HYPERSPECTRAL PLANT DISEASE FORECASTING USING GENERATIVE ADVERSARIAL NETWORKS

Alina Förster[1], Jens Behley[1], Jan Behmann[2], Ribana Roscher[1]

[1] Institute of Geodesy and Geoinformation, University of Bonn
[2] INRES-Plant Diseases and Plant Protection, University of Bonn

ABSTRACT

With a limited amount of arable land, increasing demand for food induced by growth in population can only be meet with more effective crop production and more resistant plants. Since crop plants are exposed to many different stress factors, it is relevant to investigate those factors as well as their behavior and reactions. One of the most severe stress factors are diseases, resulting in a high loss of cultivated plants. Our main objective is the forecasting of the spread of disease symptons on barley plants using a Cycle-Consistent Generative Adversarial Network. Our contributions are: (1) we provide a daily forecast for one week to advance research for better planning of plant protection measures, and (2) in contrast to most approaches which use only RGB images, we learn a model with hyperspectral images, providing an information-rich result. In our experiments, we analyze healthy barley leaves and leaves which were inoculated by powdery mildew. Images of the leaves were acquired daily with a hyperspectral microscope, from day 3 to day 14 after inoculation. We provide two methods for evaluating the predicted time series with respect to the reference time series.

***Index Terms*—** hyperspectral phenotyping, plant disease, barley, generative adversarial networks, deep learning

1. INTRODUCTION

To reduce yield losses of crops caused by different diseases, a regular monitoring and a timely reaction by the farmer is necessary. Consequently, it is important for both scientists and farmers to study traits of the crop and the reaction to different stress factors. To this end, understanding the plant's phenotype, i.e., the result of environmental influences such as abiotic and biotic stress factors occurring during plant's development and growth on its genome, is essential [1]. Hyperspectral imaging (HSI) sensors are especially suitable for phenotyping and are therefore widely used for measuring spectral information to characterize different plant traits [2, 3, 4, 5]. The interest in phenotyping increased also in the computer vision and remote sensing community due to the need and development of non-invasive sensor techniques to quantify and measure plant traits. Besides classical machine learning methods, recently also deep learning has been used for phenotyping [6, 7, 8, 9, 10].

Our contribution is the forecasting of the spread of disease symptoms on barley plants given an early stage of the disease. We use a cycle-consistent generative adversarial network (CycleGAN), a special type of a generative adversarial network (GAN), to learn daily changes of the leaf and disease from hyperspectral images. Our experiments show that the learned model is able to forecast a probable time-series of images with an advancing disease spread.

2. DATA

In our experiments we use 16 leaves of barley (Hordeum vulgare). Seven days after sowing, leaves with a length of about $10\,\mathrm{cm}$ are cut off and transferred to culture mediums of phyto agar. A fraction of the leaves is inoculated with fresh spores of Blumeria graminis f.sp. hordei isolate K1 causing powdery mildew. Images were acquired with an HSI microscope achieving a spectral resolution of up to $2.3\,\mathrm{nm}$ in the range of $420\,\mathrm{nm} - 830\,\mathrm{nm}$. For more information, we refer to Kuska et al. [2].

The measurements of inoculated leaves were taken daily from day 3 to day 14 after inoculation (dai), from which dai 4 to 11 are considered for our experiments. For healthy leaves, dai 6 is missing in the provided data set. Due to movement and deformation of the leaves between measurements, we do not have pixel correspondences over time resulting in a non-aligned, unpaired image data set.

3. THEORY

3.1. Generative Adversarial Network

A generative adversarial network (GAN) [11] consists of two neural networks, namely the generator $\mathcal{G}$ and the discriminator $\mathcal{D}$. For imaging tasks, it uses two image domains $\mathcal{X}$ and $\mathcal{Y}$, where $\mathcal{X}$ contains the training data $\boldsymbol{x}$ and $\mathcal{Y}$ contains noise vectors $\boldsymbol{y}$. Based on a noise vector drawn from the distribution $p_{\boldsymbol{y}}(\boldsymbol{y})$, the generator creates realistic images pretending to come from the same distribution as the training data $p_{\text{data}}(\boldsymbol{x})$. Its goal is to fool the discriminator being no longer able to distinguish between a training images $\boldsymbol{x}$ and generated

images expressed by $\log(1 - \mathcal{D}(\mathcal{G}(\boldsymbol{x})))$. The discriminator's task is to distinguish between real images from domain $\mathcal{X}$ and generated images. The generator $\mathcal{G}$ has no access to training data, but learns to generate realistic images through backpropagation. Consequently, the respective objective function is given by the adversarial loss:

$$\begin{aligned}\min_{\mathcal{G}} \max_{\mathcal{D}} \mathcal{V}(\mathcal{D}, \mathcal{G}) = \mathbb{E}_{\boldsymbol{x} \sim p_{\text{data}}(\boldsymbol{x})} \left[\log \mathcal{D}(\boldsymbol{x})\right] \\ + \mathbb{E}_{\boldsymbol{y} \sim p_{\boldsymbol{y}}(\boldsymbol{y})} \left[\log\left(1 - \mathcal{D}(\mathcal{G}(\boldsymbol{y}))\right)\right].\end{aligned} \tag{1}$$

3.2. CycleGAN

To deal with unpaired data of two image domains $\mathcal{X}$ and $\mathcal{Y}$ with both including real images, CycleGAN consists of two generators $\mathcal{G}$ and $\mathcal{F}$. Whereas the generator $\mathcal{G}$ maps images of $\mathcal{X}$ to $\mathcal{Y}$, the generator $\mathcal{F}$ creates images of $\mathcal{Y}$ based on images in $\mathcal{X}$. Thus, $\mathcal{F}$ is the inverse mapping of $\mathcal{G}$. As a consequence of two generators, two independent discriminators are required: $\mathcal{D}_{\mathcal{X}}$ distinguishes between real images $\boldsymbol{x}$ and generated images $\mathcal{F}(\boldsymbol{y})$, and $\mathcal{D}_{\mathcal{Y}}$ differentiates between real images $\boldsymbol{y}$ and generated images $\mathcal{G}(\boldsymbol{x})$. The adversarial loss $\mathcal{L}_{\text{GAN}}$ is defined separately for both generators similar to Eq. (1).

We follow the procedure of [12] and introduce two cycle consistency losses to stabilize the training. A cycle consistency loss ensures a small difference between an image in domain $\mathcal{X}$ and the generated image in domain $\mathcal{Y}$ transferred back to domain $\mathcal{X}$, and vice versa. The complete cycle consistency loss results in:

$$\begin{aligned}\mathcal{L}_{\text{cyc}}(\mathcal{G}, \mathcal{F}) = \mathbb{E}_{\boldsymbol{x} \sim p_{\text{data}}(\boldsymbol{x})} \left[\|\mathcal{F}(\mathcal{G}(\boldsymbol{x})) - \boldsymbol{x}\|_1\right] \\ + \mathbb{E}_{\boldsymbol{y} \sim p_{\text{data}}(\boldsymbol{y})} \left[\|\mathcal{G}(\mathcal{F}(\boldsymbol{y})) - \boldsymbol{y}\|_1\right].\end{aligned} \tag{2}$$

Consequently, the definition of the full objective function is:

$$\begin{aligned}\mathcal{L}(\mathcal{G}, \mathcal{F}, \mathcal{D}_{\mathcal{X}}, \mathcal{D}_{\mathcal{Y}}) = &\mathcal{L}_{\text{GAN}}(\mathcal{G}, \mathcal{D}_{\mathcal{Y}}, \mathcal{X}, \mathcal{Y}) \\ &+ \mathcal{L}_{\text{GAN}}(\mathcal{F}, \mathcal{D}_{\mathcal{X}}, \mathcal{X}, \mathcal{Y}) + \lambda \mathcal{L}_{\text{cyc}}(\mathcal{G}, \mathcal{F}),\end{aligned} \tag{3}$$

where λ describes the relative importance of Eq. (2). Finally, Eq. (3) is minimized for both generators and maximized for both discriminators alternatively.

4. EXPERIMENTAL SETUP

4.1. Data Preprocessing

For pre-processing, we follow the procedure of Kuska et al. [2]. We removed noisy bands at the beginning and the end of the spectral range resulting in 420 bands between $450\,\text{nm}$ and $780\,\text{nm}$. For our experiments, we considered 50 equally sampled wavelengths. For training and testing, we extracted image patches of the size 100×100 pixels. We collected 3500 image patches with healthy leaves and 1750 image patches showing inoculated leaves. For data augmentation, we mirrored all patches and get a final training data set of 7000 patches. For evaluating our forecasting approach, we aligned the images over time using a feature based method, and selected 180 patches with successful alignment, containing both healthy and inoculated leaves. The input to our approach consists of patches of dai 4 only. Please note that this procedure needs human interaction and is therefore only used for evaluation purposes. We extended the data by the information of the dai by appending each image patch with a matrix in the 3rd dimension containing the value of the dai. This is done for both training and testing data. Domain $\mathcal{X}$ includes patches from dai 4 to 10, domain $\mathcal{Y}$ consists of patches from dai 5 to 11.

4.2. Training and Testing the Network

We use an implementation[1] based on the work of Zhu et al. [12]. In order to be able to use hyperspectral images, we have adapted the given implementation accordingly. We use the forward generator $\mathcal{G}$ which has learned the daily change of the data. First, we feed an input test image of dai 4 to the learned model and generate an image of dai 5. To generate an image of dai 6 based on the test image, we use the generated image of dai 5 as input to the learned model. This step is repeated until dai 11. After each forecasting step, we correct incorrect estimated dai information in the last band to the true dai information which we assume to be known. Though, this step can be replaced by an automatic procedure which is able to quantify the strength of the disease.

4.3. Evaluation

Since the trained generative model is not limited to generate only one possible forecast, standard evaluation methods with a pixel-wise comparison are not applicable. Therefore, we apply the following two evaluation criteria: First, in the visual evaluation for inoculated images we use RGB visualization and compare the spread of the disease to those of the reference images including appearance of the disease spots like shape, color, and intensity. Moreover, we analyze the leaf itself. Second, we investigate the temporal progression of hyperspectral spectra for diseased pixels, where we manually selected three pixels in each image. We apply a 5×5 window around the chosen pixels in which we compute the average spectrum. Moreover, we averaged all spectra in one image.

5. RESULTS

Fig. 1 shows two representative examples of the future prediction for test images with inoculated leaves, where row 1 and 2 shows the first example, and row 3 and 4 shows the second example. The respective upper row shows the future prediction based on the input of dai 4. The results show a daily spread of the disease in the generated time series. The

[1] Available at `https://github.com/leehomyc/cyclegan-1`

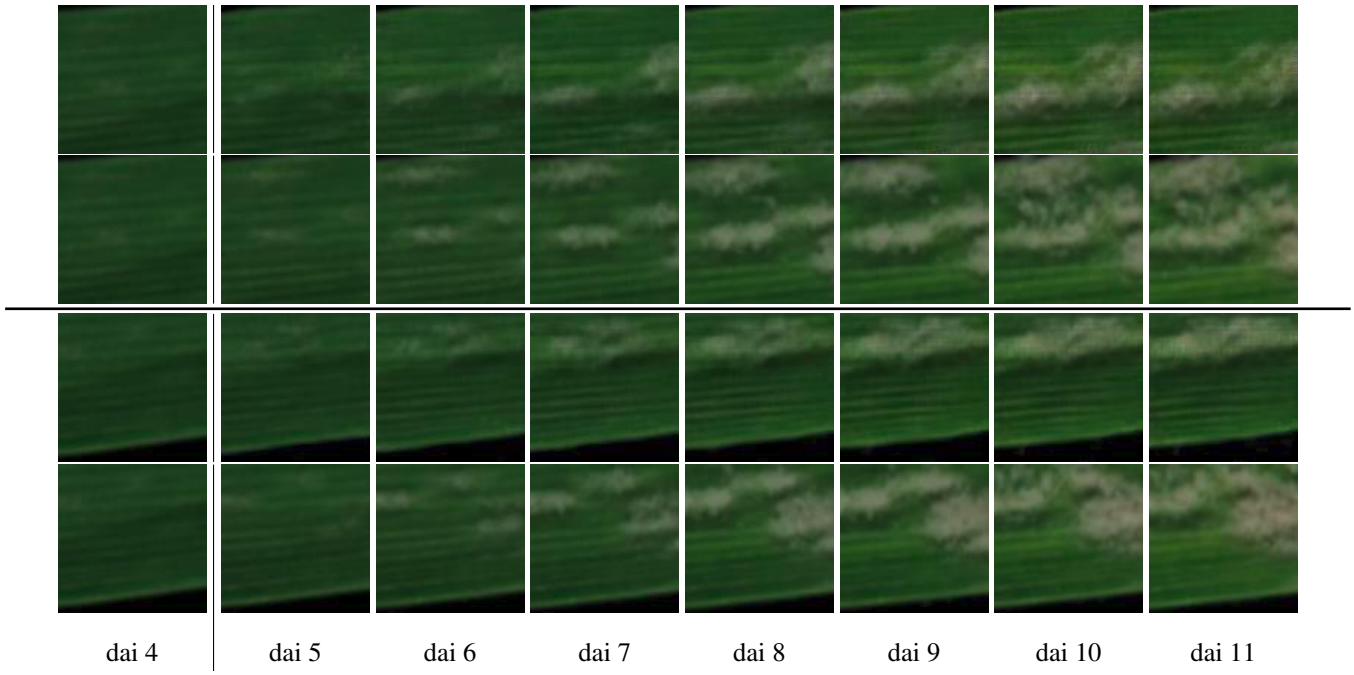

Fig. 1: RGB visualization of two representative examples of the future prediction of hyperspectral input images of dai 4 (example 1: top two rows, example 2: bottom two rows). Row 1 and 3 show the future prediction generated by the learned model and row 2 and 4 show the corresponding reference.

disease appearance shows a similar behavior in comparison to the reference. For the first example, we observe a slight decrease of the disease from dai 10 to dai 11, which usually happens between dai 9 and 10 for older disease spots. The second example shows similar results as the first one. However, in contrast to the first one, we cannot detect a decrease within the last two days. For both predicted time series, the leaf around the disease spots gets brighter over time. This finding coincides with the reference time series. We observe an underestimation of the spread in all test examples. A possible explanation is the neglecting of newly emerged disease hot spots in the reference after dai 4, what may not have been learned by the model. Another possibility is the insufficient amount of training data, which may have led to a non representative model.

The temporal progression of the reflectance spectrum for the first presented time series is given in Fig. 2. We expect a daily increase in the reflectance in the visible part of the spectrum as a result of the disease, since white powdery mildew leads to a high reflectance within the first few days covering the healthy leaf and to the withdrawing of nutrients from the leaf towards the end causing in a breakdown of pigments. This process can be seen in the left figure representing the reflectance spectra of the reference time series. The figure shows a daily increase of the reflectance values in the visible part of the spectrum, which is most visible between dai 5 and 6, 6 and 7, and 9 and 10. Also the near infrared (NIR) spectrum changes over time, however, it is not consistent through all evaluated test samples. In most cases, the reflectances of the NIR also increase, but in some cases the NIR decrease over time, or increase and decrease alternatively. The latter case is given in the illustrated time series. However, this behavior is not unusual for plants [2].

The figure on the right shows the progression of the predicted time series depicting also a typical behavior. The higher increase of the visible part of the spectrum within dai 5 and 6, 7 and 8, and 8 and 9 indicates a higher spread of the disease within these days. This finding coincides with the illustrated time series. Furthermore, we recognize the spectral decrease within the last two days that is consistent with the visual observations. Yet, the reflectance spectrum gets rougher every day, especially in the NIR spectrum. We suspect that noise in the input image is intensified by the prediction.

6. CONCLUSION

We proposed an approach to forecast the spread of powdery mildew on barley leaves using hyperspectral images. To this end, we applied cycle-consistent adversarial networks. Our experiments show that the model is able to learn the spread of the disease, which is shown both visually and in the respective reflectance spectra. Moreover, our model achieves suitable results for a prediction over seven days.

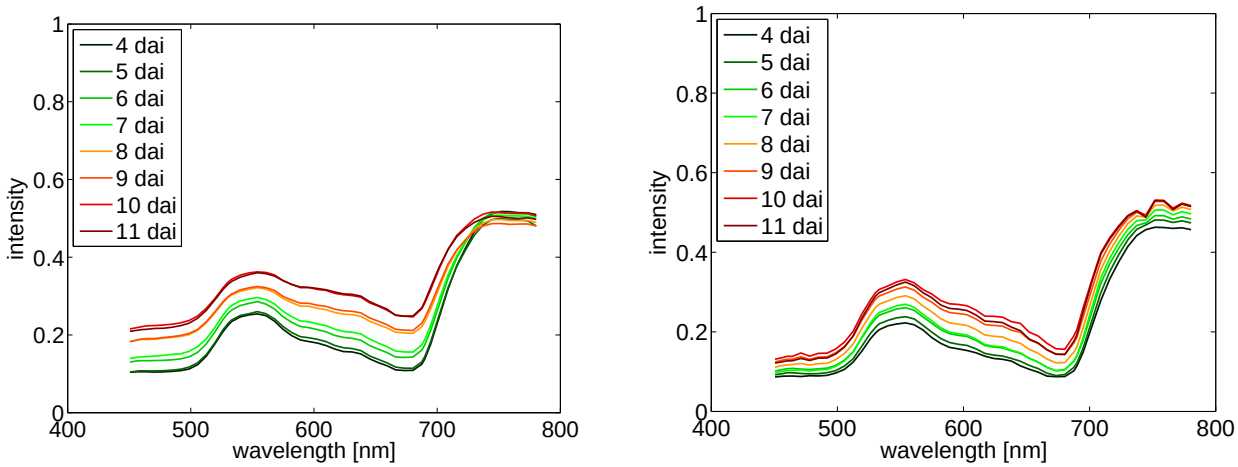

Fig. 2: Temporal progression of the intensity spectrum of selected pixels for the reference time series of row 2 in Fig. 1 (left) and for the generated time series of row 1 (right).

7. ACKNOWLEDGMENTS

We would like to thank Dr. Matheus T. Kuska and Dr. Anne-Katrin Mahlein for providing the hyperspectral data.

8. REFERENCES

[1] A. Walter, F. Liebisch, and A. Hund, "Plant phenotyping: from bean weighing to image analysis," *Plant methods*, vol. 11, no. 1, pp. 14, 2015.

[2] M. Kuska, M. Wahabzada, M. Leucker, H.-W. Dehne, K. Kersting, E.-Ch. Oerke, U. Steiner, and A.-K. Mahlein, "Hyperspectral phenotyping on the microscopic scale: towards automated characterization of plant-pathogen interactions," *Plant Methods*, vol. 11, no. 1, pp. 28, 2015.

[3] M. Wahabzada, A.-K. Mahlein, Ch. Bauckhage, U. Steiner, E.-Ch. Oerke, and K. Kersting, "Metro maps of plant disease dynamics—automated mining of differences using hyperspectral images," *PLoS One*, vol. 10, no. 1, pp. e0116902, 2015.

[4] M. Wahabzada, A.-K. Mahlein, Ch. Bauckhage, U. Steiner, E.-C. Oerke, and K. Kersting, "Plant phenotyping using probabilistic topic models: uncovering the hyperspectral language of plants," *Scientific reports*, vol. 6, pp. 22482, 2016.

[5] C. R. Yendrek, T. Tomaz, Ch. M Montes, Y. Cao, A. M. Morse, P. J. Brown, L. M. McIntyre, A. Leakey, and E. A. Ainsworth, "High-throughput phenotyping of maize leaf physiological and biochemical traits using hyperspectral reflectance," *Plant physiology*, vol. 173, no. 1, pp. 614–626, 2017.

[6] S. Ghosal, D. Blystone, A. K. Singh, B. Ganapathysubramanian, A. Singh, and S. Sarkar, "An explainable deep machine vision framework for plant stress phenotyping," *Proc. of the National Academy of Sciences*, vol. 115, no. 18, pp. 4613–4618, 2018.

[7] J. R Ubbens and I. Stavness, "Deep plant phenomics: a deep learning platform for complex plant phenotyping tasks," *Frontiers in plant science*, vol. 8, pp. 1190, 2017.

[8] Ch. DeChant, T. Wiesner-Hanks, S. Chen, E. L. Stewart, J. Yosinski, M. A. Gore, R. J. Nelson, and H. Lipson, "Automated identification of northern leaf blight-infected maize plants from field imagery using deep learning," *Phytopathology*, vol. 107, no. 11, pp. 1426–1432, 2017.

[9] Sh. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Frontiers in plant science*, vol. 7, pp. 1419, 2016.

[10] S. H. Lee, Y. L. Chang, C. S. Chan, and P. Remagnino, "Plant identification system based on a convolutional neural network for the lifeclef 2016 plant classification task.," in *CLEF (Working Notes)*, 2016, pp. 502–510.

[11] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.

[12] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. of the IEEE International Conference on Computer Vision*, 2017, pp. 2223–2232.