# Floralens: A Deep Learning Model for the Portuguese Native Flora

António Filgueiras<sup>1</sup>, Eduardo R. B. Marques<sup>1,2\*</sup>, Luís M. B. Lopes<sup>1,2</sup>, Miguel Marques<sup>1</sup>, Hugo Silva<sup>1</sup>

 ${\bf 1}$  Department of Computer Science, Faculty of Sciences, University of Porto ${\bf 2}$  CRACS/INESC-TEC

¤Current Address: Rua do Campo Alegre 1055, 4169-007 Porto, Portugal \* ebmarques@fc.up.pt

# Abstract

Machine-learning techniques, especially deep convolutional neural networks, are pivotal for image-based identification of biological species in many Citizen Science platforms. In this paper, we describe the construction of a dataset for the Portuguese native flora based on publicly available research-grade datasets, and the derivation of a high-accuracy model from it using off-the-shelf deep convolutional neural networks. We anchored the dataset in high-quality data provided by Sociedade Portuguesa de Botânica and added further sampled data from research-grade datasets available from GBIF. We find that with a careful dataset design, off-the-shelf machine-learning cloud services such as Google's AutoML Vision produce accurate models, with results comparable to those of Pl@ntNet, a state-of-the-art citizen science platform. The best model we derived, dubbed Floralens, has been integrated into the public website of Project Biolens, where we gather models for other taxa as well. The dataset used to train the model is also publicly available on Zenodo.



Fig 1. Detail of the FloraOn web application.

# 1 Introduction

The improvements in processing speed, storage capacity, and imaging sensors for mobile devices paved the way for Citizen Science [1] applications and Web services that allow amateur enthusiasts to participate in science projects. One successful case study is nature observation, specifically the photographic recording of animals, plants, and fungi in their natural habitats. Besides storing these observations, some platforms use deep-learning models to provide automatic taxonomic identification from user-provided images [2–6]. The data gathered by such projects is valuable for scientists, from hardcore taxonomists to ecologists studying the impact of human activity on biodiversity [7,8].

In project Biolens [9], we are interested in creating lightweight identification models for web and mobile apps, possibly working offline (e.g., in the field) and having the option of not immediately sharing data (e.g., for research projects). Most citizen science platforms do not provide this level of flexibility. Moreover, a detailed description of their methodological approach is lacking in the literature (cf. Section 2), making it difficult to use this accumulated experience for new projects. We, therefore, decided to independently develop a streamlined methodology to build datasets for native Portuguese species covering different taxa and to derive accurate CNN-based models using off-the-shelf machine-learning tools. Currently, Biolens includes four models: Lepilens and Mothlens (for butterflies and moths, together covering the order Lepidoptera); Dragonlens (for dragonflies and damselflies, covering the order Odonata), and, the latest addition, Floralens (covering the Plantae kingdom).

This paper describes the derivation of Floralens, a high-accuracy machine-learning model for automatic taxonomic identification of the Portuguese native flora. The work is anchored on the FloraOn dataset provided by the Sociedade Portuguesa de Botânica [10], available online via a web application (Figure 1) and as a contributed dataset in the Global Biodiversity Information Facility (GBIF) [11]. This dataset contains relatively few images per species but those provided are of very high quality and identified by experienced taxonomists. We use this list of Portuguese native species as our reference to build the dataset.

The methodology used to derive all the Biolens models has two core traits: (a) the use of data from research-grade public repositories for dataset construction, and (b) the use of Google's AutoML Vision (GAMLV) to derive the actual models. This methodology is briefly described in a short scientific outreach article (in Portuguese) [12], and, more thoroughly, in MSc theses [13, 14] (both covering different

stages of the work on Floralens), and a BSc project report [15] (covering Lepilens). Our evaluation of Floralens shows that, with a carefully designed dataset, current off-the-shelf machine-learning cloud-based services output models whose performance rivals or exceeds that obtained with models provided by state-of-the-art citizen science projects. Besides the baseline computer vision model, we study the impact of using multiple images by specimen and the geographical distribution of species on the model's accuracy. We also compare the Floralens model to those provided by the state-of-the-art platform Pl@ntnet.

The main contributions of this paper are, thus, as follows:

- a dataset for the Portuguese Flora available from Zenodo [16];
- a high-accuracy CNN-based model for the Portuguese native flora publicly available via web and mobile applications;
- a thorough quantitative evaluation of the model and comparison with the state-of-the-art platform Pl@ntNet;
- a complete set of Python notebooks used to perform the analysis presented in this paper, also available from Zenodo.

The remainder of this paper is structured as follows. Section 2 describes the current state-of-the-art regarding automatic taxonomic identification based on deep learning. Section 3 describes the construction of the dataset used in this study, the model's generation using GAMLV, and the evaluation metrics used to gauge its accuracy. Section 4 describes the results obtained with the model and its refinements, such as using multiple images of the same specimen for classification and ecological indicators such as species' geographical distribution. A detailed performance comparison with Pl@ntNet is also presented. Section 5 describes the software artifacts and datasets produced in the scope of this work. Finally, Section 6 discusses the main findings of this study and puts forward future research goals.

#### 2 State-of-the-Art

The advent of deep learning allowed the development of the first tools for automatically identifying plant species from input images [17–19]. Their continuous refinements quickly reached a point where automatic identification rivaled those made by specialists [20], thus attesting to the transformative role of AI in this field [21]. Nowadays, models based on deep learning are central tools in major citizen science platforms such as iNaturalist [2], Observation.org [3], and Pl@ntNet [6] and are integrated into web and mobile applications. Most of these models are derived using Convolutional Neural Networks [22,23] and, recently, Vision Transformers [24–26]. Citizen science platforms generate important artifacts such as curated datasets, often made available to the public through biodiversity data portals, notably GBIF [11]. These datasets enable the development of other ML models, as is the case of Floralens that, in addition to FloraOn, uses data sampled from GBIF datasets provided by iNaturalist, Observation.org, and Pl@ntNet (cf. Section 3). However, the dataset construction and model derivation methodologies used by these citizen science platforms are not documented in detail in the literature: [2] makes reference to machine learning from an end-user perspective, some details can be found online in scattered form [27]; [6] provides a basic summary of the Pl@ntNet methodology to derive models, the information is also given in slightly more extended form in [28]; [3] is a short abstract concerning their mobile app ObsIdentify [29], and both the app and the site interface are enabled by an API [21,30] whose underlying model is not explained. Moreover, extended access to the classification services in these platforms for research is hampered by the fact that in some cases the service is not exposed by the API or, commonly, requires the payment of a significant fee or subscription for large numbers of requests.

Deriving a CNN-based model or other ML models often requires expert knowledge. non-trivial configuration aspects, writing specialized code, and a heavyweight computational infrastructure for storing data and model training. The strain is more acute when the amount of data involved is non-negligible, as in the case of Floralens where we use approximately 300,000 images to derive a model. AutoML cloud services address these needs, of which GAMLV is an example. Like Google, other major cloud providers have MLaaS offerings, e.g., Amazon Rekognition [31], Apple Create ML [32], and Azure AutoML [33]. These services allow the automation of several aspects of model derivation backed by computational infrastructures including special-purpose hardware such as GPUs or TPUs, necessary to derive a model in a reasonable amount of time. Models derived using these tools are referenced in the literature of several areas of knowledge, e.g., biology [34], medicine [35, 36], agriculture [37], engineering [38]. Comparative results between GAMLV and one or more competitor platforms are provided in some of these works [34-36], generally illustrating good performance by GAMLV-derived models. In this work, we used GAMLV since we had free access to the toolkit and computational infrastructure via Google Cloud's Research Credits program.

This is a rapidly evolving research area stoked by developments in machine learning and available computational power and infrastructures. A recent trend is using metadata associated with observations to improve the identification precision. Pl@ntNet, for example, developed regional models by dividing the globe into several biogeographic domains [39] following published regional floras such as the WCVP/Kew [40] and introduced metadata providing information on the anatomic part of the plant depicted in an image, e.g., flower, leaf, stem. A similar approach is proposed by Rzanny et al. [41], where datasets with images taken from 5 different predefined perspectives per specimen are used. The authors then derive models for each perspective and fuse their outputs by averaging them across species to improve global precision. Mäder et al. [5] (Flora-Incognita) use a CNN to produce a model that classifies images of German native plants. A model that outputs the odds of finding a species in a given spatial and time context is also produced using a deep feedforward network. The outputs of these models are then fused using a recurrent neural network. More recently, Brun et al. [42] (FlorID) use data-efficient image transformers trained with images of specimens with associated locations, timestamps, and plant parts (e.g., vegetative, flower, fruit, trunk). As in Flora-Incognita, they use location (including elevation) and time attributes (season) to produce ecological predictor models that output the species most likely to be observed in a given spatial and temporal setting. They fuse the outputs of these ecological predictors with the image-based models to improve global precision. Despite these developments, some problems remain such as the apparent difficulty that these models have in identifying endangered species [43]. Floralens is closer in philosophy to regional flora projects such as Flora-Incognita [5] and FlorID [42].

Another trend that has generated considerable interest is the generation of global models based on extreme datasets such as iNat Challenge [44] and PlantCLEF/LifeCLEF [45, 46], consisting of millions of images representing tens of thousands of species out of the world's  $\sim$ 300K.

		Raw data				Dataset		
Source	#I	%I	#S	$\geq 50~{\rm I}$	#I	%I	#S	
FloraOn	22,869	0.5	2,127	0	15,191	5.1	1,397	
iNaturalist	2,753,167	66.2	2,066	1,431	90,127	30.6	1,358	
Observation.org	823,389	19.8	1,816	1,114	85,746	29.2	1,093	
Pl@ntNet	$515,\!950$	12.4	$1,\!495$	735	$102,\!537$	34.9	$1,\!373$	
Total	4,154,895		2,539	$1,\!678$	293,601		1,678	

**Table 1.** Raw data and derived dataset after sampling (#I: image count; #S: species count;  $\geq 50$  I: species with more than 50 images).

# 3 Methods

In this section we describe the creation of the Floralens dataset, the derivation of the Floralens model, and the evaluation metrics used to assess the latter's performance.

#### 3.1 Data

Our universe of species was the FloraOn catalog as of November 2021. It covers 2,712 species of native Portuguese flora. We started the construction of the dataset by gathering a large collection of images for each species in our list. These were extracted from curated datasets available from GBIF. We then sampled this collection so that each species would be represented by at least 50 and, at most, 200 images in the dataset. The goal was to ensure that a minimum number of images was available for each species and that a species would not be over-represented as it would introduce bias in the models derived from the dataset. These bounds were defined based on our prior experience building datasets for other biological taxa (cf. Section 5).

The FloraOn repository contains geo-referenced records of Flora species with associated images. The image data is relatively broad in scope, as it covers 78% of the entire catalog (2,127 out of 2,712 species), but has a limited volume: on average there are just 11 images per species, and, unsurprisingly, our 50-image lower bound threshold is not met for a single species. Hence, to adequately populate the Floralens dataset, we retrieved the FloraOn images but also considered image data from three publicly-available datasets stored and made available at GBIF by three citizen science platforms: iNaturalist [47], Observation.org [48], and Pl@ntNet [49]. This approach is used by other platforms when building their datasets, e.g., FlorID [42].

The GBIF datasets provide validated observation data and associated images, originally submitted by users of the respective platforms. However, the validation process differs among data sources, as discussed further in this section. For each dataset and each species in our universe, we used GBIF portal queries to obtain observation records in the Darwin Core Archive format [50]. Each record corresponds to an observation of a specimen, typically made in the wild by citizen scientists or experts, accompanied by its taxonomic identification, date, geographical location, and one or more images. The GBIF portal queries were parameterized to cover continental Europe, including Great Britain, Ireland, and the Mediterranean islands. We did this as a preliminary analysis suggested that observations from Portugal or even the Iberian Peninsula would yield limited data in terms of volume and variety. Fortunately, many species of the Portuguese flora are widely distributed.

The raw data, from all image sources, is listed in Table 1 (left), along with the characterization of the Floralens dataset (right) that results from sampling the raw data. The corresponding histograms, depicting the number of species versus the number of images, are illustrated in Figure 2. In the raw data, more than 4 million images were

available for consideration, covering 2,539 species (93% of the FloraOn catalog). Only 0.5% of these images are from FloraOn, and approximately two-thirds are taken from iNaturalist. Moreover, only 1,678 species reached our lower bound threshold of 50 images (61% of the FloraOn catalog). This scarcity for some species can be due to subjective issues like the visual attractiveness of the plant, e.g., having a showy flower, or it can be a real effect, reflecting its rare status in the wild. This effect makes the raw data distribution long-tailed (Figure 2a, shown in logarithmic scale).

The Floralens dataset was derived by sampling the raw data as follows. First, we filtered out species with less than 50 images. Then, for each of the remaining species, we sampled up to 200 images from the datasets, prioritizing data sources in the following order: (1) FloraOn; (2) Pl@ntNet; (3) Observation.org, and; (4) iNaturalist. That is, for the 50-200 image target per species, we use up as many images as possible from FloraOn first, then from Pl@ntNet, and so on.



(a) Raw data.
(b) Floralens dataset.
Fig 2. Dataset histograms (x-axis: number of images; y-axis: number of species).

This source-based prioritization intends to define a dataset where images are less prone to identification errors. It takes into account the curation processes associated with each data source. FloraOn is curated by botanists and features high-quality images. These often feature subtle details that help secure the identification of a species. Pl@ntNet data goes through a curation process that involves machine learning, contributors' reputation scores, and geo-based species verification [49]. Observation.org data can result from automatic validation through image recognition coupled with a check for other approved observations in the geographical vicinity, or through an expert volunteer when automated validation fails [48, 51]. Finally, iNaturalist identifications result from a crowd-sourcing effort whereby a "research-grade" identification for an image can be obtained from the consensus of just two citizen scientists [47, 52].

We used images without distinction of plant features as this metadata was not available from the curated datasets. Similarly, we do not attempt to gauge "image quality". The citizen-science platforms provide guidelines for submitting acceptable quality images, and their curation processes mostly filter inadequate items (e.g., out of focus, under/over-exposure, presence of significant digital noise, insufficient resolution, subject not sufficiently separated from surrounding vegetation).

This procedure resulted in the Floralens dataset, a collection of approximately 300K labeled images covering 1,678 species. As illustrated in Figure 2b, there are 200 images or very close to it for most of the species. The image count is 200 for 67% (1,128) of the species, 150 or higher for 79% (1,323), and 100 or higher for 86% (1,449). The data source prioritization scheme lead to a more significant fraction of FloraOn images in comparison to the raw data (the fraction grows from 0.5 to 5.1%) and, also, to a relatively even distribution of images from iNaturalist, Observation.org, and Pl@ntnet (the corresponding fractions are 30.7, 29.3, and 34.9%).

#### 3.2 Model

The process of deriving an image classification model using GAMLV is illustrated in Figure 3. Overall, it comprises three stages: (1) preparing the data set for training; (2) training the model, and (3) deploying the model onto a cloud server or (using a suitable format) onto edge devices. GAMLV essentially requires the user to focus on the dataset preparation (1), given that training (2) and deployment (3) merely require simple high-level options by the user and are otherwise automated [53,54]. The interaction with GAMLV can be conducted via a browser with a simple user interface, as we illustrate partially in this section (cf. Figure 4), or programmatically using Google Cloud APIs (e.g., in Python).



Fig 3. Model derivation using GAMLV.

The first step requires the user to load the dataset images onto a storage bucket, in this case, provided by the Google Cloud Storage service (GCS), along with a simple CSV file. The latter lists the GCS image URIs and associates each URI to a ground truth label (the name of the species in the image) and to either the train, validation, or test subset.

We fed GAMLV with train, validation, and test splits over the Floralens dataset, corresponding to fractions of 80%, 10%, and 10%. As usual, the train split is used to adjust CNN parameters during the training process in an iterative feedback loop, the validation split is used to measure the progress and convergence of that training process, and the test split is used merely for evaluating the model after training. The splits, with the image counts and data source provenance detailed in Table 2, resulted from a random selection of images for each species. Since the selection process is random and given the volume of images at stake, the overall fraction of images of each data source in each split closely matches that of the overall dataset.

Data source	Train	Valid.	Test	Total	
FloraOn	12,175	1,466	1,550	15,191	(5%)
iNaturalist	$72,\!174$	8,924	9,029	90,127	(31%)
Observation.org	68,614	$8,\!603$	$^{8,529}$	85,746	(29%)
Pl@ntNet	81,918	10,367	10,252	$102,\!537$	(35%)
All	234,881 (80%)	29,360	(10%) 29,360	(10%) 293,601	(100%)

Table 2. Train, validation, and test splits over the Floralens dataset.

In [14], we consider other strategies for defining these splits. In particular, we explored approaches that prefer specific data sources for the validation/test splits. We found that a random split, besides preserving a roughly similar fraction of images per data source in each split, results in models with better performance (contrast the results in Section 4 with those in [14]). In any case, prioritizing particular data sources (e.g. FloraOn or Pl@ntnet) for validation/test splits over others had little impact on model performance.

1	Define your model         Model name *         floralens         Octow hosted         Host your model on Google Cloud for online predictions         Image: Continue         Continue         Continue         Optimize model for         Goal       Package size         Accuracy       Latency for Google Pixel 2				Use your model			
2				redictions	TF Lite Export your model as a TF Lite package to run your model on edge or mobile devices.	TensorFlow.js Export your model as a TensorFlow,js package to run you model in the browser and in Node.js.		
				Latency for Google Pixel 2				
	<ul> <li>Higher accuracy</li> </ul>	6 MB	Higher	360 ms	<b>S</b>			
	O Best trad off	Best trade- 3.2 MB off	Medium         150 ms         Core ML           Lower         56 ms         Export a .mlmodel file to run your model on iOS and macOS devices.	150 ms	Core ML	Container		
	C Faster predictio	0.6 MB		Export a .mlmodel file to run your model on iOS and macOS devices.	Export your model as a TF Saved Model to run on a Docker contair			
	Please note that prediction latency estimates are for guidance only. Actual latency will depend on your network connectivity.           CONTINUE							
3	Set a node	hour budget	t		પ્યું			
	Enter the maximum number of node hours you want to spend training your model. We recommend using <u>16 node hours</u> for your dataset. However, you can train for as little as 1 node hours. You may also eligible to train with free node hours. Pricing guide [2]			nt to spend training your model. Iset. However, you can train for rain with free node hours.	Coral Export your model to run on Edge TPU-based devices.			
	Set your budget * 16 node hours			rs				
	Estimated completion date: Jul 25, 2023 3 AM GMT+1							

(a) Training parameters.(b) Deployment options.Fig 4. GAMLV interface for model training and deployment.

Once the dataset is imported onto AutoML, training may proceed, requiring only the user to make high-level choices for the type of model to be generated and the maximum training time, as illustrated in Figure 4a. Since we wish to use the model as part of web or mobile applications (cf. Section 5) rather than deploying it on a Google Cloud server, we select the option "Edge" model. We also toggle the option for a model that favors accuracy over latency among the three available choices. The maximum training time is specified as a "node hours" budget, where nodes are virtual machines used during training.

GAMLV required 4 node hours to complete the training of the CNN with the Floralens dataset. The service operates as a "black box". Thus, it is not possible to discern what goes on during training. For instance, no exact details or configuration options are provided for the training infrastructure (e.g., in terms of virtual machines, GPUs, or TPUs) and it is not possible to track details regarding the training process (e.g., how the model converges over time).

Once training is completed, a model can be deployed in one of several formats amenable for integration with a local application (Figure 4b). The formats include the standard SavedModel format used by TensorFlow, but also others like TF Lite [55], a lightweight TensorFlow format for use in resource-constrained hosts (e.g., mobile and embedded devices), or TFJS [56], for use in web browsers or Javascript programs. We use the TF Lite and TFJS variants in the software artifacts described in Section 5.

The model obtained by GAMLV is a 65-layer deep CNN, with the structure partially illustrated in Figure 5 for the TF Lite version, in terms of the input/initial layers (a), intermediate layers (b), and final/output layers (c). Each box in the image represents a type of neuron in the CNN, working as a function that takes a multi-dimensional (e.g., 2D, 3D, ...) vector, known as a tensor, and produces an output vector (output tensor). This function is parameterized by internal weights repeatedly adjusted using techniques such as back-propagation as part of an iterative training process. CNNs use a particular



family of functions called convolutions, which are especially suited for detecting image features (e.g., edges). In the simplest type of convolution, operating over 2D matrices, each position in the output vector called a feature map, is the dot product of a sliding window over the input matrix with a filter defined by the internal weights of the neuron. For details, see for instance Chapter 9 of [57].

The TF Lite version differs from the standard TensorFlow model only in terms of post-training optimizations like quantization (conversion of floating point weights to an integer scale) that enable the model to be interpreted faster with little degradation in accuracy [55]. As shown in Figure 5, the input layer takes a  $224 \times 224 \times 3$  tensor, corresponding to a  $224 \times 224$  (typically resized to conform with the dimensions of the CNN input) image with 3 RGB channels, with 8-bit values per color channel. The intermediate layers use several types of convolutions in a repeating pattern. The final layers include the derivation of a 1,280-feature map (a vector summarising all the captured image features [57]). This feature map is then passed to a fully connected soft-max activation function that produces the final classification vector with the label probabilities, 1,678 of them in line with the number of species covered. In this context, fully connected means that every value of the feature map, combined with the internal weights, is taken into account to calculate the value of each output position. The soft-max activation function is used to produce the neural network's final output which, in this case, is a probability distribution [57]. As expected for a probability distribution, the sum of all the probabilities in the output vector equals 1. Note that we always get such a vector, even for images outside the domain in question (in our case the native Portuguese flora). If the model is certain of a particular species, we get a high probability value for that species and low values for all the others. It can happen of course that the model is less certain, for instance when the probabilities assigned to the most high-ranked species are close (e.g., two or more very similar species). Finally, when the model fails to make any meaningful identification we get low probability values for all species.

The CNN architectures at stake are picked from the MnasNet family [58], developed with mobile and embedded devices in mind. The high-level choice between models offered by GAMLV (c.f. Figure 4a) corresponds to three different MnasNet

instantiations that do not differ in structure, just in the density of connections between layers and the number of internal weights.

#### 3.3 Evaluation Metrics

The performance of the Floralens model was evaluated over several test datasets of observations using the following metrics. Precision is the ratio of true positives (TP) relative to the total number of positives (TP + FP). A positive (identification) occurs when the classification score (a probability) returned by the model equals or exceeds a confidence level set as the threshold for analysis. Recall is the ratio of true positives relative to the total number of true examples (TP + FN). The confidence level is relevant for practical uses of a model. Using a low confidence level means that false positives are tolerable for the intended use, and we will get lower precision. On the other hand, a value too high for the confidence level will result in more false negatives, hence lower recall.

$$Precision = \frac{TP}{TP + FP} \qquad Recall = \frac{TP}{TP + FN}$$

Top-1 is the fraction of test images that the model correctly classified by the label with rank 1 (the highest-scoring label). Top-5 is similar to Top-1 but accounts for test images with a rank lower or equal to 5 (the 5 highest-scoring labels). We also use a variant of the Mean Reciprocal Rank (MRR) for test images of rank less or equal to 5. These are defined as follows:

$$Q(r_l) = \{t \in T \mid rank(t) \le r_l\}$$

Top-1 = 
$$\frac{|Q(1)|}{|T|}$$
 Top-5 =  $\frac{|Q(5)|}{|T|}$  MRR =  $\frac{1}{|T|} \sum_{t \in Q(5)} \frac{1}{rank(t)}$ 

where T is, as above, the set of all test observations (|T| = 29,360 as given in Table 2), rank(t) is the rank of the ground truth label returned by the model for the test observation t, and  $Q(r_l)$  is the subset of T that contains test observations with rank less or equal to a limit  $r_l$ .

### 4 Results

In this section, we present the results of multiple experiments conducted to characterize the Floralens model's performance and compare it to state-of-the-art models. In the following subsections, larger values indicate better model performance.

#### 4.1 Base Results

Figure 6 shows the results for precision and recall for the Floralens model applied to the test set given in Table 2, more precisely the macro-average of precision and recall values for all species. The area-under-curve (AUC) for the precision-recall correlation (in 6a), also known as the average precision, is 0.72 (the maximum value would be 1.0). Putting the confidence levels in perspective (in 6b) we see that precision and recall are approximately equal to 0.7 for a confidence level of 0.2. For a confidence level of 0.5 precision equals 0.85 and recall equals 0.53. Overall, the results indicate a reasonable predictive power for the Floralens model.



(a) Precision-Recall curve.(b) ValFig 6. FLTS results: precision and recall.

Figure 7 shows the Top-1, Top-5, and MRR results for FLTS in two plots: (7a) overall results, and per data source used in constructing the dataset (FloraOn, iNaturalist, Observation.org, and Pl@ntNet) and; (7b) boxplots displaying the distribution of the metric results over the universe of species, i.e., values for the metrics at stake for each species. The overall results again indicate reasonably good predictive power: 0.67 for Top-1 (i.e., roughly two-thirds of the FLTS images are correctly classified with rank 1), 0.86 for Top-5, and 0.75 for MRR; the results also indicate a relatively homogeneous predictive power across all data sources and metrics, as the maximum difference in values between them does not exceed 0.06: a Top-1 value of 0.64 for Observation.org vs. corresponding values of 0.70 for FloraOn and iNaturalist. Regarding the distribution of results per species, the median values for Top-1, Top-5, and MRR are respectively 0.70, 0.89, and 0.75. Looking at the first quartile, we also get indicators for reasonable predictive power: only 25 % (resp. 75 %) of species have a Top-1 value of 0.55 or lower (resp. or higher), and; the first quartile values for Top-5 and MRR are 0.80 and 0.65.

In Figure 8 we present two types of complementary results. The plots detail the Top-1, Top-5 and MRR grouped by: (8a) species growth type and; (8b) the special



(a) Overall results and results per data source.



(b) Distribution of results per species. Fig 7. FLTS results: Top-1, Top-5 and MRR

categories of species that are endangered (part of the IUCN red list), endemic, protected, or rare. In both cases, the FloraOn site [10] is the reference for the presented subsets of species. The number of species in each subset is given between parenthesis in the figures (e.g., 42 species of ferns). The y-axis reference ticks correspond to the overall FLTS values for the Top-1 (0.67), MRR (0.75), and Top-5 (0.86) metrics. Regarding species growth type, we can observe that grass species have the worst performance, as they tend to be hard to distinguish visually, and tree species are also slightly below the overall FTLS performance. As for the special categories, none of them stand out below or above the overall FTLS values. This is a good trait for the endemic species since these can be challenging for models to detect [43].



(a) Results according to species growth form.





Fig 8. Results according to species growth form and special categories.

#### 4.2 PlantCLEF and Wikipedia Test Sets

We considered two additional test sets: a random sample of 10,000 labeled images from the PlantCLEF'22-23 [46,59] competition, and a sample of close to 1,500 images from Wikipedia. The PlantCLEF data we used was only a tiny sample of the entire "trusted" training set of PlantCLEF [60] that comprises approximately 2.9 million images covering 80,000 plant species. The repository is trusted as the image labels were obtained from academic sources or collaborative platforms like Pl@ntNet or iNaturalist. Our subset was built by first filtering out species not covered by the Floralens model, obtaining data for 1,593 (out of 1,678) species, and then randomly sampling 10,000 images.

As for the Wikipedia test set, the images were identified through the Wikimedia REST API search functionality [61]. For each species in the Floralens domain, we used the species name as the keyword for a REST API search. Among other information, the

search result typically yields a reference to an image stored at Wikipedia, which we then considered for addition to the test set. After obtaining the images, we filtered out illustrations, herbarium specimens, and duplicates (associated with more than one species). Duplicates typically arise because the search may yield an image of a different species in the same genus, e.g., if the target species' name does not have a Wikipedia page. Through this process, we obtained a dataset of 1,351 images for an equal number of species (one image per species). Compared to the PlantCLEF test set, the identifications associated with these images are less reliable as they result from an uncontrolled crowd-sourced effort with no specific directives for validation.

Figure 9a shows the results of the Floralens model for these test sets considered in terms of the Top-1, Top-5, and MRR metrics for species prediction. We also recall the overall FTLS results (from Figure 7a) for easy comparison. We can observe that the PlantCLEF and Wikipedia test set results are marginally lower than those obtained for the FLTS (0.02 to 0.03 in all metrics). The results suggest that the Floralens model performs well with datasets other than our base test suite.

#### 4.3 Genus and Family Identification Results

Figures 9b and 9c show the results of genus and family identification for all datasets. The classification score for a genus/family is obtained by summing the scores for all the species in that genus/family output by the model. In [14] we created a model specific for genus identification but found that the results were essentially the same as obtained using this aggregation method. Genus results outperform species results (in Figure 9a) by 0.09 to 0.14 for Top-1, 0.05 to 0.07 for Top-5, and 0.07 to 0.11 for MRR. In turn, the family results outperform genus results by 0.07 to 0.08 for Top-1, 0.04 to 0.05 for Top-5, and 0.06 to 0.08 for MRR. As with species, the genus and family results are homogenous across all datasets. The more significant variation is observed for the genus results of the Wikipedia test set, e.g., the Top-1 result is 0.14 higher than that for the species. This might be explained by the fact that even when the image on the page for a species is wrongly labeled, Wikipedia does manage to provide one for a plant of the same genus (and by implication of the same family).

#### 4.4 Classification using Multiple Images

It is often the case that an observation of a plant in the wild has multiple images associated with it to enable better identification. In this section, we take advantage of this fact and compute an aggregate score for each such observation based on the scores returned by Floralens for each image of a given specimen. We apply this operation to all observations with two or more associated images in the FLTS.

We identified 1185 such observations, corresponding to a total of 2480 images. More precisely, most observations have 2 images, 1090 out of 1185 (91%), and the remaining 95 observations have 3 or more images (82 have 3 images, 12 have 4 images, and a standalone observation has 6 images). The observations at stake are grouped by the GBIF identifier of images in the iNaturalist, Observation.org, and Pl@ntNet datasets. We cannot determine which images are related to the same observation for images obtained from FloraOn. Our analysis was not extended to the PlantCLEF dataset for the same reason. In the Wikipedia dataset, there is only one image per species.

For the 1185 observations of single specimens, we calculated an aggregate classification score by taking the average of the scores of the individual images. Figure 10 compares the standalone image classification (the baseline FLTS results) vs. the aggregate classification of observations in terms of Top-1, Top-5, and MRR. We observe clear score improvements: +0.13 for Top-1 (0.67 to 0.80), +0.08 for Top-5 (0.86 to 0.94), and +0.12 for MRR (0.74 to 0.86). The figure also depicts (rightmost) the









results for the (comparatively few) 95 observations (9% of the cases) with 3 or more images showing even more pronounced improvements.

#### 4.5 Classification using Geographical Location

Some automated classification platforms use the geographical location associated with observations to improve the accuracy of the results. For instance, Pl@ntNet provides a parameter for the geographical zone (e.g., Southwestern Europe) and filters out species that do not occur in that region from the classification results, regardless of the output of the image classification model. We analyzed the impact of this strategy by making use of two items of information: (1) detailed FloraOn records of the geographical location of observations in Continental Portugal (excluding the Azores and Madeira archipelagos), and (2) the geographical location associated with FLTS images obtained from GBIF.

FloraOn geographic distribution data is based on the Military Grid Reference



Fig 10. Results for classification using multiple images.

System (MGRS) grid zone with a resolution of  $10 \times 10$  km. For any given species, it keeps a set of MGRS grid elements for which observations have been reported. We collected this data for almost all species covered by the Floralens dataset – 1654 out of 1678 (98%). Next, we identified all images related to observations in Continental Portugal in the FLTS through GBIF. These amount to 3172 images covering 770 species. We analyzed the concordance between geographical locations and GBIF, by calculating for each image the shortest distance between the GBIF coordinates and the FloraOn grid for the ground truth species at stake. The cumulative distribution function (CDF) of the shortest distances is shown in Figure 11a. Note that the distances in the x-axis have a logarithmic scale. The results indicate good agreement between FloraOn and image coordinates: the 95-th, 97-th, and 99-th percentiles of the distribution of distances are 18.9 km, 24.7 km, and 42.4 km.

When we supplied the model with an image taken at a given location, it replied with a ranking of suggested species. For each species in the ranking, we calculated the distance d of the observation's location to the closest known observation of that species. We discarded all species from the ranking output by the model such that d > D, where D is a parameter. We applied this geographical filter for varying values of the distance threshold D from 20 to 200 km with increments of 10 km. Note that the granularity of the MGRS grid used in FloraOn is 10km, thus justifying the choice of 20 km as the lowest value for D. To gauge the impact of the filter, we calculated the variations in the Top-1, Top-5, and MRR metrics versus the respective values without applying the filter. With no filter, we obtain 0.74 for Top-1, 0.89 for Top-5, and 0.80 for MRR. The variations are shown in Figure 11b. As illustrated, we observe small positive variations, allowing us to conclude that using a geographical filter is beneficial if not significantly so. The highest variations are in the case of D = 40 km: +0.04 for Top-1, +0.02 for Top-5, and +0.03 for MRR. As D grows larger, the impact of the filter becomes less



(a) CDF plot for the shortest distance between GBIF coordinates and FloraOn grid records.





(b) Top-1, Top-5, and MRR variations for different values of *D* in the geographical location filter.

(c) Images with altered ground truth ranks for different values of D in the geographical location filter

Fig 11. Results for geographic data filter.

significant. The operation of the filter is also illustrated in Figure 11c in terms of the fraction of images whose classification results have been affected by the filter. The plot distinguishes the following two cases: (1) the ground truth (species) rank is improved, and; (2) the ground truth is filtered out from the results. For D = 20km we obtain the maximum fraction of ground truth rank improvements and filtered results, respectively 12% and 4%. As D increases, the fraction of affected images decreases to less than 5%. Note that the filter for D = 20km is too aggressive, wrongly removing species from the rankings and hindering any gains from improving the ground truth's ranking position. This effect explains why the best Top-1, Top-5 and MRR results for the filtered model are obtained for D = 40km.

#### 4.6 Comparative Pl@ntNet API Results

To compare Floralens with the current state-of-the-art in automatic identification of plants, and given the restrictions in accessing the APIs for most of the platforms of interest (c.f. discussion in Section 2), we opted to use Pl@ntNet for which we were generously granted free extended access. Recent work has identified Pl@ntNet as one of the best-performing platforms for this purpose [43,62] and, therefore, a good reference point for our evaluation.

The Pl@ntNet API [63] is a web service providing access to the same visual identification models used by state-of-the-art Pl@ntNet apps [6]. The API returns a set of ranked species for a given image for two models for worldwide flora: a so-called "legacy" model from 2022 (henceforth  $PN^{22}$ ) generated using CNN, and; a more recent model announced in July 2023 [26], generated using Vision Transformers, henceforth  $PN^{23}$ . Through the API, it is also possible to filter results from the  $PN^{23}$  model so that only species occurring in a specific biogeographic region are included. One such region is Southwestern Europe including Portugal allowing the most head-to-head comparison between Floralens and Pl@ntNet possible. These results are identified by  $PN^{23F}$ .



Fig 12. Pl@ntNet API: comparative MRR values.

In Figure 12, we compare the MRR values obtained for Floralens, PN<sup>22</sup>, PN<sup>23</sup> and PN<sup>23F</sup> for all test datasets (FLTS, PlantCLEF, and Wikipedia). We also depict separate results gauging the impact of Pl@ntNet images: the MRR for the FLTS subset excluding Pl@ntNet images (c.f. "FLTS excl. PN" series) and for the FLTS subset containing only Pl@ntNet images (c.f. "FLTS only PN" series). Recall that images from Pl@ntNet are used to train our model, which is also (of course) the case for Pl@ntNet models. In particular, part of the Pl@ntNet images we use for testing may have been used to train the Pl@ntNet models, i.e. Pl@ntNet images in our test split could be part of Pl@ntNet train splits. Hence, the set of FLTS test images excluding Pl@ntNet may in this sense be a fairer comparison between Floralens and Pl@ntNet.

Regarding the overall test datasets, we can draw the following main conclusions: (1) Floralens performs better than  $PN^{22}$  by +0.07 (0.75 vs. 0.68) and has almost the same performance as  $PN^{22}$  for PlantCLEF (0.73 vs. 0.72), and Wikipedia (0.72 vs 0.73); (2) Floralens performs worse than both  $PN^{23}$  and  $PN^{23F}$  by -0.05 for FLTS (0.75 vs. 0.80), -0.06 for PlantCLEF (0.73 vs. 0.79) and -0.06/-0.07 for Wikipedia (0.72 vs. 0.78 and 0.72 vs. 0.79), and (3) the Southwestern Europe species filter associated with  $PN^{23F}$  has little impact on the results, compared to  $PN^{23}$ .

Focusing now on the impact of Pl@ntNet images on the FLTS results, we can see it is significant for the Pl@ntNet models, unlike Floralens which has homogeneous performance across different data sources in the FLTS as discussed previously (c.f. Figure 7a). It is clear that the Pl@ntNet models perform much better for Pl@ntNet images alone ("FLTS only PN" series), which amount to 35% of the entire FLTS (c.f. Table 2 in Section 3), than for the other images in the FLTS ("FLTS excl. PN" series): by +0.14 in the case of PN<sup>22</sup> (0.77 vs. 0.63) and +0.06 in the case of PN<sup>23</sup>/PN<sup>23F</sup> (0.84 vs. 0.78). In contrast, the corresponding Floralens values are very close for these two sets: 0.75 and 0.74. Given Pl@ntNet models' results seem to be biased for Pl@ntNet images (conforming to our initial hypothesis), FLTS images excluding Pl@ntNet may provide a more balanced comparative analysis. FLTS performs better than PN<sup>22</sup> as discussed above (0.75 vs. 0.68), but even more so for non-Pl@ntNet images (0.75 vs. 0.63, a difference of +0.12); as for PN<sup>23</sup>/PN<sup>23F</sup> models, the difference is much smaller for non-Pl@ntNet images (0.75 vs 0.78, a difference of +0.03) than for Pl@ntNet images (0.74 vs. 0.84, a difference of +0.10).

Figure 13 presents comparative MRR results again for FTLS, but grouped by (13a) species growth type and (13b) the special categories of species as analysed previously for Floralens. For a more balanced comparison, in line with the above discussion, we leave out FLTS images obtained from Pl@ntNet. The results are generally on par with the general trends already discussed. Floralens outperforms  $PN^{22}$  by a factor +0.09 (endemic species 13b) or higher up to +0.16 (rare species). In turn, it is outperformed by  $PN^{23}/PN^{23F}$  by a factor between -0.01 (rare species in 13b) and -0.10 (tree species in 13a). The cases where Floralens fares more favorably compared with  $PN^{23}/PN^{23F}$  are herb species regarding growth type (0.75 vs. 0.78), and rare species (0.73 vs 0.74) regarding special species' statuses.



(a) Results according to species growth form.



(b) Results for special categories.

Fig 13. Pl@ntNet API: comparative MRR values grouped by species growth form and special categories.

# 5 Data and Software Artifacts

In this section, we briefly describe the Biolens software platform, in which the Floralens has been made publicly available, as well as other artifacts provided to the community.

#### 5.1 Biolens Website

The Floralens model has been integrated into the Biolens project website [9]. As illustrated in the screenshots of Figure 14, the functionality is quite simple: (a) users submit photos of interest, and (b) obtain suggestions for identifications, with a textual indication of the model's confidence – high (> 70%), medium (70% - 40%) and low (< 40%). Up to 5 species are listed by the app, as long as the model outputs a minimum score of 15% for each identification. The Biolens website is hosted by a small virtual machine that requires just 2 CPU cores and 8 GB of RAM. With this configuration, invoking the model takes an average of 900 milliseconds per image, calculated directly from approximately 29,000 requests to the server between April 2021 and April 2025. As for the model itself, it takes only 8 MB of disk storage. The server configuration is lightweight as we use the TF Lite variant of the Floralens model running on the server side (similarly to the other models for other taxa also hosted on the site). The user's browser merely displays the server's results; it does not host the model.



(a) Image uploading by the user.(b) SuggestFig 14. Biolens – web application screenshots.

(b) Suggested automatic identification.

#### 5.2 Biolens App

We also recently developed a prototype version of a mobile application that can run on Android and iOS devices. The Android version is available for download at the Biolens website. A few screenshots of the application are shown in Figure 15. The functionality is similar to that of the Biolens website, but customized for a mobile application context: users can take photos of specimens on the fly and obtain instant identification suggestions without an Internet connection. All Biolens models are bundled within the app and, thus, are evaluated *in loco* on the mobile device. The identification information, the date, the current geographical location, and optional user annotations, are recorded in association with each photo. According to the reference latency values provided by Google AutoML for devices like Google Pixel 2, Samsung Galaxy S7, or iPhone X, invoking the model takes less than 100 milliseconds or lower in modern mobile devices. This is consistent with our experience and much faster than in the case of the Biolens server discussed above. The server's configuration is quite modest compared to present-day mobile devices.



(a) Model selection.(b) Image list.Fig 15. Biolens – mobile application screenshots.

(c) Details for an image.

#### 5.3 Dataset and Results

The dataset and the Python notebooks with all the code used for the results of Section 4 are available publicly from Zenodo [16]. The dataset contains the mapping between the image labels (ground truth), the image URLs from which they were retrieved, URLs for a site we maintain where all images are also stored, and GBIF identifiers when applicable (all images except those obtained from FloraOn). Ground truth and URLs are also available for the PlantCLEF and Wikipedia datasets. The Top-5 results and the corresponding confidence levels for the Floralens model and the three Pl@ntNet model variations are included for all datasets.

## 6 Discussion

In this paper, we describe the construction of a dataset for the Portuguese flora and the derivation of a deep-learning model for the automatic identification of the species therein. The universe of species was taken from the FloraOn dataset, provided by the Sociedade Portuguesa de Botânica and compiled exclusively by specialists. The dataset was constructed based on high-quality data from several research-grade datasets available via GBIF. Besides FloraOn these include: iNaturalist, Pl@ntNet, and Observation.org. The dataset is available publicly on Zenodo [16]. The Floralens model was derived from this dataset using GAMLV which provides users with tools to derive models from datasets using off-the-shelf convolutional deep neural networks. The model is available online at the BioLens Project website and can also be used as part of the Biolens mobile application (e.g., offline, in the field).

The baseline Floralens model has good predictive power, with an AUC metric value of 0.72 and, for a reference confidence level of 0.5, values for precision and recall of 0.85and 0.53, respectively. It also features a relatively homogeneous predictive power across all the dataset's data sources. We measured values of 0.72 for Top-1, 0.88 for Top-5 and 0.78 for MRR. An interesting result concerns the model's accuracy in identifying endangered, rare, protected and endemic species. Some research (e.g., [43]) suggests that current platforms struggle for these classes. The results for Floralens over these categories show that the model maintained its nominal predictive power. Using multiple images of the same specimen significantly improved the model's accuracy for all metrics: 0.80 for Top-1, 0.94 for Top-5, and 0.86 for MRR. This agrees with the findings reported by other projects (e.g., [6,41]). Introducing filters based on species' geographical distribution also improved, albeit less significantly, the model's accuracy. The results testify to the accuracy and completeness of the FloraOn dataset's geographical data. Finally, we note that the accuracy of the Floralens model is significantly better than Pl@ntnet's legacy model (PN<sup>22</sup>) which uses the same technology for model derivation, namely Convolutional Neural Networks. Moreover, the difference in performance between Floralens and the new Pl@ntnet models (PN<sup>23</sup>and PN<sup>23F</sup>), that use Vision Transformers for model derivation, is much smaller. In this case, Pl@ntnet's better accuracy might be attributed to using Vision Transformers. In other words, for the same model derivation technology, Floralens clearly outperforms Pl@ntnet for the Portuguese flora.

As for future work, we aim to improve the model's accuracy by increasing the number of observations in the dataset that feature multiple images of the same specimen, eventually tagged by plant part. This will involve a concerted effort with users of the Floralens app and specialists, combined with extracting more multi-image observations from thus far unexplored publicly available datasets, e.g., Flora Incognita. Including images from Morocco and Algeria will help enhance the dataset with images from species in common with the Mediterranean flora. We also want to address limitations that arise from the inconsistent use of taxonomic names and synonyms. For example, the species listed in FloraOn as *Atractulis gummifera* is now known as Chamaeleon gummifer. Although not ubiquitous, it is widespread in the Mediterranean region [64]. It is not included in the Floralens dataset as there weren't enough  $(\geq 50)$ images available. However, a recent query in GBIF for *Chamaeleon qummifer* yields more than enough images to include the species in the Floralens dataset in a future update. Another thread worth pursuing is using Vision Transformers to derive our models. As noted in Section 2 this is a relatively recent technology when applied to computer vision problems which appears to have significant advantages over traditional approaches such as CNNs. We also plan to continue preliminary work on using image similarity models [14]. These may provide an alternative way to clinch an identification when traditional computer vision models yield low-confidence results. Hybrid

classification models that combine both approaches are a possibility worth checking. Finally, more work is also required on the Biolens mobile app to improve its usability, optimize resource usage and integrate it with existing Citizen Science platforms. Acknowledgements. The authors thank Miguel Porto and Henrique Alves for their help with the FloraOn dataset and numerous insights, and Hugo Gresse, Pierre Bonnet, and Mathias Chouet for kindly giving us extended access to the Pl@ntNet API. This work was partially funded by projects SafeCities and Augmanity

(POCI-01-0247-FEDER-041435 and -046103, through COMPETE 2020 and Portugal 2020), by project UIDB/50014/2020 (Fundação para a Ciência e Tecnologia), and a cloud resource credits through the Google Cloud Research Credits program.

## References

- Bonney R, Cooper CB, Dickinson J, Kelling S, Phillips T, Rosenberg KV, et al. Citizen Science: A Developing Tool for Expanding Science Knowledge and Scientific Literacy. BioScience. 2009;59(11):977–984.
- 2. Altrudi S. Connecting to nature through tech? The case of the iNaturalist app. Convergence. 2021;27(1):124–141.
- 3. Schermer M, Hogeweg L. Supporting citizen scientists with automatic species identification using deep learning image recognition models. Biodiversity Information Science and Standards. 2018;.
- 4. Wäldchen J, Mäder P. Machine learning for image based species identification. Methods in Ecology and Evolution. 2018;9(11):2216–2225.
- 5. Mäder P, Boho D, Rzanny M, Seeland M, Wittich HC, Deggelmann A, et al. The Flora Incognita app-interactive plant species identification. Methods in Ecology and Evolution. 2021;12(7):1335–1342.
- Affouard A, Goëau H, Bonnet P, Lombardo JC, Joly A. Pl@ntNet app in the era of deep learning. In: International Conference on Learning Representations. Toulon, France; 2017.
- Tuia D, Kellenberger B, Beery S, Costelloe BR, Zuffi S, Risse B, et al. Perspectives in machine learning for wildlife conservation. Nature Communications. 2022;13(1):792.
- 8. Christin S, Hervet E, Lecomte N. Applications for deep learning in ecology. Methods in Ecology and Evolution. 2019;10(10):1632–1644.
- 9. Biolens; 2020. https://rubisco.dcc.fc.up.pt/biolens.
- 10. Flora-On: Flora de Portugal Interactiva; 2014. https://www.flora-on.pt.
- 11. Robertson T, Döring M, Guralnick R, Bloom D, Wieczorek J, Braak K, et al. The GBIF integrated publishing toolkit: facilitating the efficient publishing of biodiversity data on the Internet. PLOS One. 2014;9(8).
- 12. Lopes LMB, Marques ERB, Mamede T, Filgueiras A, Marques M, Coutinho M. Identificação Taxonómica em Biologia Usando Inteligência Artificial; 2022. Available from: https://rce.casadasciencias.org/rceapp/art/2022/050/.
- Marques M. A Portuguese Flora Identification Tool Using Deep Learning. Masters thesis, Faculty of Sciences, University of Porto; 2021. https://hdl.handle.net/10216/130189.

- Filgueiras A. Floralens: a deep learning model for portuguese flora. Masters thesis, Faculty of Sciences, University of Porto; 2022. https://hdl.handle.net/10216/145701.
- 15. Mamede T. On using Deep Learning for Automatic Taxonomic Identification of Butterflies. BSC project report, Faculty of Sciences, University of Porto; 2020. https:

//www.dcc.fc.up.pt/~edrdo/supervision/tmamede\_lepidoptera.pdf.

- Filgueiras A, Marques ERB, Lopes LMB, Marques M. The Floralens Dataset for Portuguese Flora; 2024. https://doi.org/10.5281/zenodo.10639701.
- Lee SH, Chan CS, Wilkin P, Remagnino P. Deep-plant: Plant identification with convolutional neural networks. In: IEEE International Conference on Image Processing; 2015. p. 452–456.
- Heredia I. Large-scale plant classification with deep neural networks. In: Computing Frontiers Conference; 2017. p. 259–262.
- 19. Sun Y, Liu Y, Wang G, Zhang H. Deep learning for plant identification in natural environment. Computational Intelligence and Neuroscience. 2017;.
- Bonnet P, Goëau H, Hang ST, Lasseck M, Sulc M, Malécot V, et al. 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; p. 131–149.
- 21. Observation.org Explanation NIA; 2022. https://observation.org/pages/nia-explain/.
- Fukushima K. Visual Feature Extraction by a Multilayered Network of Analog Threshold Elements. IEEE Transactions on Systems Science and Cybernetics. 1969;5(4):322–333.
- Fukushima K. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. Biological Cybernetics. 1980;36(4):193–202.
- Kolesnikov A, Dosovitskiy A, Weissenborn D, Heigold G, Uszkoreit J, Beyer L, et al. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In: International Conference on Learning Representations. Virtual Conference; 2021.
- 25. Bai Y, Mei J, Yuille AL, Xie C. Are Transformers more robust than CNNs? In: Advances in Neural Information Processing Systems; 2021. p. 26831–26843.
- 26. Pl@ntNet. Pl@ntNet news Covering all countries floras and new identification AI; 2023. https://plantnet.org/en/2023/07/05/ covering-all-countries-floras-new-identification-ai/.
- 27. iNaturalist Computer Vision Explorations; 2024. https://www.inaturalist.org/pages/computer\_vision\_demo.
- Joly A, Bonnet P, Affouard A, Lombardo JC, Goëau H. Pl@ntnet-my business. In: International Conference on Multimedia. ACM; 2017. p. 551–555.
- 29. Observation.org ObsIdentify; 2022. https://observation.org/apps/obsidentify/.

- Naturalis Biodiversity Center Nature Identification API v2; 2022. https://multi-source.docs.biodiversityanalysis.eu/index.html.
- 31. Amazon Rekognition; 2024. https://aws.amazon.com/rekognition/.
- 32. Apple Create ML; 2024. https://developer.apple.com/documentation/createml.
- 33. Automated Machine Learning (AutoML) Microsoft Azure; 2024. https: //azure.microsoft.com/en-us/solutions/automated-machine-learning/.
- 34. Choudhary M, Sentil S, Jones JB, Paret ML. Non-coding deep learning models for tomato biotic and abiotic stress classification using microscopic images. Frontiers in Plant Science. 2023;14:1292643.
- Korot E, Guan Z, Ferraz D, Wagner SK, Zhang G, Liu X, et al. Code-free deep learning for multi-modality medical image classification. Nature Machine Intelligence. 2021;3(4):288–298.
- 36. Borkowski AA, Wilson CP, Borkowski SA, Thomas LB, Deland LA, Grewe SJ, et al. Google Auto ML versus Apple Create ML for histopathologic cancer diagnosis; which algorithms are better? arXiv preprint arXiv:190308057. 2019;.
- Malounas I, Lentzou D, Xanthopoulos G, Fountas S. Testing the suitability of automated machine learning, hyperspectral imaging and CIELAB color space for proximal in situ fertilization level classification. Smart Agricultural Technology. 2024;8:100437.
- 38. Liang D, Xue F. Integrating automated machine learning and interpretability analysis in architecture, engineering and construction industry: A case of identifying failure modes of reinforced concrete shear walls. Computers in Industry. 2023;147:103883.
- Justamante A, Joly A, Lombardo JC, Robert F, Chouet M, Liñán S, et al.. AI-GeoSpecies: integrate artificial intelligence into your citizen science app; 2023. https://doi.org/10.5281/zenodo.7657594.
- WCVP: World Checklist of Vascular Plants; 2023. http://sftp.kew.org/pub/data-repositories/WCVP/.
- Rzanny M, Mäder P, Deggelmann A, Chen M, Wäldchen J. Flowers, leaves or both? How to obtain suitable images for automated plant identification. Plant Methods. 2019;15(1):1–11.
- 42. Brun P, de Witte L, Popp MR, Zurell D, Karger DN, Descombes P, et al.. Florid
   a Nationwide Identification Service for Plants from Photos and Habitat
  Information; 2024. Available from: https://ssrn.com/abstract=4830448.
- 43. Berjano, R and Lopez-Tirado, J and Martin-Escobar, I and Martinez-Sagarra, G and Nieto-Lugilde, D and Sanchez-Romero, J and De La Estrella, M. Mind your app: Could plant ID applications lead to an increase in extinction risk? Phytotaxa. 2023;609(1):65–68.
- 44. Van Horn G, Mac Aodha O, Song Y, Cui Y, Sun C, Shepard A, et al. The iNaturalist species classification and detection dataset. In: IEEE Conference on Computer Vision and Pattern Recognition; 2018. p. 8769–8778.

- 45. Goëau H, Bonnet P, Joly A. Plant identification based on noisy web data: the amazing performance of deep learning (LifeCLEF 2017). In: Conference and Labs of the Evaluation Forum; 2017.
- Goëau H, Bonnet P, Joly A. Overview of PlantCLEF 2022: Image-based plant identification at global scale. In: Conference and Labs of the Evaluation Forum. vol. 3180; 2022. p. 1916–1928.
- 47. iNaturalist contributors, iNaturalist (2022). iNaturalist Research-grade Observations; 2023. https://doi.org/10.15468/ab3s5x.
- 48. de Vries H, Lemmens M. Observation.org Nature data from around the World; 2023. https://doi.org/10.15468/5nilie.
- Affouard A, Joly A, Lombardo JC, Champ J, Goeau H, Bonnet P. Pl@ntNet observations. Version 1.2. Pl@ntNet; 2023. https://doi.org/10.15468/gtebaa.
- 50. Robertson T, Döring M, Guralnick R, Bloom D, Wieczorek J, Braak K, et al. The GBIF integrated publishing toolkit: facilitating the efficient publishing of biodiversity data on the internet. PLOS One. 2014;9(8).
- 51. Observation.org Validation; 2023. https://observation.org/pages/validation/.
- 52. What is the data quality assessment and how do observations qualify to become "Research Grade"?; 2023. https://www.inaturalist.org/pages/help#quality.
- 53. Bisong E. In: Google AutoML: Cloud Vision. Apress; 2019. p. 581–598.
- 54. AutoML Vision Documentation; 2023. https://cloud.google.com/vision/automl/docs/.
- 55. TensorFlow Lite, ML for Mobile and Edge Devices; 2023. https://www.tensorflow.org/lite/.
- 56. TensorFlow.js, Machine Learning for Javascript developers; 2023. https://www.tensorflow.org/js/.
- 57. Goodfellow I, Bengio Y, Courville A. Deep Learning. MIT Press; 2016.
- 58. Tan M, Chen B, Pang R, Vasudevan V, Sandler M, Howard A, et al. MnasNet: Platform-Aware Neural Architecture Search for Mobile. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition; 2019. p. 2815–2823.
- 59. PlantCLEF2022, Image-based plant identification at global scale; 2023. https://www.imageclef.org/PlantCLEF2022.
- PlantCLEF'22 trusted training set; 2022. https://lab.plantnet.org/LifeCLEF/PlantCLEF2022/train.
- 61. Wikimedia REST API; 2023. https://www.mediawiki.org/wiki/Wikimedia\_REST\_API.
- 62. Hart, A G and Bosley, H and Hooper, C and Perry, J and Sellors-Moore, J and Moore, O and Goodenough, A E . Assessing the accuracy of free automated plant identification applications. People and Nature. 2023;5(3):929–937.
- 63. Pl@ntNet. Pl@ntNet API for developers; 2023. https://my.plantnet.org.

64. Royal Botanical Gardens, Kew: Plants of the World Online: Chamaeleon gummifer; 2023. https://powo.science.kew.org/taxon/urn:lsid:ipni.org: names:192416-1.