



OPEN ACCESS

EDITED BY

Ruslan Kalendar,
University of Helsinki, Finland

REVIEWED BY

Pushkar Gole,
University of Delhi, India
Sunil Gc,
North Dakota State University, United States

*CORRESPONDENCE

Ali Beikmohammadi
✉ beikmohammadi@dsv.su.se

[†]These authors have contributed
equally to this work and share
first authorship

RECEIVED 22 April 2025

ACCEPTED 04 July 2025

PUBLISHED 05 August 2025

CITATION

Lapkovskis A, Nefedova N and
Beikmohammadi A (2025) Automatic
fused multimodal deep learning
for plant identification.
Front. Plant Sci. 16:1616020.
doi: 10.3389/fpls.2025.1616020

COPYRIGHT

© 2025 Lapkovskis, Nefedova and
Beikmohammadi. This is an open-access article
distributed under the terms of the [Creative
Commons Attribution License \(CC BY\)](#). The
use, distribution or reproduction in other
forums is permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original publication in
this journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted
which does not comply with these terms.

Automatic fused multimodal deep learning for plant identification

Alfreds Lapkovskis[†], Natalia Nefedova[†] and Ali Beikmohammadi*

Department of Computer and Systems Sciences, Stockholm University, Stockholm, Sweden

Introduction: Plant classification is vital for ecological conservation and agricultural productivity, enhancing our understanding of plant growth dynamics and aiding species preservation. The advent of deep learning (DL) techniques has revolutionized this field by enabling autonomous feature extraction, significantly reducing the dependence on manual expertise. However, conventional DL models often rely solely on single data sources, failing to capture the full biological diversity of plant species comprehensively. Recent research has turned to multimodal learning to overcome this limitation by integrating multiple data types, which enriches the representation of plant characteristics. This shift introduces the challenge of determining the optimal point for modality fusion.

Methods: In this paper, we introduce a pioneering multimodal DL-based approach for plant classification with automatic modality fusion. Utilizing the multimodal fusion architecture search, our method integrates images from multiple plant organs—flowers, leaves, fruits, and stems—into a cohesive model. To address the lack of multimodal datasets, we contributed Multimodal-PlantCLEF, a restructured version of the PlantCLEF2015 dataset tailored for multimodal tasks.

Results: Our method achieves 82.61% accuracy on 979 classes of Multimodal-PlantCLEF, outperforming late fusion by 10.33%. Through the incorporation of multimodal dropout, our approach demonstrates strong robustness to missing modalities. We validate our model against established benchmarks using standard performance metrics and McNemar's test, further underscoring its superiority.

Discussion: The proposed model surpasses state-of-the-art methods, highlighting the effectiveness of multimodality and an optimal fusion strategy. Our findings open a promising direction in future plant classification research.

KEYWORDS

plant identification, plant phenotyping, multimodal learning, fusion automation, multimodal fusion, architecture search, neural architecture search, multimodal dataset

1 Introduction

Plant classification is among the most significant tasks for agriculture and ecology. It facilitates the preservation of plant species and enhances understanding of their growth dynamics, thus protecting the environment (Zhang et al., 2012). Typically, plants can be categorized into additional specific groupings, such as weeds, invasive species, or plants exhibiting diseases and conditions. According to Dyrmann et al. (2016), between 23 and 71 percent of yield can be lost due to uncontrolled weeds, highlighting the necessity of accurately understanding weed species for more precise herbicide application. Therefore, accurate plant identification is crucial in preventing crop losses and avoiding inaccurate and unnecessary pesticide usage (Meshram et al., 2021).

Manual plant identification typically relies on leaf and flower features and demands profound domain expertise, alongside substantial allocation of time and financial resources (Beikmohammadi et al., 2022). Moreover, Saleem et al. (2018) suggest that the extensive diversity of plant species amplifies the complexity of laboratory classification. Given these challenges, there has been a shift toward automated methods. These approaches, leveraging machine learning (ML) and computer vision (Wäldchen et al., 2018), aim to predict plant types and minimize the reliance on manual skill and resources.

In this regard, some studies have explored traditional ML algorithms for plant classification (Gao et al., 2018; Saleem et al., 2018; Srivastava and Kapil, 2024; Dubey et al., 2025; Yang et al., 2025). However, these algorithms rely on the creation of hand-crafted features, a process heavily dependent on human expertise (Haichen et al., 2021). This human involvement also introduces the potential for biased assumptions and poses challenges in manually identifying suitable features for visual classification. For instance, a classifier based on leaf teeth proves ineffective for species lacking prominent leaf teeth (Wäldchen et al., 2018). Similarly, classifiers relying on leaf contours struggle with species exhibiting similar leaf shapes (Liu et al., 2016).

Recognizing these challenges, numerous studies indicate the superior performance of deep learning (DL) models compared to traditional ones (Nhan et al., 2020; Kolhar and Jagtap, 2021; Kaya et al., 2019). Consequently, researchers have recently adopted DL techniques to develop more effective models for plant identification (Beikmohammadi and Faez, 2018; Beikmohammadi et al., 2022; Espejo-Garcia et al., 2020; Ghazi et al., 2017; Ghosh et al., 2022; Kaya et al., 2019; Lee et al., 2018; Liu et al., 2022; Tan et al., 2020; Dalvi et al., 2024; Diwedi et al., 2024; Kavitha et al., 2023; Kiran and Kaur, 2025; Sharma and Vardhan, 2025; Gowthaman and Das, 2025). Such models, typically employing convolutional neural networks (CNNs), can extract features themselves without the need for explicit feature engineering (Nhan et al., 2020; Saleem et al., 2018; Wäldchen et al., 2018). However, DL models introduce

challenges in engineering model architectures, a task that demands expertise and extensive experimentation and is susceptible to errors (Elsken et al., 2019). Addressing this, neural architecture search (NAS) automates the design of neural architectures. NAS has demonstrated remarkable performance, surpassing manually designed architectures across various ML tasks, notably in image classification (Liu et al., 2021). These methods have also been effectively applied in plant identification (Umamageswari et al., 2023; Sun et al., 2022).

However, typically, both traditional ML and DL models developed for plant classification tasks are constrained to a single data source, often leaf or whole plant images. From a biological standpoint, a single organ is insufficient for classification (Nhan et al., 2020), as variations in appearance can occur within the same species due to various factors, while different species may exhibit similar features. Moreover, using a whole plant image is insufficient, as different organs vary in scale, and capturing all their details in a single image is impractical (Wäldchen et al., 2018). In response to this limitation, very recent studies have delved into the application of multimodal learning techniques (de Lutio et al., 2021; Liu et al., 2016; Nhan et al., 2020; Salve et al., 2018; Hoang Trong et al., 2020; Wang et al., 2022; Zhou et al., 2021), which integrate diverse data sources to provide a comprehensive representation of phenomena. Particularly, Nhan et al. (2020) illustrate that leveraging images from multiple plant organs outperforms reliance on a single organ, in line with botanical insights (Goëau et al., 2015). Wäldchen et al. (2018) underscore the emerging trend of multi-organ-based plant identification, indicating promising accuracy improvements due to the diverse plant viewpoints.

In multimodal learning, the fusion of modalities is recognized as a critical challenge (Barua et al., 2023; Zhang et al., 2020). Various fusion strategies outlined in the literature include early, intermediate, late, and hybrid fusions (Boulahia et al., 2021). Early fusion integrates modalities before feature extraction, such as combining multiple 2D images into a single tensor. Intermediate fusion extracts features from each modality separately and then merges them, offering deeper insights. Late fusion combines modalities at the decision level, often through averaging. The hybrid approach mixes these strategies for optimal results.

Among these fusion strategies, late fusion emerges as the most prevalent in the observed plant classification literature, presumably due to its simplicity and adaptability (Baltrusaitis et al., 2019). However, the choice of a specific strategy relies on the discretion of the model developer (Xu et al., 2021), which can introduce bias and lead to a suboptimal architecture.

To address these issues, we propose an automatic multimodal fusion approach utilizing images from four distinct plant organs—flowers, leaves, fruits, and stems. Following Nhan et al. (2020), we refer to these organs as *modalities*¹. While all organs are represented as RGB images, each encapsulates a unique set of biological features, reflecting the fundamental property of multimodality—*complementarity* (Lahat et al., 2015). Furthermore, unlike multi-view methods (Seeland and Mäder, 2021), our model explicitly

Abbreviations: DL, deep learning; ES, early stopping; IRB, inverted residual block; LR, learning rate; MD, multimodal dropout; MFAS, multimodal fusion architecture search; ML, machine learning; MSE, mean squared error; NAS, neural architecture search; SMBO, sequential model-based optimization.

¹ In this study, we use terms *modality*, *organ*, and *input* interchangeably.

requires distinct plant organs as inputs. Additionally, unlike certain multi-organ methods (Lee et al., 2018), our model has a fixed set of inputs, with each input corresponding exclusively to a specific organ. Thus, we suggest the term *multimodality* aligns more accurately with our approach.

In contrast to previous multimodal and multi-view plant classification studies [e.g., (Ghazi et al., 2017; Lee et al., 2018; de Lutio et al., 2021)], our model is constructed automatically, enabling the discovery of more optimal and efficient architectures. Additionally, the discovery process itself is accelerated and partially parallelized. While earlier works employing NAS [e.g., (Umamageswari et al., 2023; Sun et al., 2022)] have not focused on the multimodal setting, we adopt an algorithm specifically tailored for multimodal problems. As a result, our approach effectively integrates the strengths of both multimodal modeling and NAS. This leads to a compact model with a significantly smaller parameter count, facilitating deployment on resource-limited devices, such as smartphones. By delivering fast, accurate plant identification directly in the field, our method empowers farmers, ecologists, and citizen scientists with actionable insights for agricultural and environmental decision-making.

In addition, existing plant classification datasets are predominantly designed for unimodal tasks, which poses a significant challenge for developing and evaluating multimodal approaches. To address this limitation, we introduce a data preprocessing pipeline that transforms an existing unimodal dataset, namely PlantCLEF2015 (Joly et al., 2015), into a multimodal dataset comprising combinations of plant organ images. This dataset, which we call *Multimodal-PlantCLEF*, supports the development of models with a fixed number of inputs, each corresponding to a specific plant organ.

Our contribution is fourfold:

1. We propose a novel data preprocessing approach to convert an unimodal plant classification dataset into a multimodal one and apply it to transform PlantCLEF2015 into Multimodal-PlantCLEF for use in this study.
2. We propose a novel automatic fused multimodal DL approach for the first time in the context of plant classification. To do so, we first train an unimodal model for each modality using the MobileNetV3Small pre-trained model. Then, we apply a modified multimodal fusion architecture search algorithm (MFAS) (Perez-Rua et al., 2019) to automatically fuse these unimodal models.
3. We evaluate the proposed model against an established baseline in the form of late fusion using an averaging strategy (Baltrušaitis et al., 2019). We utilize standard performance metrics and McNemar's statistical test (Dietterich, 1998). The results demonstrate that our automated fusion approach enables the construction of a more effective multimodal DL model for plant identification, outperforming the baseline.
4. Additionally, we assess the proposed model on all subsets of plant organs, revealing its robustness to missing modalities when trained with a multimodal dropout technique (Cheerla and Gevaert, 2019).

The source code of our work is available online at the GitHub repository².

The remainder of this paper is organized as follows. Section 2 introduces our proposed method, outlines the dataset and data preprocessing, and describes model evaluation. Section 3 presents the obtained results. Section 4 discusses the acquired results in relation to other research. Section 5 concludes the paper and highlights the directions for future research.

2 Materials and methods

2.1 Fusion automation algorithm selection

Different layers of DL models represent different levels of abstraction, and the highest levels are not necessarily the most suitable for fusion (Perez-Rua et al., 2019). Following this insight, to automate the development of a multimodal model architecture, we explore the use of NAS algorithms specifically designed for multimodal frameworks. These algorithms offer a promising solution for automatically identifying the optimal point of fusion rather than relying on manual determination. In this regard, Perez-Rua et al. (2019) introduce MFAS, a multimodal NAS algorithm. Central to their approach is the assumption that each modality possesses a distinct pre-trained model, substantially reducing the search space by maintaining these models static during the search process. The MFAS algorithm iteratively seeks an optimal joint architecture by progressively merging individual pre-trained models at different layers. A notable advantage of this methodology lies in its focus on exclusively training fusion layers, resulting in significant computational time savings.

In another work, Xu et al. (2021) present MUFASA, an advanced multimodal NAS algorithm. One of the standout features of this algorithm is its comprehensive approach, which involves searching for optimal architectures not only for the entire fusion architecture but also for each modality individually, all while considering various fusion strategies. Unlike unimodal NAS methods, MUFASA addresses the whole architecture, leveraging the understanding of its multimodal nature. Furthermore, unlike MFAS, MUFASA addresses the architectures of individual modalities while considering their interdependencies. This unique approach positions MUFASA as potentially more powerful in tackling challenges in multimodal fusion.

While MUFASA demonstrates superior potential, it comes with a notable drawback: its high computational demands. Xu et al. (2021) indicate that achieving state-of-the-art performance on widely used academic datasets would necessitate roughly two CPU years of computational time. In contrast, MFAS has achieved high accuracy, completing within 150 hours of four P100 GPU time on a large-scale image dataset and much faster on simpler datasets (Perez-Rua et al., 2019). This difference can be attributed, at least, to the substantial training requirements of

² GitHub repository: <https://github.com/AlfredsLapkovskis/MultimodalPlantClassifier>

MUFASA, which involves optimizing a larger number of weights and evaluating a greater number of configurations. Therefore, in this study, we select MFAS, as it is considered more suitable for efficiently searching for optimal multimodal fusion architectures.

2.2 Proposed method

The MFAS algorithm employed in this work requires a pre-trained unimodal model for each modality. Therefore, the method is initiated with the creation of these models.

2.2.1 Unimodal models

To construct an unimodal model for each modality represented by different plant organs, we employ a transfer learning technique. There are two primary approaches to applying transfer learning. One involves utilizing pre-trained model weights to extract features from a dataset for subsequent classification, while the other entails full or partial updates of pre-trained weights using a new dataset. The latter is known as fine-tuning (Espejo-Garcia et al., 2020). Initially, we adopted the former approach, pre-training only the appended top layers of each unimodal model. For models where fine-tuning yielded improved performance, we selectively proceeded with this technique.

We explored MobileNetV3Small (Howard et al., 2019), DenseNet121 (Huang et al., 2017), InceptionV3 (Szegedy et al., 2016), and NasNetMobile (Zoph et al., 2018) as candidate architectures for a unimodal base model, and observed that MobileNetV3Small and DenseNet121 demonstrated superior performance; however, MobileNetV3Small also has a low number of parameters. Therefore, we employ MobileNetV3Small as the base model, utilizing weights pre-trained on the ImageNet dataset, for our transfer learning approach. This model is compatible with RGB images and an input size close to $256 \times 256 \times 3$.

2.2.2 Multimodal fusion architecture search

We leverage the MFAS algorithm, which utilizes a pre-trained model $\mathbf{f}^{(i)}(\mathbf{x}^{(i)}) = \hat{y}_i$ for each modality $i \in \{1, \dots, m\}$. Here, \hat{y}_i denotes an approximation of the true label y , derived from the input \mathbf{x} specific to modality i . Each model $\mathbf{f}^{(i)}$ with size n_i is composed of layers $\mathbf{f}_l^{(i)}$, where $l \in \{1, \dots, n_i\}$, such that for l -th layer, $\mathbf{x}_l^{(i)} = (\mathbf{f}_l^{(i)} \circ \mathbf{f}_{l-1}^{(i)} \circ \dots \circ \mathbf{f}_1^{(i)})(\mathbf{x}^{(i)})$ represents features considered for fusion with features from other modalities. The objective of the algorithm is to find optimal combinations of such features to fuse, along with determining the properties of such fusion.

To do so, fusion is performed through another model \mathbf{h} , whose layers are defined as in Equation 1:

$$\mathbf{h}_1 = \sigma_{\gamma_1^a} \left(\mathbf{W}_1 \begin{bmatrix} \mathbf{x}_{\gamma_1^1}^{(1)} \\ \vdots \\ \mathbf{x}_{\gamma_1^m}^{(m)} \end{bmatrix} \right), \mathbf{h}_l = \sigma_{\gamma_l^a} \left(\mathbf{W}_l \begin{bmatrix} \mathbf{x}_{\gamma_l^1}^{(1)} \\ \vdots \\ \mathbf{x}_{\gamma_l^m}^{(m)} \\ \mathbf{h}_{l-1} \end{bmatrix} \right), l > 1 \quad (1)$$

where for l -th layer of \mathbf{h} , $\sigma_{\gamma_l^a}$ denotes an activation function, \mathbf{W}_l is a trainable weight matrix and $\gamma_l = (\gamma_l^1, \dots, \gamma_l^m, \gamma_l^a)$ is a tuple with indices of features from each modality and an index of an activation function. The maximum number of fusion layers \mathbf{h}_l is denoted by $L: l \in \{1, \dots, L\}$, and is a hyperparameter defined prior to execution of the algorithm. A complete fusion configuration of a particular instance of \mathbf{h} with L layers is defined by a vector of tuples $[\gamma_l]_{l \in \{1, \dots, L\}}$, while a set of all possible tuples with L layers is denoted by Γ_L . A list of possible activation functions with size k and possible modality layers used in \mathbf{h}_l are also hyperparameters that can be implementation-specific.

Given this setup, the algorithm spans a large search space of size $(n_1 \times \dots \times n_m \times k)^L$. Since exploring even a small portion of these configurations manually is infeasible, Perez-Rua et al. (2019) have integrated a sequential model-based optimization (SMBO) method into their framework. They argue this approach methodically explores the search space by progressively introducing new configurations, a process which has proven to yield architectures that perform comparably to those identified through direct methods.

Although our approach to fusion automation is grounded in the methodologies outlined in (Perez-Rua et al., 2019), it also incorporates several modifications while maintaining certain similarities. In the following, we provide further elaboration on the specific details of our implementation.

Unimodal models: We utilize the pre-trained MobileNetV3Small model for each modality.

Search space: In this paper, the search space is constrained to the following parameters:

- Maximum of four fusion layers (i.e., $L = 4$).
- Two possible activation functions: ReLU and sigmoid (i.e., $k = 2$).
- 6 fusible layers in each unimodal model: the 1st activation, the 1st, 6th, and 11th inverted residual blocks, and our two appended dense layers (i.e., $n_i = 6, \forall i$).

This results in $6^4 \times 2 = 2592$ different configurations for a single fusion layer and forms a search space of size $2592^4 \approx 4.51 \times 10^{13}$.

Generating model configurations: In this work, the sampled architectures, denoted by \mathcal{S} , are reused across iterations. During the progression from $l = 1$ to $l = L$, newly generated layer configurations are either appended to the sampled configurations or used to replace existing l -th layers in \mathcal{S} .

Building architectures: We construct multimodal architectures based on fusion configurations. Since the outputs of layers from unimodal models may have different numbers of dimensions, we apply global average pooling to each multidimensional output from MobileNetV3Small, similarly to the original approach. Subsequently, the output of each fused layer is concatenated and connected with a dense layer, incorporating an activation function specified in the model configuration. In the case of an intermediate fusion layer ($l > 1$), the previous fusion layer is also connected to the dense layer. Each multimodal model is finalized with a classifier layer employing a softmax activation function.

Weight sharing: In addition to sharing weights between fusion layers that have the same indices and identical weight matrix sizes, our approach also takes into account the activation functions of the fusion layers. This adjustment acknowledges the significant impact that activation functions can have on the behavior of weights. In our model, weights are shared across all configurations.

Storing results: Similarly to the original algorithm, if the same configuration is visited twice, the best result is retained.

Surrogate: We utilize the surrogate described in Perez-Rua et al. (2018), since it has proven to be effective in this context³. This surrogate comprises an embedding layer with zero masking, yielding vectors with a length of 100; an LSTM layer with 100 neurons; and a regression layer with a single neuron and sigmoid activation. The model is compiled with an Adam optimizer with a learning rate (LR) of 0.001, and mean squared error (MSE) loss. Each update is executed for 50 epochs with a batch size of 64.

In this study, we create batches for the surrogate by applying right zero padding to configurations. Moreover, our approach slightly diverges from the original by updating the surrogate each time with the entire dataset of results.

Temperature-based sampling: Following Perez-Rua et al. (2018)⁴, the probability of sampling an architecture i with a score a_i is $p_i = a_i / \sum_j a_j$. However, when incorporating a temperature factor t , the probability of sampling i is given by $p_i^{1/t} / \sum_j p_j^{1/t}$. Thus, the larger the t , the more stochastic the sampling becomes. Similar to Perez-Rua et al. (2019), we employ inverse exponential temperature scheduling, as it has shown to perform well in this context too. Therefore, at each step s of the MFAS algorithm, the temperature is computed as $(t_{\max} - t_{\min})e^{-(s/d)^2} + t_{\min}$, where d is the decay rate.

2.3 Dataset

We selected the PlantCLEF2015 dataset (Joly et al., 2015) for our proposed plant classification model because it contains images of different plant organs, including flowers, leaves, fruits, and stems. The dataset provides a large and diverse collection that ensures a sufficient quantity of images for each organ type as depicted in Figure 1a.

This dataset follows a long-tailed distribution, as depicted in Figure 1b. This class imbalance in the data underscores the necessity of employing additional techniques to mitigate its potentially adverse effects on model learning. Furthermore, Figure 2 illustrates the distribution of organ images across classes. Notably, a considerable number of classes entirely lack certain organs. This observation highlights the necessity for models to handle missing modalities effectively, using tailored training approaches.

Moreover, it is evident that flowers, being a discriminative organ (Nhan et al., 2020), are available for most classes, followed by leaves and fruits, although with a noticeable scarcity in many

classes. This scarcity is even more pronounced in stems, being the least discriminative and the least available organ across classes.

Finally, the images in the dataset are derived from observations—individual plant specimens. It is essential to use this information to split the data for training, validation, and evaluation based on observations rather than images, in order to prevent potential bias from exposing similar images of the same observation across different data splits. However, achieving well-balanced splits in this context is non-trivial, as each observation can have an arbitrary number of images, as illustrated in Figure 3. Our approach to data splitting under these circumstances is detailed in Section 2.3.1.

2.3.1 Data preprocessing

The original PlantCLEF2015 dataset and its associated evaluation protocols are not suitable for the development of our proposed multimodal model. These protocols are designed for models that accept either a single image or an observation comprising an arbitrary number of images from various plant organs. In contrast, our model requires a fixed input format consisting of four specific plant organ images. This fundamental mismatch necessitates a custom strategy to reorganize the original dataset into fixed combinations of plant organ images, which inherently precludes the use of standard PlantCLEF2015 evaluation protocols. Given these constraints, modifying the dataset is both a justified and essential step. It is important to note that all multimodal approaches deal with this issue in plant identification. For more information, see Section 4.4. In the following, we detail our preprocessing pipeline for transforming PlantCLEF2015 into Multimodal-PlantCLEF. The preprocessing steps are visualized in Figure 4.

1. Combining the original splits: The original data splits of PlantCLEF2015 are poorly balanced with respect to plant organs; therefore, we merge them into a unified dataset to generate our own splits subsequently.

2. Collecting metadata: Initially, we extract image metadata for flowers, leaves, fruits, and stems from the dataset, resulting in a total of 67812 records.

3. Grouping class observations: We group the metadata of different organ images by their observation identifiers for each class separately.

4. Filtering organ images and classes: We eliminate specific organ images from a class if their total count in this class is less than 3, as such images cannot be split between training, validation, and test sets. Subsequently, we remove any observations that are left empty. Finally, we discard classes with fewer than 3 observations, since we split the data by observations. This filtering results in 919 classes for flowers, 726 for leaves, 500 for fruits, and 314 for stems, yielding a total of 979 unique classes.

5. Splitting the data: We split the data by observations; however, achieving balanced splits is challenging as (i) we aim to ensure a specific proportion of observations in each data split (i.e., 60%:20%:20% for training, validation, and test sets respectively) to maximize the diversity of data within them, and, simultaneously, (ii) we seek to ensure the same proportions of each plant organ's

³ See Perez-Rua et al. (2018) for a complete specification of the embedding layer, LSTM configuration, training regime, and loss settings.

⁴ See Perez-Rua et al. (2018) for a thorough derivation of the temperature schedule and its impact on sampling stochasticity.

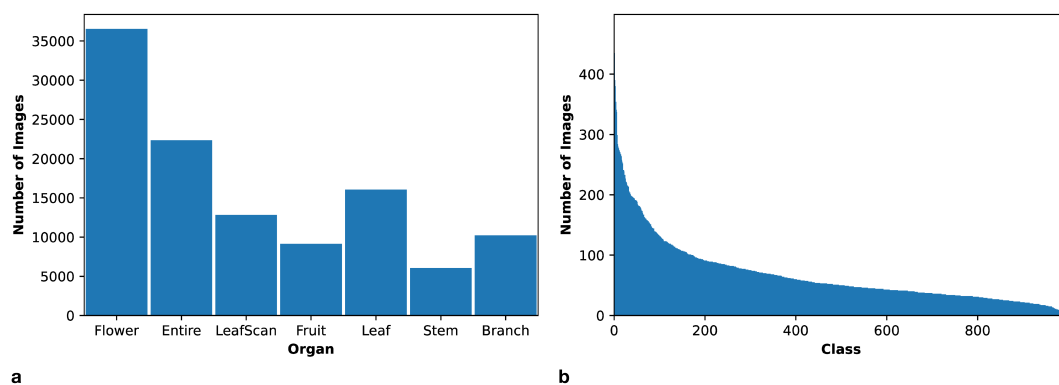


FIGURE 1

PlantCLEF2015 dataset summary. (a) Distribution of plant organs in the dataset. (b) Class distribution in the dataset, considering only the relevant organs: flowers, leaves, fruits, and stems.

images in the splits to achieve certain split sizes and balance in terms of organs. These objectives are inherently contradictory, as observations contain varying numbers of organ images (Figure 3). Consequently, we frame the splitting process as a constrained optimization problem, which we solve for *each class separately* to ensure stratification by labels:

$$\begin{aligned} & \underset{\mathbf{x}_s}{\text{minimize}} \quad \underbrace{\sum_{s \in \mathcal{S}} (\|\mathbf{x}_s\|_1 - \lambda_s N)^2}_{\text{optimize observation counts}} + \underbrace{\sum_{o \in \mathcal{O}} \sum_{s \in \mathcal{S}} (\mathbf{c}_o^T \mathbf{x}_s - \lambda_s \|\mathbf{c}_o\|_1)^2}_{\text{optimize organ counts}} \quad (2) \\ & \text{subject to} \quad \sum_{s \in \mathcal{S}} \mathbf{x}_s = \mathbf{1} \end{aligned}$$

where N is the number of observations, \mathcal{S} is a set of data splits, \mathcal{O} is the set of organ types, $\mathbf{c}_o \in \mathbb{N}^N$ is a vector containing the counts of organ o images for each observation, $\mathbf{x}_s \in \{0, 1\}^N$ is a decision vector for a split s , where the i -th element is 1 if the i -th observation belongs to split s , and $\lambda_s \in [0, 1]$ is the desired size of split s (0.6 for training and 0.2 for validation and test). Here in Equation 2, the first term of the problem minimizes the divergence between the actual and intended number of observations in each split. The second term minimizes the divergence between the actual and intended counts of each organ's images in each split. The problem's constraint ensures that each observation is placed only in one data split.

After solving this problem, we assign observations to splits according to the resulting vectors \mathbf{x}_s . If an organ split in a particular class ends up empty because all images were assigned to other splits,

we manually transfer a single image from the largest split to ensure that training, validation, and testing can be performed for that class. This correction is practically expected when a class has, for example, 3–4 images of the organ. Overall, this procedure produces reasonably balanced splits, as shown in Figure 5. Table 1 summarizes the total counts of observations and organ images across the splits.

6. Saving unimodal datasets: We save each data split separately for each modality to enable pre-training of unimodal models for MFAS. Each dataset consists of shuffled images and their corresponding labels. We map the labels to values in the range [0,978], as the original label values are impractical for use. We convert each image to RGB, resize to 256×256 resolution, and encode in JPEG format. This standardization ensures uniform model input dimensions and reduces computational and I/O load. We choose a square shape for resizing due to the observation that the average aspect ratio of images in the dataset closely approximates 1. During data loading, we also normalize pixel channel values to the range $[-1, 1]$ to meet the input requirements of MobileNetV3Small.

7. Saving multimodal datasets: In this step, we prepare the data for training the multimodal architecture. To accomplish this, the four modalities must be combined. Each class in the splits encompasses a varying number of images of distinct organs. A straightforward approach would entail generating all possible combinations of *flower* \times *leaf* \times *fruit* \times *stem* for each class. However, given the unequal distribution of organ images across

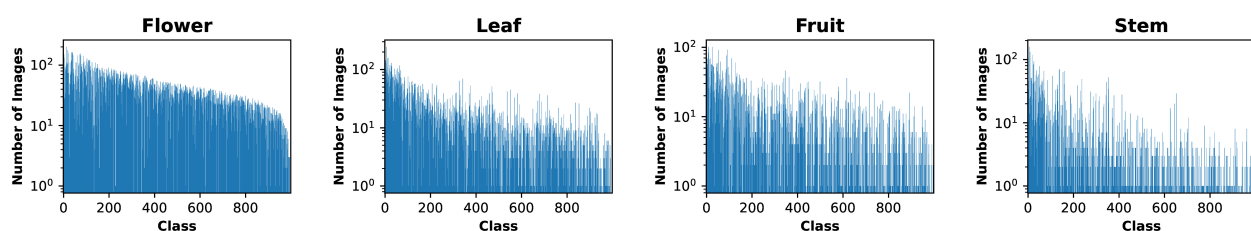


FIGURE 2

Distribution of plant organs among classes in PlantCLEF2015.

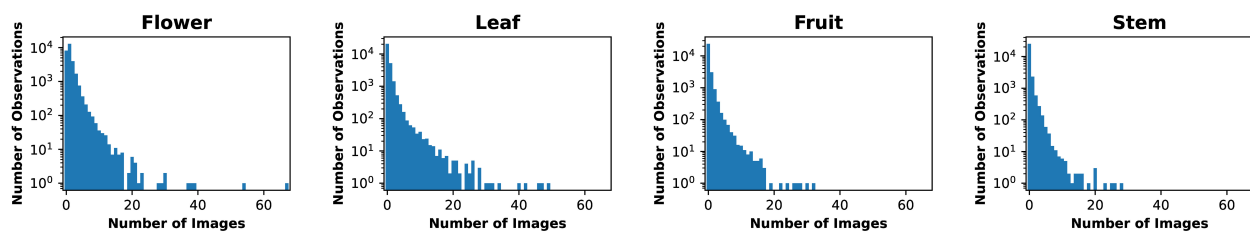


FIGURE 3
Distribution of plant organs among observations in PlantCLEF2015.

classes, this would yield numerous similar tuples, potentially burdening computational resources without significant learning gains. Instead, we only generate $N = \max(n_{\text{flower}}, n_{\text{leaf}}, n_{\text{fruit}}, n_{\text{stem}})$ random image combinations per class, where n_m represents the number of images of modality m in the class. We create random vectors $\mathbf{x}^{(m)}$ for each modality m by permuting sequences of modality images and then repeating their elements to match the length of N . Then, we generate sets of key-value pairs $R_i = \{(m, x_i^{(m)}) | \forall m\}$ for our multimodal dataset, where $i \in \{1, \dots, N\}$. We shuffle these sets and store them together with the corresponding labels to subsequently use them for our multimodal models. If a class entirely lacks a modality m , the respective vector $\mathbf{x}^{(m)}$ is not constructed, and $(m, x_i^{(m)})$ is omitted in records R_i for that class. The formats of stored images and labels follow the specifications in step 6.

2.4 Proposed model setup and configuration

2.4.1 Unimodal model training

We systematically evaluate different training hyperparameters and the impact of fine-tuning on unimodal models by progressively unfreezing inverted residual blocks (IRBs). As a result, each base MobileNetV3Small model is appended with an intermediate dense layer (with batch normalization) and a classifier layer. Initially, we train only these appended layers using the default Adam optimizer with an exponentially decaying LR and early stopping (ES), after which the best weights are restored based on validation loss. Subsequently, we fine-tune from the top layer down to and including the 4th IRB for the leaf and fruit models, and all layers

for the flower model. Fine-tuning did not yield performance improvements for the stem model, so it is omitted there. During fine-tuning, ES is delayed until the 40th epoch. Table 2 provides the complete list of unimodal model hyperparameters.

To minimize the impact of class imbalance, we employ a weighted cross-entropy loss (Equation 3)

$$\mathcal{L} = -\frac{1}{N_b} \sum_{i=1}^{N_b} \sum_{c \in C} w_c y_{i,c} \log(\hat{y}_{i,c}), \quad (3)$$

where N_b is a batch size, C is a set of all classes, $y_{i,c}$ is a binary indicator, $\hat{y}_{i,c}$ is predicted class probability, and w_c is a weight for a class $c \in C$. The indicator $y_{i,c}$ is computed according to Equation 4

$$y_{i,c} = \begin{cases} 1 & \text{if } c \text{ is a true class of } i \\ 0 & \text{otherwise} \end{cases}, \quad (4)$$

and the weight w_c as shown in Equation 5:

$$w_c = \frac{N}{|C|} \times N_c, \quad (5)$$

where N_c is the number of instances of a class c within the training data, and N is the total number of instances.

Additionally, we shuffle the data at each epoch and randomly apply image augmentations that have been shown to enhance performance, as outlined in Table 3.

To fully utilize the available data, we train two versions of each model: one on the training set only, and another on the combined training and validation sets. The first model version is used during MFAS execution (Section 2.4.2) and while tuning the final model (Section 2.4.3). Once hyperparameter tuning is complete, the validation set is no longer required; so, we merge it with the

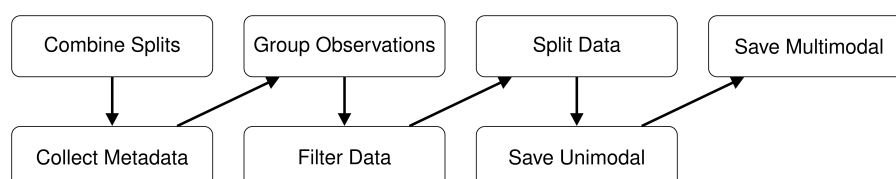
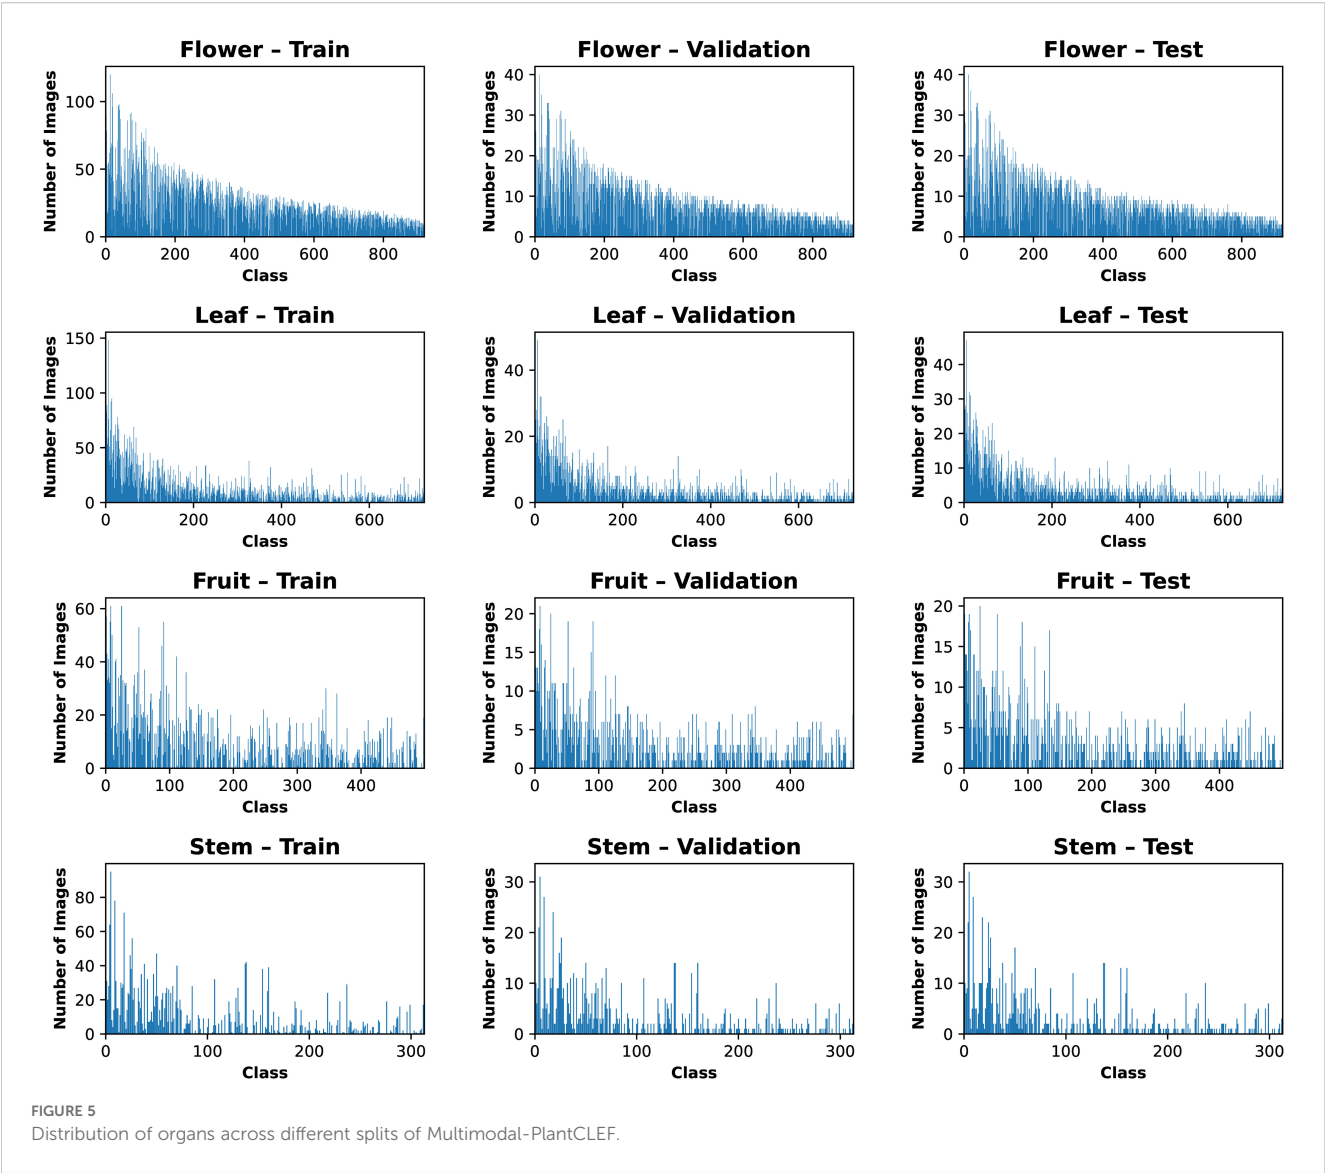


FIGURE 4
Data preprocessing steps to convert PlantCLEF2015 into standalone unimodal datasets and Multimodal-PlantCLEF.



training set and retrain the models on this combined data for evaluation purposes. Without a validation set, we cannot apply ES based on validation performance; therefore, we determine an appropriate number of training epochs based on the first model version, increasing it slightly to account for the additional data. The number of epochs used for retraining on the combined sets is provided in Table 2.

TABLE 1 Counts of observations and organs in Multimodal-PlantCLEF data splits.

Splits:	Training	Validation	Test
Observations	16737 (59.21%)	5788 (20.48%)	5740 (20.31%)
Flowers	21859 (60.09%)	7273 (19.99%)	7243 (19.91%)
Leaves	9357 (59.84%)	3172 (20.29%)	3107 (19.87%)
Fruits	5096 (59.97%)	1678 (19.75%)	1724 (20.29%)
Stems	3196 (59.05%)	1119 (20.68%)	1097 (20.27%)

2.4.2 Multimodal fusion architecture search execution

In the MFAS algorithm procedure, each multimodal architecture is trained using the Adam optimizer with a learning rate of 0.001 and a weighted cross-entropy loss, where class weights are calculated according to Equation 5. The number of neurons per fusion layer is set to 64. Each architecture is trained for 2 epochs with a batch size of 256. As batching significantly impacts training speed, we cache batches and shuffle them each epoch in buffers of 12, rather than creating new random batches each time. Given the limited number of epochs per architecture during the algorithm’s execution, this approach offers a favorable trade-off between maintaining batch randomness and optimizing training speed. Given the imbalanced nature of the dataset, we score the architectures using the $F1_{macro}$ metric on the validation set. We prefer $F1_{macro}$ over $F1_{weighted}$ as it treats classes uniformly and is independent of the number of samples in each class. Due to that, $F1_{macro}$ also allows a more straightforward comparison of performance among different subsets of modalities (Section 2.5.3).

TABLE 2 Hyperparameters of unimodal models.

Models:	Flower	Leaf	Fruit	Stem
Optimizer	Adam	Adam	Adam	Adam
Initial LR	10^{-3}	5^{-4}	5^{-4}	10^{-3}
LR decay rate	0.95	0.95	0.95	0.95
LR decay steps	200	200	200	200
Batch size	256	256	256	256
Epochs	1000	1000	1000	1000
ES patience	10	10	10	10
Classifier L1	10^{-5}	0	0	0
Classifier L2	10^{-5}	10^{-3}	10^{-3}	10^{-2}
Classifier dropout	0.4	0.1	0.2	0.0
Intermediate neurons	1024	1024	512	1024
Intermediate layer L2	0	10^{-5}	0	0
Fine-tuning:				
Fine-tuning	True	True	True	False
Initial LR	10^{-4}	10^{-4}	8^{-5}	—
ES from epoch	40	40	40	—
Fine-tuned layers	All	Up to 4th IRB	Up to 4th IRB	—
Retraining on combined training and validation sets:				
ES	False	False	False	False
Epochs	160	140	120	110
Fine-tuning epochs	160	120	180	—

Hyperparameters remain unchanged during fine-tuning or retraining on combined datasets unless specified otherwise.

Similar to [Perez-Rua et al. \(2019\)](#), we initialize our temperature scheduler with $t_{max} = 10$, $t_{min} = 0.2$, and temperature decay $d = 4$.

In this study, we execute the MFAS algorithm for 5 iterations, with 4 progression levels each, and set the number of sampled architectures to 50. It is important to note that exploring all 2592 initial architectures at the start of the algorithm is time-consuming; therefore, they were evaluated in 20 parallel batches. Subsequently, the algorithm was restarted with the obtained results and a pre-trained surrogate, continuing from the second progression level with the 50 best architectures sampled from the first level.

TABLE 3 Image augmentations for unimodal models.

Models:	Flower	Leaf	Fruit	Stem
Contrast	True	False	True	False
Horizontal flip	True	True	True	True
Vertical flip	False	True	False	False

Contrast adjustment computes a mean μ for each image channel c and sets it for each pixel as $(c - \mu) \times k + \mu$, where $k \sim U(0.75, 1.25)$.

2.4.3 Final model training

To identify the final architecture after executing the MFAS algorithm, we select the 10 most performant configurations and train each for 100 epochs. We employ ES with a patience of 10 and an initial LR of 0.001, which decays exponentially, reaching 95% of its value every 200 steps. We also apply batch normalization after each fusion layer, while retaining the other hyperparameters used during the architecture search.

Subsequently, we carefully tune the hyperparameters of the best architecture. Initially, the optimal number of neurons per layer is investigated. As suggested by [Perez-Rua et al. \(2019\)](#), small weight matrices are utilized during the algorithm to enhance its speed and reduce memory consumption. However, when focusing on a single architecture, this limitation is no longer applicable. Next, we test various configurations of LR, its decay, dropout rates, and other regularization techniques to identify the optimal setup. Additionally, we explore fine-tuning this model to potentially enhance its performance further.

After hyperparameter tuning, we combine the training and validation sets and retrain the model. We also use the corresponding versions of the unimodal models trained on the merged sets. To enhance the model's robustness to missing modalities, we implement the multimodal dropout (MD) technique proposed by [Cheerla and Gevaert \(2019\)](#). Inspired by a regular dropout, MD involves dropping entire modalities during training. In our case, this means replacing the pixel values of an organ image with zeros. This encourages the model to build representations robust to missing modalities. Each modality of a sample is dropped with a similar probability. We utilize a low dropout rate of 0.125 because many classes are completely devoid of certain modalities. Given the challenge of validating this technique during training, we train two versions of the final model: one with MD and one without. Subsequently, we evaluate both.

2.5 Model evaluation

2.5.1 Establishing baseline

We establish the baseline for the proposed model as a late fusion of all unimodal models. This serves as a straightforward approach to fusion. This enables the demonstration that the performance difference between the final model and the baseline is attributed to the unique fusion configuration of the final model.

We implement the late fusion using the averaging strategy ([Baltrusaitis et al., 2019](#)). In this approach, the final prediction for an instance x is calculated according to [Equation 6](#)

$$\hat{c} = \arg \max_{c \in C} \left[\frac{1}{|M_x|} \sum_{m \in M_x} p_m(c|x) \right], \quad (6)$$

where C is a set of all class labels, M_x is a set of unimodal models corresponding to modalities of x , and p_m denotes the probability of c predicted by a model m , given x . Note that in cases where an instance x lacks a modality, the corresponding model is not included in M_x .

2.5.2 Comparison with the baseline

The evaluation of the proposed model relies on standard performance metrics and comparison against the baseline to determine if the automatic fusion setting of the model leads to improvements. Subsequently, we apply McNemar's test (Dietterich, 1998) to ascertain if there exists a statistically significant difference between the two models.

To be more specific, in this study, we utilize accuracy, precision, recall, and F1 score as performance metrics, as specified in Equations 7–10:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}, \quad (7)$$

$$Precision = \frac{TP}{TP + FP}, \quad (8)$$

$$Recall = \frac{TP}{TP + FN}, \quad (9)$$

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall}, \quad (10)$$

where with respect to each class, *TP* (true positives) represents the number of correctly identified instances, *TN* (true negatives) denotes the number of correctly rejected instances, *FP* (false positives) signifies the number of incorrectly identified instances, and *FN* (false negatives) indicates the number of incorrectly rejected instances.

For each class, we compute macro averages of each metric, excluding accuracy. We omit micro averaging since it accumulates all *TP*, *TN*, *FP*, and *FN* values before calculating the metric, resulting in a value equivalent to accuracy. In contrast, macro averaging calculates each metric for individual classes and then averages these values across all classes. The difference between accuracy and macro-averaged metrics provides insight into the performance variation across classes. In conjunction with these metrics, top-5 and top-10 accuracies are also calculated.

2.5.3 Robustness to missing modalities

We collect the same metrics discussed in Section 2.5.2 for individual unimodal models, to enable a comparison between unimodal models and multimodal models. Moreover, we collect the mentioned metrics for the proposed model and the baseline on different subsets of modalities. This allows us to evaluate the robustness of the proposed model in the absence of certain modalities. To evaluate the statistical significance of the differences between these metrics, we apply McNemar's test.

3 Results

This section presents the findings from our comprehensive evaluation of the proposed model. We have structured our results to provide a clear and systematic presentation under various conditions and compared them to an established baseline.

3.1 Performance of unimodal models

After the training procedure, we evaluated the unimodal models. Table 4 illustrates the resulting performance metrics.

According to Table 4: (i) It is evident that the fruit and flower modalities exhibit higher scores across all metrics compared to leaves and especially stems. This is in line with expectations, considering that stems are the least discriminative organs (Nhan et al., 2020). (ii) Accuracy demonstrates higher values compared to macro metrics. This discrepancy should be attributed to differences in performance among classes. The dataset's imbalance has affected the overall performance. (iii) Despite that, similar yet lower precision and recall values suggest that models' performance is consistent on more common classes but poorer on less common. (iv) Relatively high top-*N* accuracies indicate that the models are typically close to identifying the correct class.

3.2 Finding final architecture

As a result of the search for an optimal fusion point, we sampled and trained the 10 best configurations, which are presented in Table 5. Figure 6 illustrates that all architectures achieved approximately 47%–60% in validation $F1_{macro}$ scores. Notably, the second-best architecture identified by the MFAS algorithm (ranked 2nd in Table 5) proved to be the most effective in this test, achieving a validation $F1_{macro}$ of 60.35%. Consequently, we have selected this architecture for further hyperparameter tuning.

3.3 Final model

To train the final model, we adjust various hyperparameters of the second architecture from Table 5. This architecture is visualized in Figure 7, and the full list of chosen hyperparameters is presented in Table 6. Surprisingly, neither full randomization nor augmentations led to improved performance during multimodal model training. As a result, we continue using cached batches shuffled in buffers of 12 for each epoch, as done during the MFAS execution. Additionally, fine-tuning the model did not enhance generalization and is therefore omitted.

With all the aforementioned adjustments, the model reached a validation $F1_{macro}$ score of 68.62%. Figure 8 illustrates the increase

TABLE 4 Performance metrics for unimodal models.

Modalities:	Flower	Leaf	Fruit	Stem
Accuracy	0.6796	0.4294	0.4936	0.2698
Top-5 accuracy	0.8597	0.6508	0.7129	0.4622
Top-10 accuracy	0.9082	0.7293	0.7773	0.5652
Precision _{macro}	0.5998	0.2304	0.3215	0.1298
Recall _{macro}	0.5794	0.2334	0.3444	0.1449
F1 _{macro}	0.5702	0.2181	0.3143	0.1247

TABLE 5 Best multimodal fusion configurations sampled by MFAS.

Nr.	1st fusion	2nd fusion	3rd fusion	4th fusion	F1 _{macro}
1	ReLU(B1,A,A,O)	σ (B11,B1,B6,B6)	σ (B6,B6,B11,A)	ReLU(I,I,I,B6)	0.4610
2	σ (B11,B1,B11,O)	ReLU(O,O,B6,A)	ReLU(O,B11,B6,B11)	ReLU(I,I,O,B1)	0.4561
3	ReLU(I,I,I,B11)	—	—	—	0.4469
4	ReLU(B11,B6,O,B6)	ReLU(O,B11,A,B1)	σ (B6,B6,B11,A)	ReLU(I,I,I,B6)	0.4468
5	ReLU(B11,B6,O,B6)	ReLU(O,B11,A,B1)	σ (B6,B11,B11,B11)	ReLU(I,I,I,B6)	0.4468
6	ReLU(B11,B6,O,B6)	σ (B11,B1,B6,B6)	σ (B6,B6,B11,A)	ReLU(I,I,I,B6)	0.4459
7	ReLU(B1,B1,O,B11)	σ (I,B11,B6,I)	ReLU(B1,B1,I,B11)	ReLU(I,I,A,B6)	0.4434
8	ReLU(B11,B6,O,B6)	ReLU(A,A,O,A)	σ (B6,B6,B11,A)	ReLU(I,I,I,B6)	0.4421
9	σ (B11,B1,B11,O)	ReLU(O,O,B6,A)	ReLU(O,B11,B6,B11)	ReLU(I,O,I,B6)	0.4418
10	ReLU(B6,B6,O,B11)	σ (I,B11,B6,I)	ReLU(B1,B1,I,B11)	ReLU(I,I,A,B6)	0.4387

ReLU and σ denote activation functions of fusion layers, whereas their parameters denote layers fused from flower, leaf, fruit and stem unimodal models, respectively. A – first activation, I – intermediate dense layer, O – output layer, BN – N-th inverted residual block. The F1_{macro} values were achieved by the models after 2 epochs during the search process.

in F1_{macro} for this configuration following hyperparameter tuning. To prepare the model for evaluation, we merge the training and validation sets and train two versions of the model—one with MD and one without it.

3.4 Comparison with the baseline

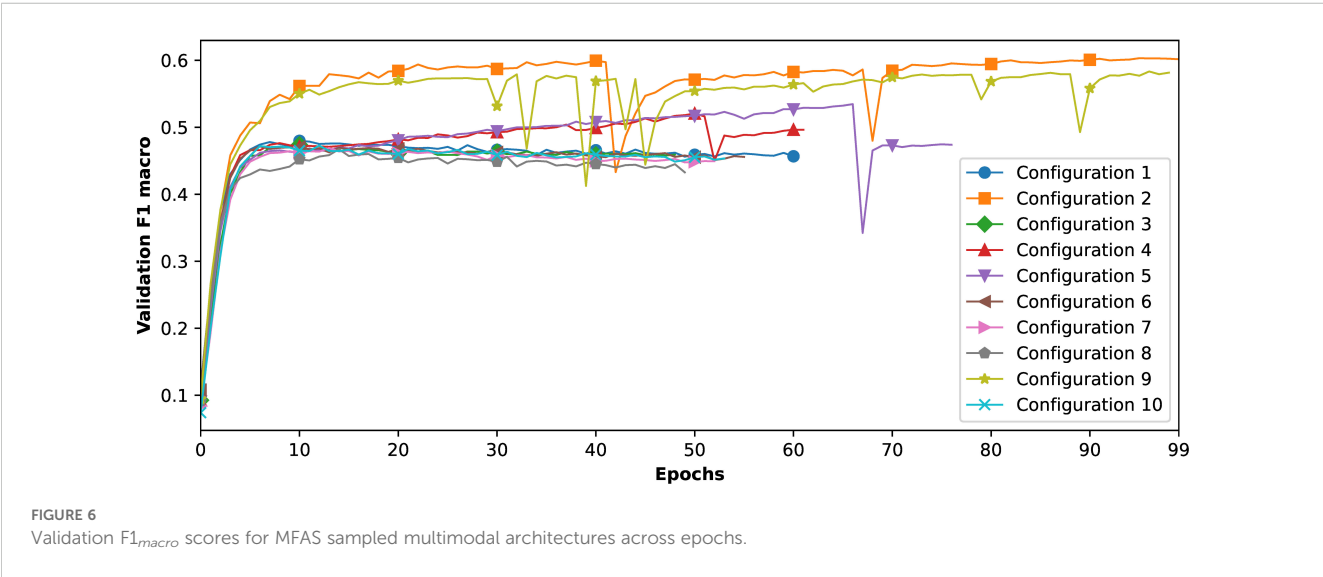
We collect the performance metrics for both versions of the proposed model and the baseline on the test set, which are displayed in Table 7.

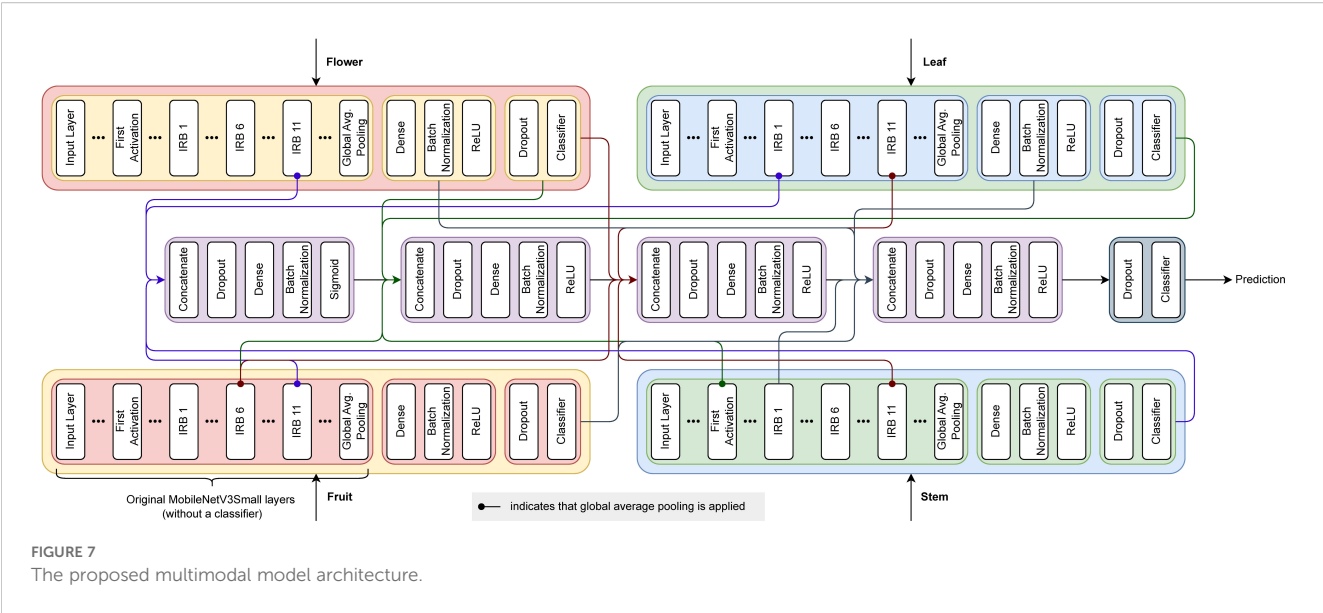
These results indicate that both versions of the proposed model significantly outperform the baseline across all explored metrics, suggesting that our automated fused model is substantially more effective and reliable for plant classification on the given dataset compared to a straightforward late fusion approach. The improvement is particularly notable in the model without MD,

achieving, for instance, +0.1459 in recall and +0.1371 in F1 score. Conversely, a more modest improvement is observed in the model with MD, with, for example, +0.0846 in recall and +0.08 in F1 score. This suggests that, overall, the model without MD generalizes better when evaluated over the entire test set, where modalities in instances are missing only if they are entirely absent in their corresponding classes across all splits. However, in real-world scenarios, the available set of modalities may differ from those represented in certain classes in our dataset. For this reason, we assess the models on subsets of modalities in Section 3.5.

3.4.1 McNemar’s test results

We also conducted McNemar’s tests to compare the performance of two model pairs: (i) the proposed model versus the baseline, and (ii) the proposed model with MD versus the baseline. The tests yielded statistics of $\chi^2_1 = 534.12$ and $\chi^2_1 = 228.74$, both of which are statistically significant ($p < 0.001$). These results





indicate a significant difference in performance metrics between each version of our proposed model and the baseline.

3.5 Comparison on subsets of modalities

One can observe from Tables 4, 7 that both our proposed model and late fusion significantly outperform each individual unimodal model. However, these observations do not fully consider the fact

that the unimodal models are evaluated on isolated unimodal datasets, whereas the multimodal model is evaluated on a multimodal dataset containing missing modalities and duplications. To address this and in order to have a fair comparison, we evaluate both our proposed models and the unimodal models on isolated modalities from the multimodal dataset. Furthermore, while cases where a single or all organs are available represent two extremes, there exists a considerable number of image combinations within the dataset where 2–3 modalities are available. Therefore, it is essential to evaluate the performance of our proposed model and the baseline on all possible subsets of 2–3 modalities. Finally, we assess our models in cases where all modalities are available. Table 8 summarizes the results of these evaluations.

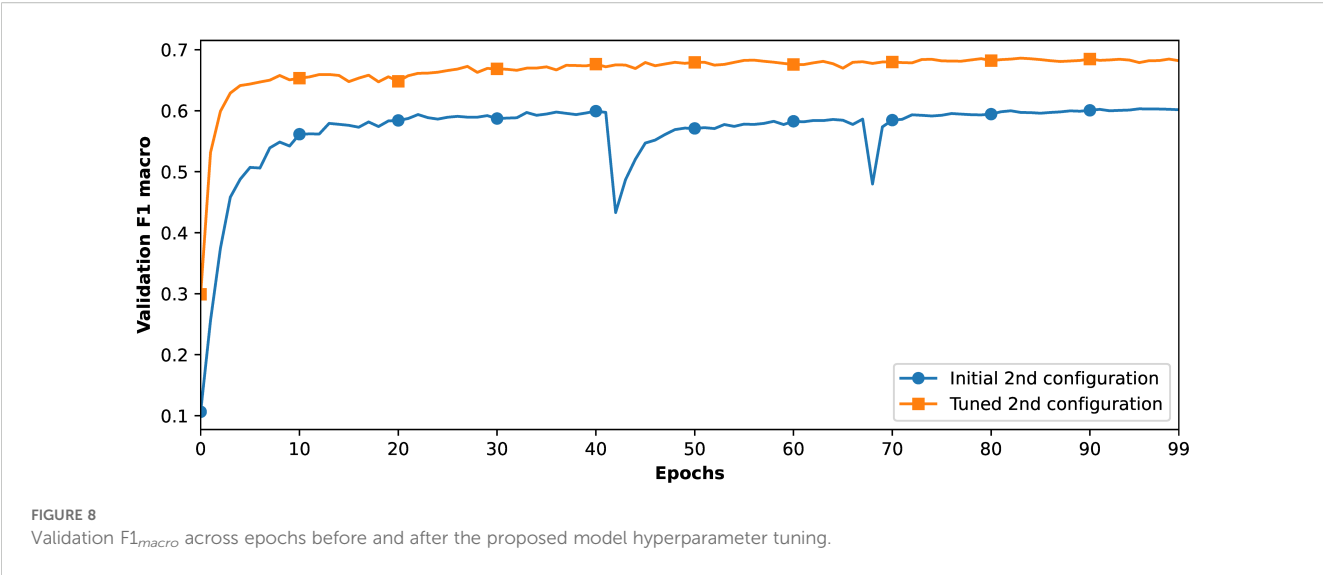
It is important to note that the metrics in Table 8 account for the absence of target modalities by excluding instances that miss any of the required modalities.

The results indicate that the baseline significantly outperforms our proposed model without MD in most cases. The proposed model exhibits superior metrics only in 3 out of 4 scenarios involving 3 modalities and when all modalities are available; however, one of the cases with 3 modalities lacks statistical significance, and another shows lower significance compared to other comparisons. Notably, there are instances where our model achieves a 0% F1 score, specifically in cases involving only fruits and/or stems, while subsets that include leaves demonstrate higher scores, and those incorporating flowers yield significantly better results than any other subsets. The distribution of scores among the unimodal models suggests a different trend, with the informal notation $F1_{flower} > F1_{fruit} > F1_{leaf} > F1_{stem}$ in unimodal models versus $F1_{flower} > F1_{leaf} > F1_{stem} \geq F1_{fruit}$ in the proposed model, implying that the proposed model may be sensitive to the imbalance in modalities, leaning more toward learning from the more common modalities. Consequently, to enhance performance on subsets with less-represented modalities, techniques such as

TABLE 6 Hyperparameters of the proposed model.

Hyperparameter	Value
Optimizer	Adam
Initial LR	5^{-4}
LR decay rate	0.9
LR decay steps	200
Batch size	256
Epochs	100
ES patience	10
Classifier dropout	0.4
Fusion dropouts	[0.0, 0.0, 0.0, 0.4]
Fusion neurons	[512, 512, 512, 512]
MD	0.0 or 0.125
Fine-tuning	False
Retraining on combined training and validation sets:	
ES	False
Epochs	100

Hyperparameters remain unchanged during retraining on combined datasets unless specified otherwise.



augmenting the data with additional instances of underrepresented modalities or applying weighted loss based on the modalities of a predicted instance should be considered. Overall, the findings suggest that the proposed model without MD is more effective when the majority of modalities are available; however, since it consistently trains with the maximum number of available modalities, this approach ultimately leads to reduced robustness when modalities are missing. The model never observes instances with missing modalities, and therefore, cannot generalize well to such cases.

Conversely, the proposed model with MD significantly outperforms the baseline in the majority of cases, indicating the effectiveness of the MD technique in enhancing robustness to missing modalities. Surprisingly, its performance with a single flower modality is very similar to that of the baseline and even better with the leaf modality, despite the baselines, in this case, being standalone unimodal models specifically designed for individual modalities. This observation highlights the model's ability to resist noise from idle modality-specific parts. Notably, this version of our model surpasses the one without MD in all cases except for the scenario where all modalities are available. This suggests that the robustness to missing modalities afforded by MD comes at a slight cost to performance across all modalities.

TABLE 7 Performance metrics for the proposed model and baseline.

Metric	Proposed model	Proposed model with MD	Baseline
Accuracy	0.8261	0.7856	0.7228
Top-5 accuracy	0.9546	0.9219	0.8982
Top-10 accuracy	0.9754	0.9528	0.9420
Precision _{macro}	0.7838	0.7397	0.6652
Recall _{macro}	0.7743	0.7130	0.6284
F1 _{macro}	0.7595	0.7024	0.6224

Similarly, to a regular dropout, MD appears to prevent complex co-adaptations (Hinton et al., 2012) to all available modalities, thus improving generalization to missing modalities. However, the model without MD consistently trains on all available modalities simultaneously, which facilitates the learning of interconnections between them. The model with MD rarely encounters the full set of modalities during training, resulting in reduced performance due to less-learned correlations between modalities.

Overall, our approach outperforms the late fusion of unimodal models. However, if robustness to missing modalities is crucial, it is essential to employ MD or similar techniques. Conversely, to maximize performance when the full set of available modalities is present, the model should consistently train with all the available modalities visible.

4 Discussion

4.1 Accuracy

Our proposed multimodal DL model, which utilizes plant organ images fused through the MFAS algorithm, demonstrates high effectiveness in automating plant classification, achieving an accuracy of 82.61% and significantly surpassing the defined baseline. Table 9 indicates that our model achieves the highest accuracy among similar research.

To be more specific, for example, Ghazi et al. (2017) employed all organs from the PlantCLEF2015 dataset. Along with advanced scoring methods, they achieved an accuracy of 80.18%. Similarly, de Lutio et al. (2021) sampled a large dataset comprising 56608 high-quality images of 977 species from the iNaturalist database. Utilizing images and spatio-temporal context, their multimodal model achieved accuracies of 79.12% without satellite imagery and 79.73% with it. Furthermore, Lee et al. (2018), also using the PlantCLEF2015 dataset, achieved an accuracy of 68.5%. Ge et al. (2016) and Nguyen et al. (2016) extracted flower species from the PlantCLEF2015 dataset and achieved an accuracy of 52.1% and

TABLE 8 Comparison of $F1_{macro}$ for the proposed and baseline models on subsets of modalities.

Modalities	# of Predictions	Proposed Model	Proposed Model with MD	Baseline
Flower	8097	0.3315**	0.5512**	0.5663
Leaf	6714	0.0535**	0.2133*	0.2012
Fruit	4614	0.0000**	0.1259**	0.2894
Stem	3531	0.0000**	0.0374**	0.1203
Flower, Leaf	6325	0.5120**	0.6285**	0.5159
Flower, Fruit	4357	0.2987**	0.5332	0.4574
Flower, Stem	3229	0.3384**	0.5684**	0.3655
Leaf, Fruit	4130	0.0604**	0.3665	0.3512
Leaf, Stem	3420	0.0674**	0.3588**	0.2686
Fruit, Stem	2675	0.0000**	0.2933**	0.3370
Flower, Fruit, Leaf	3887	0.5814	0.6004**	0.4787
Flower, Stem, Leaf	3129	0.5641*	0.6534**	0.4248
Flower, Fruit, Stem	2450	0.4540**	0.6114**	0.4197
Leaf, Fruit, Stem	2611	0.1289**	0.4921**	0.4092
Flower, Leaf, Fruit, Stem	2397	0.7197**	0.6768**	0.4665

* and ** denote statistical significance, with $p < 0.05$ and $p < 0.001$, respectively, as determined by McNemar’s test when compared to the baseline.

67.45% respectively. These results are significantly lower than ours, underscoring the effectiveness of our approach.

Extending the related work, we might observe some other multimodal studies achieving better results; however, this is often due to less complex experimental setups. For instance, Zhang et al. (2012) achieved an accuracy of 93.23%, but on a dataset comprising very small images. Similarly, Liu et al. (2016) achieved top-1, top-5 and top-10 accuracies of 71.8%, 91.2% and 96.4% without geographical data, and 50%, 100% and 100% with it, respectively. However, their dataset included only 50 species, each with 10 leaf and 10 flower images, collected in controlled conditions. Likewise, Salve et al. (2018) demonstrated a high GAR score, but their dataset consisted of only 60 species with 10 images each, all collected in laboratory conditions. Furthermore, Nhan et al. (2020) achieved a maximum 98.8% accuracy by utilizing all PlantCLEF2015 organs and 91.4–98.0% depending on model configurations, on only flowers,

leaves, fruits, and stems, but they sampled only 50 species from the dataset, each containing all organs, thus not experiencing missing modalities as we did. Seeland and Mäder (2021) achieved a maximum accuracy of 94.25% on the PlantCLEF2015 dataset; however, their results were based on a very small selection of images and classes.

In addition to these comparisons, our proposed method demonstrates robustness to missing modalities, as evidenced by its comparison with the baseline.

4.2 Computational cost

While our approach involves much *one-time* computational cost during training to identify an optimal architecture, the resulting model is highly efficient during inference. With only 10 million parameters, our model is lightweight and easily deployable

TABLE 9 Comparison with similar studies.

Study	Dataset	Classes	Modalities	Method	Accuracy
Proposed method	Multimodal-PlantCLEF	979	Flower, leaf, fruit, stem	MFAS using MobileNetV3Small	82.61%
de Lutio et al. (2021)	iNaturalist	977	Plant images, metadata, satellite imagery	ResNet	79.73%
Lee et al. (2018)	PlantCLEF2015	1000	Plant organ	HGO-CNN + Plant-StructNet	68.5%
Ghazi et al. (2017)	PlantCLEF2015	1000	Plant organ	GoogleNet + VGG	80.18%
Nguyen et al. (2016)	Modified PlantCLEF2015	967	Flower	GoogleNet	67.45%
Ge et al. (2016)	Modified PlantCLEF2015	967	Flower	MixDCNN	52.1%

to mobile devices. This parameter count could be further reduced by employing fewer fusion layers if needed.

In comparison, de Lutio et al. (2021) utilize ResNet-50 with 25 million parameters, Lee et al. (2018) design a custom architecture exceeding 50 million parameters, and Ghazi et al. (2017) employ VGG and GoogleNet, which contain 138 million and 6.8 million parameters, respectively. Contrarily, our model is significantly smaller, more efficient, and more suitable for deployment on energy-constrained devices. The reduced parameter count not only lowers memory and storage requirements but also enables faster inference and reduced power consumption, which are critical for edge computing environments such as smartphones. This underscores the importance of multimodality and an optimal fusion strategy.

4.3 Search efficiency

The MFAS algorithm incorporates several optimizations—including surrogate modeling, weight sharing, reduced model complexity, and limiting training to a small number of epochs per configuration—that enable efficient exploration of a large number of architectures. The primary bottleneck in the original procedure occurs during the first iteration, where all generated configurations must be trained to provide data for the surrogate model. To address this, we divided the initial architectures into 20 batches and trained them in parallel.

This approach allowed the algorithm to evaluate up to $2592 + 50 \times 19 = 3542$ unique architectures within a time frame roughly equivalent to training a single multimodal model for $((2592 / 20) + 50 \times 19) \times 2 = 2159.2$ epochs. Even if the first iteration were performed sequentially, the process would involve 7084 epochs—a feasible one-time computational cost for identifying an optimal architecture. Moreover, these requirements could be further reduced by decreasing the number of layers considered for fusion, the number of iterations performed or the number of architectures explored.

In comparison, manual search for an optimal configuration would require investigating a search space of size 4.51×10^{13} , making the MFAS algorithm not only efficient but also practical for large-scale multimodal learning tasks.

In this study, we employed MFAS as it is a well-established multimodal NAS algorithm that satisfied our efficiency requirements when compared to alternative methods [e.g., (Xu et al., 2021)]. By being the first to apply such methods to plant classification, we establish a benchmark for future research in this area. Future studies in plant classification can investigate other multimodal NAS approaches to further enhance training efficiency.

4.4 Limitations

In this study, we modify the PlantCLEF2015 dataset, which is a common practice in the literature (Seeland and Mäder, 2021; Nhan et al., 2020; Nguyen et al., 2016; Ge et al., 2016). This modification was necessary due to the lack of multimodal datasets and the inherent limitations of the original PlantCLEF2015 dataset, which is unsuitable for tasks requiring fixed inputs of specific plant organs.

To address this gap and promote comparability and reproducibility in future research, we contribute a preprocessing pipeline that transforms PlantCLEF2015 into Multimodal-PlantCLEF.

5 Conclusions

This study has addressed a critical task in agriculture and ecology: plant classification. By proposing a novel approach in this domain—a multimodal DL model utilizing four plant organs automatically fused via the MFAS algorithm, the study has demonstrated high performance, outperforming other state-of-the-art models despite the smaller size of the model. This underscores the effectiveness of multimodality and an optimal fusion strategy. Moreover, with the MD technique, the model exhibits robustness to missing modalities, even when only a single modality is available. Additionally, we contributed Multimodal-PlantCLEF, a restructured version of the PlantCLEF2015 dataset tailored for multimodal tasks, to support further research in this area. We believe that this proposed approach opens up a promising direction in plant classification research. This highlights the need for further exploration through the development of more sophisticated algorithms capable of handling larger numbers of species, thereby unlocking the full potential of this approach. Future research could also consider incorporating a multimodal fusion of vision transformers instead of CNNs for plant classification.

Data availability statement

Publicly available datasets were analyzed in this study. The source code of our work is available online at the GitHub repository: <https://github.com/AlfredsLapkovskis/MultimodalPlantClassifier>.

Author contributions

AL: Methodology, Data curation, Writing – review & editing, Conceptualization, Validation, Investigation, Writing – original draft, Formal analysis, Software, Visualization. NN: Visualization, Formal analysis, Data curation, Writing – original draft, Validation, Investigation, Writing – review & editing, Conceptualization, Software, Methodology. AB: Resources, Writing – original draft, Funding acquisition, Project administration, Conceptualization, Supervision, Methodology, Writing – review & editing.

Funding

The author(s) declare that financial support was received for the research and/or publication of this article. The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) at Chalmers Centre for Computational Science and Engineering (C3SE), partially funded by the Swedish Research Council through grant agreement no. 2022-06725.

Acknowledgments

A preliminary version of this manuscript has been made openly available as a preprint on arXiv (Lapkovskis et al., 2024).

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

- Baltrusaitis, T., Ahuja, C., and Morency, L.-P. (2019). Multimodal machine learning: A survey and taxonomy. *IEEE Trans. Pattern Anal. Mach. Intell.* 41, 423–443. doi: 10.1109/TPAMI.2018.2798607
- Barua, A., Ahmed, M. U., and Begum, S. (2023). A systematic literature review on multimodal machine learning: Applications, challenges, gaps and future directions. *IEEE Access* 11, 14804–14831. doi: 10.1109/ACCESS.2023.3243854
- Beikmohammadi, A., and Faez, K. (2018). “). Leaf classification for plant recognition with deep transfer learning,” in *2018 4th Iranian Conference on Signal Processing and Intelligent Systems (ICSPIS)*, (Tehran, Iran: IEEE). p.21–26. doi: 10.1109/ICSPIS.2018.8700547
- Beikmohammadi, A., Faez, K., and Motalebi, A. (2022). SWP-leafNET: A novel multistage approach for plant leaf identification based on deep CNN. *Expert Syst. Appl.* 202, 117470. doi: 10.1016/j.eswa.2022.117470
- Boulahia, S. Y., Amamra, A., Madi, M. R., and Daikh, S. (2021). Early, intermediate and late fusion strategies for robust deep learning-based multimodal action recognition. *Mach. Vision Appl.* 32, 121. doi: 10.1007/s00138-021-01249-8
- Cheerla, A., and Gevaert, O. (2019). Deep learning with multimodal representation for pancreatic prognosis prediction. *Bioinformatics* 35, i446–i454. doi: 10.1093/bioinformatics/btz342
- Dalvi, P., Kalbande, D. R., Rathod, S. S., Dalvi, H., and Agarwal, A. (2024). Multi-attribute deep CNN-based approach for detecting medicinal plants and their use for skin diseases. *IEEE Trans. Artif. Intell.* 6 (3), 710–724. doi: 10.1109/TAI.2024.3491938
- de Lutio, R., She, Y., D’Aronco, S., Russo, S., Brun, P., Wegner, J. D., et al. (2021). Digital taxonomist: Identifying plant species in community scientists’ photographs. *ISPRS J. Photogrammetry Remote Sens.* 182, 112–121. doi: 10.1016/j.isprs.2021.10.002
- Dietterich, T. G. (1998). Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Comput.* 10, 1895–1923. doi: 10.1162/089976698300017197
- Diwedi, H. K., Misra, A., and Tiwari, A. K. (2024). CNN-based medicinal plant identification and classification using optimized SVM. *Multimedia Tools Appl.* 83, 33823–33853. doi: 10.1007/s11042-023-16733-8
- Dubey, A. K., Thanikkal, J. G., Sharma, P., and Shukla, M. K. (2025). A unique morpho-feature extraction algorithm for medicinal plant identification. *Expert Syst.* 42, e13663. doi: 10.1111/essy.13663
- Dyrmann, M., Karstoft, H., and Midtby, H. S. (2016). Plant species classification using deep convolutional neural network. *Biosyst. Eng.* 151, 72–80. doi: 10.1016/j.biosystemseng.2016.08.024
- Elksen, T., Metzen, J. H., and Hutter, F. (2019). Neural architecture search: A survey. *J. Mach. Learn. Res.* 20, 1–21.
- Espejo-Garcia, B., Mylonas, N., Athanasakos, L., Fountas, S., and Vasilakoglou, I. (2020). Towards weeds identification assistance through transfer learning. *Comput. Electron. Agric.* 171, 105306. doi: 10.1016/j.compag.2020.105306
- Gao, J., Nuytens, D., Lootens, P., He, Y., and Pieters, J. G. (2018). Recognising weeds in a maize crop using a random forest machine-learning algorithm and near-infrared snapshot mosaic hyperspectral imagery. *Biosyst. Eng.* 170, 39–50. doi: 10.1016/j.biosystemseng.2018.03.006
- Ge, Z., Bewley, A., McCool, C., Corke, P., Upcroft, B., and Sanderson, C. (2016). “Fine-grained classification via mixture of deep convolutional neural networks,” in *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)* (Lake Placid, NY, USA: IEEE), 1–6.
- Ghazi, M. M., Yanikoglu, B., and Aptoula, E. (2017). Plant identification using deep neural networks via optimization of transfer learning parameters. *Neurocomputing* 235, 228–235. doi: 10.1016/j.neucom.2017.01.018
- Ghosh, S., Singh, A., Jhanjhi, N., Masud, M., and Aljahdali, S. (2022). SVM and KNN based CNN architectures for plant classification. *Computers Materials Continua* 71, 4257–4274. doi: 10.32604/cmc.2022.023414
- Goëau, H., Bonnet, P., and Joly, A. (2015). “LifeCLEF plant identification task 2015,” in *CLEF: Conference and Labs of the Evaluation Forum*, (Toulouse, France: CEUR-WS.org) Vol. 1391. 1–15.
- Gowthaman, S., and Das, A. (2025). Plant leaf identification using feature fusion of wavelet scattering network and CNN with PCA classifier. *IEEE Access*, 13, 11594–11608. doi: 10.1109/ACCESS.2025.3528992
- Haichen, J., Qingrui, C., and Zheng Guang, L. (2021). “Weeds and crops classification using deep convolutional neural network,” in *Proceedings of the 3rd International Conference on Control and Computer Vision*, vol. 20. (Association for Computing Machinery, New York, NY, USA), 40–44. ICCCV. doi: 10.1145/3425577.3425585
- Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. R. (2012). Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*.
- Hoang Trong, V., Gwang-hyun, Y., Thanh Vu, D., and Jin-young, K. (2020). Late fusion of multimodal deep neural networks for weeds classification. *Comput. Electron. Agric.* 175, 105506. doi: 10.1016/j.compag.2020.105506
- Howard, A., Sandler, M., Chu, G., Chen, L.-C., Chen, B., Tan, M., et al. (2019). “Searching for mobilenetv3,” in *Proceedings of the IEEE/CVF international conference on computer vision*. (Seoul, South Korea: IEEE) 1314–1324.
- Huang, G., Liu, Z., van der Maaten, L., and Weinberger, K. Q. (2017). “Densely connected convolutional networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*. (Honolulu, HI, USA: IEEE) 4700–4708.
- Joly, A., Goëau, H., Glotin, H., Spampinato, C., Bonnet, P., Vellinga, W.-P., et al. (2015). “LifeCLEF 2015: multimedia life species identification challenges,” in *Experimental IR Meets Multilinguality, Multimodality, and Interaction: 6th International Conference of the CLEF Association, CLEF’15, Toulouse, France, September 8–11, 2015, Proceedings*, vol. 6. (Cham: Springer), 462–483.
- Kavitha, S., Kumar, T. S., Naresh, E., Kalmani, V. H., Baman, K. D., and Pareek, P. K. (2023). Medicinal plant identification in real-time using deep learning model. *SN Comput. Sci.* 5, 73. doi: 10.1007/s42979-023-02398-5
- Kaya, A., Keceli, A. S., Catal, C., Yalic, H. Y., Temucin, H., and Tekinerdogan, B. (2019). Analysis of transfer learning for deep neural network based plant classification models. *Comput. Electron. Agric.* 158, 20–29. doi: 10.1016/j.compag.2019.01.041
- Kiran, K., and Kaur, A. (2025). “Deep learning-based plant species identification using vgg16 and transfer learning: Enhancing accuracy in biodiversity and agricultural applications,” in *2025 3rd International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT)* (Bengaluru, India: IEEE), 1538–1543.
- Kolhar, S., and Jagtap, J. (2021). Plant trait estimation and classification studies in plant phenotyping using machine vision - a review. *Inf. Process. Agric.* 10, 114–135. doi: 10.1016/j.inpa.2021.02.006
- Lahat, D., Adali, T., and Jutten, C. (2015). Multimodal data fusion: an overview of methods, challenges, and prospects. *Proc. IEEE* 103, 1449–1477. doi: 10.1109/JPROC.2015.2460697

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

Publisher’s note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Lapkovskis, A., Nefedova, N., and Beikmohammadi, A. (2024). Automatic fused multimodal deep learning for plant identification. *arXiv preprint arXiv:2406.01455*.
- Lee, S. H., Chan, C. S., and Remagnino, P. (2018). Multi-organ plant classification based on convolutional and recurrent neural networks. *IEEE Trans. Image Process.* 27, 4287–4301. doi: 10.1109/TIP.2018.2836321
- Liu, J.-C., Chiang, C.-Y., and Chen, S. (2016). “Image-based plant recognition by fusion of multimodal information,” in *2016 10th International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing (IMIS)*, (Fukuoka, Japan: IEEE) 5–11. doi: 10.1109/IMIS.2016.60
- Liu, Y., Sun, Y., Xue, B., Zhang, M., Yen, G. G., and Tan, K. C. (2021). A survey on evolutionary neural architecture search. *IEEE Trans. Neural Networks Learn. Syst.* 34, 550–570. doi: 10.1109/TNNLS.2021.3100554
- Liu, K.-H., Yang, M.-H., Huang, S.-T., and Lin, C. (2022). Plant species classification based on hyperspectral imaging via a lightweight convolutional neural network model. *Front. Plant Sci.* 13. doi: 10.3389/fpls.2022.855660
- Meshram, V., Patil, K., Meshram, V., Hanchate, D., and Ramkteke, S. (2021). Machine learning in agriculture domain: A state-of-art survey. *Artif. Intell. Life Sci.* 1, 100010. doi: 10.1016/j.ailsci.2021.100010
- Nguyen, T. T. N., Le, V., Le, T., Hai, V., Pantuwong, N., and Yagi, Y. (2016). “Flower species identification using deep convolutional neural networks,” in *AUN/SEED-Net Regional Conference for Computer and Information Engineering*. 1–6.
- Nhan, N. T. T., Le, T.-L., and Hai, V. (2020). “Do we need multiple organs for plant identification?,” in *2020 International Conference on Multimedia Analysis and Pattern Recognition (MAPR)*, (Ha Noi, Vietnam: IEEE) 1–6. doi: 10.1109/MAPR49794.2020.9237787
- Perez-Rua, J.-M., Baccouche, M., and Pateux, S. (2018). Efficient progressive neural architecture search. *arXiv preprint arXiv:1808.00391*.
- Perez-Rua, J.-M., Vielzeuf, V., Pateux, S., Baccouche, M., and Jurie, F. (2019). “MFAS: Multimodal fusion architecture search,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, (Long Beach, CA, USA: IEEE) 6959–6968. doi: 10.1109/CVPR.2019.00713
- Saleem, G., Akhtar, M., Ahmed, N., and Qureshi, W. (2018). Automated analysis of visual leaf shape features for plant classification. *Comput. Electron. Agric.* 157, 270–280. doi: 10.1016/j.compag.2018.12.038
- Salve, P., Yannawar, P., and Sardesai, M. M. (2018). Multimodal plant recognition through hybrid feature fusion technique using imaging and non-imaging hyperspectral data. *J. King Saud Univ. - Comput. Inf. Sci.* 34, 1361–1369. doi: 10.1016/j.jksuci.2018.09.018
- Seeland, M., and Mäder, P. (2021). Multi-view classification with convolutional neural networks. *PLoS One* 16, e0245230. doi: 10.1371/journal.pone.0245230
- Sharma, S., and Vardhan, M. (2025). Aelnet: Attention-based enhanced local and global features network for medicinal leaf and plant classification. *Comput. Biol. Med.* 184, 109447. doi: 10.1016/j.combiomed.2024.109447
- Srivastava, V. P., and Kapil, (2024). Semi supervised k-svcr for multi-class classification. *Multimedia Tools Appl.* 84, 1–17. doi: 10.1007/s11042-024-19228-2
- Sun, X., Li, G., Qu, P., Xie, X., Pan, X., and Zhang, W. (2022). Research on plant disease identification based on cnn. *Cogn. Robotics* 2, 155–163. doi: 10.1016/j.cogr.2022.07.001
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. (2016). “Rethinking the inception architecture for computer vision,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*. (Las Vegas, NV, USA: IEEE) 2818–2826.
- Tan, J. W., Chang, S.-W., Abdul-Kareem, S., Yap, H. J., and Yong, K.-T. (2020). Deep learning for plant species classification using leaf vein morphometric. *IEEE/ACM Trans. Comput. Biol. Bioinf.* 17, 82–90. doi: 10.1109/TCBB.2018.2848653
- Umapageswari, A., Bharathiraja, N., and Irene, D. S. (2023). A novel fuzzy c-means based chameleon swarm algorithm for segmentation and progressive neural architecture search for plant disease classification. *ICT Express* 9, 160–167. doi: 10.1016/j.ict.2021.08.019
- Wäldchen, J., Rzanny, M., Seeland, M., and Mäder, P. (2018). Automated plant species identification—trends and future directions. *PLoS Comput. Biol.* 14, 1–19. doi: 10.1371/journal.pcbi.1005993
- Wang, Y., Chen, Y., and Wang, D. (2022). Recognition of multi-modal fusion images with irregular interference. *PeerJ Comput. Sci.* 8, e1018. doi: 10.7717/peerj-cs.1018
- Xu, Z., So, D. R., and Dai, A. M. (2021). “Mufasa: Multimodal fusion architecture search for electronic health records,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, (AAAI) Vol. 35. 10532–10540.
- Yang, C., Lyu, W., Yu, Q., Jiang, Y., and Zheng, Z. (2025). Learning a discriminative region descriptor for fine-grained cultivar identification. *Comput. Electron. Agric.* 229, 109700. doi: 10.1016/j.compag.2024.109700
- Zhang, C., Yang, Z., He, X., and Deng, L. (2020). Multimodal intelligence: Representation learning, information fusion, and applications. *IEEE J. Selected Topics Signal Process.* 14 (3), 1–1. doi: 10.1109/JSTSP.2020.2987728
- Zhang, S.-W., Zhao, M.-R., and Wang, X.-F. (2012). “Plant classification based on multilinear independent component analysis,” in *Advanced Intelligent Computing Theories and Applications. With Aspects of Artificial Intelligence*. Eds. D.-S. Huang, Y. Gan, P. Gupta and M. M. Gromiha (Springer Berlin Heidelberg, Berlin, Heidelberg), 484–490.
- Zhou, J., Li, J., Wang, C., Wu, H., Zhao, C., and Teng, G. (2021). Crop disease identification and interpretation method based on multimodal deep learning. *Comput. Electron. Agric.* 189. doi: 10.1016/j.compag.2021.106408
- Zoph, B., Vasudevan, V., Shlens, J., and Le, Q. V. (2018). “Learning transferable architectures for scalable image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*. (Salt Lake City, UT, USA: IEEE) 8697–8710.