

EMERGENCE OUT OF INTERACTION: A PHENOTYPE BASED MODEL OF SPECIES EVOLUTION

George Kampis¹ and Laszlo Gulyas²

¹*History and Philosophy of Science, Eötvös University, Budapest,
P.O.Box 32, H-1518, Hungary
kampis@hps.elte.hu*

²*Computer and Automation Research Institute,
Hungarian Academy of Sciences, Budapest, Hungary
gulya@omega.ailab.sztaki.hu*

ABSTRACT

We introduce an agent-based simulation to approach the problem of persistent species evolution in artificial evolutionary systems. Our simulation is based on the theory of ‘fat’ phenotype as applied to generative evolutionary forces. We present a model as well as computational results to illustrate how ‘fat’ phenotypes can generate a changing interaction space which defines new selection forces that recursively give rise to new species. We discuss prospects for radical evolutionary technology and for emergent synthesis.

Keywords: evolution; phenotype; interaction; species emergence; agent-based simulation

1. INTRODUCTION

The fundamental problem we consider is disarmingly simple: how is it possible to produce sustained evolution in an artificial system? This is one of the most profound current challenges for Artificial Life modelling (Kampis 2002, 2003, Holland 2003). Intricately bound to this problem, there is a conceptual one. There can be various definitions of evolution, which range from genetic change to adaptation and to the growth of morphological complexity. Of these, and the many others possible, our approach utilizes one of the most fundamental evolutionary concepts, that of species. We consider the challenge of sustained evolution as the problem of producing new species, i.e. reproductively isolated and functionally different sub-populations. We present both a theoretical model and a computer simulation.

2. THE THEORETICAL FRAMEWORK

The model applies the changing dynamics of phenotype-to-phenotype interactions in a population of sexually reproducing agents.

The approach is based on the evolutionary application of “fat” interactions in natural causal processes (Kampis 2002a,b). The causal interactional framework provides a natural tool for discussing the problem of the production of species. Stable species emergence must be an inherently ecological problem: multiple species with identical needs (i.e. which occupy the same niche) tend to compete and cease to coexist. As a consequence of this situation, the problem of species evolution is closely related to the emergence and maintenance of different ecological niches (Lewontin 1983, Laland et al. 2000, Odling-Smee et al. 2003). In a series of recent writings and talks (Kampis and Gulyas 2003, Kampis 2003a,b) we developed the idea that the dynamics of niche emergence can be guided by recursive changes in phenotype-based interaction space.

In a purely genotype-based evolution model, the transformation dynamics is subjected to identical rules over the entire evolution period – rules describing the process of gene selection (i.e. replicator selection). In such a framework, the direct modeling of the genetic emergence of species is very difficult, if not impossible, to achieve, in lack of flexible new selection forces introduced artificially from the outside. Phenotype-based evolution, in contrast, does not suffer from the same difficulty. Here, selection processes are exerted via the phenotype (i.e. the interactor), which is less rigidly defined. Making selection forces depend on interactors, we can try to exploit the richer dynamics of the latter, to produce genuine and ecologically meaningful new selection forces endogenously.

In a real-world evolutionary system, the relation between the “hard” genotype, considered in the first part of the above paragraph, and the “soft” phenotype, considered second, is specified by a combination of ontogenesis and ecological context. Of these, in the present model, we focus on the second.

We view the interaction space of the phenotypes as an evolutionary product, and we attempt to grasp its origin. The idea is best illustrated on a simple example such as sexual selection. What matters as a relevant selection variable when two sexual organisms are paired is a function of the two individuals’ relations. The variable in one organism is chosen by the other organism’s variables that interact with it. Mating can occur when the values taken by these two sets of variables fit. To use a toy example, if the female prefers large male antlers (variable 1) and the male possesses the large antlers (variable 2), then a fit or match is possible, and reproduction will occur, which propagates both the genotypes and the phenotypes. As readily seen from this example, such a match is “groundless”, however, in that it is based on the relational properties of the two phenotypes alone. This characteristic of sexual selection is, we suggest, a suitable metaphor for more general ecological and evolutionary interactions.

The metaphor we wish to exploit is that ecological evolution is as groundless as sexual selection, and can be bootstrapped by similar mechanisms that keep it in motion. In this way it might be possible to efficiently ‘scan’ a large design space (Dennett 1995, Kampis 2002a) which is otherwise inaccessible for ‘total synthesis’ (Kampis 2002b).

3. BASICS OF THE MODEL

In the rest of the paper, we study the phenotype-based evolutionary dynamics of sexual selection. Our motivation was the hypothesis, corroborated by the results to be presented, that in a sexual selection-based model species, evolution can transform and finally split the population when genetic mutations produce individuals with a new phenotype, such that it redefines the relevant interactions. An example is a new female that prefers body size and not antlers, to stay with our example.

With the introduction of such “dissenter” individuals, silent phenotype traits in the other individuals (i.e., body size, in the given example) become suddenly activated and become part of a changed ecological interaction space.

Note the fact (which will be of primary importance), that all individuals are simultaneously effected by the transition. This cannot be otherwise, because they are all potential mating partners of the dissenter. As a result, the old global sexual selection pressure is now supplemented with a new one, that arises spontaneously, endogenously, and within a fully sympatric (i.e. spatially coextensive) population. This new sexual selection process can potentially lead to the development of a new best match, and, consequently, to a new, sexually reproducing, stable

sub-population, which reproductively isolated from the rest of the original population.

In the computational part, we present and discuss an agent-based simulation developed by the authors to illustrate the feasibility of this approach. The agent-based framework is provided by the REPAST environment (Cederman and Gulyas 2001, Repast 2003). We performed experiments with populations of several hundred or a thousand individuals. Our population consists of sexually reproducing gender-less individual agents (modeled on organisms such as snails) whose ecological properties are represented by a phenotype vector. The intended interpretation is that components of the phenotype vector stand for currently active (or “turned on”) ecological interactions. Mating success is introduced in this system as a function of a similarity measure, defined as a distance metric over the phenotype vector pairs. To establish a basic evolutionary setting with variability, population turnover and overlapping generations, every individual was equipped with a minimal “physiology” that requires it to eat food (supplied externally in the form of energy), and to undergo aging, leading ultimately to death (modeled as a progressive failing of the efficiency of energy processing).

Reproduction was represented as the spawning of new agents, accompanied by a crossing over and mutation, executed directly on the phenotype vectors (in other words, the underlying ontogeny is trivial). In the current version of the simulation, the only biologically (i.e. in the course of a run, internally) adjustable parameters are the phenotype vectors.




Fig.1. Dimension change prohibited, the population tends to develop into one single stable species, characterized by a self-developed center and a normal distribution.

Phenotype vectors are understood in the model as variable length records that remain fixed during the

lifetime of an agent but have plasticity otherwise. Interaction change is represented as the change of the dimensionality of the phenotype vector at the birth of a new “dissenter” agent. The semantics of interaction change implies that this transformation of dimensionality, introduced in one individual, is swept instantaneously across the whole population, corresponding to the global nature of the concept of interaction space, as discussed above.

If interaction dimension change is prohibited, the population in the model tends to develop into one stable species under a wide range of parameters, if started from a seed of agents with randomly selected phenotypes. The stable population is characterized by a self-developed center and a normal distribution (see Figure 1.) The internal introduction of new interaction dimensions, on the other hand, facilitates the spontaneous development of more species, i.e. sub-populations with different emergent foci, reproductively isolated from each other.

A key issue in the technology of modeling interaction change is the assignment of phenotype traits to new interaction dimensions. We performed experiments with several assignments, among them a type-based deterministic and a non-type based random, and a modulo rule based assignment, the meaning of which is explained at the end of Section 4. New phenotypes are ‘stretched’ along the interaction dimension, which introduces new variance. We found that the emergence of new species was effected but did not crucially depend on the choice of a particular ‘stretch’ solution applied. At the end of the paper we discuss the feasibility of natural mechanisms to achieve the same results in biologically meaningful conditions.

4. SPECIFICATION OF THE MODEL

Our experiments were conducted using a basic evolution engine that maintained a stable population. The engine defines a usual evolutionary setting and provides a sympatric environment without a spatial component. The engine simulates a partial artificial ecology with one resource, namely, energy. Each organism has an equal chance to ‘eat’ in every time step. This fact introduces an implicit competition for energy, which leads to density-dependent effects. Genotypes represented with a higher number of individuals will necessarily obtain a larger share, even if all other things equal.

The emergent model is built on top of the evolution engine. The latter consists of a population of agents and a non-spatial environment shared by them. Each agent has its phenotype represented as a vector of integers from the interval $[V_{min}, V_{max}]$. (The length of the vector is always identical for all agents.) In addition to their phenotypes, agents only have the minimum of properties: age and accumulated energy.

The evolution engine uses a quasi-parallel activation regime. Every agent gets activated exactly once per every time step, in a dynamically randomized order. The agents’

activity consists of three steps. These are: energy intake, energy consumption, and reproduction.

The agent first seeks E_{in} units of energy from the shared environment, and, depending on the amount available, it receives e_{in} units (possibly 0). The efficiency of energy intake decays with age:

$$e_{accumulated} = e_{accumulated} + e_{in} \cdot (E_{discounting})^{age},$$

where $0 < E_{discounting} < 1$. Next, the agent consumes $E_{consumption}$ units of its accumulated energy. If the agent does not have the sufficient amount available, it dies.

Surviving agents attempt to reproduce with probability $P_{encounter}$. The updating of the shared environment, which means the addition of $E_{increase}$ units to the energy pool, completes an iteration step.

When reproducing, the agent picks a random mate from a list of potential partners, limited to individuals similar enough to bear an offspring with the given agent. Similarity is measured using Euclidean distance between the agents’ phenotype vectors. An important advantage of this metric is that it is dimension independent. Therefore, it allows for sexual selection to occur in an identical way between all conceivable pairs of phenotypes, being of arbitrary dimensionality.

Given two parents, and similarity d defined as above, the number of offspring is $M_{const} + (M_{limit} - d) \cdot M_{slope}$ (for $d > M_{limit}$). The new agents inherit their parent’s phenotypes, except for mutation and crossing over that occur with probabilities $P_{mutation}$ and $P_{crossing}$, respectively, per gene. Mutation shifts the value of a gene by a random value in $[-V_{mutation}, +V_{mutation}]$. If the mutated value falls outside $[V_{min}, V_{max}]$, the offspring is dropped.

A new phenotype dimension is introduced with probability P_{change} per offspring. When such an event occurs, a new trait slot is added to the end of the agents’ phenotype vectors. The particular value an agent receives depends on the used ‘stretch’ method. The type-based method calculates this value from the agent’s old phenotype traits, whereas the non-type based method relaxes this condition. For simplicity, only the value v of the last dimension is used in the type-based procedure:

$$v_{new} = V_{min} + (v \cdot V_{stretch}) \bmod (V_{max} - V_{min} + 1),$$

where $V_{stretch}$ is a positive parameter. The type-independent method selects a uniform random value from $[V_{min}, V_{max}]$. The meaning of the expressions ‘type-based’ and ‘non-type-based’ should be clear from the procedure: the first generates identical new phenotypes for each genotype, the latter permits phenotypes and genotypes to be linked by many-to-one mappings.

5. RESULTS

We summarize our main results by first pointing out the observed dramatic difference between the ‘flatline’ behavior of the basic engine and the divergent, rich species production behavior of same system with interaction change. If interaction change is introduced, a formerly stable convergent species becomes more extended in

property space, and finally splits, giving rise to two or three new stable populations, or species. The process can repeat itself several times, ultimately producing a host of new, smaller species, all functionally and reproductively distinct from each other. This, together with the competitive forces from the density effects (discussed in 4.) yields a steadily, but slowly increasing number of species as illustrated on Figure 2.




Fig.2. The evolution of species. The graph shows the average number of species (over 10 runs) versus time using the default settings. Error bars show minimum and maximum values, respectively.




Fig.3. Spontaneously evolved species in a Fruchterman-Reingold (1991) plot after 6000 steps. Disconnected components denote sub-populations in reproductive isolation.

The qualitative process is as follows. The introduction of new dimensions alters the property distribution of phenotypes, so that their average distance increases. This makes it possible for several individuals to escape the attraction of the sexual selection force that maintains the original species and keeps it confined within a small radius. At the same time, along the new dimension some new close fits can emerge that diminish the distance

among these individuals, which are far from the center of the original species. New mating centers are defined, and reproduction makes them stronger. This can happen with a high probability, if the distance within these new germs is smaller than the average distance within the old species. Finally, with the death of individuals which were close enough to both centers, so as to reproduce with the individuals from both groups (and to maintain the whole population as a single species), the original species splits into two (or several) parts. These parts will never united again, because their intermediates are strongly selected against when the radius of the new species decreases (Figure 3.).

Prior to testing these consequences of interaction change, we conducted several experiments to test the underlying evolution engine. Our results show that, left alone, the engine maintains a single species under a very wide range of parameters. We applied conditions where a constant energy influx constrained the maximum average number of individuals maintained in the model. Under such conditions, new species can emerge at the indirect cost of others. Together with the density-dependent effects of the energy uptake algorithm (see Section 4.) this fact poses a slight limitation on the possible emergent processes we are interested in. On the other hand, maintaining a stable and realistic minimal ecosystem was considered fundamental to the study of species emergence. Detailed results about the convergence and stability properties of the evolution engine can be found in (Kampis and Gulyas 2004, Gulyas 2004).

V_{min}	0	M_{limit}	15
V_{max}	100	M_{slope}	0
$P_{encounter}$	0.1	M_{const}	1
$P_{crossing}$	0.2	$E_{consumption}$	5
$P_{mutation}$	0.1	E_{in}	10
$V_{mutation}$	2	$E_{discounting}$	0.9
		$E_{increase}$	1000

Table 1: Default parameter settings.

Let us now focus on our key findings again, the effect of interaction change. All experiments were performed with populations of several hundred or a thousand individuals. A typical run was started with phenotypes of 5 dimensions. During the simulation process this number increased to about 50 in a few thousand time steps. To ensure the robustness of our results we conducted extensive experiments with various parameter combinations. Specifically, we varied each parameter in a wide range, while keeping the rest at its default value.

As discussed earlier, the normal functioning of the evolution engine leads to a fast convergence to a single species. Sill, as shown on Figure 2., the gradual introduction of new phenotype interaction dimensions results in a growing number of reproductively isolated groups. As the energy-system of the underlying evolution

engine can support a certain number of individuals only, there is a natural upper limit to this number. Except for this factor, the interaction change as simulated in this model yields a persistent evolution of new species.

P_{change}	0.0075
'Stretch' method	Type-based / Type-independent
$V_{stretch}$ (Only used with the type-based method.)	1

Table 2. Default values for the interaction change parameters.

There are three parameters related to interaction change as summarized in Table 2. We experimented with changing them one by one. Changing P_{change} between 0.0005 and 0.001 yielded similar results: the number of species shows a monotonic growth. The curve gets steeper for higher values of the parameter, while the number of species after 12000 iterations falls between 5-15 depending on it. The value of $V_{stretch}$ was varied between 1 and 20, with similar resulting behavior in all cases, although higher parameter values yield higher variance. Finally, the type-independent 'stretch' method also reproduces our qualitative results, although with much greater variance (see Figure 4).




Fig. 3. The evolution of species with a type-independent 'stretch' method. The average number of species (over 10 runs) versus time, using default settings. Error bars show minimum and maximum values.

7. DISCUSSION AND FURTHER RESEARCH

Using a simple toy example our model demonstrates the validity of the in-principle claim that changing interaction fields can lead to emergent effects producing sustained evolution. More detailed sensitivity analysis is presented elsewhere (Kampis and Gulyas 2004, Gulyas 2004).

The presented model clearly lacks biological feasibility (dimensions are only added, never dropped; properties are assigned too radically in the whole population, etc.), yet we maintain that there exists several biologically relevant mechanisms that can produce similar results. One candidate is the existence of phenocopies, another is the

switching of food or mating preference as it occurs in actual populations. However, the computational modeling of these mechanisms would require the use of epigenetic and ecological theory, as well as detailed phenotype-to-environment mappings, and goes way beyond what is attempted here.

Future research should in particular be concerned with the extension of our results for systems with a real ecology, i.e. a variety of producers and consumers that do not stand in (explicit or implicit) competition as in the present model.

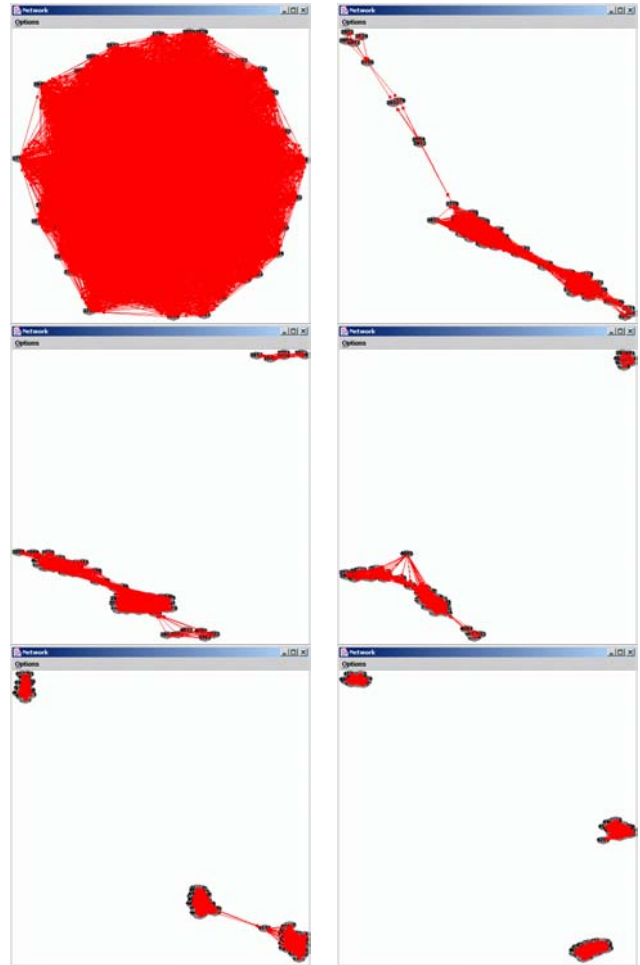


Fig. 5. Stages of the development of a population splitting into several "species", i.e. reproductively isolated sub-populations. Lines mark mutual reproduction ability between pairs of individuals represented as dots.

8. TOWARDS EVOLUTIONARY TECHNOLOGY

Finally, we discuss the relevance of our model in the modelling and simulation of emergent systems in general, and for the applications to technical as well as biological systems.

Agent-based emergence is a much studied phenomenon. We argue that interaction based emergence

is a powerful concept that can serve in this context as a template for a rich family of emergent processes, with applications in various fields. Our model and the underlying theory together suggest that in order to achieve open development in an evolving system one has to leave the world of entirely computational approaches. This recognition is in line with a number of pioneering tendencies for the building of hybrid systems, where the computational processes are complemented by physical realizations with an autonomous role.

Recent experiments, such as J. Pollack's Golem system (Lipson and Pollack 2000) utilize this idea in very a limited form. As R. Brooks notes (2001), it is often unclear what role exactly is played by the physical realization. Behind the realization-centered approach there is a strong intuition that actual realization is the 'computationally most effective' way of representing a real-world system. However, beyond this slogan, the exact potential that lies in a realization remains a question.

Our experiments suggest that the scope of realizations may be in their causal potential: realizations are causal, and causal processes can use a full 'bandwidth' of material interactions, which is beyond the grasp of any computational model. The model presented may illustrate that biological effects built on causal material properties can play an essential role in evolution, and perhaps also in a variety of other emergent processes that withstood formalization. The question, where to go from here, is beyond the present paper. It should be clear that the model presented is but a first metaphor which, when successful, might be used as a stepping stone for future models and realizations.

ACKNOWLEDGMENT

Part of the work reported here was carried out during the first author's stays in the School of Knowledge Science, JAIST, Japan, and in Kalamazoo College, MI, USA. The hospitality of these institutions, as well as the personal support of Professor S. Kunifuji (JAIST) and Professor Peter Erdi (Kalamazoo) is gratefully acknowledged. Computer simulations were done on the BeoWulf cluster of the Center for Complex Systems Studies, Physics Department, Kalamazoo College, as well as on the SUN E10K, E15K system of Hungary's NIF Supercomputing Center in Budapest.

REFERENCES

Brooks, R. (2001): The Relationship Between Matter and Life, *Nature* 409, 409-411.

Cederman, L-E. and Gulyas, L. (2001): Tutorial Sequence for RePast: An Iterated Prisoner's Dilemma Game, <http://www.cbrss.harvard.edu/events/cccsn/tutorial.pdf>

Fruchterman, T.M.J., and Reingold, E.M. 1991. Graph Drawing by Force-directed Placement, *Software-Practice and Experience*, 21(11):1129—1164.

Gulyas, L. (2004): Summary of EvoTech-V v5.0.8, Technical Report, Eötvös University. <http://hps.elte.hu/~gk/EvoTech/Summary.pdf>

Holland, J.H. (2003): Challenges for Agent-based Modeling, <http://www.nd.edu/~swarm03/Keynote/keynote.html>

Kampis, G. (2002a): A Causal Model of Evolution, *Proceedings of SEAL 02* (4th Asia-Pacific Conference on Simulated Evolution And Learning), Singapore, pp 836-840.

Kampis, G. (2002b): Towards an Evolutionary Technology (Japanese, in press), English version: www.jaist.ac.jp/~g-kampis/EvoTech/Towards.html

Kampis, G. (2003a): A Causal Model of Evolution, invited lecture at the Center for the Study of Complex Systems, U. Michigan, Ann Arbor, 2003. Feb. 20., and Department of Philosophy, U. Cincinnati, Feb, 24. www.pscs.umich.edu/calendar/abstractsw03.html#Kampis www.jaist.ac.jp/~g-kampis/USA/Causal_Model.html

Kampis, G. (2003b): Complexity and Causality, (Complex Systems 7., JAIST report), Japan Advanced Institute for Science and Technology, Tatsunokuchi, pp. 80.

Kampis, G. and Gulyas, L. (2003): Causal Structures in Embodied Systems, *ERCIM News* 53, April 2003, 23-25. www.jaist.ac.jp/~g-kampis/ERCIM/Kampis-Gulyas.html

Kampis, G. and Gulyas, L. (2004): Sustained Evolution from Changing Interaction, *Alife IX*, Boston, accepted. <http://hps.elte.hu/~kampis/EvoTech/Sustained.pdf>

Laland, K.N.; Odling-Smee, J.; and Feldman, M.W. (2000): Niche Construction, Biological Evolution and Cultural Change, *Behavioral and Brain Sciences*, 23(1):131—146.

Lewontin, R. (1983) The Organism as the Subject and Object of Evolution. *Scientia*, 118:65—82.

Lipson, H. and Pollack, J. B. (2000): Automatic Design and Manufacture of Robotic Lifeforms, *Nature* 406, pp. 974-978.

Odling-Smee, J.; Laland, K.N.; and Feldman, M.W. (2003): *Niche Construction: The Neglected Process in Evolution*, Princeton, NJ: Princeton UP.

Repast (2003): <http://repast.sourceforge.net/>