



Infrared radiometric image classification and segmentation of cloud structures using a deep-learning framework from ground-based infrared thermal camera observations

Kélian Sommer¹, Wassim Kabalan², and Romain Brunet³

¹Laboratoire Univers et Particules de Montpellier, Université de Montpellier, CNRS, Montpellier, France

²Astroparticule et Cosmologie, Université Paris Cité, CNRS, 75013 Paris, France

³Aix-Marseille Université, CNRS, CNES, LAM, Marseille, France

Correspondence: Kélian Sommer (kelian.sommer@umontpellier.fr), Wassim Kabalan (wassim@apc.in2p3.fr), and Romain Brunet (romain.brunet@lam.fr)

Received: 12 January 2024 – Discussion started: 22 March 2024

Revised: 28 July 2024 – Accepted: 29 August 2024 – Published: 12 May 2025

Abstract. Infrared thermal cameras offer reliable means of assessing atmospheric conditions by measuring the downward radiance from the sky, facilitating their usage in cloud monitoring endeavors. The precise identification and detection of clouds in images pose great challenges stemming from the indistinct boundaries inherent to cloud formations. Various methodologies for segmentation have been previously suggested. Most of them rely on color as the distinguishing criterion for cloud identification in the visible spectral domain and thus lack the ability to detect cloud structures in gray-scaled images with satisfying accuracy. In this work, we propose a new complete deep-learning framework to perform image classification and segmentation with convolutional neural networks. We demonstrate the effectiveness of this technique by conducting a series of tests and validations based on self-captured infrared sky images. Our findings reveal that the models can effectively differentiate between image types and accurately capture detailed cloud structure information at the pixel level, even when trained with a single binary ground-truth mask per input sample. The classifier model achieves an excellent accuracy of 99 % in image type distinction, while the segmentation model attains a mean pixel accuracy of 95 % in our dataset. We emphasize that our framework exhibits strong viability and can be used for infrared thermal ground-based cloud monitoring operations over extended durations. We expect to take advantage of this framework for astronomical applications by providing cloud cover selection criteria for ground-based photometric observations within the StarDICE experiment.

1 Introduction

Accurate and continuous monitoring of cloud properties contributes to a profound understanding of atmospheric processes and their subsequent impacts on various Earth systems (Liou, 1992). It provides essential insights for weather predictions and climate dynamics (Hu et al., 2004; Petzold et al., 2015). Observation methods can be divided into two primary distinct categories: downward satellite-based observations (Roy et al., 2017; Martin, 2008) and upward ground-based observations with all-sky cameras, lidar, radar, and other instruments (Wilczak et al., 1996). The principal aim of satellite-based observations is to investigate the upper regions of clouds, facilitating the examination and analysis of global atmospheric patterns and climate conditions over expansive geographical areas (Schiffer and Rossow, 1983; Boers et al., 2006; Geer et al., 2017; Várnai and Marshak, 2018). In contrast, ground-based cloud observation excels in the surveillance of localized regions, furnishing valuable data pertaining to the lower segments of clouds by giving information on cloud altitude, cloud extent, and cloud typology (Bower et al., 2000; Zhou et al., 2019). A combination of these two measurement techniques enhances our overall comprehension of cloud behavior (Mokhov and Schlesinger, 1994; Schreiner et al., 1993; Yamashita and Yoshimura, 2012; Yoshimura and Yamashita, 2013).

Ground-based observations have been extensively used in recent years and have become a viable means to detect, study, and identify cloud formations (Paczyński, 2000,

Skidmore et al., 2008; Tzoumanikas et al., 2016; Ugolnikov et al., 2017; Mommert, 2020; Tzoumanikas et al., 2016; Román et al., 2022). As technological evolution has ushered in a new era of monitoring methodologies (Mandat et al., 2014), the utilization of infrared thermal cameras has emerged as a promising avenue for atmospheric investigations through precise radiometric measurements (Szejwach, 1982; Shaw and Nugent, 2013; Liandrat et al., 2017b; Lopez et al., 2017; Klebe et al., 2014; Nikolenko and Maslov, 2021).

Because of their practical use, high sensitivity, low cost, operating range, and wide field of view (FOV) (Rogalski, 2011; Rogalski and Chrzanowski, 2014; Kimata, 2018), uncooled infrared thermal cameras are particularly useful for medicine (Ring and Ammer, 2012); agriculture (Ishimwe et al., 2014); aerial (Wilczak et al., 1996); defense (Gallo et al., 1993; Akula et al., 2011); surveillance (Wong et al., 2009); weather forecasting (Sun et al., 2008; Liandrat et al., 2017a); or even astronomical applications to determine the cloud cover fraction during operations and, therefore, to assess the quality of scientific observations (Sebag et al., 2010; Lewis et al., 2010; Klebe et al., 2012, 2014; Reil et al., 2014). Indeed, uncooled infrared micro-bolometer array sensors working in the 8–14 μm spectral band can directly detect the longwave infrared (LWIR) thermal emission of both clouds and the atmospheric background, excluding the scattered light of the sun or starlight (Houghton and Lee, 1972). These LWIR sensors are able to provide high-contrast images and allow fine radiometric measurements to detect low-emissivity cirrus clouds (Lewis et al., 2010; Shaw and Nugent, 2013).

Over recent years, multiple automatic ground-based observation systems have been developed. For example, the infrared cloud imager (ICI; see Thuraiajah and Shaw, 2005) can detect clouds and assess cloud coverage both in daylight and during the nighttime with a dedicated infrared sensor. Sharma et al. (2015) designed an instrument to detect the cloud infrared radiation to be used in the search for a potential site for India's National Large Optical Telescope project. The development of the Radiometric All-Sky Infrared Camera (RASICAM, referenced in Lewis et al., 2010, and Reil et al., 2014) was aimed at enabling automated, real-time quantitative evaluation of nighttime sky conditions for the Dark Energy Survey (Dark Energy Survey Collaboration et al., 2016). This particular camera is designed to detect, locate, and analyze the motion and properties of thin, high-altitude cirrus clouds and contrails by measuring their brightness temperature against the sky background. The All-Sky Infrared Visible Analyzer (ASIVA) is a similar instrument whose primary goal is to provide radiometrically calibrated imagery in the LWIR band to estimate fractional sky cover and sky or cloud brightness temperature, emissivity, and cloud height (Klebe et al., 2014). The ASC-200 system (Wang et al., 2021b) combines information from two all-sky

cameras facing the sky operating in both the visible spectrum (450–650 nm) and the LWIR band.

As next-generation cosmological surveys require more demanding precision in terms of photometric observations – implying better characterization of the atmosphere – monitoring telescope instrument FOVs with LWIR thermal cameras may provide significant assets with regard to (i) classifying observation quality in real time, (ii) evaluating potential cloud coverage (Smith and Toumi, 2008; Liandrat et al., 2017b; Aebi et al., 2018; Wang et al., 2021b), and (iii) estimating precipitable water vapor (PWV) content (Kelsey et al., 2022; Hack et al., 2023; Salamalikis et al., 2023).

In this study, we plan to address the first objective. We use an LWIR thermal infrared camera with a specifically chosen narrower FOV that aims to image the surrounding area of the StarDICE telescope FOV. The StarDICE metrology experiment (Betoule et al., 2022) aims to measure CALSPEC (Bohlin, 2014) spectrophotometric standard star absolute fluxes at the 0.1 % relative uncertainty level. Enhanced characterization of atmospheric conditions is required to reach the target sensitivity (Hazenbergh, 2019). As a preliminary step, basic knowledge of the atmospheric conditions in the telescope FOV may provide valuable insights into the quality of spectrophotometric measurements. However, these kinds of infrared instruments operate at high frame rates and produce considerable amounts of data, making them extremely difficult for human observers to analyze. Therefore, to determine cloud presence in infrared images, deep convolutional neural networks (CNNs) appear to be a viable approach to process images in real time. Multiple models relying on CNNs have been developed, such as CloudSegNet (Dev et al., 2019a), CloudU-Net (Shi et al., 2021b), CloudU-Netv2 (Shi et al., 2021a), SegCloud (Xie et al., 2020), TransCloudSeg (Liu et al., 2022), CloudDeepLabV3 (Li et al., 2023), ACLNet (Makwana et al., 2022), DeepCloud (Ye et al., 2017), CloudRAEDnet (Shi et al., 2022), DMNet (Zhao et al., 2022), and DPNet (Zhang et al., 2022). Nonetheless, these methodologies exclusively address RGB-colored images (Li et al., 2011; Dev et al., 2016). Colors or hue provides the essential information for segmentation (especially red and blue channels). In the case of LWIR thermal images, we implement a model capable of achieving comparable accuracy for single-channel gray-scaled images. Inspired by their large successes in image classification and structure detection for various computer vision tasks, we propose a dedicated deep-learning framework. Our approach is specifically designed for gray-scaled infrared images and consists of (i) classifying images (e.g., detecting if any cloud is present in the image and discriminating between clear and cloudy images) and (ii) identifying cloud structure (e.g., generating a pixel-based probabilistic segmentation map and verifying if the CCD camera FOV is impacted).

The remainder of the paper is structured as follows. The background of the scientific context and related works is presented in Sect. 2. Section 3 details the experimental setup and

dataset. Section 4 introduces the proposed framework, describing deep-learning architectures and training procedures. Experimental results and comparisons with other datasets are provided in Sect. 5. Relevant matters and future perspectives are discussed in Sect. 6. Section 7 summarizes the main results and finally concludes the paper.

2 Background

2.1 Motivation

StarDICE represents one of the initiatives focused on creating a measurement process that bridges the gap between laboratory flux standards (such as silicon photodiodes calibrated by NIST; see Larason and Houston, 2008) and the stars found in the CALSPEC library of spectrophotometric references (Bohlin et al., 2020). Since type-Ia supernovae (SNe Ia) and most astronomical surveys rely on the calibration of these standard stars for their measurements (Bohlin et al., 2011; Conley et al., 2011; Rubin et al., 2015; Scolnic et al., 2015; Currie et al., 2020; Brout et al., 2022; Rubin et al., 2022), successfully establishing this connection with high precision effectively addresses the calibration challenge associated with the Hubble diagram for cosmology and the study of dark energy driving the accelerated expansion of the universe (see Goobar and Leibundgut, 2011, for a review of SNe Ia in cosmology).

The StarDICE proposal relies on the near-field calibration of a stable light source (Betoule et al., 2022). It serves as a distant in situ reference for a compact astronomical telescope. One of the largest remaining sources of systematic uncertainty when observing stellar sources from the ground is the Earth's atmospheric transmission (Stubbs and Tonry, 2012; Stubbs and Brown, 2015; Li et al., 2016). It is dependent on many environmental conditions and processes, including absorption and scattering by molecular constituents (O_2 , O_3 , and others), absorption by PWV, scattering by aerosols, and shadowing by larger ice crystals and water droplets in clouds that is independent of wavelength and responsible for gray extinction (Burke et al., 2010, 2017). Current atmospheric transmission or extinction models do not integrate the possible impact of clouds. Indeed, the formation of thin clouds through the condensation of water droplets and ice can result in clouds that are extremely faint and that cannot be perceived in the visible spectrum with the naked eye. These clouds often exhibit complex spatial structures, as demonstrated in Burke et al. (2013). Clouds passing through the photometric instrument's FOV result in an attenuation of stellar flux. Previously, this issue was addressed by incorporating a gray extinction correction, involving the fitting of an empirical normalization parameter for each observation. Nevertheless, as highlighted by Burke et al. (2010), this approach has been proven to be insufficient due to calibration limitations arising from the dynamic and evolving

nature of cloud cover conditions. To tackle this challenge in the StarDICE experiment, our solution involves employing an infrared thermal camera. This specialized equipment offers high-sensitivity radiometric measurements, capturing the sky radiance within the atmosphere's transparency window (10–12 μm). With the help of additional cloud spatial structure identification analytical capabilities, this instrument may be the key to assessing photometric observation quality and label science images with superior state-of-the-art accuracy. The primary initial objective is to generate a catalog of optical exposures from the telescope that are suitable for extracting stellar photometric flux and conducting subsequent analysis.

2.2 Related work

In recent years, numerous cloud sky–cloud segmentation algorithms have been introduced, along with the increased development of all-sky ground-based cloud monitoring stations (Long et al., 2006; Yang et al., 2012; Krauz et al., 2020; Fa et al., 2019; Mommert, 2020; Li et al., 2022). Indeed, cloud segmentation is a big challenge for remote sensing applications as clouds come in various shapes and forms. Most modern approaches aim to use computer vision algorithms and train them based on very specific publicly available cloud image databases such as SWIMSEG (Dev et al., 2016), SWINSEG (Dev et al., 2019b, 2017), SWINySEG (Dev et al., 2019a), WSISEG (Xie et al., 2020), HYTA (Li et al., 2011), and TLCDD (TLCDD, 2022). Many proposed solutions are focused on visible RGB images. CloudSegNet (Dev et al., 2019a) is a lightweight deep-learning encoder–decoder network that detects clouds in daytime and nighttime visible color images. CloudU-Net (Shi et al., 2021b) modifies CloudSegNet architecture by adding dilated convolution, skip connection, and fully connected conditional random field (CRF; see McCallum, 2012) layers to demonstrate better segmentation performance overall. It uses the powerful U-Net architecture (Ronneberger et al., 2015) originally applied to medical image segmentation. CloudU-Netv2 (Shi et al., 2021a) replaces the upsampling in CloudU-Net with bilinear upsampling, improves the discrimination ability of feature representation, and uses the rectified Adam optimizer (RAadm is a variant of the Adam stochastic optimizer (Kingma and Ba, 2014) that introduces a term to rectify the variance of the adaptive learning rate; see Liu et al., 2019). SegCloud (Xie et al., 2020) has been trained based on 400 images and possesses a symmetric encoder–decoder structure and outputs low- and high-level cloud feature maps at the same resolution as input images. TransCloudSeg (Liu et al., 2022) addresses the loss of global information due to the limited receptive field size of the filters in CNNs by proposing a hybrid model containing both the CNN and a transformer (Vaswani et al., 2017) as the encoders to obtain different features. CloudDeepLabV3+ (Li et al., 2023) designs a lightweight ground-based cloud image

adaptive segmentation method that integrates multi-scale feature aggregation and multi-level attention feature enhancement. ACLNet (Makwana et al., 2022) uses EfficientNet-B0 as the backbone, atrous spatial pyramid pooling (ASPP; see Chen et al., 2017) to learn at multiple receptive fields, and a global attention module (GAM; see Liu et al., 2021) to extract fine-grained details from the image. It provides a lower error rate, higher recall, and higher $F1$ score than state-of-the-art cloud segmentation models. DeepCloud (Ye et al., 2017) uses the method of Fisher vector encoding, which is applied to execute the spatial feature aggregation and high-dimensional feature mapping of the raw deep convolutional features. CloudRAEDnet (Shi et al., 2022) proposes a residual attention-based encoder–decoder network and trains it based on the SWINySEG dataset.

The majority of these models are typically structured using an encoder–decoder architecture, which is the primary innovation brought forth by incorporating CNNs (O’Shea and Nash, 2015). The encoder is tailored to acquire representational features, facilitating the extraction of semantic information while the decoder reconstructs these representational features into the segmentation mask, allowing for pixel-level classification (Badrinarayanan et al., 2017; Alzubaidi et al., 2021).

Others have proposed solutions for all-sky infrared image classification. Liu et al. (2011) apply pre-processing steps (smoothing noise reduction, enhancement through top-hat transformation and high-pass filtering, and edge detection) before extracting features that are useful for distinguishing cirriform, cumuliform, and waveform clouds. A simple rectangle method as a supervised classifier is applied. They find a 90 % agreement between a priori classification carried out manually by visual inspection and their algorithm based on 277 images. Sun et al. (2011) suggested (i) a method for determining clear-sky radiance thresholds, (ii) a cloud identification method made up of a combined threshold and texture method, and (iii) an algorithm to retrieve cloud base height from downwelling infrared radiance. They showed that structural features are better than texture features in classifying clouds. Luo et al. (2018) proposed a three-step process: (i) pre-processing, (ii) feature extraction, and (iii) a classification method to group images into five cloud categories (stratiform, cumuliform, waveform, cirriform, and clear) based on manifold and texture features using a support vector machine (SVM; see Cortes and Vapnik, 1995). Their experimental results demonstrate a higher recognition rate, with an increase of 2 %–10 % based on ground-based infrared image datasets. These methods classify clouds into separate categories based on their typology. Until now, all the previously examined approaches, while effective within their specific domains, proved to be unsuccessful when applied to our particular use case. Therefore, we propose a new deep-learning framework based on a linear classifier and U-Net architectures to identify cloud images and to detect cloud structures in real time.

3 Experimental setup and datasets

3.1 Description of the instrument

Our instrument is an infrared thermal camera similar to the Thurairajah and Shaw (2005) device – specifically the FLIR Tau2 – which operates in the LWIR band, covering the 8–14 μm range. It features a focal plane array (FPA) consisting of 640×512 uncooled micro-bolometers, capturing images at a frame rate of 8.33 Hz. The camera is paired with a 60 mm lens at $f\#1.25$ aperture, resulting in a 10.4×8.3 degree² imaging area. The primary purpose of deploying this instrument on the equatorial mount adjacent to the StarDICE photometric telescope is to continuously assess the atmospheric conditions (specifically gray extinction) within the line of sight of the visible CCD camera during observations. The IR instrument FOV is chosen to be larger than the telescope FOV (0.5×0.5 degree²) to anticipate the movement of clouds in the smaller FOV of interest. Through meticulous calibration, radiative transfer calculations, and data analysis using simulations, we can extract valuable information about the sky to monitor real-time atmospheric conditions. In Figs. 1 and 2, we show the instrument mounted on a prototype equatorial mount with the necessary command and control equipment and one sample gray-scale image, respectively. We also monitor the surrounding and internal temperatures of the camera in real time to correct for temperature-related variations in sensor response. The device is controlled and commanded via the ThermalCapture Grabber USB 2.0 interface, which grants access to full 14-bit radiometric raw data. We have developed an open-access Python program, available on GitHub (https://github.com/Kelian98/tau2_thermalcapture, last access: 1 January 2024; <https://doi.org/10.5281/zenodo.15311830>, Sommer, 2025), to control the camera’s functions and capture images. These images are saved in FITS format (Wells et al., 1981). In this study, we only consider raw analog-to-digital unit (ADU) images for simplification purposes, but the method would be identical with radiometrically calibrated images.

3.2 Datasets and pre-processing

A substantial quantity of images is essential for the effective training and testing of both the classifier and segmentation algorithms. Our dataset comprises LWIR sky images that we captured ourselves. It encompasses a total of 3400 cloudy and clear-sky images for the classifier and 4445 sky images with cloudy-only images for the segmentation algorithm and their associated ground-truth masks. By ground truth, we mean the empirically and manually generated masks. To speed up computations and minimize memory consumption, we downsampled the original-sized images into 160×128 resolution. Cloudy-sky images were collected during a two-night period in early 2023 at Prades-le-Lez ($43^{\circ}41'51''\text{N}$, $3^{\circ}51'53''\text{E}$) during highly variable weather conditions. Con-




Figure 1. Infrared instrument installed on the equatorial table of the prototype experiment at Prades-le-Lez (France), next to the CMOS camera and telescope performing photometric measurements of stars.




Figure 2. Original gray-scale raw infrared thermal image in ADU (640×512 pixels).

versely, cloud-free images were obtained over a shorter period a month later.

To compensate for the lack of cloud-free images and to prevent potential biases in training due to data imbalance, we generated synthetic cloud-free images to create a composite dataset containing as many images as the cloudy dataset. These synthetic images replicate realistic observations by simulating 2D horizontal gradients, mimicking the increase in sky downwelling radiance as the camera's field of view tilts toward high zenith angles (i.e., low-elevation angles).

Realistic sources of noises affecting uncooled infrared thermal cameras are introduced, including read noise, fixed-pattern noise, sky noise, and the narcissus effect. This addition ensures that the spatial noise in the synthetic images closely resembles that of actual cloud-free images. Figure 3 illustrates a typical cloud-free image alongside a synthetically generated one, with the spatial noise indicated for each. It is worth noting that the absolute ADU value has no impact as the data are normalized before training.

All images and masks are visually inspected. Samples presenting artifacts such as tree branches from surroundings or buildings in the FOV corners are discarded. As the camera acquisition frame rate enables us to get ~ 8 images per second, the pre-processing algorithm included constraints on consecutive image selection based on their time series. Selected frames are taken with at least 2 s between each other to introduce a wider range of displayed clouds.

Ground-truth masks identifying cloud structure in cloud images were manually created through multiple distinct steps of non-linear stretching procedures using Astropy (Astropy Collaboration, 2013; The Astropy Collaboration, 2018) methods for each image in the dataset. They consist of a Boolean 2D array of the same image size, where “true” identified pixels represent cloud pixels, and “false” identified pixels represent clear-sky areas. This step has been partially automated. Binary masks that do not capture the cloud structure sufficiently have been kept aside to test segmentation model performance. Figure 5 depicts three raw images with their associated manually generated ground-truth cloud masks for training purposes.

Furthermore, we performed multiple random augmentations (e.g. flip, shear, rotate, shift, and zoom) on each original image to artificially enlarge the size of each dataset and to reduce overfitting (Perez and Wang, 2017; Mikołajczyk and Grochowski, 2018; Yang et al., 2022). All augmented images are produced through the random sequential applications of these five distinct operations to initial images. These operations are executed with a random varying degree of intensity contained in specific ranges. Random rotations are applied within an amplitude ranging from -45° to $+45^\circ$. Shear is introduced with a random magnitude ranging from -0.2 to $+0.2$. Shifting operations are applied to between 0 and 50 pixels in terms of both width and height to avoid the generation of unrealistic symmetric structures. Zoom operation is applied within the range of 1 to 3. No other transformation, such as histogram equalization or contrast enhancement, is applied to prevent any bias or alteration in the segmentation performance. After the selection and augmentation procedures, we conducted a visual examination of all the created sky or cloud images to ensure that they appeared to be realistic. Since all the parameters in the image augmentation process undergo controlled adjustments, our generated images closely mirror authentic sky and/or cloud scenes. Datasets for both models are split into training and validation subsets, with ratios of 80 % and 20 %, respectively.

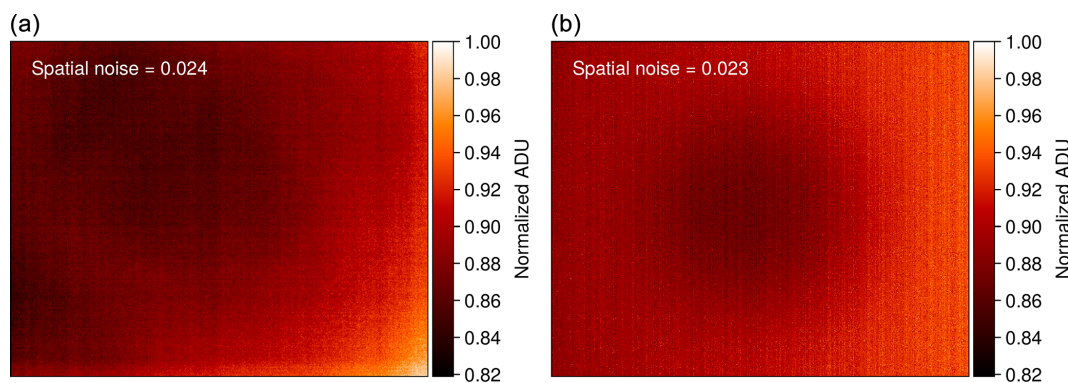


Figure 3. Comparison of a real observed clear-sky image (a) and a synthetically generated realistic image (b). Spatial noise is indicated in the top-right corner of each image. Synthetic images demonstrate high fidelity with regard to overall spatial noise.

4 Methodology

4.1 Overall framework

In this section, we outline the architectural designs of two distinct deep-learning models tailored to automate classification and segmentation tasks. On the one hand, we implement a linear classifier for image classification, whose specific goal is to discriminate between cloud-free (photometric) and cloudy images (non-photometric). On the other hand, the segmentation for cloud structure detection is performed via an optimized U-Net model (Ronneberger et al., 2015) based on pre-classified cloudy images. The output probability map can later be thresholded according to the user's needs to produce the desired predicted binary pixel segmentation map and to allow us to obtain finer details regarding the photometric state of the field at the pixel level. Figure 4 illustrates the proposed deep-learning framework compared to conventional segmentation algorithms.

4.2 Image classification

For our image classification model, we used a RidgeClassifier from scikit-learn (Pedregosa et al., 2011) to classify images as pure sky (clear) or cloudy. The RidgeClassifier is a linear classification model that employs ridge regression (Hoerl and Kennard, 1970), a technique that introduces a regularization term. This regularization helps in addressing multicollinearity, improving the model's stability and robustness, particularly in scenarios with high-dimensional data or collinear predictors. By balancing the trade-off between bias and variance, the RidgeClassifier effectively minimizes overfitting, making it a suitable choice for our image classification task. The training process for the model involves using a dataset consisting of 2720 images, where half of them are cloud infrared images and the other half are cloud-free infrared images, each paired with appropriate ground-truth labels. The datasets are deliberately balanced to avoid any biases in the model that might favor one class over the other.

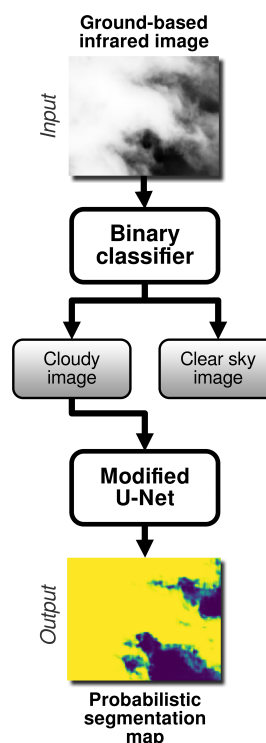


Figure 4. Schematic diagram of the framework proposed in this work. An original infrared image goes through the classifier and is labeled as cloudy or clear. Then, the modified U-Net segmentation model identifies cloud structure in the image to finally produce a probabilistic segmentation map that is used to produce a reliable metric for our application.

4.3 Image segmentation

For cloud structure identification, we adopted the U-Net architecture due to its proven efficiency in semantic segmentation tasks (Ronneberger et al., 2015). The U-Net model comprises an encoder and a decoder, which facilitate the capturing of context-rich features and precise delineation

of cloud structures. The encoder employs convolutions and max-pooling layers to progressively downsample the input image, thereby capturing high-level features. These features are then decoded using up-convolutions and skip connections, enabling the accurate reconstruction of the segmented cloud structures. Figure 5 illustrates some examples of infrared cloud images alongside their corresponding ground-truth masks and predictions. Figure 6 depicts the architecture of the segmentation model.

4.3.1 Encoder block

The encoder block of the segmentation model consists of four sets of double-convolution blocks (hereafter DoubleConv) and four max-pooling layers. A normalized and binned radiometric image of a fixed input size (160×128 pixels) is fed into the model. The DoubleConv contains two sequential convolutional layers, each followed by a rectified linear unit (ReLU) activation function (Agarap, 2018). The initial DoubleConv block applies a set of learnable filters to the input image, extracting low-level features. Subsequent DoubleConv blocks increase the complexity of the learned features by applying convolutions to the feature map generated by the previous layer, creating a hierarchy of increasingly abstract features. Following each set of DoubleConv blocks, a max-pooling layer is applied to downsample the feature map, reducing its spatial dimensions while retaining the most salient information. The architecture follows a pattern of decreasing spatial dimensions while increasing the feature depth as we move through the encoder, with the specified channels at each level being 128, 64, 32, and 16, respectively.

4.3.2 Decoder block

The decoder block also comprises four sets of DoubleConv blocks, mirroring the encoder structure in reverse order (so, in our case, 16, 32, 64, and 128). In contrast to the encoder's sequence, which involves a max-pooling layer following each DoubleConv block for downsampling, the decoder employs an upsampling layer preceding each DoubleConv block. The upsampling operation effectively increases the spatial dimensions of the feature map, preparing it for concatenation with the corresponding, non-downsampled feature map from the encoder provided by the skip connections. Post-concatenation, the DoubleConv block is applied to process the merged feature map. This upsampling followed by a convolution is also known as convolutional transposition or ConvTranspose. These skip connections ensure coherent and effective feature fusion. This structural configuration is essential for seamlessly integrating both local and global contextual information, thereby improving the accuracy of segmentation.

4.3.3 Model output

The image segmentation model generates a probabilistic mask, assigning a probability value to each pixel, indicating its likelihood of being associated with the cloud category. Using an array of 160×128 sigmoid functions, the model produces a continuous probability range between 0 and 1 for individual pixels.

4.3.4 Fine-tuning the U-Net model

In the original U-Net paper from Ronneberger et al. (2015), a basic convolutional block was interposed between the encoder and decoder, functioning as a bottleneck to refine feature maps before their upscaling in the decoding path. Yet, through empirical analysis, we identified that this bottleneck was not necessary for our data processing. While many U-Net adaptations utilize a basic convolutional block for both encoding and decoding, our findings indicated that, during training, the double-convolution blocks outperformed the simple-convolution approach.

4.4 Training procedure and implementation details

The training process comprises two distinct phases, addressing the classifier model and the U-Net segmentation model. The loss function employed for training is the binary cross-entropy, which quantifies the difference between predicted probabilities and actual binary class labels for each instance in the dataset. Mathematically, given an instance's true binary label y (0 or 1) and the predicted probability p of it belonging to class 1, the binary cross-entropy loss \mathcal{L} is calculated as follows:

$$\mathcal{L} = -\frac{1}{N} \sum_i y_i \cdot \log(f_w(x_i)) + (1 - y_i) \cdot \log(1 - f_w(x_i)), \quad (1)$$

where \mathcal{L} is the binary cross-entropy loss. N is the total number of instances in the dataset, i is the index representing an individual instance, y_i is the i th true binary label (0 or 1), and $f_w(x_i)$ is the predicted probability that belongs to class 1 based on the model with parameters w . The goal of training is to minimize this loss function by adjusting the model parameter weights w to better align the predicted probabilities $f_w(x_i)$ with the true labels y_i .

Both the classifier and the segmentation algorithms are implemented using the Python programming language, with the aid of the Flax package (Heek et al., 2023), a neural network library that is part of the JAX ecosystem (Bradbury et al., 2018). Training is conducted on the GPU cluster infrastructure of the MESO@LR (<https://meso-lr.umontpellier.fr/>, last access: 1 January 2024) high-performance computing center, utilizing an NVIDIA Quadro RTX 6000. To expedite computations and to encapsulate the global trend, images are normalized and downsampled to the fixed resolution of 160×128 . The models are trained using the Adam optimizer

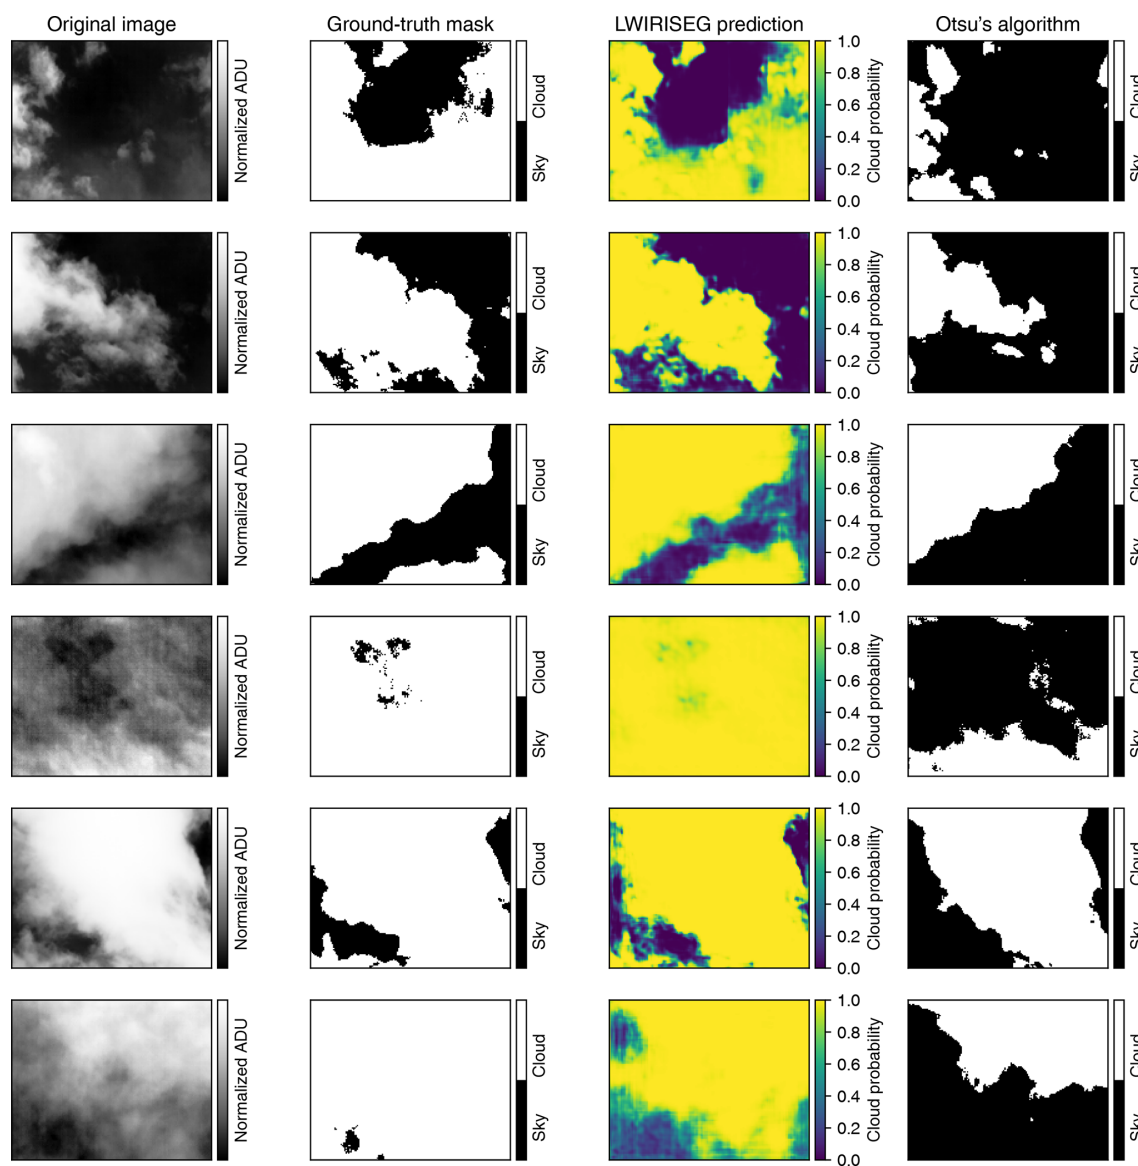


Figure 5. Examples of results with the segmentation model and Otsu's method. Each line represents a different image. The segmentation model displays well-defined cloud structure edges and gives better results than the ground-truth masks and Otsu's algorithm.

(Kingma and Ba, 2014), with a batch size of 64 images. The learning rate is initiated at $\lambda = 10^{-3}$ and decreases with a cosine learning rate decay function (Loshchilov and Hutter, 2017). To prevent overfitting and expedite the training process, an early stopping mechanism is employed, which halts the training if the loss value does not exhibit a decline below a certain threshold after 15 epochs.

During the training of the U-Net model, the hyperparameter-tuning process is carried out to identify the optimal architecture configuration. This process is facilitated by the Optuna framework (Akiba et al., 2019), which employs a sampling strategy algorithm to explore various configurations. The configurations tested cover a range of different filter numbers and sizes in the convolutional

layers of the U-Net. Among the numerous configurations tested, the architecture with channels specified as 128, 64, 32, and 16 for the encoder and decoder blocks achieved the lowest loss, indicating superior performance in segmenting cloud structures.

Furthermore, a pruning strategy is integrated within the Optuna framework to curtail the exploration of sub-optimal configurations early in the training process, thereby significantly reducing the computational resources and time required for the hyperparameter-tuning process. This strategy employs a Median Pruner, which ceases the training of trials exhibiting performance lower than the median performance of completed trials (He et al., 2018; Vadera and Ameen, 2020). The results of the hyperparameter-tuning process re-

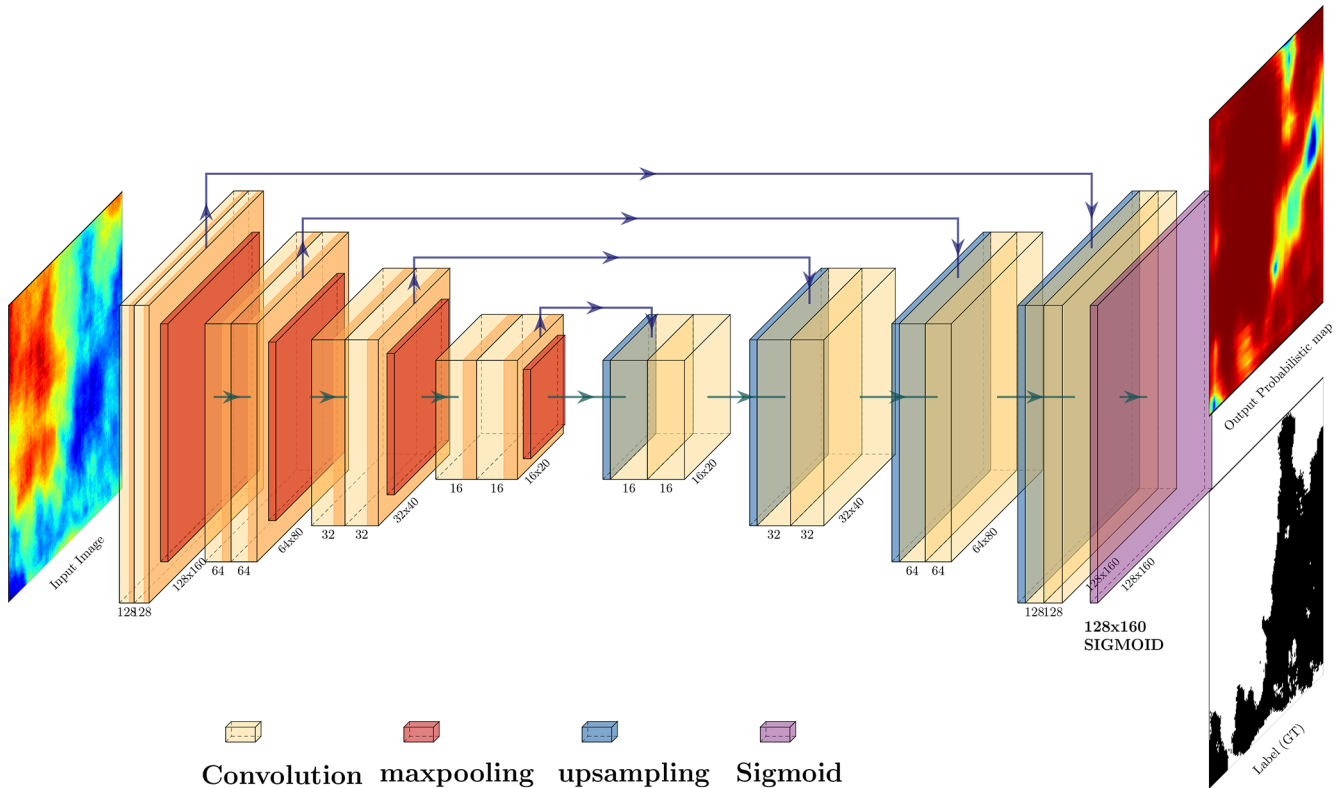


Figure 6. Schematic diagram of the proposed U-Net-based segmentation model architecture. Yellow boxes represent convolutions. Each double convolution is followed by a rectified linear unit (ReLU) activation function. Each convolution size is indicated on the lower right. The input image is a 160×128 gray-scale image. The output image is a probabilistic mask prediction of pixels being cloudy or clear. Arrows represent operations, specified by the legend – notably, green arrows represent convolutions, while purple ones represent skip connections. Tensor dimensions at the output of each block are specified.

veal that the architecture with the specified channels of 128, 64, 32, and 16 outperforms others in terms of loss minimization.

5 Experiments

5.1 Performance metrics

In order to evaluate the performance of the proposed models, we adopt several metrics: accuracy (A), precision (P), recall (R), $F1$ score ($F1$), the area under the curve (AUC), and the binary cross-entropy loss \mathcal{L} defined in Eq. (1). All of these metrics are defined in the equations below.

$$A = \frac{TP + TN}{TP + TN + FP + FN} \quad (2)$$

$$P = \frac{TP}{TP + FP} \quad (3)$$

$$R = \frac{TP}{TP + FN} \quad (4)$$

$$F1 = \frac{2 \cdot P \cdot R}{P + R} \quad (5)$$

$$AUC = \int_0^1 R(FPR^{-1}(t)) dt \quad (6)$$

In the above, true positive (TP) denotes the number of correctly classified positive instances, false positive (FP) denotes the number of negative instances that were incorrectly classified as positive, false negative (FN) denotes the number of positive instances that were incorrectly classified as negative, and true negative (TN) denotes the number of correctly classified negative instances; false-positive rate (FPR) measures the model's ability to incorrectly identify negative instances as positive among all actual negatives and is calculated as $FPR = FP / (FP + TN)$.

5.2 Results

5.2.1 Classifier

We measure the performance of our models using precision and recall metrics. In this context, precision measures the proportion of correctly predicted cloudy images among all

Table 1. Evaluation metrics for the proposed segmentation model based on publicly available state-of-the-art datasets. Note that RGB color images are transformed into gray-scale images as the IRIS-CloudDeep segmentation model is optimized for these types of data. The best values are denoted in bold font (*A* denotes accuracy, *BC* loss denotes binary cross-entropy loss, *AUC* denotes area under the curve). *FS* describes from-scratch training, whereas *FT* means fine-tuning training.

Datasets	<i>A</i> [%]	<i>BC</i> loss	<i>AUC</i>
LWIRISEG	95.17	0.1142	0.98
SWIMSEG	57.51	6.4259	53.32
SWIMSEG FT	84.56	0.3425	84.30
SWIMSEG FS	88.65	0.2680	95.76
SWINSEG	67.06	4.7272	69.55
SWINSEG FT	91.64	0.1999	91.29
SWINSEG FS	93.25	0.1670	93.64

images classified as cloudy. It reflects the classifier’s ability to minimize false positives, where a false positive is an instance of predicting an image to contain clouds when it does not. On the other hand, the recall metric quantifies the proportion of actual cloudy images that are correctly identified by the classifier, addressing its capacity to reduce false negatives. In our case, a false-negative classification refers to an image that contains clouds but that is not recognized as such by the classifier. Table 2 presents the results of the chosen model for the validation subset. All metrics prove the model effectiveness in classifying images, with accuracy, precision, recall, and *F1* score all being above 99 %. The *AUC* value of 0.99 portrays the classifier as robust. The training process on the entire dataset with 10-fold cross-validation takes under 5 min of computing on an average desktop computer.

5.2.2 Segmentation

The results presented in Table 1 demonstrate the efficacy of our segmentation model, achieving an accuracy of 94.64 % and an *AUC* value of 0.97. Figure 7 depicts the binary cross-entropy loss as a function of training epochs. The shape of the decay in these curves aligns with anticipated training patterns, confirming the model’s normal training behavior. The loss stabilizes around the 300-iteration mark, serving as a benchmark for the model’s application.

Figure 8 illustrates the resulting receiver operating characteristic (ROC) curve. As in Dev et al. (2019a), we adopt a threshold of 0.5, nearly balancing true- and false-positive rates. However, users can adjust this threshold to meet specific requirements regarding true-positive rates (TPRs) and/or FPRs. Figure 5 shows the results for some images of the validation subset. We find excellent segmentation of cloud structure in the infrared images.




Figure 7. Training and testing losses of the segmentation model over epochs. The blue curve is for the training subset, whereas the red curve is for the testing subset. After 300 epochs, both curves asymptotically stabilize at a value of approximately 0.10.

Table 2. Evaluation metrics for the proposed classification models and comparison between segmentation methods (*A* denotes accuracy, *P* denotes precision, *R* denotes recall, *F1* denotes *F1* score, and *AUC* denotes area under the curve). The best values are denoted in bold font.

	<i>A</i> [%]	<i>P</i> [%]	<i>R</i> [%]	<i>F1</i> [%]	<i>AUC</i>
Classification models					
Ridge regression	99.26	99.31	99.28	99.28	0.99
Logistic regression	94.85	94.95	95.25	94.85	0.94
Perceptron	93.14	93.27	93.86	93.12	0.93
SVM	91.35	91.52	92.50	91.31	0.91
Segmentation methods					
LWIRISEG	95.17	96.54	97.89	97.21	0.98
Otsu’s algorithm	59.16	100.00	52.53	68.88	0.76

5.3 Application with cloud counting

In our methodological framework, the initial step involves the application of Otsu’s thresholding (Otsu, 1979) to transform infrared sky images into a binary format. This effectively segregates the cloud features from the background, providing a simplified representation where clouds are distinctly highlighted. Following this, the connected-component labeling technique, as implemented by the function `skimage.measure.label`, is employed. This function discerns connected regions within the binary array, where connectivity is defined by the presence of adjacent pixels sharing the same value.

This labeling process assigns a unique identifier to each contiguous cloud region, thus enabling an accurate enumeration of individual cloud formations. By comparing this automated count to visual assessments, our analysis revealed a consistent accuracy in the segmented-image counts. Figure 9 shows examples of cloud counting for different images along with the original, ground-truth and probabilistic segmentation images. The segmented images frequently provided




Figure 8. ROC curve of the segmentation and classification models. The AUC value represents the area under the curve.

counts that closely matched visual estimations, surpassing the performance of raw and binary-mask-derived counts, particularly in scenarios where images suffered from high noise levels. Such robustness underscores the advantage of our segmentation approach in providing reliable cloud quantification.

6 Discussion

6.1 Limited benchmarking comparison

Comparing machine-learning models optimized for different types of input data can be meaningful in certain contexts, but it requires careful consideration. The most critical factor is the nature of the data. State-of-the-art models presented in Sect. 2 are optimized for different types of data (three-channel RGB images). It may seem incoherent to compare the performance with single-channel gray-scale infrared images. Indeed, these data types have distinct characteristics, and models may perform differently based on these differences. Some models (Sun et al., 2011; Liu et al., 2011; Luo et al., 2018) have been proposed to target infrared images with categorization tasks. Considering whether the models can be adapted or fine-tuned to work with both RGB and infrared data is challenging. This might involve multi-modal learning approaches (Liu et al., 2018; Li et al., 2020; Wei et al., 2023) or transfer-learning techniques (Manzo and Pellino, 2021; Wang et al., 2021a; Zhou et al., 2021), which are not within the intended purpose of this work.

Nevertheless, we attempt to evaluate the robustness of our segmentation model by testing its ability to generalize to other datasets, including SWIMSEG (Dev et al., 2016), SWINSEG (Dev et al., 2019b, 2017), and WSISEG (Xie et al., 2020).

We transform RGB images into gray-scaled images with the OpenCV (Bradski, 2000) color conversion method

COLOR_RGB2GRAY, defined by the following equation:

$$\text{RGB} \rightarrow \text{gray} = 0.299 \cdot R + 0.587 \cdot G + 0.114 \cdot B, \quad (7)$$

where R, G, and B are, respectively, red, green, and blue channels of the input color image. Metrics for each dataset are summarized in Table 1.

While the results indicate accurate recognition of most cloud structures in the images by our models, applying a model trained on our dataset directly to another dataset yields less satisfactory performance due to the sub-optimal transformation of color images to gray scale. The method's efficiency is hindered by the strong blue color channels resulting from Rayleigh scattering (Bates, 1984), particularly affecting its performance based on publicly available datasets transformed in this manner. Further efforts are required to enhance the conversion of RGB color to gray-scale images, aiming for a comparable contrast to infrared thermal images. Still, our framework demonstrates satisfactory results when exclusively trained on the modified images (FS in Table 1) as opposed to our original gray-scale images.

6.2 Comparison of segmentation predictions against ground-truth binary mask

The segmentation model occasionally exhibited superior performance compared to the ground-truth masks, as evidenced by examples in Fig. 5. Notably, in this instance, the ground-truth mask failed to identify certain sky patches on the left side of the image. Conversely, the segmentation model's prediction demonstrated a non-zero probability of the existence of sky patches in those areas. This discrepancy arises from the model's utilization of a non-linear mapping technique employing a multi-layer perceptron with ReLU activations. This mapping aims to transform a normalized continuous pixel array into a binary pixel array. Consequently, regions containing sky pixels (denoted by low pixel values) possess the potential to be assigned a low probability value for cloud presence, even when the ground-truth mask assigns certainty (a value of 1) to those patches as cloud-covered areas.

6.3 Comparison of segmentation model with Otsu's method

To validate its effectiveness, the segmentation model is evaluated against the conventional Otsu algorithm with the validation subset from our own LWIRISEG dataset. Otsu's method consists of an adaptive thresholding algorithm that automatically computes the threshold from the image histogram distribution without parameters, supervision, or any prior information (Otsu, 1979). Figure 5 depicts some typical comparison results between the two methods and the ground-truth masks given to the deep-learning model for training. The metrics defined in Sect. 5.1 are computed with Otsu's algorithm and are presented in Table 2. This demonstrates that the proposed deep-learning modified U-Net model performs

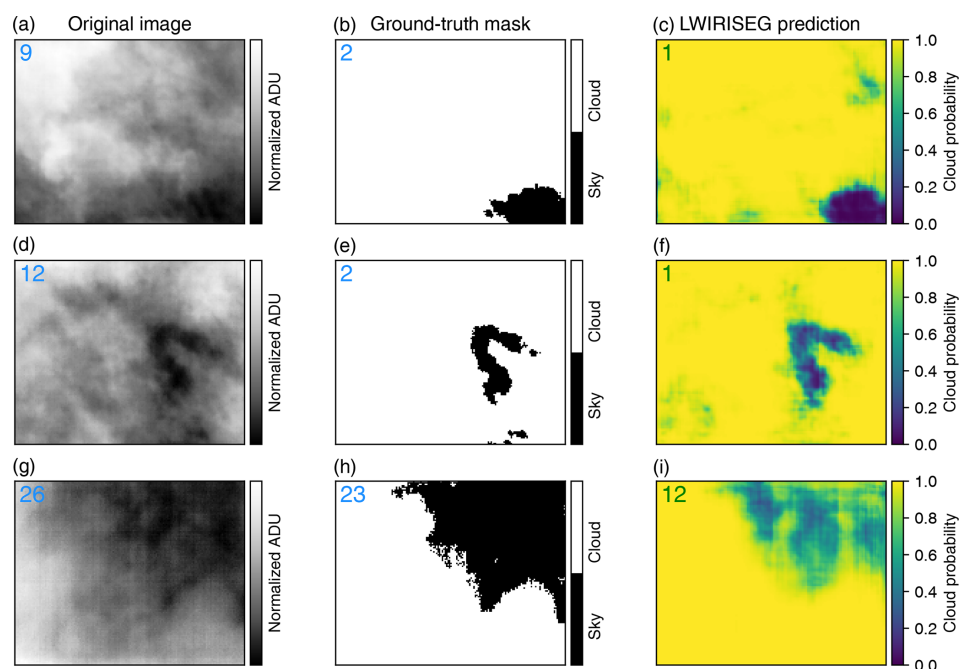


Figure 9. Examples of cloud counting for different input images. The number of automatically identified clouds is shown in the top-left corner of each subplot. For the upper and center rows (a–f), the segmentation map allows the computation of the most accurate result. Panels (g)–(i) depict an example where the computation fails based on the original image, the ground-truth mask, and the probabilistic segmentation mask.

significantly better than Otsu’s algorithm, with mean pixel accuracies being 95.17 % and 59.16 %, respectively. Perfect precision of 100 % means that Otsu’s algorithm does not produce any false positives, implying that it is overly conservative in making positive predictions. As noted by Xie et al. (2020), the primary reason for the sub-par performance of Otsu’s algorithm is its reliance on pixels of the same class having similar gray values, which contradicts the characteristics exhibited by clouds. These experimental findings validate the effectiveness of the segmentation model, highlighting its practical significance for upcoming observations.

6.4 Comparison between linear and non-linear methods for classification

In our analysis of infrared sky images, we assessed the suitability of linear classifiers against more complex deep-learning models. Employing linear classifiers such as SVM, logistic regression, perceptron, and ridge regression, optimized via stochastic gradient descent (SGD) with l2 regularizations, we aimed to prevent overfitting and maintain model simplicity. Among these, ridge regression emerged as the top performer.

Dimensionality reduction techniques supported these findings. Principal component analysis (PCA) demonstrated that two principal components could explain a significant portion of the variance (96.1 %), implying that the data are almost linearly separable. However, non-linear dimensionality reduction through uniform manifold approximation and

projection (UMAP) perfectly segregated cloudy from cloud-free images, indicating that, while the data are nearly linearly separable, non-linear methods provide strong separation, as shown in Fig. 10.

In summary, for the task at hand, deep learning models such as ResNet (He et al., 2015) may seem excessive. A well-tuned linear model, particularly ridge regression, is equally effective, if not more so, due to its interpretability and simplicity. The near-linear nature of the data suggests that simpler models could suffice for such classification challenges.

6.5 Future perspectives

The framework established in this paper is one subpart of the StarDICE data processing operations. This will serve the rest of the analysis by identifying and classifying the quality of photometric exposures performed in parallel.

The work undertaken in this paper will be used and the associated module will include the following operations for analysis: (i) classifying infrared sky images obtained by the imaging system in real time, (ii) analyzing cloud-labeled sky images and deriving the corresponding cloud structure and cover using the segmentation algorithm, and (iii) generating alerts or flags in accordance with the results.

As mentioned in Sect. 6.1, we were able to train a model based on our gray-scale images and another model based on RGB images, and both models produced great results. However, using a model trained on gray-scale images to predict masks for RGB images or even RGB images transformed

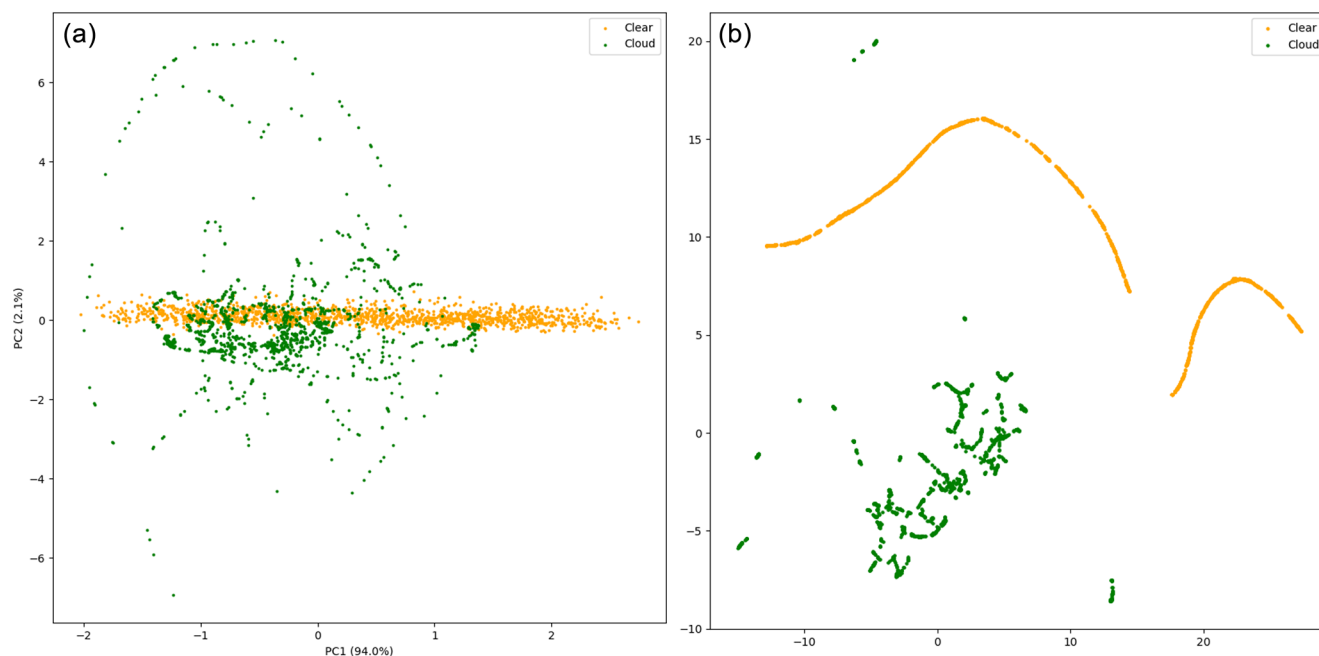


Figure 10. (a) First two principal components of the PCA of the entire dataset, representing 96.1 % of the variance, with the label “clear” referring to the cloud-free images. (b) Two-dimensional representation of UMAP of the same dataset.

to gray scale yielded poor results. Additionally, training a model on all the data resulted in sub-optimal performance. As a future endeavor, we can explore the usage of a multi-modal deep-learning model that can work with both RGB and gray-scale images.

Improving the accuracy and robustness of our framework could involve further training on a larger dataset in more variable conditions. With the upcoming remote-operation capabilities of the telescope system expected to yield a substantial volume of data next year, we anticipate capturing a broader spectrum of sky atmospheric conditions. In this case, a single network could predict two types of outputs (e.g. pixel segmentation map and a metric describing image quality). Still, additional effort will be necessary to categorize the images based on the varying cloud coverage types.

Finally, the standard U-Net model lacks temporal correlation. No information about the displacement of the cloud is taken into account. We could gain in model accuracy by incorporating the temporal information using recurrent neural network (RNN)-based algorithms (Sherstinsky, 2020) or the temporal U-Net to effectively model temporal information in sequences (Funke et al., 2023).

7 Conclusion

In this paper, we proposed a deep-learning framework for the classification and segmentation of ground-based infrared thermal images. As far as we know, it is the first framework that attempts to apply two sequential models for com-

plementary tasks on single-channel gray-scaled infrared images. Specifically, we presented the linear classifier and the U-Net-based segmentation model tailored to extract cloud structures from pre-identified cloud images both during the day or at night. The segmentation model provides the capability to identify clear-sky portions in infrared images, creating a catalog of optical images suitable for photometric measurements and analysis. Extensive experimental results on a combination of self-acquired data and transformed publicly available datasets have demonstrated the effectiveness and performance of the proposed framework. We successfully increased the size of training, testing, and validation subsets with random application of augmentation methods. We developed an accurate simulation tool to produce realistic clear-sky images. Some limitations are due to the low number of strictly different images in various conditions and errors introduced by ground-truth masks incorrectly labeled manually. Nevertheless, we demonstrated that the segmentation model can rectify poor ground-truth masks and sometimes produce better results. In the future, additional data will be collected by the infrared instrument, capturing various weather conditions. The framework may be re-trained based on heavier datasets, which will probably increase its accuracy. Furthermore, if enough data are collected with many different cloud categories and proven-to-be-accurate radiometric calibration, we will be able to expand the segmentation model to perform cloud typology through multi-label segmentation. The framework established in this work will serve as a basis for the sky quality assessment and further analysis for the StarDICE metrology experiment.

Code and data availability. The code for controlling the FLIR Tau2 camera is publicly available online at <https://doi.org/10.5281/zenodo.15311830> (Sommer, 2025). The source code for the work presented in this study is publicly available online at <https://doi.org/10.5281/zenodo.15316607> (Sommer et al., 2025). Datasets and other supporting materials will be made available from the corresponding author Kélian Sommer (kelian.sommer@umontpellier.fr) and Wassim Kabalan (was-sim@apc.in2p3.fr) upon request. Some of the results in this paper have been derived using *astropy* (<https://www.astropy.org/>, last access: 1 January 2024; Astropy Collaboration, 2013; The Astropy Collaboration, 2018), *flax* (<http://github.com/google/flax>, last access: 1 January 2024; Heek et al., 2023), *jax* (<http://github.com/google/jax>, last access: 1 January 2024; Bradbury et al., 2018), *matplotlib* (<https://matplotlib.org/>, last access: 1 January 2024; Hunter, 2007), *numpy* (<https://numpy.org/>, last access: 1 January 2024; Harris et al., 2020), *opencv* (<https://github.com/opencv/opencv>, last access: 1 January 2024; Itseez, 2015), *pandas* (<https://doi.org/10.5281/zenodo.3509134>, Pandas Development Team, 2020), *scipy* (<https://scipy.org/>, last access: 1 January 2024; Virtanen et al., 2020) and *scikit-learn* (<https://scikit-learn.org/stable/>, last access: 1 January 2024; Pedregosa et al., 2011).

Author contributions. KS conceived the instrument, collected data, pre-processed the dataset, and created ground-truth masks for segmentation and realistic synthetic data for classification. WK, RB, and KS designed the framework. WK, RB, and KS performed the experiments. KS and WK wrote the paper and collected the relevant literature. KS, WK, and RB revised the paper.

Competing interests. The contact author has declared that none of the authors has any competing interests.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. While Copernicus Publications makes every effort to include appropriate place names, the final responsibility lies with the authors.

Acknowledgements. We thank the Laboratoire Univers et Particules de Montpellier (CNRS – UMR 5299) for making the publication of this paper possible. We also thank Alexandre Boucaud, Johann Cohen-Tanugi and Bertrand Plez for their thoughtful comments and constructive suggestions on the revision of the paper. This work has been realized with the support of MESO@LR-Platform at the University of Montpellier. English grammar and syntax have been corrected with the help of AI tools, including DeepL (<https://www.deepl.com/translator>, last access: 1 January 2024) and OpenAI ChatGPT (<https://chat.openai.com/>, last access: 1 January 2024).

Financial support. This research has been supported by the Centre National de la Recherche Scientifique (part of regular IN2P3 activities).

Review statement. This paper was edited by Diego Loyola and reviewed by two anonymous referees.

References

- Aebi, C., Gröbner, J., and Kämpfer, N.: Cloud fraction determined by thermal infrared and visible all-sky cameras, *Atmos. Meas. Tech.*, 11, 5549–5563, <https://doi.org/10.5194/amt-11-5549-2018>, 2018.
- Agarap, A. F.: Deep learning using rectified linear units (relu), *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.1803.08375>, 22 March 2018.
- Akiba, T., Sano, S., Yanase, T., Ohta, T., and Koyama, M.: Optuna: A Next-generation Hyperparameter Optimization Framework, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.1907.10902>, 25 July 2019.
- Akula, A., Ghosh, R., and Sardana, H.: Thermal imaging and its application in defence systems, *AIP Conf. Proc.*, 1391, 333–335, 2011.
- Alzubaidi, L., Zhang, J., Humaidi, A. J., Al-Dujaili, A. Q., Duan, Y., Al-Shamma, O., Santamaría, J. I., Fadhel, M. A., Al-Amidie, M., and Farhan, L.: Review of deep learning: concepts, CNN architectures, challenges, applications, future directions, *Journal of Big Data*, 8, 53, <https://doi.org/10.1186/s40537-021-00444-8>, 2021.
- Astropy Collaboration, T.: Astropy: A community Python package for astronomy, *Astron. Astrophys.*, 558, A33, <https://doi.org/10.1051/0004-6361/201322068>, 2013.
- Badrinarayanan, V., Kendall, A., and Cipolla, R.: Segnet: A deep convolutional encoder-decoder architecture for image segmentation, *IEEE T. Pattern Anal.*, 39, 2481–2495, 2017.
- Bates, D.: Rayleigh scattering by air, *Planet. Space Sci.*, 32, 785–790, 1984.
- Betoule, M., Antier, S., Bertin, E., Éric Blanc, P., Bongard, S., Cohen Tanugi, J., Dagoret-Campagne, S., Feinstein, F., Hardin, D., Juramy, C., Le Guillou, L., Le Van Suu, A., Moniez, M., Neveu, J., Nuss, É., Plez, B., Regnault, N., Sepulveda, E., Sommer, K., Souverin, T., and Wang, X. F.: StarDICE I: sensor calibration bench and absolute photometric calibration of a Sony IMX411 sensor, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.2211.04913>, 9 November 2022.
- Boers, R., Acarreta, J. R., and Gras, J. L.: Satellite monitoring of the first indirect aerosol effect: Retrieval of the droplet concentration of water clouds, *J. Geophys. Res.-Atmos.*, 111, D22208, <https://doi.org/10.1029/2005JD006838>, 2006.
- Bohlin, R. C.: Hubble Space Telescopecalspec Flux Standards: Sirius (And Vega), *Astron. J.*, 147, 127, <https://doi.org/10.1088/0004-6256/147/6/127>, 2014.
- Bohlin, R. C., Gordon, K. D., Rieke, G. H., Ardila, D., Carey, S., Deustua, S., Engelbracht, C., Ferguson, H. C., Flanagan, K., Kalirai, J., Meixner, M., Noriega-Crespo, A., Su, K. Y. L., and Tremblay, P.-E.: Absolute Flux Calibration Of The Irac Instrument On The Spitzer Space Telescope Using Hub-

- ble Space Telescope Flux Standards, *Astron. J.*, 141, 173, <https://doi.org/10.1088/0004-6256/141/5/173>, 2011.
- Bohlin, R. C., Hubeny, I., and Rauch, T.: New Grids of Pure-hydrogen White Dwarf NLTE Model Atmospheres and the HST/STIS Flux Calibration, *Astron. J.*, 160, 21, <https://doi.org/10.3847/1538-3881/ab94b4>, 2020.
- Bower, K. N., Choularton, T. W., Gallagher, M. W., Beswick, K. M., Flynn, M. J., Allen, A. G., Davison, B. M., James, J. D., Robertson, L., Harrison, R. M., Hewitt, C. N., Cape, J. N., McFadyen, G. G., Milford, C., Sutton, M. A., Martinsson, B. G., Frank, G., Swietlicki, E., Zhou, J., Berg, O. H., Montes, B., Paspaspiropoulos, G., Hansson, H.-C., Leck, C., Kulmala, M., Aalto, P., Väkevä, M., Berner, A., Bizjak, M., Fuzzi, S., Laj, P., Facchini, M.-C., Orsi, G., Ricci, L., Nielsen, M., Allan, B. J., Coe, H., McFiggans, G., Plane, J. M. C., Collett Jr., J. L., Moore, K. F., and Sherman, D. E.: ACE-2 HILLCLOUD. An overview of the ACE-2 ground-based cloud experiment, *Tellus B*, 52, 750–778, 2000.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q.: JAX: composable transformations of Python+NumPy programs, GitHub [code], <http://github.com/google/jax> (last access: 1 January 2024), 2018.
- Bradski, G.: The OpenCV Library, Dr. Dobb's Journal: Software Tools for the Professional Programmer, Miller Freeman Inc., 25, 120–123, 2000.
- Brout, D., Taylor, G., Scolnic, D., Wood, C. M., Rose, B. M., Vincenzi, M., Dwomoh, A., Lidman, C., Riess, A., Ali, N., Qu, H., and Dai, M.: The Pantheon+ Analysis: SuperCal-fragilistic Cross Calibration, Retrained SALT2 Light-curve Model, and Calibration Systematic Uncertainty, *Astrophys. J.*, 938, 111, <https://doi.org/10.3847/1538-4357/ac8bcc>, 2022.
- Burke, D. L., Axelrod, T., Blondin, S., Claver, C., Željko Ivezić, Jones, L., Saha, A., Smith, A., Smith, R. C., and Stubbs, C. W.: Precision Determination Of Atmospheric Extinction At Optical And Near-Infrared Wavelengths, *Astrophys. J.*, 720, 811, <https://doi.org/10.1088/0004-637X/720/1/811>, 2010.
- Burke, D. L., Saha, A., Claver, J., Axelrod, T., Claver, C., DePoy, D., Ivezić, Ž., Jones, L., Smith, R. C., and Stubbs, C. W.: All-Weather Calibration Of Wide-Field Optical And Nir Surveys, *Astron. J.*, 147, 19, <https://doi.org/10.1088/0004-6256/147/1/19>, 2013.
- Burke, D. L., Rykoff, E. S., Allam, S., Annis, J., Bechtol, K., Bernstein, G. M., Drlica-Wagner, A., Finley, D. A., Gruendl, R. A., James, D. J., Kent, S., Kessler, R., Kuhlmann, S., Lasker, J., Li, T. S., Scolnic, D., Smith, J., Tucker, D. L., Wester, W., Yanny, B., Abbott, T. M. C., Abdalla, F. B., Benoit-Lévy, A., Bertin, E., Rosell, A. C., Kind, M. C., Carretero, J., Cunha, C. E., D'Andrea, C. B., da Costa, L. N., Desai, S., Diehl, H. T., Doel, P., Estrada, J., García-Bellido, J., Gruen, D., Gutierrez, G., Honscheid, K., Kuehn, K., Kuropatkin, N., Maia, M. A. G., March, M., Marshall, J. L., Melchior, P., Menanteau, F., Miquel, R., Plazas, A. A., Sako, M., Sanchez, E., Scarpine, V., Schindler, R., Sevilla-Noarbe, I., Smith, M., Smith, R. C., Soares-Santos, M., Sobreira, F., Suchyta, E., Tarle, G., and Walker, A. R.: Forward Global Photometric Calibration of the Dark Energy Survey, *Astron. J.*, 155, 41, <https://doi.org/10.3847/1538-3881/aa9f22>, 2017.
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., and Yuille, A. L.: DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs, *IEEE T. Pattern Anal.*, 40, 834–848, 2017.
- Conley, A., Guy, J., Sullivan, M., Regnault, N., Astier, P., Balland, C., Basa, S., Carlberg, R. G., Fouchez, D., Hardin, D., Hook, I. M., Howell, D. A., Pain, R., Palanque-Delabrouille, N., Perrett, K. M., Pritchett, C. J., Rich, J., Ruhlmann-Kleider, V., Balam, D., Baumont, S., Ellis, R. S., Fabbro, S., Fakhouri, H. K., Fourmanoit, N., González-Gaitán, S., Graham, M. L., Hudson, M. J., Hsiao, E., Kronborg, T., Lidman, C., Mourao, A. M., Neill, J. D., Perlmutter, S., Ripoché, P., Suzuki, N., and Walker, E. S.: Supernova Constraints and Systematic Uncertainties from the First Three Years of the Supernova Legacy Survey, *American Astronomical Society*, 192, 1, <https://doi.org/10.1088/0067-0049/192/1/1>, 2011.
- Cortes, C. and Vapnik, V.: Support-vector networks, *Mach. Learn.*, 20, 273–297, 1995.
- Currie, M., Rubin, D., Aldering, G., Deustua, S., Fruchter, A., and Perlmutter, S.: Evaluating the Calibration of SN Ia Anchor Datasets with a Bayesian Hierarchical Model, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.2007.02458>, 5 July 2020.
- Dark Energy Survey Collaboration, Abbott, T., Abdalla, F. B., Aleksić, J., Allam, S., Amara, A., Bacon, D., Balbinot, E., Banerji, M., Bechtol, K., Benoit-Lévy, A., Bernstein, G. M., Bertin, E., Blazek, J., Bonnett, C., Bridle, S., Brooks, D., Brunner, R. J., Buckley-Geer, E., Burke, D. L., Caminha, G. B., Capozzi, D., Carlsen, J., Carnero-Rosell, A., Carollo, M., Carrasco-Kind, M., Carretero, J., Castander, F. J., Clerkin, L., Collett, T., Conselice, C., Crocce, M., Cunha, C. E., D'Andrea, C. B., da Costa, L. N., Davis, T. M., Desai, S., Diehl, H. T., Dietrich, J. P., Dodelson, S., Doel, P., Drlica-Wagner, A., Estrada, J., Etherington, J., Evrard, A. E., Fabbri, J., Finley, D. A., Flaugher, B., Foley, R. J., Fosalba, P., Frieman, J., García-Bellido, J., Gaztanaga, E., Gerdes, D. W., Giannantonio, T., Goldstein, D. A., Gruen, D., Gruendl, R. A., Guarnieri, P., Gutierrez, G., Hartley, W., Honscheid, K., Jain, B., James, D. J., Jeltama, T., Jouvel, S., Kessler, R., King, A., Kirk, D., Kron, R., Kuehn, K., Kuropatkin, N., Lahav, O., Li, T. S., Lima, M., Lin, H., Maia, M. A. G., Makler, M., Manera, M., Maraston, C., Marshall, J. L., Martini, P., McMahon, R. G., Melchior, P., Merson, A., Miller, C. J., Miquel, R., Mohr, J. J., Morice-Atkinson, X., Naidoo, K., Neilsen, E., Nichol, R. C., Nord, B., Ogando, R., Ostrovski, F., Palmese, A., Papadopoulos, A., Peiris, H. V., Peoples, J., Percival, W. J., Plazas, A. A., Reed, S. L., Refregier, A., Romer, A. K., Roodman, A., Ross, A., Roza, E., Rykoff, E. S., Sadeh, I., Sako, M., Sánchez, C., Sanchez, E., Santiago, B., Scarpine, V., Schubnell, M., Sevilla-Noarbe, I., Sheldon, E., Smith, M., Smith, R. C., Soares-Santos, M., Sobreira, F., Soumagnac, M., Suchyta, E., Sullivan, M., Swanson, M., Tarle, G., Thaler, J., Thomas, D., Thomas, R. C., Tucker, D., Vieira, J. D., Vikram, V., Walker, A. R., Wechsler, R. H., Weller, J., Wester, W., Whiteway, L., Wilcox, H., Yanny, B., Zhang, Y., and Zuntz, J.: The Dark Energy Survey: more than dark energy – an overview, *Mon. Not. R. Astron. Soc.*, 460, 1270–1299, <https://doi.org/10.1093/mnras/stw641>, 2016.
- Dev, S., Lee, Y. H., and Winkler, S.: Color-based Segmentation of Sky/Cloud Images From Ground-based Cameras, *IEEE J. Sel. Top. Appl.*, 10, 231–242, 2016.
- Dev, S., Savoy, F. M., Lee, Y. H., and Winkler, S.: Night-time sky/cloud image segmentation, *IEEE Press*, 345–349, <https://doi.org/10.1109/ICIP.2017.8296300>, 2017.

- Dev, S., Nautiyal, A., Lee, Y. H., and Winkler, S.: CloudSegNet: A deep network for nychthemeron cloud image segmentation, *IEEE Geosci. Remote S.*, 16, 1814–1818, 2019a.
- Dev, S., Savoy, F., Lee, Y. H., and Winkler, S.: Singapore Whole sky Nighttime Image SEGmentation Database, *IEEE Dataport* [data set], <https://doi.org/10.21227/jsf0-ga67>, 2019b.
- Fa, T., Xie, W., Yiren, W., and Xia, Y.: Development of an all-sky imaging system for cloud cover assessment, *Appl. Optics*, 58, 5516, <https://doi.org/10.1364/AO.58.005516>, 2019.
- Funke, I., Rivoir, D., Krell, S., and Speidel, S.: TUNeS: A Temporal U-Net with Self-Attention for Video-based Surgical Phase Recognition, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.2307.09997>, 19 July 2023.
- Gallo, M. A., Willits, D. S., Lubke, R. A., and Thiede, E. C.: Low-cost uncooled IR sensor for battlefield surveillance, in: *Infrared Technology XIX*, SPIE, vol. 2020, 351–362, 1993.
- Geer, A., Baordo, F., Bormann, N., Chambon, P., English, S., Kazumori, M., Lawrence, H., Lean, P., Lonitz, K., and Lupu, C.: The growing impact of satellite observations sensitive to humidity, cloud and precipitation, *Q. J. Roy. Meteor. Soc.*, 143, 3189–3206, 2017.
- Goobar, A. and Leibundgut, B.: Supernova Cosmology: Legacy and Future, *Annu. Rev. Nucl. Part. S.*, 61, 251–279, <https://doi.org/10.1146/annurev-nucl-102010-130434>, 2011.
- Hack, E. D., Pauliquevis, T., Barbosa, H. M. J., Yamasoe, M. A., Klebe, D., and Correia, A. L.: Precipitable water vapor retrievals using a ground-based infrared sky camera in subtropical South America, *Atmos. Meas. Tech.*, 16, 1263–1278, <https://doi.org/10.5194/amt-16-1263-2023>, 2023.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Rio, J. F., Wiebe, M., Peterson, P., Gerard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E.: Array programming with NumPy, *Nature*, 585, 357–362, <https://doi.org/10.1038/s41586-020-2649-2>, 2020.
- Hazenber, F.: Calibration photométrique des supernovae de type Ia pour la caractérisation de l'énergie noire avec l'expérience StarDICE, PhD thesis, Sorbonne université, <http://www.theses.fr/2019SORUS142> (last access: 1 January 2024), 2019.
- He, K., Zhang, X., Ren, S., and Sun, J.: Deep Residual Learning for Image Recognition, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.1512.03385>, 10 December 2015.
- He, Y., Liu, P., Wang, Z., Hu, Z., and Yang, Y.: Filter Pruning via Geometric Median for Deep Convolutional Neural Networks Acceleration, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.1811.00250>, 1 November 2018.
- Heek, J., Levskaya, A., Oliver, A., Ritter, M., Rondepierre, B., Steiner, A., and van Zee, M.: Flax: A neural network library and ecosystem for JAX, *GitHub [code]*, <http://github.com/google/flax> (last access: 1 January 2024), 2023.
- Hoerl, A. E. and Kennard, R. W.: Ridge regression: Biased estimation for nonorthogonal problems, *Technometrics*, 12, 55–67, 1970.
- Houghton, J. T. and Lee, A. C. L.: Atmospheric Transmission in the 10–12 μm Window, *Nature Physical Science*, 238, 117–118, <https://doi.org/10.1038/physci238117a0>, 1972.
- Hu, Y., Wielicki, B. A., Yang, P., Stackhouse, P. W., Lin, B., and Young, D. F.: Application of deep convective cloud albedo observation to satellite-based study of the terrestrial atmosphere: Monitoring the stability of spaceborne measurements and assessing absorption anomaly, *IEEE T. Geosci. Remote*, 42, 2594–2599, 2004.
- Hunter, J. D.: Matplotlib: A 2D Graphics Environment, *Comput. Sci. Eng.*, 9, 90–95, <https://doi.org/10.1109/MCSE.2007.55>, 2007.
- Ishimwe, R., Abutaleb, K., Ahmed, F., Ishimwe, R., Abutaleb, K., and Ahmed, F.: Applications of thermal imaging in agriculture – A review, *Adv. Remote Sens.*, 3, 128–140, <https://doi.org/10.4236/ars.2014.33011>, 2014.
- Itseez: Open Source Computer Vision Library, *GitHub [code]*, <https://github.com/itseez/opencv> (last access: 1 January 2024), 2015.
- Kelsey, V., Riley, S., and Minschwaner, K.: Atmospheric precipitable water vapor and its correlation with clear-sky infrared temperature observations, *Atmos. Meas. Tech.*, 15, 1563–1576, <https://doi.org/10.5194/amt-15-1563-2022>, 2022.
- Kimata, M.: Uncooled infrared focal plane arrays, *IEEJ T. Electr. Electr.*, 13, 4–12, <https://doi.org/10.1002/tee.22563>, 2018.
- Kingma, D. P. and Ba, J.: Adam: A method for stochastic optimization, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.1412.6980>, 22 December 2014.
- Klebe, D., Sebag, J., Blatherwick, R. D., and Zimmer, P. C.: All-Sky Mid-Infrared Imagery to Characterize Sky Conditions and Improve Astronomical Observational Performance, *Publ. Astron. Soc. Pac.*, 124, 1309, <https://doi.org/10.1086/668866>, 2012.
- Klebe, D. I., Blatherwick, R. D., and Morris, V. R.: Ground-based all-sky mid-infrared and visible imagery for purposes of characterizing cloud properties, *Atmos. Meas. Tech.*, 7, 637–645, <https://doi.org/10.5194/amt-7-637-2014>, 2014.
- Krauz, L., Janout, P., Blažek, M., and Páta, P.: Assessing Cloud Segmentation in the Chromacity Diagram of All-Sky Images, *Remote Sens.*, 12, 1902, <https://doi.org/10.3390/rs12111902>, 2020.
- Larason, T. and Houston, J.: Spectroradiometric Detector Measurements: Ultraviolet, Visible, and Near Infrared Detectors for Spectral Power, *NIST Special Publication 250-41*, https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=841061 (last access: 1 January 2024), 2008.
- Lewis, P. M., Rogers, H., and Schindler, R. H.: A radiometric all-sky infrared camera (RASICAM) for DES/CTIO, in: *Ground-based and Airborne Instrumentation for Astronomy III*, SPIE, vol. 7735, 1307–1318, 2010.
- Li, M., Liu, S., and Zhang, Z.: Deep tensor fusion network for multimodal ground-based cloud classification in weather station networks, *Ad Hoc Netw.*, 96, 101991, <https://doi.org/10.1016/j.adhoc.2019.101991>, 2020.
- Li, Q., Lu, W., and Yang, J.: A Hybrid Thresholding Algorithm for Cloud Detection on Ground-Based Color Images, *J. Atmos. Ocean. Tech.*, 28, 1286–1296, <https://doi.org/10.1175/JTECH-D-11-00009.1>, 2011.
- Li, S., Wang, M., Wu, J., Sun, S., and Zhuang, Z.: Cloud-DeepLabV3+: a lightweight ground-based cloud segmentation method based on multi-scale feature aggregation and multi-level attention feature enhancement, *Int. J. Remote Sens.*, 44, 4836–4856, <https://doi.org/10.1080/01431161.2023.2240034>, 2023.

- Li, T. S., DePoy, D. L., Marshall, J. L., Tucker, D., Kessler, R., Anis, J., Bernstein, G. M., Boada, S., Burke, D. L., Finley, D. A., James, D. J., Kent, S., Lin, H., Marriner, J., Mondrik, N., Nagasawa, D., Rykoff, E. S., Scolnic, D., Walker, A. R., Wester, W., Abbott, T. M. C., Allam, S., Benoit-Lévy, A., Bertin, E., Brooks, D., Capozzi, D., Rosell, A. C., Kind, M. C., Carretero, J., Croce, M., Cunha, C. E., D'Andrea, C. B., da Costa, L. N., Desai, S., Diehl, H. T., Doel, P., Flaughner, B., Fosalba, P., Frieman, J., Gaztanaga, E., Goldstein, D. A., Gruen, D., Gruendl, R. A., Gutierrez, G., Honscheid, K., Kuehn, K., Kuropatkin, N., Maia, M. A. G., Melchior, P., Miller, C. J., Miquel, R., Mohr, J. J., Neilsen, E., Nichol, R. C., Nord, B., Ogando, R., Plazas, A. A., Romer, A. K., Roodman, A., Sako, M., Sanchez, E., Scarpine, V., Schubnell, M., Sevilla-Noarbe, I., Smith, R. C., Soares-Santos, M., Sobreira, F., Suchyta, E., Tarle, G., Thomas, D., and Vikram, V.: Assessment Of Systematic Chromatic Errors That Impact Sub-1 % Photometric Precision In Large-Area Sky Surveys, *Astron. J.*, 151, 157, <https://doi.org/10.3847/0004-6256/151/6/157>, 2016.
- Li, X., Wang, B., Qiu, B., and Wu, C.: An all-sky camera image classification method using cloud cover features, *Atmos. Meas. Tech.*, 15, 3629–3639, <https://doi.org/10.5194/amt-15-3629-2022>, 2022.
- Liandrat, O., Cros, S., Braun, A., Saint-Antonin, L., Decroix, J., and Schmutz, N.: Cloud cover forecast from a ground-based all sky infrared thermal camera, in: *Remote Sensing of Clouds and the Atmosphere XXII*, Proc. SPIE 10424, 10 pp., <https://doi.org/10.1117/12.2278636>, 2017a.
- Liandrat, O., Cros, S., Braun, A., Saint-Antonin, L., Decroix, J., and Schmutz, N.: Cloud cover forecast from a ground-based all sky infrared thermal camera, in: *Remote Sensing of Clouds and the Atmosphere XXII*, SPIE, vol. 10424, 19–31, 2017b.
- Liou, K.-N.: *Radiation and cloud processes in the atmosphere. Theory, observation, and modeling*, Oxford University Press, New York, ISBN 978-0-19-504910-7, 1992.
- Liu, L., Sun, X., Chen, F., Zhao, S., and Gao, T.: Cloud Classification Based on Structure Features of Infrared Images, *J. Atmos. Ocean. Tech.*, 28, 410–417, <https://doi.org/10.1175/2010JTECHA1385.1>, 2011.
- Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., and Han, J.: On the variance of the adaptive learning rate and beyond, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.1908.03265>, 8 August 2019.
- Liu, S., Li, M., Zhang, Z., Xiao, B., and Cao, X.: Multimodal ground-based cloud classification using joint fusion convolutional neural network, *Remote Sens.*, 10, 822, <https://doi.org/10.3390/rs10060822>, 2018.
- Liu, S., Zhang, J., Zhang, Z., Cao, X., and Durrani, T. S.: TransCloudSeg: Ground-Based Cloud Image Segmentation With Transformer, *IEEE J. Sel. Top. Appl.*, 15, 6121–6132, <https://doi.org/10.1109/JSTARS.2022.3194316>, 2022.
- Liu, Y., Shao, Z., and Hoffmann, N.: Global Attention Module: Better Exploiting Non-local Similarity in Deep CNNs, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.2112.05561>, 10 December 2021.
- Long, C. N., Sabburg, J. M., Calbó, J., and Pagès, D.: Retrieving Cloud Characteristics from Ground-Based Daytime Color All-Sky Images, *J. Atmos. Ocean. Tech.*, 23, 633–652, <https://doi.org/10.1175/JTECH1875.1>, 2006.
- Lopez, T., Antoine, R., Baratoux, D., and Rabinowicz, M.: Contribution of thermal infrared images on the understanding of the subsurface/atmosphere exchanges on Earth., in: *EGU General Assembly 2018*, 8–13 April 2018, Vienna, Austria, EGU General Assembly Conference Abstracts, 19, EGU2017-11811, 2017.
- Loshchilov, I. and Hutter, F.: SGDR: Stochastic Gradient Descent with Warm Restarts, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.1608.03983>, 23 February 2017.
- Luo, Q., Meng, Y., Liu, L., Zhao, X., and Zhou, Z.: Cloud classification of ground-based infrared images combining manifold and texture features, *Atmos. Meas. Tech.*, 11, 5351–5361, <https://doi.org/10.5194/amt-11-5351-2018>, 2018.
- Makwana, D., Nag, S., Susladkar, O., Deshmukh, G., Teja R, S. C., Mittal, S., and Mohan, C. K.: ACLNet: an attention and clustering-based cloud segmentation network, *Remote Sens. Lett.*, 13, 865–875, 2022.
- Mandat, D., Pech, M., Hrabovsky, M., Schovanek, P., Palatka, M., Travnicek, P., Prouza, M., and Ebr, J.: All Sky Camera instrument for night sky monitoring, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.1402.4762>, 19 February 2014.
- Manzo, M. and Pellino, S.: Voting in transfer learning system for ground-based cloud classification, *Machine Learning and Knowledge Extraction*, 3, 542–553, 2021.
- Martin, R. V.: Satellite remote sensing of surface air quality, *Atmos. Environ.*, 42, 7823–7843, 2008.
- McCallum, A.: Efficiently inducing features of conditional random fields, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.1212.2504>, 19 October 2012.
- Mikołajczyk, A. and Grochowski, M.: Data augmentation for improving deep learning in image classification problem, in: *2018 International Interdisciplinary PhD Workshop (IIPHDW)*, 9–12 May 2018, Swinoujscie, Poland, 117–122, <https://doi.org/10.1109/IIPHDW.2018.8388338>, 2018.
- Mokhov, I. L. and Schlesinger, M. E.: Analysis of global cloudiness: 2. Comparison of ground-based and satellite-based cloud climatologies, *J. Geophys. Res.-Atmos.*, 99, 17045–17065, 1994.
- Mommert, M.: Cloud Identification from All-sky Camera Data with Machine Learning, *Astrophys. J.*, 159, 178, <https://doi.org/10.3847/1538-3881/ab744f>, 2020.
- Nikolenko, I. and Maslov, I.: Infrared (thermal) camera for monitoring the state of the atmosphere above the sea horizon of the Simeiz Observatory INASAN, *INASAN Science Reports*, 6, 85–87, 2021.
- O'Shea, K. and Nash, R.: An Introduction to Convolutional Neural Networks, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.1511.08458>, 26 November 2015.
- Otsu, N.: A Threshold Selection Method from Gray-Level Histograms, *IEEE T. Syst. Man Cybern.*, 9, 62–66, <https://doi.org/10.1109/TSMC.1979.4310076>, 1979.
- Paczyński, B.: Monitoring All Sky for Variability, *Publ. Astron. Soc. Pac.*, 112, 1281, <https://doi.org/10.1086/316623>, 2000.
- Pandas Development Team: pandas-dev/pandas: Pandas, Zenodo [code], <https://doi.org/10.5281/zenodo.3509134>, 2020.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E.: Scikit-learn: Machine Learning in Python, *J. Mach. Learn. Res.*, 12, 2825–2830, 2011.

- Perez, L. and Wang, J.: The Effectiveness of Data Augmentation in Image Classification using Deep Learning, arXiv [preprint], <https://doi.org/10.48550/arXiv.1712.04621>, 13 December 2017.
- Petzold, A., Marshall, J., Nédélec, P., Smit, H. G. J., Frieß, U., Flaud, J.-M., Wahner, A., Cammas, J.-P., Volz-Thomas, A., and the IAGOS Team: Global-scale atmosphere monitoring by in-service aircraft—current achievements and future prospects of the European Research Infrastructure IAGOS, *Tellus B*, 67, 28452, <https://doi.org/10.3402/tellusb.v67.28452>, 2015.
- Reil, K., Lewis, P., Schindler, R., and Zhang, Z.: An update on the status and performance of the Radiometric All-Sky Infrared Camera (RASICAM), in: *Observatory Operations: Strategies, Processes, and Systems V*, SPIE, vol. 9149, 321–331, 2014.
- Ring, E. and Ammer, K.: Infrared thermal imaging in medicine, *Physiol. Meas.*, 33, R33–R46, 2012.
- Rogalski, A.: Recent progress in infrared detector technologies, *Infrared Phys. Techn.*, 54, 136–154, <https://doi.org/10.1016/j.infrared.2010.12.003>, 2011.
- Rogalski, A. and Chrzanowski, K.: Infrared Devices And Techniques (Revision), *Metrol. Meas. Syst.*, 21, 565–618, <https://doi.org/10.2478/mms-2014-0057>, 2014.
- Román, R., Antuña-Sánchez, J. C., Cachorro, V. E., Toledano, C., Torres, B., Mateos, D., Fuertes, D., López, C., González, R., Lapiónok, T., Herreras-Giralda, M., Dubovik, O., and de Frutos, Á. M.: Retrieval of aerosol properties using relative radiance measurements from an all-sky camera, *Atmos. Meas. Tech.*, 15, 407–433, <https://doi.org/10.5194/amt-15-407-2022>, 2022.
- Ronneberger, O., Fischer, P., and Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation, arXiv [preprint], <https://doi.org/10.48550/arXiv.1505.04597>, 18 May 2015.
- Roy, P., Behera, M., and Srivastav, S.: Satellite remote sensing: sensors, applications and techniques, *P. Natl. A. Sci. India A*, 87, 465–472, 2017.
- Rubin, D., Aldering, G., Amanullah, R., Barbary, K., Dawson, K. S., Deustua, S., Faccioli, L., Fadeyev, V., Fakhouri, H. K., Fruchter, A. S., Gladders, M. D., de Jong, R. S., Koekemoer, A., Krechmer, E., Lidman, C., Meyers, J., Nordin, J., Perlmutter, S., Ripoche, P., Schlegel, D. J., Spadafora, A., and Suzuki, N.: A Calibration Of Nicmos Camera 2 For Low Count Rates, *Astron. J.*, 149, 159, <https://doi.org/10.1088/0004-6256/149/5/159>, 2015.
- Rubin, D., Aldering, G., Antilogus, P., Aragon, C., Bailey, S., Batai, C., Bongard, S., Boone, K., Buton, C., Copin, Y., Dixon, S., Fouchez, D., Gangler, E., Gupta, R., Hayden, B., Hillebrandt, W., Kim, A. G., Kowalski, M., Kuesters, D., Leget, P. F., Mondon, F., Nordin, J., Pain, R., Pecontal, E., Pereira, R., Perlmutter, S., Ponder, K. A., Rabinowitz, D., Rigault, M., Runge, K., Saunders, C., Smadja, G., Suzuki, N., Tao, C., Taubenberger, S., Thomas, R. C., Vincenzi, M., and Factory, T. N. S.: The LSST DESC DC2 Simulated Sky Survey, arXiv [preprint], <https://doi.org/10.48550/arXiv.2205.01116>, 21 June 2022.
- Salamalikis, V., Tzoumanikas, P., Argiriou, A. A., and Kazantzidis, A.: Estimation of Precipitable Water Using Thermal Infrared Images, *Environmental Sciences Proceedings*, 26, 33, <https://doi.org/10.3390/envirosciproc2023026033>, 2023.
- Schiffer, R. A. and Rossow, W. B.: The International Satellite Cloud Climatology Project (ISCCP): The first project of the world climate research programme, *B. Am. Meteorol. Soc.*, 64, 779–784, 1983.
- Schreiner, A. J., Unger, D. A., Menzel, W. P., Ellrod, G. P., Strabala, K. I., and Pellet, J. L.: A comparison of ground and satellite observations of cloud cover, *B. Am. Meteorol. Soc.*, 74, 1851–1862, 1993.
- Scolnic, D., Casertano, S., Riess, A., Rest, A., Schlafly, E., Foley, R. J., Finkbeiner, D., Tang, C., Burgett, W. S., Chambers, K. C., Draper, P. W., Flewelling, H., Hodapp, K. W., Huber, M. E., Kaiser, N., Kudritzki, R. P., Magnier, E. A., Metcalfe, N., and Stubbs, C. W.: Supercal: Cross-Calibration Of Multiple Photometric Systems To Improve Cosmological Measurements With Type Ia Supernovae, *Astrophys. J.*, 815, 117, <https://doi.org/10.1088/0004-637X/815/2/117>, 2015.
- Sebag, J., Andrew, J., Klebe, D., and Blatherwick, R.: LSST all-sky IR camera cloud monitoring test results, *Proc. SPIE*, 7733, 773348, <https://doi.org/10.1117/12.856337>, 2010.
- Sharma, T., Parihar, P., and Kemkar, M.: All sky scanning cloud monitor for NLOT site survey, *J. Phys. Conf. Ser.*, 595, 012032, <https://doi.org/10.1088/1742-6596/595/1/012032>, 2015.
- Shaw, J. A. and Nugent, P. W.: Physics principles in radiometric infrared imaging of clouds in the atmosphere, *Eur. J. Phys.*, 34, S111–S121, <https://doi.org/10.1088/0143-0807/34/6/s111>, 2013.
- Sherstinsky, A.: Fundamentals of Recurrent Neural Network (RNN) and Long Short-Term Memory (LSTM) network, *Physica D*, 404, 132306, <https://doi.org/10.1016/j.physd.2019.132306>, 2020.
- Shi, C., Zhou, Y., and Qiu, B.: CloudU-Netv2: A cloud segmentation method for ground-based cloud images based on deep learning, *Neural Process. Lett.*, 53, 2715–2728, 2021a.
- Shi, C., Zhou, Y., Qiu, B., Guo, D., and Li, M.: CloudU-Net: A Deep Convolutional Neural Network Architecture for Daytime and Nighttime Cloud Images' Segmentation, *IEEE Geosci. Remote S.*, 18, 1688–1692, <https://doi.org/10.1109/LGRS.2020.3009227>, 2021b.
- Shi, C., Zhou, Y., and Qiu, B.: CloudRaednet: residual attention-based encoder–decoder network for ground-based cloud images segmentation in nychthemeron, *Int. J. Remote Sens.*, 43, 2059–2075, 2022.
- Skidmore, W., Els, S., Travouillon, T., Riddle, R., Seguel, J., Bustos, E., Walker, D., and Blum, R.: Using All Sky Cameras to determine cloud statistics for the Thirty Meter Telescope candidate sites, in: *Ground-based and Airborne Telescopes II*, SPIE, vol. 7012, 862–870, 2008.
- Smith, S. and Toumi, R.: Measuring Cloud Cover and Brightness Temperature with a Ground-Based Thermal Infrared Camera, *J. Appl. Meteorol. Clim.*, 47, 683–693, <https://doi.org/10.1175/2007JAMC1615.1>, 2008.
- Sommer, K.: Kelian98/tau2_thermalcapture: Release for Zenodo (main), Zenodo [code], <https://doi.org/10.5281/zenodo.15311830>, 2025.
- Sommer, K., Brunet, R., and Kabalan, W.: ASKabalan/infrared-cloud-detection: v0.1.0 (v0.1.0), Zenodo [code], <https://doi.org/10.5281/zenodo.15316607>, 2025.
- Stubbs, C. W. and Brown, Y. J.: Precise astronomical flux calibration and its impact on studying the nature of the dark energy, *Mod. Phys. Lett. A*, 30, 1530030, <https://doi.org/10.1142/S021773231530030X>, 2015.
- Stubbs, C. W. and Tonry, J. L.: The Calibration of Astronomical Observations Made Through the Atmosphere, arXiv [preprint], <https://doi.org/10.48550/arXiv.1206.6695>, 28 June 2012.

- Sun, X., Liu, L., and Zhao, S.: Whole Sky Infrared Remote Sensing of Cloud, *Proced. Earth Plan. Sc.*, 2, 278–283, <https://doi.org/10.1016/j.proeps.2011.09.044>, 2011.
- Sun, X.-J., Gao, T.-C., Zhai, D.-L., Zhao, S.-J., and Lian, J.-G.: Whole sky infrared cloud measuring system based on the uncooled infrared focal plane array, *Infrared and Laser Engineering*, 37, 761–764, 2008.
- Szejwach, G.: Determination of Semi-Transparent Cirrus Cloud Temperature from Infrared Radiances: Application to METEOSAT, *J. Appl. Meteorol. Clim.*, 21, 384–393, [https://doi.org/10.1175/1520-0450\(1982\)021<0384:DOSTCC>2.0.CO;2](https://doi.org/10.1175/1520-0450(1982)021<0384:DOSTCC>2.0.CO;2), 1982.
- The Astropy Collaboration: Astropy: A community Python package for astronomy, *Astron. Astrophys.*, 558, A33, <https://doi.org/10.3847/1538-3881/aac387>, 2018.
- Thuraiajah, B. and Shaw, J.: Cloud statistics measured with the infrared cloud imager (ICI), *IEEE T. Geosci. Remote*, 43, 2000–2007, <https://doi.org/10.1109/TGRS.2005.853716>, 2005.
- TLCDD: shuangliutjnu/TJNU-Large-Scale-Cloud-Detection-Database: v1.0.0, Zenodo [code], <https://doi.org/10.5281/zenodo.6464743>, 2022.
- Tzoumanikas, P., Nikitidou, E., Bais, A., and Kazantzidis, A.: The effect of clouds on surface solar irradiance, based on data from an all-sky imaging system, *Renew. Energ.*, 95, 314–322, 2016.
- Ugolnikov, O. S., Galkin, A. A., Pilgaev, S. V., and Roldugin, A. V.: Noctilucous cloud particle size determination based on multi-wavelength all-sky analysis, *Planet. Space Sci.*, 146, 10–19, 2017.
- Vadera, S. and Ameen, S.: Methods for Pruning Deep Neural Networks, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.2011.00241>, 31 October 2020.
- Várnai, T. and Marshak, A.: Satellite observations of cloud-related variations in aerosol properties, *Atmosphere*, 9, 430, <https://doi.org/10.3390/atmos9110430>, 2018.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I.: Attention Is All You Need, in: *Advances in Neural Information Processing Systems*, edited by: Guyon, I., Von Luxburg, U., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., Curran Associates, Inc., https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf (last access: 1 January 2024), 2017.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, *Nat. Methods*, 17, 261–272, <https://doi.org/10.1038/s41592-019-0686-2>, 2020.
- Wang, M., Zhuang, Z., Wang, K., Zhou, S., and Liu, Z.: Intelligent classification of ground-based visible cloud images using a transfer convolutional neural network and fine-tuning, *Opt. Express*, 29, 41176–41190, <https://doi.org/10.1364/OE.442455>, 2021a.
- Wang, Y., Liu, D., Xie, W., Yang, M., Gao, Z., Ling, X., Huang, Y., Li, C., Liu, Y., and Xia, Y.: Day and Night Clouds Detection Using a Thermal-Infrared All-Sky-View Camera, *Remote Sens.*, 13, 1852, <https://doi.org/10.3390/rs13091852>, 2021b.
- Wei, L., Zhu, T., Guo, Y., and Ni, C.: MMST: A Multi-Modal Ground-Based Cloud Image Classification Method, *Sensors*, 23, 4222, <https://doi.org/10.3390/s23094222>, 2023.
- Wells, D. C., Greisen, E. W., and Harten, R. H.: FITS – a Flexible Image Transport System, *Astronomy and Astrophysics Supplement*, 44, 363–370, 1981.
- Wilczak, J., Gossard, E., Neff, W., and Eberhard, W.: Ground-based remote sensing of the atmospheric boundary layer: 25 years of progress, *Bound.-Lay. Meteorol.*, 78, 321–349, 1996.
- Wong, W. K., Tan, P. N., Loo, C. K., and Lim, W. S.: An effective surveillance system using thermal camera, in: *2009 International Conference on Signal Acquisition and Processing*, 3–5 April 2009, Kuala Lumpur, Malaysia, 13–17, <https://doi.org/10.1109/ICSAP.2009.12>, 2009.
- Xie, W., Liu, D., Yang, M., Chen, S., Wang, B., Wang, Z., Xia, Y., Liu, Y., Wang, Y., and Zhang, C.: SegCloud: a novel cloud image segmentation model using a deep convolutional neural network for ground-based all-sky-view camera observation, *Atmos. Meas. Tech.*, 13, 1953–1961, <https://doi.org/10.5194/amt-13-1953-2020>, 2020.
- Yamashita, M. and Yoshimura, M.: Ground-based cloud observation for satellite-based cloud discrimination and its validation, *Int. Arch. Photogramm.*, 39, 137–140, 2012.
- Yang, J., Lu, W., Ma, Y., and Yao, W.: An Automated Cirrus Cloud Detection Method for a Ground-Based Cloud Image, *J. Atmos. Ocean. Tech.*, 29, 527–537, <https://doi.org/10.1175/JTECH-D-11-00002.1>, 2012.
- Yang, S., Xiao, W., Zhang, M., Guo, S., Zhao, J., and Shen, F.: Image Data Augmentation for Deep Learning: A Survey, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.2204.08610>, 19 April 2022.
- Ye, L., Cao, Z., and Xiao, Y.: DeepCloud: Ground-Based Cloud Image Categorization Using Deep Convolutional Features, *IEEE T. Geosci. Remote*, 55, 5729–5740, <https://doi.org/10.1109/TGRS.2017.2712809>, 2017.
- Yoshimura, M. and Yamashita, M.: Contribution of ground-based cloud observation to satellite-based cloud discrimination, *J. Environ. Sci. Eng. A*, 2, 379–382, 2013.
- Zhang, Z., Yang, S., Liu, S., Cao, X., and Durrani, T. S.: Ground-Based Remote Sensing Cloud Detection Using Dual Pyramid Network and Encoder–Decoder Constraint, *IEEE T. Geosci. Remote*, 60, 1–10, <https://doi.org/10.1109/TGRS.2022.3163917>, 2022.
- Zhao, C., Zhang, X., Luo, H., Zhong, S., Tang, L., Peng, J., and Fan, J.: Detail-Aware Multiscale Context Fusion Network for Cloud Detection, *IEEE Geosci. Remote S.*, 19, 1–5, <https://doi.org/10.1109/LGRS.2022.3207426>, 2022.
- Zhou, Q., Zhang, Y., Li, B., Li, L., Feng, J., Jia, S., Lv, S., Tao, F., and Guo, J.: Cloud-base and cloud-top heights determined from a ground-based cloud radar in Beijing, China, *Atmos. Environ.*, 201, 381–390, 2019.
- Zhou, Z., Zhang, F., Xiao, H., Wang, F., Hong, X., Wu, K., and Zhang, J.: A novel ground-based cloud image segmentation method by using deep transfer learning, *IEEE Geosci. Remote S.*, 19, 1–5, 2021.