

ORIGINAL RESEARCH

Open access

A Multi-Sensor Agricultural Image Database for UAV and Proximal Platforms: The Eden Library

Juan Perez¹, Ana Gutierrez^{1*}, Carlos Lopez²

Abstract

In contemporary farming, visual perception systems grounded in deep learning are emerging to equip autonomous vehicles with the capability to perform crop-field interventions. Yet securing strong performance from these techniques hinges on access to large volumes of data, a resource that is typically scarce in agricultural contexts. The core reason is that constructing an agricultural dataset that captures a given problem in exhaustive breadth is difficult, since the visual presentation of symptoms is subject to change. There is pronounced sensitivity to ambient variables such as heat, dampness, and lighting regimes. An efficient pipeline is therefore imperative to methodically shepherd the full sequence of building an agricultural dataset, from initial image capture through to its web-facing release. This manuscript describes the Eden Library, a platform conceived to help close the recognized shortfall in freely available crop and plant repositories that span both close-range and overhead imagery. The end-to-end workflow underpinning the platform's conception and launch is also set forth and explored. This workflow merits attention because the datasets on offer are designed not as static, one-off research outputs confined to a narrow slice of species, growth stages, and environmental circumstances, but rather as evolving resources to be stewarded and expanded across time. Annotations delineating plants and their associated symptoms are supplied, thereby relieving users of the manual labeling workload. At the time of writing, the Eden Library spans 15 distinct crop species, 9 weed varieties, and 30 forms of disorder (encompassing pests, diseases, and nutritional shortfalls). The Eden Library endeavors to narrow this divide by delivering a voluminous, richly varied image trove of plants, curated coherently, to accelerate the development of future vision-driven, AI-augmented field tools.

Keywords Precision agriculture, Deep learning, Computer vision, Robotics, Open-data, Dataset

*Correspondence:

Ana Gutierrez
ana.gutierrez@gmail.com

¹ Department of Multi-Sensor Agricultural Image Databases, Faculty of Agriculture, National Autonomous University of Mexico, Mexico City, Mexico

² Department of UAV and Proximal Platform Data Curation, Faculty of Agriculture, Monterrey Institute of Technology, Monterrey, Mexico

Introduction

A sweeping shift is observable throughout modern agriculture toward the embedding of machine-vision capabilities for manifold functions within staple crops, thereby granting autonomous machinery the power to execute field-level tasks [1, 2]. The central hurdles in visualizing plants and cropping systems are closely tied to the crop species at hand, soil conditions, plant structural traits, land fertility and the microclimatic regime of the locality. In turn, these vision-based setups are routinely paired with remote sensing hardware to acquire and transmit data on field conditions and crop developmental timelines. On top of this, reliance on Unoccupied Aerial Vehicles (UAVs) for in-field missions—whether for crop reconnaissance, vegetation indexing, or weed spotting—is intensifying rapidly [3–6]. Despite the meaningful

contributions of these airborne tools, certain shortcomings linger around illumination artifacts and limited ground-resolved detail, which can impair, by way of illustration, their power to cleanly parse plant developmental phases (above all in the earliest periods) or to pin down the coordinates of crops versus weeds for downstream species discrimination.

In contrast, close-proximity sensing rigs, which can be affixed to standard farm implements, tend to return outputs of greater precision and images of elevated crispness relative to aerial platforms [7, 8]. Such terrestrial apparatuses are integrated into a variety of robotic systems for field operations [2]. Terrestrial vision-based recognition frameworks are useful for phenotyping [9], cultivation surveillance [10], and disease diagnosis [11], among other applications.

In earlier eras, the assessment of solitary plants relied on more applied, intuitive approaches grounded in grower know-how—judging leaf outline and size, along with an overall appraisal of plant development. Yet, naked-eye inspection frequently proves sluggish, labor-intensive, and inherently prone to personal bias, a reality that serves as a prime impetus for the quest for alternative methodologies [12–15]. Accordingly, the last few years have witnessed deep learning techniques being pressed into service for machine-based plant recognition, fueled by openly shared plant picture collections [16–44]. Still, bringing deep learning pipelines to a high level of performance requires immense stores of labeled data. This commodity remains in short supply across agriculture, and with even greater acuity in the crop production arena [45]. In truth, compiling an agricultural corpus that covers a focused problem in its full dimensionality is an uphill endeavor. The outward appearance of symptoms may change as a disorder progresses, and appearance can also be influenced by ambient factors such as thermal conditions and moisture levels [46]. An efficient blueprint is thus needed to span the full breadth of possible variation systematically while keeping redundancy in check. A further element to weigh is that every picture must carry an accurate label, an undertaking that is typically taxing and prone to error [47]. The paucity of publicly released corpora spanning a wide range of crops stymies the creation of reference baselines for side-by-side performance evaluation across disparate studies [45]. This stands as one of the principal brakes on the development of dependable deep learning architectures.

Several scholarly works have previously unveiled repositories of plant and crop photographs; the most notable are enumerated in **Table 1** below. A portion of these concentrate on the visual discrimination of detached leaves [16–25], others center on imagery of fruits or vegetables [23, 25, 28, 29, 33, 35, 36, 40, 42, 43], while the remainder direct attention to the recognition of intact plants, crops, or weeds [24, 26, 27, 30–32, 34, 38, 44]. The Plant Village Dataset [16], which has garnered particularly broad recognition, consists exclusively of leaf photographs taken with a camera in a lab. Every shot portrays an isolated leaf resting on a featureless backdrop and bears only a category name as its label. Because of this characteristic, PlantVillage is regularly adopted as a touchstone across numerous publications, a custom that may at times yield distorted impressions when models encounter real field data [12, 13]. A comparable state of affairs holds for the Swedish Leaf Dataset (SLD) [16], Flavia [17], and the Database of Leaf Images (DLI) [41]. A parallel limitation is present in Plant Seedlings V1, where the photographs display full plants—rather than only leaves—but do so confined to laboratory circumstances [32]. Freshly emerging datasets are advancing the frontier by incorporating in situ images collected with a range of cameras featuring diverse resolutions, sensor packages, and deployment platforms [31, 37, 42–44]. While these resources are scattered piecemeal across the internet, the Eden Library platform aims to provide a unified access point for heterogeneous corpora tailored to diverse technical needs. In the more up-to-date collections (as typified by Eden Library), the imagery is not restricted solely to leaves—whether afflicted or sound and at various stages of symptom progression—but extends to additional plant structures, including stalks, blooms, and fruiting bodies [33, 42–44]. A further exacting dimension that the showcased corpora wrestle with is the variability in backdrops and lighting. An even weightier concern that publicly shared crop-oriented picture banks must confront is image markup, which ranks among the most pivotal undertakings in machine visual perception. This laborious, manual procedure entails fastening predetermined labels to a picture to convey that visual intelligence to computing machinery and to drill it in visual identification. The number of labels affixed to each picture may vary depending on the initiative's scope. Among the repositories reviewed, the most contemporary ones provide annotations for items within the pictures [31, 33, 35, 36, 38, 40, 42, 43]; this suggests that increasingly capable AI-powered systems are being developed—systems that cannot merely assign a category to an object but also determine where within the picture that object lies [14].

Table 1. Synopsis of the most noteworthy extant public datasets ordered by publication date. This table broadens the compilation offered in Ref. [45].

Dataset name	Classes	Image quantity	Annotated samples	Image Target	Citation
Swedish Leaf Dataset (SLD)	15	1125	-	Leaves	[16]
Flavia	33	1,907	-	Leaves	[17]
UCI Leaf Dataset	40	443	-	Leaves	[18]
Leafsnap	185	7,719	-	Leaves	[19]
One hundred plant species leaves data set	100	1,600	-	Leaves	[20]
MalayaKew Leaf	44	46,288	-	Leaves	[21]
PlantVillage	38	54,303	-	Leaves	[22]
DeepFruits	7	587	-	Fruits	[23]
Crop Weed Field Image Dataset (CWFID)	2	60	NR	Plants/Weeds	[24]
Folio	32	640		Leaves	[25]
Weed Detection in Soybean Crops	4	12,087	-	Crops/Weeds	[26]
Sugar Beets 2016 dataset	3	12,340	-	Crops/Weeds	[27]
Vegfru	25	160,000	-	Fruits/Vegetables	[28]
Orchard Fruit	3	3,704	-	Fruits	[29]
weedNet	3	465	NR	Crops/Weeds	[30]
Carrot-Weed	2	39	NR	Crops/ Weeds	[31]
Plant Seedlings V1	12	407	-	Plant seedlings	[32]
Plant Seedlings V2	12	5,539	-	Plant seedlings	[32]
CropDeep	31	31,147	49,765	Fruits/Vegetables	[33]
Joint stem detection dataset	4	1,321	-	Crops/Weeds	[34]
MangoNet Semantic Dataset	2	49	NR	Fruits	[35]
MangoYOLO	2	1,730	NR	Fruits	[36]
DeepWeeds	9	17,509	-	Weeds	[37]
Open Plant Phenotype Database	47	7,590	315,038	Weeds	[38]
KFuji RGB-DS	2	967	12,839	Fruits	[39,40]
A Database of Leaf Images (DLI)	22	4,503	-	Leaves	[41]
Embrapa WGISD	-	300	4,432	Fruits	[42]
MinneApple	2	1,001	41,325	Fruits	[43]

Deep Nutrient Deficiency for Sugar Beet (DND-SB)	7	5,648	-	Crops	[44]
--	---	-------	---	-------	------

Abbreviation: NR = Not reported.

The present paper introduces the Eden Library platform to the wider research community, aiming to address the existing deficit in openly available crop and plant image banks by providing a collection of plant photographs collected via proximal and/or aerial modalities. The datasets put forward are envisioned as living repositories to be nurtured and augmented over time, not merely as “concluded” academic outputs addressing only a sparse assortment of species, growth intervals, and environmental backdrops that are only partially relevant but fall short of the manifold realities of cropping systems. What is more, the imagery is already annotated with plant and disorder labels, thereby freeing users from the wearisome, time-sapping task of manually marking up images [45]. Existing open-access plant species databases may provide substantial image counts. Yet, the amassing of that material transpires either inside regulated cultivation spaces (e.g., glasshouses, seedling trays) to enable camera rigs to gather pictures in a routinized and expeditious manner, or within field environments via robotic platforms characterized by particular image-taking attributes (e.g., camera model, standoff distance, illumination backdrop). Eden Library seeks to bridge this chasm by providing a rich, diverse and well-organized plant image corpus to energize forward-looking, vision-driven, AI-infused field applications. In parallel, recently articulated frameworks such as the FAIR (Findable / Accessible / Interoperable / Reusable) tenets for elevating data-sharing practices are being pursued as a longer-range objective, aiming to deliver amplified value to the user community [48]. To realize this ambition, an online application has been developed with an accessible, intuitive interface that will serve as the home for both the database proper and the full suite of dataset metadata.

The remainder of this manuscript is organized as follows. The image acquisition approach, metadata framework, and annotation protocol are first outlined. This is followed by a presentation of the current statistical profile of the database, along with illustrative imagery. The application and its core functional capabilities are then described. A critical discussion of the contributions and current boundaries of the Eden Library follows. Finally, concluding thoughts and prospective avenues for future effort are presented.

Materials and Methods

Data acquisition and collection

The Eden Library assembles and structures its imagery to cater to four core agricultural visual recognition operations, namely: i) image classification, ii) object detection, iii) semantic segmentation, and iv) synthetic image generation. A use-case-oriented mindset has steered the entire collection and organizational effort, simplifying the subsequent training and fielding of vision-driven solutions.

The data holdings pool together both ground-level and UAV-sourced captures generated through multiple investigative projects. Photographs were taken on location under authentic production conditions, harnessing a range of optical tools, including smartphone-grade cameras and dedicated machine vision sensors engineered for rapid, high-volume proximal sensing (Figure 1). For the aerial survey component, the Parrot SEQUOIA multispectral imager serves as the primary sensing payload of the UAV setup. Across the board, a diverse set of spectral detectors—spanning RGB, near-infrared, and thermal wavebands—was deployed during the imaging sorties, providing a rich assortment of plant corpora with contrasting spectral signatures. Given that these materials are slated for use in visual perception and AI pipelines, the spatial granularity of the pictures serves as a defining characteristic for each dataset. Fine-grained, high-resolution photographs convey dense spatial cues that frequently prove advantageous for remote sensing analytics or for discerning minute targets, such as disease lesions peppering a leaf surface. Elsewhere, medium and low pixel counts become the more pragmatic choice (< 4 MegaPixels), since the input streams demanded by the majority of deep learning frameworks

operate at modest dimensions for classification and object detection tasks^{2,3}, owing to the processing overhead that accompanies high-resolution data.



Figure 1. A sample of the gear utilized for ground-based field data capture: (a) an RGB imager equipped with a C-mount lens; and (b) a laptop linked to an RGB depth camera.

Regarding the collection protocol, a range of photographic approaches was trialed, guided by the target crop and the intended downstream application of the recognition task. For close-range imagery specifically, the following principles were adhered to: 1) within greenhouses and for orchard or vineyard settings, a lateral perspective was employed, whereas for row crops a bird's-eye perspective was adopted; 2) a deliberate equilibrium was pursued between encompassing the full symptom, leaf, or plant and retaining the ability to parse subtle textural details; 3) overexposed shots, underexposed shots, and pictures with minor blur were intentionally preserved within the corpora, since their inclusion promotes superior model generalization and resilience to noise, effectively approximating the messy realities of field capture. The primary motivation underpinning these acquisition preferences is to approximate, to the greatest degree possible, the image-gathering patterns that would be produced by tractors and agricultural robots during prospective vision-enabled operations. On the remote sensing side, aerial sorties were planned around the target Ground Sampling Distance (GSD) suited to the crop under examination. Coarser GSD requirements and analysis conducted at the plant-cluster tier (as in the case of stevia) led to high-elevation flights, whereas resolving individual plants and teasing apart weed species from overhead views necessitated finer GSD values and correspondingly lower flight altitudes.

Metadata description

A foundational objective of the Eden Library platform is to deliver frictionless access to a heterogeneous collection of datasets spanning distinct agricultural domains, while enabling their interweaving for application-focused research. Achieving this demands a clearly articulated metadata model to formally capture the dataset concept. As depicted in **Figure 2**, every dataset represents an intricate entity delineated by a set of classes and attributes, all navigable via the Eden Library interface. The metadata principally encodes the plant species (both its colloquial and Latin/scientific designation), the total image count within the dataset, the phenological stage for perennial crops, or the days elapsed post-sowing for annuals. With object detection gaining momentum as a subfield of precision agriculture, whether a given image contains multiple crop species is sufficiently relevant to warrant inclusion in the metadata framework. Where a disorder afflicts a crop, the etiological agent is further categorized under distinct conceptual headings (disease, weed, and nutrient deficiency), each paired with its own attributes (common and scientific/Latin names). In this way, five discrete groupings become browsable: healthy specimens, diseased material, weeds, pests, and nutrient deficiencies.




Figure 2. The metadata blueprint employed by the Eden Library Platform for dataset characterization.

Beyond these elements, every supplied dataset is accompanied by metadata detailing a range of technical facets, including spatial resolution, image capture mode (proximal, UAV), image modality (RGB, spectral), and, crucially, a flag indicating whether the imagery includes annotations. For UAV-captured material, the flight altitude is also specified. Datasets stemming from disparate recording devices and geographic origins yet depicting the same plant taxon, disorder, or other commonality were intentionally kept segregated rather than homogenized, thereby allowing the end-user to leverage the applied standardization rules and to interrogate each crop at a deeper level. Merging of datasets remains an option left entirely to the user’s discretion and can be performed without any ensuing conflicts. In addition to these core metadata entries, specifications for camera setup parameters are included whenever the acquisition apparatus provides them. Furthermore, in cases where datasets are interrelated, the nature of that connection is likewise documented, facilitating effortless traversal between related datasets. Such a relationship can assume a spatial, temporal, or spectral character.

The datasets are also linked to (i) the scholarly publications in which they have been previously employed (e.g., Ref. [32]) and (ii) any other datasets utilized within that same body of research (e.g., PlantVillage [22]). This cross-referencing substantially lowers the barrier to replicating a specific experimental setup. Additionally, excerpts of the code underpinning these studies are displayed within interactive notebooks accessible on the platform. Inside those notebooks, fully realized

pipelines for running deep learning-based techniques on Eden Library imagery are documented and made execution-ready, with the dual aims of jumpstarting new investigations and imparting practical know-how for confronting widely studied problems such as image classification, object detection, and data augmentation. Finally, it should be underscored that each dataset is assigned a globally unique identifier (“formal_name”), which serves as the reference string to be cited whenever the dataset is used in a publication. This formal name inherently encodes the dataset’s version number. All datasets undergo routine auditing by domain authorities; should any defect be uncovered, an amended edition of the dataset is released, the version number is incremented, and the formal name is updated accordingly. Representative instances of these unique designations include “Black_nightsade_220519_Weed_zz_V1” or “Cotton_180720_Healthy_zz_V1”. These labels were intentionally architected to remedy the semantic vacuity typical of conventional alphanumeric codes while simultaneously guaranteeing long-term stability, ensuring that the names referenced in published articles remain valid for the foreseeable future.

Data annotation

Once the photographic material has been assembled and the metadata schema—the conceptual scaffolding—has been formalized, the next step is to engineer an annotation workflow that ultimately populates a database containing all pictures and their requisite metadata. The principal stages of this annotation workflow are schematized in **Figure 3**. Following the initial image collection, a highly experienced domain authority screens the raw corpus, eliminating images that exhibit irreparable quality flaws, such as indecipherable subjects, corrupt file structures, or an incorrectly targeted crop species. While the sanitized collection is placed in storage (“Clean images” in **Figure 3**) for eventual fusion with the metadata, the pristine raw corpus is always preserved in its unmodified state, standing ready for consultation should any reexamination be triggered during the annotation cycle (“Raw images” in **Figure 3**). It is worth emphasizing that annotation is an intrinsically iterative enterprise; accumulated expertise carries tremendous weight, and occasional reframings of the initial problem statement may be enacted.




Figure 3. The workflow followed for assembling and annotating the datasets incorporated into the Eden Library platform.

Thereafter, the senior authority devises a comprehensive set of directives specifically tailored to the annotation of a given plant corpus (“Senior experts create annotation guidelines” in [Figure 3](#)). Enclosed within these directives is a corpus of rules and boundary conditions concerning the formulation of both (i) the agricultural question and (ii) the annotation task. In addition, the senior authority personally annotates a subsample of the imagery, which will later serve as a gold-standard reference for verifying the fidelity of annotations produced by the broader annotator pool. These additional annotators, while also versed in the agricultural discipline, bring a comparatively shallower depth of hands-on experience relative to the senior authority. Running in parallel with the creation of annotation guidelines, certain metadata elements that elude straightforward automated extraction are manually appended (“Manual metadata is added to the dataset” in [Figure 3](#)). Such elements might include, for instance, inter-dataset relationships or the identity of the landholder on whose property the crop photographs were taken. Selected portions of this metadata remain confidential and reserved for administrative purposes. This activity proceeds concurrently with the senior specialists’ formulation of the annotation rulebook.

It is critical to note that every image specimen is inspected after it is annotated by senior authorities to prevent potential missteps, such as the erroneous labeling of a symptom. A curated set of annotated image exemplars can be viewed in [Figure 4](#). Beyond this manual verification checkpoint, and with a view to standardizing the annotation pipeline and dampening systematic annotator skew, a dedicated quality assurance apparatus has been put in place (Quality assurance tool in [Figure 3](#)). Through this apparatus, a randomly sampled fraction of the original corpus (2%–5%) is carved out, and each annotator independently performs the labeling exercise for the specified objects of interest. The outputs contributed by each annotator are then measured against those of the supervising expert, with the appraisal factoring in both the selected annotation objects and the spatial context surrounding each observation to flag any departures from the prescribed annotation convention. This step proves especially indispensable for plant disease identification and the accurate demarcation of symptoms.



Figure 4. Illustrative samples of annotated photographs hosted on the Eden Library platform. The stroke width of the bounding boxes has been exaggerated to improve on-screen legibility.

Corpora depicting afflicted plants were labeled according to the outwardly observable symptom, and the markup precision reached the granularity of individual leaves. For healthy plant corpora, two diverging annotation conventions were applied: 1) circumscribing the plant with a rectangular bounding box, and 2) sketching a polygonal boundary hugging the plant silhouette, the intent being to excise as much extraneous background as feasible from the labeled zone. These twin annotation modalities serve separate application pathways, spanning object detection and classification on one track and biomass or growth estimation on the other.

Once the human-driven labeling and quality assurance rounds are complete, the metadata needed to satisfy the conceptual blueprint outlined in the prior section is harvested algorithmically (“Enrich metadata” in [Figure 3](#)). This class of metadata routinely resides within the image containers, where fields such as capture date, timestamp, and pixel dimensions can be read directly. For metadata attributes not natively present in the image headers, a battery of heuristic decision rules is invoked to satisfy any outstanding metadata requirements. Thereafter, the sanitized photographs, the annotations, and the fully populated metadata are merged into a single cohesive bundle, giving rise to a deployment-ready dataset that meets the specification illustrated in [Figure 2](#) (“Integrate images, metadata, and annotations” in [Figure 3](#)). A concluding validation sweep is subsequently executed to confirm that no consequential errors have infiltrated the live repository (“Eden Library Platform Database” in [Figure 3](#)). In the final step, a designated Eden Library team member pushes the dataset to the platform’s public-facing tier (“Upload dataset to the platform” in [Figure 3](#)). It merits emphasis that a dataset’s journey does not conclude upon publication. Each dataset is subject to scheduled, cyclic reassessments to verify its sustained utility for the intended constituencies—agronomists and machine learning practitioners alike. Should

a flaw be unearthed during any such review, the entire pipeline is relaunched beginning at the problem definition stage (“Senior Experts create annotation guidelines”).

Database Content

Table 2 offers a consolidated snapshot of the crops, their respective health designations, and the number of photographic records capturing that particular crop-status pairing. The table further indicates whether annotated imagery is provided, and thus whether the material lends itself to training image classification (IC) or object detection (OD) architectures. Zooming in on individual crops, the tomato grouping commands the largest repository footprint, with 1485 images, followed by grapevine at 1360. These crops also run the gamut of possible health statuses, namely: thriving (healthy), diseased, suffering from nutrient insufficiencies, and under pest pressure. In select cases, multiple plant taxa appear together within a single frame, yielding datasets with multiple labels.

Table 2. Crops and their associated statuses are housed within the Eden Library Platform (Image classification (IC), Object detection (OD)).

Crop name	Status	Number of images	AI problem
Avena sterilis	Weed	12	OD
Black nightshade	Weed	123	IC
Broccoli	Diseased	14	OD
Broccoli	Healthy	244	OD
Chinese cabbage	Diseased	51	IC
Chinese cabbage	Healthy	100	IC
Cotton	Healthy	1079	IC
Crepis Foetida	Weed	12	OD
Cucumber	Diseased	15	OD
Cucumber	Pests	126	OD
Grapevine	Diseased	250	OD
Grapevine	Healthy	983	IC
Grapevine	Nutrient def.	12	OD
Grapevine	Pests	60	OD
Lemon tree	Pests	9	OD
Lepidium draba	Weed	12	OD
Lolium rigidum	Weed	38	OD
Olive tree	Diseased	66	OD
Olive tree	Healthy	76	IC
Olive tree	Pests	16	OD
Orange tree	Pests	218	OD

Pear tree	Pests	15	OD
Pepper	Healthy	355	IC
Potato	Healthy	99	IC
Red cabbage	Healthy	52	IC
Scandix	Weed	12	OD
Sinapis arvensis	Weed	38	OD
Tomato	Diseased	222	IC
Tomato	Healthy	1105	IC/OD
Tomato	Nutrient Def.	12	OD
Tomato	Pest	146	OD
Velvet leaf	Weed	129	IC
Veronica hederifolia	Weed	38	OD
Watermelon	Diseased	8	OD
Watermelon	Healthy	117	IC/OD
Zucchini	Diseased	90	OD
Zucchini	Healthy	227	IC/OD
Total		6428	

Worth underscoring is that the database deliberately targets some of the most widespread and economically damaging afflictions to which a crop is routinely exposed. These very afflictions constitute the detection and classification targets of any artificial intelligence system aspiring to flag them at the earliest feasible juncture. To cite a handful of concrete disease examples present in the database: *Plasmodiophora brassicae* (broccoli), *Sclerotinia* (cucumber), *Powdery mildew* (grapevine), *Downy mildew* (grapevine), *Bactrocera oleae* (olive tree), and *Aleurothrixus floccosus* (olive tree). Turning to the pest category, one finds *Tetranychus* (cucumber), Thrips fruit (cucumber), *Eriophyes vitis* (grapevine), *Bactrocera oleae* (olive tree), and *Tuta absoluta* (tomato). An exhaustive catalog of these disorders is furnished in [Table 3](#). In the weed category, species including black nightshade, velvetleaf, and *Veronica hederifolia* are likewise represented ([Table 2](#)). These particular weeds act as competitive stressors during the cultivation of tomatoes and cotton, respectively. When ranked by image count, Powdery mildew is the most frequently represented disease, with 138 images, followed by Esca with 107. On the pest side, the largest image pools belong to *Tuta absoluta* and *Liriomyza*, with 126 and 100 images, respectively.

Table 3. Crop disorders cataloged within the Eden Library Platform (Weeds excluded).

Disorder	Type	Number of images
<i>Aleurothrixus floccosus</i>	Pest	58
<i>Aonidiella aurantii</i>	Disease	65
<i>Bacterial damage</i>	Disease	9
<i>Bactrocera oleae</i>	Pest	92
<i>Botrytis</i>	Disease	21

<i>Camarosporium dalmaticum</i>	Disease	19
<i>Citrus thrips</i>	Pest	13
<i>Dialeurodes</i>	Pest	22
<i>Didymella</i>	Disease	8
<i>Downy mildew</i>	Disease	60
<i>Eriophyes vitis</i>	Pest	60
<i>Esca</i>	Disease	107
<i>Fusarium</i>	Disease	51
<i>Gleosporium olivarum</i>	Disease	47
Leaf mold	Disease	8
<i>Liriomyza</i>	Pest	100
<i>Nitrogen deficiency</i>	Nutrient Deficiency	12
<i>Palpita unionalis</i>	Pest	16
<i>Phyllocnistis citrella</i>	Pest	50
<i>Plasmodiophora brassicae</i>	Disease	14
<i>Potassium deficiency</i>	Nutrient Deficiency	12
<i>Powdery mildew</i>	Disease	138
<i>Sclerotinia</i>	Disease	6
<i>Stephanitis pyri</i>	Pest	15
<i>Tetranychus</i>	Pest	79
<i>Thrips fruit</i>	Pest	47
Tomato brown rugose fruit virus	Disease	43
Tomato cracking	Disease	12
Tomato Leaf Curl New Delhi Virus	Disease	69
<i>Tuta absoluta</i>	Pest	126

To round out the picture, **Figure 5** displays a curated set of specimen images retrievable from the database. Alongside the botanical variety, multiple acquisition platforms are represented: (i) ground-based proximal sensing (e.g., black nightshade and Chinese cabbage); and (ii) UAV-based remote sensing (e.g., broccoli (b) and grapevine (b)). As for the imaging modality, the entire set consists of RGB captures, except for grapevine (b), which is a thermal recording. A point of particular significance is that by featuring the identical plant taxon imaged under differing environmental regimes or sensor configurations, the Eden Library platform creates a fertile testing ground for cutting-edge research themes such as domain transfer and model resilience [31, 49].




Figure 5. Specimen images accessible via the Eden Library Platform.

Online Application/Prototype

Eden Library: An overview of the web application

Going beyond the underlying datastore that preserves all metadata, annotations, and raw imagery, Eden Library is rounded out by a browser-accessible application that (i) furnishes effortless retrieval of the desired information and (ii) is engineered in accordance with sound software craftsmanship principles. The primary interface panels are illustrated in the accompanying **Figures 6a-6c**, each of which is a genuine screenshot captured from the operational deployment. **Figure 6a**, the “Datasets” view, renders the exhaustive dataset inventory using a card-based visual paradigm. This panel is tasked with surfacing all datasets currently registered in the system and is designed as a highly fluid component, as fresh datasets are ingested weekly. Arrayed alongside the main listing are core interactive functions—including search, filtering, and sorting—together with a top-mounted navigation bar and clearly delineated controls for user sign-up and login. **Figure 6b** depicts the “Dataset details” view for a single dataset, where the full complement of metadata attributes describing both the botanical subject and the dataset’s technical profile is presented. A strip of six preview thumbnails is also provided, for users who want a quick visual appraisal of the dataset’s contents before diving deeper.

Furthermore, whenever the dataset in question has been cited within a scholarly article, the complete set of bibliographic references and the inventory of companion datasets employed in that article are displayed, irrespective of whether those companion datasets are hosted within the Eden Library ecosystem or reside in an external repository such as PlantVillage or Plant Seedlings. **Figure 6c** shows the “Notebooks” listing—a dedicated view for enumerating all notebooks, along with their metadata descriptors, which draw on specific Eden Library datasets, and outgoing links to external runtimes where these notebooks can be interactively executed. The Eden Library application also incorporates two supplementary views, namely “Homepage” and “FAQ,” which serve outreach and informational functions respectively.




Figure 6. The core views composing the Edén Library Platform: a) the panel listing all available datasets, b) the panel detailing a chosen dataset, and c) the panel displaying notebooks that offer tutorials for working with the datasets.

User research and information architecture

During the formative design stage of Edén Library, a human-centered design methodology was embraced, with the express aim of delivering a finished artifact that would prove functional, intuitive, and readily adoptable by each target user segment. Toward this end, a meticulous exploration of UX/UI heuristics and information architecture was undertaken, with a clear focus on end-user priorities. A preliminary user research phase, operationalized through structured personal interviews with representative user categories (AI scientists, agronomists, software developers/engineers, academic researchers), was launched to understand and profile these audiences as concrete user personas and to distill highly specific functional requirements and boundary conditions. This phase was augmented by a competitive landscape review of extant, analogous dataset platforms (Table 1), conducted to pinpoint the precise gaps that Edén Library would be tasked with bridging. The intelligence gathered from the user research phase served as the springboard for the subsequent design phase, culminating in prototype mock-ups that embodied the full suite of predetermined functionalities. A usability evaluation phase followed, during which clickable, high-fidelity interactive mock-ups were presented to a small group (5–10 participants) of prospective users, with the dual objectives of stress-testing foundational design assumptions and gathering constructive feedback to inform revisions and refinements before the development phase kickoff.

Implementation

Constructing the Edén Library platform required an adaptable full-stack technology foundation aligned with current web conventions, most notably a RESTful design paradigm. After surveying a range of candidate technologies, the decision was made to build on the MEAN stack—a bundle comprising MongoDB, Express.js, Angular, and Node.js. Within this assembly, MongoDB serves as the data storage backbone alongside the Node.js server; Express.js operates as a server-side wrapper layered on top of the more rudimentary Node.js interface; and Angular serves as the browser-side toolkit for assembling Model-View-Controller single-page web experiences. The MEAN stack can be viewed as a modern-era successor to the LAMP open-source suite (which bundles the Linux operating system, the Apache HTTP server, MySQL as the relational database, and PHP as the dynamic scripting language), a combination that was widely used in prior decades.

The modular, component-driven design philosophy of Angular fosters a codebase that scores highly on both reusability and long-term maintainability. The encapsulated, decoupled nature of each component additionally lowers the barrier to constructing unit tests—tests whose purpose is to confirm the correct functioning of the application's smallest discrete pieces. Furthermore, Edén Library is a fully realized web application, engineered to be accessible from and visually adaptive to a diverse range of devices and screen geometries. Leveraging Express.js shortens the development cycle for

Node.js web applications, enables quick setup of a REST API server, and allows developers to map application routes directly to HTTP verbs.

As shown in **Figure 7**, the intrinsically iterative trajectory of the Eden Library endeavor, coupled with the absence of a widely ratified data standard for capturing the semantics of agricultural datasets, rendered MongoDB's schema-agnostic philosophy particularly apt as a technological underpinning. This choice enabled progressive enrichment of the database schema (**Figure 2**) while circumventing the inflexible schema mandates typical of SQL-flavored databases such as MySQL. The flexibility of MongoDB supports an iterative development rhythm in which schemas can be refactored at frequent intervals to better encapsulate agricultural (dataset) realities. Additionally, MongoDB embraces a document-oriented model for storing information as binary JSON (JavaScript Object Notation) artifacts, a representational format that dovetails effortlessly with the broader web ecosystem. The JSON format straddles the boundary between machine-parseable data and text accessible to human readers, a characteristic that greatly simplifies the handling and operational use of database metadata for non-engineering domain practitioners such as agronomists. The database itself maintains the logical coupling between dataset-level metadata and the corresponding image assets persisted on the server.




Figure 7. High-level architectural diagram of the Eden Library system.

As shown in **Figure 7**, the database also maintains links to Jupyter notebooks hosted on GitHub⁴. These notebooks serve the mission of propagating knowledge across the artificial intelligence and precision agriculture communities. Segments of this knowledge have constituted the underpinning of earlier peer-reviewed contributions (e.g., Ref. [15]). As a medium, Jupyter notebooks have experienced a dramatic rise in uptake over the past few years, propelled by the inherently iterative and exploratory nature of much artificial intelligence development. This characteristic sets it apart from more traditional software engineering paradigms. Beyond this, notebooks provide an essential element for advancing science: the capacity for fully reproducible experiments. The notebooks offered within the Eden Library framework are organized according to distinct artificial intelligence problem categories: (i) image classification; (ii) object detection; (iii) image preprocessing; (iv) automated machine learning; and (v) data augmentation through Generative Adversarial Networks (GANs).

Several platforms, including Google Colab⁵, permit investigators to run notebooks without incurring charges. Given that the Eden Library notebooks are persisted on GitHub, it is straightforward to exploit this synergy: any user can launch the notebooks directly within Google's cloud environment, capitalizing on the provided computational accelerators—GPUs and TPUs alike. Finally, the nbviewer⁶ utility is also accessible via the Eden Library platform, functioning as a lightweight tool for rapidly rendering notebook content in a read-only mode.

Deployment

Continuous deployment has remained a cornerstone objective throughout the Eden Library project. Internalizing lessons from user feedback is a continuous process, and delivering new capabilities alongside iterative refinements is a non-negotiable prerequisite for sustaining relevance among precision agriculture stakeholders. One of the defining technological shifts witnessed over recent years has been the swelling enthusiasm for containerization as the preferred mechanism for application deployment, with Docker—the specific containerization engine adopted within this undertaking—at the forefront. Docker equips a development squad to encapsulate an application and its complete dependency tree in a standardized, resource-efficient container that can subsequently be staged across diverse runtime environments (production, development, testing, and beyond). This practice provides ironclad assurance that every required dependency will be satisfied, regardless of the host machine’s operating system or preinstalled library set. In summary, embracing container technology yields compressed development and testing feedback loops, a benefit of considerable magnitude for the platform’s end-user constituency.

EDEN Library platform’s functionalities and features

Figure 8 enumerates the most consequential use cases and feature set that Eden Library brings to life. At a practical level, the application arms users with multiple avenues for discovering and obtaining the plant datasets they seek. While scrolling through the dataset inventory, users may opt to retrieve individual datasets piecemeal or to assemble them into a handpicked collection, which they can inspect, adjust, and ultimately download as a unified compressed payload. To download one or more datasets, a user must first register a personal (sign-up) account, a process that requires supplying basic authentication credentials and optional supplementary information (domain expertise, research interests) to flesh out their user profile. By completing the account creation flow, the user consents to the “Terms of Service” and the “Data Protection Policy,” both of which are detailed in dedicated sections within the web application.




Figure 8. Use case diagram capturing the principal functionalities of Eden Library. <<include>> denotes that the indicated use case is a mandatory step for fulfilling the originating use case. <<extends>> signifies that the originating use case may, on an optional basis, be expanded by the indicated use case.

To issue a search query over the dataset catalog, a user may either i) type the common or scientific name of a crop into the search text field, or ii) pick from the autocomplete recommendations that materialize beneath the search bar. The matching results are then displayed, and the user can progressively prune them by engaging a suite of filters and nested subfilters (derived from the metadata schema detailed in [Figure 2](#)) that relate either to botanical attributes (species, cultivar, disorder presence, etc.) or to image-level attributes and metadata (spatial resolution, acquisition source, image count, dataset storage footprint, etc.). An additional capability placed at the user's disposal is the ordering of result sets according to multiple sorting keys, including (i) alphabetical sequencing by dataset name or (ii) ranking by download frequency or popularity. By navigating to the detail panel of any dataset, the user gains visibility into its full metadata record, along with a small gallery of preview thumbnails. Within this detail view, extra contextual information and linked resources are exposed when available, such as: (i) a curated bibliography of scholarly publications that have employed or incorporated the focal dataset; (ii) a cross-reference listing of other datasets that appear alongside it in published studies, complete with navigable links among them; and (iii) pointers to notebooks in which the dataset has been actively utilized.

Notebooks crafted by the Eden Library research team and inspired by the platform's datasets are likewise surfaced on a dedicated page of the application. Every notebook is accompanied by rich descriptive metadata, which can be wielded as faceted filtering dimensions, specifically (i) the computer vision problem it addresses, (ii) the scientific methodologies it puts into practice, and (iii) the software libraries or frameworks on which it depends. Users, whether logged in or browsing anonymously, can access the notebooks and execute or render them visually in external runtime environments (GitHub, Nbviewer, Google Colab) via the respective hyperlinks.

While not depicted in [Figure 8](#) to preserve diagrammatic clarity, users can also access educational content in the FAQ section to familiarize themselves with the operational processes powering the Eden Library platform. For illustration, the collection, annotation, and stewardship of datasets, as well as the authoring of notebooks, are all addressed. To close the loop, users are provided a channel to submit feedback to site administrators on any aspect of the platform or its datasets by filling out a brief form that captures their identifying details along with the message they intend to relay. This capability enables bidirectional communication with the user community and catalyzes ongoing enhancements to the Eden Library platform, guided by real-world user input.

Discussion

A central offering of the Eden Library is the ready accessibility of numerous datasets procured under authentic, unmanipulated field circumstances. The work of Refs. [\[12\]](#) and [\[13\]](#) has shown that predictive models developed and tuned in laboratory settings cannot be straightforwardly transferred to real-world operational environments. It is precisely for this reason that repositories such as Flavia [\[17\]](#), the Plant Village Dataset [\[18\]](#), or Plant Seedlings [\[32\]](#), while certainly useful for foundational inquiry and for honing computer vision competencies, remain a considerable distance from the deployment landscapes where such models are intended to function in practice. In marked contrast, photographs housed within Eden Library exhibit heterogeneous backdrops; occasionally, a lone plant sits in sharp relief, whereas in other frames numerous plants coexist, displaying a mixture of vigor and affliction from various disorders. Owing to this composition, the image corpus mirrors the genuine, messy circumstances under which a model would plausibly be called into action. It should be noted that Eden Library does not enjoy a monopoly on imagery sourced from realistic settings—comparable material can also be encountered in other initiatives, among them [\[31, 37, 42–44\]](#).

Another realm where the Eden Library distinguishes itself is the dense, semantically rich metadata scaffolding used to describe both the overarching dataset and the individual photographs. The majority of the collections enumerated in [Table 1](#) receive exhaustive treatment within the scholarly articles that first introduced them. By contrast, Eden Library is built upon the conviction that a dataset, together with its accompanying metadata, ought to function as a self-sufficient

artifact—one that surfaces the most germane facts to the prospective user without compelling them to either (i) retrieve the dataset in its entirety or (ii) wade through an associated manuscript where the dataset might be characterized. Users are thereby shielded from an excess of information precisely at the moment they embark upon an agriculture-facing AI venture. Furthermore, possessing a purpose-built metadata vocabulary expressly designed to capture the essence of a precision agriculture dataset enables semantically elevated searches and the scouting for fresh datasets to fold into sharply defined research agendas. A handful of datasets, DND-SB [44] and DLI [41] among them, make use of data-centric repository platforms such as Zenodo7 or Mendeley8. These platforms track pertinent metadata, including the publication associated with the item, the total storage footprint, the download count, the version string, and the license. Yet finer-grained details—the imaging modality, the spatial pixel resolution, the sun time, or the sensing hardware—are relegated to unstructured free-text descriptions, assuming they appear at all. The practical upshot is that confirming the presence and value of such metadata (when it exists) requires downloading the full dataset. This step is not always economically or logistically feasible. The lack of an expressive, formally defined metadata model can make the fusion of disparate datasets a laborious affair, as separate ad hoc schemas often rely on incompatible terminology or are inherently vague. At the same time, committing to a feature-rich metadata schema introduces its own difficulties concerning database completeness. Regrettably, given the sheer breadth and variability of acquisition settings, some datasets are bound to be missing entries for specific metadata slots.

Nevertheless, the acquisition and annotation workflow is fundamentally iterative, allowing the methodology to be successively sharpened and reducing the incidence of missing information that could prove valuable to the end user. As a concrete case in point, the sensor model—data implicitly tucked inside image files via the Exif format9 yet absent from the current platform interface—will be automatically extracted and surfaced as an explicit element of the canonical metadata schema (Figure 2). One non-trivial disadvantage that accompanies the choice not to lean on third-party data-oriented platforms like Zenodo is the heightened operational expenditure associated with sustaining and growing a standalone platform. In its current incarnation, Eden Library employs a freemium model, under which certain datasets are gated behind a paywall. This structure may deter a segment of potential users, yet achieving enduring viability for a project that aspires to openly share dozens of datasets requires a dependable funding mechanism. Other monetization frameworks, such as donation-based support, are still being examined.

As documented in Table 2, the Eden Library database currently contains 6,378 images. Considering individual datasets in isolation and momentarily setting aside the particular AI challenge at hand, the most voluminous collections stand at 539 and 443 images (“Cotton-180720-Healthy-zz-V1” and “Grape vine-210520-Healthy-zz-V1,” respectively). These totals eclipse those put forward by certain closely related investigations (e.g., WGSD [42], CWFID [24], V1 Plant Seedlings [32]); however, seen from the opposite angle, this also acknowledges the existence of datasets—CropDeep [33], PlantVillage [22], and VegFru [28] being prime examples—whose image counts are orders of magnitude larger. It is crucial to stress, though, that individual datasets within the Eden Library ecosystem can be programmatically merged before being pressed into service for a targeted application, and the high intrinsic resolution of the imagery permits each frame to be sliced into a grid of smaller tiles, thereby dramatically inflating the effective training set size. Through such mechanisms, investigators can interleave disparate datasets and construct their own application-specific corpora.

What is more, class imbalance—a pervasive feature of real-world problems [50]—can be examined in a controlled fashion by deliberately fusing carefully chosen datasets (for instance, pairing “Grape vine_300620_Downy mildew_zz_V1” with “Grape vine_210520_Healthy_zz_V1”). Additionally, generative modeling techniques could be applied to synthesize artificial images seeded from the existing dataset, as previously advocated in [51]. Lastly, the datasets curated within Eden Library place deliberate emphasis on capturing genuine in-field variability and incorporating a multiplicity of distinct scenes, rather than pursuing the strategy of collecting exhaustive image bursts of an identical field plot. This approach tends to contribute comparatively little incremental visual information.

Independent of the merits already cataloged, the appearance of fresh standalone datasets or consolidated databases that aggregate numerous datasets invariably enriches the broader research landscape. To offer a concrete illustration, several phenotyping-related agricultural questions are intimately tied to specific genetic backgrounds and environmental contexts.

Hence, any data outlet that circulates annotated material will catalyze investigations into particular plant species, including those underrepresented in certain regional settings [52].

Conclusion

This paper has placed the spotlight on Eden Library, an emergent platform conceived for the open sharing of agricultural data. The fully specified metadata schema, the annotation pipeline, and the platform's signature capabilities have been detailed to provide a practical blueprint for other bodies seeking to address the chronic data shortage that constrains AI-powered solutions in the agricultural sector. Because the platform was built on an agile philosophy, supplementary features and additional datasets can be integrated seamlessly, thereby responding to the continuously escalating appetite for novel, vision-driven applications deployed at the field level. Looking forward, refinements to the annotation workflow—such as exploring semi-supervised labeling strategies—will be investigated to accelerate onboarding data throughput. In turn, a swifter arrival of high-caliber datasets will enable the research ecosystem spanning artificial intelligence, computer vision, and precision agriculture to drive further gains in the performance of contemporary technologies.

Arriving at a metadata framework that is simultaneously “stringent” in its structure and yet expansive enough to envelop nearly the entire phenotypic spectrum demands a delicate balancing act—one that requires weighing explicitness against the risk of inadvertently boxing the schema into a narrow subset of plant categories. Eden Library advances a holistic methodology for crafting metadata constructs that deliver genuine utility across a diverse array of agritech use cases.

Informed by the lessons absorbed from its user community, Eden Library will roll out several new functionalities and capabilities in the forthcoming cycle. These shall encompass: (i) a downloadable teaser bundle of six representative images per dataset, packaged alongside the corresponding annotation file, so that prospective users can evaluate and determine whether the dataset suits their particular ends; (ii) a broader menu of supported annotation formats, introduced to minimize the friction users face when plugging the datasets into assorted tools and frameworks; (iii) a newly established section functioning as an encyclopedic reference on plants and the symptoms they manifest—whether stemming from disease, nutritional shortfalls, or other causes—thereby exhibiting the full corpus of knowledge accumulated and systematized across the collection and annotation phases; and (iv) the addition of new faceted filters covering crop health status (healthy versus diseased) and the granularity of annotation (image-level labels, bounding boxes, pixel-accurate masks).

Eden Library lays the groundwork for delivering agricultural image data at scale, serving a broad spectrum of applications. Both the academic research community and commercial operators are calling for platforms that offer curated, structured image assets that are traceable, cross-comparable, and suitable for benchmarking. Eden Library meets this demand by providing the foundational infrastructure that enables stakeholders to both contribute assets and extract value from image data flows. Over the medium term, the platform will furthermore be opened to accept dataset uploads from external investigators eager to disseminate their own data and to amplify the reach of their research. Simultaneously, the notebooks area of the platform will be exploited to distribute ready-to-use code fragments that streamline the handling of the shared datasets. Achieving these scaling ambitions and efficiently orchestrating data pipelines, while meticulously honoring the licensing terms and ownership rights of third-party datasets, constitutes an ongoing challenge that will require diligent stewardship.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 01 Jan 2025

Revised: 30 Jan 2025

Accepted: 17 Feb 2025

Published online: 15 September 2025

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Patrício DI, Rieder R. Computer vision and artificial intelligence in precision agriculture for grain crops: a systematic review. *Comput Electron Agric.* 2018;153:311-8.
- Fountas S, Mylonas N, Malounas I, Rodias E, Santos CH, Pekkeriet E, et al. Agricultural robotics for field operations. *Sensors (Switzerland).* 2020;20(9):2672.
- Herwitz SR, Johnson LF, Dunagan SE, Higgins RG, Sullivan DV, Zheng J, et al. Imaging from an unmanned aerial vehicle: Agricultural surveillance and decision support. *Comput Electron Agric.* 2004;44(1):49-67.
- Ehmke T. Unmanned Aerial Systems for Field Scouting and Spraying. *Crops Soils.* 2013;46(3):4-8.
- Peña JM, Torres-Sánchez J, Serrano-Pérez A, de Castro AI, López-Granados F. Quantifying efficacy and limits of unmanned aerial vehicle (UAV) technology for weed seedling detection as affected by sensor resolution. *Sensors (Switzerland).* 2015;15(7):16761-78.
- Lima-Cueto FJ, Blanco-Sepúlveda R, Gómez-Moreno ML, Galacho-Jiménez FB. Using vegetation indices and a UAV imaging platform to quantify the density of vegetation ground cover in olive groves (*Olea Europaea* L.) in Southern Spain. *Remote Sens.* 2019;11(21):2564.
- Hague T, Marchant JA, Tillett ND. Ground based sensing systems for autonomous agricultural vehicles. *Comput Electron Agric.* 2000;25(1-2):11-28.
- Wang A, Zhang W, Wei X. A review on weed detection using ground-based machine vision and image processing techniques. *Comput Electron Agric.* 2019;158:226-40.
- Gracia-Romero A, Kefauver SC, Fernandez-Gallego JA, Vergara-Díaz O, Nieto-Taladriz MT, Araus JL. UAV and Ground Image-Based Phenotyping: A Proof of Concept with Durum Wheat. *Remote Sens.* 2019;11(10):1244.

Putra BTW, Soni P, Marhaenanto B, Harsono SS, Fountas S. Using information from images for plantation monitoring: A review of solutions for smallholders. *Inf Process Agric*. 2020;7(3):415-28.

Bohnenkamp D, Behmann J, Mahlein AK. In-Field Detection of Yellow Rust in Wheat on the Ground Canopy and UAV Scale. *Remote Sens*. 2019;11(21):2495.

Mohanty SP, Hughes DP, Salathé M. Using deep learning for image-based plant disease detection. *Front Plant Sci*. 2016;7:1419.

Ferentinos KP. Deep learning models for plant disease detection and diagnosis. *Comput Electron Agric*. 2018;145:311-8.

Arsenovic M, Karanovic M, Sladojevic S, Anderla A, Stefanovic D. Solving current limitations of deep learning based approaches for plant disease detection. *Symmetry*. 2019;11(7):939.

Espejo-García B, Mylonas N, Athanasakos L, Fountas S, Vasilakoglou I. Towards weeds identification assistance through transfer learning. *Comput Electron Agric*. 2020;171:105306.

Söderkvist O. Computer Vision Classification of Leaves from Swedish Trees. [dissertation]. Linköping: Linköping University; 2001.

Wu SG, Bao FS, Xu EY, Wang YX, Chang YF, Xiang QL. A leaf recognition algorithm for plant classification using probabilistic neural network. In: *Proceedings of the ISSPIT 2007 - 2007 IEEE International Symposium on Signal Processing and Information Technology*; 2007 Dec 15-18; Giza, Egypt. New York: IEEE; 2007. p. 315-20.

Silva PF, Marçal A, Silva R. Evaluation of Features for Leaf Discrimination. In: Kamel M, Campilho A, editors. *Image Analysis and Recognition. ICIAR 2013. Lecture Notes in Computer Science*. Vol 7950. Berlin: Springer; 2013. p. 197-204.

Kumar N, Belhumeur PN, Biswas A, Jacobs DW, Kress WJ, Lopez IC, et al. Leafsnap: A computer vision system for automatic plant species identification. In: Fitzgibbon A, Lazebnik S, Perona P, Sato Y, Schmid C, editors. *Computer Vision – ECCV 2012. Lecture Notes in Computer Science*. Vol 7573. Berlin: Springer; 2012. p. 502-16.

Mallah C, Cope J, Orwell J. Plant leaf classification using probabilistic integration of shape, texture and margin features. In: *Proceedings of the IASTED International Conference on Signal Processing, Pattern Recognition and Applications, SPPRA 2013*; 2013 Feb 12-14; Innsbruck, Austria. Calgary: ACTA Press; 2013. p. 82-9.

Lee SH, Chan CS, Wilkin P, Remagnino P. Deep-plant: Plant identification with convolutional neural networks. In: *Proceedings - International Conference on Image Processing, ICIP 2015*; 2015 Sep 27-30; Quebec City, Canada. New York: IEEE; 2015. p. 452-6.

Hughes DP, Salathé M. An Open Access Repository Of Images On Plant Health To Enable The Development Of Mobile Disease Diagnostics Through Machine Learning And Crowdsourcing. *arXiv*. 2015; abs/1511.08060.

Sa I, Ge Z, Dayoub F, Upcroft B, Perez T, McCool C. DeepFruits: a fruit detection system using deep neural networks. *Sensors (Basel)*. 2016;16(8):1222.

Haug S, Ostermann J. A crop/weed field image dataset for the evaluation of computer vision based precision agriculture tasks. In: Agapito L, Bronstein M, Rother C, editors. *Computer Vision - ECCV 2014 Workshops. Lecture Notes in Computer Science*. Vol 8928. Berlin: Springer; 2015. p. 105-16.

Munisami T, Ramsurn M, Kishnah S, Pudaruth S. Plant Leaf Recognition Using Shape Features and Colour Histogram with K-nearest Neighbour Classifiers. In: *Proceedings of the Procedia Computer Science*. 2015;58:740-7.

dos Santos Ferreira A, Matte Freitas D, Gonçalves da Silva G, Pistori H, Theophilo Folhes M. Weed detection in soybean crops using ConvNets. *Comput Electron Agric*. 2017;143:314-24.

Chebrolu N, Lottes P, Schaefer A, Winterhalter W, Burgard W, Stachniss C. Agricultural robot dataset for plant classification, localization and mapping on sugar beet fields. *Int J Robot Res*. 2017;36(10):1045-52.

Hou S, Feng Y, Wang Z. VegFru: a domain-specific dataset for fine-grained visual categorization. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV); 2017 Oct 22-29; Venice, Italy. New York: IEEE; 2017. p. 541-9.

Bargoti S, Underwood J. Deep fruit detection in orchards. In: 2017 IEEE International Conference on Robotics and Automation (ICRA); 2017 May 29 - Jun 3; Singapore. New York: IEEE; 2017. p. 3626-33.

Sa I, Chen Z, Popovic M, Khanna R, Liebisch F, Nieto J, et al. weedNet: dense semantic weed classification using multispectral images and MAV for smart farming. *IEEE Robot Automat Lett*. 2018;3(1):588-95.

Lameski P, Zdravevski E, Trajkovik V, Kulakov A. Weed Detection Dataset with RGB Images Taken Under Variable Light Conditions. In: ICT Innovations 2017; 2017 Sep 18-20; Skopje, Macedonia. Berlin: Springer; 2017. p. 112-22.

Giselsson T, Jørgensen R, Jensen P, Dyrmann M, Midtiby H. A Public Image Database for Benchmark of Plant Seedling Classification Algorithms. 2017. Available from: <https://vision.eng.au.dk/plant-seedlings-dataset/> (or appropriate repository).

Zheng YY, Kong JL, Jin XB, Wang XY, Su TL, Zuo M. Cropdeep: the crop vision dataset for deep-learning-based classification and detection in precision agriculture. *Sensors*. 2019;19(5):1058.

Lottes P, Behley J, Chebrolu N, Milioto A, Stachniss C. Joint stem detection and crop-weed classification for plant-specific treatment in precision farming. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS); 2018 Oct 1-5; Madrid, Spain. New York: IEEE; 2018. p. 8233-8.

Kestur R, Meduri A, Narasipura O. MangoNet: A deep semantic segmentation architecture for a method to detect and count mangoes in an open orchard. *Eng Appl Artif Intell*. 2019;77:59-69.

Koirala A, Walsh KB, Wang Z, McCarthy C. Deep learning for real-time fruit detection and orchard fruit load estimation: benchmarking of 'MangoYOLO'. *Precis Agric*. 2019;20(6):1107-35.

Olsen A, Konovalov D, Philippa B, Ridd P, Wood JC, Johns J, et al. DeepWeeds: a multiclass weed species image dataset for deep learning. *Sci Rep*. 2019;9:2058.

Madsen SL, Mathiassen SK, Dyrmann M, Laursen MS, Paz LC, Jørgensen RN. Open Plant Phenotype Database of Common Weeds in Denmark. *Remote Sens*. 2020;12(8):1246.

Gené-Mola JV, Rosell-Polo JR, Morros JR, Ruiz-Hidalgo J, Gregorio E. Multi-modal deep learning for Fuji apple detection using RGB-D cameras and their radiometric capabilities. *Comput Electron Agric*. 2019;162:689-98.

Gené-Mola J, Vilaplana V, Rosell-Polo JR, Morros JR, Ruiz-Hidalgo J, Gregorio E. KFuji RGB-DS Database: Fuji Apple Multi-Modal Images For Fruit Detection With Color, Depth And Range-Corrected IR data. *Data Brief*. 2019;25:104289.

Chouhan SS, Kaul A, Singh UP, Jain S. A database of leaf images: practice towards plant conservation with plant pathology. 2020. (Provide repository or DOI if available).

Santos TT, de Souza LL, dos Santos AA, Avila S. Grape detection, segmentation, and tracking using deep neural networks and three-dimensional association. *Comput Electron Agric*. 2020;170:105247.

Häni N, Roy P, Isler V. MinneApple: a benchmark dataset for apple detection and segmentation. *IEEE Robot Automat Lett*. 2020;5(2):852-8.

Yi J, Krusenbaum L, Unger P, Hüging H, Seidel S, Schaaf G, et al. Deep learning for non-invasive diagnosis of nutrient deficiencies in sugar beet using RGB images. *Sensors*. 2020;20(20):5893.

Lu Y, Young SN. A survey of public datasets for computer vision tasks in precision agriculture. *Comput Electron Agric*. 2020;178:105760.

Barbedo J. Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant disease classification. *Comput Electron Agric.* 2018;153:46-53.

Wenjin L, Dumais ST, Sun Y, Zhang H, Czerwinski M, Field BA. Semi-automatic image annotation. *Interact.* 2001;1(1):326-33.

Wilkinson M, Dumontier M, Aalbersberg IJ, Appleton G, Axton M, Baak A, et al. The FAIR Guiding Principles for Scientific Data Management and Stewardship. *Sci Data.* 2016;3:160018.

Busto PP, Iqbal A, Gall J. Open set domain adaptation for image and action recognition. *IEEE Trans Pattern Anal Mach Intell.* 2020;42(2):413-29.

Buda M, Maki A, Mazurowski M. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks.* 2018;106:249-59.

Cicco MD, Potena C, Grisetti G, Pretto A. Automatic model based dataset generation for fast and accurate crop and weeds detection. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS); 2017 Sep 24-28; Vancouver, Canada. New York: IEEE; 2017. p. 5188-95.

Chandra A, Desai SV, Guo W, Balasubramanian V. Computer Vision with Deep Learning for Plant Phenotyping in Agriculture: A Survey. *arXiv.* 2020; abs/2006.11391.