

Fungi Recognition: A Practical Use Case

Milan Sulc
CTU in Prague

sulcmila@cmp.felk.cvut.cz

Lukas Picek
University of West Bohemia

picekl@kky.zcu.cz

Jiri Matas
CTU in Prague

matas@cmp.felk.cvut.cz

Thomas S. Jeppesen
Global Biodiversity Information Facility

tsjeppesen@gbif.org

Jacob Heilmann-Clausen
University of Copenhagen

jheilmann-clausen@snm.ku.dk

Abstract

The paper presents a system for visual recognition of 1394 fungi species based on deep convolutional neural networks and its deployment in a citizen-science project. The system allows users to automatically identify observed specimens, while providing valuable data to biologists and computer vision researchers. The underlying classification method scored first in the FGVCx Fungi Classification Kaggle competition organized in connection with the Fine-Grained Visual Categorization (FGVC) workshop at CVPR 2018. We describe our winning submission and evaluate all technicalities that increased the recognition scores, and discuss the issues related to deployment of the system via the web- and mobile- interfaces.

1. Introduction

The collection of data on appearance and occurrence of species and its annotation are crucial pillars for biological research focusing on biodiversity, climate change and species extinction [6, 18]. Involvement of citizen communities is a cost effective approach to large scale data collection. Species observation datasets collected by the broader public have already proven to add significant value for understanding both basic and more applied aspects of mycology (e.g. [2, 34]), and by improving data quality and participation in such programs, the research potential will increase. Citizen-science contributions provide about 50% of all data accessible through the Global Biodiversity Information Facility [3]. However, the data has a strong taxonomic bias towards birds and mammals [31], leaving data gaps in taxonomic groups such as fungi and insect.

Correct species identification is a challenge in citizen-science projects focusing on biodiversity. Some projects handle the issue by simply reducing complexity in the species identification process, e.g. by merging species into

multitaxa indicator groups (e.g. [9]), by focusing only on a subset of easily identifiable species or by involving human expert validators in the identification process. Other projects involve citizen-science communities in the data validation process. For instance, iNaturalist [1] regards observations as having research grade if three independent users have verified a suggested taxon ID based on an uploaded photo. Automatic image-based species identification can act both as a supplement or alternative to these approaches.

We present a computer vision system for recognition of fungi "in the wild", achieving best results in a Kaggle competition organized with the Fine-Grained Categorization Workshop at CVPR 2018, and further application of this system to assist a citizen-science community and help mycologists increase the involvement of citizens in data collection.

Applications for image-based mushroom recognition are reviewed in Section 2.1. To the best of our knowledge, our system recognizes the largest number of species, and it is the first image-based fungi recognition system to assist citizen-scientists and mycologists in identification and collection of observations.

From the computer vision perspective, the application of the system to citizen-science data collection creates a valuable continuous stream of labeled examples for a challenging fine-grained visual classification task. The increasing amount of labeled data will allow us to improve the classification baselines and to study other interesting problems, such as fungi phenotyping, location-based estimation of categorical prior, etc.

The system described here has a big potential to increase human involvement with nature by providing an real-time electronic identification tool, that can support learning in an intuitive manner, much like children learn from their parents by asking simple and naive questions that are addressed in a simple way. By linking the system to an existing my-

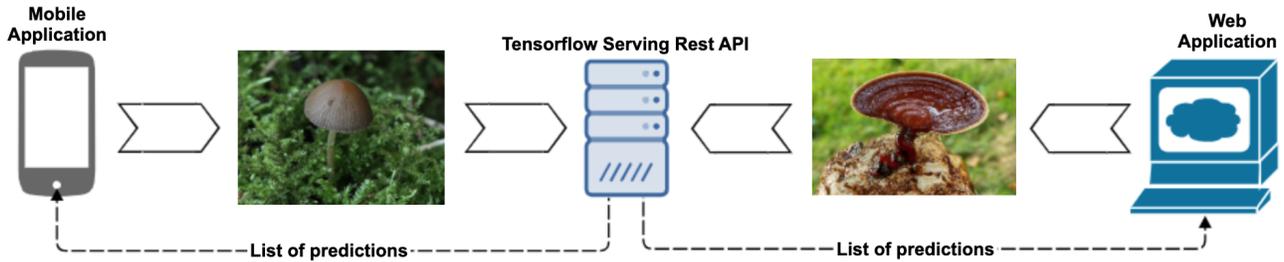


Figure 1: The fungi recognition serving pipeline.

ecological platform involving validation by the community, as is the case in the Danish Fungal Atlas [4, 8, 14], a supervised machine learning system with human in the loop is created.

2. Related work

2.1. Fungi Recognition

Several mobile applications for fungi identification include a computer vision classification system. Only few have positive user reviews on the identification results. Examples of apps with positive user reviews are:

- Mushroom Identifier¹ with 1M+ downloads and a review score of 3.9/5, recognizing 900 mushroom species,
- Mushrooms App² with 0.5M+ downloads and a review score of 4.4/5, recognizing 200 mushroom species.

De Vooren et al. [32] published an image analysis tool for mushroom cultivars identification in 1992, analyzing morphological characters like length, width and other shape descriptors.

Computer vision may also be used for classification of microscopy images of fungal spores. Tahir et al. [29] and Zielinski et al. [38] introduce datasets of microscopy images of fungal infections and propose methods to speed up medical diagnosis, allowing to avoid additional expensive biochemical tests.

2.2. Crowd-based Image Collection and Identification

The **Global Biodiversity Information Facility (GBIF)** [10] is the largest index of biodiversity data in the world. GBIF is organized as a network involving 58 participating countries and 38 organisations (mainly international) publishing more than 45,000 biodiversity datasets under open

¹<https://play.google.com/store/apps/details?id=com.pingou.champignonuf> Accessed on 2019-10-11

²<https://play.google.com/store/apps/details?id=basinac.aplikacenahouby> Accessed on 2019-10-11

source licenses. The index contains more than 1.3 billion species occurrence records of which more than 47 million include images. With the recent advances in the use of machine vision in biodiversity related technology, GBIF intends to facilitate collaborations in this field, promote responsible data use and good citation practices. GBIF has the potential to play an active role in preparing training datasets and make them accessible under open source licenses [24].

iNaturalist [16] is a pioneering crowd-based platform allowing citizens and experts to upload and categorize observations of the world fauna, flora and fungi. All annotated data are directly uploaded to GBIF once verified by three independent users. iNaturalist covers more than 238,000 species through almost 28 million observations.

Wild Me is a non-profit organization that aims to combat extinction with citizen-science and artificial intelligence. Their projects using computer vision [22] to boost detection and identification include: **Flukebook**, a collaboration system to collect citizen observations of dolphins and whales and to identify individuals, and **GiraffeSpotter**, a photo-identification database of giraffe encounters.

The **Danish Fungal Atlas (SvampeAtlas)** [4, 8, 14] involves more than 1000 volunteers who have contributed approximately 500,000 quality-checked observations of fungi. More than 270,000 old fungal records were imported into the project database which now contains more than 800,000 quality-checked fungal records. The project has resulted in a greatly improved knowledge of Denmark's fungi. More than 180 basidiomycetes³ have been added to the list of known Danish species, and several species that were considered extinct have been re-discovered. At the same time, a number of search and assistance functions have been developed that present common knowledge about the individual species of fungi, which makes it much easier to include knowledge of endangered species in the nature management and decision making.

All validated records are published to the Global Biodiversity Information Facility [10] on a weekly basis. Since

³Microscopic spore-producing structure found on the hymenophore of fruiting bodies.



Figure 2: Examples from the FGVCx Fungi training set.

2017, the Danish Fungal Atlas has had interactive validation of fungal records. When a user submits a record, a probability score is calculated for the accuracy of the identification. This score ranges from 1 to 100. The calculation includes:

1. The rarity of the species (# approved records).
2. The geographical distribution of the species.
3. Phenology of the species (e.g. many mycorrhizal fungi have a low probability score in spring).
4. User’s previous approved identifications of the same species.
5. Nr. of species within the morphological group the user has correctly identified in the past.
6. Confidence indicated by the user: Certain: 100%, Probable: 50%, Possible: 10%.

Subsequently, other users may agree on the identification, increasing the identification score in accordance with the principles 4–6, or propose alternative identifications. The identification with the highest score is highlighted, alternative identifications and their scores are also visible to logged-in users. In the search results, the probability score is displayed in three general categories:

1. Approved (score above 80) with 3 stars.
2. Likely (score between 50 and 80) with 2 stars.
3. Suggestion (score below 50) with 1 star.

A group of taxonomic experts (validators) are monitoring data in the Danish Mushroom Atlas. These have the power to approve findings regardless of the score in the interactive validation. This can be relevant for discoveries of new species, for very rare species and for records of species where special experience or sequencing of genetic material (DNA) is required for a safe identification. Expert-validated findings are marked with a small microscope icon.

2.3. Fine-grained Image Classification

The task of image-based fungi recognition is a fine-grained visual classification (or categorization) problem.

Fine-grained image classification went through significant improvements with the emergence of very deep convolutional neural networks (CNNs) and the success of Krizhevsky’s CNN [20] in the ImageNet ILSVRC-12 competition. The ImageNet dataset itself contains a number of species categories, mainly animals. Convolutional Neural Networks performed well in other fine-grained species identification tasks, including plant species classification [11, 12], dog classification [19], bird classification [35, 36], or classification of species in general [33].

3. Image Recognition Methodology

3.1. FGVCx Fungi Dataset

The FGVCx Fungi Classification Challenge provided an image dataset, that covers 1394 fungal species and is split into a training set with 85578 images, a validation set with 4182 images and a competition test set with 9758 images without publicly available labels. Examples from the training set are shown in Figure 2. There is a substantial change of categorical priors $p(k)$ between the training set and the validation set: The distribution of images per class is highly unbalanced in the training set, while the validation set distribution is uniform.

3.2. Convolutional Neural Networks

Following the advances in deep learning for fine-grained image classification, we decided to approach fungi recognition with Convolutional Neural Networks. For the FGVCx Fungi Classification challenge, we trained an ensemble of Inception-v4 and Inception-ResNet-v2 networks [28], inspired by the winning submission in the ExpertLifeCLEF plant identification challenge 2018 [11].

We trained an ensemble of 6 models listed in Table 1. All networks were trained using the Tensorflow Slim⁴ framework. We used Polyak averaging [23], keeping shadow variables with exponential moving averages of the trained variables. Hyper-parameters used during training were set

⁴<https://github.com/tensorflow/models/tree/master/research/slim>

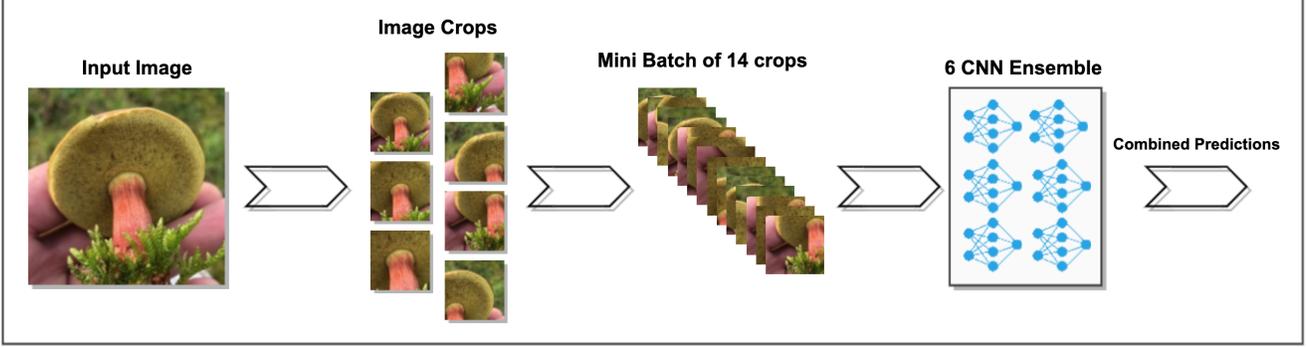


Figure 3: Predictions combined from an ensemble of 6 CNNs with test-time image augmentation (crops, mirrors).

CNN	Architecture	Input Size	Finetuned from
#1	Inception-v4	299x299	ImageNet 2012
#2	Inception-v4	299x299	LifeCLEF 2018
#3	Inception-v4 "x2"	598x598	ImageNet 2012
#4	Inception-v4 "x2"	598x598	LifeCLEF 2018
#5	Inc.-ResNet-v2	299x299	ImageNet 2012
#6	Inc.-ResNet-v2	299x299	LifeCLEF 2018

Table 1: Models trained for the FGVCx Fungi classification competition.

as follows - optimizer: RMSprop, batch size: 32, initial learning rate: 0.01, learning rate decay: exponential/staircase with decay factor 0.94, weight decay: 0.00004, moving average decay: 0.999. All six fine-tuned networks are publicly available⁵.

3.3. Adjusting Predictions by Class Priors

Let us assume that the classifier trained by cross-entropy minimization learns to estimate the posterior probabilities, i.e. $f_{\text{CNN}}(k|x) \approx p(k|x)$. If the class prior probabilities $p_e(k)$ change, the posterior probabilities should change as well. The topic of adjusting CNN predictions to new priors is discussed in [7, 25, 27]: in the case when the new class priors $p_e(k)$ are known, the new posterior $p_e(k|x)$ can be computed as:

$$\begin{aligned}
 p_e(k|\mathbf{x}_i) &= p(k|\mathbf{x}_i) \frac{p_e(k)p(\mathbf{x}_i)}{p(k)p_e(\mathbf{x}_i)} = \\
 &= \frac{p(k|\mathbf{x}_i) \frac{p_e(k)}{p(k)}}{\sum_{j=1}^K p(j|\mathbf{x}_i) \frac{p_e(j)}{p(j)}} \propto p(k|\mathbf{x}_i) \frac{p_e(k)}{p(k)}, \quad (1)
 \end{aligned}$$

⁵<https://github.com/sulc/fungi-recognition>

where we used $\sum_{k=1}^K p_e(k|\mathbf{x}_i) = 1$ to get rid of the unknown probabilities $p(\mathbf{x}_i), p_e(\mathbf{x}_i)$.

While other works [7, 25, 27] focus on estimating new unknown priors $p_e(k)$, we assume that the uniform distribution $p_e(k) = \frac{1}{K}$ is given, as it is the case of the FGVCx Fungi validation set (see Section 3.1). Then:

$$p_e(k|\mathbf{x}_i) \propto \frac{p(k|\mathbf{x}_i)}{p(k)}. \quad (2)$$

3.4. Test-time Image Augmentation

We considered the following 14 image augmentations at test time: The original image; additional 6 crops of the original image with 80% (central crop) and 60% (central crop + 4 corner crops) of the original image width/height; and the mirrored versions of the 7 foregoing augmentations. All augmentations are then resized to square inputs using bilinear interpolation.

Predictions from all augmentations are then combined by averaging (sum) or mode of the predicted classes. The pipeline is illustrated in Figure 3.

4. Online Fungi Classification Service

In order to provide a flexible and scalable image-based fungi identification service for the Danish Fungal Atlas, we created a recognition server based on the open-source TensorFlow Serving [21] framework. The server currently uses one of our pretrained models, the framework allows to deploy several models at the same time. No test-time augmentations are currently used in order to prevent server overload.

The pipeline is visualized in Figure 1: The web- and mobile apps query the recognition server via Representational State Transfer (REST) API. The server feeds the query image into the Convolutional Network and responds with the

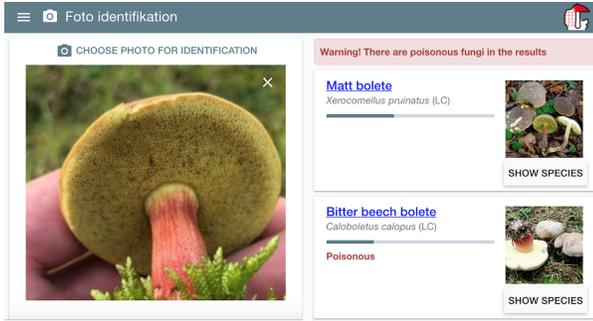


Figure 4: Screenshot from the web-based recognition app (<https://svampe.databasen.org/imagevision>).

list of predicted species probabilities. The apps then display a shortlist of the most likely species for the query. The observation is also uploaded into the Danish Fungal Atlas database. The user can manually inspect the proposed species and select the best result for annotation of the fungus observation. Screenshots of the web and mobile interfaces are shown in Figure 4 and Figure 6 respectively.

Observations uploaded into the Danish Fungal Atlas database and the proposed species identifications are then verified by the community. Images with verified species labels will be used to further fine-tune the recognition system.

5. Results

First, in Section 5.1, we evaluate the accuracy of our models on the validation set before and after applying "tricks" like test-time augmentation, ensembling, or adjusting predictions to new class priors. Second, the official challenge results are summarized in Section 5.2. And last, Section 5.3 presents the first results of the integration of the classification service into the Danish Fungal Atlas.

5.1. FGVCx Fungi Validation Dataset

Let us first validate the CNNs from Section 3.2 on the FGVCx Fungi validation set. Table 2 compares the six trained CNN models before applying additional tricks, with 1 forward pass (central crop, 80%) per image. We will continue the validation experiments with CNN 1, i.e. Inception-v4 pre-trained from an ImageNet checkpoint, which achieved the best validation accuracy.

The test-time pre-processing of the image input makes a noticeable difference. Table 4 shows the difference in accuracy for different sizes of central crop of the original image.

The advantage of adjusting the predictions with the new categorical prior is shown in Figure 5: at the end of training the accuracy increases by 3.8%, from 48.8% to 52.6%.

CNN	Acc. (%)	R@5 (%)
#1 Inception-v4 (ImageNet)	48.8	77.0
#2 Inception-v4 (LifeCLEF)	48.5	75.8
#3 Inception-v4 "x2" (ImageNet)	48.6	76.6
#4 Inception-v4 "x2" (LifeCLEF)	48.8	76.2
#5 Inc.-ResNet-v2 (ImageNet)	47.7	76.0
#6 Inc.-ResNet-v2 (LifeCLEF)	47.4	75.8
Inception-v4 [5]	44.7	73.5

Table 2: Accuracy and Recall@5 of individual networks (central crop, 80%) on the FGVCx Fungi validation set.

Central crop	Accuracy (%)	Recall@5 (%)
100%	45.9	75.1
80%	48.8	77.0
60%	48.6	76.3
40%	43.1	69.3

Table 3: Inception-v4 (finetuned from the ImageNet checkpoint) with differently sized central crops. Top-1 Accuracy and Recall@5 on the FGVCx Fungi validation set.

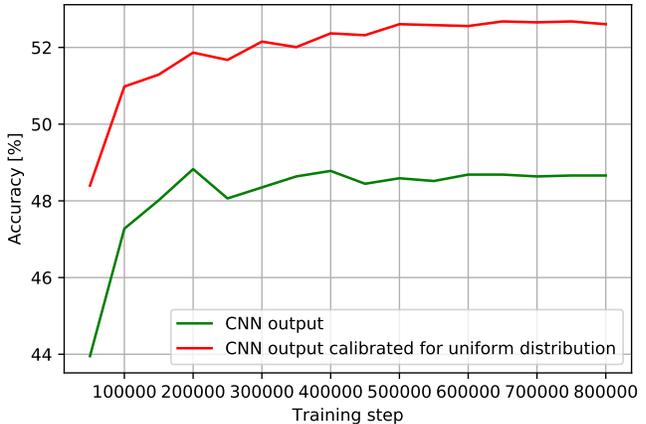


Figure 5: Accuracy of Inception-v4 (finetuned from ImageNet checkpoint) on the FGVCx Fungi validation set, before (green) and after (red) adjusting the predictions by $p_e(k)$.

5.2. FGVCx Fungi Competition

The competition test dataset on Kaggle was divided into two parts - public and private. Public results were calculated with approximately 70% of the test data and results were visible to all participants. The rest of the data was used

#CNNs	Crops	Pool	Accuracy (%)	
			Baseline	Known $p_e(k)$
1	1	–	48.8	52.6
1	14	sum	51.8	56.0
6	1	sum	54.1	58.5
6	14	sum	54.2	60.3
6	14	mode	54.2	59.1

Table 4: Top-1 recognition accuracy on the FGVCx Fungi validation set: single CNN (#1) vs. ensemble (#1,...,#6) and single central crop (1) vs. multiple crops (14). Predictions from ensembles and crops were combined by averaging (sum) or by choosing the most common top prediction (mode). Results are shown both before and after adjusting the predictions by known $p_e(k)$.

#	Team Name	Recal@3 Error (%)	
		Private Score	Public Score
1	(ours)	21.197	20.772
2	digitalspecialists	23.188	23.471
3	Val An	25.091	25.213
4	DL Analytics	28.341	26.853
5	Invincibles	28.751	28.493
6	Tian Xi	32.235	31.636
7	Igor Krashenyi	32.616	34.164
8	wakaka	42.219	41.339
9	George Yu	47.621	47.113
10	Xinshao	67.837	67.509

Table 5: Results of the top ten teams in FGVCx Fungi Classification Challenge. Private and public scores are Recall@3 errors (smaller is better). Source:

<http://kaggle.com/c/fungi-challenge-fgvc-2018/leaderboard>

for final competition evaluation to avoid any possible bias towards performance on the test images.

We chose our best performing system, i.e. the ensemble of the 6 finetuned CNNs with 14 crops per test image and with predictions adjusted to new class priors, for the final submission to Kaggle. The accumulation of predictions was done by the mode from top species per prediction had better preliminary scores on the public Kaggle test set.

Our submission to the challenge achieved the best scores in terms of Recall@3 error both in the public and private leaderboard. The Recall@3 error is defined as follows: for each image, if the ground truth label is found among the top 3 predicted labels, the error is 0, otherwise it is 1. The final score is the error averaged across all images. The results of

the top 10 teams are listed in Table 5.

5.3. Results of the Online Classifier

The experts behind the Danish Fungal Atlas have been highly impressed by the performance of the system⁶; in the application, the results of the system are referred to as AI suggested species. This has been confirmed by a data evaluation where 5760 records have been submitted for automatic recognition, of which only 904 (16 %) were not approved by community- or expert validation. This is a far better performance than most non-expert users in the system. Almost two thirds (64 %) of the approved species identifications were based on the highest ranking AI suggesting species ID, while another 7 % were based on the second highest ranking AI suggested species ID and another 6 % were based on top 3-5 suggestions.

It has not been possible to collect data on identification attempts where no useful match was returned from the AI, and the user therefore picked a taxon name not in the top 10 AI results. However, users generally stated that this rarely happened. So far the system has been tested by 652 users, each submitting between one and 526 records. For users submitting more than ten records the accuracy in terms of correct identifications guided by the system varied from 17% to 100%, pointing to quite considerable differences in how well different users have been able to identify the correct species using the system. Hence, the tool is not fully reliable, but helps the non-expert users to gain better identification skills. The accuracy was variable among the fungal morphogroups defined in the fungal atlas, varying from 24 % to 100 % for groups with more than 10 records. The accuracy was tightly correlated with the obtained morphogroup user score based on the algorithms deployed in the Danish Fungal Atlas to support community validation.

The operators of Danish Fungal Atlas also received positive feedback from several users about the new AI-identification feature.

Within the first month the server has been running, more than 20,000 images have been submitted for recognition. Note that the mobile app with the image recognition feature has only been published at the time of this paper submission, and therefore, we expect an increasing number of recognition queries.

6. Conclusions

The work described the development of a fungi recognition system: from design and validation through winning a computer vision Kaggle challenge to a final application helping citizen-scientists to identify species of observed specimen and motivating their contributions to a citizen-science project.

⁶Personal communication with the Danish Fungal Atlas.

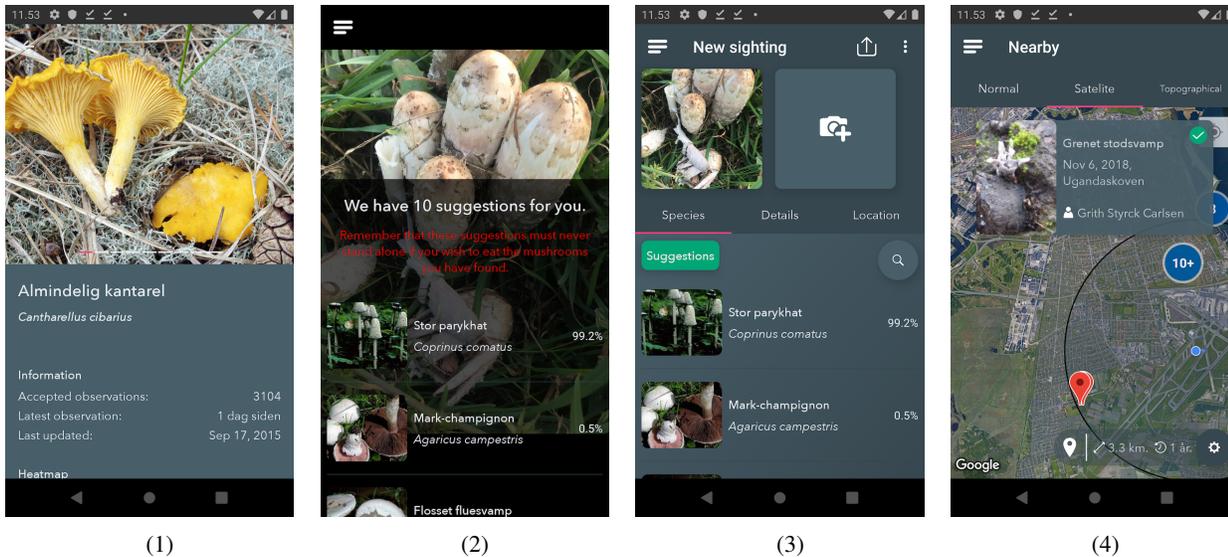


Figure 6: Screenshots from the Android app showing (1) A detailed description of selected species, (2,3) Image based recognition suggesting species for a query image, (4) Map with nearby observations.

Evaluation on the validation set in Section 5.1 showed the effect of calibrating outputs to new a-priori probabilities, test-time data augmentation and ensembles: together, these "tricks" increased the recognition accuracy by almost 12%, and helped us to score 1st in the FGVCx Fungi Classification competition hosted on Kaggle, achieving 79% Recall@3.

Integration of the image recognition system into Danish Fungal Atlas makes community-based fungi observation identification easier: from the first 592 approved annotations, 89% were based on the top-2 predictions of our model.

Cross science efforts such as the collaboration described here can develop tools for citizen-scientists that improve their skills and the quality of the data they generate. Along with data generated by DNA sequencing this may help lowering the taxonomic bias in the biodiversity information data available in the future.

Future work

The server-based inference allows computation of accurate predictions with good response time, and it motivates users to upload images. On-device mobile inference would also allow real-time recognition in areas with limited access to mobile data. Inference on mobile devices would, however, require decreasing model size and complexity. Possible directions for future work include applying efficient architectures [15, 26, 30], weight pruning and quantization [13, 17, 37].

Deeper integration into mycological information sys-

tems may allow on-line learning of the classifier. Extending the collaboration with more mycological institutes or information systems may help to improve the system even further, as it would learn from all available data.

As species distribution differs based on geographical locations and local environment, estimating the priors for different locations may be used to calibrate the predictions for observations with GPS information.

Acknowledgements

MS was supported by CTU student grant SGS17/185/OHK3/3T/13. LP was supported by the Ministry of Education, Youth and Sports of the Czech Republic project No. LO1506, and by the grant of the UWB project No. SGS-2019-027. JM was supported by OP VVV project CZ.02.1.01/0.0/0.0/16_019/0000765 Research Center for Informatics. The Danish Fungal Atlas was supported by Aage V. Jensen Naturfond.

References

- [1] inaturalist. <http://www.inaturalist.org>. Accessed: 2019-12-3.
- [2] C. Andrew, E. Heegaard, P. M. Kirk, C. Bessler, J. Heilmann-Clausen, I. Krisai-Greilhuber, T. W. Kuyper, B. Senn-Irlet, U. Bntgen, J. Diez, S. Egli, A. C. Gange, R. Halvorsen, K. Hiland, J. Nordn, F. Rusten, L. Boddy, and H. Kausrud. Big data integration: Pan-european fungal species observations' assembly for addressing contemporary questions in ecology and global change biology. *Fungal Biology Reviews*, 31(2):88 – 98, 2017.

- [3] M. Chandler, L. See, K. Copas, A. M. Z. Bonde, B. C. López, F. Danielsen, J. K. Legind, S. Masinde, A. J. Miller-Rushing, G. Newman, A. Rosemartin, and E. Turak. Contribution of citizen science towards international biodiversity monitoring. *Biological Conservation*, 213:280–294, 2017.
- [4] Danish Mycological Society. <https://svampe.databasen.org>, 2009.
- [5] Danish Mycological Society, Google Research. Danish Mycological Society fungi embedding model. TensorFlow hub. <https://doi.org/10.26161/mpbc-q144>.
- [6] J. L. Dickinson, J. Shirk, D. Bonter, R. Bonney, R. L. Crain, J. Martin, T. Phillips, and K. Purcell. The current state of citizen science as a tool for ecological research and public engagement. *Frontiers in Ecology and the Environment*, 10(6):291–297, 2012.
- [7] M. C. Du Plessis and M. Sugiyama. Semi-supervised learning of class balance under class-prior change by distribution matching. *Neural Networks*, 50:110–119, 2014.
- [8] T. G. Frslev, J. Heilmann-Clausen, C. Lange, T. Lsse, J. H. Petersen, U. Schting, T. S. Jeppesen, and J. Vesterholt. Danish mycological society, fungal records database, 2019.
- [9] J. Geldmann, J. Heilmann-Clausen, T. E. Holm, I. Levinsky, B. Markussen, K. Olsen, C. Rahbek, and A. P. Ttrup. What determines spatial bias in citizen science? exploring four recording schemes with different proficiency requirements. *Diversity and Distributions*, 22(11):1139–1149, 2016.
- [10] Global Biodiversity Information Facility. <https://www.gbif.org>, 1999.
- [11] H. Goëau, P. Bonnet, and A. Joly. Overview of expertlifecycle 2018: how far automated identification systems are from the best experts? In *CLEF working notes 2018*, 2018.
- [12] H. Goëau, P. Bonnet, and A. Joly. Overview of lifecycle plant identification task 2019: diving into data deficient tropical countries. In *CLEF working notes 2019*, 2019.
- [13] S. Han, H. Mao, and W. J. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
- [14] J. Heilmann-Clausen, H. H. Bruun, R. Ejrnæs, T. G. Frslev, T. Lsse, and J. H. Petersen. How citizen science boosted primary knowledge on fungal biodiversity in denmark. *Biological Conservation*, 237:366 – 372, 2019.
- [15] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [16] iNaturalist.org. inaturalist research-grade observations, 2019.
- [17] B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. Howard, H. Adam, and D. Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2704–2713, 2018.
- [18] W. Jetz, M. A. McGeoch, R. Guralnick, S. Ferrier, J. Beck, M. J. Costello, M. Fernandez, G. N. Geller, P. Keil, C. Merow, C. Meyer, F. E. Muller-Karger, H. M. Pereira, E. C. Regan, D. S. Schmeller, and E. Turak. Essential biodiversity variables for mapping and monitoring species populations. *Nature Ecology & Evolution*, 3(4):539–551, 2019.
- [19] A. Khosla, N. Jayadevaprakash, B. Yao, and L. Fei-Fei. Novel dataset for fine-grained image categorization. In *First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition*, Colorado Springs, CO, June 2011.
- [20] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [21] C. Olston, N. Fiedel, K. Gorovoy, J. Harmsen, L. Lao, F. Li, V. Rajashekhar, S. Ramesh, and J. Soyke. Tensorflow-serving: Flexible, high-performance ML serving. *arXiv preprint arXiv:1712.06139*, 2017.
- [22] J. Parham, C. Stewart, J. Crall, D. Rubenstein, J. Holmberg, and T. Berger-Wolf. An animal detection pipeline for identification. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1075–1083. IEEE, 2018.
- [23] B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- [24] T. Robertson, S. Belongie, H. Adam, C. Kaeser-Chen, C. Zhang, K. C. Tan, Y. Liu, D. Brul, C. Deltheil, S. Loarie, G. V. Horn, O. M. Aodha, S. Beery, P. Perona, K. Copas, and J. ThomasWaller. Training machines to identify species using gbif-mediated datasets. *Biodiversity Information Science and Standards*, 3:e37230, 2019.
- [25] M. Saerens, P. Latinne, and C. Decaestecker. Adjusting the outputs of a classifier to new a priori probabilities: a simple procedure. *Neural computation*, 14(1):21–41, 2002.
- [26] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018.
- [27] M. Sulc and J. Matas. Improving cnn classifiers by estimating test-time priors. *arXiv preprint arXiv:1805.08235v2*, 2019.
- [28] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [29] M. W. Tahir, N. A. Zaidi, A. A. Rao, R. Blank, M. J. Vellekoop, and W. Lang. A fungus spores dataset and a convolutional neural network based approach for fungus detection. *IEEE Transactions on NanoBioscience*, 17(3):281–290, July 2018.
- [30] M. Tan and Q. V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019.
- [31] J. Troudet, P. Grandcolas, A. Blin, R. Vignes-Lebbe, and F. Legendre. Taxonomic bias in biodiversity data and societal preferences. *Scientific Reports*, 7(1):9132, 2017.
- [32] J. van de Vooren, G. Polder, and G. van der Heijden. Identification of mushroom cultivars using image analysis. 1992.

- [33] G. Van Horn, O. Mac Aodha, Y. Song, Y. Cui, C. Sun, A. Shepard, H. Adam, P. Perona, and S. Belongie. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8769–8778, 2018.
- [34] A. J. van Strien, M. Boomsluitter, M. E. Noordeloos, R. J. T. Verweij, and T. W. Kuyper. Woodland ectomycorrhizal fungi benefit from large-scale reduction in nitrogen deposition in the netherlands. *Journal of Applied Ecology*, 55(1):290–298, 2018.
- [35] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- [36] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010.
- [37] J. Wu, C. Leng, Y. Wang, Q. Hu, and J. Cheng. Quantized convolutional neural networks for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4820–4828, 2016.
- [38] B. Zielinski, A. Sroka-Oleksiak, D. Rymarczyk, A. Piekarczyk, and M. Brzywczy-Wloch. Deep learning approach to description and classification of fungi microscopic images. *CoRR*, abs/1906.09449, 2019.