

ARTICLE

Open Access

Combining computer vision and deep learning to enable ultra-scale aerial phenotyping and precision agriculture: A case study of lettuce production

Alan Bauer^{1,2,3}, Aaron George Bostrom¹, Joshua Ball¹, Christopher Applegate¹, Tao Cheng⁴, Stephen Laycock³, Sergio Moreno Rojas⁵, Jacob Kirwan⁵ and Ji Zhou^{1,2,3} 

Abstract

Aerial imagery is regularly used by crop researchers, growers and farmers to monitor crops during the growing season. To extract meaningful information from large-scale aerial images collected from the field, high-throughput phenotypic analysis solutions are required, which not only produce high-quality measures of key crop traits, but also support professionals to make prompt and reliable crop management decisions. Here, we report AirSurf, an automated and open-source analytic platform that combines modern computer vision, up-to-date machine learning, and modular software engineering in order to measure yield-related phenotypes from ultra-large aerial imagery. To quantify millions of in-field lettuces acquired by fixed-wing light aircrafts equipped with normalised difference vegetation index (NDVI) sensors, we customised AirSurf by combining computer vision algorithms and a deep-learning classifier trained with over 100,000 labelled lettuce signals. The tailored platform, *AirSurf-Lettuce*, is capable of scoring and categorising iceberg lettuces with high accuracy (>98%). Furthermore, novel analysis functions have been developed to map lettuce size distribution across the field, based on which associated global positioning system (GPS) tagged harvest regions have been identified to enable growers and farmers to conduct precision agricultural practises in order to improve the actual yield as well as crop marketability before the harvest.

Introduction

As an important source of vitamins, minerals, and trace elements, leaf vegetables play crucial roles in human nutrition¹. Lettuce (*Lactuca sativa* L.), one of the most popular staple vegetable foods, has a wide range of tastes and textures cultivated for diverse customer needs². Recent research also indicates that lettuce consumption has positive effects on the reduction of cardiovascular disease and chronic conditions due to its rich nutrients

such as vitamin A, *Beta-carotene*, folate, and iron content³. While lettuce is an important and nutritional crop, fluctuating environments can increase the fragility of its production⁴. For example, the bad weather in Spain in early 2017 led to retail prices of lettuce products nearly tripled in UK supermarkets⁵. Severe weather not only causes supply shortage, but also affects crop quality. According to previous studies on lettuce growth and development^{6,7}, young plants at newly planted phase (i.e., from cotyledons unfolded to three true leaves stage) require cool and damp weather after the transplantation from the greenhouse to the field, whereas lettuce leaves can rapidly become bitter and inedible if the growth is accelerated by high ambient temperature at the head maturity phase (i.e., the growth stage before flowering).

Correspondence: Ji Zhou (Ji.Zhou@njau.edu.cn) (Ji.Zhou@earlham.ac.uk) or Jacob Kirwan (Jacob.Kirwan@GS-Growers.com)

¹Earlham Institute, Norwich Research Park, Norwich NR4 7UZ, UK

²Plant Phenomics Research Center, China-UK Plant Phenomics Research Centre, Nanjing Agricultural University, Nanjing 210095 Jiangsu, China

Full list of author information is available at the end of the article.

These authors contributed equally: Alan Bauer, Aaron George Bostrom

© The Author(s) 2019



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

Because of the dynamic nature of lettuce production, the actual yield of lettuces in commercial operations is only around 70–80% of the planted quantity⁸. Hence, to ensure the consistency of supply and quality, it is important for growers and farmers to closely monitor their crops during key growth stages, so that prompt and reliable crop management decisions can be made under changeable agricultural conditions⁹.

Aerial field phenotyping has become a popular approach for monitoring crops in recent years. Because it can acquire a large number of crop imagery in field experiments using visible, thermal, and multi-spectrum sensors, it has been widely applied to breeding, farming and crop research¹⁰. To ensure high-quality aerial image acquisition, the flight route and altitude need to be pre-determined together with the selection of appropriate imaging sensors¹¹. For example, for physiological traits such as vegetative greenness and canopy structure, a high-definition RGB camera is sufficient; however, many vegetation indices rely on multi- and hyper-spectral imaging sensors to assess important traits such as biomass, stress level, and yield potential¹². Recently, with the development of image stitching algorithms and orthomosaic generation methods, very detailed crop images can be collected by unmanned aerial vehicles (UAVs) and fixed-wing light aircrafts, which can enable high-quality field phenotyping and high-throughput phenotypic analysis¹³.

To extract meaningful phenotypic information from large-scale image datasets, a variety of computer vision¹⁴, machine learning (ML), and deep learning (DL) approaches¹⁵ have been utilised. In recent years, much attention has been paid to ML/DL techniques, based on which computational algorithms and learning models were built to accomplish tasks such as vision-based feature selection, image object classification, and pattern prediction^{16–18}. With adequate training data, suitable learning algorithms, and well-defined predictive outcomes, the integration of computer vision, ML/DL, and newly emerged analytic solutions (e.g., distributed computing) could lead to a step change for plant phenomics research in the near future¹⁹.

In this article, we present a new analytic platform called AirSurf developed for ultra-scale aerial phenotyping and yield-related phenotypic analysis. The software platform is open-source and combines tasks such as normalised difference vegetation index (NDVI) aerial imagery for data collection, computer vision for image processing, deep learning (i.e., convolutional neural networks, CNNs) for crop counting, and supervised machine learning for crop quality assessment. AirSurf was customised for commercial lettuce production so that it could be used to analyse millions of lettuces across the field. We call the tailored software platform “AirSurf-Lettuce” (AirSurf-L), which embeds a CNN model trained with over 100,000 labelled

lettuce signals to measure lettuce heads and their plantation layouts based on ultra-large NDVI images. After scoring lettuce, unsupervised ML algorithms were used to classify lettuce heads into three size categories (i.e., small, medium and large) for assessing lettuce quality. To connect phenotypic analysis with marketability and crop management decisions, a novel function has been developed in AirSurf-L to associate global positioning system (GPS) coordinates in a given field with the in-field lettuce size distribution, based on which efficient harvesting strategies could be formed to increase marketable yield.

Materials and methods

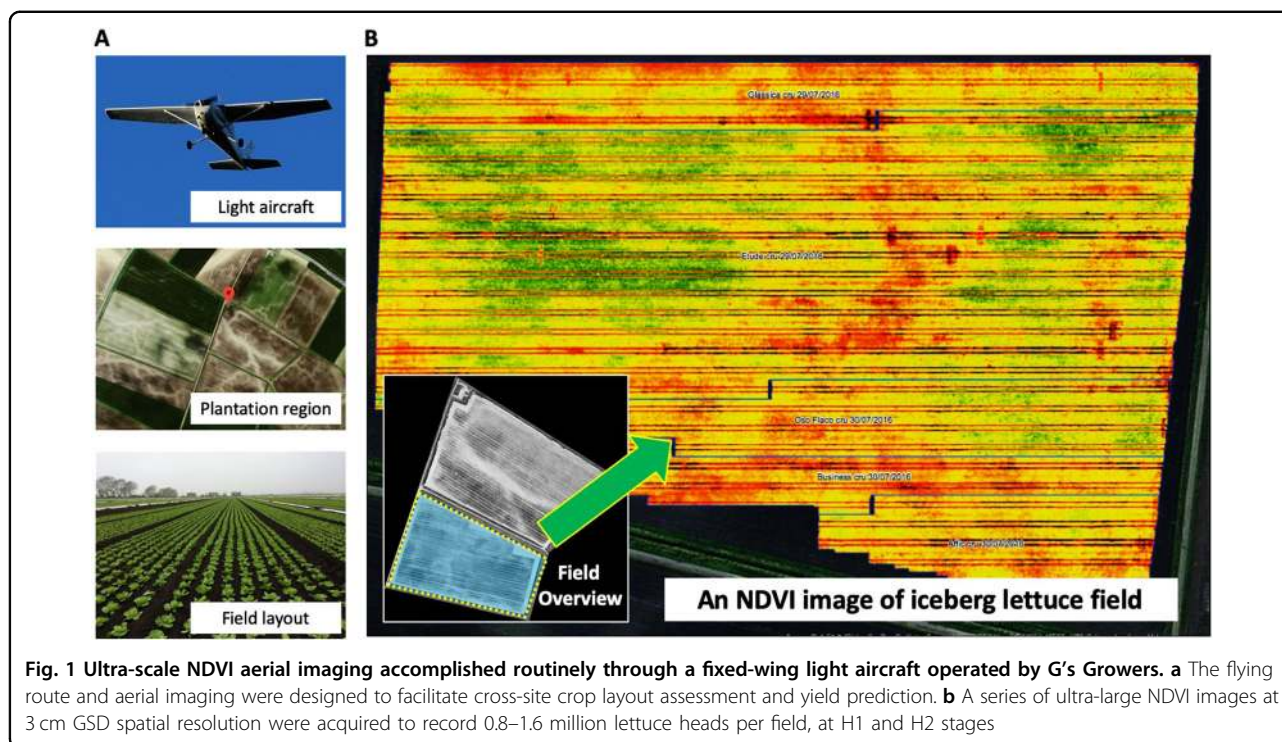
NDVI aerial imaging and experimental fields

NDVI correlates well with leaf area index and biomass²⁰ and hence was chosen for yield-related field phenotyping. The imaging sensor used is an industrial standard camera, as previously described²¹. The aerial imaging was carried out by a ‘Sky Arrow’ light aircraft, the lightest weight class (Very Light Aircraft, VLA) of any commercial aircraft, which allowed the pilot to fly with very little fuel, less than an average farm vehicle. Using VLA at 1000 feet (around 305m) in the sky, vast areas can be covered at a flight speed of 180–200 km/h, during which the NDVI sensor can gather ultra-scale crop imagery to cover four or five fields in a single flight.

The ultra-large aerial NDVI imagery was acquired routinely (i.e., four-five times per season) by G’s Growers, the second largest vegetable grower in the UK. The flying route and the imaging protocol were designed to facilitate cross-site crop assessment and yield prediction (Fig. 1a). In this study, we used a series of collected ultra-large NDVI images (1.5–2 GB per image) at 3 cm ground sample distance (GSD) spatial resolution, for iceberg lettuces at H1 and H2 stages (i.e., moderate compact and crushable head), before lettuce leaves were largely overlapped. Experimental fields in the study were all located near Ely, Cambridgeshire UK, ranging from 10 to 20 hectares, with between 800,000 and 1.6 million lettuce heads in a single field. One field (Field A, Fig. 1b) planted with around 1 million lettuce heads was used to explain the analysis workflow and associated algorithms of AirSurf-L in the following sections. A high-level manual yield counting was conducted by G’s growers’ field specialists during the harvest, which was used to verify and improve the platform. Lettuces in subsections randomly selected from Field A were scored manually by laboratory technicians at Norwich Research Park and then used as training datasets for establishing the deep learning model.

Data construction for training and testing

To generate sound datasets for ML-based image analysis, we randomly selected 60 patches (i.e., subsections) of the field of varying sizes, each containing between 300 and



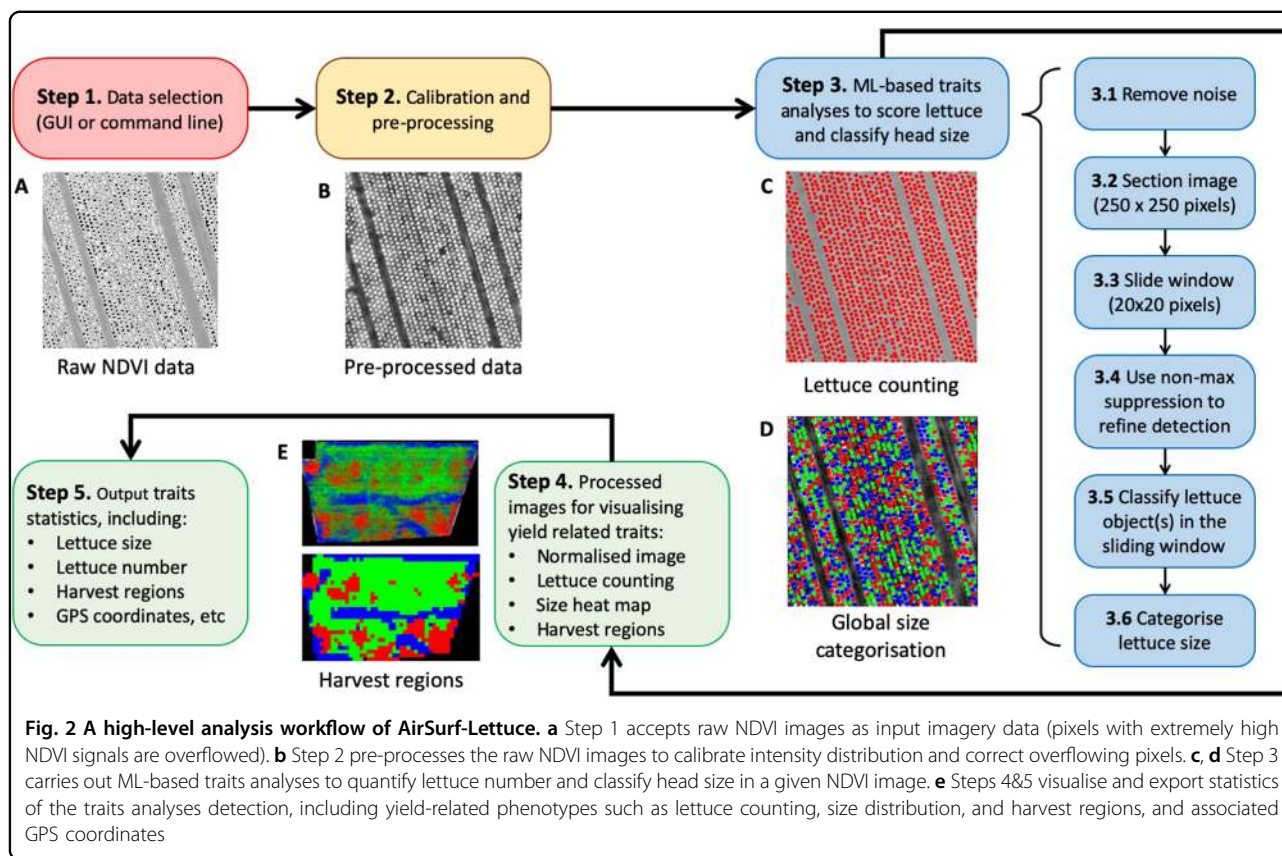
1000 lettuce heads. We then manually labelled each lettuce in the selected patches with a red dot (Supplementary Fig. 1). Each labelled lettuce, i.e., a red dot, is identified by a 20×20 pixel bounding box that can enclose a single lettuce head. We used these bounding boxes, as well as images that did not correspond to lettuce heads, to train a CNN classifier to recognise and separate millions of lettuces in the plantation region. The pixels contained within a bounding box were also used for defining lettuce size. A training dataset with over 100,000 20×20 pixel labelled bounding boxes has been created, amongst which 50% are lettuces and the remaining are background signals such as soil, edges of the field, and other non-lettuce objects. Following a standard CNN segmentation approach²², we designed a non-overlapping sliding window function to go through the whole field to separate foreground and background signals (i.e., splitting lettuce and non-lettuce objects). Training and testing datasets are equally balanced. Validation sets are used alongside training sets to verify the performance of the model, which can prevent overfitting in model training and allow us to fine-tune hyperparameters of different learning layers²³.

The analysis workflow of AirSurf-Lettuce

The analysis of yield-related phenotypes was based on NDVI signals of iceberg lettuces across the field. Figure 2 shows a high-level analysis workflow of AirSurf-L, which consists of five steps: data input, image

calibration and pre-processing, ML-based traits analyses, results visualisation, and quantifications of yield-related phenotypes. *Step 1* accepts raw NDVI images as grey-level imagery datasets. As pixels with extremely high NDVI signals usually have overflowed intensity values (i.e., black pixels in Fig. 2a), a pre-processing step (*Step 2*) is designed to calibrate raw NDVI images, so that intensity distribution can be normalised to correct overflowing pixels. At this step, an algorithm called contrast limited adaptive histogram equalisation (CLAHE)²⁴ is applied to increase the contrast between the foreground (i.e., lettuces) and background (e.g., soils) in a given NDVI image (Fig. 2b). Supplementary File S1 provides pseudo code and explanations of the image calibration and pre-processing step to ensure high-quality inputs of the learning model.

Step 3 carries out ML-based traits analyses that quantify lettuce number (Fig. 2c) and classify head size (Fig. 2d). It includes six steps: removing noise signals, partitioning a given image into sections (250×250 pixels) for local analysis, producing a sliding window (20×20 pixels) to traverse within a sectioned image, using non-max suppression to detect lettuces, and classifying recognised lettuces into three size categories (i.e., small, medium and large). The analysis result is visualised in *Step 4*, where lettuce counting, size distribution map, and GPS-tagged harvest regions are saved as a series of processed images (Fig. 2e). At the final step (*Step 5*), statistics of yield-related traits are exported to a comma-separated



values (CSV) file, including lettuce counts per field, lettuce size distribution, lettuce number and size measures within GPS grids, harvest regions, and associated GPS coordinates (Supplementary File S2). To enable users to carry out the above analysis workflow easily, a graphical user interface (GUI) software application has been developed.

AirSurf-Lettuce GUI

The GUI of AirSurf-L (Fig. 3) was developed using the native python GUI package, Tkinter²⁵, which allows the software application to be executed on different operating systems such as Windows and Mac OS (note: we only provided a packaged.exe executable file on the GitHub, see availability and requirements). Following the systems design described previously²⁶, the GUI uses an easy-to-follow approach to implement the phenotypic analysis workflow. The GUI window is divided into two parts: input section and display section. In the input section (dash rectangle coloured red in Fig. 3), a user needs to firstly load an NDVI image, which will be displayed instantaneously in the display section (dash rectangle coloured green), in the *original* tab. Secondly, the user needs to enter GPS coordinates of the field (i.e., the top left corner of the input image, which can be retrieved

from the metadata or Google Maps). Thirdly, the user is required to define the rotation value of the input image (in degrees) in comparison with the north geographical direction, so that GPS calculation can be standardised. Then, the user can tell the software whether the input image contains overflowed NDVI signals; if so, an extra calibration process will be triggered (Fig. 2, Step 2). Finally, after entering a small number of input parameters, the user can click the Start button to initiate the analysis workflow.

The GUI software follows each step described in the analysis workflow to accomplish automated phenotypic analysis. When a step is finished, an associated processed image will be displayed in the display section, showing the intermediate result of the analysis. Four processed images will be presented in the GUI window: a pre-processed and calibrated image (Fig. 3a, in the *normalised* tab), an image after lettuce scoring (Fig. 3b, in the *counts* tab), an image after size categorisation (Fig. 3c, in the *size distribution* tab), and a processed image after identifying harvest regions and their associated GPS coordinates (Fig. 3d, in the *harvest regions* tab). All processed images are saved in a result folder, along with a CSV file that contains analysis results (Supplementary File S2).

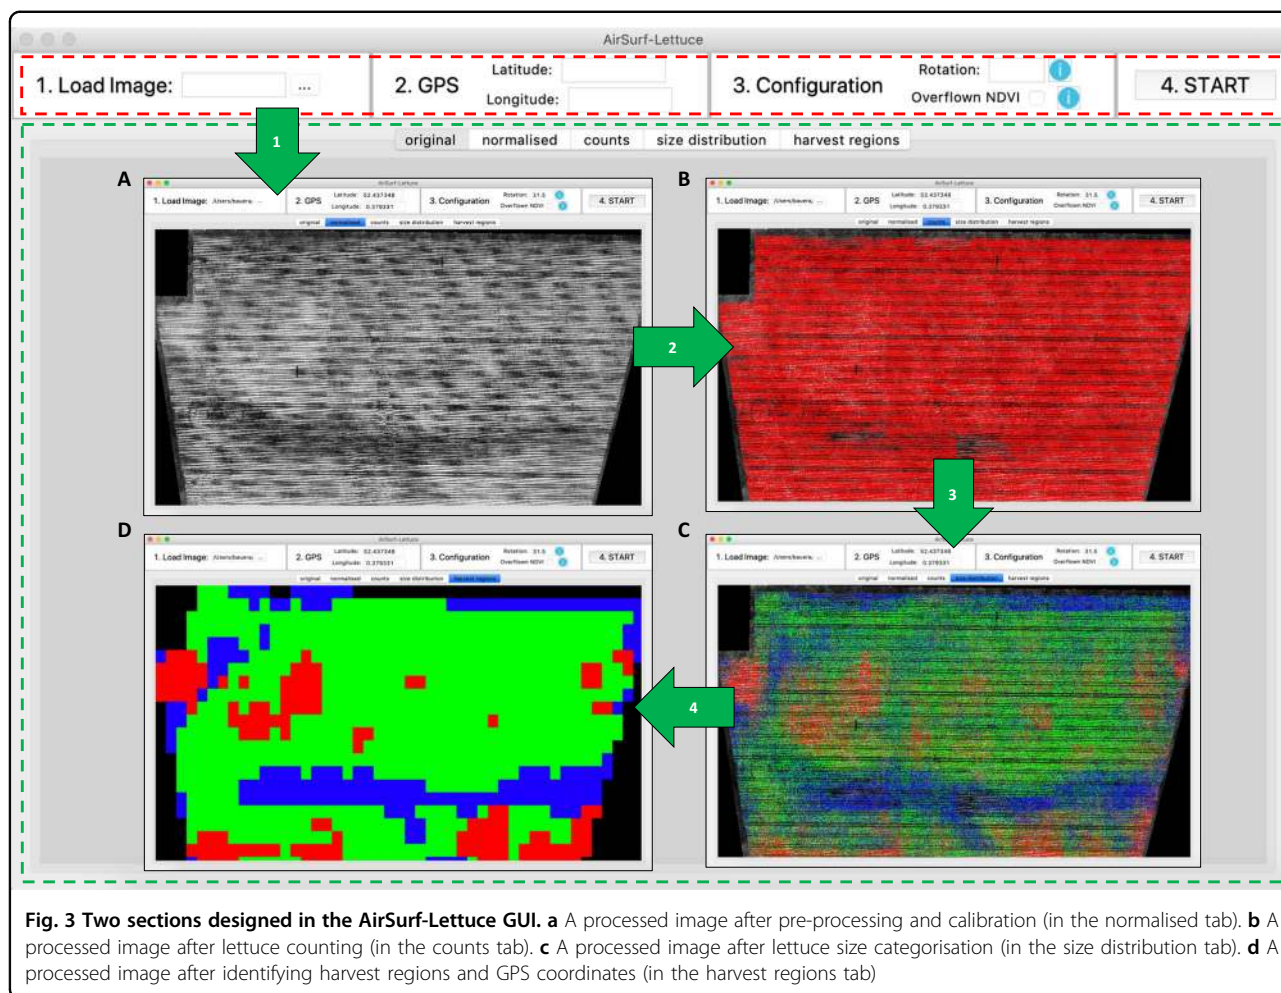


Fig. 3 Two sections designed in the AirSurf-Lettuce GUI. **a** A processed image after pre-processing and calibration (in the normalised tab). **b** A processed image after lettuce counting (in the counts tab). **c** A processed image after lettuce size categorisation (in the size distribution tab). **d** A processed image after identifying harvest regions and GPS coordinates (in the harvest regions tab)

Neural network architecture

Similar to AlexNet²⁷, a CNN-based learning architecture was established using the labelled training datasets. Figure 4a demonstrates the architecture of the CNN model, including (1) a convolutional (Conv2D) layer with 32 filters and a 3×3 kernel, with a rectified linear unit (ReLU) as the activation function, and batch normalisation to accelerate the learning process to enable higher learning rates²⁸; (2) the same block is then repeated together with a max pooling layer to down-sample input using a 2×2 kernel based on the assumption that useful input features could be contained in sub-regions; (3) after that, a second convolutional block is constructed, consisting of a Conv2D layer with 64 filters, a 3×3 kernel, a ReLU activation, and batch normalisation; (4) finally, this block is repeated, followed by another max pooling layer (with a 2×2 kernel) to complete the learning procedure. After the convolutional layers, learning layers are connected to a fully connected layer of size 512, which is followed by a dropout layer with a 50% chance. To complete the learning architecture, a binary output generates the probability of whether a given

bounding box (20×20 pixels) contains a lettuce signal. If the probability equals or is close to 100%, it indicates that it is highly likely that the bounding box contains a complete lettuce head (Fig. 4b). The above architecture is commonly applied to vision-based object detection problems²⁹. The training and validation accuracy and loss curves are reported in Fig. 4c, showing that the model converges in only 10 epochs. More importantly, to avoid overfitting, the stopping criterion was designed to guarantee the validation accuracy is higher than the training accuracy, ensuring the generalisation of the learning model. To avoid the overfitting issue of our model, the labelled data was also divided equally into train and validation sets when training the model.

The architecture is shallower than AlexNet and other modern deep learning architectures for a number of reasons: (1) the size of our dataset is relatively small for establishing very deep learning networks; (2) our target is a binary classification problem (i.e., whether or not a given bounding box contains a whole lettuce head), different from ImageNet classification tasks; (3) larger and deeper

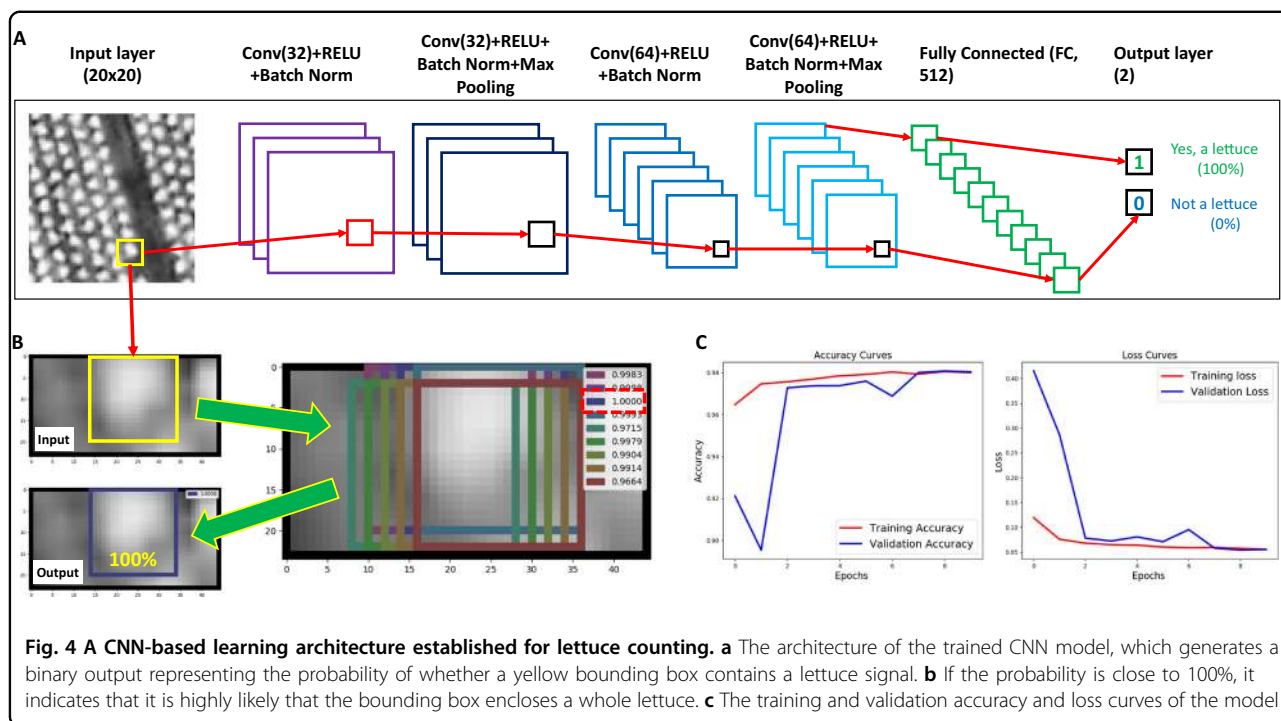


Fig. 4 A CNN-based learning architecture established for lettuce counting. **a** The architecture of the trained CNN model, which generates a binary output representing the probability of whether a yellow bounding box contains a lettuce signal. **b** If the probability is close to 100%, it indicates that it is highly likely that the bounding box encloses a whole lettuce. **c** The training and validation accuracy and loss curves of the model

neural networks require more time to train, which can be slower to execute and not feasible for prompt crop management decision required in precision agriculture.

Size categorisation algorithm

After AirSurf-L identifies bounding boxes containing lettuce heads, we employed an unsupervised ML approach to categorise lettuce into three sizes: small, medium and large. The algorithm can be easily changed to classify more size categories, if required. Pixels in the bounding box region are extracted and then NDVI values of all the pixels are put into bins. The histogram included 10 bins that spread across the NDVI value range (0–255). We included two important aspects when categorising lettuce sizes: (1) lower NDVI surrounding values do not determine the lettuce size; (2) higher NDVI values are more important for the size categorisation. As such, a geometric pattern of NDVI values for each bin was created, i.e., 64, 128, 160, 192, 208, 224, 232, 240, 244, 248, 250, 252, 253, and 254. With these cut-off values, most of the background pixels were captured in the first two bins, along with the increasing weight when values approach 255.

Having transformed the pixel regions into a series of bins, we were able to compare different regions and cluster them into three size groups using k-means clustering with the *k* value set to three. Then, clustering results are sorted through calculating the dot product between the weight vector and the cluster count vector (based on the number of bins). These sorted values determine which clustering result corresponds to which

size, which are then applied to each lettuce detected in the field. Three colours are used to indicate size categories: blue for small, green for medium, and red for large (Fig. 3c).

Results

Counting lettuces with a CNN classifier

After a CNN classifier was trained and the phenotypic analysis algorithms were implemented in the AirSurf-L, we used the software to recognise and classify lettuces in a series of ultra-large NDVI images. Initially, a broad range of sizes and orientations of lettuces with varying intensities were captured; however, the software failed to recognise lettuces in very bright regions and overly count lettuces in very dark regions (Fig. 5a), e.g., around 50,000 lettuces were wrongly detected in a one-million-head field (5% counting error). We found that this problem was caused by the trained CNN classifier, because a lettuce head is extremely tiny in an orthomosaic image (e.g., 11330 × 6600 pixels for a 7-hectare field when GSD is 3 cm, which contains over half million lettuces) under varied lighting conditions. To resolve this issue, we have designed a two-step solution: (1) sectioning the whole image into many 250 × 250 pixels sub-images, and (2) using a fix-sized bounding box (20 × 20 pixels) as a sliding window (with a stepping parameter of 5 pixels to reduce the computational complexity) to prune the detected lettuce objects in each 250 × 250 sub-image.

Another reason that caused the misdetection is due to overlapped lettuces as they could be detected repeatedly

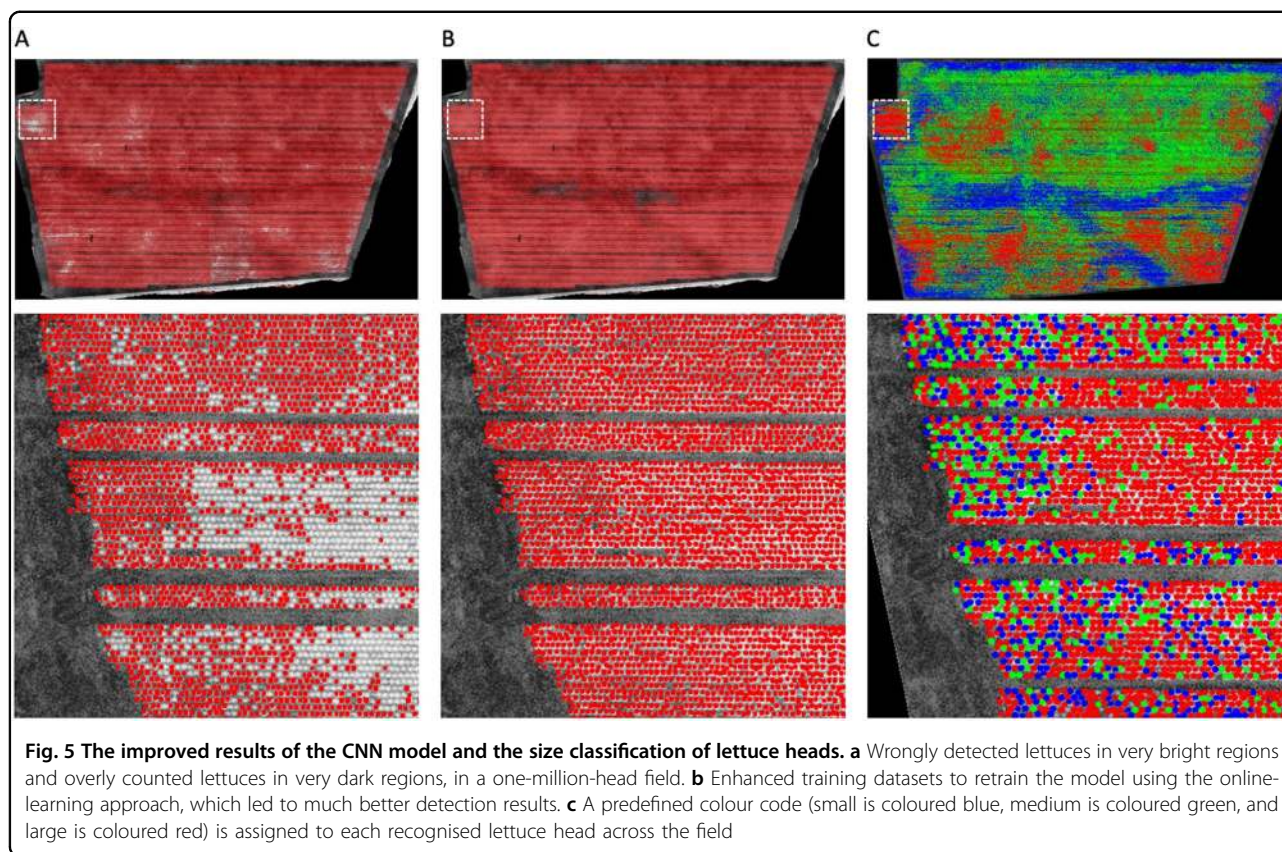


Fig. 5 The improved results of the CNN model and the size classification of lettuce heads. **a** Wrongly detected lettuces in very bright regions and overly counted lettuces in very dark regions, in a one-million-head field. **b** Enhanced training datasets to retrain the model using the online-learning approach, which led to much better detection results. **c** A predefined colour code (small is coloured blue, medium is coloured green, and large is coloured red) is assigned to each recognised lettuce head across the field

by the CNN classifier in a sub-image. Hence, we employed a non-maximum suppression (NMS) algorithm³⁰ to rectify the detection. NMS uses probabilities to order the detected lettuce objects. After the 20×20 sliding window is performed and many small patches have been identified, the NMS algorithm computes an overlap coefficient to determine how to retain these patches. As lettuces are relatively well-spaced in the field, patches (i.e., bounding boxes) enclosing a complete lettuce signal are retained, whereas partially covered signals will be removed. To select the best overlap parameter computed by the NMS, a gradient descent method is formulated and explained in Supplementary File S3.

Improved CNN classifier and the size categorisation

Besides the improved vision-based object detection, we also enhanced the training datasets by manually labelling an additional 500 lettuce signals within very bright or very dark regions. Then, newly labelled data was inserted into the training datasets to retrain the model through the online-learning approach³¹. The improved CNN model (see GitHub repository and Supplementary File S4) was tested on different experimental fields again and has dramatically enhanced the accuracy of lettuce detection (Fig. 5b).

Identified lettuces are individually analysed to determine their associated size category. The size classification is based on intensity and contrast values enclosed by the 20×20 bounding boxes, which is computed using the dot product of the histogram of pixel intensities and a weighted vector towards more pixel-based contrast values. The assumption of this design is that higher NDVI signals likely correlate with higher vegetation indices and hence bigger lettuce heads. The categorisation result of all lettuce heads is clustered into three size groups. Each lettuce is then coloured with a predefined colour code (Fig. 5c).

A GPS-tagged harvest map

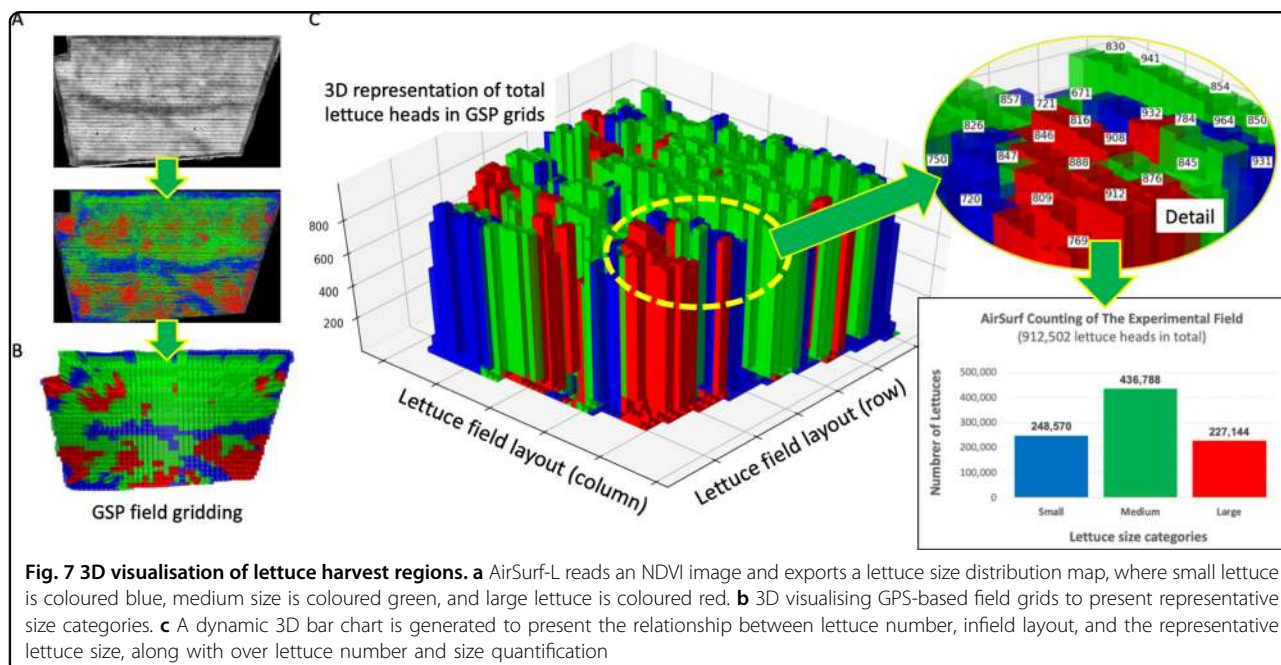
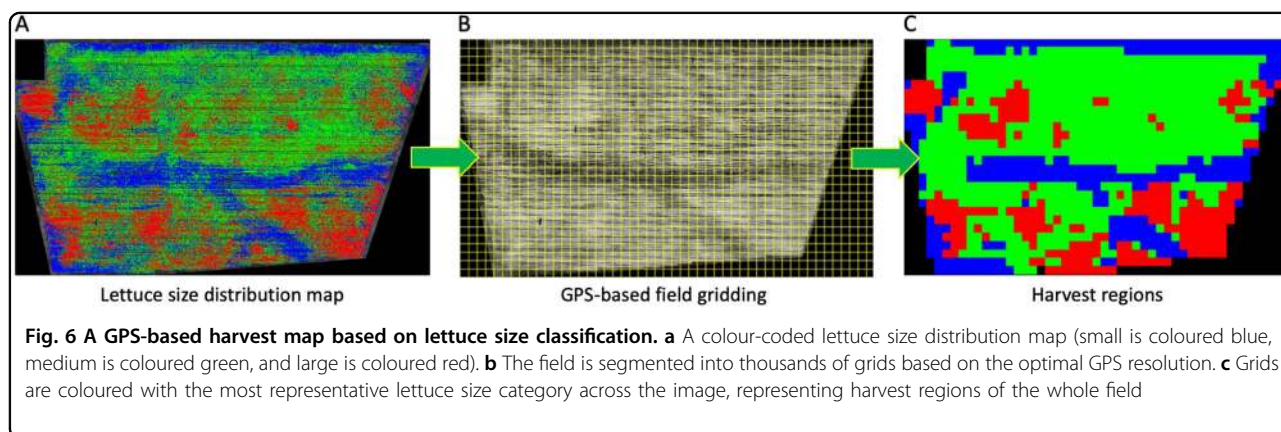
The final phase of the phenotypic analysis is to define harvest regions based on different sizes of lettuces. Using the size distribution map (Fig. 6a), the field is firstly segmented into many small grids based on the optimal GPS resolution determined by the altitude of the aerial imagery (3 cm GSD, in our case), as well as the size of the harvester machinery used by the grower. After dividing the field into thousands of grids (Fig. 6b), GPS coordinates of each grid are computed and each grid is then coloured with the most representative size category. By combining all coloured grids, a GPS-tagged harvest map is produced, representing harvest regions of the

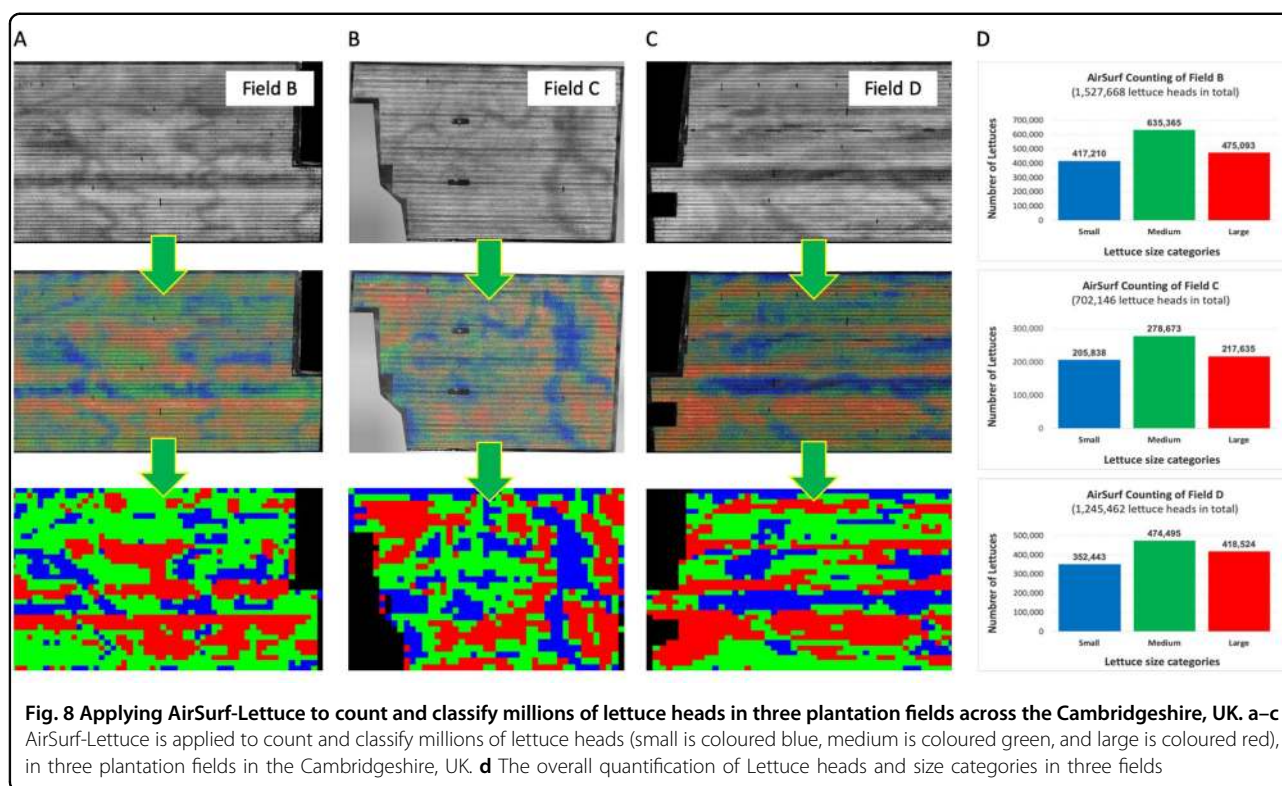
whole field (Fig. 6c). The harvest map can be used for designing harvesting strategies such as guiding a harvester to collect desired sized lettuces or arranging logistics based on the lettuce number and associated size counting. To facilitate precision agricultural practices, a result file (Supplementary File S2) is also generated by AirSurf-L at the end of the analysis workflow, containing information of each harvest grid, the associated GPS coordinates, lettuce size and number counting in each grid. To satisfy different needs for dissimilar requirements, the size of GPS-based harvest grids can be modified manually in the software.

3D visualisation for the harvesting strategy

Figure 7 uses Python-based 3D Matplotlib library³² to show the GPS-tagged harvest map. When AirSurf-L

reads an NDVI image, it computes the number of lettuce heads and associated size categories on the image (Fig. 7a). Then, by combining GPS-based field grids with the representative lettuce size in these grids (Fig. 7b), we produced a dynamic 3D bar chart script (Supplementary File S5) to present the lettuce number using the z axis, infield harvest regions (both columns and rows) using both x and y axes, and the representative lettuce size using the predefined colours (Fig. 7c). Through the 3D plot, users can zoom into any sub-region of the field to check lettuce number and representative size so that a precise harvesting strategy can be planned accordingly. The overall lettuce number and size counting of the experiment field can also be calculated.





Validation of AirSurf-Lettuce

To verify AirSurf-L and the soundness the algorithm, we have applied the platform to count and classify lettuce heads in three unseen experimental fields in Cambridgeshire, UK (Fig. 8a–c). These fields contain between 700,000 and 1,500,000 lettuces and are located in different sites around the county. Traits such as the number of lettuces per field quantified by the platform (Fig. 8d) were compared with industrial estimates, showing a low error in lettuce counting (<5% difference). Besides the field-level comparison, we also randomly selected different sizes of subsections in an experiment field to evaluate AirSurf-L. We split these subsections into three sets (i.e., 36 small regions, 21 large regions, and 57 mixed regions), where the small regions have less than 400 lettuces, the large ones contain greater than 900 lettuces heads, and mixed regions contain a variety of lettuce heads. After that, laboratory technicians manually counted lettuce heads within these regions. The correlation between the manual and automated lettuce counting shows that, for the small regions, the correlation between the human and automatic counting is approximately 2% ($R^2 = 0.978$); for the large regions, the value is around 0.8% ($R^2 = 0.988$); and for mixed regions, the R^2 correlation is over 0.9997. Supplementary Fig. 2 and Supplementary File S6 show the correlations between human and automatic counting for all three region groups.

Discussion

Traditionally, measuring in-field crops on a large scale is very time-consuming and labour-intensive. It often requires destructive techniques, potentially error-prone manual counting, or estimates of traits that are key to yield production or crop quality³³. Recent advances in ML/DL and CV techniques have led to an explosion of plant phenomics, which has rapidly improved our abilities in mining phenotypic information from large and complicated phenotyping datasets³⁴. New data-driven analytic approaches are also changing plant phenomics – collecting big data (i.e., phenotyping) is no longer the bottleneck, instead how to extract biologically relevant information (i.e., phenotypic analysis) from big data has become the current challenge³⁵. Hence, along with the development of aerial imaging and remote sensing technologies, it has become increasingly noticeable that the integration of scalable data collection, high-throughput phenotypic analysis, and yield predictive modelling is key to future crop research and precision agriculture³⁶.

AirSurf-L introduced here has addressed a specific challenge in ultra-scale aerial phenotyping and precision agricultural management through combining aerial NDVI imagery, CV, ML/DL, and software engineering, with commercial lettuce production. The platform automates the quantification of millions of lettuces across the field, which allows us to connect research-based phenotypic

analysis with real-world agriculture problems. As a cross-disciplinary project, we have chosen an agile R&D method, because technologies and requirements were constantly changing during the project. The results generated by AirSurf-L show a strong correlation between automatic counting and specialist scoring ($R^2 = 0.98$). Hence, we are confident that AirSurf-L is capable of assisting fresh vegetable growers and farmers with their large-scale field phenotyping needs as well as yield-related trait analysis.

Commercial impacts

Commercially, lettuce production offers an attractive economic profitability in comparison to many other Agri-Food businesses³⁷. To date, lettuce businesses are worth billions of dollars and employ hundreds of thousands of permanent and seasonal workers globally. European vegetable growers alone produced 2.95 million tonnes of lettuce (and chicory) in 2016, a total annual value of €2.5 billion³⁸. Further down the fresh produce supply chain, the planning and efficiency of many essential crop production activities are largely dependent on crop maturity date and the marketability of crops (i.e., the crop quality)³⁹. Marketing activities such as logistics, trading, and product marketing need to be organised several weeks before the harvest; moreover, the booking and reservation of crop distribution, agricultural equipment, and associated commercial plans with retailers also need to be determined beforehand⁴⁰. By doing so, crop can be harvested at the right time, with maximised yield⁴¹. Our work contributes directly to lettuce production through improving the actual yield of lettuces and providing reliable quantification of crop quality (e.g., lettuce size), both of which are key factors for crop production, marketing, and supply chain management.

Machine learning and computer vision in plant phenomics

Another aim of this work is to further ML- and CV-based software solutions in plant phenomics research. High-throughput plant phenotyping is a fast-growing research domain, covering many disciplines, from plant breeding, cultivation, remote sensing, to computing sciences⁴². The modular software development allows us to test and embed different open-source learning architectures⁴³ (e.g., through the TensorFlow frameworks) and CV algorithms⁴⁴ (e.g., Scikit-Image libraries) in AirSurf-L. Notably, it is worth pointing out that we have learned a number of lessons when applying ML/DL and CV in phenotypic analysis: (1) learning algorithms could perform badly if training datasets are not well-labelled and insufficient; (2) although ML/DL algorithms specialised in segmentation and classification if target objects are well-defined, there is still a big gap between object recognition and traits measurement; (3) meaningful phenotypic

analysis not only requires sufficient biological understanding to define target traits in a logical manner, but also needs bespoke algorithms to engineer features so that traits can be quantitatively described. Hence, in plant phenomics research, biological questions, analytical solutions, and software implementation need to be considered collectively in order to address approaching challenges.

Limitations and further development of the platform

Besides the promising phenotypic analysis results presented in this article, there are still limitations of the platform need to be considered: (1) AirSurf-L has been tested with top-view iceberg lettuces mainly at H1 and H2 stages, which means that analysis error could increase if there are too many overlaps between lettuce heads, e.g., from H3 stage onwards. (2) As AirSurf-L has only been tested with NDVI imagery, it is important to add new functions to the platform to incorporate other vegetation indices acquired by multi- and hyper-spectrum imaging sensors. (3) As precision agriculture management decisions are normally based on imagery, soil and climate conditions, AirSurf's results will be more reliable, if we could include soil information for each harvest region and field-level climate conditions. So, results can be compared between sites in multiple years. A potential approach is to incorporate ground-based phenotyping systems such as CropSight⁴⁵ to feed environment data to the analysis. (4) The method was tested and validated in lettuce fields in a number of geographic locations following a standard aerial imaging procedure, data collected from different sites via varied aerial imaging strategies (e.g., different angles, altitudes and GSD) could improve the soundness and compatibility of the platform. (5) Key features were constructed by learning algorithms instead of engineered, which make learning models vulnerable when facing up to totally undefined datasets. Hence, ML/DL based phenotypic analysis algorithms need to update with new labelled training data for new crop species. (6) For a field of approximately one million lettuces, it takes about 4 h to analyse a field on a decent computer (2.5 GHz Intel Core i7, 8GB memory). Most of the computational time is occupied by the learning model to identify individual lettuces, which can be improved by parallel computing or GPU (graphic processing unit) processing to speed up the analysis.

Prospects for crop research and precision agriculture

Together with recent advances in multi-scale remote sensing and phenotyping data management^{45–47}, the platform could be relatively easily expanded to incorporate other crop species such as wheat and rice by retraining the learning model with additional datasets. By doing so, AirSurf could be developed into a more

comprehensive analytic platform that will bring great significance to crop production and marketable yield for the Agri-Food sector. For example, the plant density of wheat and rice is closely related to the yield due to its influences on the allocation of water, light and fertilisers, which cannot be quantified using ground-based RGB imagery⁴⁸. Hence, utilising the ultra-scale NDVI aerial imagery and related object recognition methods embedded in AirSurf-L, the platform is likely to benefit the assessment of sowing performance, emergence rate, and plant distribution. Then, through a multi-scale phenotyping approach (i.e. integrating ground-based workstations), breeders and crop researchers could make early predictions of the grain yield of crop genotypes in field experiments.

From a precision agriculture perspective, monitoring individual plant such as a lettuce head can enable accurate monitoring of crops during key growth stages across a plantation site. It can provide growers with the real number of crops in the field, based on which yield for harvest availability can be quantified instead of estimated. The calculation of crops can also lead to accurate agricultural inputs, facilitating automated variable-rate application of fertiliser, weed control, and pesticides through tractor software system with a more precise crop distribution map⁴⁹. Furthermore, the close monitoring of key yield-related traits can be used to guide farmers and growers to reduce variability of agrichemical applications and irrigation in different fields, increasing harvest yield and better operating profit margin⁵⁰. Finally, new analytic platforms such as the AirSurf-L platform shall largely embed in routine agricultural activities, so that no major extra costs are required, making new Agri-Tech solutions more adoptive by the Agri-Food sector.

Conclusions

AirSurf-Lettuce automatically measures in-field iceberg lettuces using ultra-scale NDVI aerial images, with a focus on yield-related traits such as lettuce number, size categories, field size distribution, and GPS-tagged harvest regions. The analysis results are close to the manual counting and can be used to improve the actual yield. By monitoring millions of lettuces in the field, we demonstrate the significant value of AirSurf-L in ultra-scale field phenotyping, precise harvest strategies, and crop marketability before the harvest. We believe that our algorithm design, software implementation, the application of ML/DL and CV algorithms, and cross-disciplinary R&D activities will be highly valuable for future plant phenomics research that are destined to be more challenging. With continuous R&D work, we are confident that the platform has great potential to support the Agri-Food sector with a smart and precise crop surveillance

approach of vegetable crops and therefore lead to better crop management decisions.

Availability and requirements

Project name: AirSurf-Lettuce with G's Growers

Project home page: <https://github.com/Crop-Phenomics-Group/Airsurf-Lettuce>

Source code: <https://github.com/Crop-Phenomics-Group/AirSurf-Lettuce/>

GUI software: <https://github.com/Crop-Phenomics-Group/AirSurf-Lettuce/releases>

Operating system(s): platform independent

Programming language: Python 3.6

Requirements: Keras, TensorFlow, Skimage, and Numpy.

License: BSD-3-Clause available at <https://opensource.org/licenses/BSD-3-Clause>

Availability of supporting data

The datasets supporting the results presented here is available at <https://github.com/Crop-Phenomics-Group/Airsurf-Lettuce/releases>. Source code and other supporting data are also openly available in the GitHub repository.

Acknowledgements

We thank all members of the Zhou laboratory at EI and Nanjing Agricultural University for fruitful discussions and cross-country collaborations. We thank researchers at John Innes Centre and UEA for constructive suggestions. We gratefully acknowledge the support of NVIDIA Corporation with the award of the Quadro GPU used for this research. J.Z. was partially funded by UKRI Biotechnology and Biological Sciences Research Council's (BBSRC) Designing Future Wheat Cross-institute Strategic Programme (BB/P016855/1) to Graham Moore, BBS/E/T/000PR9785 to J.Z.; J.B. were partially supported by the Core Strategic Programme Grant (BB/CSP17270/1) at the Earlham Institute; A.G.B. and C.A. were also partially supported by G's Growers's industrial fund awarded to J.Z.; A.B. was partially supported by the Newton UK-China Agri-Tech Network+Grant (GP131JZ1G) awarded to J.Z.

Author details

¹Earlham Institute, Norwich Research Park, Norwich NR4 7UZ, UK. ²Plant Phenomics Research Center, China-UK Plant Phenomics Research Centre, Nanjing Agricultural University, Nanjing 210095 Jiangsu, China. ³School of Computing Sciences, University of East Anglia, Norwich Research Park, Norwich NR4 7TJ, UK. ⁴National Engineering and Technology Center for Information Agriculture, MARA Key Laboratory for Crop System Analysis and Decision Making, Jiangsu Key Laboratory for Information Agriculture, Nanjing Agricultural University, Nanjing 210095 Jiangsu, China. ⁵G's Growers Limited, Ely, Cambridgeshire CB7 5TZ, UK

Author contributions

J.Z., A.G.B., A.B. and J.K. wrote the manuscript, S.M.R. performed the NDVI imaging. J.K. and S.M.R. provided harvest information and biological expertise. J.Z., A.G.B. and C.A. designed the analysis algorithms. A.B., A.G.B., C.A. and J.Z. implemented the software solution. A.G.B. and A.B. built the deep learning models with J.Z.'s supervision. J.B. packaged the GUI executables. J.Z., A.B., A.G.B., C.A., S.L. and J.B. tested the software and performed the data analysis. All authors read and approved the final manuscript.

Conflict of interest

The authors declare that they have no conflict of interest.

Supplementary Information accompanies this paper at (<https://doi.org/10.1038/s41438-019-0151-5>).

Received: 8 January 2019 Revised: 1 April 2019 Accepted: 8 April 2019
Published online: 01 June 2019

References

- Gubbi, J., Buyya, R., Marusic, S. & Palaniswami, M. Internet of Things (IoT): a vision, architectural elements, and future directions. *Futur Gener. Comput. Syst.* **29**, 1645–1660 (2013).
- Mou, B. Mutations in lettuce improvement. *Int. J. Plant Genom.* **2011**, 1–7 (2011).
- Koudela, M. & Petříková, K. Nutrients content and yield in selected cultivars of leaf lettuce (*Lactuca sativa* L. var *crispa*). *Hortic. Sci.* **35**, 99–106 (2008).
- Tei, F., Scaife, A. & Aikman, D. Growth of lettuce, onion, and red beet. 1. *Growth Anal. light Intercept., Radiat. use Effic. Ann. Bot.* **78**, 633–643 (1996).
- BBC. Iceberg lettuces and broccoli rationed as vegetable crisis hits supermarkets. BBC News. 2017;1–3. Published online 3 February 2017. <https://www.bbc.com/news/uk-38851097>.
- Jie, H. Lee Sing Kong. Growth and photosynthetic characteristics of lettuce (*Lactuca sativa* L.) under fluctuating hot ambient temperatures with the manipulation of cool root-zone temperature. *J. Plant Physiol.* **152**, 387–391 (1998).
- Johkan, M., Shoji, K. & Goto, F. Hashida S nosuke, Yoshihara T. Blue light-emitting diode irradiation of seedlings improves seedling quality and growth after transplanting in red leaf lettuce. *HortScience* **45**, 1809–1814 (2010).
- Ogbodo, E. N., Okorie, P. O. & Utobo, E. B. Growth and Yield of Lettuce (*Lactuca sativa* L.) At Abakaliki Agro-Ecological Zone of Southeastern Nigeria. *World J. Agric. Sci.* **6**, 141–148 (2010).
- Tudela, J. A., Hernández, N., Pérez-Vicente, A. & Gil, M. I. Growing season climates affect quality of fresh-cut lettuce. *Postharvest Biol. Technol.* **123**, 60–68 (2017).
- Sankaran, S. et al. Low-altitude, high-resolution aerial imaging systems for row and field crop phenotyping: a review. *Eur. J. Agron.* **70**, 112–123 (2015).
- Duan T., et al. Comparison of ground cover estimates from experiment plots in cotton, sorghum and sugarcane based on images and orthomosaics captured by UAV. *Funct. Plant Biol.* <https://doi.org/10.1071/FP16123> (2016).
- Meyer, G. E. & Neto, J. C. Verification of color vegetation indices for automated crop imaging applications. *Comput. Electron. Agric.* **63**, 282–293 (2008).
- Yang, G. et al. Unmanned aerial vehicle remote sensing for field-based crop phenotyping: current status and perspectives. *Front Plant Sci.* **8**, 1–27 (2017).
- Rousseau D., Dee H., Pridmore T. Imaging Methods for Phenotyping of Plant Traits. In *Phenomics in Crop Plants: Trends, Options and Limitations*. (ed. Kumar S.) 61–74 (Springer India, 2015).
- Tsaftaris, S. A., Minervini, M. & Schar, H. Machine learning for plant phenotyping needs image processing. *Trends Plant Sci.* **21**, 989–991 (2016).
- Burridge, J., Jochua, C. N., Bucksch, A. & Lynch, J. P. Legume shovelomics: High-Throughput phenotyping of common bean (*Phaseolus vulgaris* L.) and cowpea (*Vigna unguiculata* subsp. *unguiculata*) root architecture in the field. *F. Crop. Res.* **192**, 21–32 (2016).
- Bradley, D. et al. Evolution of flower color pattern through selection on regulatory small RNAs. *Science* **358**, 925–928 (2017).
- Angermueller, C., Pärnamaa, T., Parts, L. & Stegle, O. Deep learning for computational biology. *Mol. Syst. Biol.* **12**, 878 (2016).
- Großkinsky, D. K. et al. Phenotyping in the fields: Dissecting the genetics of quantitative traits and digital farming. *New Phytol.* **207**, 950–952 (2015).
- Zaman-Allah, M. et al. Unmanned aerial platform-based multi-spectral imaging for field phenotyping of maize. *Plant Methods* **11**, 35 (2015).
- Simms, D. M., Waine, T. W., Taylor, J. C. & Juniper, G. R. The application of time-series MODIS NDVI profiles for the acquisition of crop information across Afghanistan. *Int. J. Remote Sens.* **35**, 6234–6254 (2014).
- Pound, M. P. et al. Deep machine learning provides state-of-the-art performance in image-based plant phenotyping. *Gigascience* **6**, 1–10 (2017).
- Rosebrock A. *Deep Learning for Computer Vision with Python*. <https://doi.org/10.1017/CBO9781107415324.004> (2017).
- Reza, A. M. Realization of the contrast limited adaptive histogram equalization (CLAHE) for real-time image enhancement. *J. VLSI Signal Process. Syst. Signal Image Video Technol.* **38**, 35–44 (2004).
- Shipman J. W. Tkinter 8.5 reference: a GUI for Python. New Mexico. <http://www.nmt.edu/tcc/help/pubs/tkinter/> (2013).
- Zhou, J. et al. Leaf-GP: An open and automated software application for measuring growth phenotypes for arabidopsis and wheat. *Plant Methods* **13**, 1–31 (2017).
- Alex K, Ilya S, Geoffrey E. H. ImageNet Classification with Deep Convolutional Neural Networks. In: *Neural information processing systems*. Neural Information Processing Systems (NIPS): Lake Tahoe, Nevada, 2012, pp 1097–1105.
- Choromanska A., Henaff M., Mathieu M., Arous G Ben & LeCun Y. The Loss Surface of Multilayer Networks. *arXiv14120233* [cs] (2014).
- Matsugu, M., Mori, K., Mitari, Y. & Kaneda, Y. Subject independent facial expression recognition with robust face detection using a convolutional neural network. *Neural Netw.* **16**, 555–559 (2003).
- Neubeck, A. & Van Gool, L. Efficient non-maximum suppression. *Proc. Int. Conf. Pattern Recognit.* **3**, 850–855 (2006).
- Hazan, E. Introduction to online convex optimization. *Found. Trends Optim.* **2**, 151–166 (2015).
- Hunter, J. D. Matplotlib: a 2D graphics environment. *Comput. Sci. Eng.* **9**, 90–95 (2007).
- Furbank, R. T. et al. Phenomics – technologies to relieve the phenotyping bottleneck. *Trends Plant Sci.* **16**, 635–644 (2011).
- Tardieu, F., Cabrera-Bosquet, L., Pridmore, T. & Bennett, M. Plant phenomics, from sensors to knowledge. *Curr. Biol.* **27**, R770–R783 (2017).
- ZHOU, J. et al. Plant phenomics: history, present status and challenges. *J. Nanjing Agric. Univ.* **41**, 580–588 (2018).
- Cendrero-Mateo M. P., et al. Field phenotyping: challenges and opportunities. *Terr. Ecosyst. Res. Infrastruct.* 53–80 (2017).
- Galinato, S. P. & Miles, C. A. Economic profitability of growing lettuce and tomato in western washington under high tunnel and open-field production systems. *HortTechnology* **23**, 453–461 (2013).
- Mulderij R. OVERVIEW GLOBAL LETTUCE MARKET. Tholen. <http://www.freshplaza.com/article/163938/OVERVIEW-GLOBAL-LETTUCE-MARKET> (2016).
- Shukla M, Jharkharia S. *Agri-fresh produce supply chain management: A state-of-the-art literature review*. <https://doi.org/10.1108/01443571311295608> (2013).
- Lamsal, K., Jones, P. C. & Thomas, B. W. Harvest logistics in agricultural systems with multiple, independent producers and no on-farm storage. *Comput. Ind. Eng.* **91**, 129–138 (2016).
- Schiefer, G. New technologies and their impact on the agri-food sector: an economists view. *Comput. Electron. Agric.* **43**, 163–172 (2004).
- Singh, A., Ganapathysubramanian, B., Singh, A. K. & Sarkar, S. Machine learning for high-throughput stress phenotyping in plants. *Trends Plant Sci.* **21**, 110–124 (2016).
- Pedregosa, F. et al. Scikit-learn: machine learning in python. *J. Mach. Learn Res.* **12**, 2825–2830 (2011).
- van der Walt, S. et al. Scikit-image: image processing in Python. *PeerJ* **2**, 1–18 (2014).
- Reynolds, D. et al. CropSight: a scalable and open-source information management system for distributed plant phenotyping and IoT-based crop management. *Gigascience* **8**, 1–35 (2019).
- Reynolds D., et al. What is cost-efficient phenotyping? Optimizing costs for different scenarios. *Plant Sci.* <https://doi.org/10.1016/j.plantsci.2018.06.015> (2018).
- Neveu, P. et al. Methods Dealing with multi-source and multi-scale information in plant phenomics: the ontology-driven Phenotyping Hybrid Information System. *New Phytol.* **221**, 588–601 (2018).
- Liu, S., Baret, F., Andrieu, B., Burger, P. & Hemmerlé, M. Estimation of wheat plant density at early stages using high resolution imagery. *Front Plant Sci.* **8**, 1–10 (2017).
- Frasconi, C. et al. An automatic machine able to perform variable rate application of flame weeding: design and assembly. *Chem. Eng. Trans.* **58**, 301–306 (2017).
- Brann D., Extension Grain Specialist, Soil Environmental Sciences, Virginia Technology. *A Comprehensive Approach Precision Farming: Management is the KEY*. Virginia Cooperative Extension (2009).