THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Finely-grained annotated datasets for image-based plant phenotyping

Citation for published version:

Minervini, M, Fischbach, A, Scharr, H & Tsafaris, SA 2016, 'Finely-grained annotated datasets for image-based plant phenotyping', *Pattern Recognition Letters*, vol. 81, pp. 80-89.
<https://doi.org/10.1016/j.patrec.2015.10.013>

Digital Object Identifier (DOI):

[10.1016/j.patrec.2015.10.013](https://doi.org/10.1016/j.patrec.2015.10.013)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Pattern Recognition Letters

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Graphical Abstract

Finely-grained annotated datasets for image-based plant phenotyping

Massimo Minervini, Andreas Fischbach, Hanno Scharr, Sotirios A. Tsaftaris

In this paper we present a collection of benchmark datasets for the development and evaluation of computer vision and machine learning algorithms in the context of plant phenotyping. We provide annotated imaging data and suggest suitable evaluation criteria for plant/leaf segmentation, detection, tracking as well as classification and regression problems. The Figure symbolically depicts the data available together with ground truth segmentations and further annotations and metadata.

Research Highlights

- First comprehensive annotated datasets for computer vision tasks in plant phenotyping.
- Publicly available data and evaluation criteria for eight challenging tasks.
- Tasks include fine-grained categorization of age, developmental stage, and cultivars.
- Example test cases and results on plant and leaf-wise segmentation and leaf counting.

Finely-grained annotated datasets for image-based plant phenotyping

Massimo Minervini^{a,**}, Andreas Fischbach^b, Hanno Scharf^b, Sotirios A. Tsaftaris^{a,c}

^aPattern Recognition and Image Analysis Research Unit, IMT Institute for Advanced Studies, 55100 Lucca, Italy

^bInstitute of Bio- and Geosciences: Plant Sciences (IBG-2), Forschungszentrum Jülich GmbH, 52425 Jülich, Germany

^cInstitute for Digital Communications, School of Engineering, The University of Edinburgh, Edinburgh EH9 3JL, UK

ABSTRACT

Image-based approaches to plant phenotyping are gaining momentum providing fertile ground for several interesting vision tasks where fine-grained categorization is necessary, such as leaf segmentation among a variety of cultivars, and cultivar (or mutant) identification. However, benchmark data focusing on typical imaging situations and vision tasks are still lacking, making it difficult to compare existing methodologies. This paper describes a collection of benchmark datasets of raw and annotated top-view color images of rosette plants. We briefly describe plant material, imaging setup and procedures for different experiments: one with various cultivars of *Arabidopsis* and one with tobacco undergoing different treatments. We proceed to define a set of computer vision and classification tasks and provide accompanying datasets and annotations based on our raw data. We describe the annotation process performed by experts and discuss appropriate evaluation criteria. We also offer exemplary use cases and results on some tasks obtained with parts of these data. We hope with the release of this rigorous dataset collection to invigorate the development of algorithms in the context of plant phenotyping but also provide new interesting datasets for the general computer vision community to experiment on. Data are publicly available at <http://www.plant-phenotyping.org/datasets>.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

The study of the phenotype expressed by cultivars (or mutants) of the same plant species under different environmental conditions is central to our understanding of plant function. Identifying and evaluating a plant's actual phenotype, is relevant to, e.g., seed production and plant breeders. The phenotype relies on the fine-grained categorization of a plant's properties: e.g., how many leaves, of which architecture, visual age or maturity level, to which cultivar a plant is similar.

Previously, such categorization was annotated manually by experts, but recently image-based approaches are gaining momentum. In the last decades several approaches have been proposed to measure visual traits of plants in an automated fashion (Granier et al., 2006; Walter et al., 2007; Jansen et al., 2009; Nagel et al., 2012) together with customized image processing pipelines to analyze aspects of the acquired image data (Augustin et al., 2015; De Vylder et al., 2012; Hartmann et al.,

2011; van der Heijden et al., 2012; Walter et al., 2007; Yin et al., 2014; Müller-Linow et al., 2015). However, several of the computer vision tasks encountered (for example leaf segmentation and counting) are particularly challenging and remain unsolved. In fact, most experts now agree that lack of reliable and automated algorithms to analyze these vast datasets forms a new bottleneck in our understanding of plant biology and function (Minervini et al., 2015b). We must accelerate the development and deployment of such computer vision algorithms, since according to the Food and Agriculture Organization of the United Nations (FAO), large scale experiments in plant phenotyping are a key factor in meeting agricultural needs of the future, one of which is feeding 11 billion people by 2050.

One of the factors that could accelerate the development of better algorithms and their consistent and systematic evaluation is the availability of benchmark data focusing on typical imaging situations and tasks in plant phenotyping. Till now, despite the 20 year history of imaging plants, a comprehensive collection of benchmark datasets for image-based non-destructive plant phenotyping is still lacking, with the only exception of the recent multi-modality dataset in (Cruz et al., 2015). While plant

^{**}Corresponding author.

e-mail: m.minervini@imtlucca.it (Massimo Minervini)

Fig. 1: Closeups of examples plants. Left column: Young plants with few leaves. Middle: Rosettes with elongated petioles (Arabidopsis). Right: Compact rosettes with or without rich structure within single leaves (tobacco).

related datasets exist for leaf or flower recognition (Goëau et al., 2012; Silva et al., 2013; Wu et al., 2007; Nilsback and Zisserman, 2010), these datasets were obtained in an uncontrolled or destructive manner and not in a phenotyping context.

Here we present a collection of raw and annotated images of the most frequently used rosette model plants (Arabidopsis and tobacco). We briefly describe plant material, environmental conditions, and imaging setup and procedures, which led to the collection of color images showing top-down views. Shape variability in these images is high, cmp. Figure 1.

They show compact plants, plants with long petioles, overlapping and non-overlapping leaves, different leaf numbers, etc., such that a wide range of shapes typically found in young dicotyledon plants is covered (cmp. e.g. growth stages ‘1.’ or ‘30’ of oilseed rape, sunflower, or beet in (Meier, 2001)). Using these raw images we describe a collection of datasets, with appropriate expert annotations and metadata, for a series of computer vision tasks. We emphasize tasks such as plant/leaf detection, segmentation, and tracking, leaf counting, boundary estimation, and general regression and classification. Figure 2 shows examples of such annotations. For consistency in evaluation methodologies we also discuss appropriate criteria. To offer exemplary use cases we also show results on plant and

Fig. 2: Example annotations superimposed on an image of Arabidopsis (Col-0, wild-type), acquired 17 days after germination. Shown are: the plant bounding box (green), leaf bounding boxes (yellow), the plant mask (outlined in purple), three leaf masks (blue overlays), and five leaf centers (red dots).

leaf segmentation and counting. While not all raw images have been annotated, our annotation is continuing and the datasets size is increasing. We devised and released an annotation tool (Minervini et al., 2015a) in order to facilitate future annotation. This article will serve as the reference point for describing the process and tasks for which the data can be used.

This article is organized as follows: Section 2 presents data collection; Section 3 presents datasets and annotations for a variety of vision tasks; Section 4 outlines exemplary use cases and results; and Section 5 offers discussion and conclusions.

2. Imaging setup and data

Non-invasive plant investigations are performed using various modalities with spatial scales varying from the microscopic to large outdoor fields (Chavarria-Krauser et al., 2008; Bergsträsser et al., 2015; Müller-Linow et al., 2015). Typical problems in measuring a plant’s visible properties comprise measuring size, shape, color or spectral reflection, and other structural and functional traits of whole plants, their organs, or plant populations. Plants are not static, but self-changing systems with complexity in shape and appearance increasing over time. For example, a single leaf appears in the scene and can grow several cm^2 , i.e. orders of magnitude change in size till it stops growing. Leaves can grow rapidly, in the range of hours, with the whole plant changing over days or even months, in which the surrounding environmental (as well as measurement) conditions may also change. Biologists grow model plants, such as Arabidopsis (*Arabidopsis thaliana*) and tobacco (*Nicotiana tabacum*), in controlled environments and monitor and record the phenotype. Such experiments are fundamental and ubiquitous, and recovering the phenotype implies that computer vision algorithms must deal with the complexity of the plant, the experiment and the environmental conditions.

Starting from such typical experiments, we devised imaging apparatuses and setups to acquire three imaging datasets (Fig-

(a) Ara2012

(b) Ara2013 (Canon)

(c) Tobacco

Fig. 3: Example raw images (left) and close-ups showing individual plant subjects (right) of: (a) *Arabidopsis* (Col-0, wild-type), (b) 5 different cultivars of *Arabidopsis*, and (c) tobacco undergoing different treatments. (a)-(b) *Arabidopsis* and (c) tobacco were imaged with the setups shown in Figure 4 and 5, respectively.

ure 3), which form the basis of this work. They were acquired in two different labs with highly diverse equipment, as shown in Figures 4 and 5. *Arabidopsis* images were acquired using two setups for investigating affordable hardware for plant phenotyping. Tobacco images were acquired using a robotic imaging system for the investigation of automated plant treatment optimization by an artificial cognitive system. Images differ in resolution, fidelity and scene complexity, with plants appearing in isolation or in trays, and with subjects belonging to different cultivars (mutants), or undergoing different treatments (see examples in Figure 3). The vision tasks required to estimate several phenotyping parameters include detection, segmentation, and tracking across time, for plants and individual leaves.

Prior to discussing the datasets (Section 3) and the computer vision tasks they can be used for (Section 4), we briefly describe below how plants were grown and imaged, with additional details available in Scharr et al. (2014).

2.1. *Arabidopsis* imaging with an affordable setting

Arabidopsis images were acquired in two data collections: in June 2012 and in September-October 2013, hereafter named

Fig. 4: Affordable acquisition setup for ‘Ara2012’ and ‘Ara2013’.

Ara2012 and *Ara2013*, respectively, both consisting of top-view time-lapse images of *Arabidopsis thaliana* rosettes. Experiment composition is summarized in Table 1. Plants were grown in individual pots with randomized arrangement to eliminate possible bias in the results due to variations in watering or lighting conditions. No treatments were performed.

The imaging setup (cf. Figure 4) consisted of a growth shelf and an automated affordable sensing system (Tsaftaris and Noutsos, 2009), to acquire and send images of the scene via a wireless connection to a receiving computer (<http://www.phenotiki.com>). Example images are shown in Figure 3a-b, illustrating plant arrangement and scene complexity.

Plants were illuminated artificially and controllably to emulate daylight imaging in a fixed day cycle. Camera sensors were positioned between the lights, approximately 1 m above the plants, and operated as intervalometers, i.e. acquiring images at preset times with preset imaging conditions (such as focus, exposure, field of view). Images were captured during day time only for a period of time (cf. Table 1). Two cameras were used: a 7 Mpixel consumer grade camera (Canon PowerShot SD1000), shorthand as *Canon*, and an even lower cost system based on the Raspberry Pi (<http://www.raspberrypi.org>), Model B, with the Raspberry Pi 5 Mpixel camera module, shorthand as *Rpi*. *Ara2012* used only Canon but *Ara2013* used both cameras, permitting the evaluation of algorithms on images of the same scene but with different image quality. Additionally for *Ara2013*, multiple focus images were acquired. All acquired images (width $\times$ height for Canon: 3108 $\times$ 2324 pixels; for Rpi: 2592 $\times$ 1944 pixels) were stored as lossless PNG files.

2.2. Tobacco image dataset

The tobacco images were acquired in the context of the European project ‘Gardening with a cognitive system’ (GARNICS, <http://www.garnics.eu>). The GARNICS project aimed at 3D sensing of plant growth and building perceptual representations for learning the links to actions of a robot gardener. This robot must deal with the inherent complexity of plants, which changes over time. Actions performed at plants (like watering), will have strongly delayed effects, making monitoring and controlling plants a difficult perception-action problem.

2.2.1. Imaging setup

The setup allowed for looking at plants from different poses (see Figure 5). We release top views of the plants only, as top

Table 1: Summary of information of the Arabidopsis and tobacco raw imaging data.

Experiment	Subjects	Wild-types	Mutants	Period	Total images	Image resolution	Field of View	Plant resolution
Arabidopsis (Ara2012)	19	Col-0	No	3 weeks	150	7 MPixel	large	0.25 MPixel
Arabidopsis (Ara2013, Canon)	24	Col-0	Yes (4)	7 weeks	4186	7 MPixel	large	0.25 MPixel
Arabidopsis (Ara2013, Rasp. Pi)	24	Col-0	Yes (4)	7 weeks	1951	5 MPixel	large	0.06 MPixel
Tobacco (23.01.2012)	20	Samsun	No	18 days	34560	5 MPixel	single	same
Tobacco (16.02.2012)	20	Samsun	No	20 days	38400	5 MPixel	single	same
Tobacco (15.05.2012)	20	Samsun	No	18 days	34560	5 MPixel	single	same
Tobacco (10.08.2012)	20	Samsun	No	30 days	57600	5 MPixel	single	same

Fig. 5: Hardware setup of the GARNICS robot gardener. (A) The robot arm, in the lab environment. (B) Camera head with illumination and watering system. (C) Workspace with light, temperature, and humidity sensors.

view images are used in many plant screening applications. Notice, however, that a setup performing only this type of acquisition can be considerably simpler and more affordable than the setup used here. Beneath the robot and environment sensors the implemented system featured: a) a watering and nutrient solution dispensing system for treatment application; b) high power white LED illumination (switched off for imaging); and c) low power white fluorescence illumination used for imaging.

The robot head consisted of two stereo camera systems (4× Point-Grey Grasshopper, 2448×2048, pixel size 3.45 μm), black-and-white and color, and high quality lenses (Schneider Kreuznach Xenoplan 1.4/17-0903). We added lightweight white and NIR LED light sources to the camera head. Using this setup, each plant was imaged separately from different but fixed poses. In addition, for each pose small baseline stereo image pairs were captured using each single camera by a suitable robot movement, allowing for 3D reconstruction of the plant.

Data released here stems from experiments aiming at acquiring training data for the robot gardener. Images were acquired every hour in a 24/7 manner for up to 30 days. More details on the four experiments can be found in Table 1. Overview images

of the plants from these experiments are shown in Figure 3c.

2.2.2. Plant material and growing conditions

Tobacco plants (*Nicotiana tabacum* cv. *Samsun*) were grown in 7×7 cm pots under constant light conditions with a 16h/8h day/night rhythm. Water was provided by the robot system every two hours. Nutrients were applied either manually twice a week or every other hour by the robot system. In GARNICS treatments were selected to produce training data for a cognitive system. The amounts of water and nutrient solution are therefore well adapted to the soil substrate such that the plants show distinguishable performance of generally well growing plants. Finding an optimal treatment was left to the system.

Images show growth stages from germination well into the leaf development stage, i.e. starting observations at growth stage 09 and stopping at stage 1009 to 1013 (according to the extended BBCH-scale presented in (CORESTA, 2009)), due to size restrictions. Figure 3c shows images of the final growth stages. For further details on treatments, environmental conditions, and acquisition times we refer to Scharr et al. (2014).

3. Annotated datasets

In order to provide specialized standalone datasets for a number of computer vision tasks, parts of the raw image data described in Section 2 have been carefully annotated by experts. Note that not all data have been annotated (thousands of raw images are available) and annotation is a continuing process. However, even augmented versions of the datasets in the future will follow the structure described here. Among those that have been annotated, not all are publicly available, to permit future competitions and challenges based on data which are blind to participants. Other computer vision competitions, such as PASCAL VOC (Everingham et al., 2010), and biologically inspired challenges follow this strategy. In addition, for each specific task the datasets are considered standalone and as such the user is unaware of other information (e.g., mutant type for leaf segmentation). This limits domain and prior knowledge and should lead to methodologies that are more robust to the changing and complex morphology of plants.

In the following paragraphs we describe first the semantic hierarchy considered, the manual annotation procedure, and then proceed in detailing available subsets of the data for each computer vision and machine learning task, defining annotations and appropriate metadata, and suitable evaluation criteria.

Fig. 6: Hierarchy of relationships among data, metadata, and annotations. In parentheses we provide examples of annotation variables, and we also provide pictorial examples of imaging data and annotations such as segmentation masks, bounding boxes, and leaf boundaries. Gray boxes denote metadata.

Dataset size refers to the current state of annotation as to the current date.

3.1. Overview of semantic hierarchy

Each experiment has generated a vast amount of imaging data with Arabidopsis experiments showing tray images whilst tobacco individual plants. Our internal database and annotation strategy follows the hierarchy visible in Figure 6. These original images are higher in our semantic hierarchy.

Gray boxes in Figure 6 denote related annotated metadata: experiment type, mutant type, camera used, acquisition time, experimental treatment, segmentation difficulty, etc. Non-shaded boxes denote imaging and image level annotations. Note that an experiment may contain both tray and individual plant images such as Arabidopsis for example. However, this is not a rule: for example, for tobacco datasets tray images are not available, and for Arabidopsis experiments no treatment was performed. To construct each of the standalone datasets described below, we trace information in this hierarchy and provide related metadata and annotations wherever appropriate.

3.2. Expert segmentations

A significant number of object-based annotations, e.g., bounding boxes, can be obtained computationally on the basis of pixel-level segmentation masks of plants and leaves, respectively, which have been manually annotated by experts. Here we describe how we obtained the latter and next we detail the level of annotation for each task.

Annotation consisted of three steps. First, we obtained a binary segmentation of the plant objects in the scene in a computer-aided fashion. For Arabidopsis, we used the approach based on active contours described by Minervini et al. (2014), while for tobacco, a simple color-based approach for

(a) Ara2012

(b) Ara2013 (Canon)

(c) Tobacco

Fig. 7: Examples of single plant images at different developmental stages with the corresponding ground truth leaf labeling denoted by color.

plant segmentation was used. The result of this segmentation was manually refined using raster graphics editing software, to ensure that all the visible part of the shoot is included in the plant mask and that the background (earth, moss, etc.) is excluded. Next, within the binary mask of each plant, we delineated individual leaves (including both the petiole and the blade) completely manually. A pixel with black color denotes background, while all other colors are used to uniquely identify leaves of the plants in the scene. Across the frames of the time-lapse sequence, we consistently used the same color code to label occurrences of the same leaf. To reduce observer variability and increase accuracy, the labeling process involved always two annotators: one annotating the dataset and one inspecting the other. For future extensions of the datasets, the annotation of additional images is supported by a tool that we recently released for semi-automated leaf segmentation and annotation (Minervini et al., 2015a). Figure 7 shows examples of plant images from the datasets, with corresponding pixel level annotation masks.

On a secondary inspection of the data, additional categorical *qualitative annotations* were recorded by annotators such as: estimate of segmentation difficulty (in the 1[easy]-5[hard] scale), plant appears in focus, leaves appear in vertical positions which is typical in tobacco (due to the so called nastic movements), plant is occluded by another one (when pots are placed close by), and scene contains complexities (water in the background, green moss on soil, debris or damage on leaves).

3.3. Computer vision tasks and datasets

Fine-grained information to be extracted from images is ubiquitous in plant phenotyping, since we do have to deal with how different mutants or treatments affect plant shape and characteristics. Even the same plant will have leaves of different shape and size according to their maturity, as it can also be seen in Figure 7. In some cases the phenotype of a mutant is not known, and typically researchers assign qualitative characteristics, which is simple when gross phenotype differences are evident (e.g., radically different leaf shape). However, it is when such differences are subtle that the ability to extract fine-grained information directly from images would make a tremendous impact: it would permit biologists to identify small traits to be explored further. Computer vision and machine learning could enable biologists to assess a new mutant’s phenotype by evaluating how similar it is to known lines of cultivars and mutants in a quantitative fashion.

Doing this in a fully automated fashion is fertile ground for a series of interesting vision tasks, which we outline below together with descriptions of appropriate for the task datasets and evaluation criteria. To test and develop robust solutions for a realistic range of conditions, we constructed the datasets by carefully choosing images according to our qualitative annotations (Section 3.2), and also allowing several challenging situations to occur by design in the experiments. For the evaluation criteria we took inspiration from standard criteria used in computer vision and also devised some that are relevant to phenotyping. Since each task is different and the number of available annotated images is increasing, we refrain from supposing a certain split into training, validation and testing sets—a procedure typically adopted when evaluating learning-based approaches. However, our recent experience in organizing the leaf segmentation (Section 3.3.3) and leaf counting (Section 3.3.5) contests (see Section 4) using a subset of the annotated data presented here showed that the number of available images and criteria used allowed computer vision scientists to work successfully on these problems (see e.g. (Scharf et al., 2015; Pape and Klukas, 2014, 2015; Giuffrida et al., 2015)).

Our annotations were initialized by computationally post-processing the expert-delineated leaf masks as discussed in Section 3.2. Afterwards, everything was verified (and if necessary corrected) by experts to ensure high quality and integrity. Note that for each task we outline the size (as in current number of annotated images) of the corresponding dataset.

Hereafter, to ease presentation we will denote as *leaf mask* the pixel-level leaf masks (including the visible part of a leaf blade and the corresponding petiole) and as *plant mask* the pixel-level binary mask obtained by the union of all individual leaf labels within a plant. To reduce storage all images are lossless compressed PNG files.

3.3.1. Plant detection and localization

Task: As Figure 3 shows, plants can be arranged in a grid, either fixed in position in specialized trays or not and can even be touching each other, the latter leading to a non trivial detection problem. Thus, a first vision task encountered is plant detection and localization of individual plants, in the form of

bounding boxes. Excitingly, this problem is also encountered even outside the context of phenotyping, e.g., in precision agriculture for detecting crops (van Henten et al., 2009) albeit in a more complicated setting. Here learning-based object detection approaches are ideally suited to help.

Dataset: We derived bounding boxes for each plant as the smallest bounding box enclosing the plant mask, processing each plant individually in a tray image. An additional 5% of the box size was considered to include a larger aspect of the scene. The dataset consists of 76 tray images, as 22 from Ara2012, 27 from Ara2013 (Canon), and 27 from Ara2013 (Rpi). For each tray image, a comma-separated value (CSV) file is available, reporting for each plant the corner pixel coordinates of its bounding box.

Evaluation criteria: Following upon the ubiquitous PASCAL VOC challenge Everingham et al. (2010), we suggest the bounding box overlap ratio criterion $a_o = \text{area}(B^p \cap B^{gt}) / \text{area}(B^p \cup B^{gt})$, between predicted B^p and ground truth B^{gt} bounding boxes.

3.3.2. Plant segmentation

Task: Plant biomass is an important plant breeding trait because it reflects overall plant performance. In images of rosette plants usually it is measured as projected leaf area (PLA), i.e. effectively the number of plant pixels. Finding PLA translates to the segmentation of plant from background. In simple cases this can be solved by color thresholding and other unsupervised segmentation approaches (Walter et al., 2007; De Vylder et al., 2012), but when scene complexity is high (non-smooth background, non-uniform lighting, plant overlap, presence of moss on soil) sophisticated learning-based algorithms are necessary (Perina et al., 2010; Minervini et al., 2014).

Dataset: The dataset consists of 15 tray images from Ara2012, 27 from Ara2013 (Canon), and 27 from Ara2013 (Rpi). For each image a corresponding black (background) and white (foreground) mask encoded as an indexed image provides pixel-level information on the location of plant objects.

Evaluation criteria: Several segmentation criteria exist and we suggest Dice coefficient, precision, and recall, since they have been used throughout image analysis and are common in plant imaging as well (Minervini et al., 2014). Among those, the Dice Similarity Coefficient, DSC (%) = $2|P^{gt} \cap P^{ar}| / (|P^{gt}| + |P^{ar}|)$, measures the degree of overlap among ground truth P^{gt} and algorithmic result P^{ar} binary segmentation masks, where $|\cdot|$ denotes set cardinality. We also suggest the Modified Hausdorff Distance (MHD) (Dubuisson and Jain, 1994):

$$\text{MHD} = \max \left\{ \frac{1}{|P^{gt}|} \sum_{a \in P^{gt}} \min_{b \in P^{ar}} \|a - b\|, \frac{1}{|P^{ar}|} \sum_{b \in P^{ar}} \min_{a \in P^{gt}} \|a - b\| \right\},$$

where $\|\cdot\|$ denotes Euclidean distance. The MHD has large discriminatory power, is robust to noise and is easy to interpret (it can be expressed in units of length, e.g., millimeter).

3.3.3. Leaf segmentation

Task: In rosette plants when leaves are highly overlapping, PLA may not be an accurate estimator of plant biomass any more and individual leaf segmentation is necessary. When individual leaves are segmented, distributions of leaf size can

highlight the rate of growth of new leaves with respect to old ones. However, leaf segmentation, a multi-instance segmentation problem (He and Gould, 2014), is particularly challenging since most leaves within the same plant may share appearance and shape, but can also appear severely overlapping. To complicate matters even more, plant morphology changes radically between mutants, in response to treatment, and as plants grow. Self-occlusion, shadows, leaf hairs, leaf color variations, and others add complexity. Image quality is a factor as well, so low resolution and out-of-focus (as it could occur in portions of Ara2012 and Ara2013 datasets) affect leaf segmentation accuracy. On the other hand, high resolution as in Tobacco, does introduce computational challenges.

Dataset: We used leaf masks but without temporal label consistency. 120 from Ara2012, 165 from Ara2013 (Canon), and 62 from Tobacco, images of single plants appearing centered are used. For each plant, annotations are provided in the form of indexed images (a color palette is embedded for visualization) the same dimensions of the originals. We use one label per leaf, starting from ‘1’ up to the maximum number of leaves, with ‘0’ denoting background.

Evaluation criteria: Several segmentation criteria are available for comparing between ground truth and algorithmic outcomes (Mezaris et al., 2003; Unnikrishnan et al., 2007). We suggest SymmetricBestDice, the symmetric average Dice score among all objects (leaves), where for each input label the ground truth label yielding maximum Dice is used for averaging. Best Dice (BD) is defined as:

$$\text{BD}(L^a, L^b) = \frac{1}{M} \sum_{i=1}^M \max_{1 \leq j \leq N} \frac{2|L_i^a \cap L_j^b|}{|L_i^a| + |L_j^b|},$$

where $|\cdot|$ denotes leaf area (number of pixels), and L_i^a for $1 \leq i \leq M$ and L_j^b for $1 \leq j \leq N$ are sets of leaf object segments belonging to leaf segmentations L^a and L^b , respectively. SymmetricBestDice (SBD) is then:

$$\text{SBD}(L^{ar}, L^{gt}) = \min \left\{ \text{BD}(L^{ar}, L^{gt}), \text{BD}(L^{gt}, L^{ar}) \right\}, \quad (1)$$

where L^{gt} is the ground truth and L^{ar} the algorithmic result.

3.3.4. Leaf detection

Task: In an image analysis pipeline, image-based leaf detection could serve to initialize other processes (segmentation or tracking). Due to size differences, shape and appearance similarities, and heavy occlusions, leaf detection is a complex task, and can benefit from approaches in computer vision of detecting overlapping objects in medicine, transportation, and surveillance (Arteta et al., 2013; Wohllhart et al., 2013).

Dataset: On the basis of leaf masks, we extracted for each individual leaf the smallest rectangular bounding box (possibly rotated with respect to the image coordinate system) enclosing the mask of that leaf. (In our definition such box would contain both leaf petiole and blade wherever visible and applicable.) The dataset consists of individual plant images, 120 from Ara2012, 165 from Ara2013 (Canon), and 62 from Tobacco, and for each image, a CSV file storing per row the leaf index and the coordinates of each bounding box, with as many rows

as number of leaves. Note that our annotation does include the petiole (‘leaf stalk’) in Arabidopsis.

Evaluation criteria: Number of accurate detections and their accuracy evaluated with overlap measures (cf. Section 3.3.1).

3.3.5. Leaf counting

Task: From a phenotyping perspective the number of leaves is directly related to yield potential, drought tolerance, and flowering time (Granier et al., 2006; Méndez-Vigo et al., 2010). From a computer vision perspective, it can also be used to constrain leaf detection or leaf segmentation algorithms (Nieuwenhuis et al., 2013). To this date, user interaction is required and leaf count comes as a by-product of leaf segmentation. Learning-based counting techniques could help here (Lempitsky and Zisserman, 2010; Fiaschi et al., 2013). We should note that counting leaves has additional challenges w.r.t. cell/car/people counting since in these classical counting applications all objects share similar shape and size but may have different appearance (e.g., car and people have different colors) may not occlude each other (e.g., cells usually touch but do not overlap). Instead in leaf counting, leaves may heavily overlap, and have different shape and size (mature vs. younger leaves).

Dataset: On the basis of leaf masks we extracted for each leaf the distance transform weighted-center of mass and also the center of mass. When these disagree significantly (above a threshold), and if any of the centers lie outside the binary shape, it indicates a highly asymmetric leaf (e.g., due to heavy occlusion or orientation vertical to the imaging axis) and the annotator was prompted to select a center. Leaf centers in our definition are located in the middle of the visible part of leaf blades. The dataset consists of individual plant images and accompanying binary images containing the centroids for each leaf as a single pixel. (This requires larger storage but we find it more appealing than storing centers in CSV files.) Overall 120 from Ara2012, 165 from Ara2013 (Canon), and 62 from Tobacco, raw and equally numbered annotation images are provided. A CSV file listing image names and number of leaves is also provided, for convenience of approaches that solve directly the regression problem.

Evaluation criteria: Here, we suggest: a) the difference between number of leaves in algorithm’s result and ground truth $\text{DiffFGLLabels} = \#L^{ar} - \#L^{gt}$, and b) AbsDiffFGLLabels , the absolute value of DiffFGLLabels . Note that these criteria do not take into account good localization and while count maybe correct it may not correspond to actual leaves. Alternatively, count via detection measures can be adopted (see Section 3.3.4).

3.3.6. Leaf tracking

Task: Finding growth curves of individual leaves helps us understand how a plant (or a cultivar) is growing or the effects of treatments and stresses: for example, Clauw et al. (2015) found that drought differentially affects leaves. This growth curve usually follows an exponential relationship with time (Richards, 1959; Walter and Schurr, 1999), and frequent imaging can capture small differences. This implies the precise segmentation and temporal tracking of each leaf (Yin et al., 2014).

Dataset: Building upon the leaf masks, the dataset consists of sequentially numbered PNG files of raw individual im-

ages and annotations. We provide 4 stacks of 13 images each from Ara2012 and 8 stacks of 17 images each from Ara2013 (Canon). Leaf-level segmentations are provided with leaves having the same label index throughout the sequence to ensure temporal consistency. Also, we release for each image a corresponding CSV file, where bounding box definitions are provided as previously described. Note that this dataset can also be used for leaf or plant segmentation with additional temporal information for example for joint segmentation and tracking.

Evaluation criteria: We recommend the protocol of Nawaz and Cavallaro (2013), which builds on overlap criteria, and the code available from the authors. When leaves are vertical in the imaging axis (due to severe nastic movements), overlap criteria may be unable to assign proper correspondences. This may cause lack of label consistency of a leaf across time in an algorithmic result, but it can be easily seen in individual leaf growth curves. Quantitatively, they can be detected by multiple local hypothesis testing to identify structural breaks on the growth parameters (Dellen et al., 2015).

3.3.7. Boundary estimation

Task: Some approaches to multi-instance (or multi-label) segmentation rely on accurate boundary detection, which is used for example to initialize template-based models (He and Gould, 2014; Yin et al., 2014). When image contrast and resolution are adequate, for example in Tobacco, classical edge detection works sufficiently. However, when images are partially out-of-focus and of lower resolution (due to a larger field of view) as is the case of the Arabidopsis data, learning-based methods to boundary estimation (i.e. a learned edge detector as in Martin et al. (2004)) have been shown to perform better (Pape and Klukas, 2015).

Dataset: Using leaf masks, we isolated each leaf label and found its perimeter, to produce an indexed labeled image where ‘0’ is background, ‘1’ denotes a boundary between plant and background, and ‘2’ denotes a boundary between overlapping leaves (which can be more than 1-pixel thick.). This separation may facilitate the training of specialized boundary detectors. The dataset consists of pairs of plant images and these indexed images with boundary annotations, as 120 from Ara2012, 165 from Ara2013 (Canon), and 62 from Tobacco.

Evaluation criteria: Typical criteria such as precision and recall are suggested and those not penalizing small local misalignment, which are suited for evaluating performance of boundaries between leaves, e.g., the MHD (Section 3.3.2) and its learning-based simplification in Minervini et al. (2013).

3.3.8. Classification and regression

Task: While phenotyping typically occurs in forward hypothesis testing scenarios, recently reverse hypothesis and association studies have received attention and fine-grained categorization is particularly useful. In this case phenotyping traits are recorded and are correlated with genotyping information to identify relationships among them (data mining) (O’Malley and Ecker, 2010). This is particularly useful in the case of treatments, cross hybridizations, and other processes that may affect directly or indirectly (e.g., via silencing and other epigenetic functions) the genetic code of plants (To and Kim, 2014).

Therefore, given a plant image, it is of great interest to characterize the plant, i.e. plant age or development stage, find other cultivars possessing similar traits, what possible treatment it has undergone. This may be done at plant level or may be even possible at leaf level. When these characteristics are distinct this is considered as a recognition and classification problem. However, if we are interested in percentile similarities, likelihood and regression frameworks can be employed. A similar problem in the context of precision agriculture is weed vs. crop classification, whereupon the interest is to separate a valid crop from a weed in images, usually acquired by automated robotic mechanisms (Haug et al., 2014). Here we consider and provide datasets for three cases: mutant recognition (classification), age regression, and treatment recognition (classification).

Dataset: For mutant and treatment recognition we release individual plant images and a text list denoting per each row image name, genotype, treatment type. For treatment type classification 62 data from Tobacco are currently available. For mutant classification 165 from Ara2013 (Canon) and 165 from Ara2013 (Rpi) are currently available. We assume the only input to be image data with mutant and treatment type to be predicted values. For age regression, we release individual plant images and a text list denoting per each row image name, mutant type, treatment type, and the age in hours of the plant since germination. For this task 165 images from Ara2013 (Canon), 165 from Ara2013 (Rpi), and 62 from Tobacco are currently available. Notice that the age regression task when mutant or treatment information is not available or treatment is not fixed is extremely difficult since each mutant has different growth rates in response to treatment, so in some sense the algorithm must be able to infer by appearance the unknown mutant type. For simplicity, we assume that inputs to this learning problem are: images and mutant/treatment type information to predict age.

Evaluation criteria: Due to the diversity of problems considered in this category, for classification problems precision and recall criteria are recommended, and for regression problems mean absolute error and mean squared error between predicted and ground-truth measures are encouraged.

4. Examples of use cases

This section offers examples of how parts of the datasets outlined in the previous section have been used in developing and evaluating computer vision algorithms.

For plant segmentation from complicated background using parts of the Ara2012 dataset, Minervini et al. (2014) used a new multi-channel active contour based on Chan and Vese (2001) with probabilistic priors on plant appearance. The images used included examples of heavy growth of moss in the pot, water in the tray, discoloring of plants due to temporary drought, and pot movement. On average, DSC $\approx 97\%$ (cf. Section 3.3.2) was observed for segmentation accuracy when using color and texture features and a Gaussian mixture model to describe plant appearance. In the same work, rosettes were tracked over time using overlap criteria and nearest neighbor rules. Nearest neighbor rules and a k -means clustering on the Euclidean coordinates of plant pixels were also used to resolve overlapping plants. Ex-

Fig. 8: Two examples of plant and leaf segmentation. (a) Original images, (b) ground truth leaf masks, (c) leaf segmentations obtained by (Pape and Klukas, 2014), and (d) plant segmentations obtained by (Minervini et al., 2014).

amples are shown in Figure 8d, where DSC values of 98% and 97% are reported for top and bottom images, respectively.

A specially formatted part of the data presented here was released to support the Leaf Segmentation Challenge (<http://www.plant-phenotyping.org/CVPPP2014-challenge>), of the Computer Vision Problems in Plant Phenotyping (CVPPP) workshops 2014 and 2015. CVPPP 2014 was held in conjunction with the European Conference on Computer Vision (ECCV), in Zürich, Switzerland, in September 2014; CVPPP 2015 in conjunction with the British Machine Vision Conference (BMVC), in Swansea, UK, in September 2015. Plant images were considered separately, ignoring temporal correspondence. Images contained instances with well separated leaves and simple background, but also more difficult examples with many leaf occlusions, complex leaf shapes, varying backgrounds, or plant objects not well in focus. Individual plant images and ground truth segmentations were used consisting of Ara2012, Ara2013 (Canon), and Tobacco datasets. The datasets were split into training and testing sets for the challenge. Respectively: for training, 128, 31, and 27 images and corresponding annotations were released; for testing, 33, 9, and 56 images were released (ground truth was available only to the organizers). Within this challenge, solutions were evaluated on the basis of plant segmentation, individual leaf segmentation, and leaf counting with the criteria outlined in Sections 3.3.3 and 3.3.5. Table 2 summarizes overall performance among participants that completed the challenge and the leading approaches of Pape and Klukas (2014, 2015). Although plant segmentation accuracy (DSC of plant mask) was somewhat acceptable, when considering leaf segmentation (SymmetricBestDice, Eq. (1)) and leaf count criteria poor performance was observed across the algorithms in LSC 2014, indicating the challenging nature of the problem. Some of these issues are evident also in the examples of Figure 8c: compared to the ground truth, the approach of Pape and Klukas (2014) tends to over-segment (when moss is present, e.g., the top image), separate petioles, and merge leaves. Both images present $\text{AbsDiffFGLLabels} = 2$, so count is off but for different reasons. In order to distinguish between foreground-background segmentation errors and leaf-wise segmentation errors, ground truth foreground-background

masks were provided for LSC 2015. While the average SymmetricBestDice improved to 71.3%, this performance is still quite low, which indicates that for leaf-wise segmentation reliable and highly accurate tools are still not available.

A comparative study of improved versions of the algorithms devised and tested in the context of the LSC 2014 can be found in Scharr et al. (2015). Within an in-depth analysis and discussion of the results are offered.

We also explored learning directly leaf count density based on leaf center annotations, following the approach by Lempitsky and Zisserman (2010). We used a subset of 84 images from the Ara2012 dataset, with images depicting leaf occlusions, some being out of focus, and some showing moss in the pot. The overall number of leaves per plant varied from 12 to 19. Based on leaf center annotations (Section 3.3.5), the goal was to learn via appropriate losses, density functions, the integration of which provides object counts. We extracted from the green color channel dense SIFT descriptors (Lowe, 2004) in 20 of the 84 images, with 7 SIFT bins and fixed orientations. Subsequently, we quantized the SIFT space using k -means clustering to create a codebook of size 1500. Using this codebook, we learned a linear transformation of the feature representation on the codebook approximating the density function at each pixel on 64 training images (Lempitsky and Zisserman, 2010). Testing on 33 images of Ara2012 of the LSC dataset, obtained an average $\text{AbsDiffFGLLabels} = 2.36(2.9)$ (cf. Section 3.3.5). Among those, in 67% of images we observed either no counting error or the count was off by 1 or 2 leaves at most. In comparison, referring to findings of Pape and Klukas (2014) on the same testing data, counting via segmentation performed with $\text{AbsDiffFGLLabels} = 2.2(1.3)$. For the example of Figure 8, by learning, the count was off only by 1 leaf for both test images. While these results are preliminary they demonstrate the promise of learning-based leaf count estimation.

These early findings illustrate the complexity of the problems at hand. In fact, the leaf counting problem motivated a new challenge, the leaf counting challenge (LCC) for the 2015 edition of CVPPP. The leading methods are presented in (Giuffrida et al., 2015; Pape and Klukas, 2015). The approaches, while they do rely on global regression, they are different as to the features used. Pape and Klukas (2015) extract features out of the foreground mask (e.g., perimeter, area, etc.), which after feature selection are used in a regression framework to learn leaf count per image. Several regression functions are evaluated such as

Table 2: Average performance among all participants compared to the leading approaches on the Leaf Segmentation Challenge (LSC) of CVPPP 2014 and CVPPP 2015. In parentheses standard deviations among participants for the first column, and among the three different datasets for the second column. *Note that ground truth foreground-background masks were made available in LSC 2015.

	Overall LSC 2014	Leader LSC 2014	Leader LSC 2015
SymmetricBestDice [%]	48.7 (12.8)	62.6 (19.0)	71.3 (15.1)
DSC [%]	83.1 (10.9)	95.3 (10.1)	—*
AbsDiffFGLLabels	6.3 (4.36)	2.4 (2.1)	1.1 (1.2)
DiffFGLLabels	1.3 (4.4)	-1.9 (2.7)	-0.3 (1.6)

Table 3: Performance of the leading approaches on the Leaf Counting Challenge (LCC) of CVPPP 2015. In parentheses standard deviations among the three different datasets.

LCC 2015	(Giuffrida et al., 2015)	(Pape and Klukas, 2015)
AbsDiffFGLLabels	1.43 (1.51)	1.1 (1.2)
DiffFGLLabels	-0.51 (2.02)	-0.3 (1.6)

support vector regression, random forest regression, and others. Different features were used for each dataset. On the other hand, the solution proposed by Giuffrida et al. (2015) learns features directly from the data in an unsupervised fashion to build a codebook. From the images, image patches are extracted and projected on the codebook. Responses are max pooled, leading to a global per image descriptor. Leaf count is estimated from this descriptor using support vector regression. Summary results are shown in Table 3. We see that results clearly improved with respect to the best LSC 2014 results (Table 2).

5. Discussion and Conclusions

This paper describes the first ever collection of image datasets of growing Arabidopsis and tobacco plants in a plant phenotyping context, together with annotations for a series of computer vision and learning tasks at different levels of granularity. Plant phenotyping is central to the understanding of plant function and is a tool that can enable us to meet agricultural demands of the future. Computer vision and machine learning are ideally suited to help, and this collection of datasets is intended to promote the exploration of plant phenotyping problems.

As a benefit to the scientific community, we will be continuously releasing in the public domain specially formatted datasets and annotations, accompanied by appropriate functions implementing evaluation criteria. We will be accepting contributions and corrections from the scientific community in an effort to keep these data curated. This paper will serve as the reference document on the structure and importance of these datasets. In order to facilitate future annotation by us and the broad community, following ideas from other domains (Gior-dano et al., 2015), we have already developed annotation tools (Minervini et al., 2015a). Implementing in the future such tools in web-based environments would permit the crowd-sourcing of annotations (Kavasidis et al., 2014; Goff et al., 2011). We will be also investigating additional evaluation criteria as we obtain feedback from the community.

We do hope that our publicly available datasets and future augmented versions will be adopted by the broad computer vision community as well (as with the PASCAL (Everingham et al., 2010), or the biologically focused Broad Bioimage Benchmark Collection (Ljosa et al., 2012)). Our datasets can be used to learn suitable image statistics (Heiler and Schnörr, 2005), adapt and test counting algorithms with (Fiaschi et al., 2013) and without temporal information (Arteta et al., 2013; Lempitsky and Zisserman, 2010), segmentation algorithms (Chan and Vese, 2001), multi-label segmentation (Nieuwenhuis et al., 2013; He and Gould, 2014) or detection (Barinova et al., 2012) approaches, and others. Additional depth information

as can be computed from a pair of images with different focus (Favaro and Soatto, 2007) of the Arabidopsis dataset or stereo images (Biskup et al., 2007) of the Tobacco dataset, to be released in the future, may further facilitate segmentation (Dellen et al., 2011). Using the presented datasets, image-based plant phenotyping will evolve in parallel to (and benefit from) advances in computer vision, by tracking the performance of approaches referencing these data. More importantly, it will also introduce this societally important application to a wider audience.

Acknowledgments

Tobacco research has received funding from EU’s 7th Framework Programme (FP7/2007-2013) under grant no. 247947 (GARNICS). Part of this work was performed within the German-Plant-Phenotyping Network, which is funded by the German Federal Ministry of Education and Research (project identification no. 031A053). Arabidopsis research was partially supported by a Marie Curie Action: “Reintegration Grant” (grant no. 256534) of the EU’s 7th Framework Programme.

The authors would like to thank Prof. Pierdomenico Perata and his group from Scuola Superiore Sant’Anna, Pisa, Italy, for providing plant samples and instructions on growth conditions of Arabidopsis. Finally, they also thank several annotators who have contributed to this project: Fabiana Zollo, Ines Dedovic, Mario Valerio Giuffrida, and Vasileios Sevetlidis.

References

- Arteta, C., Lempitsky, V., Noble, J.A., Zisserman, A., 2013. Learning to detect partially overlapping instances, in: CVPR, IEEE. pp. 3230–3237.
- Augustin, M., Haxhimusa, Y., Busch, W., Kropatsch, W.G., 2015. Image-based phenotyping of the mature Arabidopsis shoot system, in: Computer Vision - ECCV 2014 Workshops. Springer, pp. 231–246.
- Barinova, O., Lempitsky, V.S., Kohli, P., 2012. On detection of multiple object instances using hough transforms. IEEE Transactions on Pattern Analysis and Machine Intelligence 34, 1773–1784.
- Bergsträsser, S., Fanourakis, D., Schmittgen, S., Cendrero-Mateo, M., Jansen, M., Scharr, H., Rascher, U., 2015. HyperART: non-invasive quantification of leaf traits using hyperspectral absorption-reflectance-transmittance imaging. Plant Methods 11.
- Biskup, B., Scharr, H., Schurr, U., Rascher, U., 2007. A stereo imaging system for measuring structural parameters of plant canopies. Plant, Cell and Environment 30, 1299–1308.
- Chan, T., Vese, L., 2001. Active contours without edges. IEEE Transactions on Image Processing 10, 266–277.
- Chavarria-Krauser, A., Nagel, K., Palme, K., Schurr, U., Walter, A., Scharr, H., 2008. Spatio-temporal quantification of differential growth processes in root growth zones based on a novel combination of image sequence processing and refined concepts describing curvature production. New Phytologist 177, 811–821.
- Clauw, P., Coppens, F., De Beuf, K., Dhondt, S., Van Daele, T., Maleux, K., Storme, V., Clement, L., Gonzalez, N., Inze, D., 2015. Leaf responses to mild drought stress in natural variants of Arabidopsis thaliana. Plant Physiology .
- CORESTA, 2009. A scale for coding growth stages in tobacco crops.
- Cruz, J.A., Yin, X., Liu, X., Imran, S.M., Morris, D.D., Kramer, D.M., Chen, J., 2015. Multi-modality imagery database for plant phenotyping. Machine Vision and Applications In press.
- De Vylder, J., Vandenbussche, F.J., Hu, Y., Philips, W., Van Der Straeten, D., 2012. Rosette Tracker: an open source image analysis tool for automatic quantification of genotype effects. Plant Physiology 160, 1149–1159.
- Dellen, B., Alenyà, G., Foix, S., Torras, C., 2011. Segmenting color images into surface patches by exploiting sparse depth data, in: WACV, pp. 591–598.

- Dellen, B., Scharr, H., Torras, C., 2015. Growth signatures of rosette plants from time-lapse video. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* PP.
- Dubuisson, M.P., Jain, A.K., 1994. A modified Hausdorff distance for object matching, in: 12th IAPR Int. Conf. on Pattern Recognition, pp. 566–568.
- Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A., 2010. The Pascal Visual Object Classes (VOC) challenge. *International Journal of Computer Vision* 88, 303–338.
- Favaro, P., Soatto, S., 2007. 3-D Shape Estimation and Image Restoration: Exploiting Defocus and Motion-Blur. Springer-Verlag.
- Fiaschi, L., Gregor, K., Afonso, B., Zlatić, M., Hamprecht, F.A., 2013. Keeping count: Leveraging temporal context to count heavily overlapping objects, in: International Symposium on Biomedical Imaging (ISBI), pp. 656–659.
- Giordano, D., Kavasidis, I., Palazzo, S., Spampinato, C., 2015. Nonparametric label propagation using mutual local similarity in nearest neighbors. *Computer Vision and Image Understanding* 131, 116–127.
- Giuffrida, M.V., Minervini, M., Tsafaris, S.A., 2015. Learning to count leaves in rosette plants, in: Proceedings of the Computer Vision Problems in Plant Phenotyping (CVPPP) Workshop, BMVA Press, pp. 1.1–1.13.
- Goëau, H., Bonnet, P., Joly, A., Yahiaoui, I., Barthélémy, D., Nozha, B., Molino, J.F., 2012. The ImageCLEF 2012 plant identification task, in: CLEF, pp. 1–25.
- Goff, S.A., et al., 2011. The iPlant collaborative: Cyberinfrastructure for plant biology. *Frontiers in Plant Science* 2.
- Granier, C., Aguirrezabal, L., Chenu, K., Cookson, S.J., Dauzat, M., Hamard, P., Thioux, J.J., Rolland, G., Bouchier-Combaud, S., Lebaudy, A., Muller, B., Simonneau, T., Tardieu, F., 2006. PHENOPSIS, an automated platform for reproducible phenotyping of plant responses to soil water deficit in *Arabidopsis thaliana* permitted the identification of an accession with low sensitivity to soil water deficit. *New Phytologist* 169, 623–635.
- Hartmann, A., Czauderna, T., Hoffmann, R., Stein, N., Schreiber, F., 2011. HT-Pheno: An image analysis pipeline for high-throughput plant phenotyping. *BMC Bioinformatics* 12, 148.
- Haug, S., Michaels, A., Biber, P., Ostermann, J., 2014. Plant classification system for crop/weed discrimination without segmentation, in: IEEE Winter Conference on Applications of Computer Vision (WACV), pp. 1142–1149.
- He, X., Gould, S., 2014. An exemplar-based CRF for multi-instance object segmentation, in: CVPR, IEEE, pp. 296–303.
- van der Heijden, G., Song, Y., Horgan, G., Polder, G., Dieleman, A., Bink, M., Palloix, A., van Eeuwijk, F., Glasbey, C., 2012. SPICY: towards automated phenotyping of large pepper plants in the greenhouse. *Functional Plant Biology* 39, 870–877.
- Heiler, M., Schnörr, C., 2005. Natural image statistics for natural image segmentation. *International Journal of Computer Vision* 63, 5–19.
- van Henten, E., Goense, D., Lokhorst, C. (Eds.), 2009. Precision Agriculture '09. Wageningen Academic.
- Jansen, M., Gilmer, F., Biskup, B., Nagel, K., Rascher, U., Fischbach, A., Briem, S., Dreissen, G., Tittmann, S., Braun, S., Jaeger, I.D., Metzclaff, M., Schurr, U., Scharr, H., Walter, A., 2009. Simultaneous phenotyping of leaf growth and chlorophyll fluorescence via GROWSCREEN FLUORO allows detection of stress tolerance in *Arabidopsis thaliana* and other rosette plants. *Functional Plant Biology, Special Issue: Plant Phenomics* 36, 902–914.
- Kavasidis, I., Palazzo, S., Di Salvo, R., Giordano, D., Spampinato, C., 2014. An innovative web-based collaborative platform for video annotation. *Multimedia Tools and Applications* 70, 413–432.
- Lempitsky, V.S., Zisserman, A., 2010. Learning to count objects in images, in: Advances in Neural Information Processing Systems, pp. 1324–1332.
- Ljosa, V., Sokolnicki, K.L., Carpenter, A.E., 2012. Annotated high-throughput microscopy image sets for validation. *Nature Methods* 9, 637–637.
- Lowe, D.G., 2004. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60, 91–110.
- Martin, D., Fowlkes, C., Malik, J., 2004. Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 530–549.
- Meier, U. (Ed.), 2001. Growth stages of mono- and dicotyledonous plants. 2 ed., Federal Biological Research Centre for Agriculture and Forestry.
- Méndez-Vigo, B., de Andrés, M.T., Ramiro, M., Martínez-Zapater, J.M., Alonso-Blanco, C., 2010. Temporal analysis of natural variation for the rate of leaf production and its relationship with flowering initiation in *Arabidopsis thaliana*. *Journal of Experimental Botany* 61, 1611–1623.
- Mezaris, V., Kompatsiaris, I., Strintzis, M., 2003. Still image objective segmentation evaluation using ground truth, in: COST 276 Workshop, pp. 9–14.
- Minervini, M., Abdelsamea, M.M., Tsafaris, S.A., 2014. Image-based plant phenotyping with incremental learning and active contours. *Ecological Informatics* 23, 35–48.
- Minervini, M., Giuffrida, M.V., Tsafaris, S.A., 2015a. An interactive tool for semi-automated leaf annotation, in: Proceedings of the Computer Vision Problems in Plant Phenotyping Workshop, BMVA Press, pp. 6.1–6.13.
- Minervini, M., Rusu, C., Tsafaris, S.A., 2013. Learning computationally efficient approximations of complex image segmentation metrics, in: International Symposium on Image and Signal Processing and Analysis, pp. 60–65.
- Minervini, M., Scharr, H., Tsafaris, S.A., 2015b. Image analysis: The new bottleneck in plant phenotyping. *IEEE Signal Process. Mag.* 32, 126–131.
- Müller-Linow, M., Pinto-Espinosa, F., Scharr, H., Rascher, U., 2015. The leaf angle distribution of natural plant populations: assessing the canopy with a novel software tool. *Plant Methods* 11.
- Nagel, K., Putz, A., Gilmer, F., Heinz, K., Fischbach, A., Pfeifer, J., Faget, M., Blossfeld, S., Ernst, M., Dimaki, C., Kastenholz, B., Kleinert, A.K., Galinski, A., Scharr, H., Fiorani, F., Schurr, U., 2012. GROWSCREEN-Rhizo is a novel phenotyping robot enabling simultaneous measurements of root and shoot growth for plants grown in soil-filled rhizotrons. *Functional Plant Biology* 39, 891–904.
- Nawaz, T.H., Cavallaro, A., 2013. A protocol for evaluating video trackers under real-world conditions. *IEEE Trans. Image Process.* 22, 1354–1361.
- Nieuwenhuis, C., Toeppe, E., Cremers, D., 2013. A survey and comparison of discrete and continuous multi-label optimization approaches for the Potts model. *International Journal of Computer Vision* 104, 223–240.
- Nilsback, M.E., Zisserman, A., 2010. Delving deeper into the whorl of flower segmentation. *Image and Vision Computing* 28, 1049–1062.
- O'Malley, R.C., Ecker, J.R., 2010. Linking genotype to phenotype using the *Arabidopsis* unimutant collection. *The Plant Journal* 61, 928–940.
- Pape, J.M., Klukas, C., 2014. 3-D histogram-based segmentation and leaf detection for rosette plants, in: Computer Vision - ECCV 2014 Workshops. Springer, pp. 61–74.
- Pape, J.M., Klukas, C., 2015. Utilizing machine learning approaches to improve prediction of leaf counts and individual leaf segmentation of rosette plants, in: Proceedings of the Computer Vision Problems in Plant Phenotyping Workshop, BMVA Press, pp. 3.1–3.12.
- Perina, A., Cristani, M., Murino, V., 2010. 2LDA: Segmentation for recognition, in: International Conference on Pattern Recognition, pp. 995–998.
- Richards, F.J., 1959. A flexible growth function for empirical use. *Journal of Experimental Botany* 10, 290–301.
- Scharr, H., Minervini, M., Fischbach, A., Tsafaris, S.A., 2014. Annotated image datasets of rosette plants. Technical Report FZJ-2014-03837. Forschungszentrum Jülich. URL: <http://hdl.handle.net/2128/5848>.
- Scharr, H., Minervini, M., French, A.P., Klukas, C., Kramer, D.M., Liu, X., Muntión, I.L., Pape, J.M., Polder, G., Vukadinovic, D., Yin, X., Tsafaris, S.A., 2015. Leaf segmentation in plant phenotyping: A collation study. *Machine Vision and Applications* In press.
- Silva, P.F., Marçal, A.R., da Silva, R.M.A., 2013. Evaluation of features for leaf discrimination, in: Image Analysis and Recognition, pp. 197–204.
- To, T.K., Kim, J.M., 2014. Epigenetic regulation of gene responsiveness in *Arabidopsis*. *Frontiers in Plant Science* 4.
- Tsafaris, S.A., Noutsos, C., 2009. Plant phenotyping with low cost digital cameras and image analytics, in: Information Technologies in Environmental Engineering, pp. 238–251.
- Unnikrishnan, R., Pantofaru, C., Hebert, M., 2007. Toward objective evaluation of image segmentation algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29, 929–944.
- Walter, A., Scharr, H., Gilmer, F., Zierer, R., Nagel, K.A., Ernst, M., Wiese, A., Virnich, O., Christ, M.M., Uhlig, B., Jünger, S., Schurr, U., 2007. Dynamics of seedling growth acclimation towards altered light conditions can be quantified via GROWSCREEN: a setup and procedure designed for rapid optical phenotyping of different plant species. *New Phytologist* 174, 447–455.
- Walter, A., Schurr, U., 1999. The modular character of growth in *Nicotiana tabacum* plants under steady-state nutrition. *J. Exp. Bot.* 50, 1169–1177.
- Wohlfahrt, P., Donoser, M., Roth, P.M., Bischof, H., 2013. Detecting partially occluded objects with an implicit shape model random field, in: Asian Conference on Computer Vision (ACCV), pp. 302–315.
- Wu, S., Bao, F., Xu, E., Wang, Y.X., Chang, Y.F., Xiang, Q.L., 2007. A leaf recognition algorithm for plant classification using probabilistic neural network, in: IEEE Int. Symp. Signal. Proc. Inf. Tech. (ISSPIT), pp. 11–16.
- Yin, X., Liu, X., Chen, J., Kramer, D.M., 2014. Multi-leaf tracking from fluorescence plant videos, in: Int. Conf. on Image Processing, pp. 408–412.