



HAL
open science

Cooperative learning of Pl@ntNet's Artificial Intelligence algorithm: how does it work and how can we improve it?

Tanguy Lefort, Antoine Affouard, Benjamin Charlier, Jean-Christophe Lombardo, Mathias Chouet, Hervé Goëau, Joseph Salmon, Pierre Bonnet, Alexis Joly

► To cite this version:

Tanguy Lefort, Antoine Affouard, Benjamin Charlier, Jean-Christophe Lombardo, Mathias Chouet, et al.. Cooperative learning of Pl@ntNet's Artificial Intelligence algorithm: how does it work and how can we improve it?. 2024. hal-04603038

HAL Id: hal-04603038

<https://hal.science/hal-04603038>

Preprint submitted on 6 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Cooperative learning of Pl@ntNet's Artificial Intelligence algorithm: how does it work and how can we improve it?

Tanguy Lefort¹ Antoine Affouard² Benjamin Charlier³
Jean-Christophe Lombardo⁴ Mathias Chouet⁵ Hervé Goëau⁶
Joseph Salmon⁷ Pierre Bonnet⁸
Alexis Joly⁹

¹Univ. Montpellier, CNRS, IMAG, Inria, LIRMM, France tanguy.lefort@umontpellier.fr

²Inria, CIRAD, Montpellier, France antoine.affouard@cirad.fr

³Univ. Montpellier, CNRS, IMAG, France benjamin.charlier@umontpellier.fr

⁴Inria, LIRMM, France, jean-christophe.lombardo@inria.fr

⁵CIRAD, AMAP, Montpellier, France, mathias.chouet@cirad.fr

⁶CIRAD, AMAP, Montpellier, France, herve.goeau@cirad.fr

⁷Univ. Montpellier, CNRS, IMAG, Institut Universitaire de France (IUF), joseph.salmon@umontpellier.fr

⁸CIRAD, AMAP, Montpellier, France, pierre.bonnet@cirad.fr

⁹Inria, LIRMM, France, alexis.joly@inria.fr

Running Headline Citizen science for plant identification

Acknowledgments This work was funded by the French National Research Agency (ANR) through the grant Pl@ntAgroEco 22-PEAE0009, granted access to the HPC resources of IDRIS under the allocation A0151011389 made by GENCI, and funded by the Chaire IA CaMeLOt (ANR-20-CHIA-0001-01).

Data Availability The dataset is available at <https://doi.org/10.5281/zenodo.10782465>

Conflict of Interest The authors declare no conflicts of interest

Author Contributions T. Lefort, A. Affouard, A. Joly, B. Charlier, P. Bonnet and J. Salmon conceived the ideas and designed the evaluation methodology; A. Affouard and M. Chouet are the main developers of Pl@ntNet's backend ; T. Lefort and A. Affouard collected the evaluation data used in this paper; T. Lefort re-implemented Pl@ntNet's algorithm in python and conducted the evaluation ; J-C. Lombardo, H. Goëau and A. Joly conceived and trained Pl@ntNet's AI model; T. Lefort, B. Charlier, A. Joly and J. Salmon analyzed the outcomes of the study. All authors contributed critically to the drafts and gave final approval for publication.

Abstract

1. Deep learning models for plant species identification rely on large annotated datasets. The PI@ntNet system enables global data collection by allowing users to upload and annotate plant observations, leading to noisy labels due to diverse user skills. Achieving consensus is crucial for training, but the vast scale of collected data (number of observations, users and species) makes traditional label aggregation strategies challenging. Existing methods either retain all observations, resulting in noisy training data or selectively keep those with sufficient votes, discarding valuable information. Additionally, as many species are rarely observed, user expertise can not be evaluated as an inter-user agreement: otherwise, botanical experts would have a lower weight in the AI training step than the average user.
2. Our proposed label aggregation strategy aims to cooperatively train plant identification AI models. This strategy estimates user expertise as a trust score per user based on their ability to identify plant species from crowdsourced data. The trust score is recursively estimated from correctly identified species given the current estimated labels. This interpretable score exploits botanical experts knowledge and the heterogeneity of users. Subsequently, our strategy removes unreliable observations but retains those with limited trusted annotations, unlike other approaches.
3. We evaluate PI@ntNet's strategy on a newly released large subset of the PI@ntNet database focused on European flora, comprising over 6M observations and 800K users. This anonymized dataset of votes and observations is released openly at <https://doi.org/10.5281/zenodo.10782465>. We demonstrate that estimating users skills based on the diversity of their expertise enhances labeling performance.
4. Our findings emphasize the synergy of human annotation and data filtering in improving AI performance for a refined training dataset. We explore incorporating AI-based votes alongside human input in the label aggregation. This can further enhance human-AI interactions to detect unreliable observations (even with few votes).

Keywords: crowdsourcing, botanical skills, human-AI interaction, label aggregation, PI@ntNet, plant identification

1 Introduction

Computer vision models are a great aid in plant species recognition in the field (Vidal et al., 2021; Borowiec et al., 2022; Mäder et al., 2021). However, to train them we need large annotated datasets. These datasets are often created thanks to citizen science approaches, collecting both reliable and useful information (E. D. Brown and Williams, 2019). Among existing plant recognition applications, the PI@ntNet citizen science platform (Affouard et al., 2017) enables global data collection by allowing users to upload and annotate plant observations (Bonnet et al., 2020).

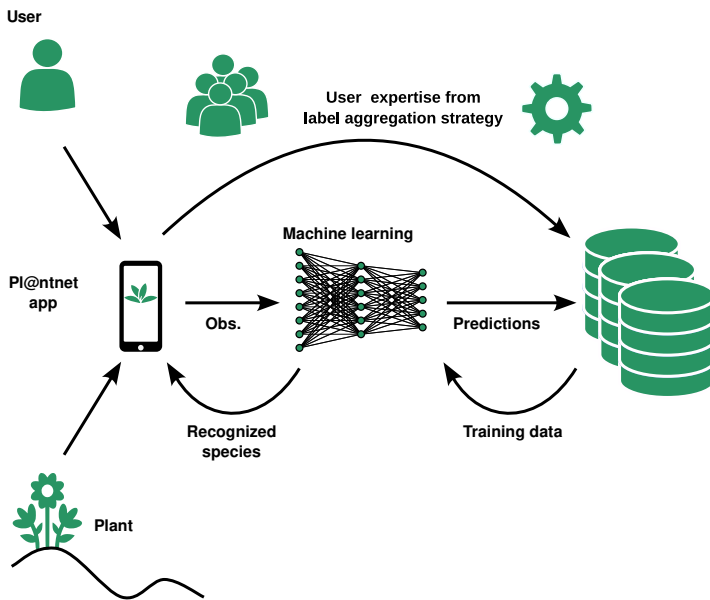


FIG. 1. PI@ntNet system of human-AI interaction for plant species recognition. Users take their plant observations in the PI@ntNet application. A prediction is output by the AI model. Users can validate the prediction or propose another species. The whole votes collection is used to evaluate user expertise (see Algorithm 1) and actively revise observations identifications.

At the time of writing, this participatory approach has resulted in the collection of over 20 million observations (images or group of images of the same plant), belonging to almost 46 000 species, by more than 6 million users worldwide. In total, more than 25 million of images are shared in these observations. The collaborative process of PI@ntNet is synthesized in Fig. 1. The AI model interacts with the human decision by proposing possible species given an observation. For each returned species, using a similarity search, the PI@ntNet system also shows similar pictures from the database. This lets users visu-

ally check that their observation is likely to belong to a predicted species given the most similar observations. For instance, such a visual control can help to compare two plants at various growth stages.

Plant species identification is a task that requires skills to recognize morphological traits (shapes, measurements, environments and specific characteristics). A large number of users with diverse skills have participated in gathering plant observations and helped improve the training dataset of our computer vision model. Their participation is based on votes that they can cast on others' observations, or by the initial species determination of their observation. The quality of each vote is then processed by the algorithm presented in Section 2.2.

Other citizen science projects such as iNaturalist (Van Horn et al., 2018) or eBird (Sullivan et al., 2009) use a similar approach to collect data, but differ in their label aggregation strategy. The iNaturalist project, with more than 2.5 million users, records the votes at different taxonomic levels. The resulting label is the aggregation of at least two votes on a species-level identification (or coarser or finer taxonomic level). A taxon requires at least a two-thirds agreements among identifiers and all users have the same weight in the decision-making. Over time, a taxon can be further refined by the community, debated or revoked. eBird handles taxon quality control by using a checklist in each region for observers. Quality control on the checklist is performed and, combined with user knowledge – number of species and checklist submitted, number of flagged observations, discussions among local experts – the species observation is accepted. The eBird project also showed that monitoring species accumulation from observers can help to sort their skills (Kelling et al., 2015). While they consider the species accumulation by hours spent on each collected observation, we propose a strategy that takes into account the entire history of observations of the observer.

In this article, we present the PI@ntNet label aggregation strategy. Using a new large-scale dataset of more than 6 million observations and 800 thousand users, we show that our strategy can improve the quality of the collected data, without removing every observation that was only labeled by single users. Finally, aggregated labels are used in practice to train an AI model. We explore how the information contained in the AI predictions can be integrated into the label aggregation strategy to generate new votes and help control data quality. By using the model's predictions within the label aggregation, the goal is to correct possible mistakes from non expert users without contradicting botanical experts.

2 Methods

2.1 Dataset and notation

To compare the different label aggregation strategies on large-scale datasets, we introduce a subset of the PI@ntNet database focused on Southwestern European flora observations – Balears, Corsica, France, Portugal, Sardegna and Spain – from 2017 to October 2023. In total, 9 005 108 votes are cast by $n_{\text{user}} = 823\,251$ users on 6 699 593 observations after two cleaning steps on the voted species. The first one is a filtering step. We only keep the votes with plant species belonging to the World Checklist of Vascular Plants (WCVP) (Govaerts, 2023). For the second step, according to Kew’s Royal Botanical Garden, we matched synonyms to their backbone species if the species is part of the *k-southwestern-europe* checklist from Plants of the World Online (POWO, 2024) (POWO) system. Note that there are plant species listed in the accepted species from WCVP that are not in the *k-southwestern-europe* POWO checklist. As there is a possible taxon ambiguity in this case – multiple species possible for a given synonym depending on the referential – we leave the proposed label untouched. The dataset is available at <https://zenodo.org/records/10782465>.

Notation In the following, denote K the number of species within the dataset. We index the observations by $i \in [n_{\bullet}] = \{1, \dots, n_{\bullet}\}$ where \mathcal{D}_{\bullet} is the considered dataset composed of n_{\bullet} observations and their associated votes. For example, the full southwestern European flora dataset from PI@ntNet of $n_{\text{SWE}} = 6\,699\,593$ observations is denoted \mathcal{D}_{SWE} . Other subsets are presented in Section 2.3. We write \mathcal{U} the set of all users. Each user u has a unique identifier used as an index, and we denote \mathcal{U}_i the set of users that have voted on observation i – *i.e.* $\mathcal{U} = \cup_{i \in [n_{\text{SWE}}]} \mathcal{U}_i$. The vote of user u on observation i is denoted $y_i^u \in [K]$. Estimated labels are denoted $\hat{y}_i \in [K]$. Each observation i is created by an author u stored in $\text{Author}(i)$.

2.2 Proposed label aggregation strategy

Algorithm 1 Pl@ntNet iterative weighted majority vote

Input: Votes as $(u, y_i^u)_{i \in [n_{\text{SWE}}], u \in [n_{\text{user}}]}$ for each observation i and user u answering the voted species y_i^u , accuracy threshold θ_{acc} , confidence threshold θ_{conf} , weight function f , initial weight $\gamma > 0$

Output: Estimated labels \hat{y}_i and validity indicator s_i for each observation i

- 1: Initialize user weights as $w_u = \gamma$ for each user $u \in [n_{\text{user}}]$
- 2: **while** not converged **do**
- 3: Get current estimated labels with a weighted majority vote

$$\forall i \in [n_{\text{SWE}}], \hat{y}_i = \arg \max_{k \in [K]} \sum_{u \in \mathcal{U}_i} w_u \mathbb{1}(y_i^u = k)$$

- 4: **for** each observation $i \in [n_{\text{SWE}}]$ **do**
- 5: Compute label confidence: $\text{conf}_i(\hat{y}_i) = \sum_{u \in \mathcal{U}_i} w_u \mathbb{1}(y_i^u = \hat{y}_i)$
- 6: Compute label accuracy: $\text{acc}_i(\hat{y}_i) = \text{conf}_i(\hat{y}_i) / \sum_{k \in [K]} \text{conf}_i(k)$
- 7: Compute validity indicator: $s_i = \mathbb{1}(\text{acc}_i(\hat{y}_i) \geq \theta_{\text{acc}} \text{ and } \text{conf}_i(\hat{y}_i) \geq \theta_{\text{conf}})$
- 8: **end for**
- 9: **for** each user $u \in [n_{\text{user}}]$ **do**
- 10: Compute the number of valid identified species for authoring observations:

$$n_u^{\text{author}} = |\{y_i^u \in [K] \mid y_i^u = \hat{y}_i, s_i = 1, \text{Author}(i) = u\}|$$

- 11: Compute the number of identified species by voting on other’s observations:

$$n_u^{\text{vote}} = |\{y_i^u \in [K] \mid y_i^u = \hat{y}_i, \text{Author}(i) \neq u\}|$$

- 12: Compute the rounding number of identified species per user:

$$n_u = \text{Round} \left(n_u^{\text{author}} + \frac{1}{10} n_u^{\text{vote}} \right)$$

- 13: Transform number of estimated species per user into trust score: $w_u = f(n_u)$
 - 14: **end for**
 - 15: **end while**
-

Pl@ntNet label aggregation strategy relies on estimating the number of correctly identified species for each user. Similar to other strategies, we rely on an EM-based iterative procedure (Dempster, Laird, and Rubin, 1977) to estimate consecutively the users’ skills and each observation’s species. The detailed iterative algorithm is provided in Algorithm 1 and available at <https://github.com/peerannot/peerannot>. The label aggregation strategy generates a trust indicator (s_i) on the observation that can reveal whether an observation is valid or not. Notice that in Algorithm 1 we value 10 times more authored observations than voting on other’s observations – if a user proposes a new observation

with a label (species name) it is more useful than proposing a label by clicking. Indeed, being on the field leads to more information on the environment and a better determination of the species. Finally, note that an identified species is exclusively identified as author – part of n_u^{author} in Algorithm 1) – or as click – part of n_u^{vote} – to avoid redundant skills. The final number of species identified by users is the aggregation of these two terms: $n_u = \text{Round} \left(n_u^{\text{author}} + \frac{1}{10} n_u^{\text{vote}} \right)$.

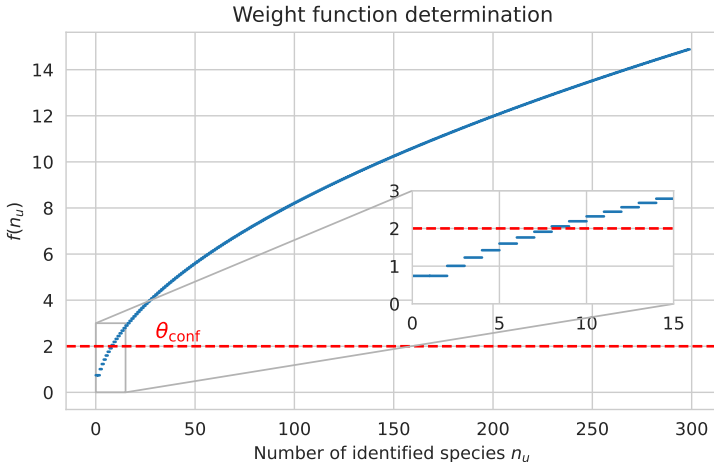


FIG. 2. Weight function in Equation (1) used to map the number of identified species to a trust score in the PI@ntNet label aggregation strategy. A new user starts with a weight of $f(0) = f(1) = \gamma \simeq 0.74$. The user confidence threshold $\theta_{\text{conf}} = 2$ requires a user to have identified at least $n_u = 8$ species to become **self-validating**. The parameters $\alpha = 0.5$, $\beta = 0.2$ and $\gamma \simeq 0.74$ are used in practice.

The weight function f shown in Fig. 2 is a non-decreasing function that maps the number of identified species n_u to a trust score in the form of:

$$w_u = f(n_u) = n_u^\alpha - n_u^\beta + \gamma, \quad (1)$$

where $\alpha, \beta \in \mathbb{R}_+^*$ are hyperparameters that were calibrated internally to fit prior knowledge and $\gamma > 0$ is the constant representing the initial weight of each user. In practice, we use $\alpha = 0.5$, $\beta = 0.2$ and $\gamma = \log(2.1) \simeq 0.74$ in the weight function. This function is sub-linear ($\mathcal{O}(\sqrt{n_u})$) but with two different behaviors. The goal of Equation (1) is to separate new users from experts and then help sort multiple experts. This is modeled by the two behaviors of the weight function. In the first part which corresponds to new users with low n_u , the term in the power of β decreases the weight. We chose an initial weight $w_u = \gamma$ such that a user has a weight equal to 1 (rounding to two decimals) with

two different identifications. This separates the users who only come once to test the application from others. In the second part with a higher number of identified species, the term to the power of β becomes negligible and we tend to the square root function. The sub-linear scale allows for reducing discrepancies between people who have identified a comparable number of species (and thus have presumably comparable expertise). As for the two thresholds that control the level of uncertainty accepted for a given label, they are set to $\theta_{\text{conf}} = 2$ to control the total weight on an observation and $\theta_{\text{acc}} = 0.7$ to control the agreement between users given their expertise.

Users are said **self-validating** when they are trusted enough so that their proposed label single-handedly makes an observation valid ($s_i = 1$). From Algorithm 1, we see that this is verified when their weight w_u is greater than the level θ_{conf} . Indeed, with a single label we obtain $\text{conf}_i(\hat{y}_i) = w_u > \theta_{\text{conf}}$ and $\text{acc}_i(\hat{y}_i) = 1 > \theta_{\text{acc}}$. In practice, this means that an experienced user who has collected enough weight can validate any observation without any other user’s vote. Note that this identification can later be invalidated by other users with enough weight thanks to the accuracy threshold θ_{acc} .

2.3 Evaluation against other aggregation strategies

Existing aggregation strategies. Plant species label aggregation is a challenging task due to the large number of species $K = 11\,425$. Hence, many classical strategies in the label aggregation literature such as Dawid and Skene’s (Dawid and Skene, 1979) and other variations (Passonneau and Carpenter, 2014; Sinha, Rao, and Balasubramanian, 2018) are not applicable as they require estimating a $K \times K$ confusion matrix for each user. For the considered dataset \mathcal{D}_{SWE} , this would result in $11\,425^2 \times 823\,251 \approx 10^{14}$ parameters to estimate. Similar issues occur for other label aggregation strategies (Whitehill et al., 2009; Hovy et al., 2013; Q. Ma and Olshevsky, 2020). We do not consider deep-learning-based crowdsourcing strategies as Rodrigues and Pereira (2018) and Chu, J. Ma, and Wang (2021) or Lefort et al. (2022) as they require training a neural network from crowdsourced labels, but do not output aggregated labels on the training set. In the PI@ntNet application, we need to propose one or multiple species for each observation to users. To overcome these issues, we consider the following label aggregation strategies that can scale with K and the number of users:

- **Majority Vote (MV)** (James, 1998): it selects the most answered label¹ and is the

¹Ties are broken at a random – creating sometimes some variability in the labeling process.

most common aggregation strategy. More formally, given an observation i :

$$\text{MV}(i, \{y_i^u\}_u) = \arg \max_{k \in [K]} \sum_{u \in \mathcal{U}_i} \mathbb{1}(y_i^u = k) .$$

- **Worker Agreement With Aggregate (WAWA)** (Limited, 2021): this strategy, also known as the inter-rater agreement, weights each user by how much they agree with the MV labels on average. More formally, given an observation i :

$$\begin{aligned} \text{WAWA}(i, \mathcal{D}_{\text{SWE}}) &= \arg \max_{k \in [K]} \sum_{u \in \mathcal{U}_i} w_u \mathbb{1}(y_i^u = k) \\ \text{with } w_u &= \frac{1}{|\{y_{i'}^u\}_{i'}|} \sum_{i'=1}^{n_{\text{SWE}}} \mathbb{1}(y_{i'}^u = \text{MV}(i', \{y_{i'}^u\}_u)) . \end{aligned}$$

As there is no observation filter for the MV and WAWA, we consider that for all observation i , $s_i = 1$ for these two strategies.

- **TwoThird**: The TwoThird aggregation generates a label for observations with at least two votes. The estimated label represents the one with at least two-thirds of the majority in agreement. Every user has the same weight in the aggregation. It is part of the iNaturalist’s label aggregation system (Van Horn et al., 2018). More formally:

$$\begin{aligned} \text{TwoThird}(i, \{y_i^u\}_u) &= \begin{cases} \text{MV}(i, \{y_i^u\}_u) & \text{if } s_i = 1 \\ \text{undefined} & \text{otherwise} \end{cases} \\ \text{with } s_i &= \mathbb{1} \left(\max_{k \in [K]} \frac{1}{|\mathcal{U}_i|} \sum_{u \in \mathcal{U}_i} \mathbb{1}(y_i^u = k) \geq \frac{2}{3} \right) . \end{aligned}$$

Creation of an evaluation set in a crowdsourcing setting. To evaluate the performance of a label aggregation strategy, it is necessary to know the ground truth on a subset of the data. However, in the context of crowdsourced data, there is no known truth for the observations. The sheer volume of data makes it impossible to ask botanical experts to create such ground truth for the whole database. Moreover, identifying species from images is much less accurate than identifying them in the field, due to the partial information contained in the image (Experts, 2018).

Instead of asking experts to label a subset of the data, we rather identify botanical experts in the PI@ntNet user database and consider their determinations as ground truth. We asked botanical-known experts to reference other experts who could have a PI@ntNet

account to create a list of expert users. To these we have added TelaBotanica (Heaton, Millerand, and Proulx, 2010) users with registered confirmed botanical experience from their account and that are also Pl@ntNet users that participated in the South-Western Europe flora subset. In total, 98 Pl@ntNet users were identified as botanical experts. Observation with at least one vote from one of these experts constitute our test set denoted $\mathcal{D}_{\text{expert}}$. The answers of these experts are considered ground truth labels and used to evaluate strategies’ performance. Despite our selection process of supposedly ‘indisputable’ experts, a few observations in the test set still end up with contradictory labels (4 observations in total). As they represent a very small percentage, we simply removed them from $\mathcal{D}_{\text{expert}}$.

Our evaluation set $\mathcal{D}_{\text{expert}}$ is finally composed of 26 811 observations which received at least one vote from one of the experts. Despite the large number of users, not all observations obtain multiple annotations. Indeed, 310 564 users were single-time voters (meaning they interacted with the system only once). The lack of votes is a large component of difficulty in the Pl@ntNet database, as there is a high imbalance of the distribution of votes between observations as represented in Fig. 4b. There is a high concentration of votes for a small percentage of the observations as shown in Fig. 4a. Of these evaluation data, 17 125 received more than two identifications and are stored in $\mathcal{D}_{\text{multiple votes}}$. Then, 1 263 have more than two votes with at least one disagreement between users and are stored in $\mathcal{D}_{\text{disagreement}}$. Fig. 3 shows the distribution of observations from \mathcal{D}_{SWE} to the finer and more ambiguous $\mathcal{D}_{\text{disagreement}}$.

Evaluation metric. To evaluate the label aggregation strategies, we use the following label recovery accuracy computed on the evaluation datasets:

$$\text{Acc}(\hat{y}, y; \mathcal{D}_{\bullet}) = \frac{1}{n_{\bullet}} \sum_{i=1}^{n_{\bullet}} \mathbb{1}(\hat{y}_i = y_i) \mathbb{1}(s_i = 1) ,$$

with $\hat{y} = (\hat{y}_i)_i$ the estimated labels on $\mathcal{D}_{\bullet} \in \{\mathcal{D}_{\text{expert}}, \mathcal{D}_{\text{multiple votes}}, \mathcal{D}_{\text{disagreement}}\}$, $y = (y_i)_i$ the associated experts labels, considered as ground truth. When the aggregation strategy indicates the observation as invalid ($s_i = 0$ for Pl@ntNet and TwoThird), this metric considers the sample as wrongly classified. Precision and recall scores are also computed to respectively measure the correctness of the observations indicated as valid and the ability to recover the ground truth observations in the valid set. We take into account the species imbalance by using a macro-average for these metrics. This treats rare species as

PI@ntnet South-Western Europe flora dataset

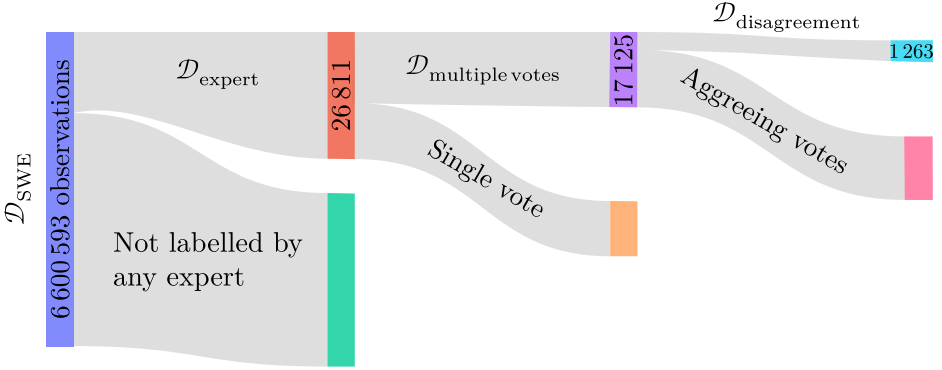


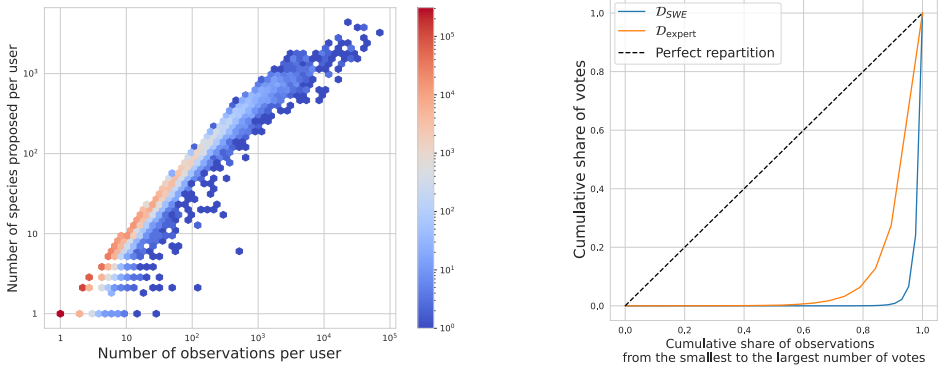
FIG. 3. Log-scales distribution of the observations in the South-West European Flora subset from the PI@ntNet database. Note that the (sub-)datasets introduced are nested: $\mathcal{D}_{\text{SWE}} \supset \mathcal{D}_{\text{expert}} \supset \mathcal{D}_{\text{multiple votes}} \supset \mathcal{D}_{\text{disagreement}}$. $\mathcal{D}_{\text{expert}}$ and the following subsets contain observations that received at least one vote from one of the experts.

equally important to common ones. Denoting respectively TP_k , FP_k and FN_k the true positives, false positives and false negatives related to the species k , the macro averaged precision and recall write

$$\text{Precision}_{\text{macro}} = \frac{1}{K} \sum_{k=1}^K \frac{\text{TP}_k}{\text{TP}_k + \text{FP}_k} \quad \text{and} \quad \text{Recall}_{\text{macro}} = \frac{1}{K} \sum_{k=1}^K \frac{\text{TP}_k}{\text{TP}_k + \text{FN}_k} .$$

As both the PI@ntNet and the TwoThird strategies can invalidate some of the observations, we also compute the proportion of observations removed from the whole dataset (whereas previous metrics are computed on the evaluation dataset). This complementary metric allows measuring the proportion of samples "lost" for the training of the AI model after the aggregation step. In practice, filtering data might remove some noisy samples from the dataset. Yet, in general, the more samples are filtered, the fewer ones to train the neural network training. Finally, we also consider the proportion of species retrieved by the aggregation strategies on $\mathcal{D}_{\text{expert}}$, $\mathcal{D}_{\text{multiple votes}}$ and $\mathcal{D}_{\text{disagreement}}$. This is a critical consideration because if a species identified by experts is absent from the aggregated data, the neural network trained on this data will be unable to make predictions for that very species.

We evaluate the label recovery Acc of each strategy on $\mathcal{D}_{\text{expert}}$, $\mathcal{D}_{\text{multiple votes}}$ and $\mathcal{D}_{\text{disagreement}}$ (see also Fig. 3): the test set where experts have provided at least one vote ($\mathcal{D}_{\text{expert}}$), its subset of observations with at least 2 votes and one from an expert



(A) Relationship between the number of observations per user and the variety of species proposed per user. Each point represents a concentration of users in the SWE flora subset. 310,564 users proposed a single vote.

(B) Lorenz curves representing the imbalance distribution of the number of votes in the South-West European Flora subset from the PI@ntNet database. This imbalance is mitigated but kept in the created test set.

FIG. 4. PI@ntNet activity summary in the SWE flora subset. (A): The majority of users have proposed a small number of observations and species. However, some users have proposed a large number of observations and species. (B): In a perfectly balanced dataset, the Lorenz curve would be the diagonal – 50% of the votes would be for 50% of the observations. In practice, there is a high imbalance of the distribution of votes between observations – 80% of the observations are represented by 10% of votes.

($\mathcal{D}_{\text{multiple votes}}$) and its subset of observations with at least 2 votes, one from an expert, and one disagreement ($\mathcal{D}_{\text{disagreement}}$). The latter is the most challenging as it contains the observations with the most ambiguity. We selected these subsets to investigate the label aggregation strategies’ performance depending on the ambiguity level.

2.4 Taking into account AI votes

While we restricted ourselves to the SWE subset, PI@ntNet’s data is collected internationally. The more correctly identified observations are added to the training set, the better the prediction of the trained model for end-users. This classifier is trained from valid observations and aggregated labels (see Fig. 1). Note that, in addition to Algorithm 1 and the filter on species names, more pre-processing are implemented for better performance (Affouard et al., 2017), such as additional rejection class (*e.g.* non-plant observations), malformed observations (multiple images of different species in a single observation). At the time of writing, the model in use in PI@ntNet is DINOv2 (Oquab et al., 2024) a transformer-based network. This network is based on contrastive learning (Waida et al., 2023), and

represents similar images as close embedding to learn similar features for similar observations and then uses supervised learning to fine tune the model. Several transformations are performed during training such as data augmentation (Suorong et al., 2023), data standardization and label smoothing (Szegedy et al., 2016). However, note that some observations from \mathcal{D}_{SWE} have been processed by an earlier version of PI@ntNet’s AI: either an InceptionV3 (Szegedy et al., 2015) or a BEIT (Bao et al., 2021) classifier. We can use the classifiers to generate votes. For an observation i , the AI vote is denoted $y_i^{\text{AI}} \in [K]$. The probability output in the classifier’s predicted species is denoted $\mathbb{P}(y_i^{\text{AI}})$.

If we consider the trained model as any other user, denoted as **AI as user**, the same label aggregation strategies as in Section 2.2 are available. However, with the PI@ntNet aggregation algorithm, the AI weight increases drastically and overpowers human users (see Section 3.1). This would mean the next PI@ntNet model is mostly trained on the predictions of the previous one. This defeats the purpose of a cooperative active learning system and the human-AI interaction. It would result in a dangerous feedback loop, and possible mode collapse. Thus, we explore alternative ways of integrating the AI votes in the aggregation algorithm:

- **AI as user:** This is the naive approach we just described. The AI is considered as any other user in the database. The total number of users is thus raised to $n_{\text{user}} + 1$.
- **Fixed weight AI:** Give a fixed weight $w_{\text{AI}} = 1.7 > 0$ to the AI. The weight is below the threshold θ_{conf} so that it can not self-validate its predictions. The confidence writes

$$\text{conf}_i(\hat{y}_i) = \sum_{u \in \mathcal{U}_i} w_u \mathbb{1}(y_i^u = \hat{y}_i) + w_{\text{AI}} \mathbb{1}(y_i^{\text{AI}} = k) . \quad (2)$$

The final estimated label becomes

$$\hat{y}_i = \arg \max_{k \in [K]} \sum_{u \in \mathcal{U}_i} w_u \mathbb{1}(y_i^u = k) + w_{\text{AI}} \mathbb{1}(y_i^{\text{AI}} = k) . \quad (3)$$

- **Invalidating AI:** The AI is considered as a user with a fixed weight and can only participate in invalidating identifications *i.e.* have $s_i = 0$. This translates as the confidence updated as in Equation (2) but the final Weighted MV remains unchanged from Algorithm 1.
- **Confident AI:** The AI is considered a user with a fixed weight and can only participate if the confidence in its prediction $\mathbb{P}(y_i^{\text{AI}})$ is over a threshold $\theta_{\text{score}} \in [0, 1]$.

The confidence writes

$$\text{conf}_i(\hat{y}_i) = \sum_{u \in \mathcal{U}_i} w_u \mathbb{1}(y_i^u = \hat{y}_i) + w_{\text{AI}} \mathbb{1}(y_i^{\text{AI}} = \hat{y}_i, \mathbb{P}(y_i^{\text{AI}}) \geq \theta_{\text{score}}) . \quad (4)$$

The final estimated label becomes

$$\hat{y}_i = \arg \max_{k \in [K]} \sum_{u \in \mathcal{U}_i} w_u \mathbb{1}(y_i^u = k) + w_{\text{AI}} \mathbb{1}(y_i^{\text{AI}} = k, \mathbb{P}(y_i^{\text{AI}}) \geq \theta_{\text{score}}) . \quad (5)$$

On the choice of the AI weight. The AI has a fixed weight $w_{\text{AI}} > 0$ for the **Fixed weight AI**, the **Invalidating AI** and the **Confident AI** strategies. The choice of this weight must meet several constraints. First, we would like to avoid the AI votes to be **self-validating** as it would validate all the AI predictions on a large part of the database, thus we must have $w_{\text{AI}} < \theta_{\text{conf}}$ in Algorithm 1. We also want the AI votes to help clean the database by invalidating some observations from low-weight users (with weight $0 < w_{\text{low}} \leq \theta_{\text{conf}}$). Thus $w_{\text{low}} / (w_{\text{low}} + w_{\text{AI}}) < \theta_{\text{acc}}$. Hence, our constraints read:

$$\begin{cases} w_{\text{AI}} < \theta_{\text{conf}} \\ \frac{w_{\text{low}}}{w_{\text{low}} + w_{\text{AI}}} < \theta_{\text{acc}} \end{cases} . \quad (6)$$

Taking the extreme case where a user becomes self-validating: $w_{\text{low}} = \theta_{\text{conf}}$, we obtain that $w_{\text{AI}} > \theta_{\text{conf}} \left(\frac{1 - \theta_{\text{acc}}}{\theta_{\text{acc}}} \right)$. And using the first condition in Equation (6), we obtain the bounds

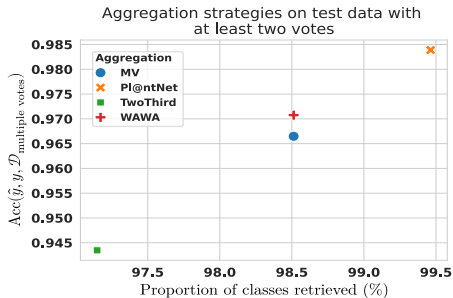
$$\theta_{\text{conf}} \left(\frac{1 - \theta_{\text{acc}}}{\theta_{\text{acc}}} \right) < w_{\text{AI}} < \theta_{\text{conf}} \quad (\iff 0.85 < w_{\text{AI}} < 2) . \quad (7)$$

As more than a million observations from our dataset only have two votes, one way to choose the AI weight is to consider that the AI can invalidate two erroneous non-experts that would both have just enough weights to make the observation valid: $1.95 = w_{\text{low}} < 2$. Then, the AI weight should be greater than their cumulated confidence: $w_{\text{AI}} > 2w_{\text{low}} \left(\frac{1 - \theta_{\text{acc}}}{\theta_{\text{acc}}} \right)$. We finally take the upper rounded value $w_{\text{AI}} = 1.70$ (which satisfies Equation (7)).


3 Results

3.1 Label aggregation performance comparison

Accuracy of the aggregation strategies. We begin by evaluating the accuracy of the label aggregation strategies on the set of observations labeled by experts, $\mathcal{D}_{\text{expert}}$.



(A) Accuracy on $\mathcal{D}_{\text{multiple votes}}$ w.r.t. to the proportion of classes recovered

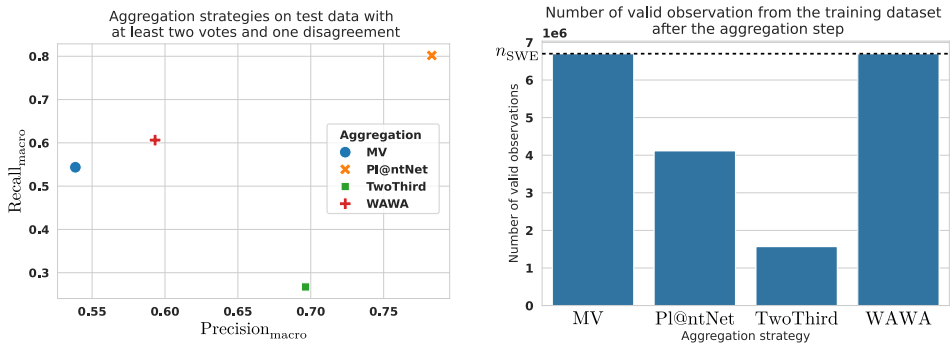


(B) Accuracy on $\mathcal{D}_{\text{disagreement}}$ w.r.t. the proportion of classes recovered

FIG. 5. Accuracy of the aggregation strategies w.r.t. the proportion of classes (species) retrieved on subsets with at least two votes – either agreeing (A) or with at least one disagreeing vote (B). The PI@ntNet aggregation is more accurate, especially in a highly ambiguous setting (B). The TwoThird data filter highly impacts how many classes are kept in the dataset and the overall accuracy in both settings. WAWA and MV perform similarly with a benefit for WAWA when skill evaluation is needed.

Fig. 5 shows how many predicted labels match the experts answers on $\mathcal{D}_{\text{multiple votes}}$ and $\mathcal{D}_{\text{disagreement}}$. More importantly, we compare this quantity with the proportion of species retrieved by the aggregation strategy. We observe that the data filtering from the TwoThird strategy – requiring at least two third of agreements – highly degrades performance with respect to other votes strategies. On $\mathcal{D}_{\text{expert}}$, MV reaches 97% of accuracy, WAWA 98%, TwoThird 60% and PI@ntNet 99%. To differentiate between the best-performing strategies, we need to look at more ambiguous observations like those in $\mathcal{D}_{\text{multiple votes}}$ and $\mathcal{D}_{\text{disagreement}}$. In highly ambiguous frameworks, the WAWA strategy outperforms the MV one. However, overall the PI@ntNet aggregation is more often in adequation with the experts and retrieves almost 90% of plant species identified by experts in highly ambiguous datasets against 73% for WAWA, 71% for MV and only 41% for TwoThird.

Precision and recall. To better evaluate each aggregation strategy, we compute the macro precision and recall metrics for each species. Results are shown in Fig. 6a. The observations filter ($s_i = 0$) for the TwoThird strategy highly impacts its ability to identify most of the positive observations for a given species. While this agreement threshold filter is created to keep as few noisy samples as possible in research-graded (data quality indicator for research database usage in TwoThird) observations, TwoThird obtains better precision than MV and WAWA but PI@ntNet’s precision shows significant improvement. WAWA strategy outperforms a naive MV aggregation showing that, indeed, weighing



(A) Precision and recall of label aggregation strategies on $\mathcal{D}_{\text{disagreement}}$.

(B) Number of observations in \mathcal{D}_{SWE} indicated as valid for training ($s_i = 1$).

FIG. 6. (A): TwoThird strategy has better precision than MV and WAWA strategies, with lower recall because of the heavy filter on the validity of observations. PI@ntNet aggregation strategy obtains best precision and recall and outperforms other strategies. (B): TwoThird performance drop can be explained in part by the high proportion of data considered invalid. Note that MV and WAWA strategies do not invalidate any observation, hence keeping potentially mislabeled or low-quality observations. PI@ntNet achieves a balance between filtering out observations and achieving high performance.

users can lead to better performance. PI@ntNet strategy outperforms all others by several orders of magnitude. Weighing users based on their number of identified species is both interpretable and effective. The observation filter does not negatively impact the recall.

Volume of valid data. The community labels are aggregated to generate training data for the AI model. The more data the better, however, we need to filter out observations with low visual quality or potentially mislabeled. This is the reason for the validity indicator s_i in the TwoThird and PI@ntNet strategies. On \mathcal{D}_{SWE} , Fig. 6b shows how much data is kept for later training. MV and WAWA keep all proposed observation for training – including potential noisy ones. TwoThird filters out most observations to keep nearly 1.5 million (representing 23.43% of the total observations). PI@ntNet finds an improved balance between filtering invalid observations and keeping enough data for training.

Qualitative results on PI@ntNet observation filter. In this section, we show some examples of observations invalidated by the PI@ntNet strategy (see Fig. 7). Invalid observations often come from the lack of user participation with other’s observations. Causes of disagreements from users can occur from a multitude of factors – blurriness, multiple species in the same observation, the distance from the plant does not allow precise iden-



FIG. 7. Examples of invalid ($s_i = 0$) and valid ($s_i = 1$) observations using the PI@ntNet strategy described in Algorithm 1.

tification, *etc.* Valid observations, as shown in the second row of Fig. 7 are zoomed in on the plant’s flower, leaf or organ to help the identification process.

3.2 Aggregation considering AI vote

The current trained neural network model in PI@ntNet’s system can make predictions based on its training on the PI@ntNet database (across different floras). We compare the four following strategies – **AI as user**, **fixed weight AI**, **invalidating AI** and **confident AI**, presented in Section 2.4 to integrate the AI vote into the PI@ntNet label aggregation strategy. For the confident AI strategy, we evaluate multiple thresholds θ_{score} . Note that if $\theta_{\text{score}} = 0$ the AI votes for all observations and if $\theta_{\text{score}} = 1$ the AI does not vote and we recover the performance of the current PI@ntNet aggregation strategy presented in Algorithm 1. We see in Fig. 8 that the confident AI strategy with $\theta_{\text{score}} = 0.7$ seems to perform best and keep the most data in both \mathcal{D}_{SWE} and $\mathcal{D}_{\text{expert}}$.

4 Discussion

We demonstrated that collaborative identification of plant species can effectively be used to obtain expert level labels. Releasing a large subset of millions of observations and thousands of users from the PI@ntNet organization, we investigate a label aggregation strategy that weighs user answers based on their estimated number of species correctly identified without using prior expert knowledge. Many strategies used previously either do not scale to the magnitude of the current databases – either PI@ntNet, iNaturalist or

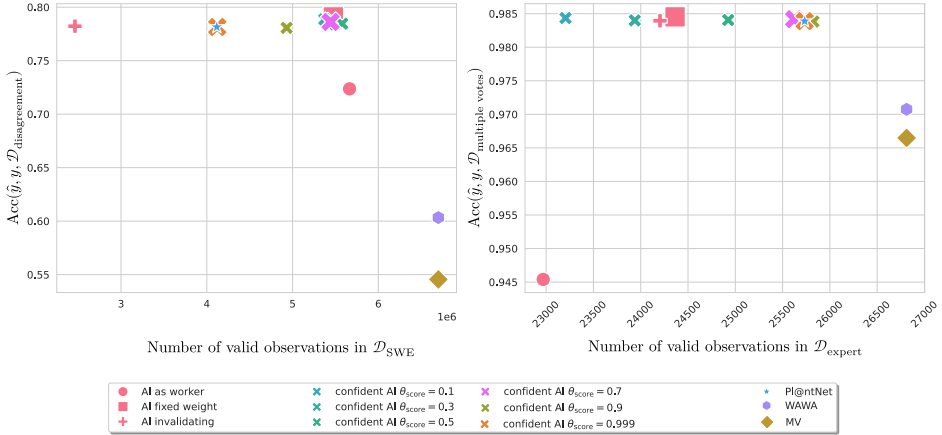


FIG. 8. Performance in label recovery and number of observations marked as valid depending on how the AI vote is integrated. MV, WAWA and PI@ntNet strategy without AI vote are used as reference. The best-performing strategy overall is **confident AI** with $\theta_{\text{score}} = 0.7$. We also see that when θ_{score} tends to 1, we recover the vanilla PI@ntNet aggregation strategy.

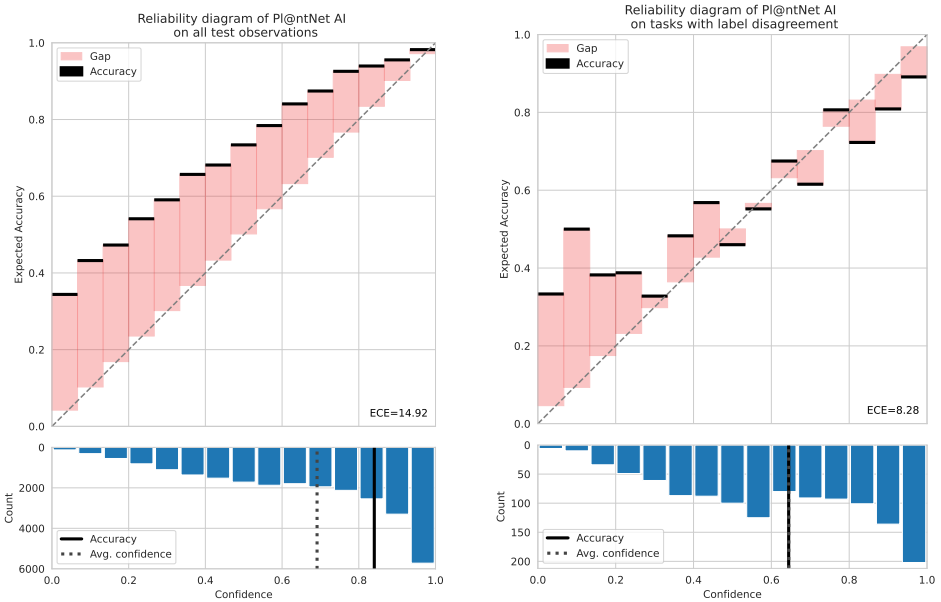
eBird – or are outperformed by our aggregation.

Our strategy weighs users based on the number of correctly identified species. This weight is interpretable and shows the diversity of the user’s skill set. It can be directly applied to other crowdsourced frameworks with a high number of classes like TwoThird or eBird. The values for both hyperparameters θ_{conf} and θ_{acc} which respectively handle the cumulated weight on observation and the agreement level for the given label can be applied as is.

Note that PI@ntNet’s label control system heavily rests on visual analysis of observations and inter-user agreements. Additional metadata such as geolocation, date, phenological stage or visual description can be registered in PI@ntNet and help identify the plant’s species but are currently not directly taken into account for user evaluation. Such information – in particular spatial information – could also be used to generate more interaction between users and collect more votes through possible common interests. In addition, users are helped by the system with images similar to the identification proposed in a given checklist. The additional information could guide users in their vote – for example by notifying a possible incoherence between the current botanical knowledge on a species and the metadata entered (such as the altitude, the distance to the sea, a species not known to survive in a given area).

As for the inclusion of the AI vote, some concerns should be raised. First, as the AI

model is trained from the aggregated labels and observations, integrating its vote should not make the AI predictions run out of control. If we consider the AI as a user, as we are in iterative training, the system fails to learn from the human labels. However, using the AI vote to invalidate the data with a fixed weight can help clean the database, and with enough weight other users can switch its validity back. However, this would not help in switching the wrong label. To do so, we investigate in Section 3.2 to only consider a fixed weight label with enough confidence from the AI model. We observe that this strategy leads to better performance. As we use the output probabilities we should discuss the calibration of our network too.



(A) Reliability diagram of PI@ntNet AI on $\mathcal{D}_{\text{expert}}$

(B) Reliability diagram of PI@ntNet AI on $\mathcal{D}_{\text{disagreement}}$

FIG. 9. Reliability diagrams for the PI@ntNet AI on the expert dataset and on the more ambiguous $\mathcal{D}_{\text{disagreement}}$ subset. The AI is overall underconfident (A). However, on more ambiguous observations it is overconfident for observations leading to high predicted probabilities (B).

Calibration is the measure of how close the output confidence is to the true probability (Niculescu-Mizil and Caruana, 2005). Currently, the PI@ntNet AI is not calibrated using post-processing methods (Platt, 1999; Guo et al., 2017). We discuss hereafter the calibration of the current AI model and possible guidelines for further integration of AI votes.

From Fig. 9a, we see that currently, PI@ntNet AI is underconfident. Meaning that it consistently underestimates its confidence and outputs to users more uncertainty than it should. One factor that can influence the results is that the calibration is computed on the test set where experts either authored or voted on observations. Botanical experts have more experience with taking pictures of plants and better equipment than the average citizen. Thus, the observation quality – and subsequently the probability distribution output by the AI – can be biased. Another factor known for leading to such suboptimal predictions is the data augmentation (Kapoor et al., 2022). As the model trains on multiple versions of each original sample with multiple distortions, these variations can become unrepresentative of the underlying sample distribution and cause unnecessary prediction difficulties. The data augmentation is used to mitigate the species imbalance of the database.

However, this imbalance is also known to lead to miscalibrations in predictions (Ao, Rueger, and Siddharthan, 2023). On Fig. 9b, we see that for ambiguous observations (where users disagree), the AI is overconfident in its highest predictions – which represents half of the dataset – and underconfident in the other half. These different calibration behaviors inform us that, if a given strategy should incorporate the AI votes in the label aggregation based on the output probabilities, we need to be able to rely on such probabilities. Therefore, even if the confident AI strategy leads to the best performance in Section 2.4, it should not be used directly without recalibration of the model – using for example temperature scaling (Guo et al., 2017). In future work, more study is needed to investigate the confidence gap of the model and the observations’ ambiguity from users’ labels. The current large-scaled and interpretable aggregation strategy from PI@ntNet already outperforms others without the AI votes.

Statement on inclusion

We affirm our commitment to promoting diversity and inclusivity in scientific research. Our collected crowdsourced data brings together a wide range of participants. We actively encourage and welcome involvement from individuals of diverse backgrounds, expertise, and perspectives, recognizing the value of their contributions in advancing ecological research and promoting a more comprehensive understanding of plant biodiversity.

References

- Affouard, A. et al. (2017). “PI@ ntnet app in the era of deep learning”. In: *ICLR: International Conference on Learning Representations*.
- Ao, S., S. Rueger, and A. Siddharthan (2023). “Two sides of miscalibration: identifying over and under-confidence prediction for network calibration”. In: *Uncertainty in Artificial Intelligence*. PMLR, pp. 77–87.
- Bao, H. et al. (2021). “Beit: Bert pre-training of image transformers”. In: *arXiv preprint arXiv:2106.08254*.
- Bonnet, P. et al. (2020). “How citizen scientists contribute to monitor protected areas thanks to automatic plant identification tools”. In: *Ecological Solutions and Evidence* 1.2, e12023.
- Borowiec, M. L. et al. (2022). “Deep learning as a tool for ecology and evolution”. In: *Methods in Ecology and Evolution* 13.8, pp. 1640–1660.
- Brown, E. D. and B. K. Williams (2019). “The potential for citizen science to produce reliable and useful information in ecology”. In: *Conservation Biology* 33.3, pp. 561–569.
- Chu, Z., J. Ma, and H. Wang (2021). “Learning from Crowds by Modeling Common Confusions.” In: *AAAI*, pp. 5832–5840.
- Dawid, A. P. and A. M. Skene (1979). “Maximum Likelihood Estimation of Observer Error-Rates Using the EM Algorithm”. In: *J. R. Stat. Soc. Ser. C. Appl. Stat.* 28.1, pp. 20–28.
- Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). “Maximum likelihood from incomplete data via the EM algorithm”. In: *J. R. Stat. Soc. Ser. B Stat. Methodol.* 39.1, pp. 1–22.
- Experts, D. L. T. C. F. (2018). “Plant Identification: Experts vs. Machines in the Era of Deep Learning”. In: *Multimedia Tools and Applications for Environmental & Biodiversity Informatics*, p. 131.
- Govaerts, R. (2023). *The World Checklist of Vascular Plants (WCVP)*. Checklist dataset. Accessed via GBIF.org on 2024-01-30. DOI: 10.15468/6h8ucr.
- Guo, C. et al. (2017). “On calibration of modern neural networks”. In: *ICML*, p. 1321.
- Heaton, L., F. Millerand, and S. Proulx (2010). “ñ Tela Botanica ž : une fertilisation croisée des amateurs et des experts”. In: *Hermès, La Revue* 57.2, pp. 61–68.
- Hovy, D. et al. (2013). “Learning Whom to Trust with MACE”. In: *Proceedings of NAACL-HLT 2013*.
- James, G. (1998). “Majority vote classifiers: theory and applications”. PhD thesis. Stanford University.
- Kapoor, S. et al. (2022). *On Uncertainty, Tempering, and Data Augmentation in Bayesian Classification*. arXiv: 2203.16481 [cs.LG].
- Kelling, S. et al. (Oct. 2015). “Can Observation Skills of Citizen Scientists Be Estimated Using Species Accumulation Curves?” In: *PLOS ONE* 10.10, pp. 1–20.
- Lefort, T. et al. (2022). “Identify ambiguous tasks combining crowdsourced labels by weighting Areas Under the Margin”. In: *arXiv preprint arXiv:2209.15380*.
- Limited, A. (2021). *Calculating Worker Agreement with Aggregate (Wawa)*. URL: <https://success.appen.com/hc/en-us/articles/202703205-Calculating-Worker-Agreement-with-Aggregate-Wawa->.
- Ma, Q. and A. Olshevsky (2020). “Adversarial crowdsourcing through robust rank-one matrix completion”. In: *NeurIPS*. Vol. 33, pp. 21841–21852.

- Mäder, P. et al. (2021). “The flora incognita app—interactive plant species identification”. In: *Methods in Ecology and Evolution* 12.7, pp. 1335–1342.
- Niculescu-Mizil, A. and R. Caruana (2005). “Predicting good probabilities with supervised learning”. In: *Proceedings of the 22nd international conference on Machine learning*, pp. 625–632.
- Oquab, M. et al. (2024). *DINOv2: Learning Robust Visual Features without Supervision*. arXiv: 2304.07193 [cs.CV].
- Passonneau, R. J. and B. Carpenter (2014). “The Benefits of a Model of Annotation”. In: *Transactions of the Association for Computational Linguistics* 2, pp. 311–326.
- Platt, J. (1999). “Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods”. In: *Advances in large margin classifiers* 10.3, pp. 61–74.
- POWO (2024). *Plants of the World Online*. Published on the Internet. Facilitated by the Royal Botanic Gardens, Kew. URL: <http://www.plantsoftheworldonline.org/>.
- Rodrigues, F. and F. Pereira (2018). “Deep learning from crowds”. In: *AAAI*. Vol. 32.
- Sinha, V. B., S. Rao, and V. N. Balasubramanian (2018). “Fast Dawid-Skene: A fast vote aggregation scheme for sentiment classification”. In: *arXiv preprint arXiv:1803.02781*.
- Sullivan, B. et al. (2009). “eBird: A citizen-based bird observation network in the biological sciences”. In: *Biological conservation* 142.10, pp. 2282–2292.
- Suorong, Y. et al. (2023). *Image Data Augmentation for Deep Learning: A Survey*. arXiv: 2204.08610 [cs.CV].
- Szegedy, C. et al. (2015). *Rethinking the Inception Architecture for Computer Vision*. arXiv: 1512.00567 [cs.CV].
- (2016). “Rethinking the inception architecture for computer vision”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826.
- Van Horn, G. et al. (2018). “The inaturalist species classification and detection dataset”. In: *CVPR*, pp. 8769–8778.
- Vidal, M. et al. (2021). “Perspectives on individual animal identification from biology and computer vision”. In: *Integrative and comparative biology* 61.3, pp. 900–916.
- Waida, H. et al. (2023). *Towards Understanding the Mechanism of Contrastive Learning via Similarity Structure: A Theoretical Analysis*. arXiv: 2304.00395 [cs.LG].
- Whitehill, J. et al. (2009). “Whose Vote Should Count More: Optimal Integration of Labels from Labelers of Unknown Expertise”. In: *NeurIPS*. Vol. 22. (Visited on 2021).