Artificial-Life Ecosystems: What are they and what could they become?

Alan Dorin[1], Kevin B. Korb[1] and Volker Grimm[2]

[1]Faculty of Information Technology, Monash University, Clayton, Australia 3800
[2]UFZ, Center for Environmental Research - UFZ, Department of Ecological Modelling, Permoserstr. 15 04318 Leipzig, Germany
alan.dorin@infotech.monash.edu.au

Abstract

This paper summarises the history of the terms *ecology* and *ecosystem*, before examining their application in the early and recent literature of A-Life agent-based software simulation. It investigates trends in A-Life that have led to a predominance of simulations incorporating artificial evolution acting on generic agents, but lacking a level of detail that would allow the emergence of phenomena relating to the transfer and transformation of energy and matter between the virtual abiotic environment and biota. Implications of these characteristics for the relevance of A-Life's virtual ecosystem models to Ecology are discussed. We argue a position that the inclusion of low-level representations of energetics, matter and evolution, in concert with pattern-oriented modelling techniques from Ecology for model validation, will improve the relevance of A-Life models to Ecology. We also suggest two methods that may allows us to meet this goal: artificial evolution can be employed as a mechanism for automating pattern-oriented ecological modelling from the level of individual species up to that of the ecosystem, or it may be employed to explore general principles of ecosystem behaviour over evolutionary time periods.

Introduction

As even a cursory survey of the early and current literature reveals, within the fields of Artificial Life and Ecological Modelling, agent (individual)-based virtual ecosystem model construction has been widely practiced ([1-4] are some early examples from A-Life, also see [5], for surveys of ecological examples [6], and more recently [7, 8]). Of course there is overlap between Ecological Modelling and A-Life publications in this regard, but a careful elucidation of the differences between the historical and current trends in the fields' approaches to virtual ecosystem construction allows us to recommend a mechanism for overcoming some of their limitations. We suggest this may be achieved by melding the approaches of both fields into models that explicitly represent energetics, matter (chemical stoichiometry) and evolution within a single simulation framework. The need for including these three frameworks was noted in the literature some time ago [9, 10]. Additionally, any purportedly descriptive simulations must be validated against ecological data.[1] One method to achieve this is through *pattern-oriented modelling* [11], a technique summarised below.

We will discuss two under-explored ways in which hybrid A-Life/Ecology ecosystem simulations of this kind may be built. Firstly, models of energy and matter transfer adopted from A-Life's artificial chemistry simulations may be incorporated within generic ecosystem simulations. In this context, artificial evolution may be employed to select agent parameters producing general patterns that may be validated against ecological field data, without regard for the behaviours of particular species or habitats. This would allow for studies of the generic properties of ecosystems.

Secondly, artificial evolution may be used to select parameters for pattern-oriented modelling from the level of specific species up to the level of specific ecosystems. Once a set of patterns has been matched, the evolution algorithm can be disabled and the ecosystem simulation may be used to answer questions concerning that specific ecosystem over sub-evolutionary time periods (or over longer periods, disregarding the effects of evolution).

By adopting such approaches it is possible to extend the range of questions that may be answered by ecosystem simulations for both Ecology, by locating parameters that match field data, and A-Life, by answering questions of ecological relevance whilst permitting exploration of the general properties of ecosystem behaviour outside the familiar domain of evolution.

Before investigating the application of ideas from Ecology and A-Life to the construction of virtual ecosystems, we shall give a brief overview of significant and relevant stages in the development of these fields.

Ecology and the Ecosystem

Ernst Haeckel coined the term *ecology* in *Generelle Morphologie der Organismen* (1866) to give form to the study of Natural History in the context of Darwin's ideas that organisms must struggle for survival. Ecology was to be the study of animals, their relationships amongst themselves, with plants and with the inorganic environment that affected their

[1] Some A-Life researchers will feel that there is no need for A-Life models to reflect reality in the way this paper proposes. It is true that many A-Life models are interesting regardless of their ability to represent reality. However, this paper examines how A-Life and Ecology may be of mutual benefit to one another. Hence we discuss ways of improving the correspondence between virtual and real ecosystems.

survival and reproduction (see [12], p207). Sixty years later South African ecologist Phillips, championing the view of another ecologist, Clements [13], insisted that a collection of plants and animals that had come into a harmonic relationship with one another and their habitat through succession to climax could be seen quite literally as a *Complex Organism* [14]. Phillips viewed the process of succession as a kind of ontogeny for his *Biotic Communities*, basing his ideas for the wholeness of this community on the holistic philosophy of Smuts. His aim was, in part, to unify Botany and Zoology under a new banner.

The ecologist Tansley, unhappy with Philips' argument, chimed into the debate and countered the use of *Complex Organism* by coining the word *ecosystem* in his retort, *The Use and Abuse of Vegetation Concepts and Terms* (1935). Several alternatives to *ecosystem* have been offered (e.g. *biogeocenosis*, *microcosm*, *epimorph*, *elementary landscape*, *microlandscape*, *biosystem*, *holocoen*, *biochora*, *ecotope*, *geocenosis*, *facies*, *epifacies*, *diatope* and *bioecos* [15]), each with a slightly different slant. However in the UK, Europe, Australia, the USA and many other research communities, Tansley's term and its designated focus have stuck [12]. This is true not only in science but also in politics, philosophy and even in marketing and popular culture.

Tansley's aim for the term was to give expression to a physical system that could legitimately take its place alongside those studied by Physics. Its components were animals, plants *and* abiotic material. He called attention to the significance of the exchange of materials and energy between organisms and the abiotic environment.

It is worth noting that Tansley was not concerned with Systems Theory, a field that came to the fore only after WWII. However his term's natural fit to this mould may be a part of the reason why the idea gathered popularity in the post-war years. The preference for *ecosystem* by U.S. ecologist Odum in the editions of his textbook *Fundamentals of Ecology* played a significant role in the term's post-war success also. He refined his definition for the term across the three editions of his text. In the third (1971) he wrote, "Any unit that includes all of the organisms… in a given area interacting with the physical environment so that a flow of energy leads to clearly defined trophic structure, biotic diversity, and material cycles (i.e. exchange of materials between living and non-living parts) within the system is an ecological system or ecosystem" [16], and established for a generation of ecologists the importance of the relationships between the Earth's biotic and abiotic components and the processes by which they exchange materials and energy.

Artificial Life and the (Virtual) Ecosystem

When building models we necessarily abstract away detail and represent only what we believe is responsible for determining the behaviour we wish to study or predict. Both the subjects of study and the decisions made regarding the level of abstraction can enlighten us about the different perspectives ecosystem modellers have adopted.

Langton's call for research to explore "life-as-it-could-be" was in keeping with John von Neumann's original interest in abstract self-reproducing computational systems. The idea that life might be some property of form, independent of matter, however much it is debated philosophically, has set the stage for explorations into virtual ecosystems within the field of A-Life that are more often generic than representative of particular organisms, species or their abiotic habitats. Below we shall discuss several of the early A-Life virtual ecosystems to highlight the ways in which this has been evident. The proffered research interests of these systems' creators parallel those from *Systems Ecology*, "The goal of Systems Ecology is an ecosystem phenomenology that does not necessarily require detailed information about individual species" [17]. Ray's *Tierra* [4], Yaeger's *Polyworld* [18], Holland's *Echo* [2] and Packard's *Bugs* [3] ecosystems (to list a few from the early days) fall within this category of what might be labeled *Generic Virtual Ecosystems*. Within a mini-review of individual based models in Ecology conducted in 1999, Grimm called for similar studies of the generic properties of ecosystems by ecologists [6].

Given the early interest in artificially evolving ecosystem models by A-Life researchers (shown below), it is ironic that Systems Ecologists at around the time of A-Life's inception wrote, "In spite of some attempts to address evolution in the systems literature, evolution is not well represented or adequately incorporated" [9]. The authors continue,

> …incorporating evolutionary theories into systems ecology is an underexplored but potentially fruitful avenue. Reiners (1986) has proposed that unifying ecosystem ecology requires at least three separate but complementary theoretical frameworks: energetics, matter (stoichiometry), and some aspect of population interactions or ecosystem "connectedness".

In keeping with Reiners' view, Loehle and Pechmann argue that evolution theory can provide the third of these frameworks. As these three authors highlight, whilst it is (of course) possible to study ecosystems without simultaneous reference to all three frameworks, numerous cases stand as examples in which the interaction of ideas from all three enhances our level of understanding. We treat each of these frameworks below in the context of virtual ecosystems in A-Life and Ecology. These two fields have many simulations that address the above concerns. As far back as 1999, Grimm requested that "individual-based modelling must refer to the framework of classical theoretical ecology" [6]. We maintain that it would be beneficial to construct complete simulations that encompass all three of the (sub-)frameworks suggested by Reiners, Loehle and Pechmann. Finally, this paper explains how pattern-oriented modelling may validate these well-rounded virtual ecosystems against their real counterparts — an issue that must be addressed to satisfy the demands of Ecology, and one that A-Lifer's should take seriously if they wish even their generic models to be pertinent.

Energetics

Energetics in A-Life's Virtual Ecosystems

At least since Odum, energetics has played a significant, even defining role in Ecology. It is considered by some to be the best-developed aspect of ecosystem ecology [10]. Energy flows within an ecosystem give rise to various well-studied phenomena including trophic levels, food chains and webs, productivities, and efficiencies. In addition, processes for energy flow are one determining factor for organism evolutionary adaptation (e.g., metabolism, organism morphology and locomotion). Ecology's focus on energetics has not been duplicated in A-Life's studies of virtual ecosystems, an issue flagged as early as 1994 by Lindgren and Nordahl in a paper that describes a similarly limited model [19]. (Also see [20] for a model explicitly modelling food webs.)

Energy enters A-Life's virtual ecosystems in different ways. In some cases it makes a "magical" appearance at a specified rate and is acquired by agents without taking the form of virtual matter at any stage (e.g., [21] in which a virtual Sun shines on the space).

In their model employing *Echo*, Forrest and Jones do not include anything that an Ecologist would recognise as an energy model [2]. As a substitute, they model different types of matter that agents must collect to persist and reproduce. This will be discussed in the following section.

In Ray's *Tierra* energy is equated with CPU cycles in which to execute instructions [4]. CPU cycles are consumed by an agent to reshape its local environment by executing instructions which may, in some cases, involve writing new instructions into daughter memory cells. The focus is not on how these CPU cycles are converted into particular types of instructions (all CPU cycles are equal, and all instructions require them in equal measure). However, improved copying efficiency can be achieved by minimising energy expenditure (CPU usage). Hence there is selection pressure acting on Tierra agents to be short and reliable.

The model described by Cooper and Ofria employs *Avida* [22] and is similar to *Tierra* in some respects, however organisms "metabolise" (see the section below on *Matter* for a discussion of the use of this term) different resources to gain benefit from them by performing computations. There is inter-agent competition for these resources, which in their model, are in limited supply. Hence, to some extent this model represents the acquisition of energy by organisms. The simulation *Cosmos* by Taylor and Hallam [23] was based on *Tierra*. Its authors explicitly note *Tierra's* problematic "energy for free" and partially rectify this in their own system by requiring *Cosmos* programs to capture energy and store it, in order to use it to perform useful work.

PolyWorld incorporates a different energy model than those listed above [18]. An initial dose of energy is allocated to newborns from their parents at birth. To persist and act, agents require an amount of energy that depends on their physiology. Agent size dictates energy storage capacity. Agents maintain this store by eating plant-like food that grows in the environment at a regulated rate, or by consuming one another. Agent bodies have a food energy value separate from their current health value. Although they may die in a fight or for lack of health, this ensures that their body provides energy to a predator or scavenger.

Even in *PolyWorld* the mechanisms by which energy is extracted from materials do not emerge from the system. They are simplified to a single, high-level, hard-coded behaviour — agents just "eat". Therefore, although the agents must optimise their use of energy, the model cannot tell us about the way in which energy is transformed from one form to another across trophic levels, from sunlight to chemical energy stored in plant or animal biomass, to heat or even to kinetic energy.

The energy models employed within EVOLVE IV [1] and the simulation framework introduced by us (Dorin and Korb) [24] aimed to facilitate a greater range of emergent properties relating to energetics than any of the models above. These authors explicitly wished to explore the impact of energetics on ecosystem behaviour and consequently settled on the provision of an artificial chemistry that lies at the heart of their simulations.

We do *not* mean to imply that simple models of energetics render a simulation useless (far from it). This simplicity however limits their utility when addressing the specific problems posed by ecologists concerning particular species in particular habitats and the manner in which ecosystems give rise to the transformation of energy from one type to another — an issue that has concerned ecologists for decades. (See [6] for a discussion of the ways and degree to which these concerns have been expressed in the individual-based models of Ecology.)

Simplified models of energetics leave questions concerning trophic levels, productivities etc. essentially untouched. Simulations that incorporate low-level mechanisms for the storage and retrieval of energy by manufacture and dissociation of various chemical bonds allow for the emergence of interactions relating to agent short-term resource collection and utilization. They also provide a means to study the evolution of these survival strategies. In addition, they allow study of modes and pathways of energy transformation at the level of the ecosystem.

Matter

Matter in A-Life's Virtual Ecosystems

Since Tansley's paper, ecosystems have been identified as entities whose significant components are biotic *and* abiotic. Odum later made it clear that the cycles of materials through ecosystems were significant in shaping them as a whole, even in defining them. This interest in biogeochemical cycles has not shaped the history of Ecosystem Ecology as closely as has energetics, but it provides a valuable means to understand ecosystems [10]. For instance: knowledge of the similarities and differences between organisms' chemical composition can inform us about their evolution and their relationships to specific habitats; an organism's ability to extract elements from complex materials to construct biomass is an important aspect of its physiology and morphology; organisms excrete waste and accumulate matter in their biomass, impacting on the relative abundance of materials in the environment and its suitability as habitat for other species. In all of these ways and many others, the properties of matter play *the* dominant role.

Some of A-Life's simulations incorporate impassable barriers [18] or explore collective construction using abiotic building material [25]. These materials play no role in the construction of agent bodies and therefore do not constitute an ecologically relevant model of matter. For a model to assist exploration in the areas listed above, it must include mechanisms for combining simple materials into more complex aggregates, and for their dissociation — an artificial chemistry. Models lacking this facility preclude the emergence of organism strategies for biosynthesis and decay, a field of study that has frequently been addressed within Ecology. Biogeochemical cycles just haven't been the focus of A-Life's virtual ecosystems.

When A-Life writers say their agents have a model "metabolism" all many of them mean is that agents persist by decrementing an energy counter at each time step (e.g., *PolyWorld*). This is a poor model of metabolism! There is usually no transformation of matter from "food" (itself a high-level abstraction) to different types of biomass or waste. Indeed, waste is hardly ever mentioned. Exceptions to this trend have appeared, including the artificial chemistry-based model ecosystems EVOLVE IV used to examine the impact of waste accumulation on mobile agents [26], and the system we proposed (Dorin and Korb) [24], mentioned above.

Our model includes decomposition as one of many chemical transformations that emerge naturally from the artificial chemistry. Patterns of material cycling through the biota and abiotic environment likewise emerge from this simulation.

Forrest and Jones' simulation [2] allows for simple material cycling through agent bodies. Materials are collected by the agents and stored for a time before being released back into the environment when the agent dies. However these materials are not used in any way to construct complex aggregates and there is no way for decomposer agents to emerge.

Without low-level representations of the type found in A-Life's artificial chemistries (such as those surveyed in [27]), virtual ecosystems cannot tell us about the emergence of chemical cycles between the abiotic environment and the biota. As we have noted, since Odum the emergence of these cycles and their impact on ecosystem behaviour has been a significant, defining issue in Ecology. In addition, as discussed in the previous section, the mechanism by which organisms store and extract energy is through the transformation of matter. Hence, the two birds of matter and energy can be killed with the single stone of artificial chemistry. Without this stone, A-Life's virtual ecosystems realise only a (small) fraction of their potential.

Evolution

Evolution in A-Life's Virtual Ecosystems

If there are two widespread traits of A-Life's virtual ecosystems in addition to a lack of interest in energy and matter transformation, perhaps they are these: (i) a focus on studying the general properties of the evolutionary process itself, (ii) the use of the artificial evolutionary process to select parameters that allow agents to meet the requirements of existence and replication in a dynamic virtual environment.

Ray's *Tierra* was constructed, "with hand-crafted organisms already capable of replication and open-ended evolution... to generate increasing diversity and complexity" [4]. The evolutionary process was needed to drive change and maintain order in *Tierra* as much as it was the focus of the researcher's attention. Ray's system was intended to model replicators and their interactions within the digital medium, not as representatives of particular biological organisms or their habitats. Ten years later, Ray's focus continued to be the evolutionary process, his aim being to, "use organic life as a source of ideas on how to create a richer evolutionary process in the digital medium" [28].

Yaeger states as one of his aims for *PolyWorld*, "create artificial life that is as close as possible to real life, by combining as many critical components of real life as possible in an artificial system" [18]. Like Ray, he includes in his essential traits of living systems their ability to replicate and evolve. Although his study is a level of abstraction above *Tierra* (his virtual organisms have behavioural primitives including the ability to eat, mate, move, see etc.), and is more literal in its representation of agents and space, still the system is generic — it aims to tell the researcher about the evolution of behaviour appropriate to *PolyWorld* agents in the virtual world, rather than about the behaviour of specific species in real habitats.

Packard's *Bugs* [3] are similar in some respects to agents from *PolyWorld*. The independent bug agents roam a space and search for food that allows them to sustain themselves and reproduce.

Forrest and Jones write of the *Echo* platform, "Echo is intended to capture important generic properties of ecological systems, and not necessarily to model any particular ecology in detail" [2]. They go on to explain that if their software could be correctly validated, it could inform us about the impact of evolution on ecosystem behaviour. For them this is perhaps the most important contribution such models may make.

There are countless much more recent examples of artificial evolution implemented in virtual ecosystems to study the evolution of altruism, parental investment, group selection, aging, epidemiology etc. (e.g., [21, 29, 30]). This path is well trodden by A-Life researchers.

Why incorporate artificial evolution?

In addition to the obvious and widely stated interest in studying the evolutionary process itself, the difficulty in writing and manually parameterising software to establish ongoing, dynamic but stable relationships between software agents in a virtual space motivates the simulation of evolution. The Genetic Algorithm (GA) has seen wide application for the solution of optimization problems and this motivates its use in virtual ecosystems also. Handcrafted systems may operate dynamically for a while, however adjustment of parameter values is needed to keep agents from extinction and virtual ecosystems from collapse.

Returning to a familiar example, Yaeger implemented in *PolyWorld* a mode to kick-start the evolutionary process and overcome the problem of parameter specification for creating viable agents. At first he imposes an external fitness function to optimise agents until they evolve parameters to successfully sustain themselves and replicate independently.

Similarly, Ray seeds *Tierra* with a hand-designed replicator. Evolution then takes over, altering the copying algorithm and enhancing its efficiency. Again, evolution optimises the agents in the current, dynamic environment.

The implications of employing digital evolution

Artificial evolution is currently limited in its ability to accurately mimic the detailed behaviour of real evolution. The application of artificial evolution as an instigator of agent change in A-Life simulations, without matching agent parameters to field data, further ensures that the simulations of ecosystems remain generic. In contrast, especially during the period to 1999 [6], Ecology's agent-based models have tended to be non-evolutionary and focused on modelling specific aspects of specific species (for instance the behaviour and resulting distribution of lynx [31]). Typical A-Life models are of limited relevance to this form of Ecology.[2] In the context of conventional ecological modelling, evolutionary algorithms are not sufficiently true-to-life to replicate the evolutionary path that gave rise to specific creatures. Additionally, in many simulations the agent and environmental models capture only a few traits of the real systems they represent. All of these simplifications to models of highly sensitive, non-linear systems ensure that whilst digital evolution may optimise agent parameters for survival in virtual environments, the values of these parameters will quite likely *not* reflect values that may be obtained from field data. In fact, the parameters often do not reflect any aspect of any real system [2]. Obviously, this does not render existing models useless — it merely renders them inapplicable where ecological problems need to be addressed concerning the evolution of specific species in specific habitats.

Pattern-Oriented Modelling

Artificial Evolution for Pattern-Oriented Modelling

Despite the tenuous connections between the specifics of artificial and real evolution, we suggest that the former may be included in ecological models for the benefit of Ecology and A-Life alike, by assisting in the process of *Pattern-Oriented Modelling* (POM) [11, 32]. POM is a general strategy for developing, validating, and parameterising ecological models: "In POM, we explicitly follow the basic research program of science: the explanation of observed patterns. Patterns are defining characteristics of a system and often, therefore, indicators of essential underlying processes and structures. Patterns contain information on the internal organization of a system, but in a *coded* form. The purpose of POM is to *decode* this information" ([11], p. 987). Patterns can thus also be viewed as *regularities* or *signals*; in economics, often the similar notion of *stylized facts* is used [33].

Multiple patterns observed in real systems are used in POM in three ways: they indicate which state variables a model should have so that in principle the same patterns can emerge in the model; they are used to select the most appropriate alternative sub-model representing a certain process; and they can be used to determine entire sets of unknown parameters (i.e. for *inverse modelling*).

Basically, POM reminds system modellers of the very basic principle of science: if we focus on single patterns to build a model we risk producing one that provides a good fit to data for the wrong reasons (i.e., mechanisms) instead of a model that is structurally realistic and captures the essentials of the system's organisation. Thus, the more pattern matches between a model and field data and the more variety these exhibit, the more certain the correspondence between model and reality.

Ideally, simulation patterns match at multiple levels of organization to the system being modelled. For example, patterns might be matched at the level of an individual organism's behaviour, the dynamics of a population, and at the level of an entire ecosystem. Commonly cited A-Life models (for instance models of animal flocking such as Reynolds' *boids* [34]) do not seek this *multi-level* pattern match since their aim is to locate *any* parameters or agent behaviours that produce a desired global behaviour. Yet if A-Life models are to be relevant to Ecology, Ecologists rightfully insist that multi-level pattern matches are required for validation.

Pattern-Oriented Modelling of Evolutionary Processes

As described earlier, artificial evolution may optimise agent parameters in virtual ecosystems enabling agents to survive and replicate. If these models also include low-level representations of energetics and matter, the simulations may inform us about *general* principles of ecosystem behaviour of relevance to Ecology, including, for example, how *physical ecosystem engineering*[3] impacts on habitats, the number of niches, the number of trophic levels in an ecosystem, species diversity, ecosystem resilience and stability etc. over evolutionary time periods. Data regarding the impact of physical ecosystem engineers on ecosystems is available [35-37] and could form the basis for pattern-oriented modelling, even if the specifics of species fell beneath the level of abstraction of the model. To the authors' knowledge no generic simulations exploring these properties of ecosystems over evolutionary time periods have yet been devised. In fact, we believe that as yet there are no simulations investigating the emergence of ecosystem engineers at all. Do they evolve readily and under what circumstances? Is there a basic organizational property of ecosystems that requires them?

This domain is ideally suited to a simulation of the form we describe, since it requires detailed models of the transformation of the abiotic and biotic environment and its impact on the evolution of biota.

[2] In Systems Ecology (as in A-Life) the need to model specific species and habitats does not apply (see above). Grimm has also argued that the general principles of ecosystem behaviour *are* (or at least *should be*) of interest to Ecologists [6].

[3] To varying extents, all organisms are physical ecosystem engineers. Physical ecosystem engineers physically alter the biotic or abiotic environment and thereby control or modulate the availability of resources to (or forces acting on) other organisms. These physical changes destroy, maintain or create habitat for other organisms [36]. Their presence is often a key factor in ecosystem behaviour. A tree is an example of a significant physical ecosystem engineer: it provides habitat for mosses, insects and birds; its roots trap soil and leaf matter, altering the impact of wind and water erosion; its branches harbour larvae or tadpoles within pools etc. Coral produces reefs, wombats dig holes, lyrebirds and blackbirds sift leaf litter. These species (and humans!) are physical ecosystem engineers that have a large impact on organisms around them.

Pattern-Oriented Modelling with Evolutionary Parameter Optimisation

POM has already been applied within ecology to validate models of specific species against their real counterparts. For instance, the technique matched the behaviour of a specific species (lynx) to an agent-based model [31]. Whereas in the POM of evolutionary processes we proposed to ignore the specific behaviours at the level of the individual organism (these fall beneath the level of detail incorporated into a generic simulation), in this case the simulation is intended to be accurate in its specific detail from the level of the individual agent up to the level of the ecosystem. In this second method, the aim is not just to find general parameters that create a virtual ecosystem that behaves like a real one, but also to locate agent parameters matching observed organism patterns so that the model correctly reflects the behaviour of real species.

At the point where parameters for agents give rise to realistic patterns of behaviour, and this has been shown to give rise to patterns matching the behaviour of the ecosystem, artificial evolution can be *disengaged*. It has served its purpose as a means for automatically parameterising the model. The simulation can then be used to answer ecological questions by researchers who wish (for whatever reason!) to disregard evolutionary effects.

If this approach is to succeed, the pattern-matching process must be automated so that evolution (which acts *within* the virtual ecosystem in this instance) can play its role as an optimization algorithm. Field data at various ecological levels must be available for this process to work. We identified (above) studies of physical ecosystem engineers as a source of data for such an implementation.

A difficulty may arise when attempting to match the behaviour of several simulated species in the virtual ecosystem simultaneously, and without allowing the dynamical whole of which they are a part to collapse. A mode in which an externally defined fitness function operates (akin to that discussed above for *PolyWorld*) may assist when virtual ecosystem collapse appears imminent, by ensuring population levels, diversity or some other measure of virtual ecosystem state does not fall outside a specified range.

There is evidence in the ecological literature that a GA is capable of correctly parameterising a model based on field data. Strand et al have successfully used this technique to generate parameters for a model of fish behaviour [38], suggesting not only the feasibility of the approach, but also its acceptability to Ecologists.

Discussion and Conclusions

We do not expect there to be significant debate concerning our position on the abundance or value of evolution-focussed A-Life virtual ecosystems. However we have also argued the position that while some examples of individual-based simulations incorporating energetics and matter have been created, this area is under-explored in A-Life's virtual ecosystems, even whilst it may be a component of the field's artificial chemistry models. Our overview of several early models from the A-Life literature highlights the initial trends in the field but as we have seen, even the updates to these models and many recent virtual ecosystems are blinkered when it comes to energy and matter. This state of affairs must be rectified if our models are to be of relevance to Ecology, as it has been traditionally understood, and as it is now being practiced.

By incorporating low-level models of the three frameworks of energetics, matter and evolution into virtual ecosystems we open significant pathways for ecologically relevant research in A-Life, particularly in domains where the interactions of the abiotic and biotic play a role. To our minds, and those of Reiners [10], Loehle and Pechmann [9], the interactions of these frameworks is a core concern.

We have indicated two strategies for employing these frameworks within virtual ecosystems. The first of these requires the application of the evolutionary process in much the same way as it has been traditionally applied within A-Life: as a means to dynamically adjust agent parameter values to support their viability and reproduction within the virtual environment. These simulations will allow us to determine generic properties of ecosystems only if the models are validated against field data employing a technique such as pattern-oriented modelling. This must ensure multi-level correspondence between simulation and reality, even if the level of abstraction of the model does not reach down to the detailed simulation of specific real species.

The second approach we suggest employs artificial evolution to match simulation patterns against data gathered from the level of specific species up to data concerning specific ecosystems. Once the parameters of the system have been optimised so as to reproduce the patterns observed in field data, the evolution algorithm is turned off. The model may then be employed to answer questions relating to the specific ecosystem and species that it represents. Unfortunately it may not then be used to study the evolution of these specific species in specific environments. This is a shortcoming of the artificial evolution algorithm (it does not model real evolution in detail) that would be worth overcoming.

It remains to be seen whether artificial evolution will be suited to pattern matching across several levels of an entire dynamical system consisting of hundreds of components. However, the widespread success of the algorithm in the solution of complex optimization problems bodes well for the approach. Additionally, the technique has already been employed successfully in Ecology to match a single species' behaviour to that of software agents and we have indicated some preliminary strategies that may assist where multiple species are interacting in a simulation. The value of correctly parameterised, detailed models of ecosystems ensures the importance of attacking this problem in the future.

References

1. Brewster, J.J. and M. Conrad. *Computer Experiments on the Development of Niche Specialization in an Artificial Ecosystem*. In *Congress on Evolutionary Computation*. 1999. Washington DC: IEEE: p. 444-451.
2. Forrest, S. and T. Jones. *Modelling Adaptive Systems with Echo*. In *Complex Systems: Mechanisms of Adaptation*. 1994: IOS Press: p. 3-21.

3. Packard, N.H. *Intrinsic Adpatation in a Simple Model for Evolution*. In *Artificial Life, SFI Studies in the Sciences of Complexity*. 1988: Addison-Wesley: p. 141-155.
4. Ray, T.S. *An approach to the synthesis of life*. In *Artificial Life II*. 1990. Santa Fe, New Mexico: Addison Wesley: p. 371-408.
5. Conrad, M. and H.H. Pattee, *Evolution Experiments with an Artificial Ecosystem.* Journal of Theoretical Biology, 1970. **28**: p. 393-409.
6. Grimm, V., *Ten years of individual-based modelling in ecology: what have we learned and what could we learn in the future?* Ecological Modelling, 1999. **115**: p. 129-148.
7. DeAngelis, D.L. and W.M. Mooij, *Individual-based modelling of ecological and evolutionary processes.* Annu. Rev. Ecol. Evol. Syst., 2005. **36**: p. 147-168.
8. Grimm, V. and S.F. Railsback, *Individual-based Modeling and Ecology*. Princeton Series in Theoretical and Computational Biology. 2005: Princeton University Press. 428.
9. Loehle, C. and J.H.K. Pechmann, *Evolution: The Missing Ingredient in Systems Ecology.* American Naturalist, 1988. **132**(6): p. 884-899.
10. Reiners, W.A., *Complementary Models for Ecosystems.* The American Naturalist, 1986. **127**(1): p. 59-73.
11. Grimm, V., et al., *Pattern-Oriented Modelling of Agent-Based Complex Systems: Lessons from Ecology.* Science, 2005. **310**: p. 987-991.
12. Golley, F.B., *A History of the Ecosystem Concept in Ecology: more than the sum of the parts*. 1993, New Haven and London: Yale University Press. 254.
13. Clements, F.E., *Preface*, in *Plant Succession: An Analysis of the Development of Vegetation*. 1916, Carnegie Institution of Washington: Washington D.C. p. 1-7.
14. Phillips, J., *The Biotic Community.* The Journal of Ecology, 1931. **19**(1): p. 1-24.
15. Sukachev, V. and N. Dylis, *Fundamentals of Forest Biogeocoenology*. 1964, Edinburgh and London: Oliver and Boyd.
16. Bergandi, D., *Reductionist Holism: An oxymoron or a philosophical chimera of Eugene Odum's Systems Ecology?* Ludus Vitalis: Journal of Philosophy and Life Sciences 3, 1995: p. 145-178.
17. Ulanowicz, R.E., *Growth and development: ecosystems phenomenology*. 1986, New York: Springer-Verlag. 203.
18. Yaeger, L., *Computational Genetics, Physiology, Metabolism, Neural Systems, Learning, Vision and Behavior or Polyworld: Life in a New Context*, in *Artificial Life III*. 1992, Addison-Wesley: City. p. 263-298.
19. Lindgren, K. and M.G. Nordahl, *Cooperation and Community Structure in Artificial Ecosystems.* Artificial Life, 1994. **1**(1/2): p. 15-37.
20. Lindgren, K. and M.G. Nordahl. *Artificial Food Webs*. In *Artificial Life III*. 1994: Addison-Wesley: p. 73-104.
21. Dorin, A. *A Co-Evolutionary Epidemiological Model for Artificial Life and Death*. In *8th European Conference on Artificial Life*. 2005. Canterbury, UK: Springer: p. 775-784.
22. Cooper, T.F. and C. Ofria. *Evolution of Stable Ecosystems in Populations of Digital Organisms*. In *Artificial Life VIII*. 2003. Sydney, Australia: MIT Press: p. 227-232.
23. Taylor, T. and J. Hallam. *Studying Evolution with Self-Replicating Computer Programs*. In *Fourth European Conference on Artificial Life*. 1997: MIT Press: p. 550-559.
24. Dorin, A. and K. Korb. *Building Artificial Ecosystems from Artificial Chemistry*. In *9th European Conference on Artificial Life*. 2007. Lisbon: Springer-Verlag: p. 103-112.
25. Theraulaz, G. and E. Bonabeau, *Modelling the Collective Building of Complex Architectures in Social Insects with Lattice Swarms.* Journal of Theoretical Biology, 1995. **177**(4): p. 381-400.
26. Brewster, J.J., R.G. Reynolds, and M.A. Brockmeyer, *Not In My Backyard: A simulation of the effects of agent mobility on environmental poisoning*, in *Proceedings of the 2002 Congress on Evolutionary Computing*. 2002: City. p. 849-854.
27. Dittrich, P., J. Ziegler, and W. Banzhaf, *Artificial Chemistries - A Review.* Artificial Life, 2001. **7**(3): p. 225-276.
28. Ray, T.S. and J.F. Hart. *Evolution of Differentiation in Multithreaded Digital Organisms*. In *Artificial Life VII, Proceedings of the Seventh International Conference on Artificial Life*. 2000: MIT Press: p. 132-140.
29. Mascaro, S., K.B. Korb, and A.E. Nicholson. *ALife Investigation of Parental Investment in Reproductive Strategies*. In *Eighth International Conference on Artificial Life (ALife VIII)*. 2003: MIT Press: p. 358-361.
30. Woodberry, O., K.B. Korb, and A.E. Nicholson, *The Evolution of Aging*, in *Proceedings of the Australian Conference on Artificial Life (ACAL 2005)*. 2005: City. p. 319-333.
31. Kramer-Schadt, S., et al., *Patterns for parameters in simulation models.* Ecological Modelling, 2007. **204**: p. 553-556.
32. Grimm, V. and U. Berger, *Seeing the forest for the trees, and vice versa: pattern-oriented ecological modelling*, in *Handbook of scaling methods in aquatic ecology: measurement, analysis, simulation*, L. Seuront and P.G. Strutton, Editors. 2003, CRC Press: Boca Raton. p. 411-428.
33. Kaldor, N., *Capital Accumulation and Economic Growth*, in *The Theory of Capital, Reprint*, F. Lutz and D. Hague, Editors. 1961/1968, Macmillan: London. p. 177-222.
34. Reynolds, C.W., *Flocks, Herds and Schools: A Distributed Behavioural Model.* Proceedings of SIGGRAPH '87, 1987. **21**(4): p. 25-34.
35. Badano, E.I. and L.A. Cavieres, *Ecosystem engineering across ecosystems: do engineer species sharing common*

features have generalized or idiosyncratic effects on species diversity? Journal of Biogeography, 2006. **33**: p. 304-313.

36. Gutiérrez, J.L. and C.G. Jones, *Physical Ecosystem Engineers as Agents of Biogeochemical Heterogeneity*. BioScience, 2006. **56**(3): p. 227-237.
37. Jones, C.G., J.H. Lawton, and M. Shachak, *Positive and negative effects of organisms as physical ecosystem engineers.* Ecology, 1997. **78**(7): p. 1946-1957.
38. Strand, E., G. Huse, and J. Giske, *Artificial Evolution of Life History and Behavior*. The American Naturalist, 2002. **159**(6): p. 624-644.