

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/324600101>

# A Robot to Shape your Natural Plant: The Machine Learning Approach to Model and Control Bio-Hybrid Systems

Article · July 2018

DOI: 10.1145/3205455.3205516

CITATIONS

0

READS

78

8 authors, including:



**Mostafa Wahby**

Universität zu Lübeck

12 PUBLICATIONS 69 CITATIONS

[SEE PROFILE](#)



**Mary Katherine Heinrich**

Universität zu Lübeck

12 PUBLICATIONS 27 CITATIONS

[SEE PROFILE](#)



**Daniel Hofstadler**

Karl-Franzens-Universität Graz

13 PUBLICATIONS 74 CITATIONS

[SEE PROFILE](#)



**Payam Zahadat**

Karl-Franzens-Universität Graz

49 PUBLICATIONS 208 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



Evolutionary Computation [View project](#)



subCULTron [View project](#)

# A Robot to Shape your Natural Plant: The Machine Learning Approach to Model and Control Bio-Hybrid Systems

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

Mary Katherine Heinrich  
Centre for IT and Architecture  
Royal Danish Academy  
Copenhagen, Denmark

Daniel Nicolas Hofstadler  
Artificial Life Lab  
University of Graz  
Graz, Austria

Payam Zahadat  
Artificial Life Lab  
University of Graz, Austria

Sebastian Risi  
Robotics, Evolution and Art Lab  
ITU Copenhagen, Denmark

Phil Ayres  
Centre for IT and Architecture  
Royal Danish Academy, Denmark

Thomas Schmickl  
Artificial Life Lab  
University of Graz, Austria

Heiko Hamann  
Institute of Computer Engineering  
University of Lübeck, Germany

## ABSTRACT

Bio-hybrid systems—close couplings of natural organisms with technology—are high potential and still underexplored. In existing work, robots have mostly influenced group behaviors of animals. We explore the possibilities of mixing robots with natural plants, merging useful attributes. Significant synergies arise by combining the plants' ability to efficiently produce shaped material and the robots' ability to extend sensing and decision-making behaviors. However, programming robots to control plant motion and shape requires good knowledge of complex plant behaviors. Therefore, we use machine learning to create a holistic plant model and evolve robot controllers. As a benchmark task we choose obstacle avoidance. We use computer vision to construct a model of plant stem stiffening and motion dynamics by training an LSTM network. The LSTM network acts as a forward model predicting change in the plant, driving the evolution of neural network robot controllers. The evolved controllers augment the plants' natural light-finding and tissue-stiffening behaviors to avoid obstacles and grow desired shapes. We successfully verify the robot controllers and bio-hybrid behavior in reality, with a physical setup and actual plants.

## CCS CONCEPTS

• **Theory of computation** → **Evolutionary algorithms**; • **Computing methodologies** → **Neural networks**; • **Computer systems organization** → **Evolutionary robotics**;

## KEYWORDS

bio-hybrid system, natural plants, decision making, evolutionary robotics, genetic algorithms, neural networks

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*GECCO '18, July 15–19, 2018, Kyoto, Japan*

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5618-3/18/07...\$15.00

<https://doi.org/10.1145/3205455.3205516>

## ACM Reference Format:

Mostafa Wahby, Mary Katherine Heinrich, Daniel Nicolas Hofstadler, Payam Zahadat, Sebastian Risi, Phil Ayres, Thomas Schmickl, and Heiko Hamann. 2018. A Robot to Shape your Natural Plant: The Machine Learning Approach to Model and Control Bio-Hybrid Systems. In *GECCO '18: Genetic and Evolutionary Computation Conference, July 15–19, 2018, Kyoto, Japan*, Jennifer B. Sartor, Theo D'Hondt, and Wolfgang De Meuter (Eds.). ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3205455.3205516>

## 1 INTRODUCTION

Recent developments in additive manufacturing (3D printing) and robotics open up techniques to produce objects of increasing size and variety, such as mugs, chairs, or even houses. Research on complex systems and evolvable hardware could interpret this production process as a growth process, such that printing an object like a house could be adaptive to unanticipated changes in the structure or environment. As an objective of the project *flora robotica* [11, 12] we investigate methods to conduct additive manufacturing with bio-hybrids—that is, mixed societies of robotic and biological systems. Our objective is to use natural plants to grow desired shapes by controlling them with robotic devices. From the perspective of developmental systems, we replace artificial substrates with a natural system, both in terms of models and physical implementation. We can then exploit features of natural plants, such as adaptive behavior and the (almost free) addition of material by growth.

We expect challenges due to the real-life complexity of biological systems, and their combination with evolutionary robotics to automatically generate appropriate robot controllers. A downside of natural growth is its slow speed, which requires simulation of the growth process for genetic algorithms to be applied. Another challenge is that holistic models of plant growth are not readily available, so we develop our own task-specific models. In summary, we realize a truly interdisciplinary approach with a rather complex tool chain, using evolution and machine learning to control plant *phototropism*—the directional behavior of motion and irreversible growth towards light.

First, instead of relying on a designed plant model, our stem stiffening and motion model is learned from experiment data recording plant behavior in the presence of certain light stimuli patterns. We

aim to capture the complex temporal dynamics of plant stiffening and motion through a particular class of recurrent neural networks called Long Short-Term Memory (LSTM) [13]. The hypothesis is that this approach will allow the model to capture the dynamics of a particular plant to a high enough degree to serve as a forward model that can guide the evolutionary search. Given this plant model we then apply methods of evolutionary robotics to evolve in simulation controllers of dynamic light stimuli for the given task of obstacle avoidance. Finally, we address the challenge of the reality gap by showing transferability of the simulated evolved controllers back to the real world.

Our focus is on setting up this rather complex toolchain, so the complexity of the task is relaxed in this early stage of research in the field. The task is to grow a plant collision-free around an obstacle and towards a target (bio-hybrid obstacle avoidance). Even this simple task brings added complication to obstacle avoidance, as lower parts of the plant still need to avoid collisions later in the run, and we cannot only focus on control of the plant tip.

In our evolutionary robotics approach we evolve artificial neural networks (ANN), which may seem at first glance an overly complex tool for this supposedly simple control task. However, we want to evolve controllers that are adaptive to the environment and to configurable tasks. Additionally, ANN is one toolchain approach that enables scaling up to more complex plant-control problems (e.g., 3D shapes, multiple stimuli) in the future. The workflow of our approach starts from preliminary plant experiments to gather data about how the plant behaves in general. The data is used to train an LSTM network that we use as simulator in our evolutionary runs. We evolve ANNs as controllers of light stimuli, which later in our reality experiments control the behavior of the real plant.

## 2 BACKGROUND AND RELATED WORK

Forming bio-hybrid societies by bringing biological and artificial agents together is a growing field. Robots can interact with natural organisms, both adapting their own behaviors and influencing those of the living system. Several mixed societies have been built where autonomous robots influence the behavior of groups of animals [10, 25]. While animals are very mobile, plants are more limited in motion, growing and adapting over time. In our previous work, we show that robots can closely interact with plants to change their environmental stimuli according to desires of humans [14, 22, 23]. As robot controllers and hardware can be designed to interact with their surroundings, they can meaningfully be combined with plants, extending their natural capabilities to grow efficiently in dynamic environments and adapt to external changes [8].

*Common bean plant: behavior and growth.* A relatively fast growing plant, the common bean (*Phaseolus vulgaris* L.) grows 3 cm/day on average [4]. Like many plants, common beans grow toward (blue) light [7] through the *phototropism* behavior, in constant balance with other competing growth behaviors. Beans in particular dramatically display *circumnutation*, a winding behavior for attachment and climbing [4]. During plant growth, new cells replicate at the tips, and older tissues gradually stiffen until they reach their final size and maturity. Incoming light adds a directional bias to winding, but only when this bias persists will the impact be irreversible and manifested as permanent curvature in the stem.

*Modeling plants.* Many models exist in plant science literature, with focus on particular aspects of plants and complex details of the biological system (e.g., [2]). Plant growth has also been a source of inspiration for several abstract models in computer science and artificial life. A prominent example is L-systems [16], where formal languages are interpreted to generate structures through a set of production rules. Branching mechanisms in plants inspire generative models adaptive to dynamics in the environment [26], and are abstractly modeled using polygon meshes of trees [27]. In this paper we develop a model of plant growth through a class of recurrent neural networks called LSTM [13] (see below), based on experimental data gathered from real plants.

*Long Short-Term Memory.* LSTMs [13] are a special class of recurrent neural networks that have been shown to effectively learn sequential patterns in a variety of domains [9, 21]. As plant growth is essentially a sequence of changes in plant tissue morphology, the LSTM can directly be applied. In LSTMs the recurrent layers normally found in recurrent ANNs are replaced by purpose-built memory cells, with content controlled by different gate types (input, forget, and output). The outputs of an LSTM memory block are sent back to block inputs and gates through recurrent connections. For a more detailed description of LSTMs see [13]. LSTMs have shown promise in plant classification by taking into account plant growth over time [18], but to the best of our knowledge they have not yet been applied to learn a plant growth forward model.

*Evolutionary Computation.* Evolutionary approaches have been applied to many areas of robotics, including design of robot controllers [3]. The approach we utilize, NEAT (NeuroEvolution of Augmenting Topologies) [20], is an evolutionary algorithm that evolves ANNs incrementally from simple initial networks while preserving the diversity of the evolutionary population. Evolution of robot controllers can follow an embodied approach [24], meaning that the controller is evolved directly in the real hardware. Another approach is to evolve controllers in simulation based on relevant models, and then transfer the evolved controller to the real hardware. While the former approach can be time-consuming and costly, the latter can suffer from the reality gap problem [15], meaning that the evolved controller performs poorly on the real hardware due to unknown limitations of the models.

## 3 METHODS

Our machine learning approach to shaping natural plants follows the methods below. First, preliminary data collection experiments in the bio-hybrid setup record plant growth patterns in reaction to light stimuli. Recorded images are processed to build a stem shape dataset. This data is used to train an LSTM in a supervised way, to simulate plant stem stiffening and motion under any given sequence of the light stimuli. Finally, the LSTM network is used as a forward model to evolve controllers in simulation, for the task of steering and shaping a plant, to avoid hitting obstacles and reach desired targets by exploiting stem stiffening phenomena.

### 3.1 Bio-hybrid setup

Following our approach in [14, 22, 23], the bio-hybrid setup is enclosed in a commercial grow box of dimensions  $120 \times 120 \times 200$  cm<sup>3</sup> in width, depth and height. The grow box interior is lined

in matte black board for a consistent background, diminishing light reflections. Freshly germinated beans ('Saxa' variety<sup>1</sup>) are placed in commercial soil in 1.5l-pots, aligned with the grow box back midpoint. The centralized robotic element consists of the following: two NeoPixel LED strips, a Raspberry Pi camera module, an LED light-bulb, and a Raspberry Pi. A NeoPixel strip contains 144 RGB LEDs, with peak-emission at wavelengths  $\lambda_{\max}$  630, 530, and 475 nm respectively. Each LED can emit up to 18 lumens at full power, consuming 0.24 W. The NeoPixel LED strips are coiled into cylindrical shapes and fastened to the grow box back wall, 30 cm above the soil and 35 cm to either side. The camera module faces the plant at a height of 32 cm and distance of 74 cm. The LED light-bulb is used as a flash when photographing, at 80 cm above the ground and centered over the pot. The Raspberry Pi runs background processes<sup>2</sup> to administrate plant experiments, including synchronizing flashes, capturing photos, extracting plant stem data, running ANNs, controlling light sources, and uploading data to a Network-attached storage device (NAS).<sup>3</sup>

## 3.2 Model setup

**3.2.1 Dataset experiments.** Our plant model is derived from our previous dataset experiments with real plants, in a bio-hybrid setup. These include six repetitions with a simplistic, non-reactive controller [22], and three repetitions with a closed-loop adaptive controller [14]. The open-loop controller switches light sources in regular six hour intervals, and the closed-loop controller switches according to plant tip position. In each experiment the plant is photographed every five minutes. The plants show influence by both growth and motion, with substantial motion horizontally. They also show variance between the behaviors of individual plants. Observing variance in plant experiments is a well-known phenomenon in plant science, which requires high numbers of repetitions. However, in the context of this research, where the focus is on evolutionary computation and robotics, such high overheads for experiments are infeasible. Instead we test our approach based on an engineering perspective by testing whether the model, that results from these experiments, helps us to successfully control a plant. We also test if the evolved controllers are able to perform properly with such dynamic and unexpected plant behavior.

**3.2.2 Stem motion tracking.** We describe our computer vision method for stem motion tracking. The 10-point description of stem geometry forms the basis of training data for our LSTM-based *Stem stiffening and motion model*. We process images from the dataset experiments described above, to record a 10-point  $xy$  description of the full stem at each timestep, representing its phototropic motion and stiffening dynamics. 10 points is sufficient to capture curvature details within the growth area of the current setup. The images are sampled<sup>4</sup> at 1/8 resolution and processed according to the following method, both for the dataset experiments described above, and for the reality gap experiments detailed in Sec. 4.2. Before processing the dataset experiment images, a set of images is compiled showing

the setup without a plant. The setup images include all states of the controller and any slight variations in lighting conditions. The set of images is sampled,<sup>4</sup> isolating the green RGB channel value at each pixel position  $(i, j)$  and remapping it onto the domain  $[0, 1]$ , forming sequence  $\Lambda$  containing a matrix  $M$  of green values for each image. To represent the interval of possible green channel values present in the setup, matrices  $L$  and  $H$  are constructed by

$$L_{i,j} = \min_{M \in \Lambda} (M_{i,j}), \quad H_{i,j} = \max_{M \in \Lambda} (M_{i,j}). \quad (1)$$

After constructing the setup matrices, dataset experiment images of plants are processed. The green channel value is isolated for each pixel  $(i, j)$ , remapped to the domain  $[0, 1]$ , and saved into the matrix  $R$ . Pixels within a window are identified as containing plant material if:  $(R_{i,j} < L_{i,j} - \theta_1) \vee (R_{i,j} > H_{i,j} + \theta_1)$ , for threshold  $\theta_1 = 0.2$ . Each identified plant pixel is extracted to set  $P$ , and their  $(x_p, y_p)$  coordinate positions are used to identify two possible locations of the plant's tip. In order to locate the growth tip  $g = (x^g, y^g)$ , plant pixels are compared to the globally defined anchor  $a = (x^a, y^a)$ , representing the position where the plant stem emerges from the soil. Two possible  $xy$  growth tip positions are identified (corner point  $c$  as the furthest from  $a$  in  $xy$ , and high point  $h$  as the same in  $y$  only) and one selected as  $g_n$  based on Euclidean distance ( $d$ ) from  $g$  in the prior timestep, such that

$$c = \arg \max_{(x_p, y_p) \in P} |x^a - x_p| + |y^a - y_p|, \quad (2)$$

$$h = \arg \max_{(x_p, y_p) \in P} |y^a - y_p|, \quad (3)$$

$$g_n = \begin{cases} h, & \text{if } d(g_{n-1}, h) < d(g_{n-1}, c) \\ c, & \text{otherwise} \end{cases}. \quad (4)$$

The remaining intermediate points  $((x^{S_2}, y^{S_2}), \dots, (x^{S_9}, y^{S_9}))$  describing the stem are preliminarily identified from set  $P$ , and then smoothed while preserving the representation of stiffening dynamics. For these eight points, the  $y^i$  are distributed evenly between the tip and anchor, as

$$y^{S_i} = \frac{i}{9} |y^a - y^g| + y^a, \quad (5)$$

and  $x^{S_i}$  are set as the averaged  $x_p$  for pixels in set  $P$  that have  $y_p$  within threshold  $\theta_2$  of the respective  $y^{S_i}$ , such that

$$x^{S_i} = \overline{x_p} : \forall x_p \in P : |y_p - y^{S_i}| < \theta_2, \quad (6)$$

where  $\theta_2 = 30$  pixels. In this way, a 10-point description  $S$  of the full stem  $S = (a, (x^{S_2}, y^{S_2}), \dots, (x^{S_9}, y^{S_9}), g)$  is defined. Due to minor variations in images caused by shadows and light reflections on the stem, this 10-point detection contains some errors. We address these errors using a simple algorithm based on Smoothing via Iterative Averaging (SIA) [17], which preserves the key topological features of the curve being smoothed. For each point in  $S_n$ , our algorithm utilizes the equation

$$(x^{S_i}, y^{S_i}) = \left( \frac{1}{2}(x^{i-1} + x^{i+1}), \frac{1}{2}(y^{i-1} + y^{i+1}) \right) \quad (7)$$

iteratively, according to the following steps: 1) for  $i \in \{2, 4, 6, 8\}$ , apply eq. 7, 2) for  $i \in \{3, 5, 7, 9\}$ , apply eq. 7, 3) for  $i \in \{2, 4, 6, 8\}$ , apply eq. 7. In this way, the intermediate stem points are smoothed with the SIA-based process, while the tip and anchor remain unchanged. The newly smoothed sequences  $S$  are converted to cm

<sup>1</sup>See our previous work [14] for all product specifications in this section.

<sup>2</sup>Managed by Systemd, system and service manager for Linux operating systems.

<sup>3</sup>The ZeroMQ (<http://zeromq.org/>) library is used for communication among these.

<sup>4</sup>Sampling was duplicated in two platforms: Python, utilizing the OpenCV library; and IronPython using Grasshopper libraries pertaining to computer vision.

and scaled to match physical setup dimensions. The anchors are then unified to standardize the data, by translating  $(x^{S_i}, y^{S_i})$  points in  $S_n$  according to  $((x^{S_i}, y^{S_i}) + (a - \bar{a} : \forall a \in S_n))$ . The resulting data is reformatted to sequence  $\Psi$  of 18-dimensional vectors  $\psi_j = (x_j^{S_2}, y_j^{S_2}, \dots, x_j^{S_9}, y_j^{S_9}, x_j^g, y_j^g)$ , without the now redundant anchor values. These vectors are the basis for regression data for our LSTM-based *Stem stiffening and motion model*, described below.

**3.2.3 LSTM trained as Stem stiffening and motion model.** Building a holistic model of plant stem dynamics is a complex task [2] that would benefit from deep learning. However there is a lack of existing data, and the substantial overhead associated with plant experiments makes it infeasible to obtain large amounts of new data (many plants can be grown in parallel but controlled light conditions, monitoring, and tracking are costly). Therefore, having obtained a small amount of data from real plant experiments—described above—we develop a method to artificially expand that data, avoiding overfitting when training the LSTM.

*Preparation of stem data for regression.* After manually removing data in  $xy$ -areas that are too sparsely populated to provide reliable data (mostly in zones far from the origin, where only one plant of nine may have reached by coincidence), we process the experiment motion tracking data in two ways to expand the set. Firstly we add noise, to reduce the tendency of overfitting, and secondly we add a generic model, such that the typical plant behavior is dominantly represented in the data for the LSTM.

In order to add normal distribution noise, in addition to the experiment data in sequence  $\Psi$ , we define noisy data in sequences  $(\Psi_1, \dots, \Psi_n)$ , where  $\psi_{n_j} = ((x_j^{S_2})^{\Psi_n}, (y_j^{S_2})^{\Psi_n}, \dots, (x_j^{S_9})^{\Psi_n}, (y_j^{S_9})^{\Psi_n}, (x_j^g)^{\Psi_n}, (y_j^g)^{\Psi_n})$ . The noise values applied to each growth tip  $(x_j^g, y_j^g)$  in  $\Psi$  are computed according to the mean  $\mu$  and standard deviation  $\sigma$  of a finite quantity ( $\theta_3$ ) of the closest growth tips  $(x_i^g, y_i^g)$  that have the same light condition  $b$ . To calculate this, for each growth tip  $(x_j^g, y_j^g)$ , all other growth tips in  $\Psi$  sharing the same light condition  $b$  are placed into sequence  $\mathbf{dist}^j$ , and are then sorted according to their Euclidean distance from the respective tip at  $m = j$ . The  $\theta_3$  closest tips for each respective tip are defined as  $W_\Psi$ , such that

$$w_j = (x_i^g, y_i^g) \in \Psi \quad | \quad n \leq \theta_3 \in \mathbf{dist}_n^j, \quad (8)$$

where  $\theta_3 = 100$ . The mean  $\mu$  for noise is calculated as

$$\mu(x_j^g) = \frac{1}{|W_\Psi|} \sum_{j=1}^{|W_\Psi|} w_j x_i^g, \quad (9)$$

with a symmetrical equation for  $\mu(y_j^g)$ , and standard deviation  $\sigma$  as

$$\sigma^2(x_j^g) = \frac{1}{|W_\Psi|} \sum_{j=1}^{|W_\Psi|} w_j \left( x_i^g - \mu(x_j^g) \right)^2, \quad (10)$$

with a symmetrical equation for  $\sigma^2(y_j^g)$ . The noisy data in each new sequence  $\Psi_n$  is calculated by first defining the noisy growth tips and then defining the noisy intermediate points in relation to the tip output. The new noisy tips  $((x_j^g)^{\Psi_n}, (y_j^g)^{\Psi_n})$  in  $\Psi_n$  are calculated using normal distribution noise, according to the existing

tips  $(x_j^g, y_j^g)$ , and  $\mu$  and  $\sigma$  values scaled by factor  $\omega$  such that

$$(x_j^g)^{\Psi_n} = x_j^g + \mathcal{N}\left(\mu(x_j^g), \sigma(x_j^g)\omega\right), \quad (11)$$

with symmetrical equation for  $(y_j^g)^{\Psi_n}$ , where scaling factor  $\omega = 0.1$ . The noisy intermediate points  $((x_j^{S_i})^{\Psi_n}, (y_j^{S_i})^{\Psi_n})$  in  $\Psi_n$ , are calculated according to the noisy tips  $((x_j^g)^{\Psi_n}, (y_j^g)^{\Psi_n})$ , and are scaled by the values  $\omega_2(x^{S_i}), \omega_2(y^{S_i})$ . The noise values are generated through an artificial mean  $\mu_2$  and standard deviation  $\sigma_2$ , defined according to the calculated standard deviation and the generated change in position of the noisy growth tips, such that

$$\mu_2\left((x_j^{S_i})^{\Psi_n}\right) = x_j^{S_i} + \left((x_j^g)^{\Psi_n} - x_j^g\right)\omega_2(x^{S_i}), \quad (12)$$

$$\sigma_2\left((x_j^{S_i})^{\Psi_n}\right) = \left(\sigma(x_j^g)\omega\right)\omega_2(x^{S_i}), \quad (13)$$

with symmetrical equations for  $\mu_2\left((y_j^{S_i})^{\Psi_n}\right), \sigma_2\left((y_j^{S_i})^{\Psi_n}\right)$ , where scaling factors  $\omega_2(x^{S_i}), \omega_2(y^{S_i})$  are defined according to the extents in  $\Psi$  of the respective intermediate point  $(x^{S_i}, y^{S_i})$  in comparison to the extents of growth tip  $(x^g, y^g)$ . These are defined as

$$\omega_2(x^{S_i}) = \left| \min_{x_i^{S_i} \in \Psi} (x_i^{S_i}) - \max_{x_i^{S_i} \in \Psi} (x_i^{S_i}) \right| \cdot \left| \min_{x_i^g \in \Psi} (x_i^g) - \max_{x_i^g \in \Psi} (x_i^g) \right|^{-1}, \quad (14)$$

with a symmetrical equation for  $\omega_2(y^{S_i})$ . In new noisy data  $\Psi_n$ , the intermediate points  $((x_j^{S_i})^{\Psi_n}, (y_j^{S_i})^{\Psi_n})$  are defined using normal distribution noise, according to  $\mu_2$  and  $\sigma_2$  and scaled by  $\omega$ , such that

$$(x_j^{S_i})^{\Psi_n} = x_j^{S_i} + \mathcal{N}\left(\mu_2((x_j^{S_i})^{\Psi_n}), \sigma_2((x_j^{S_i})^{\Psi_n})\omega\right), \quad (15)$$

with a symmetrical equation for  $(y_j^{S_i})^{\Psi_n}$ . In the methods implementation described in this paper, we conduct three runs of equations 11–15 and their respective symmetries, generating three unique sequences of noisy data:  $\Psi_1, \Psi_2, \Psi_3$ .

In order to add a generic plant model, we manually select experiment data associated with the natural plants' smoothest and least noisy movements (identified by observation), and then follow a data-driven approach. We reinforce these generic movements as dominant by adding additional translations of them, distributed over small  $xy$  distances. In addition to experiment data  $\Psi$  and noisy data  $\Psi_n$ , we define new generic model data in sequence  $\Psi_\Phi$ , with each vector  $\psi_{\Phi_j}$  structured as those in  $\Psi$ , defined as:

$$(x_j^{\Psi_\Phi}, y_j^{\Psi_\Phi}) \in \psi_{\Phi_j} = \left( (x_j \in \psi_j) \pm 10^\lambda, (y_j \in \psi_j) \pm 10^\lambda \right), \quad (16)$$

where  $\lambda = (-3, \dots, -6)$ , generating 64 new  $xy$  translations in  $\Psi_\Phi$ .

The data sequences  $\Psi, \Psi_n$ , and  $\Psi_\Phi$  are combined to form  $\Psi^*$ . Vectors in  $\Psi^*$  are then mirrored across the  $x$ -axis, as we assume the targeted plant behavior to lack left-right bias. This also doubles the quantity of data. Then  $\Psi^*$  is reformatted according to timestep, such that each new vector  $\psi_j^*$  is composed of 18 dimensions representing the current  $xy$  stem position, 18 dimensions representing the next stem position, and one dimension representing the Boolean light condition. Vectors are removed if they 1) contain duplicate stems at the current and next positions, or 2) if the  $xy$  change is greater than  $20\times$  the average  $xy$  in  $\Psi^*$ . We end up with a  $\Psi^*$  dataset containing 101,162 vectors.




Figure 1: LSTM-based model training.

*Training procedure.* In order to obtain a holistic plant model, we train the LSTM using Keras [6], a high-level wrapper of TensorFlow [1]. The data in  $\Psi^*$  is formatted as described above, in vectors containing nine 2D stem points at a given time step and at the subsequent timestep, together with current light conditions (left/right light on). The LSTM network has 19 input units (current nine points and light condition), 50 LSTM memory blocks, and 18 output units (next nine points). We shuffle the vectors  $\psi_j^*$  and split them into training (70%), validation (20%), and testing set (10%). We train the LSTM network with the training set in sequence of batches of size  $N = 30$  for 200 epochs at a steady learning rate of 0.001 using Adam optimizer [1]. The training loss  $L_t$  is the mean absolute error (MAE), defined as

$$L_t = \frac{1}{N} \sum_{i=1}^N \frac{1}{18} \sum_{j=2}^{10} |x_p^j - x_t^j| + |y_p^j - y_t^j|, \quad (17)$$

where  $x_t^j$  and  $y_t^j$  are the true  $xy$  coordinate values of the stem point  $j$ , and  $x_p^j$  and  $y_p^j$  are the corresponding predicted coordinates. The validation dataset is used to track the training progress through validation loss  $L_v$  (calculated similarly to  $L_t$  but not in batches). An early stopping callback is implemented to prevent overfitting by tracking  $L_v$  and stop the training process with patience of ten epochs (i.e., if the  $L_v$  stops improving for ten epochs). As seen in Fig. 1, the training process stops at the 27th epoch when  $L_v$  stopped improving for ten epochs at  $L_t = 1.56 \times 10^{-3}$  and  $L_v = 1.55 \times 10^{-3}$ . Then, we calculate the MAE for each of the three datasets when used as input to the network. The error values for the training, validation and testing datasets are  $1.55 \times 10^{-3}$ ,  $1.55 \times 10^{-3}$ , and  $1.44 \times 10^{-3}$  respectively. On average, the error is  $\approx 1$  mm at each coordinate value, showing that the resulting model represents plant behavior closely<sup>5</sup>.

### 3.3 Controller setup

Our controller is an ANN operating two light sources. The input to the ANN at each time step is 1) a set of 10-points (20 real numbers) representing the current plant position and shape, as described above, 2) the current coordinates of the target (2 real numbers), and 3) coordinates of the obstacle (4 real numbers). We have two setups: in silico (simulation) and in vivo ('wet' setup with plant and hardware). In silico, the 10 points are directly generated using the *stem stiffening and motion model*. In vivo, a camera and computer

vision detects the actual plant and forms the corresponding 10 points. The output of the ANN is the control triggering light sources for stimuli.

*3.3.1 Task: Obstacle avoidance by shaping the plant.* The controller has to shape the plant appropriately by navigating it around a virtual obstacle to then reach a target area (radius is 2 cm). The plant should not touch the obstacle with any part of its body throughout the experiment. Since the obstacle is virtual, it neither casts a shadow, nor does it give other physical cues (e.g., a mechanical barrier) that would allow the plant to avoid or grow around it by itself. We perform the obstacle avoidance task in two different experiment settings. In the first experiment (*left target experiment*), a fixed target is located at 5.12 cm to the left of the plant anchor and 17.9 cm above it. We evaluate the controller in four different scenarios where a rectangular obstacle ( $7 \times 3$  cm<sup>2</sup>) is centered at four different locations. In the first scenario the obstacle is centered  $\approx 8.24$  cm left of the plant anchor point and at a height of 8.8 cm. In the second scenario, the obstacle is 2.67 cm further to the right (closer to the plant), making the task more challenging. In the third scenario, the obstacle is an additional 2.67 cm further to the right, making it impossible for the plant to reach the target. Finally, the obstacle is an additional 5.33 cm further to the right, this time clearing the area enough for the plant to again reach the target. In the second experiment (*middle target experiment*), a fixed target is located above the plant anchor at a height of 17.9 cm. Here, we have only two scenarios. In the first scenario, the obstacle is centered at  $\approx 3$  cm right of the plant anchor and at a height of 8.8 cm. In the second scenario, the obstacle is centered at  $\approx 3$  cm left of the plant anchor and at a height of 8.8 cm. Hence, the controller requires different strategies to control the plant for different target/obstacle configurations, which makes the task more challenging. In addition, the plant stiffens only over time, requiring the plant tip to be guided in wide deviations from the plant's ending configuration.

*3.3.2 Evolutionary approach.* We use MultiNEAT [5], a portable library that provides Python bindings to NEAT [20], to evolve ANN controllers. We use the NEAT parameters set from our previous work [14, 22, 23]. We follow a step-wise simulation approach, where the stem description  $S_t = (x_t^a, y_t^a, x_t^{S_2}, y_t^{S_2}, \dots, x_t^{S_g}, y_t^{S_g}, x_t^g, y_t^g)$ , the target position  $x_t^*$ , and the coordinates of an obstacle  $x_i^o$  are input to the ANN at each time step  $t$ . The output of the network ( $C_t$ ) regulates the light settings. If  $C_t \leq 0.5$ , it triggers the left light source, otherwise, the right. The current plant condition and light setting  $(x, C)_t$  impact plant behavior during that time step. The *LSTM stem stiffening and motion model* is used to predict the next plant stem  $S_{t+1}$  accordingly. For experiments in reality, an image of the plant is processed to determine  $S_{t+1}$  (see Sec. 3.2.2). The simulation is stopped when the tip  $y_t^g$  value is equivalent to  $\approx 21$  cm or once the plant touches an obstacle. Beans require  $\approx 72$  hours to grow that high. This overhead is relatively manageable, and allows enough growth to exploit stem stiffening and avoid obstacles.

At the simulation end (at  $t = f$ ), performance of the ANN controller is evaluated using a behavioral fitness function  $F$  (according to the classification in [19]). Plant motion is rewarded by measuring the distance traveled by tip  $g$  towards the target along both  $x$  and  $y$  axes as  $x_r = |x^*| - |x^* - x_f^g|$ ,  $y_r = |y^*| - |y^* - y_f^g|$ , where  $(x^*, y^*)$

<sup>5</sup>Find a video at: <https://vimeo.com/265144652>

is the target position. The fitness  $F$  is then calculated by

$$F = \frac{x_r + y_r}{|x^*| + |y^*|}, \quad (18)$$

where  $|x^*| + |y^*|$  is the theoretical best fitness value the controller can achieve. If the controller is evaluated at different scenarios, then its fitness value is the average of all evaluations.

## 4 RESULTS

Based on the *stem stiffening and motion model* (see Sec. 3.2.3), we evolve the robotic controllers and evaluate their performance in simulation. Next, we transfer the fittest controllers to reality and investigate the extent of the reality gap.

### 4.1 Evolving controllers in simulation

First, we report the results of the *left target* experiment. The boxplots in Fig. 2(a) and function boxplots in Fig. 2(b) show the performance of 20 independent evolutionary runs, 1000 generations each. The best fitness per generation for all evolutionary runs is considered. Notice the steady increase in median until convergence is reached around the 500th generation. In this experiment the controller is evaluated according to four scenarios (see Sec. 3.3.1). According to the behavior of one of the best controllers (fitness of 82.5%), the controller is able to determine whether or not it needs to exploit the stem's natural stiffening over time, in order to avoid hitting the obstacle. In case there is a possibility to hit an obstacle (e.g., second scenario), the controller steers the plant away into the opposite direction of the obstacle, long enough to obtain sufficient stiffness at the lower parts of the stem, see Fig. 2(d), then steers the plant back towards the target area, see Fig. 2(e). In case the obstacle is not blocking the way (e.g., forth scenario), the controller leads the plant directly towards the target area (i.e., no stiffening is necessary). The behavior in all scenarios can be seen in the video<sup>6</sup>.

Second, we report the results of the *middle target* experiment. Here, we also have 20 evolutionary runs, 1000 generations each as shown in Figs. 2(f) and (g). In contrast to the previous experiment, the convergence is reached earlier, around the 350th generation. This indicates that the task here is easier by comparison. The expected behavior of the evolved controller is to first steer the plant away from the target while the stem stiffens, then steer the plant back to reach the target while avoiding the obstacle—as in the *left target* experiment. However, the controller here (fitness of 97.3%) steers the plant to the obstacle side near the target, see Fig. 2(h), then switches between the two light sources until the plant obtains enough height and stiffness without hitting the obstacle, see Fig. 2(i). Finally it steers the plant tip towards the target, see Fig. 2(j).

### 4.2 Performance of controllers in reality


To test controller performance in the physical world, we examine whether it can guide a natural bean plant around a virtual obstacle without colliding, and reach the target area. This addresses the reality-gap problem [15], which states that controllers evolved in simulation do not always transfer to a real setup, because the simulation is limited in principle. To test the reality gap, we use the setup described in Sec. 3.1. Computer vision detects the stem,

feeding into the ANN evolved in simulation, which controls light stimuli provided to the real bean plant. The controller we select is evolved in the *left target* experiment, with the obstacle centered  $\approx 6.3$  cm left of the plant anchor, 12.5 cm above the soil. The bio-hybrid setup completed the task with the real plant, although the plant's circumnutation behavior brought it in close proximity to the obstacle (see Fig. 3 and video<sup>6</sup>). In the experiment, the controller initially maintained the right light, guiding the plant away from the obstacle for  $\approx 37$  h, until the plant was 7.4 cm right of the plant origin and 15.4 cm above the soil, see Fig. 3(b). Then it switched to the left light, quickly bringing the plant tip to the opposite side, while the stem retained some stiffness. After less than 2 h, the plant tip is roughly in the center, with a pronounced curve in the stem, see Fig. 3(c), as lower tissues already stiffen. Then follows a phase of quick light alterations, as the controller guides the plant tip close to the obstacle edge, leaving it near to the target after clearing the obstacle. The left light is then triggered for another 37 h, successfully guiding the plant to the target, while the curvature of the stiffened stem allows it to entirely avoid the obstacle. The controller's effectiveness in exploiting the plant's stiffening behavior is seen by comparing this result to those of [14]. Here, stiffening has resulted in noticeable stem curvature, while in [14] the stems have a straight shape, even after being steered to targets on opposing sides. The evolved controller together with the real plant achieves 92.4% fitness, see Fig. 3(d). The experiment was repeated two further times, achieving fitnesses of 92.0% and 87.7%. In the latter, the bean grows abnormally. It is significantly slower (by half) than the others, which have comparable growth speeds to those in [14, 22, 23]. While both other experiments last 75 h, this bean only grows above the obstacle in hour 84. However, the controller behaves correctly, and the bean approaches the target until hour 198, when its preexisting anomalies cause collapse. We record the fitness when the plant tip reaches the target area, because it is difficult to hit a single point, considering the stiffened tissues. However, it is possible to achieve higher fitness values, as a maximum fitness of 99.3% was later observed in the second experiment.

## 5 DISCUSSION AND CONCLUSION

Following the objective of using natural plants to do additive manufacturing and to implement a real-world, tunable developmental system, we have set up a toolchain to shape natural plants in an evolutionary robotics approach. We acquired data about the growth and motion of a plant, trained a state-of-the-art LSTM as plant model, evolved robot controllers using the LSTM as simulator, and successfully tested these controllers for the reality gap. Our focus was on the delicate interplay of plant motion and tissue-stiffening to shape plants around obstacles with collision-free control. Early on, the plant motion has to be controlled strategically to provoke the correct stiffened shape later. We call this particular phenomenon 'embodied memory' because the plant tip motion and the orientation of the whole plant during the experiment is integrated over time and partially reflected in the final stiffened shape of the plant. This is particularly different from other tasks in evolutionary robotics, such as the navigation of a mobile robot, where the full history (robot trajectory) has only a minor and indirect influence on the task completion. The task could arguably be compared to the

<sup>6</sup>Find a video at: <https://vimeo.com/265144652>




**Figure 2: Performance of the evolutionary process over generations for 20 evolutionary runs.**

control of a robot arm where joints closer to the base lose their flexibility over time. In comparison to our previous work [14], where controllers were evolved to guide the plant tip into randomly generated targets, our target control behavior is evidently more complex, because the evolutionary process required 300 more generations till

convergence. The controller here needs to be aware of the whole plant body instead of only the tip, in order to be able to avoid hitting the obstacle at any point along the stem.

A key achievement of this work is the successful application of methods from machine learning (LSTM network) to create a holistic





**Figure 3: Sequence of images showing the course of a reality-gap experiment.** The yellow circles on top indicate which light is on (filled: on, empty: off). The larger filled red circle is the target area and the gray rectangle is the obstacle the plant is not allowed to touch.

plant model. Unfortunately, such models representing the plant's macroscopic reactions to stimuli are not readily available from plant science. We have shown that with data from a few generic plant experiments, a sufficient model can be obtained. However, the limited availability of data is a challenge as common in machine learning and especially deep learning. Growing plants as such can be parallelized but considerable costs are added by controlling the light conditions and tracking, hence data is sparse. We have reported our approach to data augmentation, which may also have some potential to scale up.

The presented methodology with heavy use of machine learning techniques has potential to scale up to more desirable growth tasks that go beyond mere obstacle avoiding. Options are to grow plant patterns on meter-scales or more, to grow and control multiple plants within the same area, and to grow also 3D-patterns. Besides their natural aesthetics these grown shapes may also have functionality, for example, as architectural artifacts (green walls, roofs, etc.). Therefore, we plan to make the transition to a 3D setup in future work, where we can grow more complex shapes, such as spirals, geometrical objects, or even writing. Controlling multiple plants concurrently will also add complexity, especially once we allow them to interact. We plan to automatically braid plants, use them to change material properties in construction, to investigate different plant species, and to grow complex structures, such as meshes or even benches. In addition, we investigate options to use phytosensing (i.e., using plants as sensors) that could help to implement synergistic robot-plant interactions. Hence, the presented machine learning approach of shaping plants opens doors for autonomous bio-hybrid systems with promising applications.

## ACKNOWLEDGMENTS

Project 'flora robotica' has received funding from the European Union's Horizon 2020 research and innovation program under the FET grant agreement, no. 640959.

## REFERENCES

- [1] M. Abadi et al. 2016. TensorFlow: A System for Large-Scale Machine Learning. In *OSDI*, Vol. 16. 265–283.
- [2] R. Bastien et al. 2015. A Unified Model of Shoot Tropism in Plants: Photo-, Gravitational and Propioception. *PLoS Comput Biol* 11, 2 (2015).
- [3] J. C. Bongard. 2013. Evolutionary robotics. *Commun. ACM* 56, 8 (2013), 74–83.
- [4] O. E. Checa et al. 2008. Mapping QTL for climbing ability and component traits in common bean (*Phaseolus vulgaris* L.). *Molecular Breeding* 22, 2 (2008).
- [5] P. Chervenski et al. 2012. MultiNEAT. (2012). <http://www.multineat.com/>
- [6] F. Chollet et al. 2015. Keras. (2015). <https://keras.io/>
- [7] J. M. Christie et al. 2013. Shoot phototropism in higher plants: New light through old concepts. *American Journal of Botany* 100, 1 (2013), 35–46.
- [8] P. C. Garzón and F. Keijzer. 2011. Plants: Adaptive behavior, root-brains, and minimal cognition. *Adaptive Behavior* 19, 3 (2011), 155–171.
- [9] A. Graves et al. 2013. Speech recognition with deep recurrent neural networks. In *IEEE Int. Conf. on Acoustics, speech and signal processing (ICASSP)*. IEEE.
- [10] J. Halloy et al. 2007. Social Integration of Robots into Groups of Cockroaches to Control Self-Organized Choices. *Science* 318, 5853 (Nov. 2007), 1155–1158.
- [11] H. Hamann et al. 2015. *flora robotica* – Mixed Societies of Symbiotic Robot-Plant Bio-Hybrids. In *Proc. of IEEE Symposium on Computational Intelligence*. IEEE.
- [12] H. Hamann et al. 2017. *Flora robotica*—An Architectural System Combining Living Natural Plants and Distributed Robots. *arXiv:1709.04291* (2017).
- [13] S. Hochreiter and J. Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [14] D. N. Hofstadler et al. 2017. Evolved Control of Natural Plants: Crossing the Reality Gap for User-Defined Steering of Growth and Motion. *ACM Trans. Auton. Adapt. Syst. (TAAS)* 12, 3 (2017), 15:1–15:24.
- [15] S. Koos et al. 2013. The transferability approach: Crossing the reality gap in evolutionary robotics. *IEEE Trans. on Evo. Comp.* 17, 1 (2013), 122–145.
- [16] A. Lindenmayer. 1975. Developmental algorithms for multicellular organisms: A survey of L-systems. *Journal of Theoretical Biology* 54, 1 (1975), 3–22.
- [17] M. Mansouryar et al. 2012. Smoothing via iterative averaging (SIA) a basic technique for line smoothing. *Int. Jo. of Computer and Electrical Eng.* 4, 3 (2012).
- [18] S. T. Namin et al. 2017. Deep Phenotyping: Deep Learning For Temporal Phenotype/Genotype Classification. *bioRxiv* (2017).
- [19] A. L. Nelson et al. 2009. Fitness functions in evolutionary robotics: A survey and analysis. *Robotics and Autonomous Systems* 57 (2009), 345–370.
- [20] K. O. Stanley et al. 2004. Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research* 21, 1 (2004), 63–100.
- [21] I. Sutskever et al. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*. 3104–3112.
- [22] M. Wahby et al. 2015. Evolution of Controllers for Robot-Plant Bio-Hybrids: A Simple Case Study Using a Model of Plant Growth and Motion. In *Proc. of the 25th Workshop on Computational Intelligence*. KIT Scientific Publishing, 67–86.
- [23] M. Wahby et al. 2016. An evolutionary robotics approach to the control of plant growth and motion. In *IEEE 10th Int. Conf. on Self-Adaptive and Self-Organizing Systems (SASO)*. IEEE, 21–30.
- [24] R. A. Watson et al. 2002. Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Sys.* 39, 1 (2002).
- [25] P. Zahadat et al. 2014. Social Adaptation of Robots for Modulating Self-Organization in Animal Societies. In *IEEE Int. Conf. on Self-Adaptive and Self-Organizing Systems Workshops (SASOW)*. 55–60.
- [26] P. Zahadat et al. 2017. Vascular Morphogenesis Controller: A Generative Model for Developing Morphology of Artificial Structures. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '17)*. 163–170.
- [27] A. Zamuda et al. 2014. Vectorized Procedural Models for Animated Trees Reconstruction using Differential Evolution. *Information Sciences* 278 (2014), 1–21.