

## Research Article

**Cite this article:** Schmidt-Lebuhn AN, Bell M, Eckelmann C, Evans D, Glanznig A, Li R, Mitchell A, Mitchell-Storey T, Newton M, O'Duibhir L, Southerton R, Thomas E, and Wu H (2024). WeedScan, a weed reporting system for Australia using an image classification model for identification. *Invasive Plant Sci. Manag.* doi: [10.1017/inp.2024.19](https://doi.org/10.1017/inp.2024.19)

Received: 9 February 2024

Revised: 15 May 2024

Accepted: 11 June 2024

### Associate Editor:

Catherine S Jarnevich, United States Geological Survey

### Keywords:

Computer vision; identification; surveillance; weed management


### Corresponding author:

Alexander N. Schmidt-Lebuhn;  
Email: [alexander.schmidt-lebuhn@csiro.au](mailto:alexander.schmidt-lebuhn@csiro.au)

© The Author(s), 2024. Published by Cambridge University Press on behalf of Weed Science Society of America. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.



# WeedScan, a weed reporting system for Australia using an image classification model for identification

Alexander N. Schmidt-Lebuhn<sup>1</sup> , Matt Bell<sup>2</sup>, Carsten Eckelmann<sup>3</sup>, Dane Evans<sup>2</sup>, Andreas Glanznig<sup>4</sup>, Rongxin Li<sup>5</sup>, Andrew Mitchell<sup>6</sup>, Tomas Mitchell-Storey<sup>7</sup>, Michael Newton<sup>8</sup>, Liam O'Duibhir<sup>3</sup>, Richard Southerton<sup>9</sup>, Emily Thomas<sup>10</sup> and Hanwen Wu<sup>11</sup>

<sup>1</sup>Senior Research Scientist, Centre for Australian National Biodiversity Research (a joint venture of Parks Australia and CSIRO), Canberra ACT, Australia; <sup>2</sup>Software Developer, 2pi Software, Bega, NSW, Australia; <sup>3</sup>Director, 2pi Software, Bega, NSW, Australia; <sup>4</sup>CEO, Centre for Invasive Species Solutions, University of Canberra, Bruce, ACT, Australia; <sup>5</sup>Senior Research Scientist, CSIRO Data61, Marsfield, NSW, Australia; <sup>6</sup>Former Science Leader Automated Weed Identification Project, Centre for Invasive Species Solutions, University of Canberra, Bruce, ACT, Australia; <sup>7</sup>WeedScan National Coordinator, Centre for Invasive Species Solutions, University of Canberra, Bruce, ACT, Australia; <sup>8</sup>Director, NewtonGreen Technologies, The Junction, NSW, Australia; <sup>9</sup>Former Field and Photography Technician, Centre for Australian National Biodiversity Research (a joint venture of Parks Australia and CSIRO), Canberra, ACT, Australia; <sup>10</sup>State Priority Weed Coordinator, New South Wales Government Department of Primary Industries and Regional Development, Albury, NSW, Australia and <sup>11</sup>Principal Research Scientist, New South Wales Government Department of Primary Industries and Regional Development, Wagga Wagga, NSW, Australia

## Abstract

Fast and efficient identification is critical for reducing the likelihood of weed establishment and for appropriately managing established weeds. Traditional identification tools require either knowledge of technical morphological terminology or time-consuming image matching by the user. In recent years, deep learning computer vision models have become mature enough to enable automatic identification. The major remaining bottlenecks are the availability of a sufficient number of high-quality, reliably identified training images and the user-friendly, mobile operationalization of the technology. Here, we present the first weed identification and reporting app and website for all of Australia. It includes an image classification model covering more than 400 species of weeds and some Australian native relatives, with a focus on emerging biosecurity threats and spreading weeds that can still be eradicated or contained. It links the user to additional information provided by state and territory governments, flags species that are locally reportable or notifiable, and allows the creation of observation records in a central database. State and local weed officers can create notification profiles to be alerted of relevant weed observations in their area. We discuss the background of the WeedScan project, the approach taken in design and software development, the photo library used for training the WeedScan image classifier, the model itself and its accuracy, and technical challenges and how these were overcome.

## Introduction

Weeds compete with cultivated plants such as crops and displace native vegetation, in extreme cases transforming entire ecosystems (Brandt et al. 2021). In Australia, weeds are estimated to cause annual losses of ca. AU\$5 billion to the economy (McLeod 2018). They also have major impacts on biodiversity and threatened species, with invasion of northern Australia by introduced grasses and establishment of escaped garden plants listed as Key Threatening Processes under Commonwealth environmental laws (DCCEEW 2021).

Many weed species introduced to Australia continue to spread into new, previously unaffected areas or may only begin to do so after significant lag times (Konowalik and Kolanowska 2018; Osunkoya et al. 2021). More than 5900 nonnative plant species in Australia have weed histories overseas and are likely to naturalize in Australia given the right conditions (Randall 2007). This large reservoir of potential future weeds is the source for about 20 new introduced plants naturalizing in the environment each year (Dodd et al. 2015). Other plant species are absent from Australia but are known to have the potential for significant adverse impacts should they ever establish, for example, field horsetail (*Equisetum arvense* L.). Therefore, efficient and reliable identification and reporting are critical if infestations are to be discovered early enough to enable local eradication or containment.



### Management Implications

WeedScan is the first Australian weed identification and reporting system. It enables users to identify more than 400 weeds using an image classification model and comparison against illustrated species profiles. It also connects the user to weed management information provided by the Australian states and territories and Weeds Australia and informs them if a given weed is reportable or notifiable in their jurisdiction. At the press of a button, the user can create an observation record for the WeedScan database. Weed officers can set up alert profiles to be notified whenever weeds of interest are reported from their area. This allows them to check the image attached to the record and to visit its coordinates to eradicate or contain the infestation. Efficient identification and early notification of weeds will increase the likelihood of timely management action and will reduce weed impacts. As observation records collected by WeedScan users accumulate, the data will increase our understanding of weed occurrences and the dynamics of their spread.

Efficient weed identification is constrained by the limited number of weed experts and taxonomists (Morton 2008) and the trade-offs inherent in different identification tools used by nonexperts. Text-based identification keys require the user to have at least some understanding of specialist morphological terminology, or if they employ vernacular terminology, risk being imprecise and potentially confusing to more expert users. Picture guides are necessarily restricted in their taxonomic coverage. They also require either lengthy browsing while users try to match example photos to the weeds they are trying to identify or sufficient taxonomic expertise allowing users to narrow down the species they need to compare before even starting the identification process.

Deep learning and computer vision are increasingly applied to species identification (Wäldchen and Mäder 2018a, 2018b). However, there is a noticeable gap in the development of systems specifically tailored for weed management. While many potentially relevant models or training sets are published in academic papers (Olsen *et al.* 2019; Rai *et al.* 2023; Tang *et al.* 2017; Wang *et al.* 2022; Wu *et al.* 2021), these studies often lack a mechanism for operationalization, that is, for getting models into the hands of end-users.

The most impactful operationalized examples of species identification using computer vision are mobile apps targeted at the biodiversity-interested public, for example, Seek by iNaturalist ([https://www.inaturalist.org/pages/seek\\_app](https://www.inaturalist.org/pages/seek_app)), PictureThis (<https://www.picturethisai.com>), or PlantSnap (<https://www.plantsnap.com>). The scope of the models integrated into these applications is broad but also opportunistic, in the sense that they were trained to identify only the species for which a sufficient number of training images happened to be available, as opposed to a comprehensive target list of species. Because these apps were not originally intended to be used for the reporting of biosecurity threats, they also do not provide straightforward functionality for alerting responsible authorities of weed infestations and doing so confidentially. However, efforts have been made to build alert systems into databases with citizen science contributions such as the Atlas of Living Australia (<https://www.ala.org.au/biosecurity>).

Here, we present WeedScan, a new combined weed identification and reporting system for Australia. It uses an image classification model (MobileNetV2; Sandler *et al.* 2018) for

identification, allows users to create weed records to be shared with a central database, links them to government weed management information for identified weeds, and notifies weed officers of any new records in their jurisdiction of weed species for which they created opt-in notification profiles. Users can also create WeedScan Groups to share weed records within a group and encourage collective action.

The WeedScan project was inspired by two existing smart-phone apps and one prototype. FeralScan is a system allowing users to exchange information on feral animals in Australia and to coordinate management actions such as baiting (<https://www.feralscan.org.au>). It benefits from the ease of identification of feral animals (e.g., deer or rabbits). The WeedWise app maintained by the New South Wales state government allows users to report infestations, prompting weed officers to visit the reported geolocation and eradicate the weed where indicated (<https://weeds.dpi.nsw.gov.au>). However, WeedWise lacks an identification function, requiring users to already know what weeds they need to report and how to recognize them. Finally, in 2018, Australia's national science agency CSIRO co-invested with Microsoft to produce an iOS app prototype embedding an image classification model for the mobile identification of biosecurity threats. It also created geocoded and time-stamped observation records at the press of a button. The prototype sparked a collaboration between CSIRO and the Australian Government Department of Agriculture, Fisheries and Forestry on mobile identification of biosecurity threats at the border.

In designing WeedScan, the project partners built on these prior tools to combine their various functions identified as most relevant to weed management. The goal was to build the first national weed identification and reporting system that would use computer vision to aid the users with identification, allow them to create observation records, alert weed officers of priority weed records in their areas of responsibility, connect users to weed management information, and connect users with each other. The present paper summarizes the approach taken in the design and software development of WeedScan, the collation of training images, the selection of model architecture for and training of the image classification model, and what can be learned for similar future projects operationalizing identification through computer vision and for the future of WeedScan itself. We also provide information on the scope, accuracy, and limitations of WeedScan's image classification model.

## Materials and Methods

### Priority List

State and territory governments of Australia nominated priority weeds to be included in the identification and reporting mechanisms of WeedScan. Most nominations are declared weeds relevant to each state. Some nominated weeds are rare in or absent from Australia, with the rationale that reporting them through the WeedScan app could lead to early intervention and successful eradication or containment, making them thus of interest to biosecurity authorities at the regional, state, or national level.

Other nominations constituted common and widespread weeds such as capeweed [*Arctotheca calendula* (L.) Levyns] and sowthistle (*Sonchus oleraceus* L.), which are of interest to farmers and agricultural consultants. Their inclusion also facilitated field testing of the app during development. The final priority list comprised 336 entries. However, the number of distinct priority taxa was 325 (Supplementary Material 1), because some of them

were taxonomically nested (e.g., an entire genus and one or several individual species of the same genus).

### Training Images

Scarcity of reliably identified, high-quality training images is frequently a major barrier to the training of computer vision models (Olsen et al. 2019; Wang et al. 2022). Between 2020 and 2023, we built a library of reliably identified weed images from a variety of sources. These images included preexisting, personal photo collections of team members, with the most significant contributed by coauthor AM; photos taken by coauthors during dedicated fieldwork for the WeedScan project in New South Wales (2021, 2022), Tasmania (2022), Northern Territory (2022), Queensland (2022), South Australia (2022), and Victoria (2023); photos taken by coauthor ANS-L during travel in Europe; and image donations from project partners such as states and territories of Australia, nongovernment organizations, and the public.

During dedicated fieldwork in Australia, photos were taken using a Panasonic DC-FZ10002 and a Panasonic DC-G9. Photos were taken to ideally include only a single species of interest and exclude off-target material, including, where possible, by reducing the depth of field to blur out the background of the image. Photos were taken from different orientations but only from the position of a standing or kneeling observer, to produce training images appropriate to the intended application of the WeedScan end-user pointing a smartphone at an unknown weed. Photos covered only mature weed specimens, either flowering, fruiting, or sterile, but not seedlings.

Despite the breadth of contributions to our image library, it was difficult to obtain sufficient training images of sufficient quality for some of the prioritized weed species due to their absence from or rarity in Australia, because they were accordingly underrepresented among plants encountered during fieldwork and in domestic image donations.

To supplement the project's own training images for these difficult-to-obtain weeds, we therefore queried iNaturalist for WeedScan priority species. To automate the process, we used an R 4.2.1 (R Core Team 2016) script published by Guillermo Huerta Ramos ([https://github.com/ghuertaramos/Inat\\_Images](https://github.com/ghuertaramos/Inat_Images), accessed: July 2022), which in turn makes use of the rinat library (Barve and Hart 2023). We restricted our queries to Research Grade images, but nonetheless obtained more than 137,000 images. However, it is not so much an individual image as the observation to which it belongs that is Research Grade, and consequently many images were unsuitable, for example, out-of-focus images, photos showing a mixture of plants, habitat photos, or even photos of people and field vehicles. We therefore rigorously filtered the images over subsequent months, retaining only ca. 30% to 40% of the images of each species that fulfilled our criteria for inclusion in training of the image classification model.

At the end of the project, our image library comprised approximately 195,000 weed images, of which 120,207 were selected for the final training set, for a median of 177 and a mean of 246 images per class.

### Image Classification Model

Our image classification models were trained using the MMClassification package (MMClassification Contributors 2020) on the CSIRO high-performance computing cluster Bracewell (CSIRO 2022). Due to limitations in memory and processing power on our target deployment platforms, we assessed the suitability of several lightweight networks at the beginning of our project, namely

MobileNetV2 and MobileNetV3, which include the configurations MobileNetV3-Large and MobileNetV3-Small (Howard et al. 2019; Sandler et al. 2018). These lightweight networks prioritize reducing model size and computational cost while maintaining good accuracy, making them suitable for mobile and embedded devices with resource constraints. Although significantly larger than those lightweight networks, ResNet-18 (He et al. 2015) was also tested as a performance baseline.

MobileNetV2 and MobileNetV3-Large were chosen to be further evaluated with experiments, along with ResNet-18. The experiments were performed in the initial phase of the project. At that stage, we had collected 8,952 images belonging to 18 species. We used 7,161 images for training and the remainder for validation and testing.

With MobileNetV3-Large and ResNet-18, we achieved top-1 accuracies of 20.0% and 53.7%, respectively, after 50 epochs of training with 224 iterations per epoch (losses at the last epoch of training were 2.62 and 1.56, respectively). Such poor results indicated that they were unlikely to be suitable for our data and task.

MobileNetV2 introduced improvements over the original MobileNet in the architecture design, such as inverted residuals and linear bottlenecks. It offers various hyperparameter configurations to trade off between model size and performance. Despite the reported improvements of MobileNetV3 over MobileNetV2, our preliminary evaluation suggested that the models' relative performances of MobileNetV2 and MobileNetV3 were likely task dependent. In addition, the size of the MobileNetV2 model was acceptable for our deployment target (i.e., standard smartphones). We therefore reached a decision to employ MobileNetV2.

During the course of the project, despite the drastically increased difficulty due to increasing similarities between classes, our incrementally improved models based on the MobileNetV2 architecture continued to perform well.

Pretrained models can be used to shorten the training process, because they already "understand" image features such as edges and corners. A MobileNetV2 model trained on the dataset for the ImageNet Large Scale Visual Recognition Challenge 2012 (Russakovsky et al. 2015) was adopted as our pretrained model. This dataset, often referred to as ImageNet-1, was released in 2012 and contains 1,281,167 training images and 100,000 test images belonging to 1,000 object categories such as animals, cars, or household items (Krizhevsky et al. 2012; Russakovsky et al. 2015). In contrast to other versions, each image has a single label associated with it.

Our images were center cropped, resized to 480 by 480 pixels, and normalized to the same means and standard deviations of the three color channels of the ImageNet dataset. Training images were augmented by random translation and random horizontal flip. We did not use color augmentation or rotation or vertical flip, because the dataset included species with zygomorphic flowers. Ten percent of the training library was used as the test set.

We fine-tuned the key parameters of the model. The Stochastic Gradient Descent optimizer was used to minimize the cross-entropy loss function, including a momentum factor of 0.9 and a weight decay factor of 0.000025. The initial learning rate was set to 0.025 and was reduced by 2% at each epoch. It is worth noting that the weight decay, also known as L2 regularization, was significantly smaller than commonly used values. We evaluated performance based on top-1 and top-5 accuracy, that is, respectively, the percentage of test images for which the highest-ranked suggestion of the model was correct and the percentage of test images for which the correct answer was among the five highest-ranking suggestions.

With the 488 species classifier (our final model, see Results and Discussion), the model's evolution was deemed to have converged at the 220th epoch. However, it was allowed to run to our predetermined maximum iterations of 300 epochs.

During the course of the project, the model was expanded iteratively, starting with 18 classes, followed by 57, 136, 257, and 488 in the final model iteration. This allowed us to test the integration of the model into the WeedScan website and app and conduct field testing throughout. The trained model was a file in Python Pickle format that was first converted into the cross-platform format ONNX (<https://github.com/onnx/onnx>) and then into ONNX Runtime (ORT) to optimize it for use in mobile devices.

In addition to validation during training, and to assess whether the final WeedScan model performed comparably to other compact, mobile species identification models, we benchmarked it against the “small inception tf1” CoreML model included in Seek v2.9.1 (<https://github.com/inaturalist/SeekReactNative/releases/tag/v2.9.1-138>, accessed: July 26, 2023). Because the two models serve complementary purposes and have accordingly different sets of classes, they share only 40 common species. We assembled a test set of 109 images that had not been part of either model's training library and represented 27 of the shared species (available on the CSIRO Data Access Portal at DOI: [10.25919/rcj3-5m69](https://doi.org/10.25919/rcj3-5m69)). Tests of both models were automated using Python scripts that read all images from a directory. In both cases, we noted only whether the top-1 identification suggestion was correct or not.

### Website and Database

The WeedScan web application was developed in Microsoft Visual C#, Razor Pages, Dapper, and Microsoft .NET 6. All data are stored in a managed Microsoft SQL Server database. Both the web application and database are hosted in Microsoft Azure. The application and database were developed by New South Wales Government Department of Primary Industries and Regional Development (NSW DPIRD) and extended and maintained by NewtonGreen Technologies. The WeedScan web application exposes secure, authenticated, Representational State Transfer APIs that are used by the WeedScan mobile applications and to share WeedScan data with Atlas of Living Australia (ALA) and others.

### Mobile App

The WeedScan smartphone app was developed by 2pi Software (Bega, Australia) in React Native, a cross-platform language that ensured that only one code base would have to be managed for both iOS and Android. Integration of the image classification model, however, required relevant functions to be written in the device-native languages Java and Swift. Early development focused on achieving this integration, followed by weed record creation and communications with the WeedScan backend.

Most early testing was done on Samsung phones and tablets running Android, because that operating system allows easier deployment of prototypes as Android Package Kit installation files. Since March 22, 2023, the WeedScan app has also been tested on iOS through Apple's TestFlight app. Whenever a new app version was completed, it was provided first to the CSIRO team for systematic evaluation against a comprehensive list of functions to be tested and then to the broader stakeholder group to solicit feedback.

### Technical Challenges

During the course of the project, the team faced a variety of challenges, particularly in terms of managing a large image library, weed taxonomy, and software compatibility. Team members therefore adapted or developed a set of Python scripts to assist with tasks such as validating the image classification model with a folder of test images or efficiently checking thousands of donated images for corrupted file headers and special characters in file names that would disrupt model training.

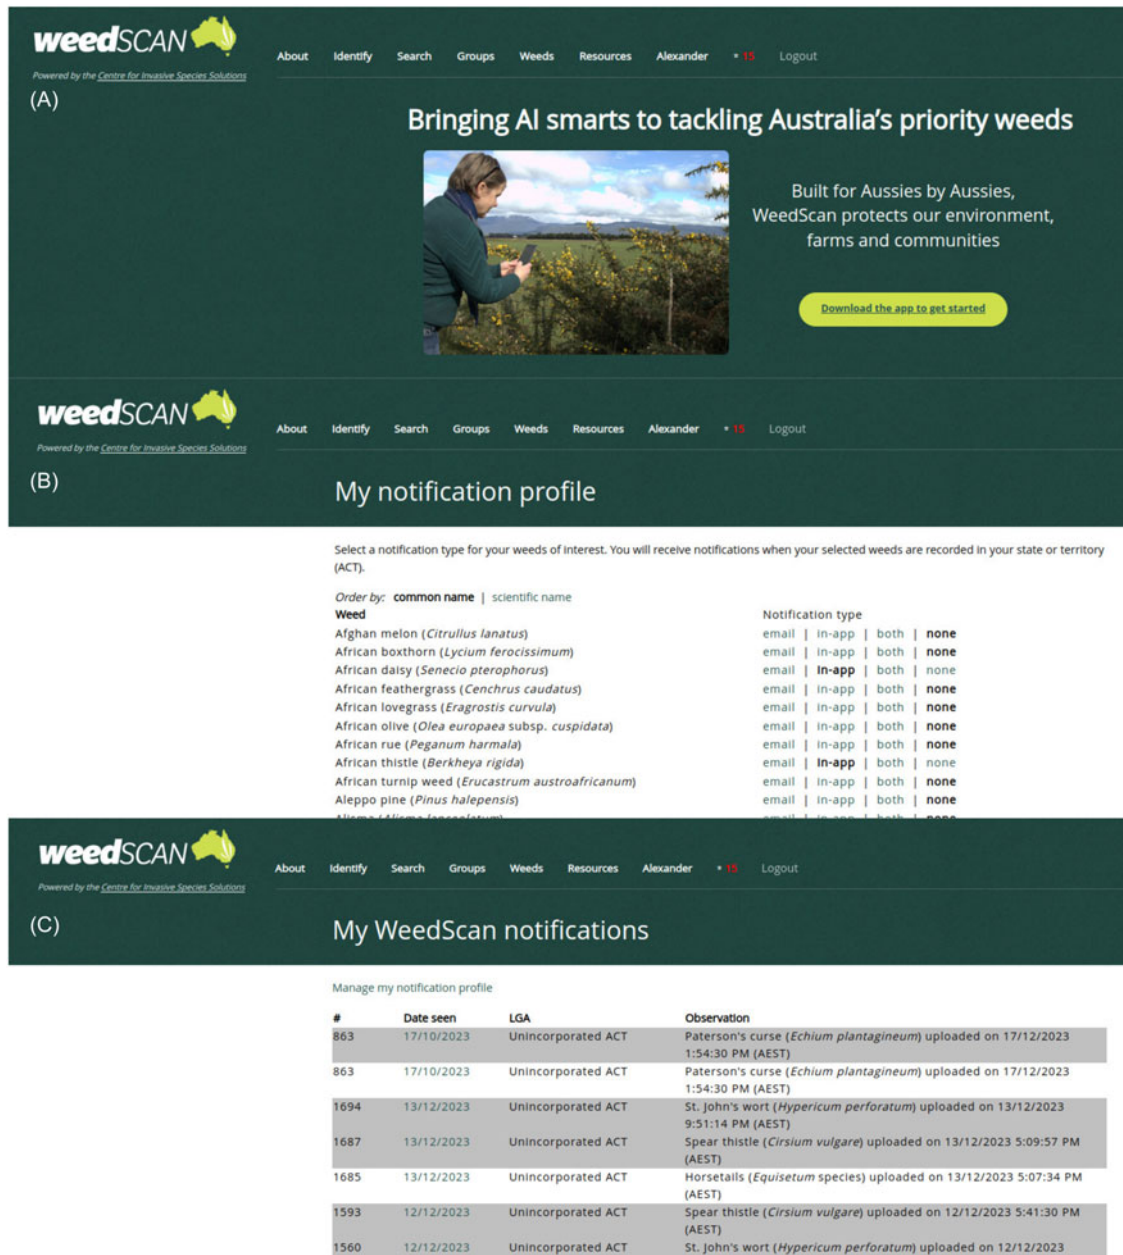
Of particular note was an upgrade of the ONNX Runtime library in October 2022, which unexpectedly was not backward compatible and therefore required the image classification model to be upgraded. Therefore, 2pi Software produced a Docker image that contains all required software dependencies at the required version to facilitate the conversion of the trained model into the ORT format required by the WeedScan app (<https://github.com/Centre-for-Invasive-Species-Solutions/pytorch-to-ort-model-converter>, accessed: June 13, 2023).

### End-User Interviews and Workshops

WeedScan governance surveys were conducted between February and March 2021. Forty-nine WeedScan governance users were interviewed to determine the desirable features for WeedScan, user types, barriers to reporting weeds, organizational use of WeedScan data/records, and the prioritization of weed groups for inclusion. The information collected was used as the backbone for building a WeedScan web application prototype. Since then, the prototype website has been repeatedly refined in response to feedback and testing by the project team. A model testing page was built in the prototype website in the administration section for people to test the integration of the image classification model. An interface was designed to allow constant uploading of new versions of the model and their integration with the WeedScan database.

The WeedScan prototype series of 17 user consultation workshops was conducted by NSW DPIRD between January and September 2022 across 16 Natural Resource Management (NRM) regions in all Australian states and territories except Western Australia. The WeedScan concept was presented, and participants were asked to complete a questionnaire either via SurveyMonkey (<https://www.surveymonkey.com>) or paper hard copy to solicit information on what functionality users would require the system to have. Feedback forms were used to query how to refine the WeedScan prototype website and guide development of the smartphone app.

A second series of 12 training workshops were conducted between February and March 2023 to train participants how to use the WeedScan website and app across 11 NRM regions covering all states and territories of Australia. This training workshop series was to demonstrate how to use the WeedScan website and smartphone app to identify and record priority weeds and, for weed officers, how to set up notification profiles to be alerted of relevant records submitted by other users. The training consisted of two parts, with the first part covering presentations and demonstrations of the WeedScan concept and the features of WeedScan such as weed identification/recording, searching weed records and profiles, groups, and weed notifications. The second part was the hands-on practical application of the WeedScan website and smartphone app using the provided



**Figure 1.** Selected screenshots of the WeedScan website. (A) Landing page with link to app stores. (B) Weed officers can set up notification profiles to specify what weed species they want to be informed of when users create records of them in their area. (C) Notification list of a weed officer. Notifications marked in white have been read.

photos of priority weeds and sometimes live plant samples brought by the participants. At the end of the workshops, participants were asked to provide feedback on WeedScan after the training and hands-on practice and suggest how it can be further improved.

## Results and Discussion

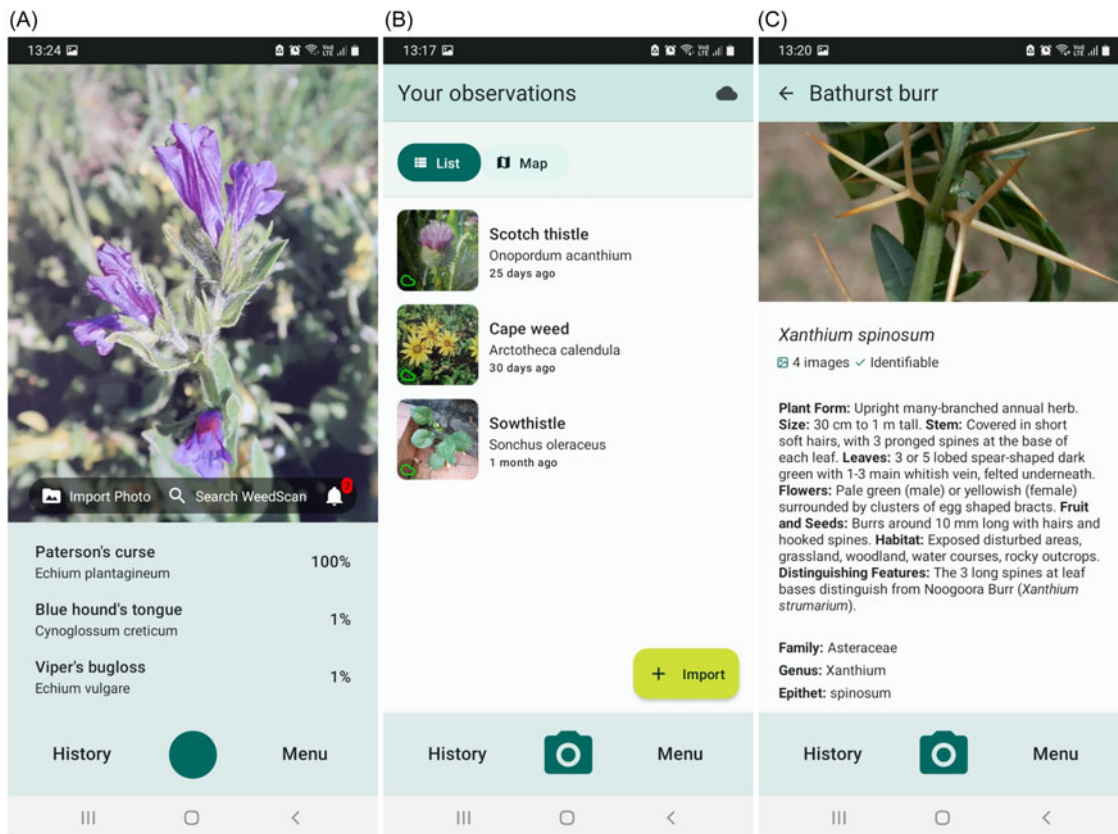
### Image Classification Model

The final WeedScan image classifier comprising 488 classes achieved a top-1 accuracy of 95.38% and a top-5 accuracy of 99.09% after 300 epochs of training. In the benchmarking test on the overlap between WeedScan and Seek, the WeedScan model achieved a top-1 accuracy of 93.58%, and the Seek model achieved

90.83%, demonstrating comparable accuracy across the limited overlap in species between the two models that could be tested.

### Website and Database

The WeedScan web application is publicly available at <https://weedscan.org.au>. The landing page (Figure 1A) provides a short summary of the system and an introductory video. On the identification and record creation interface, the user can upload a single static photo and receive an identification suggestion from the WeedScan model. The user can then create a weed record, attach additional images to it, and provide the record's geolocation by moving a pin on a map. Records can be submitted as public or private, with only administrators able to access private records. Registered users can add comments to a record and override the



**Figure 2.** Selected screens of the WeedScan mobile app. (A) Interactive identification screen showing constantly updated suggestions at the bottom. A weed record can be created instantly by pressing the button at the bottom of the screen. (B) List of records created by the user. (C) Partial view of weed profile that can be used to check an identification suggestion for correctness. The bottom of the profile shows links to additional information on state government websites or Weeds Australia. Example images showing different parts of the plant can be accessed by tapping the main image and swiping.

identification suggestion by selecting any other species from the list covered by WeedScan. Any user can set a record to private, meaning that it is only visible to administrators or weed officers, not to the public.

The WeedScan website's Search interface retrieves all public records for a species or a geographic area specified by the user. Records can be displayed in tabular format, on a map, or both. Users can also search the weed list and access weed profiles and links to information provided by Australian states and territories without first identifying a weed. Finally, they can create and join Groups of other WeedScan users that allow the sharing of observations, for example, between members of a volunteer group based around a local nature reserve.

Additional website functions are available to narrower user groups such as administrators. Weed officers can sign up to receive notifications of weeds recorded by the public (Figure 1B) and can examine their notifications lists and directly access the records that triggered the alerts (Figure 1C).

### Smartphone App

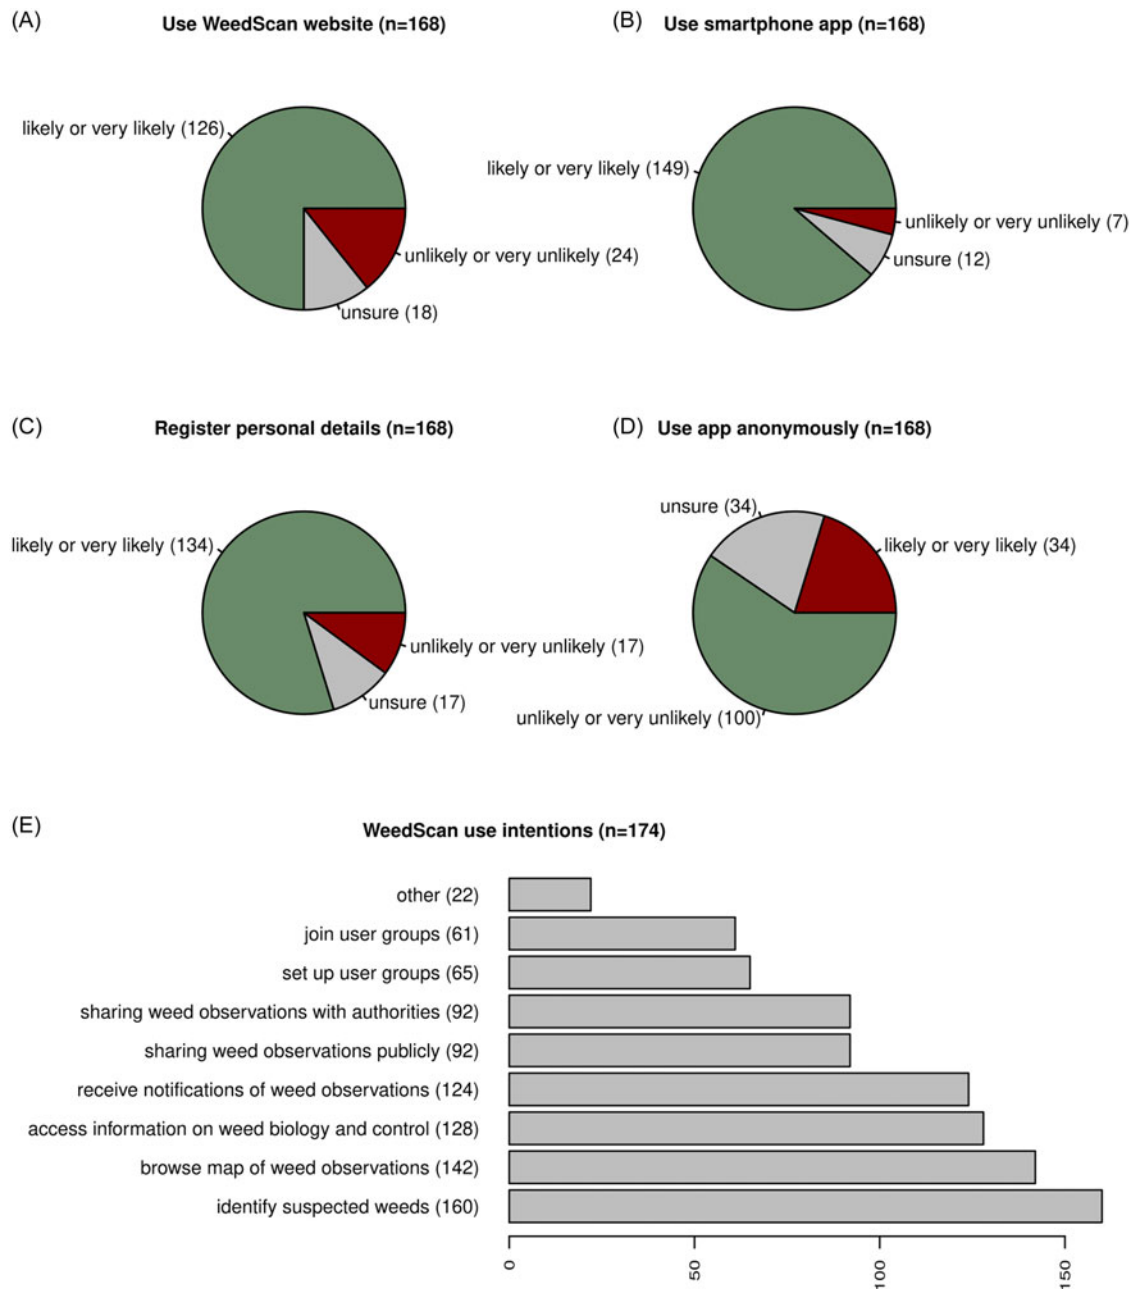
The mobile WeedScan app was published to the App Store in May 2023 (<https://apps.apple.com/au/app/weedscan/id6446067750>) and to the Google Play Store in June 2023 (<https://play.google.com/store/apps/details?id=com.weedscan>).

The app provides three options for the identification of weeds through the image classification model: (1) Interactive mode,

where frames from the device's video feed are provided to the model and displayed identification suggestions change in real time (Figure 2A). This allows the user to explore how suggestions and probabilities change when the model is shown different parts of the plant, angles, or light conditions. (2) Native camera mode, where a single, static but higher-quality image is requested from the device's built-in camera app. (3) Image import mode, allowing users to identify photos from their devices, that is, stored in their image galleries outside WeedScan.

At the press of a button, a draft weed record is created that includes an identification estimate, geolocation, and date and time. The app provides three different options for inferring geolocation coordinates while creating a weed record. In interactive mode, the current location of the device is used. When creating a record from a photo taken by the built-in camera app or from one stored in the device's gallery, the user can either select a location on a map or extract coordinates from the image metadata, if present.

Like the WeedScan website, the app allows registered users to add optional comments to their weed records and override the image classification model's identification suggestion or flag a record as private before finalizing it. The app also includes interfaces for searching the WeedScan database and displaying results in list or map views, for displaying the weed list and individual weed profiles with links to additional information on state or territory websites (Figure 2C), and for displaying notifications to users who are weed officers.



**Figure 3.** Selected results of end-user consultation workshops that informed the design of WeedScan. (A) Intention of potential end-users to access WeedScan through its website. (B) Intention to access WeedScan through the smartphone app. (C) Willingness of potential end-users to register a user account with personal details such as an email address. (D) Interest in ability to create weed observation records anonymously. (E) Expression of interest in various use cases of WeedScan.

As the image classification model is built into the mobile app, identification works even when the user is outside network reception. The app stores a user’s own weed records locally and uploads them to the WeedScan database when connection is reestablished. This also allows syncing of records a user created while logged in between different devices (Figure 2B). Users can locally delete their own records to free up space in the app.

**End-User Consultation and Feedback**

A total of 227 participants attended the first series of user consultation workshops. Of these, 178 (78.4%) completed the questionnaire. Feedback was reviewed by NSW DPIRD and

synthesized into a report informing the development of WeedScan. Selected responses are visualized in Figure 3.

Most surveyed users were likely or highly likely to use WeedScan to identify weeds (92%) and create records that capture the exact location of a weed (89%). However, some were concerned about their privacy, with 81% of respondents indicating they would create a WeedScan account or make a record that could result in follow-up contact from government weed authorities. There was also some reluctance to record weeds located on one’s own property, with only 61% of respondents likely to do so, possibly due to concerns over repercussions. The ability to use WeedScan anonymously, by not registering or logging in, satisfies the privacy concerns of some hesitant users, with 20% of respondents likely to

use this feature. Additionally, messaging informing users that a new or important weed has been identified would encourage 91% of people surveyed to create a record.

Approximately 35% of questionnaire respondents indicated that they were interested in setting up or joining a WeedScan Group to share weed observations. Group functionality was thought to be useful for Landcare groups and to assist weed professionals with managing these groups. However, beyond this, there was no clear consensus to define what a group is and how it could be used, suggesting a multiplicity of potential applications. Therefore, WeedScan currently has a simple group feature that facilitates sharing records at a local level, which will be augmented over time. Overall, the feedback on the WeedScan concept and prototype was very positive. However, the following areas were frequently raised and discussed across the workshop series:

- the limited number of weed species included in WeedScan so far;
- accuracy of the identification model leading to potential creation of incorrect records;
- availability of resources to follow up on weed records;
- availability of resources to maintain WeedScan;
- the public's expectation for government/weed officers to action records;
- adoption of WeedScan due to "app fatigue";
- potential regulatory consequences for reporting prohibited matter;
- simplified species profile description;
- ability to document and/or record investigation/follow-up/verification of records;
- ability to modify records with incorrect species identification;
- providing guidelines on how to take good photos for accurate identification;
- sharing of records with other databases; and
- alternative formats to present and export spatial data.

A total of 213 participants attended the second series of training workshops. Most of the attendees were weed/biosecurity officers, NRM officers, researchers, Landcare members, weed spotters, and interested members of the public. Indigenous stakeholders were engaged wherever possible, as in a workshop organized in Darwin, where two First Nations people attended.

The WeedScan concept received overwhelming support from workshop participants, as it addressed the need of identification without expert knowledge and recording, reporting, and notification. Many valuable comments were received, and feedback was implemented subsequently during the refinements of both the WeedScan app and web application.

### Performance of the Image Classification Model

At 95% top-1 accuracy, the performance of our model is comparable to that of other weed identification models, including models trained on much smaller numbers of species. For example, an image classifier based on LeNet-5 and trained on 820 images representing seedlings of three weed species and soybean [*Glycine max* (L.) Merr.] achieved 92.9% accuracy (Tang *et al.* 2017). Two image classifiers, Inception-v3 and ResNet-50, trained on the DeepWeeds dataset of 17,509 images representing eight weed species and diverse off-target images were reported to have accuracies of 95.1% and 95.7%, respectively (Olsen *et al.* 2019). Three object detection models, YOLOv3, YOLOv5, and Faster R-

CNN, trained on the Weed25 dataset of 14,035 images representing 25 species achieved average accuracy of detection of 91.8%, 92.4%, and 92.2%, respectively (Wang *et al.* 2022). A broader overview of deep learning approaches to weed identification and their performance was provided by Hu *et al.* (2021).

WeedScan is not intended to compete with other mobile identification apps but to serve weed identification and reporting in Australia. We conducted a benchmarking test against the Seek model only to assess whether we had achieved accuracy comparable to that of other compact models used in mobile apps. Another caveat of the benchmarking test is that the two models have very different complements of classes. This both restricted the number of species we were able to conduct the test on and introduced potential biases where classes were taxonomically and visually closer in one model but sparsely covered in the other. For example, WeedScan includes only bigleaf periwinkle (*Vinca major* L.), whereas Seek includes both that species and similar common periwinkle (*Vinca minor* L.). This means that WeedScan is unlikely to make mistakes when presented only with images of shared *V. major*, whereas Seek may confuse the two. Common viper's bugloss (*Echium vulgare* L.) and salvation jane (*Echium plantagineum* L.) represent an example of the more common inverse scenario, where WeedScan has to differentiate both species, but Seek only knows the former, as its scope is all of macroscopic life, and it therefore covers fewer weedy vascular plants than WeedScan. Despite these caveats, the test demonstrated that the WeedScan model performs approximately in line with the one mobile species identification model with which we were able to make a direct comparison.

A key challenge with the use of the WeedScan image classification model by the public is the management of expectations. Some weeds can only be reliably identified (even by the best taxonomist) if key characters are visible, which may depend on life stage. This means that even weed species that the model has been trained on cannot always be identified correctly, for example, if they are seedlings, because the training library only contained images of mature specimens. Users also have to become familiar with the camera view required for good identification results in terms of how close they have to be to the weed of interest, the need to avoid strong light/shade contrasts, and so on.

A broader problem is that of false positives; some weed species have similar-looking native or cultivated relatives. Although some of these were included in the training library, it was not feasible to do so for all such species, especially in larger genera such as *Senecio* L., which includes dozens of native Australian species. Although not as concerning as false negatives, which may lead to infestations being overlooked until the time window for eradication has closed, naively accepted false positives have the potential to frustrate users, waste the time of weed officers, and misdirect management actions.

It is therefore critical that additional lines of defense are used to guard against false positives. In the case of WeedScan, the user identifying and potentially recording weeds is urged to check weed descriptions and example images that are displayed when tapping a species name, for example, in the list of identification suggestions in the mobile app's interactive mode or on the website. The user should create a record only when convinced that the identification makes sense and, if logged in, can also override the model's suggestion. As a third check, weed officers who receive notifications can examine the attached image to confirm whether the model's identification is plausible. Together, these checks can minimize wasted effort.

One concern expressed by participants in the end-user consultation workshops was the relatively limited number of weed



species included in the first iteration of WeedScan. We envisage that expansion of the scope will be a priority in future development stages. The image classification model already covers more species than nominated for the priority list, many of them common weeds included to reduce the risk of false positives.

### Conclusions and Future Directions

WeedScan is Australia's first computer vision weed identification and reporting product and system. Nationwide consultations highlighted strong demand for this product, and consequently, we expect solid uptake by users such as weed officers and Landcare and community group members. One month after its launch, the app had been downloaded more than 10,000 times (AG, personal observation, January 31, 2024). As this iteration is a Minimum Viable Product, it will be upgraded in the future to add new weed species to the AI model, add new features of the WeedScan Groups element, and enable WeedScan to better interact with other weed citizen science products, such as the Atlas of Living Australia's weed biocontrol hub that enables citizens to record observations of biocontrol agents for 29 weed species.

**Supplementary material.** To view supplementary material for this article, please visit <https://doi.org/10.1017/inp.2024.19>

**Data availability statement.** The WeedScan image classification model and benchmarking image set are available on the CSIRO Data Access Portal at DOI: [10.25919/rcj3-5m69](https://doi.org/10.25919/rcj3-5m69). The training image library can be made available for research purposes upon reasonable request to the Centre for Invasive Species Solutions.

**Acknowledgments.** The project team acknowledges the contributions of Elissa Van Oosterhout, Pete Turner, Stephen Johnson, and Peter West (NSW DPIRD) and Richard Price (CISS) to the success of WeedScan. The project team is grateful for photo donations from the governments of New South Wales, Victoria, and the Northern Territory, from the Kimberley Weed Spotters Network, and from R. G. and Fiona J. Richardson (Specialists in Plant Publications), and to John Hosking for providing photos of *Leucanthemum vulgare* for the benchmarking test. Graham Charles and Bernie Dominiak reviewed a pre-submission version of the article.

**Funding statement.** The WeedScan project was funded by the Australian government through the National Landcare Program, and its foundational partners are the Centre for Invasive Species Solutions, CSIRO, the Atlas of Living Australia, and the governments of New South Wales, Queensland, South Australia, and Victoria. WeedScan is also supported by the Tasmanian Government and Meat & Livestock Australia.

**Competing interests.** The authors declare no conflicts of interest.

### References

[DCCEEW] Australian Government Department for Climate Change, Energy, the Environment and Water Weeds in Australia (2021) Weeds in Australia. <https://www.dccceew.gov.au/environment/invasive-species/weeds>. Accessed: January 19, 2024

Barve V, Hart E (2023) rinat: Access "iNaturalist" Data Through APIs. [github.com/ropensci/rinat](https://github.com/ropensci/rinat). Accessed: July 6, 2022

Brandt AJ, Bellingham PJ, Duncan RP, Etherington TR, Fridley JD, Howell CJ, Hulme PE, Jo I, McGlone MS, Richardson SJ, Sullivan JJ, Williams PA, Peltzer DA (2021) Naturalised plants transform the composition and function of the New Zealand flora. *Biol Invasions* 23:351–366

CSIRO (2022) Bracewell Delivers Deep Learning. <https://www.csiro.au/en/research/technology-space/it/bracewell>. Accessed: July 20, 2023

Dodd AJ, Burgman MA, McCarthy MA, Ainsworth N (2015) The changing patterns of plant naturalization in Australia. *Divers Distrib* 21:1038–1050

He K, Zhang X, Ren S, Sun J (2015) Deep residual learning for image recognition. arXiv:1512.03385

Howard A, Sandler M, Chu G, Chen L-C, Chen B, Tan M, Wang W, Zhu Y, Pang R, Vasudevan V, Le QV, Adam H (2019) Searching for MobileNetV3. arXiv:1905.02244

Hu K, Wang Z, Coleman G, Bender A, Yao T, Zeng S, Song D, Schumann A, Walsh M (2021) Deep learning techniques for in-crop weed identification: a review. arXiv:2103.14872

Konowalik K, Kolanowska M (2018) Climatic niche shift and possible future spread of the invasive South African Orchid *Disa bracteata* in Australia and adjacent areas. *PeerJ* 6:e6107

Krizhevsky A, Sutskever I, Hinton GE (2012) ImageNet classification with deep convolutional neural networks. *Adv Neural Inf Process Syst* 25. [https://papers.nips.cc/paper\\_files/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45b-Abstract.html](https://papers.nips.cc/paper_files/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45b-Abstract.html). Accessed: February 2, 2024

McLeod R (2018) Annual Costs of Weeds in Australia. Centre for Invasive Species Solutions. Sydney, NSW, Australia: eSYS Development. [invasives.com.au/wp-content/uploads/2019/01/Cost-of-weeds-report.pdf](https://invasives.com.au/wp-content/uploads/2019/01/Cost-of-weeds-report.pdf). Accessed: January 19, 2024

MMClassification Contributors (2020) OpenMMLab's Image Classification Toolbox and Benchmark. [github.com/open-mmlab/mmlclassification](https://github.com/open-mmlab/mmlclassification). Accessed: July 18, 2023

Morton SJ (2008) A review of the National Weed Detection project. Pages 340–342 in *Proceedings of the 16th Australian Weeds Conference*, Cairns Convention Centre, North Queensland, Australia, May 18–22, 2008

Olsen A, Konovalov DA, Philippa B, Ridd P, Wood JC, Johns J, Banks W, Girgenti B, Kenny O, Whinney J, Calvert B, Azghadi MR, White RD (2019) DeepWeeds: a multiclass weed species image dataset for deep learning. *Sci Rep* 9:2058

Osunkoya OO, Lock CB, Dhileepan K, Buru JC (2021) Lag times and invasion dynamics of established and emerging weeds: insights from herbarium records of Queensland, Australia. *Biol Invasions* 23:3383–3408

Rai N, Zhang Y, Ram BG, Schumacher L, Yellavajjala RK, Bajwa S, Sun X (2023) Applications of deep learning in precision weed management: a review. *Comput Electron Agric* 206:107698

Randall RP (2007) The Introduced Flora of Australia and Its Weed Status. Department of Agriculture and Food, Western Australia. Glen Osmond, SA: CRC for Australian Weed Management. [weeds.org.au/wp-content/uploads/2023/05/2007.-The-introduced-flora-of-Australia-and-its-weed-status.pdf](https://weeds.org.au/wp-content/uploads/2023/05/2007.-The-introduced-flora-of-Australia-and-its-weed-status.pdf). Accessed: January 19, 2024

R Core Team (2016) R: A Language and Environment for Statistical Computing. Vienna: R Foundation for Statistical Computing. [www.r-project.org](http://www.r-project.org). Accessed: June 23, 2022

Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, Huang Z, Karpathy A, Khosla A, Bernstein M, Berg AC, Fei-Fei L (2015) ImageNet large scale visual recognition challenge. arXiv:1409.0575

Sandler M, Howard A, Zhu M, Zhmoginov A, Chen L (2018) MobileNetV2: efficient mobile convolutional neural networks for mobile vision applications. Pages 4510–4520 in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. [http://arxiv.org/abs/1801.04381](https://arxiv.org/abs/1801.04381). Accessed: July 18, 2023

Tang J, Wang D, Zhang Z, He L, Xin J, Xu Y (2017) Weed identification based on K-means feature learning combined with convolutional neural network. *Comput Electron Agric* 135:63–70

Wäldchen J, Mäder P (2018a) Machine learning for image based species identification. *Methods Ecol Evol* 9:2216–2225

Wäldchen J, Mäder P (2018b) Plant species identification using computer vision techniques: a systematic literature review. *Arch Comput Methods Eng* 25:507–543

Wang P, Tang Y, Luo F, Wang L, Li C, Niu Q, Li H (2022) Weed25: a deep learning dataset for weed identification. *Front Plant Sci* 13:1053329

Wu Z, Chen Y, Zhao B, Kang X, Ding Y (2021) Review of weed detection methods based on computer vision. *Sensors* 21:3647