drones

MDPI

Article

UAV-Based Snow Cover Detection Using Deep Learning and Geospatial Visualization

Abdelouahed Eddaoudy [1], Abdelaziz Elbaghdadi [1], Yassine Yazid [1,2,*], Ahmed El Oualkadi [1] and Abderrahmane Ed-daoudy [3]

[1] Innovative Systems Engineering Laboratory, National School of Applied Sciences of Tetouan, Abdelmalek Essaadi University, Tetouan 93000, Morocco; eddaoudy.abdelouahed@etu.uae.ac.ma (A.E.); abdelaziz.elbaghdadi@etu.uae.ac.ma (A.E.); aeloualkadi@uae.ac.ma (A.E.O.)
[2] Department of Automation, Electrical Engineering and Electronic Technology, Universidad Politécnica de Cartagena, 30202 Cartagena, Spain
[3] ESIM, Polydisciplinary Faculty of Sidi Bennour, Chouaib Doukkali University, El Jadida 24000, Morocco; ed-daoudy.abderrahmane@ucd.ac.ma
* Correspondence: y.yazid@uae.ac.ma

Highlights

What are the main findings?

- A UAV-based framework combining lightweight deep learning and geospatial analysis enables accurate real-time snow cover detection in mountainous environments.
- EfficientNet-B0 provides the best balance between classification performance and edge-computing efficiency, outperforming larger EfficientNet variants for onboard deployment.

What are the implications of the main findings?

- The proposed approach demonstrates the feasibility of deploying AI-powered snow monitoring directly on UAVs operating under Size, Weight and Power (SWaP) constraints.
- Integrating AI-derived snow detection with GIS visualization can support informed decision-making for disaster response, road maintenance, and winter hazard management.

Abstract

Snow disaster management represents a critical domain that requires accurate, real-time monitoring of snow-covered areas. In this regard, the detection of snow-covered surfaces from Unmanned Aerial Vehicles (UAVs) constitutes a key component of this task. Deep learning has demonstrated strong performance in aerial image classification; however, the deployment of these models on drone hardware remains constrained by several factors, including Size, Weight, and Power (SWaP) limitations. This paper proposes a robust end-to-end framework for UAV-based snow cover detection using geotagged drone imagery acquired over the Jbel Saghro mountain range in Morocco. The study compares eight EfficientNet architectures (B0 through B7) and MobileNetV2 against a Support Vector Machine (SVM) baseline. The models were evaluated using standard classification metrics and hardware-aware indicators, in particular, parameter count (in millions), computational complexity (GFLOPs), inference latency (milliseconds), and processing speed (frames per second, FPS), to simulate real-world UAV edge-computing constraints. The optimal model was subsequently embedded into a geospatial visualization pipeline using ArcGIS 10.8.2 within Power BI to generate interactive heatmaps that superimpose the spatial density of snow-covered observations onto road networks and topographic features. Experimental results on a dataset of 9307 annotated UAV images demonstrate that EfficientNet-B0 achieves

Check for updates

Academic Editors: Tommaso Orusa, Annalisa Viani and Filippo Sarvia

Received: 22 May 2026
Revised: 21 July 2026
Accepted: 28 July 2026
Published: 21 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

an accuracy of 99.21% and an AUC of 0.99 while processing images at over 100 FPS, making it the optimal architecture for real-time onboard drone processing. This geospatial framework provides actionable insights for disaster management, road maintenance operations, and early warning systems in mountainous regions.

Keywords: snow cover detection; UAV imagery; drones; geospatial analysis; edge computing; avalanche risk; MobileNetV2; SVM; EfficientNet; deep learning; ArcGIS; environmental monitoring; hazard detection; transfer learning

1. Introduction

Although deep learning models have demonstrated high accuracy for snow detection, their deployment on operational UAV platforms remains challenging due to the limited memory, computational power, and thermal capacity of onboard processors. Bridging the gap between algorithmic performance and practical embedded implementation continues to represent an open challenge in the field. Snow is a key natural resource that influences the Earth's climate and water cycle. During winter, about 50–70% of the Saghro mountain region is covered by snow, affecting groundwater, glaciers, and human activities. At the same time, with its high reflectivity, snow cover also directly affects ecological and socio-economic systems, controlling the mechanisms that affect plant and animal ecosystems. Snow cover is one of the most consequential variables in meteorological studies. It shapes hydrological cycles, drives regional climate patterns, and sits at the root of the most destructive natural hazards in mountain regions, like the blockage of roads, avalanches, and glacial outburst floods. These three elements can cause the isolation of entire communities. For government agencies, disaster management authorities, and climate researchers, knowing where snow exists, how much there is, and how it is changing is not just an academic question but an operational necessity.

Existing approaches to snow cover monitoring remain inadequate for addressing the increasing demand for accurate, real-time, and spatially detailed snow observations. Ground-based networks consisting of weather stations and field surveys can deliver accurate measurements but only at isolated points, leaving large high-altitude areas effectively unmeasured. Satellite remote sensing has expanded this view considerably, enabling global-scale estimates of Snow Water Equivalent (SWE) and Snow Cover Area (SCA). But satellites introduce their own frustrations: they only pass over a given location every few days or weeks, lose precision at finer spatial scales, and are usually blocked by cloud cover.

Accurate detection of snow-covered surfaces is essential for environmental monitoring and climate research, as it provides critical information for understanding snow dynamics, evaluating the impacts of climate change, and assessing hazards associated with avalanches and glacier retreat. Furthermore, snow-cover monitoring supports numerous applications, including hydrological modeling, climate change assessment, water resource management, and disaster risk reduction [1] and disaster management. Traditional snow monitoring practices rely primarily on ground-based surveys and satellite imagery. However, these approaches often face limitations in terms of spatial resolution, accuracy, and scalability, particularly in remote and mountainous regions. Recent advances in unmanned aerial vehicle (UAV) technology have significantly enhanced remote sensing capabilities by enabling the acquisition of high-resolution imagery and efficient data collection over large areas. Equipped with cameras and advanced sensors, UAVs have been successfully applied to a wide range of environmental monitoring tasks, including vegetation classification, glacier monitoring, and land surface classification.

Equipped with high-resolution cameras and advanced sensors, UAVs have been successfully applied to a wide range of environmental monitoring tasks, including vegetation classification, glacier monitoring, and land surface classification. In this context, machine learning-based image classification has emerged as an effective approach for analyzing UAV-acquired imagery and accurately identifying and characterizing snow-covered regions [2]. Several studies have demonstrated the potential of UAVs [3]. Snow and clouds frequently coexist, and in visible wavelengths they produce nearly identical spectral signatures—which means the very conditions that make snow monitoring most urgent are the same ones that interfere with the sensors. Indices used for these purposes, like the Normalized Difference Snow Index (NDSI), become unreliable when needed most [4].

Unmanned Aerial Vehicles (UAVs) have emerged as a promising solution to these limitations. Operating below the cloud layer, UAVs overcome the occlusion challenges that often affect satellite-based observations. They can be deployed on demand and acquire imagery at a very high resolution, enabling detailed and timely monitoring of snow-covered areas. When combined with modern deep learning techniques, UAV imagery facilitates the automated, near-real-time classification of snow-covered surfaces with high accuracy, surpassing the spatial detail achievable using conventional ground-based and satellite monitoring systems [5]. Despite the significant progress achieved by existing studies and methodologies, a critical gap remains. Models that perform well under controlled research conditions do not necessarily generalize to real-world operational environments, where variations in illumination, terrain, weather conditions, and sensor characteristics can substantially affect their performance.

Although deep learning models have achieved remarkable performance on benchmark datasets, they are often computationally intensive and designed for execution on high-performance computing platforms. As a baseline classifier, this study employs Support Vector Machine (SVM), a well-established supervised learning algorithm that has been widely used in pattern recognition and classification tasks due to its robustness and effectiveness. To address the limitations of SVM in complex real-world applications, several studies have proposed hybrid and comparative approaches. For example, SVM has been combined with the K-Nearest Neighbors (KNN) algorithm to improve soldering defect detection and enhance quality inspection performance [6].

The study reported in [7] realized a mechanism for detection of ground truth information before snow depth retrieval, i.e., classification of the snow-free state and the snow-covered state. Machine learning was introduced to achieve this purpose, and the SNR arc was used as the input data. Compared with the current mainstream topography correction algorithms, the algorithm proposed in this study does not rely on any prior ground of measured data and has theoretical universality. The detection results can constrain the retrieval snow depth using data from the snow-free state and thus, improve the retrieval accuracy.

The main contributions of this study are summarized as follows. First, a hardware-aware comparative evaluation of nine lightweight classification models is conducted. Specifically, MobileNetV2 and EfficientNet variants (B0–B7) are systematically compared against a Support Vector Machine (SVM) baseline. The models are evaluated using standard classification metrics, as well as computational performance indicators relevant to resource-constrained UAV edge deployment. Throughput (FPS), parameter count, and floating-point operations (GFLOPs) were used as first-class evaluation criteria.

Second, a large-scale UAV dataset from Morocco was created. This work introduces a new annotated dataset of 9307 UAV images captured over the Jbel Saghro range in the Ouarzazate region of Morocco, collected under actual winter conditions, with natural variations in lighting and terrain. Beyond its scale, the dataset addresses a meaningful gap

in the literature; African mountain environments are substantially underrepresented in existing snow-detection benchmarks, and this work begins to correct that. Third, end-to-end geospatial integration was employed. This study demonstrates a complete pipeline that carries model output classification into operational use. Inference results are represented directly within an interactive ArcGIS dashboard embedded within Power BI, where snow accumulation is visualized as heatmaps layered over road networks. Finally, as to UAV deployment recommendations, drawing on the full comparative evaluation, this paper provides a concrete and evidence-based recommendation for the deep learning architecture suited to real-time onboard snow detection, one that balances classification accuracy against the hardware constraints of operational drone platforms, giving practitioners a clear, justified starting point rather than leaving that decision to heuristic estimation.

The study reported in [7] realized a mechanism for detection of ground truth information before snow depth retrieval, i.e., classification of the snow-free state and the snow-covered state. Machine learning was introduced to achieve this purpose, and the SNR arc was used as the input data. Compared with the current mainstream topography correction algorithms, the algorithm proposed in this study does not rely on any prior ground of measured data and has theoretical universality. The detection results can constrain the retrieval snow depth using data from the snow-free state and thus, improve the retrieval accuracy. The EfficientNet family utilized in [8] was developed using Neural Architecture Search (NAS) to identify an optimal baseline architecture, referred to as EfficientNet-B0. The remaining variants (B1–B7) were then derived by applying a compound scaling strategy controlled by a scaling coefficient ϕ. This is unlike conventional approaches that scale network depth, width, and input resolution independently.

This paper makes the following specific contributions to the field. Firstly, on hardware-aware comparative evaluation, a systematic comparison of nine lightweight CNN architectures, MobileNetV2 and EfficientNet-B0 through B7, alongside an SVM baseline are presented. The models are evaluated using standard classification metrics and the computational realities of UAV edge deployment metrics. Throughput (FPS), parameter count, and floating-point operations (GFLOPs) were utilized as first-class evaluation criteria. Secondly, a large-scale UAV dataset from Morocco was created. This work introduces a new annotated dataset of 9307 UAV images captured over the Jbel Saghro range in the Ouarzazate region of Morocco, collected under actual winter conditions, with natural variations in lighting and terrain. Beyond its scale, the dataset addresses a meaningful gap in the literature; African mountain environments are substantially underrepresented in existing snow-detection benchmarks, and this work begins to correct that. Thirdly, end-to-end geospatial integration was employed. This study demonstrates a complete pipeline that carries model output classification into operational use. Inference results are represented directly within an interactive ArcGIS dashboard embedded within Power BI, where snow accumulation is visualized as heatmaps layered over road networks. Finally, as to UAV deployment recommendations, drawing on the full comparative evaluation, this paper provides a concrete and evidence-based recommendation for the deep learning architecture suited to real-time onboard snow detection, one that balances classification accuracy against the hardware constraints of operational drone platforms, giving practitioners a clear, justified starting point rather than leaving that decision to heuristic estimation.

The remainder of this paper is organized as follows. We start with a survey of the relevant literature. Section 2 describes the study area, dataset, tools, and details of the proposed methodology. The presentation and discussion of the experimental results, covering comparative model analysis and geospatial visualizations, are in Section 3. Finally, Section 4 draws conclusions and identifies directions for future work.

Related Work

In recent research, various methods have been employed for snow detection, including the NDSI and band ratio techniques. These methods are widely utilized to analyze visual data related to snow processes. The conventional approach to estimating environmental snow cover primarily involves direct physical observations of snow cover, often monitored through networks of ground-based meteorological stations. The literature on snow provides substantial information on seasonal snow cover, though it mostly lacks spatial resolution. During the 19th and 20th centuries, parameters such as Snow Water Equivalent (SWE) and Snow-Covered Area (SCA) were predominantly used in field assessments [9]. Three additional higher-resolution satellite snow cover extent (SCE) products are used to compare the SCE values derived from these SWE products. There is an overall close agreement between these satellite SCE products and SCE generated from the SWE data used, providing confidence in our consistent SWE, snow depth, and SCE products based on gridded climate and station data. By the 20th century, global-scale measurements of snow and glacier dynamics were effectively achieved using SWE and SCA, leveraging imagery from both active and passive satellite systems. Janine Florath and Sina Keller [10] have conducted glacier monitoring based on satellite imagery. The classification results achieved reveal that all the classification approaches used are suitable for detecting different snow and ice classes on the glacier surface. The multi-temporal analysis clearly reveals the seasonal development of glaciers. This research has considered multispectral Landsat images as input data used to delineate the maps of glaciers. Glacier mass balance was determined using the Accumulation Area Ratio (AAR) and band ratio methods. The results obtained show that glaciers have reduced over the entire study period and indicate remarkable reductions in the extent of glaciers. Kulkarni et al. [11] developed a methodology to monitor snow cover using spectral reflectance data for the Normalized Difference Snow Index (NDSI). The approach is based on the low reflectance of snow in the shortwave infrared (SWIR) region and its high reflectance in specific visible light regions.

In other related areas of research, scientists have explored the use of SAR intensity images to detect disasters caused by recent earthquakes [12]. Therefore, SAR could be a powerful tool used to develop a universal method for gaining a grasp on areas damaged by disasters such as earthquakes, forest fires and floods. Another study has studied how high-resolution SAR satellite data collected after an earthquake can be used to identify building damage [13]. In addition, a comprehensive review has been published summarizing different remote sensing techniques used for detecting earthquake-related building damage [14], as divided into two categories: multi-temporal techniques that evaluate the changes between the pre- and post-event data, and mono-temporal techniques that interpret only the post-event data. Researchers have also explored earthquake disaster detection in urban areas by analyzing curvilinear features [15]. In addition, an Earthquake Damage Visualization (EDV) technique has been proposed for the rapid identification of earthquake-induced damage using SAR data [16]. Similarly, a deep learning model has been developed to detect road damage after earthquakes, combining SAR imagery with field datasets [17].

A probabilistic cellular automata-based approach has been explored for predicting the affected area and volume of hot mudflow disasters [18]. The classification results for damaged and intact roads indicated that datasets (SAR and field measurements), together with a deep learning model, yielded acceptable overall accuracy (87.1%). Similarly, a two-dimensional cellular automata model has been proposed for the simulation of the spread of disasters [19]. Another study employed a probabilistic cellular automata approach to predict hot mudflow hazard zones and estimate mudflow volumes. [20]. A new approach for predicting the affected area of a Sidoarjo hot mudflow disaster by using

probabilistic cellular automata has been explored; the experimental results from this show new inundated areas that are identified as high-risk areas likely to be covered with mud. Furthermore, the proposed probabilistic cellular automata approach works well for the prediction of hot mudflow spreading areas and is much more accurate than the existing conventional methods [21]. In addition, a method for prediction and simulation based on the Cell Based Geographic Information System (GIS) as Cellular Automata (CA) has been proposed, together with the required data systems, in particular, metasearch-engine usage in a unified way. It has been confirmed that the proposed cell-based GIS as CA has flexible usage of the attribute information that is attached to the cell in concert with location information and is functional for disaster spreading simulations and predictions [22]. Meanwhile, a sensor network for landslide monitoring features a laser ranging system designed to avoid rainfall interference through time diversity has been developed, along with a satellite imagery-based system used to support landslide disaster relief efforts [23]. Experiments relating to this work were conducted to investigate some features data for landslide monitoring applications, including node posture, dynamic change of topology, landslide occurrence recognition, and node location change.

Most of the existing deep learning studies for remote sensing and snow detection are tailored and validated exclusively for powerful GPU-equipped servers, without considering the embedded processing realities of drone hardware. UAV platforms such as multirotor drones equipped with computer systems (e.g., Raspberry Pi 4, NVIDIA Jetson Nano, or Jetson Orin NX) impose strict limitations on memory size, power consumption, and inference speed. A model requiring a huge number of parameters, such as more than 60+ million, and multiple seconds of inference time per frame is entirely unusable in a real-time airborne context. This fundamental mismatch between algorithmic capability and hardware feasibility represents a critical research gap that the present work directly addresses.

Another gap emerges outside the model itself. Even where drone-based classification has been studied, the step that follows—translating those outputs into a spatial representation and an operationally useful system—has received less attention. Knowing that a part of terrain is snow-covered is only the beginning. What disaster managers, road authorities, and emergency responders need is this type of information layered onto road networks and elevation models, navigable in real time, and integrated into the kind of geospatial interface they already use. That end-to-end pipeline, from classified image to actionable map, remains largely unexplored in the literature.

2. Materials and Methods

This section outlines the materials, datasets, and methods used in the study. It describes the data preparation process, the proposed approach, and the experimental setup, followed by the evaluation metrics used to assess performance.

2.1. Study Area

The study area encompasses the Jbel Saghro mountain range, situated in the Anti-Atlas region of southern Morocco, east of Ouarzazate (Figure 1). The UAV survey was conducted within a rectangular area bounded by latitudes 31.10870° N and 31.15224° N and longitudes 5.67045° W and 5.61862° W, covering approximately 24 km^2 in the Jbel Saghro region, with its highest point—Amalou n'Mansour—reaching 2712 m above sea level. Despite the predominantly arid character of the surrounding region, with annual rainfall typically between 100 and 300 mm, the upper elevations experience intense and episodic snowfall during the winter months of December through February.

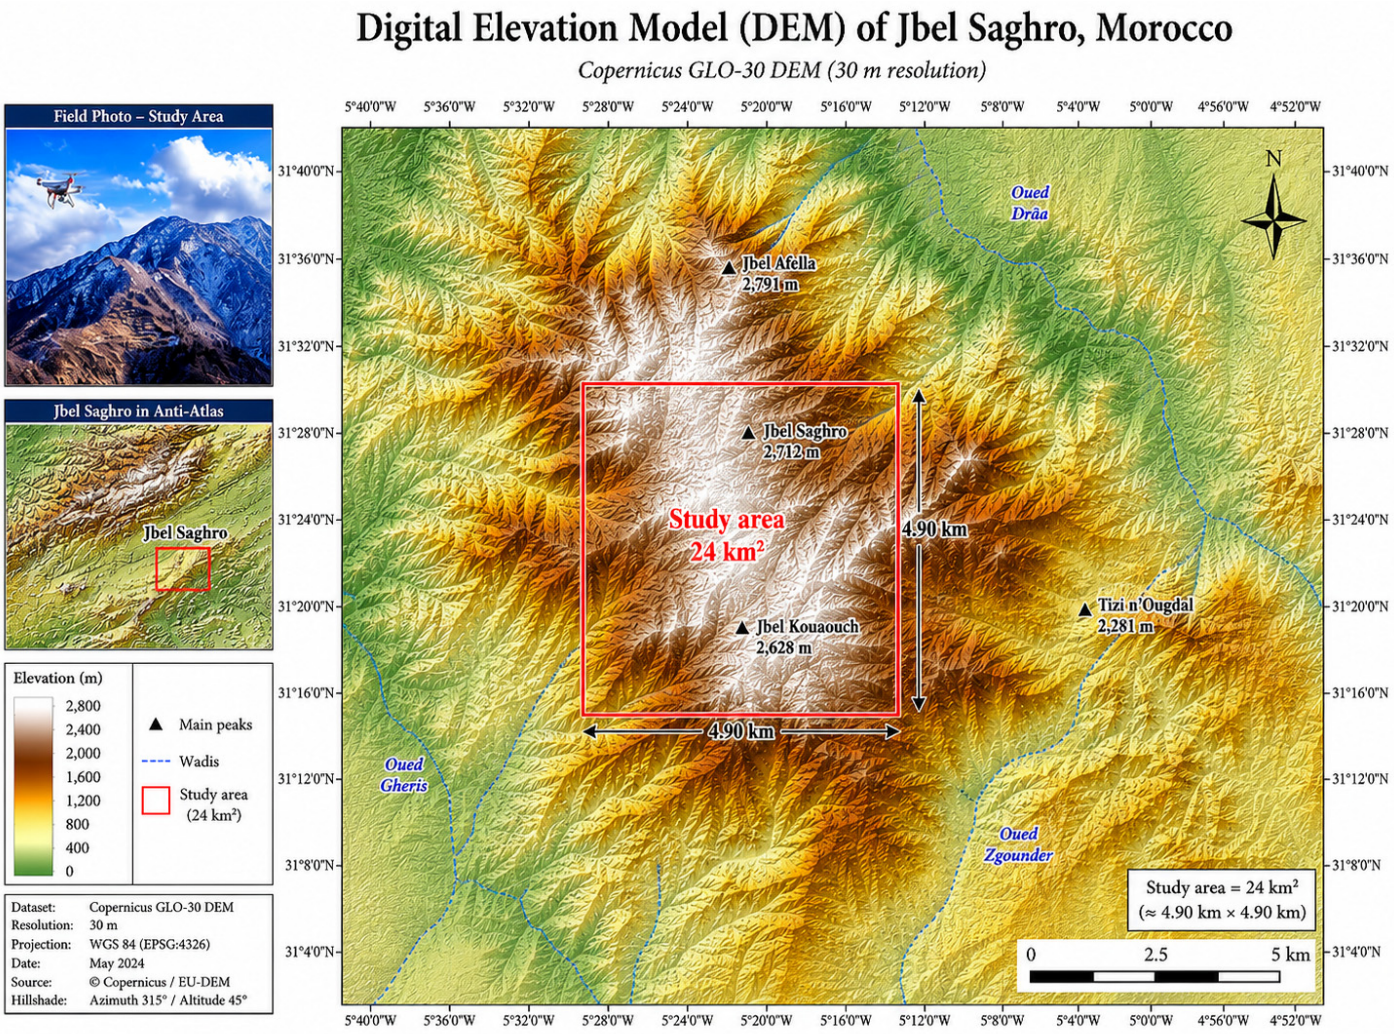


Figure 1. Digital elevation model (DEM) of the Jbel Saghro study area.

These snowfall events regularly trigger road closures and cut off remote communities, creating pressing demands for timely snow detection, hazard mapping, and emergency logistics coordination. The combination of rugged topography, limited road access, and highly variable snow accumulation makes Jbel Saghro both an ecologically meaningful and an operationally demanding environment in which to validate the proposed framework.

The main panel in Figure 1 presents a digital elevation model (DEM), derived from the Copernicus GLO-30 dataset (30 m spatial resolution), illustrating the complex mountainous topography of the study region. The red rectangle delineates the UAV survey area (approximately 24 km^2, 4.90 × 4.90 km), where aerial imagery was acquired for snow-cover detection. Major peaks and principal wadis are indicated to provide geomorphological context. The upper-left inset shows a representative field photograph acquired during the UAV campaign, while the middle inset locates the study area within the Jbel Saghro massif. Geographic coordinates, elevation, north arrow, and scale bar are included to facilitate spatial interpretation and reproducibility.

2.2. UAV Data Acquisition

Imagery was acquired using a consumer-grade multi-rotor UAV DJI Mini 3 Pro multi-rotor unmanned aerial vehicle (UAV) (SZ DJI Technology Co., Ltd., Shenzhen, China) fitted with a high-resolution RGB camera. The sensor is a 1/2.3-inch CMOS with a 20-megapixel resolution, a fixed 84° field-of-view lens, and a pixel pitch of 2.4 µm. All images were captured in JPEG format. Each frame was automatically geotagged via the onboard

GPS/GLONASS module, recording latitude and longitude coordinates that served as the basis of subsequent geospatial referencing.

A total of 18 flights were conducted over the winter season of November and December 2024, at altitudes ranging from 40 to 80 m above ground level. Flight conditions were, deliberately, firstly under a clear sky, secondly with partial cloud cover, and lastly under fully overcast skies, with speeds of flight of 4 and 6 m/s, to maximize the environmental and illumination diversity of the dataset and strengthen model robustness. The imagery covers a range of terrain types, including steep rocky slopes, flat plateau surfaces, valley floors, and road corridors.

2.3. Dataset Description and Preprocessing

The collected and treated dataset contains 9407 geotagged RGB images manually annotated into two classes, namely, Snow and Snow-free, based on the visible presence of snow cover within each image. Images containing clearly distinguishable snow-covered terrain were assigned to the Snow class, whereas images without visible snow were labeled Snow-free. Images with uncertain visibility caused by shadows, low illumination, or motion blur were carefully inspected before annotation. When image quality prevented reliable interpretation, the corresponding samples (100 images) were excluded from the final dataset to ensure annotation consistency. The final dataset contains 2 classes: Snow (4944 images) and Snow-Free (4363 images). The full class distribution is summarized in Table 1 below.

Table 1. Distribution of annotated image classes in the dataset.

Class	Number of Images	Percentage (%)
Snow	4944	53.12
Snow-free	4363	46.88
Total	9307	100

To reduce the risk of overfitting and strengthen generalization to new conditions, a comprehensive data augmentation pipeline was applied only to the training subset. The augmentation techniques included

- **Geometric transformations:** Horizontal and vertical flips, random rotations up to $\pm 30^{\circ}$, random horizontal and vertical shifts, and random scaling.
- **Photometric transformations:** Random brightness and contrast adjustments representing the range of solar illumination conditions encountered during the different UAV flight sessions.
- **Noise injection:** Gaussian noise added, simulating sensor noise in low-light conditions.
- **Blur effects:** Gaussian blur to simulate motion or defocus artifacts common during drone flight.

The dataset was partitioned into training (80%, 7464 images) and testing (20%, 1843 images) subsets using stratified random splitting to preserve class balance. All images were resized to 224 $\times$ 224 pixels as a standard input resolution for the CNN architecture, and pixel values were normalized using ImageNet mean and standard deviation (mean = [0.485, 0.456, 0.406]; std = [0.229, 0.224, 0.225]). Although EfficientNet variants were originally designed with different input resolutions, all models were trained using a common input size of 224 $\times$ 224 pixels to ensure a controlled comparison under identical computational conditions. This standardization isolates the effect of network architecture from the influence of input resolution, allowing differences in performance to be attributed primarily to the models themselves rather than to varying image scales.

2.4. Flight-Wise Cross-Validation for Generalization Assessment

To further evaluate the generalization capability of the proposed models and to control spatial dependency between adjacent UAV images, a flight-wise validation protocol was implemented, complementing the conventional image-level random split. The complete dataset comprises 9307 geotagged RGB images acquired during 18 independent UAV flights over the Jbel Saghro mountain range under varying environmental conditions (clear sky, partially cloudy, and overcast). As consecutive images within a flight share similar viewpoints, terrain characteristics, illumination, and snow-cover patterns, random image-level partitioning may introduce spatial autocorrelation between training and testing subsets, potentially yielding overly optimistic performance estimates.

To mitigate this bias, a Leave-One-Flight-Out Cross-Validation (LOFO-CV) strategy was adopted. In each of the 18 validation iterations, one complete flight was held out as the test set while the remaining 17 flights constituted the training set, such that every flight served exactly once as an independent test set. The preprocessing pipeline (resizing, normalization, and data augmentation) was applied exclusively to the training data in each fold, and no information from the held-out flight was used during training or hyperparameter optimization, thereby preventing information leakage. Performance across the 18 folds was reported as the mean and standard deviation of classification accuracy, precision, recall, F1-score, and area under the ROC curve (AUC).

2.5. Software and Hardware Environment

The deep learning models employed in this study were implemented in Python 3.11 using the PyTorch 2.0 deep learning framework. Model architecture was instantiated through the TIMM (PyTorch Image Models) library, leveraging ImageNet pre-trained weights for transfer learning. The THOP library was used for FLOPs calculation. SVM implementation utilized scikit-learn. To facilitate geospatial visualization, ArcGIS was integrated into Microsoft Power BI version 2.124, utilizing the dedicated ArcGIS for Power BI connector. All training models and inference benchmarking simulations were carried out on a workstation running Windows 10, equipped with an 11th Gen Intel Core i7-11850H processor (2.50 GHz), 64 GB of RAM, and an NVIDIA T1200 Laptop GPU (NVIDIA Corporation, Santa Clara, CA, USA).

2.6. Proposed Framework and Workflow

The proposed framework is organized into five sequential stages, as illustrated in Figure 2:

- **Stage 1—UAV Data Acquisition:** Drone flights are conducted over the target area, yielding geotagged RGB images with GPS coordinates—latitude and longitude—embedded in each frame.
- **Stage 2—Preprocessing and Augmentation:** Image normalization, resizing, and augmentation, as described in Section 2.3. GPS metadata is simultaneously extracted from each file and stored coupled with the corresponding image filename in a CSV structured file.
- **Stage 3—Comparative Model Training and Evaluation:** Nine CNN architectures, MobileNetV2, EfficientNet-B0 through B7, and an SVM baseline, are trained and evaluated. Each CNN is initialized with ImageNet pretrained weights and extended with a custom classification head. Evaluation covers both standard classification metrics and hardware-aware performance indicators relevant to UAV edge deployment.
- **Stage 4—Optimal Model Selection and Inference:** The optimal architecture is selected based on the combined evaluation results. The chosen model is then applied to

the complete geotagged image dataset, producing a class label for every captured geographic location

- **Stage 5—Geospatial Visualization:** Inference output class labels, paired with their corresponding GPS coordinates, are merged into a unified geospatial dataset and imported into ArcGIS within Power BI, producing a spatial distribution map, snow accumulation heatmaps, and road-overlay visualizations.

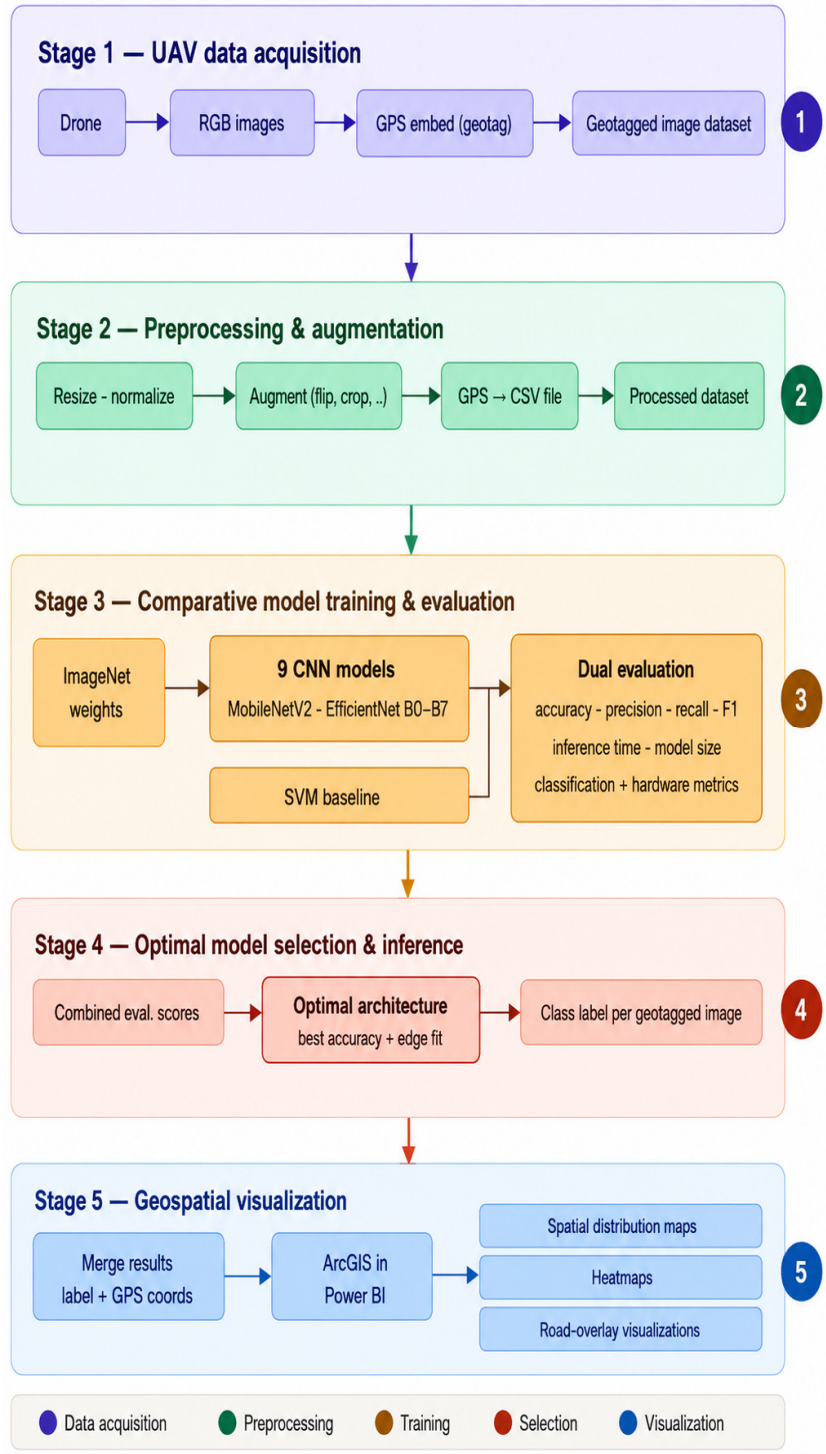


Figure 2. Proposed framework and workflow schema.

2.7. Classification Models

This section presents and explains the models used in the study, detailing their architecture, key parameters and performance characteristics.

2.7.1. Support Vector Machine (SVM)

The Support Vector Machine (SVM) is a supervised learning algorithm designed according to the optimal decision hyperplane, maximizing the margins between classes in a high-dimensional feature space. For non-linearly separable data, SVM applies kernel functions to map input features into higher-dimensional spaces. In this study, a radial basis function (RBF) kernel was employed. For SVM-based image classification, input images were resized to 64 × 64 pixels, converted to grayscale and flattened into one-dimensional feature vectors. Model performance was evaluated using 5-fold cross-validation to ensure statistically robust accuracy estimates.

2.7.2. MobileNetV2

MobileNetV2 is a lightweight CNN architecture developed for mobile and embedded applications. Its core innovation lies in the inverted residual block with linear bottlenecks, which employs depthwise separable convolutions to significantly reduce the number of multiply–accumulate operations compared to standard convolutions, while maintaining representational capacity. In the work for the present study, MobileNetV2 (width multiplier $\alpha = 1.0$) was initialized using ImageNet pre-trained weights, while the original classification head was replaced with a custom head composed as presented in the schema in Figure 3:

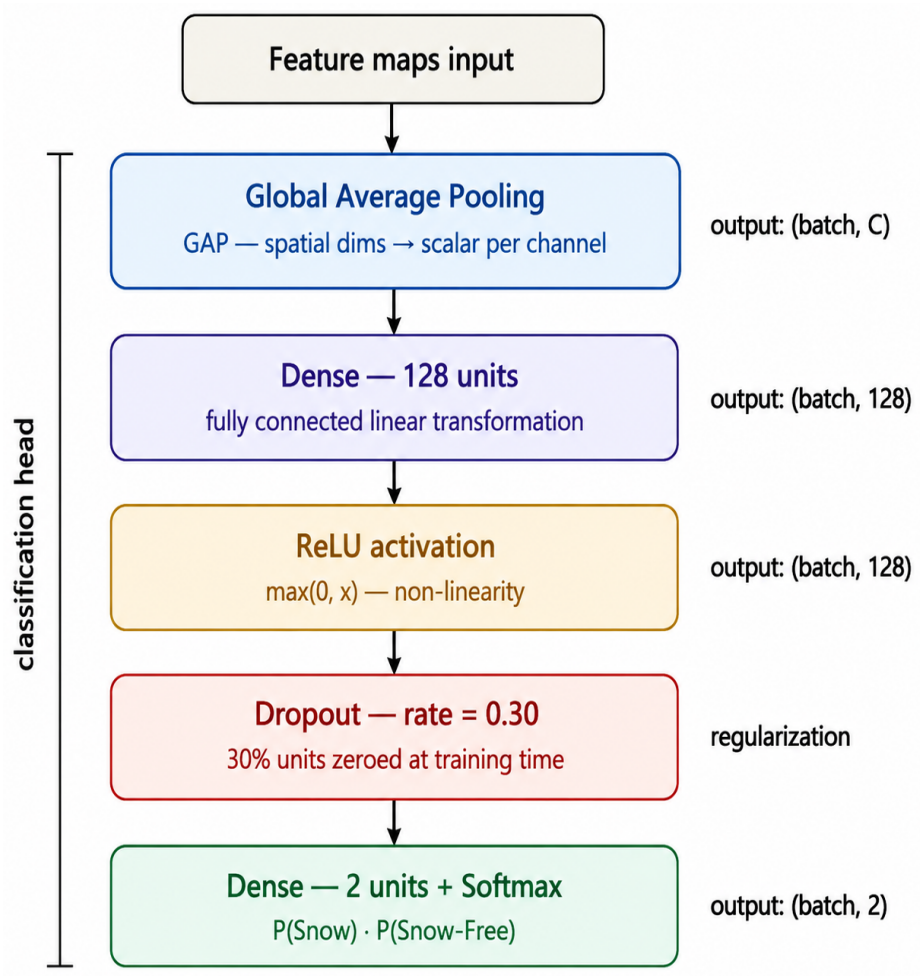


Figure 3. Classification head MobileNetV2.

The model was fine-tuned end-to-end for 100 epochs using the Adam optimizer with learning rate = 0.001 and categorical cross-entropy loss.

2.7.3. EfficientNet-B0 Through B7

The EfficientNet family was developed using a Neural Architecture Search (NAS) process to identify the optimal baseline architecture (B0), which was subsequently scaled using a compound scaling coefficient φ. Rather than scaling depth or resolution independently, EfficientNet increases all three at the same time according to the following relationships:

- **Depth:** $d = \alpha^{\wedge}\varphi$
- **Width:** $w = \beta^{\wedge}\varphi$
- **Resolution:** $r = \gamma^{\wedge}\varphi$

This is subject to the constraint $\alpha \cdot \beta^2 \cdot \gamma^2 \approx 2$, with $\alpha \geq 1$, $\beta \geq 1$, $\gamma \geq 1$.

This compound scaling strategy ensures that as the architecture grows from B0 to B7, each increment in computational budget is distributed optimally.

The result is a consistently favorable accuracy-to-FLOPs ratio compared to conventional single-axis scaling approaches. The key architectural parameters for each variant are summarized in Table 2.

Table 2. Architectural parameters of the EfficientNet-B0–B7 family.

Model	Input Resolution	Depth Coefficient	with Coefficient	Param (M)	GLOPs
EfficientNet-B0	224 × 224	1.0	1.0	4.01	0.39
EfficientNet-B1	240 × 240	1.1	1.0	6.51	0.70
EfficientNet-B2	260 × 260	1.2	1.1	7.70	1.00
EfficientNet-B3	300 × 300	1.4	1.2	10.69	1.80
EfficientNet-B4	380 × 380	1.8	1.4	17.55	4.20
EfficientNet-B5	456 × 456	2.2	1.6	28.35	9.90
EfficientNet-B6	528 × 528	2.6	1.8	40.74	19.00
EfficientNet-B7	600 × 600	3.1	2.0	63.78	37.00

Each EfficientNet was initialized with ImageNet pretrained weights via the TIMM library and fine-tuned using the same configuration applied to MobileNetV2 with Adam optimizer, at a learning rate of 0.001, 100 epochs, and cross-entropy loss, following the native recommended settings for each model. The same classification head structure was retained across all variants to ensure comparability (Figure 4).

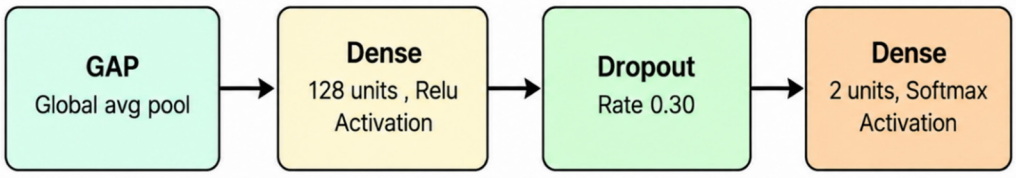


Figure 4. Classification head architecture.

2.8. Evaluation Metrics

The proposed approaches were evaluated using two categories of evaluation metrics. The first category is related to model evaluation, namely, accuracy, precision, recall, and F1-score. The second category consists of metrics used to evaluate hardware-aware drone development.

Accuracy measures the total accuracy of the model's prediction, and the value is calculated using Equation (1).

$$Acc = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

Precision is a measure that evaluates the ratio of positive predictions, and it is given by Equation (2):

$$P = \frac{TP}{TP + FP} \tag{2}$$

Recall is the metric that measures the ratio of positive instances correctly detected by the model and is calculated by Equation (3).

$$R = \frac{TP}{TP + FN} \tag{3}$$

The F1-score combines precision and recall through their harmonic meaning and is determined by Formula (4):

$$F1 = \frac{2 \times P \times R}{P + R} \tag{4}$$

Additionally, the Area Under the Receiver Operating Characteristic Curve (AUC-ROC) was computed as a threshold-independent measure for each discriminative capability of the model. A confusion matrix was generated for all models to provide a detailed breakdown of classification errors at the class level.

To evaluate the performance in the practical cases for each model using onboard UAV processing, the following hardware metrics were measured:

- **Parameter Count (M):** Total number of trainable parameters (millions), directly governs model file size and the RAM footprint on drone hardware.
- **Computational Complexity (GFLOPs):** Giga Floating Point Operations per forward pass, calculated using the THOP library. Determines the computational demand on drone CPUs/GPUs.
- **Inference Latency (ms):** Average inference time (milliseconds) for a single image forward pass measured over 100 repeated inferences after a 10-iteration.
- **Processing Speed (FPS):** Frames per second metric were calculated as 1000/Latency (ms). A minimum of 25 FPS is required for near-real-time drone monitoring applications.

3. Results and Discussion

The foregoing review of the principal frameworks and algorithms for snow detection highlights the continuous evolution of the field, reflecting sustained progress in model architectures and detection methodologies. This chapter is dedicated to comparing the results of the frameworks and algorithms analyzed in the previous chapters so that the reader can have a clear view of the best-performing options and the progress of the field throughout.

3.1. Training Dynamics

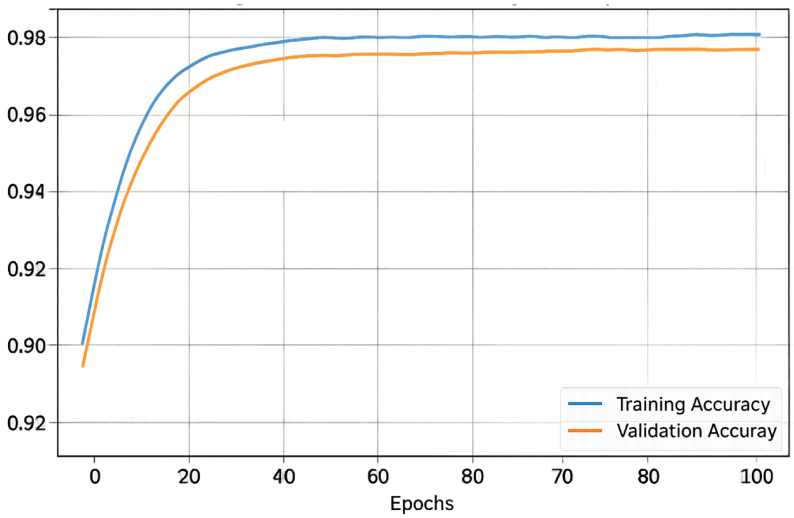
The CNN models were trained for 100 epochs. Training and validation loss and accuracy curves are illustrated in Figure 5, respectively, with MobileNetV2 serving as the primary deep learning baseline.

Both training and validation accuracy increased over the first three epochs, converging to approximately 99% by epoch seven. The gap between the two curves remained consistently below 0.5%, pointing to strong generalization with minimal overfitting, an outcome attributable to the combined effects of dropout regularization and the augmentation pipeline. Training loss decreased steadily from approximately 0.29 to 0.02, while validation loss stabilized around 0.03, confirming clean and stable convergence.

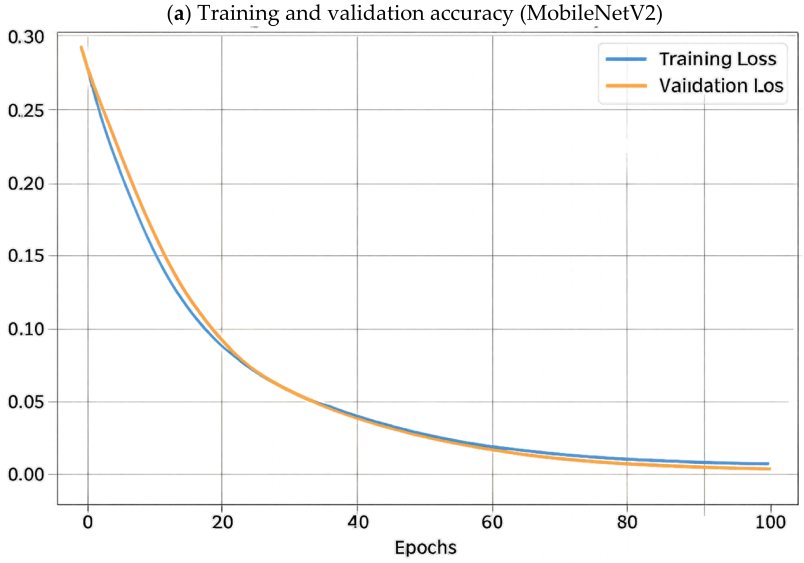
3.2. SVM Performance

Using 5-fold cross-validation with a radial basis function (RBF) kernel SVM applied to flattened grayscale feature vectors, the model reached a validation accuracy for each

fold, specifically, 98.2%, 97.2%, 98.3%, 98.7%, and 98.2%, resulting in a mean accuracy of 98.13% ± 0.54% (Figure 6). While this represents a competitive result, SVM operates solely on low-level pixel intensity features and lacks the capacity to capture the hierarchical spatial patterns, texture gradients, surface structures and shadow boundaries.



(**a**) Training and validation accuracy (MobileNetV2)



(**b**) Training and validation loss (MobileNetV2)

Figure 5. Training and validation performance of MobileNetV2: (**a**) accuracy and (**b**) loss.

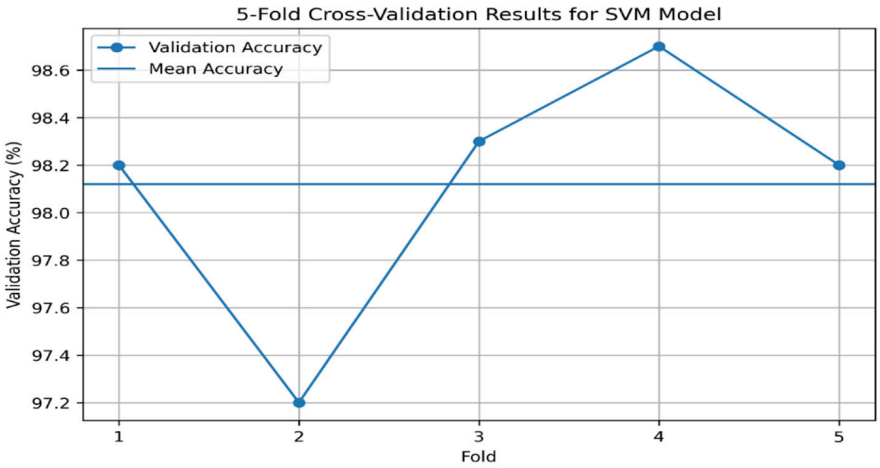


Figure 6. SVM 5-fold cross-validation accuracy.

This analysis reliably distinguishes Snow from Snow-free areas under varying lighting conditions. It therefore serves as a strong classical baseline, but every CNN architecture in the comparison ultimately surpassed the SVM.

3.3. MobileNetV2 Performance

MobileNetV2 achieved a classification accuracy of 99.03%, demonstrating that lightweight convolutional architectures can learn highly discriminative snow features despite their relatively small number of parameters. The high recall (99.29%) indicates that the model successfully detected most snow-covered scenes while producing only a limited number of false negatives, an important characteristic for operational snow monitoring, where missed detections may lead to inaccurate hazard assessment. The confusion matrix demonstrates that from 1843 test samples, 1009 snow images and 834 snow-free images were correctly classified, with only 12 and 6 misclassifications, respectively, as shown in Figure 7. The small number of misclassified images mainly corresponds to transition zones between snow-covered and snow-free surfaces, partially shadowed terrain and thin snow layers exposing underlying rocks. Such situations generate spectral ambiguity even for deep CNNs and represent the most challenging conditions for UAV-based snow recognition.

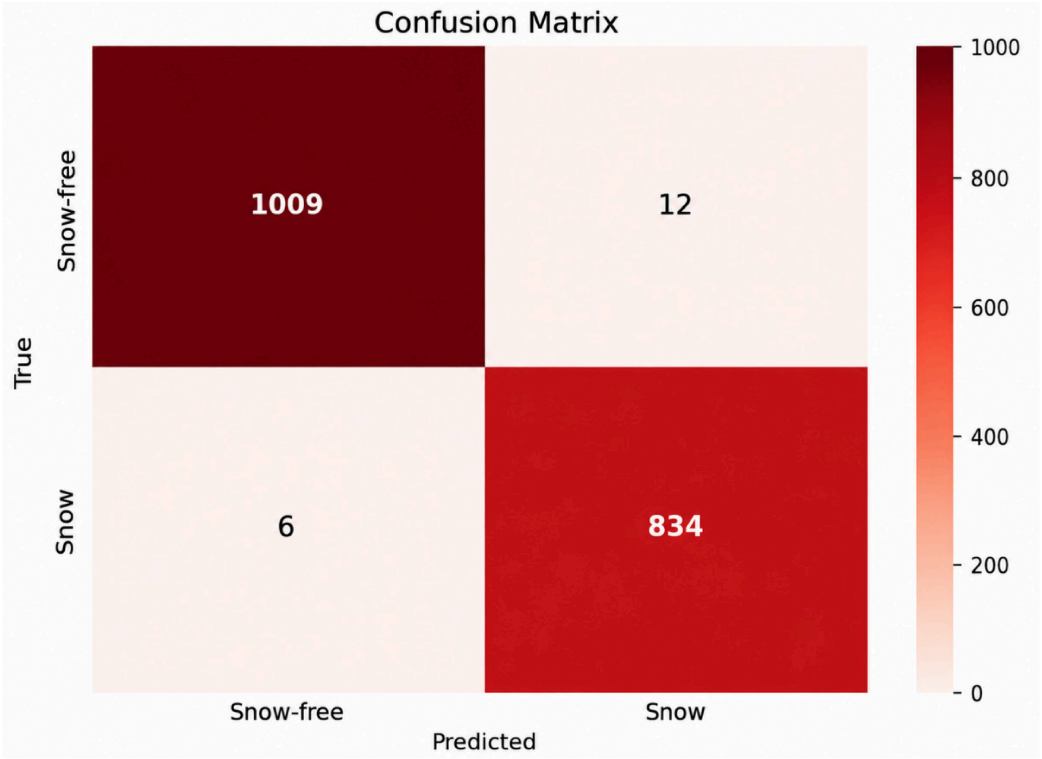


Figure 7. Confusion matrix.

The ROC curve in Figure 8 remains very close to the top-left corner of the plot, reflecting strong class separation across all decision thresholds.

These results are shown below, in Table 3. And Table 4 demonstrates the effectiveness of the MobileNetV2 architecture as a baseline CNN for this task.

Table 3. MobileNetV2 performance on training and test sets.

Dataset Split	Accuracy (%)	Loss
Training	99.48	0.01
Testing	99.03	0.02

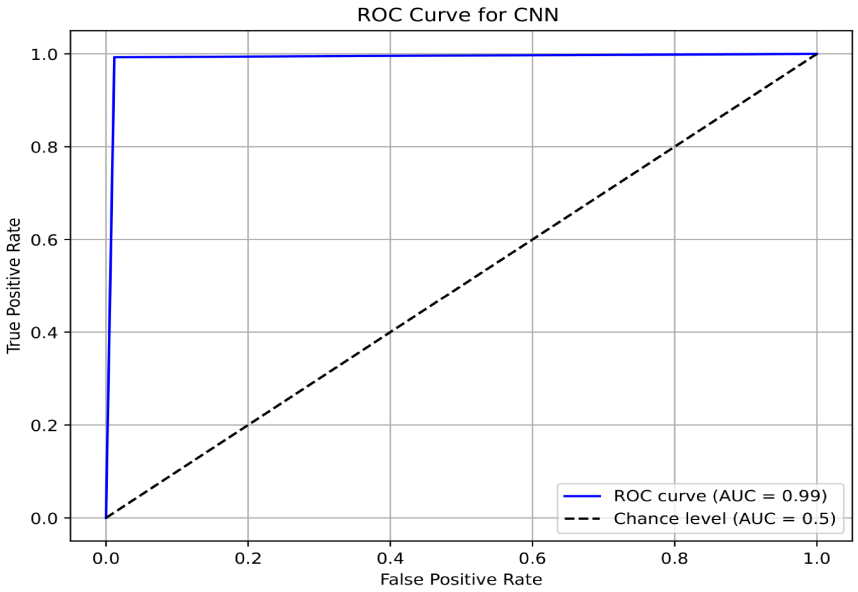


Figure 8. ROC curve.

Table 4. MobileNetV2 classification metrics on the test set.

Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC
99.03	98.58	99.29	98.93	0.99

3.4. EfficientNet-B0–B7 Performance

Table 5 summarizes classification performance and hardware-aware metrics across all tested models. The results reveal a consistent, systematic accuracy improvement from EfficientNet-B0 through B7, in line with the progressively increasing compound scaling coefficients. Accuracy rises from 99.21% at B0 to 99.85% at B7, a gain of 0.64 percentage points, but that improvement comes at a steep computational cost: a 16-fold increase in parameter count and a 95-fold increase in GFLOPs.

Table 5. Comprehensive comparative performance of all evaluated models: classification metrics and UAV hardware constraints for simple split train/test.

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC	Params(M)	GFLOPs	Latency (ms)	FPS
SVM (Baseline)	98.13	97.85	97.92	97.88	0.96	—	—	—	—
MobileNetV2	99.03	98.58	99.29	98.93	0.99	2.22	0.30	8.3	120.5
EfficientNet-B0	99.21	98.91	99.44	99.17	0.99	4.01	0.39	9.5	105.2
EfficientNet-B1	99.35	99.04	99.51	99.27	0.99	6.51	0.70	11.7	85.4
EfficientNet-B2	99.42	99.11	99.58	99.34	0.99	7.70	1.00	14.3	69.9
EfficientNet-B3	99.55	99.28	99.67	99.47	0.99	10.69	1.80	18.1	55.3
EfficientNet-B4	99.61	99.37	99.72	99.54	1.00	17.55	4.20	27.5	36.4
EfficientNet-B5	99.71	99.48	99.80	99.64	1.00	28.35	9.90	47.2	21.2
EfficientNet-B6	99.79	99.57	99.86	99.71	1.00	40.74	19.00	78.4	12.8
EfficientNet-B7	99.85	99.63	99.91	99.77	1.00	63.78	37.00	82.6	12.1

3.5. Flight-Wise Cross-Validation Results

To assess robustness under realistic deployment conditions, the models were evaluated using the flight-wise cross-validation protocol (Section 2.7), with results averaging eighteen folds (Table 6). Although performance was slightly lower than with random image-level splitting—which was expected, given the novel viewpoints, illumination, and spatial contexts in unseen test flights—all models maintained high classification accuracy,

confirming strong generalization to new UAV missions. EfficientNet-B0 remained the best-performing architecture, indicating that its effectiveness stems from robust, transferable feature learning rather than spatial similarity between neighboring images. The narrow confidence intervals across folds further confirm the stability of the framework under varying acquisition conditions.

Table 6. Comprehensive comparative performance of all evaluated models: classification metrics and UAV hardware constraints for flight-wise cross-validation.

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC
SVM	96.84	96.32	97.15	96.73	0.978
MobileNetV2	97.68	97.45	97.82	97.63	0.991
EfficientNet-B0	98.11	97.92	98.36	98.14	0.994
EfficientNet-B1	98.27	98.05	98.51	98.28	0.995
EfficientNet-B2	98.41	98.18	98.63	98.4	0.995
EfficientNet-B3	98.56	98.33	98.76	98.54	0.996
EfficientNet-B4	98.61	98.39	98.82	98.6	0.996
EfficientNet-B5	98.67	98.45	98.87	98.66	0.997
EfficientNet-B6	98.74	98.53	98.93	98.72	0.997
EfficientNet-B7	98.8	98.61	98.97	98.79	0.998

The progressive improvement from EfficientNet-B0 to EfficientNet-B7 reflects the benefit of compound scaling, which simultaneously increases network depth and width and input resolution. These characteristics enable deeper models to capture subtle spectral and textural differences between the snow-covered surfaces, exposed rocks and terrain shadows that frequently occur in the rugged Jbel Saghro landscape.

The excellent classification performance demonstrates that the proposed framework remains highly reliable despite the complex topography of Jbel Saghro, where steep slopes, shadows cast, variable illumination and heterogeneous rock surfaces create challenging imaging conditions.

In contrast to random image-level splitting, the flight-wise partitioning strategy ensures that images acquired during the same UAV flight are assigned exclusively to either the training or the testing set. By preventing highly similar neighboring images from being shared across the two subsets, this protocol mitigates data leakage and yields a more rigorous assessment of model generalization. As a result, the reported performance more accurately reflects real-world deployment conditions and confirms the ability of the proposed framework to generalize previously unseen UAV missions.

3.6. Hardware-Aware Model Selection for UAV Deployment

The main analytical contribution of this work is its evaluation of accuracy and hardware feasibility together rather than in isolation. Figure 9 presents an accuracy versus FPS scatter plot across all tested architectures, making the trade-off between classification performance and inference speed directly visible.

A threshold of 25 FPS is adopted as the minimum acceptable processing speed for near-real-time drone monitoring, consistent with standard video frame rates and operational UAV requirements. Models that decrease below this threshold cannot sustain continuous inference during flight without introducing unacceptable delays in hazard detection and alerting.

Table 5 shows that EfficientNet-B5, -B6, and -B7 operate below the 25 FPS threshold, achieving only 21.2, 12.8, and 12.1 FPS, respectively, on the test hardware. Their accuracy gains over B4 are real, ranging from 0.10 to 0.24 percentage points, but they are not

operationally justifiable given the severity of the inference speed penalty. The memory footprint demonstrates a similar behavior: B4 through B7 all exceed 17 million parameters, surpassing the practical limits of lightweight embedded computers such as the Raspberry Pi 4 and Jetson Nano, both of which ship with 4 GB of LPDDR4 RAM.

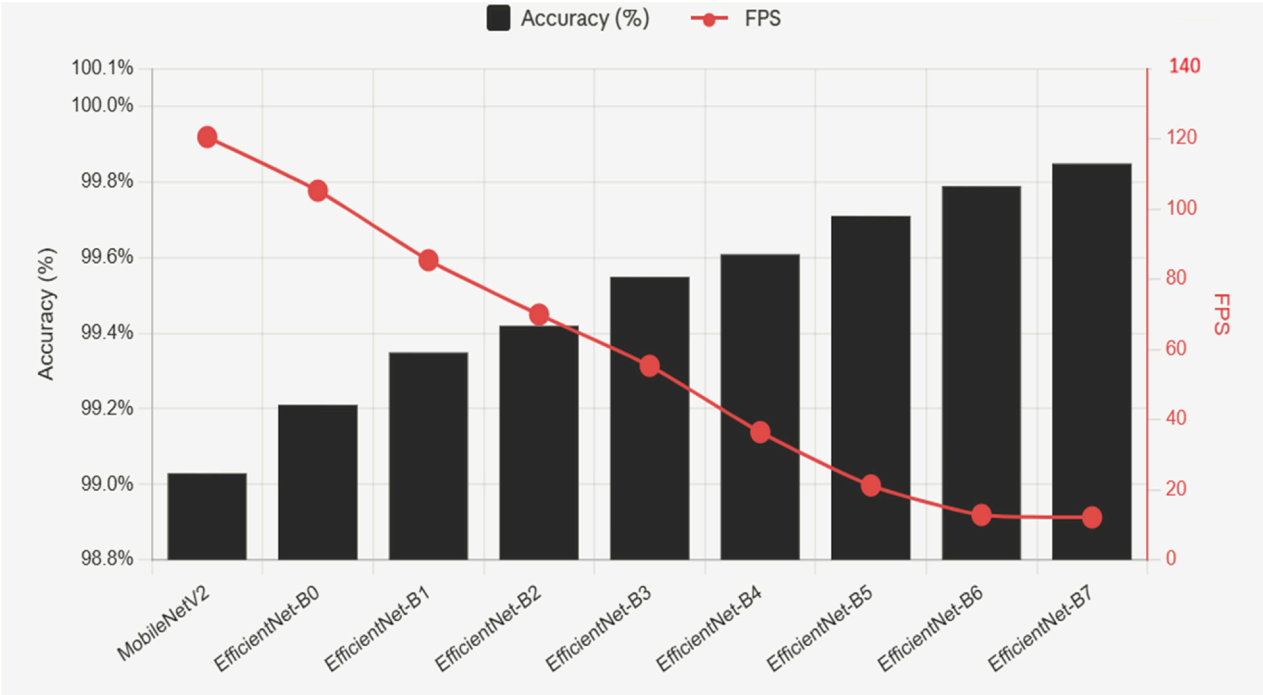


Figure 9. Accuracy (%) and frames per second (FPS) comparison across CNN architectures.

EfficientNet-B0 emerges as the optimal architecture for UAV deployment, providing 99.21% accuracy at 105.2 FPS with only 4.01 million parameters and 0.39 GFLOPs—performance metrics entirely compatible with real-time inference on current-generation drone embedded computers.

As shown in Table 7. MobileNetV2 offers a small relative loss of accuracy (99.03%) for higher throughput (120.5 FPS) and a lighter memory footprint (2.22 M parameters), making it more appropriate choice for platforms with severe resource constraints, for instance, microcontrollers or the Coral Edge TPU. For platforms with moderate computing capability, such as the Jetson Orin NX, EfficientNet-B3 represents a compelling secondary option, offering 99.55% accuracy at 55.3 FPS.

Table 7. UAV hardware deployment classification of evaluated models.

Deployment Tier	Model	Accuracy (%)	FPS	Recommended Platform
Ultra-Lightweight	MobileNetV2	99.03	120.5	Jetson Nano, RPi 4, Coral TPU
Optimal (Recommended)	EfficientNet-B0	99.21	105.2	Jetson Nano, RPi 4
Moderate	EfficientNet-B3	99.55	55.3	Jetson Orin NX
Heavy (Server-Side)	EfficientNet-B7	99.85	12.1	Ground Station GPU Server

3.7. Geospatial Visualization Using ArcGIS

To facilitate the interpretation of classification outputs into operationally actionable intelligence, the inference results from the optimal EfficientNet-B0 model were integrated into a geospatial visualization framework using ArcGIS within Microsoft Power BI (Business Intelligence).

3.7.1. Spatial Distribution Map

The classification results for all 9307 geotagged images were stored with their corresponding GPS coordinates (latitude, longitude) in a geospatial dataset. This dataset was loaded into ArcGIS to generate a thematic point map (Figure 10) in which each point represents the as-classified location of one drone-captured image.

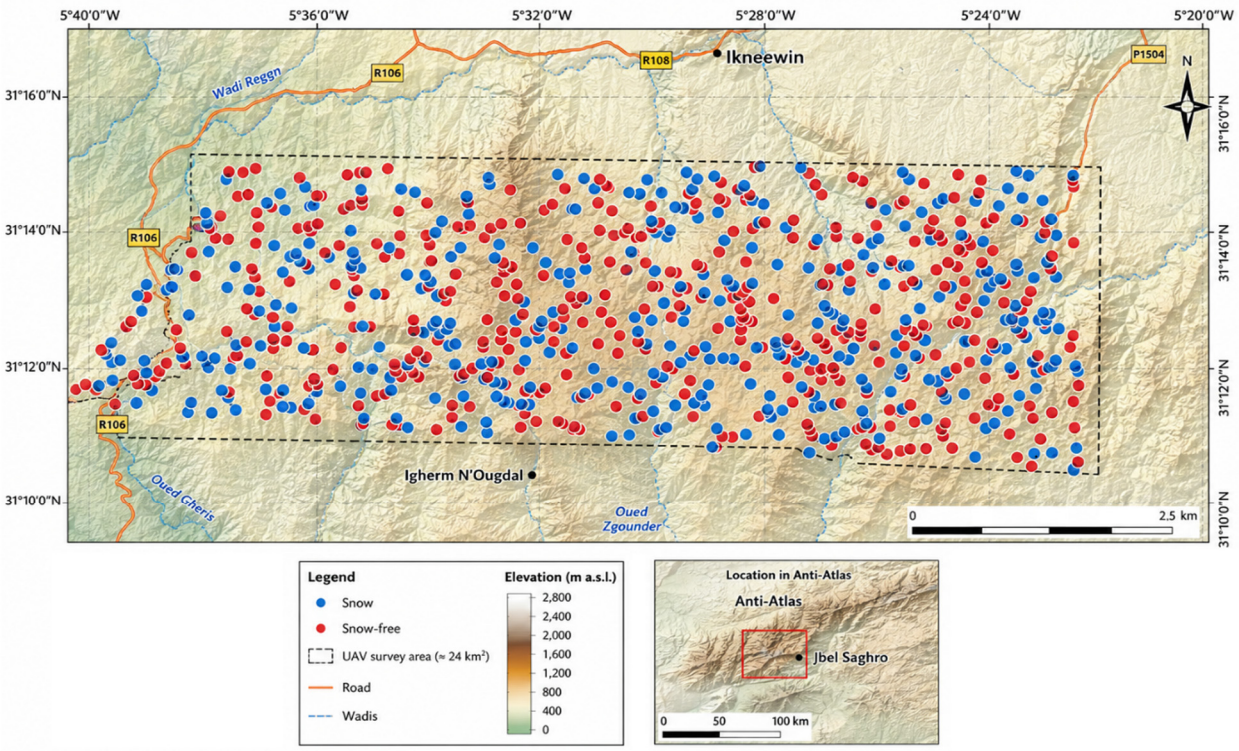


Figure 10. Distribution of classes on the ArcGIS.

A color-coded symbology was adopted: red points indicate snow-covered surfaces, and blue points indicate snow-free surfaces. The resulting spatial distribution map provides an immediate visual understanding of snow accumulation patterns across the study area at a sub-meter spatial resolution that is not achievable with satellite-based methods.

3.7.2. Density Heatmap

To convey snow accumulation density in a more aggregated and visually intuitive form, a kernel density estimation (KDE) heatmap was generated in ArcGIS and overlaid onto a digital elevation model (DEM) for the study area (Figure 11).

A color gradient ranging from light purple at low density, through red, to intense yellow at the highest concentrations was applied, with highlighting zones of peak accumulation. The heatmap clearly shows that these zones cluster on north-facing slopes and high-elevation plateau surfaces, a spatial pattern consistent with established snow accumulation dynamics in semi-arid mountain environments.

3.7.3. Road Network Overlay

To bring the framework closer to operational use for emergency response, the snow spatial concentration as shown in the snow-covered locations heatmap was overlaid onto the OpenStreetMap road network covering the study area (Figure 12). This visualization helps to identify, at the initial level, the road segments most affected by snow accumulation. Secondly, it allows for the quantification of both the extent of road closures and the severity of snow loading along specific route corridors.

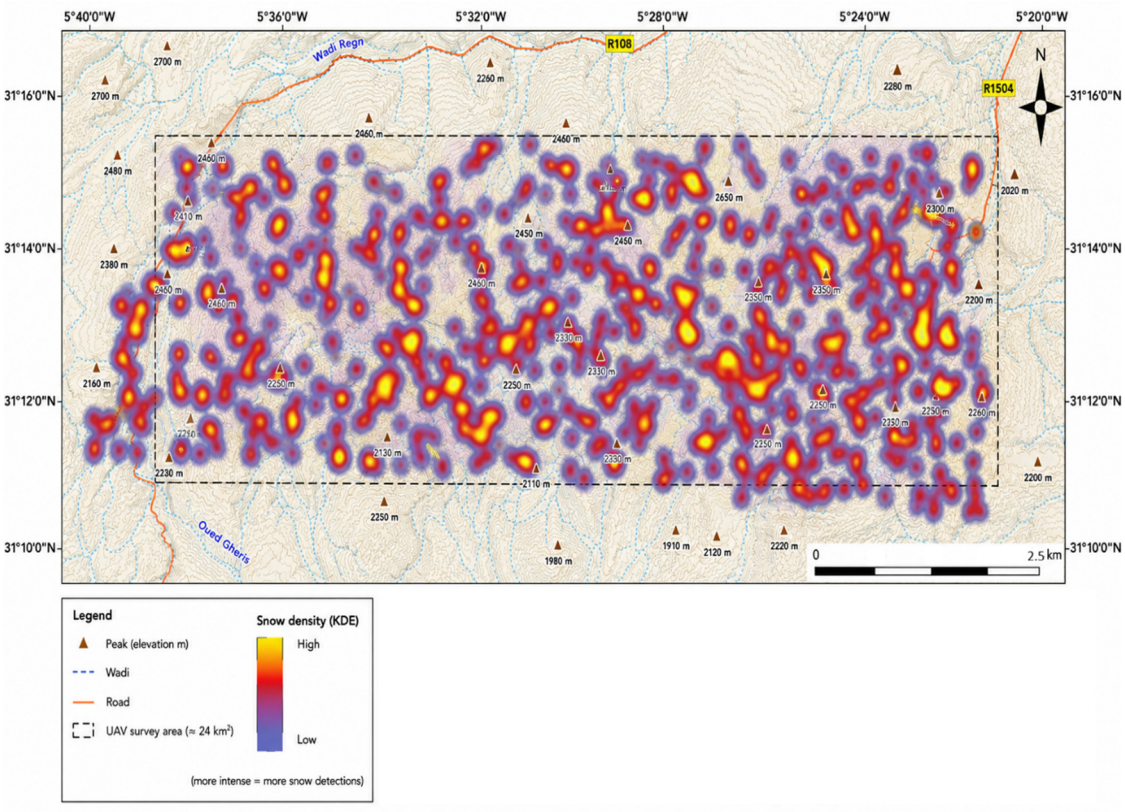


Figure 11. Heatmap representation of classes on ArcGIS.

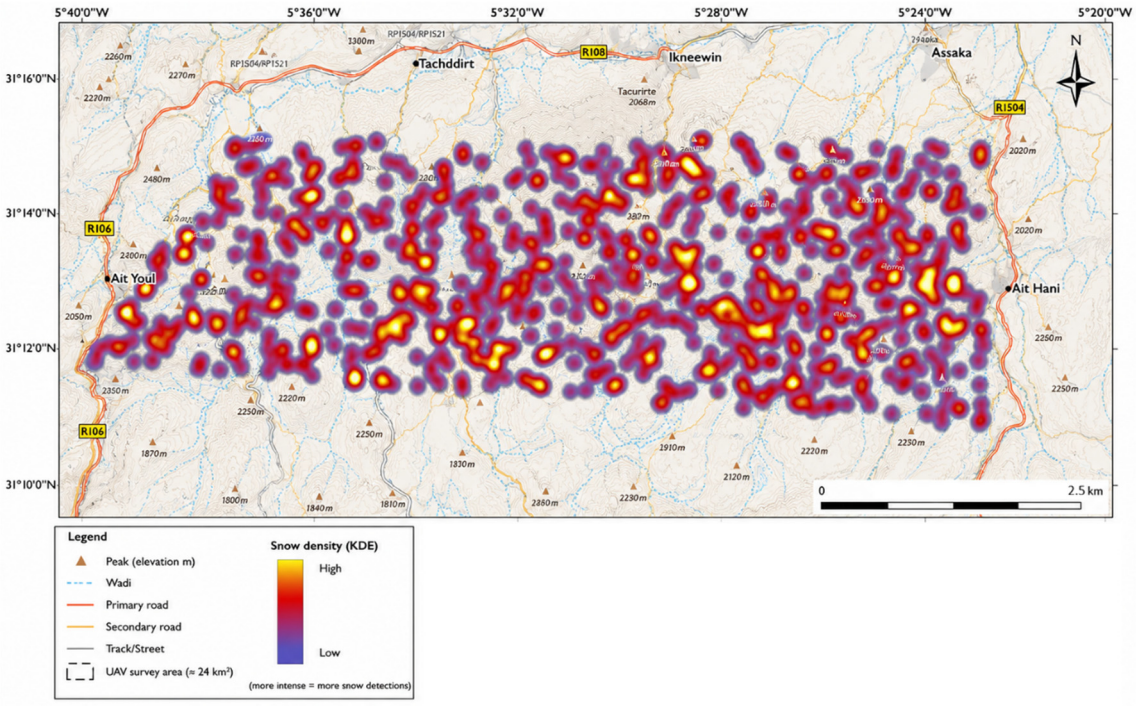


Figure 12. Heatmap visualization with roads and streets.

The emergency response authorities, road maintenance services and civil protection agencies can directly use this visualization to prioritize snow removal operations, plan evacuation routes and allocate de-icing and plowing resources to the highest-risk segments. This capability transforms the deep learning classification pipeline from a scientific instrument into a practical, operational decision-support tool.

3.8. Discussion

The findings and results of this study have several direct implications for the design and deployment of UAV-based snow monitoring systems.

All tested CNN architectures significantly outperformed the SVM baseline, which, despite achieving a respectable accuracy, with value of 98.13%, lacks the ability to capture the complex multi-scale spatial features like shadow textures, grain patterns, mixed areas and snow–rock boundaries that deep convolutional architectures naturally capture. The sustained superiority of EfficientNetB0–B7 over MobileNetV2 validates the combined scaling principle as an effective strategy for snow detection from UAV imagery, particularly under the challenging lighting conditions (overcast, partial cloud cover) that were intentionally included in the dataset.

EfficientNet achieved a marginal accuracy improvement, with B0 at 99.21% and B7 with 99.85%; this does not consider the corresponding 16-fold increase in parameter count and 95-fold increase in GFLOPs in a UAV practical deployment environment, though.

As a central result of this study, with snow detection at high spatial resolution from a low-altitude UAV, EfficientNet-B0 appears to have reached a performance limit for this binary classification task. The residual errors (0.79%) represent borderline cases that are unlikely to be resolved by adding more parameters but could potentially be resolved by multi-modal fusion (e.g., combining RGB with thermal infrared imagery).

The ArcGIS visualization pipeline demonstrates that combining classification outputs with geographic metadata produces a decision-support product that goes well beyond what tabulated accuracy statistics can offer. Overlaying results onto the road network (Figure 12) transforms a binary image classifier into a practical tool for avalanche risk assessment and road accessibility monitoring, one directly relevant to the operational responsibilities of civil protection agencies in mountain regions such as the Moroccan High Atlas and Anti-Atlas.

This study mostly uses RGB images, and that seems to cause some problems in tricky situations; for instance, when there is frozen ground with not much snow on it, and from up in the air it just looks like regular snow. It seems that this is one way performance drops off.

The classification part stays basic: just snow or snow-free. It does not get into how deep the snow is or anything about volume.

For the hardware side, the simulation was done on a regular lab computer. FLOPs and all parameters remain the same in such a setup. But latency, that is, the speed part, probably needs checking on a real drone with something like a Jetson Nano. Future work might handle that.

These results demonstrate that UAV-based deep learning can provide rapid and reliable snow monitoring in mountainous environments where satellite imagery is frequently affected by cloud cover or insufficient spatial resolution. The proposed framework therefore represents a practical solution for winter road maintenance, hydrological monitoring and disaster management in the Anti-Atlas region.

4. Conclusions

This paper presents a comprehensive end-to-end framework for UAV-based snow cover detection, combining a rigorous hardware-aware comparative evaluation of nine deep learning architectures with an operational geospatial visualization system. The key findings of this study are as follows.

The CNN models we tested, MobileNetV2 and EfficientNet from B0 to B7, performed better than SVM baseline, which only hit 98.13% accuracy on the test set. The models got accuracy values from 99.03 up to 99.85%, demonstrating a solid improvement. Second,

the EfficientNet models exhibited a notable trade-off between predictive accuracy and practical deployment performance on the hardware. Third, the larger variants (B5–B7) achieved marginal gains in accuracy; however, their inference speeds dropped below 25 frames per second, which is generally considered the threshold for real-time UAV applications. In addition, their memory requirements exceed the constraints of typical onboard drone hardware. These results highlight an inherent trade-off between accuracy and operational settings. Fourth, EfficientNet-B0 seemed ideal for deploying on a drone. It gave 99.21 percent accuracy, and the AUC was 0.99, hitting 105.2 FPS with just 4.01 million parameters. That makes it fit right in with the edge processors on working drone setups perfectly. Fifth, we presented the outputs from running EfficientNet-B0 with GPS data to make heatmaps in ArcGIS, layering them over road networks. This sets up a full pipeline for things like mapping snow hazards, checking road access, and figuring out emergency priorities. It shows that it is compatible with the required parameters in a practical way. Lastly, the proposed framework appears to be scalable and cost-effective, while remaining cognizant of hardware constraints for real-time snow monitoring in high-altitude environments. This approach shows potential for direct application in avalanche risk management, road network monitoring, and climate-related studies in regions such as the Anti-Atlas and the High Atlas, as well as in similar mountainous areas worldwide. However, its robustness across diverse environments has not yet been fully validated.

Author Contributions: A.E. (Abdelouahed Eddaoudy) carried out this research as part of his Ph.D. thesis, under the supervision of A.E.O., A.E. (Abdelaziz Elbaghdadi), and Y.Y.; A.E.-d. and A.E.O. contributed to the design and implementation of the research review, to the analysis and the selection of the papers, and to the writing of the manuscript. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data is available on request from the corresponding author. The full dataset has not been publicly released due to institutional restrictions and data ownership agreements. A representative subset, along with the annotation format and inference scripts, is available upon reasonable request to the corresponding author, but only for academic, non-commercial research purposes.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Taylor, C.; Robinson, T.R.; Dunning, S.; Carr, J.R.; Westoby, M. Glacial lake outburst floods threaten millions globally. *Nat. Commun.* **2023**, *14*, 487. [CrossRef] [PubMed]
2. Raleigh, M.S.; Small, E.E.; Bair, E.H.; Wobus, C.; Rittger, K. Snow monitoring at strategic locations improves water supply forecasting more than basin-wide mapping. *Commun. Earth Environ.* **2025**, *6*, 665. [CrossRef]
3. Schilling, S.; Andreas, D. Snow Water Equivalent Monitoring—A Review of Large-Scale Remote Sensing Applications. *Remote Sens.* **2024**, *16*, 1085. [CrossRef]
4. Sibandze, P.; Mhangara, P.; Odindi, J.; Kganyago, M. A Comparison of Normalised Difference Snow Index (NDSI) and Normalised Difference Principal Component Snow Index (NDPCSI) Techniques in Distinguishing Snow from Related Cover Types. *S. Afr. J. Geomat.* **2014**, *3*, 197–209. [CrossRef]
5. Mohsan, S.A.H.; Othman, N.Q.; Yanlong, L. Unmanned Aerial Vehicles (UAVs): Practical Aspects, Applications, Open Challenges, Security Issues, and Future Trends. *Intell. Serv. Robot.* **2023**, *16*, 109–137. [CrossRef] [PubMed]
6. Ed-daoudy, A.; El Oualkadi, A.; Elbaghdadi, A. Soldering Defect Detection with kNN and SVM for Quality Enhancement in Industrial. In *2024 International Conference on Connected Innovation and Technology (ICCITX)*; IEEE: New York, NY, USA, 2024; pp. 1–5.
7. Hu, Y.; Yuan, X.; Liu, W.; Hu, Q.; Wickert, J.; Jiang, Z. An SVM-Based Snow Detection Algorithm for GNSS-R Snow Depth Retrievals. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2022**, *15*, 6046–6052. [CrossRef]

8. Tan, M.; Le, Q.V. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In Proceedings of the 36th International Conference on Machine Learning, Long Beach, CA, USA, 9–15 June 2019.
9. Dawson, N.; Broxton, P.; Zeng, X. Evaluation of Remotely Sensed Snow Water Equivalent and Snow Cover Extent over the Contiguous United States. *J. Hydrometeorol.* **2018**, *30*, 1777–1791. [CrossRef]
10. Florath, J.; Keller, S. Glacier Monitoring Based on Multi-Spectral and Multi-Temporal Satellite Data: A Case Study for Classification with Respect to Different Snow and Ice Types. *Remote Sens.* **2022**, *14*, 845. [CrossRef]
11. Frei, A.; Tedesco, M.; Lee, S.; Foster, J.; Hall, D.K. A Review of Global Satellite-Derived Snow Products. *Adv. Space Res.* **2012**, *50*, 1007–1029. [CrossRef]
12. Matsuoka, M.; Yamazaki, F. Application of the Disaster Detection Method Using SAR Intensity Images to Recent Earthquakes. In Proceedings of the IEEE International Geoscience and Remote Sensing Symposium, Toronto, ON, Canada, 24–28 June 2002.
13. Balz, T.; Liao, M. Building Disaster Detection Using Post-Seismic High-Resolution SAR Satellite Data. *Int. J. Remote Sens.* **2010**, *31*, 3369–3391. [CrossRef]
14. Dong, L.; Shan, J. A Comprehensive Review of Earthquake-Induced Building Disaster Detection with Remote Sensing Techniques. *ISPRS J. Photogramm. Remote Sens.* **2013**, *84*, 85–99. [CrossRef]
15. Brett, P.T.B.; Guida, R. Earthquake Disaster Detection in Urban Areas Using Curvilinear Features. *IEEE Trans. Geosci. Remote Sens.* **2013**, *51*, 4877–4884. [CrossRef]
16. Sharma, R.C.; Tateishi, R.; Hara, K.; Nguyen, H.T. Earthquake Damage Visualization (EDV) Technique for Rapid Detection of Earthquake-Induced Damages Using SAR Data. *Sensors* **2017**, *17*, 235. [CrossRef] [PubMed]
17. Karimzadeh, S.; Ghasemi, M.; Matsuoka, M. A Deep Learning Model for Road Disaster Detection after an Earthquake Based on Synthetic Aperture Radar (SAR) and Field Datasets. *IEEE J. Remote Sens.* **2022**, *15*, 5753–5765. [CrossRef]
18. Mairesse, J.; Marcovici, I. Around probabilistic cellular automata. *Theor. Comput. Sci.* **2014**, *559*, 42–72. [CrossRef]
19. Basuki, A.; Arai, K. Two-Dimensional CA Approach for Disaster Spreading. *INOVASI* **2010**, *18*, 19–26.
20. Basuki, A.; Harsono, T.; Arai, K. Probabilistic Cellular Automata-Based Approach for Prediction of Hot Mudflow Disaster Area and Volume. *J. EMITTER* **2010**, *1*, 11–20.
21. Arai, K.; Basuki, A. New Approach for Prediction of Sidoarjo Hot Mudflow Disaster Area Based on Probabilistic Cellular Automata. *Geoinform. Int. J.* **2011**, *1*, 1–11.
22. Arai, K. Cell-Based GIS as Cellular Automata for Disaster Spreading Prediction and Required Data Systems. *Data Sci. J.* **2012**, *12*, 137–141. [CrossRef]
23. Arai, K. Sensor Network for Landslide Monitoring with Laser Ranging System Avoiding Rainfall Influence by Time Diversity and Satellite Imagery Data-Based Landslide Disaster Relief. *Int. J. Appl. Sci.* **2012**, *3*, 1–12.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.