Computers and Electronics in Agriculture 127 (2016) 418–424

Contents lists available at ScienceDirect

Computers and Electronics in Agriculture

journal homepage: www.elsevier.com/locate/compag

Original papers

Deep learning for plant identification using vein morphological patterns

Guillermo L. Grinblat, Lucas C. Uzal *, Mónica G. Larese, Pablo M. Granitto

CIFASIS, French Argentine International Center for Information and Systems Sciences, UNR-CONICET, Argentina

ARTICLE INFO

Article history:
Received 20 November 2015
Received in revised form 9 June 2016
Accepted 5 July 2016

Keywords:
Deep learning
Machine vision
Automatic plant identification
Leaf vein image

ABSTRACT

We propose using a deep convolutional neural network (CNN) for the problem of plant identification from leaf vein patterns. In particular, we consider classifying three different legume species: white bean, red bean and soybean. The introduction of a CNN avoids the use of handcrafted feature extractors as it is standard in state of the art pipeline. Furthermore, this deep learning approach significantly improves the accuracy of the referred pipeline. We also show that the reported accuracy is reached by increasing the model depth. Finally, by analyzing the resulting models with a simple visualization technique, we are able to unveil relevant vein patterns.

1. Introduction

Nowadays, in many typical applications of machine vision there is a tendency to replace classical techniques with deep learning algorithms (LeCun et al., 2015). In deep learning, handcrafted feature extractors are unnecessary: typically, classification results are better than those obtained with classical techniques. Some successful examples can be found in Krizhevsky et al. (2012), Cireşan et al. (2013), and Taigman et al. (2014).

Deep learning refers to training neural network architectures composed of several nonlinear processing layers. The success of deep learning is based on new model regularization techniques (Srivastava et al., 2014), improved nonlinearities design (Dahl et al., 2013), and current hardware capabilities, among others. In particular, for Machine Vision tasks, the success of deep learning is based on convolutional neural networks (CNN, LeCun et al., 1990) which have become the standard neural network variant for image processing (LeCun et al., 2015).

There are many agricultural problems currently addressed by classical machine vision techniques that may benefit from using a deep learning approach. We consider in this paper a successful example of this behavior by applying deep learning to automatic plant identification.

Automatic plant identification constitutes a challenging problem that has received increasing attention in recent years, in particular for identification based on leaf image analysis. Much of this work makes use leaf features that humans can perceive. The goal of automatization in this case is to avoid the use of human experts handling huge catalogs of plant species, and to reduce classification time. Some works are focused on leaf shape (Agarwal et al., 2006; Camargo Neto et al., 2006; Chaki and Parekh, 2012; Du et al., 2007; Gwo et al., 2013; Im et al., 1998; Solé-Casals et al., 2008), some use shape and texture (Husin et al., 2012), while others consider color and texture (Pydipati et al., 2006).

Recently, however, more attention has been payed to vein morphological patterns as a leaf fingerprint. A clear correlation has been established between vein characteristics and some properties of the leaf (such as damage and drought tolerance, among others) (Sack et al., 2008; Scoffoni et al., 2011). This suggests that vein morphology carries information suitable for plant classification when shape, color or texture differences are unobservable, as in the case of trying to separate different cultivars from the same species. This kind of features may not be easily spotted by a human observer, and automated recognition becomes indispensable.

Following this premise, Larese et al. (2014a) applied computer vision techniques to extract several vein morphological measures, and showed that it is possible to separate three different plant species by using only the extracted information and supervised machine learning algorithms. In a later work (Larese et al., 2014b), they used similar techniques to reach some degree of discrimination between plants belonging to different cultivars of the same species.

In this work we discuss the use of deep learning models for interesting agricultural problems. As a working example, we apply this new paradigm to the problem of plant identification based on vein morphology. We show that the application of a standard deep convolutional network yields better results than those obtained

* Corresponding author.
E-mail address: uzal@cifasis-conicet.gov.ar (L.C. Uzal).

http://dx.doi.org/10.1016/j.compag.2016.07.003
0168-1699/© 2016 Elsevier B.V. All rights reserved.

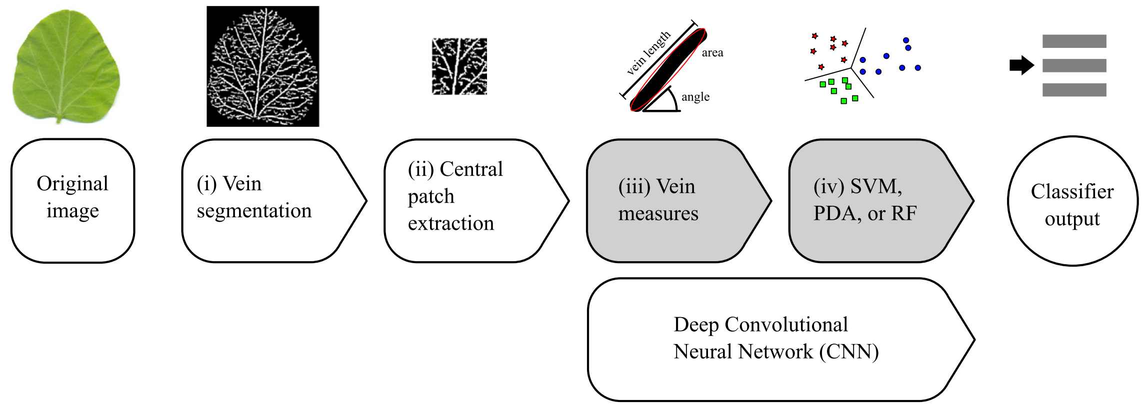

Fig. 1. Adopted pipeline (as in Larese et al. (2014b)). In this work, the grayed stages were replaced by a deep convolutional network. Stages (i) and (ii) were kept in order to allow a fair comparison with Larese et al. results. By design, these two stages filter color and leaf shape information. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

with a standard machine vision pipeline. Furthermore, the utilization of a simple model visualization technique allows us to identify meaningful vein patterns. The obtained results on plant classification from leaf vein morphology are not only valuable by themselves, but as a first step for motivating further research on the use of deep learning in agriculture.

The rest of the paper is organized as follows. In Section 2 we review the task-specific approach for the problem at hand as proposed by Larese et al. (2014b). In Section 3 we introduce the proposed deep learning methodology and explain the performed experiments in Section 4. Results are presented in Section 5, and then in Section 6 we show which patterns were deemed relevant for classification. Finally, we draw some conclusions in Section 8. In Appendix A, more detailed information about data acquisition and processing can be found.

2. Task-specific approach

Many successful Machine Learning applications make intensive use of specific knowledge about the task provided by human experts. In this section we summarize the approach considered in Larese et al. (2014b) for plant classification based on leaf veins, which makes use of expert knowledge.

The processing pipeline is divided into four stages as shown in Fig. 1. The starting point is the set of images of first foliage leaves acquired with a standard flatbed scanner (see Appendix A for more details). These images are processed according to the following stages:

(i) Vein Segmentation: first, an unconstrained version of the Hit or Miss Transform (UHMT) (Soille, 1999) is applied in order to extract vein morphological patterns. The output of this transform is a binary image—it therefore eliminates color information.

(ii) Central patch extraction: a central patch (100×100 pixels) of the binary image is cropped and the rest of the image is discarded. The purpose is to eliminate possible influences of the leaf shape.

(iii) Vein measures: at this stage, a set of features of interest was extracted with the help of LEAF GUI (Price et al., 2011). This set includes measures such as the total number of veins, total number of nodes, and mean vein width, among others.

(iv) Classification. Three different Machine Learning algorithms were tested: Support Vector Machines (SVM), Penalized Discriminant Analysis (PDA) and Random Forests (RF) (Hastie et al., 2009). These models were trained using the features obtained in the previous step.

There are two main observations about this pipeline. First, in order to highlight different levels of vein details, Larese et al. applied the UHMT to resized versions of the leaf image. The scale factors considered were 100% (no resize), 80%, and 60%. The processed images were resized back to the original size. With these three output images, two alternatives were studied. In the first one, a single combined image was obtained by adding them. For the second alternative, the three output images and the combined one were kept. We will refer to these two setups as S1 and S2 respectively. Fig. 2 shows some example images after stage (ii) for the S2 setup. The S1 setup correspond to selecting only the first column for each sample.

The second point to notice is that stage (iii) is the only one that requires specific domain knowledge. All considered measures can be automatically extracted but were specifically designed by experts to characterize vein patterns.

Also, it is important to remark the difference in the number of features extracted at stage (iii). Larese et al. extracted 52 features from each patch image. This means 52 features in the S1 setup, in contrast to 208 features in the S2 setup.

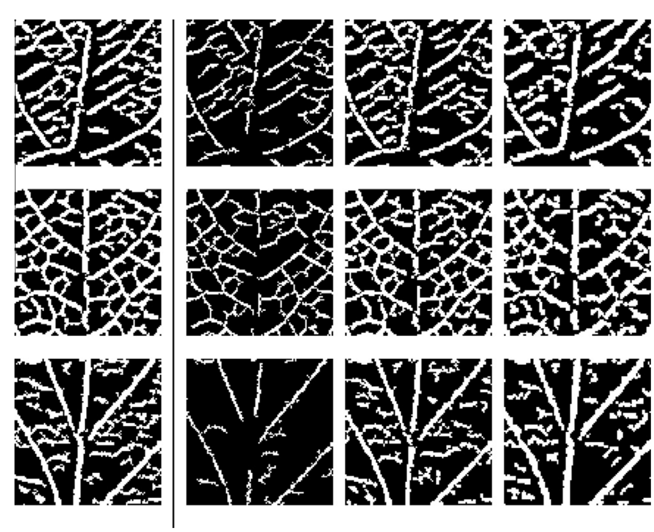

Fig. 2. Image samples obtained after processing stages (i) and (ii). The first column corresponds to preprocessing S1, while setup S2 is formed by all columns. These images are the input to the CNN.