

Planteome & BisQue: Automating Image Annotation with Ontologies using Deep-Learning Networks

Dimitrios Trigkakis, Sinisa Todorovic

Dept. of Electrical Engineering & Computer Science, Oregon State University, Corvallis, OR, USA

Justin Preece, Austin Meier, Justin Elser, Pankaj Jaiswal

Dept. of Botany & Plant Pathology, Oregon State University, Corvallis, OR, USA

Kris Kvilekval, Dmitry Fedorov, B.S. Manjunath

Department of Electrical and Computer Engineering, University of California, Santa Barbara, Santa Barbara, CA, USA

Keywords—*image analysis; segmentation; ontology; annotation; machine learning, deep learning, convolutional neural networks*

I. INTRODUCTION

The field of computer vision has recently experienced tremendous progress due to advances in deep learning [1]. This development holds particular promise in applications for plant research, due to a significant increase in the scale of image data harvesting and a strong field-driven interest in the automated processing of observable phenotypes and visible traits within agronomically important species.

Parallel developments have occurred in semantic computing; for example, new ontologies have been initiated to capture plant traits and disease indicators [2]. When these ontologies are combined with existing segmentation capabilities [3], it is possible to conceptualize software applications that give researchers the ability to analyze large quantities of plant phenotype image data, and to auto-annotate that data with meaningful, computable semantic terminology.

We have previously reported on a software application that integrates segmentation and ontologies, but lacked the ability to manage very high-resolution images, and also lacked a database platform to allow for high-volume storage requirements. We have also previously reported our migration of the AISO [4] user-guided segmentation feature to a BisQue (Bio-Image Semantic Query User Environment) [5] module to take advantage of its increased power, ability to scale, secure data management environment, and collaborative software ecosystem.

Neither AISO nor our initial BisQue implementation possessed a machine-learning component for interpreting (parts of) images. Plant researchers could benefit greatly from a trained classification model that predicts image annotations with a high degree of accuracy. We have therefore implemented two deep-learning prototypes: a coarse classification module for plant object identification (*i.e.* flower, fruit) and a fine-grained classification module that focuses on plant traits (*e.g.* reticulate vs. parallel venation, tip

shape). Both classification models return results mapped to ontology terms as a form of annotation enrichment. This current version of the Planteome Deep Segmenter module [6] combines image classification with optional guided segmentation and ontology annotation. We have most recently run the module on local Planteome BisQue client services, and are currently working with CyVerse [7] to install a hosted version on their BisQue client service.

II. RELEVANT WORK

The state-of-the-art models for classifying images typically use a hierarchy of filter convolutions for extracting discriminative image features from raw pixels and then map these deep image features to class predictions using a hierarchy of linear regressions. The former computing module is called convolutional layers, and the latter, fully connected layers. A number of convolutional layers followed by fully-connected layers constitutes a typical deep convolutional neural network (CNN). The literature presents a host of CNNs that are successful in image classification on big datasets, including AlexNet [8], ResNet [9], GoogLeNet [10], VGGNet [11] etc. These CNNs differ in the number of convolutional and fully-connected layers used, and their connectivity. Most of them are too complex and require a huge amount of training data, which is not suitable for our domain where expert annotations of images are usually scarce. In our work, we train ResNet [9] models for the task of image classification, for both the coarse classification of plant images and fine-grained leaf trait classification. We use segmentation with the dynamic graph-cuts algorithm [12] as an optional pre-processing step. While CNNs aimed directly at semantic segmentation of images have also been studied in the deep learning literature [13][14], they typically require a huge amount of pixel-wise ground-truth annotations, which are not available in our domain. Therefore, we have decided to use the pre-processing segmentation step for delineating important image parts, and then apply to the selected image segments our ResNet for coarse and fine classification.

Previous work for classification of leaves, like LeafSnap [15], attempted to identify specific species of leaves. In contrast with that work, our goal is to identify specific leaf traits that can be mapped to pre-existing ontology terms for any leaf species, even when the input image contains a leaf that has not been seen in the training dataset. We are not aware of any other work that attempts to classify leaf traits with ontological labelling in a given image.

III. IMPLEMENTATION

The latest iteration of our module on the BisQue platform preserves and extends the functionality of both AISO and our original BisQue segmentation efforts. The input consists of optional freeline user-guided markup on a selected image along with an interface for specifying module behavior. The markup defines the foreground and background of the image. The user then has the ability to perform exact object segmentation on any visual element of interest utilizing a Dynamic Graph Cuts algorithm. Our interface allows varying the segmentation quality in order to allow users to control the computational overhead of the segmentation.

The BisQue platform provides access to an online image library, along with annotations associated with graphical object on the images. The platform provides the ability to analyze segmentation results from previous module outputs in the form of overlaid graphical annotations in an image.

We augment the segmentation capabilities of the module with machine learning algorithms that can automatically extract semantic information from either the entire input image, or a segmented part of it. To this end we introduce two tasks: coarse plant classification and fine-grained leaf part classification. For the first task, we want to identify the class of a plant image, categorizing it to one of five categories, namely ‘leaf’, ‘fruit’, ‘flower’, ‘stem’ and ‘entire plant’. The segmentation and classification functionality of our module is shown in Figure 1(a-b).



Figure 1a. User guided markup on an image uploaded to the database. Markup includes foreground and background annotations for identifying the plant part of interest.




Figure 1b. Classification results over the segmented part of the image, accompanied by the ontology PO-term for the identified plant. The user can follow the provided link to the corresponding ontology for more information about the identified plant term.



Figure 2. The same image can be used for segmentation of different plant parts. The segmented part, identified as a leaf, can be further classified using the fine-grained leaf classification model.

For the second task, given an image of a leaf, we want to simultaneously classify multiple leaf characteristics, like leaf tip shape and leaf venation as shown in Fig. 3.

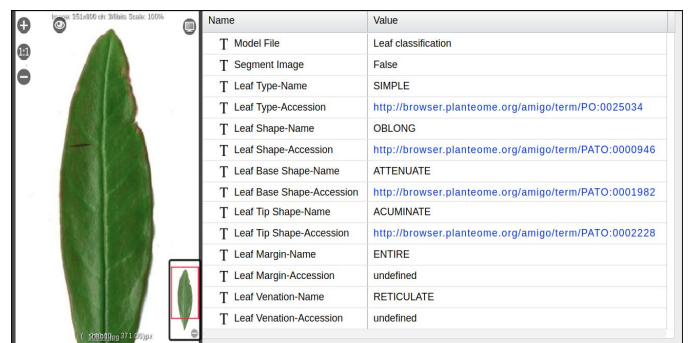


Figure 3. Fine-grained leaf classification. Each leaf characteristic is classified, and a link is provided for the corresponding ontology term.

The above functionality can be combined with the segmentation capabilities of our module, to produce classification of specific parts of the image. The user interface allows the segmentation and classification to be enabled separately or in combination. To enable better module portability, we isolate our deep learning framework within a virtual environment that uses PyTorch [16] and opencv [17] for all image operations.

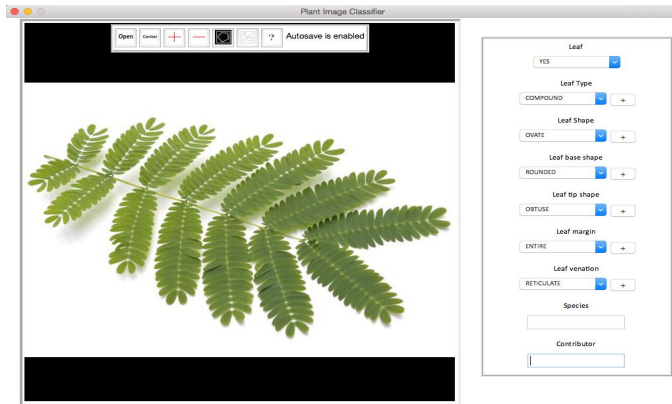


Figure 4. Desktop annotation tool, used for labelling fine-grained leaf characteristics. On the right, each classification field can be expanded so that we can find the corresponding field’s classification label. The toolbar helps with image manipulation in order to identify these classification labels.

For our task of supervised classification of images, we make use of the ImageClef 2013 plant identification dataset [18]. The dataset contains labels for the five classes in the aforementioned coarse plant classification task. We augment the dataset with a selection of images collected by the Jaiswal lab. For fine-grained leaf classification, we needed additional labels for each leaf characteristic to be classified. To that end, we developed a desktop annotator tool (Fig 4) [19], which enabled fast annotations of leaf characteristics based on reference images tagged with the appropriate terms. The annotation process using this tool resulted 549 fully annotated leaf images, selected from the subset of the aforementioned dataset when the leaves were presented clearly on a sheet (leaf-on-sheet pre-annotated category). Generalizing to natural images of leaves requires the use of the segmentation capabilities of our module to isolate a leaf in an image.

The deep learning models we employ are based on the ResNet architecture, a type of convolutional neural network. For the coarse classification task, we train ResNet-50 on the augmented ImageClef 2013 dataset, using the category labels provided in the dataset. For fine-grained classification, we find that ResNet-50 overfits the training set (i.e., cannot generalize well to new images that are not in the training set), so we reduce the number of model parameters by using ResNet-18. The model acts as a feature extractor, on top of which we

append parallel fully connected layers, one for each leaf trait. In this way we obtain multiple predictions for each image. The selection of the appropriate model is done through the interface, where a multiple choice menu allows the user to select either the coarse or the fine-grained classification pre-trained model to use in the module.

For each classification result provided by our models, a corresponding PO (Plant Ontology) or PATO (Phenotype And Trait Ontology) term is supplied that connects the class of an object with the correct ontology term. These terms can be accessed online through a supplied hyperlink, which maps all classes to their ontology through the Planteome database of ontologies.

IV. RESULTS

We report our classification results in Tables 1 and 2. In Table 1 we show the confusion matrix for the five categories. In a confusion matrix, the labels at the top row represent the predictions of the classifier, while the ground truth labels are represented in the left column. For example, the entry ‘21’ in the second row represents the fact that 21 testing samples were predicted as flower, but were actually of the fruit category. The diagonal entries represent predictions that were actually correct (the predicted and ground truth labels match). where correct predictions correspond to the diagonal entries. We obtain an accuracy of 91%. Our results for fine-grained classification of leaf traits are shown in Table 2. Each column represents a category, where each category contains a varying number of classes, so that multiple classification results have been grouped together in different classification tasks. In Table 2, we also compare with a baseline approach which represents random guessing over the classes of each category. The results reported were produced by k-fold cross validation, where 20% of the entire dataset of 549 leaves was used for the classifier’s testing accuracy, while the rest of the dataset was used for training the classifier. We repeat this process seven times on different, non-overlapping subsets of the dataset and average the testing set accuracy. The ‘Testing accuracy’ row contains these averaged testing accuracies after the aforementioned process has been completed.

	Fruit	Flower	Leaf	Stem	Entire
Fruit	1896	21	13	2	7
Flower	24	405	58	2	4
Leaf	46	69	1062	1	33
Stem	9	7	1	596	2
Entire	64	18	99	4	648

Table 1. Coarse classification task: Confusion matrix for the training set. Classification accuracy on the testing set is 91%. The table’s cells represent a sample from the testing set. That cell’s row label (leftmost column), shows the actual classification label. The cell’s column label (topmost row), shows the predicted classification label.

pave the way for future developments and more sophisticated learning models. Processing of high resolution digital images and management of high-throughput databases fulfill user needs in an era where more data becomes available as computational costs decrease dramatically. Future work will focus on online retraining of deep learning models, assembly and import of new models for new tasks, synchronous segmentation, more robust ontology metadata services, deeper semantic search integration with the BisQue platform, and public deployment and release of the software and source code (upon peer-reviewed publication; manuscript and hosted site are currently in preparation). We anticipate that further performance improvements could be made by expanding the initial training dataset of images with unannotated examples from the ImageClef dataset and then conducting semi-supervised training of our ResNet model for image classification.

ACKNOWLEDGMENTS

We acknowledge and thank the following curators for testing and using the desktop leaf annotation tool to generate annotated image data for model training purposes: Laurel Cooper, Noor Al-Bader, Austin Meier, and Parul Gupta. We would like to thank Philip Daly and Blake L. Joyce for their assistance and dedication in hosting the Planteome BisQue module on their infrastructure, via the CyVerse Extended Collaborative Support program. We thank the National Science Foundation for project support under the following awards: DBI-0735191, DBI-1265383 (CyVerse), 1340112 (Planteome, OSU), ABI-1356750 (BisQue, UCSB), and SI2-SSI 1650972 (LIMPID, UCSB).

	Leaf type	Leaf shape	Leaf base shape	Leaf tip shape	Leaf margin	Leaf venation
Testing Accuracy	98.2%	41.7%	37.1%	29.9%	51.2%	96.0%
Random Guessing	50%	3.1%	10%	5.5%	2.6%	50%

Table 2. Fine-grained leaf trait classification task: Classification accuracy on the testing set for each leaf trait. Unlike Table 1, we only show the actual accuracy percentages on the test set for the six classification fields.

V. DISCUSSION

User enhancements to our module interface have included differential markup line coloring (as shown in Fig. 2), free-line markup input, classification model selection and segmentation options. Our module not only provides segmentation of image parts that can later be annotated with ontology terms, but also the mapping of such segments to a predefined set of classes that already contain the corresponding ontology terms, providing users with the information related to the ontology of the selected segment. The ontology terms link directly to the planteome ontology database to provide all relevant information that the ontology term is associated with. Our module also allows the seamless integration of new deep models through the same interface, for other tasks related to plant classification.

VI. CONCLUSIONS

We have developed an effective suite of guided segmentation and auto-annotation tools that can be used for online image analysis. Integration of deep learning models with segmentation and ontology annotation functionality have provided a strong foundation for plant image analysis, and

REFERENCES

- [1] X. W. Chen and X. Lin, "Big Data Deep Learning: Challenges and Perspectives," in *IEEE Access*, vol. 2, pp. 514-525, 2014. doi: 10.1109/ACCESS.2014.2325029
- [2] Planteome by <http://planteome.org> is licensed under a Creative Commons Attribution 4.0 International License. Based on a work at planteome.org.
- [3] Vignesh Jagadeesh, Utkarsh Gaur, "Graph Cuts-based Image Matting / Segmentation", unpublished. Online BisQue module: http://bisque.iplantcollaborative.org/module_service/ImageMatting/
- [4] Lingtla N*, Preece J*, Todorovic S, Cooper L, Moore L, Jaiswal P. 2014. AISO: Annotation of Image Segments with Ontologies. *Journal of Biomedical Semantics*. 5:50. (*: Co-authors)
- [5] Kristian Kvilekval, Dmitry Fedorov, Boguslaw Obara, Ambuj Singh and B.S. Manjunath, "Bisque: A Platform for Bioimage Analysis and Management", *Bioinformatics*, vol. 26, no. 4, pp. 544-552, Feb. 2010.
- [6] Planteome Deep Segmenter software repository: <https://github.com/Planteome/planteome-deep-segmenter>
- [7] Merchant, Nirav, et al., "The iPlant Collaborative: Cyberinfrastructure for Enabling Data to Discovery for the Life Sciences," *PLOS Biology* (2016), doi: 10.1371/journal.pbio.1002342.
- [8] Imagenet classification with deep convolutional neural networks A Krizhevsky, I Sutskever, GE Hinton *Advances in neural information processing systems*, 1097-1105

- [9] Deep Residual Learning for Image Recognition, K He, X Zhang, S Ren, J Sun Computer Vision and Pattern Recognition (CVPR), 2016
- [10] Going Deeper with Convolutions, Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, Andrew Rabinovich (CVPR), 2015
- [11] Karen Simonyan, Andrew Zisserman: Very Deep Convolutional Networks for Large-Scale Image Recognition (ICLR), 2015
- [12] Boykov Y., Kolmogorov V. (2001) An Experimental Comparison of Min-cut/Max-flow Algorithms for Energy Minimization in Vision. In: Figueiredo M., Zerubia J., Jain A.K. (eds) Energy Minimization Methods in Computer Vision and Pattern Recognition. EMMCVPR 2001. Lecture Notes in Computer Science, vol 2134. Springer, Berlin, Heidelberg
- [13] Fully Convolutional Models for Semantic Segmentation Jonathan Long*, Evan Shelhamer*, Trevor Darrell CVPR 2015 arXiv:1411.4038
- [14] "DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs" Liang-Chieh Chen*, George Papandreou*, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille (*equal contribution) arXiv preprint, 2016
- [15] Kumar N. et al. (2012) Leafsnap: A Computer Vision System for Automatic Plant Species Identification. In: Fitzgibbon A., Lazebnik S., Perona P., Sato Y., Schmid C. (eds) Computer Vision – ECCV 2012. Lecture Notes in Computer Science, vol 7573. Springer, Berlin, Heidelberg
- [16] <https://pytorch.org/>, PyTorch framework for deep learning
- [17] <https://opencv.org/>, OpenCV library for computer vision
- [18] <http://www.imageclef.org/2013/plant>, ImageClef 2013 plant identification challenge
- [19] Plant Image Desktop Classifier software repository:
<https://github.com/Planteome/plant-image-desktop-classifier>