

Open Dimensional Dynamical Systems

Russell K. Standish
High Performance Computing Support Unit
The University of New South Wales,
Sydney, 2052

March 30, 2000

Abstract

The theme of this issue of *Artificial Life* is that traditional Dynamical Systems theory is inadequate for many problems in the study of Complex Systems. In this paper, I argue that the finite dimensional nature of traditional dynamical systems is inappropriate for evolutionary systems, but that simply embedding the dynamics into an infinite dimensional space overly complicates the mathematics and obscures important features of the complex system. Instead, I believe a new branch of mathematics is needed, which I term *open dimensional dynamical systems*, systems whose dimensionality is finite at any point in time, but whose dimension varies through time.

1 Introduction

The theme of this issue of *Artificial Life* is that traditional Dynamical Systems theory is inadequate for many problems in the study of Complex Systems. In this paper, I argue that the finite dimensional nature of traditional dynamical systems is inappropriate for evolutionary systems, but that simply embedding the dynamics into an infinite dimensional space overly complicates the mathematics and obscures important features of the complex system. Instead, I believe a new branch of mathematics is needed, which I term *open dimensional dynamical systems*, systems whose dimensionality is finite at any point in time, but whose dimension varies through time.

For concreteness, I will discuss a model evolutionary system that runs in a simulator called *Ecolab* (and called the Ecolab model[16]). Another couple of models have also been examined, a generalised Newman model[12] and an economics model based of von Neumann Technology, called *Econolab*[11, 13]. The Ecolab model is archetypal of these models.

One of the first concepts to disappear in open dimensional systems is that of the *attractor*. An attractor in a finite dimensional dynamical system is a stable set of points, i.e. has an open neighbourhood whose trajectories converge

to the attractor. However, if there is the possibility for arbitrary degrees of freedom to be added to the system at any time, then a degree of freedom that destabilises the attractor will be added at some random point in the system's history. Instead, a new concept appears to be relevant, that of the *pattern attractor*[4]. This pattern is reflected in the statistics of the process of jumping from one fleeting attractor basin to the next. In the case of the Ecolab model, the long term behaviour is to self-organise to a critical state, giving rise to power law distributions of species lifetimes.

2 Ecolab and the Dynamics of Evolution

This section outlines a model of an evolving ecology[16, 10] based on a generalised Lotka-Volterra equation, which is perhaps the simplest ecological model to use.

$$\dot{n}_i = r_i n_i + \sum_{j=1}^{n_{sp}} \beta_{ij} n_i n_j \quad (1)$$

Here \mathbf{r} is the difference between the birth rate and death rate for each species, in the absence of competition or symbiosis. β is the interaction term between species, with the diagonal terms referring to the species' self limitation, which is related in a simple way to the carrying capacity K_i for that species in the environment by $K_i = -r_i/\beta_{ii}$. In the literature (eg Strobeck[17], Case[3]) the interaction terms are expressed in a normalised form, $\alpha_{ij} = -K_i/r_i\beta_{ij}$, and $\alpha_{ii} = 1$ by definition. \mathbf{n} is the species density.

These equations are simulated on a simulator called *Ecolab*. [15] The vectors \mathbf{n} and \mathbf{r} are stored as dynamic arrays, the size of which (i.e. the system dimension) can change in time. The interaction array is stored in row/column sparse form, consisting of the four arrays `diag`, `val`, `row` and `col`. Equation (1) can be written as:

```
tmp[row] = beta.val * n[beta.col];
n += (r + beta.diag + tmp) * n;
```

This code makes up the `generate` operator in the Ecolab system. Other operators include `compact`, which removes species that have become extinct from the system (to optimise computational performance) and `mutate`, which adds a certain number of new species to the system, according to a specific algorithm to be discussed later. The operators can be called from a scripting language called TCL[9], that allows different types of experiments to be performed without recompiling the code.

Before discussing the mutation algorithm in more detail, equation (1) must be analysed to determine the conditions β must satisfy for the system to be real, and also to determine the different regimes of dynamics, from the linear (stable equilibrium) case, to limit cycles and chaos to the actual breakdown of the ecosystem.

2.1 Linear Analysis

Linear analysis starts with the fixed point of equation (1)

$$\hat{\mathbf{n}} = -\boldsymbol{\beta}^{-1}\mathbf{r}, \quad (2)$$

where $\dot{\mathbf{n}} = 0$. There is precisely one fixed point in the interior of the space of population densities (i.e. \mathbf{n} such that $n_i > 0$) provided that all components of $\hat{\mathbf{n}}$ are positive, giving rise to the following inequalities:

$$\hat{n}_i = (\boldsymbol{\beta}^{-1}\mathbf{r})_i > 0, \quad \forall i \quad (3)$$

This interior space is denoted $\mathbb{R}_+^{n_{\text{sp}}}$ mathematically.

There may also be fixed points on the boundary of $\mathbb{R}_+^{n_{\text{sp}}}$, where one or more components of \mathbf{n} are zero (corresponding to an extinct species). This is because the subecology with the living species only (i.e. with the extinct species removed) is equivalent to the full system.

The *stability* of this point is related to the negative definiteness of derivative of $\dot{\mathbf{n}}$ at $\hat{\mathbf{n}}$. The components of the derivative are given by

$$\frac{\partial \dot{n}_i}{\partial n_j} = \delta_{ij} \left(r_i + \sum_k \beta_{ik} n_k \right) + \beta_{ij} n_i \quad (4)$$

Substituting eq (2) gives

$$\left. \frac{\partial \dot{n}_i}{\partial n_j} \right|_{\hat{\mathbf{n}}} = -\beta_{ij} (\boldsymbol{\beta}^{-1}\mathbf{r})_i \quad (5)$$

Stability of the fixed point requires that this matrix should be negative definite. Since the $(\boldsymbol{\beta}^{-1}\mathbf{r})_i$ are all negative by virtue of (3), this is equivalent to $\boldsymbol{\beta}$ being negative definite, or equivalently, that its n_{sp} eigenvalues all have negative real part. Taken together with the inequalities (3), this implies that $2n_{\text{sp}}$ inequalities must be satisfied for the fixed point to be stable. This point was made by Strobeck[17], in a slightly different form. (Note that Strobeck implicitly assumes that $\sum_i r_i \hat{n}_i / K_i > 0$, so comes to the conclusion that $2n_{\text{sp}} - 1$ conditions are required.) If one were to randomly pick coefficients for a Lotka-Volterra system, then it has a probability of $4^{-n_{\text{sp}}}$ of being stable, i.e. one expects ecosystems to become more unstable as the number of species increases[8].

2.2 Permanence

Whilst stability is a nice mathematical property, it has rather less relevance when it comes to real ecologies. For example the traditional predator-prey system studied by Lotka and Volterra has a limit cycle. The fixed point is decidedly unstable, yet the ecology is *permanent* in the sense that both species' densities are larger than some threshold value for all time. Hofbauer et al. [5] and Law and Blackford[7] discuss the concept of *permanence* in Lotka-Volterra

systems, which is the property that there is a compact absorbing set $\mathcal{M} \subset \mathbb{R}_+^{n_{\text{sp}}}$ *i.e* once a trajectory of the system has entered \mathcal{M} , it remains in \mathcal{M} . They derive a sufficient condition for permanence due to Jansen[6] of the form:

$$\sum_i p_i f_i(\hat{\mathbf{n}}_B) = \sum_i p_i (r_i - \sum_j \beta_{ij} \hat{n}_{Bj}) > 0, \quad \exists p_i > 0 \quad (6)$$

for every $\hat{\mathbf{n}}_B$ equilibrium points lying on the boundary ($\hat{n}_{Bi} = 0 \quad \exists i$), provided the system is *bounded* (or equivalently *dissipative*). This condition is more general than stability of the equilibrium — the latter condition implies that a local neighbourhood of the equilibrium is an absorbing set. Also, the averaging property of Lotka-Volterra systems implies that the equilibrium must lie in the positive cone $\mathbb{R}_+^{n_{\text{sp}}}$. So (3) must still hold for permanence.

Consider the boundary points $\hat{\mathbf{n}}_B$ that are missing a single species i . Then Jansen's condition for these boundary points is

$$r_i - \sum_j \beta_{ij} \hat{n}_{Bj} > 0. \quad (7)$$

This set of conditions is linearly independent. Let the number of such boundary points be denoted by $n_B \leq n_{\text{sp}}$. Then the set of conditions (6) will have rank $n_B \leq \nu \leq n_{\text{sp}}$ (the number of linearly independent conditions), so the system has at most probability $2^{-n_{\text{sp}}-\nu}$ of satisfying Jansen's permanence condition if the coefficients are chosen uniformly at random. As stability is also sufficient for permanence, the probability lies between $4^{-n_{\text{sp}}}$ and $2^{-n_{\text{sp}}-\nu}$.

Another rather important property is *resistance to invasion*. [3] Consider a boundary equilibrium $\hat{\mathbf{n}}_B$. If it is proof against invasion from the missing species, then the full system cannot be permanent. For the boundary points that miss a single species, this implies that condition (7) is necessarily satisfied for permanence, along with (3). The probability of permanence is then bounded above by $2^{-n_{\text{sp}}-n_B}$.

Thus whilst a randomly selected ecology is more likely to be permanent than to have a stable equilibrium, the likelihood decreases exponentially with increase in species number.

2.3 Boundedness

It is necessary that the ecology be *bounded*, ie that $\sum n_i < N \quad \exists N, \quad \forall t > 0$. This requires

$$\sum_i \dot{n}_i = \mathbf{r} \cdot \mathbf{n} + \mathbf{n} \cdot \boldsymbol{\beta} \mathbf{n} < 0, \quad \forall \mathbf{n} : \sum_i n_i > N \quad \exists N \quad (8)$$

As \mathbf{n} becomes large in any direction, this functional is dominated by the quadratic term, so this implies that

$$\mathbf{n} \cdot \boldsymbol{\beta} \mathbf{n} \leq 0 \quad \forall \mathbf{n} : n_i > 0. \quad (9)$$

If strict equality holds, then $\mathbf{r} \cdot \mathbf{n} < 0$. Negative definiteness of β is sufficient, but not necessary for this condition. Another sufficient condition is to require $\forall i, j, \beta_{ii} < 0$ and $\beta_{ij} + \beta_{ji} + \beta_{ii} \leq 0$, which is used in the current study. This condition is satisfied by the Predator-Prey equations, and so allows multi-trophic systems to be built. Its main advantage is its simplicity of implementation, along with the range of interesting (i.e. non limit point) behaviour it encompasses.

2.4 Mutation

Adding mutation involves adding an additional operator to equation (1)

$$\dot{\mathbf{n}} = \mathbf{r} * \mathbf{n} + \mathbf{n} * \beta \mathbf{n} + \text{mutate}(\mu, \mathbf{r}, \mathbf{n}) \quad (10)$$

where $*$ refers to elementwise multiplication.

The mutation operator must generate new degrees of freedom $i > n_{\text{sp}}$ (where n_{sp} is the number species currently in the ecology), somehow defining the new ecological coefficients $\{r_i | i > n_{\text{sp}}\}, \{\beta_{ij} | i > n_{\text{sp}} \text{ or } j > n_{\text{sp}}\}$ from the previous state of the system. In reality, there is another layer (hidden in equation (1)) called the genotypic layer, where each organism has a definite genotype. There is a specific map from the genotypic layer to the space of ecological coefficients (hereafter called the phenotypic layer) called the *embryology*. Then the mutation operator is a convolution of the genetic algorithm operations operating at the genotypic layer, with the embryology.

A few studies, including Ray’s Tierra world, do this with an explicit mapping from the genotype to some particular organism property (e.g. interpreted as machine language instructions, or as weight in a neural net). These organisms then interact with one another to determine the population dynamics. In this model, however, we are doing away with the organismal layer, and so an explicit embryology is impossible. The only possibility left is to use a statistical model of embryology. The mapping between genotype space and the population parameters \mathbf{r}, β is expected to look like a rugged landscape, however, if two genotypes are close together (in a Hamming sense) then one might expect that the phenotypes are likely to be similar, as would the population parameters. This I call *random embryology with locality*. Here, we tend to idealise genotypes as bit strings, although strings over an arbitrary alphabet (eg the four DNA bases ACGT) can equally be considered.¹

In the simple case of point mutations, the probability $P(x)$ of any child lying distance x in genotype space from its parent follows a Poisson distribution, as this is the distribution of the number of bit flips, or deletions that might occur with a point mutation. Random embryology with locality implies that the phenotypic parameters are distributed randomly about the parent species, with a standard deviation that depends monotonically on the genotypic displacement. The simplest such model is to distribute the phenotypic parameters in a Gaussian fashion about the parent’s values, with standard deviation proportional

¹The Hamming distance is the number of bits (bases) that differ between the two strings. So for example if a single bit has been removed from one string, the Hamming distance is one.




Figure 1: Probability distribution of the relation between genotype difference and the corresponding phenotype difference

to the genotypic displacement. This constant of proportionality can be conflated with the species' intrinsic mutation rate, to give rise another phenotypic parameter μ . It is assumed that the probability of a mutation generating a previously existing species is negligible, and can be ignored. We also need another arbitrary parameter ρ , "species radius", which can be understood as the minimum genotypic distance separating species, conflated with the same constant of proportionality as μ .

We may represent the Ecolab embryology as a probability distribution $f(p, g) = \sqrt{\frac{2}{\pi}} \frac{\mu e^{-\left(\frac{\mu p}{2g}\right)^2}}{g}$, where $p = |r_i - r_j|/|r_i|$ or $p = |\beta_{ik} - \beta_{jk}|/|\beta_{ik}|$ is the distance between two species' phenotypic parameters, and g is the difference between the two genotypes. Figure 1 shows the general form of this probability distribution.

Figure 2 shows the probability distribution of a mutant phenotypical coefficient about that of its parent's value. This is given by

$$\int_0^\infty \sqrt{\frac{2}{\pi}} \frac{e^{-g/\mu - \left(\frac{\mu p}{2g}\right)^2}}{g} \mu dg. \quad (11)$$

In summary, the mutation algorithm is as follows:

1. The number of mutant species arising from species i within a timestep is $\mu_i r_i n_i / \rho$. This number is rounded stochastically to the nearest integer, e.g. 0.25 is rounded up to 1 25% of the time and down to 0 75% of the time.




Figure 2: The probability distribution of a mutant phenotypical coefficient about that of its parent's value. This has been calculated by numerical integration from formula (11). Note that the curve actually diverges at 0.

2. Roll a random number from a Poisson distribution $e^{-x/\mu+\rho}$ to determine the standard deviation σ of phenotypic variation.
3. Vary \mathbf{r} according to a Gaussian distribution about the parents' values, with σr_0 as the standard deviation, where r_0 is the range of values that \mathbf{r} is initialised to, ie $r_0 = \max_i r_i|_{t=0} - \min_i r_i|_{t=0}$.
4. The diagonal part of $\boldsymbol{\beta}$ must be negative, so vary $\boldsymbol{\beta}$ according to a log-normal distribution. This means that if the old value is β , the new value becomes $\beta' = -\exp(-\ln(\beta) + \sigma)$. These values cannot arbitrarily approach 0, however, as this would imply that some species make arbitrarily small demands on the environment, and will become infinite in number. In Ecolab, the diagonal interactions terms prevented from becoming larger than $-r/(.1 * \text{INT_MAX})$, where r is the corresponding growth rate for the new species.
5. The off diagonal components of $\boldsymbol{\beta}$, are varied in a similar fashion to \mathbf{r} . However new connections are added, or old ones removed according to $[1/p]$, where $p \in (-2, 2)$ is chosen from a uniform distribution. The values on the new connections are chosen from the same initial distribution that the off diagonal values where originally set with, ie the range $\min_{i \neq j} \beta_{ij}|_{t=0}$ to $\max_{i \neq j} \beta_{ij}|_{t=0}$. Since condition (9) is computationally expensive, we use a slightly stronger criterion that is sufficient, computationally tractable yet still allows "interesting" non-definite matrix behaviour namely that the sum $\beta_{ij} + \beta_{ji}$ should be non positive.
6. $\boldsymbol{\mu}$ must be positive, so should evolve according to the log-normal distribution like the diagonal components of $\boldsymbol{\beta}$. Similar to $\boldsymbol{\beta}$, it is a catastrophe to allow $\boldsymbol{\mu}$ to become arbitrarily large. In the real world, mutation normally exists at some fixed background rate — species can reduce the level of mutation by improving their genetic repair algorithms. In Ecolab, this ceiling on $\boldsymbol{\mu}$ is given by the `mutation(random,maxval)` variable.

2.5 Typical Results

Figure 3 shows the time behaviour for the number of species in the ecosystem for a typical run. The phenotypic parameters were seeded randomly in the ranges $-0.005 \leq \mathbf{r} \leq 0.01$, $-5 \times 10^{-5} \leq \boldsymbol{\beta}_{\text{diag}} \leq -1 \times 10^{-4}$, $-0.001 \leq \boldsymbol{\beta}_{\text{offdiag}} \leq 0.001$ and $0 \leq \boldsymbol{\mu} \leq 0.09$. The \mathbf{r} and $\boldsymbol{\beta}$ values were chosen so that several hundred individuals will be supported in the case of a single species system, and the offdiagonal terms large enough to permit interesting interactions between species, but not so large that the system collapsed to zero immediately. ρ was set at 10^4 , which was chosen by examining the histogram of differences between all the species. If ρ was too small, then a species' mutant offspring would be too similar to its parent to be really a new species. This shows up as a peak at small separation values of the histogram, which shouldn't be there according to the law of competitive exclusion.




Figure 3: n_{sp} as a function of time step




Figure 4: Distribution of species lifetimes in Ecolab.

The system rapidly evolves to one of the fixed points (by a massive extinction event!) with a negative definite β . Over time, mutations build up in the system, decreasing the stability of the system. What then follows are periods of episodic extinctions, and system growth through speciation. This is an example of *self organised criticality*[1], and gives rise to power law behaviour.

Figure 4 shows the lifetime distribution for Ecolab which has a slope of -2 for lifetimes less than 100, but -1 for larger lifetimes. At still larger times ($\tau \gg 0.1/\mu$), the distribution turns over, decaying exponentially. Previously published versions of this graph [10] only show the smaller lifetime behaviour. This is a classic demonstration of a power law[2], where the small time behaviour actually corresponds to the “churn” of nonadaptive species. At the very large times ($o(1/\mu)$), maximally adaptive organisms will be replaced by other organism just by pure chance, hence one expects an exponential drop off in lifetimes greater than about $1/\mu$.

3 Conclusion

Ecolab, Econolab and the generalised Newman model are all open dimensional dynamic systems that share a common blue print, and whose behaviour is dominated in the long run by a balance between mutation increasing the system’s complexity, and the inherent instability of the complexity breaking the system into a simpler structure. They all share a common statistical structure of the long term dynamics (in Econolab’s case, no simulations have yet been per-

formed, so this is only an expectation). They also share the property that their dynamics are defined on a positive cone ($\mathbf{n} \in \mathbb{R}_+^{n_{\text{sp}}}$), which is another difference with traditional finite dimensional dynamic systems.

One might speculate whether other classes of open dimensional systems might be possible. Stephan Halloy² noted that natural complex systems fell into two classes, those that tended to a power law, and those that tended to a log-normal distribution. Any system he examined that departed from either of these two forms had some external “driving force” explaining the departure. An example might be the distribution of plant species in crop land — the skew being explained by modern intensive agricultural practices.

It is clear from this that much work needs to be done to formulate a mathematical theory underpinning this type of system. The Ecolab type model may well be a useful example model to explore such a theory.

References

- [1] P. Bak, C. Tang, and K. Wiesenfeld. Self-organised criticality. *Phys. Rev. A*, 38:364, 1988.
- [2] Per Bak and Kim Sneppen. Punctuated equilibrium and criticality in a simple model of evolution. *Phys. Rev. Lett.*, 71:4083, 1993.
- [3] T. J. Case. Invasion resistance, species build-up and community collapse in metapopulation models with interspecies competition. *Bio. J. Linnean Soc.*, 42:239–266, 1991.
- [4] Stephan Halloy. A theoretical framework for abundance distributions in complex systems. In Standish et al. [14]. also in *Complexity International*, 6 <http://www.csu.edu.au/ci>.
- [5] J. Hofbauer, V. Hutson, and W. Jansen. Coexistence for systems governed by difference equations of Lotka-Volterra type. *J. Math. Biol.*, 25:553–570, 1987.
- [6] W. Jansen. A permanence theorem for replicator and Lotka-Volterra systems. *J. Math. Biol.*, 25:411–422, 1987.
- [7] Richard Law and Jerry C. Blackford. Self-assembling food webs: A global viewpoint of coexistence of species in Lotka-Volterra communities. *Ecology*, 73:567–578, 1992.
- [8] Robert M. May. *Stability and Complexity in Model Ecosystems*. Princeton University Press, Princeton, New Jersey, 1974.
- [9] J. K. Ousterhout. *TCL and the Tk Toolkit*. Addison-Wesley, 1994.

²presented verbally at Complex Systems '98, Sydney, Dec. 1998

- [10] R. K. Standish. Ecolab: Where to now? In *Complex Systems: From Local Interactions to Global Phenomena*, pages 263–271. IOS Press, Amsterdam, 1996. also *Complexity International*, **3**, <http://www.csu.edu.au/ci>.
- [11] R. K. Standish. The role of innovation within economics. In W. Barnett, C. Chiarella, S. Keen, R. Marks, and H. Schnabl, editors, *Commerce, Complexity and Evolution*, volume 11 of *International Symposia in Economic Theory and Econometrics*. Cambridge UP, 1999.
- [12] R. K. Standish. Statistics of certain models of evolution. *Phys. Rev. E*, 59:1545–1550, 1999.
- [13] Russell Standish. Econolab. In Standish et al. [14]. also in *Complexity International*, **6** <http://www.csu.edu.au/ci>.
- [14] Russell Standish, Bruce Henry, Simon Watt, Robert marks, Robert Stocker, David Green, Steve Keen, and Terry Bossomaier, editors. *Complex Systems '98 — Complexity Between the Ecos: From Ecology to Economics*. Complexity Online, <http://life.csu.edu.au/complex>, 1998. also in *Complexity International*, **6** <http://www.csu.edu.au/ci>.
- [15] Russell K. Standish. Ecolab documentation. Available at <http://parallel.acsu.unsw.edu.au/rks/ecolab.html>.
- [16] Russell K. Standish. Population models with random embryologies as a paradigm for evolution. In *Complex Systems: Mechanism of Adaptation*. IOS Press, Amsterdam, 1994. also *Complexity International*, **2**, <http://www.csu.edu.au/ci>.
- [17] C. Strobeck. n species competition. *Ecology*, 54:650–654, 1973.