

Overview of FungiCLEF 2023: Fungi Recognition Beyond 1/0 Cost

Lukáš Pícek^{1,*}, Milan Šulc², Rail Chamidullin¹ and Jiří Matas³

¹Department of Cybernetics, Faculty of Applied Sciences, University of West Bohemia, Czech Republic

²Second Foundation, Czech Republic

³The Center for Machine Perception Dept. of Cybernetics, FEE, Czech Technical University in Prague, Czech Republic

Abstract

Computer vision systems for fungi recognition aid mycologists, researchers, and enthusiasts in the efficient identification of mushroom species. FungiCLEF 2023, the second edition of the fungi recognition challenge at LifeCLEF, builds upon the Danish Fungi 2020 dataset and upon its predecessor by presenting several recognition tasks differing in the cost functions corresponding to different practical scenarios, including poisonous/edible decision making or discovering unseen species. With practical applications in mind, the 2023 challenge only accepted submissions with model size under 1GB. The competition received 16 final submissions from 3 teams. This overview paper provides a detailed description of the challenge data and tasks, a review of the submitted methods, and a discussion of the results.

Keywords

LifeCLEF, FungiCLEF, fine grained visual categorization, metadata, open-set recognition, fungi, species identification, machine learning, computer vision, classification

1. Introduction

Computer vision systems for fungi recognition [1, 2] aid mycologists, researchers, and enthusiasts in the efficient identification of mushroom species: A process that could take even hours of searching in the literature is now a question of seconds. For example, with the fungi identification service provided by the Atlas of Danish Fungi [3], users only take a picture of their observation, and the system instantly proposes a list of predicted species for the observation, allowing the user to manually verify the prediction by comparing the observation with photos of the species and with the matching description. Species recognition presents a challenging fine-grained recognition problem with a high number of categories (species) with high similarities between categories – some of them even genetically related – and high intra-class variances at the same time – as the observations depend on a number of factors, such as genotype, age, time of year, local conditions, etc. As part of LifeCLEF 2023 [4, 5], and FGVC-CVPR workshops, FungiCLEF 2023 follows on the previous edition [6, 7] with new test data described in Section 3.1, constraints on the models' memory footprint and a set of species recognition scenarios with different loss functions, described later in Section 3.2.

CLEF 2023: Conference and Labs of the Evaluation Forum, September 18–21, 2023, Thessaloniki, Greece

*Corresponding author.

✉ picekl@kky.zcu.cz (L. Pícek)

🆔 0000-0002-6041-9722 (L. Pícek); 0000-0002-6321-0131 (M. Šulc); 0000-0003-1728-8939 (R. Chamidullin);

0000-0003-0863-4844 (J. Matas)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

2. Related Work

With advances in computer vision and machine learning, fine-grained categorization problems gained significant attention due to the challenges of discriminating subtle differences among visually similar objects. Image-based species recognition is a challenging case of fine-grained classification due to high inter-class similarities, high intra-class variances, large numbers of categories, etc.

Image-based Fungi Recognition: Computer vision processing for fungi recognition dates back to analysis tools from morphological characteristics by De Vooren et al. [8] in 1992. More recently, Tahir et al. [9], and Zielinski et al. [10] introduced datasets of microscopy images of fungal infections for medical diagnosis. Fine-grained fungi recognition "*in the wild*" has been the aim of the FGVCx Fungi classification challenge 2018. More recently, data from the Atlas of Danish Fungi were used to build the Danish Fungi 2020 dataset [11], which is further described in Section 3.1 and which was used in the FungiCLEF 2022 and FungiCLEF 2023 competitions. All contributions to FungiCLEF 2022 [12] were based on modern Convolutional Neural Network (CNN) or transformer-inspired architectures, such as Metaformer [13], Swin Transformer [14], and BEiT [15]. The best performing teams used ensembles of both CNNs and Transformers. The winning team [16] achieved 80.43% accuracy with a combination of ConvNext-large [17] and MetaFormer. The results were often improved by combining predictions belonging to the same observation and by both training-time and test-time data augmentations. We have seen successful applications of different loss functions targeting the unbalanced class distributions - namely the Seesaw loss [18], Focal loss [19], Arcface loss [20], Sub-Center loss [21] and Adaptive Margin [22]. A number of methods have been experimented to utilize the observation metadata, including hand-crafted encoding of the metadata into feature vectors, as well as encoding of the metadata with a multilingual BERT model [23] and RoBERTa [24]. The metadata were then combined with image features extracted from a CNN or Transformer image classifier, or directly used as an input to Metaformer.

Image-based Plant Recognition: While plants (*Plantae*) and *fungi* are different kingdoms, both are mainly immobile and have similarities in general morphology and growth habitat. From the technical point of view, one can thus consider image based recognition of plant and fungi species to be related tasks. Compared to fungi recognition, a higher number of plant recognition datasets are available to the scientific community: Ranging from datasets of specific plant organs like leaves [25, 26] and flowers [27], to general plant recognition "*in the wild*" [28, 29]. Plant recognition from more constrained visual inputs, such as images of tree bark or leaves scanned on a white background, were addressed both with a wide range of hand-crafted features - such as shape [30, 26], texture [31], as well as with deep learning [32]. Plant recognition "*in the wild*" without major constraints on the visual inputs is largely dominated by deep learning approaches, as shown in the previous PlantCLEF competitions [28, 33, 29] and the best performing methods [34, 35, 36]. Moreover, "Experts vs. Machines" experiments performed in [37, 33] showed that the accuracy of the best deep convolutional networks was on the human level¹.

¹In the experiment, only image data were provided to the human experts, not the physical specimen the experts are often used to examine in practice.

3. Challenge Description

Automatic recognition of species at scale, such as in popular citizen-science projects [1, 2], requires efficient prediction on limited resources. In practice, species identification typically depends not solely on the visual observation of the specimen but also on other information available to the observer, e.g., habitat, substrate, location, and time. Thanks to rich metadata, precise annotations, and baselines available to all competitors, the challenge aims at providing a major benchmark for combining visual observations with other observed information. Interestingly, such information might be already included in the observation images (see Figure 1). Additionally, the 2023 edition of the FungiCLEF competition considers decision processes for different usage scenarios, which go beyond the commonly assumed 0/1 cost function – e.g., cost for misclassification of edible and poisonous mushrooms is an important practical aspect to be evaluated.

3.1. Dataset

The FungiCLEF 2023 dataset is based on data collected through the Atlas of Danish Fungi mobile (iOS and Android) and Web applications. All fungi specimen observation had to pass the expert validation process, therefore guaranteeing high-quality labels. Besides high-quality labels, rich observation metadata about habitat, substrate, edibility, time, location, EXIF, etc., is provided. For training the DanishFungi 2020 (DF20) dataset [11] has been provided. The DF20 contains 295,938 images – newly split into 177,170 observations – belonging to 1,604 species. The validation and test datasets were constructed from all observations submitted to the Atlas of Danish Fungi in 2021, for which expert-verified species labels are available covering observations collected across all substrate and habitat types.



Figure 1: Selected observation from the FungiCLEF 2023 private test dataset. ©Peter Stidsen.

Both datasets (Validation and Public Test) cover a similar number of observations, images, and species. The validation set contains 30,131 observations with 60,832 images belonging to 2,713 species: 1,084 known from the training set and 1,629 unknown species. The public test set contains 30,130 observations with 60,225 images belonging to 2,650 species: 1,085 known from the training set and 1,565 unknown species. In both datasets, roughly 30% of the observations captured unknown species. The private test set, used for the official leaderboard after the challenge deadline, consists of 45,021 observations with 91,231 images and corresponding metadata. The dataset statistics are listed in Table 1.

Table 1

FungiCLEF 2023 dataset statistics (taxonomic coverage) for each split and subset.

Subset	Species	Known Species	Unknown Species	Images	Observations
Training	1,604	1,604	–	295,938	177,170
Validation	2,713	1,084	1,629	60,832	30,131
Public Test	2,650	1,085	1,565	60,225	30,130
Private Test	3,299	1,116	2,183	91,231	45,021

3.2. Recognition Scenarios

Given the set of visual fungi observations and corresponding metadata, the goal is to create return a ranked list of predicted species for each observation (one or more photographs of the same individual + geographical location). To focus the challenge on models practically applicable e.g. on mobile devices, as opposed too large ensembles, models in the official benchmark have to fit a limit for memory footprint of 1 GB (in the ONNX format).

The FungiCLEF 2023 recognition challenge used several metrics representing different decision scenarios, where the goal is to minimize the empirical loss L for decisions $q(x)$ over observations x and true labels y , given a cost function $W(y, q(x))$.

$$L = \sum_i W(k_i, q(x_i)) \quad (1)$$

Different recognition scenarios and their cost function $W(y, q(x))$ are described together with their motivation in the following subsections.

Standard classification with "unknown" category. The first metric is the standard classification accuracy, i.e. the average correctness of the predicted class. All species not represented in the training set should correctly be classified as an "unknown" category. The decision function is simple: for each observation is simply represented by an identity matrix, i.e.

$$W_1(y, q(x)) = \begin{cases} 0 & \text{if } q(x) = y \\ 1 & \text{otherwise} \end{cases} \quad (2)$$

Cost for confusing edible species for poisonous and vice versa. Let us have a function d that indicates dangerous (poisonous) species as $d(y) = 1$ if species y is poisonous, and

$d(y) = 0$ otherwise. Let us denote c_{PSC} the cost for poisonous species confusion (if a poisonous observation was misclassified as edible) and c_{ESC} the cost for edible species confusion (if an edible observation was misclassified as poisonous).

$$W_2(y, q(x)) = \begin{cases} 0 & \text{if } d(y) = d(q(y)) \\ c_{\text{PSC}} & \text{if } d(y) = 1 \text{ and } d(q(y)) = 0 \\ c_{\text{ESC}} & \text{otherwise} \end{cases} \quad (3)$$

For the benchmark, we set $c_{\text{ESC}} = 1$ and $c_{\text{PSC}} = 100$.

A user-focused loss composed of both the classification error and the poisonous/edible confusion Assuming the user is interested both in the species classification as well as in recognition of poisonousness, the third cost function simply combines W_1 and W_2 :

$$W_3(y, q(x)) = W_1(y, q(x)) + W_2(y, q(x)) \quad (4)$$

Cost for missing "unknown" species is higher; misclassifying for "unknown" is cheaper than confusing species. When collecting data for biological research, as common for citizen-science projects like the Atlas of Danish Fungi, missing an observation of a new/unknown species comes with a higher cost $\alpha > 1$. On the other hand, checking a known species misclassified as "unknown" may come with a lower cost $0 < \beta < 1$ than trusting misclassification for a different species. The decision cost for each observation can be represented by the following cost matrix (where last row corresponds to unknown species, and last column to observations classified as unknown):

$$W_4 = \begin{bmatrix} 0 & \mathbf{1} & \beta \\ \mathbf{1} & \cdot & \beta \\ \alpha & \alpha & 0 \end{bmatrix} \quad (5)$$

For the benchmark, we set $\alpha = 10$ and $\beta = 0.1$.

3.3. Timeline

The FungiCLEF 2023 competition was announced alongside the release of the dataset on February 14, 2023, through the LifeCLEF, HuggingFace, and FGVC challenge pages, welcoming registrations and participation from anyone interested. The competition spanned approximately 3.5 months, with a deadline set for May 24. Unlike the previous year, the test data was kept secret. Participants were allowed to submit one submission a day using the HuggingFace evaluation platform and the provided public test set. Two weeks before the deadline, the submission limit was increased to five submissions per day. The final evaluation was done over the submitted "code+models." To simplify this process, a dockerizer sample "code+model submission" was provided one week prior to the deadline, allowing all participants to submit their "models" easily. After the submission phase finished, participants could submit post-competition entries to evaluate their ablation studies.

3.4. Working Notes

Participants were strongly encouraged to include both code and a technical report (Working Notes) containing all the necessary details to reproduce the results of their submissions. The Working Note papers submitted by participants underwent a rigorous review process, with 2–3 reviewers possessing a solid publication history in the fields of Computer Vision and Machine Learning. This ensured a sufficient level of reproducibility and maintained a high standard of quality. The review process followed a single-blind approach, and participants were given the opportunity to provide up to two rebuttals to address any concerns or comments raised during the review.

4. Participants and Methods

Twelve teams participated in the FungiCLEF 2023 challenge; four provided their models for a private evaluation, and three submitted working notes. This section provides an overview of the methods and systems employed by the participating teams, with further elaboration available in the individual working notes submitted by the participants [38, 39, 40].

meng18 [38]: The team achieved the best results in the competition, combined visual information with metadata using MetaFormer [13]. To tackle class imbalance, the authors utilized the Seesaw loss [18]. On top of that, this contribution introduced two novelties: an entropy-guided recognition of unknown species in the open-set recognition scenario and an additional poisonous-classification loss to enhance the identification of poisonous species.

stefanwolf [39], scoring second in the competition leaderboard, participated with a model based on Swin Transformer V2 [14]. A combination of predictions for images belonging to the same observation is done on the feature vector level before the classification head, following the last year’s practice of Wolf and Beyerer [41]. To cope with the long-tailed class distribution, authors use the data resampling scheme by Gupta et al. [42].

word2vector [40]: The authors, scoring third in the competition leaderboard, based their final submission on the VOLO [43] backbone architecture without any utilisation of metadata (as methods experimented in [40] did not bring an improvement in results). The class imbalance is tackled by optimising the Seesaw loss [18]. Additional improvements were achieved by rich data augmentation using the Albumentations library [44], including several mixing augmentation methods.

5. Challenge Results

The official challenge results, based on Track 1 to Track 4 Metrics, are displayed in Figure 2. Apart from that, we show the public leaderboard scores in Figure 3. The best-performing team on the private test set – *meng18* – achieved best scores in all measured scenarios. More precisely this team achieved F_1^m score of 57.15%, Poisonous → Edible confusion rate of 5.31%

and Edible → Poisonous confusion rate of 2.05%. Interestingly, none of the teams that submitted working notes optimized decision-making for each of the five tasks.

Edible ↔ Poisonous Species Confusion: Overall, all participants achieved relatively small Poisonous → Edible species confusion (2–2.5%). However the Edible → Poisonous confusion still remains relatively high.




Figure 2: Official FungiCLEF 2023 competition results – 4 teams. Private Leaderboard.




Figure 3: FungiCLEF 2023 competition results – Top10 teams on public test set. Public Leaderboard.

6. Conclusions

This paper presents an overview and results evaluation of the second edition of the FungiCLEF challenge organized in conjunction with the CLEF LifeCLEF lab [5], and CVPR-FGVC10 – The Tenth Workshop on Fine-Grained Visual Categorization organized within the CVPR conference. FungiCLEF 2023 followed up on the previous edition [12] and introduced two novel practical aspects: limitation of the model size (1GB), and different cost functions depending on the considered application.

While the best performing team introduced an additional loss towards enhancing identification of poisonous species, none of the teams optimized their models specifically for all considered cost functions. The number of contributions decreased noticeably from the previous edition of FungiCLEF. The limitation of model size, practically checked by submitting the code with trained models, may have discouraged a number of participants. On the other hand, all submitted contributions focused on the practical aspects of the recognition tasks, rather than scaling-up impractically large ensembles of machine learning models. The results suggest Fungi Recognition, as well as recognition with different application-motivated cost functions, remain an open problems with great room for improvements in the future.

Acknowledgments

RCH, LP, and JM were supported by the Technology Agency of the Czech Republic, project No. SS05010008. RCH was supported by the UWB grant, project No. SGS-2022-017. We acknowledge the funding of the challenge awards provided by the Visual Recognition Group at the Dept. of Cybernetics, Faculty of Electrical Engineering, Czech Technical University in Prague.

References

- [1] M. Sulc, L. Picek, J. Matas, T. Jeppesen, J. Heilmann-Clausen, Fungi recognition: A practical use case, in: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2020, pp. 2316–2324.
- [2] L. Picek, M. Šulc, J. Matas, J. Heilmann-Clausen, T. S. Jeppesen, E. Lind, Automatic fungi recognition: Deep learning meets mycology, *Sensors* 22 (2022) 633.
- [3] T. G. Frøslev, J. Heilmann-Clausen, C. Lange, T. Læssøe, J. H. Petersen, U. Søchting, T. S. Jeppesen, J. Vesterholt, Danish mycological society, fungal records database (2019).
- [4] A. Joly, C. Botella, L. Picek, S. Kahl, H. Goëau, B. Deneu, D. Marcos, J. Estopinan, R. Chamidullin, M. Šulc, M. Hruz, M. Servajean, B. Kellenberger, E. Cole, H. Glotin, et al., Overview of lifeclef 2023: evaluation of ai models for the identification and prediction of birds, plants, snakes and fungi, in: Experimental IR Meets Multilinguality, Multimodality, and Interaction: 14th International Conference of the CLEF Association, CLEF 2023, Thessaloniki, Greece, September 18–23, 2023, Proceedings, Springer, 2023.
- [5] A. Joly, H. Goëau, S. Kahl, L. Picek, C. Botella, D. Marcos, M. Šulc, M. Hruz, T. Lorieul, S. S. Moussi, et al., Lifeclef 2023 teaser: Species identification and prediction challenges, in: European Conference on Information Retrieval, Springer, 2023, pp. 568–576.

- [6] A. Joly, H. Goëau, S. Kahl, L. Picek, T. Lorieul, E. Cole, B. Deneu, M. Servajean, A. Durso, H. Glotin, et al., Overview of lifeclef 2022: an evaluation of machine-learning based species identification and species distribution prediction, in: *Experimental IR Meets Multilinguality, Multimodality, and Interaction: 13th International Conference of the CLEF Association, CLEF 2022, Bologna, Italy, September 5–8, 2022, Proceedings*, Springer, 2022, pp. 257–285.
- [7] L. Picek, M. Šulc, J. Matas, J. Heilmann-Clausen, Overview of fungiclef 2022: Fungi recognition as an open set classification problem (2022).
- [8] J. Van De Vooren, G. Polder, G. van der Heijden, Identification of mushroom cultivars using image analysis, *Transactions of the ASAE* 35 (1992) 347–350.
- [9] M. W. Tahir, N. A. Zaidi, A. A. Rao, R. Blank, M. J. Vellekoop, W. Lang, A fungus spores dataset and a convolutional neural network based approach for fungus detection, *IEEE transactions on nanobioscience* 17 (2018) 281–290.
- [10] B. Zielinski, A. Sroka-Oleksiak, D. Rymarczyk, A. Piekarczyk, M. Brzychczy-Wloch, Deep learning approach to description and classification of fungi microscopic images, *CoRR abs/1906.09449* (2019). URL: <http://arxiv.org/abs/1906.09449>. arXiv:1906.09449.
- [11] L. Picek, M. Šulc, J. Matas, T. S. Jeppesen, J. Heilmann-Clausen, T. Læssøe, T. Frøslev, Danish fungi 2020-not just another image recognition dataset, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2022*, pp. 1525–1535.
- [12] L. Picek, M. Šulc, J. Matas, J. Heilmann-Clausen, Overview of fungiclef 2022: Fungi recognition as an open set classification problem, *CEUR-WS*, 2022.
- [13] Q. Diao, Y. Jiang, B. Wen, J. Sun, Z. Yuan, Metaformer: A unified meta framework for fine-grained recognition, *arXiv preprint arXiv:2203.02751* (2022).
- [14] Z. Liu, H. Hu, Y. Lin, Z. Yao, Z. Xie, Y. Wei, J. Ning, Y. Cao, Z. Zhang, L. Dong, et al., Swin transformer v2: Scaling up capacity and resolution, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2022*, pp. 12009–12019.
- [15] H. Bao, L. Dong, F. Wei, Beit: Bert pre-training of image transformers, *arXiv preprint arXiv:2106.08254* (2021).
- [16] Z. Xiong, Y. Ruan, Y. Hu, Y. Zhang, Y. Zhu, S. Guo, W. Zhu, B. Han, An empirical study for fine-grained fungi recognition with transformer and convnet, in: *Working Notes of CLEF 2022 - Conference and Labs of the Evaluation Forum, 2022*.
- [17] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, S. Xie, A convnet for the 2020s, *arXiv preprint arXiv:2201.03545* (2022).
- [18] J. Wang, W. Zhang, Y. Zang, Y. Cao, J. Pang, T. Gong, K. Chen, Z. Liu, C. C. Loy, D. Lin, Seesaw loss for long-tailed instance segmentation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021*, pp. 9695–9704.
- [19] T.-Y. Lin, P. Goyal, R. Girshick, K. He, P. Dollár, Focal loss for dense object detection, in: *Proceedings of the IEEE international conference on computer vision, 2017*, pp. 2980–2988.
- [20] J. Deng, J. Guo, N. Xue, S. Zafeiriou, Arcface: Additive angular margin loss for deep face recognition, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2019*, pp. 4690–4699.
- [21] J. Deng, J. Guo, T. Liu, M. Gong, S. Zafeiriou, Sub-center arcface: Boosting face recognition by large-scale noisy web faces, in: *European Conference on Computer Vision, Springer, 2020*, pp. 741–757.

- [22] H. Liu, X. Zhu, Z. Lei, S. Z. Li, Adaptiveface: Adaptive margin and sampling for face recognition, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 11947–11956.
- [23] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, arXiv preprint arXiv:1810.04805 (2018).
- [24] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, Roberta: A robustly optimized bert pretraining approach, arXiv preprint arXiv:1907.11692 (2019).
- [25] P. Novotný, T. Suk, Leaf recognition of woody species in central europe, Biosystems Engineering 115 (2013) 444–452.
- [26] N. Kumar, P. N. Belhumeur, A. Biswas, D. W. Jacobs, W. J. Kress, I. C. Lopez, J. V. Soares, Leaf-snap: A computer vision system for automatic plant species identification, in: Computer Vision–ECCV 2012, Springer, 2012, pp. 502–516.
- [27] FGVCx flower classification challenge 2018, <https://sites.google.com/view/fgvc5/competitions/fgvcx/flowers>, URL: <https://sites.google.com/view/fgvc5/competitions/fgvcx/flowers>, accessed: 2023-07-4.
- [28] H. Goëau, P. Bonnet, A. Joly, Plant identification based on noisy web data: the amazing performance of deep learning (lifeclef 2017), in: CLEF task overview 2017, CLEF: Conference and Labs of the Evaluation Forum, Sep. 2017, Dublin, Ireland., 2017.
- [29] H. Goëau, P. Bonnet, A. Joly, Overview of lifeclef plant identification task 2019: diving into data deficient tropical countries, in: CLEF task overview 2019, CLEF: Conference and Labs of the Evaluation Forum, Sep. 2019, Lugano, Switzerland., 2019.
- [30] A. Kadir, L. E. Nugroho, A. Susanto, P. I. Santosa, A comparative experiment of several shape methods in recognizing plants., International Journal of Computer Science & Information Technology 3 (2011).
- [31] M. Šulc, J. Matas, Texture-based leaf identification, in: Computer Vision - ECCV 2014 Workshop, volume 8928 of LNCS, Springer, 2015, pp. 185–200. doi:10.1007/978-3-319-16220-1_14.
- [32] M. Šulc, J. Matas, Fine-grained recognition of plants from images, Plant Methods 13 (2017) 1–14.
- [33] P. Bonnet, H. Goëau, S. T. Hang, M. Lasseck, M. Šulc, V. Malécot, P. Jauzein, J.-C. Melet, C. You, A. Joly, Plant identification: Experts vs. machines in the era of deep learning: deep learning techniques challenge flora experts, Multimedia tools and applications for environmental & biodiversity informatics (2018) 131–149.
- [34] M. Lasseck, Image-based plant species identification with deep convolutional neural networks., in: CLEF (Working Notes), 2017.
- [35] M. Sulc, L. Pícek, J. Matas, Plant recognition by inception networks with test-time class prior estimation., in: CLEF (Working Notes), 2018.
- [36] L. Pícek, M. Sulc, J. Matas, Recognition of the amazonian flora by inception networks with test-time class prior estimation., CLEF (Working Notes) 188 (2019).
- [37] H. Goëau, A. Joly, P. Bonnet, M. Lasseck, M. Šulc, S. T. Hang, Deep learning for plant identification: how the web can compete with human experts, Biodiversity Information Science and Standards 2 (2018) e25637.
- [38] H. Ren, H. Jiang, W. Luo, M. Meng, T. Zhang, Entropy-guided open-set fine-grained fungi

- recognition, in: Working Notes of CLEF 2023 - Conference and Labs of the Evaluation Forum, 2023.
- [39] S. Wolf, J. Beyerer, It is all about the metrics: Optimizing transformer-based fungi-classification for multiple application-oriented metrics, in: Working Notes of CLEF 2023 - Conference and Labs of the Evaluation Forum, 2023.
 - [40] F. Hu, P. Wang, Y. Li, C. Duan, Z. Zhu, Y. Li, X.-S. Wei, A deep learning based solution to fungiclef2023, in: Working Notes of CLEF 2023 - Conference and Labs of the Evaluation Forum, 2023.
 - [41] S. Wolf, J. Beyerer, Transformer-based fine-grained fungi classification in an open-set scenario, Working Notes of CLEF (2022).
 - [42] A. Gupta, P. Dollar, R. Girshick, Lvis: A dataset for large vocabulary instance segmentation, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2019, pp. 5356–5364.
 - [43] L. Yuan, Q. Hou, Z. Jiang, J. Feng, S. Yan, Volo: Vision outlooker for visual recognition, IEEE transactions on pattern analysis and machine intelligence 45 (2022) 6575–6586.
 - [44] A. Buslaev, V. I. Iglovikov, E. Khvedchenya, A. Parinov, M. Druzhinin, A. A. Kalinin, Albumentations: fast and flexible image augmentations, Information 11 (2020) 125.