

An *a contrario* approach for plant disease detection

Rebecca Leygonie Sylvain Lobry Laurent Wendling

Université Paris Cité, LIPADE, F-75006 Paris, France

Rebecca.leygonie@etu.u-paris.fr

1. Introduction

Plant disease management is vital for agriculture [2, 10, 7]. The use of deep learning, and in particular computer vision models, is becoming increasingly common for the detection of plant pathologies, but comes up against a major problem: the lack of annotated data [5] which results in model overfitting. In addition, when the presence or visual nature of anomalies is uncertain, it is difficult to use learning-based methods to differentiate between normal and abnormal images. To overcome this problem, *a contrario* approaches have proved successful, in particular for locating anomalies in images [9]. These approaches, based on a reference model representing normality, evaluate the probability of an observation being normal or abnormal [6, 1].

We present an *a contrario* method that exploits the heatmaps of an explicability model to detect the essential features of an image. The model is trained to predict the age of a cauliflower, and its heat maps are used to determine whether an image is abnormal by comparing it with a reference distribution. The underlying assumption is that the important characteristics of an abnormal plant differ from those of a normal plant. Our method requires a sample of normal images, without any assumptions about other potential anomalies. For each new time series of images, a cauliflower at different ages, we calculate a score by comparing the probability that the heatmaps of cauliflowers at different ages match those of normal cauliflowers, and then compare it to a predefined threshold.

2. Methodology

2.1. Problem definition

Let $X = [x_1, x_2, \dots, x_T]$ be a temporal series of T images of size $(L \times L)$ pixels. Our goal is to detect whether X contains an anomaly using an *a contrario* approach based on heatmaps obtained by an explainability model. To achieve this, we train a predictive model on an auxiliary task. We propose to apply an explainability model, Grad-

CAM [8] that highlights the important areas of the image through heatmaps, indicating the regions that contributed to the model’s prediction. We obtain the temporal series of heatmaps H generated by Grad-CAM and calculate a score $S(H)$ to determine it contains an anomaly or not.

2.2. Score calculation

Heatmaps provide information about the areas of importance in images. We assume that these areas are different for normal and abnormal observations. Our objective is then to study the difference in distribution between an observation and a reference distribution constructed from so-called normal images. To do this, we calculate a score which allows us to assess for each observation the extent to which pixel values differ from pixel values in images from the reference distribution. As the value of a pixel is representative of its importance in the image, comparing the heatmap to a reference distribution means comparing whether the image’s areas of importance are the same as for the reference images, which are assumed to be normal.

Let H be a temporal series of T heatmaps, such that

$$H = [h_1, h_2, \dots, h_T]$$

where each heatmap is a matrix $h_{t(i,j)}$ for $i, j \in [1, L]^2$ and $t \in [1, T]$ corresponding to the image x_t .

We aim to calculate a score to determine if h_t is abnormal by estimating the probability that each pixel is drawn from the pixel distribution obtained from heatmaps of normal images at time t of the image time series of the training and validation datasets.

In other words, for an image of timestep t , for each pixel $h_{t(i,j)}$, we calculate the probability that it is drawn from the reference distribution denoted

$$\tilde{H}_{(i,j)}^t = [\tilde{h}_{1(i,j)}^t, \tilde{h}_{2(i,j)}^t, \tilde{h}_{3(i,j)}^t, \dots, \tilde{h}_{n(i,j)}^t],$$

which represents the pixel values $\tilde{h}_{(i,j)}^t$ of n heatmaps obtained from normal images at timestep t .

We compute $A(h_t)$, the average of the probabilities that each pixel is drawn from the reference distribution of pixels i, j at timestep t :

$$A(h_t) = \frac{\sum_{i,j=1}^L P(h_{t(i,j)} | \tilde{H}_{(i,j)}^t)}{L^2}. \quad (1)$$

For each time series H of T heatmaps, we calculate a score $S(H)$:

$$S(H) = \frac{\sum_{t=1}^T A(h_t)}{T}, \quad (2)$$

which represents the normality score of a heatmap time series. Please note that the score is not a probability, but rather an arithmetic mean of probabilities.

2.3. Anomaly Detection

To detect anomalies, we set a threshold to which we compare the score calculated for each new observation. The score represents how closely the observation belongs to the reference distribution. If the score is below the threshold, indicating a significant deviation from expected distributions, the observation is considered as an anomaly. The choice of threshold depends on the specific application and the desired trade-off between false positives and false negatives. Indeed, depending on the context, it may be constraining to predict that an observation is abnormal when it is not (false positive), or inversely to miss an abnormal observation by detecting it as normal (false negative).

3. Experiments and results

We are working with the GrowliFlower dataset [4], which contains georeferenced time series based on drone images of two cauliflower fields acquired in 2020 and 2021. We use a subset of the dataset called *GrowliflowerR*, which contains RGB orthophotos and phenotypic characteristics (in-situ data) collected on 740 plants, including information on whether cauliflowers are affected by a disease or defect. For the purposes of this study, we consider all defects, whatever their nature, to be anomalies. Our auxiliary task is to perform age classification, defined by the difference in days between the date the photo was taken and the date the cauliflower was planted. Ages range from 1 to 93 days.

In this study, we trained a *ResNet18* model [3] to predict the age of cauliflowers. The model showed promising performance on the test data, producing a mean square error (MSE) of 30.76 and an R^2 score of 0.95.

Once the age prediction model has been trained, we apply Grad-CAM to the predictions made by the model, highlighting areas of importance. We retrieve the heatmaps to apply our anomaly detection method.

We collect all normal observations from the training and validation datasets and define our age-based reference distributions. The dataset on which we aim to detect anomalies includes the *Test* dataset, as well as all the abnormal

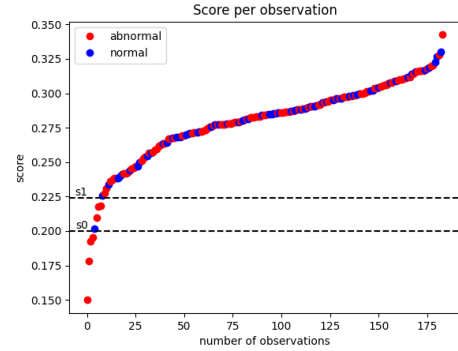


Figure 1. Score per observation. Red dots correspond to abnormal observations (which have a comment in the in-situ data) and blue dots are normal observations (with no comment in the in-situ data).

data from the *Train* and *Validation* datasets. We calculated a score for each cauliflower using the method described in Section 2 shown in Figure 1.

Our study is based on the hypothesis that the characteristics of importance are not the same for normal and abnormal cauliflowers, and that this information can be found in the heatmaps obtained by Grad-CAM.

Our approach does not result in a constant number of false alarms. However, we show that it is possible to set a small threshold below which all observations are considered anomalies (see s0 in Figure 1). This finding is important, as it highlights the possibility of using a predefined threshold to successfully identify a proportion of anomalies.

4. Conclusion

In this study, we proposed a decision-support model for anomaly detection, applied to a time series dataset of cauliflower images. The proposed model leverages spatial information learned from an auxiliary task and does not require prior assumptions about the presence of anomalies in the data. The only prerequisite is to have a sample of healthy data. The objective is to identify a subset of anomalies (setting a low threshold to minimize the number of false alarms), thereby assisting farmers in understanding plant phenotyping based on a more or less trivial learning task.

In the dataset under study, the anomalies associated with cauliflower do not exhibit temporal continuity. Consequently, while our current work deals with time series of images, we treat each image independently due to this lack of continuous anomalies. In our future research, we want to work with a dataset that exhibits continuous anomalies from a certain time step t , and exploit the temporal aspect of the images.

References

- [1] Agnes Desolneux, Lionel Moisan, and Jean-Michel Morel. *From gestalt theory to image analysis: a probabilistic approach*, volume 34. Springer Science & Business Media, 2007.
- [2] William E Fry. *Principles of plant disease management*. Academic Press, 2012.
- [3] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the Institute of Electrical and Electronics Engineers (IEEE) conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [4] Jana Kierdorf, Laura Verena Junker-Frohn, Mike Delaney, Mariele Donoso Olave, Andreas Burkart, Hannah Jaenicke, Onno Muller, Uwe Rascher, and Ribana Roscher. Growliflower: An image time-series dataset for growth analysis of cauliflower. *Journal of Field Robotics*, 40(2):173–192, 2023.
- [5] Jiajia Li, Dong Chen, Xinda Qi, Zhaojian Li, Yanbo Huang, Daniel Morris, and Xiaobo Tan. Label-efficient learning in agriculture: A comprehensive review. *arXiv preprint arXiv:2305.14691*, 2023.
- [6] David Lowe. *Perceptual organization and visual recognition*, volume 5. Springer Science & Business Media, 2012.
- [7] Serge Savary, Andrea Ficke, Jean-Noël Aubertot, and Clayton Hollier. Crop losses due to diseases and their implications for global food production losses and food security. *Food security*, 4(4):519–537, 2012.
- [8] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the international conference on computer vision*, pages 618–626, 2017.
- [9] Matias Tailanian, Pablo Musé, and Álvaro Pardo. A multi-scale a contrario method for unsupervised image anomaly detection. In *2021 20th Institute of Electrical and Electronics Engineers (IEEE) International Conference on Machine Learning and Applications (ICMLA)*, pages 179–184. IEEE, 2021.
- [10] Jan C Zadoks and Richard D et al. Schein. Epidemiology and plant disease management. *Epidemiology and plant disease management.*, 1979.