



Beneficial Catastrophes: Leveraging Abiotic Constraints through Environment-Driven Evolutionary Selection

Kevin Godin-Dubois, Sylvain Cussat-Blanc, Yves Duthen

► To cite this version:

Kevin Godin-Dubois, Sylvain Cussat-Blanc, Yves Duthen. Beneficial Catastrophes: Leveraging Abiotic Constraints through Environment-Driven Evolutionary Selection. 2020 IEEE Symposium Series on Computational Intelligence (SSCI), Dec 2020, Canberra, Canada. pp.94-101, <10.1109/SSCI47803.2020.9308411>. <hal-04793535>

HAL Id: hal-04793535

<https://hal.science/hal-04793535v1>

Submitted on 20 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Beneficial Catastrophes: Leveraging Abiotic Constraints through Environment-Driven Evolutionary Selection

Kevin Godin-Dubois, Sylvain Cussat-Blanc, Yves Duthen

University of Toulouse

IRIT - CNRS UMR 5505

31042 Toulouse, France

{kevin.dubois, sylvain.cussat-blanc, yves.duthen}@irit.fr

Abstract—Decades worth of experiments on the generation of virtual creatures have led to a large corpus of morphological and behavioral controllers. Yet despite these extensive researches, the role of the environment in the emergence of complexity is often experimenter-dependent. In this work, we present a framework in which environmental variables are the sole control mechanism for directing the evolution of an autonomous population of artificial creatures. The Environment-Driven Evolutionary Selection (EDeNS) algorithm automatically and simultaneously explores multiple alternatives, initially identical but with potentially differing dynamics with respect to their abiotic constraints. By exposing populations to different sets of constraints, they are subjected to divergent fitness functions both across the simulation space and between successive elementary steps. This framework is applied on a system composed of artificial plants autonomously reproducing in a 2D environment subjected to three factors: topography, temperature and hygrometry. We show that some of the populations obtained via EDeNS exhibit increased capabilities to invade foreign environments over populations evolved in a hospitable, constant environment. Moreover, the data thus collected highlighted two fundamental advantages of the automated exploration of abiotic constraints: the positive effects of catastrophic trimming and the unbiased designing of dynamical environments.

Index Terms—Evolutionary computing, Algorithm, Dynamic environments, Complex systems

I. CONTEXT

The generation of virtual creatures has received much attention in the late decades both in the context of isolated individuals and ecosystems. However, much of the focus in these works has been put on the capabilities of the individuals and the means by which they could perform specific tasks. In these cases, the role of the creature’s environment is often that of a physical container which stores uniformly available resources. Similarly, especially in the case of single-individual evolution where simulation time is limited, these environments are generally static.

For instance the seminal work of [15], in which an individual’s surroundings were solely comprised of its opponent, the target cube and the ground, was expended upon in [11] by evaluating similar creatures in more complex contexts. There, the authors show the stair-climbing, gap-crossing and even skating capabilities of their resulting creatures. As can

be expected, these exhibit more complex behavior than their ancestral counterparts by being evolved in a more demanding context, even though the resulting set of abilities is explicitly specified by the experimenters. Similar remarks can be made for the creatures developed in [3], which also manage to climb stairs, or the plant-looking demeanor of [16]’s Vascular Morphogenesis Controller.

When considering ecosystems of vegetal creatures, instances of more heterogeneous conditions can be found as in the work of [4] where the use of non-flat topography increased the inter-individual competition by providing a built-in bias. Additionally, the presence of varying types of constraints in [1] was shown to reproduce a distribution of survival strategy with similar characteristics to that of natural plants. In a similar fashion, dynamical conditions were experimented upon in [5] with random directions and intensity of wind promoting the emergence of robust morphologies in the evolution of a forest-like ecosystem.

Digital ecosystems are one of the area with the highest emphasis on varying environments, most notably with the Avida platform. The main advantage of such systems is that their lightweights allows for the monitoring of population-level trends over sufficiently large amounts of generations. Such a sample of individuals makes it possible to study the effect of dynamical constraints on the evolutionary process as was done in [2], [10], where different sets of rewards were tested, and [12] which introduced cataclysmic episodes. In particular, these works have highlighted the benefits one can draw from changing conditions namely increased robustness and valley-crossing [14].

In this paper, we describe a framework which not only alleviates the need for a human designer to devise fair but challenging environments for their creatures to evolve in, but also allows for a broader exploration of the set of viable ecosystems. Human interaction is only required at the initial stages to devise appropriate boundaries for the environmental dynamics. This is obtained by relying on a population-based evaluation coupled with an environment-centered selection as detailed in the following section.

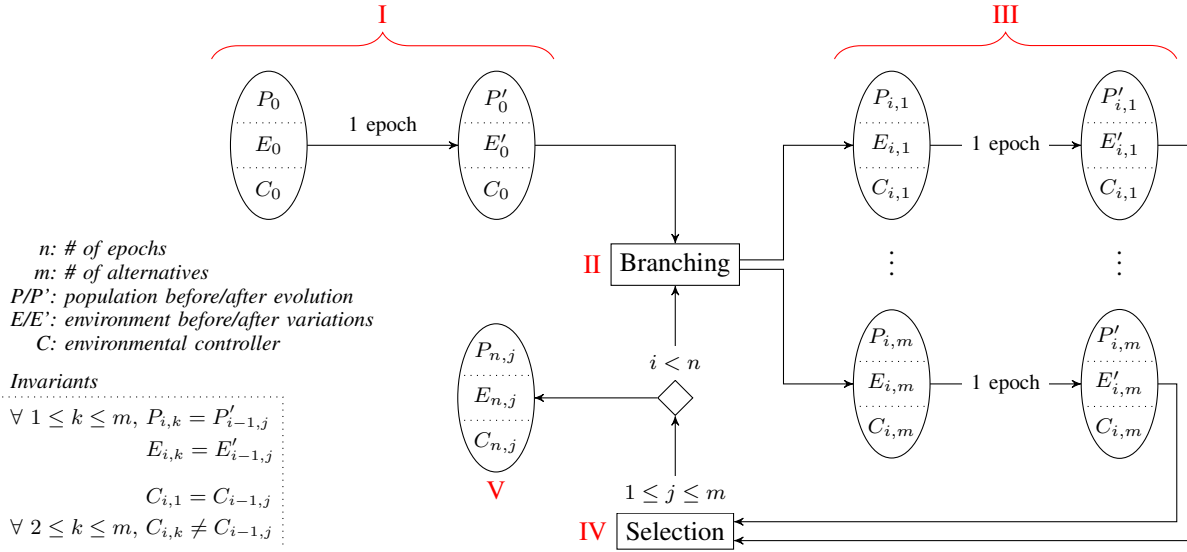


Fig. 1: Methodology of the Environment-Driven Evolutionary Selection. Major steps (in red) are detailed below.

II. EDENS FRAMEWORK

In order to explore long evolutionary periods through varying environments, traditional methods of optimization or open-ended evolution were deemed ill-suited to the task. Indeed, the objective is neither to produce a genotype perfectly adapted to a set of conditions, which by definition remain fluid, nor to subject a single population to dynamical constraints, which would more often than not result in extinctions.

To tackle this problem, we designed a framework inspired by the $1 + \lambda$ evolutionary strategy in which the concurrent evaluations can be seen as *alternative* futures, given an identical starting point. Thus, to reach a reasonably large number of generations, we approximate a single, continuous evolution by a succession of smaller simulations as summarized by figure 1. Each so-called epoch allows the populations to adapt and diversify, while the Selection-Branching bottleneck provides direction for the evolutionary process. In this instance, $n = 250$ epochs, each of length $e = 4$ years, are performed with a branching factor $m = 10$.

Step I, Initialisation

The first ecosystem is seeded with a mostly uniform population and a neutral environment i.e. with average figures across all its dimensions. The controller is also neutral ensuring constant conditions for this phase. The individuals are left free to colonize the area and reach a viable state.

Step II, Branching

The branching procedure uses as input a given ecosystem (resulting from Step I or IV) and produces m alternatives. The contents of the provided simulation are deep-copied as-is into each alternative. The environmental controller, however, is subjected to mutation, except in the case of the first alternative ($C_{i,1}$ in the figure). This implies that the difference between populations at the beginning and end of an epoch only result

from differences in the dynamics they were subjected to. Keeping the first alternative as-is also allows for reduced risks of suffering a complete extinction as, given that this alternative was viable at the previous epoch, it should still be.

Step III, Evolution

Each alternative is allowed to run for a full epoch without external feedback of any kind. Evolution is thus only guided by the combination of fitness functions resulting from the interaction between the individuals and the environmental dynamics. Indeed, as we will see in the following section, the state of the abiotic component has profound impacts on the metabolism and reproductive capabilities of our plants thus changing conditions imposes changes in their survival strategies.

Step IV, Selection

After all m alternatives have run their full course, they are evaluated by the external (human-designed) set of fitness functions. The resulting ecosystems are thus selected based on the state they have reached thanks to the *indirect* control of the environment. The ecosystem thus singled out is designed as the active alternative and is used to generate the next batch as seen in Step II.

Step V, Termination

After $n - 1$ such epochs have been performed, the algorithm terminates and the final active alternative is obtained. From this point, one has access not only to a population which has endured for a large amount of time in the face of changing environmental conditions but also to the complete timeline of its evolution. Indeed by unrolling the succession of selected alternatives, one can build a continuous set of “realities” which provide the complete historical background of this evolution, including the response to environmental perturbation and the specific sets of favored or shunned conditions.

III. IMPLEMENTATION ON ARTIFICIAL PLANTS

The system on which we deployed EDEnS is identical to that described in details in [8], a brief overview of which is given below, including this contribution's extensions.¹

Individuals are virtual plants whose body plans are generated through a step-wise L-System: rules are applied one at a time, depending on resource levels and available space. Their metabolism requires water uptake and photosynthesis to produce glucose which, in addition to water, feeds both the growth process and the reproduction-oriented organs. Indeed, thanks to a Bail-Out Crossover [7], individuals have a genetic control over which other individuals they want to mate with, thus allowing species barriers to spontaneously emerge and limiting the risks of nonviable offspring. Reproductive maturity is reached when a flower organ accumulates a sufficient amount of biomass, upon which potential mates are searched in the surrounding area. Successful mating results in the replacement of the corresponding pistil by a fruit containing the result of multiple independent crossover operations.

These creatures inhabit a 100-meters wide environment, partitioned into 1m patches which have local values for each of the three studied variables: topography t , temperature h and hygrometry w . These impact the behavior and welfare of plants in a number of ways with the former mainly acting as a reproduction and migration inhibitor. Water availability has obvious consequences on the plants' metabolism while external heat acts on uptake efficiency, evapo-transpiration and wastes production. Baseline conditions correspond to a flat ground with a hygrometry of 0.5 and a temperature of 10°C.

A. Environmental controller

The basis of the framework presented in this work is the use of automated evolution of environmental constraints and, as such, the hand-written equations used in the previous setting cannot be applied here. They have been replaced by a Cartesian Genetic Program (CGP) which was shown to be a powerful tool for generating complex functions [13].

For each patch in the environment, the CGP is fed its normalized abscissa x , current value for the three variables (t, h, w) and a sinusoidal input D of 1-year period thus designed to prevent abrupt transition for one year to another. The outputs correspond to the theoretical values for the same environmental variables at the next time step with two additional points requiring detailing. First, topography is only updated once per day (i.e. every 10 steps) to reduce the computational cost resulting from updating each plants' altitude. Second, in order to prevent excessively brutal transitions (e.g. from 40°C to -20°C in 1/10th of a day), an inertia factor $\alpha = 0.95$ is used to transform the theoretical output \hat{v}_{t+1} into the actual value v_{t+1} as given by $v_{t+1} = \alpha v_t + (1 - \alpha) \hat{v}_{t+1}$. Such a smoothing operator is of particular interest when the controllers use functions with rapidly changing values such as `rand` or `step`.

¹The code for this experiment is available at <https://github.com/kgd-al/ReusWorld/releases/tag/alife2020>.

Randomly generated CGPs were composed of 100 internal nodes and designed to produce 0 across all outputs until the first mutation operator is applied, thus guarantying initially hospitable conditions for the primordial plants. While this implied artificially setting one internal node to the `zero` function and connecting it to the three outputs, all other 99 internal nodes were left untouched thereby retaining a large variability in the subsequently available dynamics.

Mutations are performed according to the *Accumulation* operator defined in [9] where point-mutations are iteratively effected until one which has an impact on the network's output is found. With such a difference in mutation rates between expressed and reserve DNA, unused portions of the genome are explored much more quickly. This bias towards mutation of reserve DNA was deemed a necessity when considering that, initially, it comprises more than 98% of mutable values.

B. Fitness functions

Step IV of the EDEnS algorithm (selection) is performed through a set of fitness functions from which a pareto front is derived. These can be classified into two categories: the functions maintaining acceptable conditions and those aiming at specific features.

Two such control functions were used: C_P that rewards population size when inside [500;2500] (a fair density given the environment's width) and C_T that penalizes simulation time so that computationally heavy alternatives are naturally trimmed out. These, however, can be seen as safeguards against the vagaries of evolution as they only sustain the process and do not contribute to it.

The other two are focused on promoting the emergence of particular patterns. The first, F_S , uses the species information extracted by APOGeT (described in [6]), to promote speciation patterns. This is achieved by favoring large variance in the genetic distance between individuals from different evolutionary niches. The second, F_D operates at the allelic level by rewarding alternatives where the genepool² is the most different between the beginning and end of a given epoch. This acts as a disturbance factor, both promoting optimization in the current direction and exploration of different regions of the genetic space.

The selected alternative is obtained by random picking one item from the pareto front obtained after evaluating all m alternatives.

IV. EXPERIMENT ON INVASIVE CAPABILITIES

To determine whether individuals resulting from EDEnS gained an advantage in terms of invasive capabilities, we compared the performances of a group e of 10 replicates against a smaller group c of 5 controls. The populations of e were obtained with the same initial parameters, with the exception of the random number generator used to control the selection (Step IV). The control group c results from evolutions of identical duration (1K years) in the constant and hospitable initial environment.

²The set of all alleles. Decimal values are first quantized

Environment	Temperature	Hygrometry
Neutral	10°C	.5
Hot	40°C	-
Cold	-20°C	-
Dry(1/4)	-	.25
Dry(1/8)	-	.125
Periodic(1/2)	.5D	-
Periodic(3/4)	.75D	-
Periodic(1)	D	-

TABLE I: Environmental variables for the different tests. Unspecified values are identical to the Neutral state

For all evolutions, we define c_i and e_i as the final ecosystem produced by the i^{th} run of the control and evolved group, respectively. When the distinction becomes necessary, the exponents p , e , c denote the population, environment (variables) and controller, respectively. To determine whether or not the use of EDEnS resulted in “better” populations we subjected each runs to multiple types of evaluation falling into two categories: Plants-versus-Environment (PvE) and Plants-versus-Plants (PvP).

A. Plants versus Environments

In this context, “abiotic robustness” is evaluated as the capacity to survive, or even thrive, in a multitude of abiotic constraints. The test environment is composed of the ecosystem under investigation side-by-side with an empty environment, controlled by trivial, hand-crafted equations. Three dimensions of variations have been used, in addition to the neutral environment, as summarized by table I. While the Neutral state corresponds to the same conditions as those under which the control group has been evolved, we also explore resilience to different sets of temperatures (Hot and Cold), of water availability (Dry) and of seasonal patterns of different amplitude using the built-in 1-year periodicity of input D (Periodic). Additionally, once the simulation is populated, all plants are translated so that the right most coordinate of the population’s bounding box is placed exactly at the coordinate 0 (i.e. the starting line) thus ensuring fair initial conditions.

The objective, in this evaluation, is for a plant to “reach” the right-most environmental boundary which can only be done through successive generations colonizing, and surviving in, the foreign conditions. In order to prevent needlessly long evaluation times, we provide an external source of stress in the form of an increasingly large area utterly inimical to life. Starting at the left edge, this region is increased by 5% of the total range every year. In addition to reducing computational costs, this ensures that we measure the adaptive capacities of the population, as this does not leave enough time for any meaningful evolution to occur. The score derived from such an evaluation is thus the normalized distance covered on the foreign soil with 1 indicating that the goal has been reached.

As can be gathered from figure 2, the control (on the left) and evolved groups showed diverging sets of capabilities. While, in c , colonizing the familiar Neutral environment is trivially done by all but c_1 it is almost the full extent of their




Fig. 2: Synthesis of the scores obtained by each population on the PvE evaluation. The area indicates the portion of foreign environment covered by the end of the allotted time. Asterisks indicate cases of complete colonizations.

resiliency with only two success under different abiotic constraints. On the other hand, e showed an altogether different trend: one half exhibits fair adaptive capacity while the other fails to reach the goal in all test environments. Additionally, one can see the peculiar case of e_8 which, being embedded in an environment with a hostile right-most region, cannot compete in any meaningful way in this particular evaluation.

Nonetheless, despite this broad variability of abiotic robustness in the evolved group, it was found to out-perform the control group (Mann-Whitney, p-value < 0.05, outlier e_8 was not considered).

B. Plants versus Plants

Another facet we investigated is that of “biotic robustness”, that is the capacity for a population to out-perform a wide variety of peers. With this in mind, we subjected each population to pair-wise evaluations in two types of scenarios: Merged and Contiguous.

In the first type, both competing populations A^p and B^p are placed in the same environment, the contents of which are defined by A^e and A^c , for a maximal duration of 104 years. In practice, this implies that the foreign population B^p is tasked with surviving in the local conditions of A^p which should prove to be difficult given the obvious asymmetrical bias in favor of the former. The second type involves placing both A and B ecosystems next to each other and normalizing all topographies to zero so as to ensure no impassible boundaries exist. In this case both populations can persist in their own local niche with those exhibiting sufficiently efficient invasive behavior potentially colonizing the other population’s domain.

To determine which population is victorious we define, for each individual, a coefficient τ which accounts for its population of origin. That is individuals from A^p and B^p are arbitrarily assigned $\tau = 1$ and $\tau = 0$, respectively. Given that, in this model, plants can hybridize with members of a different species, both population could merge back into




Fig. 3: Aggregated scores in both types of PvP evaluation. One can see that the control groups fares better when out-reproducing the competition is a viable solution while the evolved group shows a stronger capacity for invasion.

a single genepool. To keep track of the most contributing original population, the τ parameter of an offspring is defined by the average of that of its parents. The raw score of biotic robustness, for population A^p , given a final population P , is thus given by:

$$S = \frac{1}{|P|} \sum_{p \in P} \tau_p \quad (1)$$

This computation, however, does not take into account the initial distribution of population and thus a normalized version is used which, for an initial proportion S_0 , a final score S_f and a difference $d = S_f - S_0$, is given by:

$$\hat{S} = \begin{cases} \frac{d}{1-S_0} & \text{if } d \geq 0 \\ \frac{d}{S_0} & \text{if } d < 0 \end{cases} \quad (2)$$

This provides a more intelligible interpretation of the relative capabilities of both population with a score of +100% indicating a perfect colonization and -100% a perfect defeat.

The aggregated results (fig. 3) once more show diverging behavioral trends, although less strongly than in the PvE case. Indeed this type of evaluation essentially being a zero-sum game, these histograms highlight the relative strength of each population. At first glance one can see that runs from the control group are more suited to the Merged type of PvP evaluation, in which the inter-plant competition is most intense. Indeed, in this scenario, being able to out-reproduce the invading population is a straightforward strategy for which member of c have optimized because of the constant conditions they were evolved in. To test which group fares best, we performed Wilcoxon tests on the respective rates of success which confirmed the advantage of c over e (p-value < 0.05).

Conversely, results are more mitigated in the case of the Contiguous evaluation, due to the need of competing against both the foreign population and its environment. As a matter of fact, while there is a clear champion in the form of e_5 , no statistically significant difference was found between the performances of both groups.

V. DETAILS OF INDIVIDUAL STRATEGIES

In order to gain a better understanding of the kind of events that produced and characterize the individuals from e , we briefly detail, in this section, the specifics of one strongly performing ecosystem, e_5 , and of the least successful e_4 .

A. The case of e_5

This particular population showed great proficiency in all evaluation types, indeed being the champion in the PvE and Contiguous PvP evaluations. In figure 4, some of the most important dimensions of e_5 's evolution are presented.

Panel 4a consists of the plant population where we can see fluctuations of varying intensities, the largest being labeled above the graph. The bottom row comprises the per-patch values of the environment in which e_5 evolved with the topographical variations on the left and the temperature dynamics on the right. Sub-figure 4b shows the evolution of the genetic field controlling the optimal temperature, each slice corresponding to the distribution of the population's alleles at the end of the given epoch.

By observing all four pictures simultaneously, we can see the impact of environmental variations on the dynamics of the population, both at the individual and genetic levels. Indeed, the large dips in population size are quickly compensated by a renewal of reproduction through which newly available areas are recolonized. This however, is not without impact on other portions of the simulation such as the explored regions of the genetic space. In the case of the optimal temperature, one can see that each such transitions results either in tightening of allelic diversity (markers A,B,C) or in a sudden transition from one region to another (D).

As can be seen in the environmental dynamics, this is caused by variations in the plants' external conditions: A and B result from a transition from a smooth topography to a random one, C from a seasonal pattern of temperature to desert-like conditions and D from a two-fold variation where heterogeneous, albeit constant, heat levels further diverged before "switching".




Fig. 4: Dynamics of evolved best (e_5). a) variations in plant population. b) temporal histograms of allelic density. c,d) Per-patch variations of topography and temperature, respectively.

Thus in the case of e_5 , catastrophic environmental variations induced *beneficial* trimming in the population by selecting those which exhibited the most robust behavior.

B. The case of e_4

By contrast, the evolutionary trajectory of e_4 was found different on two major points: population size and environmental harshness. With respect to the former, we observed that the average size is about a third that of e_5 , which given the well known effect of population size on genetic diversity and competitive success, can go some way to explaining this run's poorer performances.

On the latter point, we noted that environment-induced stress was of a lower magnitude: while some epochs contain random dynamics or even harsh transitions, these are much sparser with larger periods of relative calm. For instance, during the first 400 years, e_4 has only encountered two different patterns of temperature whereas, by this time, e_5 is closer to twenty. Thus, in this specific instance, we can note that poor results are observed alongside small population subjected to lenient conditions.

VI. ANALYSIS OF EVOLUTIONARY DYNAMICS

We have seen that populations from the evolved group exhibited greater cross-environment resilience than those obtained in the control setting and observed in more details the specifics of the runs. These however do not provide the large-scale viewpoint required to determine what advantages the use of EDeNS brought. Thus the following section is devoted to investigating the evolutionary dynamics, notably with respect to the different types of robustness.

A. Population-level observations

First and foremost we observe a positive correlation between all three types of evaluation³. Summarized in table II, are the Spearman correlations for the variables with the most salient results where SPF is the number of seed per fruit (i.e. litter size), σ_T the acceptable temperature range and \bar{P}_l the average population count. Comparisons can be made between the trends obtained when considering all 15 evolutions as a

³0.94 for both PvP. 0.64 and 0.71 for PvE versus Merged and Contiguous, respectively. All p-values < 0.05.

Variable	<i>c</i> & <i>e</i>			<i>e</i> only		
	PvE	PvP		PvE	PvP	
		M	C		M	C
<i>SPF</i>	-.682	-.314	-.496	-.782	-.782	-.867
σ_T	.457	.043	.386	.782	.745	.709
\bar{P}_l	.696	.589	.421	.697	.648	.697
P_r	.757	.625	.682	.685	.697	.733
P_r^-	.793	.589	.679	.770	.661	.733
\bar{P}_r	.564	.125	.511	.891	.758	.794
\bar{P}_r^-	.529	.193	.571	.782	.927	.927

TABLE II: Spearman correlations between scores and population variables. Grayed-out values are not significant under a p-value < 0.05 threshold.

single group or when only focusing on the population resulting from EDnS, in all scenarios.

For instance we can see a fairly high, negative, correlation between the litter size and reproductive success but only inside the evolved group *e*. This trend is most notable in the Contiguous evaluation where high birth rates are not the privileged mechanism by which colonization is achieved. Similarly, we can observe that the introduction of temperature-related stress induced the expected response of increasing the range of tolerable temperatures, although not too strongly due to the adverse effects produced by large values. We can also formally observe that the assertion previously made on population size \bar{P}_l was found correct although less strongly so in the case of the Merged naturalisation.

The lower portion of the table is concerned with measuring the effect of the so-called perturbations. That is, given $P_l(t)$ the population size at time *t*, we define:

$$d_e = P_l(4(e - \nu)) - P_l(4(e + \nu)) \quad P_r = \sum_{e=1}^{249} \frac{|d_e|}{249}$$

$$d_e^- = \begin{cases} d_e & \text{if } d_e < 0 \\ 0 & \text{otherwise} \end{cases} \quad P_r^- = -\sum_{e=1}^{249} \frac{d_e^-}{249}$$

$$\bar{P}_r = \frac{P_r}{\bar{P}_l} \quad \bar{P}_r^- = \frac{P_r^-}{\bar{P}_l}$$

with $\nu = 0.25$, i.e. a quarter of a year. The first metric, P_r , is the most straightforward: it computes the average variation of population size around each epoch. For members of *e*, this gives a notion of how stressful their successive environments where perceived. We also apply it to members of *c* to define the level of “background noise” which is to say the amount of perturbation measured when none should be expected. The other three metrics are variations around P_r , with P_r^- only measuring *negative* perturbations that is those in which the population size was drastically reduced⁴. \bar{P}_r and \bar{P}_r^- are normalized by \bar{P}_l to untangle the effect of perturbations themselves with that of population size.

⁴ P_r^+ was also tested and only found mildly significant.




Fig. 5: Example of extreme variations. The second epoch imposes maximal-range temperature alternance which hinders, but does not destroy, the population.

Though we obtain fairly high correlations for the two non-normalized metrics in all instances, results are more telling when relying on their normalized counter-parts. Indeed we have already seen that population size has an effect on both robustnesses, thus partially explaining the mild results. When decoupled from such an effect, however, we see a very strong response: 0.89 for normalized perturbations and 0.92 for normalized catastrophes when linking with the PvE and PvP scores, respectively.

This implies that the most successful populations were also those that were subjected to strong cataclysmic transitions, a result in concordance with the observations made on e_5 .

B. Resilience and Creativity

Given that, every epoch, the plant populations strive to reach the optimal phenotype for the set of environmental constraints they are subjected to, even infrequent, mild variations of controller could be enough to disturb the fragile balance thus reached. However, when performing a classification of the different types of environments encountered by all runs (across the main timeline) we observed that neutral environments were a relatively rare occurrence (11.5%, 27% and 18.7% for the topography, temperature and hygrometry, respectively). No link was found between the amount of neutral states and scores, thereby showing that poorly fairing runs were not doing so for lack of environment-induced stress.

In fact the opposite occurred: 3/4 of the strongest perturbations (the top 5%) were catastrophic trimmings. While some of these transitions took very simple forms such as mild temperatures suddenly raising all the way up to 40°C, others were much more intricate. For instance, consider figure 5 where the temperature initially varies between -20 and -10°C with a half-year period. The change in controller results in a much wider amplitude which leads to a near-extinction, partially compensated by the increase in available water.

Another area in which the evolutionary process proved surprisingly resourceful is that of the exploration of counter-intuitive equation sets. Indeed, while it might be reasonably expected that smooth, periodic functions would result in favorable environmental dynamics, more exotic functions have nonetheless been successfully used throughout these runs. Consider the example in figure 6: one can see that the initially continuous conditions are replaced by a chaotic alternation of values induced by a massive use of the `rand` function.




Fig. 6: Creative compositions of the `rand` function. All three environmental variables depend on unpredictable dynamics resulting in a massive, albeit not fatal, drop in population.

Formally the equation set for the second epoch is:

$$\hat{t} = \text{rand}; \quad \hat{h} = x(\text{rand} < x); \quad \hat{w} = h - t$$

where x is the normalized patch abscissa ($\in [0, 1]$). Thanks to the mediating effect of the inertia factor α , these variations are not instantaneous by nonetheless result in unpredictable conditions to which the plants needs to adapt. This high level of stress has a very marked impact upon the population size with a survival rate of only 7.5%. However, the selection process, as defined in this implementation of the EDEnS algorithm, rejects un-populated alternatives. It is thus possible to explore such sets of constraints, with potentially catastrophic effects, with only limited risk of reaching a complete extinction (i.e. across all alternatives).

VII. CONCLUSION

In this article we introduced a novel paradigm in which the evolution of a population is indirectly controlled by external sources of stress in a manner hand-crafted environments would struggle to achieve. The EDEnS framework produced populations that fared better, in terms of foreign environments colonization, than those of the control group.

By exploring the dynamics of successful communities we observed that catastrophic events resulted in evolutionary transitions which allowed the emergence of more robust demeanor. Indeed the bottleneck effect, strongly identified for instance with e_5 , allowed selective trimming of less resilient variants, thereby freeing space for their more robust alternatives. It remains unclear, however, whether these populations were robust because of the cataclysms or if these cataclysms were selected because the populations could cope with them. Further studies are thus warranted to determine whether the causality is in the expected direction (i.e. adversity causing resilience).

As a straightforward extension of this work, we can note that no mention of the phylogenetic dynamics were made due to a striking lack of speciation. Direct methodologies could be devised for favoring the emergence and maintaining of multiple, potentially cohabiting species. This could be addressed either by providing a larger environment, thus allowing different regions to be differently inhabited, at the price of more computationally costly simulations. Another direction would be to artificially create isolated “islands” thus effectively enforcing impassible geological barriers. An additional CGP could be used to control the migration rates

between different such islands, the pattern of which would allow insights into the mechanisms of similar phenomenon in biological life.

Additionally, given the loose requirements of this framework (self-sustaining populations, mutable environmental controller and a set of objectives) deployment on a broad range of problems is possible. One direct application, currently being investigated, is the manner in which predation could emerge in land-based motile creatures in response e.g. to famine. Another direction would be to use EDEnS as an interpolator: given an initial and target ecosystems, one could explore the succession of evolutionary events that led from the former to the latter, both in artificial and biological life.

REFERENCES

- [1] S. Bornhofen, S. Barot, and C. Lattaud. The evolution of CSR life-history strategies in a plant model with explicit physiology and architecture. *Ecological Modelling*, 222(1):1–10, January 2011.
- [2] Rosangela Canino-Koning, Michael J Wiser, and Charles Ofria. The Evolution of Evolvability : Changing Environments Promote Rapid Adaptation in Digital Organisms. *Proceedings of the European Conference on Artificial Life*, pages 268–275, 2016.
- [3] René Doursat and Carlos Sánchez. Growing Fine-Grained Multicellular Robots. *Soft Robotics*, 1(2):110–121, June 2014.
- [4] Marc Ebner. Evolution and Growth of Virtual Plants. In *Advances in Artificial Life. 7th European Conference on Artificial Life*, volume 2801, pages 228–237. 2003.
- [5] Christophe Eloy, Meriem Fournier, André Lacointe, and Bruno Moulia. Wind loads and competition for light sculpt trees into self-similar structures. *Nature Communications*, 8(1):1014, December 2017.
- [6] Kevin Godin-Dubois, Sylvain Cussat-Blanc, and Yves Duthen. APOGeT: Automated Phylogeny Over Geological Timescales. In *ALIFE 2019 (MethAL workshop)*, page in press, 2019.
- [7] Kevin Godin-Dubois, Sylvain Cussat-Blanc, and Yves Duthen. Self-sustainability Challenges of Plants Colonization Strategies in Virtual 3D Environments. In Paul Kaufmann and Pedro A Castillo, editors, *Applications of Evolutionary Computation*, pages 377–392. Springer International Publishing, Cham, 2019.
- [8] Kevin Godin-Dubois, Sylvain Cussat-Blanc, and Yves Duthen. Speciation under Changing Environments. In *ALIFE 19*, volume 31, pages 349–356, Cambridge, MA, 2019. MIT Press.
- [9] Brian W. Goldman and William F. Punch. Reducing Wasted Evaluations in Cartesian Genetic Programming. pages 61–72. 2013.
- [10] Alexander Lalejini, Michael J. Wiser, and Charles Ofria. Gene duplications drive the evolution of complex traits and regulation. In *Proceedings of the 14th European Conference on Artificial Life ECAL 2017*, pages 257–264, Cambridge, MA, September 2017. MIT Press.
- [11] Nicolas Lassabe, Herve Luga, and Yves Duthen. A New Step for Artificial Creatures. In *2007 IEEE Symposium on Artificial Life*, pages 243–250. IEEE, April 2007.
- [12] Tian-tong Luo, Lise Heier, Zaki Ahmad Khan, Faraz Hasan, Trond Reitan, Abdool S. Yasseen, Zi-xuan Xie, Jian-long Zhu, and Gabriel Yedid. Examining Community Stability in the Face of Mass Extinction in Communities of Digital Organisms. *Artificial Life*, 24(4):250–276, March 2019.
- [13] Julian F. Miller. Cartesian Genetic Programming. In *Natural Computing Series*, volume 43, pages 17–34. 2011.
- [14] Joshua R. Nahum, Jevin West, Benjamin M. Althouse, Luis Zaman, Charles Ofria, and Benjamin Kerr. Improved adaptation in exogenously and endogenously changing environments. In *Proceedings of the 14th European Conference on Artificial Life ECAL 2017*, pages 306–313, Cambridge, MA, September 2017. MIT Press.
- [15] Karl Sims. Evolving 3D Morphology and Behavior by Competition. *Artificial Life*, 1(4):353–372, 1994.
- [16] Payam Zahadat, Daniel Nicolas Hofstadler, and Thomas Schmickl. Vascular morphogenesis controller: A Generative Model For Developing Morphology. In *Proceedings of the Genetic and Evolutionary Computation Conference on - GECCO '17*, pages 163–170, New York, New York, USA, 2017. ACM Press.