
Reinforcement Learning for Sustainable Agriculture

Jonathan Binas¹ Leonie Luginbuehl² Yoshua Bengio¹

1. Summary

The growing population and the changing climate will push modern agriculture to its limits in an increasing number of regions on earth. Establishing next-generation sustainable food supply systems will mean producing more food on less arable land, while keeping the environmental impact to a minimum.

Modern machine learning methods have achieved super-human performance on a variety of tasks, simply learning from the outcomes of their actions. We propose a path towards more sustainable agriculture, considering plant development an optimization problem with respect to certain parameters, such as yield and environmental impact, which can be optimized in an automated way. Specifically, we propose to use reinforcement learning to autonomously explore and learn ways of influencing the development of certain types of plants, controlling environmental parameters, such as irrigation or nutrient supply, and receiving sensory feedback, such as camera images, humidity, and moisture measurements. The trained system will thus be able to provide instructions for optimal treatment of a local population of plants, based on non-invasive measurements, such as imaging.

Unlike genetic methods for plant development optimization, which typically target one or at most a few aspects of the development process, machine learning methods can take into account a large number of high-dimensional variables. We expect both greenhouse and outdoor production environments to benefit considerably from such fine-grained control systems. In particular, a learned model may be used to achieve yield maximization under challenging environmental conditions, such as limited space, extreme temperature, or water supply limitations.

2. Background

The upcoming agricultural challenges demand innovation both at the genetic level (that is, which genetic factors determine crop productivity and robustness to environmental stress,) as well as the ecological level (that is, which condi-

tions lead to environments for optimal plant development.) While genetic methods typically only target one particular aspect of plant development, a combination of ecological factors can influence development as a whole. While both approaches are complementary, we focus on the latter. Successful and sustainable ways of improving yield under limited environmental resources will impact agriculture, and thus humanity, on a global scale.

As an example scenario, consider the case of a prolonged regional drought, where irrigation becomes necessary to keep the land arable. The amount of water that can be supplied might be limited, and a fixed amount might need to be smartly distributed over the full life cycle of an organism. Similarly, the amount of fertilizer applied throughout the life cycle of a plant should be kept to a minimum to protect the environment. It is not clear, however, what the best timing for application would be: Is it best to apply a high concentration at germination, should small amounts be applied continuously, or should application rates increase with maturity of the plant? While best-practice knowledge naturally exists among farmers, it is not clear how or whether specific protocols transfer well to new scenarios, such as droughts, extreme temperature or precipitation, shifted seasons, etc. In general, it is not clear whether specific protocols are optimal in that they provide the best possible conditions for a plant, or whether more sophisticated application schedules could lead to better results. Moreover, in modern agriculture, any treatment is typically applied to the whole field, based on some average assessment. However, it is likely that different parts of a population, or even individual organisms, would benefit from individualized treatment, such as provision of specific nutrients, based on their actual individual condition, rather than the average state of the culture. We intend to tackle these points using modern machine learning methods.

Optimizing agriculture with machine learning. Deep Learning methods deliver better-than-human performance in various tasks, such as image analysis, game-playing, and control (LeCun et al., 2015; Mnih et al., 2015). It is to be expected that similarly, in agriculture, such methods can be used as decision support systems and smart controllers, which rival human performance. The power of machine learning systems comes with their ability to learn from real-world data. Rather than carrying out hard-coded behavior, such algorithms learn by themselves, through large num-

¹Mila, Montreal ²University of Cambridge, UK. Correspondence to: Jonathan Binas <jbinas@gmail.com>. CCAI workshop at the 36th International Conference on Machine Learning, Long Beach. Copyright 2019 by the authors.

bers of experiments, to optimize a particular outcome. In particular, we propose here to use reinforcement learning methods to learn to control plant development in such a way that a given quantity, say yield, is maximized given certain environmental and resource constraints. For instance, this could mean dividing a fixed amount of water into individual allotments over the organisms lifespan. It could mean to maximize yield while minimizing fertilizer administration, given certain environmental parameters, such as temperature, moisture, and the chemical composition of the soil. Unlike traditional control methods, deep learning systems are able to make sense of high-dimensional sensory input, such as image data. As a result, a trained system can infer the plant's constitution from visual data and other non-invasive measurements with high accuracy, and adjust the current administration schedule accordingly, even on the level of individual plants.

Previous applications of machine learning in agriculture include the creation of detailed computer models and simulations of plant organisms (Marshall-Colon et al., 2017), yield prediction through fruit counting (Ramos et al., 2017) or satellite imaging (Pantazi et al., 2016a), disease detection (Chung et al., 2016) or parasite discovery (Ebrahimi et al., 2017) based on visual data, and weed detection and characterization (Pantazi et al., 2016b). To our knowledge, however, machine learning techniques have not been directly applied to the control of essential agricultural supply channels, such as irrigation, fertilization, and provision of other nutrients.

3. Reinforcement learning in biological environments

We propose an approach involving both a physical and a modeling component, where an agent learns to control a number of parameters affecting plant development through reinforcement learning (Sutton et al., 1998). Rather than interacting with a virtual environment, the agent controls an array (minibatch) of growth chambers, equipped with sensors and actuators to enable precise control of the environmental conditions of the organisms at hand. Besides a




Figure 1. Illustration of the approach. After training in a controlled environment, the learned model can be used to provide optimal treatment recommendations in the field.

control unit and a camera, a growth chamber should contain sensing equipment for temperature, humidity, moisture, and more specialized chemical sensors, as well as heating and cooling elements, tunable light sources, controllable water inlets, and nutrient supply units. The agent model is provided with sensor readings as observations and the learned policy emits actions controlling the actuators of the growth chamber. A reward signal is specified based on the particular objective considered, and could involve quantities, such as yield (volume, weight) after a given period, growth rate, fertilizer usage, or water requirements. Given the lengthy nature of actual plant development trials and the relatively high sample complexity of many reinforcement learning algorithms, it appears sensible to pre-train the model on a simulated organism (Marshall-Colon et al., 2017) before deploying it to the actual physical environment.

As two key scenarios, we propose to apply reinforcement learning to fertilizer management and water management, where the algorithm learns the best possible distribution of a resource (nutrients, water) over time, given a set of environmental and plant conditions. The automated system will, through many parallel experiments, learn at which points in time it is best to provide certain amounts of the resource to the plant, such production is maximized, while minimizing overall usage of the resource. The sensory data might be augmented with virtual data, such as a noisy forecast of future conditions, similar to what a weather forecast would provide. The highly detailed plant-level data (or at least data from a local region within a field) should enable the algorithm to make better decisions regarding optimal treatment of individual organisms or local populations than what could be achieved on a per-field level. The ability of deep learning methods to take into account large amounts of high-dimensional sensory data should enable more precise predictions than what could be achieved with a small number of coarse measurements, such as temperature and precipitation.

We envision to operate several controlled growth chambers in parallel to gather enough data for the algorithms to learn. Initial experiments could be based on *Arabidopsis thaliana* as a model species, whose small size and rapid life cycle make it an ideal candidate for affordable, high-turnover experiments, and properties of which transfer reasonably well to other species (Rensink & Buell, 2004). Initial nutrient supply experiments could focus on nitrogen and phosphate. The minimization of nitrogen fertilizers is of utter importance to environmental protection. We propose to scale the approach to monocots, such as rice, maize, or wheat, once the experimental setup is stable. These species require more space and more time to grow, however, they are more directly relevant to questions of food security, and corresponding model organisms, such as *Brachypodium distachyon* and *Setaria viridis* are sufficiently easy to grow.

References

- Chung, C.-L., Huang, K.-J., Chen, S.-Y., Lai, M.-H., Chen, Y.-C., and Kuo, Y.-F. Detecting bakanae disease in rice seedlings by machine vision. *Computers and electronics in agriculture*, 121:404–411, 2016.
- Ebrahimi, M., Khoshtaghaza, M., Minaei, S., and Jamshidi, B. Vision-based pest detection based on svm classification method. *Computers and Electronics in Agriculture*, 137:52–58, 2017.
- LeCun, Y., Bengio, Y., and Hinton, G. Deep learning. *nature*, 521(7553):436, 2015.
- Marshall-Colon, A., Long, S. P., Allen, D. K., Allen, G., Beard, D. A., Benes, B., Von Caemmerer, S., Christensen, A., Cox, D. J., Hart, J. C., et al. Crops in silico: generating virtual crops using an integrative and multi-scale modeling platform. *Frontiers in plant science*, 8:786, 2017.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Pantazi, X. E., Moshou, D., Alexandridis, T., Whetton, R., and Mouazen, A. M. Wheat yield prediction using machine learning and advanced sensing techniques. *Computers and Electronics in Agriculture*, 121:57–65, 2016a.
- Pantazi, X.-E., Moshou, D., and Bravo, C. Active learning system for weed species recognition based on hyperspectral sensing. *Biosystems Engineering*, 146:193–202, 2016b.
- Ramos, P., Prieto, F. A., Montoya, E., and Oliveros, C. E. Automatic fruit count on coffee branches using computer vision. *Computers and Electronics in Agriculture*, 137: 9–22, 2017.
- Rensink, W. and Buell, C. R. Arabidopsis to rice. applying knowledge from a weed to enhance our understanding of a crop species. *Plant Physiology*, 135(2):622–629, 2004.
- Sutton, R. S., Barto, A. G., et al. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.