

Self-sustainability Challenges of Plants Colonization Strategies in Virtual 3D Environments

Kevin Godin-Dubois¹, Sylvain Cussat-Blanc¹, and Yves Duthen¹

University of Toulouse, IRT - CNRS UMR 5505, 2 rue du Doyen Gabriel Marty,
31042 Toulouse, France {kevin.dubois, sylvain.cussat-blanc,
yves.duthen}@irit.fr

Abstract. The Biosphere is a bountiful source of inspiration for the biologically inclined scientist, though one may be seized by the twists and turns of its complexity. Artificial Life emerged from the conundrum of condensing this overwhelming intricacy into a tractable volume of data.

To tackle the distant challenge of studying the long-term dynamics of artificial ecosystems, we focused in this work our efforts on plant-plant interactions in a simplified 3D setting. Through an extension of K. Sims' directed graphs, we devised a polyvalent genotype for artificial plants development. These individuals compete and collaborate with one another in a shared plot of earth subjected to dynamically changing environmental conditions. We illustrate and analyze how the use of multi-objective fitnesses generated a panel of diverse morphologies and strategies. Furthermore, we identify two driving forces of the emerge of self-reproduction and investigate their effect on self-sustainability.

Keywords: Artificial plants · Ecosystems · Autonomous reproduction · Self-sustainability

1 Introduction

Natural ecosystems are staggering by virtue of their intricate interactions networks and seemingly endless detail levels. Adaptation, co-evolution and arming races are a small subset of the myriad evolutionary strategies observed in the world. In order to translate this onto an artificial medium, one has to address three key issues. The first is finding functional, adaptive and yet computationally tractable representations for both the genomic and phenotypic aspects of the 'creatures'. Then, genetic variation and discrimination methods, the keystones of Darwin's natural selection process, ought to be devised. Ultimately, environmental pressure, whether coming from the inanimate (abiotic) or living (biotic) surroundings, should be strong enough so that diversity can emerge without excessively narrowing the viable behaviors range.

A number of models have been devised to tackle the problem of artificial morphogenesis starting with the L-Systems which successfully generated plant

morphologies in both 2-D [1] and 3-D environments [2], as well as mobile creatures [3]. In a similar fashion, the directed graphs from [4], define distinct organs and the relationship between them, allowing for very compact encoding of motile morphologies. Closer to biology is the use of Genetic Regulatory Networks to control cellular building blocks. By using such a finer-grained representation it was shown possible to evolve cell clusters towards specific shapes [5] or generate creatures with self-organized organs in non-trivial environments [6].

With rules as simple as the one in the Game of Life [7], elaborate ecosystems emerged with no reliance on evolutionary techniques. Other contributions simulated self-reproducing computer programs into a 2D memory grid where the organisms had their duplication mechanism embedded into their life-cycle [8]. Further complexification of the phenotypic space led to ecosystems such as [9,10] where the environment was a 2D continuous grid and the ‘animals’ had to manage their reproduction cycle by actively searching for mates in their surroundings.

In case of co-evolution or competition (e.g. [11]), there is a biotic component to the ecosystem in the sense that from the viewpoint of an individual, every other entity is a hard-to-control part of its environment. When considering the abiotic, i.e. non-living, component of earth’s ecosystems, it has been shown that, with only water availability and temperatures, it is possible to model most of the biodiversity observed in nature [12]. The same holds true for artificial simulations where the impact of environmental factors can be a driving force for the speciation process [2]. Additionally, dynamically changing the local constraints pushes individuals out of local minimums and promotes adaptability [13].

Our objective is to design a virtual ecosystem which tackles all three points by modeling both plants and animals in a shared environment thus inducing complex interactions and survival/reproduction strategies starting, in this paper, with the growth of vegetals. To this end we first detail, in section 2, the use of both an extended version of the directed graphs (hereafter named ‘Graph-tals’) from [4], which produces functional morphologies in the face of unknown (a)biotic constraints, and a novel reproduction scheme with self-controlled speciation capabilities. We then move on to exploring, in section 3, various outcomes stemming from this implementation in terms of growth and adaptation strategies as well as the impact of the evolution process itself before highlighting the necessary extensions to this model.

2 Self-reproducing vegetals

Starting from the breakdown of the various components of an individual plants genome (section 2.1), we expand upon how organs interact with one another to promote the whole organisms well-being and survival capacities (section 2.2). A detailed presentation of a speciation-oriented approach to self-controlled, self-reproduction follows in section 2.3. We conclude this overview of the model by the external components of the system, i.e. the environment (section 2.4) and the englobing ecosystem (section 2.5).

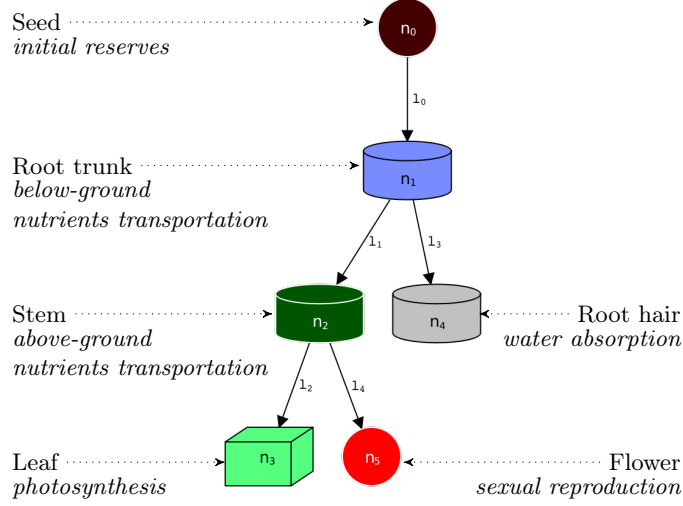


Fig. 1: Base graph in all following experiments

2.1 Graphtal fields

As in [14], our genotype is composed of both structural (figure 1) and behavioral instructions. Each node n_i describes an abstract sub-organ and is composed of an *id* and a set of parameters n_i^d specifying its skill which in turn defines its shape (sphere, box, cylinder). Furthermore, it also codes for the initial dimensions, density, color and maximal growth factor of the corresponding organ.

Behavior is controlled by two tuples $\mathbf{A}, \mathbf{S} \in [0, 1]^E$ with E the number of elements (limited to water and glucose in this work). \mathbf{A} models an organ's balance between production and consumption: a value of 0 (resp. 1) indicates a source (resp. sink) for this specific element. \mathbf{S} enables quiescent behavior by imposing a threshold below which no growth or budding actions can be performed.

A link $l_i : n_1 \rightarrow n_2$ expresses a *growth relationship* from n_1 (the parent node noted l^i) to n_2 (the 'child' node noted l^o) and, also, contains a set of parameters noted l^d describing the position on the parent's surface and the relative orientation. A special field r is used to regulate recurrent connections by imposing a maximal depth ($r > 0$) or indicate a terminal node ($r = -1$).

To easily code for the highly regular structures observed in nature, each link can select a repetition pattern from:

- *none*: no repetition
- *radial*(\mathbf{V}, N): $N-1$ copies evenly rotated around \mathbf{V}
- *random*(N, S): $N-1$ copies randomly placed

Ultimately, a handful of plant-wide parameters are stored separately. Namely the growth speed and maximal size factor, the individual's sex and two sets of parameters dW and $\{\mu, \sigma^i, \sigma^o\}$ that will be detailed more thoroughly in following sections.

2.2 Metabolism

At the beginning of a simulation, plants' seeds are filled with a limited amount of nutrients, i.e. water and glucose, to start the growth process but longer-term survival requires a strategy to maintain comfortable resource levels. Organs die from starvation if any of their nutrient reserves are exhausted and a plant is considered dead when it contains less than two organs as it can no longer sustain itself. No programmed death is implemented in either the genotype or the simulation. In these experiments, the environment contains only two types of resources from which nutrients can be extracted: water and light.

The first one is extracted through below-ground root hair and must be dispatched through the plant and more importantly to the leaves. Indeed these later require both resources to produce glucose according to Equation (1).

$$G_i = .025 * S_i * \mathbf{L} \cdot \mathbf{N}_i \quad (1)$$

where \mathbf{L} points to the sun, S_i is the leaf's photoreceptive surface and \mathbf{N}_i the \mathbf{z} axis in its local coordinates system.

Resources distribution is implemented through a decentralized mechanism which follows the gradient of nutrient. Every simulation step, organs share a portion of their reserves according to the transport equation (2) with $e \in g, w$ a nutrient.

$$d_{A \rightarrow B}^e = \frac{k_{AB}^e * stored(A) * needs(B)}{\sum_{o \in C(A)} needs(o)} \quad (2)$$

where $needs(x)$ (resp $stored(x)$) is the need (resp. stored amount) in nutrients in an organ x , $C(x)$ the set of organs connected to x and A, B two organs so that $0 \leq needs(A) < needs(B) \leq 1 \wedge B \in C(A)$. The k_{AB}^e term is a refractory coefficient dependant on the types of A and B.

2.3 Autonomous reproduction

One of the most powerful tools available to Life is its ability to adapt through the process of natural selection. Over the course of history numerous propagation scheme have been developed. In this work, we chose to focus on sexual reproduction because of its greater degree of interactions and inter-species diversity.

To this end, we improved upon the previous work in [14] by including genomic components devoted to reproduction: sex, compatibility metrics $\{\mu, \sigma_i, \sigma_o\}$ and sexual organs. These interact with one another according to algorithm 1. Given $G_1 = (N_1, L_1)$ and $G_2 = (N_2, L_2)$, two genomes, the alignment procedure creates three subsets:

$$\begin{aligned} - M_a &= \{\{l_1, l_2\}, l_1^{id} = l_2^{id} \wedge l_1^i = l_2^i \wedge l_1^o = l_2^o\} \\ - M_{ia} &= \{\{l_1, l_2\}, l_1^{id} = l_2^{id} \wedge (l_1^i \neq l_2^i \vee l_1^o \neq l_2^o)\} \\ - M_i &= \{l_1 \in L_1, \nexists l_2 \in L_2, l_1^{id} = l_2^{id}\} \cup \{l_2 \in L_2, \nexists l_1 \in L_1, l_2^{id} = l_1^{id}\} \end{aligned}$$

```

Data: P, set of plants
M ← {p ∈ P/p is male};
for m ∈ M do
  Gm ← genotype(m);
  for sm, stamen ∈ m do
    f ← random female, with pistils, in range of m;
    Gf ← genotype(f);
    A ← align(Gm, Gf);
    Cmf ← compatibility(A, Gf);
    if random toss with probability Cmf then
      delete sm;
      pf ← random pistil from f;
      pf ← Fruit(Mutate(Cross(A, Gm, Gf)));
    end
  end
end

```

Algorithm 1: Mating process




Fig. 2: Genetic compatibility function

which highlight the structural similarities between both individuals. The genetic distance d is then computed based on the average field-by-field difference of comparable node and link data, i.e. those in M_a and M_{ia} . Links in M_i are given the maximal distance of 1. The total sum is given by equation (3).

$$\begin{aligned}
 d(l_1, l_2) &= \frac{1}{3}(d(l_1^i, l_2^i) + d(l_1^o, l_2^o) + d(l_1^d, l_2^d)) \\
 D(G_1, G_2) &= |M_i| + \sum_{\{l_1, l_2\} \in M_a \cup M_{ia}} d(l_1, l_2)
 \end{aligned} \tag{3}$$

This crossover operator differs from those commonly found in the literature [4,15,6] on three points: 1) it can fail early on, 2) is biased by the *female* genome and 3) has low resistance to large structural differences. The rationale behind

point 3 is that, instead of devising a robust operator that can produce a somewhat viable offspring from two completely unrelated individuals, a minimalist alignment procedure is better suited to sexual reproduction of same species creatures in which the population is mostly homogeneous. Indeed, point 1 guarantees that the more both genomes are different the less likely it is that crossing will be attempted at all. The decision of aborting or proceeding with the reproduction is left to the female individual as, in this sexual scheme, it will have to provide the fruit in resources.

Having the compatibility function embedded into the genome gives each species a segregation scheme which is of utmost importance as the number of nodes/links is not fixed, thus requiring the optimal genetic distance to be adapted as evolution goes. Furthermore having both in-/out-breeding coefficients allow for the specification of the search spaces with adaptive plants accepting a broader range of incoming genetic material while more conservative ones could instead focus on controlled inbreeding to solidify their alleles.

2.4 Environment

Selective pressure should emerge on its own from the interaction between self-interested individuals. However the environment could play a pivotal role in guiding the complexification process whether by imposing harsh restrictions on the amount of available resources or by displaying a range of (from a creature's point of view) semi-random dynamics.

For this work, the only varying part of the environment is the sun, which produces season cycles through the following set of parameters: day (100 simulation steps) and year (300 days) lengths, latitude ($\pi/4$) and declination ($\pi/8$), i.e. typical values for a temperate climate. Additionally, the ground was a $W = 10\text{m}$ flat square with a constant supply of water.

2.5 Ecosystem

A complete ecosystem is composed of both a description of the environment (currently not evolved) and a set of plants 'templates'. Each of these latter represents whole species whose strategies for survival will be pitted against one another in their shared piece of earth.

The procedure to translate these templates into a densely populated ecosystem is straightforward. First the environment is divided into as many cells as the requested number of plants (100 in this experiment) and the largest seed size is tested against half the cell size. If this fails, the whole ecosystem is deemed non-viable and the simulation is aborted, thus preventing plants from having too large initial reserves. Otherwise, each cell is subdivided once more in four and a plant is placed in a single subcell with a random genome from the set of templates and a random vertical rotation. This leaves enough room for autonomous reproduction to place offspring even when the initial population has not entirely died out.

In the current experimental settings, the number of plants is set to 100 and only one species is considered. Furthermore, every random number used during the simulation (plants position, rotations, iterations, etc.) is generated from a fixed seed provided by the genotype (not evolved but randomly set).

3 Colonization dynamics

3.1 Evolution protocol

This work comes within the scope of studying long-term evolutionary trends in elaborate 3D ecosystems. However, evolving, from scratch, such systems with a non-trivial degree of complexity would require a prohibitive amount of computational resources. Stemming from this intent, the following experiment was designed to generate usable individuals to seed an environment with. Viable plants would thus have to develop strategies to both survive and reproduce so that their genetic material does not die off.

The evolution protocol relied on evolution programming where plants' genomes underwent single point, equiprobable, mutation on all of the fields mentioned in section 2.1. Evaluating a genotype implies populating an empty environment as described in section 2.5 and then stepping back for a maximum of $N = 60000$ simulation steps (2 simulated years) to see whether autonomous dynamics would emerge.

In order to limit the search space to the genetic fields of the plants, the environment was kept constant in all runs. As we aimed for both efficiency and diversity, we devised a range of fitness functions F_* as described below.

$$\begin{aligned}\nu &= \frac{1}{NP} \\ F_b &= \nu \sum_{t \in N} \sum_{p \in P} biomass(p, t) \\ F_p &= \nu \sum_{t \in N} \sum_{p \in P} production(t, p) \\ F_c &= \frac{\nu}{W^2} \sum_{t \in N} surface(t) \\ F_a &= \nu \sum_{p \in P} lifespan(p) * 2^{-\alpha_p}\end{aligned}$$

Where $surface(t)$ corresponds to the total surface covered by plants at time t and α_p is designed to provide a smoother gradient towards reproduction. Plainly put, these aim at producing plants which are: F_b) large, F_p) many-leaved, F_c) wide, F_a) fast reproducers. Given that every fitness is likely to be exploited into non-desired behaviors, a fifth one F_m is introduced that evaluates genomes on all four criteria at the same time.

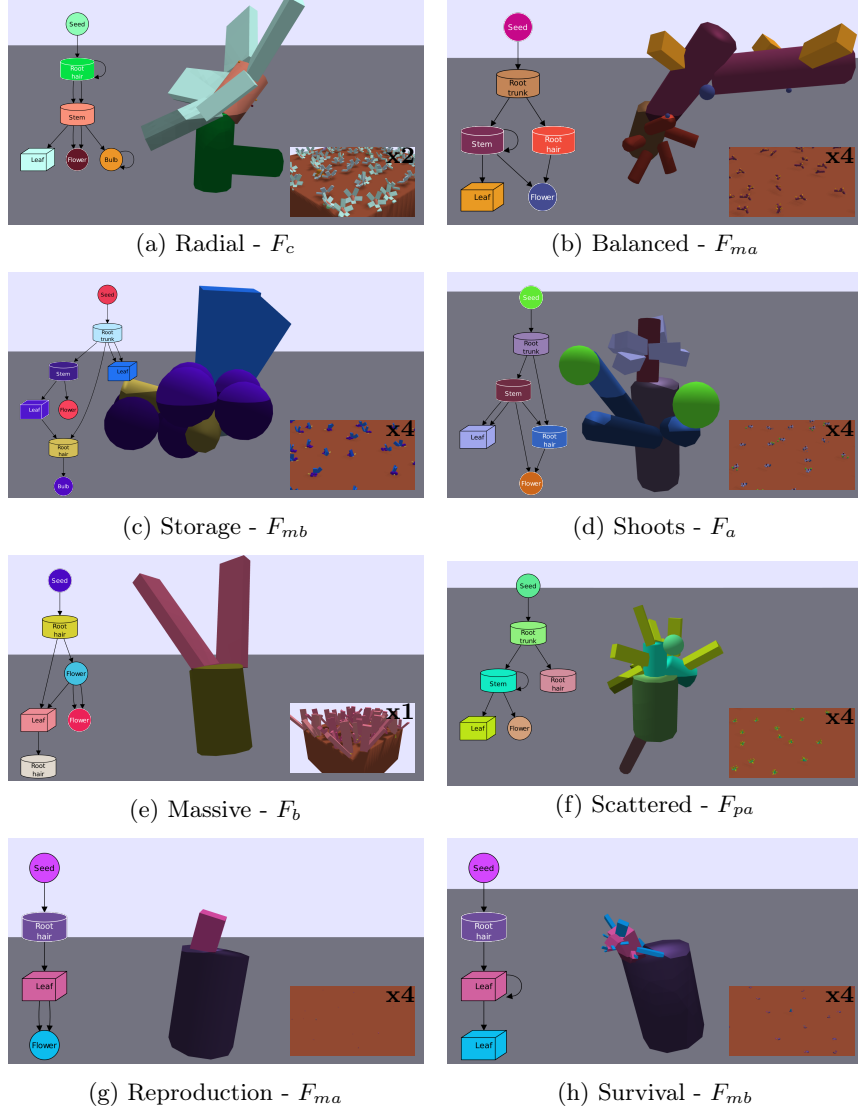


Fig. 3: Examples of the morphologies developed. (a) to (g) are at the 20th day and (h) is at the 30th.

From left to right: Genome, Single individual, Ecosystem.

The fitness that produced this individual is indicated in the caption with F_{ma} indicating the age criterion of the multi-objective fitness F_m .

Videos of these individuals' full ecosystem can be seen on <https://vimeo.com/album/5075632>.

Furthermore, in order to prevent local optimum a novelty metric is used as proposed in [16]. An individual’s ‘footprint’, i.e. its synthetic behavioral description, is (F_a, F_b, F_p, R, G, S) , with R the number of successful autonomous reproductions, G the number of autonomous generations and S the seed size.

An autonomous *reproduction*, in this context, is defined as two individuals (m, p) embedded in a simulation deciding on generating an offspring through the process described in algorithm 1. The autonomous *generation* g_c of such an offspring is $g_c = \max(g_m, g_p) + 1$, where g_m (resp. g_p) is the autonomous *generation* of the mother m (resp. father p).

In each scenario, plants are evaluated on two to five criteria using a tournament selection where 3 participants are randomly selected from the population and compete on a random objective as described in [6].

Ten runs per fitness were dispatched on a cluster of Bi-Intel(r) IVYBRIDGE 2,8 Ghz 10-cores and were re-launched as soon as they completed an evolution (250 generations) with a maximal, total, duration of five hours.

3.2 Morphologies

While three out the five fitnesses performed an average of two evolutions in the given time frame ($F_a = 2.4, F_c = 2.5, F_m = 1.8$), the remaining two behaved very differently: while F_b produced 8.4 ‘champions’ per run, F_p did not manage to bring a single one to the 250 generations threshold (min=93, max=199). This can be explained by observing the evaluation times of those final individuals which range from 12 milliseconds up to 10+ minutes.

In order to gain a further grasp on the situation, we manually examined the phenotype of the 40 best champions (out of a total of 215) that is 5 for every single objective fitness and another 5 for each criterion in F_m . Summarized in figure 3 are the morphologies of the most interesting creatures.

As one can see these evolutions produced very different strategies to cope with the environment and their respective fitness. Variation in the sun’s position and the plants’ relative orientation led to either having large leaves so that production is maximized during short favorable moments (3a, 3c, 3e), or numerous, evenly spread, leaves so that sunlight can be efficiently gathered throughout the day by different parts of the plant (3d, 3f, 3h).

Root morphology was not thoroughly investigated, due to the uniform water distribution exerting only very limited evolutionary pressure, and most individual manage with a simple root trunk connected to a handful of capillary tubes (3b, 3d, 3f). Some even went as far as to completely forsake the former (3a, 3e, 3g).

As the autonomous reproduction process starts from flowers, their growth is of utmost importance for a species’ permanence. All plants except two from figure 3 actually generate at least one such organ, though only 3g, 3e, 3a manage to bring them to maturity.




Fig. 4: Typical examples of the three strategies' dynamics. Individuals are taken from figure 3 with the red, green and blue curves corresponding to 3f, 3g and 3e

3.3 Strategies

From these morphologies and their associated dynamics graphs, we can extract three main strategies: quiescence, expansion and reproduction as illustrated in figure 4. The first one (in red) is quite straightforward in its survival method. One can see on the graphs that after a short burst of activity, early on in the simulation, this type of individuals goes into a quiescent state, keeping its metabolic value in a comfortable range so that most plants make it to the end of simulation. The expansionist (in blue) however, adopts a radically different approach: instead it tries to reach as fast as it can a mature state which can, depending on the plant, take up to a full year. This allows the ecosystem to compensate for the extremely high mortality rate: in the example depicted, 96% of the population dies in the first hundred days. Finally, the reproduction strategy (in green) relies on having the smallest possible morphology, i.e. a small seed and a single root hair directly connected to the leaf. Resources are mostly directed towards producing mature fruits as quickly as can be, thus maintaining a population in a safe range $[60, 80]$ in this case).

It is interesting to compare these behaviors with those obtained in [2] where varying environmental factors led to the emerge of the CSR triangle [17]. Indeed, a plant population under favorable conditions should evolve towards individual competition while with decreasing resources availability a slower, more conser-




Fig. 5: Repartition of average organ count per fitness across all runs.

vative, metabolism is expected. If exposed to recurrent, localized uncontrolled deaths, the ruderals would thrive with their fast life cycle and colonization approach.

The fact that all three strategies emerged within identical environments shows that, on the one hand, the genetic search space is large enough to contain very distinct viable genomes even before being subjected to an evolutionary process. On the other hand, it also warns about a possibly *too* large search space with functional genotypes separated by wide unfit combinations.

3.4 Influence of evaluation criteria

We now turn our attention to the evolution procedure itself and more specifically the contribution of our fitness functions set. The diversity of criteria used induced a similar amount of variability in the obtained genomes as one can see the range of morphologies and behavior obtained. All fitnesses, however, did not perform equally both in terms of complexity (see figure 5) and relevance. Indeed, while F_b produced plants that could grow at a sustained pace, they proved quite simplistic, morphologically speaking, with almost 60% of the champions growing less than five organs. Given that every random graphal starts with this specific amount, it shows that evolution discovered that the bigger one wants to grow, the smaller the genotype.

A similar trend can be observed in F_a , though with a slight offset caused by the necessity of having sexual organs. The global strategy for this fitness is as described in details for individual 3g: small genome, small plant, fast mating. This tendency is reversed in F_p , with no instance in the ‘Minimalist’ section of the phenotypic space. Indeed, as glucose production requires both efficiently

positioned leaves and sufficient water uptake, evolution favored genotypes with repetitive structure. On the downside, this also led to extremely long evaluation times which prevent all runs to reach 250 generations with no overwhelming advantages over the competing fitnesses.

The coverage-oriented evolutions performed by F_c led to a more balanced distribution of organ count between minimalism and over-complexification. While this is the less biologically inspired criterion, it proved more robust to being exploited by the evolutionary algorithm and, paradoxically, brought more life-like individuals about, such as 3a.

Finally, the multi-objective fitness F_m generated more all-rounder creatures, that did not suffer from over-optimization. Indeed when looking at 3c, 3b or 3f, one can observe plausible morphologies made functional by the contradictory pull of all individual fitnesses. Furthermore, it settled in a complexity landscape similar to that of F_c , though with less exploration of the uppermost region.

3.5 Reproduction

In addition to the experiment described up until now, we subjected our model to two hypotheses on the emergence on self-sustainability:

Hypothesis **F** Fruit dissemination should be supported by the system until autonomous abscission can stabilise.

Hypothesis **D** Plants should be stressed by the unavoidability of their deaths.

The former is implemented through a collection algorithm that retrieves any fruit disconnected from its plant (voluntarily, through a parent’s organ death, etc) and proceeds to its dissemination through the usual algorithm. The latter is emulated by extending the simulation duration s_d when in presence of self-reproductive behavior. That is, for a number g of autonomous generations, the allotted number of years is $s_d = \min(10, \max(g, 1) * 2)$. This allows genomes exhibiting self-reproduction capabilities to reach much higher fitness values thus increasing their chance of producing offspring in the next generation of the genetic algorithm.

We tested all four combinations with none (f,d), one (f,D and F,d) and both (F,D) of the hypotheses active in the runs to see how this would influence the capacity to develop self-reproduction (fig 6).

The initial conditions (f,d) proved quite detrimental for the emerge of self-sustainability: indeed, aside from a single outlier reaching 7 autonomous generations, 93% of the individuals obtained do not reproduce at all. Only introducing adaptive simulation duration (f,D) reduces this number to 84% and creates a secondary behavioral cluster around the 4th generation. This dynamic is inverted when only collecting immature fruits with only 54% of non-reproducing runs and 38% reaching the first generation. Finally when both features [F,D] are enabled the threshold of less than 50% of infertile individuals is crossed, albeit slightly, and peak performance is at absolute maximal across all alternatives (19th, 22nd and 30th autonomous generations obtained in the 10 years allotted time-frame).

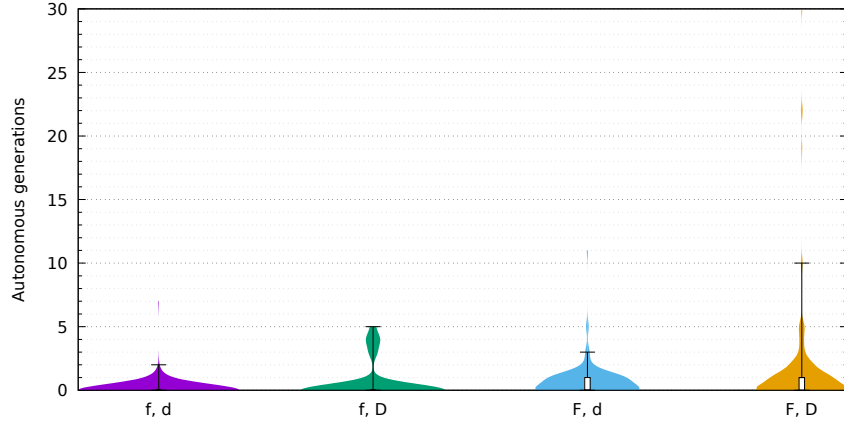


Fig. 6: Violin plots of the number of autonomous generation per run type (see the text for details). The colored area displays the kernel density overlaid with a boxplot with whiskers spanning the 95% confidence interval.

Based on the capacity to produce self-reproducing plants and the results of T-Test evaluations we can surmise that run (f,d) is outperformed by every other alternative ($p\text{-value} \leq 0.001$). Additionally, while no significant differences were detected between runs (f,D) and (F,D) the last one (F,D) shows better results ($p\text{-value} < 0.01$) than both of them.

In order to further understand the dynamics behind these differences in self-reproductive behavior, we investigated what ‘checkpoint’ individuals tended to stop at (fig 7). The different categories are

- None: No sexual organs were produced by the plants
- Flowers: Some were produced but never fecundated
- Fruits: Seeds were produced but never planted
- Repro.: Self-reproduction occurred.

At first glance, we can note different locations for the point(s)-of-failure depending on which Hypothesis was enabled. While run (f,d) seems to struggle at every checkpoint, loosing almost a third of its population each time, run (f,D) understood the importance of producing fruits but rarely found how to disseminate them into the world. Both (F,d) and (F,D) show no particular problem on this point due to the algorithm taking the lead when necessary and, instead, are clustered between individuals that do not attempt to reproduce and those that succeed.

In the end, both hypotheses are verified by this additional runs as, on the one hand, emphasizing the need to circumvent individual death promotes self-replication while, on the other hand, providing a fallback mechanism, until self-controlled abscission can stabilise, leads to fruit generation.




Fig. 7: Repartition of checkpoints for autonomous reproduction. Colors are the same as 6, for details of the checkpoints see text.

4 Conclusion

In this work, we devised a complex genetic encoding for plant morphologies derived from K. Sims' seminal work. The intricate mapping into a phenotype is highly dependent on the environmental conditions of both the abiotic (light, water) and biotic (competition) components. A self-reproduction scheme, based on our bail-out crossover algorithm, was introduced to generate self-sustaining ecosystems.

The ensuing experiment highlighted the difficulty of developing a balanced strategy between survival and reproduction with individuals mostly falling into either extremes. Furthermore, we observed that the size of the genotypic space was large enough to, in itself, create a speciation phenomena solely based on the starting point of an evolutionary run. We additionally noticed the positive impact of a multi-objective evaluation criterion to promote robustness of demeanor by simultaneously selecting from a larger range of viable mutations.

Obtaining self-sustaining individuals proved quite difficult in the current experimental settings, though we note the positive impact of providing a smoother transition towards spontaneous abscission and emulating programmed death through adaptive simulation durations. Indeed while having only limited effect on the proportion of non-reproducing individuals, this was shown to greatly reduce the production of un-fecundated flowers.

Multiple avenues of research are currently open for improving our model with regards both to plants and their environment. Indeed, although already quite complex, our graphals currently lack the means to sense, and be affected by, the temperature at ground level. As mentioned in the opening section of this paper, it is an essential component in earth's biodiversity and, thus, should be included in our future experiments. It could also be used to provide the plants

with some form of temporal perception. This would allow adaptive behaviors such as delaying the growth of a seed until a more favorable season or restarting leaf production comes ‘spring’.

The environment itself was kept mostly static for this experiment with the exception of the sun. There is ongoing development to add more diversity both topological, e.g. changing the altitude of different patches, and temporal, e.g. varying the hygrometry. These changes would induce stress in the ecosystem, thus promoting adaptation and diversity. By evolving these environments and selecting for those with the most stimulating patterns, we could obtain simulations with highly adapted yet clustered populations.

A further step in this direction would bring about co-evolution and competition by increasing the number of species used to seed the environment. We could then observe the richness of interaction between multiple strands of individuals, both plants and animal, on a long time-scale (as introduced in <https://vimeo.com/godinduboisalife/futureworks>). The use of graphpal to grow plants being straightforward to extend to animals, as this technique is extensively used in the literature.

5 Acknowledgments

This work was performed using HPC resources from CALMIP (Grant P16043) and the Bullet Physics SDK <http://bulletphysics.org>.

References

1. Prusinkiewicz, P., Lindenmayer, A.: The Algorithmic Beauty of Plants, The Virtual Laboratory, vol. 54. Springer New York, New York, NY (1990)
2. Bornhofen, S., Barot, S., Lattaud, C.: The evolution of CSR life-history strategies in a plant model with explicit physiology and architecture. *Ecological Modelling* 222(1), 1–10 (jan 2011)
3. Hornby, G.S., Pollack, J.B.: Evolving L-systems to generate virtual creatures. *Computers and Graphics (Pergamon)* 25(6), 1041–1048 (2001)
4. Sims, K.: Evolving 3D Morphology and Behavior by Competition. *Artificial Life* 1(4), 353–372 (1994)
5. Joachimczak, M., Wróbel, B.: Evo-devo in silico-a Model of a Gene Network Regulating Multicellular Development in 3D Space with Artificial Physics. *Alife* pp. 297–304 (2008)
6. Disset, J., Cussat-Blanc, S., Duthen, Y.: Evolved Development Strategies of Artificial Multicellular Organisms. In: 15th International Symposium on the Synthesis and Simulation of Living Systems (ALIFE XV 2016). pp. pp. 1–8. The MIT Press, Cancun, Mexico (2016), <https://hal.archives-ouvertes.fr/hal-01511892>
7. Gardner, M.: Mathematical games: The fantastic combinations of John Conway’s new solitaire game "life". *Scientific American* (1970)
8. Adami, C., Brown, T., Kellogg, W.K.: Evolutionary Learning in the 2D Artificial Life System Avida. In: *Artificial Life IV*. MIT Press (1994)

9. Metivier, M., Lattaud, C., Heudin, J.c., Universitaire, L.: A Stress-based Speciation Model in LifeDrop characters. *Artificial Life* (2002), <http://www.alife.org/alife8/proceedings/sub265.pdf>
10. Ventrella, J.: GenePool: Exploring the interaction between natural selection and sexual selection. *Artificial Life Models in Software* pp. 81–96 (2005)
11. Miconi, T.: Evosphere: Evolutionary dynamics in a population of fighting virtual creatures. In: 2008 IEEE Congress on Evolutionary Computation, CEC 2008. pp. 3066–3073 (2008)
12. Woodward, F.I., Williams, B.G.: Climate and plant distribution at global and local scales. *Vegetatio* 69(1-3), 189–197 (1987)
13. Canino-koning, R., Wiser, M.J., Ofria, C.: The Evolution of Evolvability : Changing Environments Promote Rapid Adaptation in Digital Organisms. *Proceedings of the European Conference on Artificial Life* pp. 268–275 (2016), https://www.mitpressjournals.org/doi/abs/10.1162/ecal_a_0047
14. Dubois, K., Cussat-Blanc, S., Duthen, Y.: Towards an Artificial Polytrophic Ecosystem. In: *Morphogenetic Engineering Workshop, at the European Conference on Artificial Life (ECAL) 2017 September 4* (2017), http://doursat.free.fr/docs/MEW2017_program_and_abstracts.pdf#page=4
15. Bonfim, D.M., de Castro, L.N.: Frankstree: A genetic programming approach to evolve derived bracketed l-systems. In: Wang, L., Chen, K., Ong, Y.S. (eds.) *Advances in Natural Computation*. pp. 1275–1278. Springer Berlin Heidelberg, Berlin, Heidelberg (2005)
16. Lehman, J., Stanley, K.O.: Exploiting open-endedness to solve problems through the search for novelty. *Artificial Life* (2008), http://eplex.cs.ucf.edu/papers/lehman_alife08.pdf
17. Grime, J.P.: Evidence for the Existence of Three Primary Strategies in Plants and Its Relevance to Ecological and Evolutionary Theory. *The American Naturalist* 111(982), 1169–1194 (1977)