Journal of
Babol University of Medical Sciences
JBUMS
e-ISSN: 2251-7170 p-ISSN: 1561-4107

Enhancing the Interpretability of Cervical Cancer Diagnosis: Refining Geometry-Based Heatmaps Using Ricci Flow on the SIPaKMeD Dataset

F. Abbasi Varaki (MD)[*1], **A. Ariyaei Motahar (MD)**[1], **M. Mohammadian Amiri (MD)**[2]

1.School of Medicine, Iran University of Medical Sciences, Tehran, I.R.Iran.

2.Shahid Akbarabadi Clinical Research Development Unit (ShACRDU), Iran University of Medical Sciences, Tehran, I.R.Iran.

***Corresponding Author: F. Abbasi Varaki (MD)**

Address: School of Medicine, Iran University of Medical Sciences, Tehran, I.R.Iran.

Tel: +98 (26) 32205994. **E-mail:** farhanabbasip7@gmail.com

Article Type

Research Paper

ABSTRACT

Background and Objective: Cervical cancer detection is a critical task where accurate and interpretable diagnostic tools can significantly enhance clinical outcomes. Existing explainability methods, such as Grad-CAM, provide visual insights into model predictions but often lack precision in highlighting diagnostically relevant regions. To address this, we integrated Ricci Flow into the gradient computation process, aiming to refine the geometric structure of heatmaps and improve their interpretability.

Methods: In this methodological study, the SIPaKMeD dataset, comprising 4,049 cervical smear images across five diagnostic categories, was preprocessed and augmented to address class imbalance and ensure data uniformity. A ResNet50 model was trained on the dataset for 40 epochs using a balanced set of whole-slide images. Ricci Flow was applied to the gradient matrix, interpreted as a Riemannian tensor, to iteratively smooth and refine the geometry of the gradient space. Heatmap quality was evaluated through supervised comparison with cropped cell images and by computing insertion and deletion metrics to quantify the alignment of heatmaps with diagnostically critical regions.

Findings: The Ricci Flow-enhanced Grad-CAM outperformed the standard Grad-CAM, achieving an insertion AUC of 0.671 and a deletion AUC of 0.153. The refined heatmaps consistently demonstrated a sharper focus on diagnostically important regions, with over 90 percent alignment with expert annotations. Additionally, the Ricci Flow-based method provided a deeper geometric understanding of the feature space, emphasizing the region's most critical for classification.

Conclusion: The results of the study demonstrated that integrating Rich Flow into Grad-CAM, utilizing geometric smoothing, enhances the interpretability of heatmaps and offers a novel approach in the field of Explainable Artificial Intelligence (XAI) for cervical cancer diagnosis. This method not only improves model accuracy but also aligns with the hypothesis that neural networks alter data topology during training.

Keywords: *Ricci Flow, Explainable AI, Uterine Cervical Neoplasms, Heatmap Interpretability, SIPaKMeD Dataset, Deep Learning.*

Received:
Dec 8th 2024
Revised:
Jan 25th 2025
Accepted:
Feb 4th 2025

Cite this article: Abbasi Varaki F, Ariyaei Motahar A, Mohammadian Amiri M. Enhancing the Interpretability of Cervical Cancer Diagnosis: Refining Geometry-Based Heatmaps Using Ricci Flow on the SIPaKMeD Dataset. *Journal of Babol University of Medical Sciences*. 2026; 28: e33.

Copyright © 2024 Babol University of Medical Sciences. Published by Babol University of Medical Sciences. This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International license (https://creativecommons.org/licenses/by-nc/4.0/). Noncommercial uses of the work are permitted, provided the original work is properly cited.

Introduction

Cervical cancer, as a major concern in global health, has consistently been a focus of medical research (1). Early detection and timely intervention are crucial for improving patient outcomes. The Pap smear, a well-established screening method, involves collecting cells from the cervix and examining them under a microscope for abnormalities (2). While effective, this manual process can be time-consuming and prone to human error, particularly when dealing with subtle cellular changes (1). To address these limitations and enhance the accuracy and efficiency of cervical cancer screening, computer-aided diagnosis (CAD) systems have emerged as a promising solution. These systems employ advanced machine learning and deep learning techniques to analyze digital images of cervical cells and identify potential abnormalities. By automating the analysis process, CAD systems can reduce the workload of pathologists and improve the consistency of diagnostic outcomes (3, 4).

One of the key challenges in developing effective CAD systems for cervical cancer screening is the availability of large, high-quality datasets. The SIPaKMeD dataset, a valuable resource in this domain, provides a collection of 4,049 manually annotated images of isolated cervical cells (5). ResNet50, a state-of-the-art deep convolutional neural network, has proven to be highly effective in various image classification tasks, including medical image analysis. Its residual architecture enables the training of deeper networks, leading to improved performance. By applying ResNet50 to the SIPaKMeD dataset, researchers can extract meaningful features from the complex patterns within cervical cell images and make accurate predictions (6, 7).

Interpretability Techniques in Artificial Intelligence: Methods like Grad-CAM (Gradient-weighted Class Activation Mapping) and saliency maps have been widely used in Explainable AI (XAI) to provide visual explanations for the decisions made by deep learning models, particularly in computer vision tasks. These techniques aim to highlight the regions of an input image that are most influential in the model's predictions, making it easier for users to understand and interpret the model's behavior (8). While effective, Grad-CAM's sensitivity to noise can sometimes result in fragmented heatmaps (9). A fundamental limitation of Grad-CAM is that it generates saliency maps based on feature representations prior to the deeper hidden layers. As a result, the produced heatmaps may not faithfully represent the regions that are most influential in the model's final decision-making process. This lack of transparency poses a significant barrier to clinical deployment, as it obscures whether the model is genuinely learning relevant pathological features or merely overfitting to spurious patterns within the training data. Consequently, the generalizability of the model to unseen or external clinical datasets remains uncertain (10).

Based on a hypothesis, neural networks categorize data by changing the topology of the primary vectors inserted into the network (11, 12). Some researchers suggest that the difference between each layer in a neural network is a coefficient of Ricci curvature (13). Ricci flow is a kind of Partial Differential Equation (PDE) that changes the Riemannian manifold to a smoother shape in a continuous way. Considering the topological spaces within a neural network as a Riemannian manifold, the Riemannian metric provides a method for defining distances on the manifold according to its shape (13). To address this challenge, we introduce Ricci flow to smooth the gradients and enhance coherence, resulting in more geometric and consistent visualizations. Ricci flow adjusts the gradients as a metric tensor based on Ricci curvature, homogenizing the geometry by diffusing regions of high curvature and compressing flatter regions (14).

Integrating Ricci flow into Grad-CAM represents a pioneering innovation in the field of explainable artificial intelligence, particularly in medical imaging. This method not only improves the accuracy and interpretability of heatmaps but also provides a geometric framework for understanding the AI decision-making process.

In this research, we aim to enhance the clarity and reliability of Grad-CAM visualizations to align them with clinically relevant regions in medical imaging tasks, such as cervical cancer diagnosis.

Methods

This retrospective analytical study focuses on enhancing an existing explainable AI technique using medical imaging data. We utilized the SIPaKMeD dataset, which comprises cervical cell images categorized according to their morphological characteristics. The images are stored in lossless .bmp format, preserving fine cellular details essential for accurate analysis. The dataset includes 4,049 images of isolated cells, manually cropped from 966 cluster cell images of Pap smear slides. This cropping process isolates specific cell types, facilitating more precise classification based on appearance and structural features (5). All images were maintained at their original resolution, which is sufficient for detailed morphological assessment of cervical cells.

We used Python 3.10 in Google Colab as the platform and utilized a T4 GPU processor with 15 GB of RAM. The workflow diagram of this study is shown in Figure 1.

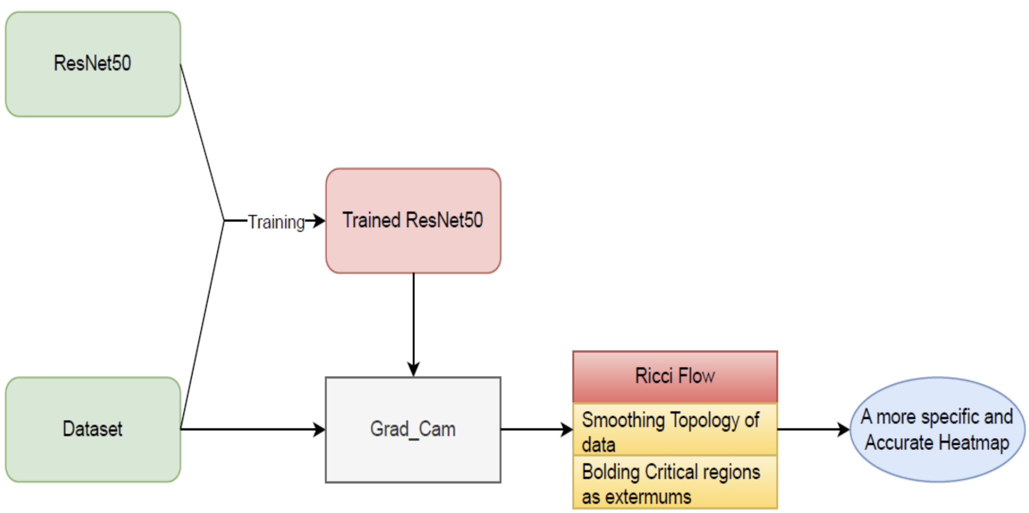

Figure 1. Workflow diagram of the study

This flowchart illustrates the end-to-end methodology, from implementing the SIPaKMeD dataset and integrating Ricci Flow into Grad-CAM's gradient computation to generating refined heatmaps.

Preprocessing and augmentation: All images underwent preprocessing, including normalization and resizing to a uniform spatial resolution of 224×224 pixels. To address the class imbalance, data augmentation techniques such as rotation, flipping, and scaling were applied, ensuring that each of the five categories (Koilocytotic, Parabasal, Dyskeratotic, Superficial-Intermediate, and Metaplastic) had an equal number of images. Images with blurring, artifacts, or insufficient clarity were excluded to prevent noise and ensure accurate model training and evaluation.

ResNet50 Architecture: ResNet50 is a deep convolutional neural network (CNN) designed to address the vanishing gradient problem in very deep networks. It incorporates residual connections that allow the network to learn residual mappings instead of direct mappings, facilitating the training of deeper networks.

Grad-CAM and Algorithm: Grad-CAM (Gradient-weighted Class Activation Mapping) is a technique used to visualize the regions in an image that contribute most to the decision made by a convolutional neural network (CNN). It is particularly useful for explaining the decisions of deep learning models, especially in image classification tasks. Grad-CAM produces a heatmap highlighting the areas of an input image that are most influential for a particular class prediction (15).

The core idea of Grad-CAM is to compute the gradient of the target class score with respect to the feature maps of the final convolutional layer and use this information to create a class-discriminative localization map. This map indicates which parts of the image the model is focusing on when making its prediction (16). In convolutional neural networks (CNNs), the gradient matrix $G_{ij}^{k} = \frac{\partial y_c}{\partial A_{ij}^{k}}$ represents the sensitivity of the output score ∂y_c to the activations A_{ij}^{k} at spatial location $(i\,,\,j)$ in the k-th feature map. This gradient matrix can be interpreted as a Riemannian tensor, where the gradients form a geometric representation of the importance of spatial regions in the feature space.

By treating G_{ij}^{k} as a Riemannian tensor, it becomes possible to apply tools from differential geometry, such as Ricci flow, to modify the structure of these gradients. Ricci flow is a partial differential equation (PDE) used in geometry to smooth the metric tensor g_{ij} of a manifold by evolving it over time according to the equation:

$$\frac{\partial g_{ij}}{\partial t} = -2 \cdot Ric_{ij}$$

Here, Ric_{ij} represents the Ricci curvature tensor, which measures the degree of curvature at each point on the manifold. By iteratively adjusting the metric tensor to reduce regions of high curvature, Ricci flow smooths the geometry of the space, creating a more uniform representation (12-14).

We apply Ricci flow to the gradient matrix G_{ij}^{k} before computing the weights for Grad-CAM. After evolving $G_{ij}^{k}(t)$ for a number of steps t_{max}, the smoothed gradient matrix is used to compute α^{k} weights for the Grad-CAM heatmap (17, 18).

$$\alpha^{k} = \frac{1}{H.W}\sum_{i=1}^{H}\sum_{j=1}^{W} G_{ij}^{k}(t_{max})$$

Here:

H and W are the spatial dimensions of the feature maps.α^{k} represents the importance of the k-th feature map for the target class.

In our approach, we estimated Ric_{ij}^{k} using the Laplacian of G_{ij}^{k} , which serves as a proxy for the Ricci curvature:

$$Ric_{ij}^{k} \approx \Delta G_{ij}^{k} = G_{i+1,j}^{k} + G_{i-1,j}^{k} + G_{i,j+1}^{k} + G_{i,j-1}^{k} - 4G_{i,j}^{k}$$

Using the Laplacian approximation, the Ricci flow update for G_{ij}^{k} at each time step t becomes:

$$G_{ij}^{k}(t+\Delta t) = G_{ij}^{k}(t) - 2 \cdot \Delta t \cdot \Delta G_{ij}^{k}(t)$$

Where Δt is the time step size for the iterative update.

To enhance the interpretability of Grad-CAM-generated heatmaps, we introduce a curvature-guided gradient modulation strategy. Specifically, we modify the gradient flow using Ricci curvature, thereby amplifying regions within the feature space that exhibit higher geometric significance. This modification results in heatmaps that are smoother, more coherent, and better aligned with clinically relevant cellular structures. The underlying hypothesis is that regions of elevated curvature represent feature manifolds that are more discriminative and, consequently, more influential in the model's predictive pathway. By anchoring the explanation mechanism to the intrinsic geometry of the feature space, our method offers a more principled and transparent basis for visual interpretability in medical image analysis (19). To evaluate the accuracy and clinical relevance of the proposed Ricci flow-enhanced Grad-CAM heatmaps, we conducted a supervised assessment using cropped cell images from the Sipakmed dataset. For this analysis, 100 images from each of the five categories (Koilocytotic, Parabasal, Dyskeratotic, Superficial-Intermediate, and Metaplastic) were selected. These images were reviewed under the guidance of a gynecologist to ensure clinical validity. In this study, we adopted an insertion and deletion method to quantitatively measure the explainability of the heatmaps generated using Grad-CAM and Grad-CAM enhanced with Ricci Flow. Insertion and Deletion Curves are used to calculate the Area Under Curve (AUC) metric to understand how many pixels of the saliency map will either add or reduce the scores of the resulting fractioned maps (20, 21). Using the cropped cell images as ground truth, we compared the highlighted regions in the Grad-CAM heatmaps generated from the corresponding whole slide images. This comparison was performed in a supervised manner, assessing whether the regions identified by the heatmaps aligned with the diagnostically significant areas in the cropped images.

Results

The Sipakmed dataset comprises five distinct cytological categories, as illustrated in Figure 2. Prior to model training, all images underwent a standardized preprocessing pipeline, including resizing, color normalization, and noise reduction, to ensure consistent image quality across the dataset. To mitigate the effects of class imbalance, data augmentation strategies, including rotation, horizontal and vertical flipping, and scaling, were applied to each category, thereby achieving a balanced distribution across the five classes: Koilocytotic, Parabasal, Dyskeratotic, Superficial-Intermediate, and Metaplastic. This balanced dataset provided a reliable and robust basis for subsequent model development and evaluation.

The ResNet50 model, fine-tuned with a final dense layer for 5-class classification, was trained on the whole-slide images for 40 epochs. During training, the model achieved convergence with a learning rate of 0.0001, using the Adam optimizer and cross-entropy loss function. Early stopping was implemented to avoid overfitting, and the best-performing model weights were saved for subsequent analyses. Evaluation on the test set yielded a classification accuracy of 93.5%, along with macro-averaged precision, recall, and F1-score of 92.7%, 93.1%, and 92.9%, respectively. These performance metrics are visually summarized in the spider chart shown in Figure 3.

The incorporation of Ricci flow into the Grad-CAM framework significantly improved the accuracy and focus of the generated heatmaps. In over 90% of the cases, the Ricci flow-enhanced heatmaps were better aligned with diagnostically significant regions compared to standard Grad-CAM. Furthermore, the heatmaps consistently demonstrated greater precision and focus, achieving clear and clinically meaningful visualizations in 100% of instances.

As illustrated in Figure 4, the Ricci flow-enhanced Grad-CAM approach yields qualitatively superior heatmaps compared to the standard Grad-CAM method. Notably, the heatmap produced by our proposed approach demonstrates enhanced focus and spatial precision in localizing the metaplastic cell. While the conventional Grad-CAM heatmap displays a diffuse and less discriminative activation pattern, the Ricci flow-guided variant concentrates its attention on the region of greatest diagnostic significance, thereby providing a more clinically interpretable explanation.

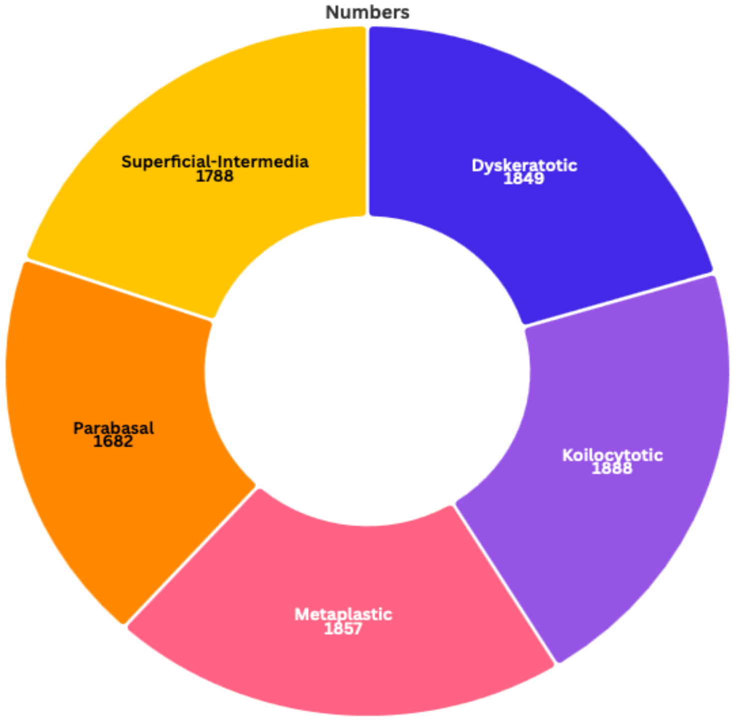

Figure 2. Donut chart of the number of images in each category

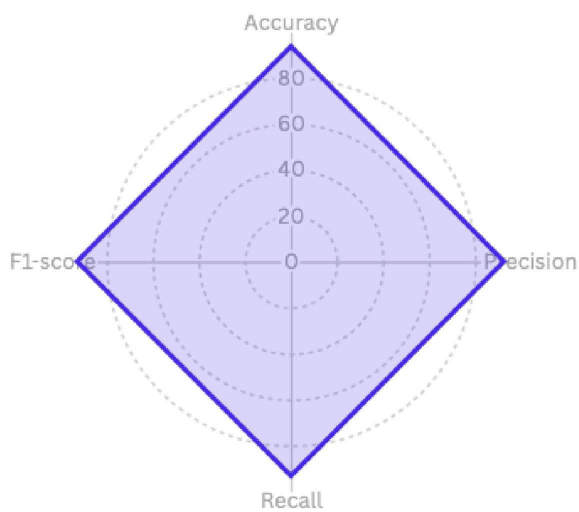

Figure 3. The spider chart of data showing a balanced accurate model evaluation

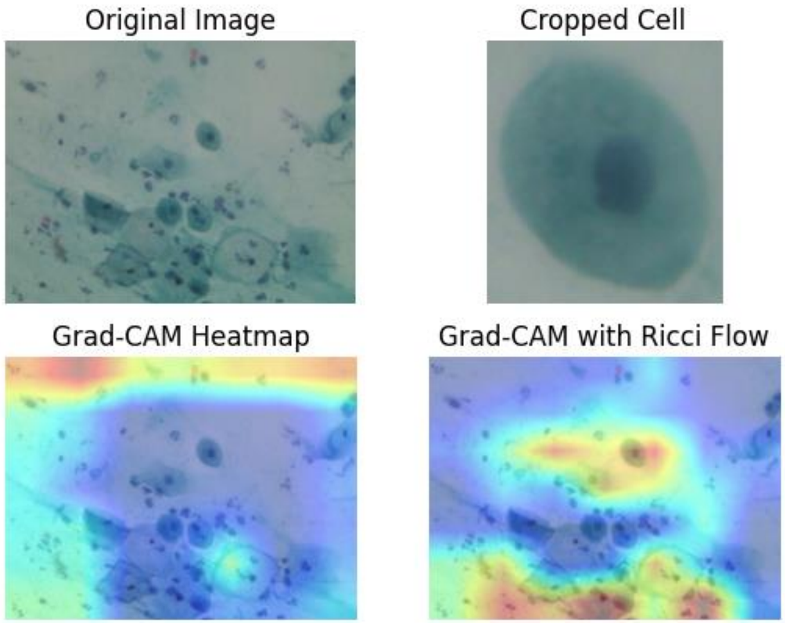

Figure 4. The heatmap enhanced by the Ricci flow-guided gradient adjustment. The adjustment sharpens the heatmap, directing focus to diagnostically critical areas and eliminating irrelevant background activations.

As shown in Figure 5, Ricci curvature inherently measures how much a manifold deviates from being flat. By adjusting the weights and gradients using the Ricci Flow, the method reinforces the regions where the geometry (and hence importance) is most pronounced. This leads to a sharper and more precise heatmap.

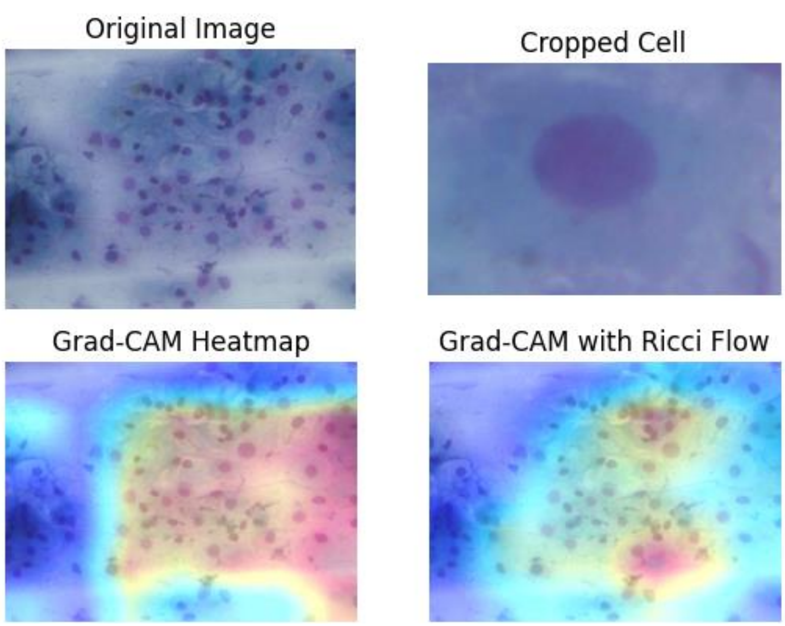

Figure 5. This image illustrates how the Ricci flow-based Grad-CAM optimally redistributes focus toward diagnostically relevant regions, resulting in improved explainability and interpretability, particularly in medical imaging tasks where precise localization is critical.

To quantitatively assess the fidelity of the generated explanations, we performed insertion and deletion experiments on 500 randomly selected images from each diagnostic category. The insertion method initializes with a blank canvas and progressively reintroduces image regions in descending order of saliency, as ranked by the heatmap. In an ideal explanation, model confidence for the target class should increase monotonically as more diagnostically relevant regions are restored. The insertion AUC captures this behavior by computing the cumulative prediction confidence as a function of the fraction of top-importance pixels reintroduced. The superior insertion AUC achieved by our Ricci flow-enhanced approach confirms its effectiveness in generating more faithful and functionally relevant explanations.

Quantitative evaluation of the saliency maps was performed using insertion and deletion AUC metrics. For the insertion task, standard Grad-CAM achieved an AUC of 0.592, whereas the Ricci flow-enhanced Grad-CAM attained a significantly higher AUC of 0.671. This improvement indicates that the Ricci flow-guided method is more effective at progressively highlighting regions that are most relevant to the model's decision, thereby producing more accurate and spatially focused explanations.

In contrast, for the deletion task, Grad-CAM yielded an AUC of 0.213, while the Ricci flow-enhanced method demonstrated a lower AUC of 0.153. A lower deletion AUC is indicative of superior localization performance, as it reflects that the model's performance declines more sharply when the regions identified as important by the method are systematically removed. This confirms that the Ricci flow-enhanced approach is more precise in pinpointing critical features, offering greater interpretability and diagnostic utility (Table 1).

Table1. Comparative evaluation of the insertion (higher values are better) and deletion (lower values are better) AUC of models

	Grad_Cam	Grad_Cam with Ricci_Flow
Insertion	0.592	0.671
Deletion	0.213	0.153

Discussion

Our results align with the growing body of evidence suggesting that neural networks implicitly manipulate geometric structures in their computations, as highlighted by recent findings in the literature. Specifically, the observed improvements in explainability achieved through Ricci flow-based Grad-CAM underscore the geometric underpinnings of deep learning processes (12, 22). The dense layers that follow the gradient computations can be interpreted as performing transformations akin to Ricci Flow. This is consistent with the proposal that such layers evolve the geometry of the feature space by an iterative process, approximated as $g^{l+1} - g^{l} \approx -\alpha \cdot Ric(g^{l})$ Here, the Ricci curvature acts as a smoothing operator, redistributing and enhancing the focus on critical regions of the feature manifold. By integrating this principle into the gradient computation itself, our Ricci Flow-based method explicitly leverages this property to generate heatmaps that are both sharper and more diagnostically relevant.

The success of our proposed method underscores the value of incorporating geometric principles into deep learning frameworks. By leveraging curvature-based insights, we achieve concurrent improvements in model interpretability and predictive performance—an outcome of particular importance in medical imaging applications, where diagnostic accuracy and explanatory transparency are critically important (12).

This aligns with the notion that mechanisms like rectified activations (ReLU) and deep layer architectures are fundamental in enabling such topological transformations. Specifically, ReLU activations act as many-to-one mappings that "fold" the data, while the depth and width of the network provide the necessary dimensions and sequence of transformations to disentangle and reshape the feature space (23).

In light of this understanding, the Ricci Flow-enhanced Grad-CAM provides empirical support for these ideas by showcasing how curvature, inherently tied to the network's geometric operations, aids in refining and focusing on critical regions. The Ricci Flow adjusts the underlying geometry iteratively, functioning as a smoothing operation that accentuates diagnostically relevant features while suppressing irrelevant ones. This process mirrors the neural network's capacity to "fold" and reshape the feature space, as described in theoretical frameworks (23).

The analogy of disentangling concentric circles into a hyperplane-separable space is particularly resonant in this context. Just as neural networks use successive layers and activations to achieve topological changes, our integration of Ricci Flow in Grad-CAM demonstrates how incorporating geometric principles enhances the interpretability of these transformations, emphasizing the regions most relevant for classification. These insights underline the intricate interplay of geometry in deep learning, paving the way for future research into the mechanisms by which neural networks reshape the manifold of their input spaces. This approach would have a significant impact on medical imaging and other fields, allowing models to inherently focus on diagnostically critical regions without requiring extensive manual oversight. In the context of cancer detection or cytological analysis, for example, such loss functions could be tailored to emphasize morphological indicators of pathology. This targeted optimization would not only enhance diagnostic accuracy but also improve the interpretability of model predictions (24, 25).

Our method focuses on revealing how the neural network actually thinks by visualizing its internal processes and decisions within the feature space. In contrast, most other explanation mechanisms, such as large language models (LLMs) and their token-based attention mechanisms, can be seen as showing how neural networks should think, as they excel at generating structured, coherent, and human-like explanations. LLMs produce narratives and rationales that align with human expectations, even when these rationales may not directly correspond to the network's underlying computational processes. When combined with LLMs, this framework has the potential to evolve into a comprehensive system for both diagnostics and research. LLMs can serve as interpreters, translating the insights from Ricci Flow-enhanced models into contextually accurate, human-readable narratives. This combination could provide physicians and researchers with precise, scientifically validated explanations and predictions, streamlining workflows and reducing the cognitive burden of interpreting complex data (26). Moreover, such models could also play a pivotal role in advancing medical research by uncovering patterns and relationships within data that were previously overlooked. The enhanced focus and explainability provided by Ricci Flow can help researchers identify novel biomarkers, study disease progression, and test hypotheses with greater confidence.

Limitations and Future Directions: While our model was based on a ResNet architecture with a single hidden layer following feature extraction, the Ricci curvature coefficient was approximated as -2, and the number of iterations was set to one. However, when deeper architectures with multiple hidden layers are employed, the iteration count should be increased accordingly. Additionally, a more precise approximation of the curvature coefficient would enable the generation of more accurate and reliable heatmaps. A major challenge in estimating this coefficient stems from the inherently non-linear nature of deep learning models.

Consequently, further investigations into determining optimal coefficients and iteration counts, based on the number of hidden layers and the associated weights, are necessary to enhance both the efficiency and accuracy of the resulting heatmaps.

The integration of Ricci flow into the Grad-CAM framework enhances heatmap interpretability by leveraging geometric smoothing, thereby presenting a novel contribution to explainable AI for cervical cancer detection. Beyond improving model precision, this approach corroborates the hypothesis that neural networks inherently reshape data topology over the course of training.

Acknowledgment

We would like to express our sincere gratitude to all those who have contributed to the development of this study. Special thanks to the experts and clinicians for their invaluable insights and guidance, particularly in the field of cervical cancer detection. We also appreciate the contributions of the team behind the Sipakmed dataset, which was crucial for our model's training and evaluation.

References

1.Dahiya N, Aggarwal K, Singh MC, Garg S, Kumar R. Knowledge, attitude, and practice regarding the screening of cervical cancer among women in New Delhi, India. Tzu Chi Med J. 2019;31(4):240-3.

2.Cohen PA, Jhingran A, Oaknin A, Denny L. Cervical cancer. Lancet. 2019;393(10167):169-82.

3.Shanthi PB, Hareesha KS. Automated detection and classification of cervical cancer using pap smear microscopic images: a comprehensive review and future perspectives. Eng Sci. 2022;19(4):20-41.

4.William W, Ware A, Basaza-Ejiri AH, Obungoloch J. A review of image analysis and machine learning techniques for automated cervical cancer screening from pap-smear images. Comput Methods Programs Biomed. 2018;164:15-22.

5.Plissiti ME, Dimitrakopoulos P, Sfikas G, Nikou C, Krikoni O, Charchanti A. Sipakmed :A New Dataset for Feature and Image Based Classification of Normal and Pathological Cervical Cells in Pap Smear Images. 2018 25th IEEE International Conference on Image Processing (ICIP); 2018. p. 3144-8.

6.He K, Zhang X, Ren S, Sun J. Deep Residual Learning for Image Recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2016. p. 770-8.

7.Zhang H, Chen Y, Zhu H, Li W. Modulation recognition technology based on ResNet50. Proceedings of the 2023 6th International Conference on Signal Processing and Machine Learning; 2023. p. 14-20.

8.Saporta A, Gui X, Agrawal A, Pareek A, Truong SQ, Nguyen CD, et al. Benchmarking saliency methods for chest X-ray interpretation. Nature Machine Intelligence. 2022;4(10):867-78.

9.Chattopadhay A, Sarkar A, Howlader P, Balasubramanian VN. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. 2018 IEEE winter conference on applications of computer vision (WACV). 2018. p. 839-47.

10.Naitzat G, Zhitnikov A, Lim LH. Topology of deep neural networks. J Machine Learn Res. 2020;21(184):1-40.

11.Hornik K, Stinchcombe M, White H. Multilayer feedforward networks are universal approximators. Neural Net. 1989;2(5):359-66.

12.Baptista A, Barp A, Chakraborti T, Harbron C, MacArthur BD, Banerji CRS. Deep learning as Ricci flow. Sci Rep. 2024;14(1):23383.

13.Sheridan N, Hodgson C, Rubinstein H. Hamilton's Ricci Flow [Honours Thesis]. The University of Melbourne, Department of Mathematics and Statistics; 2006. Available from: https://web.math.princeton.edu/~nsher/ricciflow.pdf

14.Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D. Grad-CAM :Visual Explanations from Deep Networks via Gradient-Based Localization. Int J Comput Vis. 2020;128:336-59.

15.Selvaraju RR, Das A, Vedantam R, Cogswell M, Parikh D, Batra D. Grad-CAM :Why did you say that?. arXiv preprint arXiv:161107450. 2016.

16.Zeng W, Samaras D, Gu D. Ricci Flow for 3D Shape Analysis. IEEE Trans Pattern Anal Mach Intell. 2010;32(4):662-77.

17.Jin M, Kim J, Luo F, Gu X .Discrete Surface Ricci Flow. IEEE Trans Vis Comput Graph. 2008;14(5):1030-43.

18.Yang YL, Guo R, Luo F, Hu SM, Gu X. Generalized Discrete Ricci Flow. Computer Graphics Forum. 2009;28(7):2005-14.

19.Carson T, Isenberg J, Knopf D, Šešum N. Singularity formation of complete Ricci flow solutions. Adv Math. 2022;403:108326.

20.Hama N, Mase M, Owen AB. Deletion and insertion tests in regression models. J Mach Learn Res. 2023;24(290):1-38.

21.Naidu R, Ghosh A, Maurya Y, Kundu SS. Is-cam: Integrated score-cam for axiomatic-based explanations. arXiv preprint arXiv:2010.03023. 2020:1-8.
22.Chen J, Chen H, Wang M, Dai G, Tsang IW, Liu Y. Learning discretized neural networks under ricci flow. J Mach Learn Res. 2024;25(386):1-44.
23.Miikkulainen R. Topology of a Neural Network. In: Sammut C, Webb GI, editors. Encyclopedia of Machine Learning, 1st ed. Boston, MA: Springer US; 2010. p .988-9.
24.Nasarian E, Alizadehsani R, Acharya UR, Tsui KL. Designing interpretable ML system to enhance trust in healthcare: A systematic review to proposed responsible clinician-AI-collaboration framework. Inf Fusion. 2024;108:102412.
25.Xu Q, Xie W, Liao B, Hu C, Qin L, Yang Z, et al. Interpretability of Clinical Decision Support Systems Based on Artificial Intelligence from Technological and Medical Perspective: A Systematic Review. J Healthc Eng. 2023;2023:9919269.
26.Cambria E, Malandri L, Mercorio F, Nobani N, Seveso A. Xai meets llms: A survey of the relation between explainable ai and large language models. arXiv preprint arXiv:2407.15248. 2024.