM. After applying ambiguity-guided EAT, the attention map (fourth column) was calculated for inference using the SNGP-EAT version of each model with white regions indicating the eliminated tiles. Example tiles (all from the selected patch) visualized in the fifth column include: a randomly selected tile with low ambiguity and a high-attention weight (red border), a tile with low ambiguity and a low-attention weight (purple border), and four randomly selected tiles with high ambiguity (removed by EAT, blue border). The pathologist annotated these tiles as “good indication”, “cannot determine”, and “not relevant”, respectively, demonstrating strong alignment with the TRUECAM-derived interpretability.



Extended Data Figure 9: Comparison of attention efficiency and inference speed across different model settings. **a**, Comparison of attention efficiency for traditional attention (i.e., MIL), ambiguity score, and TRUECAM (i.e., attention after EAT) with respect to the LUAD slide shown in Extended Data Fig. 7. P represents the proportion of tiles with the lowest attention scores removed. In the “MIL” setting, attention scores for Inception-v3 and two foundation models correspond to the tile probability and corresponding ABMIL attention scores, respectively, whereas scores in the setting of “Ambiguity” for Inception-v3 and two foundation models correspond to the ambiguity scores from their own models and AutoGluon models, respectively. **b**, Comparison of attention efficiency for traditional attention (i.e., MIL), ambiguity score, and TRUECAM (i.e., attention after EAT) with respect to the LUSC slide shown in Extended Data Fig. 8. **c**, Comparison of end-to-end tile inference speed. **d**, Comparison of slide-level inference speed for ABMIL-based classification models using pre-extracted tile-level representations from foundation models. Inference speed was measured using an NVIDIA GeForce 4090 (24GB) with a batch size of 512 and 32-bit floating-point precision. Benchmarking excluded the preprocessing and load time with 10 warm-up iterations followed by 30 benchmark iterations. D, Deterministic; MIL, multiple instance learning; ABMIL, attention-based multiple instance learning; SNGP, spectral-normalized neural Gaussian process; EAT, elimination of ambiguous tiles.