



Disease detection on the plant leaves by deep learning

P. Goncharov, G. Ososkov, A. Nechaevskiy,
A. Uzhinskiy, I. Nestsiaenia

Problem

Crop losses are a major threat to the wellbeing of rural families, to the economy and governments, and to food security worldwide. Quality of available data about the impact of plant diseases is variable, patchy and often missing, particularly for smallholders, who produce the majority of the world's food.

- CIP, the international research center with an historical mandate for potato, estimates **15% production losses each year due to late blight**
- USAblight (a national project on late blight on potato and tomato) says that (annual) **global losses 'exceed US\$6.7 billion'**.

As agriculture struggles to support the rapidly growing global population, plant disease reduces the production and quality of food, fibre and biofuel crops. Losses may be catastrophic or chronic, but on average account for **32% of the production of the six most important food crops**.

Prof. Dr. David Guest. Faculty of Agriculture, Food and Natural Resources, The University of Sydney

Globally, about 16% of all crops are lost to plant diseases each year.

Dr. Caitilyn Allen Department of Plant Pathology, University of Wisconsin–Madison

Our goal

Increasing number of smartphones and advances in deep learning field opens new opportunities in the crop diseases detection.

The idea is:

- **to create the plant disease detection platform (PDDP)** that will use modern organization and deep learning technologies **to provide new level of service to farmer's community;**
- provide **open access** to our image database;
- **share the code** of the models via Github;
- as end-product we are going **to develop the mobile application** allowing users to **send photos and text description of sick plants and get the cause of the illness.**

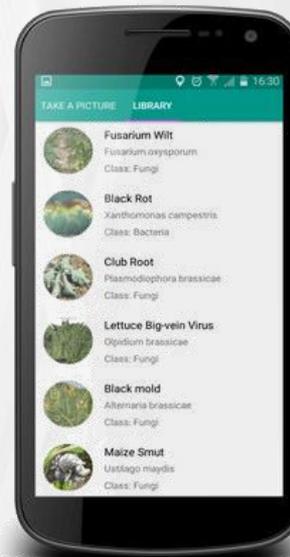


Sure, we are not the first

Probably, the most famous mobile application for plants disease detection is Plantix (plantix.net). Currently Plantix can detect more than 300 diseases.

The quality of plantix detection is hard to measure, but we made a special study processing different types of images from our self-collected database with grape diseases (Esca, Black rot, Chlorosis and Mildew).

- identification of the plants type is rather good: 60 of 70 images (87%) were recognized as grapes;
- the disease detection ability is rather limited: **less than 10% images had right disease at the top of suggestions.**



PlantVillage dataset

[\[https://arxiv.org/ftp/arxiv/papers/1511/1511.08060.pdf\]](https://arxiv.org/ftp/arxiv/papers/1511/1511.08060.pdf)

[\[https://github.com/spMohanty/PlantVillage-Dataset\]](https://github.com/spMohanty/PlantVillage-Dataset)

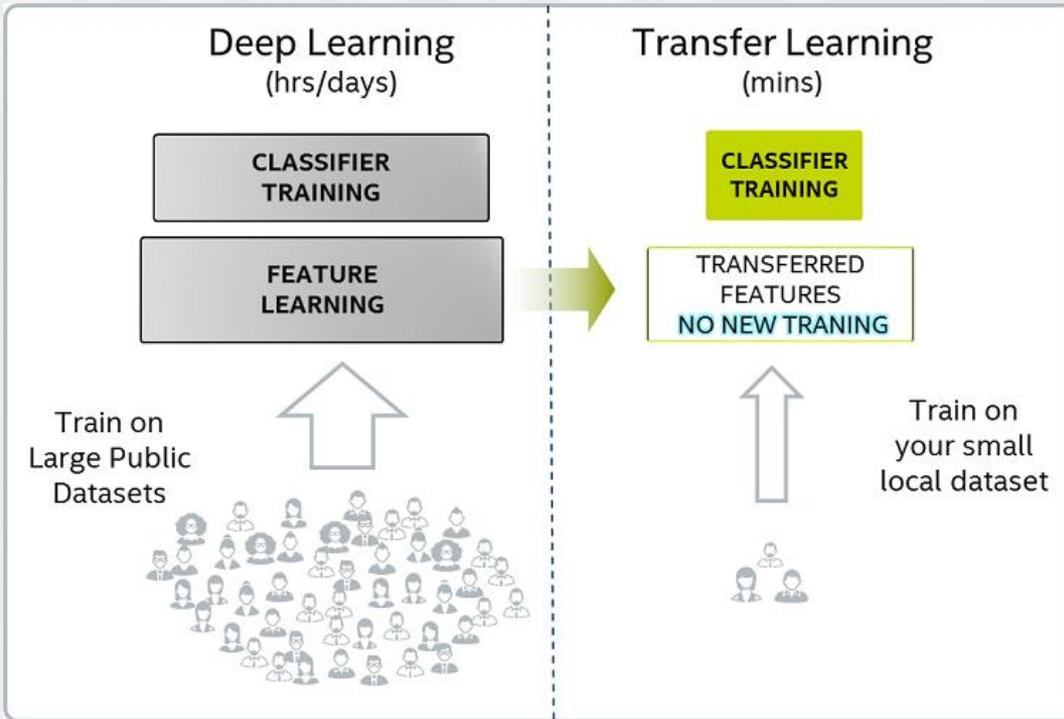
[\[https://www.crowdai.org/challenges/1\]](https://www.crowdai.org/challenges/1)

PlantVillage statistics:

- 54,306 images;
- 14 crop species;
- with 26 diseases (or healthy);
- 38 classes of both healthy and diseased leaves



Transfer learning



Steps:

1. **find a deep network pretrained** on a big dataset;
2. remove classification layer, add global average pooling followed by layer with 256 ReLU neurons with dropout of 0.5, append the classification layer;
3. **finetune** the choosed deep classifier **on the PlantVillage grapes images**;
4. **evaluate** it **on a test subset** of images, collected **from the Internet**.

We compared four models, which weights were formed to solve the ILSVRC 2015 (ImageNet Large Scale Visual Recognition Challenge), they are: **VGG19**, **InceptionV3**, **ResNet50** and **Xception**.

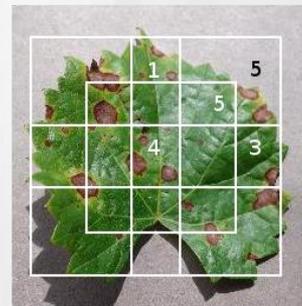
Transfer learning: unsuccessful results

ResNet50 architecture showed the best result comparing with the other ones:

- accuracy on a test subset of the PlantVillage grapes data – **99.4%**;
- accuracy on 30 images collected from the Internet – **48%**.

What we did to somehow improve the result:

- we unfroze 39 layers from the end of the network;
- applied a strong data augmentation;
- also, we supposed that only central part of a leaf is required to recognize the disease, we tried to expand our dataset by using only parts of initial images.



In spite of all, we couldn't reach over 70% of testing accuracy.



It turns out that, when we crop images, we lose the information about location of spots and it produces very noisy data. Besides some parts of diseased leaves looks healthy and the automatic split to the parts produces incorrect examples.

Data is a key

Attentively look on this picture. First row – PlantVillage images, second row – real-life images from the Internet.
Do you see anything strange? **HINT:** background



There are two different distributions of the data!

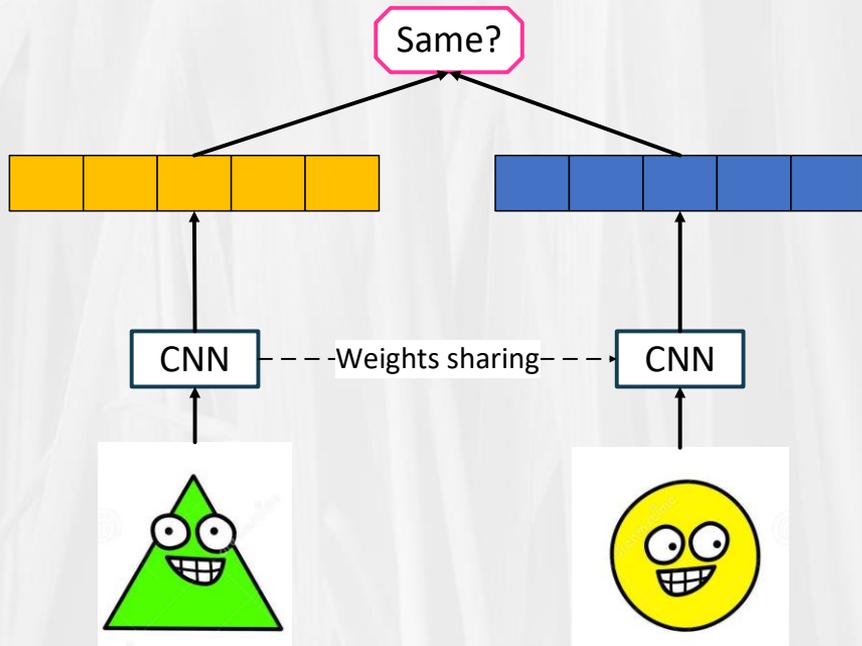
We decided to create our own dataset!

PDD dataset (<http://pdd.jinr.ru/db/>):

- started from **grapes images**;
- **4 diseases** – Esca, Black rot, Chlorosis, Mildew;
- **133 healthy images**;
- in total – **313 samples** with the size of 256x256;
- we are going add more crops to the dataset.

But 313 images is too small even for transfer learning!

Siamese network



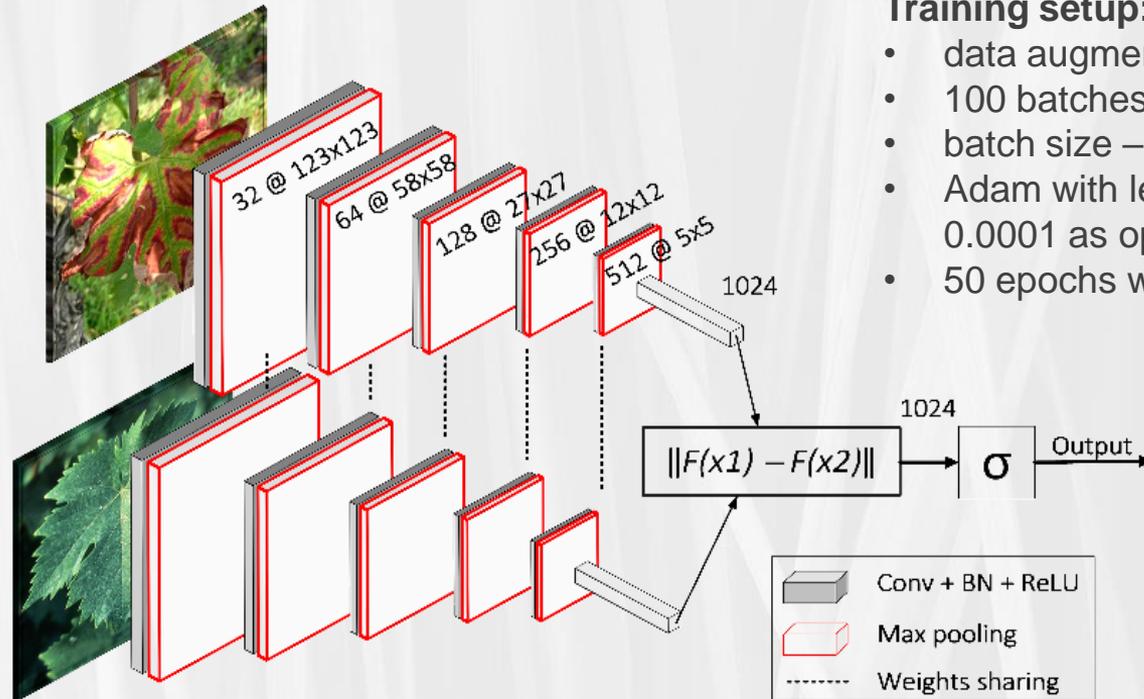
Siamese network consists of twin networks joined by the **similarity layer** with energy function at the top.

Weights of twins are tied, thus the result is invariant and in addition guarantees that very similar images can not be in very different locations in feature space.

The similarity layer determines some distance metric between so-called embeddings, i.e. high-level features representations of input pair of images.

Training on pairs is that there are quadratically many possible pairs of images to train the model on, making it hard to overfit.

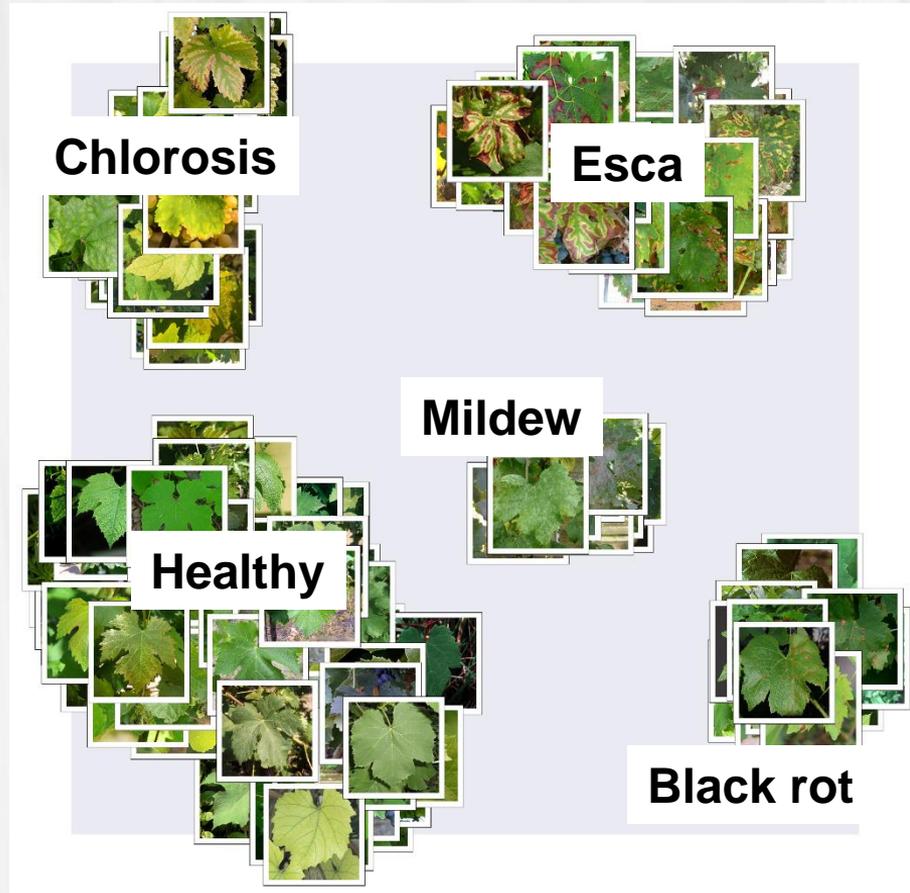
Siamese network: architecture and training setup



Training setup:

- data augmentation;
- 100 batches per epoch;
- batch size – 32;
- Adam with learning rate equals to 0.0001 as optimizer;
- 50 epochs with early stopping;

Visualization of the data embeddings



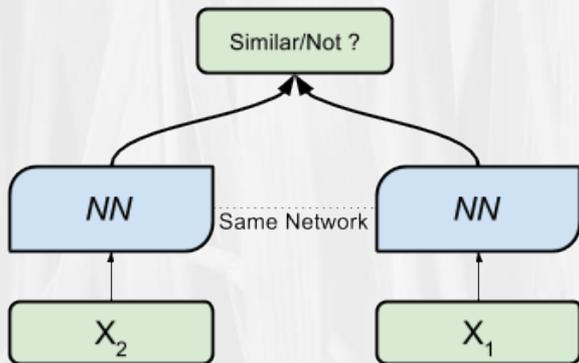
After training the siamese network, we used the high-level representations extracted by one of the twins to train the T-SNE to visualize the data distribution.

We extracted two components to plot them in 2D space.

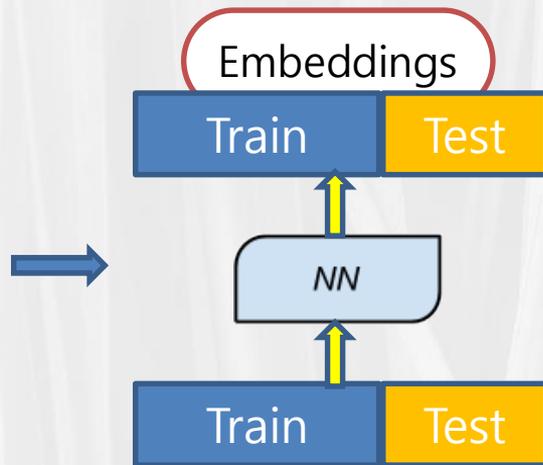
One can see that there are five separate clusters – one per each class. We have signed each cluster with the name of particular class. Although, there are a few points, which wrongly got into the different set.

K-Nearest Neighbors as classifier

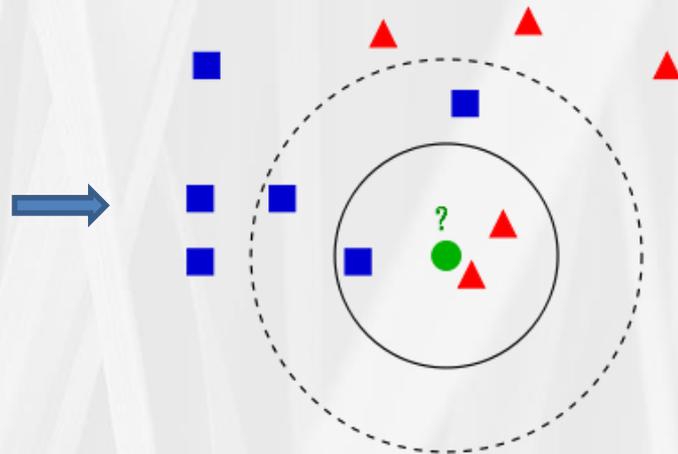
Step 1: Train the siamese network



Step 2: Use one of twins to get embeddings



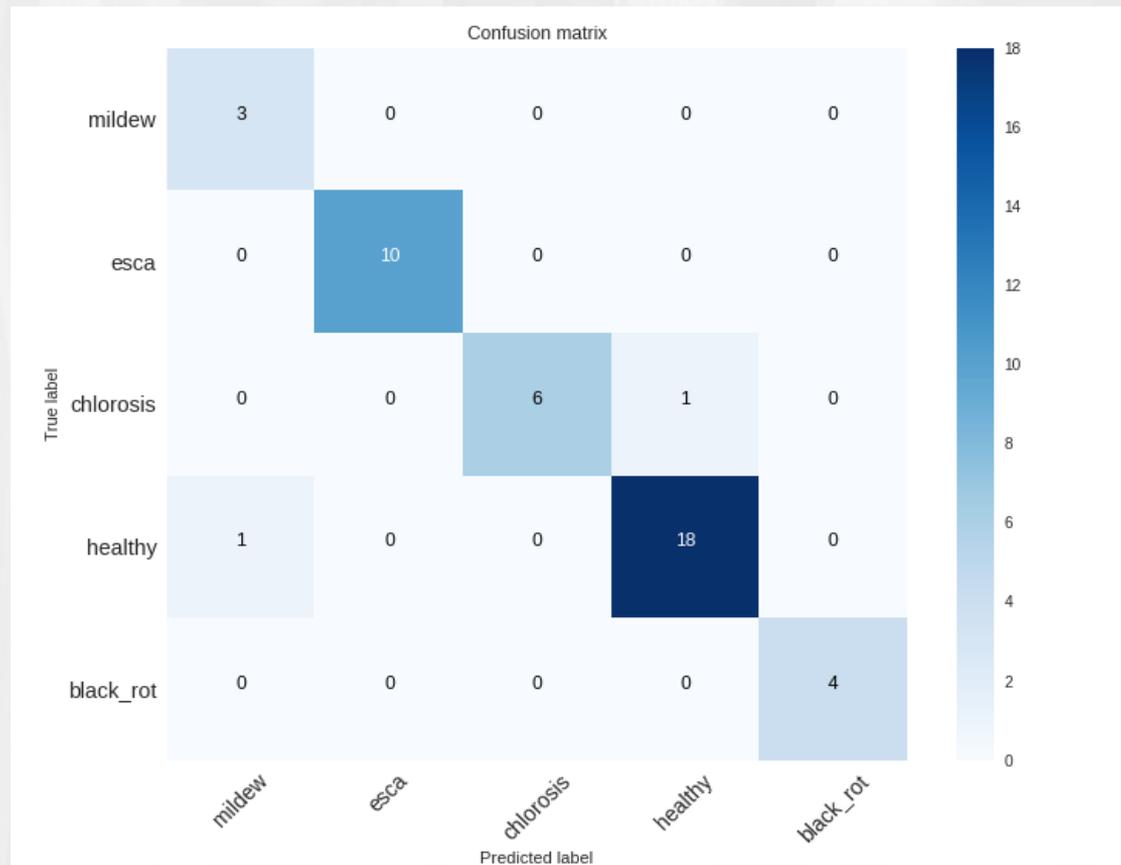
Step 3: Train and evaluate KNN with 1 neighbor



Instead of using classical KNN we formulate the problem of finding nearest neighbors as matrix multiplication of the two l2-normalized matrices of train and validation features. For details see <https://openreview.net/pdf?id=SJTQLdqIq>

Results: confusion matrix

This matrix shows how many class i objects were recognized as class j objects



Examples of the incorrect predictions

True – chlorosis



Predicted – healthy

True – healthy



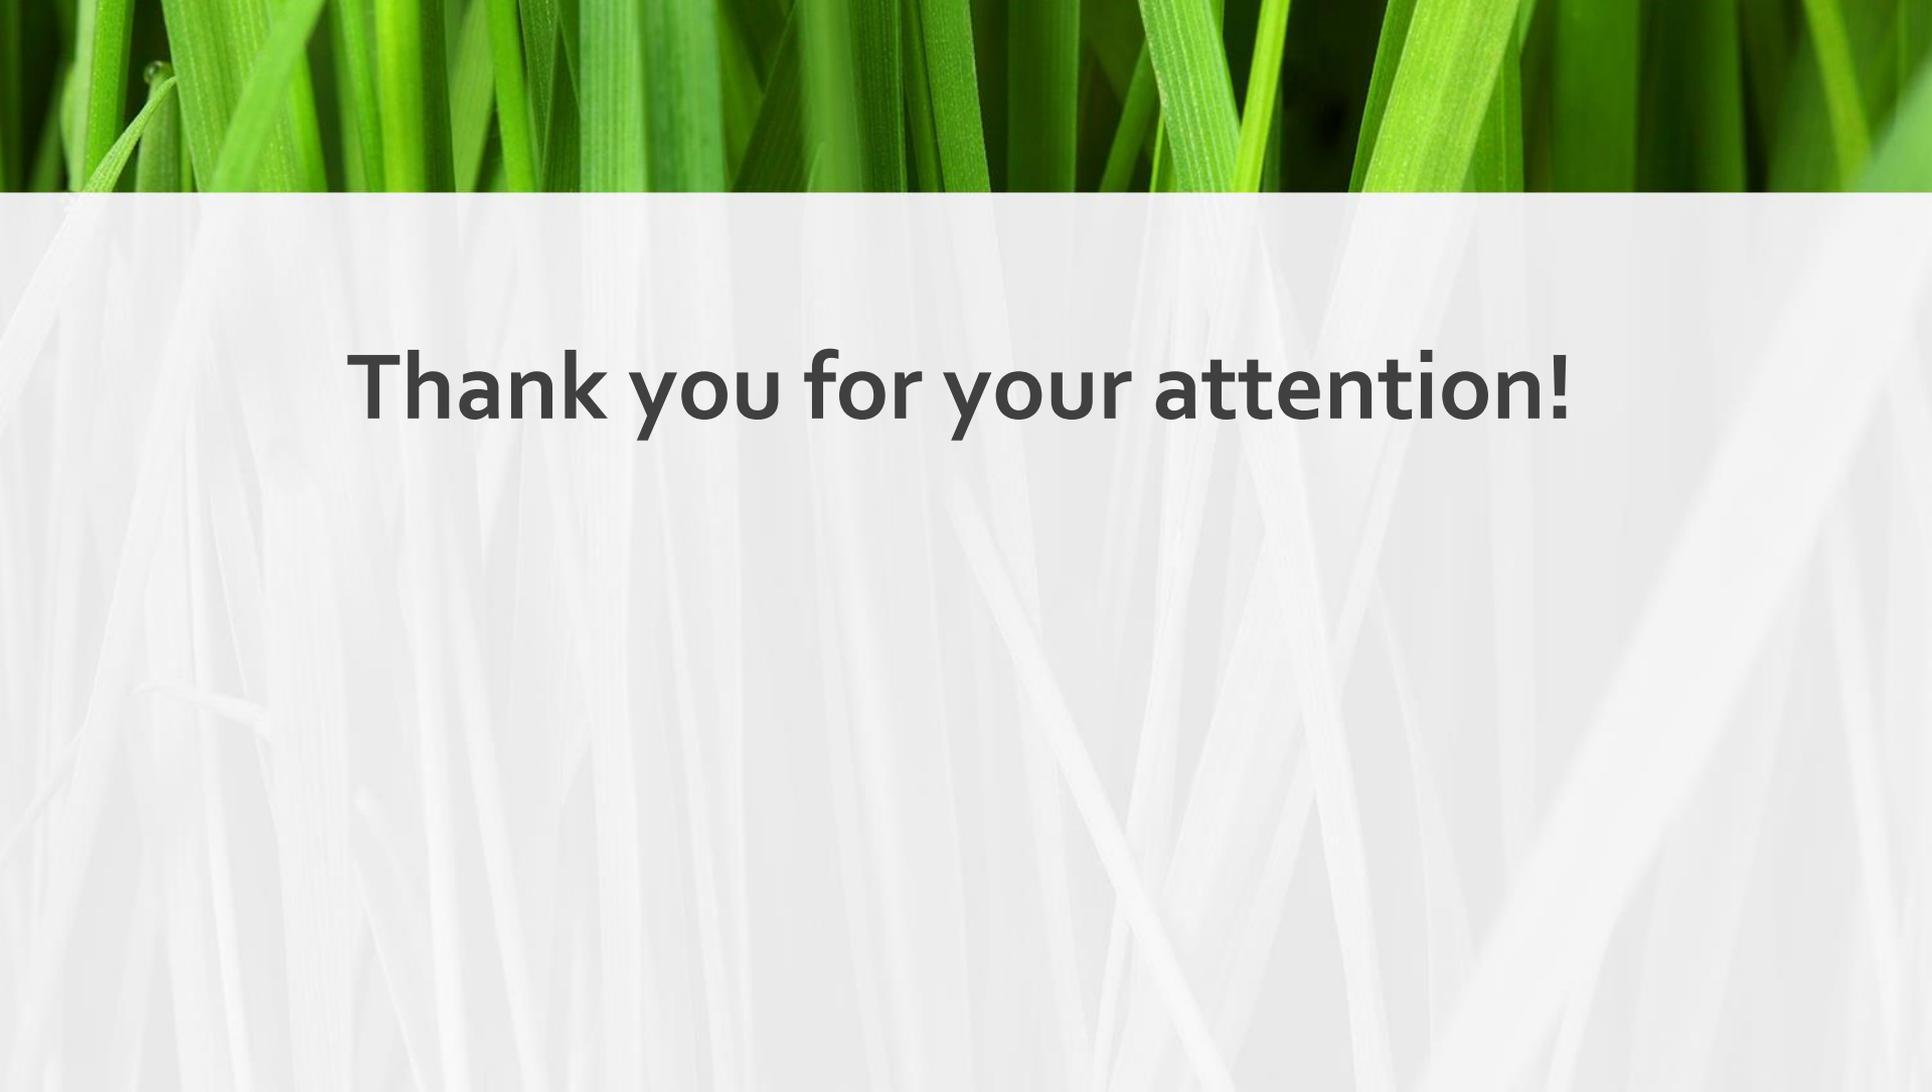
Predicted – mildew

Outlook and useful links

- Siamese neural networks are very perspective research field. We are going to use them as a basic deep learning architecture for PDDP.
- Since their power consists in seeking for differences between classes, we are going to add more classes to the train dataset soon.
- Also we are going to try a triplet loss function to train siamese network instead of binary cross-entropy.
- We will try to replace our siamese twins with the pretrained networks as part of transfer learning strategy.

Useful links:

- web site of the project: <http://pdd.jinr.ru/>
- link to the github repository: <https://github.com/Kaliostrogoblin/PDD>



Thank you for your attention!