AI

Article

Deep Learning-Based Snake Species Identification for Enhanced Snakebite Management

Mohamed Iguernane [1,†], Mourad Ouzziki [2,†], Youssef Es-Saady [2,3,*], Mohamed El Hajji [3,4], Aziza Lansari [5] and Abdellah Bouazza [5]

[1] ISIMA Laboratory, Polydisciplinary Faculty of Taroudant, Ibnou Zohr University, Taroudant 83000, Morocco; m.iguernane@uiz.ac.ma
[2] Polydisciplinary Faculty of Taroudant, Ibnou Zohr University, Taroudant 83000, Morocco; mourad.ouzziki@edu.uiz.ac.ma
[3] IRF-SIC Laboratory, Ibnou Zohr University, Agadir 80000, Morocco; m.elhajji@crmefsm.ac.ma
[4] Regional Center for Education and Training Professions-Souss Massa, Agadir 80000, Morocco
[5] 2GBEI Laboratory, Polydisciplinary Faculty of Taroudant, Ibnou Zohr University, Taroudant 83000, Morocco; a.lansari@uiz.ac.ma (A.L.); a.bouazza@uiz.ac.ma (A.B.)
* Correspondence: y.essaady@uiz.ac.ma
† These authors contributed equally to this work.

Abstract: Accuratesnake species identification is essential for effective snakebite management, particularly in regions like Morocco, where approximately 400 snakebite incidents are reported annually, with a case fatality rate of 7.2%. Identifying venomous snakes promptly can significantly improve treatment outcomes by enabling the timely administration of specific antivenoms. However, the absence of comprehensive databases and rapid identification tools for Moroccan snake species poses challenges to effective clinical responses. This study presents a deep learning-based approach for the automated identification of Moroccan snake species. Several architectures, including VGG-19, VGG-16, and EfficientNet B0, were evaluated for their classification performance. EfficientNet B0 emerged as the most effective model, achieving an accuracy of 92.23% and an F1-score of 93.67%. After training on the SnakeCLEF 2021 dataset and fine-tuning with a specialized local dataset, the model attained a validation accuracy of 94% and an F1-score of 95.86%. To ensure practical applicability, the final model was deployed on a web platform, enabling the rapid and accurate identification of snake species via image uploads. This platform serves as a valuable tool for healthcare professionals and the general public, facilitating improved clinical response and educational awareness. This study highlights the potential of AI-driven solutions to address challenges in snakebite identification and management, offering a scalable approach for regions with limited resources and high snakebite prevalence.

Keywords: snake species identification; artificial intelligence (AI); deep learning; snakebite management; Moroccan snakes; healthcare technology; transfer learning; fine-tuning

Academic Editor: Tin-Chih Toly Chen

Received: 29 November 2024
Revised: 18 January 2025
Accepted: 19 January 2025
Published: 21 January 2025

Citation: Iguernane, M.; Ouzziki, M.; Es-Saady, Y.; El Hajji, M.; Lansari, A.; Bouazza, A. Deep Learning-Based Snake Species Identification for Enhanced Snakebite Management. *AI* **2025**, *6*, 21. https://doi.org/10.3390/ai6020021

1. Introduction

Snakebite envenomation is a significant public health concern, with the World Health Organization (WHO) estimating that 1.8 to 5.4 million people are bitten by snakes annually, resulting in 1.8 to 2.7 million cases of envenomation. These incidents cause approximately 81,410 to 137,880 deaths per year, while three times as many people suffer amputations or permanent disabilities due to snakebites [1]. In Morocco, around 400 snakebite cases are reported annually, with an incidence rate of 0.1 to 1.4 per 100,000 inhabitants and a case fatality rate of 7.2% [2,3]. Accurate identification of the snake species responsible for a bite is critical to selecting the appropriate antivenom and ensuring effective treatment [4]. Although protocols in Morocco emphasize this importance [2], several challenges hinder the process: Healthcare professionals often lack the necessary training, which limits the effectiveness of treatment and highlights the need for educational and technological improvements. Many providers struggle with taxonomically identifying the snake and instead rely on symptoms to infer the species. Molecular identification techniques, such as immunoassays that detect venom antigens, also have limitations, particularly in low-resource settings. In addition, obtaining accurate snake specimens can be difficult, as they are often damaged or misidentified, further complicating the identification process. A total of 26 distinct species of snakes are native to Morocco [5], including seven medically important venomous species. These include *Naja naje legionis* from the Elapidae family and six species from the Viperidae family, such as *Daboia mauritanica*, *Cerastes cerastes*, and *Bitis arietans*, which, together, account for 67.2% of serious envenomations [6]. Despite their clinical importance, the lack of a comprehensive database and specific identification tools for Moroccan snakes presents a significant challenge in accurate and rapid species recognition, particularly in high-risk scenarios.

In recent years, deep learning has expanded tremendously, impacting nearly every field. Its remarkable classification performance has drawn attention from researchers in diverse areas such as agriculture [7], software engineering [8], computational biology [9], healthcare [10], and medicine [11]. However, a major challenge in these fields is the limited availability of large and reliable datasets, which are essential for building high-performance predictive models using deep learning techniques. To overcome this issue, transfer learning has become a widely adopted approach. Transfer learning allows for knowledge from a model trained on a specific task to be applied to a new model tackling a related but different task. Recent studies have shown that transfer learning, particularly when applied to deep convolutional neural networks (CNNs), is highly beneficial as it reduces the need for large datasets and decreases the training time due to the model being partially pre-trained. For instance, researchers [12] have achieved excellent results using transfer learning on Caltech101 datasets. Architectures such as ResNet and InceptionResNet have demonstrated exceptional performance, with the validation accuracy surpassing 99%, excelling in precision, recall, and F1-scores. Models like VGG and EfficientNet have also performed strongly, with EfficientNet achieving a 96.89% accuracy and excellent precision. In [13], the authors applied pre-trained deep convolutional networks (ResNet50, InceptionV3, and InceptionResNetV2) to classify cellular morphological changes, achieving predictive accuracies between 95% and 97%. Gu and Lee [14] introduce an approach to transfer learning, specifically tailored to medical image analysis. Their framework, real-world feature transfer learning, leverages models pre-trained on large-scale general-purpose datasets, such as ImageNet, to classify medical images portraying conditions, such as pneumonia in X-ray images.These results underscore the effectiveness of transfer learning in overcoming data limitations, making it an attractive option for snake species identification in Morocco.

A substantial body of research has explored the application of machine learning and deep learning to automate snake species identification. Early efforts using machine learning techniques highlighted key limitations. Amir et al. [15] investigated texture-based characteristics with the Color and Edge Directivity Descriptor (CEDD), achieving accuracies of 89.22% and 87.93% with nearest neighbors and backpropagation neural networks, respectively. However, their reliance on texture features alone underscored the need to incorporate additional attributes, such as shape and color, to improve classification outcomes. With advances in deep learning, convolutional neural networks (CNNs) have demonstrated remarkable potential in automating feature extraction and enhancing the classification accuracy. Patel et al. [16] developed a mobile application to identify nine species of snakes in the Galápagos Islands, achieving a maximum accuracy of 75% using Faster R-CNN and ResNet architectures. This study highlighted the utility of mobile-based solutions for real-time field applications. The study by Binta Islam et al. [17] also demonstrates the efficacy of deep learning in addressing similar challenges. Their study tested pre-trained models like VGG16 and ResNet50, as well as a self-trained CNN, to classify snakes, lizards, and toads from camera trap images. Their results showed that transfer learning-based models significantly outperformed self-trained models, achieving accuracies of 87% (VGG16) and 86% (ResNet50), compared to 72% for self-trained CNN. Yang and Sinnott [18] developed an iOS application that used CNNs for snake classification, exploring both offline and cloud-based deployment strategies to assess performance differences in the accuracy and classification time. These findings further validate the benefits of transfer learning in ecological species classification tasks. Meanwhile, Dandeniya et al. [19] applied transfer learning to distinguish Russell's Viper from Indian Rock Python, achieving a 94.5% precision using pre-trained models and crowd-sourced data from Sri Lankan Facebook groups. These examples emphasize the adaptability of deep learning techniques to regional contexts and their potential to address data scarcity through innovative approaches.

Overall, research trends reveal the progression from traditional machine learning to advanced CNN-based architectures, emphasizing the integration of geographic data and domain-specific fine-tuning to improve accuracy and applicability. Despite these advancements, there remains a notable absence of research specifically tailored to Moroccan snake species. Existing models trained in global datasets do not capture the unique characteristics of snakes native to Morocco, limiting their functionality for local snakebite management. Previous studies have highlighted the importance of understanding the distribution and frequency of snake species of venom in different regions to effectively inform healthcare strategies [20]. In particular, a meta-analysis of global snake assemblages demonstrated significant variation in the occurrence of venomous species and individuals according to the habitat and geography, emphasizing the need for region-specific approaches in snakebite management. To address this gap, this study develops a deep learning-based framework tailored to Morocco's unique ecosystem.

The objectives of this study are, first, to evaluate various deep learning architectures, including VGG-19, VGG-16, and EfficientNet B0, for the automated identification of Moroccan snake species. Second, to assess the performance of these models by training them on the SnakeCLEF 2021 dataset and fine-tuning with a specialized local dataset to improve the classification accuracy. Finally, to deploy the most effective model as a web-based platform that enables real-time snake species identification through image uploads. As a result, this study aims to provide a practical tool that facilitates accurate snake identification, assisting healthcare professionals in managing snakebite cases and raising public awareness about snake biodiversity. The key contributions of this work are as follows:

1. **Database Construction**: We constructed a comprehensive database containing images and the scientific names of 26 snake species native to Morocco, collected from diverse habitats across the country. This dataset serves as the foundation for training and evaluating deep learning models.
2. **EfficientNet B0 Architecture**: We introduce the EfficientNet B0 model, adapted through transfer learning, which achieved a state-of-the-art accuracy of 94% on the Moroccan snake species dataset.
3. **Comparative Analysis**: We performed a comparative analysis of three CNN architectures—VGG19, VGG16, and EfficientNet B0—fine-tuned specifically for snake species identification.

4. **Web Application Development**: We developed a user-friendly web application that deploys the most effective model, enabling real-time snake species identification through image uploads.

The remainder of this paper is organized as follows: Section 2 outlines the methods and materials, including details on datasets, preprocessing, and model architectures. Section 3 presents the experimentation, results, and the development of the web application for real-time identification. Finally, Section 4 concludes the study with a summary of findings and future research directions.

2. Materials and Methods

This section outlines the methodology adopted for this study, detailing snake identification techniques, dataset characteristics, preprocessing steps, and deep learning architectures used. Insights from the literature review informed the integration of both local and global datasets to ensure robust model performance and ecological relevance, as shown in Figure 1, which provides an overview of the proposed system.

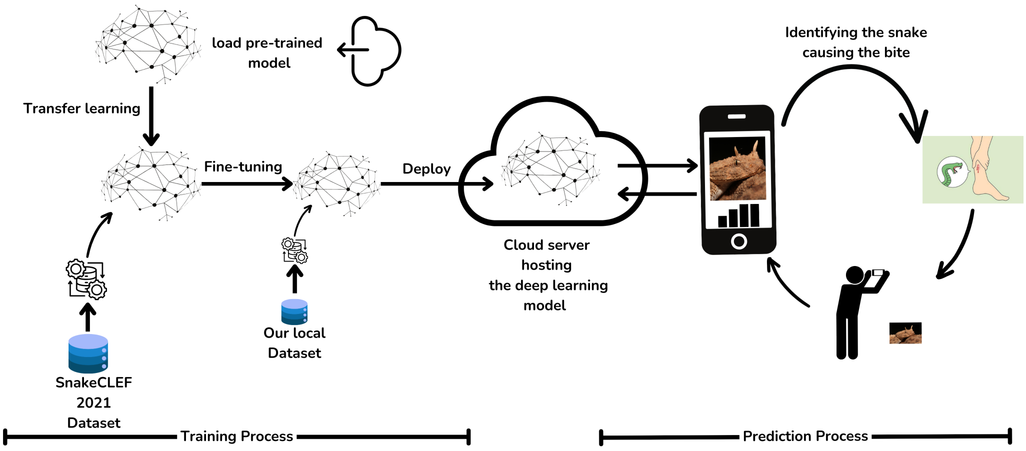

Figure 1. Overview of the proposed system.

2.1. Local Dataset Description

The local dataset comprises 3922 images representing 26 distinct snake species native to Morocco. These images were collected by herpetologists from Ibnou Zohr University (FPT-Taroudant, Morocco) between 2010 and 2023 in diverse natural habitats, including the Sahara Desert, Atlas Mountains, Rif, and Souss Valley. The dataset classifies snakes into three venom categories: venomous, non-venomous, and mildly venomous. Furthermore, it categorizes these species into eight families, encompassing all snake species identified in Morocco. This comprehensive classification provides valuable insights into biodiversity and the potential risks associated with each species [5], as summarized in Table 1.

The dataset represents a diverse distribution of classes, with some species well represented and others underrepresented. For example, medically significant venomous species (see Table 1), such as *Cerastes cerastes* and *Daboia mauritanica* (Figure 2), are more prevalent, whereas non-venomous species like *Eryx jaculus* and *Macroprotodon abubakeri* are relatively rare in Morocco. This imbalance reflects ecological realities and highlights challenges in building comprehensive classification models.

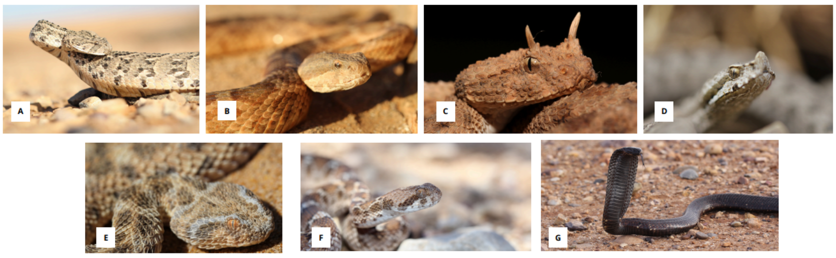

Figure 2. (**A**): *Bitis arietans*, (**B**): *Daboia mauritanica*, (**C**): *Cerastes cerastes*, (**D**): *Vipera monticola*, (**E**): *Cerastes vipera*, (**F**): *Echis pyramidum*, and (**G**): *Naja haje*. Photos by Bouazza A.

Table 1. List of Moroccan snakes with scientific, French, and English names, including venom status ($^{-}$ harmless, * mildly venomous, *** venomous) [5].

Ophidia (Snakes)	**French Name/English Name**
Family Leptotyphlopidae (Slender Blind Snakes)	
Myriopholis algeriensis (Jacquet, 1895) $^{-}$	Leptotyphlops d'Algérie/Beaked Thread-snake
Family Boidae	
Eryx jaculus (Linnaeus, 1758) $^{-}$	Boa javelot/Javeline sand boa
Family Colubridae (Colubrids)	
Hemorrhois algirus (Jan, 1863) $^{-}$	Couleuvre algire/Algerian Whip Snake
Hemorrhois hippocrepis (Linnaeus, 1758) $^{-}$	Couleuvre fer-à-cheval/Horseshoe Whip Snake
Lytorhynchus diadema (Duméril, Bibron & Duméril, 1854) $^{-}$	Couleuvre fouisseuse à diadème/Awl-headed Snake
Macroprotodon abubakeri (Wade, 2001) $^{-}$	Couleuvre à capuchons d'Abubaker /Abu Baker False Smooth Snake –
Macroprotodon brevis (Günther, 1862) $^{-}$	Couleuvre à capuchons occidentale/False Smooth Snake
Coronella girondica (Daudin, 1803) $^{-}$	Couleuvre girondine/Southern Smooth Snake
Dasypeltis sahelensis (Trape & Mané, 2006) $^{-}$	Couleuvre mangeuse d'œufs du Sahel/Sahel Egg Eater
Spalerosophis diadema (Schlegel, 1837) $^{-}$	Couleuvre à diadème de Clifford/Diadem Snake
Spalerosophis dolichospilus (Werner, 1923) $^{-}$	Couleuvre à diadème du Maghreb/Mograbin Diadem Snake
Telescopus tripolitanus (Werner, 1909) *	Couleuvre-chatde Tripolitaine/Tripolitan Cat Snake
Family Natricidae	
Natrix astreptophora (Seoane, 1884) $^{-}$	Couleuvre astreptophore/Grass Snake
Natrix maura (Linnaeus, 1758) $^{-}$	Couleuvre vipérine/Viperine Snake
Family Lamprophiidae	
Boaedon fuliginosus (Boie, 1827) $^{-}$	Couleuvre commune d'Afrique/African House Snake
Family Psammophiidae	
Malpolon insignitus (Geoffroy Saint-Hilaire, 1827) *	Couleuvre de Montpellier orientale/Eastern Montpellier Snake
Malpolon monspessulanus (Hermann, 1804) *	Couleuvre de Montpellier/Montpellier Snake
Psammophis schokari (Forsskål, 1775) *	Couleuvre de Schokar/Schokari Sand Racer
Malpolon moilensis (Reuss, 1834) *	Couleuvre de Moïla/Moila Snake
Family Elapidae (Cobras)	
Naja haje (Linnaeus, 1758) ***	Cobra d'Afrique du Nord/Egyptian cobra
Family Viperidae (Vipers)	
Vipera monticola (Saint Girons, 1953) ***	Vipère naine de l'Atlas/Atlas Dwarf Viper
Echis pyramidum (Geoffroy Saint-Hilaire, 1827) ***	Vipère des pyramides/White-Bellied Carpet Viper
Cerastes cerastes (Linnaeus, 1758) ***	Vipère à cornes/Desert Horned Viper
Cerastes vipera (Linnaeus, 1758) ***	Vipère de l'erg/Sahara Sand Viper
Bitis arietans (Merrem, 1820) ***	Vipère heurtante/Puff Adder
Daboia mauritanica (Gray, 1849) ***	Vipère de Maurétanie/Moorish viper

2.2. Global Dataset Description

The SnakeCLEF 2021 dataset provides a broader perspective, consisting of 409,679 images encompassing 772 snake species from 188 countries across all continents. Of these, 386,006 labeled images are designated for development, while 23,673 unlabeled images are reserved for testing. A train–validation split of 90% and 10% is implemented for the development data to ensure consistent species distributions. This dataset is part of the SnakeCLEF2021 species identification challenge within LifeCLEF21, by the Conference and Labs of the Evaluation Forum (CLEF), which proposes data-centric challenges aimed at identifying and predicting biodiversity [21]. The data were collected from online biodiversity platforms such as iNaturalist (https://www.inaturalist.org/, accessed on 25 July 2024), HerpMapper (https://www.herpmapper.org/, accessed on 25 July 2024), and supplemented by images scraped from Flickr (https://www.flickr.com/, accessed on 25 July 2024). Rigorous preprocessing was applied to eliminate duplicate entries and correct label inaccuracies.

The dataset is organized into clean subsets (from iNaturalist and HerpMapper) and noisy subsets (from Flickr), facilitating effective model training. However, the dataset exhibits significant class imbalance. For instance, *Thamnophis sirtalis* is represented by 22,163 images, whereas *Achalinus formosanus* has only 10 images, underscoring the challenge of achieving balanced feature learning [22]. Figure 3 illustrates a range of images from the Snake CLEF 2021 dataset.

Figure 3. Image examples from the Snake CLEF 2021 dataset.

2.3. Comparison Between the Datasets

Table 2 presents a comparative analysis of the local Moroccan dataset and the SnakeCLEF 2021 global dataset. This comparison highlights key differences in scale, scope, and use cases, illustrating their complementary roles in this study.

Table 2. Comparison between the datasets.

Dataset	# of Species	# of Images	Size (GB)	Min per Species	Max per Species	Geographical Scope
SnakeCLEF 2021 (Global Dataset)	772	386,006	66	10	22,163	Global: Multiple regions worldwide
Local Dataset	26	3922	19	16	377	Local: Moroccan ecosystems

2.4. Data Preprocessing

Data preprocessing is a fundamental step in optimizing the performance of machine learning models. A well-structured preprocessing pipeline ensures reliable and robust results. For this study, the preprocessing sequence was designed to standardize the datasets and enhance model training and generalization. All images were resized to a uniform resolution of 224 × 224 × 3 pixels to ensure compatibility with the input requirements of deep learning models. During training, a random resized crop operation was applied to introduce variability and improve the model's performance.

In contrast, validation and test images were resized consistently to maintain evaluation integrity. Data augmentation techniques, including horizontal and vertical flipping, resizing, and rotation, were applied with a 50% probability during training to mitigate overfitting and enhance model generalization. These augmentations, implemented using the Albumentations library [23], included transformations such as RandomResizedCrop, Transpose, Resize, HorizontalFlip, VerticalFlip, ShiftScaleRotate, and Normalize. Normalization of pixel values was applied to ensure consistency with the model's pre-training on the ImageNet dataset and to improve convergence during training. Specifically, pixel values for each channel RGB were increased to have means of 0.485, 0.456, and 0.406, respectively, with corresponding standard deviations of 0.229, 0.224, and 0.225.

Figure 4 presents augmented images of *Naja haje* from the local Moroccan snake dataset, providing a visual representation of the preprocessing pipeline's impact. The training images underwent various random augmentations, including horizontal and vertical flipping, rotation, and resizing. These augmentations introduce variability to the dataset, thereby improving the model's robustness by exposing it to different perspectives of the same species.

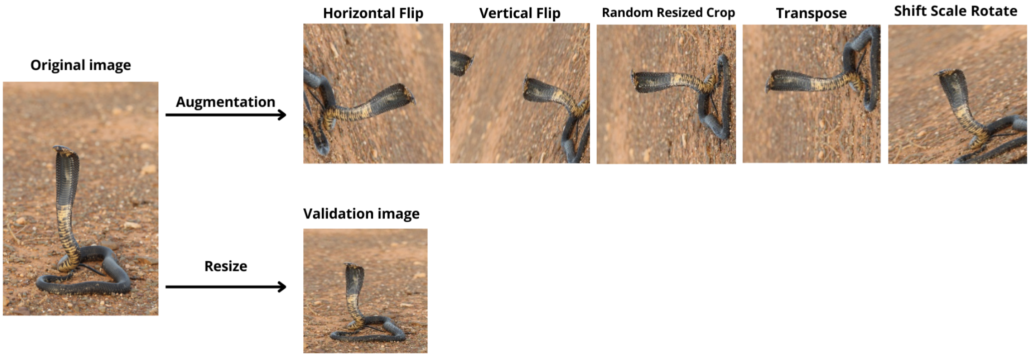

Figure 4. Albumentations-augmented images of *Naja haje* from the Moroccan local dataset.

This preprocessing step aligns the input data with the distribution expected by the pre-trained model, promoting stability and more efficient learning. This standardization improves the model's ability to generalize to unseen data. Additionally, label preprocessing was employed to format the categorical labels appropriately for multi-class classification. Using label binarization, the categorical labels were converted into binary matrices, with each class represented as a binary vector [24], enabling the model to effectively handle the classification of multiple snake species. The datasets were partitioned to ensure a balanced distribution for training and evaluation. The local dataset was split into 80% for training and 20% for testing, while the global dataset employed a 90% training and 10% testing split, maintaining consistent species distributions across subsets. These splits were designed to optimize model performance and ensure comparability between the datasets.

2.5. Overview of the Approach

This section details the methodology used to develop an accurate snake species identification model, highlighting three key areas: the selection of CNN architectures (VGG and EfficientNet), and the application of model tuning and transfer learning. Figure 5 illustrates our process of transfer learning and fine-tuning, using pre-trained CNN architectures (VGG and EfficientNet), for snake species classification.

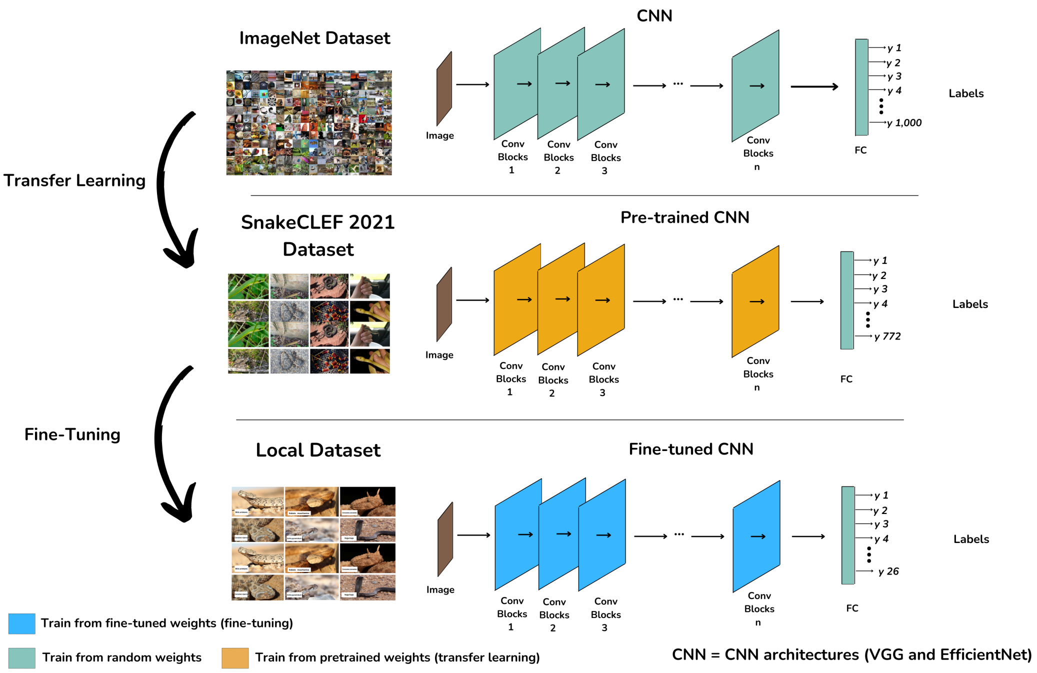

Figure 5. Transfer learning and fine-tuning process for snake species classification using pre-trained CNN architectures (VGG and EfficientNet).

2.5.1. VGG Architecture

The VGG architecture, proposed by Simonyan et al. [25], is characterized by its simplicity and depth. The VGG-16 model comprises 13 convolutional layers and 3 fully connected layers, while the VGG-19 model offers a deeper configuration with additional layers. Both architectures use 3 × 3 convolutional filters with a stride of 1, enabling the extraction of intricate image features ranging from simple edges to complex textures and shapes. A key feature of the VGG architecture is its modular design, which organizes convolutional layers into blocks, each followed by max-pooling layers. This arrangement progressively reduces spatial dimensions while preserving essential features, enabling the network to build a hierarchical understanding of input images. The final fully connected layers perform the classification based on the features extracted throughout the network. VGG's deep structure, especially in versions like VGG-16 and VGG-19, excels in recognizing complex visual patterns, as demonstrated in its success with the ImageNet dataset. These versions, comprising 16 and 19 weight layers, respectively (as detailed in Table 3 for VGG-16 and Table 4 for VGG-19), effectively capture subtle differences, enhancing the classification accuracy. Although computationally expensive, VGG remains a benchmark architecture for its high performance across a variety of image recognition tasks.

Table 3. Structure of VGG-16 model.

Stage	Operator	Resolution	#Channels	#Layers
1	Conv3 × 3	224 × 224	64	2
2	Conv3 × 3	112 × 112	128	2
3	Conv3 × 3	56 × 56	256	3
4	Conv3 × 3	28 × 28	512	3
5	Conv3 × 3	14 × 14	512	3
6	Fully connected	7 × 7	4096	3

Table 4. Structure of VGG-19 model.

Stage	Operator	Resolution	#Channels	#Layers
1	Conv3 × 3	224 × 224	64	2
2	Conv3 × 3	112 × 112	128	2
3	Conv3 × 3	56 × 56	256	4
4	Conv3 × 3	28 × 28	512	4
5	Conv3 × 3	14 × 14	512	4
6	Fully connected	7 × 7	4096	3

2.5.2. EfficientNet Architecture

EfficientNet, a convolutional neural network built on the principle of "compound scaling", optimally balances model width, depth, and resolution to improve accuracy and efficiency [26]. As illustrated in Figure 6, this approach scales these dimensions proportionally using a compound coefficient (ϕ), enabling EfficientNet to achieve superior performance with fewer parameters compared to earlier models. The architecture integrates Mobile Inverted Bottleneck (MBConv) layers and Squeeze-and-Excitation (SE) blocks, enhancing the network's capacity to focus on critical features while maintaining computational efficiency. EfficientNet's compound scaling ensures an optimal trade-off between model size and computational power, making it highly suitable for tasks requiring resource-efficient solutions.

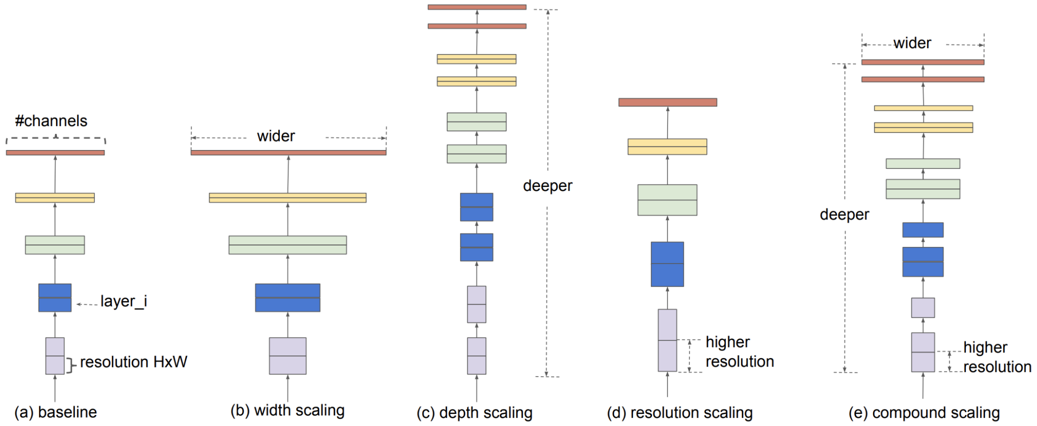

Figure 6. Different scaling methods vs compound scaling [26].

EfficientNet has demonstrated exceptional performance in snake species identification, outperforming its predecessors by achieving comparable or higher accuracy with significantly fewer parameters. For instance, while the state-of-the-art GPipe model requires 556 million parameters, EfficientNet attains a similar accuracy using only 66 million parameters. This efficiency makes EfficientNet particularly advantageous in computationally constrained environments, offering faster processing without sacrificing accuracy.

The EfficientNet family includes seven variants, EfficientNet B0 through B7, each designed with varying dimensions [26]. In this study, we used EfficientNet B0 due to its balance between computational efficiency and accuracy. Table 5 summarizes the architecture of EfficientNet B0, highlighting its ability to efficiently process images while maintaining high classification performance.

Table 5. Structure of EfficientNet-B0 model.

Stage	Operator	Resolution	#Channels	#Layers
1	Conv3 × 3	224 × 224	32	1
2	MBConv1, k3 × 3	112 × 112	16	1
3	MBConv6, k3 × 3	112 × 112	24	2
4	MBConv6, k5 × 5	56 × 56	40	2
5	MBConv6, k3 × 3	28 × 28	80	3
6	MBConv6, k5 × 5	14 × 14	112	3
7	MBConv6, k5 × 5	14 × 14	192	4
8	MBConv6, k3 × 3	7 × 7	320	1
9	Conv1 × 1 & Pooling & FC	7 × 7	1280	1

2.5.3. Transfer Learning Approach and Model Fine-Tuning

Transfer learning employs two key approaches: feature extraction and fine-tuning [27]. Feature extraction leverages fixed pre-trained weights to capture relevant features for the target task, integrating them with a task-specific classifier [28,29]. In contrast, fine-tuning involves modifying the pre-trained model's parameters to align with the specific requirements of the target task [30]. This approach is valued for its computational efficiency and its ability to facilitate model reuse.

Pre-trained weight was used from the Timm library, which includes models trained on the extensive ImageNet dataset. ImageNet is a hierarchical image database comprising over 14 million annotated images spanning 1000 categories, widely recognized for its contributions to advancing image classification and object detection tasks. To tailor the model specifically for snake species classification, we implemented a two-stage transfer learning approach. Initially, the model was pre-trained with weights from the Timm library [31], using models trained on the extensive ImageNet dataset. In the first stage, the model was trained on the global SnakeCLEF 2021 dataset, which includes 772 classes, allowing it to capture generalized image features across a wide variety of snake species. This stage enabled the model to extract relevant patterns applicable to the broader domain of snake identification. In the second stage, fine-tuning was applied to adapt the model to the local dataset. The final layer of the model was replaced with a classification head specific to the 26 local snake species. We froze the feature extraction layers and unfroze the fully connected layer to fine-tune the model using a smaller learning rate. This approach preserved the pre-trained weights while allowing the model to adjust effectively to the distinct characteristics of the local dataset. Fine-tuning in this stage ensured that the generalized features learned in the first phase were refined and better aligned with the specific requirements of the task.

This dual approach—initial training on a large, diverse dataset through transfer learning followed by task-specific fine-tuning—enabled the model to leverage the generalized features learned from ImageNet and SnakeCLEF 2021 while efficiently adapting to the unique morphological and color patterns of the local snake species. This method is hypothesized to improve both classification accuracy and computational efficiency, making it particularly suitable for addressing the challenges posed by the limited size of the local dataset.

3. Experimentation, Results, and Web Application

3.1. Experimental Setup

The experimentation involved training three model architectures, VGG19, VGG16, and EfficientNet B0, each initialized with pre-trained weights from the Timm library, based on the extensive ImageNet dataset. The training process was conducted over 50 epochs using the Adam optimizer [32] with a learning rate of 1×10^{-4}. The loss function, Categorical Cross-Entropy, was used to evaluate the error between the predicted and actual labels, with a batch size of 64 representing the number of samples processed per iteration. To enhance the training efficiency and convergence, we employed the StepLR scheduler, with a step size of 5 epochs and a gamma value of 0.1, to dynamically reduce the learning rate. All models were implemented and trained using the PyTorch framework, leveraging Python 3 for its robust libraries, including Timm, which facilitates transfer learning. The experiments were conducted on a machine equipped with a 12th Gen Intel(R) Core(TM) i7-12700KF CPU (3.60 GHz, 12 cores), 64 GB of RAM, and an NVIDIA GeForce RTX 3090 GPU with 24 GB of GDDR6X memory. This hardware configuration ensured consistent performance and efficient handling of all training and evaluation tasks.

3.2. Evaluation Metrics

Evaluation metrics are essential quantitative measures for assessing the performance and effectiveness of machine learning models. These metrics provide insights into a model's predictive capability, generalization ability, and overall quality, enabling the comparison of different models or algorithms [33]. The primary metrics used in this study are accuracy and $F1$-score, calculated using the following equations:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}; \tag{1}$$

$$F1 = \frac{2 \times TP}{2 \times TP + FP + FN}; \tag{2}$$

where TP is True Positive, FN is False Negative, FP is False Positive, and TN is True Negative.

3.3. Results and Discussion

After preparing the datasets and loading the pre-trained models, we trained and evaluated several architectures. The results are discussed below. During the training of EfficientNet B0, metrics such as the F1-score, accuracy, and loss were recorded for both the training and validation sets at each epoch. As shown in Figure 7, the training and validation losses steadily decreased and converged without significant divergence, indicating that the model avoided overfitting. The alignment between the training and validation F1-scores confirmed the model's reliability in terms of precision and recall on the validation set.

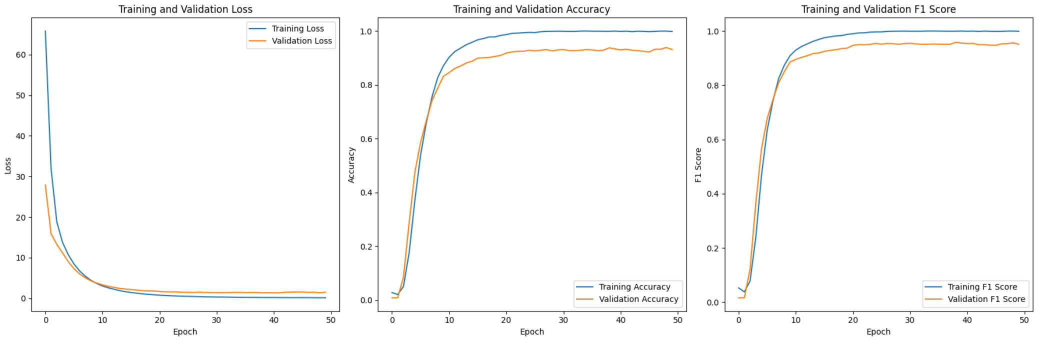

Figure 7. EfficientNet-B0 fine-tuning on local dataset.

The comparative results of the trained models are presented in Tables 6 and 7. EfficientNet B0 consistently outperformed the VGG architectures, achieving the highest accuracy of 94% and an F1-score of 95.86% after pre-training on the SnakeCLEF 2021 dataset and fine-tuning on the local dataset. This superior performance is attributed to EfficientNet B0's advanced scaling techniques, which effectively balance the model size, computational cost, and performance. This balance enables the model to generalize better to the local dataset compared to the deeper VGG models, which, despite their strong feature extraction capabilities, are prone to overfitting due to their larger number of parameters. The application of transfer learning significantly improved the performance of all the models. Pre-training on a large, diverse dataset such as SnakeCLEF 2021 provided generalized feature representations that enhanced the models' ability to adapt to the smaller, domain-specific local dataset. For example, VGG-16 exhibited notable gains in the accuracy and F1-score compared to its performance when trained solely on the local dataset, underlining the effectiveness of leveraging a large pre-trained model. EfficientNet B0 emerged as the most effective architecture for snake species identification, excelling in both accuracy and F1-score across all the training scenarios. Its efficient design, leveraging compound scaling and pre-trained weights, allowed it to outperform the VGG models while maintaining computational efficiency. These results highlight the importance of employing advanced architectures and transfer learning techniques for tasks involving limited domain-specific data.

Table 6. Comparison of performance metrics across different models trained on the local dataset only.

Model	Metric	Performance on Testing Data
VGG-19	Accuracy	0.8089
	F1-Score	0.8736
VGG-16	Accuracy	0.8429
	F1-Score	0.8960
EfficientNet B0	Accuracy	0.9223
	F1-Score	0.9367

Table 7. Comparison of performance metrics across different models trained on the SnakeCLEF 2021 dataset and then fine-tuned on a local dataset.

Model	Metric	Performance on Testing Data
VGG-19	Accuracy	0.8599
	F1-Score	0.8936
VGG-16	Accuracy	0.9019
	F1-Score	0.9247
EfficientNet B0	Accuracy	0.9400
	F1-Score	0.9586

3.4. Web Application

In order both to facilitate the practical application of the research findings and to provide a convenient tool for snake species identification, a web application has been developed which enables users to upload images and then to identify the species in question. The interface allows users to interact smoothly by dragging and dropping an image into the designated area. Upon uploading an image and clicking the "Submit" button, the application processes the input instantly using the pre-trained EfficientNet B0 model fine-tuned for Moroccan snake species.

The results section of the application displays the top five predictions, ranked by confidence scores, as shown in Figure 8. For example, the application can accurately identify species such as *Naja haje* (Figure 8) or *Bitis arietans* (Figure 9), providing detailed predictions accompanied by confidence levels for each identified species. This intuitive interface and real-time processing make the application a practical tool for researchers, herpetologists, and the general public, offering a robust solution for snake species identification.

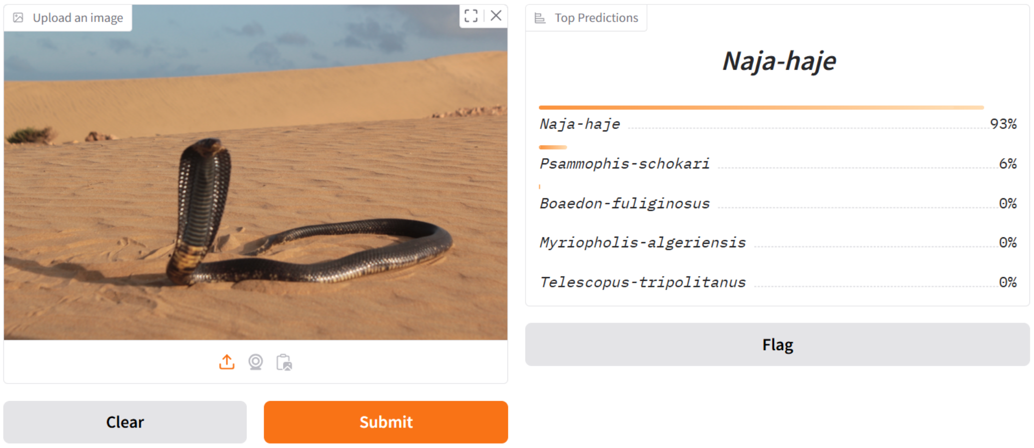

Naja-haje

Naja-haje	93%
Psammophis-schokari	6%
Boaedon-fuliginosus	0%
Myriopholis-algeriensis	0%
Telescopus-tripolitanus	0%

Figure 8. Identify the *Naja haje*.

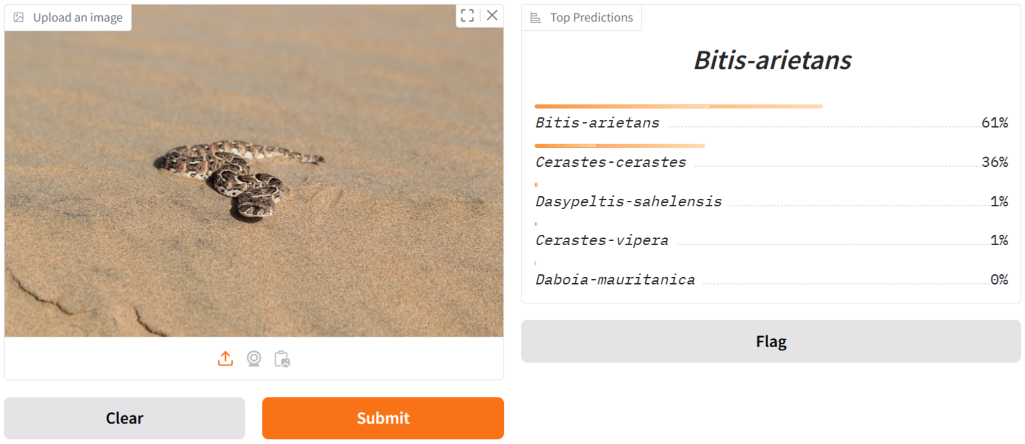

Bitis-arietans

Bitis-arietans	61%
Cerastes-cerastes	36%
Dasypeltis-sahelensis	1%
Cerastes-vipera	1%
Daboia-mauritanica	0%

Figure 9. Identify the *Bitis arietans*.

To assess the computational efficiency of our models, we tested them on the task of identifying *Naja haje* shown in Figure 8 and calculated the inference time and memory usage on a machine with an Intel(R) Core(TM) i5-8250U CPU @ 1.60 GHz 1.80 GHz. The results, presented in Table 8, show that EfficientNet B0 not only achieved the highest performance metrics but also demonstrated significantly faster inference times and a lower memory consumption compared to the VGG models. Specifically, EfficientNet B0 had an inference time of 31.72 ms and memory usage of 1409.09 MB, while VGG-16 and VGG-19 took considerably longer and required more memory.

Table 8. Comparison of inference time and occupied memory for different models.

Model	Inference Time (ms)	Occupied Memory (MB)
EfficientNet B0	31.72	1409.09
VGG-16	251.26	1847.21
VGG-19	339.88	1848.65

4. Conclusions and Future Work

Snake species identification remains a challenging task due to the diversity of species, subtle morphological differences, and imbalanced datasets. In this study, the EfficientNet B0 architecture, fine-tuned on global and local datasets, achieved a validation accuracy of 94% and an F1-score of 95.86%. These results demonstrate the potential of deep learning models to handle complex classification tasks and adapt to domain-specific challenges through transfer learning. Despite these promising outcomes, challenges remain, including the underrepresentation of certain species, the impact of low-quality images, and the presence of non-snake images in datasets, which poses an additional challenge as the model might generate False Positives or fail to indicate the absence of snakes. Furthermore, the snake involved is often not seen, adding another layer of difficulty in accurate identification. Addressing these limitations is crucial to improving the model's reliability and adaptability.

Future efforts will explore advanced architectures, such as transformers, for capturing intricate visual details, and incorporate geographic data to provide contextual support for classification. Tackling dataset imbalances and extending the system's application to global datasets will further enhance its usefulness. These advancements aim to establish a reliable tool for ecological research and medical applications, improving snakebite management and biodiversity studies.

Author Contributions: Conceptualization, M.O., M.I., Y.E.-S., M.E.H., A.B. and A.L.; methodology, M.I., M.O., Y.E.-S., M.E.H., A.B. and A.L.; software, M.O., M.I., Y.E.-S. and M.E.H.; validation, Y.E.-S., M.E.H., A.B. and A.L.; formal analysis, M.O., M.I., Y.E.-S., M.E.H., A.B. and A.L.; investigation, M.O., M.I., Y.E.-S., M.E.H., A.B. and A.L.; data curation, A.B. and A.L.; writing—original draft preparation, M.O. and M.I.; writing—review and editing, M.O., M.I., Y.E.-S., M.E.H., A.B. and A.L.; visualization, M.O., M.I., Y.E.-S., M.E.H., A.B. and A.L.; supervision, M.I., Y.E.-S., M.E.H., A.B. and A.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data supporting the reported results are not publicly available due to privacy or ethical restrictions. However, the source code used for the training, testing, and deployment of the models is available at https://github.com/MouradOuzziki/Snake-ID-DeepLearning.git, accessed on 17 January 2025.

Acknowledgments: We would like to extend our heartfelt thanks to the Moroccan Center for Polytechnical Research and Innovation (CMRPI) for its invaluable support in this research. Its contribution has been essential to the success of our research.

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

References

1. WHO. Snakebite Envenoming. 2023. Available online: https://www.who.int/news-room/fact-sheets/detail/snakebite-envenoming (accessed on 20 October 2024).
2. Moutaouakkil, Y.; El Jaoudi, R.; Tadlaoui, Y.; Adouani, B.; El Cadi, M.A.; Serragui, S.; Cherrah, Y.; Lamsaouri, J.; Bousliman, Y. Envenimations ophidiennes au Maroc: Revue de la littérature. *Toxicol. Anal. Clin.* **2018**, *30*, 246–255. [CrossRef]
3. Kasturiratne, A.; Wickremasinghe, A.R.; De Silva, N.; Gunawardena, N.K.; Pathmeswaran, A.; Premaratna, R.; Savioli, L.; Lalloo, D.G.; De Silva, H.J. The Global Burden of Snakebite: A Literature Analysis and Modelling Based on Regional Estimates of Envenoming and Deaths. *PLoS Med.* **2008**, *5*, e218. [CrossRef] [PubMed]
4. Russell, J.J.; Schoenbrunner, A.; Janis, J.E. Snake Bite Management: A Scoping Review of the Literature. *Plast. Reconstr. Surg. Glob. Open* **2021**, *9*, e3506. [CrossRef] [PubMed]
5. Bouazza, A.; El Mouden, E.; Rihane, A. Checklist of amphibians and reptiles of Morocco: A taxonomic update and standard Arabic names. *Herpetol. Notes* **2021**, *14*, 1–14.
6. Chafiq, F.; Chrouqui, N.; El Jaoudi, R.; Fekhaoui, M.; Soulaymani, A.; Rhalem, N.; Mokhtari, A.; Soulaymani-Bencheikh, R. 247. Envenomation by Daboia mauretanica Snakes in Tiznit Province, Morocco: Report of Four Cases. *Toxicon* **2012**, *60*, 223. [CrossRef]
7. Ouhami, M.; Es-saady, Y.; El Hajji, M.; Hafiane, A.; Canals, R.; Yassa, M.E. Deep Transfer Learning Models for Tomato Disease Detection. In *Image and Signal Processing*; El Moataz, A., Mammass, D., Mansouri, A., Nouboud, F., Eds.; Lecture Notes in Computer Science; Springer International Publishing: Cham, Switzerland, 2020; Volume 12119, pp. 65–73. [CrossRef]

8. Li, X.; Jiang, H.; Ren, Z.; Li, G.; Zhang, J. Deep Learning in Software Engineering. *arXiv* **2018**, arXiv:1805.04825. [CrossRef]
9. Angermueller, C.; Pärnamaa, T.; Parts, L.; Stegle, O. Deep learning for computational biology. *Mol. Syst. Biol.* **2016**, *12*, 878. [CrossRef]
10. Shamshirband, S.; Fathi, M.; Dehzangi, A.; Chronopoulos, A.T.; Alinejad-Rokny, H. A review on deep learning approaches in healthcare systems: Taxonomies, challenges, and open issues. *J. Biomed. Inform.* **2021**, *113*, 103627. [CrossRef]
11. Zhang, H.; Qie, Y. Applying Deep Learning to Medical Imaging: A Review. *Appl. Sci.* **2023**, *13*, 10521. [CrossRef]
12. Umeaduma, L.I. Survey of image classification models for transfer learning. *World J. Adv. Res. Rev.* **2024**, *21*, 373–383. [CrossRef]
13. Kensert, A.; Harrison, P.J.; Spjuth, O. Transfer Learning with Deep Convolutional Neural Networks for Classifying Cellular Morphological Changes. *SLAS Discov.* **2019**, *24*, 466–475. [CrossRef] [PubMed]
14. Gu, C.; Lee, M. Deep Transfer Learning Using Real-World Image Features for Medical Image Classification, with a Case Study on Pneumonia X-ray Images. *Bioengineering* **2024**, *11*, 406. [CrossRef] [PubMed]
15. Amir, A.; Zahri, N.A.H.; Yaakob, N.; Ahmad, R.B. Image Classification for Snake Species Using Machine Learning Techniques. In *Computational Intelligence in Information Systems*; Phon-Amnuaisuk, S., Au, T.W., Omar, S., Eds.; Advances in Intelligent Systems and Computing; Springer International Publishing: Cham, Switzerland, 2017; Volume 532, pp. 52–59. [CrossRef]
16. Patel, A.; Cheung, L.; Khatod, N.; Matijosaitiene, I.; Arteaga, A.; Gilkey, J.W. Revealing the Unknown: Real-Time Recognition of Galápagos Snake Species Using Deep Learning. *Animals* **2020**, *10*, 806. [CrossRef] [PubMed]
17. Binta Islam, S.; Valles, D.; Hibbitts, T.J.; Ryberg, W.A.; Walkup, D.K.; Forstner, M.R.J. Animal Species Recognition with Deep Convolutional Neural Networks from Ecological Camera Trap Images. *Animals* **2023**, *13*, 1526. [CrossRef]
18. Yang, Z.; Sinnott, R. Snake Detection and Classification using Deep Learning. In Proceedings of the 54th Hawaii International Conference on System Sciences, Maui, HI, USA, 4–7 January 2012. [CrossRef]
19. Dandeniya, D.; Wickramasinghe, B.; Dasanayaka, C. A Web-based Application for Snake Species Identification using Vision Transformer and CNN-based Ensemble Meta Classifier. In Proceedings of the 2022 IEEE Pune Section International Conference (PuneCon), Pune, India, 15–17 December 2022; pp. 1–5. [CrossRef]
20. Luiselli, L.; Sale, L.; Akani, G.C.; Amori, G. Venomous Snake Abundance Within Snake Species' Assemblages Worldwide. *Diversity* **2020**, *12*, 69. [CrossRef]
21. Joly, A.; Goëau, H.; Kahl, S.; Deneu, B.; Servajean, M.; Cole, E.; Picek, L.; Ruiz De Castañeda, R.; Bolon, I.; Durso, A.; et al. Overview of LifeCLEF 2020: A System-Oriented Evaluation of Automated Species Identification and Species Distribution Prediction. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction*; Arampatzis, A., Kanoulas, E., Tsikrika, T., Vrochidis, S., Joho, H., Lioma, C., Eickhoff, C., Névéol, A., Cappellato, L., Ferro, N., Eds.; Lecture Notes in Computer Science; Springer International Publishing: Cham, Switzerland, 2020; Volume 12260, pp. 342–363. [CrossRef]
22. Picek, L.; Durso, A.M.; Bolon, I.; de Castaneda, R.L.R. Overview of SnakeCLEF 2021: Automatic Snake Species Identification with Country-Level Focus. In Proceedings of the Conference and Labs of the Evaluation Forum, Bucharest, Romania, 21–24 September 2021.
23. Buslaev, A.; Parinov, A.; Khvedchenya, E.; Iglovikov, V.I.; Kalinin, A.A. Albumentations: Fast and flexible image augmentations. *Information* **2020**, *11*, 125. [CrossRef]
24. Dwaraka Srihith, I.V.; Donald, A.D.; Aditya, T.; Sai Srinivas, T.A.; Thippanna, G.; Lakshmi, V.P. Numeric Metamorphosis: Converting Categorical Features With Python. *J. Adv. Res. Mob. Comput.* **2023**, *5*, 30–35. [CrossRef]
25. Simonyan, K.; Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv* **2015**, arXiv:1409.1556. [CrossRef]
26. Tan, M.; Le, Q.V. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. *arXiv* **2020**, arXiv:1905.11946. [CrossRef]
27. Zhuang, F.; Qi, Z.; Duan, K.; Xi, D.; Zhu, Y.; Zhu, H.; Xiong, H.; He, Q. A Comprehensive Survey on Transfer Learning. *Proc. IEEE* **2021**, *109*, 43–76. [CrossRef]
28. Shi, Z.; Hao, H.; Zhao, M.; Feng, Y.; He, L.; Wang, Y.; Suzuki, K. A deep CNN based transfer learning method for false positive reduction. *Multimed. Tools Appl.* **2019**, *78*, 1017–1033. [CrossRef]
29. Miotto, R.; Wang, F.; Wang, S.; Jiang, X.; Dudley, J.T. Deep learning for healthcare: Review, opportunities and challenges. *Briefings Bioinform.* **2017**, *19*, 1236–1246. [CrossRef] [PubMed]
30. Abou Baker, N.; Zengeler, N.; Handmann, U. A Transfer Learning Evaluation of Deep Neural Networks for Image Classification. *Mach. Learn. Knowl. Extr.* **2022**, *4*, 22–41. [CrossRef]
31. Wightman, R. Pytorch Image Models. 2019. Available online: https://github.com/rwightman/pytorch-image-models (accessed on 22 November 2024).

32. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. *arXiv* **2017**, arXiv:1412.6980. [CrossRef]
33. Hossin, M.; Sulaiman, M.N. A Review on Evaluation Metrics for Data Classification Evaluations. *Int. J. Data Min. Knowl. Manag. Process.* **2015**, *5*, 1–11. [CrossRef]