Overview of SnakeCLEF 2021: Automatic Snake Species Identification with Country-Level Focus

Lukáš Picek[1], Andrew M. Durso[2], Isabelle Bolon[3] and Rafael Ruiz de Castañeda[3]

[1]*Department of Cybernetics, Faculty of Applied Sciences, University of West Bohemia, Czechia*
[2]*Department of Biological Sciences, Florida Gulf Coast University, Florida, USA*
[3]*Institute of Global Health, Department of Community Health and Medicine, University of Geneva, Switzerland*

Abstract
A robust and accurate AI-driven system as an assistance tool for snake species identification has vast potential to help lower deaths and disabilities caused by snakebites. With that in mind, we prepared the SnakeCLEF 2021: Automatic Snake Species Identification Challenge with Country-Level Focus, designed to provide an evaluation platform that can help track the performance of end-to-end AI-driven snake species recognition systems with a focus on overall country-wise performance. We have provided 386,006 photographs of 772 snake species collected in 188 countries and country-species presence mapping for the challenge. In this paper, we report 1) a description of the provided data, 2) evaluation methodology and principles, 3) an overview of the systems submitted by the participating teams, and 4) a discussion of the obtained results.

Keywords
LifeCLEF, SnakeCLEF, fine grained visual categorization, global health, epidemiology, snake bite, snake, reptile, benchmark, biodiversity, species identification, machine learning, computer vision, classification

1. Introduction

Building an automatic and robust image-based system for snake species identification is an important goal for biodiversity, conservation, and global health. With recent estimates of 81,410 - 137,880 deaths and up to three times as many victims of amputations, permanent disability and disfigurement (globally each year) caused by venomous snakebite [1], such a system has the potential to improve eco-epidemiological data and treatment outcomes (e.g. based on the specific use of antivenoms) [2, 3]. This applies especially in remote geographic areas and developing countries, where automatic snake species identification has the greatest potential to save lives.

The difficulty of snake species identification – from both a human and a machine perspective [4] – lies in the high intra-class and low inter-class variance in appearance, which may depend on geographic location, color morph, sex, or age (Figure 1 and Figure 2). At the same time, many species are visually similar to other species (e.g. mimicry [5]). Our knowledge of which snake species occur in which countries is incomplete, and it is common that most or all images of a given snake species might originate from a small handful of countries or even a single

CLEF 2021 – Conference and Labs of the Evaluation Forum, September 21–24, 2021, Bucharest, Romania
picekl@kky.zcu.cz (L. Picek)
0000-0002-6041-9722 (L. Picek); 0000-0002-3008-7763 (A. M. Durso); 0000-0001-5940-2731 (I. Bolon); 0000-0002-2287-0985 (R. R. d. Castañeda)

country [6]. Furthermore, many snake species resemble species found on other continents, with which they are entirely allopatric [7]. Knowing the geographic origin of an unidentified snake can narrow down the possible correct identifications considerably. In no location on Earth do more than 126 of the approximately 3,900 snake species co-occur [8]. Thus, regularization to all countries is a critical component of any snake identification method. In previous LifeCLEF Snake Species Identification challenges [9, 10] we measured relatively poor performance – 0.625 Macro F1 score – showing that snake identification is a task with a lot of space for improvement.

Figure 1: Variation in *Vipera berus* (European Adder) color and pattern within central Europe. Examples from Czechia, Germany, Switzerland and Poland, demonstrating different color morphs within a species. Taken from iNaturalist: *©Thorsten Stegmann, ©jandetka, ©jandetka, and ©jandetka*

Figure 2: *Naja nigricincta* from northern Namibia (left) and South Africa (right), demonstrating geographical variation within a species. Taken from iNaturalist: *©Di Franklin, and ©bryanmaritz*

2. Task description

The main goal of this challenge was to build a system that is capable of recognizing 772 snake species based on the given unseen image and relevant geographical location, with a focus on worldwide performance. Unlike the previous SnakeCLEF edition – where we used the disclosed dataset – we did not ask the participants to submit their solutions through Docker environment. Just a simple CSV file with Top1 species prediction for each image was expected.

2.1. Dataset

For this year's challenge, we have prepared a new dataset with 409,679 images belonging to 772 snake species from 188 countries and all continents (386,006 images with labels targeted for development and 23,673 images without labels for testing). In addition, we provide a simple train/val (90% / 10%) split to validate preliminary results while ensuring the same species distributions. Furthermore, we prepared a compact subset (70,208 images) for fast prototyping. The test set data consists of 23,673 images submitted to the iNaturalist platform within the first four months of 2021. Unlike in previous years, where the final testing set remained undisclosed, we provided the test data to the participants.

All data were gathered from online biodiversity platforms (i.e., iNaturalist, HerpMapper) and further extended by data scraped from Flickr. In contrast to the previous SnakeCLEF edition [10], we increased the number of images and covered countries, and filtered noisy labels and duplicated images. In addition, we defined clean (iNaturalist / HerpMapper) and noisy (Flickr) subsets within the development data. The provided dataset has a heavy long-tailed class distribution, where the most frequent species (*Thamnophis sirtalis*) is represented by 22,163 images and the least frequent by just 10 (*Achalinus formosanus*). For additional dataset parameters refer to Table 1 and Table 2.

Table 1
Details of the SnakeCLEF 2021 datasets and their comparison with previous edition.

Dataset	Species	Images	# of Countries	min per species	max per species
SnakeCLEF 2020	783	259,214	145	19	14,433
SnakeCLEF 2021	772	386,006	188	10	22,163
SnakeCLEF 2021 Comp.	768	70,208	178	1	299

Table 2
SnakeCLEF 2021 data sources and their taxonomic and geographic coverage.

Data Source	# of Species	# of Genera	# of Families	# of Images	# of Countries
iNaturalist	762	265	17	277,025	181
HerpMapper	614	244	17	58,351	98
Flickr	733	260	18	50,630	125
Total	772	269	18	386,006	188

2.1.1. Geographical Information

Considering that all snake species have distinct, largely stable geographic ranges, with a maximum of 126 species of snakes occurring within the same 50×50 km^2 area [8], geographical information plays a crucial role in correct snake species identification [11]. To evaluate this, we have gathered two levels of geographical label (i.e., country and continent) for approximately 87% of the data. We have collected observations across 188 countries and all continents. A small proportion of images (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 [12, 13]).
2. Bites from captive snakes may occur (although the identity of the snake would normally be clear in this case; e.g. [14, 15]).
3. Captive snakes sometimes escape and can found introduced populations outside their native range (e.g. [16, 17]).

Additionally, we provide a mapping matrix (MM) describing species-country presence to allow better worldwide regularization, based on the August 2020 release of The Reptile Database [18].

$$\mathrm{MM}_{cs} = \begin{cases} 1, & if \text{ species } S \in \text{ country } C \\ 0, & else \end{cases} \tag{1}$$

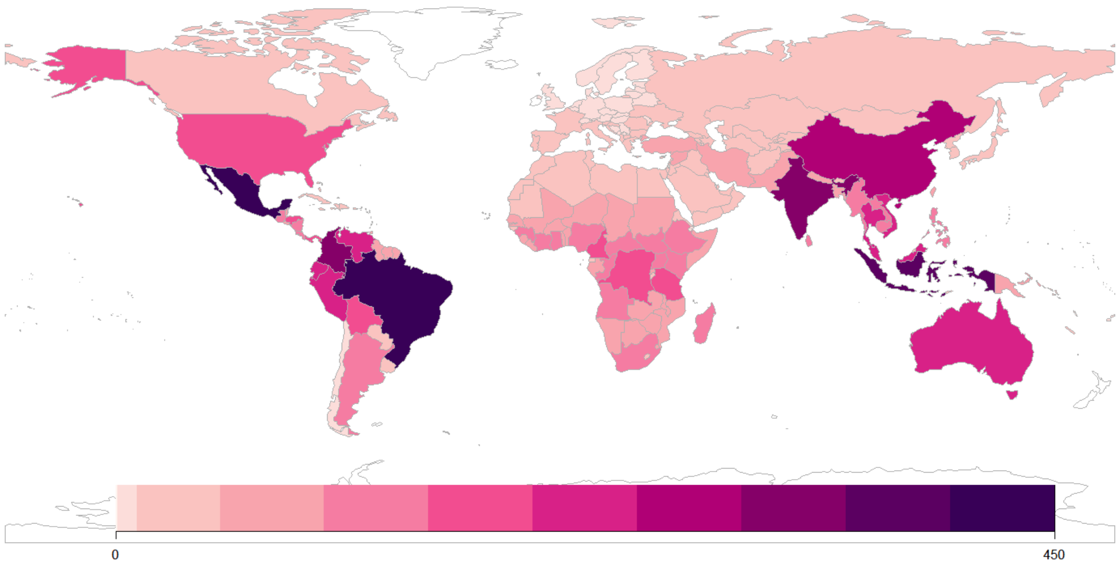

Figure 3: Worldwide snake species distribution, showing the number of species found in each country. Large countries in the tropics (Brazil, Mexico, Colombia, India, Indonesia) have more than 300 species.

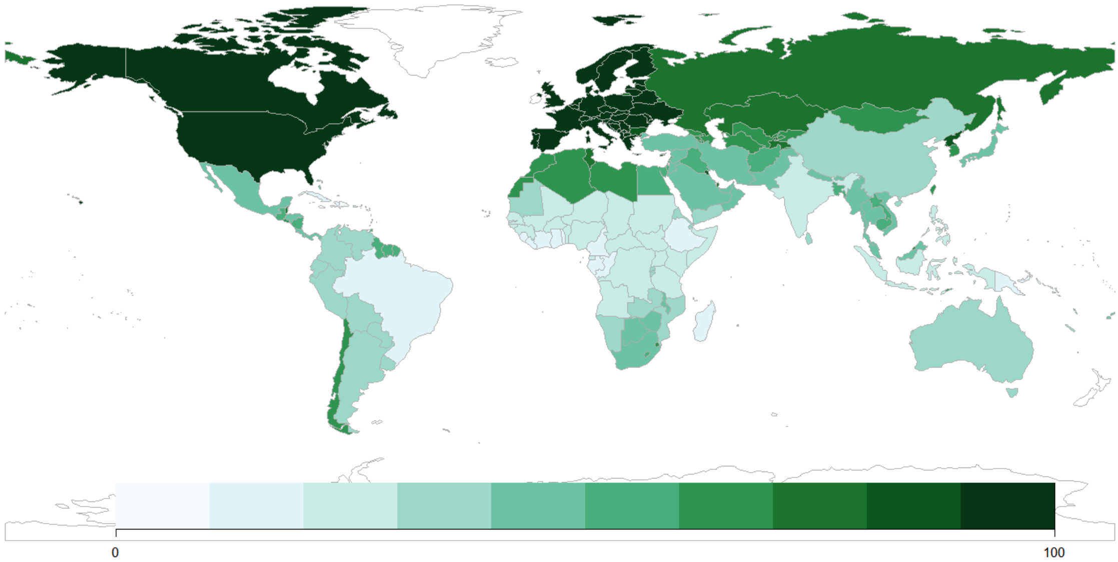

Figure 4: Percentage of snake species per country included in SnakeCLEF2021. The countries with the best coverage are in Europe, Oceania, and North America.

The vast majority (77%) of all images came from the United States and Canada, with 9% from Latin American 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 [6, 19], 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 (Figure 4). Priority areas for improvement of the training dataset in future rounds are countries with high snake species diversity and low citizen science participation, especially Indonesia, Papua New Guinea, Madagascar, and several central African and Caribbean countries (Figure 3).

2.2. Timeline

The training data were made public in February 2021 through the AICrowd challenge page, and anyone with research ambitions was able to register and participate in the competition. Releasing the test data in mid-May, we provided up to 100 days to participants to work on their submissions. The test data were released three days before the competition deadline, minimizing the possibility of manual labelling and other exploits. Each team had an opportunity to submit up to 10 submissions corresponding to different approaches or different settings of the same method. The final evaluation was done via a CSV file containing Top1 prediction for each given test image. Once the submission phase was closed (mid-June), the participants we allowed to submit so-called post-competition submissions to evaluate any interesting findings.

2.3. 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:

$$\text{Macro } F1_c = \frac{1}{N}\sum_{c=0}^{N} F1_c\,, \quad F1_c = \frac{1}{\sum_{s=1}^{k} MM_{cs}} \times \sum_{s=0}^{N} F1_s MM_{cs} \tag{2}$$

where *c* is country index, *s* is species index, $(F1_c)$ is the country performance, and MM_{cs} is the mapping matrix described in Subsection 2.1.1. To get the $F1_s$ we use following formula for each species:

$$F1_s = 2 \times \frac{P_s \times R_s}{P_s + R_s} \tag{3}$$

$$P_s = \frac{tp_s}{tp_s + fp_s}\,, \qquad R_s = \frac{tp_s}{tp_s + fn_s} \tag{4}$$

To allow deeper comparison on different levels, we also measure the Top1 Accuracy and the Macro F1 score. The Macro F1 score is calculated as the mean of all $F1_s$ scores:

$$\text{Macro F1} = \frac{1}{N}\sum_{s=0}^{N} F1_s \tag{5}$$

where *s* is the species index and *N* the number of species. Final Macro F1 is calculated by computing the F1 score for each species as the harmonic mean of the species Precision (P_s) and the Recall (R_s).

2.4. Working Notes

All participants were asked to provide a *Working Note* paper – a technical report with information needed to reproduce the results of all submissions. All submitted *Working Notes* were reviewed by 2-3 reviewers with a decent publication history and PhD in Computer Vision and Machine Learning, ensuring a sufficient level of reproducibility and quality. The review process was single-blind and offered up to two rebuttals.

3. Participants and Methods

Seven teams participated in the SnakeCLEF 2021 challenge and submitted a total of 46 runs. We have seen a vast increase in interest related to automatic snake recognition from the last year [20]. Interestingly, three participating teams are originated from India – the country with the most snakebites worldwide [21]. Most of the participants (6 out of 7) provided a technical report with a description for each run, evaluated experiments and used methods, techniques and experiments [22, 23, 24, 25, 26, 27]. Such a report had to pass a single-blind review, ensuring a sufficient level of reproducibility and quality. For all the teams, we synthesized a short description.

BME-TMIT [22]: The BME-TMIT was the only team that used a two-stage approach with detection and classification neural networks. EfficientDet [28] and EfficientNet [29] were utilized for object detection and classification, respectively. Additionally, the location metadata integration increased the F1 country by 0.089 on the test data. Based on evaluated experiments, we can conclude that object detection and the inclusion of geographical data showed significant improvement in all measured performance metrics. Utilizing that, they achieved the highest scores in all measured metrics – Macro F_{1_c} of 0.903, F_{1_c} of 0.864, and 94.94% Top1 Accuracy.)

CMP [23]: The CMP team experimented with different deep residual convolutional neural networks (i.e., ResNet [30], ResNeXt [31], and ResNeSt [32]) and different loss functions, including standard cross-entropy, weighted cross-entropy and soft F1 loss. The performed experiment showed that the standard cross-entropy loss achieved superior performance in all measured metrics on the validation set. Thus, their best method is an ensemble of two ResNeSt-200, ResNet-101, and ResNeXt-101, combining the top one predictions by majority voting strategy. Additionally, they increased the performance with mixed-precision training and by dropping the predictions of the species not occurring in the country of the given image. Interestingly, their best single model in the case of Macro F_{1_c} was fine-tuned just on the compact subset with the almost flat distribution.

FHDO-BCSG [24]: The FHDO-BCSG team utilized the EfficientNets [29] and the Vision Transformers (ViT) [33] in their experiments. In a subsequent step, they multiplied the prior probabilities of the location context with the model predictions. Without surprise, the combination of both modes achieved the best performance, more precisely a Macro F_{1_c} score of 0.829.

SSN [25]: SSN team used a classical approach with just a single ResNeXt-50-V2 optimized with Adam and plenty of image augmentations, i.e., random crop, transposition, horizontal/vertical flip, shift, scale and rotation. With such an approach, they achieved a relatively small error rate in terms of Top1 Accuracy (14.23%) but reached just the 0.724 in case of Macro F_{1_c}.

UAIC AI [26]: This team used relatively old CNN architectures GoogLeNet [34], VGG16 [35] and ResNet-18 [30]. Even though they did not achieve high scores, they helped us to understand the magnitude of the difference in performance between "pioneer" and the current state-of-the-art architectures on a long-tailed fine-grained dataset. Their best score – 0.785 Macro F_{1_c} – was achieved by the ResNet-18 architecture.

SSN-MLRG [27]: The SSN-MLRG team used the Inception-ResNet-v2 [36] as a feature extractor and concatenated extracted image features with geographic information. Such a feature vector is later forwarded into trained gradient boosting classifier. This approach achieved the worst performance in the competition (0.269 Macro F_{1_c}) and revealed the superiority of the neural network based classifiers.

Gokul: This work primarily builds on their solution around ViT (ViT-Base-16) and the CNN based ResNet101-v2 architectures [20]. An ensemble of both, with a few bells and whistles, improved the Country Based F1 score up to 0.877 (2nd place).

4. Results and Discussion

We report the achieved performance by all the collected runs in Figure 5, Figure 6, and Figure 7. The best performing model achieved an impressive Macro F_{1_c} of 0.903 while having 94.82% Top1 Accuracy and Macro F_1 of 0.855. Interestingly, the model with the highest Macro F_{1_c} was not the best in terms of Top1 Accuracy and Macro F_1. The main outcomes we can derive from the results are the following:

Object detection improves classification: Utilization of the detection network for a better region of interest selection showed a significant performance gain in the case of the winning team. However, such an approach requires additional labelling procedures and the construction of two neural network models. Furthermore, a two-stage solution might be too heavy for deployment on edge devices; thus, its usage is probably impossible.

CNN outperforms ViT in snake recognition: Similar to last year's challenge [10], all participants featured deep convolutional neural networks. Besides CNNs, Vision Transformers (ViT) [33] were utilized by two teams. Interestingly, the performance of the ViT was slightly worse, which is contradictory to their performance in fungi recognition [37], thus showing that ViT might not be the best option for all fine-grained tasks.

Geography improves classification: Same as last year, usage of geographical information improved the recognition capability. No matter which technique was used, every team that incorporated the location metadata information increased the system's performance by a significant margin, e.g., +0.089 and +0.103 Macro F_{1_c}, in the case of BME-TMIT and FHDO-BCSG respectively.

Vast increase in performance: This year we experienced a significant performance increase in all measured metrics. Comparing the top Macro F1 score achieved in 2020 (0.625) and 2021 (0.864), we can see a 2.75 times smaller error rate. This is mainly due to increasing research efforts in automatic snake species identification. With a Top1 Accuracy close to 95%, the 2021 SnakeCLEF challenge helped to build a system that has similar performance to other approaches for natural species recognition [38, 39, 40, 41].

Increased interest in automatic snake species recognition: This year the SnakeCLEF 2021 challenge attracted seven research teams from India, Czechia, Germany, Romania, and Hungary. This is so far the biggest participation in our Snake Identification challenges and even exceeds participation in other well-established LifeCLEF challenges. In 2022 we hope that interest will continue to increase.

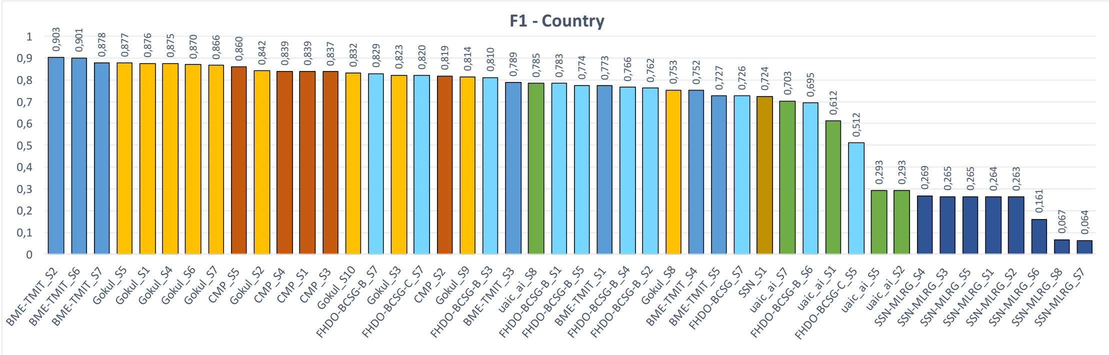

Figure 5: Official Macro $F1_c$ scores achieved by all runs to the SnakeCLEF 2021 competition.

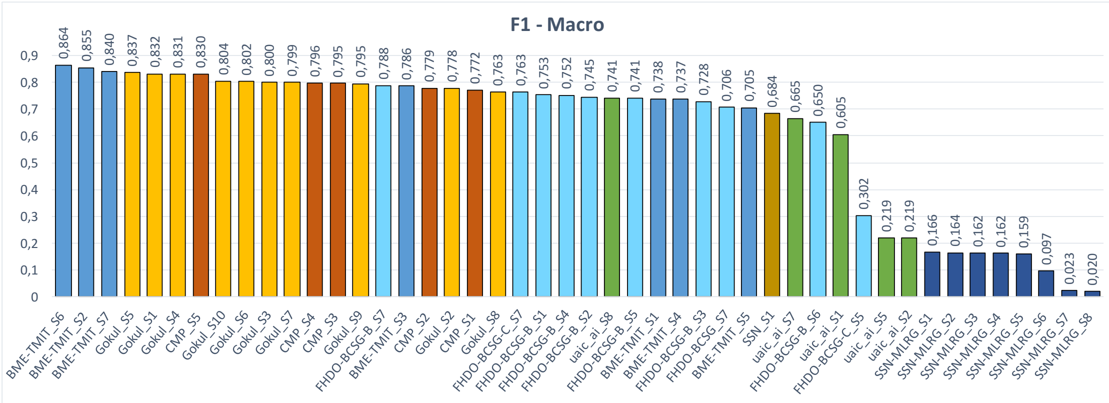

Figure 6: Official Macro F_1 scores achieved by all runs to the SnakeCLEF 2021 competition.

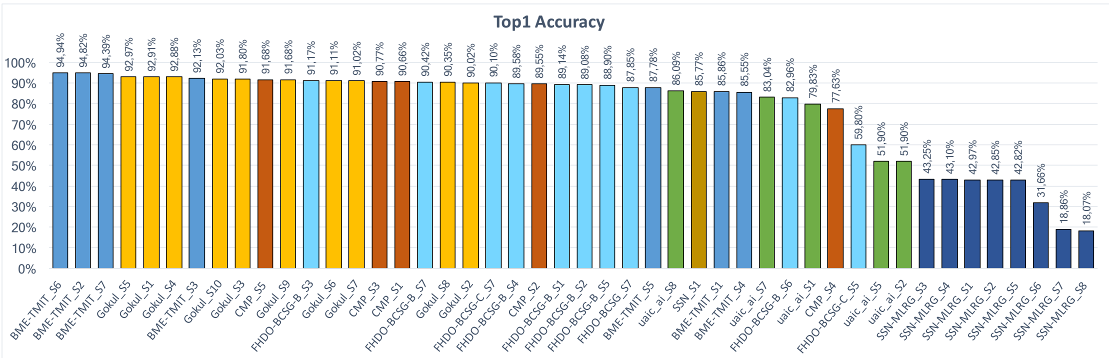

Figure 7: Official Top1 Accuracy scores achieved by all runs to the SnakeCLEF 2021 competition.

5. Conclusions and Perspectives

This paper presents an overview and results of the second edition of the SnakeCLEF challenge organized in conjunction with the Conference and Labs of the Evaluation Forum (CLEF[1]) and LifeCLEF[2] research platform [42]. This year, we based the evaluation on the worldwide species distribution. We have prepared the largest and most diverse snake image dataset to date, covering 772 snake species with 409,679 images observed across 188 countries. This dataset represents the most challenging dataset for automated snake species recognition in existence to date. For future editions, we plan to focus upon the following:

1. Extend the dataset, with new and rare species as well as reduce the bias towards North America.
2. Integrate the snake species toxicity level into the dataset and lower the possibility of medically-critical mis-prediction, i.e., confusion of venomous species with non-venomous.
3. Compare machine-learning based algorithms with human experts to better evaluate how far automated systems are from human expertise [4].

Acknowledgments

LP was supported by the UWB grant, project No. SGS-2019-027. A. M. Durso was supported by the Fondation privée des Hôpitaux Universitaires de Genève (award QS04-20). We thank the users and admins of open citizen science initiatives (iNaturalist, HerpMapper), and Flickr) for their efforts building these global datasets. We thank A. Flahault and the Fondation Louis-Jeantet, and F. Chappuis for supporting R. Ruiz de Castañeda and this research at the Institute of Global Health and at the Department of Community Health and Medicine of the University of Geneva.

References

[1] W. H. O. (WHO), Snakebite envenoming: Global situation., 2021 (accessed June 28, 2021). URL: https://www.who.int/news-room/fact-sheets/detail/snakebite-envenoming.

[2] R. Ruiz de Castañeda, A. M. Durso, N. Ray, J. L. Fernández, D. J. Williams, G. Alcoba, F. Chappuis, M. Salathé, I. Bolon, Snakebite and snake identification: empowering neglected communities and health-care providers with ai, The Lancet Digital Health 1 (2019) e202–e203.

[3] I. Bolon, A. M. Durso, S. Botero Mesa, N. Ray, G. Alcoba, F. Chappuis, R. Ruiz de Castañeda, Identifying the snake: First scoping review on practices of communities and healthcare providers confronted with snakebite across the world, PLoS one 15 (2020) e0229989.

[4] A. M. Durso, I. Bolon, A. Kleinhesselink, M. Mondardini, J. Fernandez-Marquez, F. Gutsche-Jones, C. Gwilliams, M. Tanner, C. E. Smith, W. Wüster, et al., Crowdsourcing snake identification with online communities of professional herpetologists and avocational snake enthusiasts, Royal Society open science 8 (2021) 201273.

[1] http://www.clef-initiative.eu/
[2] http://www.lifeclef.org/

[5] A. R. Davis Rabosky, C. L. Cox, D. L. Rabosky, P. O. Title, I. A. Holmes, A. Feldman, J. A. McGuire, Coral snakes predict the evolution of mimicry across new world snakes, Nature communications 7 (2016) 1–9.
[6] A. M. Durso, R. Ruiz de Castañeda, C. Montalcini, M. R. Mondardini, J. L. Fernandez-Marques, F. Grey, M. M. Müller, P. Uetz, B. M. Marshall, R. J. Gray, et al., Citizen science and online data: Opportunities and challenges for snake ecology and action against snakebite, Toxicon: X (2021) 100071.
[7] D. W. Pfennig, S. P. Mullen, Mimics without models: causes and consequences of allopatry in batesian mimicry complexes, Proceedings of the Royal Society B: Biological Sciences 277 (2010) 2577–2585.
[8] U. Roll, A. Feldman, M. Novosolov, A. Allison, A. M. Bauer, R. Bernard, M. Böhm, F. Castro-Herrera, L. Chirio, B. Collen, et al., The global distribution of tetrapods reveals a need for targeted reptile conservation, Nature Ecology & Evolution 1 (2017) 1677–1682.
[9] A. Joly, H. Goëau, S. Kahl, B. Deneu, M. Servajean, E. Cole, L. Picek, R. R. De Castaneda, I. Bolon, A. Durso, et al., Overview of lifeclef 2020: a system-oriented evaluation of automated species identification and species distribution prediction, in: International Conference of the Cross-Language Evaluation Forum for European Languages, Springer, 2020, pp. 342–363.
[10] L. Picek, R. Ruiz de Castaañeda, A. M. Durso, S. P. Mohanty, Overview of the snakeclef 2020: Automatic snake species identification challenge, in: CLEF task overview 2020, CLEF: Conference and Labs of the Evaluation Forum, Sep. 2020, Thessaloniki, Greece., 2020.
[11] H. C. Wittich, M. Seeland, J. Wäldchen, M. Rzanny, P. Mäder, Recommending plant taxa for supporting on-site species identification, BMC bioinformatics 19 (2018) 190.
[12] F. Hierink, I. Bolon, A. M. Durso, R. Ruiz de Castañeda, C. Zambrana-Torrelio, E. A. Eskew, N. Ray, Forty-four years of global trade in cites-listed snakes: Trends and implications for conservation and public health, Biological Conservation 248 (2020) 108601.
[13] D. J. Natusch, J. F. Carter, P. W. Aust, N. Van Tri, U. Tinggi, A. Riyanto, J. A. Lyons, et al., Serpent's source: Determining the source and geographic origin of traded python skins using isotopic and elemental markers, Biological Conservation 209 (2017) 406–414.
[14] A. Schaper, H. Desel, M. Ebbecke, L. D. Haro, M. Deters, H. Hentschel, M. Hermanns-Clausen, C. Langer, Bites and stings by exotic pets in europe: An 11 year analysis of 404 cases from northeastern germany and southeastern france, Clinical Toxicology 47 (2009) 39–43.
[15] B. J. Warrick, L. V. Boyer, S. A. Seifert, Non-native (exotic) snake envenomations in the us, 2005–2011, Toxins 6 (2014) 2899–2911.
[16] M. Á. Cabrera-Pérez, R. Gallo-Barneto, I. Esteve, C. Patiño-Martínez, L. F. López-Jurado, et al., The management and control of the california kingsnake in gran canaria (canary islands): project life+ lampropeltis, Aliens: The Invasive Species Bulletin 32 (????) 20–28.
[17] F. Kraus, Alien reptiles and amphibians: a scientific compendium and analysis, volume 4, Springer Science & Business Media, 2008.
[18] P. Uetz, P. Freed, J. Hošek, et al., The reptile database, 2020. URL: https://reptile-database.reptarium.cz/advanced_search.
[19] E. E. Millar, E. C. Hazell, S. Melles, The 'cottage effect'in citizen science? spatial bias in

aquatic monitoring programs, International Journal of Geographical Information Science 33 (2019) 1612–1632.
[20] A. M. Durso, G. K. Moorthy, S. P. Mohanty, I. Bolon, M. Salathé, R. Ruiz de Castañeda, Supervised learning computer vision benchmark for snake species identification from photographs: Implications for herpetology and global health, Frontiers in Artificial Intelligence 4 (2021) 17.
[21] B. Mohapatra, D. A. Warrell, W. Suraweera, P. Bhatia, N. Dhingra, R. M. Jotkar, P. S. Rodriguez, K. Mishra, R. Whitaker, P. Jha, et al., Snakebite mortality in india: a nationally representative mortality survey, PLoS Negl Trop Dis 5 (2011) e1018.
[22] R. Borsodi, D. Papp, Incorporation of object detection models and location data into snake species classification, in: Working Notes of CLEF 2021 - Conference and Labs of the Evaluation Forum, 2021.
[23] R. Chamidullin, M. Šulc, J. Matas, L. Picek, A deep learning method for visual recognitionof snake species, in: Working Notes of CLEF 2021 - Conference and Labs of the Evaluation Forum, 2021.
[24] L. Bloch, C. M. Friedrich, Efficientnets and vision transformers for snake species identification using image and location information, in: Working Notes of CLEF 2021 - Conference and Labs of the Evaluation Forum, 2021.
[25] K. Lekshmi, P. Balasundaram, G. Pradeep, S. Sekhar B, R. Kumar M, Automatic snake classification using deep learning algorithm, in: Working Notes of CLEF 2021 - Conference and Labs of the Evaluation Forum, 2021.
[26] L.-G. Coca, A.-T. Popa, R.-C. Croitoru, I. Bejan Luciana-Paraschiva, Adrian, Uaic-ai at snakeclef 2021: Impact of convolutions in snake species recognition, in: Working Notes of CLEF 2021 - Conference and Labs of the Evaluation Forum, 2021.
[27] D. Karthik, P. Mirunalini, J. Kumar, Snake species classification using transfer learning technique, in: Working Notes of CLEF 2021 - Conference and Labs of the Evaluation Forum, 2021.
[28] M. Tan, R. Pang, Q. V. Le, Efficientdet: Scalable and efficient object detection, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2020, pp. 10781–10790.
[29] M. Tan, Q. Le, Efficientnet: Rethinking model scaling for convolutional neural networks, in: International Conference on Machine Learning, PMLR, 2019, pp. 6105–6114.
[30] K. He, X. Zhang, S. Ren, J. Sun, Deep Residual Learning for Image Recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
[31] S. Xie, R. Girshick, P. Dollar, Z. Tu, K. He, Aggregated Residual Transformations for Deep Neural Networks, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.
[32] H. Zhang, C. Wu, Z. Zhang, Y. Zhu, H. Lin, Z. Zhang, Y. Sun, T. He, J. Mueller, R. Manmatha, M. Li, A. Smola, ResNeSt: Split-Attention Networks, 2020. arXiv:2004.08955.
[33] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al., An image is worth 16x16 words: Transformers for image recognition at scale, arXiv preprint arXiv:2010.11929 (2020).
[34] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke,

A. Rabinovich, Going deeper with convolutions, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2015, pp. 1–9.
[35] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, arXiv preprint arXiv:1409.1556 (2014).
[36] C. Szegedy, S. Ioffe, V. Vanhoucke, A. A. Alemi, Inception-v4, inception-resnet and the impact of residual connections on learning, in: Thirty-first AAAI conference on artificial intelligence, 2017.
[37] L. Picek, M. Šulc, J. Matas, J. Heilmann-Clausen, T. S. Jeppesen, T. Læssøe, T. Frøslev, Danish fungi 2020 – not just another image recognition dataset, 2021. arXiv:2103.10107.
[38] M. S. Norouzzadeh, A. Nguyen, M. Kosmala, A. Swanson, M. S. Palmer, C. Packer, J. Clune, Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning, Proceedings of the National Academy of Sciences 115 (2018) E5716–E5725. doi:10.1073/pnas.1719367115.
[39] M. Willi, R. T. Pitman, A. W. Cardoso, C. Locke, A. Swanson, A. Boyer, M. Veldthuis, L. Fortson, Identifying animal species in camera trap images using deep learning and citizen science, Methods in Ecology and Evolution 10 (2019) 80–91. doi:10.1111/2041-210X.13099.
[40] M. Sulc, L. Picek, J. Matas, Plant recognition by inception networks with test-time class prior estimation, in: CLEF working notes 2018, CLEF: Conference and Labs of the Evaluation Forum, Sep. 2018, Avignon, France., 2018.
[41] L. Picek, M. Sulc, J. Matas, Recognition of the amazonian flora by inception networks with test-time class prior estimation, in: CLEF working notes 2019, CLEF: Conference and Labs of the Evaluation Forum, Sep. 2019, Lugano, Switzerland., 2019.
[42] A. Joly, H. Goëau, S. Kahl, L. Picek, T. Lorieul, E. Cole, B. Deneu, M. Servajean, R. Ruiz De Castañeda, I. Bolon, H. Glotin, R. Planqué, W.-P. Vellinga, A. Dorso, H. Klinck, T. Denton, I. Eggel, P. Bonnet, H. Müller, Overview of lifeclef 2021: a system-oriented evaluation of automated species identification and species distribution prediction, in: Proceedings of the Twelfth International Conference of the CLEF Association (CLEF 2021), 2021.

A. Country Distribution