A Viskores-Enhanced Vision Transformer Framework for MRI-Based Cancer Image Classification on Frontier

Suzan Anwar
Computer Science Department
Philander Smith University
Little Rock, Arkansas, USA

Percia Burrows
Computer Science Department
Philander Smith University
Little Rock, Arkansas, USA

Ken Moreland
Oak Ridge National Laboratory
Oak Ridge, Tennessee, USA

***Abstract*—This work investigates the integration of Viskores, a performance-portable visualization framework, into a Vision Transformer (ViT)-based MRI classification pipeline executed on the Frontier supercomputer. Instead of training directly on raw magnetic resonance imaging (MRI) data, our approach first generates visualization-derived feature representations using Viskores operations, including median filtering, Sobel gradients, Laplacian filtering, texture extraction, Otsu thresholding, and morphological processing. Our ViT model is then trained using various combinations of these feature extraction operations. Both the scientific visualization and artificial intelligence training and inference were accelerated using Frontier's high-performance computing (HPC) hardware. We evaluate these representations on a four-class brain tumor MRI dataset containing 7,200 images and compare them with a conventional raw-image ViT baseline. The best-performing visualization-enhanced pipeline improves classification accuracy from 91.94% to 95.50%. The poster presents the end-to-end HPC workflow, performance comparison, and an ablation study, demonstrating the potential of integrating scientific visualization and AI within scalable, reproducible workflows for medical imaging applications.**

***Index Terms*—high-performance computing, scientific visualization, Viskores, Vision Transformer, medical imaging, brain tumor MRI, Frontier, artificial intelligence**

I. Introduction

Artificial intelligence has become increasingly important for medical image analysis, including the classification of magnetic resonance imaging data. Vision Transformers have demonstrated strong image-classification capabilities by modeling relationships between image patches. However, most medical imaging workflows train directly on raw image intensities and do not evaluate whether scientific visualization operations can create more informative representations for downstream learning.

Scientific visualization techniques traditionally support human interpretation by enhancing boundaries, gradients, textures, and regions of interest. In this work, we investigate whether these same operations can serve as an automated feature-generation stage for artificial intelligence. The study integrates Viskores with a Vision Transformer workflow on Frontier, the leadership-class supercomputer at the Oak Ridge Leadership Computing Facility.

The primary objective is to compare a conventional raw-image classification pipeline with a visualization-enhanced pipeline while maintaining the same dataset, model architecture, and evaluation procedure.

II. Dataset and Classification Task

The experiments use a brain tumor MRI dataset containing 7,200 two-dimensional images. The dataset includes four diagnostic categories: glioma, meningioma, pituitary tumor, and no tumor.

The classification task predicts one of the four categories from each MRI image. The data are partitioned into training, validation, and testing subsets while preserving class representation.

III. Visualization-Enhanced Workflow

Two experimental pipelines are evaluated.

A. Raw-Image Baseline

The baseline pipeline provides normalized MRI images directly to the Vision Transformer. This establishes the classification performance obtained without visualization-derived preprocessing.

B. Viskores-Enhanced Pipeline

The second pipeline applies scientific visualization operations before classification. The evaluated operations include:

- median filtering for noise reduction,
- Sobel gradients for edge and boundary extraction,
- Laplacian filtering for second-order structural information,
- texture extraction for local spatial variation,
- Otsu thresholding for intensity-based segmentation, and
- morphological processing for refinement of segmented structures.

The resulting representations are provided to the same Vision Transformer architecture used in the baseline. This design isolates the effect of visualization-based feature generation from changes to the underlying neural network.

In addition to the individual operators, an ablation study evaluates fourteen preprocessing configurations to determine whether individual operations or combinations provide the greatest benefit.

TABLE I
BRAIN TUMOR MRI CLASSIFICATION RESULTS

Pipeline	Accuracy	Weighted F1
Raw MRI ViT	91.94%	91.65%
Viskores Enhanced MRI ViT	95.50%	95.42%

IV. HPC EXPERIMENTAL ENVIRONMENT

The workflow was executed on the Frontier supercomputer at Oak Ridge National Laboratory. Frontier provides GPU-accelerated computing through AMD accelerators and the ROCm software ecosystem.

The computational workflow combines Viskores-based processing, PyTorch model training, and SLURM job scheduling. The experiments use consistent training settings across the raw and visualization-enhanced pipelines to support a fair comparison.

Frontier enables the preprocessing, training, and evaluation stages to be organized within a reproducible HPC workflow.

V. RESULTS

The raw-image Vision Transformer baseline achieved an accuracy of 91.94%. The strongest visualization-enhanced configuration used median filtering and achieved an accuracy of 95.50%, corresponding to an improvement of 3.56 percentage points over the baseline.

The ablation study indicates that simpler preprocessing can outperform more complex combinations. Median filtering provided the strongest result, suggesting that reducing image noise while preserving clinically relevant structures can improve the quality of representations learned by the Vision Transformer.

VI. POSTER PRESENTATION

The poster will present the two pipelines side by side, followed by examples of the Viskores-generated representations. A central workflow diagram will show the connection between MRI data, Viskores preprocessing, Vision Transformer training, and classification evaluation on Frontier.

The poster will also include a comparative results chart and an ablation-study visualization covering the evaluated preprocessing configurations. QR codes are planned to provide access to workflow documentation, selected results, and reproducibility materials.

VII. CONCLUSION

This work demonstrates that scientific visualization can serve as more than a tool for human interpretation. When integrated into an artificial intelligence workflow, visualization operations can generate representations that improve downstream classification performance.

The Viskores-enhanced Vision Transformer pipeline improved the brain tumor MRI classification accuracy by 3.56 percentage points over the baseline.

Future work will evaluate repeated training runs, additional datasets, explainability methods, expanded Viskores operators, and three-dimensional BraTS MRI classification.

ACKNOWLEDGMENT

This research was conducted as part of the U.S. Department of Energy Visiting Faculty Program at ORNL. The authors acknowledge the Oak Ridge Leadership Computing Facility and the Frontier supercomputer for supporting the computational experiments.