

Distributed Deep Learning for Remote Sensing Data Interpretation

Juan M. Haut, *Senior Member, IEEE*, Mercedes E. Paoletti, *Senior Member, IEEE*, Sergio Moreno-Álvarez, Javier Plaza, *Senior Member, IEEE*, Juan A. Rico-Gallego and Antonio Plaza, *Fellow, IEEE*

Abstract—As a newly emerging technology, deep learning is a very promising field in big data applications. Remote sensing often involves huge data volumes obtained daily by numerous in-orbit satellites. This makes it a perfect target area for data-driven applications. Nowadays, technological advances in terms of software and hardware have a noticeable impact on Earth observation applications, more specifically in remote sensing techniques and procedures, allowing for the acquisition of data sets with greater quality at higher acquisition ratios. This results in the collection of huge amounts of remotely sensed data, characterized by their large spatial resolution (in terms of the number of pixels per scene), and very high spectral dimensionality, with hundreds or even thousands of spectral bands. As a result, remote sensing instruments on space borne and airborne platforms are now generating data cubes with extremely high dimensionality, imposing several restrictions in terms of both processing runtimes and storage capacity. In this paper, we provide a comprehensive review of the state-of-the-art in deep learning for remote sensing data interpretation, analyzing the strengths and weaknesses of the most widely used techniques in the literature, as well as an exhaustive description of their parallel and distributed implementations (with particular focus on those conducted using cloud computing systems). We also provide quantitative results, offering an assessment of a deep learning technique in a specific case study (source code available: <https://github.com/mhaut/cloud-dnn-HSI>). The paper concludes with some remarks and hints about future challenges in the application of deep learning techniques to distributed remote sensing data interpretation problems. We emphasize the role of the cloud in providing a powerful architecture that is now able to manage vast amounts of remotely sensed data due to its implementation simplicity, low cost and high efficiency when compared to other parallel and distributed architectures, such as grid computing or dedicated clusters.

Index Terms—Remote sensing, deep learning, parallel and distributed architectures, cloud computing, big data.

This work was supported by Junta de Extremadura under Grant GR18060, by Spanish Ministerio de Ciencia e Innovación through project PID2019-110315RB-I00 (APRISA), by the European Union’s Horizon 2020 research and innovation program under grant agreement No. 734541 (EOXPOSURE), and by the computing facilities of Extremadura Research Centre for Advanced Technologies (CETA-CIEMAT), funded by the European Regional Development Fund (ERDF). CETA-CIEMAT belongs to CIEMAT and the Government of Spain.

J.M. Haut is with the Department of Communication and Control Systems, Higher School of Computer Engineering, National Distance Education University, 28015 Madrid, Spain (e-mail: juanmariohaut@unex.es).

M.E. Paoletti, J. Plaza and A. Plaza are with the Hyperspectral Computing Laboratory, Department of Technology of Computers and Communications, Escuela Politécnica, University of Extremadura, 10003 Cáceres, Spain (e-mail: mpaoletti@unex.es; jplaza@unex.es; aplaza@unex.es).

Sergio Moreno-Álvarez and Juan A. Rico-Gallego are with the Department of Computer Systems Engineering and Telematics, University of Extremadura, 10003 Cáceres, Spain (email: smoreno@unex.es; jarico@unex.es).

I. INTRODUCTION

A. Big remote sensing data

Remotely sensed images provide very detailed information about the surface of the Earth, which often results in very significant computational requirements. Present and future missions for Earth Observation (EO) have significantly increased the spatial, spectral and temporal resolution of existing imaging instruments, resulting in higher data volumes. Meanwhile, the variety of multi-sensor and multi-resolution acquisitions inevitably lead to the gradual acceptance of the generated data sets as “big remote sensing data” [1], not merely due to their high volume, but also due to the inherent complexity of the data, which calls for advanced processing techniques that often use external sources of information (e.g., social media data) for adequate interpretation [2].

For instance, remotely sensed hyperspectral images [3] record information using hundreds of spectral bands collected at nearly contiguous wavelengths in the electromagnetic spectrum. This significantly increases their volume with regards to other kinds of image data used in remote sensing applications, such as multi-spectral images (tens of bands), radar, or microwave data sets, creating important requirements in terms of storage and processing. In fact, there has been an exponential growth in these requirements due to recent technological advances in both the quality and number of available imaging instruments. This results from the ever-increasing number of EO missions, that are now generating a nearly continuous flow of remotely sensed data, which fostered the creation of large remote sensing data repositories that can only be exploited using adequate parallel and distributed processing techniques [4]. Only in the hyperspectral domain, the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) [5] and its new generation (AVIRIS-NG) –operated by NASA’s Jet Propulsion Laboratory– acquires almost 9 gigabytes (GBs) of data per hour. Similarly, the EO-1 Hyperion sensor acquires about 71.9 GBs of data per hour, which means over 1.6 terabytes (TBs) of data per day. Most of the EO missions that will be operational soon, e.g. the German Environmental Mapping and Analysis Program (EnMAP)¹, exhibit equal or even higher data acquisition rates. This confirms that remote sensing data has entered the “big data era” [6]. To be more specific, we summarize below the properties that qualify remote sensing data a kind of big data [7]:

- *Data volume.* As mentioned before, remote sensing data in the optical, radar and microwave domains are now

¹<http://www.enmap.org/>

characterized by their huge dimensionality, which results in the acquisition of several TBs of data per day. The total amount of data archived by the Earth Observing System Data and Information System (EOSDIS)² now exceeds 30 petabytes (PBs) of data. The amount of archived data at the China National Satellite Meteorological Center (NSMC)³ exceeds 5 PBs, and the China Center for Resources Satellite Data and Application (CCRSDA)⁴ has more than 20 PBs of remote sensing data in its archive.

- *Data variety.* According to the state of the satellite industry report⁵, there are more than 300 EO satellites currently in orbit. All of them carry at least one EO instrument, and they are able to collect and transmit data continuously to Earth stations. This means that hundreds of different kinds of remotely sensed data are transmitted (in parallel) to the respective ground receiving stations every day. In addition, since Landsat-1 was first operational in 1972, there have been more than 500 EO satellites launched into space, collecting and archiving more than 1000 different types of remote sensing data.
- *Data velocity.* With the development of satellite constellations, the satellite revisit times have gradually transitioned from months to days, hours, or even minutes. As a result, the multi-temporal resolution of remote sensing data has increased exponentially, allowing for advanced environmental monitoring and climate change applications. As mentioned before, data centers now receive an almost continuous flow of remote sensing data at ever-increasing speeds, which creates important requirements in terms of storage and analysis (particularly in the context of real-time applications).

In addition to the three aforementioned characteristics, commonly known as the “three V’s” of remote sensing big data [6], there are also other important aspects inherent to the analysis and interpretation of such data. One of the most important ones is data *heterogeneity*. Specifically, due to the existence of various satellite orbits and specifications for different sensors (with different storage formats, data projections, spatial resolutions, and revisit times), there are vast differences in the formats of the archived data, and these differences create difficulties when developing general data interpretation techniques. Currently, big remote sensing data analysis is attracting significant attention from governmental and commercial partners, as well as from academic institutions. This is because the high-level products that can be obtained from the data are useful in many relevant applications, including agriculture, disaster prevention and reduction, environmental monitoring, public safety and urban planning, among many others [8].

One of the most important remote sensing big data projects has been the EOSDIS [9], which provides end-to-end capabilities for managing NASA’s Earth science data from various sources. In Europe, the European Space Agency (ESA) has been organizing the “Big Data from Space” conference, with

²<https://earthdata.nasa.gov/eosdis>

³<https://www.nsmc.org.cn/en/>

⁴<https://data.globalchange.gov/organization/china-center-resources-satellite-data-application>

⁵<https://sia.org/news-resources/state-of-the-satellite-industry-report/>

the ultimate goal of stimulating interactions and bringing together partners and service providers willing to exploit and interpret remotely sensed big data collected from space. The Group on Earth Observations (GEO) [10], a large-scale co-operation organization, also promotes the development of big data in remote sensing applications. As for commercial applications, Google Earth⁶ is a success case of bringing remotely sensed big data to a large number of users around the world. Many remote sensing applications such as target detection, land-cover classification, spectral unmixing, pansharpening, etc. can now be developed easily by resorting to Google Earth and advanced processing algorithms. At the academic level, we have also seen important efforts in top journals launching multiple special issues devoted to the processing and analysis of remotely sensed big data [1], [6]. Despite these significant advances, the processing of remotely sensed big data still faces significant challenges that we summarize below:

- *Data integration challenges.* A unified data standard is needed for heterogeneous remote sensing data integration. This includes uniform data standards, metadata standards and image standards. However, due to the massive, multi-modal and heterogeneous nature of big remote sensing data, this is a challenging and yet unaccomplished goal.
- *Data processing challenges.* How to design computationally efficient and application-specific data processing and storage techniques (while providing unified interfaces to simplify the access to distributed collections of big remote sensing data) is also a pressing challenge. In a computing system, data transmission is generally the bottleneck due to the limited network bandwidth [8]. Also, the dependency between tasks may introduce ordering constraints, and the optimized scheduling of these tasks may be critical to achieve satisfactory processing performance.

B. Cloud computing in remote sensing

Currently, cloud computing⁷ platforms are increasingly being used to process and store remotely sensed big data in distributed architectures [7]. The explosive growth of remote sensing big data has revolutionized the way these data are managed and processed. Still, important challenges remain due to the complex management of multi-modal, multi-spectral, multi-resolution and multi-temporal remote sensing data sets, which appear in various formats and/or distributed across data centers. In this regard, cloud computing (based on virtualization technology) offers the potential to integrate computing, storage, network and other physical resources to build a virtual resource pool where advanced techniques for remote sensing data processing can be developed and deployed. In other words, the cloud provides users with services to integrate data, processing, production, computing platforms, storage, and integrated spatial analysis, so as to provide solutions in different application domains such as environmental problems,

⁶<https://www.google.com/intl/es/earth/>

⁷A formal widely accepted definition of cloud computing can be found in [11].

land-use/land-cover, urban planning and so on. Cloud computing technology also offers advanced capabilities for service-oriented and high-performance computing (HPC). The use of cloud computing for the analysis of large repositories of remote sensing images is now considered a natural solution, resulting from the evolution of techniques previously developed for other types of computing platforms, such as commodity clusters or grid environments [12].

Quite surprisingly, there are not too many works yet describing the use of cloud computing infrastructures for the implementation of remote sensing data processing techniques. This is partially due to the lack of open repositories containing labeled remote sensing images for public use, a situation that is now changing thanks to initiatives such as BigEarthNet⁸. Also, NASA and ESA are now providing large distributed repositories of remote sensing data for open use by the scientific community (e.g., the Sentinel program⁹). Due to the availability of such repositories, the development of techniques based on cloud computing for distributed processing of remote sensing images has become a very timely research line.

A relevant question at this point is whether cloud computing can become a *de-facto* architecture for remotely sensed data interpretation in future years. Our belief is that, by virtue of its elasticity and high transparency levels, cloud computing offers a truly unique paradigm for big remote sensing data processing, in which computational resources can be accommodated in the form of ubiquitous services on-demand, on a pay-per-use basis. In addition, cloud-enabled remote sensing data processing infrastructures and services can now be delivered for large-scale remote sensing data across geographically distributed data centers, which was impossible even with the most powerful compute clusters. As a result, the incorporation of the cloud to big remote sensing data processing initiatives reveals its capacity to deal with the increased computational and storage challenges introduced by modern remote sensing applications, especially when coupled with the powerful new deep learning (DL) algorithms [13] that have been shown to provide an excellent tool for information extraction from scientific data in general and from remotely sensed data sets in particular [14].

C. Paper organization and contributions

In this paper, we provide a review of the most important initiatives that have been developed so far in the use of cloud computing architectures (as compared with other HPC solutions, such as commodity clusters or grid computing platforms) for remote sensing data interpretation, with particular focus on DL techniques and their implementations in the cloud. We believe that this review is quite unique in the sense that it provides a completely new flavour with regards to other published works that focus on DL in remote sensing [15], [16], [17], big remote sensing data [1], [6], or high performance computing in remote sensing [4], [18], [19]. None of these review papers considered cloud computing as a *de-facto* architecture for big remote sensing data processing, which is now a widespread

implementation option (in particular, when computationally demanding DL algorithms are involved in the information extraction process). As a result, we believe that this review is necessary given the recent advances in processing strategies, which inevitably led to using DL algorithms in the cloud for the successful processing and interpretation of big remote sensing data sets. The remainder of the paper is structured as follows:

- Section II provides a general overview of techniques for DL in remote sensing data processing, and a taxonomy of DL architectures that have been widely used in this context.
- Section III provides a comprehensive review of available approaches for the efficient implementation of remote sensing data processing techniques based on DL algorithms in HPC architectures, including clusters, grids and cloud computing systems. This section also includes a discussion and some critical observations resulting from the in-depth analysis of the works published so far in these areas.
- Section IV focuses on cloud computing as the current *de-facto* architecture for big remote sensing data processing using DL algorithms. This section first provides some basic concepts about cloud computing. Then, it details some of the most popular frameworks and programming models that have been used for DL-based processing of remotely sensed data and other scientific applications in the cloud.
- Section V describes the most popular machine learning (ML) and DL libraries and frameworks in cloud computing environments, with particular emphasis on those that have been already used in remote sensing applications.
- Section VI provides a case study with a processing example that illustrates a representative technique for DL-based distributed processing of remotely sensed hyperspectral data in the cloud, providing also some suggestions for practical use and exploitation.
- Finally, section VII concludes the paper with some remarks and hints at plausible future research avenues.

II. DEEP LEARNING IN REMOTE SENSING

A. Remote sensing data processing

Remote sensing technology now provides high quality data from the surface of the Earth (in terms of detailed resolution, good signal-to-noise ratio, robustness to perturbations, and accurate error corrections). Advanced analysis and interpretation methods are required to extract the most useful information contained in the data. As a result, the remote sensing discipline involves many Earth science disciplines, such as meteorology, geology or ecology, as well as a variety of engineering skills to properly interpret the huge amount of remotely sensed data provided by a constellation of air/space EO instruments.

Furthermore, a wide variety of remotely sensed data can be obtained from these instruments, where optical imaging systems capturing reflected sunlight are pretty popular due to the rich information they contain, with different formats and resolutions (spatial, spectral and temporal), enabling a very

⁸<http://bigearth.net/>

⁹<https://sentinel.esa.int/>

detailed and comprehensive assessment of surface properties [20], [21]. For illustrative purposes, Table I lists some popular remote sensing instruments which have been or are currently operational.

TABLE I
LIST OF POPULAR REMOTE SENSING INSTRUMENTS.

GSD: GROUND SAMPLE DISTANCE, PAN: PANCHROMATIC, MSI: MULTISPECTRAL, HSI: HYPERSPECTRAL

Name	Type	Spatial resolution	Spectral bands	Spectral Range
RADARSAT-2 [22]	SAR	1-100 m	1 (C-band)	5.405 GHz
ALOS PALSAR [23]	SAR	6.25-12.5 m	1 (L-band)	1.27 GHz
AirMOSS [24]	SAR	100 m	1 (P-band)	0.43 GHz
TerraSAR-X [25]	SAR	0.24-40 m	1 (X-band)	19.65 GHz
Sentinel-1 [26]	SAR	5-20 m	1 (C-band)	5.404 GHz
LVIS [27]	Laser	20 m	1	1064 nm
EROS-B [28]	PAN	0.7m	1	500-900nm
GAOFEN-1 [29]	MSI	2.0-8 m	5: PAN Blue Green Red NIR	450-900 nm 450-520 nm 520-590 nm 630-690 nm 770-890 nm
GAOFEN-2 [29]	MSI	0.81-3.24 m	5: PAN Blue Green Red NIR	450-900 nm 450-520 nm 520-590 nm 630-690 nm 770-890 nm
IKONOS [30]	MSI	0.82-3.2 m	5: PAN Blue Green Red NIR	526-929 nm 445-516 nm 506-595 nm 632-698 nm 757-853 nm
WorldView-4[31]	MSI	0.3-1.2 m	5: PAN Blue Green Red NIR	450-800 nm 450-510 nm 510-580 nm 655-690 nm 780-920 nm
Sentinel-2 [32]	MSI	10-60 m	13	443-2190 nm
Sentinel-3 OLCI [33]	MSI	300-1200 m	21	400-1200 nm
Landsat-8 [34]	MSI	15-30 m	11	300-12510 nm
MODIS [35]	MSI	250 - 1000 m	36	405-14385 nm
CHRIS [36]	HSI	18-36 m	62	415-1050 nm
AVIRIS [5]	HSI	20 m	224	360-2450 nm
AVIRIS-NG [37]	HSI	0.3-4.0 m	600	380-2510 nm
ROSIS [38]	HSI	1.3 m	115	430-860 nm
CASI [39]	HSI	2.5 m	114	360-1050 nm
HYDICE [40]	HSI	1-7 m	210	400-2500 nm
HYMAP [41]	HSI	5 m	126	450-2500 nm
PRISM [42]	HSI	2.5 m	248	350-1050 nm
EnMAP [43]	HSI	30 m	228	420-2400 nm
HISUI [44]	HSI	30 m	185	400-2500 nm
DESIS [45]	HSI	30 m	180	400-1000 nm
HYPERION [46]	HSI	30 m	220	400-2500 nm
PRISMA [47]	HSI	30 m	237	400-2500 nm
SHALOM [48]	HSI	10 m	241	400-2500 nm

In fact, optical remote sensing data play an important role in many different activities [49]. On the one hand, panchromatic, standard RGB and multi-spectral products often exhibit impressive spatial resolution, which strongly facilitates the detection of contours, textures and structures within the scene. On the other hand, although the spatial resolution is partially sacrificed, hyperspectral instruments collect hundreds of narrow and near-continuous bands ranging from the VNIR to the shortwave infrared (SWIR) wavelength regions, providing very detailed spectral signatures of the materials covered by each pixel, thus allowing a more accurate and detailed analysis of the spectral features available in the data [50].

The advanced products resulting from the analysis and interpretation of remotely sensed data sets allow for the implementation of long-term development strategies in several fields, such as precision agriculture [51], urban planning [52], or desertification monitoring [53], among others.

The rich and detailed information contained in remote sensing data needs to be adequately extracted and processed to be consequently exploited at the user level. In this regard, there is

a strong demand for an accurate and computationally efficient analysis techniques, which can be categorized considering different criteria. In particular, according to their purpose [54], we can classify available techniques into the following groups:

- *Restoration and denoising methods* manage data corruption and anomalies introduced during the acquisition process that may significantly degrade the quality of the collected data [55], involving radiometric and geometric corrections, or diffuse solar radiation and management of atmospheric effects.
- *Enhancement methods* transform the captured data to increase the quality of certain features, making them more suitable to human vision skills to conduct visual analysis. This involves contrast enhancement, super-resolution and pan-sharpening, for instance.
- *Transformation methods* modify the scene content in either the spectral or spatial domain for feature extraction, image compression or filtering purposes, such as the principal component analysis (PCA), tasseled cap transformation (TCT), vegetation indices or the wavelet transform (WT).



Fig. 1. Remote sensing classification can be tackled at three different levels: pixel-level (left), object-level (center), and scene-level (right)

- *Classification methods* interpret the content of remotely sensed scenes. Three classification levels can be distinguished (see Fig. 1) [56], where pixel-level classification labels each pixel in a scene with a semantic class [57]; object-level classification seeks to recognize the elements present in the scene (by usually combining spectral-spatial features) [58]; and scene-level classification provides a global meaning (i.e., a semantic class) to the entire scene by understanding and interpreting its features [59]. Moreover, sub-pixel analyses can also be conducted by analyzing the spectral mixtures at each pixel (termed sub-pixel classification or spectral unmixing) [60], [61], [18].

Focusing on data transformation and classification approaches, ML and DL methods have provided a wide range of processing algorithms for both regression and classification of complex nonlinear systems [62], [63], [14], [16], implementing promising learning paradigms to derive information from the data. These methods range from purely unsupervised strategies to supervised ones, with a vast collection of semi-supervised and hybrid-based methods in between [64], [65], [66], [67].

For instance, k -means clustering is a popular unsupervised method that groups similar samples together by exploring similarity measures, and is able to discover underlying patterns [68]. On the contrary, the k -nearest neighbors (k NN) usually explores the similarity between samples in a supervised way [69]. Also, decision trees (DTs) [70] and random forests (RFs) [71] are supervised methods, where RFs develop multiple trees from the randomly sampled subspace of the input sample and then combine the output through a voting/maximum rule. In addition, the multinomial logistic regression (MLR) [72], Gaussian mixture models (GMMs) [73] and naive Bayes-based (NB) approaches [74] are probabilistic models that analyze the data distribution to conduct their assumptions. Hidden Markov models (HMMs) [75] and support vector machines (SVMs) [76] are accurate statistical classifiers. Particularly, the SVM is considered an efficient and stable algorithm for high-dimensional data classification. This method learns the decision hyperplane that can best separate training samples in a kernel-included high dimensional feature space. Finally, artificial neural networks (ANNs) [77] are versatile empirical-modeling algorithms composed of hierarchical layers of neurons that process input stimuli using synaptic weights, and transmit their responses through activation functions. As a result, each layer refines the neural responses to the input data, obtaining increasingly abstract representations by adjusting the model weights, which are automatically learned from the data through the forward-backward propagation mechanism to extract the most relevant information. Moreover, ANNs offer a very flexible architecture in which both the number of layers and neurons (and even the shape and direction of the connections) can be established by the programmer. In this sense, the great flexibility and automatic adjustment of neural models (without any prior knowledge about the statistical features of the data) are major advantages that have positioned ANNs as a very attractive approach, creating a sharp contrast to traditional ML techniques, which usually require careful engineering to extract complex handcrafted features, requiring specific knowledge to recognize the specific regularities present in the data.




Fig. 2. Total number of papers per year (a) and citations received by papers (b) in the area “DL in remote sensing”. Source: Web of Knowledge. Search string: AB=(“deep learning” AND “remote sensing”). Total number of results: 1621. Date of the query: February 20, 2021.

Indeed, the study and implementation of ANNs is so extensive that, within ML, the sub-field of DL has emerged as a hot topic for signal data processing [13]. Particularly, DL algorithms have gained significant popularity in remote

sensing data analysis over the past few years. For illustrative purposes, Fig. 2 shows the total number of papers per year published in this topic, and also the number of citations received, revealing an exponential increase in recent years. The figure was generated using the Web of Knowledge engine¹⁰, and the exact search string used (and the date of the query) are specified in the figure caption—where “AB” indicates that the search was conducted in the abstracts of the papers—for reproducibility purposes.

A detailed review of available approaches in this area was given in [15], which discusses how DL has been applied for remote sensing data analysis tasks such as image fusion, registration, scene classification, object detection, land use and land cover classification, segmentation, etc. Different application fields are also covered, including open challenges, and directions for future research. Also, the paper [17] focuses on remotely sensed hyperspectral data, providing a comprehensive review of methods for DL-based classification in this field and discussing the strengths and weaknesses of these methods. Accordingly, the following section includes a brief taxonomy of the main DL models developed for remote sensing data analysis.

B. Taxonomy of deep learning (DL) architectures

There is a great variety of deep models due to the great flexibility of existing architectures in terms of topology, data-path connections and types of layers. In general terms, the scientific community recognizes five models, i.e., stacked autoencoders (SAEs), deep belief networks (DBNs), recurrent neural networks (RNNs) and convolutional neural networks (CNNs), as the main architectures, from which a great variety of modified networks have been implemented [17], such as generative adversarial networks (GANs), which have gradually become a mainstream architecture in the field of remote sensing. In the following, we review these DL models.

1) *SAEs*: The auto-encoder (AE) implements an encoder-decoder structure to learn a code representation from the input data in an unsupervised way. It defines an optimization problem that attempts to minimize the reconstruction error $\|(\mathbf{W}_d \sigma(\mathbf{W}_e \mathbf{X}^T))^T - \mathbf{X}\|_2^2$ by learning the matrices of bases \mathbf{W}_d and \mathbf{W}_e , where $\mathbf{X} \in \mathbb{R}^{N \times B}$ defines the remote sensing data as an input matrix of N samples with B channels (feature space), $\mathbf{W}_e \in \mathbb{R}^{B' \times B}$ comprises the recognition weights of the encoder, which maps the input data to a code/latent representation with $B' \neq B$ features (code/latent space), and $\mathbf{W}_d \in \mathbb{R}^{B \times B'}$ comprises the generative weights of the decoder, which recovers the original feature space by reconstructing the input data. σ acts as an activation function. The SAE deepens the model by stacking several AEs [see Fig. 3(a)], where the AEs of the stacked-encoder (bottom-half) find a series of lower-dimensional features and the stacked-decoder (top-half) performs the opposite function [78].

2) *DBNs*: The DBN is a multi-layer generative model inspired by restricted Boltzmann machines (RBMs). An RBM is a two-layer stochastic network trained to minimize the input reconstruction error in a similar way as AEs (*Gibbs sampling*),

¹⁰<http://www.webofknowledge.com/>

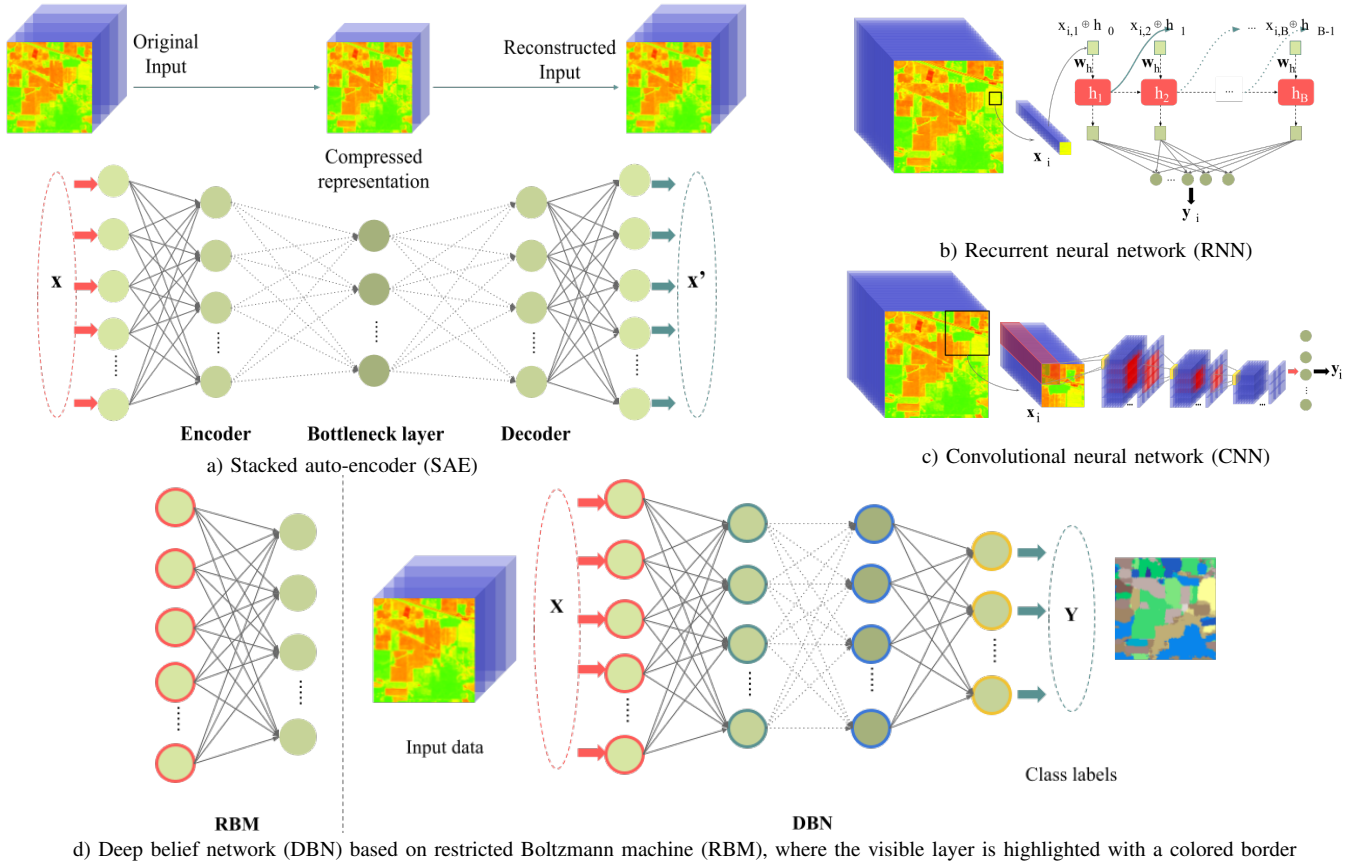


Fig. 3. Graphical illustration of traditional deep network architectures applied for processing remote sensing data cubes. The SAE (a) optimizes the reconstruction error between its vector input \mathbf{X} and its output \mathbf{X}' , where the bottleneck layer contains the latent data representation. In the traditional RNN model (b), the data cube is processed in band-by-band fashion, where the spectral signature of each pixel x_i is processed as a temporal sequence, obtaining as a result a hidden state \mathbf{h} that works as the model memory. The CNN (c) processes 3D inputs (i.e., the pixel x_i and its surrounding area) by applying multidimensional kernels that act as filters for particular features (borders, shapes, etc.), obtaining a set of feature maps with the neuronal responses to that stimuli. Finally, the DBN (d) is composed of several RBMs which process and reconstruct the input data, mimicking the SAE behaviour to learn the probability distribution of the input in an unsupervised fashion.

where the visible layer deals with the input data and the hidden layer conducts feature extraction, capturing higher-order data correlations observed in the visible units while learning a probability distribution over the input data [see Fig. 3(d)]. DBNs take advantage of RBMs, concatenating several pre-trained RBMs and refining the full-model parameters through labeled data.

3) *RNNs*: The RNN [see Fig. 3(b)] retains memory and learns data sequences by introducing loops in its connections. As a result, the neural responses in each step depend on those of the previous step by means of an internal state, which creates a sequential dependency that provides an association between the current and the previous data sample. According to which hidden states are created and how they are managed, three types of RNNs can be distinguished [79]: the simple vanilla RNN, the gate-based long short memory (LSTM), and the simplified gated recurrent unit (GRU).

4) *CNNs*: In contrast to other deep models (which were originally implemented with fully connected layers, e.g., the AE), the CNN introduces the convolution layer as a set of locally connected weights which are rearranged in an n -dimensional grid. As a result, the convolution kernels act as feature detection-extraction filters, where neuronal responses

are arranged in a feature map that not only indicates the presence of a particular stimulus detected by the kernel (edges, borders, shapes) but also the location of these stimuli in the spatial domain. This enables abstract and refined spatial relationships to be maintained and extracted within the data through a hierarchical stack of convolution layers [see Fig. 3(c)], which are combined with other layers (activation, normalization and pooling functions, for instance) to extract elaborate patterns from the raw inputs.

The flexibility of convolution kernels has demonstrated a great potential to extract any kind of feature from the raw data, without applying complex pre-processing mechanisms. Furthermore, the great architectural plasticity in terms of kernel size and grid organization (which produces 1D, 2D and 3D models), layer connections (direct, residuals, skip and short-connections), data paths settings and wide/depth configuration, along with the impressive generalization power and the ability to make strong assumptions about the input data, have established convolution-based models as the most successful and popular deep networks. In fact, these networks represent the state-of-the-art in image processing through derived network models, such as residual (ResNets), dense (DenseNets) and capsule-based (CapsNets), among many others [80], [81].

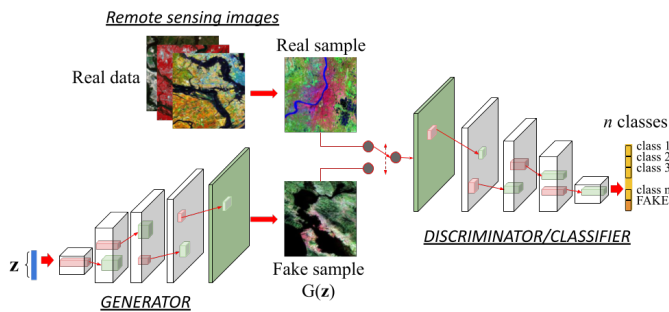


Fig. 4. Generative adversarial network (GAN) for remote sensing data processing

5) *GANs*: The aforementioned networks work as discriminative models, which map original inputs to some desired outputs (by learning conditional distributions between them) to minimize a loss function. In contrast, generative approaches (such as GANs) learn the joint probability between inputs and outputs, modeling the data distribution to generate new samples rather than just evaluating the available ones. Thus, GANs (see Fig. 4) model a data distribution from a random noise vector through an adversarial process, where two neural models, i.e. *generative* and *discriminative* networks, are simultaneously trained in competition (the former to deceive the latter, and the latter to avoid being deceived by the samples generated by the former).

III. IMPLEMENTATIONS

In typical DL models, such as those illustrated in Fig. 3, there are millions of parameters (which define the model) and large amounts of data are required to learn these parameters. This leads to a computationally intensive process in which the learning step consumes a lot of time. Therefore, it is important to come up with parallel and distributed algorithms which can run much faster and drastically reduce the training times in the context of remote sensing applications. In the following, we provide a description of different parallelization strategies for accelerating DL algorithms in remote sensing applications by resorting to three types of HPC architectures: clusters, grids and clouds. A description of the main challenges faced by these different architectures, along with a brief comparison among them is then presented. The section concludes with a discussion on their potential role in solving remote sensing problems via DL algorithms.

A. Cluster computing

In this subsection, we describe some of the most relevant approaches in the recent literature to exploit cluster computer architectures (including GPU clusters) for the efficient interpretation of remote sensing data. Fig. 5 shows the total number of papers published in this area, along with the number of citations received (according to Web of Knowledge). In the following, we discuss some of the most relevant contributions in this field.




Fig. 5. Total number of papers per year (a) and citations received by papers (b) in the area "deep learning in remote sensing using cluster computing". Source: Web of Knowledge. Search string: AB=("deep learning" AND "remote sensing" AND "cluster" AND "comput*"). Total number of results: 11. Date of the query: February 20, 2021.

As one of the most notable recent developments, the authors in [82] implement a remote sensing data flow (RESFlow) for improving DL algorithms and allowing them to perform computations on large-scale remotely sensed images. The RESFlow works by dividing the data into homogeneous partitions that can fit simple models in homogeneous (i.e., commodity cluster-based) machines. Despite its cluster-oriented nature, RESFlow uses Apache Spark (a tool that has been widely used in cloud computing, as described in section IV) to accelerate DL inference. The RESFlow incorporates a strategy to optimize resource utilization across multiple executors assigned to a single worker. The framework invokes DL inference at three stages: during deep feature extraction, deep metric mapping, and deep semantic segmentation. Resource sharing in graphical processing units (GPUs) is adopted to achieve a fully parallelized pipeline for all execution steps.

RESFlow uses Apache Spark, but there are multiple options to distribute training over cluster computing implementations. Other options include TensorFlow¹¹, PyTorch¹² or Horovod¹³. These frameworks provide multiple benefits. For example, Pytorch includes multiple extensions (as NVIDIA Apex¹⁴) to enable streamline mixed precision with distributed training, and Horovod employs efficient inter-GPU communications. Similar to Apache Spark, some cluster computing approaches take advantage of Kubernetes¹⁵ (k8s) architecture workflow, which allows for the deployment automation, scaling, and management of ML/DL applications, as described in [83].

The work in [84] exemplifies the unique advantages provided by parallel computing environments and programming techniques to solve large-scale problems, such as the training of classification algorithms for remote sensing data. Specifically, the authors demonstrate that the training of deep CNNs can be efficiently implemented using cluster computers containing a large number of GPUs. The obtained results confirm that parallel training can dramatically reduce the amount of time needed to perform the full training process, obtaining almost linear scalability in a cluster of GPUs without losing any test accuracy.

¹¹<https://www.tensorflow.org/>

¹²<https://pytorch.org/>

¹³<https://horovod.ai/>

¹⁴<https://nvidia.github.io/apex/>

¹⁵<https://kubernetes.io/es/>

At this point, we emphasize that some works do not need to exploit clusters of GPUs to conduct the desired calculations. For instance, in the work [85], the authors resort to a shared memory system with only 4 GPUs to develop a Multi-GPU Training Framework (MTFC) of a CNN for remotely sensed hyperspectral image classification. The authors first develop a parallel neighbor pixel extractor (PNPE) that generates 3D cube samples from the input data automatically. Then, they perform a series of optimizations in the MTFC, such as task division, fine-grained mapping between tasks and GPU thread blocks and shared memory usage reduction. To further improve the training speed, the authors exploit CUDA streams and multiple GPUs to train mini-batches of data samples simultaneously. The MTFC is shown to outperform popular ML frameworks such as Caffe¹⁶ and Theano¹⁷, while offering the same level of classification accuracy.

The paper [86] provides several parallelization approaches for deep neural networks (DNNs), taking into account network overheads as well as optimal resource allocation, since network communication is often slower than inter-machine communication (while some layers are more computationally expensive than others). Specifically, the authors consider a multi-modal DNN architecture and identify several strategies to optimize performance when training is accomplished on Apache Spark (this framework will be described in details in section IV). The authors compare their newly developed architecture with an equivalent DNN architecture modeled after a data parallelization approach. The experiments in the paper reveal that the way in which the model is parallelized has a very significant impact on resource allocation, and that hyperparameter tuning can significantly reduce network overheads.

Other relevant developments in this area include the paper [87], which presents parallel versions of DL techniques for dimensionality reduction of remotely sensed images, also implemented in an Apache Spark cluster. The paper [88] presents an improved version of the aforementioned development which scales even better in clusters of GPUs. The paper [89] presents a parallel and distributed DL-based spectral unmixing algorithm for remotely sensed hyperspectral data, again using Apache Spark for the implementation on a cluster computer.

B. Grid computing

Although many parallel systems are inherently homogeneous, a most recent trend in HPC systems is to use highly heterogeneous computing resources, where the heterogeneity is generally the result of technological advancement in progress of time. With increasing heterogeneity, grid computing emerged as a premier technology that could facilitate the processing of remote sensing data in heterogeneous and distributed computing platforms.

Although the grid has recently evolved into architectures with more quality of service, such as the cloud, there were several reasons for using grid computing for remote sensing

image processing when the first grid-oriented architectures appeared. First and foremost, the required computing performance may not be available locally. Also, the required performance may not be available in just one location, with a possible solution being cooperative computing. Last but not least, the required computing services may be only available in specialized centres, and in this case the solution is application-specific computing. This led to the development of some grid-based approaches that are now mostly transitioning into cloud computing implementations, as will be described in the following subsection.

For illustrative purposes, Fig. 6 shows the total number of papers published in this area, along with the number of citations received. In the following, we discuss some of the most relevant contributions focused on using grid computing platforms for solving remote sensing problems via DL algorithms.




Fig. 6. Total number of papers per year (a) and citations received by papers (b) in the area “deep learning in remote sensing using grid computing”. Source: Web of Knowledge. Search string: AB=(“deep learning” AND “remote sensing” AND “comput*” AND “grid”). Total number of results: 9. Date of the query: February 20, 2021.

The GEOGrid project was one of the first DL-oriented initiatives aimed at providing an e-Science infrastructure to the remote sensing community. It is specifically developed to integrate a wide variety of remote sensing data sets, and is accessible online as a set of services¹⁸.

The platform called grid processing on demand, shortly G-POD¹⁹, was the first one to provide a grid-based environment for processing satellite images provided by ESA, offering several image processing services and DL algorithms mainly intended for environmental studies. G-POD has been successfully applied in real applications, like flood area detection.

As an add-on to this tool, the platform for satellite imagery search and retrieval, called Ground European Network for Earth Science Inter-operations and Digital Repositories (GENESI-DR) [90], offers an advanced interface for digital data discovery and retrieval, where the original images are processed using G-POD facilities (comprising some DL algorithms). The ultimate goal of GENESI-DR was to build an open access service to digital repositories focusing on fast search, discovery, and access to remotely sensed imagery in the context of post-disaster damage assessment. Once a disaster alert has been issued, response time is critical to providing relevant damage information to analysts and/or stakeholders.

¹⁶<https://caffe.berkeleyvision.org/>

¹⁷<https://github.com/Theano/>

¹⁸<https://www.geogrid.com/>

¹⁹<https://gpod.eo.esa.int/>

In this regard, GENESI-DR provides rapid area mapping and near real-time ortho-rectification web processing services to support post-disaster damage needs.

Also, the GiSHEO²⁰ platform (on-demand Grid services for training and high education in Earth Observation) addresses an important need for specialized training services in Earth observation. Solutions were developed for data management, image processing service deployment, workflow-based service composition, and user interaction, with particular attention to services for image processing (able to exploit free image processing tools, along with some DL techniques). A special feature of the platform is the connection with the GENESI-DR catalog, which provides plenty of remote sensing data sets for free.

To conclude this subsection, it is important to emphasize that the Committee on Earth Observation Satellites (CEOS)²¹, an international coordinating body for space borne missions focused on the study of the Earth, maintains a working group on information systems and services, with the ultimate goal of promoting the development of inter-operable systems for managing EO data internationally. In this regard, several grid platforms have greatly benefited from CEOS standards when developing grid-based tools for accurate interpretation of remotely sensed data [7].

C. Cloud computing

Cloud computing solutions represent an evolution of grid-based approaches and exhibit the potential to manage and process vast amounts of remotely sensed data in fault-tolerant environments by interconnecting distributed and specialized nodes. This strategy can significantly reduce the processing costs often involved in grid computing, leading to natural and cheap solutions for remotely sensed data processing. For illustrative purposes, Fig. 7 shows the total number of papers published in this area, along with the number of citations. In the following, we review some of the most significant contributions based on cloud computing for solving remote sensing problems via DL techniques.




Fig. 7. Total number of papers per year (a) and citations received by papers (b) in the area “deep learning in remote sensing using cloud computing”. Source: Web of Knowledge. Search string: AB=(“deep learning” AND “remote sensing” AND “comput*” AND “cloud”). Total number of results: 78. Date of the query: February 20, 2021.

One of the most relevant works addressing the implementation of DL algorithms for remote sensing data analysis in the

²⁰<http://cgis.utcluj.ro/projects/gisheo>

²¹<https://ceos.org/>

cloud was presented in [91], in which the authors introduced a new cloud-based technique for spectral analysis and compression of hyperspectral images. Specifically, the authors provide a cloud implementation of the AE, a popular deep network for non-linear data compression. Apache Spark (described in detail in section IV) serves as the backbone of the cloud computing environment by connecting the available processing nodes using a master-slave architecture. The obtained results indicate that cloud computing architectures offer an adequate solution for compressing and interpreting big remotely sensed data sets.

The paper [92] proposes an acceleration method for hyperspectral image classification that exploits scheduling meta-heuristics to automatically and optimally distribute the workload across multiple computing resources on a cloud platform. A representative DL-based classification processing chain is first distributed and implemented in parallel based on the MapReduce mechanism (described in detail in section IV) on Apache Spark. The optimal execution on Spark is further formulated as a divisible scheduling framework that takes into account both task execution precedences and task divisibility when allocating the divisible and indivisible sub-tasks onto computing nodes. The scheduling results provide an optimized solution to automatic processing of big hyperspectral data on cloud environments. The experimental results demonstrate that this approach can achieve significant speedups in the classification of hyperspectral imagery on Spark, obtaining also significant scalability with regards to increasing data volumes.

The paper [93] exploited the idea that state-of-the-art DL-based algorithms and cloud computing infrastructure have become available with a great potential to revolutionize the processing of remotely sensed images. Specifically, their study evaluated, using thousands images obtained over a 12-month period, the performance of three ML and DL approaches (RFs, LSTMs, and U-Nets). The DL algorithms (LSTMs and U-Nets) were implemented using the TensorFlow framework (described in detail in section IV), while the ML-based RF utilized the Google Earth Engine platform. The study concluded that, although the use of ML/DL algorithms depends highly on the availability of labeled samples and the generalization of these methods still presents some challenges, algorithms based on ANNs can still be used in the cloud to map large geographic regions that consider a wide variety of satellite data formats.

The paper [94] uses cloud computing to make global-oriented spatio-temporal data simulations using the OpenStack management framework (described in detail in section IV). This is accomplished by resorting to discrete global grid systems (DGGs), designed to portray real-world phenomena by providing a spatio-temporal unified framework on discrete geo-spatial data structures, along with DL-based algorithm to address the challenges resulting from big remote sensing data storage, processing, and analysis.

The paper [95] presents a new parallel content-based image retrieval (CBIR) system from remotely sensed hyperspectral image repositories, implemented on a cloud computing platform. The method exploits information from spectral unmix-

ing [96] and DL to accurately retrieve hyperspectral scenes. To this end, the authors implement a distributed and DL-based unmixing method that operates on a cloud computing environment. In addition, they implement a global standard distributed repository of hyperspectral images equipped with a large spectral library in a software as a service mode (this concept will be described in detail in section IV), providing users with big hyperspectral data storage, management, and retrieval capabilities through a powerful web interface. The parallel unmixing process is then incorporated into the CBIR system to achieve a highly efficient unmixing-based content retrieval system.

Other important contributions in this area include the paper [97], which proposed a model that facilitates the utilization and performance of Apache Spark algorithms in cloud environments. Also, the paper [98] presents the architecture of InterIMAGE cloud Platform (ICP) data mining package, a distributed tool for the analysis of remotely sensed data. The paper [99] proposes an extension of Apache Hadoop (described in detail in section IV) that executes operations for processing remotely sensed images (including DL methods) in a highly distributed and efficient way. The paper [100] proposes a highly scalable and efficient segmentation model for remotely sensed images, capable of segmenting very high resolution imagery with DL algorithms. The paper [101] proposes a method for DL-based cloud computing based on image sampling, which models the remotely sensed data set to be processed as a streaming service and divides it with a Voronoi diagram. The paper [102] describes a Java software based on the MapReduce model for handling and processing remotely sensed images using DL methods. The paper [103] presents a deep method for storing images using MapReduce. The paper [104] also uses a MapReduce framework for DL-based parallel processing of remotely sensed data through Apache Hadoop. The work [105] describes a set of requirements to achieve a generalized and integrated EO information system and the associated (real-time and off-line) data processing techniques, based on DL. The paper [106] presents a new DL-based approach for distributed processing of large-scale satellite images in the cloud. The paper [107] presents a distributed spatio-temporal indexing architecture implemented on the cloud, and a distributed DL-based algorithm for improved spatial-temporal queries. The paper [108] discusses the requirements of overlapping data organization and proposes two extensions of the Hadoop File System (HDFS) –described in section IV– and the MapReduce programming model for dealing with remotely sensed data. The paper [109] describes a DL framework for the efficient analysis of large image volumes which processes daily the data obtained by NASA’s EO-1 satellite.

D. Challenges and comparison

The different distribution perspectives described in the three previous subsections exhibit numerous differences. As a summary, Table II provides an overlook of the main similarities and differences between the discussed strategies.

Regarding the cost, there are initiatives that provide free computing platform services for the scientific and research

TABLE II
COMPARISON BETWEEN CLUSTER, GRID AND CLOUD ARCHITECTURES

	Cluster computing	Grid computing	Cloud computing
Scalability	Low	Low	High
Elasticity	No	No	Yes
Heterogeneity	No	Yes	Yes
Compute capability	Cluster dependant	Grid dependant	On-demand
Business services	No	No	Yes
Private service cost	Medium	High	On-demand
Public service available	Yes	Yes	Yes
Resource handling/allocation	Centralized	Distributed	Both
Job queuing	Yes	Yes	No

communities. An example is PRACE²² (Partnership for Advanced Computing in Europe), aimed at high-impact scientific research and engineering development across different disciplines. Also, there are specific projects for different distributed computing approaches. Condor²³ was created for research and education purposes. EGI-InSPIRE²⁴ (Integrated Sustainable Pan-European Infrastructure for Researchers in Europe) was created by the European Commission for the benefit of the scientific communities within the European Research Area to exploit grid infrastructures. This project is also available for cloud computing.

Usually, cloud computing has been identified as a distributed service, which is not entirely true. Alternatives are UCI²⁵ (Unified Cloud Interface) or OpenNebula²⁶, which bring flexibility, scalability and simplicity for cloud computing management. In addition to these projects, the XSEDE²⁷ (Extreme Science and Engineering Discovery Environment) ecosystem and the EGI²⁸ (European Grid Infrastructure) provide different cost-free alternatives.

One of the most important features of the cloud (and one of the main reasons for its popularity) is the elasticity and scalability of its architecture. Since cluster/grid architectures are limited to available hardware resources, cloud computing offers the possibility to increase such resources by resorting to the elasticity property, i.e., using resources from different infrastructures. This leads to an increase in the heterogeneity of computing and communication resources, and to the use of both centralized and distributed resource handling and allocation.

Finally, another relevant feature of cloud computing is that the infrastructure does not need to use a job queue for managing executions from different users. This significantly reduces the waiting times for the execution of jobs that are needed in other architectures, such as clusters and grids.

E. Discussion

In this subsection we discuss some important aspects identified after the systematic review conducted in the previous subsections, in an attempt to answer relevant questions such as the role of parallel and distributed computing as an efficient tool for solving remote sensing problems via DL algorithms.

²²<https://prace-ri.eu/>

²³<http://www.cs.wisc.edu/condor/condor>

²⁴<http://www.egi.eu/projects/egi-inspire/>

²⁵<https://code.google.com/archive/p/unifiedcloud/>

²⁶<https://opennebula.io/>

²⁷<https://www.xsede.org/>

²⁸<https://www.egi.eu/services/cloud-compute/>

In our systematic review we have found that distributed computing technologies are highly demanded for DL-based data processing when large volumes of remotely sensed data need to be processed. Commodity cluster computers (possibly including hardware accelerators such as GPUs) [110], grid environments [111], [112], and cloud computing systems [113] have been the most demanded types of HPC platforms for big remote sensing data processing. Recently, cloud computing has become a standard for distributed processing due to its scalability, low cost, service-oriented and high-performance properties [7]. Therefore, this technology offers the potential to deal with tasks that must be accomplished over large data repositories. As result, cloud computing can be seen as the most natural solution for the analysis of large volumes of remotely sensed data, as well as an evolution of other HPC techniques (such as cluster and grid computing) that corrects their limitations and expands their possibilities.

We have also observed that there are comparatively few efforts in the literature aimed at the exploitation of cluster and grid computing infrastructure for the processing of remote sensing images as compared to cloud computing implementations. In fact, we have noticed that the cloud is now in clear expansion and routinely used to solve remote sensing problems (particularly those involving DL algorithms). This is because of the popularity of the programming languages and frameworks available for implementing DL-based algorithms in the cloud (some of these tools will be described in detail in the following section).

Another important observation arising from our literature review is that the most widely used tool for remote sensing data processing in cloud computing environments is Apache Hadoop [114], described in more details in section IV) [115], although recently Apache Spark [116] has also become a reference tool. The main difference between Spark and Hadoop is the fact that the former distributes the data in Resilient Distributed data sets (RDDs) [117] that can be managed more efficiently. As a result, the speedup that Spark can achieve with respect to Hadoop is very high. In addition, Spark provides a ML/DL library called *MLlib* [118] –described in section IV– that operates in a distributed and parallel manner, so that it provides very good performance with big remote sensing data. Although Hadoop has been widely used in the past, Spark is now a standard due to its speed and better memory usage. Our study also shows that most researchers take advantage of existing cloud computing frameworks when dealing with remote sensing data, rather than developing new ones.

In the following section we focus on cloud computing as a *de-facto* paradigm for distributed processing of remotely sensed data sets, and describe the most widely used frameworks and programming languages that have been adopted in this context (some of them already mentioned in this section), with particular emphasis on the availability of DL-oriented tools.

IV. FOCUS ON CLOUD COMPUTING

This section first introduces some basic concepts about cloud computing, and then glances at the delivery models,

available execution frameworks and programming models supporting this technology (with particular emphasis on the tools that have been specifically used in remote sensing applications).

A. Cloud computing basics

Advances in distributed computing technologies, together with the high amount of data generated and consumed by a growing number of devices (including remote sensing instruments), have leveraged the adoption of the emergent cloud computing technology. Nowadays, cloud computing has become a suitable model to cover a wide range of users needs, including data analytics, data mining, remote sensing, social media and other computational and data-intensive applications. Cloud computing is a model that provides users with ubiquitous, on-demand access to remote hardware and software resources in the form of services. This model has become a cost-effective solution that usually reduces initial investment, management and maintenance costs with respect to previous clusters and grid technologies. Cloud computing provides elastic computing, storage and networking payment services. The term elastic refers to the ability of the model to dynamically adapt to user scalability and variable workload necessities [119], [120].

At the core of cloud computing is the *virtualization technology* [121], [122]. Virtualization abstracts the underlying physical resources, such as computing, storage and networking, as a set of virtual resources, typically enclosed in the form of isolated instances called *virtual machines* (VMs). Such modular design enables some advantageous features, as replication of data instances, fault tolerance, security (by limiting the interaction between modules) and migration of instances [123].

Cloud computing determines how the virtualized resources are allocated, deployed and delivered to users. In this sense, delivery models are structured in three basic categories that describe the form in which users access the resources [124]:

- At the higher level of abstraction is *software as a service* (*SaaS*), in which the user accesses to end applications usually developed, managed and maintained by the cloud provider in its cloud infrastructures. Users access these applications via web interfaces. A common example is a CBIR systems for remote sensing data repositories.
- Platform as a service (*PaaS*) refers to the provision of development full stacks, including operating systems, libraries, management and monitoring tools. The users accessing those services are usually software developers with limited control over the underlying infrastructure, that is managed by cloud providers.
- The infrastructure as a service (*IaaS*) model lies at the lower level of abstraction, and allows users to manage an elastic infrastructure composed of a set of computational, storage and networking virtual resources. The users develop and deploy their applications on such virtual resources and manage the infrastructure.

Without neglecting its benefits, cloud computing faces an important set of challenges. The most significant is *security*,

in the sense of both ensuring the access to the data exclusively by authorized users and systems (privacy), and maintaining the integrity and availability of the data distribution (even geographically) and replication across different locations (e.g., as in different remote sensing data centers). *Performance* is an additional key factor impacting the adoption of cloud computing, specially by scientific applications that usually execute on tightly coupled high performance clusters.

Clusters and clouds have different design goals and features [125]. While the main goal of clusters is performance (supported by dedicated parallel computing resources connected with minimal latency and stable high bandwidth networks), the goal of clouds is to make available on-demand virtual resources in a elastic platform with a reasonable cost. Hence, technical management issues –derived from the shared use of physical resources by virtual machines and multiple users, dynamic on-demand scalability, data movement, virtualization overhead and workload balance in the presence of heterogeneity– impact the performance of applications deployed on the cloud infrastructure.

B. Deployment frameworks

Henceforth, we adopt the perspective of users of an IaaS model, allocating a pool of virtual resources connected by networking services to deploy and run scientific applications. As an effective management framework to deploy and run the applications we highlight OpenStack²⁹, although there are other open-source frameworks such as Apache Cloudstack³⁰ and OpenNebula that offer similar features. In the following, we review the functionality of several of the multiple services available in OpenStack.

Openstack is oriented to manage the complete life cycle of an IaaS cloud system composed of a large number of virtual resources for computing, storage and networking. The Openstack design architecture is highly modular. Each independent module implements a specific service and exposes a well-defined application programming interface (API), making the system extensible and able to support the integration of third-party services.

Some of the main software components of the OpenStack framework are outlined next and shown in Fig. 8. We structure them in three categories:

- *Computing service components* for deploying and managing virtual machine and Linux containers. These facilities are supported fundamentally by the *Nova* compute engine. Furthermore, OpenStack offers a general purpose management framework component for hardware accelerators (such as GPUs) called *Cyborg*.
- *Networking service components* with support for different network technologies and equipment. *Neutron* is the main component, and it allows to manage software-defined networks (SDNs) and attach virtual devices to ports on these networks.
- *Storage service components*, including *Cinder* component for block storage and *Swift* component that delivers

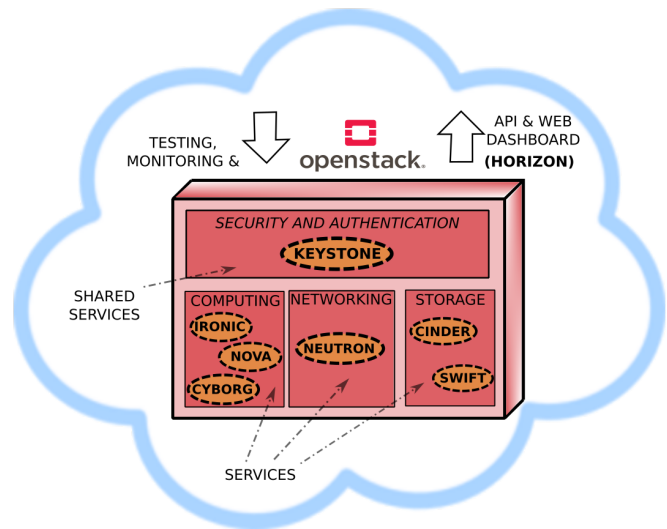


Fig. 8. OpenStack main computing, networking and storage components.

services for securely storing unstructured data as a pool of objects and files.

In addition to the aforementioned computing, networking and storage services, it is worth noting that OpenStack includes multiple software services for monitoring, development, recovery, databases, orchestration of virtualized resources, workload balance and more. Among them, for instance, *Keystone* provides services for security and authentication of users and applications. Finally, *Horizon* presents a web interface in the form of a dashboard to manage the virtual resources composing the cloud infrastructure.

C. Programming models

The term remote sensing big data was coined to refer to the increasing need of storing, managing and processing vast amounts of remotely sensed data, produced at a high rate and in a wide range of formats. Such overwhelm flow of data requires flexible parallel platforms with a large computational capacity and specific programming facilities.

In contrast to cluster computer-centric paradigms, as message passing, big data applications require a data-centric approach, in which a computational task should be deployed in a computational resource as close as possible to the data location. As a result, the movement of data through the network can be minimized. The limited capability of the network bandwidth is a factor that, together with the high volume and heterogeneity of remotely sensed data, highly affects the performance. In the following, we review some relevant programming models and application ecosystems that have been widely used in big remote sensing applications.

1) *MapReduce*: It is a programming model aimed at developing scalable and robust applications working with large data sets [126], [127]. It was originally developed by Google and, together with the distributed Google File System (GFS) [128], has been successfully used in solving numerous big remote sensing data problems (as described in section III). This model

²⁹<https://www.openstack.org>

³⁰<https://cloudstack.apache.org>

is particularly suitable for data-centric environments, such as big remote sensing data processing on cloud computing platforms [129]. MapReduce provides the developer with a simple interface based on the *map* and *reduce* functions. An application takes as an input a structured set of key-value data. The $map(k1, v1)$ function transforms the input to an intermediate set of key-value pairs in the target domain $(k2, v2)$. The MapReduce library combines the intermediate values $(v2)$ in a per-key basis. Finally, the $reduce(k2, list(v2))$ function operates on the list of values corresponding to each target key to generate the output. Usually, both *map* and *reduce* are executed by parallel tasks deployed across a set of computational resources. The model interface abstracts the complexities of the execution of the remote sensing application in the specific platform, and leaves the underlying details to the implementation.

MapReduce implementations are ultimately based on the master-worker approach, in which a *master* process manages the automatic parallelization and scheduling of the tasks on the computational nodes, and coordinates their execution. It partitions the remotely sensed data to be processed by each *map* and *reduce* tasks, and schedules those tasks on computational resources as close as possible to the data to be processed. To achieve this, it relies on GFS, which provides the locations of the blocks of data to be processed. This design improves data locality and minimizes the data transfers through the network, hence improving the performance.

Besides, MapReduce combines the output intermediate files of the *map* tasks and combines the data according to the output keys. This intermediate stage tackles the complexity of a high amount of communication and coordination tasks, that are hidden to the developer. Finally, the implementation delivers processed data chunks to the *reduce* tasks. Furthermore, MapReduce implementation promotes fault tolerance mechanisms to detect and re-execute tasks when necessary. In this sense, the communication between the deployed tasks is achieved using intermediate files.

2) *Hadoop Ecosystem*: Hadoop³¹ has been widely used to parallelize remote sensing data processing tasks (see section III). It is a popular open-source and scalable big data software framework based on the MapReduce paradigm, the HDFS and Yet Another Resource Negotiator (YARN) [115] resource manager.

HDFS basic functionality is similar to that of GFS. Data are split up in blocks of fixed size and replicated across several nodes to ensure fault tolerance and availability. Its implementation follows a master-worker approach. A *Namenode* process manages metadata (such as block mapping information) and delivers that information to the MapReduce library when requested. *DataNode* processes execute in each virtual resource and effectively store data blocks and provide data reading and writing services to applications.

MapReduce functionality is implemented by a *JobTracker* process, which receives job requests and schedules job tasks to different nodes. Each node is controlled by a *TaskTracker* process which monitors the execution and reports to the

JobTracker if a problem appears. In such case, *JobTracker* resubmits the involved tasks to the same or a different *TaskTracker*. In this sense, Hadoop decouples the MapReduce programming model and the associated resource management. YARN delivers the former services. Its internal architecture is based on three main components:

- The first component is the global per-cluster *ResourceManager* process, which accepts job submissions and allocates resources for the application. It is responsible for scheduling the application in the available resources.
- The second component is the *NodeManager*, a per-node process. It is responsible of the execution of the tasks assigned to its node.
- Finally, the *ApplicationMaster* is a per-application process that monitors the application necessities and their status along their lifecycle, negotiating resources with the *ResourceManager*.

The Hadoop framework has been enhanced with multiple tools and services forming the so-called *Hadoop Ecosystem*, which has been exploited in a variety of remote sensing applications (see section III). It includes relational databases managers (as *Hive*), NoSQL databases (as *HBase*, a column-oriented distributed database running on top of HDFS), distributed ML/DL and linear algebra solvers (as *Mahout*), real-time facilities (*Storm*), efficient alternatives to the MapReduce programming model and utilities for the orchestration of the services and components (*ZooKeeper*).

3) *Apache Spark*: Originally developed at UC Berkeley, Apache Spark [116], [130] was designed to gain velocity in the processing of big data. Although the MapReduce model adequately adapts to a large number of applications as highly parallel batch jobs, it incurs in significant latency in both interactive applications and in those with an iterative pattern of execution, in which the same data sets need to be continuously reloaded from the file system.

In this respect, one of the most relevant features of Spark is its ability to support persistent data in memory, which greatly benefits performance. This feature is implemented in the run-time system of Spark, known as *Spark Core Engine*. In addition to this module, the Spark ecosystem includes a set of utilities, as *Spark SQL*, that allows to manage structured and semi-structured data organized in columns, known as *DataFrames*, the *Spark Streaming* module for performing real-time processing on data streams (ideal for remote sensing applications with real-time constraints, as described in section III), and the *MLlib* library that includes distributed ML and DL algorithms.

Spark follows a master-worker execution model, with a *driver* node acting as the master and a set of worker nodes. The execution model is shown in Fig. 9. An application submitted to Spark starts its execution in the context of the driver node. The application creates a *SparkContext* object that transforms the sequence of operations described in the main program into an execution plan. The execution plan is represented as a directed acyclic graph (DAG), in which nodes are data elements and edges are operations on such data. *SparkContext* splits up the DAG in stages to be executed by tasks. Then, *SparkContext* negotiates the acquisition of

³¹<https://hadoop.apache.org>

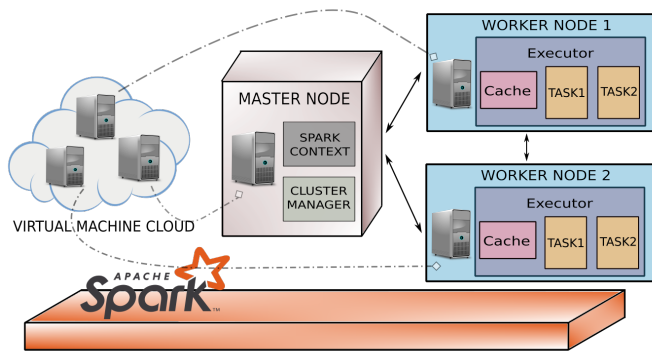


Fig. 9. Apache Spark execution model in a cloud computing platform.

resources with the *Cluster Manager*. While the normal form of executing a Spark application is to use its own native cluster manager, Spark is able to run on Hadoop clusters on top of the YARN resource manager. In any case, after the resources have been obtained, SparkContext launches an *Executor* process in each worker node. Based on the execution plan, SparkContext schedules tasks to worker nodes and coordinates its execution. Conversely, Executors are in charge of effectively executing tasks assigned to its worker nodes and providing access to data.

The *Resilient Distributed Dataset (RDD)* [117] is the main abstraction supporting the Spark model. An RDD is a read-only collection of objects partitioned and distributed across the worker nodes. The assignment of tasks to worker nodes considers data-locality, that is, the availability (or closeness) of the data to be processed by the task in the RDD partition assigned to the work node. An important feature of RDDs is that the data can be cached in memory, and hence reused in recurrent parallel operations with minimum overheads. Further, the RDD is a fault-tolerant data structure. In this sense, any operation on an RDD object is logged in such a way that, in case of node failure, the RDD can be reconstructed using the operation *lineage*. Because lineage dependencies become large and their management is time-consuming, users may decide to establish checkpoints in the execution. Additionally, the immutable nature of the RDD objects benefits another fault-tolerance mechanism, the execution of backup copies (duplicates) of running tasks if failed or straggler tasks are detected.

RDDs can be created from data structures in memory or in the file system, and also using data obtained from any Hadoop service, including the HDFS or databases as HBase. In addition to the mechanisms for creating RDDs, Spark includes a set of coarse-grained operations to process RDD data sets. These operations are structured in two types:

- The first type is given by *transformation* operations, which apply a function to a RDD and generate a transformed RDD data set as a result. In turn, transformations are classified into *narrow*, that involve data located in the worker node where the task executes the operation, and *wide*, that involve data across multiple worker nodes, and therefore, the required data are copied from other

partitions. Movement of data is coordinated by the driver. Examples of transformation operations defined in the API are *map*, *filter*, *groupbykey* and *reducebykey*.

- The second type of operations are called *Actions*. Actions are operations that retrieve non-RDD values (as statistical or processed values) from RDDs, and their value is returned to the driver program. Examples of actions are *count*, *collect*, *reduce* and *foreach*. Some of them take as an input parameter a function to be applied to the data.

Transformations are *lazily* executed, in the sense that a sequence of transformations are effectively executed when an action is performed on the transformed RDD. This mechanism improves performance in cases in which only final results (and not intermediate results of the sequence of transformations) are transferred to the driver program. Nevertheless, by default, each transformed RDD is recomputed each time an action is executed on it. The persistence mechanism of Spark allows to keep in memory the data, improving the performance of recurrent operations on the RDD (this is particularly beneficial for image processing operations involving sliding windows or kernels, which are very popular in many remote sensing applications).

Moreover, wide transformations are inefficient operations in the Spark model, consisting of independent tasks operating on their own assigned RDD partitions, because such operations require data movement. Every time a task executes an operation on remote data, SparkContext coordinates the disk input/output (I/O) and network transmissions between the involved nodes. This costly procedure is called *shuffle*. The persistence capability highly improves the performance of the shuffle operations by caching in memory the RDD data that are going to be reused in wide transformations, which is also popular in multi-scale image processing operations adopted in many remote sensing applications (see section III).

Finally, Spark includes two additional mechanisms to avoid recurrent copies of shared data between the worker nodes, called *shared variables*. The first one is called *broadcast*, and it allows to diffuse a set of values to worker nodes, that will hold a read-only copy of the data. The second one allows to maintain simple (associative) *accumulators* shared across worker nodes.

D. Cloud computing for scientific applications

Scientific applications (including remote sensing ones) exploit the low latency and dedicated resources of clusters to obtain high performance. On the contrary, cloud computing offers elastic multi-tenancy (resource time sharing) and non-predictable and unstable network facilities. Nevertheless, there is a great interest in evaluating if cost-efficient and flexible cloud platforms can execute scientific HPC applications at a reasonable level of efficiency.

The paper [131] proposes AzureMapReduce, a decentralized MapReduce implementation for Microsoft Azure cloud infrastructures, and evaluates its (weak-scale) scalability and performance with a remote sensing application. Authors claim that MapReduce applications in cloud infrastructures exhibit comparable performance to MapReduce applications executed

on traditional clusters. On the contrary, work [132] analyzes the performance of the HPC Challenge (HPCC) benchmark [133] on Amazon EC2 cloud platform, and concludes that the performance of general scientific applications on cloud infrastructures is at least one order of magnitude lower than that on clusters and supercomputers. Moreover, work [134] analyzes the performance of loosely coupled many-task scientific computing applications on four commercial cloud computing provider platforms with the same aforementioned conclusion.

The thorough study in [135] presents a run-time performance comparison of the characteristics of the Amazon EC2 cluster computing instances and a supercomputer. The paper evaluates latency and bandwidth micro-benchmarks, HPCC matrix multiplication kernels, NAS Parallel Benchmarks (NPB [136]) and four full-scale remote sensing applications used at NASA. The results show that, in one node, performances are equivalent while, in several nodes, the network overheads of the cloud computing infrastructure have a huge impact in performance.

The paper [125] evaluates Amazon IaaS services at different levels. Authors execute micro-benchmarks to extract raw performance of latency, bandwidth, memory and processing services. Furthermore, they execute the parallel High Performance Linpack (HPL) [137] benchmark to compare cluster and cloud environments. The goal was to identify advantages and limitations of cloud platforms. They conclude that I/O and network performance differences are the main factors impacting the applications performance. One of the main drawbacks detected in the cloud is the network infrastructure based on Ethernet, that is often not suitable for the necessities of HPC applications (including remote sensing ones). Paper [123] proposes the HPC2 model, that bridges the gap between cluster and cloud platforms with a set of proposals, including using Infiniband as network technology (as it is commonly the case in current remote sensing applications implemented in cloud environments).

There is a consensus in that the performance differences between platforms come from the inherent overheads in virtualization, memory, storage and I/O, and latency of the network infrastructure [138], [139]. Further, work [140] offers an extensive study of the performance of several cloud providers of public IaaS services: Amazon Elastic Computing Cloud (EC2), Microsoft Azure, Google Compute Engine (GCE) and IBM SoftLayer (SL), and it concludes that, indeed, there are substantial differences between the performance of infrastructures of different cloud providers.

To overcome the differences between cluster and cloud platforms, several works maintain that it is not enough to straightforwardly run cluster applications on cloud platforms. This is in contrast with many cloud implementations of remote sensing algorithms described in section III, which simply run available cluster-based codes in cloud environments. For instance, paper [141] proposes to slightly transform HPC applications as representative HPC kernel solvers by optimizing computational granularity, which has a high impact in scheduling and communication/computation overlapping. Additionally, they propose to transform cloud facilities to use thin virtual machines and CPU affinity mechanisms. They

conclude that, by transforming HPC applications (such as remote sensing ones) to be run in a cloud, and making clouds *HPC-aware*, the impact of the latency and multi-tenancy is significantly reduced. In this sense, paper [142] proposes to use the MapReduce with Access Patterns (MRAP) model, that extends MapReduce with usual HPC application data access pattern semantics (non-contiguous and fine grained), while taking advantage of the inherent scalability and fault-tolerance features of MapReduce. This is a promising solution to increase the performance of the MapReduce-based remote sensing implementations described in section III.

V. MACHINE AND DEEP LEARNING LIBRARIES AND FRAMEWORKS IN CLOUD COMPUTING ENVIRONMENTS

Numerous efforts have been devoted to the development and efficient execution of ML and DL applications on cloud computing infrastructures, not only as optimized libraries and services, but also as applications on top of the Spark and MapReduce programming models. In this respect, cloud providers offer several facilities for this challenging task. Among them, Amazon AWS promotes a ML platform called SageMaker³² to build and train different models, with support for TensorFlow and Spark. IBM provides tools for different frameworks, including TensorFlow and Keras³³. GCE also supports the use of TensorFlow and provides an infrastructure based on GPU computational devices. Microsoft Azure, on the other hand, bases its services on Kubernetes and allows to use accelerators for its ML/DL resources.

Nevertheless, the iterative execution pattern of ML/DL learning applications (in which the same data are recurrently operated, as in many remote sensing data processing algorithms) does not naturally adapt to established cloud computing programming models. Virtualization and network overheads are key factors impacting the efficiency of ML/DL learning applications, which are usually supported on highly optimized computing linear algebra libraries such as BLAS (Basic Linear Algebra Subprograms) [143] and high performance networks and communications based on Message Passing Interface (MPI) [144] to achieve high performance.

Additionally, ML and DL models have dramatically grown in terms of structural complexity and depth. Training models on huge data sets (such as those involved in remote sensing applications) has become a computationally very intensive (as well as a memory-consuming) task, that usually requires several days even using specialized hardware such as GPUs. To overcome this limitation, several methods for parallel training of ML and DL models have been developed. The parallelization schemes can be structured in three main groups [145], [146]:

- The first type is the *data-parallelism* scheme, in which several replicas of a model are simultaneously trained in different computational devices on disjoint partitions of the remote sensing data set.
- The second type is called *model-parallelism*, and it is used when a model overcomes the memory capacity of

³²<https://aws.amazon.com/sagemaker/>

³³<https://keras.io/>

one computational device, and then it has to be partitioned and deployed on several devices.

- Finally, the last type is the *hybrid-parallelism* scheme, that merges data and model-based approaches.

These parallelization schemes can be implemented as multiple distributed computing approaches. Table III summarizes the different characteristics of each scheme for cloud and cluster/grid computing approaches.

TABLE III
SUMMARY OF PARALLELIZATION SCHEMES FOR DISTRIBUTED COMPUTING APPROACHES.

Scheme	Cloud computing	Cluster/Grid computing
Data-parallelism	<ul style="list-style-type: none"> • Master node holds the data. • Data/computation are split. • Model is replicated. • Workers sends results to the master. 	<ul style="list-style-type: none"> • Data are partitioned over executors. • Model is replicated over executors. • Executors share intermediate results.
Model-parallelism	<ul style="list-style-type: none"> • Master node holds the data. • Data are not partitioned. • Model/computation is split. • Workers sends results to the master. 	<ul style="list-style-type: none"> • Data are replicated over executors. • Model/computation is partitioned. • Executors share layer outputs and results.
Hybrid-parallelism	<ul style="list-style-type: none"> • Master node holds the data. • Data are partitioned. • Model is split over workers. • Computation depends on data and model workload. • Workers sends results to the master. 	<ul style="list-style-type: none"> • Data are split over executors. • Model is partitioned over executors. • Computation depends on data and model workload. • Executors share intermediate results and outputs.

The rest of the section outlines the main frameworks and libraries offered by cloud providers to face the challenge of efficient training of ML and DL models when processing remotely sensed data on cloud computing infrastructures.

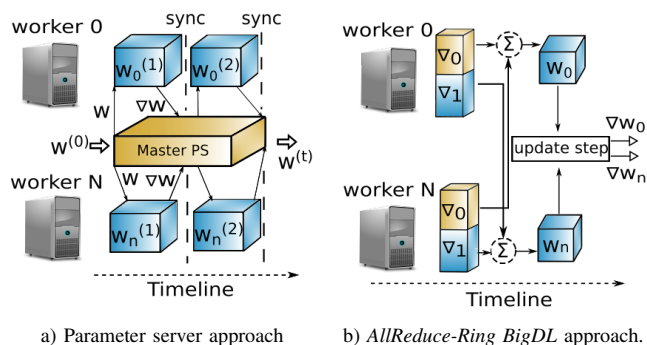


Fig. 10. Cloud computing training pipeline in the three discussed DL frameworks (*TensorFlowOnSpark*, *SparkTorch* and *BigDL*).

A. Libraries

This subsection provides an overview of some well-known ML and DL libraries that are used in cloud computing environments to build models efficiently. These libraries have been used in the past to process and accelerate remote sensing applications.

1) *Weka*: The *Weka*³⁴ library was developed at the University of Waikato [147]. It is a Java-based, open-source library that allows to build ML and DL models for several types of algorithms, including classification, clustering, data mining, etc. A relevant feature of the library is that it is multi-platform, and even runs on lightweight devices on top of the Android operating system [148]. It supports multiple programming languages with different packages and plugins,

as the *DeepLearning4J*³⁵, the *RPlugin*³⁶ or several *Python*³⁷ DL libraries. *Weka* was initially developed to offer a simple and easy to use interface. Currently, it provides a distributed version [149] implemented on top of Spark and RDDs.

2) *Mllib*: The *Mllib*³⁸ is a distributed ML/DL library which provides model training based on the data parallelism scheme [118]. It was developed in the *Scala*³⁹ programming language, and supports Java, Scala and Python programming languages. Its main features are scalability and fast implementation of numerous ML/DL algorithms, as well as linear algebra, statistics and optimization primitives. It is built on Spark, and implements efficient communication primitives for data transmissions performed by a large amount of processes, training large models using the data-parallelism scheme. Some communication primitives of special interest are the *broadcast*, that efficiently distributes data over processes training the model, and the tree-structured *aggregation* primitive, which collects processed data avoiding possible bottlenecks. *Mllib* was initially designed to efficiently operate on fully distributed environments, which entails a significant performance advantage with respect to *Weka* [150].

The execution model of *Mllib* is based on the master-worker paradigm, where the master process acts as a *parameter server*, and maintains a centralized copy of the global parameters of the model. It combines values received from the worker tasks after each training iteration. Tasks deployed on the computational resources of the platform process their assigned data set partitions and communicate results to the parameter server. The data partitions assigned to each task are processed in *batches*, being the size of each batch a key optimization parameter which directly affects the resulting accuracy of the model and the efficiency of the training [151], [152].

B. Frameworks

This subsection discusses different frameworks to develop ML/DL applications on cloud environments.

1) *TensorFlowOnSpark*: This framework combines salient features of the TensorFlow DL library with Spark and Hadoop to provide with a scalable ML/DL development and training platform [153]. It supports all TensorFlow functionalities, including asynchronous and synchronous training, data and model-based parallelism schemes, and monitoring with *TensorBoard*⁴⁰. It also enables distributed TensorFlow-based training on cloud computing clusters, with the additional goal of minimizing the amount of code refactoring required to run existing TensorFlow applications. In summary, *TensorFlowOnSpark* deploys a Spark cluster on a cloud infrastructure and provides facilities for injecting both RDD and HDFS data in the TensorFlow models executed by the tasks scheduled in the worker nodes.

³⁵<https://deeplearning4j.org/>

³⁶<https://weka.sourceforge.io/packageMetadata/RPlugin/>

³⁷<https://www.python.org/>

³⁸<https://spark.apache.org/mllib/>

³⁹<https://www.scala-lang.org/>

⁴⁰<https://www.tensorflow.org/tensorboard>

³⁴<https://weka.sourceforge.io/>

2) *SparkTorch*: This framework is intended to execute code based on the PyTorch library [154] across nodes in a Spark cluster. The distributed training works under a data-parallel paradigm, and uses both tree reductions and parameter server mechanisms to combine partial results from tasks deployed on the cloud platform. There are two main training modes available in *SparkTorch*:

- The first one is the *asynchronous* training mode, which ensures that the replicated models deployed in nodes are synchronized through each training iteration.
- The second training mode, called the *Hogwild* approach [155], allows lock-free task accessing to shared memory in order to update parameter values. This mode eliminates the overheads associated with locking. However, in this mode, a task could overwrite the progress of another tasks, a risk that the developers claim that could be assumed when the data to be accessed is sparse.

3) *BigDL*: This framework is also implemented on top of Apache Spark to run DL applications as standard Spark programs [156]. It offers support for large-scale distributed applications and provides efficient processing for data-analysis, data injection to neural network models, and distributed training or inference, using an unified pipeline. Before training, the model and RDDs are partitioned and cached in memory across the cloud resources. *BigDL* supports two parameter synchronization mechanisms. The first one maintains a centralized parameter server, and the second one uses collective operations as *AllReduce* to combine the parameters computed by tasks. Despite the fact that collective message passing primitives are not particularly suitable for the execution model of a Spark cluster, *BigDL* implements an efficient *AllReduce* algorithm using Spark primitives, allowing for the integration of DL algorithms in cloud computing environments.

It is important to note that the three aforementioned frameworks can use two different communication approaches:

- The *parameter server approach*, illustrated in Fig. 10(a) [157], consists of a centralized architecture where the computational nodes are partitioned into masters and workers. The workers maintain a workload and data partition, while the master maintains the global shared parameters. The workers communicate with the master to share the weights generated at each iteration of the model. The master is responsible for the aggregation of the global weights. In cloud computing environments, additional workers may be added or removed from the execution. This must be handled by the system, so as to switch on any new workers and send to them the appropriate computations and data partitions.
- The *AllReduce-Ring BigDL approach*, shown in Fig. 10(b), consists of a decentralized architecture where each Spark task computes its local gradients, dividing the local gradients into N partitions. Each task manages its corresponding parameter partition, which is shuffled to the corresponding task to aggregate gradients and then update the corresponding weights. Then, each task launches a Spark broadcast operation with the updated weights, so that these are read before the next step.

A discussion between the aforementioned frameworks and their pros and cons in cloud environments is needed at this point:

- In terms of applications, all of them provide full integration with Spark ML/DL algorithms. While *TensorFlowOnSpark* provides a large amount of algorithms and applications, *BigDL* includes extensive DL functionalities.
- The ease of use differs among different frameworks. While the *TensorFlowOnSpark* interface is clear and easy to use, the documentation for *BigDL* is intuitive and provides a comprehensive support for ML/DL algorithms.
- Attending to the distributed training, *SparkTorch* provides asynchronous and synchronous schemes, since *TensorFlowOnSpark* asynchronous parameter server is highly efficient. We also note that *SparkTorch* is in a premature developing phase compared to the *TensorFlowOnSpark* and *BigDL* implementations. Thus, the latter frameworks still offer notable advantages. An important aspect of *TensorFlowOnSpark* is the creation of checkpoints to recover from failures. These checkpoints are stored in the HDFS by TensorFlow. A similar point between these two last frameworks can be found in the monitoring.
- In terms of scaling and performance, *BigDL* takes a step forward the other frameworks. This is due to multiple factors. First, it provides extensive documentation to deploy ML/DL algorithms in different providers as EC2. Also, attending to the execution, it provides a synchronous SGD and an optimized *AllReduce* in the communication step. Another interesting point is that it can be used only for prediction and, hence, it can load models from different ML/DL frameworks. Finally, the aforementioned frameworks execute powerful long-running tasks, while *BigDL* uses short-running, non-blocking tasks for the model computation.

VI. CASE STUDY

This section presents a case study in which a deep neural network (implemented on the cloud) is used to process a large hyperspectral remote sensing image. As noted before, hyperspectral data cubes comprise a significant amount of information, which allows us to model the physical characteristics of the observed materials by analyzing the detailed spectral signatures (collected on a pixel-by-pixel basis), that provide rich information for land-cover analysis. When applied to hyperspectral images, classification methods suffer from important processing time and computing/storage constraints, resulting from the extremely high dimensionality of the data. Therefore, the implementation of such classifiers in cloud computing architectures is an effective solution. Here, we use a deep multi-layer perceptron (MLP) [17] as the distributed classifier, and perform classification experiments on a benchmark classification data set widely used in the hyperspectral imaging community.

The remainder of this section is organized as follows. First, we describe our cloud implementation of the MLP classifier. Then, we describe the hyperspectral image used

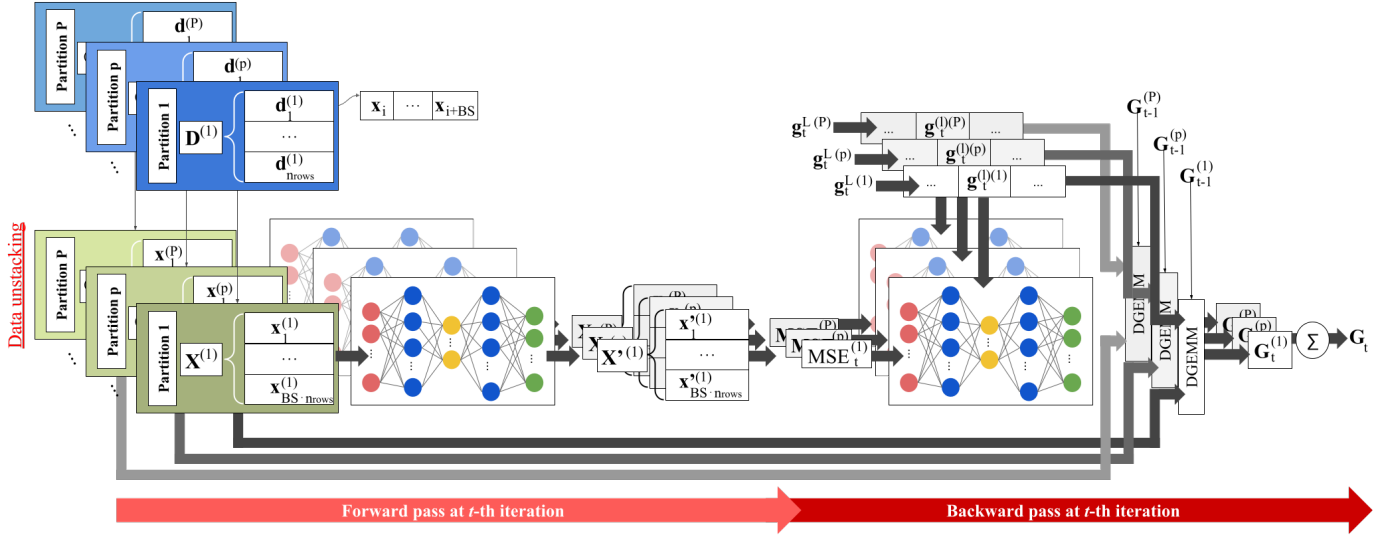


Fig. 11. Graphical overview of the forward and backward pipelines of our distributed MLP classifier

for validation purposes. Then, we provide the characteristics and configuration of the cloud computing platform used for experiments. The section ends with a detailed discussion of our conducted experiments and with some remarks on the practical utility of distributed DL algorithms in remote sensing applications.

A. Distributed multi-layer perceptron (MLP) classifier

Let us denote a hyperspectral data cube as $\mathbf{X} \in \mathbb{R}^{h \times w \times n_{bands}}$, where each data sample \mathbf{x}_i is of size $h \times w$ (being h the height and w the width in pixels of the image), and each pixel can be denoted by $\mathbf{x}_i \in \mathbb{R}^{n_{bands}} = [x_{i,1}, x_{i,2}, \dots, x_{i,n_{bands}}]$. In an MLP classifier, each layer l performs a data transformation of the weights and the input data (\mathbf{W}^l and \mathbf{x}_i) as follows:

$$\mathbf{x}_i^{l+1} = \mathcal{H}(\mathbf{x}_i^l \cdot \mathbf{W}^l + b^l) \quad (1)$$

where \mathbf{x}_i^{l+1} is a feature representation of the input data obtained by the neurons of that layer (l). Neurons are obtained as the dot product between the output from the previous layer neurons plus the bias b (shift parameter) through an activation function such as the rectified linear unit (ReLU) or the sigmoid function, among others [17]. Hence, the k th feature of the $\mathbf{x}_i^{(l+1)}$ sample can be obtained as:

$$x_{i,k}^{l+1} = \mathcal{H}\left(\sum_{j=1}^{n^{l-1}} (x_{i,j}^l \cdot w_{k,j}^l) + b^l\right) \quad (2)$$

Attending to the MLP optimization step, the optimizer tries to obtain the set of parameters \mathbf{W} and *bias* that minimize the loss error. Hence, a back-propagation step calculates the gradient of the error in order to minimize the final error. At each step, the updating process is defined as follows:

$$\mathbf{W}_{t+1} = \mathbf{W}_t + \nabla \mathbf{W} \quad (3)$$

where $\Delta \mathbf{W} = \mu_t \cdot \mathbf{p}_t$, being μ the learning rate and \mathbf{p} the descent direction of the gradient at time step t . To optimize this operation, traditional methods use the information obtained from the Hessian matrix:

$$\mathbf{H}_t \cdot \mathbf{p}_t = -\nabla E(\mathbf{X}, \mathcal{W}_t) \quad (4a)$$

$$\mathbf{p}_t = -\mathbf{H}_t^{-1} \cdot \nabla E(\mathbf{X}, \mathcal{W}_t) \quad (4b)$$

$$\mathcal{W}_{t+1} = \mathcal{W}_t - \mu_t \cdot \mathbf{H}_t^{-1} \cdot \nabla E(\mathbf{X}, \mathcal{W}_t) \quad (4c)$$

where \mathcal{W}_t are the network weights, \mathbf{H} is the Hessian matrix and $\nabla E(\mathbf{X}, \mathcal{W}_t)$ is the gradient of the error at step t . Regarding this, as the computation requirements are high, the optimization provided by the BFGS (Broyden–Fletcher–Goldfarb–Shanno) algorithm [158] is used, providing an estimation of the Hessian matrix changes:

$$\mathcal{W}_{t+1} = \mathcal{W}_t - \mu_t \cdot \mathbf{G}_t \nabla E(\mathbf{X}, \mathcal{W}_t) \quad (5)$$

with \mathbf{G}_t being the inverse Hessian approximation matrix. Since \mathbf{G} is the inverse of the Hessian matrix \mathbf{H}^{-1} , and the approximation matrix \mathbf{G} needs to be updated in each step, the BFGS method updates it using the following equation, assuming that $\mathbf{q}_t = \mathbf{H} \cdot \mathbf{p}_t$ and $\mathbf{H}^{-1} \cdot \mathbf{q}_t = \mathbf{p}_t$:

$$\mathbf{G}_{t+1} = \mathbf{G}_t + \frac{\mathbf{p}_t \cdot \mathbf{p}_t^T}{\mathbf{p}_t^T \cdot \mathbf{q}_t} - \mathbf{G}_t \cdot \frac{\mathbf{q}_t \cdot \mathbf{q}_t^T}{\mathbf{q}_t^T \cdot \mathbf{G}_t \cdot \mathbf{q}_t} \cdot \mathbf{G}_t \quad (6)$$

This strategy is quite appropriate for hyperspectral data sets, because the computation is high and the processing requires to read repetitively the original data set. In this way, a distributed cloud computing implementation can make the MLP faster and highly scalable.

Specifically, our distributed implementation reshapes the hyperspectral data from $\mathbf{X} \in \mathbb{R}^{h \times w \times n_{bands}}$ to $\mathbf{X} \in \mathbb{R}^{n_{pixels} \times n_{bands}}$. This way, each row or column collects the full

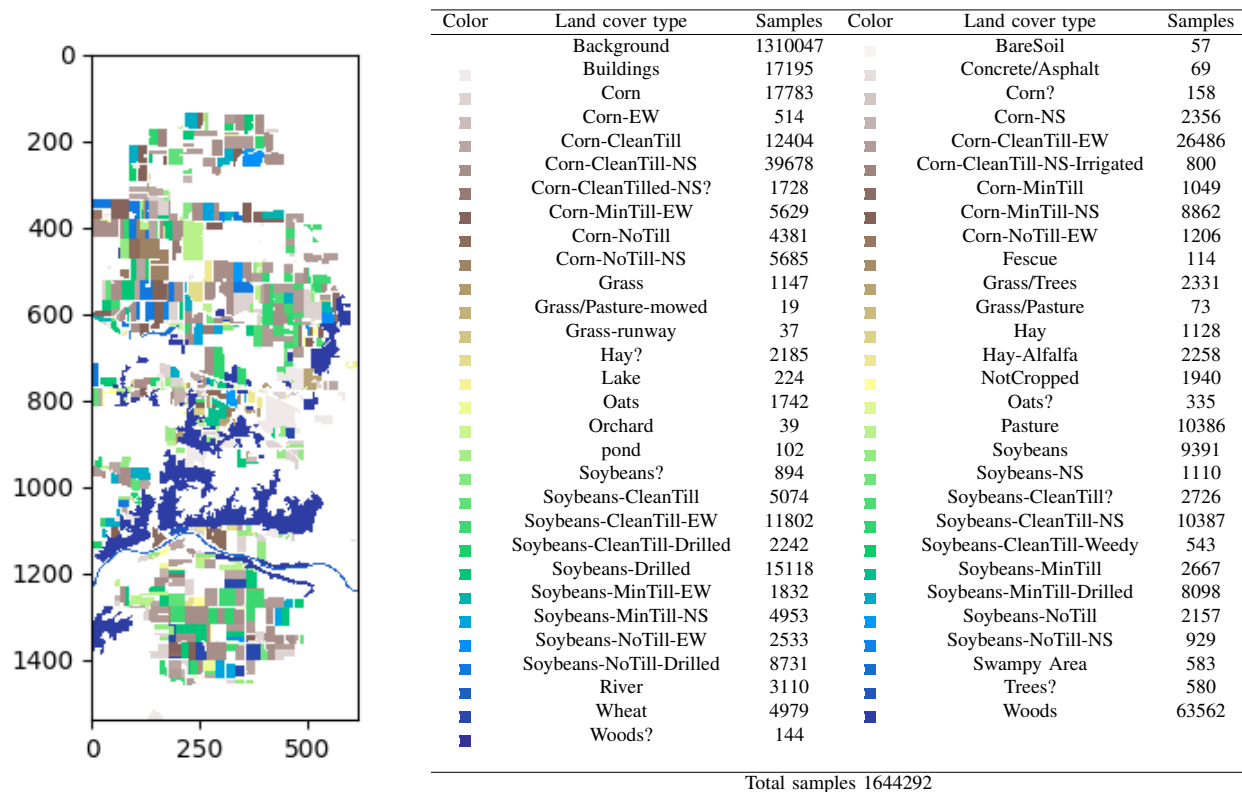


Fig. 12. Available labeled samples (and their distribution) in the AVIRIS Big Indian Pines (BIP) data set considered in experiments.

representation of a pixel. The master node reads and divides the original data set over P partitions, which are assigned to the corresponding workers in the cluster. The data are stored in each worker as an RDD. To take advantage of this, we can improve the computation time of the distributed algorithm approaches using a block size (BS) implementation. In our implementation, each data partition (row) $\mathbf{r}_j^{(p)}$ is transformed based on the minimum and maximum feature of a sample with the minimum and maximum from the column of its data partition as follows:

$$\mathbf{r}_j^{(p)} = \frac{\mathbf{r}_j^{(p)} - \mathbf{r}_{min}^{(p)}}{\mathbf{r}_{max}^{(p)} - \mathbf{r}_{min}^{(p)}} \cdot (\mathbf{x}_{max} - \mathbf{x}_{min}) + \mathbf{x}_{min} \quad (7)$$

Now, every training iteration can be performed using a forward-backward procedure. An aggregation is done after each step to compute and process the gradients and losses from workers and, hence, return a single gradient and loss. In each forward propagation, each worker forwards its corresponding data partition $\mathbf{X}^{(p)}$ through the layers. Then, gradients are computed at the back-propagation step, obtaining for each partition (p) the $\mathbf{G}_t^{(p)}$ matrix at time step t . Gradients are sent to the master node and then the computation of ΔW_t takes place. Assuming that $\mathbf{X}^{(p)} \in (\mathbb{R}^{BS \cdot n_{rows}} \times n_{bands})$ for (p) partitions, Eq. (1) is distributed as:

$$\mathbf{X}_i^{(l+1),(p)} = \mathcal{H}(\mathbf{X}^{(l),(p)} \cdot \mathbf{W}^{(l)} + b^{(l)}) \quad (8)$$

where $x_i^{(l+1),(p)}$ represents the output matrix neurons of size $(\mathbb{R}^{BS \cdot n_{rows}} \times n_{neurons}^l)$ from layer l , $x^{(l),(p)}$ is the input pixel matrix of size $(BS \cdot n_{rows})$ from the previous layer, $\mathbf{W}^{(l)}$ is

the matrix of weights, which connects neurons from previous layers with the actual one, and \mathcal{H} is the activation function (e.g., the ReLU).

Once the forward step is completed in every partition $\mathbf{X}^{(p)}$, the error loss is computed by every worker. The final error is calculated by the master as the mean of all the errors provided by the slaves. Then, the partition error is back-propagated to calculate the $\mathbf{G}_t^{(p)}$ gradient at each time step t . The gradients of each partition are computed using a parallel double-precision matrix-matrix multiplication (DGEMM), implemented in BLAS:

$$C = \alpha * \mathbf{X}^{(p)} * g_t^{L(p)} + \beta * \mathbf{G}_{t-1}^{(p)} \quad (9)$$

where α and β are regularization parameters set to $1/n_{bands}$ and 1, respectively, $g_t^{L(p)}$ is a matrix representing the neuron impact per layer $L = [l_1, l_2, \dots, l_n]$, and $\mathbf{G}_{t-1}^{(p)}$ denotes the previous gradient matrix values. The variable p , as indicated previously, represents the partition of the data. In the end, all partition gradients $\mathbf{G}_t^{(p)}$ are summed to obtain the global gradient matrix \mathbf{G}_t at the time step t .

Fig. 11 provides a graphical description of the distributed forward and backward pipelines of our distributed MLP during the training stage (considering iteration t). This is conducted after unstacking the hyperspectral samples in each distributed data partition, where each one is allocated to a different worker node.

B. Hyperspectral data set

To evaluate the performance of our cloud implementation of the MLP classifier, several experiments have been conducted over the Big Indian Pines (BIP) data set⁴¹. It was gathered by the AVIRIS sensor [159] during a flight campaign over the agricultural Indian Pines test site, in north-western Indiana. The scene was collected at the beginning of the 1992 growing season, and comprises several regular patches of different crops coupled with irregular forest and grass zones. The data cube comprises 1417×617 pixels, with a ground sampling distance of 20 mpp. Furthermore, each pixel comprises 220 channels recorded over a spectral range of 400-500 nm, with nominal spectral resolution of 10 nm. However, 20 bands [0-9, 210-219] were removed in order to avoid null, noisy and water absorption bands, keeping the remaining 200 bands for experimental purposes. The complexity of this challenging image is quite remarkable, as the pixels are very mixed due to the low spatial resolution, while the available ground-truth is composed of 58 different and highly unbalanced land-cover classes, covering only 20.33% of the samples. The size of the data set exceeds 1 GB. Fig. 12 shows the available labeled samples per class and their distribution in the scene.

C. Platform configuration

The designed experiments have been conducted over an OpenStack-based cloud infrastructure, which has been implemented onto a hardware platform composed of two \times Intel[®] Xeon[™] CPUs E5-2650v2 @2.60GHz with 8 cores (16 way multi-task processing), 160GB RAM and 600 GB of HDD SAS 10k. In this sense, within the cloud environment, 9 virtual machines (1 master instance and 8 slave instances) have been launched. Each virtual machine runs Ubuntu 20.04 as operating system, with Spark 3.0.1 and Java 9.0.4 serving as running platforms. Furthermore, the Spark framework provides the distributed *MLlib* library, which is used to support the implementation of our cloud-based MLP classifier⁴².

D. Experimental discussion

During the experimentation, the computational load of a fully connected deep network has been distributed in order to evaluate the performance of the cloud infrastructure when dealing with hyperspectral remote sensing image classification. In particular, a deep MLP has been designed to explore the impact of its computational burden over the cloud environment by involving all model parameters with all input elements in the computation of the matrix operations described by Eq. (1).

Table IV describes the architecture of the implemented MLP, specifying the number of layers and the number of neurons comprised by each layer. It is noteworthy that a fully-connected neural model entails $\sum_i N_i N_{i+1}$ trainable parameters (where N_i represents the number of nodes at the i^{th} layer with $i = [1, L - 1]$, and L is the number of layers), which in turn involves approximately $\sum_i 2N_i N_{i+1}$ floating point operations (FLOPs). In particular, the implemented MLP comprises 78,624 trainable parameters, which implies at least 157,248 FLOPs. Furthermore, in the calculation of FLOPs, the number of samples processed by the model must be taken into account.

In addition, two approaches have been followed to assess the obtained performance in terms of runtimes and speedup. Indeed, experiments have been conducted considering both different numbers of workers and different amounts of training data, where the first approach analyses the obtained runtimes and speedup by focusing on the number of workers, while the second approach evaluates the scalability of the cloud environment by focusing on the amount of data. For each experiment, we conduct 5 Monte Carlo runs and report the average results.

1) *Approach based on the cloud environment size*: The first approach distributes the MLP over the cloud infrastructure to measure the overall performance improvement of the system by evaluating its potential with a different number of workers. In this regard, a cloud environment with 1, 2, 4 and 8 working nodes has been launched. Moreover, for each configuration, different amounts of training data have been considered. Particularly, 10%, 35%, 50%, 65% and 90% of the available labeled data have been randomly selected from the BIP dataset to learn the 78,624 trainable parameters comprised by the MLP.

Fig. 13(a) provides a graphical representation of the obtained results in terms of runtime. As we can observe, for each configuration of the cloud infrastructure, five measurements have been collected, which correspond to the number of training samples. In this regard, focusing on each configuration, the obtained runtimes increase as more data are included during the training stage. This is a reasonable and expected behaviour since, with the same resources, an increase of the problem size brings an according increase in the runtime. However, comparing the runtimes among the different configurations, with only one worker the execution time soars, easily exceeding 20,000 seconds with 90% of training samples (purple bar). On the contrary, as the number of workers increases, the observed runtimes decrease significantly. In particular, with 8 workers, the highest runtime with 90% of the training samples never exceeds 4000 seconds.

These results have a clear impact on the speedup of the MLP model, which is increased as more workers are launched into the cloud environment. Fig. 13(b) provides a graphical description of the obtained speedups. Once more, for each configuration, five measurements are depicted corresponding with different training percentages. In this regard, the cloud infrastructure with 1 worker node has been considered as the baseline configuration, which exhibits a speedup of 1 for all training sizes. Therefore, the speedups of the following

TABLE IV

CONFIGURATION OF THE IMPLEMENTED MLP CLASSIFIER (NUMBER OF NEURONS PER LAYER)

Input	Hidden 1	Hidden 2	Hidden 3	Output
200	144	144	144	58

⁴¹<https://engineering.purdue.edu/biehl/MultiSpec/hyperspectral.html>

⁴²The source code used in this experimentation is currently available on <https://github.com/mhaut/cloud-dnn-HSI>

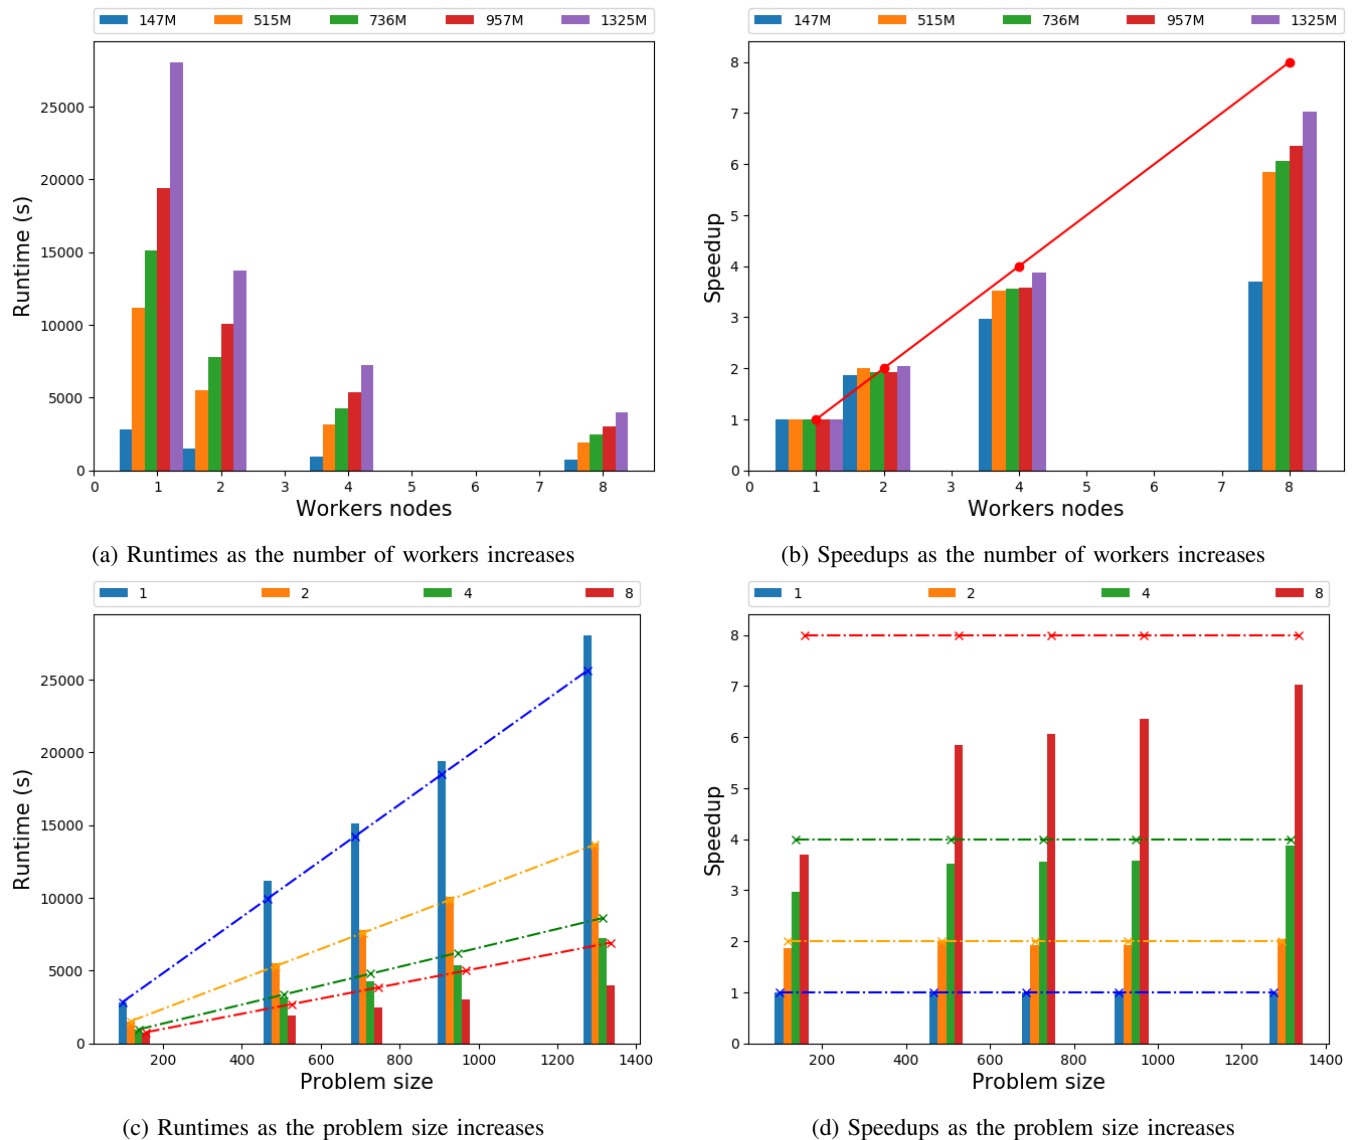


Fig. 13. Distributed MLP performance. (a) and (b) provide the runtimes and speedups as the number of workers increases, considering different amounts of training samples for each one. (c) and (d) provide the obtained runtimes and speedups as the problem size increases, evaluating the theoretical and obtained behaviour for a different number of workers. The different problem sizes result from the selection of different amounts of training data. Depicted lines provide theoretical speedups and expected runtime measurements, while bars provide the actual values obtained.

configurations have been consequently obtained. It should be noted that the improvement in speedup is not exactly the same as the theoretical speedup marked by the red line in Fig. 13(b), as communication and scheduling times inevitably affect the system performance. However, the data volume is high enough to take full advantage of the cloud environment, without the communication bottleneck preventing a good performance result in terms of runtimes. This is clearly evident in every configuration of the cloud environment. For instance, focusing on configurations with 2 and 4 workers, the speedups obtained with different training percentages are quite similar, suggesting that already with 10-50% of training data, computational resources are being optimally exploited. However, with 8 working nodes, the speedup is significantly higher when more data are processed, as computational resources are much larger, thus providing more room to exploit a bigger amount

of data (i.e., to process larger datasets).

2) *Approach based on problem size:* As mentioned above, the second approach evaluates the behaviour of the implemented cloud solution by placing the focus on the problem size, i.e., by considering different training set sizes. In this regard, the experiment attempts to measure the scalability of each of the cloud environment configurations adopted to solve the problem (with 1, 2, 4 and 8 nodes) when different amounts of labeled samples are considered at the training stage. As in the previous experiment, 10%, 35%, 50%, 65% and 90% of the available labeled samples have been randomly selected to comprise the training sets, resulting in different data sizes (in MBs), which are indicated on the x-axis of the plots reported in Figs. 13(c) and 13(d).

Fig. 13(c) provides the obtained results in terms of runtimes. Following the previous results, for each amount of training

TABLE V
OVERALL ACCURACY COMPARISON

Nodes	10%	35%	50%	65%	90%
1	63.65±0.36	77.57±0.62	81.27±0.29	83.04±0.21	85.16±0.34
2	63.50±0.28	77.67±0.63	80.99±0.14	83.25±0.26	85.39±0.08
4	63.62±0.48	77.56±0.12	81.15±0.38	83.22±0.33	85.80±0.09
8	63.24±0.51	77.76±0.34	81.25±0.18	83.44±0.11	85.40±0.36
Parallel (PyTorch)	61.05 ±1.81	77.45 ±0.49	80.50 ±0.34	82.24 ±0.38	84.37 ±0.33




Fig. 14. Runtime details of the distributed MLP as the problem size grows, with 4 (a) and 8 (b) working nodes. The different problem sizes result from the selection of different amounts of training data. Red lines provide expected measurements while blue bars provide the actual measured values.




Fig. 15. Classification map obtained by the cloud-based MLP for the BIP dataset using 50% of the available labeled data for training.

data, runtimes decrease as more workers are launched into the cloud environment. In this sense, the 1-worker configuration is the slowest one, with runtimes that are more than seven times longer than the ones obtained by the 8-workers configuration for the case with the most training data. In this sense, if we connect the top of each bar, it is particularly interesting to observe the slope of the line, where the one corresponding

to the 1-worker configuration is the steepest. It even exceeds the theoretical line, which makes it the worst configuration. On the contrary, the 4-workers and 8-workers configurations scale remarkably well. Moreover, they even give better times than theoretically expected. Figs. 14(a) and 14(b) provide the runtime details for these configurations. As it can be observed in Fig. 14(a), although the runtime increases (as expected), the curve is not entirely linear, which implies that the distribution model is more optimal when there are more data to distribute. Moreover, the distance between the obtained runtimes and the theoretical ones (which are highlighted as a red line) increases as more data are processed. Particularly, the theoretical runtimes times are 1.18 times higher than the obtained ones. This proves that the model scales appropriately with the size of the problem. This is clearly visible in Fig. 14(b), where theoretical runtimes are 1.59 times higher than those currently obtained.

Finally, Fig. 13(d) provides the obtained speedups regarding the problem size. Once more, the 1-worker configuration has been considered as the baseline where, for each size of the training set, its speedup is set to 1. Therefore, the speedups of the 2, 4 and 8-workers configurations have been consequently obtained. As in the previous plots, the theoretical speedups have been marked as dotted lines for each configuration. In this sense, the speedup exhibited by the 2-workers configuration is quite close to the theoretical one, while for the 4 and 8-workers configurations, their speedups improve as the size of the processed data grows. This is particularly evident when 8 working nodes are launched into the cloud environment. These results indicate that the computational resources provided by the 8-workers configuration exhibit great scalability, with a great potential to process bigger remote sensing datasets in order to optimise the use of the cloud environment capacities.

3) *Accuracy evaluation*: Finally, the reliability of the classification results has been measured in terms of overall accuracy (OA). In this sense, Table V provides the OAs, which have been obtained after training the deep MLP with 10%, 35%, 50%, 65% and 90% of randomly selected samples. Furthermore, for the cloud distributed MLP, configurations with 1, 2, 4 and 8 working nodes have been considered. These results have been compared with a parallel implementation based on the PyTorch framework. As we can observe, the obtained OAs improve as the MLP network is trained with more samples. Moreover, after comparing the different implementations, we can see that the obtained results are quite similar. Fig. 15 provides a classification map that has been obtained by training the cloud-based MLP with 50% of the available labeled samples. In this regard, the cloud solution not only provides an

efficient way to distribute the storage and computation load, but also reaches good performance in terms of accuracy.

To conclude this section, we emphasize that the results that have been obtained with the MLP are perfectly extrapolable to other deep networks such as CNNs. Indeed, the cloud environment can run both architectures in a distributed manner (as it is not a specialized hardware), reaching impressive performance in image analysis with CNNs as well [160].

VII. CONCLUSIONS AND FUTURE LINES

In this paper, we have presented a comprehensive review of recent efforts in parallel and distributed processing of remotely sensed images, with particular emphasis on DL-based approaches and their cloud implementation. Our review reflects the growing importance of using cloud computing techniques for distributed processing of remote sensing images, which is of great importance due to the current availability of open big remote sensing data repositories. Our review also summarized the processing tools and techniques that have been used in different remote sensing applications, which is believed to provide a useful guideline for new users that wish to develop computationally efficient techniques in this field.

We provide a case study illustrating the results obtained by a DL algorithm (implemented in the cloud) when processing a big hyperspectral image. Since hyperspectral data are characterized by their large size and complex processing and storage requirements, we believe that the results provided in our case study offer a good perspective on the possibilities of implementing DL algorithms in the cloud for addressing the processing challenges involved in the extraction of information from remotely sensed images. In the future, we will expand our study by considering the inclusion and optimization of specific accelerators (such as GPUs) in the cloud environment for DL-based remote sensing data processing and interpretation.

REFERENCES

- [1] B. Zhang, Z. Chen, D. Peng, J. A. Benediktsson, B. Liu, L. Zou, J. Li, and A. Plaza, "Remotely sensed big data: evolution in model development for information extraction [point of view]," *Proceedings of the IEEE*, vol. 107, no. 12, pp. 2294–2301, 2019.
- [2] J. Li, J. A. Benediktsson, B. Zhang, T. Yang, and A. Plaza, "Spatial technology and social media in remote sensing: A survey," *Proceedings of the IEEE*, vol. 105, no. 10, pp. 1855–1864, 2017.
- [3] C.-I. Chang, *Hyperspectral imaging: Techniques for spectral detection and classification*. Kluwer Academic/Plenum Publishers: New York., 2003.
- [4] A. Plaza, J. Plaza, A. Paz, and S. Sanchez, "Parallel Hyperspectral Image and Signal Processing," *IEEE Signal Processing Magazine*, vol. 28, pp. 196–218, 2011.
- [5] R. O. Green, M. L. Eastwood, C. M. Sarture, T. G. Chrien, M. Aronson, B. J. Chippendale, J. A. Faust, B. E. Pavri, C. J. Chovit, M. Solis *et al.*, "Imaging spectroscopy and the airborne visible/infrared imaging spectrometer (aviris)," *Remote sensing of environment*, vol. 65, no. 3, pp. 227–248, 1998.
- [6] M. Chi, A. Plaza, J. A. Benediktsson, Z. Sun, J. Shen, and Y. Zhu, "Big data for remote sensing: Challenges and opportunities," *Proceedings of the IEEE*, vol. 104, no. 11, pp. 2207–2219, 2016.
- [7] Y. M. L. Wang, J. Yang, *Cloud computing in remote sensing*. CRC Press., 2019.
- [8] C.-I. C. A. Plaza, *High Performance computing in remote sensing*. CRC Press., 2006.
- [9] J. Behnke, T. H. Watts, B. Kobler, D. Lowe, S. Fox, and R. Meyer, "Eosdis petabyte archives: tenth anniversary," in *22nd IEEE/13th NASA Goddard Conference on Mass Storage Systems and Technologies (MSSST'05)*. IEEE, 2005, pp. 81–93.
- [10] B. Ryan and D. Cripe, "The Group on Earth Observations (GEO) through 2025," in *40th COSPAR Scientific Assembly*, vol. 40, Jan. 2014, pp. A0.1–1–14.
- [11] P. M. Mell and T. Grance, "Sp 800-145. the nist definition of cloud computing," National Institute of Standards and Technology, Gaithersburg, MD, USA, Tech. Rep., 2011.
- [12] J. Plaza, R. Pérez, A. Plaza, P. Martínez, and D. Valencia, "Parallel morphological/neural processing of hyperspectral images using heterogeneous and homogeneous platforms," *Cluster Computing*, vol. 11, no. 1, pp. 17–32, 2008.
- [13] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [14] L. Zhang, L. Zhang, and B. Du, "Deep learning for remote sensing data: A technical tutorial on the state of the art," *IEEE Geoscience and Remote Sensing Magazine*, vol. 4, no. 2, pp. 22–40, 2016.
- [15] L. Ma, Y. Liu, X. Zhang, Y. Ye, G. Yin, and B. A. Johnson, "Deep learning in remote sensing applications: A meta-analysis and review," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 152, pp. 166 – 177, 2019.
- [16] X. X. Zhu, D. Tuia, L. Mou, G.-S. Xia, L. Zhang, F. Xu, and F. Fraundorfer, "Deep learning in remote sensing: A comprehensive review and list of resources," *IEEE Geoscience and Remote Sensing Magazine*, vol. 5, no. 4, pp. 8–36, 2017.
- [17] M. Paoletti, J. Haut, J. Plaza, and A. Plaza, "Deep learning classifiers for hyperspectral imaging: A review," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 158, pp. 279 – 317, 2019.
- [18] A. Plaza, Q. Du, Y.-L. Chang, and R. L. King, "High performance computing for hyperspectral remote sensing," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 4, no. 3, pp. 528–544, 2011.
- [19] C. A. Lee, S. D. Gasster, A. Plaza, C. Chang, and B. Huang, "Recent developments in high performance computing for remote sensing: A review," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 4, no. 3, pp. 508–527, 2011.
- [20] A. F. Goetz, "Three decades of hyperspectral remote sensing of the earth: A personal view," *Remote Sensing of Environment*, vol. 113, pp. S5–S16, 2009.
- [21] N. Yokoya, C. Grohnfeldt, and J. Chanussot, "Hyperspectral and multispectral data fusion: A comparative review of the recent literature," *IEEE Geoscience and Remote Sensing Magazine*, vol. 5, no. 2, pp. 29–56, 2017.
- [22] C. Liu, J. Shang, P. W. Vachon, and H. McNairn, "Multiyear crop monitoring using polarimetric radarsat-2 data," *IEEE Transactions on Geoscience and Remote sensing*, vol. 51, no. 4, pp. 2227–2240, 2012.
- [23] A. Rosenqvist, M. Shimada, N. Ito, and M. Watanabe, "Alos palsar: A pathfinder mission for global-scale monitoring of the environment," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 11, pp. 3307–3316, 2007.
- [24] M. Moghaddam, D. Entekhabi, P. R. Moorcroft, Y. Lou, E. Chapin, S. S. Saatchi, R. H. Reichle, W. T. Crow, R. H. Cuenca, A. Tabatabaeejad *et al.*, "Airborne microwave observatory of sub-canopy and subsurface (airmoss) earth venture suborbital mission overview," *AGUFM*, vol. 2015, pp. B53A–0537, 2015.
- [25] W. Pitz and D. Miller, "The terrasars-x satellite," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 48, no. 2, pp. 615–622, 2010.
- [26] D. Geudtner, R. Torres, P. Snoeij, M. Davidson, and B. Rommen, "Sentinel-1 system capabilities and applications," in *2014 IEEE Geoscience and Remote Sensing Symposium*. IEEE, 2014, pp. 1457–1460.
- [27] D. Yi, J. P. Harbeck, S. S. Manizade, N. T. Kurtz, M. Studinger, and M. Hofton, "Arctic sea ice freeboard retrieval with waveform characteristics for nasa's airborne topographic mapper (atm) and land, vegetation, and ice sensor (Ivis)," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, no. 3, pp. 1403–1410, 2014.
- [28] N. Levin, K. Johansen, J. M. Hacker, and S. Phinn, "A new source for high spatial resolution night time images—the eros-b commercial satellite," *Remote Sensing of Environment*, vol. 149, pp. 1–12, 2014.
- [29] Y. Chen, R. Fan, M. Bilal, X. Yang, J. Wang, and W. Li, "Multilevel cloud detection for high-resolution remote sensing imagery using multiple convolutional neural networks," *ISPRS International Journal of Geo-Information*, vol. 7, no. 5, p. 181, 2018.
- [30] G. Dial, H. Bowen, F. Gerlach, J. Grodecki, and R. Oleszczuk, "Ikonos satellite, imagery, and products," *Remote sensing of Environment*, vol. 88, no. 1-2, pp. 23–36, 2003.
- [31] M. Gašparović, L. Rumora, M. Miler, and D. Medak, "Effect of fusing sentinel-2 and worldview-4 imagery on the various vegetation indices," *Journal of Applied Remote Sensing*, vol. 13, no. 3, p. 036503, 2019.

TABLE VI
LIST OF ACRONYMS USED IN THIS PAPER


Acronym	Full name	Acronym	Full name
EO	Earth Observation	MTFC	Multi-GPU Training Framework
GPUs	Graphics Processing Units	PNPE	Parallel Neighbor Pixel Extractor
CPUs	Central Processing Units	G-POD	Grid Processing on Demand
AVIRIS	Airborne Visible/Infrared Imaging Spectrometer	GENESI-DR	Ground European Network for Earth Science Interoperations and Digital Repositories
AVIRIS-NG	Airborne Visible/Infrared Imaging Spectrometer New Generation	GiSHEO	Grid Services for training and High Education in Earth Observation
NASA	National Aeronautics and Space Administration	CEOS	Committee on Earth Observation Satellites
TBs	Terabytes	DGGS	Discrete Global Grid Systems
GBs	Gigabytes	CBIR	Content-based Image Retrieval
PBs	Petabytes	ICP	InterIMAGE cloud Platform
mpp	Meters Per Pixel	HDFS	Hadoop File System
EnMAP	Environmental Mapping and Analysis Program	YARN	Yet Another Resource Negotiator
EOSDIS	Earth Observing System Data and Information System	RDDS	Resilient Distributed Datasets
NSMC	China National Satellite Meteorological Center	VMs	Virtual Machine
CCRSDA	China Center for Resources Satellite Data and Application	SaaS	Software as a Service
ESA	European Space Agency	PaaS	Platform as a Service
GEO	Group on Earth Observations	IaaS	Infrastructure as a Service
HPC	High-Performance Computing	API	Application Programming Interface
SAR	Synthetic Aperture Radar	SDNs	Software Defined Networks
LIDAR	Light Detection and Ranging	DNs	Domain Name Servers
OLI	Operational Land Imager	GFS	Google File System
ASTER	Advanced Spaceborne Thermal Emission and Reflection Radiometer	DAG	Directed Acyclic Graph
VNIR	Visible and Near-Infrared	I/O	Input/Output
TIR	Thermal Infrared	HPCC	High Performance Computing Challenger
SWIR	Shortwave Infrared	AWS	Amazon Web Services
GSD	Ground Sample Distance	HPL	High Performance Linpack
PAN	Panchromatic	EC2	Elastic Computing Cloud
MSI	Multi-spectral Images	GCE	Google Compute Engine
HSI	Hyper-spectral Images	SL	IBM SoftLayer
PCA	Principal Component Analysis	MRAP	MapReduce with Access Patterns
TCT	Tasseled Cap Transformation	MPI	Message Passing Interface
WT	Wavelet Transform	BLAS	Basic Linear Algebra Subprograms
KNN	K-nearest Neighbors	ML	Machine Learning
DTs	Decision Trees	DL	Deep Learning
RFs	Random Forests	OS	Operating System
MLR	Multinomial Logistic Regression	MAP	Maximum A Posteriori
GMMs	Gaussian Mixture Models	BS	Block Size
NB	Naive Bayes-based approaches	MLP	Multi-Layer Perceptron
HMMs	Hidden Markov models	ReLU	Rectified Linear Units
SVMs	Support Vector Machines	BFGS	Broyden-Fletcher-Goldfarb-Shanno
ANNs	Artificial Neural Networks	DGEMM	Double-precision Matrix-Matrix Multiplication
SAEs	Stacked AutoEncoders	TF	TensorFlow
DBNs	Deep Belief Networks	PS	Parameter Server
RNNs	Recurrent Neuronal Networks	IP	Indian Pines
CNNs	Convolutional Neural Networks	BIP	Big Indian Pines
DNNs	Deep Neural Networks	HDD	Hard Disk Drive
RBMs	Restricted Boltzmann Machines	RAM	Random Access Memory
LSTM	Long Short Term Memory	FLOPs	Floating Point Operations
RESFlow	Remote Sensing data Flow	PRACE	Partnership for Advanced Computing in Europe
InSPIRE	Integrated Sustainable Pan-European Infrastructure for Researchers in Europe	UCI	Unified Cloud Interface
XSEDE	Extreme Science and Engineering Discovery Environment		

- [32] M. Drusch, U. Del Bello, S. Carlier, O. Colin, V. Fernandez, F. Gascon, B. Hoersch, C. Isola, P. Laberinti, P. Martimort *et al.*, "Sentinel-2: Esa's optical high-resolution mission for gmes operational services," *Remote sensing of Environment*, vol. 120, pp. 25–36, 2012.
- [33] C. Donlon, B. Berruti, A. Buongiorno, M.-H. Ferreira, P. Féménias, J. Frerick, P. Goryl, U. Klein, H. Laur, C. Mavrocordatos *et al.*, "The global monitoring for environment and security (gmes) sentinel-3 mission," *Remote Sensing of Environment*, vol. 120, pp. 37–57, 2012.
- [34] D. P. Roy, M. A. Wulder, T. R. Loveland, C. E. Woodcock, R. G. Allen, M. C. Anderson, D. Helder, J. R. Irons, D. M. Johnson, R. Kennedy *et al.*, "Landsat-8: Science and product vision for terrestrial global change research," *Remote sensing of Environment*, vol. 145, pp. 154–172, 2014.
- [35] A. Savtchenko, D. Ouzounov, S. Ahmad, J. Acker, G. Leptoukh, J. Koziana, and D. Nickless, "Terra and aqua modis products available from nasa ges daac," *Advances in Space Research*, vol. 34, no. 4, pp. 710–714, 2004.
- [36] M. A. Cutter, D. R. Lobb, and R. A. Cockshott, "Compact high resolution imaging spectrometer (chris)," *Acta Astronautica*, vol. 46, no. 2-6, pp. 263–268, 2000.
- [37] J. W. Chapman, D. R. Thompson, M. C. Helmlinger, B. D. Bue, R. O. Green, M. L. Eastwood, S. Geier, W. Olson-Duvall, and S. R. Lundeen, "Spectral and radiometric calibration of the next generation airborne visible infrared spectrometer (aviris-ng)," *Remote Sensing*, vol. 11, no. 18, p. 2129, 2019.
- [38] B. Kunkel, F. Blechinger, R. Lutz, R. Doerffer, H. van der Piepen, and M. Schroder, "ROSIS (Reflective Optics System Imaging Spectrometer) - A candidate instrument for polar platform missions," in *Proc. SPIE 0868 Optoelectronic technologies for remote sensing from space*, J. Seeley and S. Bowyer, Eds., 1988, p. 8.
- [39] S. Babey and C. Anger, "A compact airborne spectrographic imager (casi)," in *Quantitative Remote Sensing: An Economic Tool for the Nineties, Volume 1*, 1989, pp. 1028–1031.
- [40] L. J. Rickard, R. W. Basedow, E. F. Zalewski, P. R. Silverglate, and M. Landers, "Hydice: An airborne system for hyperspectral imaging," in *Imaging Spectrometry of the Terrestrial Environment*, vol. 1937. International Society for Optics and Photonics, 1993, pp. 173–180.
- [41] T. Cocks, R. Jenssen, A. Stewart, I. Wilson, and T. Shields, "The hymaptm airborne hyperspectral sensor: the system, calibration and performance," in *Proceedings of the 1st EARSeL workshop on Imaging Spectroscopy*. EARSeL, 1998, pp. 37–42.
- [42] P. Mouroulis, B. Van Gorp, R. O. Green, H. Dierssen, D. W. Wilson, M. Eastwood, J. Boardman, B.-C. Gao, D. Cohen, B. Franklin *et al.*, "Portable remote imaging spectrometer coastal ocean sensor: design, characteristics, and first flight results," *Applied optics*, vol. 53, no. 7, pp. 1363–1380, 2014.
- [43] L. Guanter, H. Kaufmann, K. Segl, S. Foerster, C. Rogass, S. Chabrilat, T. Kuester, A. Hollstein, G. Rossner, C. Chlebek *et al.*, "The enmap spaceborne imaging spectroscopy mission for earth observation," *Remote Sensing*, vol. 7, no. 7, pp. 8830–8857, 2015.
- [44] N. Yokoya and A. Iwasaki, "Hyperspectral and multispectral data fusion mission on hyperspectral imager suite (hisui)," in *2013 IEEE International Geoscience and Remote Sensing Symposium-IGARSS*. IEEE, 2013, pp. 4086–4089.
- [45] A. Eckardt, J. Horack, F. Lehmann, D. Krutz, J. Drescher, M. Whorton, and M. Soutullo, "Desis (dlr earth sensing imaging spectrometer for the iss-buses platform)," in *2015 IEEE international geoscience and remote sensing symposium (IGARSS)*. IEEE, 2015, pp. 1457–1459.
- [46] J. S. Pearlman, P. S. Barry, C. C. Segal, J. Shepanski, D. Beiso, and S. L. Carman, "Hyperion, a space-based imaging spectrometer," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 41, no. 6, pp. 1160–1173, 2003.
- [47] C. Galeazzi, A. Sacchetti, A. Cisbani, and G. Babini, "The prisma program," in *Geoscience and Remote Sensing Symposium, 2008. IGARSS 2008. IEEE International*, vol. 4. IEEE, 2008, pp. IV–105.
- [48] T. Feingersh and E. B. Dor, "Shalom—a commercial hyperspectral space mission," *Optical payloads for space missions*, pp. 247–263, 2015.
- [49] G. A. Swayze, R. N. Clark, A. F. Goetz, T. G. Chrien, and N. S.

- Gorelick, "Effects of spectrometer band pass, sampling, and signal-to-noise ratio on spectral identification using the tetracorder algorithm," *Journal of Geophysical Research: Planets*, vol. 108, no. E9, pp. 5105–5135, 2003.
- [50] A. F. H. Goetz, G. Vane, J. E. Solomon, and B. N. Rock, "Imaging spectrometry for earth remote sensing," *Science*, vol. 228, no. 4704, pp. 1147–1153, 1985.
- [51] S. K. Seelan, S. Laguette, G. M. Casady, and G. A. Seielstad, "Remote sensing applications for precision agriculture: A learning community approach," *Remote sensing of environment*, vol. 88, no. 1-2, pp. 157–169, 2003.
- [52] H. M. Pham, Y. Yamaguchi, and T. Q. Bui, "A case study on the relation between city planning and urban growth using remote sensing and spatial metrics," *Landscape and Urban Planning*, vol. 100, no. 3, pp. 223–230, 2011.
- [53] S. A. Azzouzi, A. Vidal-Pantaleoni, and H. A. Bentounes, "Desertification monitoring in biskra, algeria, with landsat imagery by means of supervised classification and change detection methods," *IEEE Access*, vol. 5, pp. 9065–9072, 2017.
- [54] X. Ceamanos and S. Valero, "Processing hyperspectral images," in *Optical Remote Sensing of Land Surface*. Elsevier, 2016, pp. 163–200.
- [55] A. Maffei, J. M. Haut, M. E. Paoletti, J. Plaza, L. Bruzzone, and A. Plaza, "A single model cnn for hyperspectral image denoising," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 4, pp. 2516–2529, 2019.
- [56] G. Cheng, X. Xie, J. Han, L. Guo, and G.-S. Xia, "Remote sensing image scene classification meets deep learning: Challenges, methods, benchmarks, and opportunities," *arXiv preprint arXiv:2005.01094*, 2020.
- [57] M. E. Paoletti, J. M. Haut, J. Plaza, and A. Plaza, "A new deep convolutional neural network for fast hyperspectral image classification," *ISPRS journal of photogrammetry and remote sensing*, vol. 145, pp. 120–147, 2018.
- [58] V. Walter, "Object-based classification of remote sensing data for change detection," *ISPRS Journal of photogrammetry and remote sensing*, vol. 58, no. 3-4, pp. 225–238, 2004.
- [59] K. Nogueira, O. A. Penatti, and J. A. Dos Santos, "Towards better exploiting convolutional neural networks for remote sensing scene classification," *Pattern Recognition*, vol. 61, pp. 539–556, 2017.
- [60] J. Plaza, A. Plaza, R. Perez, and P. Martinez, "On the Use of Small Training Sets for Neural Network-Based Characterization of Mixed Pixels in Remotely Sensed Hyperspectral Images," *Pattern Recognition*, vol. 42, pp. 3032–3045, 2009.
- [61] J. A. G. Jaramago, M. E. Paoletti, J. M. Haut, R. Fernandez-Beltran, A. Plaza, and J. Plaza, "Gpu parallel implementation of dual-depth sparse probabilistic latent semantic analysis for hyperspectral unmixing," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 12, no. 9, pp. 3156–3167, 2019.
- [62] G. Camps-Valls, "Machine learning in remote sensing data processing," in *2009 IEEE international workshop on machine learning for signal processing*. IEEE, 2009, pp. 1–6.
- [63] D. J. Lary, A. H. Alavi, A. H. Gandomi, and A. L. Walker, "Machine learning in geosciences and remote sensing," *Geoscience Frontiers*, vol. 7, no. 1, pp. 3–10, 2016.
- [64] J. M. Haut, R. Fernandez-Beltran, M. E. Paoletti, J. Plaza, A. Plaza, and F. Pla, "A new deep generative network for unsupervised remote sensing single-image super-resolution," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 11, pp. 6792–6810, 2018.
- [65] D. Tuia and G. Camps-Valls, "Semisupervised remote sensing image classification with cluster kernels," *IEEE Geoscience and Remote Sensing Letters*, vol. 6, no. 2, pp. 224–228, 2009.
- [66] D. Tuia, M. Volpi, L. Copa, M. Kanevski, and J. Munoz-Mari, "A survey of active learning algorithms for supervised remote sensing image classification," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 3, pp. 606–617, 2011.
- [67] Y. Li, K. Fu, H. Sun, and X. Sun, "An aircraft detection framework based on reinforcement learning and convolutional neural networks in remote sensing images," *Remote Sensing*, vol. 10, no. 2, p. 243, 2018.
- [68] J. M. Haut, M. Paoletti, J. Plaza, and A. Plaza, "Cloud implementation of the k-means algorithm for hyperspectral image analysis," *The Journal of Supercomputing*, vol. 73, no. 1, pp. 514–529, 2017.
- [69] E. Blanzieri and F. Melgani, "Nearest neighbor classification of remote sensing images with the maximal margin principle," *IEEE Transactions on geoscience and remote sensing*, vol. 46, no. 6, pp. 1804–1811, 2008.
- [70] S. Delalieux, B. Somers, B. Haest, T. Spanhove, J. V. Borre, and C. Múcher, "Heathland conservation status mapping through integration of hyperspectral mixture analysis and decision tree classifiers," *Remote sensing of environment*, vol. 126, pp. 222–231, 2012.
- [71] K. Y. Peerbhaya, O. Mutanga, and R. Ismail, "Random forests unsupervised classification: the detection and mapping of solanum mauritanum infestations in plantation forestry using hyperspectral data," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 8, no. 6, pp. 3107–3122, 2015.
- [72] J. M. Haut and M. E. Paoletti, "Cloud implementation of multinomial logistic regression for uav hyperspectral images," *IEEE Journal on Miniaturization for Air and Space Systems*, vol. 1, no. 3, pp. 163–171, 2020.
- [73] J. Ju, E. D. Kolaczyk, and S. Gopal, "Gaussian mixture discriminant analysis and sub-pixel land cover characterization in remote sensing," *Remote Sensing of Environment*, vol. 84, no. 4, pp. 550–560, 2003.
- [74] J. Yang, Z. Ye, X. Zhang, W. Liu, and H. Jin, "Attribute weighted naive bayes for remote sensing image classification based on cuckoo search algorithm," in *2017 International Conference on Security, Pattern Analysis, and Cybernetics (SPAC)*. IEEE, 2017, pp. 169–174.
- [75] P. B. C. Leite, R. Q. Feitosa, A. R. Formaggio, G. A. O. P. da Costa, K. Pakzad, and I. D. Sanches, "Hidden markov models for crop recognition in remote sensing image sequences," *Pattern Recognition Letters*, vol. 32, no. 1, pp. 19–26, 2011.
- [76] M. E. Paoletti, J. M. Haut, X. Tao, J. P. Miguel, and A. Plaza, "A new gpu implementation of support vector machines for fast hyperspectral image classification," *Remote Sensing*, vol. 12, no. 8, p. 1257, 2020.
- [77] S. Yang, Q. Feng, T. Liang, B. Liu, W. Zhang, and H. Xie, "Modeling grassland above-ground biomass based on artificial neural network and remote sensing in the three-river headwaters region," *Remote Sensing of Environment*, vol. 204, pp. 448–455, 2018.
- [78] Y. Chen, Z. Lin, X. Zhao, G. Wang, and Y. Gu, "Deep learning-based classification of hyperspectral data," *IEEE Journal of Selected topics in applied earth observations and remote sensing*, vol. 7, no. 6, pp. 2094–2107, 2014.
- [79] M. E. Paoletti, J. M. Haut, J. Plaza, and A. Plaza, "Scalable recurrent neural network for hyperspectral image classification," *The Journal of Supercomputing*, pp. 1–17, 2020.
- [80] M. E. Paoletti, J. M. Haut, R. Fernandez-Beltran, J. Plaza, A. J. Plaza, and F. Pla, "Deep pyramidal residual networks for spectral-spatial hyperspectral image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 2, pp. 740–754, 2018.
- [81] M. E. Paoletti, J. M. Haut, R. Fernandez-Beltran, J. Plaza, A. Plaza, J. Li, and F. Pla, "Capsule networks for hyperspectral image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 4, pp. 2145–2160, 2018.
- [82] D. Lunga, J. Gerrand, L. Yang, C. Layton, and R. Stewart, "Apache spark accelerated deep learning inference for large scale satellite image analytics," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 13, pp. 271–283, 2020.
- [83] M. Ma, H. Pouransari, D. Chao, S. Adya, S. Serrano, Y. Qin, D. Gimnichier, and D. Walsh, "Democratizing production-scale distributed deep learning," *ArXiv*, vol. abs/1811.00143, 2018.
- [84] R. Sedona, G. Cavallaro, J. Jitsev, A. Strube, M. Riedel, and J. A. Benediktsson, "Remote sensing big data classification with high performance distributed deep learning," *Remote Sensing*, vol. 11, no. 24, 2019. [Online]. Available: <https://www.mdpi.com/2072-4292/11/24/3056>
- [85] Y. Lu, K. Xie, G. Xu, H. Dong, C. Li, and T. Li, "Mtfc: A multi-gpu training framework for cube-cnn-based hyperspectral image classification," *IEEE Transactions on Emerging Topics in Computing*, pp. 1–1, 2020.
- [86] M. Aspri, G. Tsagakatakis, and P. Tsakalides, "Distributed training and inference of deep learning models for multi-modal land cover classification," *Remote Sensing*, vol. 12, no. 17, 2020. [Online]. Available: <https://www.mdpi.com/2072-4292/12/17/2670>
- [87] Z. Wu, Y. Li, A. Plaza, J. Li, F. Xiao, and Z. Wei, "Fast principal component analysis for hyperspectral imaging based on cloud computing," in *IEEE International Geoscience and Remote Sensing Symposium (IEEE IGARSS 2015), Milan, Italy*, 2015.
- [88] —, "Parallel and Distributed Dimensionality Reduction of Hyperspectral Data on Cloud Computing Architectures," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. PP, no. 99, pp. 1–9, 2016.
- [89] J. Gu, Z. Wu, Y. Li, Y. Chen, Z. Wei, and W. Wang, "Parallel Optimization of Pixel Purity Index Algorithm for Hyperspectral Unmixing Based on Spark," *2015 Third International Conference on Advanced Cloud and Big Data*, pp. 159–166, 2015.

- [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7435468>
- [90] C. Bielski, S. Gentilini, and M. Pappalardo, "Post-disaster image processing for damage analysis using genesi-dr, wps and grid computing," *Remote Sensing*, vol. 3, no. 6, pp. 1234–1250, 2011. [Online]. Available: <https://www.mdpi.com/2072-4292/3/6/1234>
- [91] J. M. Haut, J. A. Gallardo, M. E. Paoletti, G. Cavallaro, J. Plaza, A. Plaza, and M. Riedel, "Cloud deep networks for hyperspectral image analysis," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 12, pp. 9832–9848, 2019.
- [92] Z. Wu, J. Sun, Y. Zhang, Y. Zhu, J. Li, A. Plaza, J. A. Benediktsson, and Z. Wei, "Scheduling-guided automatic processing of massive hyperspectral image classification on cloud computing architectures," *IEEE Transactions on Cybernetics*, pp. 1–14, 2020.
- [93] L. Parente, E. Taquary, A. P. Silva, C. Souza, and L. Ferreira, "Next generation mapping: Combining deep learning, cloud computing, and big remote sensing data," *Remote Sensing*, vol. 11, no. 23, 2019. [Online]. Available: <https://www.mdpi.com/2072-4292/11/23/2881>
- [94] X. Yao, G. Li, J. Xia, J. Ben, Q. Cao, L. Zhao, Y. Ma, L. Zhang, and D. Zhu, "Enabling the big earth observation data via cloud computing and dggs: Opportunities and challenges," *Remote Sensing*, vol. 12, no. 1, 2020. [Online]. Available: <https://www.mdpi.com/2072-4292/12/1/62>
- [95] P. Zheng, Z. Wu, J. Sun, Y. Zhang, Y. Zhu, Y. Shen, J. Yang, Z. Wei, and A. Plaza, "A parallel unmixing-based content retrieval system for distributed hyperspectral imagery repository on cloud computing platforms," *Remote Sensing*, vol. 13, no. 2, 2021. [Online]. Available: <https://www.mdpi.com/2072-4292/13/2/176>
- [96] J. M. Bioucas-Dias, A. Plaza, N. Dobigeon, M. Parente, Q. Du, P. Gader, and J. Chanussot, "Hyperspectral unmixing overview: Geometrical, statistical and sparse regression-based approaches," *{IEEE} J. Sel. Topics Appl. Earth Observations Remote Sens.*, vol. 5, no. 2, pp. 354–379, 2012.
- [97] W. Huang, L. Meng, D. Zhang, and W. Zhang, "In-memory parallel processing of massive remotely sensed data using an apache spark on hadoop yarn model," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 10, no. 1, pp. 3–19, 2017.
- [98] V. A. Ayma, R. S. Ferreira, P. N. Happ, D. A. B. Oliveira, G. A. O. P. Costa, R. Q. Feitosa, A. Plaza, and P. Gamba, "On the architecture of a big data classification tool based on a map reduce approach for hyperspectral image analysis," in *2015 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 2015, pp. 1508–1511.
- [99] P. Bajcsy, P. Nguyen, A. Vandecreme, and M. Brady, "Spatial computations over terabyte-sized images on hadoop platforms," *Proceedings - 2014 IEEE International Conference on Big Data, IEEE Big Data 2014*, pp. 816–824, 2015.
- [100] Happ, P. N., R. S. Ferreira, G. A. O. P. Costa, R. Q. Feitosa, C. Bentes, and P. Gamba, "Towards distributed region growing image segmentation based on MapReduce," *Geoscience and Remote Sensing Symposium (IGARSS)*, no. July, pp. 4352–4355, 2015. [Online]. Available: <http://ieeexplore.ieee.org/xpls/abs/all.jsp?arnumber=7326790>
- [101] J. Xing and R. Sieber, "Sampling based image splitting in large scale distributed computing of earth observation data," *2014 IEEE Geoscience and Remote Sensing Symposium*, pp. 1409–1412, 2014. [Online]. Available: <http://www.scopus.com/inward/record.url?eid=s2-0-84911366834/&partnerID=tZOTx3y1>
- [102] X. Pan and S. Zhang, "A remote sensing image cloud processing system based on Hadoop," *Proceedings - 2012 IEEE 2nd International Conference on Cloud Computing and Intelligence Systems, IEEE CCIS 2012*, vol. 1, pp. 492–494, 2013.
- [103] Y. Liu, B. Chen, W. He, and Y. Fang, "Massive image data management using HBase and MapReduce," *International Conference on Geoinformatics*, no. 2011, 2013.
- [104] R. Rajak, D. Raveendran, M. C. Bh, and S. S. Medasani, "High Resolution Satellite Image Processing Using Hadoop Framework," *2015 IEEE International Conference on Cloud Computing in Emerging Markets (CCEM)*, pp. 16–21, 2015. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7436925>
- [105] M. M. Rathore, A. Ahmad, A. Paul, and A. Daniel, "Hadoop based real-Time Big Data Architecture for remote sensing Earth Observatory System," *6th International Conference on Computing, Communications and Networking Technologies, ICCCNT 2015*, vol. i, 2016.
- [106] W. Nina, R. Cruz, J. Serrano, J. Cuba, Y. Huaynacho, A. Mamani-Aliaga, Y. Yari, and P. Yanyachi, "A new approach to the massive processing of satellite images," *Proceedings - 2015 41st Latin American Computing Conference, CLEI 2015*, 2015.
- [107] Y. Zhong, J. Fang, and X. Zhao, "VegaIndexer: A Distributed composite index scheme for big spatio-temporal sensor data on cloud," *International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 1713–1716, 2013.
- [108] R. Kune, P. Konugurthi, A. Agarwal, R. R. Chillarige, and R. Buyya, "Xhami – extended hdfs and mapreduce interface for image processing applications," in *2015 IEEE International Conference on Cloud Computing in Emerging Markets (CCEM)*, 2015, pp. 43–51.
- [109] M. T. Patterson, N. Anderson, C. Bennett, J. Bruggemann, R. L. Grossman, M. Handy, V. Ly, D. J. Mandl, S. Pederson, J. Pivarski, R. Powell, J. Spring, W. Wells, and J. Xia, "The matsu wheel: A cloud-based framework for efficient analysis and reanalysis of earth satellite imagery," in *2016 IEEE Second International Conference on Big Data Computing Service and Applications (BigDataService)*, 2016, pp. 156–165.
- [110] A. Plaza, D. Valencia, J. Plaza, and P. Martinez, "Commodity cluster-based parallel processing of hyper spectral imagery," *Journal of Parallel and Distributed Computing*, vol. 66, no. 3, pp. 345–358, 2006.
- [111] G. Aloisio and M. Cafaro, "A dynamic earth observation system," *Parallel Computing*, vol. 29, no. 10, pp. 1357–1362, 2003.
- [112] D. Gorgan, V. Bacu, T. Stefanut, D. Rodila, and D. Mihon, "Grid based satellite image processing platform for earth observation application development," in *2009 IEEE International Workshop on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications*, 2009, pp. 247–252.
- [113] Z. Chen, N. Chen, C. Yang, and L. Di, "Cloud computing enabled web processing service for earth observation data processing," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 5, no. 6, pp. 1637–1649, 2012.
- [114] K. Shvachko, H. Kuang, S. Radia, and R. Chansler, "The Hadoop distributed file system," *2010 IEEE 26th Symposium on Mass Storage Systems and Technologies, MSST2010*, pp. 1–10, 2010.
- [115] V. Kumar Vavilapalli, A. C. Murthy, C. Douglas, S. Agarwal, M. Konar, R. Evans, T. Graves, J. Lowe, H. Shah, S. Seth, B. Saha, C. Curino, O. O 'malley, S. Radia, B. Reed, and E. Baldeschwieler, "Apache Hadoop YARN: Yet Another Resource Negotiator," *SOCC '13 Proceedings of the 4th annual Symposium on Cloud Computing*, vol. 13, pp. 1–3, 2013. [Online]. Available: <http://dx.doi.org/10.1145/2523616.2523633>
- [116] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica, "Spark : Cluster Computing with Working Sets," *HotCloud'10 Proceedings of the 2nd USENIX conference on Hot topics in cloud computing*, p. 10, 2010.
- [117] M. Zaharia, M. Chowdhury, T. Das, and A. Dave, "Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing," *NSDI'12 Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, pp. 2–2, 2012. [Online]. Available: <https://www.usenix.org/system/files/conference/nsdi12/nsdi12-final138.pdf>
- [118] X. Meng, J. Bradley, B. Yavuz, E. Sparks, S. Venkataraman, D. Liu, J. Freeman, D. Tsai, M. Amde, S. Owen, D. Xin, R. Xin, M. J. Franklin, R. Zadeh, M. Zaharia, and A. Talwalkar, "MLlib: Machine Learning in Apache Spark," *CoRR*, vol. 17, pp. 1–7, 2015. [Online]. Available: <http://arxiv.org/abs/1505.06807>
- [119] D. Marinescu, *Cloud Computing: Theory and Practice*. Elsevier Science, 2017. [Online]. Available: <https://books.google.es/books?id=O9smDwAAQBAJ>
- [120] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "A view of cloud computing," *Commun. ACM*, vol. 53, no. 4, pp. 50–58, Apr. 2010. [Online]. Available: <https://doi.org/10.1145/1721654.1721672>
- [121] J. Smith and R. Nair, "The architecture of virtual machines," *Computer*, vol. 38, pp. 32–38, 2005.
- [122] M. Pearce, S. Zeadally, and R. Hunt, "Virtualization: Issues, security threats, and solutions," *ACM Comput. Surv.*, vol. 45, no. 2, Mar. 2013. [Online]. Available: <https://doi.org/10.1145/2431211.2431216>
- [123] V. Mauch, M. Kunze, and M. Hillenbrand, "High performance cloud computing," *Future Generation Computer Systems*, vol. 29, no. 6, pp. 1408 – 1416, 2013, including Special sections: High Performance Computing in the Cloud & Resource Discovery Mechanisms for P2P Systems. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167739X12000647>
- [124] (2009) Security guidance for critical areas of focus in cloud computing v2.1. [Online]. Available: <https://cloudsecurityalliance.org/csaguide.pdf>
- [125] I. Sadooghi, J. H. Martin, T. Li, K. Brandstatter, K. Maheshwari, T. P. P. de Lacerda Ruivo, G. Garzoglio, S. Timm, Y. Zhao, and I. Raicu, "Understanding the performance and potential of cloud computing for

- scientific applications,” *IEEE Transactions on Cloud Computing*, vol. 5, no. 2, pp. 358–371, 2017.
- [126] J. Dean and S. Ghemawat, “Mapreduce: Simplified data processing on large clusters,” in *Proceedings of the 6th Conference on Symposium on Operating Systems Design Implementation - Volume 6*, ser. OSDI’04. USA: USENIX Association, 2004, p. 10.
- [127] —, “Mapreduce: Simplified data processing on large clusters,” *Commun. ACM*, vol. 51, no. 1, p. 107–113, Jan. 2008. [Online]. Available: <https://doi.org/10.1145/1327452.1327492>
- [128] S. Ghemawat, H. Gobioff, and S.-T. Leung, “The google file system,” *SIGOPS Oper. Syst. Rev.*, vol. 37, no. 5, p. 29–43, Oct. 2003. [Online]. Available: <https://doi.org/10.1145/1165389.945450>
- [129] J. Ekanayake, S. Pallickara, and G. Fox, “Mapreduce for data intensive scientific analyses,” in *2008 IEEE Fourth International Conference on eScience*, 2008, pp. 277–284.
- [130] M. Zaharia, *An Architecture for Fast and General Data Processing on Large Clusters*. Association for Computing Machinery and Morgan Claypool, 2016.
- [131] T. Gunarathne, T. Wu, J. Qiu, and G. Fox, “Mapreduce in the clouds for science,” in *2010 IEEE Second International Conference on Cloud Computing Technology and Science*, 2010, pp. 565–572.
- [132] S. Ostermann, A. Iosup, N. Yigitbasi, R. Prodan, T. Fahringer, and D. Epema, “A performance analysis of ec2 cloud computing services for scientific computing,” in *Cloud Computing*, D. R. Avresky, M. Diaz, A. Bode, B. Ciciani, and E. Dekel, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 115–131.
- [133] J. Dongarra and P. Luszczek, *HPC Challenge Benchmark*. Boston, MA: Springer US, 2011, pp. 844–850. [Online]. Available: https://doi.org/10.1007/978-0-387-09766-4_156
- [134] A. Iosup, S. Ostermann, M. N. Yigitbasi, R. Prodan, T. Fahringer, and D. Epema, “Performance analysis of cloud computing services for many-tasks scientific computing,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 22, no. 6, pp. 931–945, 2011.
- [135] P. Mehrotra, J. Djomehri, S. Heistand, R. Hood, H. Jin, A. Lazanoff, S. Saini, and R. Biswas, “Performance evaluation of amazon ec2 for nasa hpc applications,” in *Proceedings of the 3rd Workshop on Scientific Cloud Computing*, ser. ScienceCloud ’12. New York, NY, USA: Association for Computing Machinery, 2012, pp. 41–50. [Online]. Available: <https://doi.org/10.1145/2287036.2287045>
- [136] D. Bailey, E. Barszcz, J. Barton, D. Browning, R. Carter, L. Dagum, R. Fatoohi, P. Frederickson, T. Lasinski, R. Schreiber, H. Simon, V. Venkatakrisnan, and S. Weeratunga, “The nas parallel benchmarks,” *Int. J. High Perform. Comput. Appl.*, vol. 5, no. 3, pp. 63–73, Sep. 1991. [Online]. Available: <https://doi.org/10.1177/109434209100500306>
- [137] A. Petitet, R. Whaley, J. Dongarra, and A. Cleary. (2018, Dec.) Hpl – a portable implementation of the high-performance linpack benchmark for distributed-memory computers. [Online]. Available: <http://www.netlib.org/benchmark/hpl/>
- [138] G. Wang and T. S. E. Ng, “The impact of virtualization on network performance of amazon ec2 data center,” in *2010 Proceedings IEEE INFOCOM*, 2010, pp. 1–9.
- [139] J. R. Lange, K. Pedretti, P. Dinda, P. G. Bridges, C. Bae, P. Soltero, and A. Merritt, “Minimal-overhead virtualization of a large scale supercomputer,” in *Proceedings of the 7th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments*, ser. VEE ’11. New York, NY, USA: Association for Computing Machinery, 2011, pp. 169–180. [Online]. Available: <https://doi.org/10.1145/1952682.1952705>
- [140] P. Leitner and J. Cito, “Patterns in the chaos—a study of performance variation and predictability in public iaas clouds,” *ACM Trans. Internet Technol.*, vol. 16, no. 3, Apr. 2016. [Online]. Available: <https://doi.org/10.1145/2885497>
- [141] A. Gupta, P. Faraboschi, F. Gioachin, L. V. Kale, R. Kaufmann, B. Lee, V. March, D. Milojcic, and C. H. Suen, “Evaluating and improving the performance and scheduling of hpc applications in cloud,” *IEEE Transactions on Cloud Computing*, vol. 4, no. 3, pp. 307–321, 2016.
- [142] S. Sehrish, G. Mackey, P. Shang, J. Wang, and J. Bent, “Supporting hpc analytics applications with access patterns using data restructuring and data-centric scheduling techniques in mapreduce,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 24, no. 1, pp. 158–169, 2013.
- [143] L. S. Blackford, A. Petitet, R. Pozo, K. Remington, R. C. Whaley, J. Demmel, J. Dongarra, I. Duff, S. Hammarling, G. Henry *et al.*, “An updated set of basic linear algebra subprograms (blas),” *ACM Transactions on Mathematical Software*, vol. 28, no. 2, pp. 135–151, 2002.
- [144] M. P. Forum, “Mpi: A message-passing interface standard,” University of Tennessee, USA, Tech. Rep. CS-94-230, 1994.
- [145] A. Krizhevsky, “One weird trick for parallelizing convolutional neural networks,” *CoRR*, vol. abs/1404.5997, 2014.
- [146] T. Ben-Nun and T. Hoefer, “Demystifying parallel and distributed deep learning: An in-depth concurrency analysis,” *CoRR*, vol. abs/1802.09941, 2018.
- [147] E. Frank, M. A. Hall, G. Holmes, R. Kirkby, B. Pfahringer, and I. H. Witten, *Weka: A machine learning workbench for data mining*. Berlin: Springer, 2005, pp. 1305–1314. [Online]. Available: <http://researchcommons.waikato.ac.nz/handle/10289/1497>
- [148] T. Guo, “Cloud-based or on-device: An empirical study of mobile deep inference,” 2018.
- [149] A. Koliopoulos, P. Yiapanis, F. Tekiner, G. Nenadic, and J. Keane, “A parallel distributed weka framework for big data mining using spark,” in *2015 IEEE International Congress on Big Data*, 2015, pp. 9–16.
- [150] M. Assefi, E. Behraves, G. Liu, and A. P. Tafti, “Big data machine learning using apache spark mllib,” in *2017 IEEE International Conference on Big Data (Big Data)*, 2017, pp. 3492–3498.
- [151] Y. You, I. Gitman, and B. Ginsburg, “Large batch training of convolutional networks,” 2017, cite arxiv:1708.03888. [Online]. Available: <http://arxiv.org/abs/1708.03888>
- [152] S. L. Smith, P.-J. Kindermans, and Q. V. Le, “Don’t decay the learning rate, increase the batch size,” in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=B1Yy1BxCZ>
- [153] J. Dowling, “Distributed deep learning with apache spark and tensorflow,” in *Europe Spark+AI Summit*, 2018, pp. 0–48.
- [154] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [155] F. Niu, B. Recht, C. Ré, and S. Wright, “Hogwild!: A lock-free approach to parallelizing stochastic gradient descent,” *NIPS*, vol. 24, 06 2011.
- [156] J. J. Dai, Y. Wang, X. Qiu, D. Ding, Y. Zhang, Y. Wang, X. Jia, L. C. Zhang, Y. Wan, Z. Li, J. Wang, S. Huang, Z. Wu, Y. Wang, Y. Yang, B. She, D. Shi, Q. Lu, K. Huang, and G. Song, “Bigdl: A distributed deep learning framework for big data,” in *Proceedings of the ACM Symposium on Cloud Computing*, ser. SoCC’19. Association for Computing Machinery, 2019, pp. 50–60. [Online]. Available: <https://arxiv.org/pdf/1804.05839.pdf>
- [157] M. Li, D. G. Andersen, J. W. Park, A. J. Smola, A. Ahmed, V. Josifovski, J. Long, E. J. Shekita, and B.-Y. Su, “Scaling distributed machine learning with the parameter server,” in *Proceedings of the 11th USENIX Conference on Operating Systems Design and Implementation*, ser. OSDI’14. USA: USENIX Association, 2014, p. 583–598.
- [158] N. M. Nawi, M. R. Ransing, and R. S. Ransing, “An improved learning algorithm based on the broyden-fletcher-goldfarb-shanno (bfgs) method for back propagation neural networks,” in *Sixth International Conference on Intelligent Systems Design and Applications*, vol. 1, 2006, pp. 152–157.
- [159] R. O. Green, M. L. Eastwood, C. M. Sarture, T. G. Chrien, M. Aronsson, B. J. Chippendale, J. A. Faust, B. E. Pavri, C. J. Chovit, M. Solis, M. R. Olah, and O. Williams, “Imaging Spectroscopy and the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS),” *Remote Sensing of Environment*, vol. 65, no. 3, pp. 227–248, 9 1998. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0034425798000649>
- [160] S. Moreno-Álvarez, J. M. Haut, M. E. Paoletti, J. A. Rico-Gallego, J. C. Díaz-Martín, and J. Plaza, “Training deep neural networks: a static load balancing approach,” *The Journal of Supercomputing*, vol. 76, no. 12, pp. 9739–9754, 2020. [Online]. Available: <https://doi.org/10.1007/s11227-020-03200-6>



Juan M. Haut (S'17-M'19-SM'20) is an Associated Professor with the Department of Communication and Control Systems at the National Distance Education University, Madrid, Spain. Also, he was a member of the Hyperspectral Computing Laboratory (HyperComp) at the Department of Technology of Computers and Communications, University of Extremadura, where he received the B.Sc and M.Sc. degrees in computer engineering in 2011 and 2014, respectively, and the Ph.D. degree in Information Technology in 2019 supported by an University

Teacher Training Programme from the Spanish Ministry of Education.

Dr. Haut was a recipient of the Outstanding Ph.D. Award at the University of Extremadura in 2019. His research interests include remote sensing data processing and high dimensional data analysis, applying machine (deep) learning and cloud computing approaches. In this sense, he has authored/co-authored more than 30 JCR journal articles (more than 20 in IEEE journals) and 20 peer-reviewed conference proceeding papers. Some of his contributions have been recognized as hot-topic publications for their impact on the scientific community.

Also, he was a recipient of the Outstanding Paper Award in the 2019 IEEE WHISPERS conference. From his experience as a reviewer, it is worth mentioning his active collaboration in more than 10 scientific journals, such as the IEEE Transactions on Geoscience and Remote Sensing, IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, and IEEE Geoscience and Remote Sensing Letters, being awarded with the Best Reviewer recognition of the IEEE Geoscience and Remote Sensing Letters in 2018. Furthermore, he has guest-edited three special issues on hyperspectral remote sensing for different journals. He is also an Associate Editor of the IEEE Geoscience and Remote Sensing Letters and IEEE Journal on Miniaturization for Air and Space Systems.



Javier Plaza (M'09-SM'15) is a member of the Hyperspectral Computing Laboratory at the Department of Technology of Computers and Communications, University of Extremadura, where he received the M.Sc. degree in 2004 and the PhD degree in 2008, both in Computer Engineering. He was the recipient of the Outstanding Ph.D. Dissertation Award at the University of Extremadura in 2008. His main research interests comprise hyperspectral data processing and parallel computing of remote sensing data. He has authored more than 200 publications, including

over 80 JCR journal papers, 10 book chapters, and 100 peer-reviewed conference proceeding papers. He has guest edited 4 special issues on hyperspectral remote sensing for different journals. Prof. Plaza is an Associate Editor for IEEE Geoscience and Remote Sensing Letters and an Associate Editor of the IEEE Remote Sensing Code Library. He is a recipient of the Best Column Award of the IEEE Signal Processing Magazine in 2015 and the most highly cited paper (2005-2010) in the Journal of Parallel and Distributed Computing. He received best paper awards at the IEEE International Conference on Space Technology and the IEEE Symposium on Signal Processing and Information Technology. <http://www.umbc.edu/rssipl/people/jplaza>



Mercedes E. Paoletti (S'17-M'20-SM'20) received the B.Sc and M.Sc. degrees in computer engineering from the University of Extremadura, Cáceres, Spain, in 2014 and 2016, respectively. Also, she obtained her PhD degree in 2020 supported by an University Teacher Training Programme from the Spanish Ministry of Education, as a member of the Hyperspectral Computing Laboratory (HyperComp) at the Department of Technology of Computers and Communications, University of Extremadura. She is currently a researcher at the Department of

Computer Architecture, University of Málaga. Her research interests include remote sensing and analysis of very high spectral resolution with the current focus on DL and high performance computing.

She has served as a reviewer for the IEEE Transactions on Geoscience and Remote Sensing and IEEE Geoscience and Remote Sensing Letters, in which she was recognized as a best reviewer in 2019. She was also a recipient of the 2019 Outstanding Paper Award recognition in the IEEE WHISPERS 2019 conference, and a recipient of the Outstanding Ph.D. Award at the University of Extremadura in 2020.



Sergio Moreno Alvarez completed the B.Sc and M.sc. Degree in Computer Engineering at the University of Extremadura in 2017 and 2019, respectively. Currently, since October 2019 he is completing his Ph.D in the Department of Computer Systems Engineering and Telematics. As research experience, he has participated regional projects. His main interests are high performance computing, neural networks and DL. He is currently a Researcher in the School of Technology (University of Extremadura). He has published 4 JCR papers

in international journals and 3 presentations at international and national conferences.



Juan-Antonio Rico-Gallego received the Computer Science Engineering degree and the PhD degree on Computer Science from the University of Extremadura in 2002 and 2016 respectively. Formerly a software consultant, he is an associate professor at the Dept. of Computer Systems Engineering of the University of Extremadura (Spain). His research interests are in analytical performance models on heterogeneous platforms and the usage of deep and reinforcement learning techniques to HPC platforms problems, including scheduling, process deployment

and mapping, load balancing and communication modeling. Other research interest include current MPI implementations and applications. He codevelops AzequiaMPI, an efficient thread-based full MPI 1.3 standard implementation.



Antonio Plaza (M'05-SM'07-F'15) is the Head of the Hyperspectral Computing Laboratory at the Department of Technology of Computers and Communications, University of Extremadura, where he received the M.Sc. degree in 1999 and the PhD degree in 2002, both in Computer Engineering. His main research interests comprise hyperspectral data processing and parallel computing of remote sensing data. He has authored more than 600 publications, including over 300 JCR journal papers (over 220 in IEEE journals), 23 book chapters, and around

300 peer-reviewed conference proceeding papers. He has guest edited 10 special issues on hyperspectral remote sensing for different journals. Prof. Plaza is a Fellow of IEEE "for contributions to hyperspectral data processing and parallel computing of Earth observation data." He is a recipient of the recognition of Best Reviewers of the IEEE Geoscience and Remote Sensing Letters (in 2009) and a recipient of the recognition of Best Reviewers of the IEEE Transactions on Geoscience and Remote Sensing (in 2010), for which he served as Associate Editor in 2007-2012. He is also an Associate Editor for IEEE Access (receiving a recognition as outstanding Associate Editor of the journal in 2017). He was also a member of the steering committee of the IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing (JSTARS). He is a recipient of the Best Column Award of the IEEE Signal Processing Magazine in 2015, the 2013 Best Paper Award of the JSTARS journal, and the most highly cited paper (2005-2010) in the Journal of Parallel and Distributed Computing. He received best paper awards at the IEEE International Conference on Space Technology and the IEEE Symposium on Signal Processing and Information Technology. He served as the Director of Education Activities for the IEEE Geoscience and Remote Sensing Society (GRSS) in 2011-2012, and as President of the Spanish Chapter of IEEE GRSS in 2012-2016. He has reviewed more than 500 manuscripts for over 50 different journals. He served as Editor-in-Chief of the IEEE Transactions on Geoscience and Remote Sensing from 2013 to 2017, and is currently the Editor-in-Chief of the IEEE Journal on Miniaturization for Air and Space Systems (J-MASS). He has been distinguished as a Highly Cited researcher by Clarivate Analytics in 2018, 2019 and 2020. Additional information: <http://www.umbc.edu/rssi/pl/people/aplaza>