

Evosphere: evolutionary dynamics in a population of fighting virtual creatures

Thomas Miconi

Abstract—It is often suggested that traditional models of artificial evolution, based on explicit, human-defined fitness functions, are fundamentally more restricted and less creative than natural evolution, in which no such constraint exists. After a discussion and refinement of this statement, we suggest a classification of evolutionary systems according to their evolutionary “creativity”. We describe an environment, called Evosphere, in which a population of 3D creatures interact, fight with each other, and evolve freely on the surface of a “micro-planet”. We demonstrate the onset of natural selection and adaptive evolution within this virtual world, both by visual inspection and statistical analysis. We show that the introduction of reproductively isolated species enriches the dynamics of the system, leading to simple evolutionary feedbacks among species.

I. INTRODUCTION

A. Natural and artificial selection

Artificial evolution usually involves a selection process based on fixed, human-defined fitness functions. There is much speculation that the use of a pre-defined fitness function constrains the evolutionary process into a predictable pathway, thereby curtailing the creativity of evolution as it occurs in nature. In this line of reasoning, evolution based on human-defined fitness functions is sometimes equated to the process of *artificial selection*, as applied by humans to crops and animals in order to favour certain desirable characteristics; this is contrasted with “real” natural selection, which is bound neither by our human interests, nor by our lack of insight or imagination. The conclusion is that only “free”, unguided natural selection can produce the dazzling amount of complexity and diversity evident in nature. For example, Ray [8] writes:

I suggest that the most likely way to achieve complexity increase in digital evolution is through evolution by natural selection in an ecological community. No attempt should be made to provide fitness functions, or artificial selection, to guide evolution towards useful products. Rather, evolution should be free to explore the possibilities without the burden of human “guidance”.

Traditionally we have managed evolution through manipulating selective forces. In this new approach, our role is to create the conditions for complexity increase, rather than trying to guide it through artificial selection.

Ray prudently points out that “this is an interesting scientific challenge, as the conditions that generate complexity increase are unknown.”

In a similar vein, Channon [3] writes (emphasis original):

In the context of evolutionary emergence, any artificial selection system used constitutes just one of the parts of a system. Artificial selection can only select for that which is specified to. Therefore anything [unexpected] that emerges during evolution must be due to another aspect of selection, which must in turn be due to the innate dynamics of the system - natural selection.

Packard [7] has attempted to make this distinction more explicit, by introducing the concepts of *intrinsic* and *extrinsic* adaptation (emphasis original):

A common approach in modelling evolutionary processes is to regard a member of a population as fit or unfit according to how it is evaluated by a fitness function. The fitness function is a map that assigns a real number (the fitness) to every possible member of the population, and it is generally specified as an a priori feature of the model...I will call any adaptive dynamics that uses such an a priori fitness function *extrinsic adaptation*...

The biosphere does not appear to have any *a priori* fitness function defined on a space of possible organisms; in fact, one of the most amazing features of biological evolution is that the biosphere evolves automatically, with each organism (or population of an organism) adapting to an environment made up of both external natural features and all other organisms it interacts with...I will call adaptation that occurs as a result of a population of subsystems changing in response to interactions between them (without an a priori fitness function) *intrinsic adaptation*.

B. Explicit and implicit fitness functions

There may be an objection to this view. In any evolutionary simulation, it is necessary that the survival and reproduction of agents should somehow differ, and be determined by specific rules (defined *a priori* by the human programmer). If reproduction and survival were equal among all individuals, or utterly random rather than dependent on certain pre-defined parameters and events, no evolution could occur in the first place. But a set of rules that determine which individuals survive and reproduce (or not) is precisely what we call a fitness function. This being so, how could we build a system “without any fitness function”?

The issue is resolved by making a distinction between *explicit* and *implicit* fitness function. It is true that the programmer must define the rules that determine when and

how an individual dies, or reproduce, or is born. Note that these rules can be very simple, such as (say) “reproduce if a certain neuron fires and there is any agent nearby”, or random elimination, or anything else. Together, these processes define the *explicit* fitness function - the narrow specification of which conditions will immediately cause death, or reproduction, or birth. However, this under-determines the question of *which* individuals will turn out to actually fulfil these conditions. This is decided by the dynamics of the process, which occur as a (potentially unpredictable) result of the rules of the system. If the fulfilment of these conditions depends on heritable characteristics of the individuals, then these characteristics will also be subject to adaptive evolution. For example, if the explicit fitness function states that reproduction occurs whenever a certain neuron fires and a partner is nearby, and if individuals have control over their position, and if this control has a genetic component, then there will be evolutionary pressures to optimise the location of individuals in such a way as to optimise reproduction; if energy is lost by firing neurons, the same may apply to the choice of when to fire the reproductive neuron, etc.

In other words, it is true that the immediate rules of birth and death must necessarily be defined by the programmer; but if the system is sufficiently complex, then these basic, low-level rules (the explicit fitness function) are only one component of the actual fitness landscape, that is, of the actual function that maps the characteristics of an individual to its contribution to future generations (its “fitness”). Out of the interaction between these explicit fitness rules, the genetic representation (what is under genetic control) and the dynamics of the system, an *implicit* fitness function may spontaneously emerge, that may not be immediately predictable from the explicit rules of the system.

C. Different types of implicit fitness functions

Note that even if an implicit fitness function does emerge that would not be directly deductible from the explicit rules, it is perfectly possible that this function should turn out to be utterly boring and static. That is, the fitness landscape that emerges from the interactions between the explicit reproduction and replacement mechanisms, and the dynamics of the system, may well turn out to be a simple, fixed fitness landscape. If such a thing happens, then all we have done is indirectly recreating a fixed (possibly multimodal) fitness function, that the system will duly optimise. While this may be of interest in itself, this is not what the previously quoted authors had in mind.

The other, more interesting possibility is that the emergent implicit fitness landscape could be *multimodal* and *dynamic*. That is, instead of a single, stable fitness function, the system could continuously generate many local fitness functions, constantly offering new niches to which some lineages could adapt. Of course, one simple way to obtain this is to impose time-varying external conditions on the population (e.g. changing climates, varying food locations, etc.) But the most interesting case is when these fitness functions are self-generated: they arise spontaneously from the interactions

between individuals, and change over time as individuals themselves change through evolution and thereby offer new adaptive opportunities for other individuals. In this case, the population constantly generates its own implicit fitness functions.

It is clear that we are actually talking about *coevolution*. Coevolution occurs when the fitness of individuals are affected by each other’s heritable characteristics, to the effect that their fitness landscapes are effectively coupled. This is clearly a prerequisite if implicit fitness functions are to emerge from the interactions of agents.

However, it is also important to note that while coevolution is necessary for the spontaneous emergence of self-generated fitness functions, the former does not always lead to the latter: it is perfectly possible to have a coevolutionary environment (in which the fitness of individuals is affected by each other’s heritable characteristics), and yet the implicit fitness function that emerges from these interactions turns out to be fixed and unimodal. In fact, in some situations, this is precisely what we want: for example, in coevolutionary optimisation (say, designing a chess player through coevolution), we use coevolution as a substitute for a non-existent fitness function that would capture the essence of “good play”. We are not interested in mutual adaptation or ecological effects - in fact, these are seen as a nuisance that harms progress in terms of general abilities.

But in the context of artificial ecologies, it might be preferable to observe dynamic fitness landscapes, in which various implicit fitness function would emerge *in response to each other*. That is, the evolving lineages would enter into a constant cascade of innovations, adaptations and counter-adaptations, perpetually generating (and filling) new niches and thereby offering new opportunities for others. Such processes, which are readily apparent in nature, seem to capture the intent of the authors quoted above. Through such a process, artificial evolution might attempt to imitate nature’s ability to generate “endless forms most beautiful” from “simple beginnings”.

This discussion suggests a natural classification of evolutionary environments, according to their degree of evolutionary “creativity”.

- 1) Level 0: Fitness function is entirely external, and no implicit fitness function emerges beyond the explicit fitness function. Coevolutionary effects play no significant role in the fitness landscape.
- 2) Level 1: The fitness landscape is affected by coevolutionary effects, but results in a single, stable implicit fitness function; if several peaks exist, they do not emerge in response to each other, but rather correspond to various independent, pre-existing optima of the

fitness landscape.¹

- Level 1b: Similar to Level 1, but the implicit fitness function that emerges is *unpredictable* and can vary from run to run.
- 3) Level 2: The fitness landscape is affected by coevolutionary effects, and results in several implicit fitness functions that emerge in response to each other. This process goes on for a finite period of time, after which the system settles in a fixed equilibrium in which no novel implicit fitness function appears. This may include constant cycling through previously found fitness function. In short, the total number of implicit fitness function generated (and optimised) by the system is bounded over time.
 - Level 2b: Similar to Level 2, but the implicit fitness functions that emerge are unpredictable and can vary from run to run.
 - 4) Level 3: The fitness landscape is coevolutionary and the system constantly generate new implicit fitness functions which emerge in response to each other. No equilibrium is found, and novel fitness functions are constantly being generated. This corresponds to a situation of truly “open-ended evolution”.

II. THE EVOSPHERE MODEL: FIGHTING CREATURES ON A MICRO-PLANET ENVIRONMENT

Evosphere is a system in which 3D creatures interact, reproduce and evolve in a physically realistic environment. As the name indicates, this environment is based on a spherical “micro-planet” upon which a population of creatures dwell freely.

The creatures used in Evosphere are similar to those described by Sims [10], [9], with several modifications. They are articulated structures made up of rigid blocks, joined by hinge joints, and controlled by neural networks (each block has its own local neurons, and neural connections occur both within and between blocks). They possess sensors for joint angle, contact, and the relative position and distance of other creatures. Both body and neural information are fully under genetic control. The main differences with Sims include the use of standard, McCulloch-Pitts neurons rather than the complex operator neurons used by Sims, as well as a more powerful developmental system. The details of these creatures have been described in a previous publication. [6]

In addition, we have implemented *physical combat* among virtual creatures: creatures are able to trade blows with each other, resulting in physical damage, very much in a natural way. This poses the difficulty that in a collision between two rigid blocks supposedly made of the same “stuff”, the symmetry of physics implies that both should suffer

¹To distinguish between levels 0 and 1, a simple test suggests itself: let us alter the system so that, at any time, only one individual is subject to evolution (that is, the evolving genome is expressed by only one individual within the system) while the rest of the population is composed of trivial, *fixed* individuals. If the results of evolution are the same in this altered system as in the original system, then it is clear that the trajectory of the system is not dependent on coevolutionary effects.

equivalent damages (“action equals reaction”). In nature, damage is not necessarily symmetrical due to the variety of materials and structures: for example, teeth penetrate flesh more easily than the other way round. However, transposing such variety in our system would pose several issues. We solve this problem by specifying that in a collision between two limbs, each limb inflicts a damage proportional to how much this limb was moving towards the other, that is, how much this limb contributed to the occurrence and importance of the collision. The effect is that the creature that deals the blow is rewarded over the one that receives it. This system is described in full detail in a related paper. [5]

Because our environment is expected to contain many free-living creatures, we would like it to be *open*, that is, without boundaries. The imposition of barriers could introduce unpredictable effects into the dynamics of the system. While certain systems (such as Polyworld [11]) choose to embrace the existence and effects of such barriers, others (such as Geb [2]) solve the problem by using a “toroidal” 2D environment where the top connects with the bottom and the left with the right. In a 3D environment, the simplest structure that dispenses with boundaries is obviously the sphere. We therefore choose to implement our environment in the form of a sphere, or “micro-planet”. Creatures dwell on the surface of this planet, which is sufficiently large to allow for the free movement of each creature.

III. EXPERIMENTS

A. The Rules

The environment described above is very general and many rules can be implemented to determine survival, elimination and replacement of individuals. Here we use the following, simple set of rules.

Every new creature is born with a certain amount of “energy” E_0 , which is a constant of the system. At every timestep, this energy is being decreased by a (small) amount δE_{time} . Additionally, energy is also lost through physical damage: if a creature suffers damage D , its energy is decreased by an amount δE_{damage} proportional to D . The parameters are chosen in such a way that physical contact is much more damaging than the passage of time.

When a creature’s energy falls to zero or under zero, it is removed from the simulation and immediately replaced with a new creature: the population size is therefore kept constant. How the new creature is generated depends on the manner in which the removed creature met its end. If the creature has not suffered any damage in its last timestep (that is, if its demise in this timestep is entirely due to δE_{time}), then it is replaced by randomly selecting two parents from the population, and generating a new individual through recombination and mutation of these parents’ genomes.

But if this creature has suffered some damage in its last timestep, that is, if it has been “killed”, then it is replaced with a mutated clone of the creature that has inflicted this damage (its “killer”). In the unlikely case that two or more creatures have inflicted damage upon the dying creature

during the same timestep, the one that inflicted the highest damage is deemed the “killer”.

There is the question of how to insert the new creature into the ongoing simulation, while minimising disturbances. We simply “drop” the new creature, from a fixed height, over a random point of the sphere. In addition, each new creature is granted an initial “grace delay” during which it cannot suffer or inflict damage. This delay is long enough to cover the fall to the ground and a short period beyond.

From these rules we can infer certain properties of the system:

- 1) The initial amount of energy E_0 sets a maximum lifetime (equal to $E_0/\delta E_{time}$) that no creature can exceed. This maximum is that of a creature that does not undergo any contact. Any action taken by the creature can only *reduce* this lifetime (through physical damage), or leave it unaffected.² Crucially, one cannot gain energy through combat: the system does not reduce to the “multi-way” coevolution of fighters.
- 2) A creature can favour its own reproduction in two ways: by killing others, or (to a lesser extent) by avoiding damage better than others. In the first case, it can generate clones. In the second case, it will have more chances than others to be selected as a random parent for dying creatures.
- 3) Due to the manner in which damage is apportioned, any contact is risky. When making contact with another creature, if the other creature’s limb is coming towards you, you will suffer damage.
- 4) A creature that inflicts much damage upon another will only benefit if this damage is deadly. If the creature goes on to live for just another timestep, and then dies, the aggressor will not benefit.

IV. RESULTS

A. Emergence of natural selection and evolution

Natural selection and evolution consistently arise in the system, as observed both by visual inspection of creatures and by numerical assessment of the system’s behaviour (see below). The system generated a wide range of morphologies.³ Though some broad categories may be found (“worms”, two-armed “crawlers”, etc.), some creatures were altogether quite bizarre in their construction. Figure 2 shows four pairs of creatures evolved from four different worlds.

The majority of behaviours were purely motile behaviours and did not make use of sensor input: these creatures evolved efficient movements for “hitting without being hit.” However, some creatures did include sensory input, including contact sensors and detection of other individuals. For example, the

²In particular, in the early stages when damage will be infrequent, it is expected that many poor creatures will reach the maximum lifetime - simply by standing still and not having any contacting neighbour.

³All individual creatures described in this chapter had either the highest or second-highest number of children of all creatures alive at the time of observation. In open evolution, maximum number of children is the closest approximation to the notion of “champion” that exists in experiments based on fitness function or head-to-head coevolution.

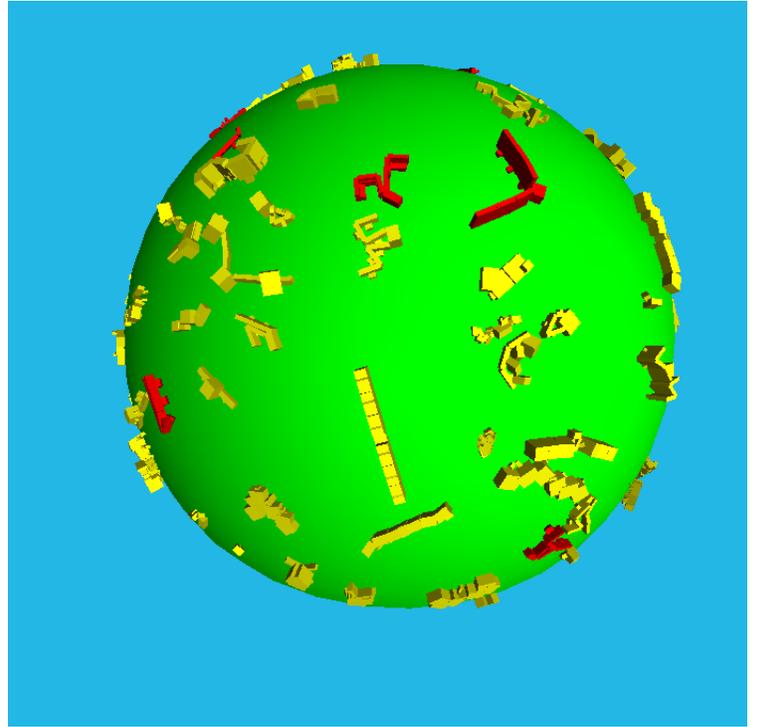


Fig. 1. Evosphere, as seen from “space”.

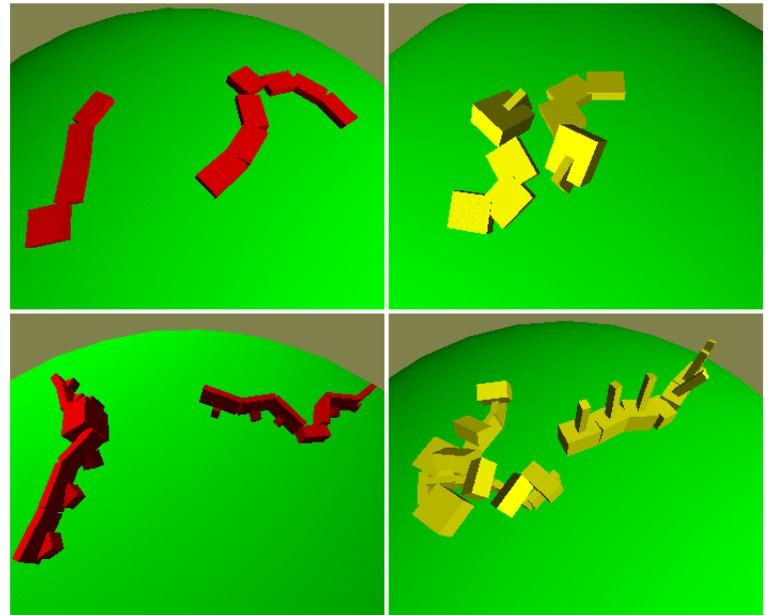


Fig. 2. Four pairs of creatures from various runs.

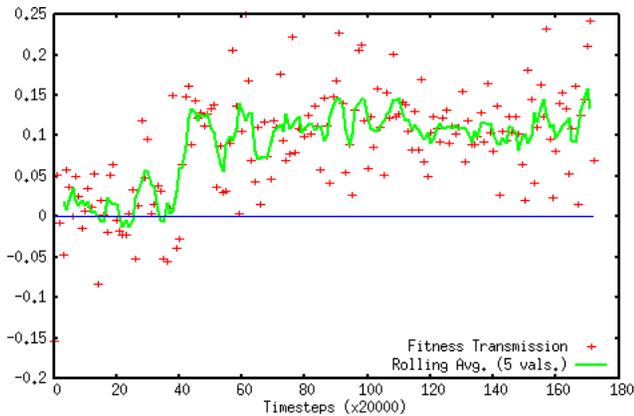


Fig. 3. Fitness transmission as a function of time, together with a rolling average over 5 values. An initial inactive phase suddenly gives way to a phase of consistently positive activity.

“crested” creature in the top-right corner of Figure 2 uses a sensor in its head limb to approach the closest other creature in a “swirling” motion (the creature moves by undulation).

In general, the dominant strategies were variants of a “roaming” behaviour: move fast in an undulating manner to hit whatever comes close.

B. Fitness transmission analysis

To confirm the onset of genuinely adaptive evolution by natural selection, we measure *fitness transmission* in the system. Fitness transmission, described in [4, Chap. 9], is a measure of evolutionary activity which seeks to address certain shortcomings of existing measures (such as Bedau-Packard statistics [1]). It essentially measures the statistical correlation between the fitness of parents and children, where “fitness” is estimated by the number of grandchildren (suitably normalised to remove dependencies).

Figure 3 shows a graph of fitness transmission as a function of time. A genealogic record of the system is divided into periods of 20000 timesteps (a rather arbitrary value corresponding to twice the maximum possible lifetime of a creature). The graph in this figure provides the fitness transmission calculated for each of these periods, as well as a rolling average over 5 successive values.

As indicated by this graph, the system usually goes through an early phase of very low or absent fitness transmission, denoting low or non-existent natural selection. In this phase, creatures are so ill-adapted that any damage they inflict or suffer arises mostly through chance rather than because of their characteristics. However, certain characteristics can be faintly advantageous, such as temporary motion as opposed to absolute immobility (depending on the presence of close neighbours).

This phase ends with the appearance of persistent, oscillating motion. As discussed in previous papers (e.g. [6]), persistent motion requires oscillating neural activation and therefore specific connection patterns. Oscillating limbs are highly advantageous for two reasons. First, they allow the

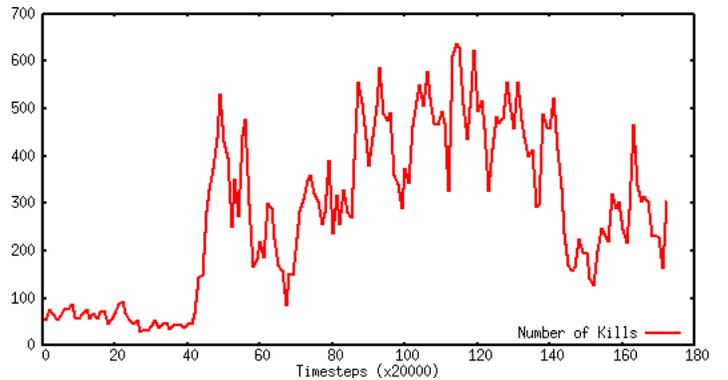


Fig. 4. Number of kills within the population of the same run as in Figure 3. A sudden rise in the number of kills coincides with the onset of positive fitness transmission.

creature to displace itself on the surface of the world, rather than remaining in its initial location. Second, constant motion of limbs implies constant possibility of inflicting damage upon other creatures. Once this innovation is introduced in the system, natural selection has something substantial to operate on (that is, fitness-impacting heritable characteristics), and Darwinian evolution begins in earnest. This is reflected in a sudden increase in the value of fitness transmission. Roaming behaviours quickly emerge, and thrive.

As an additional confirmation of fitness transmission analysis, we provide the number of killings per period of time (where the division in periods is identical to that used for fitness transmission). The onset of high fitness transmission occurs precisely at the time when the number of killings starts to increase significantly, indicating that new capacities are spreading throughout the population.

V. COEVOLUTION BETWEEN SPECIES

A. Panmictic models and diversity

While the previous experiments demonstrated the onset of adaptive evolution, the end result of this evolution tends to be predictable and uniform: large individuals equipped with oscillating limbs, roaming over the sphere and hitting anything they can. Also, while there is variability between runs, each run tends to converge towards one superior morphology, leaving little room for diversity. Thus our system (as it stands) seems to fall in class 1, according to the classification discussed above.

We note, however, that the rules for reproduction and replacement are not really favourable to diversity: we used a panmictic model, in which everyone can mate with everyone. Once a dominant strategy emerges, it can easily take over the population. Thus we ask the following question: is predictability of outcome a fundamental aspect of Evosphere, or is it simply an effect of panmictic reproduction? To answer this question, we eliminate the panmictic aspect of the system by introducing *species*, that is, reproductively isolated populations.

B. Species implementation

We want to implement species in such a way as to enforce genetic separation, but also leaving evolution as free as possible. In particular, we want species numbers to fluctuate over time, according to their respective success - to the point of going extinct if such is their fate.

To do this, we set a fixed number of species N_s . Each creature belongs to a certain species. If a creature dies (without being killed), it is replaced by a new creature from a randomly picked species. Two parents from this species are randomly selected, and the genetic operators mentioned above are applied. The resulting new creature is included in the same species as its parents. On the other hand, if a creature is killed, it is replaced by a mutated clone of its killer. The new creature (the mutated clone) is included in the same species as the killer. Thus species can increase and decrease through deaths and replacements.

If a species' population falls to 3 individuals, then this species is deemed extinct. Its remaining creature are removed, and the available slot are attributed to the largest species. The largest species is then *split* into two separate species, of equal size. One of these two daughter groups continues the split species, while the other replaces the extinct species. Thus, as each extinction creates a split, the number of species is kept constant.

C. Results

As in the panmictic model, the first behaviour to emerge is consistently a “roaming” behaviour: creatures that use oscillating limbs to move fast and hit at close-by creatures. Discovery of these behaviours is the key event that generates active natural selection, as indicated by the fitness transmission graph (compare Figure 7 and Figure 6).

This behaviour then consistently triggers an evolutionary response from other species, namely *miniaturisation*: competing species tend to get smaller in volume. The advantage of miniaturisation is obvious: the smaller you are, the lower the chance of getting hit by roamers. When the number of roamers is high enough, small creatures can obtain a strong survival advantage, which leads to increased reproductive success as long-lived individuals have more opportunities to be chosen as parents.

Figures 5-6 depict the population size, the number of kills inflicted, and average creature size for runs involving two or three species, as a function of time. In these pictures, the x -axis corresponds to “cycles”, where each cycle covers 5000 timesteps. While the dynamics of these runs vary, certain common features are readily seen: after a short initial period, one species reaches an efficient (roaming) behaviour, resulting in a high number of kills and population increase. The other species shrink rapidly, not only in numbers, but also in average size. In the end a steady state emerges, in which one species has specialised as “active” roamers, reproducing mostly through kills, while the others specialise in miniaturisation and reproduce “passively” through replacement of deaths. As indicated in the figures, extinction events

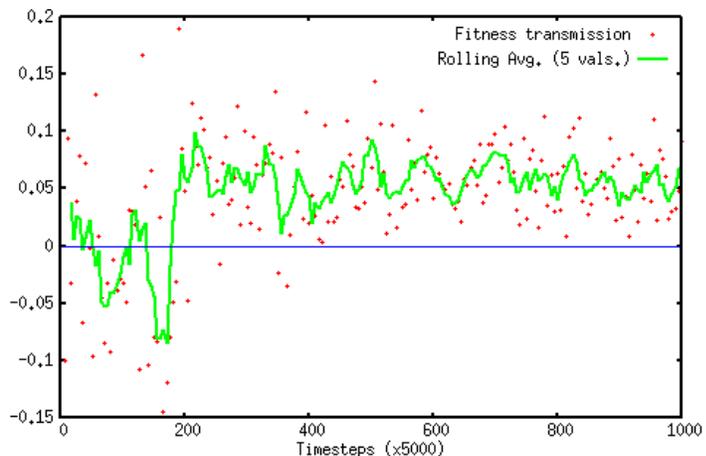


Fig. 7. Fitness transmission for the run depicted in Figure 6. The onset of positive fitness transmission coincides closely with the increases in numbers of kills and population size divergence.

may also occur. The “miniature” species seem to have much more variance in volume than the active species, which is expected given their smaller population sizes.

The division between active (killing), large, numerous species and passive (non-killing), small species was consistently observed. Furthermore, the system seems to favour the presence of only one active species. After extinction events, when the largest species (which is invariably the active one) is split into two groups, only one of these groups remains active, while the other is forced into a passive lifestyle.

The behaviour of the active species generally relies on the use of one or more undulating limbs to achieve locomotion. This can be used in several ways. One option is to emphasise locomotion, in order to cover as much ground as possible. Species following this strategy tend to evolve slim morphologies with very efficient motile behaviour. This strategy often culminates in streamlined, high-speed “worms”. Another possibility is to emphasise damage dealt on contact. These species develop morphologies that travel less efficiently, but maximise the probability and intensity of damage when colliding with other individuals.

The presence of different strategies, which emerge in response to each other (rather than being immediate adaptations to external parameters of the system) indicate that the system lies in class 2 of our classification.

VI. CONCLUSION

We have described an open environment, Evosphere, in which free-living creatures interact freely. Reproduction can occur either by killing and cloning, or simply by staying alive long enough to be chosen as a parent for replacement of dying creatures that were not directly killed. This environment is successful in generating natural selection and adaptive evolution. Adaptive evolution is demonstrated both by visual inspection of creatures, and by fitness transmission calculations.

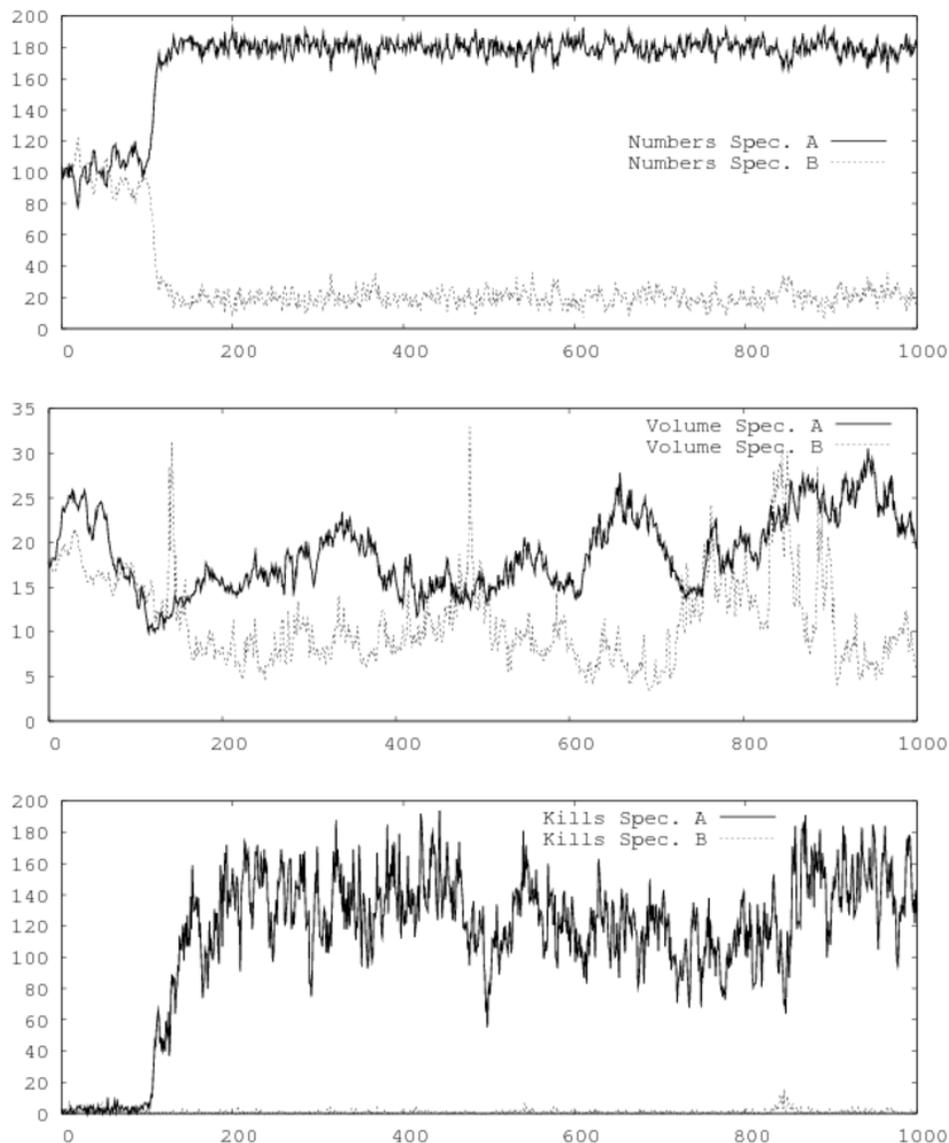


Fig. 5. From top to bottom: population sizes, average volume, and number of kills inflicted for two species A and B. The x -axis (time dimension) is graded in “cycles”, where each cycle corresponds to 5000 timesteps. Notice the initial indistinct period, followed by a division between a larger, active species and a smaller, passive species. Notice the diminution in volume of (active) species A around cycle 500: this is due to the invention of a highly successful “worm” form, which is smaller in volume than previous killer morphologies.

Furthermore, when restrictions on mating are imposed through the introduction of species, evolutionary “responses” are observed, in that one species evolves in adaptation to another one’s evolutionary changes. Thus evolving species in this system are in principle capable of creating and altering each other’s fitness landscape. While the behaviours observed so far remain simple and limited in number, this is an encouraging indication of this system’s potential for future investigations.

REFERENCES

[1] M. A. Bedau, E. Snyder, and N. H. Packard. A classification of long-term evolutionary dynamics. In C. Adami, R. K. Belew, H. Kitano, and C. E. Taylor, editors, *Artificial Life VI: Proceedings of the Sixth*

International Workshop on the Synthesis and Simulation of Living Systems, pages 228–237, Cambridge, MA, 1998. MIT Press.

- [2] A. D. Channon. Unbounded evolutionary dynamics in a system of agents that actively process and transform their environment. *Genetic Programming and Evolvable Machines*, 7(3):253–281, 2006.
- [3] A. D. Channon and R. I. Damper. Perpetuating evolutionary emergence. In R. Pfeifer, B. Blumberg, J.-A. Meyer, and S. Wilson, editors, *From Animals To Animats 5: Proceedings 5th International Conference on the Simulation of Adaptive Behavior (SAB '98)*, pages 534–539, Cambridge, MA, 1998. MIT Press.
- [4] T. Miconi. *The Road to Everywhere: Evolution, Complexity and Progress in Natural and Artificial Systems*. PhD thesis, University of Birmingham, 2007.
- [5] T. Miconi. In silicon no one can hear you scream: Evolving fighting creatures. In M. O. et al., editor, *Proceedings of the 2008 EuroGP conference*, volume 4971 of *LNCS*. Springer, 2008.
- [6] T. Miconi and A. Channon. An improved system for artificial creatures

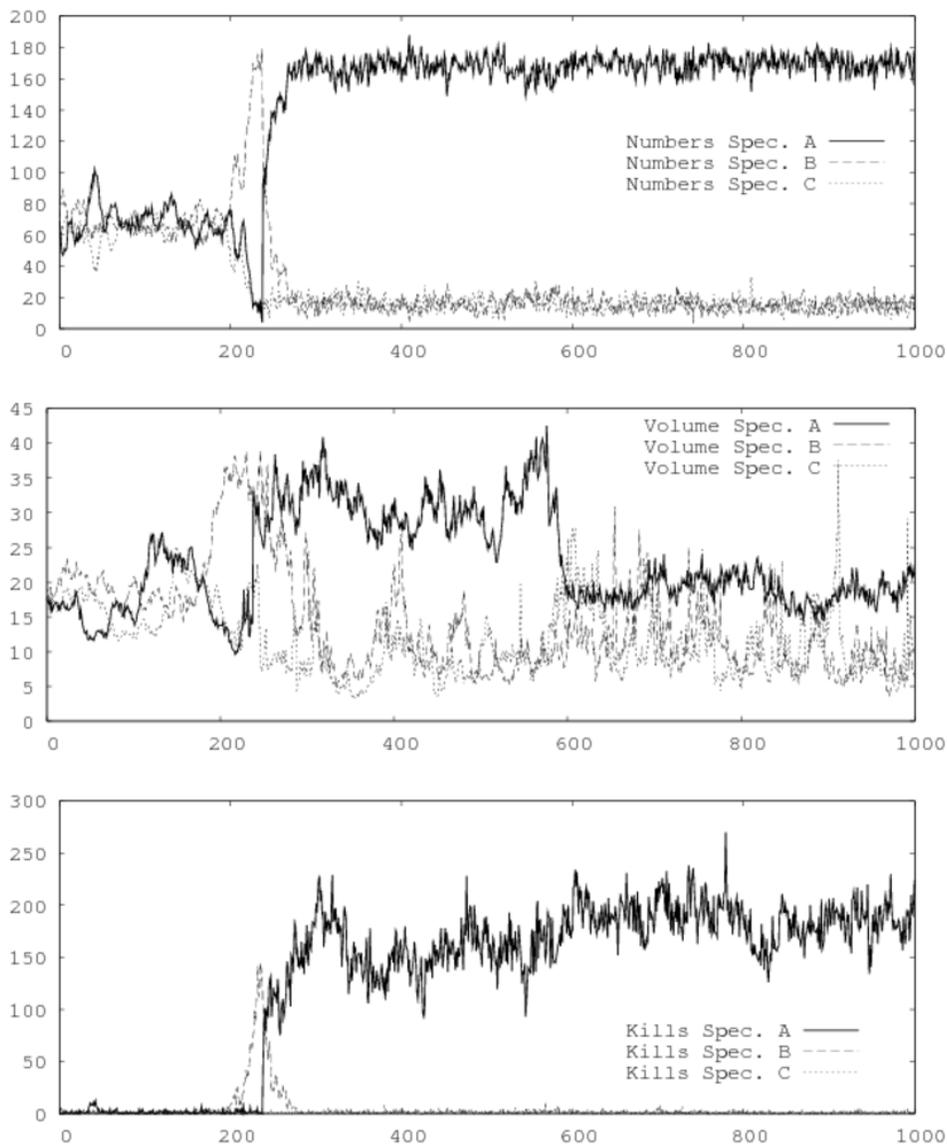


Fig. 6. From top to bottom: population sizes, average volume, and number of kills inflicted for three species A, B and C. The gradation and general features are similar to those in Figure 5. The apparent initial exchange in positions between species A and B is due to an extinction event: soon after species A discovers a working killing behaviour, species B goes extinct. Species A is split in two and one of the resulting groups is affected to replace species B. This group proves more successful and confines its sister group (still labelled “species A”) to a passive, miniature lifestyle.

evolution. In L. Rocha, M. Bedau, D. Floreano, R. Goldstone, A. Vespignani, and L. Yaeger, editors, *Procs 10th Intl Conf on Simulation and Synthesis of Living Systems (ALIFE X)*. MIT Press, 2006.

[7] N. H. Packard. Intrinsic adaptation in a simple model for evolution. In C. G. Langton, C. Taloy, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II: Procs of the workshop on the Synthesis and Simulation of Living Systems*. Addison Wesley, 1991.

[8] T. S. Ray. Artificial life. In R. Dulbecco, D. Baltimore, F. Jacob, and R. Levi-Montalcini, editors, *Frontiers of Life, Volume One: The Origins of Life*, pages 107–124. Academic Press, 2001.

[9] K. Sims. Evolving 3d morphology and behavior by competition. In R. Brooks and P. Maes, editors, *Procs 4th Intl Works on Synthesis and Simulation of Living Systems (ALIFE IV)*, pages 28–39. MIT Press, 1994.

[10] K. Sims. Evolving virtual creatures. In *SIGGRAPH 94*, pages 15–22. ACM Press, 1994.

[11] L. Yaeger. Computational genetics, physiology, metabolism, neural systems, learning, vision and behaviour or polyworld: Life in a new context. In C. G. Langton, editor, *Artificial Life III, Vol. XVII of SFI Studies in the Sciences of Complexity*, pages 263–298. Addison-Wesley, 1994.