IEEE Access
Multidisciplinary ⁝ Rapid Review ⁝ Open Access Journal

Received 20 February 2025, accepted 15 March 2025, date of publication 20 March 2025, date of current version 1 April 2025.

Digital Object Identifier 10.1109/ACCESS.2025.3553422

APPLIED RESEARCH

Cloud Detection Challenge-Methods and Results

ALESSIO BARBARO CHISARI [1], (Student Member, IEEE), LUCA GUARNERA [1], (Member, IEEE), ALESSANDRO ORTIS [1], (Senior Member, IEEE), WLADIMIRO CARLO PATATU [2], BRUNO CASELLA [3], LUCA NASO [4], GIUSEPPE PUGLISI [4], VINCENZO DEL ZOPPO [4], MARIO VALERIO GIUFFRIDA [5], AND SEBASTIANO BATTIATO [1], (Senior Member, IEEE)

[1]Image Processing Laboratory, Department of Mathematics and Computer Science, University of Catania, 95125 Catania, Italy
[2]Research and Development Business Unit, EHT S.C.p.A., 95129 Catania, Italy
[3]Alpha Research Group, University of Turin, 10124 Turin, Italy
[4]Koexai SRL, 95100 Catania, Italy
[5]School of Computer Science, University of Nottingham, NG7 2RD Nottingham, U.K.

Corresponding author: Luca Guarnera (luca.guarnera@unict.it)

This work was supported in part by Italian Ministry of the University and Research (MUR–PON) under Project E69J21011470006.

ABSTRACT Accurate cloud detection is critical for advancing atmospheric monitoring and meteorological forecasting. This paper presents the Cloud Detection Challenge, an initiative aimed at enhancing cloud detection through innovative solutions using lidar-based ceilometer data. This initiative was hosted by IEEE MetroXRAINE 2024, and 11 teams participated in this initiative. Participants were provided with a novel dataset of backscatter profiles converted into time-height plots, offering unique insights into atmospheric conditions beyond conventional imagery. Data collection employed a Lufft CHM 15k ceilometer, capturing cloud dynamics every 15 seconds located near Mt. Etna, an active volcano in Italy. The dataset includes 1568 hourly labeled backscatter profiles, serving as a benchmark for state-of-the-art deep learning models. The challenge sets a baseline performance of 89.57% accuracy, 92.73% F1-score, 89.82% precision, and 95.84% recall, inviting participants to develop models to exceed these results. Submissions proposed a wide-range of AI-based approaches, including Transformer and Convolutional Neural Network architectures, showcasing the potential of advanced image analysis techniques in lidar-based cloud detection. This paper details the challenge framework, as well as the methodologies proposed by top-performing teams, offering a comparative evaluation of their effectiveness. Our initiative advances cloud detection technologies and underscores their broader implications for environmental monitoring, agriculture, and satellite imaging. The insights and dataset presented herein lay the groundwork for future advancements in leveraging lidar data for atmospheric analysis.

INDEX TERMS Binary classification, ceilometers, cloud detection challenge, computer vision, deep learning, LIDAR.

I. INTRODUCTION

Climate and weather patterns significantly impact societies, influencing the economy and the public safety and well-being. Extreme weather events can disrupt critical infrastructures, affect transportation networks, and, in severe cases, lead to catastrophic consequences, endangering lives and property. In this context, technological tools capable of accurately detecting and forecasting atmospheric phenomena are paramount. Among these tools, lidar-based ceilometers stand out for their ability to provide precise, high-frequency vertical measurements of cloud characteristics, making them invaluable for localized atmospheric monitoring.

A ceilometer is a meteorological device used to measure cloud base height by emitting a light beam and analyzing its reflection [1], [2]. It can also assess aerosol concentration through backscatter analysis. Different ceilometer types exist, including lidar-based systems. As demonstrated in [3], lidar-based systems offer a significant advantage over traditional remote sensing methods, such as satellite photogrammetry,

The associate editor coordinating the review of this manuscript and approving it for publication was Yiqi Liu.

in different domains, when greater precision in detection and analysis is required. Lidar enables more accurate measurements of cloud height, density, and structure, in our context. In contrast, traditional remote sensing techniques often struggle with depth perception and may introduce errors in estimating cloud boundaries and vertical distributions. While satellite-based methods are useful for large-scale atmospheric monitoring, lidar remains the superior choice for advanced meteorological studies, aviation safety, and climate research where high-resolution cloud profiling is essential.

The data collected by ceilometers is typically converted into time-height plots of backscatter coefficients (commonly referred to as backscatter profiles [4]), where the x-axis represents time and the y-axis represents altitude [5]. These backscatter profiles can be analyzed to detect the presence or absence of clouds. Past studies in the literature have explored the use of deep networks [6] for analyzing lidar imagery [7]. However, the lack of publicly available datasets hinders research in the field. This prompted us to collect and release a novel dataset along with benchmark results on state-of-the-art architectures.

The *Cloud Detection Challenge*, hosted by IEEE MetroXRAINE 2024, aimed to advance cloud detection leveraging lidar-based ceilometer data. The challenge introduced a novel dataset of backscatter profiles transformed into time-height plots, inviting participants to develop state-of-the-art binary classification models for cloud detection. Unlike conventional imagery, these profiles provided unique insights into atmospheric conditions, capturing variations in cloud presence over time. Participants were tasked with surpassing our benchmark using a deep network to outperform our results in terms of accuracy, precision, recall, and F1 score [6].

For data acquisition, we employed a Lufft CHM 15k lidar-based ceilometer, which collected measurements every 15 seconds. This device is designed to determine cloud heights, penetration depths, coverage, vertical visibility, and aerosol layers. It was positioned near San Giovanni La Punta (Catania, Italy) at coordinates [37° 34' 43.997" N, 15° 6' 11.002" E]. Data acquisition spanned from January 1, 2023, to March 15, 2023. Ground-truth labels were generated using a high-resolution Weather Research and Forecasting (WRF) model specific to the ceilometer's location. The resulting dataset poses a challenging benchmark for analyzing backscatter profiles and evaluating state-of-the-art architectures for cloud detection.

Several teams from around the world participated in the challenge, presenting solutions that leveraged advanced techniques to address the dataset's unique challenges. The competition served as a platform for testing innovative methods, fostering collaboration, and advancing the field of atmospheric data analysis. This paper details the competition's key aspects, highlights the diverse approaches adopted by participating teams, and evaluates the performance of our proposed solution in comparison to others.

The paper is structured as follows: Section II shows some state-of-the-art methods in the field. Section III provides a comprehensive overview of the cloud detection challenge. Section IV details the employed dataset. Section V describes the proposed baseline for the Cloud Detection Challenge. Section VI illustrates the solutions proposed by the participants. A detailed analysis and discussion of the results are presented in the Section VII. Finally, the Section VIII concludes the paper and highlights directions for future work.

II. RELATED WORKS

This section is divided into two subsections, each addressing a different aspect of atmospheric detection. The first subsection examines various lidar-based approaches for cloud detection, outlining the key methodologies and technologies used in the literature. The second subsection focuses on advancements in ceilometer-based atmospheric monitoring, exploring data analysis techniques and models designed to interpret backscatter profiles.

A. COMPARISON OF LIDAR-BASED CLOUD DETECTION APPROACHES

Several approaches have been proposed for detecting clouds using lidar-based techniques, with significant differences in the type of instruments employed and the atmospheric layers they probe. While our instrument leverages lidar-based ceilometer systems to capture backscatter data from the atmospheric boundary layer, other studies have employed alternative lidar configurations with distinct operational principles and observational capabilities. The approach in [8] incorporates a superconducting nanowire single-photon detector (SNSPD) in high-sensitivity atmospheric lidar, enhancing detection of faint backscatter signals from high-altitude clouds with minimal noise. However, these systems primarily target the mid-to-upper atmosphere and demand sophisticated calibration to mitigate signal attenuation. In contrast, hybrid radar-lidar methods, such as [9], integrate millimeter-wave cloud radar with ground-based multi-wavelength lidar. While radar provides superior penetration depth, lidar ensures finer resolution at lower altitudes. However, these composite systems face challenges in distinguishing drizzle from cloud droplets and require co-located instrumentation, limiting their spatial flexibility. Airborne laser scanning (ALS) lidar, as demonstrated in [10], employs aircraft-mounted pulsed lasers in the near-infrared spectrum to generate high-resolution 3D cloud reconstructions. Despite its ability to capture detailed cloud morphology, ALS is constrained by operational costs, limited temporal resolution, and dependency on flight schedules. Our tool differs from these instruments by focusing on earth-to-satellite observations using lidar-based ceilometer. Ceilometers operate at a single wavelength to continuously monitor the lower atmosphere. Their primary advantage lies in their ability to provide near real-time, high-frequency sampling of the planetary boundary layer (PBL), the atmospheric region most directly influenced by human activity and surface-level meteorological processes. Unlike high-power research lidar systems, ceilometers are optimised for detecting cloud base heights and aerosol layers at altitudes ranging from a few tens of metres to several kilometres. This capability makes them particularly relevant for studying cloud formation and dynamics in urban and industrial environments, where localised emissions and thermal effects strongly modulate atmospheric composition.

B. APPROACHES IN CEILOMETER-BASED ATMOSPHERIC MONITORING

Backscatter profiles acquired by ceilometers have been shown to be highly correlated in the presence of atmospheric particulate matter, as demonstrated in previous studies [11], [12]. In fact, the efficacy of ceilometer data in detecting volcanic emissions has been notably exemplified by its successful application during the 2010 eruption of the Icelandic volcano Eyjafjallajökull [13].

Given its potential, the exploration of data analysis techniques for the analysis of ceilometer-acquired data has been a subject of discussion since the beginning of the Data Mining Project [14]. In fact, a milestone for the research community emerged from the work of Wiegner et al. [15], wherein an approach to calibrate measurements from a Jenoptik CHM 15kx ceilometer was presented. Later, Arun et al. [16] delved into the synergy between ground-based ceilometer observations and satellite data from remote sensing sources in their study. By combining these modalities, they aimed to enhance the precision of cloud detection, highlighting the evolving landscape of data fusion for atmospheric analysis.

In [17], the authors proposed a technique for detecting specific meteorological phenomena, such as fog and clouds, using a lidar-based ceilometer. The methodology involved the application of classical machine learning methods, including Support Vector Machines (SVM), as well as shallow neural networks. These techniques used raw data obtained from the ceilometer as predictive features, enabling the accurate identification of atmospheric events. Similarly, in [18], the authors undertook cloud classification by taking advantage of both ceilometer data with sky images captured by a camera. Within their study, a random forest approach was employed to perform multi-class classification, effectively discerning various cloud types. This integration of data sources facilitated comprehensive cloud identification. Sleeman et al. [19] used lidar-based ceilometer data to detect the Planetary Bounday Layer Height (PBLH) with the use of machine learning techniques. In [4], they introduced an unsupervised methodology for classifying meteorological occurrences, leveraging k-means clustering. An autoencoder was trained to learn a suitable representation of backscatter profiles, subsequently organized into clusters. While demonstrating promise, this technique was presented as a proof-of-concept. Notably, the absence of labeled data and a comprehensive evaluation hampered its full validation. Conversely, the study in [7] addressed cloud detection through Fully Convolutional Networks. In their approach, backscatter profiles were provided into their model via a *mask algorithm*, and the model was trained in a supervised fashion, as they labeled a dataset of backscatter profiles. This dataset enabled an in-depth quantitative performance analysis of their proposed methodology, setting it apart from prior works in the literature. For further reading, a comprehensive review of cloud detection, including the use of ceilometer data, can be found in [20].

The current literature introduces a variety of approaches tackling several tasks based on ceilometer data using distinct datasets. Although previous works have adopted classic machine learning approaches, the nature of the data collected in this paper (*i.e.*, backscatter profiles as shown in Figure 1) are better suited to be analysed with deep neural networks. We want to highlight the difference between our dataset and those available in the literature. Traditional remote sensing systems using lidar technology are predominantly satellite-based and acquire data from a top-down (satellite-to-earth) perspective. Although these approaches have been extensively studied and provide valuable atmospheric and surface observations, they differ fundamentally from our methodology, which adopts a bottom-up (Earth-to-satellite) perspective. This inversion in the data acquisition paradigm introduces a new and exciting aspect to cloud detection, as the dataset we propose captures atmospheric dynamics from a perspective rarely explored in the literature. Because of this unique perspective, our dataset differs significantly from existing datasets, making it new and ahead of its time. The distinct nature of these data presents new challenges and opportunities for the research community, motivating us to propose a dedicated scientific challenge. Then, we aim to promote advances in cloud detection methodologies and improve the understanding of atmospheric phenomena observed by ground-based lidar systems.

In this paper, we propose the results of Cloud Detection Challenge, in which teams from all over the world have the opportunity to get in touch with a new dataset of ceilometer backscatter profiles acquired in proximity of an active volcano. Particularly, this document presents the solution of the best teams participating in the challenge.

III. CLOUD DETECTION CHALLENGE

In this competition, participants are invited to tackle the task of binary classification focused on cloud detection using a novel dataset that was exclusively released for this challenge. The dataset was built by converting raw data collected by a lidar-based ceilometer into images. Images represent backscatter profile labeled with binary annotations indicating the presence or absence of clouds.

The goal of the proposed challenge[1] was to encourage teams around the world to develop innovative solutions

[1] https://iplab.dmi.unict.it/cloud-detection-challenge/

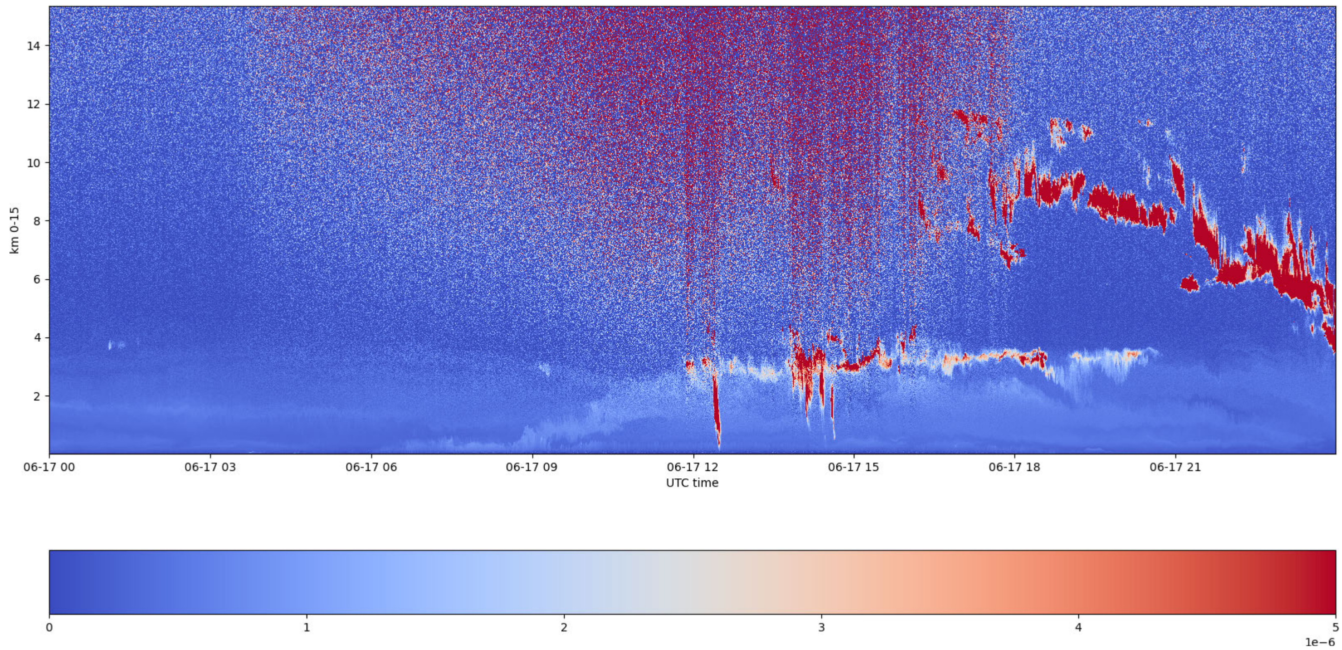

FIGURE 1. **Backscatter profile of 24-hour measurements taken on the 17th of June 2022. As explained in IV, the colour of the plot depends on the intensity of the measured particle: intense blue means absence of particulate; red means intense presence of particulate. Best viewed in colour.**

outperforming our baseline [21] performance of this new dataset. The performance of the models was evaluated using the following metrics: (i) accuracy; (ii) F1 score; (iii) precision; (iv) recall. The winner has been determined by the team with the best value across all metrics. In the event of a tie, the values of precision and recall had higher priority.

A. SIGNIFICANCE OF THE CHALLENGE

The Cloud Detection Challenge holds paramount importance in computer vision applications and image analysis. The ability to accurately identify the presence of clouds in satellite imagery or landscape photographs carries profound implications across various sectors, including meteorology, environmental monitoring, agriculture, and satellite imaging. By using a new dataset based on new types of data and, therefore, new types of images, there is an opportunity to increase research on this topic by integrating new data sources with those already known. Precise cloud classification is essential for understanding climate changes, predicting weather phenomena, and optimizing agricultural operations. The challenge not only calls for innovation in developing advanced models but also provides an opportunity to make substantial contributions to scientific and technological progress in strategic sectors dependent on image analysis accuracy. By participating in this challenge, researchers and developers can showcase their skills and abilities in computer vision, contributing to creating solutions that push beyond current technological frontiers. The results have a tangible impact on real-world applications, enhancing our understanding of the environment and supporting informed decision-making across various industries.

B. CRITERIA OF JUDGING A SUBMISSION

The evaluation of submissions in the Cloud Detection Challenge was designed to ensure a fair and comprehensive assessment of each proposed solution. Given the complexity of lidar-based backscatter data, the criteria focused on both the accuracy of predictions and the balance between the other performance metrics. This approach aimed to highlight not only the ability of the models to correctly classify cloudy vs. clear skies but also their robustness in handling edge cases and maintaining consistency across varying conditions. The following criteria outline the specific metrics and their role in determining the effectiveness of the submitted models.

1) *Classification Accuracy*: Precision in correctly distinguishing between the two classes is crucial. Accuracy will be used as a starting point to assess the overall model performance.
2) *Precision and Recall*: Precision indicates the proportion of true positive predictions among all positive predictions. Recall measures the proportion of true positive predictions among all actual positive instances. A good balance between precision and recall is desirable, but their importance may vary depending on the context of the problem.
3) *F1 Score*: The F1 score is the harmonic mean of precision and recall. This metric provides a balance between the two and can be particularly useful in cases where minimizing both false positives and false negatives is important.

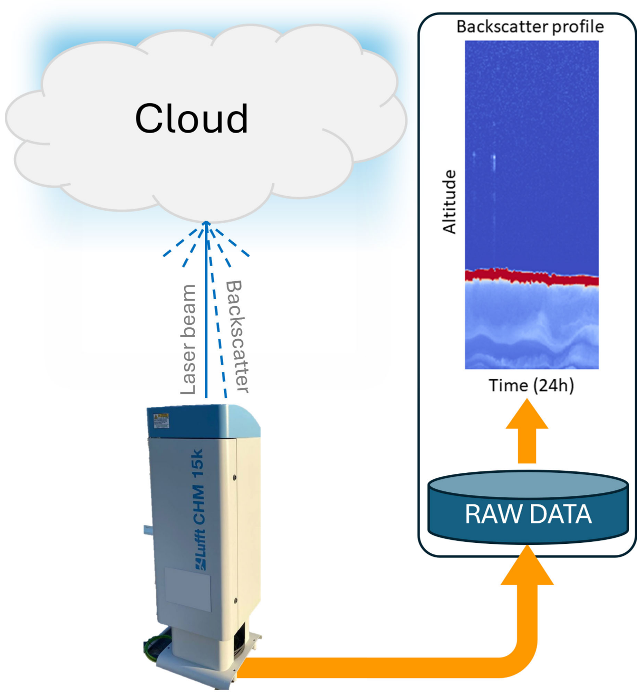

FIGURE 2. Visual representation of the data collection process. Backscatter raw data is utilized to generate a time-height plotting of backscatter coefficients (profile).

IV. DATASET

We used a Lufft CHM 15k ceilometer, a device based on *Light Detection and Ranging* (LiDAR) technology. The system operates by emitting short light pulses generated by a solid-state laser into the atmosphere. These pulses interact with aerosols, water droplets, and air molecules, resulting in scattered light. A portion of this scattered light, known as backscatter, is reflected back to the device and serves as the primary data source for processing. By measuring the time-of-flight of the laser pulses, the device determines the distance to the scattering particles. The reflected signal's altitude profile is examined to derive the raw backscatter intensity, β-raw. From this, the attenuated backscatter coefficient, β-att, is determined using a calibration constant for accuracy. The ceilometer performed measurements at intervals of 15 seconds, offering an accurate assessment of atmospheric particle density. Leveraging signal reflection, it was feasible to identify cloud cover layers. Figure 2 provides a graphical depiction of the data acquisition workflow.

Upon gathering raw data, relevant variables were normalized using a scaling factor specific to the lidar-based ceilometer. These variables were then utilized to create backscatter profiles. In these profiles, the horizontal axis represents time, while the vertical axis denotes particle altitude, based on the backscatter coefficient. The plot's color gradient reflects particle density: deep blue indicates a lack of particulates, whereas red signals high particulate presence. The numerical scale spans from 0 to $5 \cdot 10^{-6}$. As shown in Figure 1, daily backscatter profiles were generated. Each daily profile was then split in 1hr-long non-overlapping windows, producing 24 individual profiles per day. Overall, the dataset comprises 1568 images, each with dimensions of 150×1000, representing hourly measurements.

The generated profiles were annotated with the support of the *Weather Research and Forecasting* (WRF) Model. The WRF model is a high-resolution mesoscale system designed for both research and operational weather forecasting. Its workflow, displayed in Figure 3, includes two dynamic cores, a data processing pipeline, and parallel computing support. The model accommodates meteorological scales from tens of meters to several thousand kilometers. With a spatial resolution of 1×1 km, it outperforms global models like the Global Forecast System (GFS), which typically operates at 27×27 km. The GFS, provided by the *National Center for Atmospheric Research* (NCAR), served as the primary input for the WRF model [22].

The WRF model produces a netCDF file that encodes a 3D spatial grid. Example outputs are shown in Figure 3 (Visual Examples Block). Latitude and longitude are aligned with the x-axis and y-axis, while 40 pressure levels are represented on the z-axis. The shown images differ in spatial resolution and the quality of WRF outputs. The first image has a resolution of 9×9 km, meaning that each point in the spatial grid is 9 km apart from the next. This results in a relatively coarse representation of atmospheric conditions, as finer details of meteorological phenomena are not captured at this scale. The second image, on the other hand, uses a finer grid with a resolution of 3×3 km. This increased level of detail is achieved through the nesting technique within the WRF model. Nesting involves defining one or more higher-resolution grids (called nested domains) within a coarser-resolution grid (parent domain). In this case, the 3×3 km grid is nested within the parent domain of 9×9 km. During this process, the WRF model integrates the meteorological conditions from the parent domain to enhance the representation of atmospheric phenomena on a local scale, thus providing more accurate forecasts in areas of particular interest. Finally, the third image shows data obtained at a resolution of 1×1 km, which is nine times more precise than the initial resolution. At this stage, nesting is further refined by creating a third-level domain, allowing for the capture of very fine meteorological details, such as local variations in temperature, wind, and precipitation. This high-resolution output is particularly useful in contexts where it is necessary to predict small-scale atmospheric phenomena, such as urban micro-climates, intense precipitation events, or wind variations in mountainous regions. Focusing on the ceilometer's location, data for cloud presence at each pressure level was extracted. This enabled the determination of hourly cloud cover above the target area, facilitating the generation of ground-truth labels for each backscatter profile. The proposed dataset was split into train, validation and test, with the following proportions: 49% (769 samples), 21% (329) and 30% (470), respectively, with 1050 true class

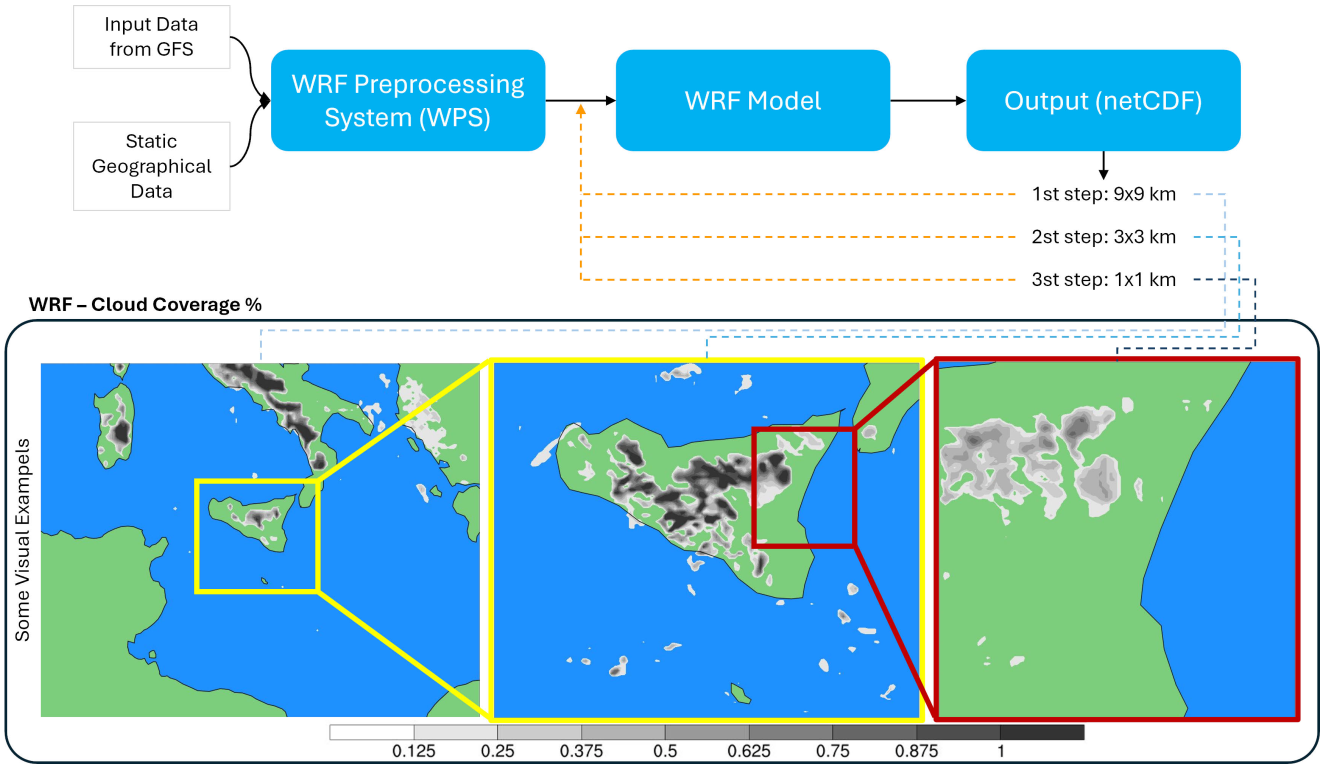

FIGURE 3. Workflow of the WRF model. Outputs from the initial two steps are inputs for the final step. Global data from the GFS forms the first input. Visual examples blocks show some outputs of the WRF model. Details in Section IV.

samples. The participants chose to use other configurations to train the proposed model, partly because the test set was released later.

V. BASELINE

ResNet-50 was chosen as the best architecture in the binary classification task, presented in [21] in which this baseline solution is presented. The dataset is divided into training and validation sets in a 70-30 ratio, and preprocessing steps include resizing, normalization, and data augmentation to enhance model generalization. The ResNet-50 model, pre-trained on ImageNet, is adapted with a custom classification head to suit the binary task.

A. METHODOLOGY

The proposed approach uses a convolutional neural network (CNN) approach for classification, leveraging the ResNet-50 architecture pretrained on ImageNet.

B. DATA PREPROCESSING

Dataset Structure: The dataset is organized into train and validation directories, each containing two subfolders (true and false). The split ratio is 70% - 30%. Transformations:

- Resized to 224 × 224 pixels.
- Resizing crop to 224 pixels.
- Random horizontal flips (probability: 50%) to increase diversity.
- Normalized using the ImageNet mean ([0.485, 0.456, 0.406]) and standard deviation ([0.229, 0.224, 0.225]).

C. MODEL ARCHITECTURE

Base Model: ResNet-50, a 50-layer deep CNN, is utilized as the backbone, as represented in Figure 4. Its pretrained weights on ImageNet facilitate robust feature extraction. Custom Head: The fully connected layers are replaced or extended to suit the binary classification task, ensuring compatibility with the dataset while preserving essential learned features.

D. TRAINING STRATEGY

Device: Training is performed on a GPU-enabled Google Colab Pro instance. Loss Function: CrossEntropyLoss is used to optimize model predictions for binary classification. Optimizer: Stochastic Gradient Descent (SGD) is employed, enabling controlled weight updates via adjustable learning rates. Metrics: Performance is evaluated using accuracy, F1-score, precision, and recall.

E. EXPERIMENTAL SETUP

Hardware: Google Colab Pro provides GPU acceleration, significantly reducing training time and computational overhead. Data Access: The dataset resides in Google Drive, mounted as a root directory within the Colab environment for seamless integration. Hyperparameters: Batch size and

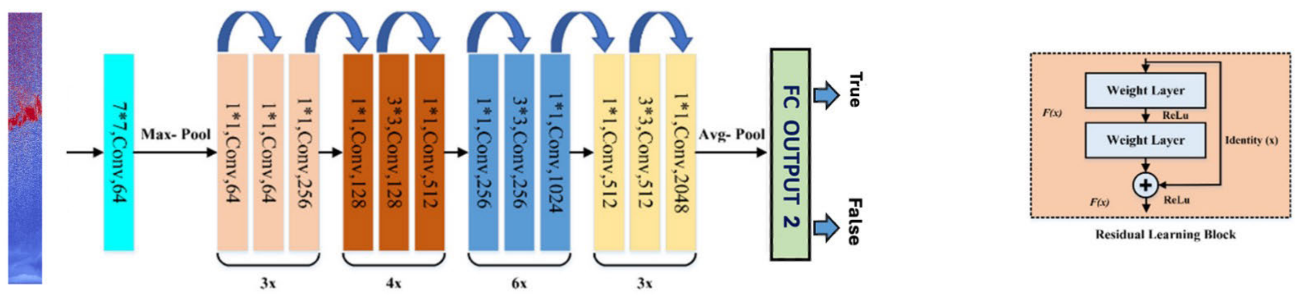

FIGURE 4. Baseline architecture for binary cloud classification: the diagram illustrates a ResNet-50 pre-trained network, employing residual blocks for deep learning of spatial and temporal features from backscatter profiles acquired via ceilometer. The final output is a binary classification, indicating 'True' or 'False' for cloud presence (figure adapted from [23]).

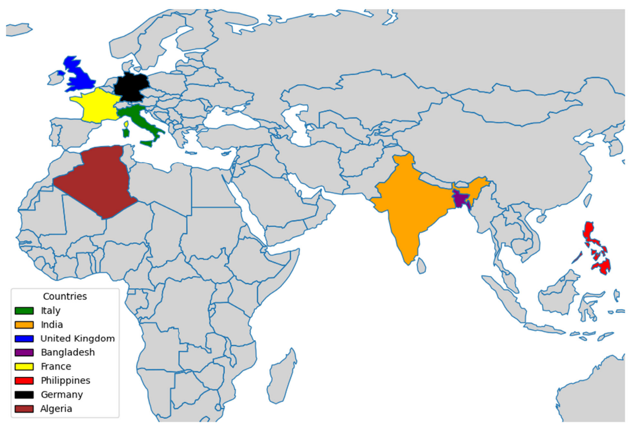

FIGURE 5. Registered participants by affiliation country.

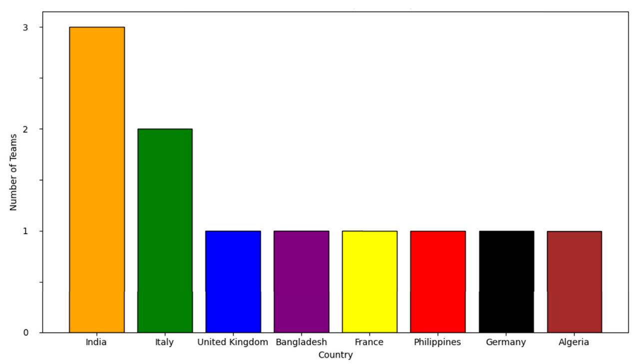

FIGURE 6. Number of participating teams per country.

learning rates are defined but may vary across iterations. The number of epochs and early stopping criteria ensure optimal convergence.

VI. PARTICIPANT SUBMISSIONS TO THE CHALLENGE

A total of 11 teams from all over the world registered for the challenge. In particular, 63.6% of the participants are single researchers, while 36.4% are research groups. Figure 5 shows the countries of origin of the teams registered to the challenge, while Figure 6 shows the number of teams from the respective country. We selected the top two solutions based on (i) the absolute values of the evaluation metrics and (ii) the best-proposed architecture. In the following subsection, we offer details of the solutions proposed by each participants, with a brief presentation, the methodology, and the adopted experimental.

A. ALPHA RESEARCH GROUP-UNVERSITY OF TURIN (UNITO)

Alpha Research Group (represented by Bruno Casella - University of Turin) used a pretrained version of the Vision Transformer (ViT) [24] as shown in Figure 7. The idea behind using a transformer architecture comes from the intrinsic nature of the dataset, as it contains both spatial and temporal features. Taking inspiration from the BEVT paper [25], which proposes a BERT [26] pretraining of Video Transformers and states that for difficult actions, the spatial priors learning should be decoupled from the temporal priors learning, the UniTO researcher hypothesized that temporal features could benefit from spatial features and vice versa.

1) METHODOLOGY

Each image is resized to 384 × 384. Training and validation data are normalized (mean and standard deviation of 50%). The pretrained ViT was trained by minimizing the binary cross-entropy loss with mini-batch gradient descent using the SGD optimizer, with a learning rate of 0.0001, momentum of 0.7, weight decay of 0.000001, and batch size was 12. An early stopping criterion with patience 1 and delta 0 was set for a maximum of 60 training epochs. As augmentation techniques, the participant adopted random horizontal flips, applied with a probability of 50%.

2) EXPERIMENTAL SETUP

The Alpha Research Group used a dedicated server with an Intel Xeon Processor (Skylake, IBRS, 8 sockets of one core) and one Tesla T4 GPU to run the experiments. PyTorch 2.0.1 was adopted as deep learning framework. Each epoch required around 2 minutes. The validation loss,

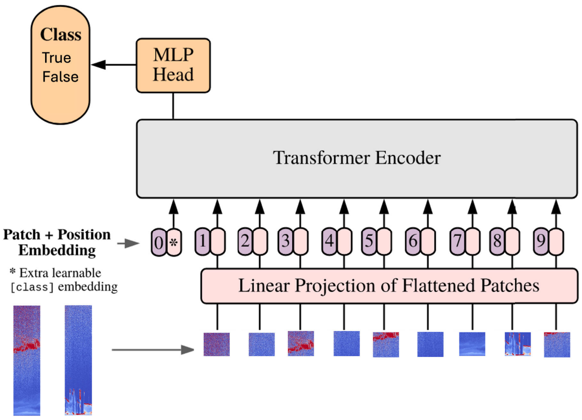

FIGURE 7. Alpha Research Group's proposed architecture.

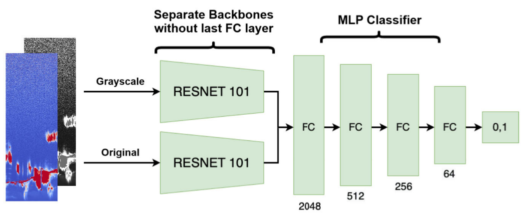

FIGURE 8. Koexai's proposed architecture.

in combination with the early stop criterion, was used as the evaluation metric.

B. KOEXAI (INDUSTRIAL SECTOR)

This section highlights Koexai's contribution to the Cloud Detection Challenge, focusing on the innovative approach employed for the automatic classification of atmospheric conditions using ceilometer data. The challenge aimed to create a robust model for accurately differentiating between cloudy and clear skies based on images created from data captured by these devices. To achieve this, Koexai designed a dual-stream deep learning architecture that processes both RGB and grayscale images to enhance feature extraction and improve classification performance.

1) DATA PREPARATION

To prepare the dataset for training and validation, the original dataset was split into 80% for training and 20% for validation. This division was performed while ensuring a balanced distribution of target classes, utilising the Bhattacharyya distance [27] to minimise biases. Although clouds typically cover only a portion of each image, labels indicating the presence or absence of clouds were provided at the image level. This highlights an opportunity for improving automatic classification by refining the dataset labelling, in addition to enhancing the classification model itself. No domain-specific transformations were applied to the dataset during this process.

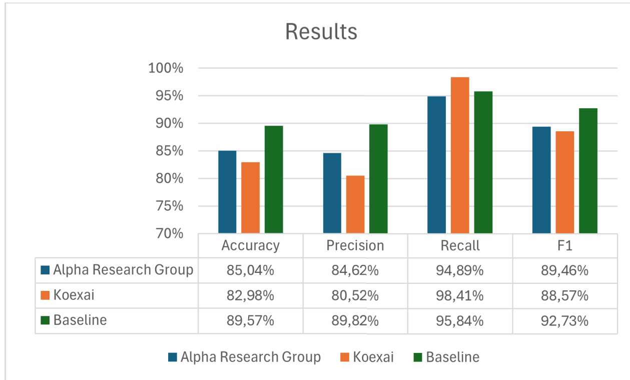

FIGURE 9. Graphical results for each team.

2) MODEL ARCHITECTURE

The proposed model utilises two ResNet-101 [28] backbones operating in parallel: one dedicated to processing RGB images and the other to handling their grayscale counterparts, as shown in Figure 8. This dual architecture allows for a more comprehensive feature extraction from the data, with the RGB stream capturing colour and texture details, while the grayscale stream emphasises structural and contrast-based attributes. The ceilometer images were re-scaled to 224×224 pixels to align with the input format of the ResNet-101 models. The last 1024 feature maps from each backbone were then average pooled and concatenated. These concatenated features were subsequently passed through a multi-layer perceptron (MLP) consisting of three fully connected hidden layers with 512, 256, and 64 neurons, employing LeakyReLU activation functions. The architecture culminates in a final sigmoid activation function for binary classification.

3) TRAINING STRATEGY

The model was implemented using PyTorch and trained for 500 epochs with the Adam optimizer [29], utilising default parameters ($\beta_1 = 0.9$ and $\beta_2 = 0.999$). To enhance training efficiency, several techniques were employed, including an adaptive learning rate, early stopping criteria, gradient clipping, and weight decay set to (1×10^{-5}). The weights of both ResNet-101 networks were initialised with pre-trained weights from ImageNet and fine-tuned separately to adapt them to the specific classification task.

4) RESULTS AND DISCUSSION

Although quantitative metrics such as accuracy and F1-score could not be computed on the test set due to challenge constraints, these metrics were assessed on the validation set. The results demonstrated that the model effectively distinguished between cloudy and clear skies, with minimal signs of overfitting. The dual-stream architecture outperformed single backbone models by capturing a wider range of visual features and leveraging information from both channels to enhance overall model quality. Koexai's innovative approach highlights the potential of dual-stream

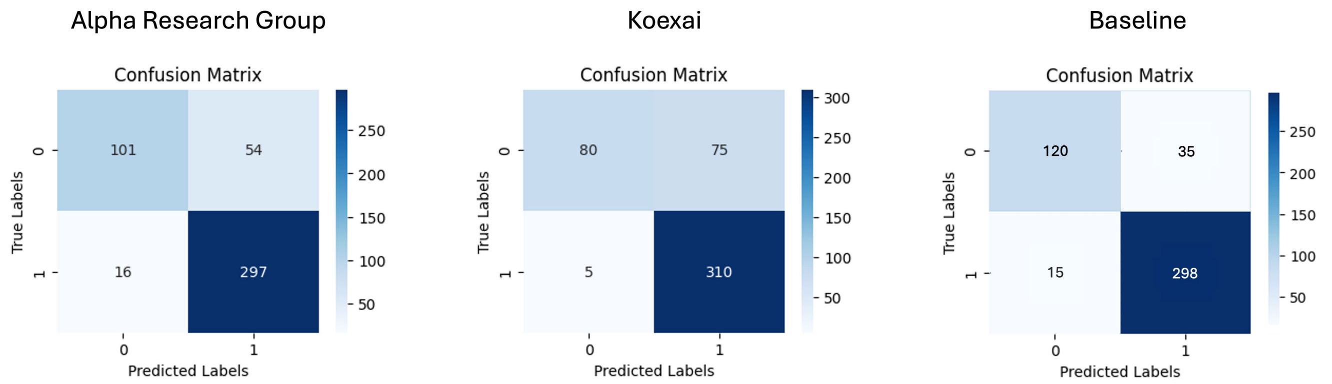

FIGURE 10. Confusion matrix for each team solution: [fist] Alpha Research Group's solution, [second] Koexai's solution, [third] Baseline solution. 0 stands for label "False", 1 for label "True."

TABLE 1. Comparison of the computational needs of each proposed solution.

	Training Time	Epochs	OS	Hardware
Alpha	20m:28s	10	Ubuntu	Intel Xeon CPU Tesla T4 GPU
Koexai	3h:59m:12s	500	Debian 12	Intel Core i5 8th CPU GTX 1080 Ti GPU
Baseline	26m:39s	60	Ubuntu	Intel Xeon CPU Tesla V100 GPU

architectures in improving cloud detection capabilities. The model's robustness suggests that it is well-suited for integration into automated weather monitoring systems and can be further adapted for other meteorological tasks, such as cloud type classification or atmospheric anomaly detection.

VII. RANKING AND DISCUSSION

The Cloud Detection Challenge aimed to foster innovative approaches for analyzing lidar-based ceilometer data, pushing the boundaries of binary cloud detection. The results obtained from participating teams reveal key insights into the efficacy of diverse architectures and methodologies when applied to a novel and complex dataset. This section provides a comprehensive analysis of the results obtained from the challenge. A focus is placed on performance metrics, architectural choices, and the broader implications of the proposed solutions. Figure 9 shows the key values for each metric for each team participating in the competition, while Figure 10 shows all confusion matrices for each proposed solution. Each Figure includes the baseline results.

A. COMPUTATIONAL NEEDS

The computational needs associated with each of the proposed approaches, as delineated in Table 1, can be characterized by several markedly distinct facets. Koexai's approach is associated with a substantially longer training time (3 hours, 59 minutes, and 12 seconds) when juxtaposed with the training durations of the Alpha Research Group and Baseline methods, which are approximately 20 minutes, 28 seconds and 26 minutes, 39 seconds, respectively. This pronounced disparity in training times is primarily attributable to the divergent training strategies adopted by the respective approaches. Specifically, Koexai's proposed method required complete training of the architecture from its initial, uninitialized state. In contrast, both the Alpha Research Group and the Baseline approaches capitalized on the benefits of employing pre-trained architectures. In addition to the differences in overall training duration, a significant discrepancy is evident in the number of training epochs implemented across the approaches. Koexai's proposed approach, which engaged in full training from scratch, was run for 500 epochs, while the Alpha Research Group approach required only 10 epochs and the Baseline approach was limited to 60 epochs. These lower epoch counts are not simply indicative of a truncated training process but rather are the result of the employment of an early stopping criterion, used to avoid overfitting during the training process. The third aspect warranting detailed consideration is the influence of the computational hardware on the overall training process. Koexai used a system configured with an Intel Core i5 8th generation CPU and an NVIDIA GTX 1080 Ti GPU. Both Alpha Research Group and the Baseline approaches were executed on systems that utilized Intel Xeon CPUs in conjunction with NVIDIA GPUs that are particularly well-suited for deep learning applications, such as the Tesla T4 GPU for the Alpha Research Group and the Tesla V100 GPU for the Baseline. Koexai, by virtue of its full-from-scratch training regimen and extended epoch count, naturally incurs a higher computational cost relative to the pre-trained models utilized by the Alpha Research Group and Baseline approaches. Moreover, the disparities in hardware further contribute to the observed variations in training durations.

B. PERFORMANCE SUMMARY

The challenge dataset provided a benchmark for state-of-the-art solutions, with the baseline model achieving 89.57% accuracy, 92.73% F1-score, 89.82% precision, and

95.84% recall. These metrics served as a reference point for evaluating the success of participant submissions.

Among the participants:

- *Koexai Team*: Surpassed the baseline in recall, indicating a strong ability to correctly identify true positive cases (cloudy skies). However, other metrics, including accuracy, F1-score, and precision, remained slightly below the baseline.
- *Alpha Research Group*: Delivered competitive results with a transformer-based approach, but the performance did not exceed the baseline in any key metric.

These outcomes highlight the difficulty of outperforming the robust baseline, which leveraged deep learning to effectively capture the temporal and spatial nuances of the ceilometer data. In terms of the various metrics, we can say that overall *accuracy* was high across all solutions, reflecting their ability to classify samples as cloudy or clear skies effectively. However, no model consistently outperformed the baseline, suggesting that while general predictions were reliable, subtle challenges such as ambiguous atmospheric conditions may have limited further improvements. *Precision* remained strong for both the baseline and participant models, indicating that most predicted cloudy conditions were indeed correct, although some solutions that prioritized recall saw a slight trade-off in precision. *Recall*, on the other hand, was a standout metric for the Koexai team, whose dual-stream CNN architecture effectively leveraged both RGB and grayscale information to enhance sensitivity to cloudy conditions, even in challenging scenarios. Lastly, the *F1-score* highlighted the baseline's strength, as it remained unbeaten, showcasing its ability to balance the correct identification of clouds while minimizing false positives.

The varied results reflect both the strengths and limitations of each approach:

- **Baseline Superiority**: The baseline model's performance underscores the effectiveness of architectures carefully tailored to the dataset's unique characteristics.
- **Koexai's Precision-Recall Tradeoff**: By excelling in recall, the Koexai team demonstrated the importance of architectural innovation for addressing false negatives. However, this came at the cost of reduced precision, suggesting areas for future improvement.
- **Challenges in Architectural Optimization**: The transformer-based solution proposed by Alpha Research Group showed the best results but struggled to outperform the baseline. This indicates the need for further refinement in handling the dataset's temporal and spatial complexities.

While the overall performance was impressive, a few limitations became apparent. For instance, differences in data preprocessing approaches, such as normalization and augmentation, had a noticeable impact on the model outcomes. Fine-tuning these steps could lead to meaningful improvements. Additionally, the complexity of certain architectures, like transformers, introduced risks of overfitting, especially given the relatively small dataset size.

C. DISCUSSION

The analysis of the results achieved by participating teams highlights the complexity of the proposed challenge and the diverse methodologies employed to tackle the binary classification of backscatter images. While all solutions demonstrated promising results, none managed to outperform the baseline metrics, underscoring the effectiveness of the baseline framework as a robust starting point.

Among the participants, the Koexai team stood out by leveraging a dual-stream architecture, which significantly improved recall. This result demonstrated the model's ability to accurately identify cloudy conditions, even in challenging scenarios. However, this improvement came at a slight cost to overall precision, indicating room for further optimization to balance these metrics. Similarly, the transformer-based approach proposed by the Alpha Research Group showcased the potential of attention mechanisms for analyzing complex data, effectively capturing both temporal and spatial features present in the backscatter profiles. Despite these strengths, the difficulty of surpassing the baseline highlights the challenges posed by the dataset's unique characteristics and the inherent complexity of backscatter data. Certain limitations were evident in the proposed solutions, such as the lack of ensemble approaches that could combine the strengths of different models and the absence of advanced strategies to handle class imbalances in the dataset.

The high accuracy and F1-scores observed across most solutions demonstrate the value of backscatter profiles for atmospheric analysis, with potential applications in areas like environmental monitoring and precision agriculture. The ability to distinguish between clear and cloudy conditions with high accuracy reinforces the role of lidar-based data in complementing traditional weather prediction methods.

Looking ahead, future iterations of the challenge could explore multi-class classification to identify specific cloud types or incorporate complementary meteorological data, such as temperature or humidity, to enhance predictive capabilities. Furthermore, leveraging semi-supervised or unsupervised learning techniques could maximize the dataset's utility and further address the challenges of class imbalance.

Overall, the results validate the competition framework as a valuable benchmark for advancing innovation in atmospheric data analysis. At the same time, they provide insightful directions for future improvements and refinements, paving the way for more robust and versatile solutions.

VIII. CONCLUSION AND FUTURE WORKS

The Cloud Detection Challenge has demonstrated the potential of lidar-based ceilometer data for advancing binary cloud detection, emphasizing both the opportunities and challenges inherent in this domain. The outcomes showcase the capability of state-of-the-art deep learning methods to extract meaningful insights from backscatter profiles, which provide unique temporal and spatial details of atmospheric conditions. While the baseline model set a high standard for accuracy, F1-score, precision, and recall, the varied performances of the participant teams highlight the complexity of the task and the room for innovation.

Despite the notable successes, the challenge underscored areas for improvement. Differences in preprocessing strategies, risks of overfitting in complex architectures, and the dataset's inherent characteristics all posed obstacles that prevented any solution from consistently outperforming the baseline. These findings suggest that future work should explore advanced preprocessing techniques, ensemble methods, and strategies to handle class imbalances effectively. Additionally, integrating complementary data sources, such as meteorological or atmospheric parameters, could further enhance model performance and applicability.

Looking forward, the insights gained from this challenge open the door to numerous exciting directions. Expanding the task to multi-class classification, incorporating additional environmental variables, and exploring semi-supervised learning could significantly enhance the versatility and robustness of these models. The advancements achieved through this competition not only contribute to the field of atmospheric monitoring but also provide a foundation for broader applications in environmental analysis and beyond. By building on these results, future efforts can continue to push the boundaries of innovation in cloud detection and related domains.

ACKNOWLEDGMENT

(Alessio Barbaro Chisari and Luca Guarnera contributed equally to this work.)

REFERENCES

[1] L. Tositti, E. Brattich, C. Cassardo, P. Morozzi, A. Bracci, A. Marinoni, S. Di Sabatino, F. Porcù, and A. Zappi, "Development and evolution of an anomalous Asian dust event across Europe in March 2020," *Atmos. Chem. Phys.*, vol. 22, no. 6, pp. 4047–4073, Mar. 2022.

[2] CHM 15k Ceilometer—User manual, Fellbach, Germany, Feb. 7, 2025.

[3] G. D. Pearse, J. P. Dash, H. J. Persson, and M. S. Watt, "Comparison of high-density LiDAR and satellite photogrammetry for forest inventory," *ISPRS J. Photogramm. Remote Sens.*, vol. 142, pp. 257–267, Aug. 2018.

[4] Michael Dammann, Ina Mattis, Michael Neitzke, and Ralf Möller, "Towards the automatic analysis of ceilometer backscattering profiles using unsupervised learning," in *Proc. NeurIPS Workshop Tackling Climate Change Mach. Learn.*, 2022, pp. 1–11.

[5] W. Kong and F. Yi, "Convective boundary layer evolution from LiDAR backscatter and its relationship with surface aerosol concentration at a location of a central China megacity," *J. Geophys. Res., Atmos.*, vol. 120, no. 15, pp. 7928–7940, Aug. 2015.

[6] C. M. Bishop and H. Bishop, *Deep Learning: Foundations and Concepts*. Cham, Switzerland: Springer, 2023.

[7] E. Cromwell and D. Flynn, "LiDAR cloud detection with fully convolutional networks," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Waikoloa, HI, USA, Jan. 2019, pp. 619–627.

[8] D. Salvoni, A. Sannino, L. Parlato, S. Amoruso, G. P. Pepe, A. Boselli, M. Ejrnaes, X. Wang, R. Cristiano, C. Zhang, and L. You, "LiDAR measurement of clouds profile with a superconducting nanowire single photon detector," in *Proc. IEEE 14th Workshop Low Temp. Electron. (WOLTE)*, Apr. 2021, pp. 1–4.

[9] W. Zou, Z. Yin, Y. Dai, Y. Chen, Z. Bu, S. Li, Y. He, X. Hu, D. Müller, T. Lu, X. Dong, and X. Wang, "Robust LiDAR-radar composite cloud boundary detection method with rainfall pixels removal," *IEEE Trans. Geosci. Remote Sens.*, vol. 62, 2024, Art. no. 4111216.

[10] W. Y. Yan, "Airborne LiDAR data artifacts: What we know thus far," *IEEE Geosci. Remote Sens. Mag.*, vol. 11, no. 3, pp. 21–45, Sep. 2023.

[11] C. Muenkel, S. Emeis, W. Mueller, and K. Schaefer, "Aerosol concentration measurements with a LiDAR ceilometer: Results of a one year measuring campaign," *Proc. SPIE*, vol. 5235, pp. 486–496, Feb. 2004.

[12] C. Münkel, N. Eresmaa, J. Räsänen, and A. Karppinen, "Retrieval of mixing height and dust concentration with LiDAR ceilometer," *Boundary-Layer Meteorol.*, vol. 124, no. 1, pp. 117–128, Jul. 2007.

[13] H. Flentje, H. Claude, T. Elste, S. Gilge, U. Köhler, C. Plass-Dülmer, W. Steinbrecht, W. Thomas, A. Werner, and W. Fricke, "The Eyjafjallajökull eruption in April 2010—Detection of volcanic plume using in-situ measurements, ozone sondes and LiDAR-ceilometer profiles," *Atmos. Chem. Phys.*, vol. 10, no. 20, pp. 10085–10092, Oct. 2010.

[14] J. Bartok, O. Habala, P. Bednar, M. Gazak, and L. Hluchý, "Data mining and integration for predicting significant meteorological phenomena," *Proc. Comput. Sci.*, vol. 1, no. 1, pp. 37–46, May 2010.

[15] M. Wiegner and A. Geiß, "Aerosol profiling with the jenoptik ceilometer CHM15kx," *Atmos. Meas. Techn.*, vol. 5, no. 8, pp. 1953–1964, Aug. 2012.

[16] S. H. Arun, S. K. Sharma, S. Chaurasia, R. Vaishnav, and R. Kumar, "Fog/low clouds detection over the Delhi Earth station using the ceilometer and the INSAT-3D/3DR satellite data," *Int. J. Remote Sens.*, vol. 39, no. 12, pp. 4130–4144, Jun. 2018.

[17] Y.-H. Kim, S.-H. Moon, and Y. Yoon, "Detection of precipitation and fog using machine learning on backscatter data from LiDAR ceilometer," *Appl. Sci.*, vol. 10, no. 18, p. 6452, Sep. 2020.

[18] J. Huertas-Tato, F. J. Rodríguez-Benítez, C. Arbizu-Barrena, R. Aler-Mur, I. Galvan-Leon, and D. Pozo-Vázquez, "Automatic cloud-type classification based on the combined use of a sky camera and a ceilometer," *J. Geophys. Res., Atmos.*, vol. 122, no. 20, pp. 11045–11061, Oct. 2017.

[19] J. Sleeman, M. Halem, Z. Yang, V. Caicedo, B. Demoz, and R. Delgado, "A deep machine learning approach for LiDAR based boundary layer height detection," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Waikoloa, HI, USA, Sep. 2020, pp. 3676–3679.

[20] S. Mahajan and B. Fataniya, "Cloud detection methodologies: Variants and development—A review," *Complex Intell. Syst.*, vol. 6, no. 2, pp. 251–261, Jul. 2020.

[21] A. B. Chisari, A. Ortis, L. Guarnera, W. C. Patatu, R. A. Giandolfo, E. Spampinato, S. Battiato, and M. V. Giuffrida, "On the cloud detection from backscattered images generated from a LiDAR-based ceilometer: Current state and opportunities," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Abu Dhabi, UAE, Oct. 2024, pp. 144–150.

[22] *NCEP GFS 0.25 Degree Global Forecast Grids Historical Archive*, Nat. Centers for Environ. Predict., Nat. Weather Service, NOAA, U.S. Dept. Commerce, Boulder, CO, USA, 2015.

[23] L. Ali, F. Alnajjar, H. A. Jassmi, M. Gocho, W. Khan, and M. A. Serhani, "Performance evaluation of deep CNN-based crack detection and localization techniques for concrete structures," *Sensors*, vol. 21, no. 5, p. 1688, Mar. 2021.

[24] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16×16 words: Transformers for image recognition at scale," 2020, *arXiv:2010.11929*.

[25] R. Wang, D. Chen, Z. Wu, Y. Chen, X. Dai, M. Liu, Y.-G. Jiang, L. Zhou, and L. Yuan, "BEVT: BERT pretraining of video transformers," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, New Orleans, LA, USA, Jun. 2022, pp. 14713–14723.

[26] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol., (NAACL-HLT)*, vol. 1, Minneapolis, MN, USA, Jun. 2019, pp. 4171–4186.

[27] A. Bhattacharyya, "On a measure of divergence between two statistical populations defined by their probability distribution," *Bull. Calcutta Math. Soc.*, vol. 35, pp. 99–110, Apr. 1943.

[28] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," 2015, *arXiv:1512.03385*.

[29] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.

ALESSIO BARBARO CHISARI (Student Member, IEEE) received the B.Sc. degree in computer engineering and the M.Sc. degree in management engineering from the University of Catania, in 2019 and 2021, respectively, where he is currently pursuing the Ph.D. degree in computer science. He is a member of the Image Processing Laboratory (IPLab). He was also a Visiting Ph.D. Student with the University of Nottingham, U.K., in 2023. He is also a Technical Business Developer and a Technical Project Manager with the Research and Development Business Unit, EHT S.C.p.A., in which he manages and checks the scientific relevance of research and development projects in the IT field. His current research interests include computer vision and deep learning techniques applied to atmospheric phenomena.

LUCA GUARNERA (Member, IEEE) was born in Catania, in October 1992. He received the Ph.D. degree in computer science, in October 2021, discussing the thesis entitled "Discovering Fingerprints for Deepfake Detection and Multimedia-Enhanced Forensic Investigations" at the Department of Mathematics and Computer Science, University of Catania. Since January 2022, he has been a Research Fellow in computer science with the University of Catania. His Ph.D. research was carried out at the University of Hertfordshire, College Lane Campus, Hatfield, U.K., under the supervision of Prof. Salvatore Livatino, working on the creation of a software for forensic ballistics analysis and firearms comparison, using the Oculus Rift S. Headset. His main research interests include computer vision, machine learning, multimedia forensics, and its related fields with a focus on the Deepfake phenomenon.

ALESSANDRO ORTIS (Senior Member, IEEE) received the Ph.D. degree in mathematics and computer science, with part of his research conducted with Imperial College London. He is currently a Researcher and a Lecturer with the University of Catania, specializing in artificial intelligence for computer vision and security. His work focuses on AI vulnerabilities, such as Deepfakes, adversarial machine learning, and biometrics. He also teaches multimedia security and biometrics with the University of Catania. He is a member of the IEEE Biometric Council, and part of the Image Processing Laboratory (IPLAB) Research Group, which is involved in international collaborations and funded projects.

WLADIMIRO CARLO PATATU is currently an Electronic Engineer with an M.B.A. and extensive experience in driving complex IT innovation projects. He was an EHT Manager with the Research and Development Business Unit, leads the team that provides scientific guidance for innovation projects. With longstanding experience as a Scientific Coordinator for complex research and development funded projects. His research focuses on leveraging innovative IT technologies, models, and methods to drive innovation across multiple sectors, including agritech, life science, smart communities, cultural heritage, smart industry, and energy and environment.

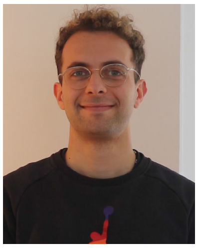

BRUNO CASELLA received the B.Sc. degree in computer engineering, in 2020, and the M.Sc. degree in data science for management from the University of Catania. He is currently pursuing the Ph.D. degree in modeling and data science with the University of Turin. His current research interests include federated learning, transfer learning, and distributed computing.

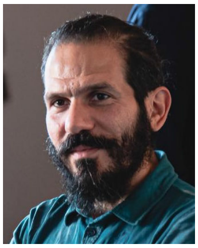

LUCA NASO received the Master of Science degree in physics from the University of Catania, in 2004, the Diploma di Licenza degree from the Class of Experimental Sciences, Scuola Superiore di Catania (SSC), in 2005, and the Ph.D. degree in astrophysics from Scuola Internazionale Superiore di Studi Avanzati (SISSA), in 2009. He was an Astrophysicist, he has conducted research at several prestigious institutions worldwide, including Trieste, Padua, Warsaw, Oxford, and Beijing, focusing on magnetic fields in compact objects. His work has led to multiple publications in peer-reviewed journals, significantly contributing to the field. In 2013, he transitioned from academic research to digital technologies, big data, and machine learning, embarking on a diverse career that includes roles, such as a Project Manager, a Product Manager, a Data Scientist, the Chief Technology Officer (CTO), the Chief Data Officer (CDO), a Startup Mentor, and a University Lecturer. Currently, he is the CEO and a Founder with Koexai SRL, an innovative startup that provides tailored data science consulting services to small and medium enterprises, including research centers, helping them harness the power of data to achieve their strategic goals. In this role, he leads a Team of data scientists, working closely with clients to understand their needs and develop effective data-driven solutions. Additionally, he is the Head of the Data Science Centre of Excellence, EHT S.C.p.A., where he oversees projects that leverage data science to drive business solutions.

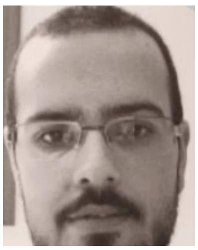

GIUSEPPE PUGLISI received the bachelor's degree in computer science, the master's degree, and the Ph.D. degree in machine learning from the University of Catania, in 2016, 2018, and 2023, respectively. He was an AWS Cloud Practitioner and a Former Member of the IPLab Research Center, Department of Mathematics and Computer Science (DMI), University of Catania, he has contributed to innovative research projects in artificial intelligence. Currently, he is a Data Scientist with Koexai SRL, where he applies his expertise in artificial intelligence and machine learning to develop advanced solutions in data science and computer vision. He is a Computer Vision Specialist with a robust foundation in machine learning techniques, including deep learning, reinforcement learning, and deep reinforcement learning. His work focuses on integrating cutting-edge algorithms and cloud technologies to enhance his offerings, driving innovation, and efficiency for clients.

VINCENZO DEL ZOPPO is a Data Scientist with more than 12 years of experience specializing in computer vision, deep learning, radar, automatic number plate recognition (ANPR), and advanced driver assistance systems (ADAS). Currently, he is the Tech Lead of Koexai SRL, an innovative startup dedicated to providing tailored data science consulting services to small and medium enterprises, including research centers. He received multiple national problem-solving awards during high school, which led to university scholarships to study physics from the Scuola Superiore di Catania, Italy. During his academic career, he was a Key Member of the UniCT Team that achieved second place in the EuroBot World Finals, in 2011. With a diverse background across industries, such as automotive, security, defence, aerospace, healthcare, and finance, he has developed a strong expertise in cutting-edge AI applications. He has contributed to applied research through several published articles and holds a patent in this field. Throughout his career, he has taken on increasing responsibilities, mentoring and leading small research teams while ensuring the successful execution of complex projects for several Fortune 500 companies. In this role, he leads a skilled team of data scientists and engineers to develop advanced, data-driven solutions that address their unique challenges, and strategic goals. His leadership ensures the delivery of high-quality results, fostering a culture of innovation, and excellence within the team.

MARIO VALERIO GIUFFRIDA is currently an Assistant Professor of computer vision with the School of Computer Science, University of Nottingham, and a member of the Computer Vision Laboratory (CVL). He was a Co-Investigator in the PhenomUK scoping project, he contributed to advancing phenotyping technologies in U.K. He has also organized international workshops on computer vision in plant phenotyping at leading international conferences, fostering collaboration, and knowledge exchange in this field. His research interests include deep learning applied to biological problems, including plant and medical imaging analysis.

SEBASTIANO BATTIATO (Senior Member, IEEE) received the degree in computer science from the University of Catania, in 1995, and the Ph.D. degree in computer science and applied mathematics from the University of Naples, in 1999. He has held various positions with the University of Catania, starting as an Assistant Professor, in 2004, then becoming an Associate Professor, in 2011, and finally a Full Professor, in 2016. Currently, he is the Deputy Rector of the Strategic Planning and Information Systems. Additionally, he is the Director and a Co-Founder of the International Computer Vision Summer School (ICVSS) and the International Forensics Summer School (IFOSS). He is highly active in research, directing the IPLab research laboratory and participating in numerous national and international projects. Throughout his career, he has been involved in leading roles within academic programs and research initiatives, including chairing the bachelor's program in computer science and serving as the Rector's delegate for Education. He has supervised numerous Ph.D. students and postdocs, edited several books, and authored more than 100 articles in various publications. His research interests include computer vision, imaging technology, and multimedia forensics. He is recognized for his contributions, receiving accolades, such as the 2017 PAMI Mark Everingham Prize and the 2011 Best Associate Editor Award from IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY. He also plays key roles in organizing international events and serves as an Associate Editor for *Journal of Electronic Imaging* (SPIE).