Differentiable-GLCM with Directional Aggregation: Improving Rotational Invariance and Feature Stability For Radiomics Features*

Haoyue Chen[1][0009-0003-8176-9021], Rosario Corso[2][0000-0001-9123-4977], Albert Comelli[3][0000-0002-9290-6103], and Anthony Yezzi[1][0000-0002-3771-4889]

[1] School of Electrical and Computer Engineering,
Georgia Institute of Technology, Atlanta, GA, USA
{hchen765, ay10}@gatech.edu
[2] Dipartimento di Matematica e Informatica,
Università degli Studi di Palermo, Palermo, Italy
rosario.corso02@unipa.it
[3] Ri.MED Foundation, Via Bandiera 11,
90133 Palermo, Italy
acomelli@fondazionerimed.com

Abstract. In our prior work, we introduced Differentiable-GLCM, a continuous-space generalized framework that computes co-occurrence via radius-weighted neighborhoods rather than discrete pixel offsets. While this approach yields smooth, differentiable matrices with improved geometric consistency over the classical Haralick GLCM, it remains inherently directionally sensitive.

In this paper, we improve the Differentiable-GLCM framework by aggregating co-occurrence information across directions to enhance rotational robustness. We evaluate three aggregation strategies: (i) discrete directional averaging, (ii) continuous angular integration yielding an isotropic representation, and (iii) a computationally efficient spatial weighting scheme that closely approximates the integrated formulation.

Experiments on rotated 2D planar and 3D volumetric medical images demonstrate marked improvements in stability and rotational consistency compared to both classical GLCM and the original directional Differentiable-GLCM. Measured via Frobenius-norm metrics, these improvements are particularly strong for higher-order texture descriptors. Overall, directionally aggregated Differentiable-GLCM offers a compact, rotation-robust, and optimization-ready texture descriptor that serves as a drop-in replacement for classical GLCM in radiomics, variational segmentation, and neural network pipelines.

Keywords: Differentiable-GLCM · Directional Aggregation · Rotational Robustness.

* Supported by the U.S. Army Research Office (W911NF-22-1-0267; A.Y., H.C.).

1 Introduction

Texture descriptors are central to medical image analysis, providing quantitative measures of tissue heterogeneity for diagnosis, grading, and outcome prediction. Among statistical texture models, the Gray Level Co-occurrence Matrix (GLCM) remains widely used in radiomics due to its ability to capture second-order spatial dependencies. However, classical GLCM relies on discrete offsets, fixed orientations, and hard quantization, which introduces rotation sensitivity and precludes differentiability, limiting robustness and integration with optimization- and learning-based models.

In our prior work, we proposed a Differentiable-GLCM framework in continuous space, where co-occurrence statistics are computed using radius-weighted neighborhoods rather than discrete pixel pairs [4]. This formulation yields smooth, differentiable co-occurrence matrices for prescribed displacement directions, improving geometric consistency and enabling gradient-based optimization. Nevertheless, the construction remains directional: each matrix is tied to a specific displacement direction, and using multiple directions results in a high-dimensional representation that is still orientation-sensitive and computationally costly.

In this paper, we introduce a directionally aggregated extension of this framework and adopt the term Differentiable-GLCM to describe both the original continuous formulation and its aggregated variants. We develop three averaging strategies: (i) discrete directional averaging over a finite set of orientations, (ii) continuous angular integration to obtain a highly isotropic co-occurrence representation, and (iii) a band-based circular approximation that preserves structural accuracy while reducing computational cost. Experiments on rotated 2D/3D medical images demonstrate consistent improvements in rotational consistency for both co-occurrence matrices and derived texture features, particularly for higher-order metrics. The resulting descriptor is compact, rotation-robust, and directly compatible with IBSI-style radiomics features, making it a practical plug-in replacement for classical GLCM.

2 Related Work

2.1 Texture Analysis

Texture analysis plays a foundational role in image processing and pattern recognition, providing quantitative descriptors of spatial arrangements and surface characteristics of intensity variations. Over the past several decades, texture features have been widely used in applications ranging from material inspection and industrial quality control to remote sensing, medical imaging, and motion analysis [3]. Numerous approaches have been proposed—including statistical, structural, model-based, and transform-domain methods—but statistical co-occurrence models remain among the most influential due to their ability to capture spatial dependencies while remaining computationally tractable.

2.2 Radiomics and Texture-Based Biomarkers

Radiomics has emerged as a powerful framework for quantitative medical image analysis, enabling high-throughput extraction of clinically meaningful features from CT, MRI, and PET scans [14, 9]. Radiomic features encompass first-order statistics, geometric shape descriptors, and a wide range of texture metrics—including those derived from GLCM—that reflect tissue microstructure and heterogeneity. When integrated with clinical and molecular data, such features can support predictive modeling for diagnosis, prognosis, and treatment response [13, 16]. A typical radiomics workflow includes image acquisition, segmentation, feature extraction, data management, and model building [12]. Robustness across scanners, protocols, and segmentations remains a central challenge, particularly for texture descriptors that are sensitive to image orientation or resampling.

2.3 Classical GLCM and Directional Aggregation Strategies

The Gray-Level Co-occurrence Matrix (GLCM), introduced by Haralick et al. [11], remains a cornerstone in texture analysis for capturing second-order statistics of gray-level pairings [15]. Despite its wide adoption in segmentation, classification, and radiomics, classical GLCM exhibits several fundamental limitations: it relies on discrete pixel offsets, operates on a small set of fixed angular directions, and lacks geometric continuity. These properties introduce rotational bias and reduce robustness, particularly in isotropic or 3D imaging scenarios.

To address these challenges, numerous studies have investigated multi-orientation or aggregated GLCM formulations. Directionally averaged GLCMs and multi-angle integration have demonstrated improvements in robustness, discriminative power, and noise resilience. For example, a multi-direction GLCM method was applied to SAR image segmentation [6], where texture similarity between superpixels was computed from several orientations, enhancing region merging under complex scattering conditions. Other aggregation strategies include principal-component or quadratic averaging of GLCM entries, which have been shown to improve texture classification accuracy by preserving richer co-occurrence structure than standard Haralick features [1].

Additional work has demonstrated the value of aggregated GLCM features in medical imaging. GLCM descriptors were combined with Local Binary Patterns (LBP) and a nonlinear interaction feature space to enhance the classification of glioma, meningioma, and pituitary tumors [8]. Aggregated and nonlinear GLCM-derived statistics were shown to substantially improve discriminative performance, achieving classification accuracies exceeding 99%. Similarly, Xian [17] used GLCM texture features averaged across four canonical directions to distinguish malignant from benign liver tumors in ultrasound images. Using a fuzzy SVM classifier, the aggregated GLCM-based system achieved an accuracy of 97%, outperforming conventional SVM approaches.

These studies collectively reinforce the importance of integrating texture information across multiple orientations. However, nearly all existing approaches remain inherently discrete, relying on a finite set of angular samples. This discrete sampling limits rotational consistency and precludes smooth, differentiable formulations suitable for optimization-based segmentation or learning frameworks. These limitations motivate the need for continuous and differentiable co-occurrence models that provide highly isotropic texture representations.

2.4 Generalized, Continuous, and Differentiable GLCM Methods

Recent work has explored generalized GLCM formulations to address the non-differentiability and grid dependence of classical counting-based approaches. In our prior work, we introduced a continuous, radius-weighted co-occurrence framework in which discrete pixel offsets are replaced by continuous displacement vectors, yielding smooth and differentiable co-occurrence matrices [4]. This formulation improves geometric consistency and enables integration with variational segmentation and learning-based models.

However, the framework remains directional, producing one co-occurrence matrix per displacement direction. As a result, residual orientation dependence persists, and the collection of directional matrices becomes high-dimensional and inefficient for radiomics pipelines and optimization-based applications. While discrete directional averaging has been considered, continuous angular integration remains largely unexplored—particularly in differentiable settings—motivating the need for a rotation-robust generalization.

3 Motivation

Although the generalized continuous GLCM framework enables differentiability and improved geometric consistency, its directional construction limits robustness and practicality in radiomics and model-based segmentation.

3.1 Rotation Robustness and Stability

Each directional matrix captures texture statistics along a single orientation and therefore remains sensitive to object pose, scanner geometry, and preprocessing-induced rotations. In medical imaging scenarios involving isotropic tissues or variable patient positioning, this sensitivity introduces undesirable variability. Aggregating co-occurrence information across orientations produces a single isotropic representation that suppresses orientation bias and reduces variance, yielding more stable texture estimates under noise, resampling, and small perturbations.

3.2 Efficiency and Radiomics Compatibility

Maintaining a full stack of directional matrices increases storage, memory usage, and differentiation cost, particularly in large 3D volumes or iterative optimization. Collapsing the orientation dimension into a single aggregated matrix substantially reduces computational overhead while preserving descriptive power. Moreover, most radiomics feature definitions standardized by the Image Biomarker Standardisation Initiative (IBSI) are based on directionally averaged GLCMs, making an aggregated formulation a direct plug-in replacement for classical GLCM.

3.3 Integration into Segmentation and Learning Frameworks

Optimization-based methods and neural networks favor compact, differentiable descriptors. A large collection of directional co-occurrence matrices complicates energy formulation and gradient computation, whereas an aggregated representation enables cleaner variational models and more efficient end-to-end learning.

These considerations motivate the development of a directionally aggregated extension of the Differentiable-GLCM framework. The aggregated formulation preserves the benefits of continuity and differentiability while mitigating residual direction dependence, resulting in a rotation-robust, computationally efficient, and model-friendly texture descriptor.

4 Method

4.1 Previous Method

In our previous work, we introduced a continuous and differentiable generalization of the classical Haralick GLCM [4]. Rather than counting gray-level pairs at a single discrete offset, our method computes co-occurrence statistics over an ℓ_2-based displaced neighborhood. Consequently, each matrix entry is formed by a spatially weighted aggregation of surrounding locations consistent with a prescribed displacement magnitude, instead of relying on a single paired pixel. This method has demonstrated strong discriminative power in extracting textural features and improved rotation robustness compared to traditional GLCM [5].

Formally, for a reference pixel at spatial position (x_i, y_j), we consider a circular neighborhood $\mathcal{N}$ of radius r and displace it by a continuous vector $\boldsymbol{\delta} = (\delta x, \delta y)$, yielding a displaced neighborhood $\hat{\mathcal{N}}$. Each point $(\hat{x}, \hat{y}) \in \hat{\mathcal{N}}$ is assigned a spatial weight $w(\hat{d})$ that depends on its Euclidean distance $\hat{d}$ from the displaced center, and the contribution of the neighborhood to the co-occurrence entry $\mathrm{GLCM}(g_1, g_2)$ is obtained by summing the normalized weights over all pixels whose quantized gray levels match (g_1, g_2). This construction replaces hard, single-offset counting with a smooth, radius-weighted aggregation that is naturally differentiable with respect to both intensities and geometry, and that improves rotational robustness by relying on an ℓ_2 ball rather than an axis-aligned window.

4.2 Discrete Averaging

The traditional Gray-Level Co-occurrence Matrix (GLCM) method, often referred to as Haralick's formulation [11], computes co-occurrence statistics along

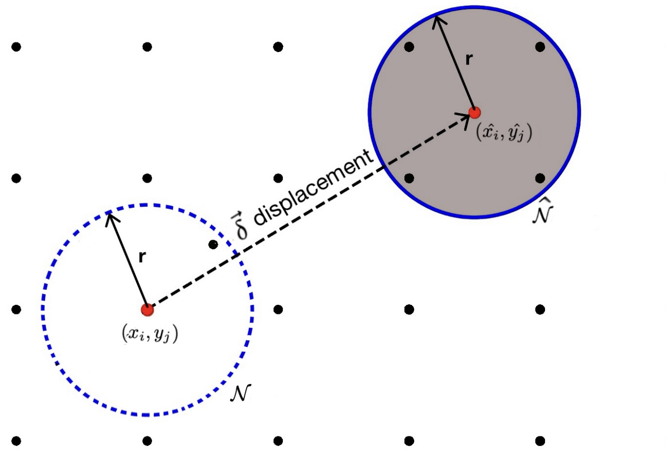

Fig. 1: Schematic illustration of our generalized GLCM construction. For each reference location, a circular neighborhood is displaced by a vector $\boldsymbol{\delta}$ and all pixels in the displaced neighborhood contribute to the co-occurrence statistics with distance-dependent weights.

a set of discrete spatial directions. In 2D, four principal directions are typically considered (0°, 45°, 90°, and 135°), while in 3D this set is commonly extended to 13 unique directions. For each direction and displacement, a GLCM tabulates how often pairs of gray levels co-occur at that relative offset.

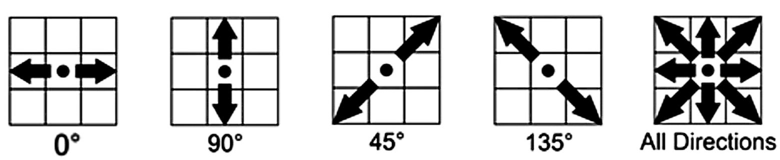

Fig. 2: Schematic illustration of the principal spatial directions used for 2D GLCM computation (0°, 45°, 90°, 135°). Directional GLCMs are first computed separately along each orientation and can then be combined by directional averaging [7].

A common strategy to reduce directional sensitivity is to compute GLCMs along multiple fixed directions and then aggregate them by arithmetic averaging. Let $\text{GLCM}_{\theta_k}(g_1, g_2)$ denote the co-occurrence matrix computed at angle θ_k for gray levels g_1 and g_2, and let N_θ be the number of directions. The directionally averaged GLCM is defined as

$$\mathrm{GLCM}_{\mathrm{avg}}(g_1, g_2) = \frac{1}{N_\theta} \sum_{k=1}^{N_\theta} \mathrm{GLCM}_{\theta_k}(g_1, g_2). \tag{1}$$

This discrete averaging reduces orientation dependence and improves rotational consistency of the derived texture features.

In our experiments, we adopt the same multi-directional scheme for both the classical (Haralick) GLCM and our generalized GLCM formulation: GLCMs are first computed along all available directions, and then their directionally averaged versions are used for a fair and consistent comparison.

4.3 Continuous Averaging

Continuous Averaging with Angular Integration To achieve comprehensive rotational robustness in the continuous co-occurrence framework, we introduce an integral-based approach that averages spatial weight contributions across all possible displacement directions.

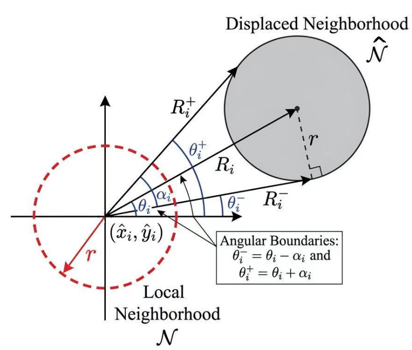

Fig. 3: Angular Integration of Weight Functions

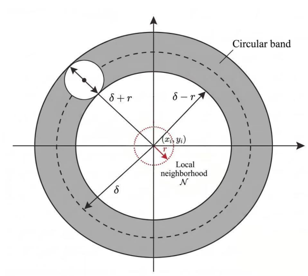

Fig. 4: Simplified Formulation of Weight Functions

For a given reference pixel at (x_i, y_i) and a contributing pixel at $(\hat{x}, \hat{y})$, the localized spatial weight $w(\hat{x}, \hat{y}, \theta)$ for a specific displacement angle θ and magnitude δ decays quadratically within a neighborhood radius r. To generalize this into a highly isotropic representation, we integrate these localized spatial relationships over the full 2π angular domain. The resulting aggregated weight distribution, $W(\hat{x}, \hat{y})$, represents the total accumulated contribution:

$$W(\hat{x}, \hat{y}) = \int_0^{2\pi} w(\hat{x}, \hat{y}, \theta) d\theta. \tag{2}$$

Because the contributing pixel only falls within the active radius r for a restricted subset of angles—defined symmetrically by a maximum angular deviation α_i—this global integral reduces to a bounded domain. Evaluating this bounded integral yields an exact, closed-form analytic solution:

$$W(\hat{x},\hat{y}) = 2\arccos\left(\frac{\rho_i^2+\delta^2-r^2}{2\delta\rho_i}\right)\left(1-\frac{\rho_i^2+\delta^2}{r^2}\right) + \frac{4\delta\rho_i}{r^2}\sqrt{1-\left(\frac{\rho_i^2+\delta^2-r^2}{2\delta\rho_i}\right)^2}, \tag{3}$$

where $\rho_i = ||(x_i, y_i) - (\hat{x}, \hat{y})||$ denotes the inter-pixel distance. If the argument of the arccosine falls outside the valid domain $[-1, 1]$, the pixel lies entirely outside the reachable annular region for any angle θ, and $W(\hat{x}, \hat{y})$ is strictly set to zero.

A comprehensive geometric derivation of the active integration domain, the angular limits, and the final analytic form is provided in Appendix A.

Continuous Averaging with Simplified Formulation of Weight Functions To achieve a computationally efficient continuous formulation, we replace the exact angular integral with a radial weighting scheme evaluated over an annular region (or "band"). As illustrated in Fig. 4, for each reference pixel (x_i, y_i) within the region of interest (ROI), this band encompasses all neighboring pixels whose radial distance approximates the prescribed displacement magnitude δ. Specifically, given a neighborhood radius r, the band spans the radial interval $[\delta - r, \delta + r]$. For each pixel $(\hat{x}, \hat{y})$ within this region, we assign a continuous spatial weight based on its radial deviation from the ideal displacement, defined as:

$$w = 1 - \frac{\left(\sqrt{(\hat{x}-x_i)^2+(\hat{y}-y_i)^2}-\delta\right)^2}{r^2} \tag{4}$$

This weight function ensures that pixels closer to the displaced center receive higher contributions, while those at the outer edges of the band contribute less. By integrating this weighting scheme across all pixels in the band, we obtain a smoothly averaged version of the GLCM, improving robustness to variations in object shape and texture.

4.4 GLCM Radiomics Feature in 3D

To evaluate the effectiveness of our continuous GLCM framework in three dimensions, we extended the 2D spatial aggregation formulation to the 3D volumetric domain. In this setting, each reference voxel (x_i, y_i, z_i) contributes to the co-occurrence matrix based on a spherical neighborhood displaced by a 3D vector $\boldsymbol{\delta}$ with fixed magnitude δ and direction parameterized over the unit sphere.

Specifically, for each reference voxel (x_i, y_i, z_i) inside the ROI, we identify all voxels $(\hat{x}, \hat{y}, \hat{z})$ such that their Euclidean distance to the reference voxel lies within the band $[\delta - r, \delta + r]$. We then assign a simplified continuous weight to each voxel using the following formula:

$$w = 1 - \frac{\left(\sqrt{(\hat{x} - x_i)^2 + (\hat{y} - y_i)^2 + (\hat{z} - z_i)^2} - \delta\right)^2}{r^2} \tag{5}$$

This function encodes the relative spatial relationship between the reference voxel and its neighbors located at a fixed radial distance, decreasing smoothly as the deviation from the ideal displacement magnitude increases.

We conducted experiments on volumetric medical image data to compare the rotational consistency of GLCMs derived using our continuous 3D volumetric formulation.

5 Results

The averaging methods were evaluated on a diverse cohort of clinical CT cases, including both brain and lung tumor datasets sourced from Radiopaedia and the Ri.MED Foundation [4, 2]. To systematically assess rotational consistency, the volumes were subjected to synthetic spatial rotations in 45° increments utilizing linear interpolation. The visualizations and quantitative metrics presented herein were derived from a representative brain CT ROI to ensure consistent geometric comparison across all aggregation strategies.

5.1 Discrete Averaging

Here we report the relative error in Fig. 5, showing that the higher-order feature metrics become considerably more stable under directional averaging than in the single-direction case.

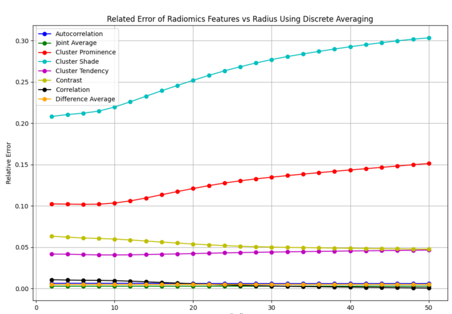

Fig. 5: Relative Error Between Original and Rotated Image at $\|\delta\|_2 = 50$ (Discrete)

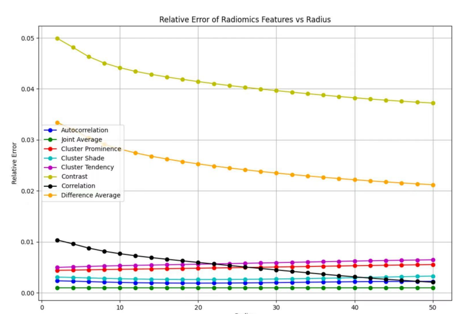

Fig. 6: Relative Error Between Original and Rotated Image at $\|\delta\|_2 = 50$ (3D)

5.2 Continuous Averaging

By applying the integrated weights in Eq. 3 and the simplified weights in Eq. 4, the resulting Frobenius norms between the GLCM matrices of the original and rotated images, together with their relative errors, are shown in Fig. 8.

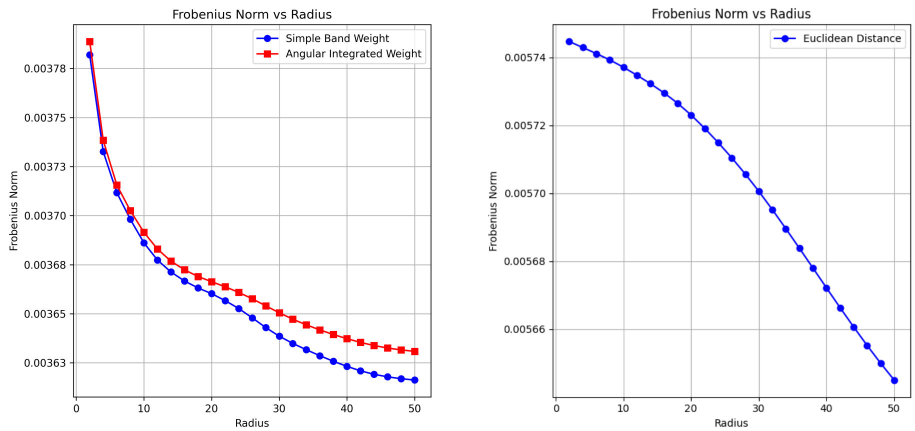

(a) Frobenius norm between original and rotated GLCMs, obtained using integrated and simplified weights.

(b) Radiomics Features vs Radius (Original) in 3D.

Fig. 7: GLCM-based stability evaluation versus neighborhood radius at displacement $\|\delta\|_2 = 50$. (a) Frobenius norm stability metric comparing integrated versus simplified weighting. (b) 3D Radiomics Features for the original image configuration.

As shown in Fig. 7a, the curves obtained with integrated and simplified weights are nearly indistinguishable, indicating that the simplified weighting scheme closely approximates the full integrated formulation. To quantify this agreement more explicitly, we also report the relative error between the two methods in the following analysis:

The resulting GLCMs preserve key structural properties while eliminating directional bias. Although the integrated-weights formulation is mathematically more accurate, our experiments show that using simplified weights does not materially affect the final texture representation. Frobenius-norm comparisons between GLCMs computed with both schemes exhibit only minor discrepancies, indicating that the simplified approach closely approximates the integrated formulation. Consequently, in practical settings where computational efficiency is critical, the simplified weighting can be used directly without significantly compromising texture characterization, while the integrated scheme remains available when maximal fidelity is required.

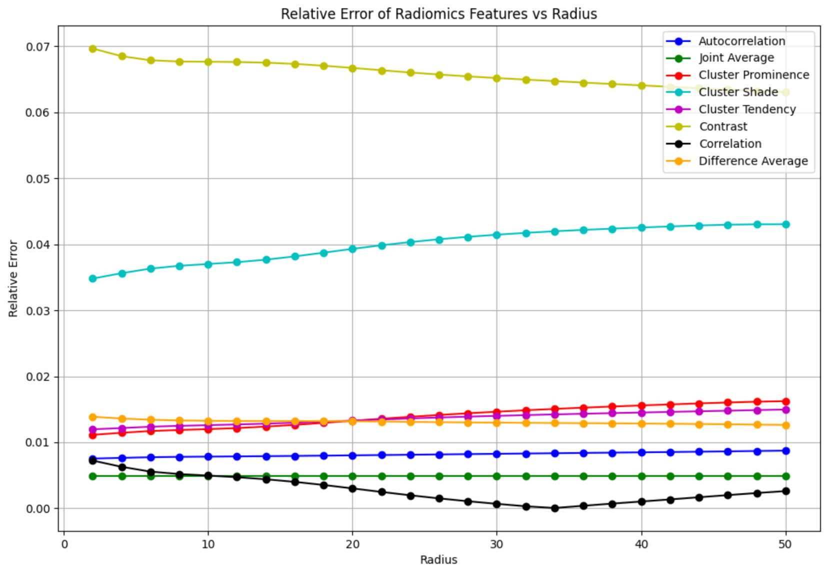

(a) Integrated Weights

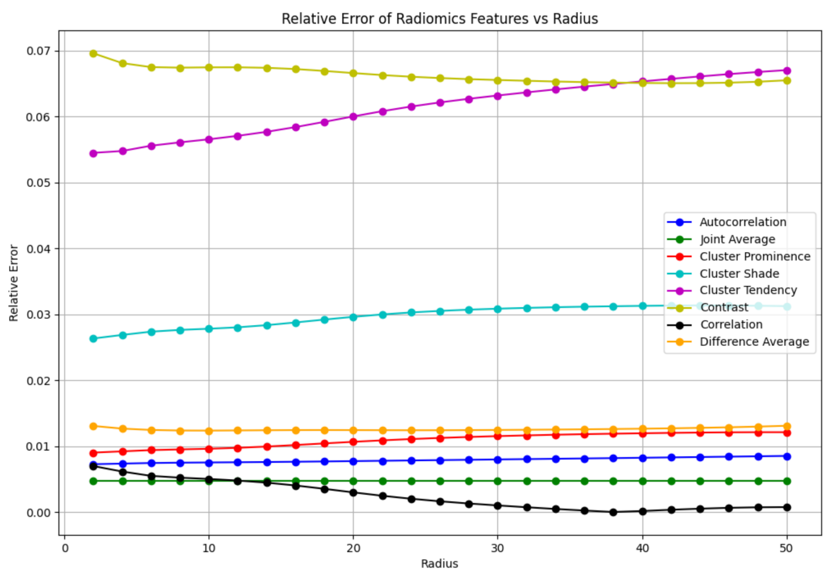

(b) Simplified Weights

Fig. 8: Relative Error Between Original and Rotated Images at $\|\delta\|_2 = 50$ Using Integrated and Simplified Weights

5.3 Results in 3D Volumetric Data

To assess the rotational invariance of the proposed method in a volumetric context, we rotate the input volume and compute the Frobenius norm (Euclidean distance) between the GLCMs derived from the original and rotated 3D data.

For this evaluation, we select a region of interest (ROI) bounded between coordinates $(100, 100, 10)$ and $(200, 200, 20)$. This ROI spans eleven slices along the z-axis, capturing a substantial three-dimensional structure. It is important to utilize a sufficiently deep volume for this assessment, as evaluating on very thin slices would nearly reduce the analysis to a purely 2D case. By comparing GLCMs computed on the original and rotated volumes in this larger domain, we aim to assess the method's ability to maintain rotational invariance when the data exhibit increased depth and more complex spatial variations.

For quantization, we set the number of gray levels to 64, a choice that has been previously tested and shown to offer the best trade-off between texture representation accuracy and computational stability across rotations. We evaluate the method using a displacement magnitude of $\|\delta\|_2 = 50$ in Fig. 7b.

The results demonstrate that the Frobenius norms of the evaluated cases align with our expectations, indicating that the proposed method performs effectively in 3D settings.

To further validate this result, we present a comparison of radiomics features for the case with displacement magnitude $\|\delta\|_2 = 50$ in Fig. 6.

6 Conclusion and Future Work

In this work, we enhance the Differentiable-GLCM framework by introducing directionally aggregated strategies—discrete averaging, continuous angular integration, and a simplified spatial weighting scheme—to achieve highly robust rotational consistency. Across all 2D and 3D experiments, these methods consistently improved the stability of co-occurrence statistics, particularly for higher-order texture descriptors. Relative error analyses demonstrate that discrete and continuous aggregations converge toward a stable, highly isotropic representation, while the simplified spatial weighting provides a highly efficient approximation without significant loss of structural fidelity.

While this study establishes the theoretical foundation and statistical robustness of the directionally aggregated Differentiable-GLCM, we acknowledge certain limitations in the current evaluation scope. A primary direction for future work is systematically benchmarking the computational cost of our continuous formulations against classical GLCM and other rotation-invariant descriptors. Furthermore, we will transition from intrinsic stability metrics to comprehensive downstream task evaluations. Because the inherent differentiability of our formulation is maintained with respect to both image intensities and region geometry, it is highly suitable for variational segmentation. This will enable the derivation of texture-driven shape gradients for active contours and 3D surfaces, allowing segmentation to be guided by local second-order structures rather than solely by intensity contrast.

To facilitate broad clinical translation and downstream classification tasks, the proposed aggregated framework will be natively integrated into the open-source PyRadiomics ecosystem [10] as a first-class feature family. This open-source extension will support robust 2D and 3D computation, including utilities for mapping physical displacements to voxel spaces to handle anisotropic resolution. Overall, directionally aggregated Differentiable-GLCM provides a rotation-robust, differentiable co-occurrence model that bridges classical texture analysis with modern optimization and learning-based methods, establishing a flexible foundation for future methodological advances and clinical applications.

References

1. Beliakov, G., James, S., Troiano, L.: Texture recognition by using glcm and various aggregation functions. In: 2008 IEEE International Conference on Fuzzy Systems (IEEE World Congress on Computational Intelligence). pp. 1472–1476 (2008). https://doi.org/10.1109/FUZZY.2008.4630566
2. Bickle, I.: Normal ct brain - pediatric. Radiopaedia.org (2017). https://doi.org/10.53347/rID-53337, `https://radiopaedia.org/cases/53337`, case study, Radiopaedia.org (Accessed on 21 Jan 2025)
3. Castellano, G., Bonilha, L., Li, L., Cendes, F.: Texture analysis of medical images. Clinical radiology **59**(12), 1061–1069 (2004)
4. Chen, H., Yezzi, A., Corso, R., Comelli, A.: A generalized gray level co-occurrence matrix for rotation-invariant texture detection in radiomics. In: Rodolà, E., Galasso, F., Masi, I. (eds.) Image Analysis and Processing - ICIAP 2025 Workshops. pp. 187–198. Springer Nature Switzerland, Cham (2026)
5. Chen, H., Yezzi, A., Corso, R., Comelli, A.: Rotation-robust radiomics: Differentiable glcm. In: 2026 IEEE Southwest Symposium on Image Analysis and Interpretation (SSIAI). pp. 37–40. IEEE, Santa Fe, NM, USA (March 2026)
6. Cui, M., Huang, Y., Wang, R., Pei, J., Huo, W., Zhang, Y., Yang, H.: A superpixel aggregation method based on multi-direction gray level co-occurrence matrix for sar image segmentation. In: Proceedings of the 2021 IEEE International Geoscience and Remote Sensing Symposium (IGARSS). pp. 3487–3490 (2021). https://doi.org/10.1109/IGARSS47720.2021.9554951
7. dGB Earth Sciences: OpendTect 6.6 Documentation: Appendix A – Attributes and Filters, Section 11.28: Texture – Directional. `https://doc.opendtect.org/6.6.0/doc/od_userdoc/content/appendix_a-attributes_and_filters.htm`, accessed: Feb. 7, 2025
8. Dheepak, G., Christaline, A.J., Vaishali, D.: Brain tumor classification: A novel approach integrating GLCM, LBP and composite features. Frontiers in Oncology **13**, 1248452 (2024). https://doi.org/10.3389/fonc.2023.1248452, `https://doi.org/10.3389/fonc.2023.1248452`
9. Gillies, R.J., Kinahan, P.E., Hricak, H.: Radiomics: Images are more than pictures, they are data. Radiology **278**(2), 563–577 (2016). https://doi.org/10.1148/radiol.2015151169, epub 2015 Nov 18
10. van Griethuysen, J.J.M., Fedorov, A., Parmar, C., Hosny, A., Aucoin, N., Narayan, V., Beets-Tan, R.G.H., Fillon-Robin, J.C., Pieper, S., Aerts, H.J.W.L.: Computational radiomics system to decode the radiographic phenotype. Cancer Research **77**(21), e104–e107 (2017). https://doi.org/10.1158/0008-5472.CAN-17-0339, `https://doi.org/10.1158/0008-5472.CAN-17-0339`

11. Haralick, R.M., Shanmugam, K., Dinstein, I.: Textural features for image classification. IEEE Transactions on Systems, Man, and Cybernetics **SMC-3**(6), 610–621 (1973). https://doi.org/10.1109/TSMC.1973.4309314
12. Kumar, V., Gu, Y., Basu, S., Berglund, A., Eschrich, S.A., Schabath, M.B., Forster, K., Aerts, H.J., Dekker, A., Fenstermacher, D., Goldgof, D.B., Hall, L.O., Lambin, P., Balagurunathan, Y., Gatenby, R.A., Gillies, R.J.: Radiomics: the process and the challenges. Magnetic Resonance Imaging **30**(9), 1234–1248 (2012). https://doi.org/https://doi.org/10.1016/j.mri.2012.06.010, https://www.sciencedirect.com/science/article/pii/S0730725X12002202, quantitative Imaging in Cancer
13. Lambin, P., Leijenaar, R.T., Deist, T.M., Peerlings, J., de Jong, E.E., van Timmeren, J., Sanduleanu, S., Larue, R.T., Even, A.J., Jochems, A., van Wijk, Y., Woodruff, H., van Soest, J., Lustberg, T., Roelofs, E., van Elmpt, W., Dekker, A., Mottaghy, F.M., Wildberger, J.E., Walsh, S.: Radiomics: the bridge between medical imaging and personalized medicine. Nature Reviews Clinical Oncology **14**, 749–762 (2017). https://doi.org/10.1038/nrclinonc.2017.141
14. Mayerhoefer, M.E., Materka, A., Langs, G., Häggström, I., Szczypiński, P., Gibbs, P., Cook, G.: Introduction to radiomics. Journal of Nuclear Medicine **61**(4), 488–495 (2020). https://doi.org/10.2967/jnumed.118.222893
15. Mohanaiah, P., Sathyanarayana, P., GuruKumar, L.: Image texture feature extraction using glcm approach. International journal of scientific and research publications **3**(5), 1–5 (2013)
16. Rizzo, S., Botta, F., Raimondi, S., Origgi, D., Fanciullo, C., Morganti, A.G., Bellomi, M.: Radiomics: the facts and the challenges of image analysis. European Radiology Experimental **2**(1), 36 (2018). https://doi.org/10.1186/s41747-018-0068-z
17. ming Xian, G.: An identification method of malignant and benign liver tumors from ultrasonography based on glcm texture features and fuzzy svm. Expert Systems with Applications **37**(10), 6737–6741 (2010). https://doi.org/https://doi.org/10.1016/j.eswa.2010.02.067, https://www.sciencedirect.com/science/article/pii/S0957417410001065

A Derivation of the Continuous Angular Integral

This appendix provides the detailed geometric derivation for the continuous angular integration introduced in Section 4.3. For a reference pixel (x_i, y_i), a displaced center is parameterized by a displacement magnitude δ and orientation θ:

$$(\hat{x}_i, \hat{y}_i) = (x_i, y_i) + \delta(\cos\theta, \sin\theta). \tag{6}$$

The spatial influence on a contributing pixel $(\hat{x}, \hat{y})$ is governed by a quadratically decaying weight function capped at a neighborhood radius r:

$$w(\hat{x}, \hat{y}, \theta) = 1 - \frac{\min\left(r^2, ||(x_i, y_i) + \delta(\cos\theta, \sin\theta) - (\hat{x}, \hat{y})||^2\right)}{r^2}. \tag{7}$$

To achieve an isotropic representation, we integrate this weight over the full 2π domain:

$$W(\hat{x}, \hat{y}) = \int_0^{2\pi} w(\hat{x}, \hat{y}, \theta) d\theta. \tag{8}$$

Because $(\hat{x}, \hat{y})$ only falls within the active radius r for a restricted subset of angles, the integral reduces to an active angular bound $[\theta_i^-, \theta_i^+]$:

$$W(\hat{x}, \hat{y}) = \int_{\theta_i^-}^{\theta_i^+} \left(1 - \frac{||(x_i, y_i) + \delta(\cos\theta, \sin\theta) - (\hat{x}, \hat{y})||^2}{r^2}\right) d\theta. \tag{9}$$

The limits of integration are defined symmetrically around a principal angle θ_i by a maximum angular deviation α_i, such that $\theta_i^- = \theta_i - \alpha_i$ and $\theta_i^+ = \theta_i + \alpha_i$. Geometrically, α_i represents the boundary where $(\hat{x}, \hat{y})$ intersects the perimeter of the displaced region. Setting the squared distance equal to r^2:

$$||(x_i, y_i) + \delta(\cos\theta_i, \sin\theta_i) - (\hat{x}, \hat{y})||^2 = r^2. \tag{10}$$

Expanding the squared L_2-norm yields:

$$||(x_i, y_i) - (\hat{x}, \hat{y})||^2 + 2\delta(\cos\theta_i, \sin\theta_i) \cdot ((x_i, y_i) - (\hat{x}, \hat{y})) + \delta^2 = r^2. \tag{11}$$

Let $\rho_i = ||(x_i, y_i) - (\hat{x}, \hat{y})||$ denote the inter-pixel distance. Applying the dot product identity $(\cos\theta_i, \sin\theta_i) \cdot ((x_i, y_i) - (\hat{x}, \hat{y})) = -\rho_i \cos\alpha_i$, the expansion simplifies to:

$$\rho_i^2 + \delta^2 - 2\delta\rho_i \cos\alpha_i = r^2. \tag{12}$$

Isolating the angular term provides the explicit solution for the maximal deviation:

$$\alpha_i = \arccos\left(\frac{\rho_i^2 + \delta^2 - r^2}{2\delta\rho_i}\right). \tag{13}$$

Using the substitution $\alpha = \theta - \theta_i$, the continuous weight integral becomes:

$$W(\hat{x}, \hat{y}) = \int_{-\alpha_i}^{\alpha_i} \left(1 - \frac{\rho_i^2 - 2\delta\rho_i \cos\alpha + \delta^2}{r^2}\right) d\alpha. \tag{14}$$

Evaluating the integral by splitting the terms gives:

$$W(\hat{x}, \hat{y}) = 2\alpha_i \left(1 - \frac{\rho_i^2 + \delta^2}{r^2}\right) + \frac{4\delta\rho_i}{r^2} \sin\alpha_i. \tag{15}$$

Using the trigonometric identity $\sin\alpha_i = \sqrt{1 - \cos^2\alpha_i}$, we substitute the previous expression for $\cos\alpha_i$:

$$\sin\alpha_i = \sqrt{1 - \left(\frac{\rho_i^2 + \delta^2 - r^2}{2\delta\rho_i}\right)^2}. \tag{16}$$

Substituting α_i and $\sin\alpha_i$ back into the evaluated integral yields the final analytic formulation:

$$W(\hat{x}, \hat{y}) = 2\arccos\left(\frac{\rho_i^2 + \delta^2 - r^2}{2\delta\rho_i}\right)\left(1 - \frac{\rho_i^2 + \delta^2}{r^2}\right) + \frac{4\delta\rho_i}{r^2}\sqrt{1 - \left(\frac{\rho_i^2 + \delta^2 - r^2}{2\delta\rho_i}\right)^2}. \tag{17}$$