A Multimodal Knowledge-enhanced Whole-slide Pathology Foundation Model

Yingxue Xu[1†], Yihui Wang[1†], Fengtao Zhou[1], Jiabo Ma[1], Shu Yang[1], Huangjing Lin[2], Xin Wang[3], Jiguang Wang[4,5], Li Liang[6,7,8], Anjia Han[9], Ronald Cheong Kin Chan[10], Hao Chen[1,4,5*]

[1]Department of Computer Science and Engineering , The Hong Kong University of Science and Technology, Hong Kong SAR, China.
[2]Department of Computer Science and Engineering , The Chinese University of Hong Kong, Hong Kong SAR, China.
[3]Department of Surgery, Prince of Wales Hospital, The Chinese University of Hong Kong, Hong Kong, China.
[4]Department of Chemical and Biological Engineering, The Hong Kong University of Science and Technology, Hong Kong SAR, China.
[5]Division of Life Science, The Hong Kong University of Science and Technology, Hong Kong SAR, China.
[6]Department of Pathology, Nanfang Hospital and School of Basic Medical Sciences, Southern Medical University, Guangzhou, China.
[7] Guangdong Provincial Key Laboratory of Molecular Tumor Pathology, Guangzhou, China.
[8] Jinfeng Laboratory, Chongqing, China.
[9]Department of Pathology, The First Affiliated Hospital , Sun Yat-sen University, Guangzhou, China.
[10]Department of Anatomical and Cellular Pathology, Prince of Wales Hospital,The Chinese University of Hong Kong, Hong Kong SAR, China.

*Corresponding author(s). E-mail(s): jhc@cse.ust.hk;
Contributing authors: {yxueb, ywangrm, fzhouaf, jmabq, syangcw}@connect.ust.hk;
hjlin@cse.cuhk.edu.hk; xinwang@cuhk.edu.hk; jgwang@ust.hk; lli@smu.edu.cn;
hananjia@mail.sysu.edu.cn; ronaldckchan@cuhk.edu.hk;
[†]These authors contributed equally to this work.

Abstract

Remarkable strides in computational pathology (CPath) have been made in the task-agnostic foundation model (FM) that advances the performance of a wide array of downstream clinical tasks. Despite the promising performance, there are still several challenges. First, prior works have resorted to either vision-only or vision-captions data, disregarding invaluable pathology reports and gene expression

profiles which respectively offer distinct knowledge for versatile clinical applications. Second, the current progress in pathology FMs predominantly concentrates on the patch level, where the restricted context of patch-level pretraining fails to capture whole-slide patterns. Here we curated the largest multimodal dataset consisting of H&E diagnostic whole slide images (WSIs) along with their associated pathology reports and RNA-Seq data, which resulted in 26,169 slide-level modality pairs from 10,275 patients across 32 cancer types. To leverage these data for CPath, we propose a novel whole-slide pretraining paradigm which injects multimodal knowledge at the whole-slide context into the pathology FM, called **M**ultimodal **S**elf-**TA**ught P**R**etraining (mSTAR). The proposed paradigm revolutionizes the workflow of pretraining for CPath, which enables the pathology FM to acquire the whole-slide context. To the best of our knowledge, this is the first attempt to incorporate multimodal knowledge at the slide level for enhancing pathology FMs, expanding the modelling context from unimodal to multimodal knowledge and from patch-level to slide-level. To systematically evaluate the capabilities of mSTAR, extensive experiments including slide-level unimodal and multimodal applications, are conducted across 7 diverse types of tasks on 43 subtasks, resulting in the largest spectrum of downstream tasks. The average performance in varying types of slide-level applications consistently demonstrates significant performance enhancements for mSTAR compared to other state-of-the-art FMs, with statistically critical differences observed. In particular, mSTAR showcases the promising superiority in multimodal capabilities attained through the integration of multimodal knowledge.

Keywords: foundation models, computational pathology, multimodal, pretraining

1 Introduction

The recent advancements in foundation models (FMs) [2, 3, 4, 5] for computational pathology (CPath) have demonstrated considerable progress in an incredibly broad spectrum of clinical tasks, such as cancer diagnosis, treatment and prognosis. Despite encouraging performance in general-purpose pathology foundation models, there are still several unresolved challenges.

First, massive multimodal data in line with clinical practices are under-utilized for pretraining, such as pathology reports and gene expression profiles. Existing pathology FMs either focus on vision-only [3] or vision-captions data [2, 4], in which captions are too short to offer sufficiently rich information although attempting to incorporate different modalities. The power of multimodal data has been repeatedly substantiated not only in the general machine learning community [6, 7] but also in the field of medical cancer research [8, 9, 10]. In the clinical workflow, as shown in **Figure** 1a, pathology reports often provide the complete and exhaustive information of whole slides, while patients' gene expression profiles offer insights of quantitative molecular dynamics that can complement the qualitative morphological view provided by a slide. The integration of these slide-level multimodal data can establish a broad and holistic perspective, thereby

undoubtedly enhancing the capabilities of pathology FMs to perform various clinical tasks.

Second, existing efforts in pathology FMs are primarily directed towards the modelling of patch/ROI-level data [2, 3, 4], leading to limited contexts for slide-level clinical applications. Conventional models typically treat individual patch images as independent samples for pretraining a patch extractor, and subsequently employ multiple instance learning (MIL) [11, 12, 13] to perform slide-level modelling based on embedded patch features. Recent concurrent work has attempted to pretrain a slide-level FM [5]. However, it is achieved by pretraining a slide aggregator on top of pre-extracted patch features with a fixed patch extractor. This way poses an inherent limitation that the upper bound of pretraining performance is inevitably constrained by the quality of patch features, thus limiting their generalization abilities.

In this study, we curated the largest multimodal dataset including three modalities: H&E diagnostic WSIs, pathology reports and RNA-Seq data, resulting in 26,169 slide-level modality pairs (22k used for pretraining) from 10,275 patients across 32 cancer types sourced from The Cancer Genome Atlas (TCGA) (**Figure** 1c-e). To leverage these multimodal data for CPath, we developed a novel whole-slide pretraining paradigm which injects multimodal knowledge into the
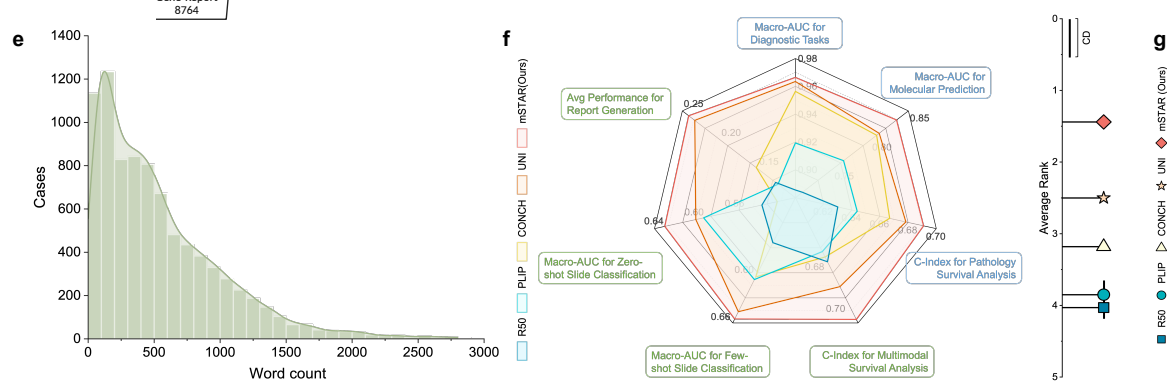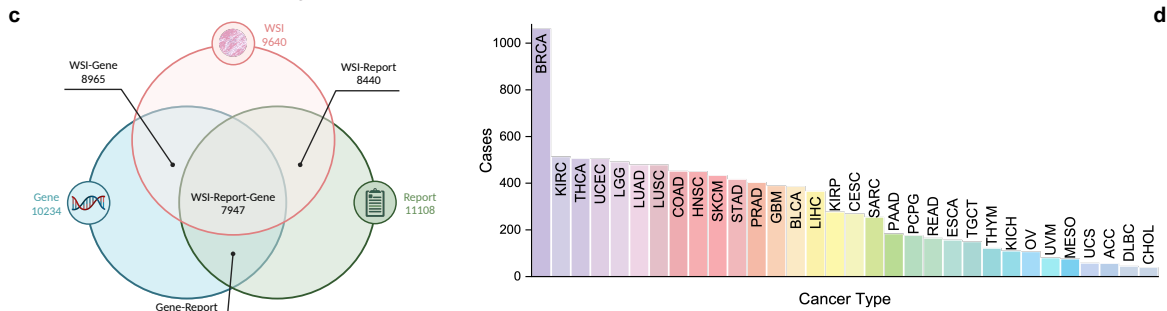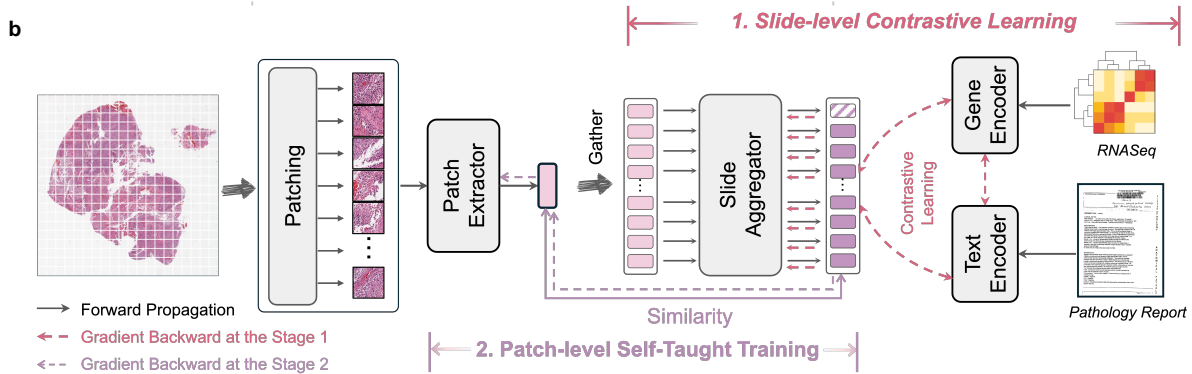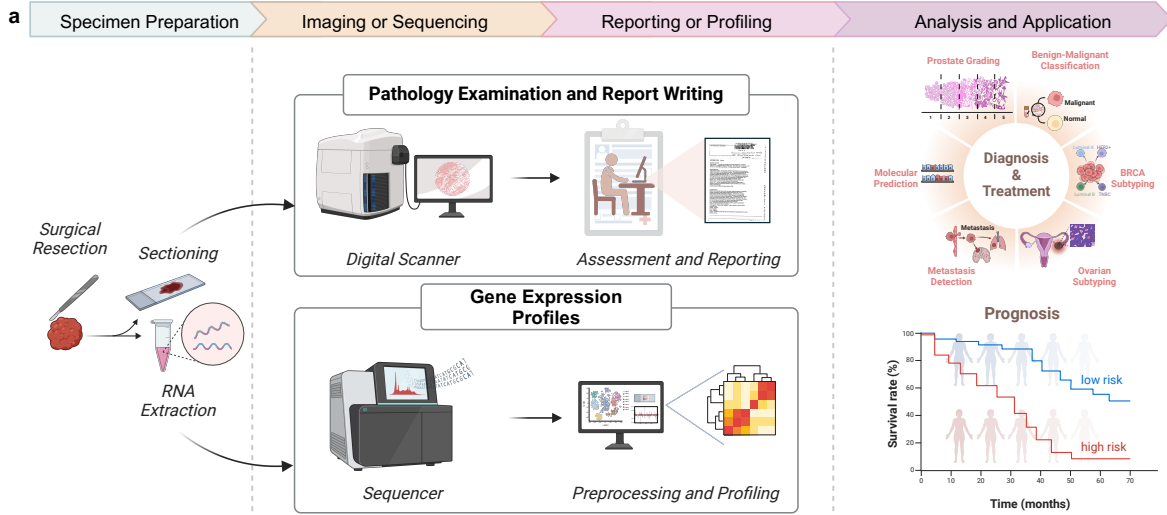
2

a Specimen Preparation | Imaging or Sequencing | Reporting or Profiling | Analysis and Application

Surgical Resection — Sectioning — RNA Extraction

Pathology Examination and Report Writing
Digital Scanner — Assessment and Reporting

Gene Expression Profiles
Sequencer — Preprocessing and Profiling

Diagnosis & Treatment: Prostate Grading, Benign-Malignant Classification, BRCA Subtyping, Ovarian Subtyping, Metastasis Detection, Molecular Prediction

Prognosis
Survival rate (%) vs Time (months): low risk, high risk

b 1. Slide-level Contrastive Learning

Patching — Patch Extractor — Gather — Slide Aggregator — Gene Encoder (RNASeq) / Text Encoder (Pathology Report)

Contrastive Learning

Similarity

2. Patch-level Self-Taught Training

→ Forward Propagation
- - → Gradient Backward at the Stage 1
- - → Gradient Backward at the Stage 2

c WSI 9640; Gene 10234; Report 11108; WSI-Gene 8965; WSI-Report 8440; WSI-Report-Gene 7947; Gene-Report 8764

d Cases vs Cancer Type: BRCA, KIRC, THCA, UCEC, LGG, LUAD, LUSC, COAD, HNSC, SKCM, STAD, PRAD, GBM, BLCA, LIHC, KIRP, CESC, SARC, PAAD, PCPG, READ, ESCA, TGCT, THYM, KICH, OV, UVM, MESO, UCS, ACC, DLBC, CHOL

e Cases vs Word count

f Radar chart: Macro-AUC for Diagnostic Tasks, Macro-AUC for Molecular Prediction, C-Index for Pathology Survival Analysis, C-Index for Multimodal Survival Analysis, Macro-AUC for Few-shot Slide Classification, Macro-AUC for Zero-shot Slide Classification, Avg Performance for Report Generation
Legend: mSTAR(Ours), UNI, CONCH, PLIP, R50

g Average Rank: mSTAR (Ours), UNI, CONCH, PLIP, R50; CD

3

Fig. 1: **Overview of the study**. **a**, The workflow in clinical practice for diagnosis, treatment and prognosis of oncology, which primarily involves three common modalities data: WSIs, pathology reports and gene expression profiles. **b**, **The overview of mSTAR paradigm**. mSTAR consists of two stages: 1) Slide-level Contrastive Learning, 2) Patch-level Self-Taught Training. **c-e**, statistics of data used in this study, including **c)** Venn Graph of cases across various modalities, **d)** the number of cases for different cancer types and **e)** the distribution of word count for pathology reports. **f**, The average performance across 7 distinct types of applications, including unimodal (blue boxes) and multimodal tasks (green boxes). **g**, Critical Difference (CD) Diagram [1] of the average rank performance across 43 diverse tasks for mSTAR and the competing pathology foundation models, in which Bonferroni-Dunn test is employed to assess whether there is a critical difference across different models on multiple datasets. In the CD diagram, a significant difference ($P < 0.001$) is indicated by the non-overlapping vertical lines of CD connecting different models.

pathology FM, termed **M**ultimodal **S**elf-**TA**ught P**R**etraining (mSTAR) (**Figure** 1b and 2), broadening the context of modelling into the slide level. To tackle the aforementioned challenges, this paradigm first pretrains a slide aggregator that absorbs multimodal knowledge via slide-level contrastive learning. This slide aggregator will act as a bridge that injects whole-slide contextual multimodal knowledge into the patch extractor at the next stage. Specifically, the slide aggregator gathers all pre-extracted patch features of an entire slide and integrates them into a slide-level representation, which then contrastively aligns with other modalities at the slide level. At the second stage, the pretrained aggregator seamlessly propagates multimodal knowledge learned at the slide-level in the first stage to the patch extractor through self-taught training. During self-taught training, the pretrained aggregator serves as the "Teacher" to supervise the pretraining of the patch extractor, which enforces patch features embedded by the patch extractor to closely resemble the ones re-embedded by the pretrained aggregator.

The proposed paradigm revolutionizes the workflow of pretraining for computational pathology, which allows the pathology FM to possess powerful whole-slide abilities. Furthermore, to the best of our knowledge, this is the first endeavor to inject multimodal knowledge at the slide-level context into pathology foundation models, broadening the scope of contextual modelling from unimodal to multimodal knowledge and from patch-level to slide-level. To systematically investigate the capabilities of mSTAR including unimodal and multimodal applications, we evaluate 7 types of applications on 43 diverse

subtasks, resulting in the largest spectrum of downstream tasks. The experimental results have demonstrated that, mSTAR showcased consistent superiority in not only unimodal but also multimodal applications on average, with statistically significant differences compared to previous state-of-the-art (SOTA) FMs (**Figure** 1f-g). These findings indicate the noteworthy power of multimodal knowledge at the whole-slide context, which endows pathology foundation models with enhanced capabilities that extend beyond the patch-level unimodal context, enabling them to achieve more generalized performance.

2 Results

2.1 The Overview of mSTAR

The proposed mSTAR aims to provide a novel whole-slide pretraining paradigm that injects multimodal knowledge into the pathology foundation model. Compared with existing pathology foundation models, mSTAR has the following innovative designs to fully unleash its power in a wide spectrum of pathological downstream tasks. First, clinical multimodal data are fully harnessed in pretraining to endow the pathology FM with multimodal knowledge for comprehensive perspectives in clinical tasks. Second, the whole-slide pretraining paradigm provides an alternative way to obtain whole-slide contexts for pathology FMs through self-taught training. To the best of our knowledge, this is the first work to inject multimodal knowledge at the whole-slide context into a pathology FM, broadening the contextual understanding for CPath from patch-level to slide-level and from unimodal to multimodal knowledge.
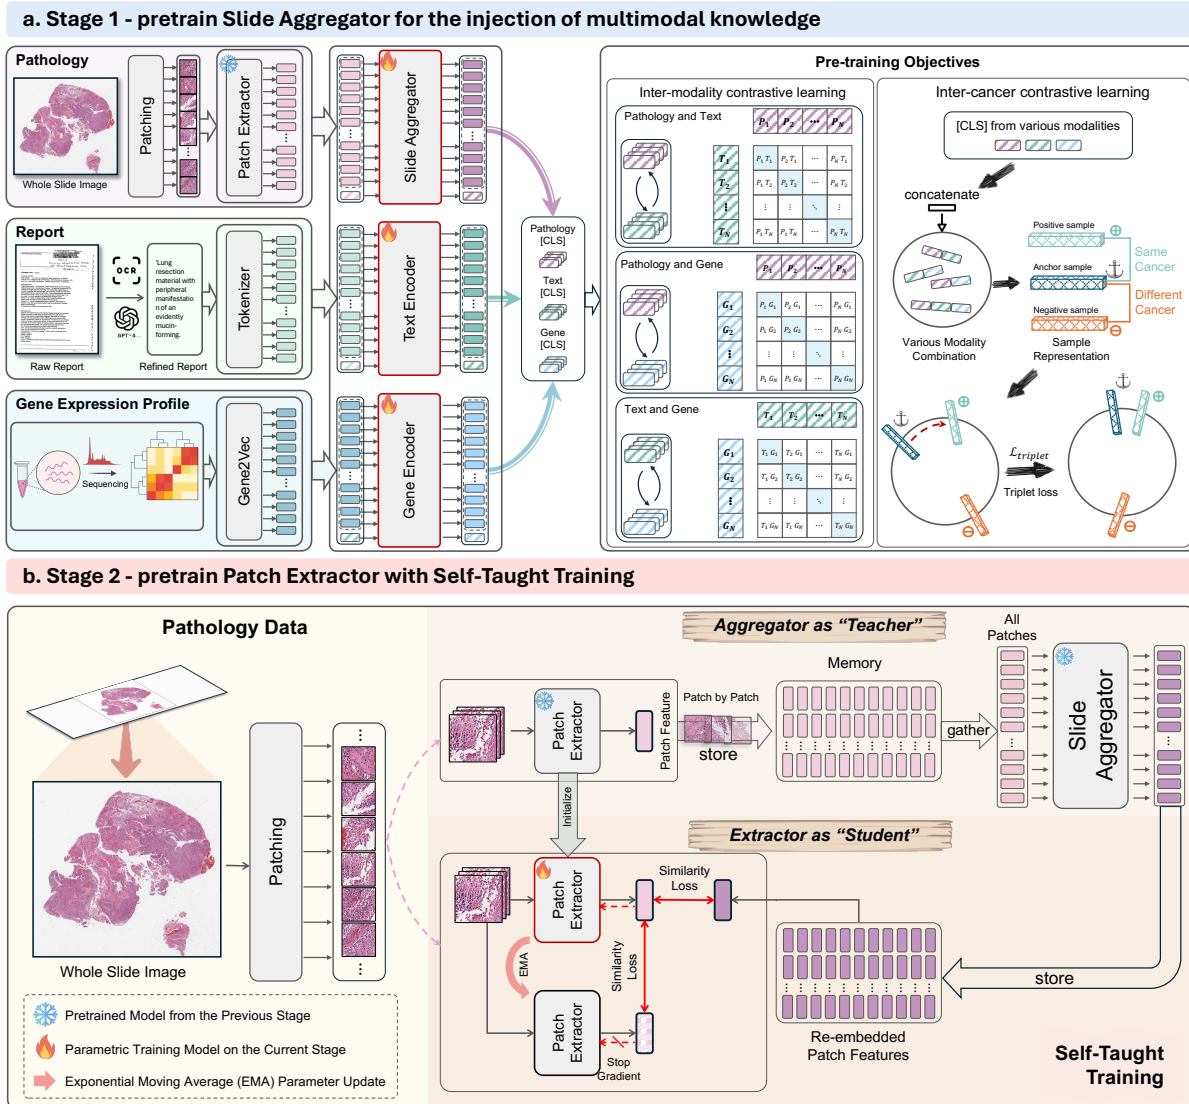
Fig. 2: **The Overview of mSTAR Pipeline**. mSTAR is a whole-slide pretraining paradigm comprising two-stage pretraining. **a**, Stage 1 aims to inject multimodal knowledge into a slide aggregator by slide-level contrastive learning among WSIs, pathology reports and gene expression data. **b**, Stage 2 aims to seamlessly propagate multimodal knowledge learned at the slide level into the patch extractor by Self-Taught training, which leverages the slide aggregator pretrained in Stage 1 as "Teacher" and enforces patch extractor to be "Student".

The overview of mSTAR is shown in **Figure 2**, consisting of two stages of pretraining.

In the first stage, the objective is to inject multimodal knowledge into the slide aggregator by slide-level contrastive learning among three modalities, i.e., WSIs, pathology reports and RNA-Seq profiles. Note that the pretrained slide aggregator will act as a bridge that propagates multimodal knowledge into the patch extractor in the next stage. To this end, as shown in **Figure 2**a, we first utilized a pretrained patch extractor, a state-of-the-art pathology foundation model named UNI [3], to encode each patch image of a slide into patch features. Then the resulting patch features are fed into a slide aggregator and integrated into a slide-level representation

which is subsequently aligned with other modalities through inter-modality contrastive learning. Furthermore, to mitigate the influence of heterogeneity across different types of cancers, the pretraining of the slide aggregator is also supervised by inter-cancer contrastive learning. This approach brings samples of the same cancer type closer together while concurrently pushing samples of different cancer types apart.

In the second stage, the pretrained slide aggregator acquiring multimodal knowledge, can serve as the teacher model to seamlessly propagate multimodal knowledge at the slide-level context into the patch extractor, called Self-Taught Training (**Figure** 2b). Specifically, the patch extractor is pretrained through encouraging the extracted patch features to be as similar as possible to those re-embedded by the pretrained aggregator. At the same time, to avoid catastrophic forgetting, we also enforce a similarity constraint between the extracted features and ones embedded by the exponential moving average (EMA) patch extractor.

With these two stages, multimodal knowledge at the whole-slide context can be seamlessly embedded into foundation models. As a result, the model acquires the ability to comprehend both patches and the entire WSI, which facilitates downstream tasks at different levels. In the end, the pathology foundation model can achieve advanced abilities with the extended context from patch-level to slide-level and from unimodal to multimodal knowledge. More details of mSTAR can be found in the Section 4.2.

2.2 Pathological Slide Classification for Diagnosis and Treatment

Slide classification holds significant relevance in clinical settings, particularly in accurate cancer diagnosis and treatment. For example, it aids in categorizing the specific subtype of cancer, which in turn guides the healthcare professionals in precise diagnosis and treatment planning. Additionally, pathologists can identify molecular markers within tissue slides, allowing them to determine the molecular subtypes of cancer. This information is of utmost importance as it forms the foundation for developing personalized treatment plans and enhancing treatment effectiveness. Therefore, the evaluation of slide-level diagnostic

tasks, such as cancer subtyping, grading, etc., is the first step of this study, followed by molecular prediction based on the analysis of tissue slides. From the perspective of technique, here we are investigating unimodal abilities through pathological slide classification. To this end, 12 diverse slide classification datasets were used for assessing the ability of decision-making diagnosis at the slide level, consisting of 6 diagnostic subtasks and 6 molecular prediction subtasks. For diagnostic tasks, the following subtasks were included: breast metastasis detection on CAMELYON [14, 15], prostate ISUP grading on PANDA [16], ovarian cancer subtyping on UBC-OCEAN [17, 18], BRCA subtyping on TCGA BRCA [19], NSCLC subtyping on TCGA NSCLC [19] and RCC subtyping on RCC-DHMC [20]. For molecular prediction, we performed the prediction of molecular subtypes ER, HER2 and PR on BCNB [21] datasets, respectively. This task also includes molecular subtypes prediction on TCGA-BRCA, TCGA-CRC and TCGA-GBMLGG datasets [19]. The details of every task are illustrated in the Section 4.3.

The evaluations are two-fold: 1) we examined patch features of different pathology foundation models in slide classification (**Figure** 3a-d), including ResNet50 (R50) [22], PLIP [2], CONCH [4], UNI [3] and our mSTAR. To obtain slide/patient-level classification decision, we adopted two commonly used MIL backbones as the slide aggregator trained from scratch, attention-based MIL (ABMIL) [11] and TransMIL [13] for integrating these extracted patch features. ABMIL is a simple yet robust MIL approach, which is usually used for evaluation in previous foundation model research [3, 4, 5]. Considering that we employed the pretrained TransMIL as the "Teacher" in the proposed method, we also present the performance of TransMIL from scratch for all extractors. 2) Since we pretrained a TransMIL in the first stage in the proposed paradigm, we additionally explore how it collaborates with the pretrained patch extractor (**Figure** 3d), leading to an advanced version (mSTAR+). We further compared the performance of different extractors between being equipped with a non-pretrained TransMIL (built from scratch) and the pretrained TransMIL+ (**Figure** 3e).

Fig. 3: **Slide Classification Performance**. **a**, Comparison of Macro-AUC between mSTAR and compared methods on 12 slide classification tasks including 6 diagnostic subtasks and 6 molecular prediction subtasks. They are evaluated with ABMIL and TransMIL trained from scratch, based on different patch extractors. "Overall" refers to the average result between ABMIL and TransMIL. **b**, The average Macro-AUC on 12 slide classification datasets. **c-d**, The slide classification performance (Macro-AUC and 95% CI) on 6 diagnostic subtasks and 6 molecular prediction subtasks with ABMIL and TransMIL, respectively. If the performance of mSTAR is the best one compared against the compared FMs, P-value would be presented. mSTAR+ represents mSTAR (patch extractor) equipped with the pretrained aggregator (TransMIL+). The colors of legends are shared across sub-figures a-d. **e**, The performance (Macro-AUC) gains of the aggregator (TransMIL) between "before" and "after" pretraining among various datasets. "Positive" values indicate performance increases compared with the ones equipped with TransMIL built from scratch, and vice versa. Results of the bar chart in orange and green background are from diagnostic tasks and molecular prediction, respectively. All of performance increases showed significant differences with $P < 0.05$.

All comparisons are based on the metric of Macro-AUC, a commonly used classification metric, which does not rely on the selection of the decision threshold and is insensitive to the sample ratio of various classes. Furthermore, to avoid the impact of potential inherent bias in various datasets and different MIL backbones, following the previous research [1], the average rank of Macro-AUC was reported as well (**Figure** 3a) and the one-sided Wilcoxon signed-rank test was

performed on various datasets to examine the statistical difference between mSTAR and other FMs.

From an overall perspective, we assessed the average rank and average Macro-AUC for mSTAR and the compared foundation models across 12 diverse subtasks. The results showed that mSTAR consistently achieved the first place across different tasks, regardless of different MIL backbones employed or when considering the overall performance. For the consistency of performance improvements, **Figure** 3c-d (**Extended Data** Table A2-A3) demonstrated that mSTAR achieved the best performance on 11 out of 12 subtasks with ABMIL and 9 out of 12 subtasks with TransMIL. When we incorporated the pretrained aggregator, TransMIL+, mSTAR+ won other SOTA FMs on 10 out of 12 subtasks (**Extended Data** Table A4).

When we delve into the specific metric of Macro-AUC, we found that mSTAR achieved the best performance with increases of $+1.3\%$ overall ($P < 0.001$) compared to the second-best model, UNI, as shown in **Figure** 3b. The upward trends of average Macro-AUC were witnessed in ABMIL and TransMIL, going from 0.905 to 0.908 ($P < 0.001$) and from 0.870 to 0.892 ($+2.2\%$, $P < 0.005$), respectively. These results indicate that mSTAR has the promising ability of slide classification in all. More detailed results for each subtask are presented as follows.

mSTAR improves Slide-level Diagnostic Tasks. Among diagnostic tasks, the consistent improvements in ABMIL become apparent, as mSTAR outperformed the second-best approach on all diagnostic datasets with significant differentiation ($P < 0.001$, **Figure** 3c). TransMIL exhibited a similar trend (**Figure** 3d), in which results on 4 out of 6 datasets indicated significant differences.

When considering the incorporation with the pretrained aggregator, mSTAR+ exhibited further advancements on various datasets and achieved the best on 5 out of 6 diagnostic tasks. For example, compared to the second-best compared FM, mSTAR+ outperformed by $+2.01\%$ on CAMELYON ($P < 0.001$). Furthermore, an interesting finding from **Figure** 3e (**Extended Data** Table A4) was that the pretrained aggregator was even able to collaborate with other patch

extractors achieving critical performance gains, although it was solely pretrained on features extracted by mSTAR. As shown in **Figure** 3e, it consistently displayed superior performance in the majority of tasks for different FMs. Among these, the pretrained aggregator achieved a notable increase of up to $+9.10\%$ ($P < 0.001$) compared to the one built from scratch, although there were slight decreases (-1.15% at most) of performance in some specific subtasks. From **Figure** 3e, we also observed that the enhancements made to CONCH, UNI, and mSTAR were not substantial, which could be attributed to the fact that their performances have already approached 100% and the room for improvement was significantly limited on these 6 subtasks.

mSTAR improves Molecular Prediction. Following the same evaluation strategy in diagnostic tasks, we also investigated its capability of molecular prediction from tissue slides.

From the perspective of consistency, mSTAR with ABMIL achieved the best performance on 5 out of 6 tasks (**Figure** 3c and **Extended Data** Table A2), in which 4 tasks exhibited a significance difference ($P < 0.001$). With TransMIL, clearer improvements (**Figure** 3d and **Extended Data** Table A3) were observed that 4 out of 6 tasks demonstrated a significant difference ($P < 0.05$). Taking a closer look at Macro-AUC, mSTAR outperformed other SOTA foundation models by $+2.23\%$ overall with a critical difference ($P < 0.001$) on average across different MIL backbones (**Extended Data** Table A2-A3). Specifically, for TransMIL, a substantial increase of $+3.71\%$ on average Macro-AUC across 6 molecular tasks was observed, although mSTAR achieved slight improvement with ABMIL ($P < 0.05$). In particular, substantial performance increases of $+4.64\%$ and $+4.21\%$ were observed on BRCA-Molecular and CRC-Molecular tasks ($P < 0.001$), respectively.

In addition, when incorporating the pretrained aggregator TransMIL+ (**Extended Data** Table A4), further improvements were achieved where mSTAR+ outperformed the compared FMs on 5 out of 6 datasets (**Figure** 3d, $P < 0.001$) by $+6.35\%$ on CRC-Molecular, $+2.93\%$ on BCNB-HER2, $+1.82\%$ on GBMLGG-IDH1 and $+1.5\%$ on BCNB-ER. From **Figure** 3e, we observed a similar trend that the majority of subtasks

demonstrated further enhancements for different extractors. An interesting finding different from that of diagnostic tasks was that performance gains on CONCH and UNI, in which there were no notable increases in diagnostic tasks, demonstrated remarkable improvements in molecular prediction (e.g., +13.10% on CRC-Molecular and +7.70% on BRCA-Molecular for CONCH, +9.00% on BRCA-Molecular and +6.29% on BCNB-ER for UNI, $P < 0.001$). This validates that the involvement of RNA-Seq data in pretraining contributes to facilitating molecular prediction, as pretraining data used in CONCH and UNI did not include gene expression profiles and our aggregator was contrastively pretrained on RNA-Seq data. It is worth noting that the pretrained aggregator provided relatively modest gains for our extractor mSTAR (+2.66% on GBMLGG-IDH1, +2.14% on CRC-Molecular and +1.73% on BCNB-HER2, $P < 0.001$). This can be explained by the fact that the knowledge regarding gene expression profiles has already been incorporated into the patch extractor mSTAR through self-taught training, which has resulted in satisfactory performance in molecular prediction.

Overall, mSTAR demonstrates consistent performance gains on both diagnostic tasks and molecular prediction at the slide-level with a significant difference, compared with prior patch-level SOTA pathology foundation models. This suggests that the modelling of whole-slide contexts is conducive to slide-level classifications. Furthermore, incorporating the pretrained aggregator is able to result in further improvements, especially on molecular prediction tasks. In particular, when using the aggregator pretrained on three modalities, those extractors that have not been exposed to gene expression profiles can acquire additional significant enhancements in molecular prediction. This can be attributed to the incorporation of RNA-Seq data during the pretraining process, thereby providing additional enhancements.

2.3 Pathological Survival Analysis for Prognosis

Prognostic analysis is an intricate clinical endeavor, necessitating a thorough analysis from a multitude of facets. In this regard, the integration of multimodal data has proven instrumental in enabling more comprehensive prognostic assessments [8, 9, 10, 23]. Therefore, it is crucial to delve into the exploration of the role of multimodal knowledge within the broader whole-slide context in enhancing prognosis estimation. In this study, we assessed a key prognostic task, overall survival (OS) prediction, on top of pathological tissue slides.

mSTAR advances Cancer Survival Prediction. We investigate the capability of predicting overall survival (OS) across 9 diverse cancers, based on the metric of C-Index, a metric commonly used in survival prediction. To avoid data leakage, the data used for evaluation were held out from pretraining data. To ensure the reliable evaluation, we excluded cancer types with fewer than 400 cases or those with too few uncensored data out of the 32 cancer datasets in TCGA [19], resulting in 9 cancer datasets: Breast Invasive Carcinoma (BRCA), Colon Adenocarcinoma and Rectum Adenocarcinoma (CRC), Glioblastoma Multiforme and Brain Lower Grade Glioma (GBMLGG), Head and Neck Squamous Cell Carcinoma (HNSC), Kidney Renal Clear Cell Carcinoma (KIRC), Lung Adenocarcinoma (LUAD), Lung Squamous Cell Carcinoma (LUSC), Skin Cutaneous Melanoma (SKCM) and Uterine Corpus Endometrial Carcinoma (UCEC). Similarly, following the settings of slide-level classification, we also used ABMIL and TransMIL for evaluating patch extractors and then the performance of the pretrained slide aggregator was assessed. More details are presented in Section 4.3.

For consistency of performance increases, mSTAR performed best compared to other foundation models, where **Figure** 4a shows that mSTAR always ranked in front of other foundation models for different MIL backbones and the overall performance. When considering each individual subtask, mSTAR achieved the best performance on 7 out of 9 datasets with ABMIL (**Figure** 4c and **Extended Data** Table A5), in which 5 tasks demonstrated a significant difference ($P < 0.01$) compared to the second-best approach. With TransMIL (**Figure** 4d and **Extended Data** Table A6), mSTAR presented more consistent improvements, which outperformed compared approaches on 8 out of 9 cancer types and
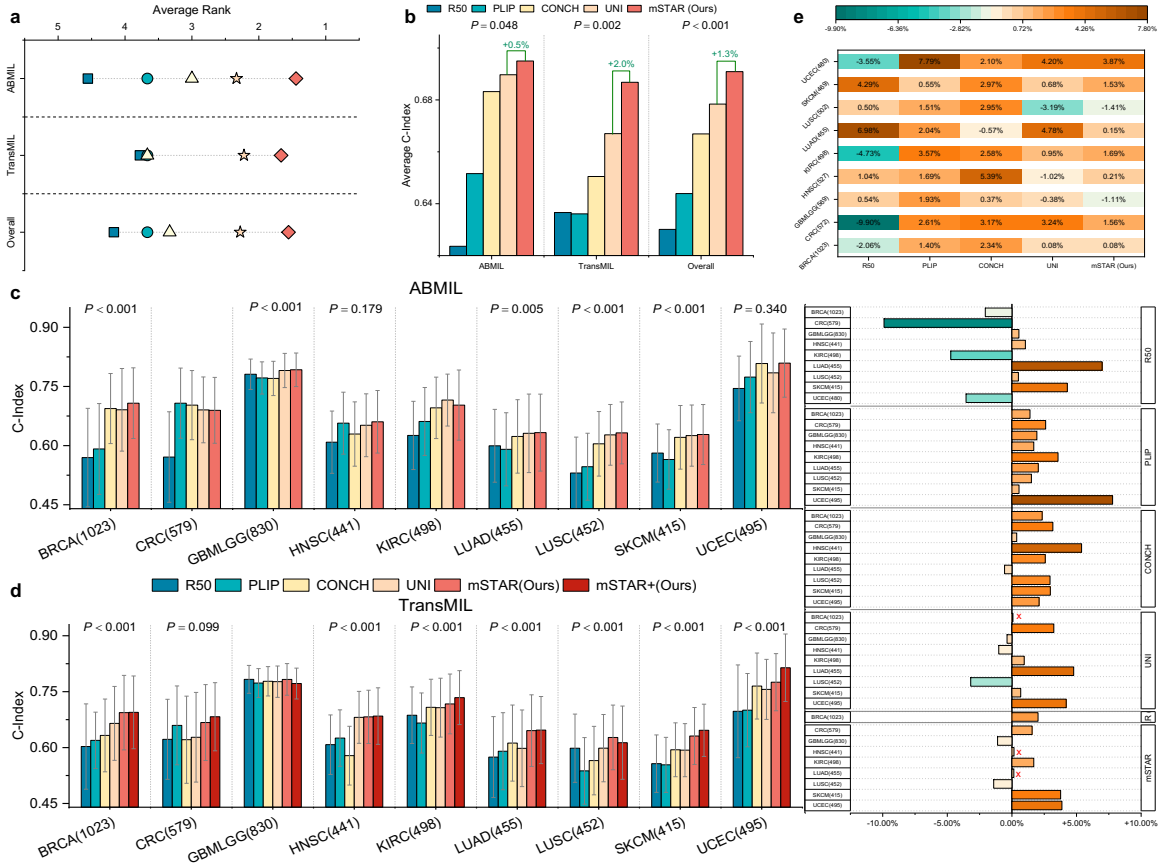
Fig. 4: **Pathological Survival Analysis Performance**. **a**, Comparison of C-Index between mSTAR and compared methods on 9 TCGA datasets. They are evaluated with ABMIL and TransMIL trained from scratch, based on different patch extractors. "Overall" refers to the average results between ABMIL and TransMIL. **b**, The average C-Index on 9 TCGA datasets. **c-d** The survival analysis performance (C-Index and 95% CI) on 9 TCGA datasets with ABMIL and TransMIL, respectively. If the performance of mSTAR is the best one compared against the compared FMs, P-value would be presented. mSTAR+ represents mSTAR (patch extractor) equipped with the pretrained aggregator (TransMIL+). The colors of legends are shared across sub-figures a-d. **e** The performance (C-Index) gains of the aggregator (TransMIL) between "before" and "after" pretraining among 9 TCGA cancer types. "Positive" values indicate performance increases compared with the ones equipped with TransMIL built from scratch, and vice versa. All of performance increases showed significant differences with $P < 0.05$, except the ones marked by ×.

showed significant improvements compared to the second-best method in 7 cancer types ($P < 0.001$).

When specifically assessing the C-Index metric, the average C-Index showcased significant performance gains compared against the second-best model (UNI) by +1.98% for TransMIL ($P < 0.005$), +0.53% for ABMIL ($P < 0.05$), and +1.30% for the average performance of different MIL backbones (Overall, $P < 0.001$). Next,

we delved into details of performances across 9 datasets, where we observed apparent improvements ($P < 0.001$) of +3.7% on SKCM, +3.34% on LUAD, +2.87% on LUSC and +1.03% on UCEC over the second-best model based on TransMIL (4d and **Extended Data** Table A6).

These results indicated the effectiveness of multimodal knowledge involved in mSTAR in extracting discriminative features for distinguishing complex survival patterns in pathological images.

We also investigated the performance of incorporating the pretrained aggregator. From observations of **Figure** 4d, mSTAR+ equipped with the pretrained aggregator pushed performance to higher levels (+5.23% on SKCM, +4.9% on UCEC, +2.96% on BRCA, +2.60% on KIRC and +2.33% on CRC, $P < 0.001$) compared to the second-performing approaches. By comparing TransMIL "before" and "after" pretraining for various extractors (**Figure** 4e, **Extended Data** Table A7), we observed that, in the majority of cases, the pretrained aggregator exhibited superior performance, surpassing the one built from scratch by a margin of up to +7.79% ($P < 0.001$). This indicates the introduction of pathology report and gene expression profiles is beneficial to survival analysis by providing additional expert knowledge and molecular information. In addition, a similar trend to that of slide classification was observed that, PLIP and CONCH that had a lack of knowledge of gene expression profiles exhibited more notable improvements than mSTAR, when collaborating with our aggregator pretrained on gene data. This supports the analysis illustrated in the results of slide classification again, in which multimodal knowledge has been transferred into the patch extractor mSTAR, leading to relatively moderate improvements when using the pretrained aggregator. In other words, mSTAR itself has absorbed the knowledge in pathology reports and gene expression profiles, which has contributed to survival prediction.

To sum up, mSTAR demonstrated similar trends on both slide classification and survival analysis with consistent improvements, which repeatedly validated its capability on unimodal pathological clinical tasks. Additional results on the pretrained aggregator, on the one hand, provided extra evidence of the significance of pathology reports and RNA-Seq data at the whole-slide context, especially on molecular prediction and survival analysis. On the other hand, it indicated that mSTAR itself could acquire multimodal knowledge at the slide level, thereby resulting in better performance in these diverse tasks.

2.4 Multimodal Capability

Multimodal contrastive pretraining naturally allows the model to be applied to multimodal downstream tasks. In this study, since mSTAR aligned multimodal data that may contribute to alleviating inter-modal heterogeneity, we first examine whether features extracted by mSTAR facilitate multimodal fusion by assessing multimodal survival analysis tasks. Next, we investigate more cross-modal abilities, including slide-level few-shot and zero-shot classification and pathological report generation.

mSTAR facilitates Multimodal Fusion. The involvement of multimodal data in the pretraining process can enhance the model's capability of capturing complex interactions across different modalities and aligning each other by contrastive learning, thereby potentially alleviating the heterogeneity of different modalities and improving its performance on a wider range of cancer types. To validate this, we replaced pathological features with ones extracted by various extractors in existing multimodal fusion models for 9 TCGA cancer survival prediction tasks, to observe the differences that would arise. Specifically, recent SOTA multimodal fusion models were employed in this study, including MCAT [8], Porpoise [24], MOTCat [9] and CMTA [10].

On the whole, mSTAR has clearly outperformed other SOTA methods by a wide margin. Considering average rank, mSTAR ranked between 1.22 and 1.67 among various fusion models and the overall rank is 1.47, which left the second-best approach UNI far behind (**Figure** 5a) ranking at 2.56 to 3.00 and 2.68 on average. For average C-Index (**Figure** 5b), mSTAR achieved consistent and notable enhancement in multimodal fusion with a significant difference, with average performance increases of +1.8% ($P < 0.001$). Among different multimodal fusion models, mSTAR outperformed the second-best UNI by +1.7% ($P = 0.02$) for MCAT, +1.1% ($P = 0.01$) for Porpoise, +2.5% ($P < 0.005$) for MOTCat and +1.9% ($P < 0.005$) for CMTA.

Across various datasets, based on MCAT, mSTAR surpassed other SOTA approaches on 7 out of 9 tasks (**Figure** 5c and **Extended Data** Table A8), especially on LUAD (+3.3%, $P < 0.001$) and BRCA (+1.9%, $P < 0.001$). In the case
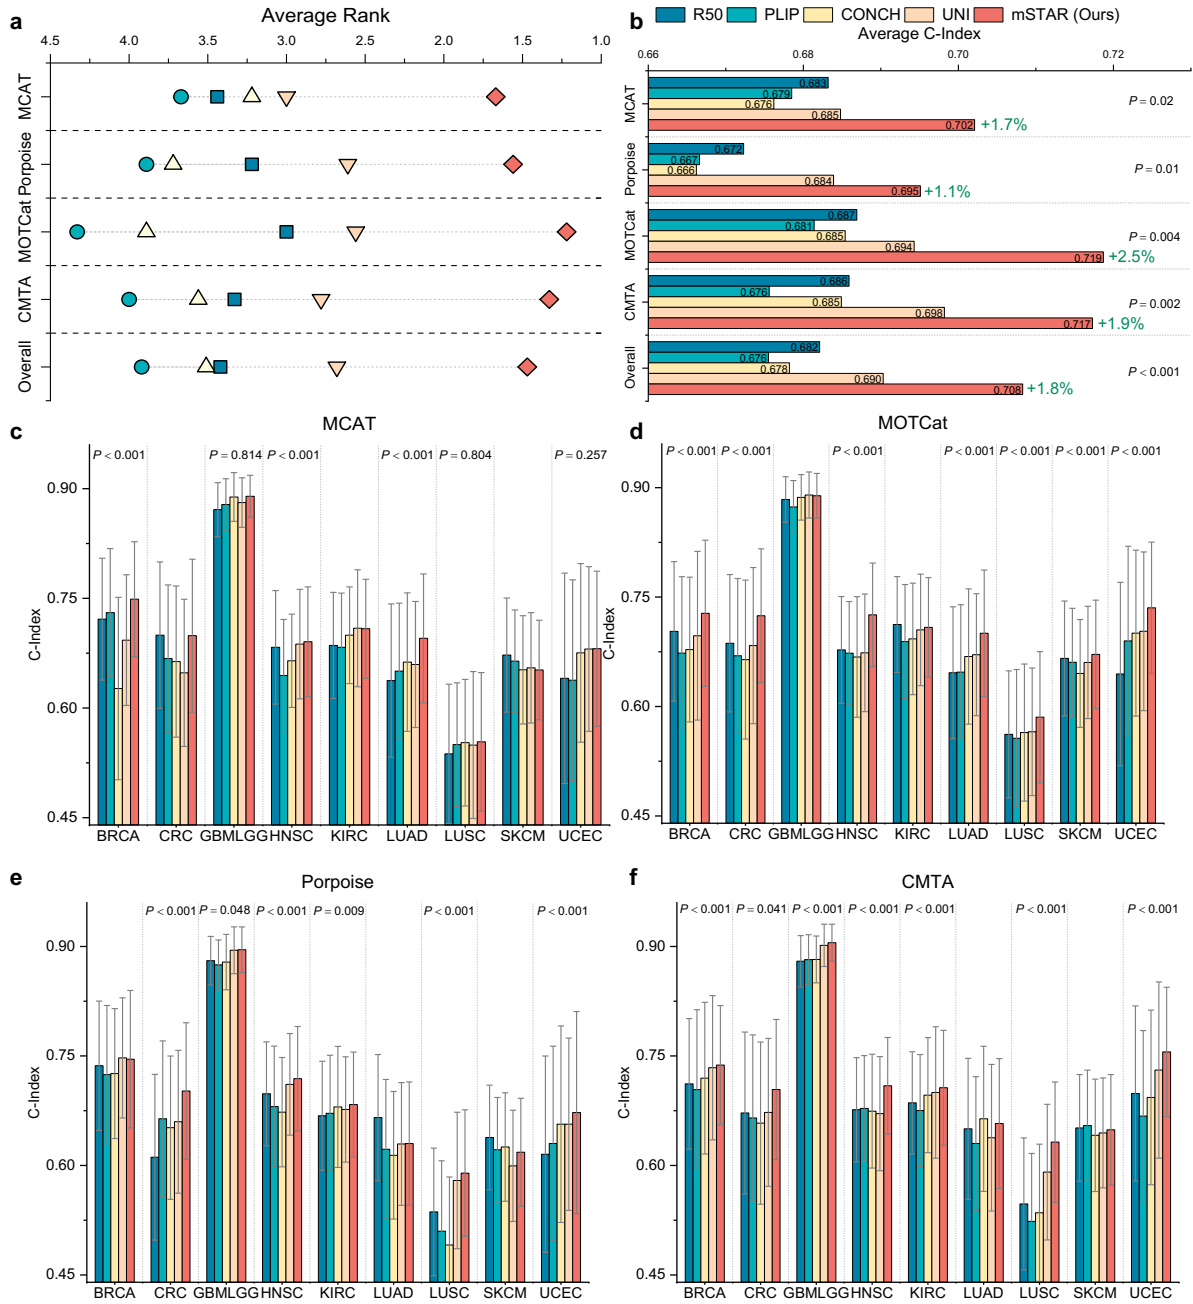
11

Fig. 5: **Multimodal Survival Analysis Performance**. **a**, Comparison of C-Index between mSTAR and compared methods on 9 TCGA datasets. They are evaluated with different multimodal fusion models, MCAT, Porpoise, MOTCat and CMTA trained from scratch based on the compared patch extractors and mSTAR. "Overall" refers to the average results among these multimodal fusion methods. **b**, The average C-Index on 9 TCGA datasets. **c-f** The survival analysis performance (C-Index and 95% CI) on 9 TCGA datasets with various multimodal fusion models, respectively. If the performance of mSTAR is the best one compared against the compared FMs, P-value would be presented. The colors of legends are shared across all sub-figures.
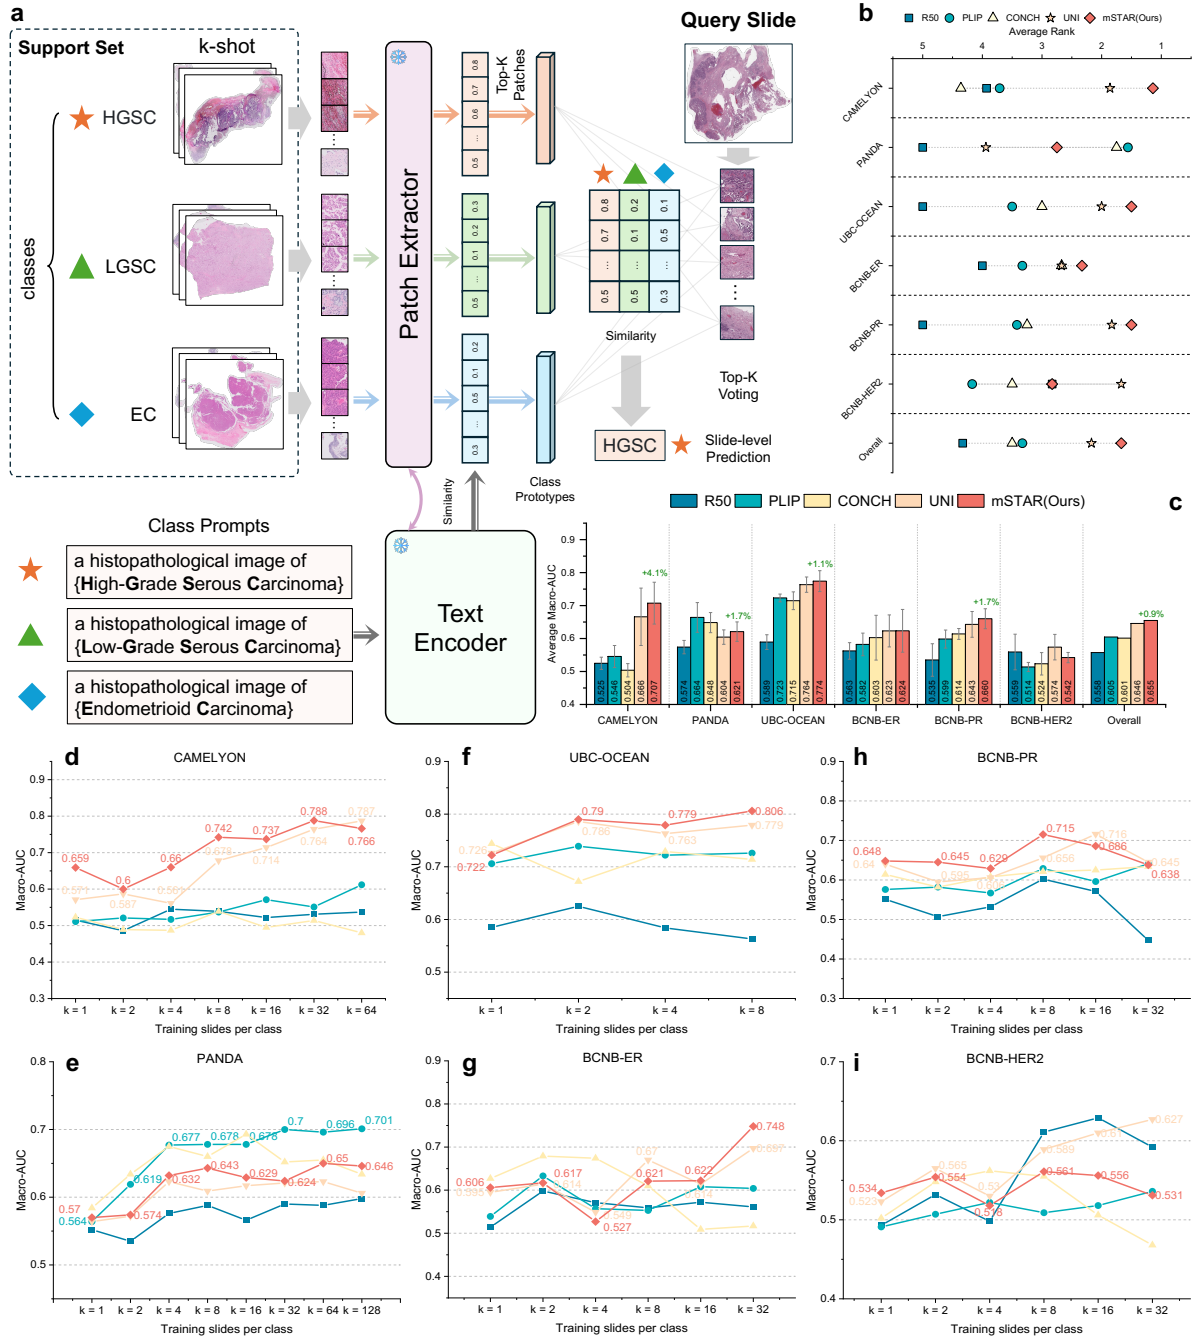
Fig. 6: **Performance of Few-shot Slide Classification**. **a**, Few-shot Slide Classification via Prototypical MI-FewShot. Top-K Patches of slides in the support set are selected to construct class prototypes for each category. They are chosen by higher similarities to the class embedding of class prompts via the pretrained text encoder. For a query slide, similarities between all patches and each class prototype are computed, and the class label is assigned to the most similar class prototype by top-K patches voting. **b**, Average Rank on Macro-AUC of compared FMs and mSTAR on 6 subtasks and overall. "Overall" refers to the average result among various subtasks. **c**, Average Macro-AUC across k-shots on 6 subtasks and overall. The numbers in green indicate the performance improvement over the vision-only FM, UNI. **d-i**, Macro-AUC across varying shots on every dataset for each FM. The detailed performance is presented for mSTAR and the best or second-best FM. The colors of legends are shared across sub-figures b-i.

13

of MOTCat, mSTAR excelled in 7 out of 9 tasks (**Figure** 5d and **Extended Data** Table A10) with performance increases of 0.6%-4.8% ($P <$ 0.001). mSTAR with Porpoise demonstrated superior performance in the majority of tasks, topping 5 out of 9 datasets (**Figure** 5e and **Extended Data** Table A9), which increased the second-best model by up to +3.8% ($P < 0.001$). For CMTA, mSTAR delivered the highest performance in 7 of 9 (**Figure** 5f and **Extended Data** Table A11), advancing the second-best one by 0.4%-4.1% ($P <$ 0.001).

In a nutshell, the remarkable increases across various datasets and diverse multimodal fusion backbone models vividly demonstrate the tremendous contributions of multimodal knowledge embedded by slide-level multimodal contrastive learning in facilitating multimodal fusion.

mSTAR advances Few-shot Slide Classification. Few-shot learning aims to achieve accurate prediction using only a small number of labeled data for each class, which is especially valuable in the context of medical scenarios where data scarcity is a common challenge. The ability of generalizing to new tasks with a few labels can be enhanced by multimodal data by offering a more comprehensive and diverse set of information. In this study, we investigated the few-shot capability of different patch extractors given different shots of labeled slides.

To ensure that the pretrained extractor can solely rely on the limited labeled slides provided, we excluded TCGA classification subtasks for evaluation, since mSTAR was pretrained on TCGA data. As a result, every patch extractor was assessed on 6 slide classification subtasks, CAMELYON, PANDA, UBC-OCEAN, BCNB-ER, BCNB-PR and BCNB-HER2. To enable the few-shot classification in the context of multiple instance learning, we proposed MI-Fewshot, as illustrated in **Figure** 6a and the Section 4.3. It utilizes a pretrained text encoder to select the most relevant patches from a limited amount of labeled data using similarity scores. These selected patches are then averaged to form the class prototype for accomplishing few-shot classification. Since MI-Fewshot requires a pretrained text encoder, we equipped those vision-only extractors, i.e., R50 and UNI, with the pretrained text encoder of mSTAR for assessment, while other extractors used their own designated text encoders.

Across 6 subtasks, mSTAR ranked the 1st place overall and achieved the best performance on 4 subtasks (**Figure** 6b and **Extended Data** Table A12), where mSTAR outperformed second-best FMs on average by a large margin (+4.1% on Camelyon, +1.0% on UBC-OCEAN, +2.0% on BCNB-PR and +0.9% for overall), as shown in **Figure** 6c (**Extended Data** Table A12). To demonstrate the benefits of multimodal knowledge, we also presented the performance increase compared to the vision-only SOTA FM, UNI. On 4 out of 6 subtasks, mSTAR exhibited better performance on average, with an additional increase on PANDA by +1.7%.

Following previous studies [4], we examined performance on varying shots of every subtask (**Figure** 6d-i and **Extended Data** Table A13-A18). We observed that mSTAR has showcased superior performance with a larger margin when given fewer slides on Camelyon, while the performance gap has gradually diminished as the number of slides increased. This indicated mSTAR performed better on the setting with fewer slides. On UBC-OCEAN, mSTAR always performed the best. For PANDA, mSTAR consistently outperformed vision-only UNI across varying shots, which also indicated the effectiveness of multimodal knowledge.

All in all, these results demonstrated mSTAR advanced generalization capabilities in label-efficient scenarios, which could be attributed to slide-level multimodal contrastive learning. Specially, by comparing with vision-only FMs, mSTAR showcased clear superiority, which indicated the extra benefit of multimodal knowledge.

mSTAR enables Zero-shot Slide Classification. In the context of open clinical practice, out-of-the-box capability of diagnostic decision-making without the need for task-specific fine-tuning is crucial for efficient utilization of multimodal models, especially in resource-constrained scenarios where the access to sufficient data and computational resources may be limited, such as underdeveloped areas. Therefore, we also evaluated zero-shot's capability of each extractor on 6 slide classification subtasks where we excluded TCGA subtasks as well, following the similar settings in few-shot's
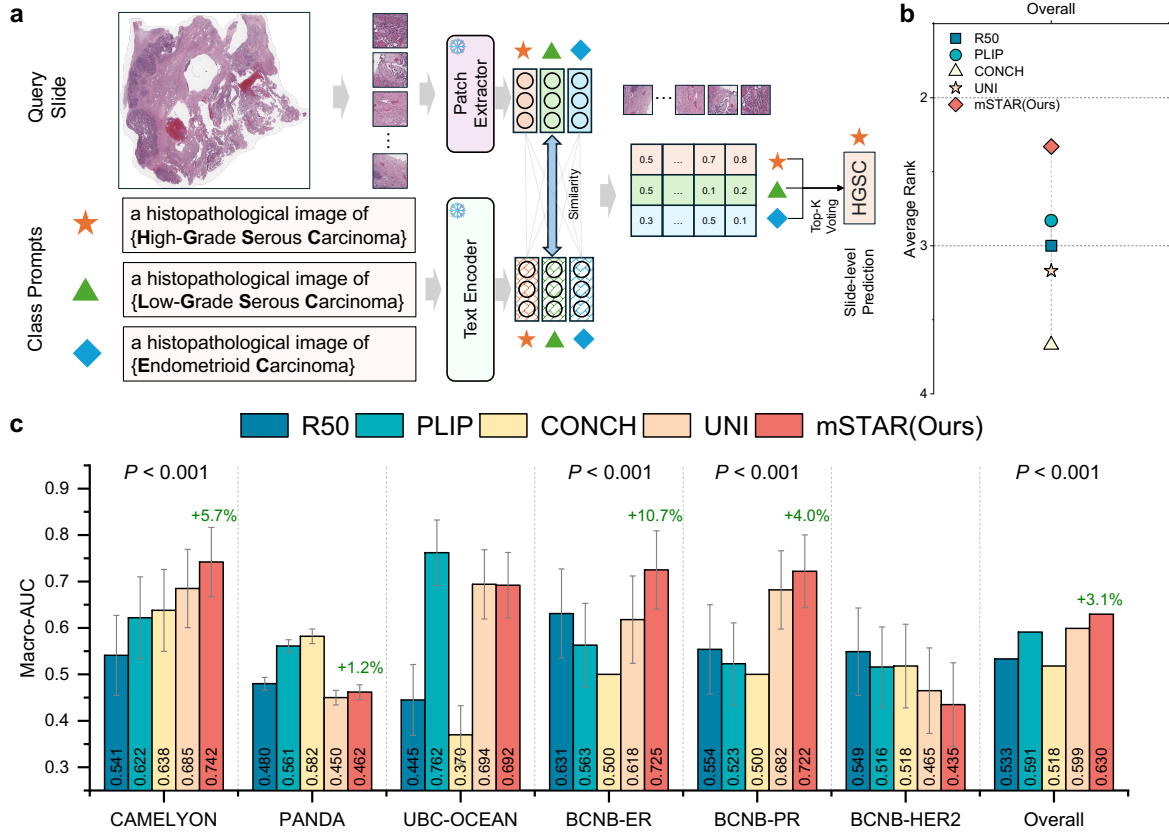
Fig. 7: **Performance of Zero-shot Slide Classification**. **a**, Zero-shot Slide Classification via MI-Zero [25]. For a query slide, similarities between all patches and each class embedding are computed, and the class label is assigned to the most similar class prototype by top-K patches voting. **b**, Average Rank on Macro-AUC of compared FMs and mSTAR across 6 subtasks. "Overall" refers to the average result among various subtasks. **c**, Macro-AUC on 6 subtasks and overall. The numbers in green indicate the improvement over the vision-only FM, UNI. The colors of legends are shared across sub-figures b-c. If the performance of mSTAR is the best one compared against the compared FMs, P-value would be presented.

evaluation. To produce slide-level predictions for patch extractors, MI-Zero [25] was adopted through top-K patches voting based on patch similarities to class prototypes embedded by the pretrained text encoders (**Figure** 7a).

Across 6 slide subtasks, mSTAR outperformed other FMs on half of the tasks and performed best on the overall result (**Figure** 7b and c). Similarly, we also compared mSTAR with the vision-only FM, UNI, to validate the effectiveness of multimodal knowledge injection. As a result, notable performance increases were observed on 4 of 6 tasks, and meanwhile mSTAR achieved a comparable performance with UNI on an extra UBC-OCEAN task. Specifically, compared to the

second-best FM, UNI, mSTAR achieved clear enhancement in these tasks by +3.1% on average ($P < 0.001$) with a significant difference. In particular, on BCNB-ER task, a remarkable rise of +9.4% ($P < 0.001$) was observed compared to the second-best FM, UNI, and we further saw performance enhancement over UNI from 0.618 to 0.725 (+10.7%, $P < 0.001$). For CAMELYON and BCNB-PR, there were notable boosts of +5.7% and +4.0% over the second-performing FM ($P < 0.001$), respectively. For PANDA, mSTAR significantly enhanced UNI from 0.45 to 0.462 ($P < 0.001$).

Overall, despite not being the best performing on all subtasks, mSTAR clearly outperformed the

15
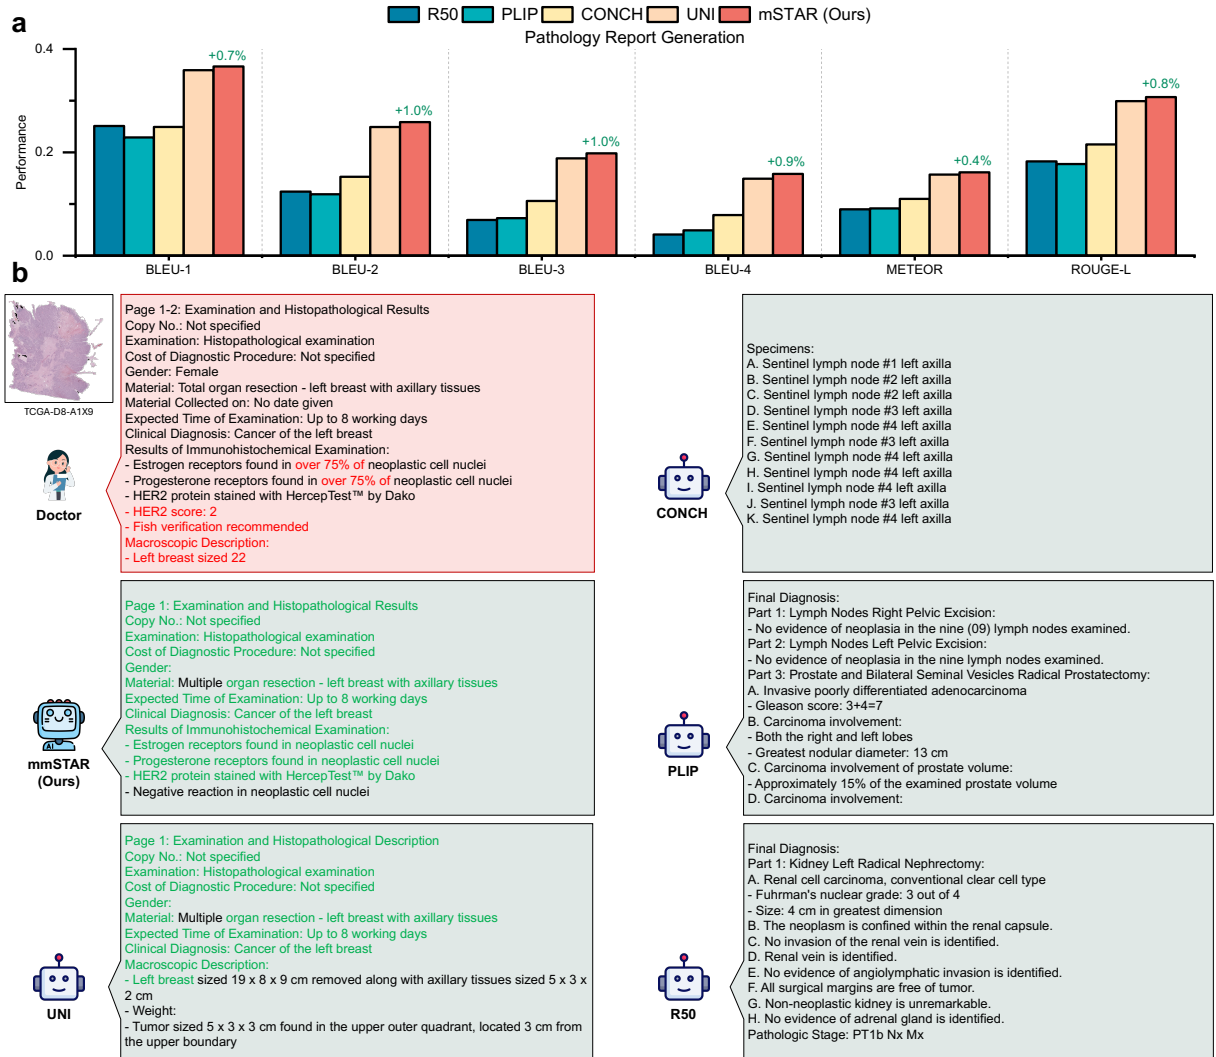
Fig. 8: **Pathology Report Generation Performance. a**, quantitative results (BLEU-1-4, METEOR and ROUGE-L) of various compared models and mSTAR. **b**, an example of report generation, where texts in red boxes are grount-truth from doctors and ones in dark green boxes are generated with pathological features from various models. The words highlighted in green are aligned with ground-truth, while ones marked in red are missing by mSTAR from ground-truth.

second-best FM with a large margin on half of the subtasks. When compared with the vision-only FM, mSTAR performed better or comparative on 5 out of 6 subtasks, which indicates the effectiveness of multimodal contrastive learning. These results showcased that mSTAR can achieve notable advancements in its out-of-the-box ability on certain subtasks. However, there is still considerable room for further enhancements. The integration of additional multimodal data

through the pretraining approach of mSTAR has the potential to enhance its zero-shot capability.

mSTAR excels in Pathological Report Generation. Automated generation of pathology reports has enormous potential in simplifying the report-writing process and reducing the workload burden on pathologists. Next, we assessed its capability of report generation on TCGA datasets. It is worth mentioning that cases used

for evaluation are completely held out from the pretraining materials. We presented both quantitative and qualitative results.

From the quantitative perspective, we evaluated multiple metrics including BLEU, METEOR and ROUGE-L to assess various aspects of the generated text, such as precision of n-grams (contiguous sequences of words), order, alignment, recall, etc. Across different metrics, mSTAR consistently outperformed the second-best approach by nearly 1% (**Figure** 8a and **Extended Data** Table A20), which indicates better quality of generated reports.

We continued to qualitatively evaluate the quality of generated reports. A case study is presented in **Figure** 8b and more cases are provided in the **Extended Data** Figure A1. From the case of **Figure** 8b, we found that mSTAR is able to more accurately generate important diagnostic facts with minimal omissions, such as clinical diagnosis. The result of **Figure** 8b generated by CONCH, PLIP and R50 contain a significant number of factual errors. Compared with the second-best approach, UNI, mSTAR is capable of generating immunohistochemical results on top of H&E WSIs, such as HER2, which may be attributed to the involvement of gene expression data in pretraining. This may align with the earlier investigation of mSTAR's superior performance in molecular prediction. From the case of **Extended Data** Figure A1a, mSTAR is able to generate the majority of key diagnostic information, while other FM just produced a few key words. The case of **Extended Data** Figure A1b demonstrated that mSTAR excelled in producing precise histologic type and grade, while the majority of the results generated by other FMs are factually incorrect. Although UNI included grading information in the case of **Extended Data** Figure A1c, the result of pathological T staging was significantly erroneous. This can be attributed to the involvement of gene expression profiles, since RNA-Seq data are known to be strong indicators of staging [26, 27], subtyping [28] and grading [29, 30]. However, all FMs are unable to provide accurate quantity estimations, such as lesion sizes and tumor location. The possible reasons for the failure of quantity estimations are that they lack information regarding the magnification and locations of target objects mentioned in a certain report. These pieces of information are likely recorded by doctors during the specimen collection process and cannot be directly inferred from the WSI.

In conclusion, the multimodal data endows mSTAR with unique knowledge from each modality, such as identifying molecular profiles, enabling it to generate more comprehensive and accurate reports. These encouraging results underscored the potential of mSTAR for report generation and its significance in pathological decision-making.

3 Discussion

In this study, we delve into how to harness the full potential of multimodal data to effectively advance the performance of the pathology foundation models and bridge the gaps of previous foundation models in multimodal capabilities. Additionally, we explored a new whole-slide pretraining paradigm for CPath, which broadened the context of modelling for better performance on slide-level tasks. We curated the largest multimodal dataset consisting of WSIs, pathology reports and RNA-Seq data, which contained over 26k public modality pairs from 10,275 patients across 32 major cancer types. Diverse experimental results on 43 subtasks demonstrated that mSTAR excelled in not only unimodal tasks but also multimodal tasks at the slide level. In particular, with the inclusion of gene expression data, mSTAR performed well in slide-level molecular prediction. Furthermore, multimodal contrastive learning facilitated multimodal fusion tasks and endowed the model with more generalized capabilities in low-shot scenarios. Simultaneously, the incorporation of clinical pathology reports in the pretraining process resulted in the advanced capability of report generation by providing the expert knowledge.

In the realm of prior investigations into pathology foundation models, two prominent categories have emerged: vision-only models [3, 31, 5] and vision-language models [4, 2]. However, these approaches fail to tap into a vast wealth of information inherent in clinical pathology reports and gene expression profiles. Pathology reports usually provide authentic expert knowledge in line with the clinical practice, while gene expression profiles always serve as robust indicators of disease status for clinical applications in diagnosis [32] and prognosis [33]. As shown in Extended Data

Table A1, the involvement of pathology reports and gene expression data can bring extra performance gains. We also found that there was still considerable room for further improvement in molecular prediction, multimodal fusion for survival analysis and report generation in existing FMs, which can benefit from the inherent knowledge in pathology reports [34] and gene expression profiles [29, 28, 30]. Simultaneously, we notice that the majority of existing FMs primarily focus on patch/ROI-level models and short texts, in which restricted contexts hinder their practical performance in slide-level clinical applications, such as survival analysis and report generation. It is worth noting that, unlike CONCH supervised by generative loss, mSTAR without involving generative components in pretraining, still demonstrated encouraging performance in report generation with producing more comprehensive texts. Recently, beyond working on small patches/ROIs, we noticed that a study attempted to work on slide-level foundation models, which pretrained the model on patch features. However, the pretrained performance significantly depends on the quality of patch features. In other words, their performance would be limited by the patch extractor. We believe that end-to-end pretraining is a promising solution in the future, while its current implementation is hindered by hardware limitations. Therefore, mSTAR bridges this gap through self-taught training to seamlessly transfer the knowledge captured by the slide aggregator into the patch extractor.

Distinct from previous researches, our study provides the uniqueness in two folds. First, our findings showcase the remarkable power of leveraging multimodal data to drive the progress of pathology foundation models, especially in enhancing multimodal capabilities. Second, we found a unified way to bridge the gap between slide-level and patch-level pretraining, bringing us closer to achieving end-to-end pretraining on raw slide data. We believe this innovative unified paradigm will revolutionize the workflow of pretraining for CPath. Moreover, this paradigm allows the injection of multimodal knowledge into pathology foundation models in an appropriate manner, which may hold the potential to harness more modalities to construct a stronger foundation model for CPath.

Although preliminary results are encouraging, this study still has several limitations. First, a key limitation is the scale of data for pretraining compared to previous works of pathology foundation models. By expanding the scale of multimodal data for pretraining, we can expect to unlock further potential for enhancing various abilities, such as multimodal capabilities. Second, we still potentially have a long way to go before achieving the true end-to-end foundation model. Before that, mSTAR will serve as an alternative solution to seamlessly bridge slide-level and patch-level pretraining. However, there are still several challenges to be further explored, such as the appropriate way to propagate the pretrained knowledge embedded in the slide aggregator and the architectural design of slide aggregator. In mSTAR, due to a large number of patches of a WSI that would lead to extremely high computational costs, we selected TransMIL with linear time complexity as the slide aggregator. However, the increase in training speed comes at the expense of sacrificing a portion of the performance. Fortunately, a multitude of innovative architectures for handling long sequences are emerging, such as Mamba [35], LongNet [36], etc, which we explore in concurrent work [37]. We believe that these new architectures will undoubtedly create new avenues in exploring more efficient and powerful pretraining paradigms for CPath. In the future, we plan to incorporate more multimodal data into pretraining, such as multi-omics data, and explore new efficient pretraining architecture.

References

[1] Janez Demšar. Statistical comparisons of classifiers over multiple data sets. *The Journal of Machine learning research*, 7:1–30, 2006.

[2] Zhi Huang, Federico Bianchi, Mert Yuksekgonul, Thomas J Montine, and James Zou. A visual–language foundation model for pathology image analysis using medical twitter. *Nature medicine*, 29(9):2307–2316, 2023.

[3] Richard J Chen, Tong Ding, Ming Y Lu, Drew FK Williamson, Guillaume Jaume, Andrew H Song, Bowen Chen, Andrew Zhang, Daniel Shao, Muhammad Shaban, et al. Towards a general-purpose foundation

model for computational pathology. *Nature Medicine*, 30(3):850–862, 2024.

[4] Ming Y Lu, Bowen Chen, Drew FK Williamson, Richard J Chen, Ivy Liang, Tong Ding, Guillaume Jaume, Igor Odintsov, Long Phi Le, Georg Gerber, et al. A visual-language foundation model for computational pathology. *Nature Medicine*, 30(3):863–874, 2024.

[5] Hanwen Xu, Naoto Usuyama, Jaspreet Bagga, Sheng Zhang, Rajesh Rao, Tristan Naumann, Cliff Wong, Zelalem Gero, Javier González, Yu Gu, et al. A whole-slide foundation model for digital pathology from real-world data. *Nature*, pages 1–8, 2024.

[6] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

[7] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multimodal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19113–19122, 2023.

[8] Richard J Chen, Ming Y Lu, Wei-Hung Weng, Tiffany Y Chen, Drew FK Williamson, Trevor Manz, Maha Shady, and Faisal Mahmood. Multimodal co-attention transformer for survival prediction in gigapixel whole slide images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4025, 2021.

[9] Yingxue Xu and Hao Chen. Multimodal optimal transport-based co-attention transformer with global structure consistency for survival prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21241–21251, 2023.

[10] Fengtao Zhou and Hao Chen. Cross-modal translation and alignment for survival analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21485–21494, 2023.

[11] Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *International conference on machine learning*, pages 2127–2136. PMLR, 2018.

[12] Ming Y Lu, Drew FK Williamson, Tiffany Y Chen, Richard J Chen, Matteo Barbieri, and Faisal Mahmood. Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature Biomedical Engineering*, 5(6):555–570, 2021.

[13] Zhuchen Shao, Hao Bian, Yang Chen, Yifeng Wang, Jian Zhang, Xiangyang Ji, et al. Transmil: Transformer based correlated multiple instance learning for whole slide image classification. *Advances in neural information processing systems*, 34:2136–2147, 2021.

[14] Babak Ehteshami Bejnordi, Mitko Veta, Paul Johannes Van Diest, Bram Van Ginneken, Nico Karssemeijer, Geert Litjens, Jeroen AWM Van Der Laak, Meyke Hermsen, Quirine F Manson, Maschenka Balkenhol, et al. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *Jama*, 318(22):2199–2210, 2017.

[15] Peter Bandi, Oscar Geessink, Quirine Manson, Marcory Van Dijk, Maschenka Balkenhol, Meyke Hermsen, Babak Ehteshami Bejnordi, Byungjae Lee, Kyunghyun Paeng, Aoxiao Zhong, et al. From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge. *IEEE transactions on medical imaging*, 38(2):550–560, 2018.

[16] Wouter Bulten, Kimmo Kartasalo, Po-Hsuan Cameron Chen, Peter Ström, Hans Pinckaers, Kunal Nagpal, Yuannan Cai, David F Steiner, Hester Van Boven, Robert Vink, et al. Artificial intelligence for diagnosis and gleason grading of prostate cancer: the panda challenge. *Nature medicine*, 28(1):154–163, 2022.

[17] Maryam Asadi-Aghbolaghi, Hossein Farahani, Allen Zhang, Ardalan Akbari, Sirim Kim, Ashley Chow, Sohier Dane, OCEAN Challenge Consortium, OTTA Consortium, David G Huntsman, et al. Machine learning-driven histotype diagnosis of ovarian carcinoma: Insights from the ocean ai challenge. *medRxiv*, pages 2024–04, 2024.

[18] Hossein Farahani, Jeffrey Boschman, David Farnell, Amirali Darbandsari, Allen Zhang,

19

Pouya Ahmadvand, Steven JM Jones, David Huntsman, Martin Köbel, C Blake Gilks, et al. Deep learning-based histotype diagnosis of ovarian carcinoma whole-slide pathology images. *Modern Pathology*, 35(12):1983–1990, 2022.

[19] John N Weinstein, Eric A Collisson, Gordon B Mills, Kenna R Shaw, Brad A Ozenberger, Kyle Ellrott, Ilya Shmulevich, Chris Sander, and Joshua M Stuart. The cancer genome atlas pan-cancer analysis project. *Nature genetics*, 45(10):1113–1120, 2013.

[20] Mengdan Zhu, Bing Ren, Ryland Richards, Matthew Suriawinata, Naofumi Tomita, and Saeed Hassanpour. Development and evaluation of a deep neural network for histologic classification of renal cell carcinoma on biopsy and surgical resection slides. *Scientific reports*, 11(1):7080, 2021.

[21] Feng Xu, Chuang Zhu, Wenqi Tang, Ying Wang, Yu Zhang, Jie Li, Hongchuan Jiang, Zhongyue Shi, Jun Liu, and Mulan Jin. Predicting axillary lymph node metastasis in early breast cancer using deep learning on primary tumor biopsy slides. *Frontiers in Oncology*, page 4133, 2021.

[22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 630–645. Springer, 2016.

[23] Yilan Zhang, Yingxue Xu, Jianqi Chen, Fengying Xie, and Hao Chen. Prototypical information bottlenecking and disentangling for multimodal cancer survival prediction. In *The Twelfth International Conference on Learning Representations*.

[24] Richard J Chen, Ming Y Lu, Drew FK Williamson, Tiffany Y Chen, Jana Lipkova, Zahra Noor, Muhammad Shaban, Maha Shady, Mane Williams, Bumjin Joo, et al. Pan-cancer integrative histology-genomic analysis via multimodal deep learning. *Cancer Cell*, 40(8):865–878, 2022.

[25] Ming Y Lu, Bowen Chen, Andrew Zhang, Drew FK Williamson, Richard J Chen, Tong Ding, Long Phi Le, Yung-Sung Chuang, and Faisal Mahmood. Visual language pre-trained multiple instance zero-shot transfer for histopathology images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19764–19775, 2023.

[26] Sudhanshu Shukla, Joseph R Evans, Rohit Malik, Felix Y Feng, Saravana M Dhanasekaran, Xuhong Cao, Guoan Chen, David G Beer, Hui Jiang, and Arul M Chinnaiyan. Development of a rna-seq based prognostic signature in lung adenocarcinoma. *JNCI: Journal of the National Cancer Institute*, 109(1):djw200, 2017.

[27] Jin-Cheng Guo, Yang Wu, Yang Chen, Feng Pan, Zhi-Yong Wu, Jia-Sheng Zhang, Jian-Yi Wu, Xiu-E Xu, Jian-Mei Zhao, En-Min Li, et al. Protein-coding genes combined with long noncoding rna as a novel transcriptome molecular staging model to predict the survival of patients with esophageal squamous cell carcinoma. *Cancer communications*, 38:1–13, 2018.

[28] Milad Mostavi, Yu-Chiao Chiu, Yufei Huang, and Yidong Chen. Convolutional neural network models for cancer type prediction based on gene expression. *BMC medical genomics*, 13:1–13, 2020.

[29] Mei Wang, Daniel Klevebring, Johan Lindberg, Kamila Czene, Henrik Grönberg, and Mattias Rantalainen. Determining breast cancer histological grade from rna-sequencing data. *Breast Cancer Research*, 18:1–13, 2016.

[30] Y Wang, B Acs, S Robertson, B Liu, Leslie Solorzano, Carolina Wählby, J Hartman, and M Rantalainen. Improved breast cancer histological grading using deep learning. *Annals of Oncology*, 33(1):89–98, 2022.

[31] Eugene Vorontsov, Alican Bozkurt, Adam Casson, George Shaikovski, Michal Zelechowski, Siqi Liu, Philippe Mathieu, Alexander van Eck, Donghun Lee, Julian Viret, et al. Virchow: A million-slide digital pathology foundation model. *arXiv preprint arXiv:2309.07778*, 2023.

[32] Mingye Hong, Shuang Tao, Ling Zhang, Li-Ting Diao, Xuanmei Huang, Shaohui Huang, Shu-Juan Xie, Zhen-Dong Xiao, and Hua Zhang. Rna sequencing: new technologies and applications in cancer research. *Journal of hematology & oncology*, 13:1–16, 2020.

[33] David G Beer, Sharon LR Kardia, Chiang-Ching Huang, Thomas J Giordano, Albert M

Levin, David E Misek, Lin Lin, Guoan Chen, Tarek G Gharib, Dafydd G Thomas, et al. Gene-expression profiles predict survival of patients with lung adenocarcinoma. *Nature medicine*, 8(8):816–824, 2002.

[34] Zhengrui Guo, Jiabo Ma, Yingxue Xu, Yihui Wang, Liansheng Wang, and Hao Chen. Histgen: Histopathology report generation via local-global feature encoding and cross-modal context interaction. *arXiv preprint arXiv:2403.05396*, 2024.

[35] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023.

[36] Jiayu Ding, Shuming Ma, Li Dong, Xingxing Zhang, Shaohan Huang, Wenhui Wang, Nanning Zheng, and Furu Wei. Longnet: Scaling transformers to 1,000,000,000 tokens. *arXiv preprint arXiv:2307.02486*, 2023.

[37] Shu Yang, Yihui Wang, and Hao Chen. Mambamil: Enhancing long sequence modeling with sequence reordering in computational pathology. *arXiv preprint arXiv:2403.06800*, 2024.

[38] Bo Li and Colin N Dewey. Rsem: accurate transcript quantification from rna-seq data with or without a reference genome. *BMC bioinformatics*, 12:1–16, 2011.

[39] Abhishek Sarkar and Matthew Stephens. Separating measurement and expression models clarifies confusion in single-cell rna sequencing analysis. *Nature genetics*, 53(6):770–777, 2021.

[40] Ashraful Haque, Jessica Engel, Sarah A Teichmann, and Tapio Lönnberg. A practical guide to single-cell rna-sequencing for biomedical research and clinical applications. *Genome medicine*, 9:1–12, 2017.

[41] Fan Yang, Wenchuan Wang, Fang Wang, Yuan Fang, Duyu Tang, Junzhou Huang, Hui Lu, and Jianhua Yao. scbert as a large-scale pretrained deep language model for cell type annotation of single-cell rna-seq data. *Nature Machine Intelligence*, 4(10):852–866, 2022.

[42] Jingcheng Du, Peilin Jia, Yulin Dai, Cui Tao, Zhongming Zhao, and Degui Zhi. Gene2vec: distributed representation of genes based on co-expression. *BMC genomics*, 20:7–15, 2019.

[43] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.

[44] Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240, 2020.

[45] Haotian Cui, Chloe Wang, Hassaan Maan, Kuan Pang, Fengning Luo, Nan Duan, and Bo Wang. scgpt: toward building a foundation model for single-cell multi-omics using generative ai. *Nature Methods*, pages 1–11, 2024.

[46] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017.

[47] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

[48] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.

[49] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022.

[50] Shekoufeh Gorgi Zadeh and Matthias Schmid. Bias in cross-entropy-based training of deep survival networks. *IEEE transactions on pattern analysis and machine intelligence*, 43(9):3126–3137, 2020.

[51] Arthur Liberzon, Chet Birger, Helga Thorvaldsdóttir, Mahmoud Ghandi, Jill P Mesirov, and Pablo Tamayo. The molecular signatures database hallmark gene set

21

collection. *Cell systems*, 1(6):417–425, 2015.

[52] Yan Wang, Wei-Lun Chao, Kilian Q Weinberger, and Laurens Van Der Maaten. Simpleshot: Revisiting nearest-neighbor classification for few-shot learning. *arXiv preprint arXiv:1911.04623*, 2019.

[53] Bradley Efron and Robert J Tibshirani. *An introduction to the bootstrap.* Chapman and Hall/CRC, 1994.

[54] Frank Wilcoxon. Individual comparisons by ranking methods. In *Breakthroughs in statistics: Methodology and distribution*, pages 196–202. Springer, 1992.

[55] Aatish Thennavan, Francisco Beca, Youli Xia, Susana Garcia-Recio, Kimberly Allison, Laura C Collins, M Tse Gary, Yunn-Yi Chen, Stuart J Schnitt, Katherine A Hoadley, et al. Molecular analysis of tcga breast cancer histologic types. *Cell genomics*, 1(3), 2021.

[56] Justin Guinney, Rodrigo Dienstmann, Xin Wang, Aurélien De Reynies, Andreas Schlicker, Charlotte Soneson, Laetitia Marisa, Paul Roepman, Gift Nyamundanda, Paolo Angelino, et al. The consensus molecular subtypes of colorectal cancer. *Nature medicine*, 21(11):1350–1356, 2015.

[57] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002.

[58] Michael Denkowski and Alon Lavie. Meteor 1.3: Automatic metric for reliable optimization and evaluation of machine translation systems. In *Proceedings of the sixth workshop on statistical machine translation*, pages 85–91, 2011.

[59] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.

4 Methods

4.1 Pretraining Dataset Curation

Data used for pretraining in this study were totally obtained from a publicly available source, the Cancer Genome Atlas Program (TCGA) [19], in which we collected 9,640 cases (11,765 slides)

of diagnostics formalin-fixed paraffin-embedded (FFPE) H&E WSIs, 11,108 pathology reports and 10,234 cases of bulk RNA-Seq data across all 32 cancer types of TCGA. After quality control, we curated a dataset with 8,440 WSI-Report pairs, 8,965 WSI-RNA-Seq pairs and 8,764 Report-RNA-Seq pairs, resulting in 26,169 modality pairs, as shown in Figure 1c. Given that numerous downstream tasks, such as survival analysis, were evaluated on TCGA data, we held out some validation and test cases. For 9 cancer datasets comprising over 400 cases, we adopted a split ratio of 7:1:2 for train-validation-test folds. For those cases involving multiple slides, we combined their patches or features into a single case for pretraining at the patient level. This ensured slides belonging to one case were included within the same fold, thereby preventing potential data leakage. We also considered label stratification for survival labels at patient-level, since we primarily evaluated the performance of survival prediction on TCGA data. Note that all cases without survival labels were used for pretraining. Details of data splitting for these 9 cancer datasets are provided in **Extended Data** Table A21. After data partitioning, we curated 22,127 modality pairs for contrastive learning, consisting of 7,083 WSI-Report pairs, 7,538 WSI-RNA-Seq pairs and 7,506 Report-RNA-Seq pairs. Among these, there were 7,947 cases with all three modalities for pretraining. For acquisition of high-quality data, we conducted the subsequent pre-processing procedures for each modality.

WSI Pre-processing. To conduct slide- (or patient-) level tasks on WSIs, our processing pipeline involved tissue segmentation, patching, and feature extraction (for pretraining aggregators and evaluation). For tissue segmentation, we employed the CLAM library [12], which performed binary thresholding on the saturation channel of a downsampled RGB slide, converted to the hue-saturation-value (HSV) color space. The resulting segmentation mask was obtained by filtering the contours based on their area. The hyperparameters of segmentation are released on our codebase. Furthermore, slides that were corrupted and those containing a small proportion of tissue region were excluded from this study. As a result, we acquired 9,608 cases of 11,727 slides for pretraining and evaluation.

To adhere to established practices of previous works [12, 3, 4], we partitioned the segmented tissue regions into 256×256 pixels patches at $20\times$-equivalent magnification without overlaps and then resized all patches to 224×224 pixels for feature extraction. Using pretrained patch extractors that were kept frozen, we pre-extracted embeddings for each patch and stored them for subsequent evaluation purposes.

Report Pre-processing. For pathology reports, we curated open-source texts from TCGA and converted them from their original PDF format to editable text format via Amazon Web Services (AWS) Optical Character Recognition (OCR) tools, resulting in 9,523 Reports. For quality control, we curated these reports via the powerful language tool, GPT-4, with appropriate prompts provided in **Extended Data** Table A22, and re-checked them manually to ensure the unchanged original intent. The statistical distribution of word counts for reports is demonstrated in **Figure** 1e, in which the majority of cases have word counts below 500.

RNA-Seq Pre-processing. We accessed RNA-Seq data of TCGA from cBioportal database, which were preprocessed and normalized using RSEM [38]. An inherent difficulty in gene expression modelling arises from the variations in absolute magnitudes observed across different sequencing protocols [39]. Therefore, we further applied a common preprocessing technique log1p transformation [40] for gene expression values. Following previous works [41], Gene2Vec [42] contributed to better representing the gene names by enforcing that words with similar meanings are assigned similar representations. Therefore, we retained genes present in the Gene2Vec vocabulary. In the end, we obtained 9,890 cases RNA-Seq data, each consisting of genes with a length of 17,425.

4.2 Pretraining Framework

To utilize multimodal knowledge at the whole-slide context for enhancing the pathology foundation model, we propose a whole-slide pretraining paradigm consisting of two-stage pretraining, as shown in **Figure** 2. In the first stage, we aim to inject multimodal knowledge into the slide aggregator by contrastive learning, including inter-modality contrastive learning (following CLIP [43]) and inter-cancer contrastive learning. In the second stage, to seamlessly propagate multimodal knowledge at the slide-level context into the patch extractor, we leverage the slide aggregator pretrained in the first stage, serving as a "Teacher" model, to supervise the pretraining of the patch extractor, termed Self-Taught training. In this way, multimodal knowledge of the whole-slide context can be injected into the pathology FM.

Stage 1 - Pretrain Slide Aggregator. In this stage, we aim to pretrain a slide aggregator that learns multimodal knowledge by contrastive learning with other modalities. Note that the pre-trained slide aggregator plays a role of "Teacher" that propagates the learned knowledge into the patch extractor at the next stage. These modules to be trained are highlighted in red boxes in the **Figure** 2a, in which we pretrain a 2-layer TransMIL [13] as the slide aggregator for WSIs, a Bert-like text encoder (following BioBert-Base-v1.2 [44]) for pathology reports, and a Performer (following scBERT [41]) for RNA-Seq data.

Given these transformer-like encoders, we need to tokenize raw data of every modality into token embeddings before feeding them into their respective encoders. For pathology, we obtained non-overlapping 224×224 patches as early mentioned, and then for every patch, we used a pretrained patch extractor, UNI [3], to extract patch features, where a patch feature was regarded as a token embedding for the slide aggregator. After gathering 4,096 patch features for the i-th patient's WSIs, $\mathcal{P}_i = \{\mathbf{p}_i^m\}_{m=1}^M$, we fed them into the slide aggregator to integrate all patch features and got a 512-dimensional pathological [CLS] token embedding $\mathbf{P}_i$ as the slide-level representation, where M is the number of patches and it was fixed into 4,096. For cases where the number of patches exceeds 4,096, a random selection of 4,096 patches is made, while for cases with fewer than 4,096 patches, padding is applied using the mean value. For those cases where one patient has more than one WSI, we simply concatenated them together. Note that all patch features were transformed into 512-dimensional features by a linear projection before being forwarded into the aggregator.

For pathology reports, we adopted the text encoder for randomly truncated 512 tokens and

outputted the report [CLS] token embedding $\mathbf{T}_i$. For cases where the length of the text is less than 512, the special token '[pad]' was padded. The RNA-Seq data was organized as a set of 2-tuple (g_i, e_i) comprising of the gene name g_i and its expression variable e_i. Following previous works [41, 45], to assure that genes with potential co-expression get close together, we employed Gene2Vec [42] to generate 200-dimensional gene embeddings for each gene name g_i. Gene expression can be viewed as the manifestation or presence of each gene, which has been well-documented within a biological system. Therefore, we applied the term-frequency-analysis method used in previous works [41, 45] to discretize the continuous expression variable e_i through binning technique. Subsequently, the discrete variable was transformed into a 200-dimensional embedding, which was then integrated into the final gene token embedding $\mathbf{g}_i$ by addition. Through forwarding the gene encoder, we can get the gene [CLS] token embedding $\mathbf{G}_i$. It is worth noting that encoder outputs from report and gene modalities were transformed into 512-dimensional features by a linear projection for contrastive learning.

To optimize the model through pretraining, we incorporate two objectives including inter-modality contrastive learning and inter-cancer contrastive learning. In the case of inter-modality contrastive learning, given the [CLS] representation of each modality, every two modalities can be paired together, which finally yielded three combinations: WSI-report $(\mathbf{P}_i, \mathbf{T}_i)$, WSI-gene $(\mathbf{P}_i, \mathbf{G}_i)$ and report-gene $(\mathbf{T}_i, \mathbf{G}_i)$. During pretraining between every modality pairs, a mini-batch consisted of N samples, e.g., $\{(\mathbf{P}_i, \mathbf{T}_i)\}_{i=1}^{N}$ for WSI-report. Contrastive learning imposes a higher similarity in modality pairs from the same sample. Take WSI-report pairs as an example, and the loss function can be formulated as

$$
\begin{aligned}
\mathcal{L}_{P-T} = &- \frac{1}{2N} \sum_{i=1}^{N} \log \frac{\exp\left(\mathbf{P}_i^{\top} \mathbf{T}_i / \tau\right)}{\sum_{j=1}^{N} \exp\left(\mathbf{P}_i^{\top} \mathbf{T}_j / \tau\right)} \\
&- \frac{1}{2N} \sum_{j=1}^{N} \log \frac{\exp\left(\mathbf{T}_j^{\top} \mathbf{P}_j / \tau\right)}{\sum_{i=1}^{N} \exp\left(\mathbf{T}_j^{\top} \mathbf{P}_i / \tau\right)}
\end{aligned} \quad (1)
$$

where τ is a scale factor of the contrastive loss and it was set by default following CLIP [43]. Similarly,

we can get $\mathcal{L}_{P-G}$ and $\mathcal{L}_{T-G}$ and finally combine them by addition.

To alleviate the heterogeneity of various cancer types, we utilized inherent cancer labels available in TCGA for the inter-cancer pretraining objective. Specifically, [CLS] tokens of available modalities (regardless of whether they involved two or three modalities) would be concatenated into a single anchor representation $\boldsymbol{a}_i$. Furthermore, positive and negative samples were obtained within the mini-batch, and they were from the same cancer and different cancers, respectively. Similarly, they were constructed in the same way by concatenating the [CLS] tokens from available modalities, leading to $\boldsymbol{a}^+$ and $\boldsymbol{a}^-$ for positive and negative samples, respectively. Subsequently, we enforced a triplet loss $\mathcal{L}_{triplet}$ for them to bring the samples of the same cancer closer than that of the negative sample:

$$
\mathcal{L}_{triplet} = \frac{1}{N} \sum_{i=1}^{N} \max\left(d(\boldsymbol{a}_i, \boldsymbol{a}^+) - d(\boldsymbol{a}_i, \boldsymbol{a}^-) + \epsilon, 0\right) \quad (2)
$$

where $\boldsymbol{a}^+$ and $\boldsymbol{a}^-$ represent the farthest positive samples and nearest negative samples within a mini-batch, respectively, following the hard sample mining technique [46]. Here we used l_2 distance for function $d(\cdot)$ and ϵ is the margin which was set 0.3 based on smoother stability of loss degradation in the training set. Through these two pretraining objectives, as a result, we can get a well-trained slide aggregator that absorbed multimodal knowledge, which would be the 'Teacher' for the patch extractor at the next stage.

Stage 2 - Pretrain Patch Extractor. Upon finishing the first stage of pretraining, we can obtain a slide aggregator incorporating multimodal knowledge by being pretrained with multimodal data. In this stage, we leverage the pretrained slide aggregator as "Teacher" to seamlessly propagate multimodal knowledge into pathological patch extractor (ViT-L [47]), as shown in **Figure** 2b, which is termed Self-Taught training. Specifically, for each WSI, we gathered their patch features $\mathcal{P}_i = \{\mathbf{p}_i^m\}_{m=1}^{M}$ of the i-th WSI and fed them into the aggregator pretrained in the previous stage, where M refers to the number of patches of this WSI. Following the setting in the previous stage, M was fixed as 4,096. In this way, every patch can be re-embedded into new

features $\hat{\mathcal{P}}_i = \{\hat{\mathbf{p}}_i^m\}_{m=1}^M$ incorporating multimodal knowledge. With these re-embedded features as the objective guidance, we can pretrain a patch extractor by enforcing the extracted patch feature to get as close as possible to the ones re-embedded by the well-trained aggregator. To achieve this, for each patch, we can query its corresponding re-embedded feature $\hat{\mathbf{p}}_i^m$ encoded by the aggregator and further tuned the extractor with a loss function that minimizes the discrepancy between patch features encoded by the patch extractor and the corresponding re-embedded features incorporating multimodal knowledge:

$$\min \sum_m^M ||f(\mathbf{p}_i^m) - \hat{\mathbf{p}}_i^m||_1 \qquad (3)$$

where $f(\cdot)$ is a linear projection for adjusting the dimension of features and it transformed them into 512-dimensional features. Additionally, to avoid the catastrophic forgetting problem, a siamese structure is employed for the patch extractor consisting of two identical branches, where the parameters of one branch are updated using gradient descent, while the parameters of the other branch are updated using an Exponential Moving Average (EMA) of the parameters from the previous branch, without any gradient updates. Afterward, we enforced a similarity constraint between the patch features $\mathbf{p}_i^m$ extracted by the branch with gradient updates and those $\overline{\mathbf{p}}_i^m$ embedded by the branch with EMA updates. In the end, we combined two objectives into a loss function for pretraining the patch extractor:

$$\min \sum_m^{M_i} \lambda \cdot ||f(\mathbf{p}_i^m) - \hat{\mathbf{p}}_i^m||_1 + (1-\lambda) \cdot ||\mathbf{p}_i^m - \overline{\mathbf{p}}_i^m||_1$$
$$(4)$$

where λ is a balancing coefficient and it was set 0.6 based on smoother stability of loss degradation. By doing this, the patch extractor was enhanced by multimodal knowledge at the whole-slide context.

4.3 Downstream Tasks

Comparisons and Baselines. To investigate the benefit of enhancing the patch extractor by incorporating multimodal knowledge at the slide level, we compared mSTAR against one general baseline and three SOTA pretrained extractors commonly used in the CPath community: (1) ResNet50 [22] pretrained on ImageNet-1K [48], a commonly used baseline in many slide-level tasks [13, 9]. (2) PLIP [2], a vision-language (V-L) architecture (CLIP [43]) pretrained on OpenPath consisting of over 200k pathological patch-caption pairs. (3) CONCH [4], a V-L CoCa [49] framework with an additional generative loss pretrained on over 1.17 million pathological patch–caption pairs. (4) UNI [3], a pure vision patch extractor pretrained on more than 100 million patches from over 100k WSIs. Through pre-extracted patch features via these encoders, we can get 1024-dimensional (1024-d) embeddings for ResNet50, UNI, and mSTAR, and 512-dimensional (512-d) embeddings for PLIP and CONCH.

Unless otherwise specified, we obtained slide-level predictions by training the widely used attention-based multiple-instance learning (ABMIL) [11], a MIL aggregator that integrates all patch features of a WSI into the slide-level representation according to attention scores. We also employed TransMIL [13] as a MIL backbone, a transformer-like MIL architecture with linear time complexity, since we leveraged TransMIL serving a "Teacher" to propagate multimodal knowledge into the "Student" patch extractor through self-taught training in this work. We additionally presented results of mSTAR equipped with the pretrained aggregator, which was contrastively pretrained features embedded by our pretrained patch extractor mSTAR, resulting in the advanced version of the proposed approach termed mSTAR+.

In the experiments of multimodal fusion, we employed 4 existing SOTA multimodal integration models, MCAT [8], Porpoise [24], CMTA [10] and MOTCat [9]. For slide-level Zero- or Few-shot Classification that requires textual class prototypes, we considered the pretrained model as a good zero- or few-shot learner. As a result, following the paradigm of CLIP, we compared against those approaches that are equipped with the text encoder by utilizing the pretrained text encoder as a good classification head, including PLIP and CONCH. For those visual-only FMs, i.e., R50 and UNI, we employed the pretrained text encoder of mSTAR to construct class prototypes. For the pathological report generation,

we trained a recent model HistGen [34] based on pathological features from various FMs.

Slide-level Unimodal and Multimodal Tasks. For these tasks, we follow the conventional two-stage MIL paradigm comprising pre-extraction of patch features as instances and the training of a MIL aggregator that integrates patch features (or instances) into a single slide-level (or bag) feature. The aggregator took all patch features of a WSI as an input and mapped them into a hidden embedding as a single slide-level representation. Subsequently, the slide-level representation was passed through a fully connected classifier head, resulting in logits. Lastly, based on logits, we performed two types of slide-level tasks including classification supervised by cross-entropy loss with slide labels , and survival prediction (an ordinal regression task) supervised by NLL loss [50] with survival labels (Overall Survival in month), ranging from various diagnosis, treatment and prognosis tasks.

Classification. We performed Breast metastasis detection on CAMELYON (CAMELYON16+17) [14, 15], Prostate ISUP grading on PANDA [16], ovarian cancer subtyping on UBC-OCEAN [17, 18], BRCA subtyping on TCGA BRCA [19], NSCLC subtyping on TCGA NSCLC [19] and RCC subtyping on RCC-DHMC [20], resulting in 6 diverse tasks. For molecular predition, we also conducted evaluations on BCNB-ER, BCNB-HER2, BCNB-PR [21], TCGA BRCA, TCGA CRC and TCGA GBMLGG datasets. For all classification tasks, we label-stratified these datasets into 7:1:2 train-validation-test folds unless explicitly stated otherwise. For datasets that included TCGA data used in pretraining, we put the pretraining cases into the training split, while ensuring that the validation and test sets were held out from pretraining sources. Following the setting of previous works [4, 3], we treated one slide as one case separately. The model performing best in the validation split was chosen to be evaluated on the test set with 1000-times bootstrapping for 95% confidence interval (CI).

For all MIL backbones, we used the same hyper-parameters set for mSTAR and the competing FMs, in which the hidden dimensions are 512 and dropout keeps $p = 0.25$ after each intermediate layer in the network for regularization. We trained each model for 30 epochs on the training split by an Adam optimizer of the learning rate of 2×10^{-4} along with a cosine learning rate scheduler. The full set of hyperparameters is summarized in **Extended Data** Table A23.

Survival Prediction. For the survival prediction task, we performed evaluations on 9 TCGA datasets with 7:3 train-test splits, where each dataset consisted of more than 400 cases and the test set was merged from validation and test folds of pretraining splits and it was totally held out from pretraining data. In particular, CRC combined COAD and READ datasets and GBMLGG consisted of GBM and LGG. Following the conventional setting [8, 9], we concatenate features of all slides belonging to a single patient as one case for the patient-level prediction. To provide robust statistical performances, we conducted a 1000 times bootstrapping evaluation on the test set for 95% CI using the model performing best on the whole test set.

Regarding the hyperparameters of MIL backbones, we set the hidden dimension as 512 and dropout $p = 0.25$. In this task, we train each MIL model for 30 epochs optimized by Adam with the learning rate of 2×10^{-4} along with a cosine learning rate scheduler. For all optimization hyperparameters, refer to **Extended Data** Table A24.

The pretraining approach based on contrastive learning demonstrates promise in alleviating the heterogeneity of different modalities, thereby contributing to multimodal data integration. To verify this potential, in addition to vision-only models, we further evaluate multimodal fusion models for this task. Given the off-the-shelf multimodal fusion approaches that integrated WSIs and RNA-Seq data for survival prediction, we replaced their pathological features with ones embedded by pretrained extractors. It is worth highlighting that the training and evaluation of multimodal models followed the same splits as that of vision-only models, and we simply discarded those cases without paired RNA-Seq data.

For the aforementioned four existing multimodal integration models, we followed their default hyperparameters for these models, and detailed hyperparameters for each model are presented in **Extended Data** Table A25-A27. For Porpoise, the input length of RNA-Seq varies across different cancer datasets in TCGA and

the hidden dimension for RNA-Seq is fixed as 25, while the hidden dimension of pathological features was first transformed into 512 and then 256. Both modality branches adopted the dropout technique with $p = 0.1$. Lastly, features from two modalities were fused into a 256-dimensional slide-level feature. For MCAT and MOTCat, the hidden dimension of features was 256 for both modalities and dropout was 0.25 for regularization. Subsequently, features from two modalities were concatenated and integrated into a 256-dimensional slide-level representation. Similarly, CMTA followed the same hyperparameters except the hidden dimension of RNA-Seq which first became 1024 and then 256. For RNA-Seq data of MCAT, CMTA and MOTCat, embeddings were defined based on 6 functional categories according to [51] provided in MCAT by default, including 1) Tumor Supression, 2) Oncogenesis, 3) Protein Kinases, 4) Cellular Differentiation, 5) Transcription, and 6) Cytokines and Growth. More training hyperparameters are provided in **Extended Data** Table A28.

Few-shot Slide Classification. We performed evaluation on the same test set of each dataset as that of Slide Classification. Since our model was pretrained on TCGA data, we excluded classification datasets involving TCGA data. In this task, the pretrained model can be considered a good few-shot learner, which allows more label-efficient tasks. We proposed a new method utilizing pretrained patch extractors and the pretrained text encoder inspired by a non-parametric few-shot approach for WSIs, MI-SimpleShot [3] to evaluate the capability of few-shot learning for pretrained models , termed as MI-FewShot. Note that test sets were excluded from pretraining sources. Within the test set, we need to construct some episodes to evaluate the few-shot capability, including establishing the support set and query set. Due to the limited number of classes, we followed the same configuration as that of CONCH for the support set, in which evaluated classes were kept the same and the number of sampled slides k per class (i.e., "k-shot") was varying for training from $k = 1, 2, 4, 8, \ldots$. To ensure the low-sample scenarios, the maximum limit for sampled examples per class was set at either 256 or the maximum power of 2 that does not exceed the number of samples in the smallest

class. These sampled slides were used for creating slide-level class prototypes of the support set. The remaining samples, excluding those included in the support set, were considered as the query set for this episode. Due to the small support sizes, we randomly sampled k slides ("shot") per class with 5 repeated experiments (called "episodes") and reported the averaged results with error bars following the previous work [4].

Based on the support set, we need to construct the class prototypes using the sampled k slides per class. To achieve this, considering the pretrained text encoder is a good classifier, we first utilized it to establish textual prototypes of every class. To this end, we constructed a pool of prompt templates (see **Extended Data** Table A30) and ensembled all the prompts within a class by using the average text embedding as the class prototype. After obtaining textual class prototypes, following SimpleShot [52], we can use multiple patches to establish slide class prototypes based on k slides for each class. Specifically, every patch feature can be obtained by the pretrained extractor from k slides per class, followed by $L2$ normalization. Then based on cosine-similarity scores between the normalized patch feature and its corresponding textual prototypes associated with a given class, the average embedding of top-K patch features was used to represent this class. Subsequently, to obtain slide-level prediction for every query slide, we first embedded every patch of the query slide into a low-dimensional feature representation using the pretrained extractor, and then every $L2$ normalized patch feature was compared to each slide class prototype described above to calculate the cosine similarity score. As a result, the slide-level decision of this query was made by majority voting of the top-K patch predictions.

Zero-shot Slide Classification. We adopted an evaluation scheme similar to that of Few-shot Slide Classification to assess the ability of zero-shot learning. For all pretrained extractors, we employed a non-parametric MI-Zero [25] that does not rely on parametric training. The ensembling prompt of templates (the same as that of few-shot classification) was used as the textual classification, which was utilized to compute the cosine similarity score with every patch feature.

27

In the end, MI-Zero made the slide-level decision for every slide in the test set based on the majority voting of top-K scores. In addition, we conducted 1000 times bootstrapping on the test set for 95% confidence interval (CI). Other training configurations were consistent with those of Few-shot Slide Classification as described earlier.

Pathological Report Generation. We performed the assessment for this task on TCGA pathological datasets. Since report generation needs more samples to ensure robust generated results, we collected all WSIs and pathology reports used in pretraining as the training set to finetune the specific model of report generation, our prior work HistGen [34]. The held-out data from pretraining sources were used as validation and test sets. Given patients' pathology features from WSIs of each FM, HistGen is able to produce a sequence of words. Specifically, given extracted pathological features from the foundation model, the encoder-decoder architecture of HistGen would encode them into the latent features for report decoding. Subsequently, these features are utilized by the text decoder to generate the report. The quality of the generated report is directly influenced by the quality of the pathological features encoded by each FM. For all optimization hyperparameters, refer to **Extended Data** Table A29.

Statistical Analysis. Unless otherwise specified, we employ non-parametric bootstrapping with 1,000 bootstrap [53] replicates to estimate 95% confidence intervals for all experiments. To assess the observed differences in performance between the two models, we utilize a one-sided Wilcoxon signed-rank test [54] for statistical significance, following the previous work [5].

Evaluation Datasets

CAMELYON [14, 15] for Breast Metastasis Detection (2 classes). This dataset comprises 399 slides from the Cancer Metastases in Lymph Nodes Challenge 2016 (CAMELYON16) [14] and 500 slides from the CAMELYON17 challenge [15], resulting in 899 slides for the breast metastasis detection of two classes ("normal" v.s. "metastasis"). After removing a corrupted slide,

we obtained a total of 898 WSIs (557 normal, 341 metastasis). For training and evaluation, we employed the label-stratified 7:1:2 train-validation-test splits (630:91:180 slides).

PANDA [16] for Prostate ISUP grading (6 classes). Derived from the PANDA challenge [16], the ISUP (International Society of Urological Pathology) grading task includes a collection of 10,616 prostate cancer core needle biopsies for prostate cancer evaluation of 6 grades (also known as "classes"). After tissue segmentation, slides with a low tumor proportion were excluded, which resulted in 10,202 slides. For training and evaluation, we label-stratified PANDA into 7:1:2 train–validation–test folds (7,143:1,019:2,040 slides).

UBC-OCEAN [17, 18] for Ovarian Cancer Subtyping (5 classes). The UBC-OCEAN (University of British Columbia - Ovarian Cancer subtypE clAssification and outlier detectioN) dataset consists of 538 slides, which aims to classify ovarian cancer subtypes into 5 categories. After performing tissue segmentation, a total of 527 slides were acquired (98 CC, 122 EC, 221 HGSC, 43 LGSC and 43 MC). The class information is presented in **Extend Data** Table A33. For training and evaluation, we label-stratified the dataset into train–validation-test folds (369:52:106 slides).

TCGA-BRCA [19]for BRCA Subtyping (2 classes). The TCGA BRCA (Breast Invasive Carcinoma) dataset included BRCA H&E diagnostic histopathology WSIs from TCGA. This dataset encompassed cases of primary IDC (Invasive Ductal Carcinoma) and ILC (Invasive Lobular Carcinoma). After excluding slides with inadequate proportional tumor, a total of 985 slides were gathered, comprising 787 IDC and 198 ILC slides. Following the splits for pretraining, which approximately yielded 7:1:2 train-validation-test folds (656:95:234 slides), we ensure validation and test sets held out from pretraining sources.

TCGA-NSCLC [19] for NSCLC Subtyping (2 classes). The TCGA NSCLC (Non-Small Cell Lung Cancer) dataset comprised NSCLC H&E diagnostic slides from TCGA, including cases of primary lung adenocarcinoma (LUAD) and lung squamous cell carcinoma (LUSC). After tissue segmentation, a total of 1,053 slides were obtained, consisting of 541 LUAD and 512 LUSC slides. Similarly, we used the same pretraining

splits train-validation-test of an approximate ratio 7:1:2 (664:100:289 slides) to avoid data contamination.

RCC-DHMC [20] for RCC Subtyping (5 classes). Renal cell carcinoma (RCC) subtyping in this study involved WSIs of Benign (29 slides), Renal Oncocytoma (ROCY, 66 slides), Chromophobe RCC (CHRCC, 23 slides), Clear Cell RCC (CCRCC, 344 slides), and Papillary RCC (PRCC, 101 slides) obtained from the Dartmouth-Hitchcock Medical Center (DHMC). For training and evaluation, we label-stratified the dataset into train–validation-test folds (392:56:115 slides).

BCNB datasets [21] for ER (2 classes), PR (2 classes) and HER2 prediciton (2 classes). The Early Breast Cancer Core-Needle Biopsy (BCNB) WSI dataset, encompasses core-needle biopsy WSIs obtained from patients diagnosed with early breast cancer. The official provides data from 1058 patients, while we collected 1038 WSIs paired with ER, PR, HER2 status after tissue extraction. For training and evaluation, the train-validation-test cohort is label-stratified in a ratio of 7:1:2, resulting in training-validation-testing of 727:103:208 slides.

TCGA-BRCA for Molecular Prediction [55] (5 classes). This dataset is derived from the TCGA BRCA dataset, consisting of Basal (94 slides), HER2 (56 slides), LumA (228 slides), LumB (127 slides) and Normal (29 slides) classes. For training and evaluation, we label-stratified the dataset into train–validation-test cohorts (381:51:102 slides).

TCGA-CRC for Molecular Prediction [56] (4 classes). The dataset used in this study is derived from the TCGA CRC (Colon Adenocarcinoma and Rectum Adenocarcinoma) dataset, which includes the Colon Adenocarcinoma (COAD) and Rectum Adenocarcinoma (READ) datasets. It comprises four consensus molecular subtypes (CMSs). To facilitate training and evaluation, we stratified the dataset based on labels into train-validation-test cohorts with proportions of 344:49:99 slides, respectively.

TCGA-GBMLGG for Glioma IDH1 Mutation Prediction (2 classes). This dataset is specifically curated for predicting glioma IDH1 mutation. It consists of 804 FFPE&E-stained WSIs from the TCGA database, with a focus on the TCGA-GBM (Glioblastoma Multiforme) and TCGA-LGG (Brain Lower Grade Glioma) subsets. It is important to note that cases without reported IDH mutation status have been excluded from the dataset. The slides in this dataset are categorized into two classes: IDH-1 Wildtype (600 slides) and IDH-1 Mutant (204 slides). To facilitate training and evaluation, the train-validation-test cohort has been carefully stratified based on labels, with a ratio of 7:1:2. Consequently, there are 563 slides for training, 80 slides for validation, and 161 slides for testing.

9 TCGA Datasets for Unimodal and Multimodal Survival Analysis. In pretraining splits, we employed case- and label-stratified 7:1:2 training-validation-test splits for 9 TCGA cancer datasets of over 400 cases. We evaluated the capability of survival analysis on the same validation and test sets totally excluded from pretraining data. More information about the 9 TCGA cancer datasets were provided in **Extended Data** Table A21. For multimodal survival analysis, we employed the same splits and simply skip the cases without paired RNA-Seq data.

TCGA Datasets for Pathological Report Generation. During pretraining, we employed training-validation-test splits for some cancer datasets of over 400 cases and other data were put into pretraining materials. Following this setting, we considered all pretraining data containing pathology reports as the training set, and the held-out validation-test sets were re-used, resulting in 7073:452:934 cases for train-validation-test splits.

Evaluation Metrics

For classification tasks, Macro-AUC and its 95% confidence interval (CI) are reported considering alleviating the impact of unbalanced data, which doesn't depend on the selection of the decision threshold and is not affected by the sample ratio of classes. For survival prediction tasks, we report the commonly used Concordance Index (C-Index) and its 95% CI, which is defined as the probability that two randomly selected individuals will have risk predictions correctly ordered. For few-shot and zero-shot classification, we also reported the same metric as the ones in classification tasks. In pathological report generation, in line with our prior studies HistGen [34], we report various metrics, BLEU@K [57], METEOR [58] and

ROUGE-L [59], to assess the accuracy of predicted captions against the ground-truth captions from different perspectives. BLEU@K measures the similarity between machine-generated text and ground truth by comparing the presence and frequency of n-grams. METEOR is a metric that evaluates precision and recall by matching unigrams while also factoring in synonyms and word variations between the original text and the reference. On the other hand, ROUGE-L measures the similarity in n-gram overlap between the generated texts and the ground truth.

Computing Software and Hardware. We conducted all experiments and analyses in this study using Python (v3.11.5) and PyTorch (v2.2.1, CUDA 11.7) (https://pytorch.org) unless stated otherwise, and these can be reproduced with open-source libraries as described below. To pretrain aggregator, the implementaion of the text encoder pretrained on PubMed was maintained by the codebase (https://github.com/dmis-lab/biobert) and its pretrained weights can be assessed in the open-source timm library from Hugging Face (https://huggingface.co). For extractor pretraining, we initialize the backbone with the pretrained weights of UNI codebase (https://github.com/mahmoodlab/uni). OpenSlide (v3.4.1) and openslide-python (v1.3.1) were utilized to support the processing of WSIs in conjunction with CLAM (https://github.com/mahmoodlab/CLAM). Implementations of other visual pretrained encoders compared in the study can be accessed through the following links: ResNet-50 pretrained on ImageNet-1K (https://github.com/mahmoodlab/CLAM), PLIP (https://github.com/PathologyFoundation/plip) and CONCH (https://github.com/mahmoodlab/CONCH). Implementations of zero-shot learning for WSIs were provided in MI-Zero (https://github.com/mahmoodlab/MI-Zero). For training MIL models for downstream tasks, we adapted the code of ABMIL from the CLAM codebase (https://github.com/mahmoodlab/CLAM) and TransMIL from its source codebase (https://github.com/szc19990412/TransMIL/tree/main). For multimodal survival prediction, we used the off-the-shelf multimodal fusion models: MCAT (https://github.com/mahmoodlab/MCAT), Porpoise (https://github.com/mahmoodlab/PORPOISE), MOTCat (https://github.com/Innse/MOTCat) and CMTA (https://github.com/FT-ZHOU-ZZZ/CMTA).

For pathological report generation, HistGen (https://github.com/dddavid4real/HistGen) is applied. We used 4×80 GB NVIDIA H800 GPUs (graphics processing unit) for pretraining aggregator and a single 80 GB NVIDIA H800 GPU for pretraining extractor. These GPUs were set up for multi-GPU, multi-node training, employing distributed data-parallel (DDP) techniques. All other experiments for downstream tasks were conducted on single 24 GB NVIDIA 3090 GPUs or single 80 GB H800 GPU.

Data availability. This study incorporates a total of TCGA [19] data across 32 cancers including WSIs (https://portal.gdc.cancer.gov/) and their reports (https://github.com/dddavid4real/HistGen). RNASeq data are collected from cBioPortal (https://www.cbioportal.org/). Extra data for evaluation are as follows: Camelyon16&17 (https://camelyon16.grand-challenge.org/Data/), PANDA (https://www.kaggle.com/c/prostate-cancer-grade-assessment/data), UBC-OCEAN (https://www.kaggle.com/competitions/UBC-OCEAN/data), BCNB (https://bcnb.grand-challenge.org/) and RCC-DHMC (https://bmirds.github.io/KidneyCancer/).

Code availability. The code and weights of mSTAR will be released in the codebase upon paper acceptance.

Ethics declarations. This study has been reviewed and approved by the Human and Artefacts Research Ethics Committee (HAREC). The protocol number is HREP-2024-0212.

Author contributions. Y.X., Y.W. and H.C. conceived the study. Y.X. designed the pretraining approach and experiments, and prepared the manuscript. Y.W. performed foundation model pretraining and conducted evaluation on downstream tasks including unimodal survival analysis, few/zero-shot classification and report generation. F.Z. conducted evaluation of slide-level subtyping, molecular prediction and multimodal fusion for survival analysis. J.M. collected the data and assisted in results analysis and coding. S.Y. participated in coding for pretraining. H.L. assisted in the experimental design. X.W., L.L., A.H. and

R.C.K.C. assisted in the design of evaluation on clinical tasks and the interpretation of experimental results. J.W. assisted in curating RNA-Seq data and offered suggestions for experimental designs. H.C. supervised the research.

Declarations. The authors have no conflicts of interest to declare.

Appendix A   Extended Data

(a) TCGA-D1-A179

TCGA-D1-A179

Doctor

Slide Disposition:
Diagnosis:
A. Uterus:
- Hysterectomy
- High-grade serous carcinoma identified
- Tumor information: Mass measuring 70 x 41 x 0.4 cm, circumferentially involving the endometrial cavity
- Tumor invasion: 0.4 cm into the myometrium
- Total myometrial thickness: 16 cm
- No involvement of the endocervix
- Lymphovascular space invasion identified
- Negative margins for tumor

B. Lymph Nodes (Right and Left Para-aortic lymphadenectomy):
- Multiple lymph nodes (6 out of 9) positive for metastatic carcinoma consistent with endometrial primary

C. Ovary and Fallopian Tube (Right Salpingo-oophorectomy):
- High-grade serous carcinoma forming

CONCH

Specimens:
A. Uterus, cervix, bilateral tubes, and ovaries
B. Left pelvic lymph nodes
C. Left pelvic lymph nodes
D. Left pelvic lymph nodes
E. Left pelvic lymph nodes
F. Left pelvic lymph nodes

Diagnosis:
A. Uterus, cervix, bilateral tubes, and ovaries:
- Endometrial adenocarcinoma, endometrioid type
- FIGO grade 3
- Nuclear grade 3
- Solid growth pattern
- Myometrial invasion is identified.
- The extent of invasion is confined to the endometrium.
- Cervix is present.
- The cervix is confined to the endometrium.

mmSTAR (Ours)

Diagnosis:
A. Uterus, bilateral ovaries, and fallopian tubes:
- Hysterectomy and bilateral salpingo-oophorectomy
- FIGO Grade II of III endometrial adenocarcinoma, endometrioid type
- Tumor information: Mass measuring 28 x 25 x 22 cm, located in the anterior uterine wall
- Tumor invasion: 2 cm into the myometrium
- Total myometrial thickness: 18 cm
- No involvement of the endocervix
- No lymphovascular space invasion identified
- Negative margins for tumor
- Bilateral ovaries and fallopian tubes are negative for tumor

PLIP

Primary Site: Cervix
Diagnosis: Uterine cervix biopsy
- Poorly differentiated and invasive squamous cell carcinoma

UNI

Tissue Description:
- A1, A2, A3, A4, A5, A6, A7, A8, A9, A10, J1, K1, L1, M1, N1
- B1, B2, B3, B4, B5, B6, B7, B8, B9, B10, B11, B12, B13, B14, B15
- C1, C2, C3, C4, C5, C6, C7, C8, C9, C10, C11, C12
- D1, D2, D3, D4, D5, D6, D7, D8, D9, D10, D11
- E1, E2, E3, E4
- F1
- G1, G2
- H1, H2, H3, H4
- IL, I2, I3

Uterus, right ovary (20 x 10 x 11 cm) with a 90 cm segment of right fallopian tube, and left ovary (20 x 10 x 08 cm) with a 20

R50

Final Diagnosis:
Part 1: Kidney Left Radical Nephrectomy:
A. Renal cell carcinoma, conventional clear cell type
- Fuhrman's nuclear grade: 3 out of 4
- Size: 4 cm in greatest dimension
B. The neoplasm is confined within the renal capsule.
C. No invasion of the renal vein is identified.
D. Renal vein is identified.
E. No evidence of angiolymphatic invasion is identified.
F. All surgical margins are free of tumor.
G. Non-neoplastic kidney is unremarkable.
H. No evidence of adrenal gland is identified.
Pathologic Stage: PT1b Nx Mx

(a) TCGA-D1-A179

(b) TCGA-BS-A0V6

TCGA-BS-A0V6

Doctor

Medical Record Number: N/A
History and Clinical Diagnosis: Endometrial adenocarcinoma, Grade 1
Postoperative Diagnosis: Pending pathology examination
Specimens Received:
A. Uterus, cervix, bilateral tubes, and ovaries
B. Left pelvic lymph node
C. Left aortic lymph node
D. Right pelvic lymph node
E. Right aortic lymph node
Diagnosis:
A. Uterus, cervix, bilateral tubes, and ovaries:
- Endometrial adenocarcinoma
- Operative Procedure: Total abdominal hysterectomy and bilateral salpingo-oophorectomy (TAH-BSO) with staging
- Histologic Type: Endometrioid
- Histologic Grade: FIGO Grade 1
- Nuclear Grade: 2
- Tumor Size: 35 cm
- Extent of Invasion: Less than 50% of myometrium, 1 mm invasion, 18 mm myometrial thickness
- Lymphovascular invasion:

CONCH

Specimens Submitted:
1. Uterus
2. Cervix
3. Bilateral fallopian tubes
4. Ovaries
5. Left external iliac lymph nodes
6. Left obturator lymph nodes
7. Left obturator lymph nodes
8. Left obturator lymph nodes
9. Left obturator lymph nodes
10. Left obturator lymph nodes
11. Left common iliac lymph nodes
12. Right common iliac lymph nodes
13. Left common iliac lymph nodes
The submitted specimens include the uterus, cervix, bilateral fallopian tubes, ovaries, and several lymph nodes from different regions of the pelvis and abdomen.

mmSTAR (Ours)

Client History/Clinical Diagnosis: Endometrial cancer
Postoperative Diagnosis: Pending pathology examination
Specimens Received:
A. Uterus, cervix, bilateral tubes, and ovaries
B. Left pelvic lymph node
C. Right pelvic lymph node
D. Left pelvic lymph node
E. Left pelvic lymph node
F. Uterus, cervix, tubes, and ovaries
Diagnosis:
A. Uterus, cervix, bilateral tubes, and ovaries:
- Endometrial adenocarcinoma, endometrioid type, FIGO grade I
- Tumor confined to the uterine fundus
- Tumor size: 50 x 30 x 25 cm
- No myometrial invasion identified
- The tumor invades the myometrium to a depth of

PLIP

Specimens:
A. Right pelvic lymph nodes
B. Left pelvic lymph nodes
C. Prostate
Gross Description:
A. Right pelvic lymph nodes:
- Received fresh and labeled with the patient's identification.
- Multiple fragments of fatty tissue measuring 06 x 06 x 05 cm.
- The specimen is bisected and entirely submitted in a single cassette.
B. Right pelvic lymph nodes:
- Received fresh and labeled with the patient's identification.
- No further information is provided.
C. Prostate:
- No specific information is provided regarding the gross description.

UNI

Page 1: Examination and Macroscopic Description
Histopathological Examination:
Material: Multiple organ resection - uterus, cervix, tubes, and ovaries.
Expected time of examination: Up to 8 working days.
Clinical Diagnosis:
Macroscopic Description:
- A segment of the large intestine measuring [unspecified] cm in length.
- Omentum tissue sized 12 x 8 x 5 cm in thickness.
- The omentum sized 12 x 10 x 10 cm in the mucosa.
- The lesion surrounds 100% of the intestine circumference and narrows its lumen.
- The endometrial cavity measures 12 x 11 cm in the lower uterine segment, located 12 cm from the

R50

Final Diagnosis:
Part 1: Kidney Left Radical Nephrectomy:
A. Renal cell carcinoma, conventional clear cell type
- Fuhrman's nuclear grade: 3 out of 4
- Size: 4 cm in greatest dimension
B. The neoplasm is confined within the renal capsule.
C. No invasion of the renal vein is identified.
D. Renal vein invasion is identified.
E. No evidence of angiolymphatic invasion is identified.
F. All surgical margins are free of tumor.
G. Non-neoplastic kidney is unremarkable.
H. No evidence of adrenal gland involvement is identified.
Pathologic Stage: PT1b Nx Mx

(b) TCGA-BS-A0V6

Doctor (TCGA-DF-A2L0)

Sample Information:
Sample Type: Tumor
Sample Preparation: Fresh frozen
Site of Primary: Uterine
Site of Tissue: Endometrial mass
Year of Sample Collection: Unknown
Age at Sample Collection: Unknown years
Sample Comments: None
Days to Procedure Date: 0
Days to Diagnosis: 9
Pathological Information:
Type of Procedure: Surgical resection
Site of Primary Histology: Endometrium
Bilateral Disease: Unknown
Tumor Size: Unknown
Histology: Endometrioid adenocarcinoma NOS
Grade/Differentiation: III
Pathological T: T1c
Pathological N: Unknown (Information not available regarding the number of nodes sampled)
Number of Nodes Positive: 17
Clinical M: M0
Histology Comments: Information not available

UNI

Sample Information:
Sample Type: Tumor
Diagnosis: Squamous cell carcinoma, not otherwise specified (NOS)
Year of Sample Collection: Unknown/not specified
Age at Sample Collection: Unknown/not specified years
Days to Procedure Date: Unknown/not specified
Days to Diagnosis: Unknown/not specified
Pathological Information:
Type of Procedure: Resection
Site of Tissue/Primary Histology: Uterus
Tumor Size: 3 cm
Histology: Endometrioid adenocarcinoma NOS
Grade/Differentiation: III
Pathological T: T3
Pathological N: Unknown/not specified
Clinical M: M0
Histology Comments: None provided
Additional Information:
Sample Type: Buffy
Year of Sample Collection: Unknown/not specified
Age at Sample Collection: Unknown/not specified years
Days to Procedure Date: Unknown/not specified
Days to Diagnosis: Unknown/not specified

PLIP

Primary Site: Cervix
Histopathology Findings:
Uterine cervix biopsy
Diagnosis: Poorly differentiated and invasive squamous cell carcinoma

mmSTAR (Ours)

Sample Information:
Sample Type: Tumor
Sample Preparation: Fresh frozen
Site of Primary: Uterine
Site of Tissue: Cervix
Year of Sample Collection: Unknown
Age at Sample Collection: Unknown years
Days to Procedure Date: Unknown
Days to Diagnosis: Unknown
Pathological Information:
Type of Procedure: Surgical resection
Site of Primary Histology: Uterus
Bilateral Disease: Unknown
Tumor Size: 55 cm
Histology: Endometrioid adenocarcinoma NOS
Grade/Differentiation: III
Pathological T: T1b
Pathological N: NX
Clinical M: M0
Histology Comments: Unknown
Additional Information:
Sample Type: Buffy
Year of Sample Collection: Unknown
Age at Sample Collection: Unknown years
Days to Procedure Date: Unknown
Days to Diagnosis: Unknown

CONCH

Specimens Submitted:
1. Uterus
2. Cervix
3. Bilateral fallopian tubes
4. Bilateral ovaries
5. Right external iliac lymph nodes
6. Left obturator lymph nodes
7. Right hypogastric lymph nodes
8. Right obturator lymph nodes
9. Right hypogastric lymph nodes
10. Right hypogastric lymph nodes
11. Right hypogastric lymph nodes
12. Right hypogastric lymph nodes
13. Right hypogastric lymph nodes
14. Right external iliac lymph nodes
15. Right obturator lymph nodes

R50

Final Diagnosis:
Part 1: Kidney Left Radical Nephrectomy:
A. Renal cell carcinoma, conventional clear cell type
- Fuhrman's nuclear grade: 3 out of 4
- Size: 4 cm in greatest dimension
B. The neoplasm is confined within the renal capsule.
C. No invasion of the renal vein is identified.
D. Renal vein invasion is identified.
E. No evidence of angiolymphatic invasion is identified.
F. All surgical margins are free of tumor.
G. Non-neoplastic kidney is unremarkable.
H. No evidence of adrenal gland involvement is identified.
Pathologic Stage: PT1b Nx Mx

(c) TCGA-DF-A2L0

Fig. A1: **Pathology Reports Generation** for several cases.

Table A1: **Comparison of the aggregator pretrained on different combinational modalities (Pathology, Reports and RNASeq) for survival analysis.** Note that pathological patch features are extracted by UNI. Average C-Index and its std across 9 TCGA survival datasets are reported.

Aggregator (TransMIL)	Pretrained Data			Avg C-Index
	Pathology	Reports	RNASeq	
w/o pretraining				0.666±0.068
w/ pretraining	✓	✓		0.679±0.057
	✓		✓	0.679±0.059
	✓	✓	✓	**0.689±0.055**

Table A2: **Macro-AUC of Slide Classification with ABMIL on 12 tasks**. Best performing model for each metric is bolded. 95% CI is included in parentheses.

Task		R50	PLIP	CONCH	UNI	mSTAR(Ours)
Diagnosis	CAMELYON	0.9236 (0.8905,0.9567)	0.9242 (0.8915,0.9569)	0.9823 (0.9654,0.9992)	0.9814 (0.9683,0.9945)	**0.9934 (0.9865,1.0000)**
	PANDA	0.8872 (0.8811,0.8933)	0.9045 (0.8988,0.9102)	0.9220 (0.9169,0.9271)	0.9455 (0.9410,0.9500)	**0.9468 (0.9423,0.9513)**
	UBC-OCEAN	0.8940 (0.8644,0.9236)	0.9453 (0.9230,0.9676)	0.9724 (0.9606,0.9842)	0.9732 (0.9612,0.9852)	**0.9764 (0.9637,0.9891)**
	TCGA-BRCA	0.7560 (0.6876,0.8244)	0.8533 (0.8047,0.9019)	0.9386 (0.9119,0.9653)	0.9391 (0.9038,0.9744)	**0.9441 (0.9110,0.9772)**
	TCGA-NSCLC	0.9111 (0.8876,0.9346)	0.9395 (0.9193,0.9597)	0.9723 (0.9613,0.9833)	0.9695 (0.9558,0.9832)	**0.9730 (0.9607,0.9853)**
	RCC-DHMC	0.9589 (0.9440,0.9738)	0.9889 (0.9820,0.9958)	0.9835 (0.9727,0.9943)	0.9889 (0.9818,0.9960)	**0.9898 (0.9837,0.9959)**
Molecular Prediction	BCNB-ER	0.7974 (0.7466,0.8482)	0.8460 (0.8058,0.8862)	0.8503 (0.8072,0.8934)	0.8886 (0.8506,0.9266)	**0.8887 (0.8519,0.9255)**
	BCNB-HER2	0.6683 (0.6075,0.7291)	0.6789 (0.6203,0.7375)	**0.7465 (0.6938,0.7992)**	0.7409 (0.6890,0.7928)	0.7419 (0.6896,0.7942)
	BCNB-PR	0.6889 (0.6289,0.7489)	0.8041 (0.7520,0.8562)	0.8092 (0.7592,0.8592)	0.8297 (0.7787,0.8807)	**0.8335 (0.7829,0.8841)**
	BRCA-Molecular	0.7017 (0.6637,0.7397)	0.7353 (0.6924,0.7782)	0.7506 (0.7161,0.7851)	0.7613 (0.7235,0.7991)	**0.7646 (0.7279,0.8013)**
	CRC-Molecular	0.6304 (0.5830,0.6778)	0.7451 (0.7016,0.7886)	0.8382 (0.7998,0.8766)	0.8548 (0.8227,0.8869)	**0.8560 (0.8227,0.8893)**
	GBMLGG-IDH1	0.8787 (0.8250,0.9324)	0.9345 (0.9022,0.9668)	0.9777 (0.9644,0.9910)	0.9891 (0.9820,0.9962)	**0.9902 (0.9831,0.9973)**
Avg Macro-AUC		0.8080	0.8583	0.8953	0.9052	**0.9082**
Avg Rank		5.00	3.88	2.75	2.29	**1.08**

Table A3: **Macro-AUC of Slide Classification with TransMIL on 12 tasks**. Best performing model for each metric is bolded. 95% CI is included in parentheses.

Task		R50	PLIP	CONCH	UNI	mSTAR(Ours)
Diagnosis	CAMELYON	0.9164 (0.8819,0.9509)	0.9161 (0.8808,0.9514)	0.9577 (0.9308,0.9846)	**0.9719 (0.9550,0.9888)**	0.9689 (0.9446,0.9932)
	PANDA	0.8870 (0.8809,0.8931)	0.8929 (0.8868,0.8990)	0.9108 (0.9053,0.9163)	0.9311 (0.9264,0.9358)	**0.9317 (0.9270,0.9364)**
	UBC-OCEAN	0.8891 (0.8536,0.9246)	0.9106 (0.8836,0.9376)	0.9589 (0.9409,0.9769)	0.9643 (0.9482,0.9804)	**0.9722 (0.9595,0.9849)**
	TCGA-BRCA	0.7553 (0.6936,0.8170)	0.8511 (0.8005,0.9017)	0.9335 (0.9061,0.9609)	0.9453 (0.9122,0.9784)	**0.9480 (0.9141,0.9819)**
	TCGA-NSCLC	0.9035 (0.8786,0.9284)	0.9321 (0.9111,0.9531)	0.9699 (0.9564,0.9834)	**0.9725 (0.9609,0.9841)**	0.9647 (0.9473,0.9821)
	RCC-DHMC	0.9674 (0.9541,0.9807)	0.9622 (0.9451,0.9793)	0.9850 (0.9768,0.9932)	0.9902 (0.9843,0.9961)	**0.9918 (0.9869,0.9967)**
Molecular Prediction	BCNB-ER	0.7508 (0.6953,0.8063)	0.8124 (0.7618,0.8630)	0.8705 (0.8297,0.9113)	0.7993 (0.7478,0.8508)	**0.8791 (0.8379,0.9203)**
	BCNB-HER2	0.4955 (0.4355,0.5555)	0.6638 (0.6028,0.7248)	0.7035 (0.6447,0.7623)	0.6916 (0.6338,0.7494)	**0.7155 (0.6606,0.7704)**
	BCNB-PR	0.6327 (0.5670,0.6984)	0.7178 (0.6586,0.7770)	**0.8062 (0.7564,0.8560)**	0.7832 (0.7293,0.8371)	0.8029 (0.7537,0.8521)
	BRCA-Molecular	0.6476 (0.6076,0.6876)	0.7174 (0.6749,0.7599)	0.6773 (0.6338,0.7208)	0.6638 (0.6222,0.7054)	**0.7638 (0.7299,0.7977)**
	CRC-Molecular	0.6214 (0.5675,0.6753)	0.6147 (0.5633,0.6661)	0.6914 (0.6422,0.7406)	0.7765 (0.7349,0.8181)	**0.8186 (0.7784,0.8588)**
	GBMLGG-IDH1	0.9259 (0.8916,0.9602)	0.9282 (0.8953,0.9611)	0.9639 (0.9463,0.9815)	**0.9640 (0.9419,0.9861)**	0.9556 (0.9276,0.9836)
Avg Macro-AUC		0.7807	0.8275	0.8682	0.8702	**0.8924**
Avg Rank		4.83	3.92	2.50	2.25	**1.50**

Table A4: **Macro-AUC of Slide Classification with the pretrained TransMIL on 12 datasets**. The performance gains of the aggregator (TransMIL) between "before" and "after" pretraining. "Positive"(+) values indicate performance increases (↑) compared with the ones equipped with TransMIL built from scratch, and vice versa. The best improvement is bolded.

Task		R50	PLIP	CONCH	UNI	mSTAR (Ours)
Diagnosis	CAMELYON	-0.16%	-1.15%	-0.50%	+1.60%	**+2.31%**
	PANDA	+0.21%	-0.61%	-0.03%	**+0.95%**	+0.86%
	UBC-OCEAN	+2.03%	**+3.19%**	+0.24%	+0.55%	-0.41%
	TCGA-BRCA	**+9.10%**	-0.84%	-0.31%	-0.91%	-0.52%
	TCGA-NSCLC	**+2.07%**	+1.18%	-1.12%	+0.26%	+1.06%
	RCC-DHMC	**+3.64%**	+1.54%	+1.54%	+1.01%	-0.05%
Molecular Prediction	BCNB-ER	+7.27%	-0.56%	-2.19%	**+6.29%**	+0.64%
	BCNB-HER2	**+12.28%**	+1.76%	+3.39%	+5.08%	+1.73%
	BCNB-PR	**+9.86%**	+0.64%	-1.70%	+1.40%	-0.92%
	BRCA-Molecular	+3.06%	-0.46%	+7.70%	**+9.00%**	-0.83%
	CRC-Molecular	-0.67%	**+14.94%**	+13.10%	+4.85%	+2.14%
	GBMLGG-IDH1	+0.75%	+1.05%	+0.27%	+2.10%	**+2.66%**
Avg increases		**+4.12%**	+1.72%	+1.70%	+2.68%	+0.72%

Table A5: **C-Index of Survival Analysis with ABMIL on 9 TCGA datasets**. Best performing model for each metric is bolded. 95% CI is included in parentheses.

Task	R50	PLIP	CONCH	UNI	mSTAR (Ours)
BRCA (1023)	0.5696 (0.4446,0.6945)	0.5915 (0.4764,0.7065)	0.6940 (0.6052,0.7828)	0.6908 (0.5860,0.7957)	**0.7076 (0.6180,0.7972)**
CRC (579)	0.5710 (0.4562,0.6857)	**0.7076 (0.6186,0.7966)**	0.7026 (0.6149,0.7902)	0.6906 (0.6071,0.7742)	0.6895 (0.6059,0.7731)
GBMLGG (830)	0.7810 (0.7425,0.8196)	0.7716 (0.7307,0.8124)	0.7704 (0.7269,0.8139)	0.7905 (0.7471,0.8339)	**0.7923 (0.7497,0.8349)**
HNSC (441)	0.6087 (0.5297,0.6878)	0.6569 (0.5783,0.7355)	0.6294 (0.5478,0.7109)	0.6516 (0.5718,0.7315)	**0.6604 (0.5810,0.7398)**
KIRC (498)	0.6258 (0.5392,0.7123)	0.6613 (0.5751,0.7474)	0.6957 (0.6178,0.7736)	**0.7155 (0.6496,0.7814)**	0.7027 (0.6137,0.7917)
LUAD (455)	0.5995 (0.5071,0.6920)	0.5906 (0.4978,0.6834)	0.6233 (0.5302,0.7163)	0.6312 (0.5316,0.7308)	**0.6329 (0.5353,0.7305)**
LUSC (452)	0.5305 (0.4395,0.6214)	0.5464 (0.4614,0.6314)	0.6045 (0.5226,0.6865)	0.6273 (0.5502,0.7044)	**0.6323 (0.5538,0.7108)**
SKCM (415)	0.5811 (0.5074,0.6547)	0.5648 (0.4893,0.6402)	0.6210 (0.5403,0.7017)	0.6254 (0.5478,0.7030)	**0.6281 (0.5520,0.7041)**
UCEC (495)	0.7450 (0.6630,0.8271)	0.7738 (0.6835,0.8640)	0.8082 (0.7080,0.9085)	0.7845 (0.6833,0.8858)	**0.8092 (0.7227,0.8957)**
Avg C-Index	0.6236	0.6516	0.6832	0.6897	**0.6950**
Avg Rank	4.56	3.67	3.00	2.33	**1.44**

Table A6: C-Index of Survival Analysis with TransMIL on 9 TCGA datasets. Best performing model for each metric is bolded. 95% CI is included in parentheses.

Task	R50	PLIP	CONCH	UNI	mSTAR (Ours)
BRCA (1023)	0.6028 (0.4883,0.7174)	0.6193 (0.5436,0.6950)	0.6326 (0.5350,0.7302)	0.6649 (0.5657,0.7641)	**0.6937 (0.6098,0.7776)**
CRC (579)	0.6222 (0.5148,0.7296)	0.6595 (0.5537,0.7654)	0.6211 (0.5042,0.7380)	0.6277 (0.5074,0.7480)	**0.6672 (0.5654,0.7690)**
GBMLGG (830)	**0.7830 (0.7456,0.8204)**	0.7731 (0.7344,0.8117)	0.7778 (0.7384,0.8172)	0.7768 (0.7347,0.8190)	0.7829 (0.7404,0.8254)
HNSC (441)	0.6078 (0.5280,0.6875)	0.6255 (0.5498,0.7011)	0.5783 (0.4995,0.6572)	0.6811 (0.6114,0.7508)	**0.6824 (0.6109,0.7539)**
KIRC (498)	0.6868 (0.6109,0.7627)	0.6659 (0.5853,0.7465)	0.7080 (0.6327,0.7833)	0.7071 (0.6281,0.7860)	**0.7171 (0.6369,0.7972)**
LUAD (455)	0.5745 (0.4660,0.6829)	0.5902 (0.4868,0.6936)	0.6121 (0.5100,0.7143)	0.5980 (0.4952,0.7009)	**0.6455 (0.5494,0.7417)**
LUSC (452)	0.5983 (0.5068,0.6899)	0.5375 (0.4481,0.6270)	0.5651 (0.4732,0.6571)	0.5985 (0.5082,0.6888)	**0.6272 (0.5404,0.7140)**
SKCM (415)	0.5569 (0.4799,0.6340)	0.5537 (0.4799,0.6276)	0.5941 (0.5218,0.6664)	0.5931 (0.5219,0.6643)	**0.6311 (0.5647,0.6976)**
UCEC (495)	0.6973 (0.5729,0.8218)	0.7004 (0.6021,0.7988)	0.7651 (0.6766,0.8537)	0.7561 (0.6756,0.8366)	**0.7754 (0.6988,0.8520)**
Avg C-Index	0.6366	0.6361	0.6505	0.6670	**0.6914**
Avg Rank	3.89	4.11	3.11	2.78	**1.11**

Table A7: C-Index of Survival Analysis with the pretrained TransMIL on 9 TCGA datasets. The performance gains of the aggregator (TransMIL) between "before"" and "after" pretraining. 'Positive'(+) values indicate performance increases (↑) compared with the ones equipped with TransMIL built from scratch, and vice versa. The best improvement is bolded.

Task	R50	PLIP	CONCH	UNI	mSTAR (Ours)
BRCA (1023)	-2.06%	+1.40%	**+2.34%**	+0.08%	+0.08%
CRC (579)	-9.90%	+2.61%	+3.17%	**+3.24%**	+1.56%
GBMLGG (830)	+0.54%	**+1.93%**	+0.37%	-0.38%	-1.11%
HNSC (441)	+1.04%	+1.69%	**+5.39%**	-1.02%	+0.21%
KIRC (498)	-4.73%	**+3.57%**	+2.58%	+0.95%	+1.69%
LUAD (455)	**+6.98%**	+2.04%	-0.57%	+4.78%	+0.15%
LUSC (452)	+0.50%	+1.51%	**+2.95%**	-3.19%	-1.41%
SKCM (415)	**+4.29%**	+0.55%	+2.97%	+0.68%	+1.53%
UCEC (495)	-3.55%	**+7.79%**	+2.10%	+4.20%	+3.87%
Avg increases	-0.77%	**+2.57%**	+2.37%	+1.04%	+0.73%

Table A8: C-Index of Multimodal Survival Analysis with MCAT on 9 TCGA datasets. Best performing model for each metric is bolded. 95% CI is included in parentheses.

Task	R50	PLIP	CONCH	UNI	mSTAR (Ours)
BRCA	0.7215 (0.6382,0.8048)	0.7304 (0.6428,0.8180)	0.6266 (0.5019,0.7513)	0.6928 (0.6034,0.7822)	**0.7489 (0.6703,0.8275)**
CRC	0.6996 (0.5994,0.7998)	0.6676 (0.5669,0.7683)	0.6635 (0.5601,0.7669)	0.6481 (0.5474,0.7488)	**0.6989 (0.5943,0.8035)**
GBMLGG	0.8713 (0.8343,0.9083)	0.8784 (0.8433,0.9135)	0.8886 (0.8553,0.9219)	0.8810 (0.8471,0.9149)	**0.8896 (0.8610,0.9182)**
HNSC	0.6830 (0.6054,0.7606)	0.6444 (0.5677,0.7211)	0.6646 (0.6008,0.7284)	0.6874 (0.6125,0.7623)	**0.6907 (0.6157,0.7657)**
KIRC	0.6856 (0.6130,0.7582)	0.6830 (0.6088,0.7572)	0.6996 (0.6335,0.7657)	**0.7091 (0.6291,0.7891)**	0.7084 (0.6407,0.7761)
LUAD	0.6375 (0.5326,0.7424)	0.6505 (0.5576,0.7434)	0.6628 (0.5681,0.7575)	0.6595 (0.5733,0.7457)	**0.6953 (0.6073,0.7833)**
LUSC	0.5373 (0.4420,0.6326)	0.5499 (0.4654,0.6344)	0.5526 (0.4660,0.6392)	0.5493 (0.4488,0.6498)	**0.5537 (0.4590,0.6484)**
SKCM	**0.6724 (0.5944,0.7504)**	0.6641 (0.5941,0.7341)	0.6522 (0.5781,0.7263)	0.6549 (0.5794,0.7304)	0.6520 (0.5840,0.7200)
UCEC	0.6406 (0.4967,0.7845)	0.6380 (0.5007,0.7753)	0.6754 (0.5532,0.7976)	0.6807 (0.5681,0.7933)	**0.6811 (0.5751,0.7871)**
Avg C-Index	0.6832	0.6785	0.6762	0.6848	**0.7021**
Avg Rank	3.44	3.67	3.22	3.00	**1.67**

Table A9: C-Index of Multimodal Survival Analysis with Porpoise on 9 TCGA datasets. Best performing model for each metric is bolded. 95% CI is included in parentheses.

	R50	PLIP	CONCH	UNI	mSTAR (Ours)
BRCA	0.7366 (0.6480,0.8252)	0.7244 (0.6297,0.8191)	0.7259 (0.6369,0.8149)	**0.7474 (0.6651,0.8297)**	0.7456 (0.6515,0.8397)
CRC	0.6113 (0.4978,0.7248)	0.6639 (0.5571,0.7707)	0.6517 (0.5535,0.7499)	0.6599 (0.5621,0.7577)	**0.7020 (0.6085,0.7955)**
GBMLGG	0.8805 (0.8472,0.9138)	0.8748 (0.8407,0.9089)	0.8786 (0.8406,0.9166)	0.8948 (0.8627,0.9269)	**0.8955 (0.8643,0.9267)**
HNSC	0.6982 (0.6271,0.7693)	0.6806 (0.5977,0.7635)	0.6730 (0.5981,0.7479)	0.7111 (0.6415,0.7807)	**0.7189 (0.6474,0.7904)**
KIRC	0.6681 (0.5934,0.7428)	0.6715 (0.5921,0.7509)	0.6802 (0.5973,0.7631)	0.6767 (0.6046,0.7488)	**0.6834 (0.6115,0.7553)**
LUAD	**0.6656 (0.5792,0.7520)**	0.6222 (0.5266,0.7178)	0.6139 (0.5265,0.7013)	0.6295 (0.5454,0.7136)	0.6300 (0.5455,0.7145)
LUSC	0.5364 (0.4490,0.6238)	0.5100 (0.4136,0.6064)	0.4911 (0.3980,0.5842)	0.5794 (0.4859,0.6729)	**0.5895 (0.5027,0.6763)**
SKCM	**0.6383 (0.5666,0.7100)**	0.6215 (0.5498,0.6932)	0.6253 (0.5510,0.6996)	0.5995 (0.5233,0.6757)	0.6182 (0.5443,0.6921)
UCEC	0.6154 (0.4809,0.7499)	0.6301 (0.4968,0.7634)	0.6565 (0.5218,0.7912)	0.6565 (0.5385,0.7745)	**0.6724 (0.5340,0.8108)**
Avg C-Index	0.6723	0.6666	0.6662	0.6839	**0.6951**
Avg Rank	3.22	3.89	3.72	2.61	**1.56**

Table A10: **C-Index of Multimodal Survival Analysis with MOTCat on 9 TCGA datasets**. Best performing model for each metric is bolded. 95% CI is included in parentheses.

	R50	PLIP	CONCH	UNI	mSTAR (Ours)
BRCA	0.7030 (0.6075,0.7985)	0.6732 (0.5685,0.7779)	0.6781 (0.5789,0.7773)	0.6970 (0.5814,0.8126)	**0.7278 (0.6276,0.8280)**
CRC	0.6867 (0.5926,0.7808)	0.6695 (0.5635,0.7755)	0.6643 (0.5553,0.7733)	0.6835 (0.5765,0.7905)	**0.7244 (0.6327,0.8161)**
GBMLGG	0.8838 (0.8526,0.9150)	0.8736 (0.8375,0.9097)	0.8867 (0.8555,0.9179)	**0.8899 (0.8583,0.9215)**	0.8889 (0.8583,0.9195)
HNSC	0.6776 (0.6045,0.7507)	0.6731 (0.6023,0.7439)	0.6679 (0.5852,0.7506)	0.6736 (0.5930,0.7542)	**0.7256 (0.6546,0.7966)**
KIRC	**0.7124 (0.6467,0.7781)**	0.6890 (0.6108,0.7672)	0.6927 (0.6163,0.7691)	0.7050 (0.6286,0.7814)	0.7085 (0.6403,0.7767)
LUAD	0.6463 (0.5559,0.7367)	0.6472 (0.5549,0.7395)	0.6685 (0.5758,0.7612)	0.6710 (0.5875,0.7545)	**0.7004 (0.6138,0.7870)**
LUSC	0.5618 (0.4748,0.6488)	0.5564 (0.4621,0.6507)	0.5643 (0.4704,0.6582)	0.5654 (0.4780,0.6528)	**0.5855 (0.4957,0.6753)**
SKCM	0.6660 (0.5874,0.7446)	0.6607 (0.5868,0.7346)	0.6454 (0.5715,0.7193)	0.6604 (0.5834,0.7374)	**0.6715 (0.5972,0.7458)**
UCEC	0.6446 (0.5190,0.7702)	0.6899 (0.5601,0.8197)	0.7006 (0.5869,0.8143)	0.7031 (0.5943,0.8119)	**0.7353 (0.6453,0.8253)**
Avg C-Index	0.6869	0.6814	0.6854	0.6943	**0.7187**
Avg Rank	3.00	4.33	3.89	2.56	**1.22**

Table A11: **C-Index of Multimodal Survival Analysis with CMTA on 9 TCGA datasets**. Best performing model for each metric is bolded. 95% CI is included in parentheses.

	R50	PLIP	CONCH	UNI	mSTAR (Ours)
BRCA	0.7118 (0.6224,0.8012)	0.7039 (0.5945,0.8133)	0.7196 (0.6157,0.8235)	0.7339 (0.6351,0.8327)	**0.7375 (0.6560,0.8190)**
CRC	0.6719 (0.5610,0.7828)	0.6651 (0.5514,0.7788)	0.6580 (0.5469,0.7691)	0.6727 (0.5714,0.7740)	**0.7040 (0.6080,0.8000)**
GBMLGG	0.8798 (0.8447,0.9149)	0.8820 (0.8479,0.9161)	0.8822 (0.8501,0.9143)	0.9016 (0.8726,0.9306)	**0.9051 (0.8798,0.9304)**
HNSC	0.6764 (0.6051,0.7477)	0.6782 (0.6059,0.7505)	0.6744 (0.5964,0.7524)	0.6710 (0.5928,0.7492)	**0.7092 (0.6430,0.7754)**
KIRC	0.6858 (0.6158,0.7558)	0.6752 (0.5984,0.7520)	0.6963 (0.6175,0.7751)	0.7000 (0.6100,0.7900)	**0.7065 (0.6279,0.7851)**
LUAD	0.6502 (0.5536,0.7468)	0.6301 (0.5386,0.7216)	**0.6639 (0.5645,0.7633)**	0.6380 (0.5376,0.7384)	0.6574 (0.5684,0.7464)
LUSC	0.5473 (0.4569,0.6377)	0.5234 (0.4303,0.6165)	0.5352 (0.4413,0.6291)	0.5909 (0.4980,0.6838)	**0.6320 (0.5497,0.7143)**
SKCM	0.6514 (0.5783,0.7245)	**0.6546 (0.5786,0.7306)**	0.6413 (0.5643,0.7183)	0.6446 (0.5693,0.7199)	0.6488 (0.5731,0.7245)
UCEC	0.6986 (0.5786,0.8186)	0.6676 (0.5504,0.7848)	0.6932 (0.5736,0.8128)	0.7307 (0.6100,0.8514)	**0.7555 (0.6669,0.8441)**
Avg C-Index	0.6859	0.6756	0.6849	0.6982	**0.7173**
Avg Rank	3.33	4.00	3.56	2.78	**1.33**

Table A12: **Average Macro-AUC of Few-shot Slide Classification across different shots on 6 datasets**. Best performing model for each metric is in bold. Mean±std is presented.

	R50	PLIP	CONCH	UNI	mSTAR(Ours)
CAMELYON	0.525±0.019	0.546±0.033	0.504±0.020	0.666±0.087	**0.707±0.063**
PANDA	0.574±0.020	**0.664±0.045**	0.648±0.031	0.604±0.022	0.621±0.030
UBC-OCEAN	0.589±0.022	0.723±0.012	0.715±0.027	0.764±0.023	**0.774±0.032**
BCNB-ER	0.563±0.025	0.582±0.034	0.603±0.068	0.623±0.049	**0.624±0.065**
BCNB-PR	0.535±0.049	0.599±0.027	0.614±0.017	0.640±0.039	**0.660±0.030**
BCNB-HER2	0.559±0.054	0.514±0.014	0.524±0.034	**0.574±0.039**	0.542±0.016
Avg Macro-AUC	0.5575	0.6046	0.6012	0.6457	**0.6548**
Avg Rank	4.33	3.33	3.50	2.17	**1.67**

Table A13: **Macro-AUC of Few-shot Slide Classification on CAMELYON across different shots**. Best performing model for each metric is bolded. Mean±std is presented.

k-shot	R50	PLIP	CONCH	UNI	mSTAR(Ours)
k = 1	0.515±0.053	0.511±0.019	0.523±0.037	0.571±0.114	**0.659±0.086**
k = 2	0.486±0.030	0.521±0.026	0.489±0.045	0.587±0.107	**0.600±0.090**
k = 4	0.545±0.035	0.517±0.037	0.487±0.067	0.561±0.075	**0.660±0.023**
k = 8	0.539±0.051	0.537±0.012	0.539±0.061	0.678±0.083	**0.742±0.051**
k = 16	0.522±0.067	0.571±0.015	0.495±0.053	0.714±0.050	**0.737±0.014**
k = 32	0.531±0.020	0.551±0.031	0.514±0.094	0.764±0.036	**0.788±0.029**
k = 64	0.537±0.083	0.612±0.152	0.480±0.136	**0.787±0.079**	0.766±0.096
Avg Macro-AUC	0.5250	0.5457	0.5039	0.6660	**0.7074**
Avg Rank	3.93	3.71	4.36	1.86	**1.14**

Table A14: **Macro-AUC of Few-shot Slide Classification on PANDA across different shots**. Best performing model for each metric is bolded. Mean±std is presented.

k-shot	R50	PLIP	CONCH	UNI	mSTAR(Ours)
k = 1	0.552±0.057	0.564±0.052	**0.584±0.064**	0.564±0.059	0.570±0.056
k = 2	0.535±0.036	0.619±0.033	**0.634±0.037**	0.572±0.040	0.574±0.057
k = 4	0.576±0.077	**0.677±0.006**	0.675±0.028	0.623±0.042	0.632±0.057
k = 8	0.588±0.014	**0.678±0.010**	0.660±0.031	0.609±0.010	0.643±0.023
k = 16	0.566±0.030	0.678±0.022	**0.693±0.018**	0.617±0.041	0.629±0.039
k = 32	0.590±0.024	**0.700±0.008**	0.652±0.026	0.621±0.034	0.624±0.035
k = 64	0.588±0.038	**0.696±0.008**	0.655±0.026	0.623±0.011	0.650±0.028
k = 128	0.598±0.005	**0.701±0.010**	0.634±0.022	0.606±0.009	0.646±0.025
Avg Macro-AUC	0.5741	**0.6641**	0.6484	0.6044	0.6210
Avg Rank	5.00	**1.56**	1.75	3.94	2.75

Table A15: **Macro-AUC of Few-shot Slide Classification on UBC-OCEAN across different shots**. Best performing model for each metric is bolded. Mean±std is presented.

k-shot	R50	PLIP	CONCH	UNI	mSTAR(Ours)
k = 1	0.585±0.069	0.706±0.036	**0.744±0.035**	0.726±0.048	0.722±0.049
k = 2	0.625±0.024	0.739±0.033	0.672±0.078	0.786±0.017	**0.790±0.019**
k = 4	0.584±0.047	0.722±0.061	0.729±0.037	0.763±0.035	**0.779±0.034**
k = 8	0.563±0.077	0.726±0.071	0.714±0.042	0.779±0.069	**0.806±0.089**
Avg Macro-AUC	0.5893	0.7232	0.7147	0.7635	**0.7742**
Avg Rank	5.00	3.50	3.00	2.00	**1.50**

Table A16: Macro-AUC of Few-shot Slide Classification on BCNB-ER across different shots. Best performing model for each metric is bolded. Mean±std is presented.

k-shot	R50	PLIP	CONCH	UNI	mSTAR(Ours)
k = 1	0.514±0.084	0.539±0.104	**0.627±0.106**	0.595±0.090	0.606±0.112
k = 2	0.598±0.048	0.633±0.061	**0.679±0.082**	0.614±0.095	0.617±0.143
k = 4	0.571±0.036	0.557±0.079	**0.674±0.110**	0.549±0.162	0.527±0.169
k = 8	0.559±0.073	0.553±0.123	0.610±0.071	**0.670±0.115**	0.621±0.116
k = 16	0.572±0.070	0.608±0.048	0.509±0.083	0.614±0.081	**0.622±0.085**
k = 32	0.561±0.054	0.604±0.069	0.517±0.096	0.697±0.047	**0.748±0.051**
Avg Macro-AUC	0.5625	0.5823	0.6027	0.6232	**0.6235**
Avg Rank	4.00	3.33	2.67	2.67	**2.33**

Table A17: Macro-AUC of Few-shot Slide Classification on BCNB-PR across different shots. Best performing model for each metric is bolded. Mean±std is presented.

k-shot	R50	PLIP	CONCH	UNI	mSTAR(Ours)
k = 1	0.551±0.035	0.576±0.069	0.614±0.096	0.640±0.096	**0.648±0.137**
k = 2	0.507±0.025	0.582±0.037	0.582±0.049	0.595±0.075	**0.645±0.071**
k = 4	0.532±0.092	0.567±0.093	0.608±0.102	0.606±0.107	**0.629±0.160**
k = 8	0.602±0.022	0.629±0.071	0.621±0.087	0.656±0.039	**0.715±0.028**
k = 16	0.571±0.120	0.596±0.058	0.625±0.101	**0.716±0.029**	0.686±0.057
k = 32	0.447±0.058	0.641±0.010	0.634±0.078	**0.645±0.149**	0.638±0.073
Avg Macro-AUC	0.5350	0.5985	0.6140	0.6430	**0.6602**
Avg Rank	5.00	3.42	3.25	1.83	**1.50**

Table A18: Macro-AUC of Few-shot Slide Classification on BCNB-HER2 across different shots. Best performing model for each metric is bolded. Mean±std is presented.

k-shot	R50	PLIP	CONCH	UNI	mSTAR(Ours)
k = 1	0.493±0.050	0.491±0.067	0.502±0.068	0.523±0.080	**0.534±0.081**
k = 2	0.532±0.027	0.507±0.033	0.548±0.032	**0.565±0.079**	0.554±0.064
k = 4	0.498±0.051	0.522±0.046	**0.562±0.060**	0.530±0.037	0.518±0.071
k = 8	**0.611±0.033**	0.509±0.029	0.555±0.062	0.589±0.053	0.561±0.041
k = 16	**0.629±0.037**	0.518±0.025	0.506±0.013	0.610±0.080	0.556±0.015
k = 32	0.592±0.056	0.536±0.029	0.468±0.019	**0.627±0.040**	0.531±0.036
Avg Macro-AUC	0.5592	0.5138	0.5235	**0.5740**	0.5423
Avg Rank	2.83	4.17	3.50	**1.67**	2.83

Table A19: **Macro-AUC of Zero-shot Slide Classification on 6 datasets**. Best performing model for each metric is in bold. 95% CI is included in parentheses.

	R50	PLIP	CONCH	UNI	mSTAR(Ours)
CAMELYON	0.541 (0.455,0.627)	0.622 (0.534,0.710)	0.638 (0.550,0.726)	0.685 (0.601,0.769)	**0.742 (0.668,0.817)**
PANDA	0.480 (0.466,0.494)	0.561 (0.547,0.575)	**0.582 (0.566,0.598)**	0.450 (0.434,0.466)	0.462 (0.446,0.478)
UBC-OCEAN	0.445 (0.369,0.521)	**0.762 (0.691,0.833)**	0.370 (0.307,0.433)	0.694 (0.620,0.769)	0.692 (0.621,0.763)
BCNB-ER	0.631 (0.535,0.727)	0.563 (0.473,0.653)	0.500 (0.500,0.500)	0.618 (0.524,0.712)	**0.725 (0.641,0.809)**
BCNB-PR	0.554 (0.458,0.650)	0.523 (0.435,0.611)	0.500 (0.500,0.500)	0.682 (0.598,0.766)	**0.722 (0.644,0.800)**
BCNB-HER2	**0.549 (0.455,0.643)**	0.516 (0.430,0.602)	0.518 (0.428,0.608)	0.465 (0.373,0.557)	0.435 (0.345,0.525)
Avg Macro-AUC	0.5333	0.5912	0.5180	0.5990	**0.6297**
Avg Rank	3.00	3.00	3.50	3.17	**2.33**

Table A20: **Performance on Pathological Report Generation**. Best performing model for each metric is in bold. 6 metrics are reported.

Metric	R50	PLIP	CONCH	UNI	mSTAR (Ours)
BLEU-1	0.2510	0.2288	0.2491	0.3590	**0.3662**
BLEU-2	0.1240	0.1189	0.1527	0.2490	**0.2585**
BLEU-3	0.0692	0.0727	0.1059	0.1884	**0.1981**
BLEU-4	0.0410	0.0491	0.0787	0.1490	**0.1583**
METEOR	0.0898	0.0915	0.1099	0.1570	**0.1613**
ROUGE-L	0.1825	0.1774	0.2153	0.2991	**0.3070**

Table A21: **Data splits (cases) of 9 cancer datasets on TCGA**. CRC=COAD+READ and GBMLGG=GBM+LGG. Apart from these specific cancer types, other TCGA datasets were used for pretraining without data partitioning.

Cancer Type	Train	Validation	Test	Total
BRCA	716	102	307	1023
CRC	405	58	116	579
GBMLGG	581	83	166	830
HNSC	308	44	89	441
KIRC	348	50	100	498
LUAD	318	45	92	455
LUSC	316	45	91	452
SKCM	290	41	84	415
UCEC	346	49	100	495

Table A22: **Prompts** used in GPT-4 for cleaning pathology reports.

Helping me to check the formatting and spelling of the supplied text, including some incorrect use of punctuation like mis-using of 'X' and 'x' and capitalization and deletion of some words of unknown meaning as well.

Table A23: **Hyperparameters of ABMIL and TransMIL used for Slide Classification**. A single 80GB NVIDIA H800 GPU was used for training.

Hyperparameter	Value
Input dim	512/1024
Hidden dim	512
Dropout	0.25
Batch size	1
Epochs	30
Optimizer	Adam
Learning rate	2e-4
Scheduler	Cosine
Weight decay	1e-5

Table A24: **Hyperparameters of ABMIL and TransMIL used for Survival Analysis**. A single 80GB NVIDIA H800 GPU was used for training.

Hyperparameter	Value
Input dim	512/1024
Hidden dim	512
Dropout	0.25
Batch size	1
Epochs	30
Optimizer	Adam
Learning rate	2e-4
Scheduler	Cosine
Weight decay	1e-5

Table A25: **Architecture Hyperparameters of Porpoise used for Multimodal Survival Analysis**. A single 80GB NVIDIA H800 GPU was used for training. D is the number of genes, varying across different datasets.

Hyperparameter	WSI	RNASeq
Input dim	512/1024	D
Hidden dim	$512 \rightarrow 256$	256
Dropout	0.1	0.1
Feature dim after fusion	256	

Table A26: **Architecture Hyperparameters of MCAT and MOTCat used for Multimodal Survival Analysis**. A single 80GB NVIDIA H800 GPU was used for training. D is the number of genes, varying across different unique functional categories and different datasets. $256 \times 2 \rightarrow 256$ refers to 256-dimensional features from two modalities were fused into a 256-dimensional slide-level feature.

Hyperparameter	WSI	RNASeq
Input dim	512/1024	$6 \times D$
Hidden dim	256	6×256
Dropout	0.25	0.25
Feature dim after fusion	$256 \times 2 \rightarrow 256$	

Table A27: **Architecture Hyperparameters of CMTA used for Multimodal Survival Analysis**.
A single 80GB NVIDIA H800 GPU was used for training. D is the number of genes, varying across
different datasets. $256 \times 2 \rightarrow 256$ refers to 256-dimensional features from two modalities were fused into
a 256-dimensional slide-level feature.

Hyperparameter	WSI	RNASeq
Input dim	512/1024	$6 \times D$
Hidden dim	256	$6 \times (1024 \rightarrow 256)$
Dropout	0.25	0.25
Feature dim after fusion	$256 \times 2 \rightarrow 256$	

Table A28: **Training Hyperparameters used in Multimodal Survival Analysis**. A single 80GB
NVIDIA H800 GPU was used for training.

Hyperparameter	Value
Batch size	1
Epochs	30
Optimizer	Adam
Learning rate	2e-4
Scheduler	Cosine
Weight decay	1e-5

Table A29: **Architecture Hyperparameters of HistGen used for Pathology Report Generation**. A single 80GB NVIDIA H800 GPU was used for training. D is the feature dimension of pathological
patch, varying across different FMs.

Hyperparameter	Value
Input dim	D
Hidden dim	512
Layers of decoder	3 (8 heads)
Dimension of cross-modal context	512×2048
Optimizer	Adam
Learning rate	1e-4
Beam size	3

Table A30: Prompt templates for all tasks that required prompts. In these templates, the placeholder "CLASSNAME" was replaced with the actual name of the class. The class prompts of CLASSNAME for each task are presented in A31-A36.

#	Prompt
1	CLASSNAME.
2	a photomicrograph showing CLASSNAME.
3	a photomicrograph of CLASSNAME.
4	an image of CLASSNAME.
5	an image showing CLASSNAME.
6	an example of CLASSNAME.
7	CLASSNAME is shown.
8	this is CLASSNAME.
9	there is CLASSNAME.
10	a histopathological image showing CLASSNAME.
11	a histopathological image of CLASSNAME.
12	a histopathological photograph of CLASSNAME.
13	a histopathological photograph showing CLASSNAME.
14	shows CLASSNAME.
15	presence of CLASSNAME.
16	CLASSNAME is present.
17	an H&E stained image of CLASSNAME.
18	an H&E stained image showing CLASSNAME.
19	an H&E image showing CLASSNAME.
20	an H&E image of CLASSNAME.
21	CLASSNAME, H&E stain.
22	CLASSNAME, H&E.
23	Close-up view of CLASSNAME.
24	Detailed image capturing CLASSNAME.
25	Microscopic analysis reveals CLASSNAME.
26	CLASSNAME under magnification.
27	The structure of CLASSNAME is visible.
28	Diagnosing CLASSNAME.
29	Typical appearance of CLASSNAME.
30	Zoomed image of CLASSNAME.
31	CLASSNAME in detail.
32	CLASSNAME at high magnification.
33	CLASSNAME in a clinical sample.
34	Clinical representation of CLASSNAME.
35	Pathological review of CLASSNAME.
36	Characteristics of CLASSNAME observed.
37	CLASSNAME identified.
38	Diagnosis: CLASSNAME.
39	Specimen showing CLASSNAME.
40	View of CLASSNAME with immunohistochemistry.
41	CLASSNAME with special staining.
42	Observation of CLASSNAME with fluorescent staining.

Table A31: Class prompts for Breast Metastasis Detection.

Task	Class	CLASSNAME
Breast Metastasis Detection	Tumor	metastatic cancer lymph node with metastatic cancer metastatic breast cancer in lymph node tumor
	Normal	non-metastatic lymph node without metastatic cancer non-metastatic lymph node normal

Table A32: Class prompts of PANDA for Prostate ISUP grading.

Task	Class	CLASSNAME
Prostate ISUP grading	0	benign tissue, healthy tissue, non-cancerous tissue, normal tissue
	1	Gleason score 6, ISUP grade 1, low-grade cancer, Gleason pattern 3+3
	2	Gleason score 7 (3+4), ISUP grade 2, moderately differentiated cancer, Gleason pattern 3+4
	3	Gleason score 7 (4+3), ISUP grade 3, moderately differentiated cancer, Gleason pattern 4+3
	4	Gleason score 8, ISUP grade 4, high-grade cancer, Gleason pattern 4+4, 3+5, or 5+3
	5	Gleason score 9 or 10, ISUP grade 5, high-grade aggressive cancer, Gleason pattern 4+5, 5+4, or 5+5

Table A33: Class prompts of UBC-OCEAN for Ovarian Cancer Subtyping.

Task	Class	CLASSNAME
Ovarian Cancer Subtyping	HGSC	high-grade serous carcinoma, ovarian high-grade serous carcinoma, high-grade serous carcinoma of ovary
	LGSC	low-grade serous carcinoma, ovarian low-grade serous carcinoma, low-grade serous carcinoma of ovary
	EC	endometrioid carcinoma, ovarian endometrioid carcinoma, endometrioid carcinoma of ovary
	CC	clear cell carcinoma, ovarian clear cell carcinoma, clear cell carcinoma of ovary
	MC	mucinous carcinoma, ovarian mucinous carcinoma, mucinous carcinoma of ovary

Table A34: Class prompts of BCNB for ER Prediction.

Task	Class	CLASSNAME
ER Prediction	Positive	ER positive, ER+, ER positive breast cancer, Estrogen Receptor positive, ER+ breast cancer, ER positive carcinoma, Estrogen Receptor positive breast cancer
	Negative	ER negative, ER-, ER negative breast cancer, Estrogen Receptor negative, ER- breast cancer, ER negative carcinoma, Estrogen Receptor negative breast cancer

Table A35: Class prompts of BCNB for PR Prediction.

Task	Class	CLASSNAME
PR Prediction	Positive	PR positive, PR+, PR positive breast cancer, Progesterone Receptor positive, PR+ breast cancer, PR positive carcinoma, Progesterone Receptor positive breast cancer
	Negative	PR negative, PR-, PR negative breast cancer, Progesterone Receptor negative, PR- breast cancer, PR negative carcinoma, Progesterone Receptor negative breast cancer

Table A36: Class prompts of BCNB for HER2 Prediction.

Task	Class	CLASSNAME
HER2 Prediction	Positive	HER2 positive, HER2+, HER2 positive breast cancer, HER2 overexpression, HER2+ breast cancer, HER2 positive carcinoma, Human Epidermal growth factor Receptor 2 positive, Human Epidermal growth factor Receptor 2 overexpression
	Negative	HER2 negative, HER2-, HER2 negative breast cancer, no HER2 overexpression, HER2- breast cancer, HER2 negative carcinoma, Human Epidermal growth factor Receptor 2 negative