

Design and Analysis of Artificial Intelligence Model for the Global Issue of Poisonous Reptile Identification

¹Manish Bhardwaj, ²Anamika Singh, ³Manish Kumar, ⁴Sumit Kumar Sharma, ⁵Jaideep Kumar

¹Department of Computer Science and Information Technology, KIET Group of Institutions, Delhi-NCR, Ghaziabad, India

²Department of Applied Science, KIET Group of Institutions, Delhi-NCR, Ghaziabad, India

³Department of Computer Science and Engineering, Ajay Kumar Garg Engineering College, Ghaziabad, India

⁴Department of Information Technology, Ajay Kumar Garg Engineering College, Ghaziabad, India

⁵Department of CSE(IOT), RKGIT, Ghaziabad, India

Abstract

It is estimated that between 80,000 and 140,000 people die year from snakebite envenoming, and another 450,000 become disabled due to this neglected tropical disease. By the year 2030, the World Health Organization hopes to have reduced this load by half. To get there, we need to improve healthcare practitioners' access to up-to-date information and diagnostic tools, as well as fill up the data gap in snake ecology and snakebite epidemiology. First, we need better taxonomic identification of snakes that can cause bites. With the availability of AI-based identification tools for other animals, there is a unique chance to apply machine learning to snake identification and the potentially fatal condition of snakebite envenoming. We trained a deep learning model using 386,006 pictures of snakes, including 198 poisonous and 574 non-venomous species across 188 nations, using the cutting-edge neural network architecture Vision Transformer. We used Flickr and the online biodiversity databases iNaturalist and HerpMapper to compile these images.

For the first time, we demonstrate that AI is capable of correctly categorizing a wide range of snakes, both venomous and nonvenomous, from throughout the world, including similar-looking species from snakebite-prone regions. For snakebite epidemiologists and healthcare practitioners, herpetologists, and the general public, this study lays the groundwork for building global, regional, or national snake identification support systems.

Keywords: World Health Organization, Healthcare, Artificial Intelligence, Computer Vision Technique, Image Processing.

INTRODUCTION

Worldwide, there are around five million cases of snakebite each year. About half of these snake bites occur in low-resource areas of Africa, Asia, and Latin America, and are caused by venomous snakes, which kill 81,000-138,000 people and cripple another 400,000 [1]. In 2019, the World Health Organization (WHO) released a plan to reduce the burden of snakebite envenoming (snakebite henceforth) by half by 2030 through the development of safe, effective, and easily accessible antivenoms. To get there, we need improved knowledge of snake diversity and distribution in snakebite endemic areas [1-3] and more accurate diagnosis of snakebite at the species level. There are more than 3,900 snake species [4] on Earth; over 700 of them are venomous, and 292 are considered MIVS by the WHO.

To guarantee that antivenoms and therapy reach those who need them most, we must accurately identify venomous and non-venomous biting snakes [2, 3, 5-8]. The inaccurate or unnecessary use of

antivenoms wastes this rare and frequently expensive treatment [8, 11], and misidentification of biting snakes can lead to the imprecise, wasteful, and potentially hazardous administration of antivenom.

It is unclear whether or not immunoassays (which identify venom antigens in snakebite victims) and other molecular snake identification techniques will be implemented in low-resource settings [2, 10]. Healthcare practitioners, who are not often trained in herpetology, often struggle with taxonomic identification of the snake when victims or family bring the snake to the health facility or provide a photo [9, 12, 13]. Doctors observe victims for signs of envenomation, from which they can deduce the sort of snake that bit them and decide how best to treat them. Direct snake identification can supplement this syndromic approach, which has its own drawbacks (such as syndromic misclassification) [5, 9, 14].

While many efforts have been made to use computer vision to identify birds (e.g., the Merlin Bird ID app can recognize over 7,500 species) and other animals like fish and butterflies [15], few projects aim to use computer vision to identify snakes, and those that do are often restricted to certain taxonomic groups or geographic areas. We created and evaluated a computer vision model to categorize a wide range of snakes, using hundreds of images of snakes collected from around the world and made available through open biodiversity platforms (iNaturalist and HerpMapper) and another internet resource (Flickr)[19, 20]. We also demonstrated the algorithm's capacity to differentiate between sympatrically similar species and its high average per-country accuracy. We also tested whether the "unreasonable effectiveness" of noisy data for finegrained recognition applies to snake recognition [21, 22] and explored the significance of geographical information in fine-tuning the accuracy of snake species identification.

1. METHODS APPLIED FOR THIS STUDY

Within the context of the snake species identification challenge SnakeCLEF2021, we employed a portion of the world's largest snake photo dataset, which is detailed in full in Durso et al., 2021 [20]. Data-driven challenges in the areas of biodiversity identification and prediction are proposed at LifeCLEF21, which is part of the Conference and Labs of the Evaluation Forum (CLEF) [23].

There are a total of 386,006 images in the training dataset, representing 772 snake species found in 188 nations and every continent except Antarctica. Eighty-seven percent of the images are from iNaturalist and HerpMapper, two online biodiversity databases. We scraped data from Flickr (accounting for 13% of the total) to further expand the dataset, especially for species that were underrepresented in the original. Despite our best efforts, Flickr proved to be an incoherent and potentially misleading data source due to the presence of mislabeled photos and images of non-snake subjects (such as captive snakes and photos of snake habitat).

772 out of 3,921 species (20%) were represented in the training set by at least ten photographs each. With 22,163 pictures representing the most common species (*Thamnophis sirtalis*), and only ten representing the least common, the dataset displays a clear long-tailed class distribution.

The 772 species are divided into 18 families (with 1–418 members in each) and 269 genera (1–32 members in each). Of the 292 known snake species, WHO classifies 198 as MIVS (medically important) [24]; the remaining 574 are not.

To compare our results with those of other studies, we used the test set from the SnakeCLEF2021 competition, which consisted of 23,673 images uploaded to the iNaturalist platform during the first four months of 2021.

Given that all snake species have different, mostly unchanging geographic ranges—a maximum of 126 snake species can reside within the same 50 x 50 km² area [25]—geographical information may play a vital role in determining the correct identification of snake species. To do this analysis, we obtained nation and continental labels for almost 87% of the Data.

In order to facilitate future research into the global performance of the AI model, we developed The Reptile Database-based Mapping Matrix (MM) characterizing country-species presence to enable improved worldwide regularization.

$$MM = \begin{cases} 1, & \text{if species } S \in \text{country } C \\ 0, & \text{else} \end{cases} \quad (1)$$

1.1. The AI Module

The AI module relies heavily on Vision Transformer (ViT), a modern neural network architecture that achieves excellent results in image categorization [30]. The ViT is distinct from convolutional neural networks since it processes images using a conventional Transformer encoder, much like how natural language processing treats a sequence of words.

This section details the entire training and grading process, including the training methodology and image enhancements. To improve the model's performance, we detail the guiding concepts that were applied. All offered metrics can be reproduced thanks to our inclusion of a link to the open-source code, trained checkpoints, and photos.

1.2. Training Techniques

Starting from a publicly available starting point (GitHub), the model underwent two rounds of refinement. The 21.07 NGC Docker container, which houses the PyTorch deep learning framework, was utilized. To ensure that the input resolutions of the pre-trained models were met, all photos were scaled down to 224×224 or 384×384 .

In the first stage, we trained the model for 50 iterations, beginning from the pre-trained ImageNet-1k checkpoint. Each image in the training set was sent 50 times, to put it another way. Stochastic Gradient Descent with a 0.9 momentum was utilized for optimization. Using error rates between the AI model's predictions and the real names of species, the optimization process Stochastic Gradient Descent iteratively adjusts the model's parameters. We scheduled the learning rate using an Adaptive Learning Rate (LR) technique, where the initial LR was set to 0.01 and then decreased by 10% every other epoch (without affecting the validation loss). Softmax Cross Entropy was used to determine the degree of damage. We used a mini-batch size of 256 to accumulate gradients for improved convergence.

The second stage involved using both the training and validation sets to fine-tune the model. Furthermore, we have swapped out the SoftMax Cross Entropy for the hard-focused Focal Loss [31]. This ensures that the most common species are not overrepresented in the training data. The model was then fine-tuned for an additional 20 epochs using the One Cycle Learning Rate Policy provided by Smith et al. (2019) [32].

1.3. Data Cleaning

We incorporated weakly labelled data from Flickr (i.e. data with a relatively high proportion of inaccurate species labels) to improve the number of samples for species with few photos in online biodiversity platforms. In fact, this method is frequently employed to increase the number of samples for rare classes and provide improved performance overall, provided that the number of wrong classifications does not overwhelm the classifier. We trained the ViT-Base/32-224 on the clean (without Flickr data) and full set using the Stage1 training approach to determine if the "unreasonable effectiveness" of the weakly labelled data for fine-grained recognition also applies to snake recognition [21, 22]. Experiments reveal that all performance indicators are enhanced by integrating noisy data for rare species (Table 1). As a result of this, we train the recognition algorithm on the Full set.

Table 1: Two Models Comparison of ViT-Base/32

| | F1-Species | C1 Accuracy-Species | F1 Country | C1 Accuracy Genus | F1-Genus |
|--------------|------------|---------------------|------------|-------------------|----------|
| Clear Set | 69.5% | 82.6% | 68.4% | 90.2% | 72.6% |
| Complete Set | 74.6% | 87.9% | 76.2% | 93.4% | 77.7% |

1.4. Data Augmentation

Several augmentation strategies from the Python Albumentations library [33] were used to prevent overfitting (where the model performs exceptionally well on training data but poorly on test data) and to boost the model's regularization capability. In order to prepare, we utilized:

- RandomResizedCrop randomly selects 70-100% of the image to crop,
- Horizontal Flip, which, with a 50% chance, inverts the image horizontally
- 50% Chance Vertical Flip, which flips the image vertically
- The image is rotated, scaled, and shifted with a 75% chance and upper bounds of 6.25%, 25%, and 45 correspondingly using the randomShiftScaleRotate method.
- On a scale from 0 to 100, JpegCompression has a 50% chance of altering the image quality,
- Blur — Image is blurred with a 7x7 linear filter and a 10% probability.
- RandomBrightnessContrast, which varies the brightness and contrast between -0.2 and 0.2 with a 30% chance,
- HueSaturationValue: modifies the input image's hue, saturation, and value with a 20% chance and arbitrary limits of -20%, -30%, and -20%;
- ImageNormalization: color values are transformed from [0,255] to [1,1] and averaged (0.5) and standard deviation (0.5) over all channels.

Figure 1 Shows the Various examples of Images after the Augmented Technique.

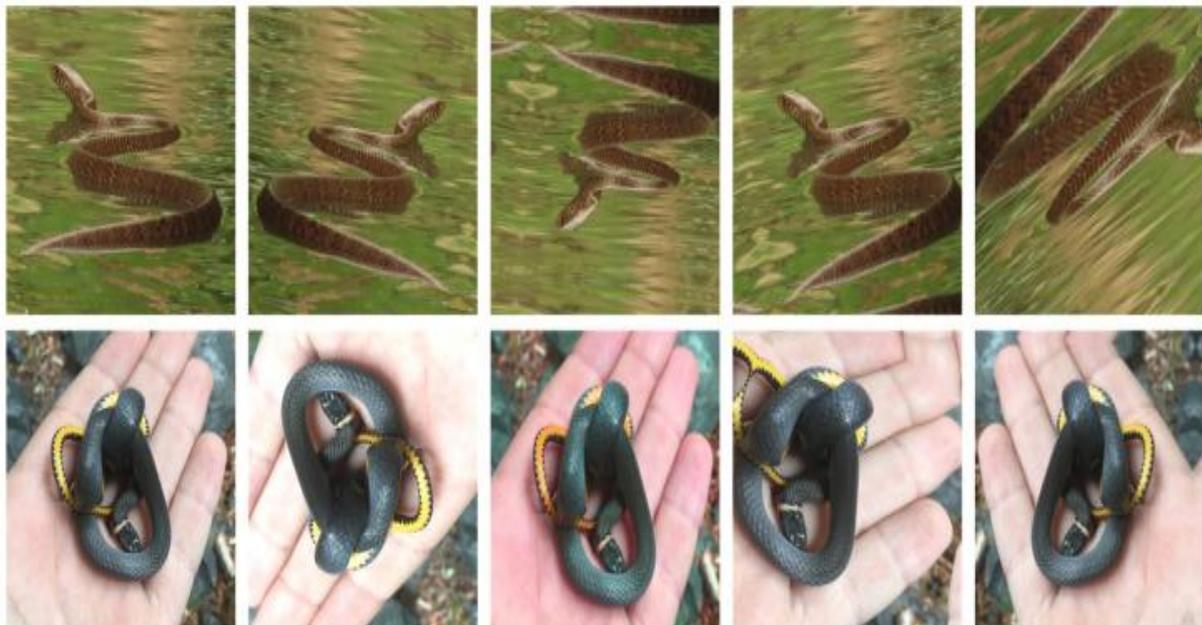


Figure 1: Some Examples of Augmented Images

1.5. Test Time Augmentation

The test-time augmentation (TTA) is a straightforward method for boosting a neural network's efficiency. To obtain the forecast, a set of enhanced photos is generated and fed to the algorithm instead of a single photograph. We make a set of four images: the original, mirrored horizontally and

vertically, and rotated through 180 degrees. The prediction is based on an average of the four possible outcomes.

By exposing the algorithm to the original image at various scales and through varying viewpoints, the recognition performance is enhanced. Table 2 displays the findings indicating the effect of the TTA.

Table 2: Comparison According to the Test Time Augmentation

| | F1-Species | C1 Accuracy-Species | F1 Country | C1 Accuracy Genus | F1-Genus |
|------------------------|------------|---------------------|------------|-------------------|----------|
| Starting Line | 88.6% | 94% | 88.1% | 97.9% | 93.4% |
| Test Time Augmentation | 88.9% | 95.4% | 89.1% | 98.2% | 92.1% |

1.6. Techniques used for the Evaluation

In order to guarantee that global performance is prioritized, we established the Macro F1 Country Performance (Macro F1C) as the primary metric. We determine it by averaging the F1 scores of each country:

$$Macro\ F1_c = \frac{1}{N} \sum_{c=1}^N F1_c, \quad F1_c = \frac{1}{\sum_{s=1}^N MM} \sum_{s=1}^N F1_s \cdot MM \tag{2}$$

Using data from The Reptile Database's August 2020 release [4], we may write: where C is the country index, S is the species index, F1C is the country performance, and MMCS is the mapping matrix representing species-country presence that permits improved worldwide regularization.

The following formula is used for each species to calculate the F1S:

$$F1_s = 2 \cdot \frac{P_s \times R_s}{P_s + R_s}, \quad P_s = \frac{tp_s}{tp_s + fp_s}, \quad R_s = \frac{tp_s}{tp_s + fn_s} \tag{3}$$

We also calculate a Top1 Accuracy (Correct Assessments/All Assessments) and a Macro F1 score to facilitate further in-depth comparisons across scales. The Macro F1 score is the average of the subscores (F1S):

$$Macro\ F1 = \frac{1}{N} \sum_{s=1}^N F1_s \tag{4}$$

where S represents the index of the species and N the total number of species. The final Macro F1 is determined by arithmetically averaging the species-specific F1 scores (Precision (P_s) and Recall (R_s)).

2. OUTPUT OF THE STUDY

2.1. Performance of the Proposed Model

We evaluated the model on a separate dataset consisting of 23,673 images. The model was optimized by applying several transformations to the data (i.e., cropping at random, flipping the data horizontally or vertically, randomly rotating the data, etc.) based on the innovative neural network architecture, Vision Transformer. The method of adding extra time to tests is also applied in a real-world setting. We also added geographical metadata, which greatly improved the system's efficiency and cut the relative error rate by 33.3%.

The algorithm successfully identifies species and genera in testing photos in Table 3. The model's C1-accuracy is 96.2%, and its macro-averaged F1 score is 92.4% (as determined by averaging the F1 scores of all species). The model has a C-1 accuracy of 99.1% and a macro-averaged F1 score of 95% when it comes to genus recognition. The average F1 score across all countries contributes to a 94.3% macro F1 performance.

Table 3: Performance of the Proposed Model

| Test Set | F1-Species | C1 Accuracy-Species | F1 Country | C1 Accuracy Genus | F1-Genus |
|---------------|------------|---------------------|------------|-------------------|----------|
| Starting Line | 92.4% | 96.2% | 94.3% | 99.1% | 95% |

2.2. Relationship Between number of Training Images per Species and F1 Score

The F1 score scales logarithmically with the number of species-specific training photos used. Species that have recently been delineated using molecular methods and geographic location rather than morphological characteristics, and which previously belonged to more widespread species complexes, such as *Agkistrodon piscivorus/A*, tend to be relatively inaccurately identified for their quantity of training data. Figure 2 Provides the distribution of Number of training Images and F1 Score.

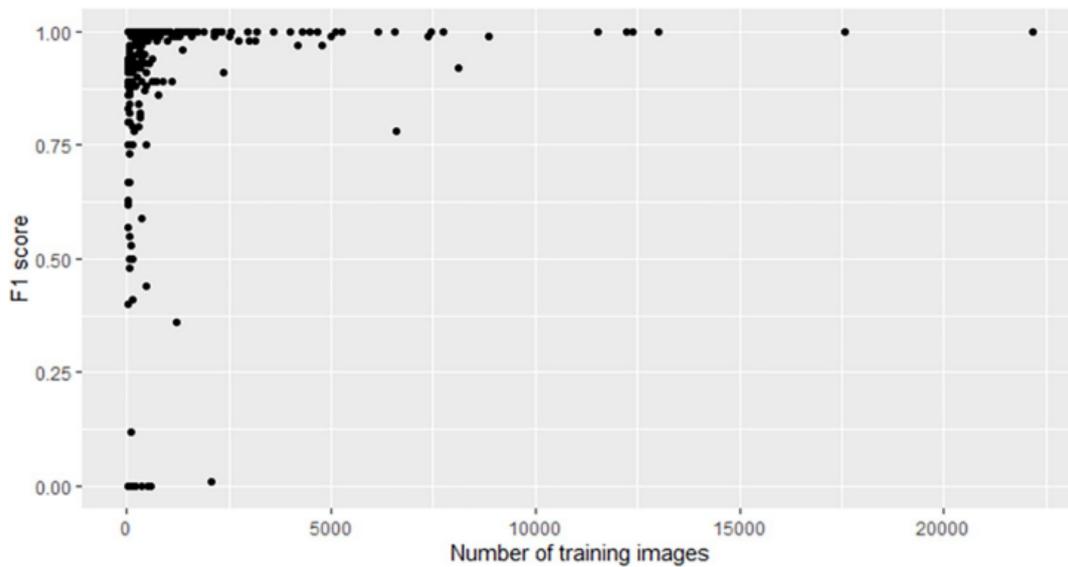


Figure 2: Graphical Representation of F1 with No. of Training image

conanti, *Boa constrictor*, *B. imperator*, and *S. deserticola*. Training images for these species complexes are considerably more likely to be mislabeled due to confusion over how best to separate the putative species, especially from photos that lack geographic locality information, as found by Durso et al. (2021) [34]. The species for which there are less training photographs will likely be misidentified as belonging to a different species, according to probabilistic analysis.

Species from a wide variety of clades and geographies are represented among those with high accuracy despite a limited amount of training data. The Bandy-bandy (*Vermicella annulata*), a burrowing elapid from Australia with striking black and white bands, was trained using just 522 photos at an F1 = 1. There are other banded snakes in Australia [35], but *V. annulata* stands out for being one of the most distinctive Australian snake species, in part because of its bodybridging defensive behavior [36]. This behavior may offer the algorithm with unique traits (e.g., certain postures) to identify. We think that include them would induce at least occasional confusion among them, reducing the F1 for *V. annulata*, because there are five other species in the genus *Vermicella* that are all similar in appearance but none of which fulfilled the minimum criterion of ten training photos. Many additional varied genera are also only represented by a single species that passed the 10-image criterion in our training set.

2.3. Proposed Model Performance at Country Level

Using a list of snake species found in each country, we analyzed how well the model performed at the national level. Above 70% in 97% of countries, 88% in 90% of countries, and 56% in 95% of

countries is the model's F1. Model coverage and performance are trade-offs that shift from region to region. Figure 3 provides the detailed performance of the model applied at the country level.

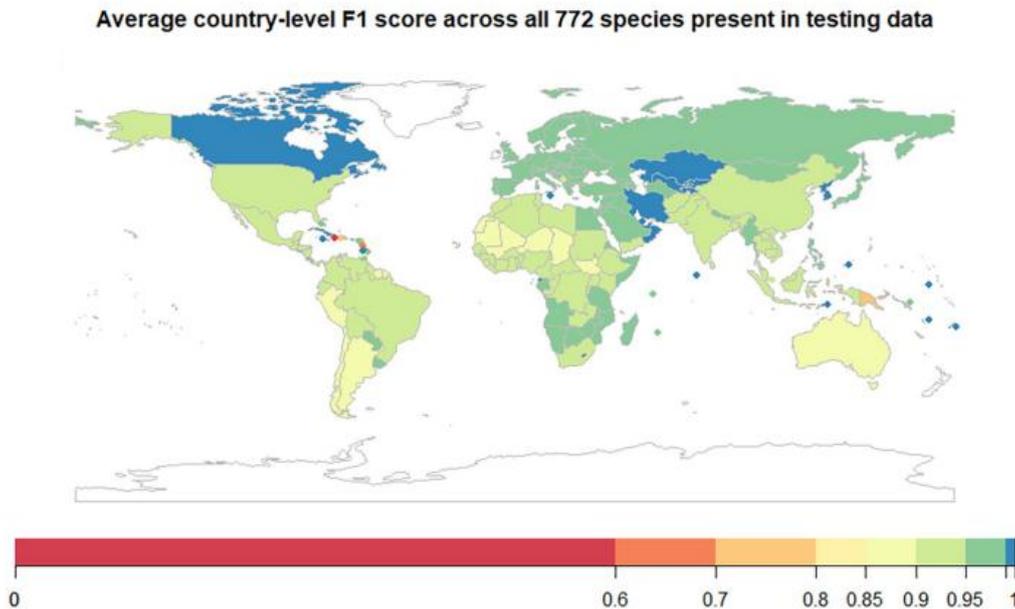


Figure 3: Average country level F1 Score across all 772 species present in testing data

Our analysis of the two continents with the largest snakebite burden revealed that countries in Asia with more coverage have higher average performance, but in Africa, performance peaks at intermediate coverage.

2.4. Comparing model performance between MIVS and non-MIVS species

There were a total of 772 species employed in the training process, 198 MIVS and 574 non-MIVS. We analyzed how well the model could distinguish between MIVS and non-MIVS species. Distributions of F1 scores were similar for MIVS and non-MIVS, with an average of 95 and a standard deviation of 17 for MIVS and 93 and 22 for non-MIVS, respectively ($p = 0.61$ for both comparisons).

3. CONCLUSION

To combat snakebite and other neglected tropical diseases in low-resource settings, we have developed and publicly released an international AI model for autonomous detection of snakes. In order to improve the health and well-being of the world's poorest people, AI's capabilities must be welcomed and used securely and fairly.

Reference

[1] Williams DJ, Faiz MA, Abela-Ridder B, Ainsworth S, Bulfone TC, Nickerson AD, et al. Strategy for a globally coordinated response to a priority neglected tropical disease: Snakebite envenoming. *PLoS Negl Trop Dis.* 2019; 13(2):e0007059. <https://doi.org/10.1371/journal.pntd.0007059> PMID: 30789906

[2] Williams HF, Layfield HJ, Vallance T, Patel K, Bicknell AB, Trim SA, et al. The urgent need to develop novel strategies for the diagnosis and treatment of snakebites. *Toxins.* 2019; 11(6):363. <https://doi.org/10.3390/toxins11060363> PMID: 31226842

[3] Laxme RS, Khochare S, de Souza HF, Ahuja B, Suranse V, Martin G, et al. Beyond the ‘big four’: Venom profiling of the medically important yet neglected Indian snakes reveals disturbing

- antivenom deficiencies. *PLoS Negl Trop Dis*. 2019; 13(12):e0007899. <https://doi.org/10.1371/journal.pntd.0007899> PMID: 31805055
- [4] Uetz P, Freed P, Hos'ek J. The Reptile Database. 2020. [cited 2022 February 02]. Available from: <http://reptile-database.reptarium.cz>.
- [5] Ariaratnam CA, Sheriff MHR, Arambepola C, Theakston RDG, Warrell DA. Syndromic approach to treatment of snake bite in Sri Lanka based on results of a prospective national hospital-based survey of patients envenomed by identified snakes. *Am J Trop Med Hyg*. 2009; 81(4):725–31. <https://doi.org/10.4269/ajtmh.2009.09-0225> PMID: 19815895
- [6] Williams DJ, Habib AG, Warrell DA. Clinical studies of the effectiveness and safety of antivenoms. *Toxicon*. 2018; 150:1–10. <https://doi.org/10.1016/j.toxicon.2018.05.001> PMID: 29746978
- [7] Ruiz de Castañeda R, Durso AM, Ray N, Fernández JL, Williams DJ, Alcoba G, et al. Snakebite and snake identification: empowering neglected communities and health-care providers with AI. *The Lancet Digital Health*. 2019; 1(5):e202–e3. [https://doi.org/10.1016/S2589-7500\(19\)30086-X](https://doi.org/10.1016/S2589-7500(19)30086-X) PMID: 33323268
- [8] Ralph R, Sharma SK, Faiz MA, Ribeiro I, Rijal S, Chappuis F, et al. The timing is right to end snakebite deaths in South Asia. *BMJ*. 2019;364. <https://doi.org/10.1136/bmj.k5317> PMID: 30670457
- [9] Bolon I, Durso AM, Botero Mesa S, Ray N, Alcoba G, Chappuis F, et al. Identifying the snake: First scoping review on practices of communities and healthcare providers confronted with snakebite across the world. *PLoS ONE*. 2020; 15(3):e0229989. <https://doi.org/10.1371/journal.pone.0229989> PMID: 32134964
- [10] Knudsen C, Ju'rgensen JA, Føns S, Haack AM, Friis RU, Dam SH, et al. Snakebite Envenoming Diagnosis and Diagnostics. *Front Immunol*. 2021; 12:1268. <https://doi.org/10.3389/fimmu.2021.661457> PMID: 33995385
- [11] Sharma SK, Alirol E, Ghimire A, Shrestha S, Jha R, Parajuli SB, et al. Acute severe anaphylaxis in Nepali patients with neurotoxic snakebite envenoming treated with the VINS polyvalent antivenom. *J Trop Med*. 2019;2019. <https://doi.org/10.1155/2019/2689171> PMID: 31205473
- [12] Inthanomchanh V, Reyer JA, Blessmen J, Phrasisombath K, Yamamoto E, Hamajima N. Assessment of knowledge about snakebite management amongst healthcare providers in the provincial and two district hospitals in Savannakhet Province, Lao PDR. *Nagoya J Med Sci*. 2017; 79(3):299. <https://doi.org/10.18999/nagjms.79.3.299> PMID: 28878435
- [13] Michael GC, Grema BA, Aliyu I, Alhaji MA, Lawal TO, Ibrahim H, et al. Knowledge of venomous snakes, snakebite first aid, treatment, and prevention among clinicians in northern Nigeria: a cross-sectional multicentre study. *Trans R Soc Trop Med Hyg*. 2018; 112(2):47–56. <https://doi.org/10.1093/trstmh/try028> PMID: 29617989
- [14] WHO. Regional Office for South-East Asia, Guidelines for the management of snakebite 2nd edition 2016. [cited 2022 February 02]. Available from: <https://apps.who.int/iris/handle/10665/249547>.
- [15] Weinstein BG. A computer vision for animal ecology. *J Anim. Ecol*. 2018; 87(3):533–45. <https://doi.org/10.1111/1365-2656.12780> PMID: 29111567
- [16] Amir A, Zahri NAH, Yaakob N, Ahmad RB. Image classification for snake species using machine learning techniques. *International Conference on Computational Intelligence in Information System*; 2016: Springer Cham; 2016. (pp. 52–59)
- [17] Patel A, Cheung L, Khatod N, Matijosaitiene I, Arteaga A, Gilkey JW. Revealing the unknown: real-time recognition of Galápagos snake species using deep learning. *Animals*. 2020; 10(5):806. <https://doi.org/10.3390/ani10050806> PMID: 32384793

- [18] Rajabizadeh M, Rezghi M. A comparative study on image-based snake identification using machine learning. *Sci Rep.* 2021; 11(1):1–16. <https://www.nature.com/articles/s41598-021-96031-1>
- [19] Geneviève LD, Ray N, Chappuis F, Alcoba G, Mondardini MR, Bolon I, et al. Participatory approaches and open data on venomous snakes: A neglected opportunity in the global snakebite crisis? *PLoS Negl Trop Dis.* 2018; 12(3):e0006162. <https://doi.org/10.1371/journal.pntd.0006162> PMID: 29518075.
- [20] Durso AM, Ruiz de Castañeda R, Montalcini C, Mondardini MR, Fernandez-Marques JL, Grey F, et al. Citizen science and online data: Opportunities and challenges for snake ecology and action against snakebite. *Toxicon: X.* 2021:100071. <https://doi.org/10.1016/j.toxcx.2021.100071> PMID: 34278294
- [21] Krause J, Sapp B, Howard A, Zhou H, Toshev A, Duerig T, et al. The unreasonable effectiveness of noisy data for fine-grained recognition. *European Conference on Computer Vision*; 2016: Springer. https://link.springer.com/chapter/10.1007/978-3-319-46487-9_19
- [22] Sun C, Shrivastava A, Singh S, Gupta A. Revisiting unreasonable effectiveness of data in deep learning era. *Proceedings of the IEEE international conference on computer vision*; 2017. https://openaccess.thecvf.com/content_ICCV_2017/papers/Sun_Revisiting_Unreasonable_Effectiveness_ICCV_2017_paper.pdf
- [23] Joly A, Goeau H, Kahl S, Picek L, Lorieul T, Cole E, et al. Overview of lifeclef 2021: An evaluation of machine-learning based species identification and species distribution prediction. *International Conference of the Cross-Language Evaluation Forum for European Languages*; 2021: Springer.
- [24] WHO. Snakebite Information and Data Platform 2021. [cited 2022 February 02]. Available from: https://www.who.int/teams/control-of-neglected-tropical-diseases/snakebite-envenoming/snakebiteinformation-and-data-platform/overview#tab=tab_1.
- [25] Roll U, Feldman A, Novosolov M, Allison A, Bauer AM, Bernard R, et al. The global distribution of tetrapods reveals a need for targeted reptile conservation. *Nat Ecol Evol.* 2017; 1(11):1677–82. <https://doi.org/10.1038/s41559-017-0332-2> PMID: 28993667
- [26] South A. rworldmap: A New R package for Mapping Global Data. *R J.* 2011; Vol. 3(1): 35–43.
- [27] Hierink F, Bolon I, Durso AM, Ruiz de Castañeda R, Zambrana-Torrel C, Eskew EA, et al. Forty-four years of global trade in CITES-listed snakes: Trends and implications for conservation and public health. *Biol Conserv.* 2020; 248:108601. <https://doi.org/10.1016/j.biocon.2020.108601>
- [28] Warrick BJ, Boyer LV, Seifert SA. Non-native (exotic) snake envenomations in the US, 2005–2011. *Toxins.* 2014; 6(10):2899–911. <https://doi.org/10.3390/toxins6102899> PMID: 25268980
- [29] Kraus F. *Alien reptiles and amphibians: a scientific compendium and analysis*: Springer Science & Business Media; 2008.
- [30] Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929.* 2020. [cited 2022 February 02]. Available from: <https://arxiv.org/abs/2010.11929>.
- [31] Lin T-Y, Goyal P, Girshick R, He K, Dollár P. Focal loss for dense object detection. *Proceedings of the IEEE international conference on computer vision*; 2017. https://openaccess.thecvf.com/content_ICCV_2017/papers/Lin_Focal_Loss_for_ICCV_2017_paper.pdf

- [32] Smith LN, Topin N. Super-convergence: Very fast training of neural networks using large learning rates. *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*; 2019: International Society for Optics and Photonics.
- [33] Buslaev A, Iglovikov VI, Khvedchenya E, Parinov A, Druzhinin M, Kalinin AA. Alumentations: fast and flexible image augmentations. *Information*. 2020; 11(2):125.
- [34] Durso AM, Moorthy GK, Mohanty SP, Bolon I, Salathe´ M, Ruiz De Castañeda R. Supervised learning computer vision benchmark for snake species identification from photographs: Implications for herpetology and global health. *Frontiers in Artificial Intelligence*. 2021; 4:17. <https://doi.org/10.3389/frai.2021.582110> PMID: 33959704
- [35] Farooq H, Uetz P. Identifying Australian snakes by color patterns. *Vertebr Zool*. 2020; 70(3):473–82.
- [36] Bustard HR. Defensive display behavior in the bandy-bandy, *Vermicella annulata* (Serpentes: Elapidae). *Herpetologica*. 1969; 25(4):319–20.
- [37] Durso AM, Bolon I, Kleinhesselink A, Mondardini M, Fernandez-Marquez J, Gutsche-Jones F, et al. Crowdsourcing snake identification with online communities of professional herpetologists and avocational snake enthusiasts. *R Soc Open Sci*. 2021; 8(1):201273. <https://doi.org/10.1098/rsos.201273> PMID: 33614073
- [38] Rathnayaka RN, Ranathunga PN, Kularatne S. Paediatric cases of Ceylon krait (*Bungarus ceylonicus*) bites and some similar looking non-venomous snakebites in Sri Lanka: Misidentification and antivenom administration. *Toxicon*. 2021; 198:143–50. <https://doi.org/10.1016/j.toxicon.2021.04.019> PMID: 33957150
- [39] Pandey DP, Pandey GS, Devkota K, Goode M. Public perceptions of snakes and snakebite management: implications for conservation and human health in southern Nepal. *J Ethnobiol Ethnomed*. 2016; 12(1):1–25. <https://doi.org/10.1186/s13002-016-0092-0> PMID: 27255454
- [40] Picek L, Durso A, Ruiz de Castañeda R, Bolon I. Overview of SnakeCLEF 2021: Automatic snake species identification with country-level focus. *Working Notes of CLEF*. 2021.