

MANIFOLD LEARNING FOR HYPERSPECTRAL IMAGES

Fethi Harkat^{* †}, Guillaume Gey^{*}, Valérie Perrier[†], Kevin Polisano[†], Tiphaine Debeuret^{*},

^{*}Detection Technology Plc, A Grid, Otakaari 5A, 02150 Espoo, Finland

[†]Univ. Grenoble Alpes, CNRS, Grenoble INP, LJK, 38000 Grenoble, France

ABSTRACT

Traditional feature extraction and projection techniques, such as Principal Component Analysis, struggle to adequately represent X-Ray Transmission (XRT) Multi-Energy (ME) images, limiting the performance of neural networks in decision-making processes. To address this issue, we propose a method that approximates the dataset topology by constructing adjacency graphs using the Uniform Manifold Approximation and Projection. This approach captures non-linear correlations within the data, significantly improving the performance of machine learning algorithms, particularly in processing Hyperspectral Images (HSI) from X-ray transmission spectroscopy. This technique not only preserves the global structure of the data but also enhances feature separability, leading to more accurate and robust classification results.

Index Terms— Hyperspectral imaging, X-ray detection, Topology, Manifold learning, Dimensionality reduction, Deep learning, Feature extraction, Computer Vision.

1. INTRODUCTION

Recent advances in Hyperspectral Images (HSI) analysis have primarily focused on reflection spectroscopy in the visible or near-infrared light domains. However, these methods often underperform in X-ray transmission (XRT) spectroscopy [2] due to the distinct noise characteristics in this energy domain, the low-energy resolution of X-ray detectors, and poor counting statistics. Traditional methods, such as the Principal Component Analysis (PCA) [3], and Non-negative Matrix Factorization (NMF) [4] have shown limitations in characterizing these datasets, which in turn affects the efficiency of neural networks. Over the past decade, manifold learning approaches have gained increasing interest for HSI analysis [5]. These approaches typically aim to project data onto a low-dimensional space to reveal intrinsic structures, reduce noise, and improve classification performance. Studies have demonstrated the effectiveness of manifold-based techniques, such as t-distributed Stochastic Neighbor Embedding (t-SNE) [6]

combined with Convolutional Neural Networks (CNN) [7], in enhancing non-linear dimensionality reduction [8] [9] and classification performance [10, 11]. In this work, we propose to apply the Uniform Manifold Approximation and Projection (UMAP) [12] [13] technique to X-ray HSI due to its ability to more accurately approximate the dataset’s topology, thus better capturing nonlinear correlations and improving the performance of subsequent machine learning algorithms.

Our detector consists of an array of semi-conductor active areas, which we will refer to as “detector pixels” to avoid confusion with image processing terminology. Each detector typically contains 512 detector pixels. The data acquisition process involves using an X-ray source to illuminate an object placed on a moving belt, while a detector captures the X-rays transmitted through the object during its motion (Fig. 1). During the process, both the source and the detector remain static. The resulting images have a shape of $[H, W, C]$ where H represents the time axis, W the detector pixel axis, and C the energy band axis.

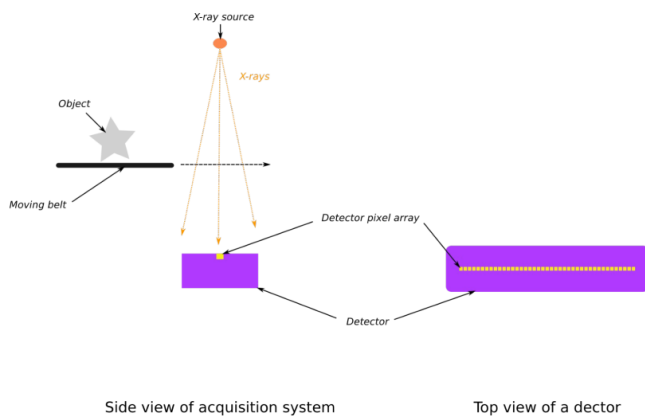


Fig. 1: Scheme of the experimental setting (left) view from aside. A top view of the detector is also given (right) to make the data collection process clearer to the reader.

This work focuses on three main objectives:

- Constructing an expressive representation of spectral data,
- Reducing the dimensionality of this representation,

- Evaluating it using deep learning, while also gaining insights into why this method is effective.

In this paper, we investigate the application of UMAP to XRT HSI data. The proposed methodology details the approaches and tools used in our analysis. In Section 2, we present the UMAP method and its application in our approach. Section 3 provides an overview of our experiments, and finally, in Section 4, we present our results and insights.

2. METHODOLOGY

2.1. UMAP: Uniform Manifold Approximation and Projection

UMAP [12] is a powerful technique for dimensionality reduction and visualization, widely used in data analysis and visualization [14, 15]. It serves as an alternative to t-SNE, with the added advantage of preserving both local and global structures in high-dimensional data. UMAP has gained popularity due to its ability to retain meaningful relationships between data points while remaining computationally efficient.

The method is based on manifold learning, which assumes that data lies on a lower-dimensional manifold embedded on a high-dimensional space. UMAP seeks to find a lower-dimensional representation that preserves the intrinsic structure of data. This is achieved by constructing a fuzzy topological representation of neighborhood relationships. Using fuzzy set theory [16], UMAP estimates the probability that data points are neighbors. The algorithm then minimizes the cross-entropy between pairwise probabilities in the high-dimensional space and their counterparts in the low-dimensional embedding, ensuring that the global structure is effectively preserved.

In our experiments, we used the Parametric UMAP [17] which employs a Neural Network (NN) to learn a mapping from the original data to its embedding while optimizing the same objective as traditional UMAP. This approach enables fast embedding of new data, as the trained network directly maps inputs without requiring re-optimization, making it particularly suitable for large or dynamic datasets.

2.2. Our approach

Our approach (Fig. 2) consists of the following steps:

1. Fit the UMAP model with images of size $[H, W, C]$ where H , W and C are respectively the width, height and number of bands of raw images, with the target dimension D , where $D \leq C$.
2. Project all images using the UMAP model onto the learnt lower-dimensional space of size D , and reshape them to their new shape $[H, W, D]$.

3. Feed the projected images into a specific CNN model to perform various machine learning tasks such as segmentation or regression.




Fig. 2: Pipeline of our proposed method.

3. NUMERICAL EXPERIMENTS

To evaluate the effectiveness of our proposed methodology, we assess its performance across three different datasets and tasks. Through these experiments, we aim to demonstrate the adaptability and efficiency of our approach for both segmentation and regression challenges in hyperspectral imaging (HSI). Below, we describe each dataset’s characteristics and the respective machine learning task.

3.1. Segmentation of Cigarettes

The Cigarettes dataset was created internally and consists of HSI images of random luggage, which may contain either a pack or a carton of cigarettes (positive samples) or none (negative samples). The objective of this experiment is to segment cigarettes in the images when they are present. Binary masks of cigarettes inside the luggage have been carefully annotated for each relevant image.



Fig. 3: Comparison between raw spectral bands (top row) and UMAP-projected bands (bottom row).

We aimed to expand our training set without the need for complex and costly image acquisition procedures. Following the Beer-Lambert law [18], which models photon attenuation in a medium, we generated synthetic training images by separately capturing luggage and cigarette images, applying white normalization, and combining them. Using this approach, we built a dataset of 320 training images and 30 synthetic test images see Fig 4. In addition, 21 real test images were acquired with manually placed cigarette cartons inside luggage to evaluate the performance of our pipeline on real data. Results for both test sets are presented in Section 4.




Fig. 4: Illustration of the fusion process: normalized luggage and cigarette images are combined to generate synthetic data.

To reduce the computational cost, the input images were cropped resized from $[16, 1000, 512]$ to $[16, 152, 128]$ using 2D average pooling. The trained Parametric UMAP then projected the downsized images onto a subspace of dimension $D = 5$ before feeding them into the CNN. A sample of these bands is shown in Figure 3.

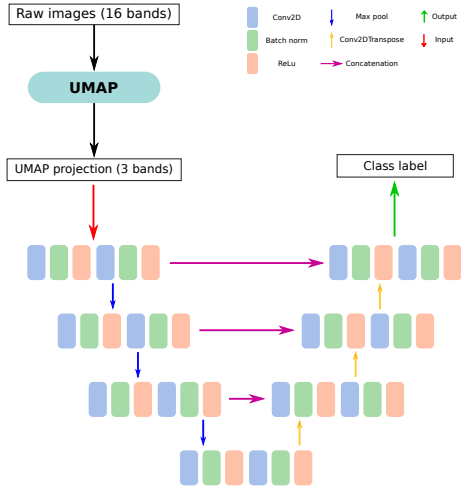


Fig. 5: The U-Net architecture consists of four downsampling blocks, each with a Double Convolution followed by Max Pooling (stride = 2), progressively increasing the number of filters from 8 to 64. The upsampling blocks include an Up-sampling (scale = 2), concatenation with the corresponding output, and a Double Convolution to refine features. The final convolution reduces the output to a single prediction channel.

For the segmentation task, a U-Net model [19] is used with Binary Cross Entropy as the loss function. The symmetric architecture (see Fig. 5) of U-Net includes an encoder, capturing context by compressing the input into feature maps, and a decoder, reconstructing the spatial dimensions to produce the segmented output. UMAP hyperparameters were left to their default values.

3.2. Chemical composition of Stones

The Stones dataset consists of individual stone images from an industrial mine, each composed of two main chemical components. Each image includes a binary mask, mass (grams), and concentration percentages of components A and B , with a third component C accounting for residual elements (see Fig. 6). The goal is to predict mass and concentrations of A and B from the HSI and binary mask. C can be derived from A and B and always has the lowest concentration, therefore it is not a relevant prediction target. As B concentration dominates, standard error averaging is unsuitable for performance evaluation, necessitating a more balanced metric. To ensure a better estimation of performances over all the targets, we therefore define the *Special Harmonic* (SH) score:

$$\text{SH-score} = M \frac{\prod_{i=1}^M S_i}{\sum_{i=1}^M S_i} \quad (1)$$

where M is the number of regression targets, and S_i the associated individual score. The SH-score as well as the S_i are bounded to $[0, 1]$, where 1 is a perfect result. S_i are defined as follows:

$$S_i = \exp\left(-\frac{\sum_{j=1}^N |p_{ij} - t_{ij}|}{\sum_{j=1}^N t_{ij}}\right), \quad (2)$$

where p_{ij} and t_{ij} respectively stands for the prediction and the target value associated to i -th regression target for j -th sample, N being the total number of samples. The rationale behind this formulation of individual score S_i is to collapse quickly in case of poor performances, thus giving a solid estimation of performance robustness through the SH-score. We decided to name this score as such because of its close resemblance with the *Harmonic Mean* stripped from its cross-products in the denominator. To avoid any unnecessary complexity in this work, we decided to use a simple sum of weighted squared errors as the minimization objective:

$$\text{Loss} = \sum_{i=1}^M w_i \left(\frac{p_i - t_i}{t_i}\right)^2 \quad (3)$$

where p_i , t_i and w_i respectively stands for the prediction, target and weight values associated to i -th regression target. The w_i have been empirically set to $\{0.5, 5, 5\}$, corresponding respectively to the mass, the material A and B targets.

The Stones dataset consists of 1,453 training images and 514 test images. To reduce computational load, 4% of the dataset is used for Parametric UMAP training, selecting a balanced subset based on mass and concentrations. The UMAP hyperparameters are set to a minimum distance of 0.5 and 50 neighbors, the configuration found to work best empirically. The minimum distance controls point spacing, while the number of neighbors balances local and global structures.

We tested UMAP projections with 10, 20, and 30 bands, feeding the results into a CNN encoder of four Inception [20]




Fig. 6: Labels of Stones: each stone has an ID along with information about its total mass and material concentrations. The material *C* is residual elements and always has a very low concentration.




Fig. 7: Customized filters used for the regression task.

blocks (Figure 7) followed by a task-specific decision layer. All 514 test images were projected using the fitted UMAP models, and results were compared to the CNN trained on 64 raw bands, 20-band PCA, and 20-band NMF.

3.3. Indian Pines dataset

The Indian Pines dataset [21], captured by the AVIRIS sensor in 1992 over Indiana, consists of a 145×145 pixel image with 220 spectral bands. It includes 16 land-cover classes, with labeled ground truth for part of the image, and is widely used to evaluate hyperspectral classification and feature extraction methods due to its high dimensionality and complexity.

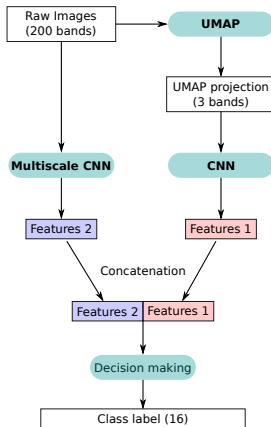


Fig. 8: Multiscale CNN model with two parallel branches and their respective inputs, the UMAP projection and the raw image.

For this experiment, we used a multiscale CNN model

[22] with two branches and two inputs (see Fig. 8). Multi-scale feature fusion help leveraging diverse spatial structures and rich texture features through extensive neighborhood associations. As before, data were projected onto a low dimension representation using UMAP, then fed into the CNN architecture. Results were compared to those obtained using t-SNE and PCA as data reduction techniques. In all cases, the target dimension was set to $D = 3$. UMAP hyperparameters were left to their default values for this experiment as well.

4. RESULTS

4.1. Cigarettes dataset

As mentioned in Section 3.1, we evaluated our approach on both real and synthetic datasets using IoU and Dice Score as segmentation metrics [23]. Each method was run 10 times to compute error bars, and the averaged results are summarized in Table 1.

Dataset	Method	Bands	Mean IoU	IoU Err	Mean Dice	Dice Err
Synthetic	Raw Data	16	0.766	0.032	0.865	0.005
	NMF	5	0.214	0.012	0.330	0.009
	PCA	5	0.417	0.012	0.572	0.005
	PCA	8	0.433	0.010	0.598	0.007
	UMAP	5	0.833	0.016	0.901	0.006
Real	Raw Data	16	0.230	0.020	0.370	0.010
	NMF	5	0.200	0.030	0.330	0.010
	PCA	5	0.180	0.030	0.300	0.010
	UMAP	5	0.320	0.040	0.490	0.010

Table 1: Comparison of segmentation performance across synthetic and real datasets for different dimensionality reduction methods.

Table 1 shows that our UMAP-based approach clearly outperforms PCA, NMF, and even the raw data, demonstrating its ability to capture non-linear spectral structures while removing noise and redundancy. PCA slightly improves over NMF, but both linear methods fall far behind UMAP, confirming that linear projections are insufficient for complex hyperspectral XRT data. Since all methods were evaluated using the same number of bands, the comparison is fair.

UMAP significantly outperforms PCA and NMF on both datasets, even with fewer bands. The lower scores on the real dataset reflect the crude approximations used to compensate for limited data. Despite this gap, the ranking of methods remains consistent, confirming the robustness of UMAP.

PCA, though popular for dimensionality reduction, may underperform in supervised tasks as it preserves variance without considering class separability. This can discard key spectral features and disrupt spatial-spectral patterns, reducing CNN performance and class-discriminative accuracy.

To better understand these limitations, we conducted a component-wise analysis of the PCA output using two complementary criteria: the explained variance ratio and the Mutual Information (MI) score [24]. The explained variance ratio indicates the amount of spectral information retained by

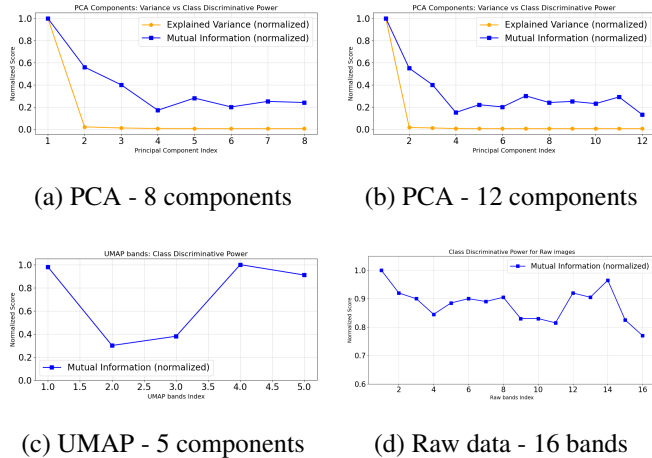


Fig. 9: Comparison of PCA, UMAP, and raw hyperspectral data. The orange color is consistently used for mutual information across all plots.

each component, while the MI score evaluates its relevance for classification by measuring the dependency between each PCA component and the class labels. This combined analysis highlights components that, despite having low variance, contain valuable discriminative information that would be overlooked by a purely variance-based selection.

Figures 9 (a) and (b) show the PCA results with 8 and 12 components, where the x-axis represents component indices and the y-axis two normalized metrics: explained variance and mutual information. Figures 9 (c) and (d) display the mutual information of the five UMAP components and the 16 raw spectral bands, respectively.

The raw data (Figure 9 (d)) contains several highly informative bands, confirming its rich spectral structure. UMAP (Figure 9 (c)) effectively condenses this information into fewer components, preserving class-relevant features through its non-linear mapping. In contrast, PCA prioritizes variance, and some low-variance components exhibit moderate mutual information, showing that variance alone does not guarantee class relevance.

Overall, UMAP and raw data provide stronger and more consistent class-discriminative information, while PCA requires careful component selection for supervised tasks.

4.2. Stones dataset

Method	Bands	Mean SH	Err SH	Mean S_A	Err S_A	Mean S_B	Err S_B
RAW	64	0.598	0.008	0.528	0.006	0.964	0.003
NMF	20	0.596	0.006	0.512	0.011	0.963	0.004
PCA	20	0.617	0.005	0.538	0.006	0.964	0.003
UMAP	10	0.641	0.007	0.559	0.008	0.972	0.002
UMAP	20	0.677	0.003	0.614	0.008	0.974	0.002
UMAP	30	0.636	0.006	0.552	0.005	0.970	0.001

Table 2: Performance metrics (mean and standard error) for different dimensionality reduction methods and band settings.

Table 2 shows that our UMAP-based dimensionality reduction consistently outperforms both traditional methods (PCA, NMF) and raw hyperspectral data. Using only 20 UMAP-derived bands yields the best results across all metrics (Mean SH = 0.677, Mean A = 0.614), surpassing the raw data baseline (0.598 and 0.528, respectively). UMAP with 10 bands underperforms due to information loss, while 30 bands slightly decrease performance, suggesting the reintroduction of noise. PCA and NMF were tested with the same number of bands (20) for a fair comparison. However, both methods fall behind UMAP, especially in Mean SH and Mean A, confirming that linear projections are insufficient for complex hyperspectral XRT data.

4.3. Indian pines dataset

Method	Accuracy (%)
UMAP-CNN	98.80
T-SNE-CNN [11]	97.89
PCA-CNN	96.79
DC-CNN [25]	95.50
DR-CNN [11]	94.93
Multiscale-CNN	87.42
1D-CNN [26]	66.36
SVM [27]	74.81
RF-200 [28]	61.89

Table 3: Indian Pines classification results.

Table 3 shows that UMAP-CNN achieves the best accuracy, surpassing both t-SNE-CNN and PCA-CNN. By capturing non-linear spectral structures that PCA cannot, UMAP provides more discriminative features with the same CNN. It is also more scalable and efficient than t-SNE, which becomes computationally prohibitive on larger datasets.

5. CONCLUSION

UMAP-based approaches consistently outperform traditional linear methods like PCA and NMF in our experiments. This is expected for X-ray transmission spectroscopy, where energy bands exhibit strong non-linear behavior, far more than in optical HSI. By approximating dataset topology through graph construction, our method captures these non-linear correlations, improving model performance. This work serves as a proof of concept, showing that accounting for non-linearities is essential in processing this type of data and motivating further research in topological data analysis for more accurate and robust representations.

6. REFERENCES

- [1] Harkat Fethi, “Manifold learning for hyperspectral images,” <https://gricad-gitlab.univ-grenoble-alpes.fr/polisank/manifold-learning-for-hyperspectral-images>, 2025, Accessed: 2025-03-20.
- [2] Xiangyu Ou, Xue Chen, Xianning Xu, Lili Xie, Xiaofeng Chen, Zhongzhu Hong, Hua Bai, Xiaowang Liu, Qiushui Chen, Lin Li, et al., “Recent development in x-ray imaging technology: Future and challenges,” *Research*, 2021.
- [3] Karl Pearson, “On lines and planes of closest fit to systems of points in space,” *Philosophical Magazine*, vol. 2, no. 11, pp. 559–572, 1901.
- [4] Daniel Lee and H Sebastian Seung, “Algorithms for non-negative matrix factorization,” *Advances in neural information processing systems*, vol. 13, 2000.
- [5] Dalton Lunga, Saurabh Prasad, Melba M. Crawford, and Okan Ersoy, “Manifold-learning-based feature extraction for classification of hyperspectral data: A review of advances in manifold learning,” *IEEE Signal Processing Magazine*, vol. 31, no. 1, pp. 55–66, 2014.
- [6] Laurens van der Maaten and Geoffrey Hinton, “Visualizing data using t-SNE,” *Journal of Machine Learning Research*, vol. 9, pp. 2579–2605, 2008.
- [7] Yann LeCun, Yoshua Bengio, et al., “Convolutional networks for images, speech, and time series,” *The handbook of brain theory and neural networks*, vol. 3361, no. 10, pp. 1995, 1995.
- [8] Carlos Oscar Sánchez Sorzano, Javier Vargas, and A Pascual Montano, “A survey of dimensionality reduction techniques,” *arXiv preprint arXiv:1403.2877*, 2014.
- [9] Weikuan Jia, Meili Sun, Jian Lian, and Sujuan Hou, “Feature dimensionality reduction: a review,” *Complex & Intelligent Systems*, vol. 8, no. 3, pp. 2663–2693, 2022.
- [10] C.M. Bachmann, T.L. Ainsworth, and R.A. Fusina, “Exploiting manifold geometry in hyperspectral imagery,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 43, no. 3, pp. 441–454, 2005.
- [11] Lianru Gao, Daixin Gu, Lina Zhuang, Jinchang Ren, Dong Yang, and Bing Zhang, “Combining t-distributed stochastic neighbor embedding with convolutional neural networks for hyperspectral image classification,” *IEEE Geoscience and Remote Sensing Letters*, vol. 17, no. 8, pp. 1368–1372, 2020.
- [12] Leland McInnes, John Healy, and James Melville, “UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction,” Sept. 2020, arXiv:1802.03426 [cs, stat].
- [13] John Healy and Leland McInnes, “Uniform manifold approximation and projection,” *Nature Reviews Methods Primers*, vol. 4, no. 1, pp. 82, 2024.
- [14] Etienne Becht, Leland McInnes, John Healy, Charles-Antoine Dutertre, Immanuel WH Kwok, Lai Guan Ng, Florent Ginhoux, and Evan W Newell, “Dimensionality reduction for visualizing single-cell data using UMAP,” *Nature biotechnology*, vol. 37, no. 1, pp. 38–44, 2019.
- [15] Alex Diaz-Papkovich, Luke Anderson-Trocmé, and Simon Gravel, “A review of UMAP in population genetics,” *Journal of Human Genetics*, vol. 66, no. 1, pp. 85–91, 2021.
- [16] Hans-Jürgen Zimmermann, *Fuzzy set theory—and its applications*, Springer Science & Business Media, 2011.
- [17] Tim Sainburg, Leland McInnes, and Timothy Q Gentner, “Parametric UMAP embeddings for representation and semisupervised learning,” *Neural Computation*, vol. 33, no. 11, pp. 2881–2907, 2021.
- [18] F.P. Miller, A.F. Vandome, and J. McBrewster, *Beer-Lambert Law*, VDM Publishing, 2009.
- [19] Reza Azad, Ehsan Khodapanah Aghdam, Amelie Rauland, Yiwei Jia, Atlas Haddadi Avval, Afshin Bozorgpour, Sanaz Karimijafarbigloo, Joseph Paul Cohen, Ehsan Adeli, and Dorit Merhof, “Medical image segmentation review: The success of u-net,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [20] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna, “Rethinking the inception architecture for computer vision,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [21] Purdue University, “Indian pines test site, nw tippecanoe county, indiana,” 1992, Hyperspectral image data collected by AVIRIS sensor, available through Purdue University’s MultiSpec software.
- [22] Shaohua Li, Yong Liu, Xiuchao Sui, Cheng Chen, Gabriel Tjio, Daniel Shu Wei Ting, and Rick Siow Mong Goh, “Multi-instance multi-scale CNN for medical image classification,” in *Medical Image Computing and Computer Assisted Intervention—MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part IV 22*. Springer, 2019, pp. 531–539.

- [23] Dominik Müller, Iñaki Soto-Rey, and Frank Kramer, “Towards a guideline for evaluation metrics in medical image segmentation,” *BMC Research Notes*, vol. 15, no. 1, pp. 210, 2022.
- [24] Jorge R Vergara and Pablo A Estévez, “A review of feature selection methods based on mutual information,” *Neural computing and applications*, vol. 24, no. 1, pp. 175–186, 2014.
- [25] Haokui Zhang, Ying Li, Yuzhu Zhang, and Qiang Shen, “Spectral-spatial classification of hyperspectral imagery using a dual-channel convolutional neural network,” *Remote sensing letters*, vol. 8, no. 5, pp. 438–447, 2017.
- [26] Yushi Chen, Hanlu Jiang, Chunyang Li, Xiuping Jia, and Pedram Ghamisi, “Deep feature extraction and classification of hyperspectral images based on convolutional neural networks,” *IEEE transactions on geoscience and remote sensing*, vol. 54, no. 10, pp. 6232–6251, 2016.
- [27] Grégoire Mercier and Marc Lennon, “Support vector machines for hyperspectral image classification with spectral-based kernels,” in *IGARSS 2003. 2003 IEEE International geoscience and remote sensing symposium. Proceedings (IEEE Cat. No. 03CH37477)*. IEEE, 2003, vol. 1, pp. 288–290.
- [28] Sveinn R Joelsson, Jon Atli Benediktsson, and Johannes R Sveinsson, “Random forest classifiers for hyperspectral data,” in *Proceedings. 2005 IEEE International Geoscience and Remote Sensing Symposium, 2005. IGARSS’05*. IEEE, 2005, vol. 1, pp. 4–pp.