

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

ISPRS Journal of Photogrammetry and Remote Sensing

journal homepage: www.elsevier.com/locate/isprsjprs

Digital taxonomist: Identifying plant species in community scientists' photographs

Riccardo de Lutio^{a,*}, Yihang She^a, Stefano D'Aronco^a, Stefania Russo^a, Philipp Brun^b, Jan D. Wegner^{a,c}, Konrad Schindler^a

^a *EcoVision Lab, Photogrammetry and Remote Sensing, ETH Zürich, Switzerland*

^b *Land Change Science, Dynamic Macroecology, WSL, Switzerland*

^c *Institute for Computational Science, University of Zurich, Switzerland*

ARTICLE INFO

Keywords:

Species recognition
Community science
Hierarchical classification
Multimodal learning

ABSTRACT

Automatic identification of plant specimens from amateur photographs could improve species range maps, thus supporting ecosystems research as well as conservation efforts. However, classifying plant specimens based on image data alone is challenging: some species exhibit large variations in visual appearance, while at the same time different species are often visually similar; additionally, species observations follow a highly imbalanced, long-tailed distribution due to differences in abundance as well as observer biases. On the other hand, most species observations are accompanied by side information about the spatial, temporal and ecological context. Moreover, biological species are not an unordered list of classes but embedded in a hierarchical taxonomic structure. We propose a multimodal deep learning model that takes into account these additional cues in a unified framework. Our *Digital Taxonomist* is able to identify plant species in photographs better than a classifier trained on the image content alone, the performance gained is over 6 percent points in terms of accuracy.

1. Introduction

Biodiversity describes the diversity of life in terms of species' numbers, similarity, abundance, and distribution across spatial scales (Barrotta and Gronda, 2020; Gaston and Spicer, 2004). Biodiversity is essential to human well-being but rapidly deteriorating worldwide in response to anthropogenic pressure (Díaz et al., 2019). To effectively conserve biodiversity, its spatio-temporal distribution needs to be well understood, which requires efficient monitoring schemes. Scientific surveys conducted at regional or country scales are, however, costly in terms of time and financial resources, as highly skilled professionals need to repeatedly examine extensive geographical areas and carefully document the encountered species.

One viable way to complement professional biodiversity monitoring is the community science approach. The community science paradigm aims at involving the general public in scientific observations and investigations, and is particularly useful in cases where the *experiment* is characterized by a large spatial and/or temporal scale (Silvertown, 2009). The community science approach has a long history in

biodiversity monitoring (Dickinson et al., 2010). For example, volunteers have participated in the annual Christmas Bird Counts of the National Audubon Society in the USA since 1900 (Butcher and Niven, 2007).

With the rise of smartphones and other portable electronic devices, community science in biodiversity monitoring has grown. Over the past decade, a multitude of smartphone apps have been released, allowing community scientists to conveniently report observations of plants and animals, and to upload images to online databases. Among the most popular of these apps is the iNaturalist (iNaturalist, 2021) initiative, with over 3 million users and more than 36 million valid observations¹ distributed across the globe.

Although data gathered with community science is extremely valuable, it poses a number of challenges that need to be solved before it can be exploited effectively. One major issue is data quality, i.e., it is generally difficult to ensure that the collected data is correct and consistent. The main reasons are that community science data (either in the form of images or simple species presence observation) (*i*) are collected by non-experts with varying training, expertise and skills, for

* Corresponding author.

E-mail address: riccardo.delutio@geod.baug.ethz.ch (R. de Lutio).

¹ A valid observation is an observation that has a date, a location, media evidence (image or sound), and has not been voted captive/cultivated.

<https://doi.org/10.1016/j.isprsjprs.2021.10.002>

Received 26 May 2021; Received in revised form 4 October 2021; Accepted 4 October 2021

Available online 25 October 2021

0924-2716/© 2021 The Author(s). Published by Elsevier B.V. on behalf of International Society for Photogrammetry and Remote Sensing, Inc. (ISPRS). This is an

open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

instance, community scientists will on average not be able to name rare species as well as specialists; (ii) often exhibit significant biases due to geographical variations in sampling effort, observation methods and traditions, as well as regional differences in infrastructure and accessibility.

In the context of biodiversity and species distribution mapping, Machine Learning (ML) can provide several tools for mitigating at least some of these limitations. For instance, the species recognition for the data collected on the field can be automatized to some extent to help the community scientist. This can either be done on-device to assist the user during data collection, as well as in a second step to assist the experts in verifying the user-supplied labels. In recent years, computer vision has made great progress, mostly due to the rise of statistical ML. In fact, the application that spearheaded this development was the classification of image content into human-defined (semantic) categories (Deng et al., 2009). It is thus natural to ask whether ML can also assist community scientists to classify their photographs into taxonomic species, helping them to correctly identify what they have observed; thus paving the way towards more accurate and larger-scale species distribution maps. Visual species recognition has been studied fairly extensively in recent years, with different image sources ranging from carefully collected zoological or botanical collections to uncontrolled outdoor and camera trap data (Khosla et al., 2011; Welinder et al., 2010; Beery et al., 2020). In this paper, we specifically focus on the case of recognising plant species in data collected via community science applications such as iNaturalist (iNaturalist, 2021) or Info Flora (Info Flora, 2021). Properties that distinguish this specific scenario from other image classification tasks include: (i) Species observation numbers show an imbalanced distribution, as some species are naturally rare or harder to find and document than others (and perhaps also less attractive to photograph), such that they are rarely observed and only a few samples are available to train an ML model; (ii) Side information is often readily available, e.g., the location and time when the image was taken are usually known, and in turn can be linked to further information like terrain maps, satellite images, etc.; (iii) Biological species are related to each other in a hierarchical manner, i.e., through a taxonomic tree,² and one can leverage these relations during both training and inference. In particular, one may assume that, at any level of the hierarchy, species in the same group are, on average, more similar than species in distinct groups (see Fig. A1).

In this study, we develop an ML model for classifying community science photographs. Our focus is on how to best *exploit side information* that comes with the actual photograph, to improve species recognition. By side information, we mean the locations and time points of the observations, as well as associated environmental variables and optical satellite imagery. Location and time are usually uploaded together with the images.³ Our model is inspired by other works such as (Chu et al., 2019; Mac Aodha et al., 2019), however, there are a few key differences: (i) we make use of additional metadata (altitude and Sentinel-2), (ii) we train the model following a late fusion strategy and (iii) we make use of the marginalisation loss (Kumar and Zheng, 2017).

Many environmental variables are publicly available, as are remote sensing images, e.g., the Sentinel-2 satellite data repositories (Copernicus open access hub, 2021). Moreover, we include the taxonomic hierarchy to improve model performance at inference time. Hierarchically structured class labels can be beneficial in two different ways: on the one hand, the hierarchy can be used as a regularisation of the model, which has been shown to improve the classification of rare classes (Turkoglu et al., 2021); on the other hand, the hierarchy can also be used at inference time to provide a prediction (at a coarser level) for species not

present in the list of the output classes. We investigate different strategies to exploit the side information and empirically compare them. We find that a model combining the community science images, spatio-temporal context, hierarchical labels and remote sensing images trained in a joint manner with a late fusion strategy performs the best. We validate the proposed method on a subset of the iNaturalist catalogue, with 56,608 observations of 977 distinct plant species, which includes observations of plant species across the territory of Switzerland.

2. Related work

2.1. Context-based modelling

Research has shown that the location context is important for modeling the distribution of species, and therefore can especially benefit fine-grained classification tasks. In (Wittich et al., 2018) the authors adopt a nearest neighbour approach to predict the possible species that a person could encounter at certain locations given the previously recorded nearby observations. Although the paper acknowledges the fact that such information can be used to help and speed up species recognition, they do not combine their method with any image-based classification model. In (Berg et al., 2014) the location and time where a photo was taken are used to define a prior distribution over bird species occurrences. An adaptive kernel density estimation is employed to construct that distribution, which is then combined with probabilistic output from a Support Vector Machine (SVM). Although the proposed method is effective when using spatial and temporal metadata to improve classification, the usage of SVM severely limited the overall performance. Novel, deep learning-based methods can achieve higher accuracies on the same dataset without spatio-temporal priors (Foret et al., 2021). With the fast advancement of deep learning, researchers have developed ways to utilise the location context with Convolutional Neural Networks (CNNs). In (Tang et al., 2015) the authors investigate how to encode the image's GPS coordinate to increase prediction accuracy. The encoding is then concatenated with the image representation from the CNN before the final (linear) classifier. The paper also investigates the impact of further map features, e.g., precipitation maps, alongside simple GPS coordinates. (Chu et al., 2019 and Mac Aodha et al., 2019) are two studies that combine deep learning and geographical information to improve species recognition accuracy. In (Chu et al., 2019) the authors propose a refinement network that merges the prediction from a CNN with a secondary network that receives as input the location where the image was taken. The weights of the CNN network are kept frozen while training the refinement module. As a second option, the paper proposes a method where the location-aware network can alter the feature extraction inside the CNN, based on the picture's location. This second technique, however, did not lead to a substantial improvement. (Mac Aodha et al., 2019) propose a slightly different solution for the same problem, in this case the network responsible for extracting the geographical prior is in fact trained separately. The problem in this case is that the dataset consists exclusively of positive labels, i.e., it contains no information where the context speaks *against* a certain species label. To overcome this, the authors propose a joint embedding loss able to deal with presence-only datasets. The difference between the two approaches is that in the former (Chu et al., 2019) the geographical network is trained to improve the image-based prediction coming from the CNN, but cannot make a meaningful prediction on its own, i.e., without the CNN; whereas in the latter work (Mac Aodha et al., 2019) the geographical network is trained separately and can also be evaluated without an image, effectively producing a species distribution map.

2.2. Hierarchical labels

Complementary to location context, structure among the species labels helps the classification task by sharing features among related (i.e., nearby) classes. In (Srivastava and Salakhutdinov, 2013), the output

² Namely, a sub-tree of the general hierarchy of (from top to bottom) kingdom, phylum, class, order, family, genus and species (Stace, 1991).

³ These parameters constitute sensitive personal information, but community scientists are usually willing to disclose them to geo-locate their observations.

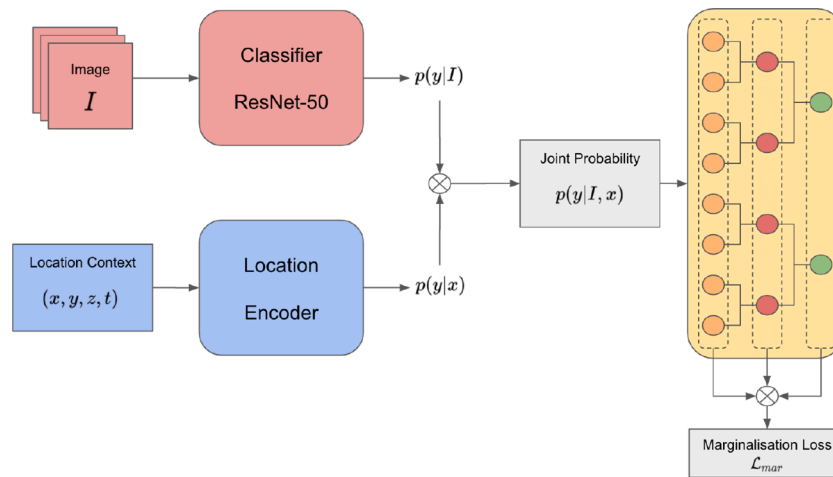


Fig. 1. Overview of our model.

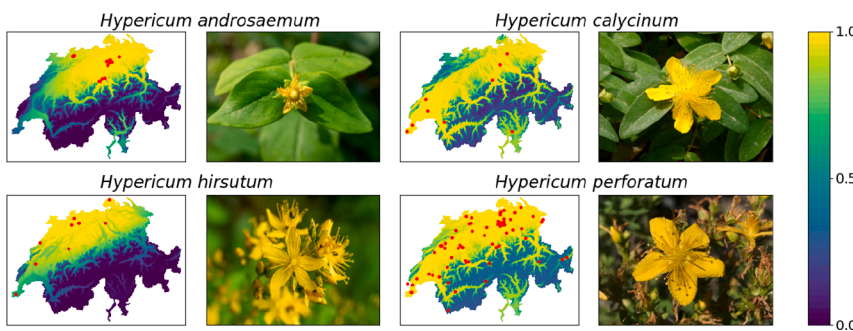


Fig. 2. Different *Hypericum* species, in order *H. androsaemum*, *H. calycinum*, *H. hirsutum* and *H. perforatum*. The present species are visually similar but have different geographical distribution ranges. For such groups of species additional spatio-temporal information can help to improve classification accuracy. For each species we visualise the probability score learned by our Location Encoder (left), the location of the training samples (red dots) and a sample image from our training set (right). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

classes are organised in a hierarchical structure, and features are transferred between related classes to inject the a priori hierarchy into the deep neural network classifier. (Yan et al., 2015) was another early work that tackled hierarchical classification in the context of visual recognition. The proposed method is limited to a 2-level hierarchy, and it is composed of two classifiers: a coarser one, which separates more easily distinguishable classes, and a finer one that resolves the more difficult cases. (Xiao et al., 2014), and more recently (Roy et al., 2018) analysed the use of hierarchical labels for visual recognition for the specific case of incremental learning. A hierarchical classifier for clothing recognition was proposed in (Kumar and Zheng, 2017). The model predicts a label hierarchy instead of a single label for the input object, by analyzing detection errors. The method exhibits good generalization capabilities also for novel clothing products that were not seen during training. In the past years, researchers have explored different ways to inject knowledge about hierarchical labels into neural networks. The authors of (Chen et al., 2018) propose a framework to predict the category scores at each hierarchy (tree) level in a top-down manner, with a multi-head network where each branch is responsible for a different level. Recently, (Dhall et al., 2020) have investigated and compared a number of strategies and loss functions to integrate hierarchical semantic structure into a CNN, including per-level classifiers, hierarchical *softmax*, and a marginalisation loss. The marginalisation loss summarizes the hierarchical information in a bottom-up manner and, although being one of the simplest approaches, emerged as one of the most effective. In (Turkoglu et al., 2021) the authors investigate the task of classifying agricultural crops from a sequence of satellite images, where the crop labels also exhibit a hierarchical structure (e.g., wheat is more similar to other cereals than to, say, orchards). They propose a convolutional recurrent architecture, where increasing depth in the spatial/convolutional dimension corresponds to a finer hierarchy level,

thus deriving higher-level features for finer classification from coarser lower-level features. The layout is specific to the recurrent setup and it is unclear how to adapt it to conventional CNNs without disrupting the feature extraction backbone.

As a general comment, we note that methods designed for hierarchical labels tend to use custom architectures and cannot easily be combined with well-known, pre-trained high-performance backbones.

3. Methodology

We now outline our proposed model for plant species classification. The model can be understood as composed of two branches: the first branch infers a probability distribution over plant species, by looking exclusively at the input image; the second branch infers another species distribution only from the auxiliary information, which is then combined with the image-based prediction to obtain a refined posterior distribution. The entire two-branch network is supervised jointly with a hierarchical loss that leverages the structure of the taxonomy.

3.1. Inference from image

Given an image I that depicts a certain plant specimen, we can use a CNN to infer its species y . The network outputs a probability distribution $p(y|I; \theta)$ over all C possible species, where θ are the learnable parameters (convolution weights). To lighten the notation we drop θ when it is clear from the context, and simply write $p(y|I)$. In our implementation we use the popular ResNet architecture (He et al., 2016), although other networks could also be employed. Our ResNet is pre-trained on ImageNet (Deng et al., 2009), a setting that has become common practice to speed up training and boost performance with limited data.

3.2. Inference from spatio-temporal context

As explained above, community science observations are often accompanied by auxiliary information, in particular spatio-temporal context, i.e., where and when the photo was taken. We denote that spatio-temporal context by the vector \mathbf{x} . The spatial information includes longitude (x), latitude (y), and altitude (z), while the day of the year t represents the temporal information.⁴ This information is typically included in the images' metadata, except for the altitude, which can be easily derived from the location given a Digital Elevation Model (DEM). The spatio-temporal context of an observation has been shown to be a useful cue for classifying species observations (see Section 2.1 and Fig. 2) – which is not surprising, as the probability of observing a certain species varies greatly across space and time.

Several methods have been proposed to merge such auxiliary information into the classification, for instance see (Chu et al., 2019; Mac Aodha et al., 2019). We will briefly describe the different strategies and highlight their pros and cons:

Early Fusion In this case the image I and auxiliary information \mathbf{x} are together fed into a model which shall directly predict $p(y|I, \mathbf{x}; \theta, \phi)$. That model is trained by minimizing a suitable loss function such as the cross-entropy between the predicted and true labels. The advantage of such an approach is that it does not impose any independence assumptions and the model can, in principle, leverage any statistical relation between y and the inputs, including correlations between I and \mathbf{x} . However, this generality comes at a price: (i) at inference time the complete auxiliary information \mathbf{x} must be fed to the model to obtain a reliable prediction, and (ii) if the training data is scarce, processing the two sources I and \mathbf{x} together increases the risk of over-fitting to spurious correlations.

Separate Training This approach, exemplified by (Mac Aodha et al., 2019), takes the opposite route and employs two completely separate networks: one “main” network processes only the image to obtain $p(y|I; \theta)$, the second “auxiliary” one processes only the side information to obtain $p(y|\mathbf{x}; \phi)$. The two networks are trained separately and produce separate scores that are only merged at inference time. This corresponds to the assumption that I and \mathbf{x} are independent, such that $p(y|I, \mathbf{x}) \propto p(y|I) \cdot p(y|\mathbf{x})$. The main advantage of this approach is a much reduced danger of over-fitting, as visual information and context are decoupled. A further advantage is that one can use additional datasets without images to train the spatio-temporal prior. On the other hand, training that prior without supporting image information can also be difficult, particularly in the common situation with presence-only annotations (Mac Aodha et al., 2019). Finally, any real correlations between \mathbf{x} and I will be lost, by construction.⁵

Late Fusion This approach, employed for instance as one of the methods in (Chu et al., 2019), constitutes a compromise between early fusion and the separate training. Separate branches are maintained for I and

\mathbf{x} . But their scores are not only combined during inference but also during training, with a joint loss function on the combined prediction $p(y|I, \mathbf{x})$. The risk of over-fitting remains low compared to early fusion, as the model admits correlations between visual and auxiliary cues only “globally”, but not between individual variables: $p(y|\mathbf{x})$ acts as a spatio-temporally varying rescaling of the image-based class scores $p(y|I)$, and vice versa. At the same time, presence-only observations do not challenge the training of the spatio-temporal prior, as the loss is computed only after including the visual information.

All the aforementioned methods are legitimate design choices, whether to prefer one or the other depends on the particular problem as well as the available data. In the experiment section, we empirically compare their performance for plant species classification. In terms of network architecture, for separate training and late fusion, the auxiliary information is first embedded into a C -dimensional vector with a fully-connected network (FCN_{context}), with C the number of classes (see Fig. 1). The FCN_{context} , with parameters ϕ , has as last layer a sigmoid, such that its output represents a presence/absence probability *per class*. Note that the sigmoid (rather than a softmax over C classes) is chosen to reflect that, at a given place and time, multiple species can be present with high probability.

3.3. Inference using auxiliary Sentinel-2 images

Finally, given that we know the location where a specific species observation was made, we can extract additional context information from remotely-sensed sources, to potentially improve species identification performance. To illustrate this, we add a Sentinel-2 image of the region around \mathbf{x} as further auxiliary data. Sentinel-2 was chosen for its potential to supplement meaningful information about the local ecosystem: it provides complete coverage of the region of interest (Switzerland). We choose to only use the 4 bands with the highest spatial resolution (10 m GSD) across the visible and infrared spectrum (ranging from 0.5 to 1.0 μm). These are commonly used to derive vegetation information and have been shown to be sufficient to derive further vegetation parameters (Lang et al., 2019).

The satellite data S is fed into the model in a similar fashion as the location context. The only difference is that the embedding of the raw data into the C -dimensional vector $p(y|S; \psi)$ is a convolutional encoder with parameters ψ (rather than a fully-connected network), to account for the nature of image data. In our implementation we use a ResNet-50.

As before, the embedded satellite imagery is combined with the other inputs according to the late fusion strategy and all three branches are trained jointly, via the merged score $p(y|I, \mathbf{x}, S)$.

3.4. Integration of taxonomic hierarchy

Hierarchical labels derived from plant taxonomy are another source of non-visual a priori information about plant species. The taxonomic hierarchy endows the output space with additional structure that may help to correctly classify plant species, especially if the training data is heavily imbalanced. Attempts to use the hierarchy rest on the assumption that closely related species in the tree have higher visual similarity than more distant ones.⁶ On the one hand, the hierarchical grouping (for instance, of many rare species into a common genus) gives rare species statistical strength, as confusing them with each other becomes cheaper than confusing them with some frequently observed species from a different genus. On the other hand, the grouping also benefits the

⁴ Thus assuming the distribution is seasonally varying but stationary over a few years.

⁵ Such patterns are likely to exist. Examples include location-specific shadows or time-dependent snow cover.

⁶ In expectation, not necessarily in every instance




Fig. 3. The idea of the marginalisation loss is to simultaneously apply a cross-entropy loss at all levels of the taxonomic hierarchy. As the output of the classifier is the probability distribution over all species, marginalising over all species within each genus yields the probability distribution over genera. This procedure can then be repeated to derive the distributions at all higher levels. The marginalisation loss is simply the sum of the intermediate losses computed at each level.

fine-grained species classification, as it favours feature sharing between adjacent classes that, by themselves, have too few samples to learn a good representation (Srivastava and Salakhutdinov, 2013). The taxonomic levels we use are, from the bottom to the top of the hierarchy: species, genus, family, order, class and phylum.

To integrate hierarchical labels, we adopt the marginalisation loss proposed in (Kumar and Zheng, 2017). As shown in Fig. 3, the output of the classifier is the probability distribution over all species. Marginalising over all species within each genus thus yields the probability distribution over genera. This procedure can then be repeated to derive the distribution over families, etc.:

$$p(y_i^l) = \sum_{j \in K_i} p(y_j^{l+1}) \quad (1)$$

where $p(y_i^l)$ is the predicted probability for the i -th label at hierarchy level l , and $p(y_j^{l+1})$ is the probability of class j at the next-coarser hierarchy level $l + 1$. With K_i we denote the set of child classes of parent class i . Based on the distribution $p(y^l)$ derived at level l , we can compute a cross-entropy loss \mathcal{L}^l for each individual level. The marginalisation loss is then simply the sum of all these intermediate losses:

$$\mathcal{L}_{mar} = \sum_l \mathcal{L}^l. \quad (2)$$




Fig. 4. Sample distribution for each species in the training dataset. Note the logarithmic scale of the y-axis. Both diagrams share the same scale.

3.5. Data preprocessing

All community science images were resized to the size of 256×256 and then centre-cropped to 224×224 . The images used for training were additionally augmented by random rotations, random horizontal flips and color-jitter, which are all standard methods to help mitigate the risk of over-fitting. Furthermore, all images were normalized according to the mean and standard deviation of the training set.

We encode the observation time, measured as day of year t , into (t_1, t_2) using the sine-cosine mapping (Mac Aodha et al., 2019), Eq. 3. In this way December 31st and January 1st are mapped close to each other, correctly accounting for the cyclic nature of the variable.

$$\begin{cases} t_1 = \sin\left(\frac{2\pi t}{365}\right) \\ t_2 = \cos\left(\frac{2\pi t}{365}\right) \end{cases} \quad (3)$$

Regarding the location coordinates, we rescale longitude, latitude and altitude separately to fit into the interval $[-1, 1]$ and denote the triple of normalised coordinates as our geo-location (x, y, z) .

Finally, the Sentinel-2 images are extracted from a cloud-free mosaic of images taken in 2020. As previously indicated, we only use the four spectral bands with a 10 m spatial resolution (R, G, B and N-IR), since they are often sufficient to derive vegetation parameters (Lang et al., 2019). From this mosaic, we extract patches of 256×256 pixels to ensure enough context (ca. 1.3 km around the sample location, see Table A.2 for a comparison of the performance with different sized patches).

3.6. Balanced sampling

We used a balanced sampling strategy, where the sampling weight of each image W_i is inversely proportional to the number of images N_{y_i} of the corresponding class y_i :

$$W_i = \frac{1}{N_{y_i}}. \quad (4)$$

This strategy will oversample the rare species from the tail of the distribution and undersample the frequent species from the head of the distribution, so as to mitigate the impact of the imbalance on the classifier. It should be noted that the effect cannot be completely removed: even when sampled with higher frequency, the few images of a rarely observed species will inevitably carry less information than the many example images of an abundant species. As a result there is no clear advantage in neither of the two approaches, making this a mere design choice. In fact, using the balanced sampling strategy, compared to the conventional training method, improves the per-class accuracy while decreasing the overall accuracy (see Table A.1). Although the differences in performance are small, we decide to prioritise the per-class accuracy since we believe it is more important for our application and

Table 1
Overview of our dataset.

Description	Images	Species
Overall	60,781	2,374
Selected	56,608	977
Unseen	1,650	330




Fig. 5. Sample distribution for all “Research Grade” iNaturalist observations in Switzerland. Note the logarithmic scale of the y-axis.

thus choose to use the balanced sampling approach.

Fig. 4 shows the number of training images per species in the training set before and after balanced sampling.

We employ a stochastic gradient descent (SGD) (Bottou et al., 2012) to optimise the parameters of our model. We set two different learning rates, a smaller one of $5 \cdot 10^{-5}$ for the pre-trained convolutional layers of the CNN, and a larger one of $2 \cdot 10^{-3}$ for the fully connected layers. These learning rates are further reduced when a plateau is reached. The batch size is fixed to 32, and all models were trained for 100 epochs. We use a cross-entropy loss for baselines that ignore the label hierarchy, and the marginalisation loss (Eq. 2) for hierarchically structured labels.

All results (unless stated otherwise) are computed with 5-fold cross-validation, stratified to ensure uniform class distribution across all folds.

4. Dataset

From the iNaturalist database (iNaturalist, 2021), we have downloaded all images⁷ of plants that are located in Switzerland and labeled as “Research Grade”. The latter constitutes the highest level of data quality, where observations meet five criteria: (1) they must include a

Table 3
Comparison of different training strategies. Note that the top-k accuracies indicate the average species-specific metrics.

Model	Accuracy (%)	Top-1 (%)	Top-3 (%)	Top-5 (%)
Separate Training	76.41	65.29	82.09	86.58
Joint Training: Early Fusion	73.65	64.84	79.60	83.87
Joint Training: Late Fusion	79.12	69.76	84.86	88.95

date, (2) a spatial geo-reference, (3) a picture (or sound, but we only focus on images in this work), (4) the subject must be a naturally living organism (not captive or cultivated), and (5) at least 2 identifiers should agree on a taxon, out of a minimum of 3 identifiers.

As shown in Table 1, a total of 60,781 images were downloaded (see Fig. 6), which represented 2,374 species. However, as seen in Fig. 5, the dataset is highly imbalanced and follows a long-tail distribution. We discard all species with < 10 images in order to ensure reliability and statistical significance of the experimental results. After this filtering we are left with 56,608 images representing 977 species. We also generated a dataset of unseen species for further experiments (see Section 5.4). These are observations of species that have fewer than 10 but more than 5 images. For each of those species, we select 5 images at random.

Besides the images, the dataset also contains non-visual information, including the additional data that we use in our model, i.e., longitude, latitude, day of the year and hierarchical labels. To obtain altitude we extract the height value corresponding to the given geo-location from the swissALTI3D DEM of the Swiss national mapping agency (Swisstopo, 2021).

5. Experimental results

5.1. Model performance

We have conducted experiments with the following models to empirically determine their performance gain: (1) **Baseline**, which corresponds to a standard ResNet50; (2) **Baseline + Location Context**, where we add the location encoder to the baseline in a late fusion setup; (3) **Baseline + Hierarchical Labels**, where we add the marginalisation loss to the baseline; and (4) **Proposed Model**, which leverages both the location context and the hierarchical labels. Here, we again use the late fusion strategy, which empirically achieved the best performance (see Table 3).



Fig. 6. Example images from our dataset.

⁷ As of November 5th, 2020.

Table 2

Ablation study of our proposed model. Note that the top-k accuracies denote the average species-specific metrics in order to give the same weight to all the species in the evaluation.

Model	Accuracy (%)	Top-1 (%)	Top-3 (%)	Top-5 (%)
Baseline	73.48	62.48	79.04	83.97
Baseline + Location Context	76.99	67.47	82.50	87.01
Baseline + Hierarchical Labels	76.30	65.49	82.20	86.81
Proposed Model	79.12	69.76	84.86	88.95

As seen in Table 2, adding either the location context or the hierarchical labels to the baseline model significantly improves the results, for all metrics. Note that we compute the top-k accuracies as the average species-specific metrics in order to give the same weight to all the species in the evaluation. Thus the overall accuracy is higher than the top-1 hit-rate due to the imbalanced nature of our dataset, which is preserved in our stratified cross-validation. Furthermore, improvements from location context and hierarchical labels are largely orthogonal, as expected, since they leverage different types of information. These results indicate a clear benefit of complementing visual cues from community science images with additional sources of information. For a more detailed ablation study of the exact contributions of every component in our location context see Section 5.5. Visual inspection of misclassified images confirms that location context helps in the case of visually similar species that occur in different geographical regions (see Fig. 7), whereas hierarchical labels help to classify species with few images.

Finally, as seen in Fig. 8 our proposed model improves over the baseline for all four ranges of species counts and the margin of improvement is largest for the tail species of the dataset with a number of images between 10 and 50. This is very useful since rare species are more commonly misidentified by community science and are particularly important for conservation purposes.

5.2. Training strategies

Table 3 compares the three different training strategies described in Section 3.2. The separate training strategy has the advantage that one can use the image classifier and get reasonable predictions even when metadata is missing. Whereas the joint training strategies should always

perform better, at the cost of being less flexible, as metadata is mandatory. Under ideal circumstances, one would also expect the early fusion strategy to perform best, as it is not subject to any factorisation constraints on $p(y|I, x)$ and can leverage the complete correlation structure. In practice, we however observe the worst performance, see Table 3. It appears that the increased model capacity leads to overfitting. The late fusion training strategy, with its restricted interaction between image and context cues, emerges as the best compromise with clearly superior performance. Separate training does bring a noticeable improvement over the baseline but does not reach the late fusion approach. Likely this is, at least in part, due to the presence-only labels hampering the learning of the prior $p(y|x)$.

5.3. Evaluation at different hierarchy levels

When using the taxonomic hierarchy during training in conjunction with the marginalisation loss, we can predict at inference time labels at different hierarchy levels. If taxonomic distance indeed correlates with similar visual features and ecological requirements (see Fig. A1), then the predictions at higher levels should be increasingly more correct. I.e., even if a specimen is assigned the wrong species label it might be assigned the correct genus label, as it is more likely to be confused with a similar species from the same genus.⁸

We have evaluated our model at all taxonomic levels that we use, see

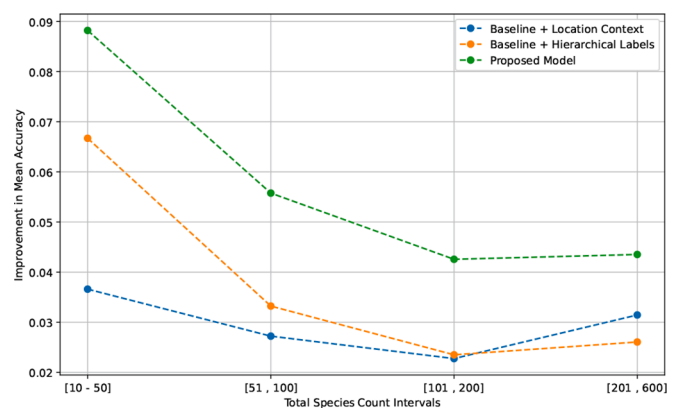


Fig. 8. Improvement in Mean Accuracy over the baseline for species with different numbers of images in the dataset.

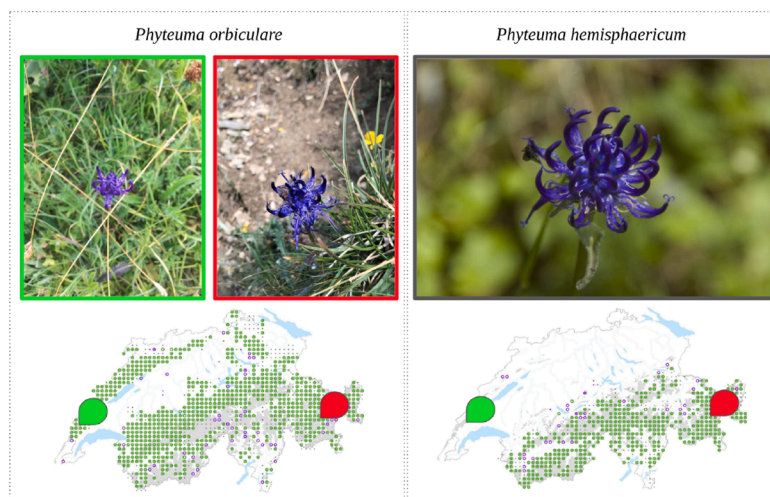


Fig. 7. Misclassification example: both images of *Phyteuma orbiculare* (Left) are misclassified as *Phyteuma hemisphaericum* (Right) by our baseline model. When including the location context, our proposed model correctly classifies the image with the green frame, whereas the image with the red frame is still misclassified. The green and red arrows indicate the locations of the respective left two images. The underlying maps are the species distribution maps downloaded from Info Flora (Info Flora, 2021). This highlights the importance of including the location information to distinguish visually similar species that have different geographical ranges. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

⁸ Note that also the chance level increases, as there are fewer possible labels.

Table 5

Results at different hierarchical levels. Note that top-3 and top-5 accuracy at phylum level are meaningless, since there are only 3 possible phyla. Note that the top-k accuracies indicate the average species-specific metrics.

Metric (%)	Species	Genus	Family	Order	Class	Phylum
Accuracy	79.02	83.39	87.26	88.54	97.24	99.89
Top-1	69.50	73.23	75.84	78.53	88.52	86.63
Top-3	84.49	85.76	87.81	89.34	98.01	100.0
Top-5	88.57	89.37	91.45	92.95	99.51	100.0

Table 4

Number of classes at each hierarchical level.

Level	Species	Genus	Family	Order	Class	Phylum
Number	977	489	121	50	8	3

Table 6

Accuracy (%) on unseen species.

Evaluation Set	Species	Genus	Family	Order	Class	Phylum
5-fold Cross-Val	79.02	83.39	87.26	88.54	97.24	99.89
Unseen Species	–	24.27	41.86	50.23	85.60	96.00

Table 5. Indeed, the performance is better for the higher levels (c.f. **Table 4**). Furthermore, higher up in the hierarchy, fewer classes are poorly represented; the long-tail distribution is less extreme.

5.4. Experiments with unseen species

Given the hierarchical labels, it is also possible to classify new species which the classifier has not seen at all during training. While the assigned species label will necessarily always be wrong, one would hope that the predictions at coarser taxonomy levels are often sensible. For this experiment, we picked 330 species that were initially discarded from our dataset for having <10 images, but for which at least 5 images are available, c.f. the “Unseen” row in **Table 1**. The corresponding results in **Table 6** confirm our intuition: while there is of course a significant performance drop compared to the trained species, it is still possible to classify unseen species into the right Genus, Family or Order with reasonable performance, well above chance level (the probability of success of a classifier that always predicts the most common class). This capability can be extremely useful in the context of community science, where the coarser labels can be used to refer examples to the right expert for classification or to detect gaps in the taxonomy lists offered to users.

5.5. Contextual information and Sentinel-2 images ablation study

To investigate the contributions of different types of contextual information, and the potential benefit of adding satellite imagery, we perform extensive ablation studies.

In **Table 7** we show the impact of the different contextual information (altitude, geo-coordinates, day of the year) on the evaluated metrics. As it can be seen they all contribute to some extent, with the altitude being one of the most important. Considering the high altitude variability of the Swiss landscape it was rather expected that the altitude could carry the most valuable information. When the full context is combined the performance metrics show a further increase meaning that the additional data carry orthogonal information.

Finally, **Table 8** displays the performance achieved with the integration of Sentinel-2 imagery. Overall, their impact turns out to be small. When naively adding the Sentinel-2 branch, performance even drops slightly, apparently due to over-fitting. By adding standard dropout regularisation (Srivastava et al., 2014) on the last fully-connected layer, we were able to remedy this behaviour and achieve a mild (but still statistically significant) performance gain. To ensure that the

Table 7

Ablation study of spatio-temporal context. Note that the top-k accuracies indicate the average species-specific metrics.

Model	Accuracy (%)	Top-1 (%)	Top-3 (%)	Top-5 (%)
Baseline	73.48	62.48	79.04	83.97
Baseline + Altitude	75.40	65.11	81.00	85.80
Baseline + Geo-coordinates	75.07	64.86	80.63	85.45
Baseline + Day of the year	75.51	65.00	80.91	85.51
Baseline + Full Location	76.99	67.47	82.50	87.01
Context				

Table 8

Results adding Sentinel-2 mosaic. Note that the top-k accuracies indicate the average species-specific metrics.

Model	Accuracy (%)	Top-1 (%)	Top-3 (%)	Top-5 (%)
Proposed model	79.12	69.76	84.86	88.95
Proposed model with Dropout	78.02	67.77	83.39	87.52
Proposed model + Sen-2	78.59	68.29	84.43	88.60
Proposed model + Sen-2 with Dropout	79.73	70.32	85.52	89.33

difference is actually caused by the satellite imagery and not the dropout, we add an additional baseline where the model without the Sentinel-2 branch is trained with drop-out. Interestingly, this even degraded the performance.

While it is promising that the much-enriched context information from the satellite image brings an improvement over the simple geo-location, that gain is relatively modest, at least with our implementation. Further research, beyond the scope of the present paper, will be needed to clarify the potential of satellite (or airborne) data as auxiliary information.

6. Conclusion

In this work, we have demonstrated that easily accessible side information can bring rather large performance gains when classifying community science photographs. We have focused on the spatio-temporal context of the observations, and have shown how it can refine the classification model by providing relevant prior knowledge regarding the distribution and occurrence of species observations. We have also briefly touched on extended radiometric context from optical satellite imagery, a direction where we see quite some potential for further research. Moreover, we have verified that exploiting the hierarchical structure of biological taxonomy not only improves the species recognition performance, but also enables more reliable predictions at coarser taxonomy levels, and even coarse classification of species not seen at all during the classifier training.

In terms of practical community science applications, our model is also a step towards a viable scheme for verifying user-supplied labels. For instance, the proposed method could provide hints to the community scientist when labelling the species, or it could facilitate the reviewing validation by experts, marking specific observations where the model disagrees with the label provided by the community scientist. Of course, these suggestions would need to be followed with care in practice to avoid creating a confirmation bias of the model. We hope that, ultimately, a larger number of correct species observations will contribute to better species distribution models, to inform biodiversity research and conservation initiatives, particularly for rare species.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Ablation studies

Fig. A1.




Fig. A1. Confusion matrix for our proposed model. The species in the rows and columns have been ordered based on their taxonomy. The hierarchy between the species is made apparent by the dendrograms. The block-like structure along the diagonal indicates that species that are close in terms of their taxonomy are misclassified for each other more often than unrelated species.

Appendix B. Confusion matrix

Tables A.1 and A.2.

Table A.1

Ablation study of balanced sampling. Note that the top-k accuracies denote the average species-specific metrics.

Model	Accuracy (%)	Top-1 (%)	Top-3 (%)	Top-5 (%)
Baseline + No Balanced Sampling	75.57	62.08	80.9	86.18
Baseline + Balanced Sampling	73.48	62.48	79.04	83.97
Proposed Model + No Balanced Sampling	80.05	69.23	86.15	90.29
Proposed Model + Balanced Sampling	79.12	69.76	84.86	88.95

Table A.2

Comparison of different extents for Sentinel-2 images. Note that the top-k accuracies indicate the average species-specific metrics.

Model	Accuracy (%)	Top-1 (%)	Top-3 (%)	Top-5 (%)
No Sent-2	79.12	69.76	84.86	88.95
Small Sent-2 (128 × 128)	79.5	69.56	84.76	88.92
Normal Sent-2 (256 × 256)	79.73	70.32	85.52	89.33
Large Sent-2 (512 × 512)	79.16	69.9	84.94	88.91

References

- Barrotta, P., Gronda, R., 2020. Controversies and Interdisciplinarity: Beyond disciplinary fragmentation for a new knowledge model. In: Ch. What is the Meaning of Biodiversity?, 16. John Benjamins Publishing Company, pp. 115–131.
- Beery, S., Cole, E., Gjoka, A., 2020. The iWildCam 2020 competition dataset. In: Proceedings, IEEE Conference on Computer Vision and Pattern Recognition Workshops.
- Berg, T., Liu, J., Woo Lee, S., Alexander, M.L., Jacobs, D.W., Belhumeur, P.N., 2014. Birdsnap: Large-scale fine-grained visual categorization of birds. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2011–2018.
- Bottou, L., 2012. Neural Networks: Tricks of the Trade, second ed. Berlin Heidelberg: Springer, pp. 421–436 (Ch. Stochastic Gradient Descent Tricks).
- Butcher, G., Niven, D., 2007. Combining data from the Christmas Bird Count and the Breeding Bird Survey to determine the continental status and trends of North America birds. Tech. rep. National Audubon Society.
- Chen, T., Wu, W., Gao, Y., Dong, L., Luo, X., Lin, L., 2018. Fine-grained representation learning and recognition by exploiting hierarchical semantic embedding. In: Proceedings of the ACM International Conference on Multimedia, pp. 2023–2031.
- Chu, G., Potetz, B., Wang, W., Howard, A., Song, Y., Brucher, F., Leung, T., Adam, H., 2019. Geo-aware networks for fine-grained recognition. In: Proceedings, IEEE International Conference on Computer Vision Workshops, pp. 247–254.
- Copernicus open access hub. <https://scihub.copernicus.eu> (last accessed on 26.05.2021).
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. ImageNet: A large-scale hierarchical image database. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 248–255.
- Dhall, A., Makarova, A., Ganea, O., Pavlo, D., Greeff, M., Krause, A., 2020. Hierarchical image classification using entailment cone embeddings. In: Proceedings, IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 836–837.
- Díaz, S., Settele, J., Brondízio, E., Ngo, H., Guèze, M., Agard, J., Armeth, A., Balvanera, P., Brauman, K., Butchart, S., Chan, K., Garibaldi, L., Ichii, K., Liu, J., Subramanian, S., Midgley, G., Miloslavich, P., Molnár, Z., Obura, D., Pfaff, A., Polasky, S., Purvis, A., Razaque, J., Reyers, B., Chowdhury, R., Shin, Y., Visseren-Hamakers, I., Willis, K., Zayas, C., 2019. Summary for policymakers of the global assessment report on biodiversity and ecosystem services, Tech. rep. Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services.
- Dickinson, J.L., Zuckerberg, B., Bonter, D.N., 2010. Citizen science as an ecological research tool: challenges and benefits. *Ann. Rev. Ecol. Evol. Systematics* 41, 149–172.
- Foret, P., Kleiner, A., Mobahi, H., Neyshabur, B., 2021. Sharpness-aware minimization for efficiently improving generalization. In: Proceedings of the International Conference on Learning Representations.
- Gaston, K.J., Spicer, J.I., 2004. *Biodiversity: An introduction*, second ed. Blackwell Publishing.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778.
- iNaturalist, <https://www.inaturalist.org> (last accessed on 26.05.2021).
- Info Flora. <https://www.infoflora.ch> last accessed on 26.05.2021.
- Khosla, A., Jayadevaprakash, N., Yao, B., Fei-Fei, L., 2011. Novel dataset for fine-grained image categorization. In: First Workshop on Fine-Grained Visual Categorization at the IEEE Conference on Computer Vision and Pattern Recognition.
- Kumar, S., Zheng, R., 2017. Hierarchical category detector for clothing recognition from visual data. In: Proceedings, IEEE International Conference on Computer Vision Workshops, pp. 2306–2312.
- Lang, N., Schindler, K., Wegner, J.D., 2019. Country-wide high-resolution vegetation height mapping with Sentinel-2. *Remote Sens. Environ.* 233, 111347.
- Mac Aodha, O., Cole, E., Perona, P., 2019. Presence-only geographical priors for fine-grained image classification. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 9596–9606.
- Roy, D., Panda, P., Roy, K., 2018. Tree-CNN: A hierarchical deep convolutional neural network for incremental learning. arXiv: 1802.05800.
- Silvertown, J., 2009. A new dawn for citizen science. *Trends Ecol. Evol.* 24 (9), 467–471.
- Srivastava, N., Salakhutdinov, R., 2013. Discriminative transfer learning with tree-based priors. In: Proceedings, Advances in Neural Information Processing Systems.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15 (56), 1929–1958.
- Stace, C.A., 1991. *Plant Taxonomy and Biosystematics*. Cambridge University Press.
- Swisstopo. <https://www.swisstopo.admin.ch/en/geodata/height/alti3d.html> (last accessed on 26.05.2021).
- Tang, K., Paluri, M., Fei-Fei, L., Fergus, R., Bourdev, L., 2015. Improving image classification with location context. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1008–1016.
- Turkoglu, M.O., D'Aronco, S., Perich, G., Liebis, F., Streit, C., Schindler, K., Wegner, J. D., 2021. Crop mapping from image time series: deep learning with multi-scale label hierarchies. arXiv:2102.08820.
- Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., Perona, P., 2010. Caltech-UCSD Birds 200. Tech. rep. California Institute of Technology.
- Wittich, H.C., Seeland, M., Wäldchen, J., Rzanny, M., Mäder, P., 2018. Recommending plant taxa for supporting on-site species identification. *BMC Bioinformatics* 19 (1), 1–17.
- Xiao, T., Zhang, J., Yang, K., Peng, Y., Zhang, Z., 2014. Error-driven incremental learning in deep convolutional neural network for large-scale image classification. In: Proceedings of the ACM International Conference on Multimedia, pp. 177–186.
- Yan, Z., Zhang, H., Piramuthu, R., Jagadeesh, V., DeCoste, D., Di, W., Yu, Y., 2015. HD-CNN: hierarchical deep convolutional neural networks for large scale visual recognition. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2740–2748.