LifeCLEF 2022 Teaser: An Evaluation of Machine-Learning Based Species Identification and Species Distribution Prediction

Alexis Joly, Hervé Goëau, Stefan Kahl, Lukáš Picek, Titouan Lorieul, Elijah Cole, Benjamin Deneu, Maximilien Servajean, Andrew Durso, Isabelle Bolon, et al.

LifeCLEF 2022 Teaser: An Evaluation of Machine-Learning Based Species Identification and Species Distribution Prediction

Alexis Joly[1(✉)], Hervé Goëau[2], Stefan Kahl[6], Lukáš Picek[10],
Titouan Lorieul[1], Elijah Cole[9], Benjamin Deneu[1],
Maximilien Servajean[7], Andrew Durso[11], Isabelle Bolon[8],
Hervé Glotin[3], Robert Planqué[4], Willem-Pier Vellinga[4], Holger Klinck[6],
Tom Denton[12], Ivan Eggel[5], Pierre Bonnet[2], Henning Müller[5],
and Milan Šulc[13]

[1] Inria, LIRMM, Univ Montpellier, CNRS, Montpellier, France
alexis.joly@inria.fr
[2] CIRAD, UMR AMAP, Montpellier, Occitanie, France
[3] Univ. Toulon, Aix Marseille Univ., CNRS, LIS, DYNI Team,
Marseille, France
[4] Xeno-canto Foundation, Amsterdam, The Netherlands
[5] HES-SO, Sierre, Switzerland
[6] KLYCCB, Cornell Lab of Ornithology, Cornell University, Ithaca, USA
[7] LIRMM, AMIS, Univ Paul Valéry Montpellier, University Montpellier, CNRS,
Montpellier, France
[8] ISG, Department of Community Health and Medicine,
UNIGE, Geneva, Switzerland
[9] Department of Computing and Mathematical Sciences, Caltech, Pasadena, USA
[10] Department of Cybernetics, FAV, University of West Bohemia, Plzen, Czechia
[11] Department of Biological Sciences, Florida Gulf Coast University,
Fort Myers, USA
[12] Google LLC, San Francisco, USA
[13] Department of Cybernetics, FEE, CTU in Prague, Prague, Czech Republic

Abstract. Building accurate knowledge of the identity, the geographic distribution and the evolution of species is essential for the sustainable development of humanity, as well as for biodiversity conservation. However, the difficulty of identifying plants, animals and fungi is hindering the aggregation of new data and knowledge. Identifying and naming living organisms is almost impossible for the general public and is often difficult even for professionals and naturalists. Bridging this gap is a key step towards enabling effective biodiversity monitoring systems. The LifeCLEF campaign, presented in this paper, has been promoting and evaluating advances in this domain since 2011. The 2022 edition proposes five data-oriented challenges related to the identification and

prediction of biodiversity: (i) PlantCLEF: very large-scale plant identification, (ii) BirdCLEF: bird species recognition in audio soundscapes, (iii) GeoLifeCLEF: remote sensing based prediction of species, (iv) SnakeCLEF: Snake Species Identification in Medically Important scenarios, and (v) FungiCLEF: Fungi recognition from images and meta-data.

1 LifeCLEF Lab Overview

Accurately identifying organisms observed in the wild is an essential step in ecological studies. Unfortunately, observing and identifying living organisms requires high levels of expertise. For instance, vascular plants alone account for more than 300,000 different species and the distinctions between them can be quite subtle. The world-wide shortage shortage of trained taxonomists and curators capable of identifying organisms has come to be known as the *taxonomic impediment*. Since the Rio Conference of 1992, it has been recognized as one of the major obstacles to the global implementation of the Convention on Biological Diversity[1]. In 2004, Gaston and O'Neill [7] discussed the potential of automated approaches for species identification. They suggested that, if the scientific community were able to (i) produce large training datasets, (ii) precisely evaluate error rates, (iii) scale up automated approaches, and (iv) detect novel species, then it would be possible to develop a generic automated species identification system that would open up new vistas for research in biology and related fields.

Since the publication of [7], automated species identification has been studied in many contexts [6,9,13,23,25,30,31,36]. This area continues to expand rapidly, particularly due to advances in deep learning [5,8,24,26,32–35]. In order to measure progress in a sustainable and repeatable way, the LifeCLEF[2] research platform was created in 2014 as a continuation and extension of the plant identification task that had been run within the ImageCLEF lab[3] since 2011 [10–12]. Since 2014, LifeCLEF expanded the challenge by considering animals in addition to plants, and including audio and video content in addition to images [14–21]. About 100-500 research groups annually register to LifeCLEF in order to either download the data, register to the mailing list or benefit from the shared evaluation tools. The number of participants who finally crossed the finish line by submitting runs was respectively: 22 in 2014, 18 in 2015, 17 in 2016, 18 in 2017, 13 in 2018, 16 in 2019, 16 in 2020, 1,022 in 2021 (including the 1,004 participants of the BirdCLEF Kaggle challenge). The 2022 edition proposes five data-oriented challenges: three in the continuity of the 2021 edition (BirdCLEF, GeoLifeCLEF and SnakeCLEF), one new challenge related to fungi recognition with a focus on the combination of visual information with meta-data and on edible vs. poisonous species (FungiCLEF), and a considerable expansion of the PlantCLEF

---

[1] https://www.cbd.int/.
[2] http://www.lifeclef.org/.
[3] http://www.imageclef.org/.

challenge towards the identification of the world's flora (about 300K species). In the following sections, we describe for each of the five challenges the motivation, the used data collection and the evaluated task.

2 PlantCLEF Challenge: Identify the World's Flora

Motivation: It is estimated that there are more than 300,000 species of vascular plants in the world. Increasing our knowledge of these species is of paramount importance for the development of human civilization (agriculture, construction, pharmacopoeia, etc.), especially in the context of the biodiversity crisis [22]. However, the burden of systematic plant identification by human experts strongly penalizes the aggregation of new data and knowledge. Since then, automatic identification has made considerable progress in recent years as highlighted during all previous editions of PlantCLEF. Deep learning techniques now seem mature enough to address the ultimate but realistic problem of global identification of plant biodiversity in spite of many problems that the data may present (a huge number of classes, very strongly unbalanced classes, partially erroneous identifications, duplications, variable visual quality, diversity of visual contents such as photos or herbarium sheets, etc.).

Data Collection: the training dataset that will be used this year can be distinguished in 2 main categories: labeled and unlabeled (i.e. with or without species labels provided and checked by humans). The labeled training dataset will be based on a dataset of more than 5M images covering more than 290k plant species based on a web crawl with Google and Bing search engines and the Encyclopedia of Life webportal. All datasets provided in previous editions of PlantCLEF can also be used and the use of external data will be encouraged, notably via the gbif-dl[4] package which facilitates the download of media data from the world's largest biodiversity database (GBIF[5]) by wrapping its public API. The unlabeled training dataset will be based on more than 9 million pictures coming from the Pl@ntNet platform [4] (associated with a pseudo-label but without human verification). Finally, the test set will be a set of tens of thousands pictures verified by world class experts related to various regions of the world and taxonomic groups.

Task Description: the task will be evaluated as a plant species retrieval task based on multi-image plant observations from the test set. The goal will be to retrieve the correct plant species among the top results of a ranked list of species returned by the evaluated system. The participants will first have access to the training set and a few months later, they will be provided with the whole test set. Semi-supervised or unsupervised approaches will be strongly encouraged and a starter package with a pre-trained model based on this type of method exploiting the unlabeled training dataset will be provided.

---

[4] https://github.com/plantnet/gbif-dl.
[5] https://www.gbif.org/.

3 BirdCLEF Challenge: Bird Species Identification in Soundscape Recordings

Motivation: Recognizing bird sounds in complex soundscapes is an important sampling tool that often helps to reduce the limitations of point counts[6]. In the future, archives of recorded soundscapes will become increasingly valuable as the habitats in which they were recorded will be lost in the near future. It is imperative to develop new technologies that can cope with the increasing amount of audio data and that can help to accelerate the process of species diversity assessments. In the past few years, deep learning approaches have transformed the field of automated soundscape analysis. Yet, when training data is sparse, detection systems struggle with the recognition of rare species. The goal of this competition is to establish training and test datasets that can serve as real-world applicable evaluation scenarios for endangered habitats and help the scientific community to advance their conservation efforts through automated bird sound recognition.

Data Collection: We will build on the experience from previous editions and adjust the overall task to encourage participants to focus on few-shot learning and task-specific model designs. We will select training and test data to suit this demand. As for previous years, Xeno-canto will be the primary source for training data, expertly annotated soundscape recordings will be used for testing. We will focus on bird species for which there is limited training data, but we will also include common species so that participants can train good recognition systems. In search of suitable test data, we will consider different data sources with varying complexity (call density, chorus, signal-to-noise ratio, anthropophony), and quality (mono and stereo recordings). We also want to focus on very specific real-world use cases (e.g., conservation efforts in Hawaii) and frame the competition based on the demand of the particular use case. Additionally, we are considering including unlabeled data to encourage self-supervised learning regimes.

Task Description: The competition will be held on Kaggle and the evaluation mode will resemble the 2021 test mode (i.e., hidden test data, code competition). We will use established metrics like F1 score and LwLRAP which reflect use cases for which precision is key and also allow organizers to assess system performance independent of fine-tuned confidence thresholds. Participants will be asked to return a list of species for short audio segments extracted from labeled soundscape data. In the past, we used 5-second segments, and we will consider increasing the duration of these context windows to better reflect the overall ground truth label distribution. However, the overall structure of the task will remain unchanged, as it provides a well-established base that has resulted in significant participation in past editions (e.g., 1,004 participants and 9,307 submissions in 2021). Again, we will strive to keep the dataset size reasonably

---

[6] e.g. some species might be oversampled or undersampled.

small ($<$50 GB) and easy to process, and we will also provide introductory code repositories and write-ups to lower the entry level of the competition.

4 GeoLifeCLEF Challenge: Species Prediction Based on Occurrence Data, Environmental Data and Remote Sensing Data

Motivation: Automatically predicting the list of species that are the most likely to be observed at a given location is useful for many scenarios in bio-diversity conservation, ecotourism, land management, etc. First of all, it allows improve species identification tools by reducing the list of candidate species that are observable at a given location (be they automated, semi-automated or based on classical field guides or flora). More generally, it facilitates biodiversity inventories through the development of location-based recommendation services (typically on mobile phones), it favours the involvement of non-expert nature observers, as well as accelerate the annotation or validation of species observed by non-experts to produce high quality datasets. Last but not least, it might serve educational purposes thanks to biodiversity discovery applications providing functionalities such as contextualized educational pathways.

Data Collection: The GeoLifeCLEF dataset (already used in 2020 and 2021) contains about 2 million observations of around 30 thousand plant and animal species. Each observation is paired with very high-resolution covariates (aerial imagery, land cover, altitude) and environmental rasters (bioclimatic variables, soil type, etc.). The dataset took months to build in its raw format ($\sim$850 GB) and we reformatted it in a more convenient and memory efficient format ($\sim$100 GB). Indeed, it has not yet been used to its full potential due (i) to the computing power required to train models on it, and, (ii) to the complexity and the wide variety of challenges of the tackled task. In 2021, the challenge focused on measuring the efficiency of remote sensing imagery to predict the presence of species at a given location. In 2022, the objective is to make this competition more realistic by changing the evaluation protocol: the models will be evaluated on new presence/absence observation data. This means that the 2022 challenge will tackle two main issues in species presence prediction: (i) taking into consideration the sampling bias due to the presence-only nature of the training observation data, and, (ii) predicting relevant sets of species present at the given location.

Task Description: Given the test set of locations (i.e. geo-coordinates) and corresponding high-resolution and environmental covariates, the goal of the task will be to return for each location a ranked list of species sorted according to the likelihood that they might have been observed at that location. The metric used will be a multi-label metric such as mean average precision (mAP).

5 SnakeCLEF Challenge: Automated Snake Species Identification with Country-Level Focus

Motivation: Developing a robust system for identifying snake species from photographs is an important goal in biodiversity and global health. With over half a million of deaths and disability from venomous snakebite annually, understanding the global distribution of more than 3,900 snake species and differentiating them from images (particularly images of low quality) will significantly improve epidemiology data and treatment outcomes. From previous editions, we learned that "machines" are capable of accurate recognition (Macro F_{1_c} >90%, Accuracy ∼95%) even in the scenarios with long-tailed class distributions and ∼800 species. Thus, testing over real Medically Important Scenarios and integrating information on species toxicity is the next step to provide a more reliable "machine" prediction.

Data Collection: The dataset used in previous editions [27,29] will be extended with new and rare species as well as with images from countries with no or just a few samples, reducing the uneven species distributions across all the countries included in the data. For testing, we tailored two sets, one for a machine evaluation and the second for the HUMAN vs AI comparison. The SnakeCLEF 2022 dataset covers 1,000 snake species on more than 500,000 images and from approximately 200 countries – adding 224 new species. In addition, we include: (i) snake species toxicity level, allowing us to research methods for lowering the possibility of medically-critical mis-prediction, i.e., confusion of venomous species with non-venomous. (ii) country-species mapping file describing species-country presence based on the The Reptile Database and allowing better worldwide regularization.

Task Description: Given the set of images and corresponding geographic locality information, the goal of the task is to return for each image a ranked list of species sorted according to the likelihood that they are in the image and might have been observed at that location and minimising the venomous/non-venomous confusion.

6 FungiCLEF Challenge: Fungi Recognition from Images and Metadata

Motivation: Automatic recognition of fungi species assists mycologists, citizen scientists and nature enthusiasts in species identification in the wild. Its availability supports the collection of valuable biodiversity data. In practice, species identification typically does not depend solely on the visual observation of the specimen but also on other information available to the observer – such as habitat, substrate, location and time. Thanks to rich metadata, precise annotations, and baselines available to all competitors, the challenge provides a benchmark for image recognition with the use of additional information. Moreover, similarly to SnakeCLEF, the toxicity of a mushroom can be crucial for the decision

of a mushroom picker. The task will explore the decision process beyond the commonly assumed 0/1 cost function.

Data Collection: The challenge dataset is based on the DF20 dataset [28], contains 295,938 training images belonging to 1,604 species observed mostly in Denmark. All training samples passed an expert validation process, guaranteeing high quality labels. Rich observation metadata about habitat, substrate, time, location, EXIF etc. are provided. The challenge comes with two different test sets: (i) The first is unique in its annotation process, as all test images belong to physical samples sent for DNA sequencing. (ii) The second, with approximately 60k images, covers the whole year and includes observations collected across all substrate and habitat types.

Task Description: Given the set of images and corresponding metadata, the goal of the task is to return for each image a ranked list of species sorted according to the likelihood of the species appearing in the image. A baseline procedure to include meta-data in the decision problem, as well as pre-trained baseline image classifiers, will be provided as part of the task description to all participants.

7 Timeline and Registration Instructions

All information about the timeline and participation in the challenges is provided on the LifeCLEF 2022 web pages [3]. The challenges themselves are ran on the AIcrowd platform [1] and the Kaggle platform [2] for the registration, the submission of runs, the display of the leaderboard, etc.

8 Conclusions and Perspectives

To fully reach its objective, an evaluation campaign such as LifeCLEF requires a long-term research effort so as to (i) encourage non-incremental contributions, (ii) measure consistent performance gaps, (iii) progressively scale-up the problem and (iv) enable the emergence of a strong community. The 2022 edition of the lab supports this vision and also includes the following innovations:

- A new task on fungi recognition from images and metadata.
- A widening of the plant task at the scale of the world flora (100K-300K species).
- The inclusion of new data for the bird task with a focus on unsupervised training, stereo audio recordings and concrete conservation use cases.
- The inclusion of presence-absence test data for the GeoLifeCLEF challenge.
- The evaluation of decision problems for poisonous and venomous species identification represents a task beyond 0/1 cost function, not represented in computer vision benchmarks.

References

1. AICrowd. https://www.aicrowd.com/
2. Kaggle. https://www.kaggle.com/
3. LifeCLEF (2022). https://www.imageclef.org/LifeCLEF2022
4. Affouard, A., Goeau, H., Bonnet, P., Lombardo, J.C., Joly, A.: Pl@ntnet app in the era of deep learning. In: 5th International Conference on Learning Representations (ICLR 2017), 24–26 April 2017, Toulon, France (2017)
5. Goëau, H., et al.: Plant Identification: Experts vs. Machines in the Era of Deep Learning. In: Joly, A., Vrochidis, S., Karatzas, K., Karppinen, A., Bonnet, P. (eds.) Multimedia Tools and Applications for Environmental & Biodiversity Informatics. MSA, pp. 131–149. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-76445-0_8
6. Cai, J., Ee, D., Pham, B., Roe, P., Zhang, J.: Sensor network for the monitoring of ecosystem: Bird species recognition. In: Intelligent Sensors, Sensor Networks and Information, 2007. ISSNIP 2007. 3rd International Conference on (2007). https://doi.org/10.1109/ISSNIP.2007.4496859
7. Gaston, K.J., O'Neill, M.A.: Automated species identification: why not? Philos. Trans. Royal Soc. London B: Biol. Sci. **359**(1444), 655–667 (2004)
8. Ghazi, M.M., Yanikoglu, B., Aptoula, E.: Plant identification using deep neural networks via optimization of transfer learning parameters. Neurocomputing **235**, 228–235 (2017)
9. Glotin, H., Clark, C., LeCun, Y., Dugan, P., Halkias, X., Sueur, J.: Proceeding 1st workshop on Machine Learning for Bioacoustics - ICML4B. ICML, Atlanta USA (2013). http://sabiod.org/ICML4B2013_book.pdf
10. Goëau, H., et al.: The imageclef 2013 plant identification task. In: CLEF task Overview 2013, CLEF: Conference and Labs of the Evaluation Forum, Sep. 2013, Valencia, Spain. Valencia (2013)
11. Goëau, H., et al.: The imageclef 2011 plant images classification task. In: CLEF task Overview 2011, CLEF: Conference and Labs of the Evaluation Forum, Sep. 2011, Amsterdam, Netherlands. (2011)
12. Goëau, H., et al.: Imageclef 2012 plant images identification task. In: CLEF Task Overview 2012, CLEF: Conference and Labs of the Evaluation Forum, Sep. 2012, Rome, Italy. Rome (2012)
13. Joly, A., et al.: Interactive plant identification based on social image data. Ecol. Inf. **23**, 22–34 (2014)
14. Joly, A., et al.: Overview of LifeCLEF 2018: a large-scale evaluation of species identification and recommendation algorithms in the era of ai. In: Jones, G.J., et al. (eds.) CLEF: Cross-Language Evaluation Forum for European Languages. Experimental IR Meets Multilinguality, Multimodality, and Interaction, vol. LNCS. Springer, Avigon, France (Sep 2018)
15. Joly, A., et al.: Overview of LifeCLEF 2019: Identification of Amazonian Plants, South & North American Birds, and Niche Prediction. In: Crestani, F., et al. (eds.) CLEF 2019 - Conference and Labs of the Evaluation Forum. Experimental IR Meets Multilinguality, Multimodality, and Interaction, vol. LNCS, pp. 387–401. Lugano, Switzerland (Sep 2019). https://doi.org/10.1007/978-3-030-28577-7_29, https://hal.umontpellier.fr/hal-02281455

16. Joly, A., et al.: LifeCLEF 2016: Multimedia Life Species Identification Challenges. In: Fuhr, N., et al. (eds.) CLEF: Cross-Language Evaluation Forum. Experimental IR Meets Multilinguality, Multimodality, and Interaction, vol. LNCS, pp. 286–310. Springer, Évora, Portugal (Sep 2016). https://doi.org/10.1007/978-3-319-44564-9_26, https://hal.archives-ouvertes.fr/hal-01373781

17. Joly, A., et al.: LifeCLEF 2017 Lab Overview: Multimedia Species Identification Challenges. In: Jones, G.J., et al. (eds.) CLEF: Cross-Language Evaluation Forum. Experimental IR Meets Multilinguality, Multimodality, and Interaction, vol. LNCS, pp. 255–274. Springer, Dublin, Ireland (Sep 2017). https://doi.org/10.1007/978-3-319-65813-1_24, https://hal.archives-ouvertes.fr/hal-01629191

18. Joly, A., et al.: LifeCLEF 2014: Multimedia Life Species Identification Challenges. In: CLEF: Cross-Language Evaluation Forum. Information Access Evaluation. Multilinguality, Multimodality, and Interaction, vol. LNCS, pp. 229–249. Springer International Publishing, Sheffield, United Kingdom (Sep 2014). https://doi.org/10.1007/978-3-319-11382-1_20, https://hal.inria.fr/hal-01075770

19. Joly, A., et al.: Lifeclef 2015: multimedia life species identification challenges. In: Experimental IR Meets Multilinguality, Multimodality, and Interaction, pp. 462–483. Springe, Chem (2015)

20. Joly, A., 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, pp. 342–363. Springer, Chem (2020)

21. Joly, A., et al.: Overview of lifeclef 2021: an evaluation of machine-learning based species identification and species distribution prediction. In: International Conference of the Cross-Language Evaluation Forum for European Languages, pp. 371–393. Springer, Chem (2021)

22. Koh, L.P., Dunn, R.R., Sodhi, N.S., Colwell, R.K., Proctor, H.C., Smith, V.S.: Species coextinctions and the biodiversity crisis. Science **305**(5690), 1632–1634 (2004)

23. Lee, D.J., Schoenberger, R.B., Shiozawa, D., Xu, X., Zhan, P.: Contour matching for a fish recognition and migration-monitoring system. In: Optics East, pp. 37–48. International Society for Optics and Photonics (2004)

24. Lee, S.H., Chan, C.S., Remagnino, P.: Multi-organ plant classification based on convolutional and recurrent neural networks. IEEE Trans. Image Process. **27**(9), 4287–4301 (2018)

25. NIPS International Conference on Neural Information Processing Scaled for Bioacoustics, from Neurons to Big Data (2013). http://sabiod.org/nips4b

26. Norouzzadeh, M.S., Morris, D., Beery, S., Joshi, N., Jojic, N., Clune, J.: A deep active learning system for species identification and counting in camera trap images. Methods Ecol. Evol. **12**(1), 150–161 (2021)

27. Picek, L., Ruiz De Castañeda, R., Durso, A.M., Sharada, P.M.: 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)

28. Picek, L., Sulc, M., Matas, J., Heilmann-Clausen, J., Jeppesen, T., Læssøe, T., Frøslev, T.: Danish fungi 2020 - not just another image recognition dataset. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (2022)

29. Picek, L., Durso, A.M., Ruiz De Castañeda, R., Bolon, I.: Overview of SnakeCLEF 2021: Automatic snake species identification with country-level focus. In: Working Notes of CLEF 2021 - Conference and Labs of the Evaluation Forum (2021)

30. Towsey, M., Planitz, B., Nantes, A., Wimmer, J., Roe, P.: A toolbox for animal call recognition. Bioacoustics **21**(2), 107–125 (2012)
31. Trifa, V.M., Kirschel, A.N., Taylor, C.E., Vallejo, E.E.: Automated species recognition of antbirds in a Mexican rainforest using hidden Markov models. J. Acoust. Soc. Am. **123**, 2424 (2008)
32. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. CVPR (2018)
33. Villon, S., Mouillot, D., Chaumont, M., Subsol, G., Claverie, T., Villéger, S.: A new method to control error rates in automated species identification with deep learning algorithms. Sci. Reports **10**(1), 1–13 (2020)
34. Wäldchen, J., Mäder, P.: Machine learning for image based species identification. Methods Ecol. Evol. **9**(11), 2216–2225 (2018)
35. Wäldchen, J., Rzanny, M., Seeland, M., Mäder, P.: Automated plant species identification-trends and future directions. PLoS Comput. Biol. **14**(4), e1005993 (2018)
36. Yu, X., Wang, J., Kays, R., Jansen, P.A., Wang, T., Huang, T.: Automated identification of animal species in camera trap images. EURASIP J. Image Video Process. **2013**(1), 1–10 (2013). https://doi.org/10.1186/1687-5281-2013-52