MULTIFUNCTIONAL PLATFORM AND MOBILE APPLICATION FOR PLANT DISEASE DETECTION

A. Uzhinskiy [a], G. Ososkov, P. Goncharov, A. Nechaevskiy

Joint Institute for Nuclear Research, 6 Joliot-Curie, Dubna, Moscow region, 141980, Russia

E-mail: [a] auzhinskiy@jinr.ru

Crop losses are the major threat to the wellbeing of rural families, to the economy and governments, and to food security worldwide. We present a multifunctional platform for plant disease detection (PDDP). PDDP consists of a set of interconnected services and tools developed, deployed, and hosted with the help of the JINR cloud infrastructure. PDDP was designed using modern organization and deep learning technologies to provide a new level of service to the farmers' community. A mobile application allowing users to send photos and text descriptions of sick plants and get the cause of the illness and treatment is part of PDDP. We collected a special database of grape, wheat and corn leaves consisting of fifteen sets of images. We tried different neural network architectures on these data and selected the best one. The architecture and basic principles of the platform and networks are described and compared with other well-known solutions.

Keywords: Siamese networks, convolutional neural networks, deep learning, plant disease detection

Alexander Uzhinskiy, Gennady Ososkov, Pavel Goncharov, Andrey Nechaevskiy

1. Introduction

Plant diseases are a serious threat to the economy and food security worldwide. According to some well-known studies, crop losses by diseases are between 10 and 30% [1]. An increasing number of smartphones and advances in deep learning can help with this problem. We started this project in 2018. By that time, there were many kinds of studies in which deep learning was used to identify plant diseases. Some of them report about a great detection level, more than 96%. Generally, researchers use transfer learning approaches and images from PlantVillage [2] (an open at that time database with 54,306 images of 14 crop species) or self-collected databases. However, there was a lack of a real application or sites where one could upload an image and get a prediction. The only mobile application we found that really could recognize plant diseases was Plantix [3]. Back to 2018, the Plantix accuracy of detection on our test subset of 70 images was over 15%.

We tried to reproduce some of the studies and get good results with detection of grape diseases on PlantVillage images – over 99% accuracy, but accuracy on a test subset from the Internet was less than 50% [4]. The problem was in the synthetic nature of PlantVillage images – same light, background and leaves orientation. We could not find any alternatives to PlantVillage, so we had to create our own database of diseased leaves. We understand that to facilitate the detection and prevention of diseases of agricultural plants we should not only develop a good model but also create all necessary environments to work with it. That is why we decided to develop a multifunctional platform that should use modern organization and deep learning technologies to provide a new level of service to the farmers' community.

2. Architecture and abilities

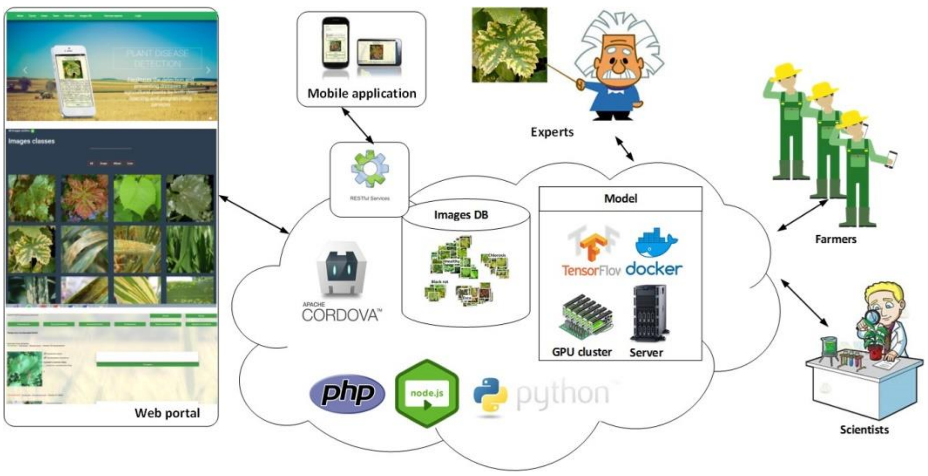

Figure 1. Architecture of the platform

PDDP consists of a set of interconnected services and tools developed, deployed and hosted at the Joint Institute for Nuclear Research cloud infrastructure [5]. It provides the necessary scalability of the solution and if some part of the platform requires more resources, they can be easily allocated.

Users communicate with PDDP through a web portal (pdd.jinr.ru), a mobile application or web services. The web portal has public and private parts to provide all necessary interfaces for work and communication to users, experts, and supervisors. The image database is open and free for download. The TensorFlow model is implemented as a Tensorflow serving in the Docker container, so it can work at the virtual server or the GPU cluster.

PDDP users can do the following: send photos and text descriptions of sick plants through the web interface or mobile application and get the cause of the illness; browse through disease descriptions and galleries of ill plants; verify that the requested disease was recognized correctly and the treatment helped.

PDDP experts can browse user requests and verify the correctness of the recognition; request addition of their image or the image from user requests to the DB; request changes of the description of the disease; request retraining of the model with new images.

PDDP supervisors can add new images to the database; initiate retraining of the model; get different statistical metrics about portal users.

Researchers can download all or only part of the base, work with the image database through the web interface or API.

3. Mobile application

PDDP users can run recognition tasks from the private or public part of the web portal, but we believe that the most convenient way is a mobile application. We developed the mobile application using Apache Cordova, so we can build it for Android, iOS, and Windows. Currently, we deployed only the Android version that can be found at Google Play under the "PDDApp" name.

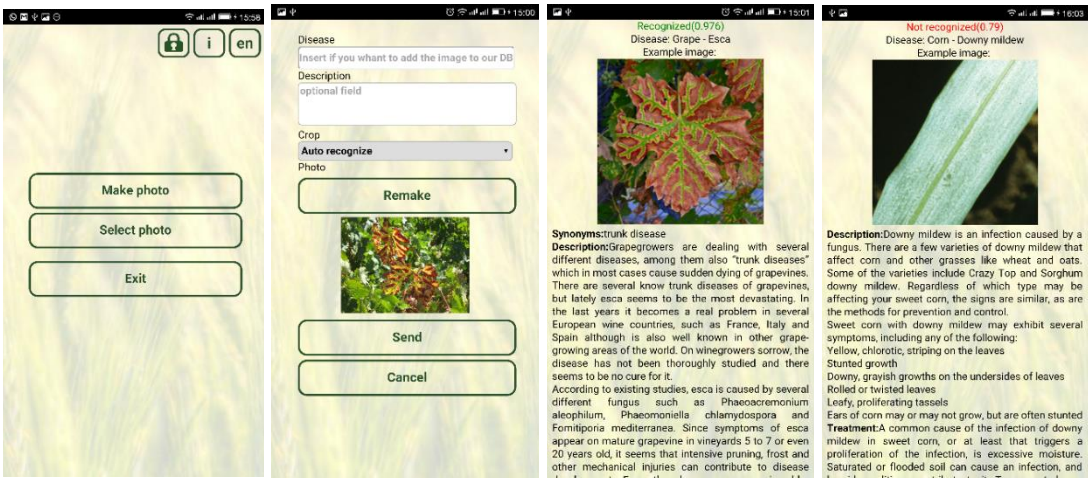

Figure 2. Examples of the PDDApp interfaces

A user has the opportunity to take a photo of the diseased plant and get a prediction for the disease and treatment suggestions. It is possible to download images if the user could not take a photo. The application requires access to the Internet to work. We have tried to run the model on the mobile device and managed to decrease the size of the model by ten times without serious accuracy loss. We are going to implement an offline mode for the application when crops and disease description data settle down.

4. Model and image database

The most popular way to deal with image classification problems in a vast majority of domains is to use a deep neural network trained on a big dataset with further fine-tuning of the chosen deep classifier on your dataset. We made our comparative study of transfer learning models that are available in open access and found out that the ResNet50 architecture reached 99.4% classification accuracy on a test subset of PlantVillage data but was stuck on our self-collected dataset with unsatisfactory 48%. We investigated the problem and discovered that it referred to the type of images used. PlantVillage photos were collected and processed under special controlled conditions, so they are rather synthetic and differ from real-life images. It gave us the idea of creating our own database. At the very beginning, our database had only 5 classes of grape leaves (healthy, esca, chlorosis, powdery mildew and black rot) – 313 images total. The only way to train a deep neural network on a small dataset is one-shot learning, in particular, Siamese networks [6]. The Siamese network consists of twin networks joined by the similarity layer with the energy function at the top. Weights of twins are tied (the same), thus, the result is invariant and guarantees that very similar images cannot be in very different locations in the feature space. The similarity layer determines some distance metric between so-called embeddings, i.e. high-level feature representations of the input pair of images. Training on pairs is more beneficial since it produces quadratically more possible pairs of images to train the model on, making it hard to overfit. From the trained one-shot model, the encoder network represented as a «shoulder» of this model or a so-called twin is extracted for further use as a feature extractor. The role of the classifier takes the k-nearest neighbors algorithm, which operates on the feature vectors - outputs of the trained twin. Cosine similarity was applied for the distance metric. The parameter K was set to 1 to be equivalent to the one-shot learning task. The classification accuracy of the model was measured on a test subset of grape images and reached 95% using all five classes [4].

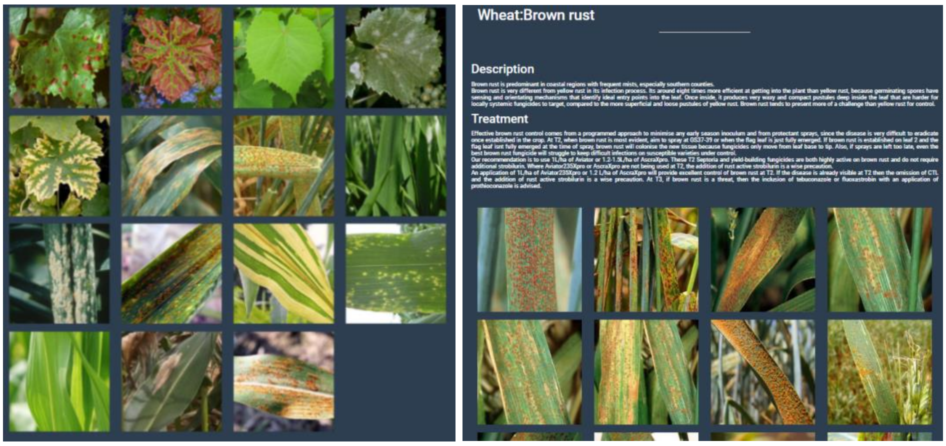

Figure 3. PDDP image database

We have expanded the PDDP database since the previous results were published. We added two other crops – wheat and corn. Each of the added crops is represented by 5 sets of images: corn (diseases: downy mildew, eyespot, northern leaf blight, southern rust and healthy) and wheat (diseases: black chaff, brown rust, powdery mildew, yellow rust and healthy). The final version of the dataset includes 15 classes with 611 leaf photos in total. After training on all 15 classes even for 150 epochs, the best version of the model obtained an 86% test classification accuracy. Probably, such a decrease in accuracy value may be caused by using KNN as a classifier. It is a well-known fact that KNN suffers from hubs when working with high-dimensional data. A hub is a node which tends to have much more in-going edges than the other nodes. To deal with hubs one can reweight all distances using special scaling parameters or simply replace KNN with another classifier.

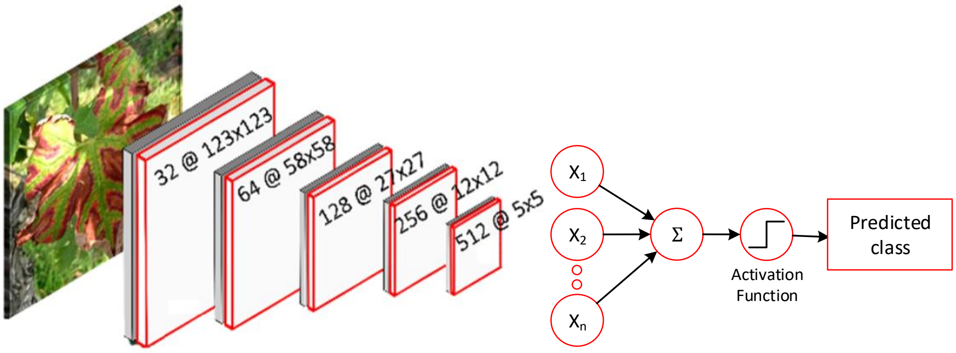

Figure 4. Best NNA architecture: one of two Siamese twins and single-layer perceptron

To improve the test classification accuracy we made a special comparative study of different types of estimators including logistic regression, support vector machines with cosine similarity as a kernel, decision tree, random forest, gradient boosting and a simple single-layer perceptron with one input and one output layer ending with softmax activation. The single-layer perceptron being trained for 100 epochs with the Adam optimizer allows us to obtain the classification accuracy equal to 95.71% on a test subset of images. The best architecture we created is presented in Figure 4.

5. Alternatives

By September 2019, the only known alternative to our solution was Plantix. Plantix models have improved a lot over the last year, and detection accuracy on the test subset of 70 images now is over 50%. Fortunately, there is no information about their models, and their image database is closed.

AutoML solutions have become increasingly popular over the past few years, helping non-machine learning experts solve problems of image recognition and classification. AutoML services allow users to upload their datasets, automatically select and train machine-learning models and provide interfaces to use models. We decided to compare our models with several commercial AutoML platforms: Google Cloud Vision [7], Microsoft Custom Vision [8], and IBM Watson Visual Recognition [9]. We created a test subset of images consisting of 30 images that were used for model training, 30 images that were not used for training and 20 images out of our crop diseases domain. The results are presented in Table 1. As one can see, our new model has a detection level similar to the models created by commercial platforms.

Table 1. Comparison of detection accuracy as the number of correctly recognized images for each group of PDDP and AutoML models (except for the last line, which shows the number of misclassified images)

	Old model	New model	Google Cloud Vision	Microsoft Custom Vision	IBM Watson Visual Recognition
Known (30)	27	29	28	29	29
Unknown (30)	20	24	22	25	25
Not in domain (20)	0	5	1	7	2

6. Acknowledgement

The reported study was funded by RFBR according to the research project № 18-07-00829.

7. Conclusion

We developed PDDP to facilitate the detection and prevention of diseases of agricultural plants. Our web portal and mobile application are ready to use. We have a database of 3 crops and 15 classes, 613 images total, that can be downloaded from pdd.jinr.ru. We developed a special Siamese transfer learning method, which leads to a significant increase in accuracy. We compared our solution with some well-known AutoML products and showed that our model detected diseases well.

We are going to expand our image DB and improve the mobile App and the web portal. We will explore other types of Siamese loss functions (triplet loss) and optimize the existing deep neural network architecture. We are working on a model for classification by text description. Currently, we support Russian and English languages. The Arabic language is also in our plans.

References

[1] S. Savary, A. Ficke, Jean-Noël Aubertot, et al., Crop losses due to diseases and their implications for global food production losses and food security//Springer Food Security, 4, pp. 519–537, 2012

[2] PlantVillage project home page [Electronic resource]: https://plantvillage.psu.edu/

[3] Plantix project home page [Electronic resource]: https://plantix.net (Accessed 1.10.2019)

[4] Goncharov, P., Ososkov, G., Nechaevskiy, et al., I. Disease Detection on the Plant Leaves by Deep Learning. International Conference on Neuroinformatics. Springer, Cham, pp. 151-159, 2018

[5] Korenkov V., Balashov N., Kutovskiy N., et al., Clouds of JINR, University of Sofia and INRNE — current state of the project, CEUR Workshop Proceedings, Vol. 2267, pp. 248-251, 2019

[6] Koch, G., Zemel, R., Salakhutdinov, R. Siamese neural networks for one-shot image recognition. In: ICML Deep Learning Workshop, Vol. 2, 2015

[7] Google Cloud Vision project home page [Electronic resource]: https://cloud.google.com/vision
[8] Microsoft Custom Vision project home page [Electronic resource]: https://www.customvision.ai/

[9] IBM Watson Visual Recognition project home page [Electronic resource]: https://www.ibm.com/watson/services/visual-recognition/ (Accessed 1.10.2019)