Learning to Correct for Bad Camera Settings in Large Scale Plant Monitoring

Zongyang Li[1], Abby Stylianou[2] and Robert Pless[2]
Donald Danforth Plant Science Center[1], George Washington University[2]
ZLi@danforthcenter.org, {astylianou|pless}@gwu.edu

Abstract: In large scale, automated image capture systems, incorrect camera settings can lead to images that are completely useless or which break assumptions made by image analysis algorithms, but there may also be sufficient data to learn to automatically correct bad data. We consider the specific problem of over- and under-saturated images in large scale plant growth monitoring, propose a generative approach that addresses these saturation issues for the calculation of plant canopy cover, and suggest future areas of research in this problem domain.

Problem Domain: In many cases, properly calibrated color images allow for simple heuristics to give accurate labels for field images. For example, a pixel-based classifier of "more green or more red" may be used to classify pixels as plant or ground. However, in realistic imaging environments that include direct sunlight, cloudy days and harsh shadows, parts of the image are too dark for accurate color assessment, or so bright that the pixel values are saturated and the color information is lost. While there are many approaches for color correction, they don't apply when data is lost due to saturation or because not enough light was captured at all. In these cases additional information may be necessary to do the correction or for accurate pixel labelling. One source of additional information comes when images are captured as part of a large scale image capture project, where there may be related images captured with *good* color calibration. One case like this is the TERRA-REF field trial [2], which captures millimeter resolution photos of 1600 plots daily over many growing seasons.

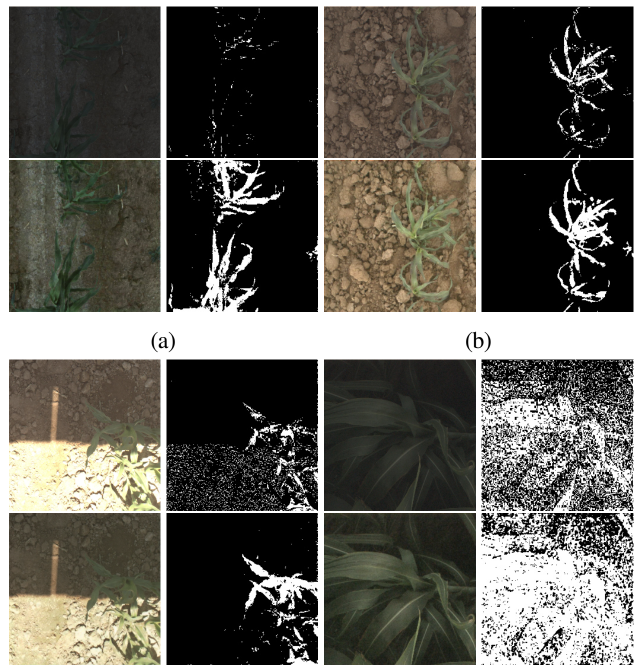

(a) (b) (c) (d)

Figure 1: Each subfigure shows the original image over the color-corrected image, and the masks derived from each. (a-c) give examples where the mask becomes better, and (d) shows a case where the mask does not.

Approach: We use the color images from the TERRA dataset to train an automatic, data driven, task specific color correction approach. We start by identifying 383,124 images on days with image capture setups that were in good conditions. We simulate bad conditions by randomly changing the image brightness (+/- 80%) and image contrast (+/- 20%). The original images are treated as a target, and a conditional GAN [1] is trained to try to generate those targets from corrupted images.

Preliminary Results and Next Steps: We apply the trained GAN to unseen images taken in bad conditions and compute canopy masks using this corrected data, with largely positive results detailed in Figure 1. Next steps include directly estimate the plant mask or to exploring richer models of simulating bad data.

References

[1] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *IEEE conference on computer vision and pattern recognition (CVPR)*, pages 1125–1134, 2017.

[2] Maxwell Burnette et al. TERRA-REF Data Processing Infrastructure. In *Proceedings of the Practice and Experience on Advanced Research Computing*, PEARC '18, pages 27:1–27:7. ACM, 2018.