

A Concept of Full Plant Morphology Modeling for Robot-Plant Bio-Hybrids

Mostafa Wahby, Julian Petzold and Heiko Hamann

Institute of Computer Engineering
University of Lübeck, Germany
wahby@iti.uni-luebeck.de

Abstract

Robot-plant bio-hybrid systems are getting increasing attention due to the wide range of applications they offer. Such synergies between robots and natural plants will allow, for example, establishing highly reliable environmental monitoring systems or growing the architecture of our future cities. We explore the latter application where robots exploit the plants' ability to produce construction material, and plants exploit the robots' sensing and computational capabilities. In our previous work, we used machine learning techniques to model plant behavior in their early life stages. We collected a 10-point plant stem description dataset and used it to train an LSTM as a forward model that predicts plant dynamics and drives the evolution of plant shaping controllers. Here, we show our vision to model plant behaviors in later stages, where full-plant morphology will be used to train state-of-the-art sequence modeling networks capable of simulating more complex plant dynamics.

Motivation

In recent years, the concept of bio-hybrid systems, ranging from remote-controlled insects (Sato et al., 2009) to mixed societies of fish and robots (Schmickl et al., 2013), has been actively researched. Often, these systems combine a robotic component with animals, but there are also systems that combine robots with natural plants (Hamann et al., 2015). Robot-plant bio-hybrids have many practical use cases, such as on-demand customizable construction (Wahby et al., 2018a). The combination of hardware and living plants allows the use of such systems as adaptive, self-repairing objects or buildings (Heinrich et al., 2019).

The robotic part of such a bio-hybrid system often provides a control system for the biological part, which serves as a building material or substrate (Hamann et al., 2015). Due to their slow growth, time scales for interactions with plants are often larger than with animals. Therefore, the natural building material must be modeled. This way, a robotic controller can be rapidly prototyped or trained without the need for time consuming plant experiments. A bio-hybrid manufacturing system's practical applications will likely use multiple plants, but a single-plant model is less complicated to generate and can be generalized to multiple plants later

on. Some Alife related approaches, e.g., L-Systems (Lindenmayer, 1975; von Mammen et al., 2011), generate structures using production rules in a formal language model.

In our previous work, we approached modeling the growth dynamics of common bean plants for plant shaping. The common bean is fast-growing and shows a pronounced directional growth towards blue light due to phototropism (Christie and Murphy, 2013). In an initial approach (Wahby et al., 2015; Hofstadler et al., 2017; Wahby et al., 2016), we mathematically modeled the motion dynamics of only the plant tip. Later, in an approach to generate a more representative model, we included 10 points along the plant stem (Wahby et al., 2018b). This was sufficient for small-scale plant shaping experiments (up to 20 cm high), but it does not scale up to include later plant growth stages and complex morphology, e.g., branching, stiffening, etc.

Plant growth is a sequential process that changes the plant morphology. Therefore, a sequence modeling and prediction approach is necessary to encode and simulate the temporal plant dynamics. This was successfully achieved by (Wahby et al., 2018b), by training a long-short term memory (LSTM) network (Gers et al., 2000). LSTMs are limited in capacity, but they tend to forget outdated information; that's why they were suitable for our earlier small-scale experiments. However, to extend this work, a transformer (Vaswani et al., 2017)—a state-of-the-art sequence modeling approach—or a recurrent highway network (Zilly et al., 2017) might be more suitable because they provide higher capacity and handle higher amounts of data in parallel. Instead of the 10-point plant description, we would like to explore using such approaches provided with full plant geometry for larger scale plant behavior prediction. In a preliminary approach we present a trained Variational Autoencoder (VAE, Kingma and Welling, 2014) that is capable of encoding the full plant geometry.

Approach

In our previous work, we combined methods from machine learning and evolutionary algorithms to control a plant shaping robot. The robot observes and records plant growth pat-

terns in reaction to light stimuli. Recorded images are processed to build a 10-point xy stem geometry description dataset, see Figure 1(a). This data is used to train an LSTM in a supervised way to simulate plant stem stiffening, growth, and motion under any given sequence of light stimuli. The LSTM is used as a holistic plant forward model to evolve controllers in simulation, for the task of steering and shaping plants to reach desired targets by exploiting stem stiffening phenomena. The fittest controllers were successfully run in reality, with a physical setup and actual plants.

Here we present our new concept: we would like to extend the model to describe more complex plant geometries (e.g., taller plants with branches). We are required to encode the full plant morphology. We present a preliminary approach, where we trained a VAE based on the recorded plant growth patterns from our previous work. The VAE network learns an abstract, encoded description of each observed input plant frame and can then be used to train a more capable, holistic plant model and to encode the observed plant frames at each time step during experiment execution.


Apart from modeling the full plant morphology, our newly proposed setup introduces more autonomy and adaptivity. As plants grow, the robot can collect new information (i.e., records more plant frames). Therefore, the VAE network can be continuously trained and enhanced on-the-fly. The holistic model will adapt to newly observed dynamics due to, for example, growth, stiffening, or branching.

Results

We resize a subset (12,500 frames) of our recorded dataset to 64×64 pixels and normalize the pixel values between zero and one; then, we used them to train a convolutional VAE (ConvVAE). Inspired by Ha and Schmidhuber (2018), we configured the ConvVAE to have a $64 \times 64 \times 3$ input tensor, four convolutional layers, and four deconvolutional layers. The convolutional layers can encode and downsample a frame into a low dimension latent vector of size 64. A latent vector can be decoded and reconstructed by passing through the four deconvolutional layers. The convolutional and deconvolutional layers use a stride of two pixels (indicating that a filter convolves around the input volume by shifting two units in each dimension). All layers use a Rectified Linear Activation (ReLU) function except the output deconvolutional layer, which uses a sigmoid function to output values between zero and one. We trained the ConvVAE for 50 epochs with a mean square distance reconstruction loss between an input dataset frame and the reconstructed frame. On batches of size 100, we obtain a total loss of value 5.68 at the last epoch. In Figure 1(b), we show a sample reconstruction of an encoded plant frame using our trained ConvVAE.¹

The ConvVAE serves as a high-dimensional generative model of the plant, providing a smooth latent vector that we

use to generate model instances. This latent vector allows us to identify and eliminate certain unwanted features of our input data, such as reflections of the robotic controller’s lighting or the plant’s pot (see Figure 1). In future work, we will also investigate whether we can use the latent vector to enhance certain image features (e.g., plant material) to provide a better plant controller.



(a) 10-point xy description (Wahby et al., 2018b) (b) A reconstructed frame using our trained ConvVAE.

Figure 1: A comparison between robot controllers inputs. In (a), we input a 10-point xy plant stem geometry, but in (b), we input an encoded representation of the whole frame.

Conclusion

In pursuit of bringing bio-hybrid construction systems to deployment outside controlled laboratory environments, we have taken steps towards modeling plants’ full morphology. Using a VAE, we encode the full plant information using a 64×64 pixel RGB-image. In previous work (Wahby et al., 2018b), the plant representation was restricted to a fixed topology of 10 connected stem points, which limited our model to the top 20 cm of the plant’s stem. The VAE is able to encode plants of arbitrary size and complexity.

In future work, we have to examine the tradeoff between the computational power and data we use to train our VAE and the amount of details our model reflects. For larger plants, we have to retain a higher resolution in the output encoding, as small but relevant features might otherwise be lost. If this is the case, our LSTM used for sequence modeling and growth prediction may not be suitable anymore. Therefore, we will examine compressive transformers (Rae et al., 2020), which could both encode the plant and model its growth dynamics, replacing both LSTM and VAE.

The extended plant morphology may provide additional feedback to a robotic plant controller, enabling methods that are able to process larger amounts of data in parallel, such as recurrent highway networks (Zilly et al., 2017). In turn, this opens up tasks that were previously impossible, such as complex plant interactions and manufacturing bifurcating 3D-objects using plants.

¹Find a video at: <https://vimeo.com/523547383>

References

- Christie, J. M. and Murphy, A. S. (2013). Shoot phototropism in higher plants: new light through old concepts. *American Journal of Botany*, 100(1):35–46.
- Gers, F. A., Schmidhuber, J. A., and Cummins, F. A. (2000). Learning to forget: Continual prediction with lstm. *Neural Computation*, 12(10):2451–2471.
- Ha, D. and Schmidhuber, J. (2018). Recurrent world models facilitate policy evolution. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NIPS 2018)*, pages 2455–2467, USA. Curran Associates Inc.
- Hamann, H. et al. (2015). *flora robotica* – mixed societies of symbiotic robot-plant bio-hybrids. In *Proceedings of IEEE Symposium on Computational Intelligence*. IEEE.
- Heinrich, M. K., von Mammen, S., Hofstadler, D. N., Wahby, M., Zahadat, P., Skrzypczak, T., Soorati, M. D., Krela, R., Kwiatkowski, W., Schmickl, T., et al. (2019). Constructing living buildings: a review of relevant technologies for a novel application of biohybrid robotics. *Journal of the Royal Society Interface*, 16(156):20190238.
- Hofstadler, D. N., Wahby, M., Heinrich, M. K., Hamann, H., Zahadat, P., Ayres, P., and Schmickl, T. (2017). Evolved control of natural plants: Crossing the reality gap for user-defined steering of growth and motion. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, 12(3):15:1–15:24.
- Kingma, D. P. and Welling, M. (2014). Auto-encoding variational bayes. In *Proceedings of the 2nd International Conference on Learning Representations (ICLR 2014)*.
- Lindenmayer, A. (1975). Developmental algorithms for multicellular organisms: a survey of L-systems. *Journal of Theoretical Biology*, 54(1):3–22.
- Rae, J. W., Potapenko, A., Jayakumar, S. M., Hillier, C., and Lillcrap, T. P. (2020). Compressive transformers for long-range sequence modelling. In *International Conference on Learning Representations*.
- Sato, H., Berry, C. W., Peeri, Y., Baghoomian, E., Casey, B. E., Lavella, G., VandenBrooks, J. M., Harrison, J., and Maharbiz, M. M. (2009). Remote radio control of insect flight. *Frontiers in Integrative Neuroscience*, 3:24.
- Schmickl, T., Bogdan, S., Correia, L., Kernbach, S., Mondada, F., Bodi, M., Gribovskiy, A., Hahshold, S., Miklic, D., Szopek, M., et al. (2013). Assisi: mixing animals with robots in a hybrid society. In *Conference on Biomimetic and Biohybrid Systems*, pages 441–443. Springer.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- von Mammen, S., Wißmeier, T., Wong, J., and Jacob, C. (2011). Artistic Exploration of the Worlds of Digital Developmental Swarms. *Leonardo*, 44(1):5–13.
- Wahby, M., Divband Soorati, M., von Mammen, S., and Hamann, H. (2015). Evolution of controllers for robot-plant bio-hybrids: a simple case study using a model of plant growth and motion. In *Proceedings of the 25th Workshop on Computational Intelligence*, pages 67–86. KIT Scientific Publishing.
- Wahby, M., Heinrich, M. K., Hofstadler, D. N., Neufeld, E., Kuksin, I., Zahadat, P., Schmickl, T., Ayres, P., and Hamann, H. (2018a). Autonomously shaping natural climbing plants: a bio-hybrid approach. *Royal Society Open Science*, 5(10):180296.
- Wahby, M., Heinrich, M. K., Hofstadler, D. N., Zahadat, P., Risi, S., Ayres, P., Schmickl, T., and Hamann, H. (2018b). A robot to shape your natural plant: the machine learning approach to model and control bio-hybrid systems. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2018)*, pages 165–172. ACM.
- Wahby, M., Hofstadler, D. N., Heinrich, M. K., Zahadat, P., and Hamann, H. (2016). An evolutionary robotics approach to the control of plant growth and motion: modeling plants and crossing the reality gap. In *Proceedings of the 10th International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2016)*, pages 21–30. IEEE.
- Zilly, J. G., Srivastava, R. K., Koutník, J., and Schmidhuber, J. (2017). Recurrent highway networks. In *International Conference on Machine Learning*, pages 4189–4198. PMLR.