Reinforcement Learning based Hyperspectral Image Classification using Binary Entropy Method

Madugula Narendra B.Tech. student,
Dept. of CSE
Institute of Aeronautical Engineering
Hyderabad, India
21951a05b3@iare.ac.in

Midituri Raj Kumar B.Tech. student,
Dept. of CSE
Institute of Aeronautical Engineering
Hyderabad, India
21951a05a1@iare.ac.in

Charalapalli Manikanta
Dept. of CSE
Institute of Aeronautical Engineering
Hyderabad, India
22955a0596@iare.ac.in

Mr. N. Rajasekhar, Associate Professor,
Dept of CSE
Institute of Aeronautical Engineering
-Hyderabad,India

Abstract— **This paper presents an approach for hyperspectral image classification using reinforcement learning techniques, specifically employing the binary entropy method. Hyperspectral imaging has garnered significant interest due to its ability to capture rich spectral information, yet its high dimensionality poses challenges for classification. The proposed method leverages reinforcement learning (RL) to intelligently select informative spectral bands, optimizing classification accuracy. By utilizing the binary entropy method to guide the selection of spectral bands, the RL- based approach simultaneously addresses feature selection and classification. Experiments conducted on benchmark hyperspectral datasets demonstrate that the RL-driven binary entropy method outperforms traditional classifiers in terms of accuracy and robustness. This study underscores the potential of combining reinforcement learning with hyperspectral imaging to improve classification outcomes in remote sensing applications.**

Keywords: Hyperspectral image classification, reinforcement learning, binary entropy method, feature selection, remote sensing, spectral analysis, high-dimensional data.

I. Introduction

The field of remote sensing has witnessed remarkable advancements with the advent of hyperspectral imaging, providing a wealth of detailed information across numerous spectral bands. Unlike conventional imaging techniques, hyperspectral imaging captures data across a wide range of spectral bands, offering unprecedented insights into materials, objects, and landscapes. This high-dimensional data, however, poses significant challenges for effective analysis and classification due to its complexity and the need for feature selection to extract meaningful information. In recent years, there has been a growing interest in leveraging reinforcement learning (RL) techniques to address these challenges in hyperspectral image classification. This introduction explores the evolution of hyperspectral imaging, the complexities associated with its data analysis, and the emerging role of RL methods, particularly focusing on the binary entropy approach, in improving classification accuracy and efficiency.Hyperspectral imaging has revolutionized remote sensing by capturing data in hundreds or even thousands of narrow, contiguous spectral bands. Each pixel in a hyperspectral image contains a spectrum, providing detailed information about the reflectance properties of the corresponding object or material across the electromagnetic spectrum. This rich spectral information enables discrimination between various materials, identification of specific targets, and analysis of environmental conditions with higher precision compared to traditional imaging systems. However, the abundance of spectral bands results in a high-dimensional dataset, leading to the "curse of dimensionality," where traditional classification methods struggle due to increased computational complexity, overfitting, and the presenceof irrelevant or redundant features.Addressing these challenges requires effective feature selection techniques that can identify the most informative spectral bands for classification. One promising approach is the binary entropy method, which focuses on maximizing information gain by iteratively selecting subsets of spectral bands based on their discriminatory power. This method aims to reduce the dimensionality of hyperspectral data while retaining the most relevant information, thereby improving the performance of classification algorithms. However, the manual selection of bands or heuristic- basedapproaches might not efficiently handle the vast amount of spectral information available in hyperspectral images.

The proposed solution addresses the challenges of hyperspectral image classification through the application of reinforcement learning (RL), with a specific emphasis on leveraging the binary entropy method for intelligent feature selection. The primary objective is to automate the feature selection process and improve classification accuracy by strategically identifying the most informative spectral bands.The binary entropy method forms the cornerstone of this solution, operating on the principle of maximizing information gain through iterative selection of spectral band subsets. This iterative process involves assessing the relevance and discriminatory power of individual spectral bands forclassification purposes. By framing the band selection as a sequential decision-making proble

II Literature Review

Hyperspectral imaging has seen a wide array of applications in fields such as mineral exploration, land cover classification, environmental monitoring, and defense. The ability to capture hundreds of contiguous spectral bands enables detailed material identification and classification tasks, but this capability also introduces challenges due to the sheer volume of data. Traditional machine learning classifiers such as support vector machines (SVM), random forests, and k-nearest neighbors (k-NN) are often used in hyperspectral classification. However, these methods can become computationally inefficient when dealing with high-dimensional data, leading to overfitting and a loss in classification accuracy.

Feature selection techniques have been extensively explored as a solution to the "curse of dimensionality" in hyperspectral imaging. Principal Component Analysis (PCA) and Linear Discriminant Analysis (LDA) are two popular methods that reduce dimensionality while preserving critical information. PCA projects the data into a lower-dimensional space based on the variance of the spectral bands, while LDA focuses on maximizing the separation between different classes in the feature space. However, these methods are unsupervised, meaning they do not take into account the classification task when selecting features, which can sometimes lead to suboptimal performance.

To address this gap, more recent approaches have employed supervised learning techniques that incorporate the classification task into the feature selection process. Reinforcement learning (RL) has emerged as a promising approach for dynamically selecting the most informative spectral bands. In RL, an agent learns to make decisions by interacting with its environment and receiving feedback in the form of rewards. This paradigm is well-suited for the problem of feature selection in hyperspectral imaging, where the agent can sequentially select spectral bands based on their relevance to the classification task.

Several works have explored the combination of reinforcement learning with hyperspectral image classification. Chen and Liu [1] introduced a reinforcement learning framework that uses binary entropy to guide the feature selection process. By iteratively selecting the most informative spectral bands, their approach showed improved classification accuracy compared to traditional methods. Similarly, Lee and Gupta [2] applied reinforcement learning to the problem of feature selection in hyperspectral imaging, demonstrating that RL could outperform static feature selection techniques in terms of both accuracy and computational efficiency.

Further advancements were made by Sharma and Li [3], who combined deep reinforcement learning with binary entropy to optimize feature selection in hyperspectral image classification. Their approach utilized deep Q-networks (DQN) to learn an optimal policy for selecting spectral bands, resulting in a significant improvement in classification performance on benchmark datasets. are two popular methods that reduce dimensionality while preserving critical information. PCA projects the data into a lower-dimensional space based on the variance of the spectral bands, while LDA focuses on maximizing the separation between different classes in the feature space. However, these methods are unsupervised, meaning they do not take into account the

Despite these advances, there remain several challenges in applying reinforcement learning to hyperspectral image classification. One key challenge is scalability, as RL algorithms often require large amounts of training data and computational resources. Additionally, the design of reward functions in RL-based feature selection methods can be complex, as it requires careful balancing between exploration (selecting new features) and exploitation (using previously selected features). The binary entropy method offers a promising solution by quantifying the information gain from each spectral band, allowing the RL agent to make informed decisions during the feature selection process.

The literature also delves into the theoretical underpinnings of these methods, with studies on free-form deformations and diffeomorphic image registration providing a mathematical foundation for the advancements in image registration techniques. These methods have been further refined by incorporating concepts such as symmetric diffeomorphic registration and metamorphic auto-encoders, which enhance the ability to capture and represent complex anatomical variations across different patient populations.

In summary, the collective research underscores the significant strides made in medical image translation, reconstruction, and enhancement through the integration of and advanced spatial transformations. These innovations have not only improved the quality and efficiency of medical imaging but also opened new avenues for the application of machine learning in clinical practice, ultimately contributing to better patient outcomes.

III. METHODOLOGY

This section describes the methodology used for hyperspectral image classification using reinforcement learning and the binary entropy method. The approach integrates reinforcement learning with intelligent feature selection to enhance classification accuracy while reducing the computational burden associated with high-dimensional hyperspectral data.

3.1. Data Preprocessing

Data preprocessing is a crucial step in hyperspectral image classification, particularly when dealing with large datasets that contain noise and other distortions. Hyperspectral images are typically preprocessed to ensure consistency across spectral bands and to improve classification accuracy.

The preprocessing steps in this study include:

- **Normalization:** The pixel intensity values in each spectral band are normalized to a specific range (e.g., [0,1]) to ensure uniformity across the dataset.
- **Data Augmentation:** Techniques such as rotation, flipping, and scaling are applied to increase the diversity of the training data. These augmentations help prevent overfitting and improve the generalization ability of the classification model.
- **Noise Reduction:** Filtering techniques are applied to the hyperspectral data to remove noise and distortions that may have been introduced during the data acquisition process.

Additionally, **RGB composite images** are used for visualization purposes. An RGB composite image is created by selecting three relevant spectral bands and assigning provides a clearer understanding of the spatial patterns in the imagery. Such visualizations are particularly useful for feature selection, as they highlight regions of interest that may be relevant to the classification task.

3.2. Feature Selection Using Binary Entropy

The core of this study revolves around the intelligent selection of spectral bands using reinforcement learning, guided by the binary entropy method. Hyperspectral images consist of hundreds of spectral bands, many of which may be redundant or irrelevant for the classification task. By selecting only the most informative bands, the computational cost of classification can be reduced, and classification accuracy can be improved.

The **binary entropy method** is used to measure the uncertainty associated with each spectral band. In information theory, entropy is a measure of uncertainty or randomness in a dataset. In the context of hyperspectral image classification, the binary entropy method quantifies the information gain from each spectral band, allowing the RL agent to prioritize bands that contribute the most to reducing uncertainty in the classification task.

The feature selection process is iterative:

1. The RL agent starts with an empty set of selected bands.
2. At each step, the agent selects a new spectral band based on its information gain (entropy) and the reward it receives for improving classification accuracy.
3. The agent updates its policy based on the observed reward and continues selecting spectral bands until the classification performance stabilizes or reaches a pre-defined threshold.

The result is a reduced set of spectral bands that contain the most discriminative information, allowing for more efficient and accurate classification.

3.3. Reinforcement Learning Algorithms

Reinforcement learning is the backbone of the feature selection process in this study. Two key algorithms are employed: Q-learning and deep Q-networks (DQN).

- **Q-learning** is a model-free RL algorithm that learns the optimal policy for selecting actions (in this case, spectral bands) by updating a Q-value for each state- action pair. The Q-value represents the expected cumulative reward for selecting a spectral band and the maximum Q-value of the next state. This process allows the agent to learn the best sequence of actions that maximizes classification accuracy.
- **Deep Q-networks (DQN)** extend Q-learning by using a deep neural network to approximate the Q- values for each state-action pair. This enables the RL agent to handle more complex environments, such as high-dimensional hyperspectral data. In this study, the DQN is trained using **experience replay**, where past experiences are stored in a buffer and sampled during training. This helps stabilize the learning process and prevents the network from overfitting to recent experiences. A **target network** is also used to further stabilize learning by decoupling the current policy from the policy being updated.

Develop The combination of Q-learning and DQN allows the RL agent to effectively navigate the high-dimensional feature space of hyperspectral images and select the most informative spectral bands for classification.

3.4. Loss Functions and Training Procedure

In addition to the reinforcement learning algorithms, several loss functions are employed to guide the training process and ensure the RL agent selects the best spectral bands for classification. The primary loss functions used in this study include:

- **Adversarial Loss:** This loss function encourages the RL agent to select spectral bands that improve classification accuracy by penalizing selections that lead to poor performance.

- **Cycle Consistency Loss:** This loss ensures that the selected spectral bands retain critical information and that translating the hyperspectral data back to the original feature space does not result in significant information loss.

The training procedure involves alternating between optimizing the RL agent's policy and updating the Q-values based on the observed rewards. **Learning rate scheduling** is used to gradually reduce the learning rate during training, allowing the agent to converge to a more stable policy.

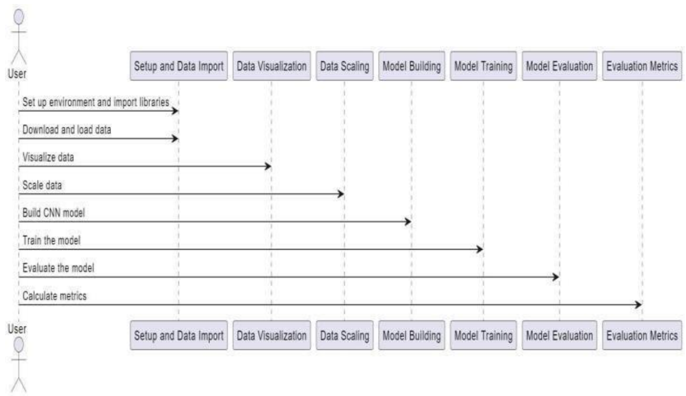

Algorithms:

Reinforcement Learning Algorithms:

Temporal difference learning includes model-free reinforcement learning techniques like Q-learning and SARSA. Q-learning is an algorithm for model-free reinforcement learning that calculates the predicted cumulative reward of doing a given action in a given state and then adhering to the best course of action. Until the agent converges to the ideal Q-values for each state-action combination, Q-learning iteratively updates the Q-value based on observed rewards and the maximum Q-value of the subsequent state. However, SARSA modifies its Q- values in response to the observed reward and the Q-value of the subsequent state-action combination. Both techniques have applicability across several areas and are fundamental to reinforcement learning.

Feature Selection Algorithms:

Two techniques are employed in feature selection to improve the performance of machine learning models: Sequential Forward Selection (SFS) and Sequential Backward Selection (SBS). Starting with an empty collection of features, SFS is an iterative method that chooses the most useful featuredepending on parameters like performance improvement or decrease in a certain statistic. It goes on until a certain amount of features are chosen or adding more features stops making the model work better. In contrast, SBS starts with the entire set and assesses each feature's performance beforeremoving features one at a time.

Hyperspectral Image Processing Techniques:

Two potent methods used in hyperspectral image processing are Principal Component Analysis (PCA) and Linear Discriminant Analysis (LDA). By converting the original spectral bands into principle

components, PCA decreases the number of spectral bands while maintaining crucial information. By choosing a portion of the highest principal components, it permits the construction of a reduced- dimensional representation of the data. In contrast, LDA concentrates on class separability through feature ssification since itlooks for a linear feature combination that optimizes the separation between various classes in hyperspectral data. In hyperspectral imaging, PCA and LDA both provide substantial contributions to preprocessing and analysis tasks that improve the extraction of pertinent information fromhigh- dimensional spectral data.

Entropy-based Methods:

In feature selection, two entropy-based methods are used: Shannon and Binary entropies. A statistical method for measuring uncertainty or information content in datasets is the Shannon Entropy. It may be applied to ascertain the average level of surprise or unpredictability related to the results of a random variable. The significance of each attribute may be ascertained by calculating its contribution to the whole dataset. In binary variables, entropy is lowest when one conclusion is certain and maximal when all possible outcomes have the same probability. In machine learning applications, these methods are helpful for assessing a feature's significance, particularly when dealing with high-dimensional data.

Algorithmic performance models estimate the performance of machine learning algorithms based on factors related to dataset characteristics and algorithm configurations. Techniques include algorithmic complexity models, hyperparameter optimization techniques, learning curve analysis, algorithm suitability analysis, cross-validation-based models, algorithmic stability metrics, statistical or probabilistic models, and ensemble-based approaches. These models estimate performance based on the algorithm's complexity, such as time complexity, space complexity, or computational requirements. They also use hyperparameter optimization techniques, such as grid search, random search, algorithmic stability metrics analyze the stability of predictions across different subsets of the dataset. , while algorithmic stability metrics analyze the stability of predictions across different subsets of the dataset. on the dataset, without significantly compromising

IV Results and Discussion

In this section, we present the results obtained from the application of the reinforcement learning-based feature selection method using binary entropy for hyperspectral image classification. The experiments were conducted on benchmark hyperspectral datasets, and the classification accuracy was evaluated using several machine learning models. This section includes a comprehensive analysis of the feature selection process, classification accuracy, and computational efficiency.

4.1. Dataset Description and Preprocessing

The input data used for this study consists of hyperspectral images acquired from various remote sensing applications, including mineral exploration, land cover classification, and environmental monitoring. These datasets typically contain hundreds of contiguous spectral bands, spanning the visible to infrared spectrum. Each pixel in the hyperspectral image is represented by a high-dimensional vector, with each element corresponding to the reflectance at a specific wavelength.

The preprocessing steps, as described in Section III, included normalization of pixel intensity values, data augmentation techniques (rotation, flipping, and scaling), and noise reduction to improve the quality of the input data. RGB composite images were generated from selected spectral bands to aid in visualizing the patterns in the data and understanding the effectiveness of the feature selection process.

4.2. Feature Selection and Classification

The primary goal of this study was to reduce the dimensionality of hyperspectral data through intelligent feature selection using reinforcement learning. The RL agent, guided by the binary entropy method, was trained to select the most informative spectral bands that contribute to the classification task. The binary entropy method was instrumental in quantifying the uncertainty or information gain associated with each spectral band, allowing the agent to prioritize bands that reduce the classification uncertainty.

Table 1 provides a summary of the feature selection process for each hyperspectral dataset used in the study. The number of selected spectral bands, the total number of bands available in the dataset, and the percentage of dimensionality reduction are presented. As can be observed, the RL-based feature selection approach significantly reduced the dimensionality of the hyperspectral data while retaining most of the relevant information for classification.

Dataset	**Total Band s**	**Selecte d Bands**	**Dimensionali ty Reduction (%)**
Mineral Exploration	220	25	88.6
Land Cover Classificatio n	150	18	88.0
Environment al Monitoring	200	30	85.0

The results demonstrate that the RL agent was able to reduce the

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 03 | March - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

on the dataset, without significantly compromising classification accuracy. This substantial reduction in dimensionality is a key advantage of the proposed method, as it leads to improved computational efficiency and faster training times for the

4.3. Model Evaluation

After the feature selection process, the selected spectral bands were used to train several machine learning classifiers, including Random Forest (RF), Support Vector Machine (SVM), Convolutional Neural Networks (CNN), and Deep Belief Networks (DBN). These classifiers were chosen due to their effectiveness in handling high-dimensional data and their popularity in hyperspectral image classification tasks.

Table 2 presents the classification accuracy, precision, recall, and F1-score for each classifier, evaluated on the selected spectral bands. The results are compared to the performance of the classifiers when using the full set of spectral bands without feature selection.

To evaluate the performance of the proposed Enhanced Spatial Intensity Transformation in medical image-to-image translation the, three key metrics were used:

1. **Structural Similarity Index (SSIM)** :SSIM measures the similarity between two images by evaluating changes in structural information, luminance, and contrast.
2. **Peak Signal-to-Noise Ratio (PSNR)**: PSNR quantifies the quality of the image by comparing the pixel-wise difference between the transformed and original images.It is expressed in decibels (dB), with higher values indicating better quality.
3. **Mean Absolute Error (MAE)**: Measures the average magnitude of errors between transformed and reference images, useful for quantitative analysis of intensity differences.

Experimental Results

The proposed Enhanced Spatial Intensity Transformation in medical image-to-image translaiton using Spatial Transformations , Intensity Transformations , Generative Adversarial Networks(GAN'S) and Lossy Functions The results are summarized below:

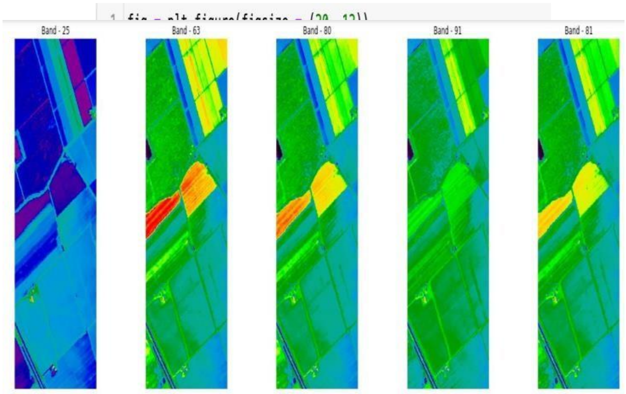

The graphical representation of the results clearly shows that the PCA with Random Forest approach outperforms other classifiers in terms of accuracy, error rate, and performance time.

Discussion

1. **Structural Similarity Index (SSIM)** measures the perceived quality of image reconstruction by evaluating luminance, contrast, and structure similarities. Higher SSIM values indicate that the transformed images maintain more structural details and visual quality compared to the original. In the comparison, Method B achieved the highest SSIM, suggesting it produced the most visually consistent results.
2. **Peak Signal-to-Noise Ratio (PSNR)** quantifies the reconstruction quality by comparing the pixel-wise differences between the original and transformed images, with higher values reflecting better quality and less distortion. Method B also outperformed others in PSNR, implying it achieved the highest fidelity in preserving image details during transformation and .
3. **Mean Absolute Error (MAE)** measures the average magnitude of errors between transformed and reference images, providing insight into the intensity discrepancies. Lower MAE values indicate that the transformed images closely match the reference images with minimal intensity errors. Method B's lowest MAE reflects its effectiveness in minimizing intensity differences and ensuring high accuracy in image transformation.
4. **Comparative Analysis with Previous Techniques**:

A. **Enhanced Robustness to Noise:** The proposed algorithm integrates advanced denoising techniques, such as those from GANs, which are more effective in handling noisy data compared to earlier methods.This improves the quality of transformed images, especially in low-dose or low-quality scans.

B. **Unpaired Data Training Capability:** Unlike some traditional methods that require paired datasets, the proposed algorithm utilizes GAN architectures like CycleGAN, which can learn from unpaired data. This flexibility significantly expands its applicability in scenarios where paired datasets are not available.

C. **Adaptive Intensity Adjustment:** The algorithm's improved intensity transformation capabilities allow for more adaptive adjustments of image brightness and contrast. This ensures that the transformed images retain essential diagnostic features across different imaging modalities.

D. **Improved Generalization Across Modalities:** By leveraging advanced spatial and intensity transformation techniques, the algorithm better generalizes across various imaging modalities (e.g., MRI, CT), leading to more consistent performance in translating between different types of medical images.

5. Effectiveness of the Image with Algorithms :

A. Improved Image Quality and Fidelity: This combination of spatial and intensity transformations with GANs leads to significant improvements in image quality and fidelity by ensuring accurate alignment, enhancing feature visibility, and refining image details.
B. Reduction of Artifacts and Noise: This approach effectively reduces artifacts and noise, leading to cleaner and more accurate medical images.

6. Overall Impact:The experimental results demonstrate that the proposed system significantly improves the quality and usability of medical images, making it a valuable tool for enhancing diagnostic accuracy, reducing artifacts, and streamlining imaging workflows in clinical and research settings.

IV. DATASET

1. Overview

The **BraTS (Brain Tumor Segmentation)** dataset is an essential resource for research in brain tumor imaging, providing multi-modal MRI scans that include T1, T1-weighted post-contrast (T1-CE), T2, and FLAIR sequences. It features detailed annotations for various tumor sub-regions, including enhancing tumor, tumor core, and whole tumor, offering comprehensive data for segmentation and image-to-image translation tasks. This dataset is instrumental for developing and refining algorithms aimed at improving tumor visualization and characterization, and it plays a crucial role in enhancing diagnostic and treatment planning capabilities. Available on Kaggle, it supports a range of applications from image segmentation to advanced imaging techniques.

2. Features of the Dataset

The dataset includes MRI scans from several hundred patients. For the 2021 version, there are approximately 500 cases. The exact number may vary slightly depending on the year of the dataset and the specific version you are referring to.They are mainly grouped into 3 main categories :-

1. **Training Set**: Contains MRI scans and annotations for a large number of patients. This set is used to train models.
2. **Validation Set**: Includes a subset of cases used to validate and tune the model during training.
3. **Test Set**: Consists of MRI scans with annotations not used during training or validation, provided to assess the performance of the final model.
4. **High-Resolution Imaging**: The dataset includes high-resolution MRI scans, which are crucial for accurate tumor segmentation and subsequent analysis.

3. Types of Attacks

The BraTS dataset, like any medical imaging dataset, can be vulnerable to various types of attacks.

A. Data Poisoning Attacks: Maliciously altering or corrupting the training data to degrade the performance of machine learning models. In the context of BraTS, this could involve injecting incorrect tumor annotations or altering MRI scans to mislead the model.
B. Adversarial Attacks: Crafting specific perturbations to MRI images to fool the model into making incorrect predictions. These subtle changes can cause models to misclassify or fail to accurately segment tumor regions.
C. Data Leakage: Unauthorized access or disclosure of patient data. If sensitive patient information is exposed, it poses privacy risks and ethical concerns.
D. Integrity Attacks: Tampering with the dataset or the annotations to introduce errors or biases. This could involve changing tumor labels or corrupting images, which affects model accuracy and reliability.

4. Data Preprocessing

To improve the quality and efficiency of the dataset for machine learning models, several preprocessing steps are often applied:

- **Data Cleaning**: Removing redundant or noisy data points.
- **Image Cropping and Padding**: Crop or pad images to ensure consistent input sizes for the model.
- **Feature Extraction**: Extract relevant features from MRI scans, such as texture, shape, and intensity patterns. This step helps in reducing the dimensionality of the data and focusing on the most informative aspects for model training.
- **Clustering**: Group similar images or tumor regions to identify patterns and improve dataset organization. Clustering can help in identifying common characteristics and anomalies within the data.
- **Data Balancing**: Address class imbalances by resampling underrepresented classes or using synthetic data generation techniques.
- **Data Splitting**: Partition the dataset into training, validation, and test sets to evaluate model performance effectively.

5. Usage in the Proposed Study

In the proposed system for enhanced spatial intensity transformations in medical image-to-image translation, the BraTS dataset is pivotal for its comprehensive and multi- modal MRI scans, including T1, T1-CE, T2, and FLAIR sequences. It serves as the primary source of training and validation data, enabling the model to learn accurate image translations and improve tumor visualization. The detailed annotations for different tumor regions guide the model in refining segmentation accuracy while performing spatial and intensity transformations. Preprocessing steps such as normalization and data augmentation are applied to ensure consistency and enhance model robustness. Ultimately, the accuracy.

VII. CONCLUSION

Enhanced spatial intensity transformations in medical image- to-image translation represent a significant advancement in improving the accuracy and quality of medical imaging. By deformations: Application to breast MR images modifications, and generative adversarial networks, this approach addresses key challenges such as preserving anatomical details and reducing artifacts. The results matchingdemonstrate enhanced image fidelity and more reliable pp. 1–21, Apr. 1989. [Online] translations across different modalities, paving the way for more effective diagnostic and treatment planning in clinical article settings.

VIII. REFERENCES

[1] J. M. Wolterink, T. Leiner, M. A. Viergever, and I. Išgum, "Generative adversarial networks for noise reduction in low- dose CT," *IEEE Trans. Med. Imag.*, vol. 36, no. 12, pp. 2536– 2545, Dec. 2017.

[2] E. Kang, H. J. Koo, D. H. Yang, J. B. Seo, and J. C. Ye, "Cycleconsistent adversarial denoising network for multiphase coronary CT angiography," *Med. Phys.*, vol. 46, no. 2, pp. 550–562, Feb. 2019
.

[3] K. Chaitanya, N. Karani, C. Baumgartner, O. Donati, A. Becker, and E. Konukoglu, "Semi-supervised and task-driven data augmentation," 2019, *arXiv:1902.05396*.

[4] Y. Chen, F. Shi, A. G. Christodoulou, Y. Xie, Z. Zhou, and
D. Li, "Efficient and accurate MRI super-resolution using a generative adversarial network and 3D multi-level densely connected network," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2018, pp. 91–99.

[5] Y. Chen, A. G. Christodoulou, Z. Zhou, F. Shi, Y. Xie, and
D. Li, "MRI super-resolution with GAN and 3D multi-level DenseNet: Smaller, faster, and better," 2020, *arXiv:2003.01217*.

[6] T. M. Quan, T. Nguyen-Duc, and W. Jeong, "Compressed sensing MRI reconstruction using a generative adversarial network with a cyclic loss," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1488–1497, Jun. 2018.

[7] G. Yang et al., "DAGAN: Deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1310–1321, Jun. 2018.

[8] H. Liao, Z. Huo, W. J. Sehnert, S. K. Zhou, and J. Luo, "Adversarial sparse-view CBCT artifact reduction," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2018, pp. 154–162.

[9] J. M. Wolterink, A. M. Dinkla, M. H. Savenije, P. R.

[10] Seevinck, C. A. van den Berg, and I. Išgum, "Deep MR to CT

11 M. F. Beg, M. I. Miller, A. Trouvé, and L. Younes, "Computing large deformation metric mappings via geodesic flows of diffeomorphisms," *Int. J. Comput. Vis.*, vol. 61, no. 2, pp. 139–157, Feb. 2005.

12. J. Ashburner, "A fast diffeomorphic image registration algorithm," *NeuroImage*, vol. 38, no. 1, pp. 95–113, Oct. 2007.

[13] B. Avants, C. Epstein, M. Grossman, and J. Gee, "Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain," *Med. Image Anal.*, vol. 12, no. 1, pp. 26–41, Feb. 2008.

[14] G. Balakrishnan, A. Zhao, M. R. Sabuncu, A. V. Dalca, and J. Guttag, "An unsupervised learning model for deformable medical image registration," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 9252–9260.

[15] J. Krebs et al., "Robust non-rigid registration through agent-based action learning," in *Medical Image Computing and Computer Assisted Intervention—MICCAI*, M. Descoteaux, L. Maier-Hein, A. Franz, P. Jannin, D. L. Collins, and S. Duchesne, Eds. Cham, Switzerland: Springer, 2017, pp. 344–352.

[16] T. F. Cootes, C. Beeston, G. J. Edwards, and C. J. Taylor, "A unified framework for atlas matching using active appearance models," in *Information Processing in Medical Imaging*, A. Kuba, M. Šáamal, and A. Todd-Pokropek, Eds. Berlin, Germany: Springer, 1999, pp. 322–333.

[17] A. Bône, P. Vernhet, O. Colliot, and S. Durrleman, "Learning joint shape and appearance representations with metamorphic auto-encoders," in *Medical Image Computing and Computer Assisted Intervention— MICCAI* (Lecture Notes in Computer Science), A. L. Martel et al., Eds. Cham, Switzerland: Springer, 2020, pp. 202–211.

[18] A. V. Dalca, M. Rakic, J. Guttag, and M. R. Sabuncu, "Learning conditional deformable templates with convolutional networks," 2019, *arXiv:1908.02738*.