EXPLOITING LOW-DIMENSIONAL MANIFOLD OF FEATURES FOR FEW-SHOT WHOLE SLIDE IMAGE CLASSIFICATION

Anonymous authors

000

001

002 003 004

005

006 007 008

010 011

012

013

014

015

016

018

019

021

023

024

027

029

030

033

035

036

037

038

040

041

042

043

045

Paper under double-blind review

ABSTRACT

Few-shot Whole Slide Image (WSI) classification is severely hampered by overfitting. We argue that this is not merely a data-scarcity issue but a fundamentally geometric problem. Grounded in the manifold hypothesis, our analysis shows that features from pathology foundation models exhibit a low-dimensional manifold geometry that is easily perturbed by downstream models. This insight reveals a key potential issue in downstream multiple instance learning models: linear layers are geometry-agnostic and, as we show empirically, can distort the manifold geometry of the features. To address this, we propose the Manifold Residual (MR) block, a plug-and-play module that is explicitly geometry-aware. The MR block reframes the linear layer as residual learning and decouples it into two pathways: (1) a fixed, random matrix serving as a geometric anchor that approximately preserves topology while also acting as a spectral shaper to sharpen the feature spectrum; and (2) a trainable, low-rank residual pathway that acts as a residual learner for task-specific adaptation, with its structural bottleneck explicitly mirroring the low effective rank of the features. This decoupling imposes a structured inductive bias and reduces learning to a simpler residual fitting task. Through extensive experiments, we demonstrate that our approach achieves state-of-the-art results with significantly fewer parameters, offering a new paradigm for few-shot WSI classification. Code will be released upon acceptance.

1 Introduction

Analysis of gigapixel Whole Slide Images (WSIs) is the gold standard of modern pathology (Chen et al., 2024; Guo et al., 2024; Lu et al., 2021; Xiong et al., 2024a;b), yet it operates under a uniquely challenging paradigm defined by two constraints. First, the gigapixel WSIs makes Multiple Instance Learning (MIL) the de facto paradigm, where a WSI is represented as a bag of patch features (Ilse et al., 2018; Shao et al., 2021; Xiong et al., 2023). Second, the expert-intensive annotation process results in datasets that are not only small (few-shot) but also weakly labeled, typically with only slide-level labels (Lu et al., 2021). This few-shot, weakly-supervised setting forces models to identify discriminative patterns from thousands of unlabeled instances using only a few slide-level labels (Guo et al., 2025b), leading to severe overfitting.

To understand the root cause of overfitting, we *look beyond the learning algorithm and into the intrinsic structure of the features themselves*. Grounded in the manifold hypothesis (Tenenbaum et al., 2000), we investigate the intrinsic geometry of the CONCH features (Lu et al., 2024), presenting a multi-faceted line of evidence that they exhibit a low-dimensional, nonlinear manifold geometry. First, spectral analysis (Skean et al., 2025) reveals a low effective rank of 29.7 (*vs.* CONCH dimension 512), confirming low dimensionality (Fig. 1(a)). Second, t-SNE visualization exposes a distinct cluster topology (Fig. 1(b)). Finally, and most critically, tangent space analysis shows a non-flat distance-dependent growth of geometric drift (Fig. 1(d), original), providing quantitative evidence of the intrinsic curvature of the manifold structure.

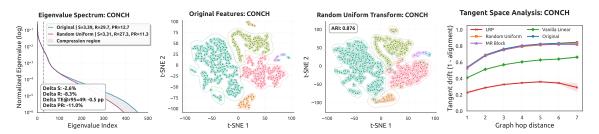


Figure 1: Geometric analysis of features reveals a low-dimensional manifold, which linear layers distort and our MR block preserves. (a) Spectral analysis confirms low-dimensionality. (b, c) t-SNE shows a cluster topology, and a random transformation preserves it. (d) Tangent space analysis visualizes this preservation: our MR block maintains the intrinsic curvature of the manifold, while a linear layer can cause distortion.

Therefore, we posit that the cause of overfitting is geometric: while features from pathology foundation models exhibit a low-dimensional manifold geometry, existing MIL models often fail to preserve this fragile structure. This is not incidental but a limitation rooted in their primary building block: the linear layer. Across the diverse MIL architectures, including ABMIL (Ilse et al., 2018), Transformers (Shao et al., 2021; Tang et al., 2024), information-bottleneck (Huang et al., 2024) and density-based (Zhu et al., 2024) approaches, the linear layer serves as the indispensable component. However, the linear layer is inherently geometry-agnostic, lacking an architectural bias to respect the manifold structure. Our tangent space analysis provides direct, empirical evidence of this destructive tendency: as shown in Fig. 1(d), a trained linear layer (green line) severely distorts the intrinsic geometry of the manifold compared to the original (blue line). Consequently, especially under the few-shot setting, these linear layers are prone to learning an overly complex mapping that does not respect the low-rank nature of features and distorts their structure, thereby discarding the geometric priors learned during large-scale pretraining. This points to a fundamental research gap: a lack of mechanisms designed to actively preserve and leverage this crucial, yet fragile, manifold structure.

To address this, we propose the Manifold Residual (MR) block, a plug-and-play block. The MR block aims to better preserve the geometric structure of the pretrained features. It reframes the mapping as residual learning and decouples it into two parallel paths: a frozen base path for geometric preservation and a trainable low-rank residual path (LRP) for task-specific adaptation. The base path employs a fixed random matrix with uniform initialization that serves as both a geometric anchor and a spectral shaper. Theoretically, random projection theory shows that such a transformation provides a geometry-preserving embedding that, with high probability, approximately preserves pairwise Euclidean distances (Johnson et al., 1984; Dasgupta & Gupta, 2003) and local geodesic structure (Baraniuk & Wakin, 2009), and acts as a near-isometry on tangent spaces under standard smoothness assumptions (Papaspiliopoulos, 2020). Empirically, our Fig. 1 shows that the geometric anchor preserves the cluster topology (c), retains the original distance—drift profile in tangent space (d), and sharpens the eigenvalue spectrum (a). The LRP employs a bottleneck structure to model the inherently low-rank task-specific residuals, explicitly mirroring the low-rank features. Although the LRP alone alters geometry to fit task-specific residuals (Fig. 1(d), red), the anchor balances this effect; together, the MR block as a whole (Fig. 1(d), purple) maintains manifold geometry while capturing task-specific signals. The final representation is the sum of the two paths, balancing structural stability with adaptivity.

The contribution of our paper can be summarized as follows:

- 1. We present quantitative evidence consistent with a low-dimensional manifold structure in foundation-model features, and identify a common failure mode: linear layers can disrupt this fragile structure.
- 2. We propose the Manifold Residual block, a theoretically grounded, plug-and-play module that mitigates manifold degradation through its geometric anchor and low-rank residual design.

3. We conduct extensive experiments, including comparative experiments, ablation studies, and sensitivity analyses, to validate both the efficacy and efficiency of our method.

2 RELATED WORK

2.1 MULTIPLE INSTANCE LEARNING

In MIL, WSIs are tessellated into patches, which are then input into a pretrained neural network to obtain features. Regarding obtaining slide-level logits using these features, MIL can be categorized into two streams: *instance-level* and *embedding-level* approaches. In the instance-level paradigm, logits are first computed for each patch and then aggregated to yield a slide-level prediction (Campanella et al., 2019; Chikontwe et al., 2020; Hou et al., 2016; Kanavati et al., 2020). Conversely, embedding-level approaches first aggregate patch features via pooling or aggregator networks to form a slide-level embedding, which is subsequently used to generate the final prediction (Guo et al., 2025a; Huang et al., 2024; Ilse et al., 2018; Lin et al., 2023; Li et al., 2024b; Lu et al., 2021; Shao et al., 2021; Tang et al., 2024; Xiong et al., 2023; Zhang et al., 2022).

2.2 PATHOLOGY FOUNDATION MODELS AND FEW-SHOT WSI CLASSIFICATION

Previously, feature extraction relied on ImageNet-pretrained ResNet (He et al., 2016). However, the domain discrepancy between natural images and WSIs often limits their efficacy in pathology domain. To address this, both vision and vision-language pathology foundation models have been developed (Xiong et al., 2025). For example, Virchow (Vorontsov et al., 2024), UNI (Chen et al., 2024), and GPFM (Ma et al., 2024) are pretrained on pathology image datasets, and VLMs, including PLIP (Huang et al., 2023), CONCH (Lu et al., 2024), mSTAR (Xu et al., 2024) and MUSK (Xiang et al., 2025), leverage image-text pairs to enhance multimodal tasks. These pathology foundation models enable few-shot weakly-supervised learning (FSWSL) of WSI classification, which is especially appealing due to the substantial costs and labor incurred by WSI annotations. FSWSL was introduced in TOP (Qu et al., 2023), and subsequent studies (Fu et al., 2024; Guo et al., 2025b; Li et al., 2024a; Qu et al., 2024; Shi et al., 2024) have further expanded on this paradigm.

2.3 GEOMETRIC DEEP LEARNING AND REPRESENTATION ANALYSIS

The manifold hypothesis, central to geometric deep learning, posits that high-dimensional data often lie on a low-dimensional manifold (Tenenbaum et al., 2000; Fefferman et al., 2016; Brahma et al., 2016). The geometric properties of this learned manifold are shaped by the pretraining objectives, for instance, contrastive methods like CLIP (Radford et al., 2021) encourage a topology of well-separated semantic clusters (Tian et al., 2021); self-distillation approaches like DINO (Oquab et al., 2024) preserve fine-grained local structures; while supervised learning often prioritizes linear separability. These manifolds can be quantitatively characterized using spectral analysis of the gram matrix, which yields metrics like the effective rank (Roy & Vetterli, 2007; Vershynin, 2009). To further probe for non-linearity, tangent space analysis directly measures the curvature of the manifold, with significant non-zero curvature being the signature of a non-linear structure (Papaspiliopoulos, 2020). Our work leverages these tools to conduct a rigorous geometric analysis of pathology features, revealing the properties our MR block is designed to preserve.

3 METHODOLOGY

We first formalize the MIL setting and introduce the geometric tools used to probe the feature manifold, and then detail the Manifold Residual block, mitigating manifold degradation while enabling parameter-efficient task-specific adaptation. Finally, we present empirical evidence and theoretical support for our MR block.

Figure 2: Illustration of the linear transformation (a), our proposed MR block (b), and the overall pipeline of MIL for WSI classification task and where our MR block takes effect (c).

3.1 PRELIMINARIES

141

142

143

144

145

146

147

148

149

150 151

152

153154155

156 157

158 159

160

161 162

163

164

165 166 167

168

169

170 171

172173

174175176

177

178

179

180

181

182

183 184

185

187

3.1.1 Multiple Instance Learning Formulation

In MIL, we first partition the WSI into N non-overlapping patches. A pretrained feature extractor then encodes each patch into a feature vector $\mathbf{f}_i \in \mathbb{R}^{d_p}$, where d_p is the output dimension. The set of these patch features, known as the bag $\mathbf{F} = \{\mathbf{f}_i\}_{i=1}^N$, is aggregated by an MIL model to produce a single bag feature $\mathbf{f} = \text{MIL}(\mathbf{F}) \in \mathbb{R}^{d_f}$, where d_f is the dimension of the bag feature. Finally, this bag feature is fed into a classifier to obtain the slide-level logits $\mathbf{p} = \text{CLS}(\mathbf{f}) \in \mathbb{R}^C$ with C being the number of classes. Notably, most modern MIL models employ attention pooling to compute \mathbf{f} , which simultaneously yields both the bag feature and interpretable importance scores for each patch. This framework is illustrated in Fig. 2 (right).

3.1.2 Spectral Analysis of Feature Representations

Spectral analysis is a classic technique used to estimate the intrinsic dimensionality of feature representations. The process begins by computing the Gram matrix $K = F_n F_n^{\top}$ from the normalized feature matrix F_n . An eigendecomposition is then performed on K to obtain the eigenvalues $\{\lambda_i(K)\}_{i=1}^N$, whose plotted distribution forms the spectral curve for visual analysis. For quantitative analysis, we normalize these eigenvalues into a probability distribution $p_i = \lambda_i(K) / \sum_j \lambda_j(K)$ and use it to compute the Von Neumann entropy $S = -\sum_i p_i \log p_i$, and finally, the effective rank $R_{\text{eff}} = \exp(S)$ (Skean et al., 2025) $(N \gg d_p)$.

3.1.3 TANGENT SPACE ANALYSIS

To probe for nonlinearity beyond the low-dimensionality, we measure the local curvature of the manifold via tangent space analysis (Zhang & Zha, 2004). This technique provides evidence of a nonlinear structure (Lim et al., 2024) by quantifying how the local geometry changes across the feature space. We approximate the geodesic paths by constructing a k-nearest-neighbor graph (Tenenbaum et al., 2000) (k=12) over the features F_n with cosine similarity and compute the local tangent space \mathcal{T}_i at each point f_i using local PCA (Kambhatla & Leen, 1997). The drift between the tangent spaces of two points, \mathcal{T}_i and \mathcal{T}_j , is then defined as:

$$\mathcal{D}(\mathcal{T}_i, \mathcal{T}_j) = 1 - \frac{1}{d_s} \left\| \mathbf{V}_i^{\mathsf{T}} \mathbf{V}_j \right\|_F^2, \tag{1}$$

where V_i and V_j are orthogonal bases of the d_s -dimensional tangent spaces, and $\|\cdot\|_F$ is the Frobenius norm.

3.2 Manifold Residual Block

The design of the MR block is an architectural response to the geometric deficiencies of vanilla linear layers. In this section, we first detail its architecture, and then its empirical motivation and theoretical guarantees.

3.2.1 ARCHITECTURE

To mitigate the geometric degradation introduced by a vanilla linear layer, we introduce the MR block, a parameter-efficient, plug-and-play substitute that is explicitly geometry-aware. As illustrated in Fig. 2(c), the MR block reframes a linear map as the sum of two pathways: a fixed geometric anchor for manifold preservation and a trainable LRP for task-specific adaptation.

Given the input $X \in \mathbb{R}^{N \times d_0}$ with the input feature dimension being d_0 , our MR block can be expressed by,

$$f_{MR}(\mathbf{X}) = GELU(\mathbf{X}\mathbf{W}_2)\mathbf{W}_1 + \mathbf{X}\mathbf{B}, \tag{2}$$

with output dimension d_1 . Here GELU(·) is the elementwise GELU activation (Hendrycks & Gimpel, 2016); the LRP uses $W_2 \in \mathbb{R}^{d_0 \times r}$ and $W_1 \in \mathbb{R}^{r \times d_1}$ with rank parameter $r \ll \min(d_0, d_1)$; the anchor uses a fixed random matrix $B \in \mathbb{R}^{d_0 \times d_1}$ initialized with Kaiming uniform (He et al., 2015). For initialization, we set $W_1 = \mathbf{0}$ and draw W_2 from Kaiming uniform, so the residual path initially contributes zero and $f_{\mathrm{MR}}(X) = XB$ at step zero. During training, the LRP is activated only if it improves the objective. Therefore, even though Fig. 1(d) shows that the tangent space drift of the LRP alone can be significant, the anchor balances this effect and the overall MR mapping better preserves manifold geometry.

Beyond preserving geometry, the MR block also reduces the number of trainable parameters. For any rank $r < (d_0d_1)/(d_0+d_1)$, the trainable parameter count of LRP $r(d_0+d_1)$ is strictly fewer than d_0d_1 parameters of a vanilla linear layer, theoretically reducing the risk of overfitting (see Sec. C for details).

3.2.2 EMPIRICAL EVIDENCE

First, we show that features are inherently low dimensional by spectral analysis (Sec. 3.1.2). The eigenvalue spectrum decays rapidly, yielding an effective rank of $29.7 \ll 512$). This provides quantitative evidence that variance is concentrated within a compact, low-dimensional structure.

Second, we provide evidence that this structure exhibits a curved, nonlinear manifold. To distinguish a curved manifold from a linear subspace, we probe nonlinearity via tangent space analysis (Sec. 3.1.3). As shown in Fig. 1(d), geometric drift increases with graph-hop distance. This nonzero curvature indicates a curved, nonlinear manifold geometry that a trained linear layer fails to preserve. This geometric fragility is the central problem the MR block is designed to address.

3.2.3 COMPONENT JUSTIFICATION AND THEORETICAL GUARANTEES

The geometric anchor preserves manifold structure that linear layers distort. The fixed random projection serves as a geometric anchor. Random projection theory shows that this transformation approximately preserves essential geometric properties (see Sec. D). Empirically, Fig. 1 empirically corroborates this: the anchor (c) approximately preserves cluster and neighborhood structure, (d) maintains the intrinsic curvature pattern in tangent space drift, and (a) sharpens the eigenvalue spectrum, improving spectral concentration.

The LRP enables task-specific adaptation. Although our LRP is similar to LoRA (Hu et al., 2022), they serve different purposes. The LRP is motivated by the low effective rank of features. Its low-rank bottleneck imposes a structured inductive bias, ensuring that task-specific adaptation remains consistent with the low-rank nature of the manifold, as evidenced in Fig. 1(a). This reduces learning to a simpler residual-fitting task.

Approximation behavior of the MR block. Our MR block can approximate any linear layer to arbitrary precision, thereby providing a worst-case performance guarantee: if the MR block cannot improve over the

Table 1: Experimental results comparing our method with baselines are presented under various shot settings. Performance metrics in which our method outperforms the baseline are highlighted in italic blue and the best metrics are in bold red. "#P" is the number of parameters (million).

| k | Methods | #P | AUC↑ | Camelyon16 F1↑ | Acc.↑ | AUC↑ | NSCLC F1↑ | Acc.↑ | AUC↑ | RCC F1↑ | Acc.↑ |
|----|---------------------|--------------|-------------------------------|---------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|------------------------|
| _ | **** | 40.5 | ' | · ' | ' | · ' | | · ' | · · | ' | |
| | ViLaMIL | 40.7 | | 0.484 _{0.091} | | | | | | | |
| | ABMIL | 0.26 | | $0.492_{0.072}$ | | | | | | | |
| | MR-ABMIL | 0.10 | | 0.438 _{0.100} | | | | | | | |
| | TransMIL | 2.41 | | 0.467 _{0.075} | | | | | | | |
| 4 | MR-TransMIL | 2.01 | | 0.509 _{0.063} | | | | | | | |
| 4 | CATE MR-CATE | 1.93 | | $0.500_{0.098}$ | | 0.00- | 0.0 | 0.000 | 0.020 | 0.0-0 | 0.0-0 |
| | | 0.84 | | 0.584 _{0.158} | | | | | | | |
| | DGRMIL MR-DGRMIL | 4.08 3.29 | | 0.594 _{0.220} | | | | | | | |
| | RRTMIL | 2.44 | | 0.557 _{0.184} | | | | | | | |
| | MR-RRTMIL | 2.04 | | $0.548_{0.052}$ $0.427_{0.036}$ | | | | | | | |
| | WIK-KKI WIIL | 2.04 | 0.4030.038 | 0.4270.036 | 0.0290.006 | 0.0990.063 | 0.0390.050 | 0.0400.050 | 0.9070.014 | 0.8230.027 | 0.8010.022 |
| | ViLaMIL | 40.7 | $0.740_{0.172}$ | $0.659_{0.221}$ | $0.718_{0.165}$ | $0.884_{0.052}$ | $0.792_{0.057}$ | $0.792_{0.057}$ | $0.968_{0.009}$ | $0.862_{0.021}$ | $0.880_{0.017}$ |
| | ABMIL | 0.26 | $0.492_{0.056}$ | $0.489_{0.030}$ | $0.597_{0.048}$ | $0.877_{0.059}$ | $0.810_{0.054}$ | $0.811_{0.054}$ | $0.918_{0.063}$ | $0.766_{0.084}$ | $0.791_{0.079}$ |
| | MR-ABMIL | 0.10 | $0.565_{0.155}$ | $0.538_{0.139}$ | $0.619_{0.087}$ | $0.940_{0.033}$ | $0.859_{0.033}$ | $0.860_{0.033}$ | $0.981_{0.008}$ | $0.876_{0.029}$ | $0.896_{0.021}$ |
| | TransMIL | 2.41 | $0.575_{0.025}$ | $0.520_{0.089}$ | $0.536_{0.074}$ | $0.879_{0.061}$ | $0.803_{0.052}$ | $0.804_{0.052}$ | $0.979_{0.006}$ | $0.865_{0.039}$ | $0.884_{0.031}$ |
| | MR-TransMIL | 2.01 | | <i>0.542</i> _{0.093} | | | | | | | |
| 8 | CATE | 1.93 | | $0.691_{0.200}$ | | | | | | | |
| | MR-CATE | 0.84 | | $0.817_{0.160}$ | | | | | | | |
| | DGRMIL | 4.08 | | $0.725_{0.178}$ | | | | | | | |
| | MR-DGRMIL | 3.29 | | $0.860_{0.035}$ | | | | | | | |
| | RRTMIL | 2.44 | | $0.626_{0.154}$ | | | | | | | |
| | MR-RRTMIL | 2.04 | $0.853_{0.163}$ | $0.792_{0.166}$ | 0.817 _{0.133} | $0.873_{0.078}$ | $0.796_{0.079}$ | $0.798_{0.078}$ | 0.978 _{0.004} | $0.863_{0.015}$ | 0.878 _{0.014} |
| | ViLaMIL | 40.7 | 0.814 _{0.231} | $0.775_{0.220}$ | $0.820_{0.150}$ | $0.928_{0.029}$ | $0.849_{0.027}$ | $0.850_{0.027}$ | 0.974 _{0.011} | $0.855_{0.030}$ | $0.879_{0.022}$ |
| | ABMIL | 0.26 | $0.732_{0.171}$ | $0.712_{0.154}$ | $0.757_{0.102}$ | $0.936_{0.024}$ | $0.878_{0.034}$ | $0.878_{0.034}$ | $0.947_{0.052}$ | $0.810_{0.078}$ | $0.828_{0.085}$ |
| | MR-ABMIL | 0.10 | $0.790_{0.231}$ | $0.745_{0.224}$ | $0.780_{0.171}$ | 0.965 _{0.012} | 0.891 _{0.010} | 0.891 _{0.010} | 0.9830.007 | 0.8910.028 | $0.911_{0.020}$ |
| | TransMIL | 2.41 | $0.697_{0.157}$ | $0.625_{0.192}$ | $0.693_{0.133}$ | $0.956_{0.017}$ | $0.882_{0.036}$ | $0.882_{0.036}$ | $0.984_{0.003}$ | $0.893_{0.016}$ | $0.909_{0.007}$ |
| | MR-TransMIL | 2.01 | 0.7930.174 | $0.751_{0.148}$ | $0.771_{0.132}$ | $0.932_{0.044}$ | $0.840_{0.061}$ | $0.841_{0.060}$ | $0.986_{0.003}$ | $0.893_{0.016}$ | 0.912 _{0.010} |
| 16 | CATE | 1.93 | $0.922_{0.031}$ | $0.875_{0.032}$ | $0.885_{0.028}$ | $0.942_{0.021}$ | $0.851_{0.020}$ | $0.852_{0.020}$ | $0.979_{0.006}$ | $0.874_{0.027}$ | $0.893_{0.020}$ |
| | MR-CATE | 0.84 | $0.924_{0.032}$ | $0.867_{0.039}$ | $0.878_{0.034}$ | $0.96I_{0.014}$ | $0.886_{0.033}$ | $0.887_{0.033}$ | 0.986 _{0.006} | 0.898 _{0.031} | $0.916_{0.023}$ |
| | DGRMIL | 4.08 | $0.885_{0.050}$ | | | | | | | $0.863_{0.010}$ | |
| | MR-DGRMIL | 3.29 | $0.923_{0.058}$ | $0.876_{0.058}$ | $0.887_{0.052}$ | $0.939_{0.017}$ | $0.866_{0.027}$ | $0.867_{0.027}$ | 0.977 _{0.008} | $0.880_{0.020}$ | $0.903_{0.015}$ |
| | RRTMIL | 2.44 | | $0.748_{0.129}$ | | | | | | | |
| | MR-RRTMIL | 2.04 | 0.937 _{0.016} | $0.890_{0.024}$ | 0.898 _{0.023} | $0.910_{0.054}$ | $0.825_{0.076}$ | $0.828_{0.071}$ | $0.978_{0.008}$ | $0.865_{0.041}$ | $0.889_{0.029}$ |

linear layer then it can match its behavior arbitrarily closely, ensuring no loss in worst-case performance. In addition, rather than directly proving our formulation, we establish a more fundamental result without activation functions. Subsequently, with the universal approximation theorem (Barron, 1993; Cybenko, 1989; Funahashi, 1989; Hornik et al., 1989; Lu & Lu, 2020), our architecture with non-linear activation function can achieve even superior approximation precision. The proof is given in Sec. E.

4 EXPERIMENTS AND RESULTS

4.1 Dataset Descriptions and Implementation Details

Datasets. We conduct extensive experiments on Camelyon16 (Litjens et al., 2018), TCGA-NSCLC and TCGA-RCC datasets. The details of the datasets used in this work are available in Sec. G.1.

Evaluation Metrics. We employ three complementary metrics for a comprehensive assessment: Area Under the Receiver Operating Characteristic Curve (AUC), F1-score, and Accuracy

Table 2: Ablation studies of our proposed method across varying shot settings on the Camelyon16, TCGA-NSCLC, and TCGA-RCC datasets. Metrics highlighted in bold red indicate the top result.

| | M-41-1- | #D | (| Camelyon16 | | | NSCLC | | RCC | | | |
|----|------------------------------|------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|------------------------|------------------------|-------------------------------|--|
| k | Methods | #P | AUC↑ | F1↑ | Acc.↑ | AUC↑ | F1↑ | Acc.↑ | AUC↑ | F1↑ | Acc.↑ | |
| | Ours | 2.01 | 0.714 _{0.149} | 0.655 _{0.121} | 0.699 _{0.073} | 0.9510.035 | 0.8610.047 | 0.8620.046 | 0.968 _{0.013} | 0.8430.054 | 0.870 _{0.044} | |
| | $+m{B}\mathrm{FT}$ | 2.54 | $0.553_{0.011}$ | $0.558_{0.027}$ | $0.606_{0.023}$ | $0.887_{0.021}$ | $0.817_{0.033}$ | $0.818_{0.033}$ | $0.964_{0.019}$ | $0.843_{0.036}$ | $0.864_{0.032}$ | |
| 8 | - $oldsymbol{B}$ | 2.01 | $0.553_{0.006}$ | $0.557_{0.016}$ | $0.584_{0.019}$ | $0.882_{0.052}$ | $0.810_{0.069}$ | $0.812_{0.066}$ | $0.976_{0.011}$ | $0.864_{0.050}$ | $0.886_{0.044}$ | |
| | - $oldsymbol{B}oldsymbol{x}$ | 2.01 | $0.563_{0.025}$ | $0.558_{0.029}$ | $0.591_{0.027}$ | $0.927_{0.027}$ | $0.847_{0.034}$ | $0.849_{0.034}$ | $0.981_{0.004}$ | $0.874_{0.037}$ | 0.893 _{0.030} | |
| | - LRP | 1.89 | $0.510_{0.177}$ | $0.494_{0.142}$ | $0.589_{0.088}$ | $0.894_{0.063}$ | $0.816_{0.058}$ | $0.817_{\scriptstyle 0.058}$ | $0.981_{0.008}$ | $0.869_{0.028}$ | $0.893_{0.021}$ | |
| | Ours | 2.01 | 0.864 _{0.117} | 0.8090.120 | 0.8190.118 | 0.960 _{0.017} | 0.887 _{0.014} | 0.887 _{0.014} | 0.9830.008 | 0.877 _{0.023} | $0.900_{0.020}$ | |
| | + $m{B}$ FT | 2.54 | $0.513_{0.094}$ | $0.472_{0.073}$ | $0.611_{0.045}$ | $0.931_{0.024}$ | $0.872_{0.012}$ | $0.872_{0.012}$ | $0.974_{0.011}$ | $0.854_{0.028}$ | $0.879_{0.024}$ | |
| 16 | - $oldsymbol{B}$ | 2.01 | $0.619_{0.136}$ | $0.528_{0.083}$ | $0.633_{0.074}$ | $0.919_{0.024}$ | $0.847_{0.024}$ | $0.848_{0.024}$ | $0.977_{0.013}$ | $0.860_{0.041}$ | $0.883_{0.036}$ | |
| | - $oldsymbol{B}oldsymbol{x}$ | 2.01 | $0.722_{0.198}$ | $0.676_{0.199}$ | $0.735_{0.130}$ | $0.955_{0.007}$ | $0.874_{0.010}$ | $0.874_{0.010}$ | $0.985_{0.004}$ | $0.884_{0.021}$ | $0.906_{0.016}$ | |
| | - LRP | 1.89 | $0.683_{0.202}$ | $0.651_{0.178}$ | $0.698_{0.128}$ | $0.937_{0.020}$ | $0.864_{\scriptstyle 0.011}$ | $0.864_{0.010}$ | $0.983_{0.006}$ | $0.870_{0.013}$ | $0.896_{0.010}$ | |

Table 3: Performance of MR-ABMIL with replaced linear layers under varying shot settings on Camelyon16, TCGA-NSCLC, and TCGA-RCC datasets. "U" and "V" denote two matrices in MR-ABMIL.

| \overline{k} | Methods | #P | Camelyon16 AUC↑ F1↑ | | Acc.↑ | $\begin{array}{cccc} & & & NSCLC \\ Acc. \uparrow & AUC \uparrow & F1 \uparrow & Acc. \uparrow \end{array}$ | | | | RCC F1↑ | | | |
|----------------|-----------------|------|---|-----------------|-----------------|---|-----------------|-------------------------------|-----------------|-----------------|-----------------|--|--|
| 4 | V + U V U | 0.18 | 0.461 _{0.137} 0.456 _{0.003} 0.528 _{0.105} | $0.439_{0.002}$ | $0.598_{0.014}$ | $0.861_{0.035}$ | $0.782_{0.032}$ | $0.784_{0.031}$ | $0.940_{0.048}$ | $0.790_{0.075}$ | $0.817_{0.073}$ | | |
| 8 | V + U V U | 0.18 | 0.565 _{0.155} 0.714 _{0.149} 0.728 _{0.191} | $0.655_{0.121}$ | $0.699_{0.073}$ | $0.951_{0.035}$ | $0.861_{0.047}$ | 0.862 _{0.046} | $0.968_{0.013}$ | $0.843_{0.054}$ | $0.870_{0.044}$ | | |

Implementation Details. To ensure a fair and robust comparison, we use a consistent set of hyper-parameters and training protocols for all methods to isolate the architectural impact of our MR block, detailed in Sec. G.2.

4.2 COMPARISON WITH STATE-OF-THE-ART METHODS

Table 1 presents the results of few-shot (k=4,8,16) setting on Camelyon16, TCGA-NSCLC, and TCGA-RCC. Three significant findings emerge from this analysis. First, models augmented with the MR block consistently outperform their corresponding baselines across datasets and shot counts, and they match or surpass the current state of the art ViLaMIL (Shi et al., 2024) while using far fewer trainable parameters. Second, our MR block improves parameter efficiency: replacing the vanilla linear layer with our MR block yields smaller models that deliver better performance across multiple MIL backbones, indicating that MR provides a beneficial low-rank inductive bias rather than simply adding capacity. Third, performance increases steadily as k grows for all methods, with MR variants showing the largest gains at moderate shots and remaining competitive at higher shots, where RCC results approach a ceiling. We also observe reduced variability across runs for several MR settings, suggesting improved training stability. Overall, MR is a plug-in that enhances accuracy, data efficiency, and robustness without increasing model size.

4.3 ABLATION STUDIES AND SENSITIVITY ANALYSIS

Ablation Study of MR Block Components. Our ablation studies in Table 2 validate our decoupled design. Removing the LRP degrades performance, confirming its necessity for adaptation. Removing the geometric anchor (-B) while retaining the residual connection is also detrimental, underscoring its dual role as a

Table 4: Experimental results on different initialization methods in MR-CATE. "K." and "X." stands for Kaiming and Xavier initialization. "U." and "N." stands for uniform and normal distribution. The number following is the input parameter for initialization. The gray rows stand for our main setting.

| 1- | Mathada | (| Camelyon16 | | | NSCLC | | RCC | | | |
|----|--|---|---|----------------------------------|--|--|---|--|--|---|--|
| k | Methods | AUC↑ | F1↑ | Acc.↑ | AUC↑ | F1↑ | Acc.↑ | AUC↑ | F1 ↑ | Acc.↑ | |
| | K. N. 0 K. U. 0 | 0 | | | | 0.00. | 0.752 _{0.055} 0.794 _{0.031} | 0.0 | 0.000 | 0.0-0 | |
| 4 | K. U. $\sqrt{5}$ X. N. 0 X. U. 0 | 0.621 _{0.169} 0.662 _{0.141} | 0.584 _{0.158} 0.488 _{0.201} | $0.623_{0.135} \\ 0.647_{0.129}$ | 0.870 _{0.039} 0.840 _{0.053} | 0.794 _{0.032} 0.763 _{0.042} | $0.794_{0.032} \\ 0.764_{0.042} \\ 0.791_{0.026}$ | 0.982 _{0.012} 0.981 _{0.014} | 0.887 _{0.020} 0.879 _{0.034} | 0.905 _{0.018} 0.901 _{0.029} | |
| | K. N. 0 K. U. 0 | | | | | | $0.870_{0.028} \\ 0.864_{0.021}$ | | | | |
| 16 | K. U. √5 X. N. 0 X. U. 0 | 0.945 _{0.015} | $0.901_{0.015}$ | 0.907 _{0.016} | $0.956_{0.020}$ | $0.885_{0.014}$ | 0.887 _{0.033} 0.885 _{0.014} 0.888 _{0.031} | 0.987 _{0.005} | $0.889_{0.022}$ | $0.909_{0.014}$ | |

Table 5: Experimental results on different pretrained feature extractors under various shot settings. "R50" stands for ResNet50, and "MR-" stands for MIL model with our MR block.

| | | | | ABN | 1IL | | | CATE | | | | | |
|-----|---------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|
| k | Methods | | NSCLC | | | RCC | | | NSCLC | | | RCC | |
| | | AUC↑ | F1↑ | Acc.↑ |
| | R50 | $0.590_{0.037}$ | 0.547 _{0.081} | $0.570_{0.046}$ | 0.7460.069 | 0.5660.090 | 0.5850.107 | 0.6490.078 | 0.574 _{0.106} | $0.600_{0.069}$ | 0.8260.110 | 0.6420.102 | $0.689_{0.070}$ |
| Q | MR-R50 | 0.6430.042 | 0.5970.033 | $0.601_{0.032}$ | $0.808_{0.045}$ | 0.6230.067 | 0.6390.083 | $0.635_{0.064}$ | $0.607_{0.057}$ | $0.614_{0.050}$ | 0.8950.036 | $0.750_{0.072}$ | $0.766_{0.060}$ |
| 0 | UNI | $0.808_{0.079}$ | $0.728_{0.062}$ | $0.731_{0.062}$ | $0.986_{0.004}$ | $0.887_{0.021}$ | $0.903_{0.018}$ | $0.827_{0.078}$ | $0.724_{0.051}$ | $0.730_{0.051}$ | $0.981_{0.006}$ | $0.878_{0.034}$ | $0.893_{0.030}$ |
| | MR-UNI | 0.913 _{0.050} | 0.817 _{0.044} | 0.818 _{0.043} | 0.986 _{0.006} | 0.912 _{0.018} | 0.923 _{0.014} | 0.888 _{0.066} | 0.806 _{0.068} | $0.808_{0.067}$ | 0.985 _{0.004} | 0.908 _{0.014} | $0.918_{0.011}$ |
| | R50 | 0.656 _{0.099} | 0.6230.090 | 0.6250.088 | 0.840 _{0.040} | 0.677 _{0.038} | 0.7090.051 | 0.6630.101 | 0.5620.128 | 0.6110.074 | $0.900_{0.080}$ | 0.747 _{0.132} | $0.780_{0.090}$ |
| 16 | MR-R50 | | | | | | | $0.663_{0.126}$ | | | | | |
| -10 | UNI | $0.911_{0.027}$ | $0.836_{0.030}$ | $0.836_{0.030}$ | $0.989_{0.003}$ | $0.892_{0.019}$ | $0.910_{0.018}$ | $0.872_{0.075}$ | $0.758_{0.092}$ | $0.769_{0.077}$ | $0.980_{0.009}$ | $0.880_{0.025}$ | $0.903_{0.016}$ |
| | MR-UNI | 0.948 _{0.028} | 0.870 _{0.036} | 0.870 _{0.036} | 0.989 _{0.004} | 0.911 _{0.024} | 0.924 _{0.018} | 0.917 _{0.046} | 0.826 _{0.060} | 0.829 _{0.056} | 0.985 _{0.006} | 0.910 _{0.019} | 0.922 _{0.015} |

geometric anchor and a spectral shaper. More importantly, making the anchor trainable causes a catastrophic performance collapse. This provides direct empirical proof of our core hypothesis: an unconstrained linear layer shatters the feature manifold, while our design preserves it.

Layer-wise Performance Analysis. We conducted a layer-wise ablation study to evaluate how our method affects performance when implemented at different layers within MIL. We selected ABMIL (Ilse et al., 2018) due to its structural simplicity, with results reported in Tables 3 and 7. This study investigates two components of ABMIL (Ilse et al., 2018), U and V (detailed in Sec. G.4), analyzing performances when selectively replacing these components with the MR block. The results demonstrate that optimal performance is achieved when both components are replaced with our MR block, and our method is robust to positions in MIL.

Choice of Initialization. We investigated the sensitivity of the MR block to the initialization scheme of B, by testing four schemes: Kaiming/Xavier Normal/Uniform. The results in Tables 4 and 8 demonstrate that our method is broadly insensitive to the specific choice of initialization. Across most settings, the performance differences between these methods are marginal, confirming that the benefits of the geometric anchor are not sensitive to initialization. Our chosen default, Kaiming Uniform with $\sqrt{5}$, consistently delivers superior performance, validating it as an effective choice without the need for additional hyper-parameter tuning.

Foundation Model Selection. We investigated whether our approach demonstrates compatibility across different embeddings. Using CONCH, UNI, and ResNet50 with ABMIL and CATE, we present results in Tables 5 and 9. Our method demonstrated consistent improvements over baselines across all model

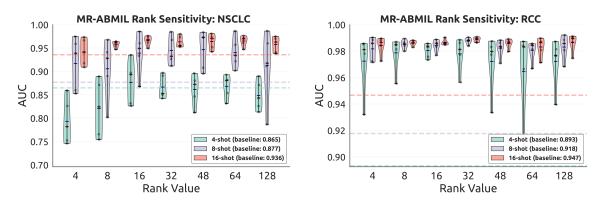


Figure 3: Sensitivity analysis of r using MR-ABMIL. Dashed lines indicate the performance of the ABMIL.

combinations, substantiating its robustness across diverse MIL and foundation models. CONCH embeddings consistently outperformed UNI and ResNet50, aligning with expectations as ResNet50 lacks pathology pretraining. The dimensional difference between UNI (1,024) and CONCH (512) likely contributes to lower performance of UNI due to inflated higher-dimensional spaces under few-shot settings.

Sensitivity Analysis of the Rank. Our sensitivity analysis on the residual rank r (Figs. 3 and 6) shows that model performance saturates around r=32. This empirically-found saturation point aligns remarkably with the theoretically-predicted effective rank of the features, confirming that a low-rank path is sufficient to capture the essential manifold information. While the performance of the simpler MR-ABMIL model peaks precisely at r=32, the more complex MR-CATE exhibits marginal gains at higher ranks, which we attribute to its capacity to model finer-grained feature interactions beyond the primary manifold structure.

4.4 MODEL INTERPRETABILITY ANALYSIS UNDER EXTREME LIMITED RESOURCE

Figure 7 illustrates the heatmaps generated by our MR-CATE, which demonstrates remarkable robustness in the extreme 2-shot setting where standard models often fail to produce meaningful heatmaps. In the original images (upper left), blue curves delineate the approximate tumor boundaries, and the corresponding heatmap is shown in the lower left panel. Notably, our method captures additional fine-grained boundaries not indicated in the original annotations, as shown in the right panel. Although all regions on the right fall within the blue-delineated tumor boundaries, our model precisely identifies the boundaries between distinct morphological patterns, demonstrating its sensitivity in capturing tumor heterogeneity.

5 CONCLUSION

In this work, we introduced a new geometric perspective on the problem of overfitting in few-shot WSI classification. We provided both quantitative and visual evidence that features from pathology foundation models exhibit a fragile, low-dimensional manifold geometry, and identified a common model failure mode: the systematic degradation of this manifold structure by the geometry-agnostic linear layers, which are central to modern MIL models. The proposed Manifold Residual block, a plug-and-play module preserves this crucial geometry via a fixed random geometric anchor while enabling parameter-efficient, task-specific adaptation through a low-rank residual path. Our extensive experimental results not only achieve state-of-the-art results but also empirically support our geometric diagnosis, thereby establishing a new, structure-aware paradigm for developing robust models with implications beyond computational pathology.

REPRODUCIBILITY STATEMENT

Code will be released upon acceptance of this work.

In addition, we point to all components needed to replicate our results. The MIL setup and geometry tools are in Secs. 3.1, 3.1.2 and 3.1.3. The MR architecture is in Sec. 3.2.1; full experimental protocols (datasets, few-shot k, backbones, metrics) and implementation details are in the Sec. G. Ablations and sensitivity studies (component/placement/initialization/rank) appear in Tables 2 to 4 and Fig. 3 and their appendix counterparts. Assumptions and complete proofs are in Secs. D and E.

REFERENCES

- Richard G Baraniuk and Michael B Wakin. Random projections of smooth manifolds. *Foundations of computational mathematics*, 9(1):51–77, 2009.
- Andrew R Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3):930–945, 1993.
- Pratik Prabhanjan Brahma, Dapeng Wu, and Yiyuan She. Why deep learning works: A manifold disentanglement perspective. *IEEE Transactions on Neural Networks and Learning Systems*, 27(10):1997–2008, 2016. doi: 10.1109/TNNLS.2015.2496947.
- Gabriele Campanella, Matthew G Hanna, Luke Geneslaw, Allen Miraflor, Vitor Werneck Krauss Silva, Klaus J Busam, Edi Brogi, Victor E Reuter, David S Klimstra, and Thomas J Fuchs. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature medicine*, 25(8):1301–1309, 2019.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9650–9660, 2021.
- Richard J Chen, Tong Ding, Ming Y Lu, Drew FK Williamson, Guillaume Jaume, Bowen Chen, Andrew Zhang, Daniel Shao, Andrew H Song, Muhammad Shaban, et al. Towards a general-purpose foundation model for computational pathology. *Nature Medicine*, 2024.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PmLR, 2020.
- Philip Chikontwe, Meejeong Kim, Soo Jeong Nam, Heounjeong Go, and Sang Hyun Park. Multiple instance learning with center embeddings for histopathology classification. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part V 23*, pp. 519–528. Springer, 2020.
- George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- Sanjoy Dasgupta and Anupam Gupta. An elementary proof of a theorem of johnson and lindenstrauss. *Random Structures & Algorithms*, 22(1):60–65, 2003.
- Charles Fefferman, Sanjoy Mitter, and Hariharan Narayanan. Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 29(4):983–1049, 2016.

Kexue Fu, Linhao Qu, Shuo Wang, Ying Xiong, Ilias Maglogiannis, Longxiang Gao, Manning Wang, et al. Fast: A dual-tier few-shot learning paradigm for whole slide image classification. *Advances in Neural Information Processing Systems*, 37:105090–105113, 2024.

Ken-Ichi Funahashi. On the approximate realization of continuous mappings by neural networks. *Neural networks*, 2(3):183–192, 1989.

Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11): 665–673, 2020.

Gabriel Goh, Nick Cammarata, Chelsea Voss, Shan Carter, Michael Petrov, Ludwig Schubert, Alec Radford, and Chris Olah. Multimodal neurons in artificial neural networks. *Distill*, 6(3):e30, 2021.

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. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 189–199. Springer,

Zhengrui Guo, Qichen Sun, Jiabo Ma, Lishuang Feng, Jinzhuo Wang, and Hao Chen. Context matters: Query-aware dynamic long sequence modeling of gigapixel images. *arXiv preprint arXiv:2501.18984*, 2025a.

Zhengrui Guo, Conghao Xiong, Jiabo Ma, Qichen Sun, Lishuang Feng, Jinzhuo Wang, and Hao Chen. Focus: Knowledge-enhanced adaptive visual compression for few-shot whole slide image classification. In

 Proceedings of the Computer Vision and Pattern Recognition Conference, pp. 15590–15600, 2025b.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034, 2015.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.

Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.

Le Hou, Dimitris Samaras, Tahsin M Kurc, Yi Gao, James E Davis, and Joel H Saltz. Patch-based convolutional neural network for whole slide tissue image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2424–2433, 2016.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *International Conference on Learning Representations*, 1(2):3, 2022.

Yanyan Huang, Weiqin Zhao, Yihang Chen, Yu Fu, and Lequan Yu. Free lunch in pathology foundation model: Task-specific model adaptation with concept-guided feature enhancement. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

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.

- Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *International conference on machine learning*, pp. 2127–2136. PMLR, 2018.
 - William B Johnson, Joram Lindenstrauss, et al. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.
 - Nandakishore Kambhatla and Todd K. Leen. Dimension reduction by local principal component analysis. *Neural Computation*, 9(7):1493–1516, 1997. doi: 10.1162/neco.1997.9.7.1493.
 - Fahdi Kanavati, Gouji Toyokawa, Seiya Momosaki, Michael Rambeau, Yuka Kozuma, Fumihiro Shoji, Koji Yamazaki, Sadanori Takeo, Osamu Iizuka, and Masayuki Tsuneki. Weakly-supervised learning for lung carcinoma classification using deep learning. *Scientific reports*, 10(1):9297, 2020.
 - Hao Li, Ying Chen, Yifei Chen, Rongshan Yu, Wenxian Yang, Liansheng Wang, Bowen Ding, and Yuchen Han. Generalizable whole slide image classification with fine-grained visual-semantic interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11398–11407, 2024a.
 - Jiawen Li, Yuxuan Chen, Hongbo Chu, Qiehe Sun, Tian Guan, Anjia Han, and Yonghong He. Dynamic graph representation with knowledge-aware attention for histopathology whole slide image analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11323–11332, 2024b.
 - Victor Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Y Zou. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. *Advances in Neural Information Processing Systems*, 35:17612–17625, 2022.
 - Uzu Lim, Harald Oberhauser, and Vidit Nanda. Tangent space and dimension estimation with the wasserstein distance. *SIAM Journal on Applied Algebra and Geometry*, 8(3):650–685, 2024.
 - Tiancheng Lin, Zhimiao Yu, Hongyu Hu, Yi Xu, and Chang-Wen Chen. Interventional bag multi-instance learning on whole-slide pathological images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19830–19839, 2023.
 - Geert Litjens, Peter Bandi, Babak Ehteshami Bejnordi, Oscar Geessink, Maschenka Balkenhol, Peter Bult, Altuna Halilovic, Meyke Hermsen, Rob van de Loo, Rob Vogels, Quirine F Manson, Nikolas Stathonikos, Alexi Baidoshvili, Paul van Diest, Carla Wauters, Marcory van Dijk, and Jeroen van der Laak. 1399 H&Estained sentinel lymph node sections of breast cancer patients: the CAMELYON dataset. *GigaScience*, 7 (6), June 2018. ISSN 2047-217X. doi: 10.1093/gigascience/giy065. URL https://academic.oup.com/gigascience/article/doi/10.1093/gigascience/giy065/5026175.
 - Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.
 - 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.
 - 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*, 30:863–874, 2024.
 - Yulong Lu and Jianfeng Lu. A universal approximation theorem of deep neural networks for expressing probability distributions. *Advances in neural information processing systems*, 33:3094–3105, 2020.

Jiabo Ma, Zhengrui Guo, Fengtao Zhou, Yihui Wang, Yingxue Xu, Yu Cai, Zhengjie Zhu, Cheng Jin, Yi Lin, Xinrui Jiang, et al. Towards a generalizable pathology foundation model via unified knowledge distillation. *arXiv preprint arXiv:2407.18449*, 2024.

- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=a68SUt6zFt. Featured Certification.
- Omiros Papaspiliopoulos. High-dimensional probability: An introduction with applications in data science, 2020.
- Linhao Qu, Kexue Fu, Manning Wang, Zhijian Song, et al. The rise of ai language pathologists: Exploring two-level prompt learning for few-shot weakly-supervised whole slide image classification. *Advances in Neural Information Processing Systems*, 36:67551–67564, 2023.
- Linhao Qu, Dingkang Yang, Dan Huang, Qinhao Guo, Rongkui Luo, Shaoting Zhang, and Xiaosong Wang. Pathology-knowledge enhanced multi-instance prompt learning for few-shot whole slide image classification. In *European Conference on Computer Vision*, pp. 196–212. Springer, 2024.
- 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*, pp. 8748–8763. PmLR, 2021.
- Maithra Raghu, Chiyuan Zhang, Jon Kleinberg, and Samy Bengio. Transfusion: Understanding transfer learning for medical imaging. *Advances in neural information processing systems*, 32, 2019.
- Olivier Roy and Martin Vetterli. The effective rank: A measure of effective dimensionality. In 2007 15th European signal processing conference, pp. 606–610. IEEE, 2007.
- 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.
- Jiangbo Shi, Chen Li, Tieliang Gong, Yefeng Zheng, and Huazhu Fu. Vila-mil: Dual-scale vision-language multiple instance learning for whole slide image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11248–11258, 2024.
- Oscar Skean, Md Rifat Arefin, Dan Zhao, Niket Nikul Patel, Jalal Naghiyev, Yann LeCun, and Ravid Shwartz-Ziv. Layer by layer: Uncovering hidden representations in language models. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=WGXb7UdvTX.
- Wenhao Tang, Fengtao Zhou, Sheng Huang, Xiang Zhu, Yi Zhang, and Bo Liu. Feature re-embedding: Towards foundation model-level performance in computational pathology. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11343–11352, 2024.
- Joshua B Tenenbaum, Vin de Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.
- Yuandong Tian, Xinlei Chen, and Surya Ganguli. Understanding self-supervised learning dynamics without contrastive pairs. In *International Conference on Machine Learning*, pp. 10268–10278. PMLR, 2021.

Roman Vershynin. High-dimensional probability, 2009.

- Eugene Vorontsov, Alican Bozkurt, Adam Casson, George Shaikovski, Michal Zelechowski, Kristen Severson, Eric Zimmermann, James Hall, Neil Tenenholtz, Nicolo Fusi, et al. A foundation model for clinical-grade computational pathology and rare cancers detection. *Nature medicine*, 30(10):2924–2935, 2024.
- Jinxi Xiang, Xiyue Wang, Xiaoming Zhang, Yinghua Xi, Feyisope Eweje, Yijiang Chen, Yuchen Li, Colin Bergstrom, Matthew Gopaulchan, Ted Kim, et al. A vision–language foundation model for precision oncology. *Nature*, pp. 1–10, 2025.
- Conghao Xiong, Hao Chen, Joseph J.Y. Sung, and Irwin King. Diagnose like a pathologist: Transformer-enabled hierarchical attention-guided multiple instance learning for whole slide image classification. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, pp. 1587–1595, 8 2023. doi: 10.24963/ijcai.2023/176. URL https://doi.org/10.24963/ijcai.2023/176. Main Track.
- Conghao Xiong, Hao Chen, Hao Zheng, Dong Wei, Yefeng Zheng, Joseph JY Sung, and Irwin King. Mome: Mixture of multimodal experts for cancer survival prediction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 318–328. Springer, 2024a.
- Conghao Xiong, Yi Lin, Hao Chen, Hao Zheng, Dong Wei, Yefeng Zheng, Joseph JY Sung, and Irwin King. Takt: Target-aware knowledge transfer for whole slide image classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 503–513. Springer, 2024b.
- Conghao Xiong, Hao Chen, and Joseph JY Sung. A survey of pathology foundation model: Progress and future directions. *arXiv* preprint *arXiv*:2504.04045, 2025.
- Yingxue Xu, Yihui Wang, Fengtao Zhou, Jiabo Ma, Shu Yang, Huangjing Lin, Xin Wang, Jiguang Wang, Li Liang, Anjia Han, et al. A multimodal knowledge-enhanced whole-slide pathology foundation model. arXiv preprint arXiv:2407.15362, 2024.
- Hongrun Zhang, Yanda Meng, Yitian Zhao, Yihong Qiao, Xiaoyun Yang, Sarah E Coupland, and Yalin Zheng. Dtfd-mil: Double-tier feature distillation multiple instance learning for histopathology whole slide image classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 18802–18812, 2022.
- Zhenyue Zhang and Hongyuan Zha. Principal manifolds and nonlinear dimensionality reduction via tangent space alignment. *SIAM journal on scientific computing*, 26(1):313–338, 2004.
- Wenhui Zhu, Xiwen Chen, Peijie Qiu, Aristeidis Sotiras, Abolfazl Razi, and Yalin Wang. Dgr-mil: Exploring diverse global representation in multiple instance learning for whole slide image classification. In *European conference on computer vision*, pp. 333–351. Springer, 2024.
- Barret Zoph, Golnaz Ghiasi, Tsung-Yi Lin, Yin Cui, Hanxiao Liu, Ekin Dogus Cubuk, and Quoc Le. Rethinking pre-training and self-training. *Advances in neural information processing systems*, 33:3833–3845, 2020.

Appendix

LLM Usage A В Spectral Analysis Results C Complexity Analysis D Proof of Random Recalibration Matrix Preserving Input Characteristics Е Proof of Approximation Capabilities of Manifold Residual Block F Additional Proofs G **Experiment Setups** Additional Experimental Results \mathbf{H} I Limitations and Future Works

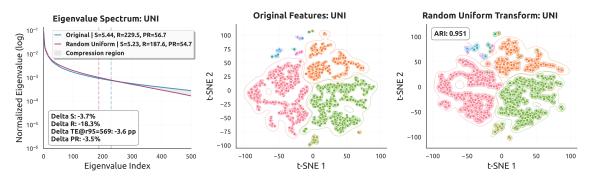


Figure 4: Random transformation preserves the intrinsic low-dimensional manifold of UNI features. (a) Spectral analysis confirms the low-rank structure. (b) t-SNE reveals the cluster topology of the manifold. (c) High ARI and consistent coloring visually and quantitatively prove this topology is robustly preserved.

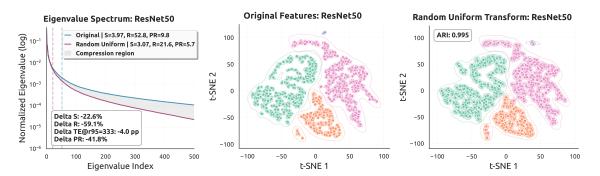


Figure 5: Random transformation preserves the intrinsic low-dimensional manifold of ResNet50 features. (a) Spectral analysis confirms the low-rank structure. (b) t-SNE reveals the cluster topology of the manifold. (c) High ARI and consistent coloring visually and quantitatively prove this topology is robustly preserved.

A LLM USAGE

The authors utilized LLMs as a writing assistant during the preparation of this manuscript. The role of it was primarily focused on tasks related to language refinement and improving the clarity of the narrative. Its application included: (1) enhancing the conciseness and academic tone of sentences and paragraphs; (2) restructuring complex sentences for better readability; and (3) brainstorming alternative phrasings for key scientific arguments.

It is important to emphasize that all core scientific ideas, experimental design, results, and conclusions were exclusively developed by the human authors. The LLM served as a sophisticated editing and brainstorming tool. All text generated or modified by the LLM was critically reviewed, edited, and ultimately approved by the authors to ensure it accurately reflected our original intent and scientific findings.

B SPECTRAL ANALYSIS RESULTS

We apply the same procedures to UNI (Chen et al., 2024) and ResNet50 (He et al., 2016) features of Camelyon16 (Litjens et al., 2018) dataset, and present the results in Figs. 4 and 5. From these figures, we can

see that features from these two feature extractors share similar characteristics as those from CONCH (Lu et al., 2024): rapid decaying spectrum and low effective rank. In addition, after applying a random uniform transformation on the features, the high ARIs and the cluster visualizations demonstrate that the random transformation can preserve the manifold structure well, both visually and quantitatively.

Interestingly, our analysis reveals a significant variance in the degree of this compression across models, which reflects their distinct, training-induced information encoding strategies. At one extreme, ImageNet-pretrained ResNet50 exhibits a dramatic rank collapse (from 52.8 to 21.6). We argue this is not a sign of superior and compact structure discovery, but a domain-induced information collapse. This aligns with foundational studies on transfer learning, which show that ImageNet priors are often suboptimal for specialized domains like pathology (Raghu et al., 2019; Zoph et al., 2020), leading to a reliance on superficial "shortcut" features rather than genuine histological patterns (Geirhos et al., 2020). At the other extreme, CONCH, trained with text-alignment, is incentivized to learn a highly abstract, semantic-driven representation that captures key diagnostic concepts (Goh et al., 2021), a behavior consistent with findings in the vision-language models (Liang et al., 2022). This results in a highly compressed, specialized feature manifold. In contrast, UNI trained via self-supervision, is driven to capture a broader spectrum of morphological diversity. Its objective of instance discrimination forces the model to learn features sensitive to fine-grained visual details (Chen et al., 2020; Caron et al., 2021), resulting in a richer, higher-dimensional (yet still low-rank) feature manifold. Crucially, despite these differences, the features from both expert PFMs (CONCH and UNI) provide a far richer and more reliable structural prior than the collapsed representation of the generalist model.

C COMPLEXITY ANALYSIS

C.1 SPACE COMPLEXITY ANALYSIS

In the linear layer, the matrix A_* contains d_0d_1 trainable parameters; however, in our approach, as B is frozen, the only trainable part is LRP, with $(d_0+d_1)r$ trainable parameters. Defining the reduction ratio $r_r=(d_0+d_1)r/(d_0d_1)$, and enforcing $r_r<1$, yields the constraint $r<(d_0d_1)/(d_0+d_1)$. If $r_r\geq 1$, the trainable parameter count in the MR block would match or exceed d_0d_1 , the trainable parameter count in A_* , which makes our MR block meaningless.

C.2 TIME COMPLEXITY ANALYSIS

We begin by analyzing a general case time complexity and then use the result to analyze our method.

In this section, we derive the computational cost of both the forward and backward passes through a single linear layer. We show that each pass scales linearly with the total number of weights d_0d_1 .

For convenience of the following derivations, we define the notations as follows,

$$M = [M_{ij}] \in \mathbb{R}^{d_0 \times d_1}, \tag{3}$$

$$x = [x_j] \in \mathbb{R}^{d_0}, \tag{4}$$

$$z = Mx = [z_i] \in \mathbb{R}^{d_1}, \tag{5}$$

$$\delta = \nabla_z L = [\delta_i] \in \mathbb{R}^{d_1}. \tag{6}$$

C.3 FORWARD PASS

To compute the pre-activation vector z, each component z_i performs a dot product between the *i*th row of M and the input x, the z_i is given by,

$$z_i = \sum_{j=1}^{d_0} M_{ij} x_j. (7)$$

Since there are d_1 such outputs and each requires d_0 multiplications (and roughly the same number of additions), the total cost of the forward pass is $\mathcal{O}(d_0d_1)$.

C.4 BACKWARD PASS

The backward pass consists of three main steps: computing the gradient with respect to the weights, propagating the gradient to the inputs, and updating the weights.

Weight Gradients The derivative of the loss L with respect to each weight M_{ij} follows by the chain rule, is given by,

$$\frac{\partial L}{\partial M_{ij}} = \frac{\partial L}{\partial z_i} \frac{\partial z_i}{\partial M_{ij}} = \delta_i \cdot x_j. \tag{8}$$

Collecting these into the full gradient matrix yields the outer product,

$$\nabla_M L = \delta x^{\top}, \tag{9}$$

which requires computing one scalar product for each of the d_0d_1 entries. Hence, the cost of this step is $\mathcal{O}(d_0d_1)$.

Input Gradients To propagate the error back into the input space, we compute,

$$\nabla_x L = M^{\top} \delta. \tag{10}$$

This is a matrix–vector product of dimensions $d_0 \times d_1$, again costing $\mathcal{O}(d_0 d_1)$.

Weight Update A typical gradient descent update modifies each entry of M by

$$M_{ij} \leftarrow M_{ij} - \eta \frac{\partial L}{\partial M_{ij}}.$$
 (11)

Updating all d_0d_1 weights elementwise thus costs

$$\mathcal{O}(d_0 d_1). \tag{12}$$

C.5 TOTAL TIME COMPLEXITY

Summing the costs of the forward pass and all backward-pass components, we obtain $\mathcal{O}(d_0d_1)$ complexity. Therefore, a single backpropagation step for one example in this layer has time complexity linear in the number of parameters d_0d_1 .

D PROOF OF RANDOM RECALIBRATION MATRIX PRESERVING INPUT CHARACTERISTICS

We categorize the geometric characteristics of the input features into variance and covariance; inner products and norms; cosine similarity; pairwise distances; condition numbers; the restricted isometry property;

subspace embeddings; manifold geometry; cluster structure; nearest-neighbor relationships; and simplex volumes. We then demonstrate that each of these properties satisfies precise invariance conditions, ensuring that the recalibration matrix preserves the underlying geometric structure.

In this proof, we take the Kaiming uniform initialization with $a = \sqrt{5}$ as an example, which is the default initialization method for linear layers.

D.1 VARIANCE AND COVARIANCE PRESERVATION

Intuitively, this result says that when you multiply your data covariance by a "random" matrix M, all of the original variance gets spread out evenly across the new d_1 dimensions, up to the constant factor $d_1/(3 d_0)$. In other words, no particular direction in the original space is favored or suppressed on average.

To derive the scaling of total variance under projection, we start from

$$\operatorname{tr}(M^{\top} \Sigma M) = \sum_{p=1}^{d_1} (M^{\top} \Sigma M)_{pp} = \sum_{p=1}^{d_1} \sum_{q=1}^{d_0} \sum_{r=1}^{d_0} M_{q,p} \Sigma_{q,r} M_{r,p}.$$
(13)

Taking expectation and using linearity gives

$$E[\operatorname{tr}(M^{\top} \Sigma M)] = \sum_{p=1}^{d_1} \sum_{q=1}^{d_0} \sum_{r=1}^{d_0} \Sigma_{q,r} E[M_{q,p} M_{r,p}].$$
(14)

Because the entries $M_{q,p}$ are independent with $E[M_{q,p}] = 0$, only terms with q = r survive:

$$E[M_{q,p} M_{r,p}] = \begin{cases} Var(M_{q,p}) = \frac{1}{3 d_0}, & q = r, \\ 0, & q \neq r. \end{cases}$$
 (15)

Substituting back, we have

$$E[\operatorname{tr}(M^{\top} \Sigma M)] = \sum_{p=1}^{d_1} \sum_{q=1}^{d_0} \Sigma_{q,q} \frac{1}{3 d_0} = \frac{d_1}{3 d_0} \sum_{q=1}^{d_0} \Sigma_{q,q} = \frac{d_1}{3 d_0} \operatorname{tr}(\Sigma).$$
 (16)

D.2 INNER PRODUCTS AND NORMS

A random projection with independent, zero-mean entries treats every coordinate of your vectors in an unbiased, "democratic" way. On average, the dot product between any two vectors u and v simply gets scaled by the factor $d_1/(3 d_0)$, but its sign and relative magnitude remain the same. Geometrically, this means angles and lengths are preserved in expectation, no particular direction in the original space is preferred, and so the overall similarity structure of the data survives the embedding up to a known scale.

To see how pairwise inner products scale, write

$$\langle M u, M v \rangle = \sum_{p=1}^{d_1} (M u)_p (M v)_p = \sum_{p=1}^{d_1} \sum_{i=1}^{d_0} \sum_{j=1}^{d_0} M_{p,i} u_i M_{p,j} v_j.$$
 (17)

Taking expectation:

$$E[\langle M u, M v \rangle] = \sum_{p,i,j} u_i v_j E[M_{p,i} M_{p,j}]. \tag{18}$$

By independence and zero mean, only i = j terms contribute:

$$E[M_{p,i} M_{p,j}] = \begin{cases} \frac{1}{3d_0}, & i = j, \\ 0, & i \neq j, \end{cases}$$
 (19)

so

$$E[\langle M u, M v \rangle] = \sum_{n=1}^{d_1} \sum_{i=1}^{d_0} u_i v_i \frac{1}{3 d_0} = \frac{d_1}{3 d_0} \langle u, v \rangle.$$
 (20)

Similarly, the norm squared follows as the special case u = v,

$$E[\|Mu\|^2] = E[\langle Mu, Mu \rangle] = \frac{d_1}{3 d_0} \|u\|^2.$$
 (21)

D.3 COSINE SIMILARITY

Cosine similarity measures only the angle between two vectors, not their lengths. Since the random projection scales all lengths and dot products by the same constant in expectation, it does not distort angles on average. In effect, no particular direction is stretched more than any other, so the typical cosine similarity between any two points remains exactly as it was before projection.

Since both inner products and norms incur the same factor $\frac{d_1}{3d_0}$ in expectation, the mean cosine similarity is exactly preserved:

$$E[\cos \theta'] = \frac{E[\langle M u, M v \rangle]}{\sqrt{E[\|M u\|^2] E[\|M v\|^2]}} = \frac{\frac{d_1}{3 d_0} \langle u, v \rangle}{\sqrt{\frac{d_1}{3 d_0} \|u\|^2 \frac{d_1}{3 d_0} \|v\|^2}} = \cos \theta.$$
 (22)

D.4 PAIRWISE DISTANCES

This property guarantees that a random projection will, with overwhelming likelihood, maintain the original distances between every pair of data points up to a small, controllable error. In effect, the projection spreads each original distance across many independent components, and by concentration of measure those components collectively reproduce the true separation very closely. As a result, one can embed high-dimensional data into a much lower-dimensional space without significantly distorting the geometric relationships, ensuring that algorithms relying on interpoint distances, such as clustering or nearest-neighbor retrieval, continue to perform reliably.

Let $v = x_i - x_j$. Then

$$||Mv||^2 = \sum_{p=1}^{d_1} \left(\sum_{q=1}^{d_0} M_{p,q} v_q\right)^2.$$
 (23)

Define for each p,

$$Y_p = \sum_{q=1}^{d_0} M_{p,q} v_q, \quad E[Y_p] = 0, \quad \text{Var}(Y_p) = \frac{\|v\|^2}{3 d_0}.$$
 (24)

By concentration of $\sum_{p} Y_{p}^{2}$, with probability $\geq 1 - \delta$,

$$(1 - \varepsilon) \|v\|^2 \le \|M v\|^2 \le (1 + \varepsilon) \|v\|^2, \tag{25}$$

provided

$$d_1 > C \varepsilon^{-2} \ln(N/\delta). \tag{26}$$

D.5 CONDITION NUMBER

This property ensures that a random projection will not substantially worsen the numerical stability of the data. In other words, the projection maintains the balance between the directions in which the data stretches the most and the least. As a result, any algorithms that depend on solving linear systems or performing matrix decompositions will continue to behave reliably after projection, because the projected matrix remains nearly as well-conditioned as the original.

For data matrix X and any unit u,

$$||X M u||^2 = u^{\top} M^{\top} X^{\top} X M u.$$
 (27)

Applying matrix concentration on XM shows that if

$$d_1 \ge C \,\varepsilon^{-2} \,d_0,\tag{28}$$

then with high probability

$$(1 - \varepsilon) u^{\mathsf{T}} X^{\mathsf{T}} X u < u^{\mathsf{T}} M^{\mathsf{T}} X^{\mathsf{T}} X M u < (1 + \varepsilon) u^{\mathsf{T}} X^{\mathsf{T}} X u, \tag{29}$$

so

$$\kappa(XM) \le \frac{1+\varepsilon}{1-\varepsilon} \,\kappa(X) \approx (1+\varepsilon) \,\kappa(X).$$
(30)

D.6 RESTRICTED ISOMETRY PROPERTY

This property ensures that a random projection acts almost like a perfect length-preserver on all sparse vectors at once. In practice, it means that if the data has an underlying sparse structure, we can compress it drastically without losing the ability to tell sparse signals apart or to reconstruct them reliably. By making the projection dimension grow in line with how many nonzeros we expect, we guard against any sparse pattern being overly distorted.

For any K-sparse x,

$$||M x||^2 = \sum_{p=1}^{d_1} \left(\sum_{q \in \text{supp}(x)} M_{p,q} x_q\right)^2.$$
 (31)

A union bound over all supports yields that if

$$d_1 \ge C \,\delta^{-2} \,K \,\ln \frac{d_0}{K},\tag{32}$$

then for all K-sparse x,

$$(1 - \delta) \|x\|^2 \le \|M x\|^2 \le (1 + \delta) \|x\|^2. \tag{33}$$

D.7 SUBSPACE EMBEDDING

This property guarantees that a random projection acts as a near-perfect isometric embedding for any fixed low-dimensional subspace. It means that all geometric relationships, lengths and angles, within that subspace are preserved almost exactly after compression. As a result, any downstream algorithm that relies on the subspace structure (for example, solving linear least-squares problems or performing spectral decompositions) will perform almost identically in the reduced space, yet with greatly reduced computational and storage costs.

For a d-dimensional subspace with basis $A \in \mathbb{R}^{d_0 \times d}$, one shows with a net argument that if

$$d_1 = O(\varepsilon^{-2} d), \tag{34}$$

then with high probability

$$(1 - \varepsilon) A^{\top} A \prec A^{\top} M^{\top} M A \prec (1 + \varepsilon) A^{\top} A. \tag{35}$$

Table 6: Experimental results comparing our method with baselines are presented under various shot settings. Performance metrics in which our method outperforms the baseline are highlighted in italic blue and the best metrics are in bold red. "#P" is the number of parameters (million).

| _ | | | | Camelyon16 | | | NSCLC | | | RCC | |
|----|---------------------|--------------|--|---------------------------------|------------------------|------------------------|-------------------------------|------------------------------------|-------------------------------|-------------------------------|-------------------------------|
| k | Methods | #P | AUC↑ | F1↑ | Acc.↑ | AUC↑ | F1↑ | Acc.↑ | AUC↑ | F1↑ | Acc.↑ |
| | ViLaMIL | 40.7 | $0.538_{0.127}$ | $0.458_{0.115}$ | $0.611_{0.078}$ | $0.776_{0.050}$ | $0.694_{0.028}$ | $0.697_{0.028}$ | $0.903_{0.043}$ | $0.745_{0.068}$ | $0.797_{0.057}$ |
| | ABMIL | 0.26 | $0.412_{0.076}$ | $0.431_{0.065}$ | $0.465_{0.062}$ | $0.774_{0.141}$ | $0.684_{0.098}$ | $0.691_{0.094}$ | $0.844_{0.073}$ | $0.632_{0.081}$ | $0.647_{0.091}$ |
| | MR-ABMIL | 0.10 | | $0.432_{0.086}$ | | | | | | | |
| | TransMIL | 2.41 | | $0.450_{0.049}$ | | | | | | | |
| _ | MR-TransMIL | 2.01 | | 0.475 _{0.089} | | | | 0.717 _{0.064} | | | |
| 2 | CATE | 1.93 | | $0.457_{0.061}$ | | | | | | | |
| | MR-CATE | 0.84 | | 0.519 _{0.161} | | | | | | | |
| | DGRMIL | 4.08 | | 0.427 _{0.037} | | | | | | | |
| | MR-DGRMIL | 3.29 | | 0.520 _{0.054} | | | | | | | |
| | RRTMIL MR-RRTMIL | 2.44 | $0.360_{0.037}$ $0.476_{0.081}$ | 0.566 _{0.060} | | | | | | | |
| | | 2.04 | 0.002 | 0.490 _{0.058} | | | | 0.552 _{0.065} | | | |
| | ViLaMIL | 40.7 | | $0.484_{0.091}$ | | | | | | $0.817_{0.038}$ | 0.0-0 |
| | ABMIL | 0.26 | | $0.492_{0.072}$ | | | | | | | |
| | MR-ABMIL | 0.10 | | 0.438 _{0.100} | | | | | | | |
| | TransMIL | 2.41 | 0 | 0.467 _{0.075} | 0.00- | 0.0-0 | 0.0-0 | 0.00 | 0.000 | 0.00 | 0.0 |
| 4 | MR-TransMIL | 2.01 | | 0.509 _{0.063} | | | | | | | |
| 4 | CATE | 1.93 | | 0.500 _{0.098} | | | | | | | |
| | MR-CATE DGRMIL | 0.84 4.08 | | 0.584 _{0.158} | | | | | | | |
| | MR-DGRMIL | 3.29 | | $0.594_{0.220}$ $0.557_{0.184}$ | | | | | | | |
| | RRTMIL | 2.44 | | $0.537_{0.184}$ $0.548_{0.052}$ | | | | | | | |
| | MR-RRTMIL | 2.04 | | 0.427 _{0.036} | | | | | | | |
| | ViLaMIL | 40.7 | 0.7400 170 | 0.659 _{0.221} | 0.7180.105 | 0.8840.050 | 0.7920.057 | 0.7920.057 | 0.9680,000 | 0.8620.001 | 0.8800.017 |
| | ABMIL | 0.26 | | $0.039_{0.221}$ $0.489_{0.030}$ | | | | | | | |
| | MR-ABMIL | 0.10 | | 0.538 _{0.139} | | | | | | | |
| | TransMIL | 2.41 | | $0.520_{0.089}$ | | | | | | | |
| | MR-TransMIL | 2.01 | | 0.5420.093 | | | | | | | |
| 8 | CATE | 1.93 | $0.798_{0.147}$ | $0.691_{0.200}$ | $0.760_{0.126}$ | $0.877_{0.057}$ | $0.780_{0.049}$ | $0.782_{0.049}$ | $0.978_{0.008}$ | $0.864_{0.030}$ | $0.883_{0.024}$ |
| | MR-CATE | 0.84 | $0.880_{0.148}$ | <i>0.817</i> _{0.160} | $0.829_{0.153}$ | $0.935_{0.053}$ | $0.846_{0.057}$ | <i>0.847</i> _{0.056} | 0.986 _{0.003} | 0.891 _{0.013} | 0.907 _{0.010} |
| | DGRMIL | 4.08 | $0.847_{0.100}$ | $0.725_{0.178}$ | | | | | | | |
| | MR-DGRMIL | 3.29 | 0.934 _{0.005} | | | | | $0.830_{0.094}$ | | | |
| | RRTMIL | 2.44 | $0.688_{0.178}$ | 0.626 _{0.154} | | | | 0.742 _{0.065} | | | |
| | MR-RRTMIL | 2.04 | 0.853 _{0.163} | 0.792 _{0.166} | 0.817 _{0.133} | 0.8/3 _{0.078} | 0.796 _{0.079} | 0.798 _{0.078} | 0.9/8 _{0.004} | 0.863 _{0.015} | 0.8/8 _{0.014} |
| | ViLaMIL | 40.7 | $0.814_{0.231}$ | $0.775_{0.220}$ | $0.820_{0.150}$ | $0.928_{0.029}$ | $0.849_{0.027}$ | $0.850_{0.027}$ | $0.974_{0.011}$ | $0.855_{0.030}$ | $0.879_{0.022}$ |
| | ABMIL | 0.26 | | $0.712_{0.154}$ | | | | | | | |
| | MR-ABMIL | 0.10 | $0.790_{0.231}$ | $0.745_{0.224}$ | $0.780_{0.171}$ | $0.965_{0.012}$ | 0.891 _{0.010} | 0.891 _{0.010} | $0.983_{0.007}$ | $0.891_{0.028}$ | $0.911_{0.020}$ |
| | TransMIL | 2.41 | | $0.625_{0.192}$ | | | | | | | |
| | MR-TransMIL | 2.01 | | 0.751 _{0.148} | | | | | | 0.893 _{0.016} | |
| 16 | CATE | 1.93 | | 0.875 _{0.032} | | | | | | | |
| | MR-CATE | 0.84 | 0.924 _{0.032} | | | | | 0.887 _{0.033} | | | |
| | DGRMIL MB DGRMII | 4.08 | 0.885 _{0.050} | | | | | 0.858 _{0.039} | | | |
| | MR-DGRMIL RRTMIL | 3.29 | $0.923_{0.058}$ $0.840_{0.169}$ | | | | | 0.867 _{0.027} | | | |
| | MR-RRTMIL | 2.04 | 0.840 _{0.169} 0.937 _{0.016} | | | | | $0.827_{0.041}$ $0.828_{0.071}$ | | | |
| | IVIA-IXIXI IVIIL | 2.04 | 0.016 | 0.0700.024 | 0.0700.023 | 0.7100.054 | 0.0230.076 | 0.0200.071 | 0.7760.008 | 0.0050.041 | 0.0070.029 |

D.8 MANIFOLD GEOMETRY

This property ensures that a random projection will faithfully reproduce the intrinsic geometry of any low-dimensional manifold embedded in high-dimensional space. Even though the data may lie on a curved, nonlinear surface, we can compress it down to far fewer dimensions and still retain nearly all of the true "geodesic" distances along that surface. The key idea is that the manifold can be approximated by a finite

Table 7: Performance of MR-ABMIL with replaced linear layers under varying shot settings on Camelyon16, TCGA-NSCLC, and TCGA-RCC datasets. "U" and "V" denote two matrices in MR-ABMIL.

| \overline{k} | Methods | #P | AUC↑ | Camelyon16 F1↑ | Acc.↑ | AUC↑ | NSCLC F1↑ | Acc.↑ | AUC↑ | RCC F1↑ | Acc.↑ |
|----------------|-----------------|------|---|-------------------|-------------------------------|-----------------|-----------------|-----------------|-------------------------------|-----------------|-----------------|
| 4 | V + U V U | 0.18 | 0.461 _{0.137} 0.456 _{0.003} 0.528 _{0.105} | $0.439_{0.002}$ | 0.598 _{0.014} | $0.861_{0.035}$ | $0.782_{0.032}$ | $0.784_{0.031}$ | $0.940_{0.048}$ | $0.790_{0.075}$ | $0.817_{0.073}$ |
| 8 | V + U V U | 0.18 | $0.565_{0.155} \\ 0.714_{0.149} \\ \textbf{0.728}_{0.191}$ | $0.655_{0.121}$ | $0.699_{0.073}$ | $0.951_{0.035}$ | $0.861_{0.047}$ | $0.862_{0.046}$ | $0.968_{0.013}$ | $0.843_{0.054}$ | $0.870_{0.044}$ |
| 16 | V + U V U | 0.18 | $0.790_{0.231} \\ 0.864_{0.117} \\ 0.825_{0.184}$ | $0.809_{0.120}$ | $0.819_{0.118}$ | $0.960_{0.017}$ | $0.887_{0.014}$ | $0.887_{0.014}$ | 0.983 _{0.008} | $0.877_{0.023}$ | $0.900_{0.020}$ |

set of local patches, and the projection preserves each patch's geometry so accurately that the entire shape remains intact. As a result, analyses that rely on the manifold's structure, such as nonlinear dimensionality reduction or manifold-based learning, remain valid and effective after compression.

Cover a d-dimensional manifold of volume V by $N \approx (V/\tau)^d$ balls. If

$$d_1 \ge C \,\varepsilon^{-2} \,d\,\ln\frac{V}{\tau},\tag{36}$$

then all geodesic distances are preserved within $1 \pm \varepsilon$.

D.9 CLUSTER LABELS

This criterion ensures that the worst-case contraction of the gap between clusters still exceeds the worst-case expansion of each cluster's size. In practice it means that random projection will not cause any overlap between clusters that were originally well separated. As a result, any clustering or classification based on distance remains unchanged, and all original labels are preserved.

Clusters of diameter D separated by Δ satisfy label preservation if

$$(1 - \varepsilon) \Delta > (1 + \varepsilon) D. \tag{37}$$

D.10 NEAREST NEIGHBORS

The nearest-neighbor property refers to the guarantee that, after random projection, each point's set of k closest points (its k-nearest-neighbor graph) remains exactly the same as in the original high-dimensional space.

Fast JL transforms M = PHD satisfy the same distortion bounds and run in time $O(d_0 \log d_0 + d_1)$; the k-NN graph is unchanged if

$$\varepsilon < \frac{\gamma}{2D}.$$
 (38)

D.11 SIMPLEX VOLUME

A simplex in this context is the most elementary convex polytope determined by one "base" point together with a set of other points that do not all lie in a lower-dimensional subspace. In two dimensions it is a triangle,

Table 8: Experimental results on different initialization methods in MR-CATE. "K." and "X." stands for Kaiming and Xavier initialization. "U." and "N." stands for uniform and normal distribution. The number following is the input parameter for initialization. The gray rows stand for our main setting.

| | 3.5.1.1 | (| Camelyon16 | | | NSCLC | | | RCC | |
|----------|------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|------------------------|-------------------------------|-------------------------------|-------------------------------|
| κ | Methods | AUC↑ | F1↑ | Acc.↑ | AUC↑ | F1↑ | Acc.↑ | AUC↑ | F1↑ | Acc.↑ |
| | K. N. 0 | $0.630_{0.161}$ | $0.495_{0.197}$ | 0.657 _{0.111} | $0.822_{0.062}$ | $0.750_{0.057}$ | $0.752_{0.055}$ | $0.980_{0.014}$ | $0.878_{0.018}$ | 0.901 _{0.018} |
| | K. U. 0 | $0.624_{0.170}$ | $0.489_{0.169}$ | $0.637_{0.105}$ | $0.865_{0.044}$ | $0.794_{0.032}$ | $0.794_{0.031}$ | $0.981_{0.014}$ | $0.882_{0.030}$ | $0.903_{0.025}$ |
| 4 | K. U. $\sqrt{5}$ | $0.621_{0.169}$ | $0.584_{0.158}$ | $0.623_{0.135}$ | $0.870_{0.039}$ | 0.794 _{0.032} | $0.794_{0.032}$ | 0.982 _{0.012} | $0.887_{0.020}$ | 0.905 _{0.018} |
| | X. N. 0 | 0.662 _{0.141} | $0.488_{0.201}$ | $0.647_{0.129}$ | $0.840_{0.053}$ | $0.763_{0.042}$ | $0.764_{0.042}$ | $0.981_{0.014}$ | $0.879_{0.034}$ | $0.901_{0.029}$ |
| | X. U. 0 | $0.611_{0.171}$ | $0.512_{0.178}$ | $0.636_{0.128}$ | $0.864_{0.046}$ | $0.791_{0.027}$ | $0.791_{0.026}$ | $0.981_{0.013}$ | $0.882_{0.029}$ | $0.902_{0.026}$ |
| | K. N. 0 | 0.895 _{0.138} | $0.781_{0.224}$ | 0.833 _{0.122} | $0.904_{0.059}$ | $0.821_{0.059}$ | $0.822_{0.059}$ | $0.985_{0.005}$ | $0.898_{0.021}$ | 0.911 _{0.019} |
| | K. U. 0 | $0.874_{0.138}$ | $0.777_{0.209}$ | $0.828_{0.113}$ | $0.919_{0.048}$ | $0.829_{0.044}$ | $0.830_{0.043}$ | $0.985_{0.005}$ | $0.891_{0.020}$ | $0.907_{0.016}$ |
| 8 | K. U. $\sqrt{5}$ | $0.880_{0.148}$ | 0.817 _{0.160} | $0.829_{0.153}$ | 0.935 _{0.053} | $0.846_{0.057}$ | $0.847_{0.056}$ | $0.986_{0.003}$ | $0.891_{0.013}$ | $0.907_{0.010}$ |
| | X. N. 0 | $0.891_{0.132}$ | $0.759_{0.211}$ | $0.812_{0.110}$ | $0.921_{0.071}$ | $0.841_{0.067}$ | $0.842_{0.067}$ | $0.986_{0.004}$ | 0.903 _{0.011} | 0.915 _{0.010} |
| | X. U. 0 | $0.856_{0.147}$ | $0.762_{0.195}$ | $0.816_{0.111}$ | 0.935 _{0.047} | 0.850 _{0.053} | $0.850_{0.052}$ | 0.987 _{0.004} | $0.900_{0.019}$ | 0.915 _{0.016} |
| | K. N. 0 | 0.938 _{0.017} | 0.8890.046 | 0.896 _{0.045} | 0.954 _{0.022} | 0.8690.028 | 0.870 _{0.028} | 0.9860.004 | 0.8820.022 | 0.896 _{0.020} |
| | K. U. 0 | $0.922_{0.022}$ | $0.875_{0.019}$ | $0.884_{0.020}$ | $0.946_{0.023}$ | $0.863_{0.022}$ | $0.864_{0.021}$ | $0.986_{0.004}$ | $0.888_{0.023}$ | $0.909_{0.020}$ |
| 16 | K. U. $\sqrt{5}$ | $0.924_{0.032}$ | $0.867_{0.039}$ | $0.878_{0.034}$ | $0.961_{0.014}$ | $0.886_{0.033}$ | $0.887_{0.033}$ | $0.986_{0.006}$ | 0.898 _{0.031} | 0.916 _{0.023} |
| | X. N. 0 | 0.945 _{0.015} | $0.901_{0.015}$ | 0.907 _{0.016} | $0.956_{0.020}$ | $0.885_{0.014}$ | $0.885_{0.014}$ | 0.987 _{0.005} | $0.889_{0.022}$ | $0.909_{0.014}$ |
| | X. U. 0 | $0.927_{0.037}$ | $0.885_{0.035}$ | $0.891_{0.034}$ | $0.958_{0.013}$ | $0.887_{0.031}$ | $0.888_{0.031}$ | 0.987 _{0.004} | $0.895_{0.031}$ | $0.912_{0.025}$ |

in three it is a tetrahedron, and in higher dimensions the generalization of those. Its volume quantifies the amount of space enclosed by those corner points. Because the projection acts almost as an exact length-preserver on each independent direction in the simplex, the total enclosed volume can only change by the small, precisely bounded factor implied by the near-isometry.

For d+1 affinely independent points,

$$Vol(\Delta) = \frac{1}{d!} \left| \det[x_1 - x_0, \dots, x_d - x_0] \right|.$$
 (39)

Since M is an approximate isometry on this d-dimensional span, its singular values lie in $[\sqrt{1-\varepsilon}, \sqrt{1+\varepsilon}]$, giving

$$(1 - \varepsilon)^{d/2} \operatorname{Vol}(\Delta) \le \operatorname{Vol}(\Delta') \le (1 + \varepsilon)^{d/2} \operatorname{Vol}(\Delta), \tag{40}$$

hence squared volume changes by at most $(1 \pm \varepsilon)^d$.

E Proof of Approximation Capabilities of Manifold Residual Block

In this section, we prove that our MR block can approximate any linear layer to arbitrary precision, thereby providing a worst-case performance guarantee: if the MR block cannot improve over the linear layer then it can match its behavior arbitrarily closely, ensuring no loss in worst-case performance. In addition, rather than directly proving our formulation, we establish a more fundamental result without activation functions. Subsequently, with the universal approximation theorem (Barron, 1993; Cybenko, 1989; Funahashi, 1989; Hornik et al., 1989; Lu & Lu, 2020), our architecture with non-linear activation function can achieve even superior approximation precision.

Theorem E.1 (Universal Approximation by Manifold Residual Block). Let $A_* \in \mathbb{R}^{d_0 \times d_1}$ have full rank $r_A = \min\{d_0, d_1\}$. Draw a frozen matrix $B \in \mathbb{R}^{d_0 \times d_1}$ with sub-Gaussian entries. Then almost surely: for every $\varepsilon > 0$ there exist integers $r \leq r_A$ and matrices W_2 and W_1 such that,

$$\|\boldsymbol{A}_* - (\boldsymbol{B} + \boldsymbol{W}_2 \, \boldsymbol{W}_1)\|_F \le \varepsilon$$
, where $\boldsymbol{W}_2 \in \mathbb{R}^{d_0 \times r}$, and $\boldsymbol{W}_1 \in \mathbb{R}^{r \times d_1}$. (41)

Table 9: Experimental results on different pretrained feature extractors under various shot settings. "R50" stands for ResNet50, and "MR-" stands for MIL model with our MR block.

| | | | | ABM | IIL | | | | | CA | ГЕ | | |
|----|---------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|
| k | Methods | | NSCLC | | | RCC | | | NSCLC | | | RCC | |
| | | AUC↑ | F1↑ | Acc.↑ |
| | R50 | 0.6160.026 | $0.535_{0.068}$ | $0.571_{0.036}$ | $0.747_{0.074}$ | $0.577_{0.075}$ | 0.6110.092 | $0.532_{0.082}$ | $0.403_{0.105}$ | $0.510_{0.033}$ | $0.778_{0.034}$ | $0.470_{0.053}$ | 0.6170.090 |
| 4 | MR-R50 | 0.6230.062 | 0.5530.117 | $0.590_{0.047}$ | 0.8020.059 | 0.6310.076 | 0.6700.066 | $0.514_{0.091}$ | <i>0.413</i> _{0.114} | $0.514_{0.042}$ | $0.786_{0.027}$ | 0.4820.089 | $0.596_{0.140}$ |
| - | UNI | $0.745_{0.085}$ | $0.685_{0.071}$ | $0.690_{0.070}$ | $0.966_{0.018}$ | $0.838_{0.043}$ | $0.861_{0.029}$ | $0.646_{0.086}$ | $0.416_{0.128}$ | $0.533_{0.076}$ | $0.905_{0.045}$ | $0.688_{0.058}$ | $0.714_{0.090}$ |
| | MR-UNI | 0.865 _{0.081} | 0.775 _{0.097} | 0.782 _{0.084} | 0.979 _{0.014} | 0.875 _{0.023} | 0.896 _{0.026} | $0.590_{0.115}$ | 0.497 _{0.127} | 0.561 _{0.097} | 0.921 _{0.045} | 0.735 _{0.097} | 0.769 _{0.091} |
| | R50 | 0.5900.037 | 0.5470.081 | $0.570_{0.046}$ | 0.7460.069 | $0.566_{0.090}$ | 0.5850.107 | $0.649_{0.078}$ | $0.574_{0.106}$ | $0.600_{0.069}$ | $0.826_{0.110}$ | $0.642_{0.102}$ | $0.689_{0.070}$ |
| 8 | MR-R50 | 0.6430.042 | 0.5970.033 | 0.6010.032 | 0.8080.045 | 0.6230.067 | 0.6390.083 | $0.635_{0.064}$ | 0.6070.057 | 0.6140.050 | 0.8950.036 | 0.7500.072 | 0.7660.060 |
| 0 | UNI | | | | | | | $0.827_{0.078}$ | | | | | |
| | MR-UNI | 0.913 _{0.050} | 0.817 _{0.044} | 0.818 _{0.043} | 0.986 _{0.006} | 0.912 _{0.018} | 0.923 _{0.014} | 0.888 _{0.066} | $0.806_{0.068}$ | $0.808_{0.067}$ | 0.985 _{0.004} | 0.908 _{0.014} | 0.918 _{0.011} |
| | R50 | 0.6560.099 | 0.6230.090 | $0.625_{0.088}$ | $0.840_{0.040}$ | $0.677_{0.038}$ | $0.709_{0.051}$ | $0.663_{0.101}$ | $0.562_{0.128}$ | $0.611_{0.074}$ | $0.900_{0.080}$ | $0.747_{0.132}$ | $0.780_{0.090}$ |
| 16 | MR-R50 | 0.6720.066 | 0.6370.050 | $0.639_{0.048}$ | $0.859_{0.021}$ | 0.7080.024 | 0.7340.018 | $0.663_{0.126}$ | $0.592_{0.151}$ | $0.630_{0.082}$ | 0.941 _{0.042} | $0.814_{0.051}$ | $0.824_{0.056}$ |
| 10 | UNI | | | | | | | $0.872_{0.075}$ | | | | | |
| | MR-UNI | 0.948 _{0.028} | 0.870 _{0.036} | 0.870 _{0.036} | 0.989 _{0.004} | 0.911 _{0.024} | 0.924 _{0.018} | 0.917 _{0.046} | 0.826 _{0.060} | 0.829 _{0.056} | 0.985 _{0.006} | 0.910 _{0.019} | 0.922 _{0.015} |

Proof. Full-Rank Guarantees and Singular-Value Bounds. We prove that B is full rank almost surely in Sec. F.1. Let $\delta, \eta \in (0, 1)$, according to Johnson–Lindenstrauss Lemma, we have,

$$(1-\delta)\|\boldsymbol{x}\|_{2} \leq \frac{1}{\sqrt{d_{0}}}\|\boldsymbol{B}\,\boldsymbol{x}\|_{2} \leq (1+\delta)\|\boldsymbol{x}\|_{2} \implies \forall \sigma(\boldsymbol{B}) \in [\sqrt{d_{0}}(1-\delta), \sqrt{d_{0}}(1+\delta)],$$
 (42)

with probability at least $1 - \eta$, if the input dimension satisfies $d_0 \ge C(d_1 + \ln(1/\eta))/\delta^2$.

Truncated SVD of the Residual. Let $E = A_* - B$. Its singular value decomposition is given as $E = U \Sigma V^{\top}$. By the Eckart-Young-Mirsky theorem, we choose $r \leq r_A$ and we have,

$$\|\boldsymbol{E} - \boldsymbol{E}_r\|_F = \sqrt{\sum_{i=r+1}^{\text{rank}(\boldsymbol{E})} \sigma_i^2(\boldsymbol{E})} \le \varepsilon, \text{ where } \boldsymbol{E}_r = \sum_{i=1}^r \sigma_i \, \boldsymbol{u}_i \, \boldsymbol{v}_i^\top.$$
 (43)

Low-rank Correction Construction. From the last step, we can set W_1 and W_2 as,

$$\boldsymbol{W}_{2} = [\boldsymbol{u}_{1}, \dots, \boldsymbol{u}_{r}] \boldsymbol{\Sigma}_{r}^{1/2}, \quad \boldsymbol{W}_{1} = \boldsymbol{\Sigma}_{r}^{1/2} [\boldsymbol{v}_{1}, \dots, \boldsymbol{v}_{r}]^{\mathsf{T}}, \text{ where } \boldsymbol{\Sigma}_{r} = \operatorname{diag}(\sigma_{1}, \dots, \sigma_{r}),$$
(44)

which satisfies $W_2W_1 = E_r$ and therefore, for sufficiently large r, we have

$$\|\mathbf{A}_* - (\mathbf{B} + \mathbf{W}_2 \mathbf{W}_1)\|_F = \|\mathbf{E} - \mathbf{E}_r\|_F < \varepsilon.$$
 (45)

F ADDITIONAL PROOFS

F.1 RANDOMLY INITIALIZED MATRIX HAS FULL RANK

Theorem F.1. Let $W \in \mathbb{R}^{d_0 \times d_1}$ have entries drawn i.i.d. from a continuous distribution (e.g. Kaiming–uniform with gain $\sqrt{5}$). Then with probability one,

$$\operatorname{rank}(\boldsymbol{W}) = \min\{d_0, d_1\}. \tag{46}$$

Proof. Set $r = \min\{d_0, d_1\}$. Define the singular set

$$S = \{ \boldsymbol{W} : \operatorname{rank}(\boldsymbol{W}) < r \}. \tag{47}$$

1175 A matrix in S must make every $r \times r$ minor vanish. Label those minors by index sets I, and write

1177
$$Z_I = \{ \boldsymbol{W} : \det_I(\boldsymbol{W}) = 0 \}. \tag{48}$$

Each \det_I is a nonzero polynomial in the entries of W, so by basic measure theory its zero set Z_I has Lebesgue measure zero. Since there are finitely many I,

$$S = \bigcup_{I} Z_{I} \tag{49}$$

is a finite union of null sets, hence also null. Finally, the distribution on entries is absolutely continuous with respect to Lebesgue measure, so

$$\Pr(\boldsymbol{W} \in S) = 0,\tag{50}$$

and therefore W has full rank r almost surely.

F.2 MAXIMUM RANK OF THE LRP

Theorem F.2 (Rank of a Low-Rank Product). Let $W_2 \in \mathbb{R}^{m \times r}$ and $W_1 \in \mathbb{R}^{r \times n}$ with $r < \min(m, n)$. Then

$$\operatorname{rank}(\boldsymbol{W}_2\boldsymbol{W}_1) \leq r. \tag{51}$$

Proof. Observe that

$$rank(\mathbf{W}_2) \le r, \qquad rank(\mathbf{W}_1) \le r, \tag{52}$$

since W_2 has only r columns and W_1 has only r rows. A standard rank-inequality for matrix products states

$$rank(AB) \le min(rank(A), rank(B)) \tag{53}$$

for any conformable A, B. Applying this with $A = W_2$ and $B = W_1$ gives

$$\operatorname{rank}(\boldsymbol{W}_2 \boldsymbol{W}_1) \leq \min(\operatorname{rank}(\boldsymbol{W}_2), \operatorname{rank}(\boldsymbol{W}_1)) \leq r. \tag{54}$$

G.1 Dataset Descriptions

EXPERIMENT SETUPS

Camelyon16 Litjens et al. (2018) comprises 397 WSIs of lymph node sections from breast cancer patients, annotated for metastatic presence. The dataset is officially partitioned into 270 training specimens (157 normal, 111 tumor-containing) and 129 testing specimens.

TCGA-RCC¹ encompasses 940 WSIs from three distinct renal cell carcinoma subtypes: Kidney Chromophobe (TGCA-KICH, 121 WSIs from 109 cases), Kidney Renal Clear Cell Carcinoma (TCGA-KIRC, 519 WSIs from 513 cases), and Kidney Renal Papillary Cell Carcinoma (TCGA-KIRP, 300 WSIs from 276 cases).

TCGA-NSCLC consists of 1,053 WSIs representing two major lung cancer histological subtypes: Lung Squamous Cell Carcinoma (TCGA-LUSC, 512 WSIs from 478 cases) and Lung Adenocarcinoma (TCGA-LUAD, 541 WSIs from 478 cases). Following the same methodology as for the RCC dataset.

¹The TCGA data used in our work is available in https://portal.gdc.cancer.gov.

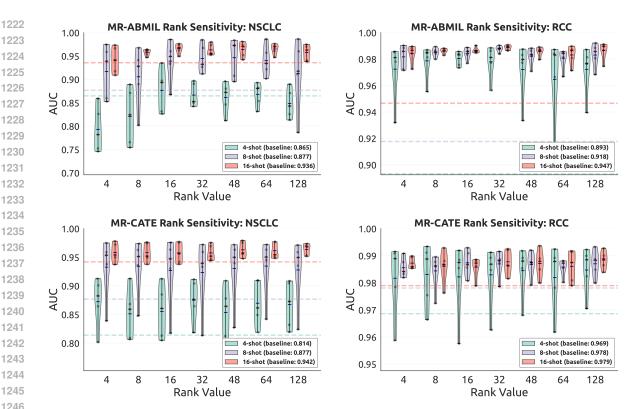


Figure 6: Sensitivity analysis of parameter *r* across multiple datasets using MR-ABMIL and MR-CATE. Dashed lines indicate the performance of the respective baseline methods (*i.e.*, ABMIL and CATE).

G.2 EXPERIMENTAL PROTOCOL AND HYPER-PARAMETER SETTING

Experimental Protocol. We benchmark our approach on established multiple instance learning frameworks: AB-MIL (Ilse et al., 2018) (attention-based method), TransMIL (Shao et al., 2021) (Transformer-based method), CATE (Huang et al., 2024) (information-theory-based method), DRGMIL (Zhu et al., 2024) (density modeling method), and RRTMIL (Tang et al., 2024) (re-embedding method). In addition, we also compare our methods with ViLaMIL (Shi et al., 2024), the recent visual language models specifically proposed for few-shot WSI classification. All implementations utilize their official codes. We employ early stopping with a patience of 20 epochs on validation loss, while enforcing a minimum training duration of 50 epochs. For our experiments, we adhere to the official training and test partitions for the Camelyon16 dataset. We further split the official training set into new training and validation sets using a 70/30 ratio. For the TCGA-NSCLC and TCGA-RCC datasets, we perform our own random split, partitioning the data into training, validation, and test sets with a 60/20/20 ratio. From these constructed training sets, we then create our few-shot scenarios by randomly sampling k WSIs per class to create five different subsets for each fold of our cross-validation. The average performance and standard deviation across these five runs is reported.

Hyper-parameters Details. To rigorously assess the universal, plug-and-play value of our block, we adopt a strict paired-comparison protocol under a unified hyperparameter setting. For each baseline architecture, we evaluate its original and MR-enhanced versions while keeping all other conditions identical, ensuring the observed performance gain is solely attributable to our block. We use the AdamW (Loshchilov & Hutter,

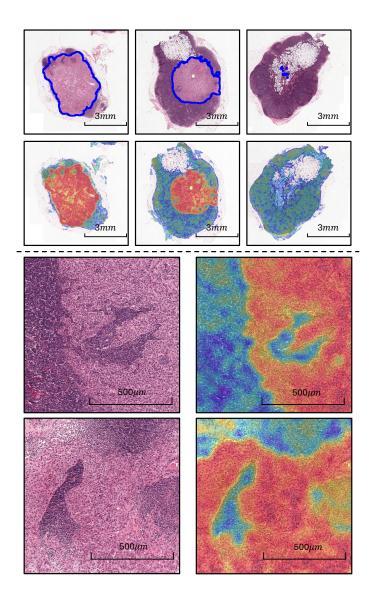


Figure 7: MR-CATE-generated heatmaps for tumor058 (Camelyon16). The left panel shows heatmaps for three biopsy regions, while the right panel displays fine-grained patches with their heatmaps. Red indicates high attention (tumor), whereas blue indicates low attention (normal tissue).

2019) optimizer with learning rate and weight decay being 5×10^{-4} and 10^{-5} . The linear scheduler is utilized for our methods with starting factor and end factor being 0.01 and 0.1, respectively. Dropout is set to 0.25. We use CONCH (Lu et al., 2024), UNI (Chen et al., 2024) and ResNet-50 (He et al., 2016) to extract features from the non-overlapping 224×224 patches, obtained from $20 \times$ magnification of the WSIs, and the resulting feature dimensions are 512, 1,024 and 1,024, respectively. The CONCH features are utilized throughout our experiments as they demonstrate superior performance than the other two features. The UNI and ResNet-50 features are only utilized for investigating the robustness of our method against features from

different foundation models. The rank r is set to 64. GELU (Hendrycks & Gimpel, 2016) is chosen as it demonstrates the most superior performance among others.

G.3 COMPUTER RESOURCES

We conduct our experiments on an NVIDIA A100 GPU and Intel(R) Xeon(R) Silver 4410Y CPU, with Ubuntu system version 22.04.5 LTS (GNU/Linux 5.15.0-136-generic x86_64). More detailed python environmental information would be released with the code.

G.4 REPLACEMENT POSITION OF THE MODELS

Preliminaries of ABMIL. ABMIL can be formulated as follows,

$$a_{k} = \frac{\exp\left\{\boldsymbol{w}^{\top}\left(\tanh(\boldsymbol{V}\boldsymbol{h}_{k}^{\top})\odot\operatorname{sigm}(\boldsymbol{U}\boldsymbol{h}_{k}^{\top})\right)\right\}}{\sum_{j=1}^{K}\exp\left\{\boldsymbol{w}^{\top}\left(\tanh(\boldsymbol{V}\boldsymbol{h}_{j}^{\top})\odot\operatorname{sigm}(\boldsymbol{U}\boldsymbol{h}_{j}^{\top})\right)\right\}}, \qquad \boldsymbol{z} = \sum_{k=1}^{K}a_{k}\boldsymbol{h}_{k},$$
(55)

where z is the slide-level feature.

ABMIL: replace both the V and U matrix with our MR block.

TransMIL²: only replace the out matrix in nystrom attention implementation with our MR block.

1336 CATE³: replace all linear layers in x_linear, interv_linear, feature output linear layer and the second linear layer in encoder_IB with our MR block.

DGRMIL⁴: replace linear layers in self.fcl of optimizer_triple and self.crossffn of DGRMIL.

1341 RRTMIL⁵: replace linear layer in self.proj of InnerAttention.

More details will be available along with the code.

H ADDITIONAL EXPERIMENTAL RESULTS

The additional experimental results are presented in Tables 7 to 9. The discussions of these tables are already available in Sec. 4.

I LIMITATIONS AND FUTURE WORKS

This work introduces an MR block to mitigate the manifold degradation of the vanilla linear layer while enabling parameter-efficient task-specific adaptation. However, our empirical evaluation has several limitations. First, we fix the rank r across all layers, a choice that is unlikely to be optimal in real-world scenarios. A more comprehensive exploration of different rank values within a single model would better characterize the trade-off between model size and accuracy. Second, our study of insertion positions is limited to the ABMIL architecture; to establish broader applicability, the MR block should also be tested in more sophisticated MIL frameworks such as TransMIL and CATE. We treat these limitations as future directions of our work.

²https://github.com/szc19990412/TransMIL

³https://github.com/HKU-MedAI/CATE

⁴https://github.com/ChongQingNoSubway/DGR-MIL

⁵https://github.com/DearCaat/RRT-MIL

We anticipate that the proposed MR block will broadly advance few-shot learning within computational pathology and beyond. Its minimal deployment requirements render it readily applicable in a wide range of settings. In computational pathology, it may enhance tasks such as few-shot report generation, survival outcome prediction, mutation inference, and treatment-response modeling, among others. More generally, the MR block could improve performance in computer vision applications, including few-shot segmentation and classification, and may even be adapted to natural-language-processing tasks. Further exploration across these domains is therefore warranted.