NATIONAL UNIVERSITY OF SCIENCE AND TECHNOLOGY POLITEHNICA OF BUCHAREST (UNSTPB)

Doctoral School of Electronics, Telecommunications and Information Technology

Decision No. 142 from 22-11-2023

Ph.D. THESIS SUMMARY

Omid GHOZATLOU

ÎNVĂȚAREA CU ESANTIOANE ADVERSARIALE PENTRU IMAGINI SATELITARE MULTISPECTRALE

LEARNING WITH ADVERSARIAL SAMPLES FOR EARTH OBSERVATION MULTISPECTRAL IMAGES

THESIS COMMITTEE

Prof. Dr. Ing. Gheorghe BREZEANU UNSTPB	President
Prof. Dr. Ing. Mihai DATCU UNSTPB	PhD Supervisor
Prof. Dr. Ing. Mihai CIUC UNSTPB	Referee
Prof. Dr. Ing. Cosmin ANCUTI University Politehnica Timisoara	Referee
Dr. Ing. Miguel HEREDIA CONDE University of Siegen	Referee

BUCHAREST 2023

Content

Chapter 1

Introduction

The term "Earth Observation" (EO) is a critical component of remote sensing (RS), enabling us to comprehend and oversee our planet's intricate and ever-changing environment. It involves acquiring information about Earth's surface and atmosphere from various sensors positioned on satellites, aircraft, drones, and ground-based platforms. In addition, it plays a crucial role in monitoring and managing a range of applications, such as agriculture, forestry, urban planning, disaster preparedness, climate monitoring, and environmental preservation.

Over the years, EO has evolved significantly, driven by technological advancements, satellite systems, and data analysis methods. Recently, the incorporation of Deep Learning (DL) techniques has revolutionized EO by extracting valuable insights from the vast and expanding collection of Earth imagesWith the continuous growth in the number of Earth-observing satellites and the availability of high-resolution, multispectral, and hyperspectral imagery, the volume of EO data is increasing at an unprecedented rate. This surge in data has presented both new challenges and opportunities for researchers, who are increasingly relying on DL as a potent tool to unlock the untapped potential within these datasets.

The integration of DL techniques with EO has expanded the frontiers for gleaning valuable insights from the extensive and ever-expanding collection of Earth images. This development has resulted in the creation of innovative applications in population monitoring, disaster management, and environmental preservation. Moreover, the utilization of EO imagery has had a societal and policy impact, driving the formulation of innovative and comprehensive regulatory frameworks aimed at balancing private and public interests in space-based EO.

1.1 Scope of the Thesis

In this thesis, we introduce new solutions to address this issue. We focus on four strategies which will be discussed in more detail in the next chapters. These solutions include Active learning (AL), Query-by-Example, Physic-aware deep models, and Generative adversarial networks (GANs). AL can address the adversarial samples issue in RS image classification by iteratively selecting the most informative samples to improve the model's performance. This approach helps the model to better understand the underlying structure and properties of the data, making it more resilient to adversarial attacks and improving its overall performance.

Another strategy we study in this thesis is Query-by-Example for RS image retrieval. This strategy aims to find the most similar image to the query and optimize the network's weights so that adversarial samples are far away from the query image in the latent space. In this way, the network will understand the similar samples to query and not be vulnerable against adversarial samples.

Another solution is guiding the DNN by using the statistical properties of data or injecting the physical properties of data into the model. In general, physic-aware deep models can address the adversarial samples issue in RS image classification and make the model more resilient to adversarial attacks by including physical knowledge in the learning process. GANs are a type of DL model that consists of two neural networks, a generator and a discriminator, which are trained together in a process known as adversarial training. The generator creates synthetic samples, while the discriminator attempts to distinguish between real and synthetic samples. GANs can be used to generate adversarial samples for RS image analysis, which can then be used to improve the robustness of the classification model against adversarial attacks.

Chapter 2

AL

ML plays a crucial role in RS image processing and has demonstrated impressive benefits for many applications such as RS image analysis. In particular, DNN have shown interesting properties on image datasets, and they are effective in extracting high-level intrinsic features and classifying complex problems. One of the most challenging issues in ML algorithms is labeling data because training a DNN is costly and necessitates a large number of training samples. In RS, obtaining labeled data can be expensive, time-consuming, or even impossible.

An effective way to deal with this issue is through AL, which can be used to leverage small amounts of labeled data. AL is an important concept in ML and has significance in RS image analysis. AL refers to a selection strategy for acquiring labeled data that allows the ML algorithm to actively select which samples to label, in order to improve model performance with less labeled data. Selectively labeling samples enhances the performance of ML algorithms on tasks such as classification or segmentation without requiring large amounts of labeled data.

This selection is done by a query strategy that offers the user the most informative or representative samples of an unlabeled data set. To retrain the model, the operator will relabel the selected samples and add them to the training set. Therefore, we train the model interactively with fewer training samples than the traditional ML passive learning methods. It also is more efficient (reducing the annotation cost) because it uses samples that are more useful for learning.

In this chapter, we use ResNet18 as a classifier to compare both classical and cutting-edge AL strategies within the same framework. Five cutting-edge deep AL methods are used, the most recent of which has never been tested on RS data before. We also propose a new AL performance metric based on the number of labeled samples. We use the random sampling strategy as a baseline and compare the performance of the strategies to the maximum accuracy achieved by the baseline. This novel metric calculates how many fewer labeled samples are needed for a given sampling strategy to achieve the same maximum accuracy as a random sampling strategy. In the end, ablation studies have been conducted on four different classifiers and AL batch query sizes. This study discusses a review of AL in RS for various applications. The challenges of combining DL and AL are then discussed, and deep AL solutions are provided. In the third section, a comparison of both classical and deep AL strategies is presented. Finally, the future scope and perspective of AL in RS are discussed.

2.1 Deep AL Challenges and Solutions

In the past decade, DL models have demonstrated remarkable prowess in signal processing. Leveraging their ability to capture hierarchical features from diverse data types—such as numerical, image, text, and audio—DL models have emerged as powerful solutions for tasks spanning recognition, regression, semi-supervised, and unsupervised problems. Distinctive traits set Deep AL (DAL) apart from traditional AL. First, DAL employs batch-based sample querying, a departure from the one-by-one query method commonly utilized in traditional AL algorithms. This shift not only reduces frequent model retraining but also addresses the limited variation in training data, a concern that can exacerbate overfitting. Second, while classical AL often employs traditional ML algorithms like SVM, DAL engages DL learners such as CNNs. This adaptation has been examined across various DL architectures, including Bayesian CNNs and stacked restricted Boltzmann machines, for image and text classification. It is noteworthy that a majority of proposed DAL methodologies consistently outperform random sampling by a substantial margin, irrespectively of the uncertainty-based strategy and classifier employed.

DAL is an innovative and intricate approach that seeks to harness the strengths of both DL and AL methodologies. Integrating AL with deep architectures, such as CNNs, presents challenges. AL techniques predominantly rely on probabilistic functions that analyze the probability distribution of a sample of belonging to existing classes. This clashes with most DL architectures, which typically yield predicted class labels instead of class probabilities, as they lack an inherent uncertainty mapping. To tackle this quandary, certain studies capitalize on the properties of Bayesian networks and conceive Bayesian Neural Networks (BNNs) to achieve probability distributions over network weights. For instance, in [51], a BNN is introduced to extract an output probability matrix from input data. This matrix facilitates the application of AL's probabilistic function, providing uncertainty estimates and a probabilistic interpretation of DL models through inferences of weight distributions. Moreover, in the domain of seismo-volcanic monitoring introduces a Bayesian temporal CNN (B-TCN) for continuous event detection and classification. This framework efficiently extracts the most uncertain events from continuous data streams. In essence, DAL represents a dynamic convergence of DL and AL methodologies, leveraging their complementary strengths. The incorporation of Bayesian Neural Networks and innovative architectures exemplifies the strides taken to surmount the challenges of combining these two paradigms, ultimately opening pathways for more informed and robust decision-making in complex domains like seismo-volcanic monitoring.

Furthermore, the realm of AL witnesses endeavors aimed at refining the AL process through the fusion of diverse criteria. A notable instance introducing an innovative AL algorithm that maximizes two selection criteria—representative and uncertainty sampling. This algorithm unfolds in two stages: firstly, an unsupervised feature learning process creates a weighted incremental dictionary based on training data, facilitating the estimation of data representativeness; subsequently, the supervised learning phase leverages this framework to estimate data uncertainty. Addressing the challenge of limited high-quality labeled training data, recent studies pivot towards semi-supervised techniques. Such techniques generate pseudo-labels from a small set of labeled examples, which in turn enriches the labeled training set. Notably, ML models can exhibit confident yet incorrect predictions, particularly when the training dataset fails to describe the latent space.

An alternative approach to tackle the scarcity of labeled data in DAL is through GANs, frequently employed to synthesize training samples. Adversarial AL gains traction within this framework. For instance [59] pioneers a feature-oriented adversarial AL strategy, utilizing high-level features from an intermediate layer of a DL classifier to establish a GAN-based acquisition heuristic. Moreover [32] introduces a novel active deep feature extraction scheme, incorporating both representative and informative criteria. An adapted adversarial autoencoder serves as the initial step to facilitate classification-specific deep feature extraction. The incorporation of dictionary learning and a multi-variance and distributional distance measure guides the selection of valuable candidate training samples.

In pursuit of enhancing DAL performance, alterations to the loss function or the deployment of deep ensemble methods come to the forefront. A supplementary module is often incorporated to enhance the target model's performance and guide DAL selections. One such approach involves predicting instances where the target model might err by introducing a loss prediction model alongside the target backbone model. Furthermore, an study constructs an auxiliary deep network for the basic learner, dedicated to learning the uncertainty of unlabeled samples in the candidate dataset. Converging the features of the original training data and those of the intermediate hidden layer of the basic learner, a fully connected network is devised, bolstered by a novel loss function tailored to the task.

In summary, researchers are innovating across various fronts to amplify the efficacy of DAL. These strategies encompass intricate combinations of selection criteria, integration of semi-supervised techniques, exploitation of GANs and adversarial AL, and the augmentation of target models through diverse supplemental modules. These multi-pronged approaches collectively contribute to the evolution of DAL, presenting a nuanced landscape of techniques to confront challenges and advance its capabilities.

2.2 Classification Performance Comparison

In this section, we present nine query strategies in detail. These strategies are mostly based on uncertainty and aim to select new data samples that maximally reduce the uncertainty in the basic learners.

{Entropy}

The Maximum Entropy selects instances based on the entropy of the model's prediction distribution. It is based on the idea that instances for which the model has high entropy, i.e. instances for which the model is uncertain about the correct label are likely to provide the most information for improving the model.

{Least Confident}

Least Confidence is an uncertainty-based AL query strategy that selects the instances for which the model is least confident. In other words, the model selects instances for which it has the highest uncertainty in its predictions. The model ranks instances by their prediction confidence scores and the instances with the lowest scores are selected for labeling.

{Margin}

Margin Sampling selects instances for labeling based on the difference between the model's maximum predicted probability and the second highest predicted probability. The idea is that instances with a high margin between the maximum and second highest predicted probability are likely to provide the most information for improving the model.

{Mean Std.}

Mean Std. is based on the standard deviation of the model's predicted probabilities for each class. It is based on the idea that instances for which the model's predicted probabilities have a high standard deviation are likely to provide the most information for improving the model. It maximizes the average of the standard deviation of the predicted probabilities over all classes.

{BALD}

Batch Bayesian AL by Disagreements (BALD) is a practical approximation for the mutual information between a batch of points and model parameters. This approximation serves as an acquisition function, enabling the joint selection of multiple informative points in the context of deep Bayesian Active Learning. BALD focuses on choosing data points expected to maximize the information gained from the model parameters, specifically the mutual information between predictions and the model posterior.

{BADGE}

Batch AL by Diverse Gradient Embeddings (BADGE) employs diverse gradient embeddings to capture the variability in a model's predictions. It selects a diverse subset of samples that spans the embedding space, measuring uncertainty as the gradient magnitude concerning parameters in the output layer. This computation is based on the most likely label according to the model. The objective of BADGE is to reduce redundancy and enhance the diversity of selected samples, contributing to the efficiency and effectiveness of the Active Learning process.

{VAAL}

Variational Adversarial AL (VAAL) is an AL algorithm that combines variational autoencoders (VAEs) and adversarial training to select informative samples for annotation. It uses the VAE to model the data distribution and generate new samples that are close to the real data, and the adversarial training to discriminate between the real and generated samples. VAAL then selects the most informative samples based on a measure of uncertainty and diversity and uses them to update the model. By combining the VAE and adversarial training, VAAL can generate diverse and informative samples, and the AL process can be more efficient and effective. VAAL has been shown to outperform several other state-of-the-art AL algorithms on various benchmark datasets.

{WAAL}

Wasserstein Adversarial AL (WAAL) offers theoretical insights by framing the interactive process in Active Learning (AL) as distribution matching, employing the Wasserstein distance. In this method, informative samples challenging for the current model to classify are chosen and incorporated into the training set to enhance model performance. The selection of the most informative samples involves a combination of uncertainty and diversity measures. The Wasserstein distance is then utilized to identify examples that maximize the dissimilarity between predicted and true distributions, making them more likely to be informative. Leveraging the Wasserstein distance allows WAAL to generate more informative and diverse samples, thereby improving the efficiency and effectiveness of the AL process. WAAL has demonstrated superior performance compared to several other state-of-the-art AL algorithms across various benchmark datasets.

{LPL}

Loss Prediction Loss (LPL) is a 2019 strategy that incorporates a small parametric module into a target model. This module is trained to predict the loss of unlabeled inputs concerning the target model. Subsequently, the module can recommend data for which the target model is prone to make incorrect predictions. The objective of this approach is to enhance the performance of the target model by leveraging the predictions of the loss prediction module.

2.3 Results

In the assessment of Active Learning (AL) strategies, ResNet is employed as the classifier, and AUBC, final accuracy, and LDN metrics are presented for both EuroSAT (RGB) and EuroSAT (MS). The dataset is evenly divided into 50% training and 50% test sets. The results are outlined in Table 1 when M=500 and Q=10000. Conventional AL strategies consistently show 1% to 3% higher performance compared to Random selection in terms of AUBC. Among them, Entropy exhibits the highest AUBC (76.01%) and final accuracy (88.57%). Mean Std. also performs notably well. DAL (Diversity in Active Learning) strategies, such as VAAL and LPL, achieve even higher AUBC values, around 80%. WAAL and BALD also outperform conventional strategies with AUBC values of 77.8% and 76.18%, respectively. In terms of final accuracy, all DAL strategies outperform conventional strategies. LPL, VAAL, and WAAL achieve final accuracies of 94.12%, 92.98%, and 92.66%, respectively—representing over 7% higher accuracy than Random after Q is exhausted. Other strategies, like BALD and BADGE, also achieve higher final accuracy, around 90% and 89%, respectively. Among conventional strategies, Entropy and Mean Std. outperform Margin and Least Confident on the EuroSAT (RGB) dataset. In summary, advanced AL strategies, especially DAL techniques, markedly enhance accuracy and efficiency across various datasets.

Table 2.2: Comparison of nine AL strategies on EuroSAT (RGB) and EuroSAT (MS).

	EuroSAT (RGB)			**EuroSAT (MS)**		
Strategy	*AUBC*	*Final ACC*	*LDN*	*AUBC*	*Final ACC*	*LDN*
Random	74.524	85.27	100	74.609	86.55	94.33
Entropy	76.013	88.57	72.6	80.925	92.88	47.81
Least Confident	74.434	86.07	78.25	80.594	93	43.91
Margin	75.163	86.61	78.94	76.172	87.9	74.81
Mean Std.	75.235	87.39	71.23	77.456	91.22	48.65
BALD	76.176	90.06	69.47	79.933	92.34	55.76
BADGE	74.957	89.01	82.78	81.326	91.74	39.31
VAAL	80.11	92.98	50.83	82.664	94.98	34.8
WAAL	77.794	92.66	64.67	84.427	94.12	20.28
LPL	79.87	94.12	51.54	85.677	96.63	22.34

As discussed earlier, the LDN metric elucidates the percentage of labeled data required to achieve the optimal performance obtained by Random sampling. Lower LDN values denote superior performance. Table 2.2 showcases the compelling efficiency of certain strategies. VAAL and LPL, for instance, necessitate just 50% of Q compared to Random, enabling them to achieve Random's peak accuracy with merely half the labeling budget. Similarly, WAAL and BALD exhibit reduced LDN values when contrasted with conventional AL strategies. When comparing the average LDN across conventional strategies (75.25%) with that of DAL strategies (63.85%), a revelation emerges: DAL strategies require around 12% fewer data to attain Random's maximum accuracy. Notably, while BADGE boasts commendable final accuracy, its high LDN and comparatively low AUBC dampen its overall performance across the entirety of the AL process.

Extending our evaluation to the EuroSAT (MS) dataset, we find that the increased information gleaned from the 13 MS bands translates to enhanced performance. Broadly speaking, AUBC and final accuracy values soar across all AL strategies when applied to the MS dataset compared to RGB. Relevant disparities arise; notably, Entropy and Least Confident achieve approximately 5% and 6% higher accuracy, respectively, on EuroSAT (MS) compared to EuroSAT (RGB). Among DAL strategies, BALD demonstrates a commendable 4% improvement, while the other strategies exhibit more than a 6% enhancement. LPL and WAAL stand out, achieving 85.68% and 84.43%, respectively—highlighting their prowess as the best and second-best strategies.

An intriguing observation arises: both Entropy and Least Confident exhibit approximately 1% higher AUBC than BALD, showcasing nuanced dynamics for EuroSAT (MS). Additionally, the final accuracy of all strategies on the MS dataset surpasses that of EuroSAT (RGB). LPL and VAAL lead the pack, securing the highest and second-highest final accuracy, respectively. The majority of conventional AL strategies, except Margin, achieve final accuracy rates exceeding 91%. This margin is not far from the average final accuracy of DAL strategies, which hovers around 94%. This robust outcome accentuates the MS dataset's capacity to empower models with more comprehensive information, culminating in superior performance.

One of the most intriguing analyses revolves around the concept of LDN, shedding light on the efficiency of WAAL and LPL. Impressively, these strategies require only a mere 20% of labeled data to achieve the same pinnacle of accuracy as Random sampling. This revelation underscores their strategic significance, particularly when confronted with budgetary constraints that limit labeling expenditures. Equally compelling are the LDN values for VAAL and BADGE, standing at 34.8% and 39.31%, respectively.

A notable highlight emerges from BADGE's performance on the EuroSAT (MS) dataset. Here, it displays an exceptional 6% improvement in AUBC and a remarkable reduction of almost 40% in LDN when compared to the RGB dataset. This underscores BADGE's remarkable potential, particularly in scenarios where labeled samples are more abundant. The comparative study of Active Learning (AL) strategy performance between EuroSAT (RGB) and EuroSAT (MS) unveils the transformative power of incorporating Multi-Spectral (MS) images. This integration effectively addresses the challenges posed by limited labeled data, resulting in an overall boost in model performance.

2.4 Conclusion

In this chapter, we conducted a comprehensive comparison of nine distinct AL strategies within a unified framework, focusing on their performance on two RS image datasets. Notably, some of these strategies were employed on RS datasets for the first time, showcasing the novelty and breadth of our exploration. Our analysis spanned two classifiers, RF and ResNet18, elucidating their performance across varying amounts of labeled samples. The findings underscore the general trend where DNNs, in particular, exhibit improved performance with a larger pool of labeled data samples, reinforcing the value of data abundance.

By comparing conventional AL strategies using ResNet and RF classifiers, we observed that while RF demonstrates stable performance-budget curves, it falls short of achieving the accuracy levels of ResNet during the AL process. An ablation study delved into the impact of the number of query samples per AL round (M) on ResNet's performance, revealing that a higher M yields more stable performance-budget curves. Moreover, our study introduced a novel metric, LDN, for evaluating AL strategies. LDN revealed that certain DAL strategies, such as LPL and WAAL, achieved the lowest requirement for labeled data to perform at the baseline level. This characteristic positions these strategies as particularly beneficial when facing limited labeling budgets.

Additionally, our study delivered a comprehensive overview of the evolution of AL in the RS domain, starting with the initial naïve AL strategies and tracing their continuous refinement over time. We highlighted the recent surge in attention towards leveraging DL in AL and elucidated the critical necessity of employing AL in the context of DL networks. We tackled challenges unique to DL, pointing to innovative solutions, such as Bayesian networks to address the non-probabilistic nature of DNNs, the amalgamation of diverse criteria to ensure informativeness and representativeness, the application of semi-supervised techniques for pseudo-labeling, and the utilization of GANs to mitigate the scarcity of labeled data. Furthermore, we advocated for the design of specialized networks for AL by modifying loss functions, showcasing their potential to enhance the overall AL process. In essence, this study constitutes a multifaceted exploration, offering insights into the effectiveness of various AL strategies on RS datasets, novel performance metrics, and innovative solutions to address challenges in integrating AL with DL networks. Our findings contribute to a deeper understanding of the synergy between AL and RS, providing valuable guidance for efficient and effective data acquisition and annotation in this dynamic domain.

Chapter 3

Query-by-Example for RS Image Retrieval

Thanks to the capability of mimicking a non-linear function of DNN, these models are able to capture the essential characteristics of images. The embedding in latent-space layers of DNN models can be taken as high-level features to comprehensively represent the visual content of RS images. As a result, DL provides the most common and useful feature extraction methods with many applications in RS image processing. The main issue with using DL is that it is greedy for data to optimize a massive number of parameters. If a DNN model is to be trained in a supervised manner, a large amount of labeled data is required.

In the RS big data era, we can easily collect a large amount of raw data, but accurately labeling oversized data becomes challenging because there exist many kinds of RS images compared with the fixed RGB format of natural images in the computer vision domain. For this reason, experts study unsupervised DL methods to tackle the data labeling problem. However, these methods can have wildly inaccurate results and need enhancement and improvement to achieve more accurate and reliable results. In recent studies, many enhancement techniques have been proposed to help or guide DL models. Physics-aware DL, combining feature extraction approaches with DNN, modifying the network architecture, and designing cost functions are some examples of these techniques, which are further described below.

In RS, experts pay attention to extracting physical information in different wavelength bands of multispectral images. In addition to the spectral features, spatial features provide important information on data. For example, authors in [139] provide a method to capture spatial detail. Therefore, a graph convolutional network with pairwise similarity constraint is proposed to address RS image retrieval. Generally, different shape features help to make up for each other's defects. As a consequence, the combination of multiple hand-crafted features often presents stronger representation ability and benefits in improving RS image retrieval.

Our work also contemplates the design of a new loss function that can cluster similar features (or embeddings) in the latent space. The main idea is used for anomaly detection, which is an unsupervised problem by nature. Image retrieval like anomaly detection can be considered a one-class classification that can be conducted by any kind of classification method. Therefore, this can be a clever solution to solve the scarcity of labeled training data. For instance, authors proposed a data enclosing-ball minimizing autoencoder for change detection. As well as, another study suggested a bidirectional GAN-based one-class classifier for network intrusion detection in which only normal data samples are used for training and both normal and anomalous data samples are used during testing to identify anomalous samples in the test set. In line with this research, we selected and utilized a one-class classification methodology. However, we enhanced the strategy by creating a new cost function to model the multispectral data for the RS image retrieval task.

In this chapter, we exploit and modify the idea of deep SVDD [2] for RS image retrieval. Deep SVDD is inspired by kernel-based one-class classification and uses a neural network to map most of the data network representations into a hypersphere. Therefore, a DNN is jointly trained to map the data into a hypersphere of minimum volume in the latent space. It is expected that relevant (similar) samples to the query are concentrated inside of the hypersphere and irrelevant samples (outliers or ambiguous) are moved away from the hypersphere. The closest embedding to the hypersphere center corresponds to the most relevant sample to the query. We modify the network's objective function to take advantage of the statistical information of data. We employ the covariance regularization in [156] to prevent an informational collapse in which the encoder produces constant or non-informative vectors. It penalizes unnecessary redundancy of the embedding in the latent space. In addition, we propose a novel cost function to minimize the volume of the hypersphere by unlocking the hypersphere's predefined center while preventing network divergence during training. It allows the hypersphere center to be free to compress the relevant samples as much as possible while driving irrelevant ones away.

Our main contributions are the following:

- To the best of our knowledge, this paper is the first to study deep SVDD (or any one-class classification method) for CBIR. Due to not being a fully supervised approach, less labeled data (only samples of one class) is needed for training and the samples of other classes can be unknown.
- Enhancing the cost function with an additional covariance term to minimize the correlation between dimensions in the latent space becomes especially relevant when dealing with high-dimensional data.
- Providing a novel cost function by unlocking the hypersphere's predefined center and considering the average of embeddings over a batch in each iteration as a moving center while preventing network divergence during training.
- Furthermore, we investigate the relationship between the entropy of images and the image retrieval performance of the enhanced methods.

3.1 Results

The 60 most relevant samples corresponding to different queries are shown in Figure 3.2. The top-left sample is the most relevant and the score is increasing from left to right. The first two rows show the results of deep SVDD for \$d=128\$. We can see that the image retrieval works well for only the See/Lake class, which has uniform distribution and the lowest average of entropy. By looking at the most relevant samples for other classes, we recognize all samples are uniform and texture-less. That means the network is able to learn only uniform samples and demonstrates the need for a more elaborated cost function to learn more complex features.

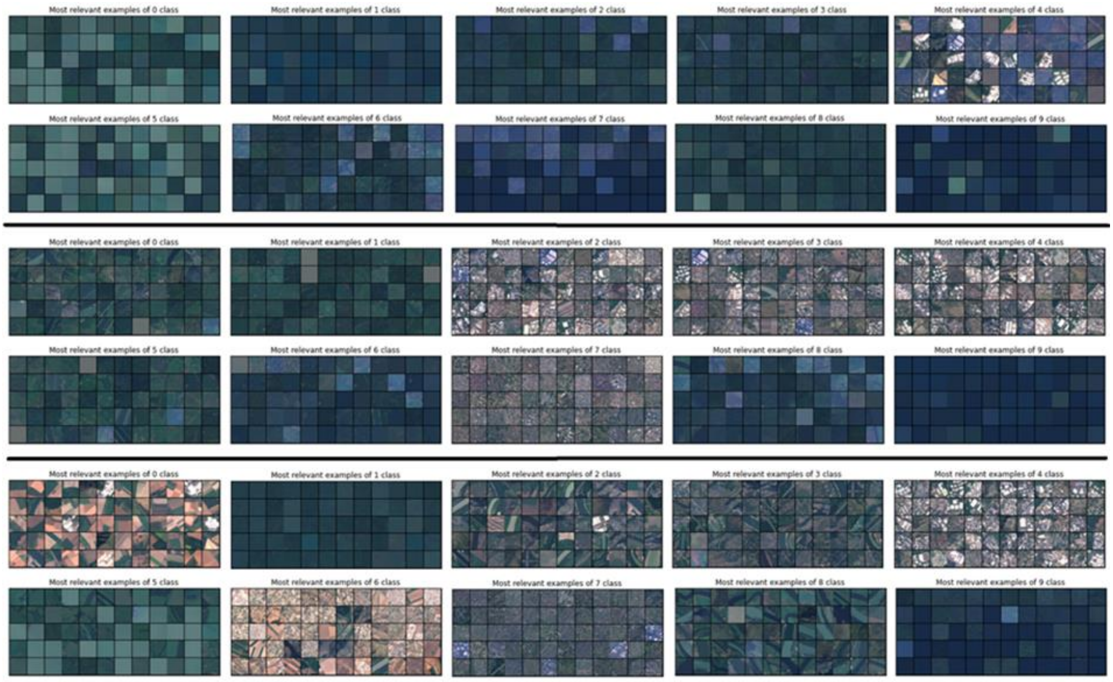

Figure 3.2: The top 60 most relevant samples are retrieved from the test set.

In contrast, most of the retrieved samples using an additional covariance term are more complex samples much closer to the human expectation for each class. The results for $\nu = 1024d$ are shown in the two middle rows of Figure 3.2. The performance of the model for Industrial and Residential queries improved more than for other classes. Also, the most relevant samples of these two classes are totally correct. In addition, retrieved samples of Forest and Pasture are much more accurate than using deep SVDD. However, some samples of Forest are retrieved for a query of Pasture, which might be because of the similar spatial structure. The last two rows indicate the results of DC-SVDD when $a=0.001$ and $b=1$. They show that the additional drifting center term improves the image retrieval of almost all classes. The samples of the Annual Crops, Permanent Crops, and River are retrieved more reliably than with previous methods. Also, the most relevant samples of Industrial and Residential classes are still totally correct. In addition, image retrieval for Highway queries has improved, despite there being some samples of Rivers in the top 60 most relevant samples. The similar spatial texture of these two classes leads the network to consider them the same class. It is worth mentioning that there are some incorrectly retrieved samples. More specifically, one can observe some retrieved Forest samples for the Sea query and Annual Crops for the Permanent Crops query. Unfortunately, no method can correctly retrieve samples of Herbaceous Vegetation query.

3.2 Conclusion

In this chapter, we have introduced two novel methods to guide DNN by considering statistical information of data. The first proposed method penalizes unnecessary redundancy and minimizes the correlation between different dimensions of embedding. The second method tries to minimize hypersphere volume considering the unblocking of the center of the hypersphere. This relaxation makes the clustering of data easier for DNN. Both proposed methods overcome the limitations of Deep SVDD in particular when the query belongs to a complex class. We have emphasized that these are very important advantages because the main objective of CBIR is to optimize the search with a minimum number of annotated images to find the most similar samples in the archive.

Moreover, we have introduced the use of deep SVDD for CBIR problems (particularly in the context of the query by example) in RS. The experimental performances of the proposed system were evaluated on an archive of 5400 images describing 10 different categories for the EuroSAT dataset and also another well-known dataset, UCM with 21 classes is employed to assess the effectiveness and robustness of the proposed methods. The results show that the proposed methods provide efficient image retrieval performance, largely surpassing the considered baseline method (SVDD). In addition, we evaluated the performance on multispectral data, attaining superior results as compared to RGB data. It is also worth noting that the proposed methods are independent of the considered DNN, and therefore, they can be used for any cost function to guide the DNN to achieve a superior representation of the data.

Chapter 4

Physic-Aware DNNs

In this section, we propose wavelet-guided DNNs by taking advantage of the properties of WST to extract invariant features for robust clustering. In the proposed method, the network is trained on samples from only one class in the training set and is evaluated on the test set including all classes. This strategy has the benefit of not requiring labeling for every class in the dataset which addresses the scarcity of training labeled data. We focus on the one-class classification task and discuss the challenges and benefits of some methods used in recent studies. One-class classification methods aim to directly learn a decision boundary with a low error when applied to unseen data.

Authors proposed a novel method, called one-class transfer learning. They took into account that in the field of computer vision, one can access labeled data from different domains that are not related to one-class classification datasets and benefit from using data from a different domain. Similarly, authors utilize the objective function inspired by information theory, which maximizes the distance between normal and anomalous data in terms of the joint distribution of images and their representation.

We exploit a WST network to guide the DNNs for one-class classification. WST provides the invariant representation which is informative because of keeping a high frequency of data. Using these properties, can cluster the normal samples and distinguish outliers. At first, WST extracts scattering coefficients from given images, and then the network is applied to them. This leads to obtaining more robust and accurate results. Furthermore, to address the problem of the insufficient amount of training set for DNNs, an adequate initial feature extraction to ease the training task is essential. A thorough evaluation confirmed superior performance and robustness both to outliers (not normal samples) in the training set and to translation and rotation of the test set. Furthermore, we investigate the relationship between the entropy of images and the guidance of WST for highly textured images.

4.1 Physics-Aware GAN for Cloud Removal

RS imagery data offers a valuable means for EO to analyze and extract information, providing insights into Earth's resources and various physical phenomena parameters. With the advancements in sensor technology for satellite imagery, access to high-quality images featuring high spatial resolution has become feasible. This capability enables the capture of detailed information in images, leading to the growth of applications like land cover classification, thematic mapping, environmental monitoring, and natural resource management. Despite these advancements, it's worth noting that statistics indicate cloud cover obscures more than half of the Earth. This is a common issue with optical RS images causing information to be obscured by clouds and their associated shadows. As a result, we cannot capture reliable information from corrupted images unless we use clear sky images at the same time that they are not available. Therefore, one reasonable solution is improving the networks by leveraging trustworthy and transparent physical properties. Our proposed method exploits spectral angular distance (SAD) to train cycle-consistent adversarial networks with illumination invariant features.

Since, we need cloud-free images, many methods are proposed to detect and remove clouds from RS images. By reviewing studies in cloud removal, we have grouped the methods into three categories. One category is multitemporal-based in which there are used multitemporal images of the same area. The drawback of this category is that on the one hand, the time interval of multitemporal images is long on the other hand the area is changing rapidly. For this reason, usually, the accuracy of the reconstructed area is low. One comprehensive study in this category proposes spatiotemporal fusion using a Poisson-adjustment method for cloud removal in multi-sensor and multi-temporal images. The approach employs a Poisson-based residual correction strategy to enhance spectral coherence between recovered and cloud-affected regions.

In the past decade, DL has played an even more significant role in RS, with a focus on DL-based approaches. GANs, a notable neural network model, utilizes two networks in an adversarial manner to generate clearer and sharper images. Nowadays various kinds of GANs have been created for different applications. Therefore, It was not overlooked by the RS experts and many studies have been published for cloud removal using GANs.

The third category is based on multisensory data fusion which is used in auxiliary penetrable modalities. A helpful auxiliary data in this field consists of SAR images. Thanks to long-wavelength SAR can penetrate through clouds in different weather with different kinds of clouds. It has caused it to be used a lot by various methods in the RS area such as data fusion and image translation to obtain and reconstruct the information in contaminated optical images. SAR-optical data fusion was proposed in to remove clouds using cycle GAN. Also, they computed cloud probability masks to model cloud coverage explicitly while reconstructing cloud-covered information.

Unlike most computer vision algorithms for dehazing which are based on image enhancement, we inject physical properties into DNNs to reconstruct the contaminated regions. Therefore, we use 12 Sentinel-2 bands in order to benefit the spectral reflectance of multispectral images. In this chapter, we extend the proposed model in [3] based on cycle-consistent GAN to remove clouds from corrupted multispectral images. The most significant advantages of this method are the elimination of the need for paired images (cloudy/cloudless) and the use of an auxiliary modality that penetrates clouds. In contrast to [3] we exploit transparent information in the dataset by translating pixel values into polar coordinates, and then we train the network with illumination invariant features to reduce the impact of clouds and shadows. The proposed method, Hybrid GAN-SAD not only achieves notable results but also increases the trustworthiness of the network by utilizing reliable physical properties.

4.2 Results

For a comprehensive comparison of results, we employed three distinct training strategies for our network. In the initial experimental setup, the network was trained exclusively using the RGB visible bands (B2, B3, B4) from the training dataset. In the second configuration, the network was trained using Near-Infrared (B8) in addition to RGB (4 bands), and in the last setup, the full complement of bands (12 bands) was utilized. The outcomes of each method are depicted in Figure 4.8. For visualization purposes, only B4, B3, and B2 images were used, and the brightness of the figure was adjusted for clarity.

The figure's first row displays samples of test images, while the second row exhibits the network output using only visible bands as implemented by Singh et al. [3]. Subsequently, the third row presents results incorporating Near-Infrared bands, and the fourth row showcases outcomes using the full spectrum of dataset bands. Furthermore, the fifth row illustrates results from the network trained on images converted to polar coordinates, representing our proposed method. Lastly, the sixth row displays cloud-free patches with a 16-day interval, serving as the ground truth. Notably, haziness was effectively removed in all methods (columns 5-7). Near-Infrared, benefiting from its ability to penetrate thin clouds, outperforms other methods in dehazing (columns 2-4). Our proposed model not only exhibits a notable impact on cloud removal but also demonstrates considerable efficacy in shadow elimination and information recovery beneath shadows, a capability not observed in other methods (columns 2-4). A closer inspection reveals that handling corrupted patches caused by dense clouds poses a challenge, and complete information recovery from the background without leveraging auxiliary data such as SAR proves to be impossible.

4.3 Conclusion

In the first section of this chapter, we demonstrate that WST considers the statistical properties of data and extracts invariant features that can help DNNs for more accurate classification, especially for more complex images. It is worth noting that complex RS images usually have big intraclass diversity and complex features which makes classification more difficult and leads to misclassification. The proposed method alleviates the need for a large training set for the DNNs because it leverages WST to achieve more stable features. The robustness to pollution in a normal training set and also robustness to transformation were discussed. In addition, it was observed that WST compensates for the limited performance of the DNNs for classes with high entropy. By improving DNN classification performance, the model is more robust against adversarial or any kind of artifacts.

Figure 4.8: Qualitative results, Row I: cloudy test samples, Row II: results of RGB, Row III: results of RGB+IR, Row IV: results of full bands, Row V: results of Hybrid GAN-SAD, Row VI: ground truth.

The second section addresses directly the physical adversarial attacks in RS images. Natural adversarial in EO such as clouds and shadows are common artifacts in RS image analysis. By injecting physical properties into the cycle-consistent GAN, we were able to convert a cloudy multispectral image to a cloudless image. To recover information beneath clouds and shadows, we create a synthetic multispectral space to obtain illumination-invariant features and train the GAN in this space. The proposed method modifies DNNs by injecting physical properties to achieve trustworthy results. As we recover real information about the background using the physical properties of data, we can trust the outputs of the model.

Chapter 5

Synthetic Image Generation by GANs

GANs, a neural network type for image generation and data synthesis, tackle adversarial challenges in RS by establishing an objective function through training data. The discriminator distinguishes real and fake images, guiding the generator to enhance its output based on current network weights. GANs are widely used for various data synthesis and manipulation tasks, particularly in the visual domain.

The ocean, crucial for Earth's climate regulation, is studied using SAR images, providing insights into oceanic processes and their climate change impact. This chapter introduces a GAN-based method for creating realistic and diverse ocean-pattern SAR images. Utilizing a style-based generator network and an adversarial discriminator network, the approach recognizes and generates intricate patterns in SAR imagery. To counter discriminator overfitting common in limited training data scenarios, the ADA mechanism is employed. Training the GAN with ADA helps the generator capture spatial and statistical characteristics of oceanic phenomena under restricted data conditions. However, utilizing GAN-generated images for RS applications requires a thorough assessment of their quality and authenticity, as well as validating the model's performance on real-world data. We examine the reliability of GAN-based generated images and propose two approaches approximating precision and recall for GANs performance. DL has become a pivotal tool in the processing of RS images, playing a vital role in EO technology. Its effectiveness in extracting semantic information and classifying high-resolution satellite images has yielded promising outcomes, demonstrating high accuracy, recall, precision, and other evaluation metrics. The increasing enthusiasm within the RS community for DL methods has spurred the creation of numerous architectures specifically designed to tackle RS problems, frequently delivering outstanding performance. These advancements have notably enhanced RS image analysis, establishing DL as an indispensable component in the field.

5.1 Why Synthetic data?

Implementing DL models for RS image processing poses a challenge due to the demand for substantial amounts of training data. Adequate training data is essential for achieving optimal performance in machine learning models, enabling them to discern underlying patterns and relationships within the data. However, acquiring a significant volume of training data can be both costly and time-intensive, and in some instances, obtaining a representative sample for a specific problem may prove challenging.

Various solutions have been proposed to tackle the data requirements issue in training DL models for RS image processing. Data augmentation techniques, such as flipping, rotation, and scaling, serve to amplify the training data by generating new samples from the existing dataset. Another strategy is transfer learning, which involves utilizing pre-trained models on extensive datasets to extract features and fine-tune the model for the specific RS problem. Active learning, conversely, entails the iterative selection of the most informative samples from the dataset to train the model, minimizing the need for a large amount of labeled data.

A particularly promising solution for data augmentation involves synthetic data generation through GANs. GANs, comprising a generator and a discriminator, engage in a competitive process to produce realistic synthetic data. The generator crafts synthetic samples, and the discriminator assesses their quality, providing feedback to enhance the generated data. Utilizing GANs for synthetic data generation offers more diverse and realistic samples compared to conventional data augmentation techniques, which rely on linear transformations. Synthetic data generated by GANs proves valuable in augmenting the training dataset, mitigating data scarcity limitations, and enhancing the overall performance of DL models in RS image processing.

Evaluating GAN-based generated images for RS applications necessitates a comprehensive assessment of their quality, authenticity, and the model's real-world performance validation. The selection of appropriate evaluation metrics plays a pivotal role in accurately gauging GAN performance. However, identifying suitable metrics for GAN-generated images proves challenging, given that traditional metrics may not entirely capture the images' quality and diversity. Typically, researchers rely on metrics such as Inception Score (IS), Frechet Inception Distance (FID), and Kernel Inception Distance (KID) to evaluate GAN performance. While these metrics offer insights into training progress, they do not necessarily correlate with real-world tasks**.**

5.2 Why Ocean image analysis?

The ocean assumes a vital role in Earth's climate regulation, influencing the intricate dynamics among the ocean, atmosphere, and other climate elements. Observing the ocean provides valuable insights into diverse processes, particularly with the aid of SAR imagery. SAR contributes detailed information on oceanic and coastal activities, fostering an improved comprehension of climate dynamics. This enhanced understanding, in turn, aids in refining climate models and supporting research, monitoring, and mitigation endeavors related to climate change. Additionally, precise measurements and observations of the ocean surface are integral for comprehending air–sea interactions and developing high-resolution climate models.

SAR plays a pivotal role in advancing our understanding of the world's oceans. One of its most significant advantages in oceanography is its all-weather capability. Unlike optical sensors, which are hindered by cloud cover and darkness, SAR can operate effectively regardless of weather conditions, providing a continuous stream of data crucial for studying dynamic ocean processes. This is particularly important for monitoring and analyzing phenomena such as ocean currents, wave patterns, and wind behavior, which are vital components of ocean circulation and climate systems. Additionally, SAR's ability to operate day and night allows researchers to capture critical nighttime events and track changes in oceanic conditions around the clock, contributing to a more comprehensive understanding of the marine environment.

In this chapter, our focus is on the analysis of Ocean SAR images. We present a novel evaluation metric specifically crafted to measure the diversity of images generated by GANs. Our methodology involves training a classification network using images generated by a GAN and subsequently evaluating its performance on a test set comprising real-world images. This evaluation metric quantifies the difference between the acquired (generated images) and the desired (real images) data distributions. By assessing the classification network's accuracy in categorizing real images, we can deduce the similarity between the generated and real images, offering a robust measure of their likeness.

To evaluate the reliability of the generated images, we utilize ResNet18 as a classifier and conduct training and testing in two distinct experimental setups. In each setup, we create two balanced datasets, one comprising real images and the other generated images. The first setup involves training the classifier on real images and testing on generated images, while the second setup operates vice versa. By comparing the classification accuracy in each setup, we approximate the precision and recall for the performance of GANs.

5.3 StyleGAN2 with ADA

In this section, we present the methodology of StyleGAN with ADA [211], which extends the original StyleGAN [212] framework with a dynamic augmentation strategy for the discriminator during training. The StyleGAN2-ADA architecture consists of a generator and a discriminator network, both of which are composed of convolutional layers. The architecture is similar to StyleGAN2, with the primary difference being the introduction of the adaptive discriminator augmentation mechanism to stabilize training when using limited data. The methodology encompasses the following steps and components:

{Mapping Network}

The mapping network is a fully connected network that maps points in the latent space to an intermediate latent space. This network is responsible for controlling the style of the generated images.

{Intermediate Latent Space}

The intermediate latent space, also referred to as the style space, serves as a mechanism to influence the style of generated images across various levels of detail.

This style space is incorporated into the generator model at multiple stages, enabling precise control over the characteristics of the generated images.

{Weight Demodulation}

Weight Demodulation, a technique introduced in StyleGAN2, reconstructs the augmentation operation from the original StyleGAN. This method is employed to regulate the style of generated images at different levels of detail.

{Noise Injection}

Noise is deliberately introduced at each point within the generator model, functioning as a source of variation. This noise contributes to the creation of stochastic variations in the generated images, introducing texture and fine details.

{Adaptive Discriminator Augmentation (ADA)}

The adaptive discriminator augmentation mechanism in StyleGAN2-ADA applies random augmentations to the input images during training, which helps prevent overfitting and stabilizes training when using limited data. The generator and discriminator networks are trained using an optimization process that minimizes their respective loss functions, such as the Wasserstein loss with gradient penalty.

Chapter 6

Conclusions

We can draw the conclusion that the problem of adversarial samples in RS image classification is a serious issue that needs to be addressed based on the proposed methods and the experimental results of the thesis. Adversarial samples can cause deep learning models to produce incorrect predictions, leading to misclassification. This issue is particularly challenging for RS applications, as it can be caused by natural phenomena such as clouds, shadows, or artifacts in the satellite images.

To address this issue, we proposed four solutions in this thesis, including AL, query-by-example, physic-aware deep models, and GANs. AL can improve the model's performance by iteratively selecting the most informative samples to better understand the underlying structure and properties of the data, making it more resilient to adversarial attacks. Query-by-Example aims to find the most similar image to the query and optimize the network's weights so that adversarial samples are far away from the query image in the latent space. Physic-aware deep models can address the adversarial samples issue in RS image classification by incorporating domain-specific knowledge and physical properties of the data into the learning process. GANs can be used to generate adversarial samples for RS image analysis, which can then be used to improve the robustness of the classification model against adversarial attacks.

Overall, this thesis contributes to the development of more accurate, robust, and efficient algorithms for analyzing satellite images and understanding various environmental and geographical phenomena. By investigating these topics, we have proposed novel solutions to address the issue of adversarial samples in RS image classification, which can help to improve the reliability and security of deep learning models in RS applications.

6.1 Original contributions

In the face of adversarial challenges inherent in RS image classification, this thesis delves into innovative strategies to enhance the robustness and reliability of DL models for EO multispectral images. Adversarial samples, stemming from intentional modifications or natural perturbations, pose a significant threat to accurate image classification in RS applications.

In the first chapter, a meticulous comparison of nine AL strategies on two RS image datasets was conducted. We bring under the same framework and benchmark all strategies, among which, the most recent ones have not yet been tested on RS data. We assess the performance of AL techniques with regard to a new metric called LDN, shedding light on the efficiency of different strategies under limited labeling budgets. This investigation not only showcased the novel application of certain strategies to AL but also emphasized the growing importance of data abundance, particularly for DNNs. The historical evolution of AL in RS was also traced, emphasizing recent advancements, challenges unique to DL, and the potential of specialized network design for AL.

Moving to the second chapter, the focus shifted to guiding DNNs through the incorporation of statistical information. Two novel methods were introduced, addressing the limitations of Deep SVDD and demonstrating their efficacy in CBIR problems. The versatility of these methods across different DNN architectures was underscored, showcasing their potential to enhance data representation in various contexts. This chapter demonstrated that pushing anomaly samples away from the normal samples in a latent space model makes the model robust against outliers and any kind of adversarial samples.

The third chapter expanded on the significance of considering statistical properties and exploiting the WST to extract invariant features for improved classification of complex RS images. The study revealed the method's effectiveness in mitigating challenges such as intraclass diversity, pollution in training sets, and robustness to transformations. Additionally, the incorporation of physical properties into a cycle-consistent GAN addressed natural adversarial elements, providing a trustable solution for image analysis in the presence of clouds and shadows.

In the fourth chapter, the utilization of StyleGAN2-ADA for synthetic image generation was explored, focusing on ocean patterns in Sentinel-1A WV SAR images. The objective is to robust a classifier and improve the classification accuracy by training the model on GAN-based generated images. The study highlighted the dependence of FID on training size, the trade-off between image quality and diversity, and the importance of a diverse dataset for effective GAN training. The evaluation metrics used provided nuanced insights into the precision, recall, and limitations of GAN-generated images.

In summary, this thesis navigates the complex landscape of EO multispectral image analysis, addressing the persistent challenge of adversarial samples. The strategies explored, spanning AL, physics-aware models, Query-by-Example for image retrieval, and synthetic image generation by GANs, collectively contribute to the development of more accurate, robust, and efficient algorithms for the intricate task of satellite image analysis in diverse environmental and geographical contexts. As we move forward, these findings pave the way for advancements in RS applications, fostering a deeper understanding of Earth's dynamics through the lens of advanced DL techniques.

.

6.2 List of original publications

List of Journal papers

I. Omid Ghozatlou, Miguel Heredia Conde, Mihai Datcu. A Review and Perspective of Active Learning for Remote Sensing Image Analysis. In IEEE Geoscience and Remote Sensing Magazine (GRSM), Under revision.

II. Omid Ghozatlou, Miguel Heredia Conde, Mihai Datcu. Query by Example in Remote Sensing Image Archive Using Enhanced Deep Support Vector Data Description. In IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, vol. 16, pp. 1197-1210, 2023, DOI: 10.1109/JSTARS.2022.3233105.

List of Conference papers

I. Omid Ghozatlou, Mihai Datcu, Bertrand Chapron, GAN-Based Ocean Pattern SAR Image Augmentation. In IEEE International Geoscience and Remote Sensing Symposium (IGARSS), pp. 4056-4059, Pasadena, CA, USA, 2023, DOI: 10.1109/IGARSS52108.2023.10283353.

II. Omid Ghozatlou, Mihai Datcu. Comparative Studies on Similarity Distances for Remote Sensing Image Classification. In IEEE 5th International Conference on Image Processing Applications and Systems (IPAS), pp. 1-6, Genova, Italy, 2022, DOI: 10.1109/IPAS55744.2022.10052824.

III. Omid Ghozatlou, Miguel Heredia Conde, Mihai Datcu. Wavelet-Guided Deep Neural Network For Robust One-Class Classification. In 12th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS), pp. 1-5, Rome, Italy, 2022, DOI: 10.1109/WHISPERS56178.2022.9955071.

IV. Omid Ghozatlou, Mihai Datcu. Hybrid GAN and Spectral Angular Distance for Cloud Removal. In IEEE International Geoscience and Remote Sensing Symposium (IGARSS), pp. 2695-2698, Brussels, Belgium, 2021, DOI: 10.1109/IGARSS47720.2021.9554891.

6.3 Perspectives for further developments

Several key perspectives for further development include advanced AL strategies that integrate domain-specific knowledge. Employ adaptive query strategies that specifically consider spectral diversity and characteristics inherent to EO images. Another avenue involves combining diverse solutions to tackle adversarial samples in RS image classification. For instance, merging AL with GANs for synthetic data generation. GANs, in this context, serve not only to estimate model prediction uncertainty but also to prioritize uncertain samples for labeling, enhancing the model's robustness and confidence.

Our primary focus for future development centers on refining synthetic image generation using GANs while accounting for the physical properties of EO data. As detailed in Chapter 5, the generated images demonstrate realism, high quality, and similarity to real data. However, in the frequency domain, distinguishing GAN-generated images from real ones is possible, indicating GANs' challenges in learning spectral distributions. Addressing this, we aim to mitigate the spectrum discrepancy by training GANs in the frequency domain and incorporating physical properties into the model. This approach aims to enhance GANs' performance by better preserving the physical characteristics of the data.

Bibliography

[2] Lukas Ruff, Robert A. Vandermeulen, Nico Gornitz, Lucas Deecke, Shoaib A. Siddiqui, Alexander Binder, Emmanuel Muller, and Marius Kloft. Deep one-class classification. 35th International Conference on Machine Learning, ICML 2018, 10:6981–6996, 2018.

[3] Praveer Singh and Nikos Komodakis. Cloud-GAN: Cloud removal for sentinel-2 imagery using a cyclic consistent generative adversarial networks. International Geoscience and Remote Sensing Symposium (IGARSS), 2018 July:1772–1775, 2018.

[32] Xue Wang, Kun Tan, Cen Pan, Jianwei Ding, Zhaoxian Liu, and Bo Han. Active deep feature extraction for hyperspectral image classification based on adversarial learning. IEEE Geoscience and Remote Sensing Letters, 19:1–5, 2022.

[51] Juan Mario Haut, Mercedes E. Paoletti, Javier Plaza, Jun Li, and Antonio Plaza. Active learning with convolutional neural networks for hyperspectral image classification using a new bayesian approach. IEEE Transactions on Geoscience and Remote Sensing, 56(11):6440–6461, 2018.

[59] Guangxing Wang and Peng Ren. Hyperspectral image classification with feature-oriented adversarial active learning. Remote Sensing, 12(23), 2020.

[139] Ushasi Chaudhuri, Biplab Banerjee, and Avik Bhattacharya. Siamese graph convolutional network for content based remote sensing image retrieval. Computer Vision and Image Understanding, 184:22–30, 2019.

[156] Adrien Bardes, Jean Ponce, and Yann LeCun. VICReg: Variance invariancecovariance regularization for self-supervised learning. In International Conference on Learning Representations, 2022.

[211] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In Advances in Neural Information Processing Systems, volume 33, pages 12104–12114. Curran Associates, Inc., 2020.

[212] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 4401–4410, 2019.