

RESEARCH ARTICLE

An artificial intelligence model to identify snakes from across the world: Opportunities and challenges for global health and herpetology

Isabelle Bolon^{1*}, Lukáš Píček^{2,3}, Andrew M. Durso⁴, Gabriel Alcoba^{1,5,6}, François Chappuis^{5,7}, Rafael Ruiz de Castañeda^{1,5}

1 Institute of Global Health, Department of Community Health and Medicine, Faculty of Medicine, University of Geneva, Geneva, Switzerland, **2** Department of Cybernetics, Faculty of Applied Sciences, University of West Bohemia, Pilsen, Czechia, **3** PiVa AI s.r.o, Plzeň, Czechia, **4** Department of Biological Sciences, Florida Gulf Coast University, Fort Myers, Florida, United States of America, **5** Division of Tropical and Humanitarian Medicine, Geneva University Hospitals and University of Geneva, Geneva, Switzerland, **6** Médecins Sans Frontières—Doctors Without Borders, Geneva, Switzerland, **7** Department of Community Health and Medicine, Faculty of Medicine, University of Geneva, Geneva, Switzerland

* Isabelle.Bolon@unige.ch



OPEN ACCESS

Citation: Bolon I, Píček L, Durso AM, Alcoba G, Chappuis F, Ruiz de Castañeda R (2022) An artificial intelligence model to identify snakes from across the world: Opportunities and challenges for global health and herpetology. *PLoS Negl Trop Dis* 16(8): e0010647. <https://doi.org/10.1371/journal.pntd.0010647>

Editor: José María Gutiérrez, Universidad de Costa Rica, COSTA RICA

Received: March 10, 2022

Accepted: July 8, 2022

Published: August 15, 2022

Copyright: © 2022 Bolon et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The training and test datasets and the model checkpoints are available via the link <https://indat.mff.cuni.cz/repository/xmlui/handle/20.500.12800/1-4773>. The source code of the model is available via the link <https://github.com/picekl/PLoSNTD-GlobalSnakeID>. We collected Flickr data in November 2018 using a python script (https://github.com/cam4ani/snakes/blob/master/get_flickr_data.ipynb), prior to the point when many users may have deleted their

Abstract

Background

Snakebite envenoming is a neglected tropical disease that kills an estimated 81,000 to 138,000 people and disables another 400,000 globally every year. The World Health Organization aims to halve this burden by 2030. To achieve this ambitious goal, we need to close the data gap in snake ecology and snakebite epidemiology and give healthcare providers up-to-date knowledge and access to better diagnostic tools. An essential first step is to improve the capacity to identify biting snakes taxonomically. The existence of AI-based identification tools for other animals offers an innovative opportunity to apply machine learning to snake identification and snakebite envenoming, a life-threatening situation.

Methodology

We developed an AI model based on Vision Transformer, a recent neural network architecture, and a comprehensive snake photo dataset of 386,006 training photos covering 198 venomous and 574 non-venomous snake species from 188 countries. We gathered photos from online biodiversity platforms (iNaturalist and HerpMapper) and a photo-sharing site (Flickr).

Principal findings

The model macro-averaged F1 score, which reflects the species-wise performance as averaging performance for each species, is 92.2%. The accuracy on a species and genus level is 96.0% and 99.0%, respectively. The average accuracy per country is 94.2%. The model

accounts or moved them to other platforms due to changes in Flickr's policies.

Funding: This work was funded by the Fondation privée des Hôpitaux Universitaires de Genève (grant number QS04-20) <https://www.fondationhug.org/>. The grant was awarded to IB and RRdC. AMD was supported by this grant. LP was supported by the project No. SGS-2022-017 of the University of West Bohemia and by the Technology Agency of the Czech Republic, project No. SS05010008. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

accurately classifies selected venomous and non-venomous lookalike species from South-east Asia and sub-Saharan Africa.

Conclusions

To our knowledge, this model's taxonomic and geographic coverage and performance are unprecedented. This model could provide high-speed and low-cost snake identification to support snakebite victims and healthcare providers in low-resource settings, as well as zoologists, conservationists, and nature lovers from across the world.

Author summary

Snakebite kills 81,000–138,000 people a year and disables another 400,000 mainly in tropical and subtropical Africa, Asia and Latin America. We must be able to correctly identify biting snakes if we want to understand the diversity of snakes that threatens different populations in endemic countries. This is key to improve snakebite epidemiological data and to ensure the appropriate distribution of antivenoms in a given country and the specific treatment of patients with these antivenoms when bitten by prevalent snakes. However, snakes are diverse and healthcare providers lack the expertise to identify them, i.e., victims carry the snake to the hospital or take photos. Here we used thousands of snake photos from across the world and computer vision to develop an AI model to classify snakes. We show for the first time that AI can accurately classify a large diversity of venomous and non-venomous snakes from across the world including lookalike species from snakebite endemic countries. This study sets the foundation for developing global, regional or national snake identification support systems for snakebite epidemiologists and healthcare providers, herpetologists, and the general public.

Introduction

More than five million snakebites occur globally every year. Venomous snakes cause about half of these bites and kill 81,000–138,000 people and disable another 400,000 in low-resource settings in Africa, Asia, and Latin America [1]. In 2019, the World Health Organization (WHO) launched a roadmap to develop safe, effective, and accessible antivenoms and halve snakebite envenoming (snakebite hereafter) burden by 2030 [1]. Achieving this also depends on improving snakebite diagnosis at the snake species level and better understanding of snake diversity and distribution in snakebite endemic areas [1–3]. Globally, there are over 3,900 snake species [4], about 700 of these are venomous and 292 are medically important venomous species (MIVS) according to WHO [1].

We must correctly identify venomous and non-venomous biting snakes to ensure appropriate distribution of antivenoms and treatment of victims where most needed [2, 3, 5–8]. Misidentification of biting snakes can result in imprecise, unnecessary, and potentially unsafe use of antivenom [9, 10], while unnecessary or incorrect use of antivenoms wastes this scarce and often expensive treatment [8, 11].

Molecular snake identification techniques (e.g., immunoassays that detect venom antigens in snakebite victims) have limitations, and their deployment in low-resource settings remains to be seen [2, 10]. When victims or relatives bring the snake to the health facility or provide a photo, healthcare providers, who are generally not trained in herpetology, often struggle to

taxonomically identify the snake [9, 12, 13]. They watch for victims' symptoms to determine the type of envenomation, infer the associated biting snake, and subsequently decide on the treatment. This syndromic approach has limitations (e.g. syndromic misclassification) and can be complemented by direct snake identification [5, 9, 14].

Computer vision techniques have been developed to identify birds (e.g., Merlin Bird ID app recognizes over 7,500 species) and other animals like fish and butterflies [15], yet few initiatives seek to identify snakes and they are limited to certain taxonomic groups or geographic areas (e.g., [16–18]). In this study, we developed and tested the performance of a computer vision model to classify a large diversity of snakes using thousands of snake photos from across the world publicly available on open biodiversity platforms (iNaturalist and HerpMapper) and another online platform (Flickr) [19, 20]. We further showed the high average per-country accuracy of this algorithm and its capability to distinguish sympatric lookalike species. We also investigated the role of geographical information in fine-tuning the accuracy of snake species identification and if the phenomenon of “unreasonable effectiveness” of noisy data for fine-grained recognition applies to snake recognition [21, 22].

Methods

Snake photo datasets

We used a subset of the world's largest snake photo dataset, described in detail in Durso et al., 2021 [20], which we provided within the snake species identification challenge Snake-CLEF2021. This challenge is part of LifeCLEF21, the Conference and Labs of the Evaluation Forum (CLEF) that proposes data-oriented challenges related to the identification and prediction of biodiversity [23].

The training dataset contains 386,006 photos belonging to 772 snake species from 188 countries and all continents except Antarctica. Most of the photos (87%) come from the online biodiversity platforms iNaturalist (www.inaturalist.org) and HerpMapper (www.herpmapper.org). For species with the fewest images, we further extended the dataset by scraping data from Flickr (13% of the total). Although we attempted to reduce this, Flickr provided us with a heterogeneous and noisy source of data, including imprecise photo labels and photos of non-snake subjects (e.g., captive snakes, photos of snake habitat).

Within the training set, 772 of the world's 3,921 species ($\pm 20\%$) had at least ten photos, our threshold for inclusion. The dataset has a marked long-tailed class distribution, where the most frequent species (*Thamnophis sirtalis*) is represented by 22,163 photos and the least frequent by just ten.

The 772 species are classified into 18 families (1–418 species per family in the training set) and 269 genera (1–32 species per genus in the training set). They include 198 out of the 292 (68%) medically important venomous snake species (MIVS) according to WHO's classification [24] and 574 non-MIVS.

For testing, we used the test set from the SnakeCLEF2021 competition with 23,673 photos submitted to the iNaturalist platform within the first four months of 2021, allowing us to compare the performance with other studies (see S1 and S2 Tables for more details on the training and testing datasets and S3 Table and S1 and S2 Figs for MIVS included in the datasets).

Considering that all snake species have distinct, largely stable geographic ranges, with a maximum of 126 species of snakes occurring within the same $50 \times 50 \text{ km}^2$ area [25], geographical information could play a crucial role in correct snake species identification. To evaluate this, we gathered two levels of geographical label (i.e., country and continent) for approximately 87% of the data.

The vast majority (77%) of photos came from the United States and Canada, with 9% from Latin America and the Caribbean, 5.7% from Europe, 4.5% from Asia, 1.8% from Africa, and 1.5% from Australia/Oceania. Bias at smaller spatial scales undoubtedly exists as well [20], largely due to where participants in citizen science projects and other snake photographers are concentrated. Nevertheless, snake species from nearly every country were represented, with 46/215 (21%) of countries having all of their snake species represented, mostly in Europe. Nearly half of all countries (106/215; 49%) had more than 50% of their snake species represented (Fig 1).

A small proportion of photos (ca. 1–2%), particularly from Flickr, show captive snakes that are kept outside of their native range (e.g., North American *Pantherophis guttatus* in Europe or Australian *Morelia viridis* in the USA). We opted to retain these for three reasons:

1. Users of an automated identification system may wish to use it on captive snakes (e.g., in the case of customs seizures) [27].
2. Bites from captive snakes may occur (although the identity of the snake would normally be clear in this case) (e.g., [28]).
3. Captive snakes sometimes escape and can establish introduced populations outside their native range (e.g., [29]).

To support further studies related to the worldwide performance of the AI model, we created a Mapping Matrix (MM) describing country-species presence to allow better worldwide regularization, based on The Reptile Database [4].

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

The Artificial Intelligence module

The main building block of the AI module is a recent neural network architecture—Vision Transformer (ViT)—with state-of-the-art image categorization performance [30]. Apart from the convolutional neural networks, the ViT avoids convolutional layers while interpreting an image as a sequence of patches and processing it by a standard Transformer encoder like natural language processing.

This section describes the full training and evaluation procedure, including the training strategy and image augmentations. We include the description of used principles that helped to increase the model performance. We include the link to the open-sourced code, trained checkpoints and images, allowing reproducibility for all provided metrics.

Training strategy

The model was initialized from a publicly available checkpoint (GitHub) and further fine-tuned in two consecutive stages. The PyTorch deep learning framework within the 21.07 NGC Docker container was used. All images were resized to the input size of 224 x 224 or 384 x 384 to match the input resolutions of the pre-trained models.

Stage1: Starting from the ImageNet-1k pre-trained checkpoint, we trained the model for 50 epochs. In other words, each image in the training set was feed-forwarded 50 times. For optimization, Stochastic Gradient Descent with momentum set to 0.9 was used. Stochastic Gradient Descent is an optimization algorithm that iteratively modifies the AI model parameters by measuring discrepancies between predicted and correct species names. We used an Adaptive Learning Rate (LR) strategy to schedule the learning rate, i.e., starting LR of 0.01 was reduced by 10% on every second epoch without validation loss reduction. The loss was calculated as

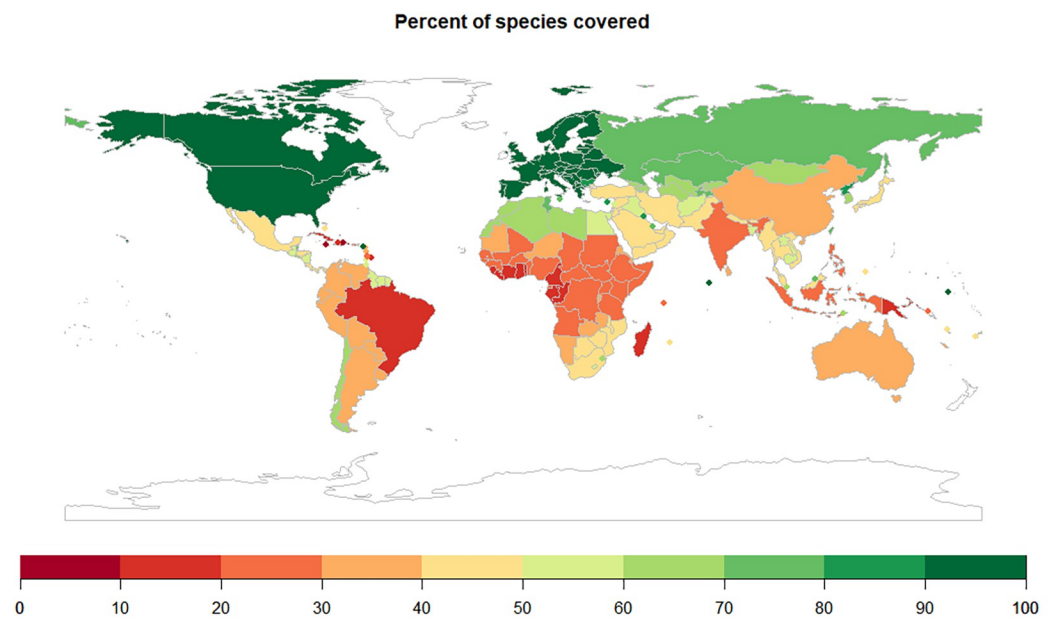


Fig 1. Percent of snake species represented in the training dataset for each country. The map was drawn using R package `rworldmap` v 1.3.6 [26]. See [S1 Data](#) for underlying values.

<https://doi.org/10.1371/journal.pntd.0010647.g001>

Softmax Cross Entropy. To allow better convergence, we accumulated gradients to match a mini-batch size of 256.

Stage2: We used both the training and validation set to fine-tune the model in the second stage. In addition, we have substituted the SoftMax Cross Entropy with the Focal Loss that focuses on the hard examples [31]. With that, we prevent the common species with the most samples from overwhelming the model during training. Next, we used the One Cycle Learning Rate Policy, proposed by Smith et al. (2019) [32], to fine-tune the model for an additional 20 epochs.

Data cleaning

To increase the number of samples for species with few images in online biodiversity platforms, we added weakly labelled data from Flickr (i.e., data with a relatively high number of incorrect species labels). This procedure is commonly used in practice to maximize the number of samples for rare classes and to allow better performance overall as long as the number of incorrect labels does not overwhelm the classifier. To test if the “unreasonable effectiveness” of the weakly labelled data for fine-grained recognition applies to snake recognition as well [21, 22], we trained the ViT-Base/32–224 on the clean (without Flickr data) and full set using the Stage1 training strategy. The experiment results show that including noisy data for rare species improves the performance in all measured metrics (Table 1). Learning from that, we use the Full set to develop the recognition system.

Table 1. Performance comparison for two ViT-Base/32 models trained on the SnakeCLEF2021 dataset and its “clean” subset.

	F1 Country	F1—Species	Top1 Accuracy—Species	F1—Genus	Top1 Accuracy—Genus
Clean set	68.6%	69.7%	82.3%	72.5%	90.0%
Full set	75.9%	74.2%	88.2%	77.9%	93.4%

<https://doi.org/10.1371/journal.pntd.0010647.t001>

Data augmentations

To prevent overfitting—a state of the model when performing nearly perfect on training data but badly on test data—and to increase the regularization capability of the model, we utilized several augmentation techniques from the Python Albumentations library [33] (Fig 2). During training, we used:

- *RandomResizedCrop*: randomly crops 70–100% from the original image,
- *Horizontal Flip*: flips the image with 50% probability,
- *Vertical Flip*: flips the image with 50% probability,
- *RandomShiftScaleRotate*: shifts, scales and rotates the image with 75% probability, and upper limits of $\pm 6.25\%$, $\pm 25\%$, and $\pm 45^\circ$ respectively,
- *JpegCompression*: changes the image quality with 50% probability on a scale of 50–100,
- *Blur*: using a 7x7 linear filter to blur image with 10% probability,
- *RandomBrightnessContrast*: adjusts the contrast and brightness with 30% probability by a random factor in a range -0.2–0.2,
- *HueSaturationValue*: changes hue, saturation and value of the input image with 20% probability and random limits of -20–20, -30–30, and -20–20%, respectively, and
- *ImageNormalization*: colour values are re-scaled from 0 – 255 to 0 – 1, and normalized by mean (0.5) and std (0.5) in all channels.




Fig 2. Examples of augmented images. The first image from the left is the original image. Top: image by chiuluan, iNaturalist 207060926 (CC-BY); Bottom: image by Alex Karasoulos, iNaturalist 207359674 (CC-BY).

<https://doi.org/10.1371/journal.pntd.0010647.g002>

Test-time augmentation

The test-time augmentation (TTA) is a simple procedure commonly used to improve the performance of the neural network. Instead of feeding the model with a single photograph to get the prediction, a batch of augmented images is created and feed-forwarded. In our case, we create a batch of 4 with (i) original image, (ii) horizontally flipped image, (iii) vertically flipped image, and (iv) image rotated by 180° . All four probabilities are averaged to get the prediction.

Such a procedure helps to improve the recognition performance by allowing the algorithm to see the original image on different scales or observed from different angles. The results showing the impact of the TTA are presented in Table 2.

Table 2. Test time augmentations experiment.

	F1 Country	F1—Species	Top1 Accuracy—Species	F1—Genus	Top1 Accuracy—Genus
Baseline	91.1%	88.8%	94.1%	93.2%	98.4%
Test-time augmentation	91.3%	89.1%	95.2%	92.5%	98.5%

<https://doi.org/10.1371/journal.pntd.0010647.t002>

Exploiting geographic information

We tested and utilised two approaches for the integration of geographic information—binary masking, which automatically removes all the irrelevant species for a country from the prediction and a simple statistical approach based on the assumption that the class posterior given the image I and location L can be estimated as:

$$P(S|I, L) = \frac{P(S|I) \cdot P(S|L)}{P(S)}, \tag{2}$$

where p(S) is the species prior in the training set, and the conditional probability P(S|L) is calculated as the relative frequency of species S within the given location.

The first scenario, where we removed all the irrelevant SoftMax values based on the species-presence knowledge, helped us to increase the performance by a significant margin; 2.93% and 3.11% in F1 Country and F1—Species, respectively (Table 3).

Evaluation Protocol

To assure focus on worldwide performance, we defined the macro F1 country performance (Macro F1_C) as the main metric. We calculate it as the mean of country F1 scores:

$$Macro\ F1_c = \frac{1}{N} \sum_{c=1}^N F1_c, \ F1_c = \frac{1}{\sum_{s=1}^N MM_{cs}} \sum_{s=1}^N F1_s \cdot MM_{cs} \tag{3}$$

where C is country index, S is species index, F1_C is the country performance, and MM_{CS} is the mapping matrix describing species-country presence that allows better worldwide regularization; extracted from the August 2020 release of The Reptile Database [4].

To get the F1_S we use the following formula for each species:

$$F1_s = 2 \cdot \frac{P_s \times R_s}{P_s + R_s}, \ P_s = \frac{tp_s}{tp_s + fp_s}, \ R_s = \frac{tp_s}{tp_s + fn_s} \tag{4}$$

To allow deeper comparison on different levels, we also measure the Top1 Accuracy ($\frac{Correct\ Assessments}{All\ Assessments}$) and the Macro F1 score. The Macro F1 score is calculated as the mean of all F1_S scores:

$$Macro\ F1 = \frac{1}{N} \sum_{s=1}^N F1_s, \tag{5}$$

where S is the species index and N the number of species. Final Macro F1 is calculated by

Table 3. Achieved performance on the SnakeCLEF2021 test set using different locational data and metadata integration method.

	F1 Country	F1—Species	Top1 Accuracy—Species	F1—Genus	Top1 Accuracy—Genus
Model with Test-time augmentation	91.3%	89.1%	95.2%	92.5%	98.5%
Country Prior	90.0%	89.1%	95.4%	91.5%	98.6%
Continent Prior	93.2%	90.7%	95.1%	95.6%	98.9%
Presence Masking	94.2%	92.2%	96.0%	94.9%	99.0%

<https://doi.org/10.1371/journal.pntd.0010647.t003>

computing the F1 score for each species as the harmonic mean of the species Precision (P_S) and the Recall (R_S).

Results

Model overall performance

We tested the performance of the model on an independent dataset comprising 23,673 photos. The model is based on the novel neural network architecture—Vision Transformer—and was optimized using various data augmentations (i.e., random crop, horizontal/vertical flip, random rotation, etc.). In addition, the test-time augmentation procedure is used in the production environment. We also incorporated geographic metadata information, which increased the system's performance by a significant margin, reducing the relative error rate by 33.3%.

The model accurately classifies testing images at the species and genus levels (Table 4). The model macro-averaged F1 score, calculated as the mean of all species F1 scores, is 92.2% and the top-1 accuracy is 96.0%. For genus recognition, the model achieves a macro-averaged F1 score of 94.9% and a 99.0% top-1 accuracy. The macro F1 country performance, calculated as the mean of country F1 scores, is 94.2%.

Table 4. Overall performance of the model.

Test Set	F1 Country	F1—Species	Top1 Accuracy—Species	F1—Genus	Top1 Accuracy—Genus
SnakeCLEF2021	94.2%	92.2%	96.0%	94.9%	99.0%

<https://doi.org/10.1371/journal.pntd.0010647.t004>

Relationship between number of training images per species and F1 score

There is a logarithmic relationship between the number of training images per species and the F1 score (Fig 3). Species that stand out as being relatively inaccurately identified for their quantity of training data are mostly those which have been recently delineated primarily using molecular methods and geographic location rather than morphological characteristics, and previously belonged to more widespread species complexes, such as *Agkistrodon piscivorus*/A.

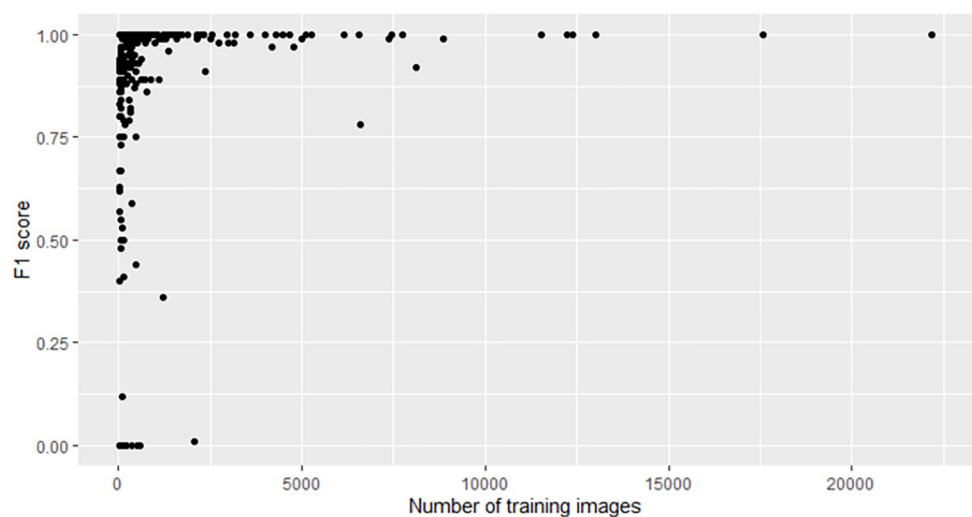


Fig 3. Distribution of F1 scores according to the number of training images per species. See S2 Data for underlying values.

<https://doi.org/10.1371/journal.pntd.0010647.g003>

conanti, *Boa constrictor*/*B. imperator*, and *Salvadora hexalepis*/*S. deserticola*. Durso et al. (2021) [34] found the same and noted that training images for these species complexes are much more likely to be mislabeled due to confusion over how best to differentiate the putative species, especially from photos that lack geographic locality information. From the probabilistic point of view, the species with fewer training images will normally be mistaken for those with more.

Species having good accuracy despite relatively low quantity of training data include diverse species from different clades and continents. For instance, Bandy-bandy (*Vermicella annulata*), an Australian burrowing elapid with bold black-and-white bands, had an F1 = 1 and only 522 training images. Although there are other banded snakes in Australia [35], *V. annulata* is among the most distinctive Australian snake species and is also notable for its body-bridging defensive behaviour [36], which might provide unique features (e.g. particular postures) that the algorithm could detect. However, there are five other species in the genus *Vermicella*, all of which are similar in appearance but none of which met the minimum threshold of ten training images, and we suspect that including these would produce at least occasional confusion among them, lowering the F1 for *V. annulata*. The same applies to many other diverse genera represented in our training data by only one species that met the 10-image threshold.

Country-level performance of the model

We assessed the country-level performance of the model, taking into account the list of snake species occurring in each country. The model's country-wise F1 is above 70% in 97% of countries, above 90% in 88% of countries, and above 95% in 56% of countries (Fig 4).

There is a trade-off between the model coverage and performance that varies by region. Focusing on the two continents with the highest snakebite burden, we found that, in Asia, countries with higher coverage have higher performance on average, whereas in Africa, performance peaks at intermediate coverage (Fig 5).

Comparing model performance between MIVS and non-MIVS species

The 772 species used to train the model include 198 MIVS and 574 non-MIVS species. We compared the model performance in identifying MIVS and non-MIVS species (Fig 6). The model performed equally at identifying MIVS and non-MIVS, with a similar F1 score distribution for each category (average \pm S.D. F1 = $95 \pm 17\%$ for MIVS vs $93 \pm 22\%$ for non-MIVS; $p = 0.61$).

Model performance in distinguishing snake lookalikes

The training dataset includes many similar-looking species, including several non-MIVS species that have evolved to mimic MIVS species. These species usually occur in the same geographic area as each other and resemble one another so closely that they are often confused, even by experts [9, 37]. We assessed the model performance in distinguishing selected look-alike species groups occurring in Southeast Asia (Fig 7) and sub-Saharan Africa (Fig 8).

In Southeast Asia, one of the most difficult medically-relevant snake identification challenges is telling apart venomous kraits (genus *Bungarus*) from non-venomous wolfsnakes (genus *Lycodon*). Confusion between these two groups of snakes has led to unnecessary or delayed use of antivenom, resulting in victim death [5, 9, 38]. Whereas >91% of community members in southern Nepal confused these two groups [39], our model accurately classified all testing photos at the genus level (although it did confuse some species within each genus). Several species in each genus remain unrepresented in the training data.

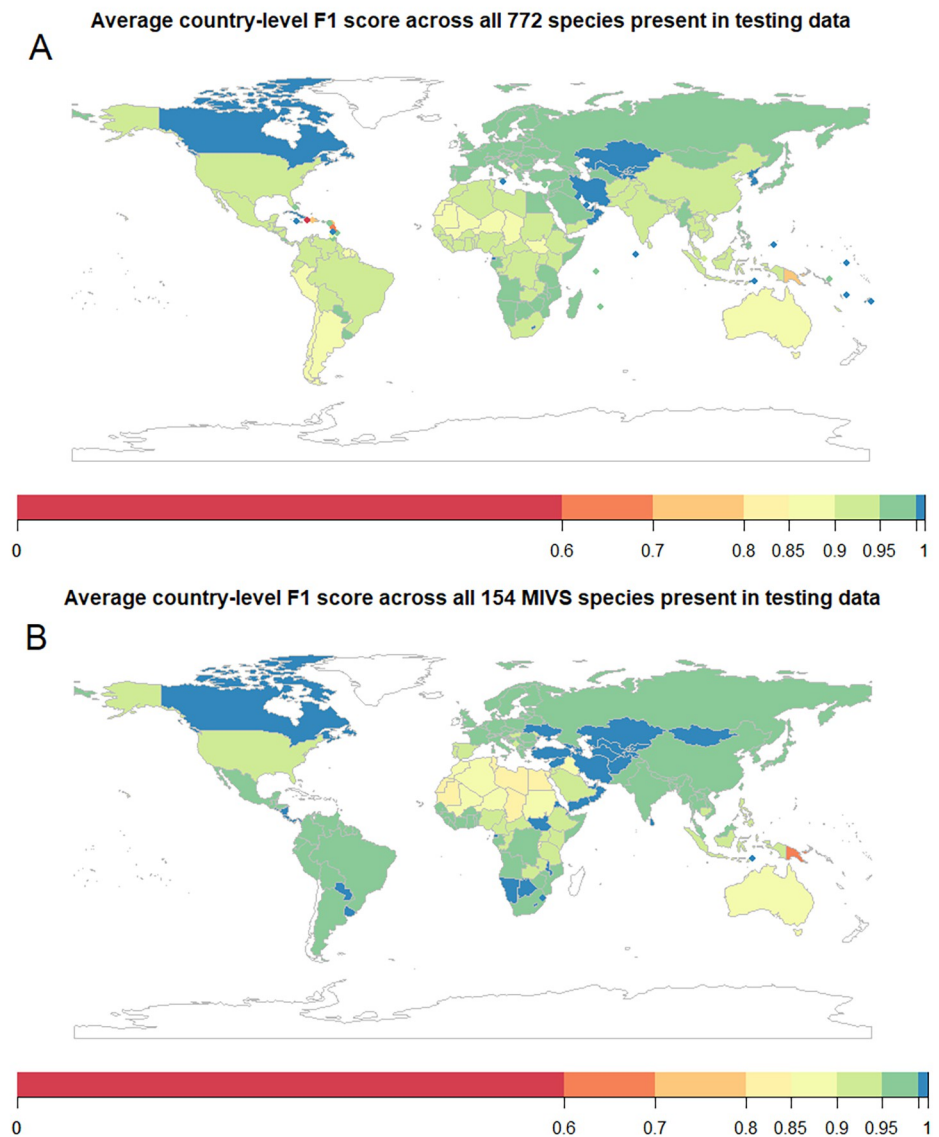


Fig 4. Maps of average F1 score by country (A All species, B MIVS only). Countries and territories with the lowest overall F1 scores for all species are Martinique (36%), Haiti (40%), Dominica (62.5%), St. Lucia (62.5%), Aruba (68%), and Papua New Guinea (74.8%). The map was drawn using R package rworldmap v 1.3.6 [26]. See [S3](#) and [S4 Data](#) for panels A and B underlying values.

<https://doi.org/10.1371/journal.pntd.0010647.g004>

In Sub-Saharan Africa, similar-looking venomous genera (*Bitis* and *Echis*) and mildly venomous (*Causus*) or non-venomous (*Dasypeltis*) lookalikes are frequently confused [37]. Our model accurately classified all testing images at the species level. However, few testing photos were available for many taxa, and several species in each genus remain unrepresented in the training data.

Discussion

This study shows for the first time that AI (i.e., computer vision and fine-grained image classification) can identify a large diversity of venomous and non-venomous snakes at a high level




Fig 5. Scatterplot of Asian and African countries by coverage and F1 score. See S5 Data for underlying values. For other continents, see S3 Fig. and interactive online version at <https://chart-studio.plotly.com/~amdurso/6/#/>.

<https://doi.org/10.1371/journal.pntd.0010647.g005>

of accuracy. The taxonomic and geographic coverage and the performance of the computer vision model proposed are unprecedented: 772 snake species and 269 genera from 188 countries on six continents and an F1-macro score of 92.2%, with a species and genus level accuracy of 96.0% and 99.0%, respectively. This study sets the foundation for developing global, regional or national snake identification systems for herpetologists, epidemiologists and healthcare providers, and the general public, and highlights the importance of AI, open data, and crowdsourcing to tackling snakebite.

For developing this new AI model, we used snake photos from nearly every country (188 out of 215). This global approach contrasts with previous studies focused on a specific country

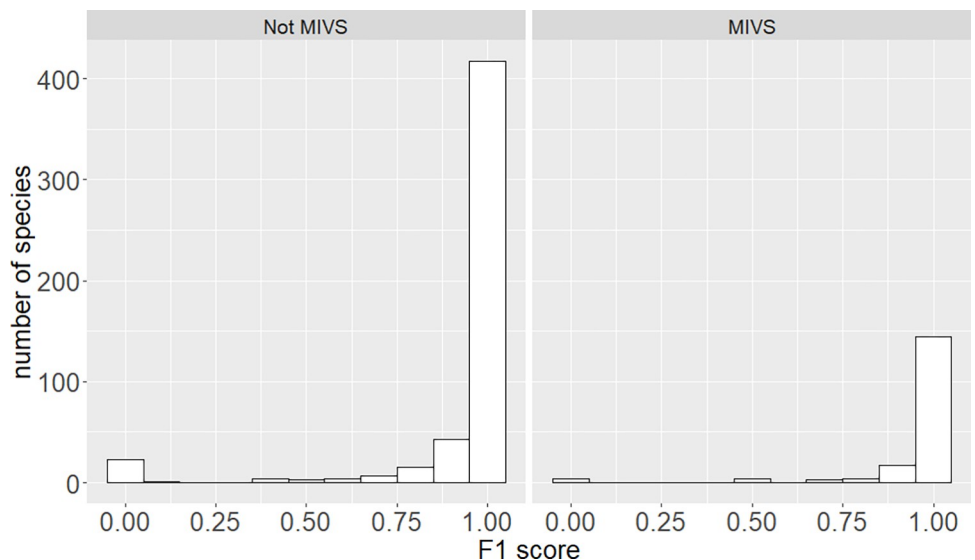


Fig 6. Comparison of F1 score distribution between non-MIVS (574 species) and MIVS (198 species). See S6 Data for underlying values.

<https://doi.org/10.1371/journal.pntd.0010647.g006>

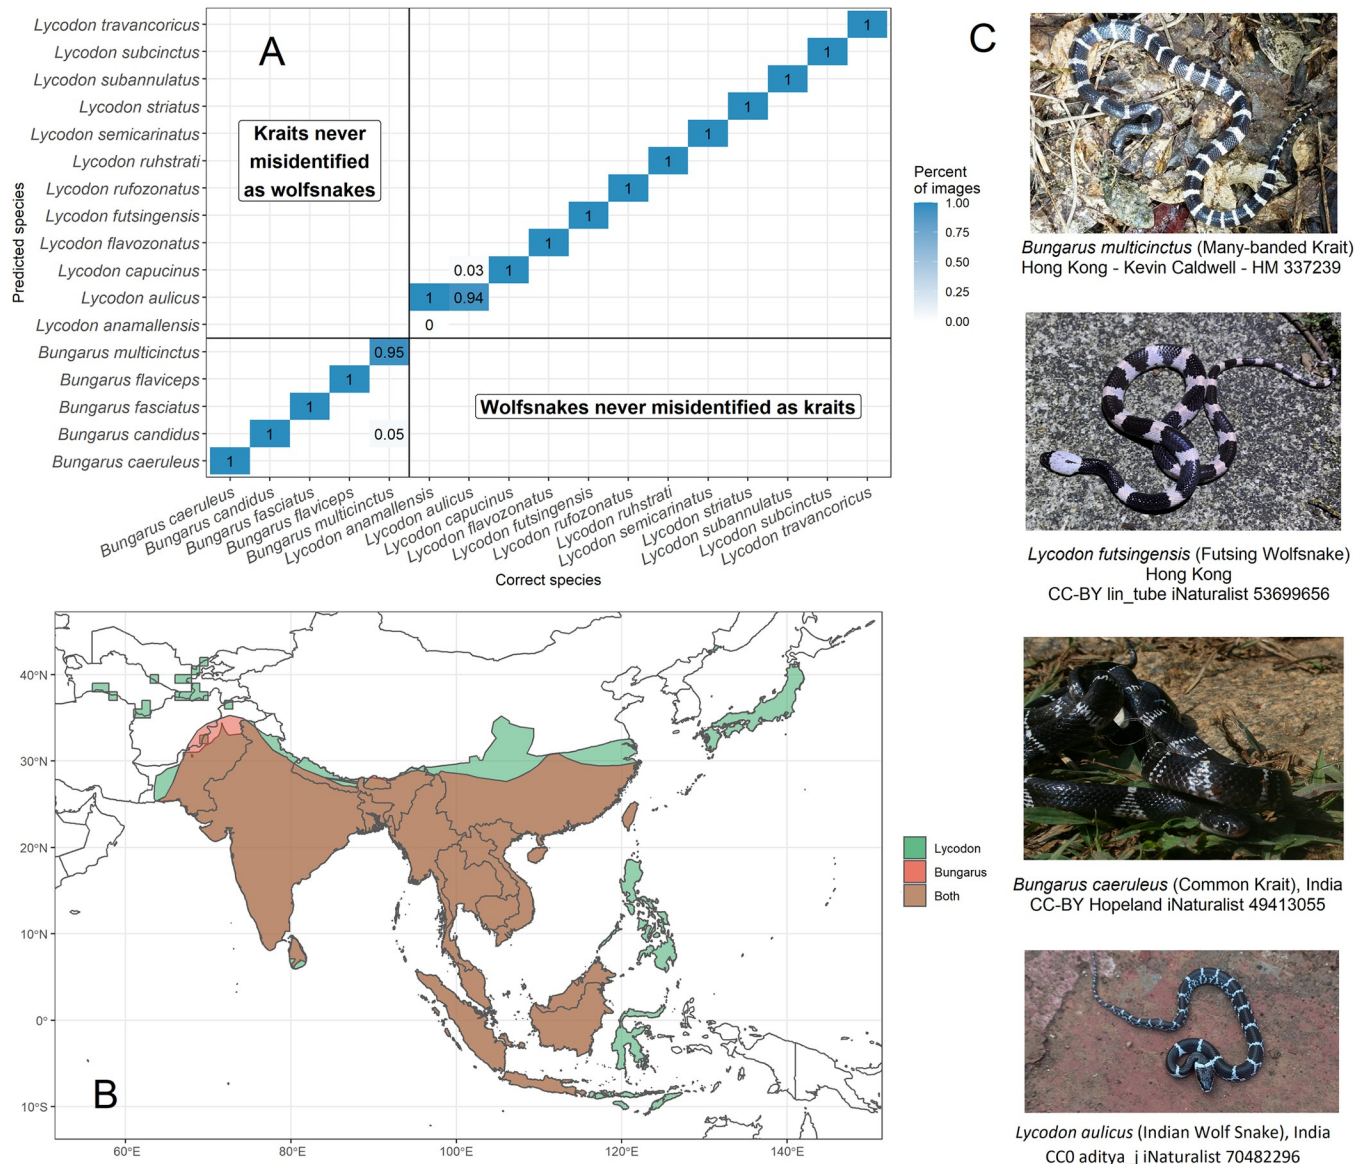


Fig 7. Computer vision model performance for identifying lookalike venomous and non-venomous snake species in Southeast Asia (*Bungarus spp.* vs. *Lycodon spp.*). All species in these genera from the training data were included. **A** Confusion matrix of the classification results of the model on lookalike snake species with the percentage (blue colour code) of correctly (diagonal cells) and incorrectly (off-diagonal cells) classified images, **B** Geographic range of the relevant snake genera based on [25], **C** Representative photos of some of the lookalike snake species tested with the model. The map was drawn using R package rworldmap v 1.3.6 [26]. See S7 Data for panel A underlying values.

<https://doi.org/10.1371/journal.pntd.0010647.g007>

(India, Indonesia, Iran, Malaysia, or Sri Lanka) (e.g., [16, 18]), a geographic area (Galapagos Islands) [17], or a limited number of species (e.g., 6 in Rajabizadeh et al. (2021) [18], 9 in Patel et al. (2020) [17], 22 in Amir et al. (2016) [16]). AI-based snake photo classification is challenging, mainly because of the large diversity of snakes and the often-limited number of photos per class. Thus collaborative approaches and open data sharing, like in biodiversity platforms such as HerpMapper and iNaturalist, are essential. These platforms also provide promising AI models for classifying snakes, although they do not deliver snake-specific performance metrics.

We developed part of our model collaborating with an international community of AI scientists via AICrowd and the SnakeCLEF2021 challenge [40], and building on solutions for

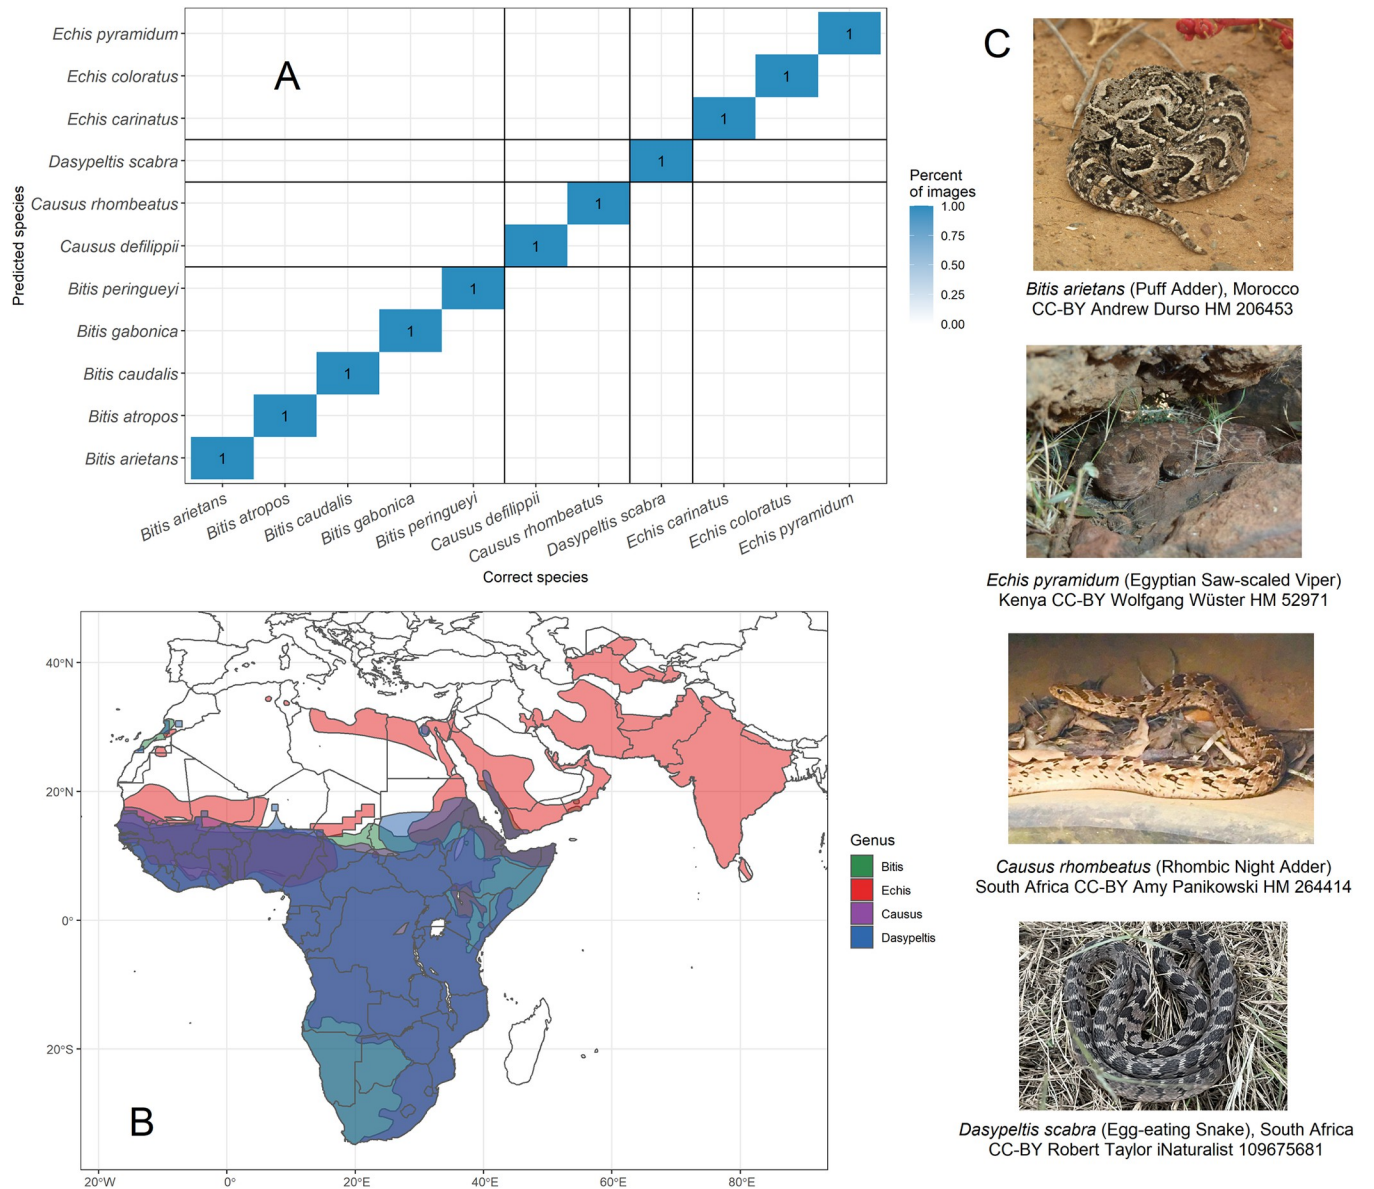


Fig 8. Computer vision model performance for identifying lookalike venomous and non-venomous snake species in Sub-Saharan Africa (*Bitis spp.* vs. *Echis spp.* vs. *Causus spp.* vs. *Dasypeltis spp.*). All species in these genera from the training data were included. **A** Confusion matrix of the classification results of the model on lookalike snake species with the percentage (blue colour code) of correctly (diagonal cells) and incorrectly (off-diagonal cells) classified images, **B** Geographic range of the relevant snake genera based on [25], **C** Representative photos of some of the lookalike snake species tested with the model. The map was drawn using R package rworldmap v 1.3.6 [26]. See S8 Data for panel A underlying values.

<https://doi.org/10.1371/journal.pntd.0010647.g008>

classifying fungi [41]. Our model is based on Vision Transformer, the state-of-the-art deep neural network, and uses simple and replicable training procedure and unique geographic data exploitation. With an F1 Country of 94.2%, F1 species of 92.2%, and Top1 accuracy of 96.0%, there is, to our knowledge, no other AI-based system offering similar performance, even for a smaller number of snake species.

The average model accuracy per country is above 90% in 88% of countries. Accurate taxonomic identification of snakes is critical to improving snake ecology and snakebite epidemiology data in endemic countries. Understanding the diversity of snakes and which snakes bite

locals is the basis for more precise and cost-effective distribution of often-limited national stock of antivenoms. This is also the rationale for WHO's "Snakebite Information and Data Platform" and the associated sub-Saharan Africa antivenom stockpile programme [24].

Snakebite is a health emergency requiring rapid transfer of the victim to a healthcare facility (e.g., neurotoxic envenomation can produce generalized paralysis, respiratory arrest, and death in few hours). Precise and rapid snake identification could help healthcare providers anticipate victim signs and symptoms, complement the commonly used syndromic approach, make good choices on the use of antivenom or plan victim transfers depending on antivenom availability. Accurate snake species identification is particularly important in snakebite endemic countries and regions where species-specific monovalent antivenoms are used to treat victims. Our AI model distinguishes lookalike venomous and non-venomous species often confused in clinical practice with potentially severe implications for the victim (e.g., wolfsnakes (*Lycodon spp.*) and kraits (*Bungarus spp.*)) [5, 9, 38]. Existing snake identification techniques (e.g. immunoassays, PCR-based tests) require laboratory capacity and lengthy and costly procedures and are thus not yet adapted to low-resource snakebite endemic countries [10]. In certain countries, healthcare providers rely on local, national and even international herpetologists for identifying snakes brought to the health facility or photographed by the victims or relatives, as recommended by WHO (e.g., [14]). Although some specific platforms provide snake identification support (e.g., Sri Lanka (www.snakesidentification.org) [42], Thailand (www.thailandsnakes.com/thailand-snake-id/) [24]), an informal and unsecured process has emerged among healthcare providers, who share snake photos with other medics or herpetologists for identification via email, WhatsApp or even Facebook. Herpetologists usually provide their expertise on a volunteer basis and their availability and reactivity may be limited. In this context, AI could deliver simple, continuous, instantaneous, and low-cost snake identification to complement herpetologists [37, 43] and support healthcare providers and snakebite victims, especially considering the rapidly growing internet and mobile technology penetration in many snakebite endemic countries [44] (e.g., India). Improving the capacity of communities at risk of snakebite, conservationists, nature lovers, among others groups, to identify and learn about snakes could both help prevent snakebite envenoming (e.g., by recognising venomous snakes) and protect snakes (e.g., fewer snakes will be killed out of fear or ignorance).

Other professional groups, such as snake rescuers, veterinarian practitioners confronted to snakebite in companion and livestock animals [45], and wildlife trade inspectors (e.g., 6.3 million snakes were traded globally between 1975 and 2018 [27]), or biology students, snake enthusiasts and travellers, could also benefit from this AI.

Our study has several limitations. First, the snake photo dataset used to train the model does not yet cover the whole species diversity of some highly biodiverse regions: the mean percentage of species covered per country in Asia is 57%, 44% for Oceania, 41% for Latin America/Caribbean, and 34% for Africa. We have shown that online communities of snake enthusiasts and herpetologists can contribute large volumes of geo-tagged snake observations [19] and that aggregating photos among online data sources brought our team to the leading edge of global snake image datasets, with almost 80% of the world's species represented by at least one image [20]. Yet, only 20% of the world's species met our 10-image threshold for inclusion in the training dataset, highlighting the extremely long-tailed distribution of photos per species. Although online data sets continue to grow, connecting with difficult-to-access online (e.g. sub-communities of Facebook, Field Herp Forum) and offline (e.g. private WhatsApp snake identification chat groups, private or natural history museum image collections) communities will allow us to progress in snake identification. The species under-represented in our snake photo dataset are shown in [S1 Fig](#) and [S3 Table](#) and the most-wanted species

globally in [20]. To help improve the model accuracy and coverage, we encourage professional herpetologists and snake enthusiasts with photos of these species or missing species to submit them to The Reptile Database [4] or the citizen science biodiversity platforms iNaturalist and HerpMapper. Second, the snake photos used to evaluate the model performance (iNaturalist photos) may be "easy" for the algorithm to correctly identify and not reflect the kinds of photos that are usually taken in the context of a snakebite event (e.g., the biting snake has been killed and its head smashed, the photo is blurred, only part of the snake's body is visible, etc.). Further research is needed to test this AI snake identification system in the field. Third, we compared a limited number of lookalike species that we selected based on our literature review [9] and discussions with clinicians (GA). Further studies need to include and more systematically compare lookalike snake species occurring at the global scale.

Conclusion

We have built and openly shared an international AI model for the automatic identification of snakes, setting the basis for further machine learning research and AI applications to tackling snakebite and other neglected tropical diseases in low-resource settings. The power of AI must be embraced and used safely and equitably to improve health and wellbeing of the poorest communities of the world.

Supporting information

S1 Table. SnakeCLEF 2021 training set. Data sources and their taxonomic and geographic coverage.
(DOCX)

S2 Table. Details of the datasets used for training and testing of the model.
(DOCX)

S3 Table. MIVS species included in the training and test datasets.
(XLSX)

S1 Fig. Number of training photos per MIVS species.
(PDF)

S2 Fig. Number of test photos per MIVS species.
(PDF)

S3 Fig. Scatterplot of countries by coverage and accuracy. See [S9 Data](#) for underlying values. See interactive online version at <https://chart-studio.plotly.com/~amdurso/6/#/>.
(TIFF)

S1 Data. CSV file containing the underlying numerical data for Fig 1.
(CSV)

S2 Data. CSV file containing the underlying numerical data for Fig 3.
(CSV)

S3 Data. CSV file containing the underlying numerical data for Fig 4A.
(CSV)

S4 Data. CSV file containing the underlying numerical data for Fig 4B.
(CSV)

S5 Data. CSV file containing the underlying numerical data for Fig 5.
(CSV)

S6 Data. CSV file containing the underlying numerical data for Fig 6.
(CSV)

S7 Data. CSV file containing the underlying numerical data for Fig 7.
(CSV)

S8 Data. CSV file containing the underlying numerical data for Fig 8.
(CSV)

S9 Data. CSV file containing the underlying numerical data for S3 Fig.
(CSV)

Acknowledgments

We thank M. Salathé (Digital Epidemiology Lab, EPFL) and S.P. Mohanty (AICrowd platform) and A. Joly (LifeCLEF platform) for hosting and supporting several snake identification challenges. We thank all participants in our AICrowd and SnakeCLEF snake identification challenges. We are grateful to the admins, curators, and community members of online biodiversity platforms (iNaturalist, HerpMapper) for their efforts in building, curating, and sharing their image libraries. We thank C. Montalcini for supporting snake photo collection and organization. We thank A. Flahault for supporting this work at the Institute of Global Health.

Author Contributions

Conceptualization: Isabelle Bolon, Lukáš Pícek, Andrew M. Durso, Gabriel Alcoba, François Chappuis, Rafael Ruiz de Castañeda.

Data curation: Lukáš Pícek, Andrew M. Durso.

Formal analysis: Isabelle Bolon, Lukáš Pícek, Andrew M. Durso, Rafael Ruiz de Castañeda.

Funding acquisition: Isabelle Bolon, Rafael Ruiz de Castañeda.

Investigation: Isabelle Bolon, Lukáš Pícek, Andrew M. Durso, Gabriel Alcoba, François Chappuis, Rafael Ruiz de Castañeda.

Methodology: Isabelle Bolon, Lukáš Pícek, Andrew M. Durso, Gabriel Alcoba, François Chappuis, Rafael Ruiz de Castañeda.

Project administration: Isabelle Bolon, Rafael Ruiz de Castañeda.

Resources: Lukáš Pícek, Andrew M. Durso.

Software: Lukáš Pícek, Andrew M. Durso.

Supervision: Isabelle Bolon, Rafael Ruiz de Castañeda.

Validation: Isabelle Bolon, Lukáš Pícek, Andrew M. Durso, Gabriel Alcoba, François Chappuis, Rafael Ruiz de Castañeda.

Visualization: Lukáš Pícek, Andrew M. Durso.

Writing – original draft: Isabelle Bolon, Lukáš Pícek, Andrew M. Durso, Rafael Ruiz de Castañeda.

Writing – review & editing: Isabelle Bolon, Lukáš Pícek, Andrew M. Durso, Gabriel Alcoba, François Chappuis, Rafael Ruiz de Castañeda.

References

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, Hošek 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. *Toxicol*. 2018; 150:1–10. <https://doi.org/10.1016/j.toxicol.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, Jürgensen 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, Goëau 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/snakebite-information-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:201011929*. 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, Igloukov VI, Khvedchenya E, Parinov A, Druzhinin M, Kalinin AA. Albuumentations: fast and flexible image augmentations. *Information*. 2020; 11(2):125.
34. Durso AM, Moorthy GK, Mohanty SP, Bolon I, Salathé 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-Marques 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.

41. Picek L, Šulc M, Matas J, Heilmann-Clausen J, Jeppesen TS, Lind E. Automatic Fungi Recognition: Deep Learning Meets Mycology. *Sensors*. 2022; 22(2):633. <https://doi.org/10.3390/s22020633> PMID: [35062595](https://pubmed.ncbi.nlm.nih.gov/35062595/)
42. Maduwage K, Karunathilake P, Gutiérrez JM. Web-based snake identification service: A successful model of snake identification in Sri Lanka. *Toxicon*. 2022; 205:24–30. <https://doi.org/10.1016/j.toxicon.2021.11.007> PMID: [34774917](https://pubmed.ncbi.nlm.nih.gov/34774917/)
43. Ruiz de Castañeda R, Grey F, Williams DJ. Citizen surveys could map snakebite risks. *Nature*. 2019; 571(7766):478–. <https://doi.org/10.1038/d41586-019-02247-7>
44. We are social 2021. [cited 2022 February 02]. Available from: <https://wearesocial.com/uk/blog/2021/01/digital-2021-uk/>.
45. Bolon I, Finat M, Herrera M, Nickerson A, Grace D, Schütte S, et al. Snakebite in domestic animals: First global scoping review. *Prev Vet Med*. 2019; 170: 104729. <https://doi.org/10.1016/j.prevetmed.2019.104729> PMID: [31421490](https://pubmed.ncbi.nlm.nih.gov/31421490/)