Fundamentals of dynamical systems theory and applications to ecological modelling

Gavin M Abernethy

Department of Computing Science and Mathematics

2023

Dynamical systems

A set of possible states together with a rule that determines the present state of the system in terms of past states.

Continuous-time dynamical systems:

Modelled by differential equations.

Example: **logistic equation**

$$\frac{\mathrm{d}x}{\mathrm{d}t} = rx(1 - x)$$

Useful as a simple model of a (normalised) ecological population with limited resources.

Discrete-time dynamical systems

Modelled by *difference* equations (*maps*).

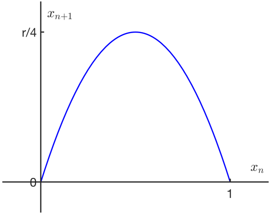

Example: logistic map

$$x_{n+1} = rx_n(1 - x_n)$$

Better represent a population that reproduces in discrete generations.
Typically richer behaviour than their continuous-time counterparts.

Example: Hénon map

$$\begin{aligned} x_{n+1} &= 1 - ax_n^2 + y_n \\ y_{n+1} &= bx_n \end{aligned}$$

Fundamental concepts of discrete dynamical systems

Given space X (e.g. $\mathbb{R}$ or $\mathbb{R}^2$), the map is encoded by a (possibly vector-valued) function $\mathbf{f} : X \to X$.

Given initial point (or *state*) $\mathbf{x}_0 \in X$, then the set of points visited:

$$\{\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n, \ldots\}$$

is the **orbit** of $\mathbf{x}_0$ under $\mathbf{f}$, where $\mathbf{x}_{n+1} = \mathbf{f}(\mathbf{x}_n)$. That is, the orbit is:

$$\{\mathbf{x}_0, \mathbf{f}(\mathbf{x}_0), \mathbf{f}^2(\mathbf{x}_0), \mathbf{f}^3(\mathbf{x}_0), \ldots\}$$

Example: the orbit of $x_0 = 0.3$ under the doubling map $f(x) = 2x$ is:

$$\{0.3,\ 0.6,\ 1.2,\ 2.4,\ 4.8, \ldots\}$$

Typically we are concerned only with the long-term (*limiting*) behaviour of the system, rather than *transient* states.

What is a fixed point of the system?

This is a state $\mathbf{x}^*$ such that $\mathbf{f}(\mathbf{x}^*) = \mathbf{x}^*$ (a point that maps to itself).

For example, for the logistic map $x_{n+1} = rx_n(1 - x_n)$, fixed points are given by setting $x_n = x_{n+1} = x$ and solving:

$$x = rx(1 - x)$$

Then either $x = 0$, or (if $r \neq 0$):

$$\frac{1}{r} = 1 - x \quad \Longrightarrow \quad x = 1 - \frac{1}{r}$$

Ecologically: zero starting population remains at zero!

Stability of fixed points

A key concern in dynamical systems is the *stability* of a fixed point. If the system is perturbed, will it return to the fixed point or not?

Consider a ball placed at the bottom of a smooth bowl, there it will remain - a fixed point of the system. If we push it slightly away, it will roll back again - a *stable* fixed point.

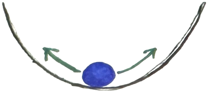

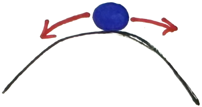

Alternatively, balance the ball atop a smooth hill. A disturbance in any direction will cause it to roll further (the perturbation is magnified) and not return. This is an *unstable* fixed point.

Stability - ecological significance

For our single population (e.g. *arvicola*) governed by the logistic map, stability has an important meaning.

If $r = 2$, the non-zero fixed point $1 - \frac{1}{r}$ (more generally, *state state* or *equilibrium*) is at $1 - \frac{1}{2} = \frac{1}{2}$ of the maximum density. But what happens if you step on one?

Can the population recover or are they doomed? An unstable steady state has little relevance to the real world where the population will constantly be subject to perturbations.

Linear stability analysis in 1D discrete-time systems

Stability criterion for fixed points of $x_{n+1} = f(x_n)$:

The fixed point $x = x^*$ is stable if: $|f'(x^*)| < 1$

Let $x_n = x^* + \epsilon$, then by the Taylor expansion about the fixed point:

$$\begin{aligned} x_{n+1} &= f(x_n) = f(x^* + \epsilon) \\ &= f(x^*) + f'(x^*)\epsilon + \frac{1}{2!} f''(x^*)\epsilon^2 + \dots \end{aligned}$$

Close to x^*, initial error ϵ is small, so the nonlinear terms $\epsilon^2, \epsilon^3, \dots$ will be relatively tiny. Thus the ratio of the new error to old is:

$$\left| \frac{x_{n+1} - x^*}{x_n - x^*} \right| \approx |f'(x^*)|$$

This is *linear stability analysis*.

Linear stability analysis in 1D discrete-time systems

Applying this to the logistic map:

$$f(x) = rx(1-x) = rx - rx^2 \quad \Longrightarrow \quad f'(x) = r - 2rx = r(1-2x)$$

So $x^* = 0$ is stable if $|r(1-0)| < 1$, that is if $0 < r < 1$.

And the non-zero fixed point $x^* = 1 - \dfrac{1}{r}$ is stable if:

$$\left| r\big(1 - 2\big(1 - \frac{1}{r}\big)\big) \right| < 1 \quad \Longrightarrow \quad \text{if } |2 - r| < 1$$

That is, if:

$$1 < r < 3$$

Stable fixed points as attractors

What happens if we start our orbit near to a stable fixed point? From the time-series, the orbit is *attracted* to the currently stable fixed point in the logistic map:

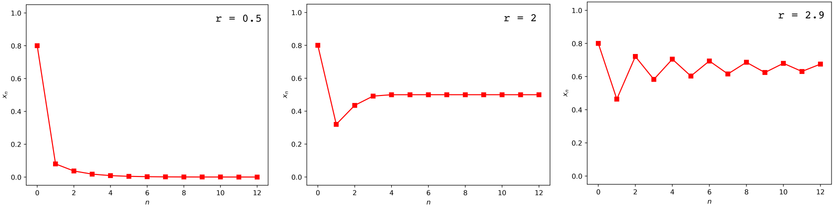

So what happens when the non-zero fixed point *also* becomes unstable at $r = 3$?

Periodic orbits and equilibria

The system instead converges to a *periodic* orbit, with period 2. Now the **attractor** (the set of points constituting the limiting behaviour) is a set of two points instead of one.

Increasing r, the system converges instead to a period-4 orbit:

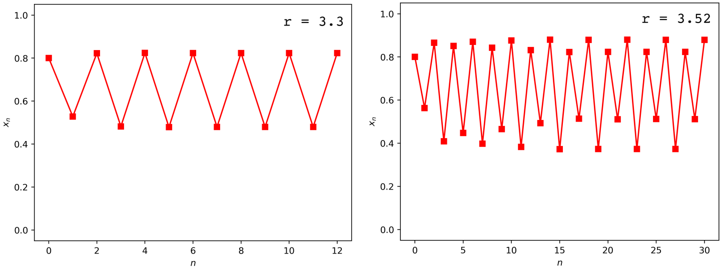

What happens if we increase r even further?

Feigenbaum diagram

Plotting the attractor (the limiting set of points) against control parameter r for the logistic map yields this famous image:

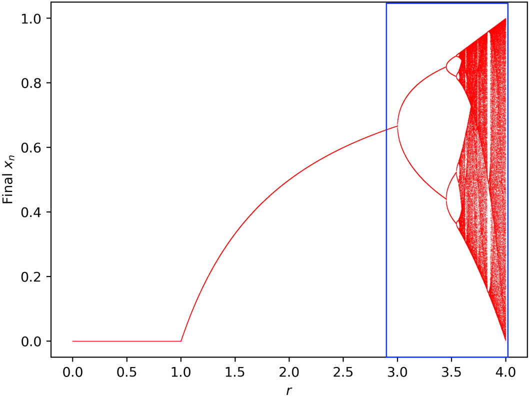

(For $r > 4$, the system is unbounded with negative populations.)

Feigenbaum diagram

This is a bifurcation diagram (*bifurcation* - a fundamental change in the limiting behaviour of the system).

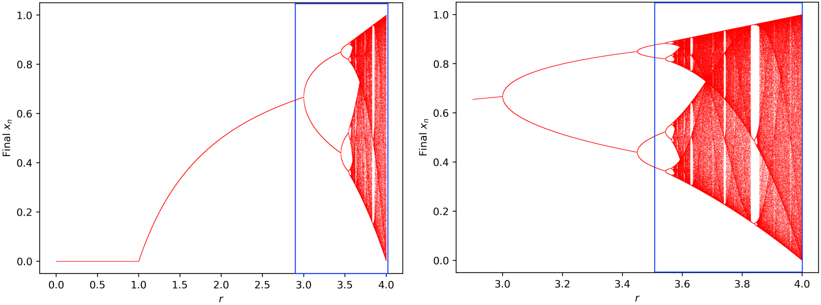

Also known as a Feigenbaum diagram after the US physicist.

Feigenbaum's constant and universality

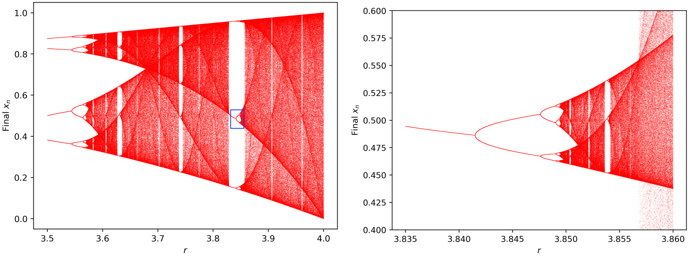

As r increases, period-doubling bifurcations occur more frequently. Let r_n be the value of r where the n^{th} bifurcation occurs. Then...

Feigenbaum's constant:

$$\delta = \lim_{n\to\infty} \frac{r_n - r_{n-1}}{r_{n+1} - r_n} = 4.6692016\ldots$$

This holds for a wider class of unimodel maps (*universality*), perhaps the most significant mathematical discovery using a calculator.

Chaotic dynamics

What is going on in the limit as the period tends to infinity?

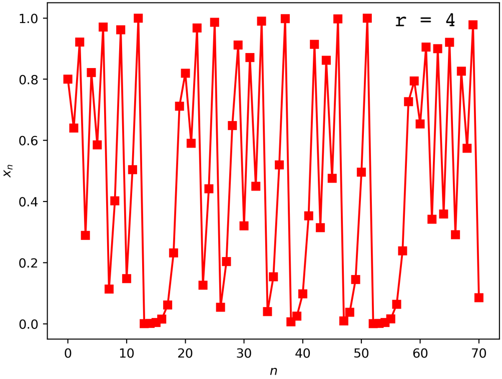

The orbit appears random, but it isn't.

Chaos: sensitivity to initial conditions

What if we start with a small difference in our initial values?

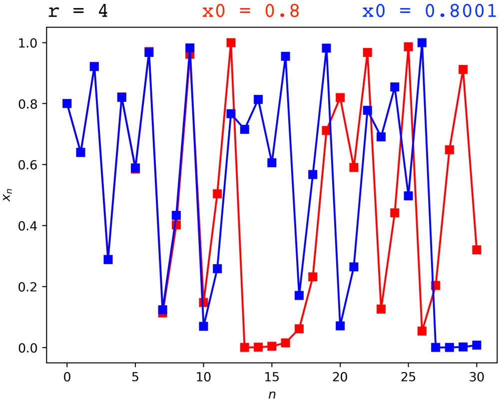

The orbits rapidly appear totally unrelated (the *butterfly effect*).

Chaos: formal definition

Generally speaking, we define a dynamical system (with parameters specified) as "chaotic" if all of the following properties hold:

- Deterministic.
- Sensitivity to initial conditions (we will soon see how to actually define this).
- Bounded (otherwise something like $x_{n+1} = 2x_n$ might fall within the definition).
- Not "asymptotically periodic" (i.e. approaching a periodic orbit).

2D ecological systems: fixed points and stability

Now consider a two-dimensional system, modelling two interacting species. For example, a Lotka-Volterra model of a prey x and predator y:

$$\begin{aligned} x_{n+1} &= F_1(x_n, y_n) = r x_n(1 - x_n) - c x_n y_n \\ y_{n+1} &= F_2(x_n, y_n) = \lambda c x_n y_n + (1 - d) y_n \end{aligned}$$

Again there may exist fixed points (x^*, y^*) where *both* populations remain unchanged. Analysing the stability concerns the Jacobian matrix:

Jacobian matrix for two-dimensional maps

$$J(x, y) = \begin{pmatrix} \frac{\partial F_1}{\partial x} & \frac{\partial F_1}{\partial y} \\ \frac{\partial F_2}{\partial x} & \frac{\partial F_2}{\partial y} \end{pmatrix}$$

Then $J(x^*, y^*)$ has two eigenvalues Λ_1, Λ_2, and the fixed point is stable and attracting if $|\Lambda_k| < 1, \; k = 1, 2$.

2D ecological systems: fixed points

Like the logistic map, there can be fixed points - where the population of prey x_n and predators y_n remains the same forever:

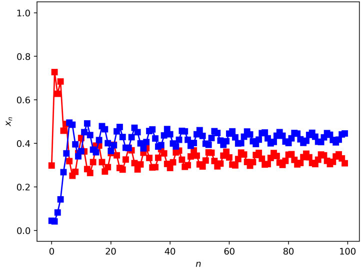

(a) Time-series

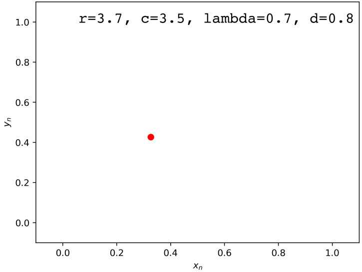

(b) Attractor

Quasiperiodic orbits in 2D

We can also obtain periodic orbits, or increasing the prey growth rate r can give a new kind of *quasiperiodic* behaviour (around the now unstable fixed point):

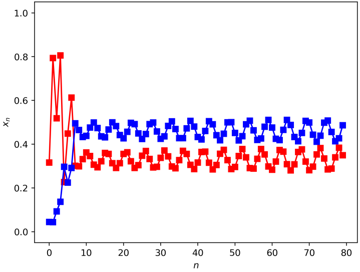

(a) Time-series

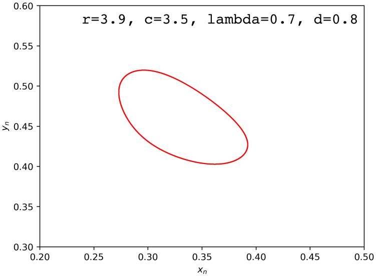

(b) Attractor

Chaotic orbits in 2D

We can again choose parameter values that yield a never-repeating chaotic sequence of points (x_n, y_n):

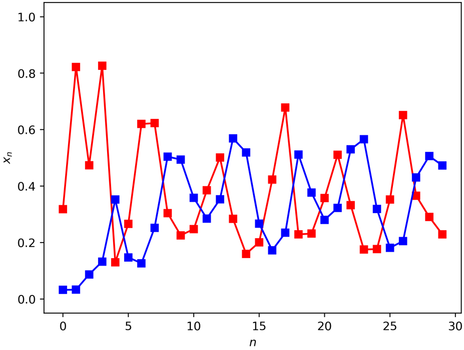

Non-chaotic orbits were converging to some steady state. What does the final set of points for this chaotic orbit look like when plotted?

Strange attractors

This is a **strange attractor**.

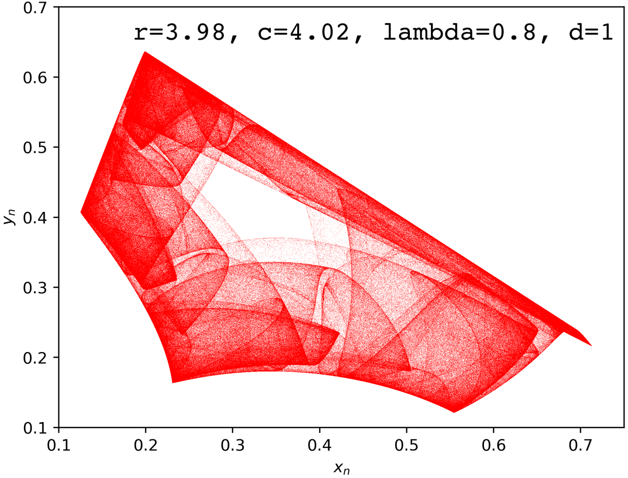

$$x_{n+1} = 3.98x_n(1 - x_n) - 4.02x_n y_n$$

$$y_{n+1} = 3.216x_n y_n$$

5 million plotted, after discarding 15 million transients.

Strange attractors

Typically when the attracting sets of chaotic orbits are plotted, the result in geometric terms is a fractal object, with non-integer dimension, infinite detail and self-similarity.

Recall the Hénon map:

$$\begin{aligned} x_{n+1} &= 1 - ax_n^2 + y_n \\ y_{n+1} &= bx_n \end{aligned}$$

and set $a = 1.4,\ b = 0.3$.

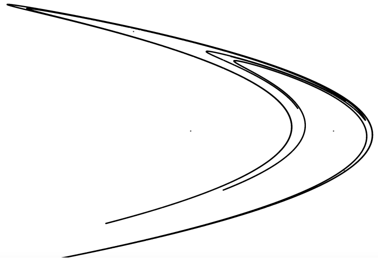

Classifying chaos

To determine if a system is chaotic, boundedness and determinism are easy to confirm. To classify "sensitivity to initial conditions", travelling along the orbit $(x_n)_n$, will a small separation between a nearby orbit be magnified or diminished? That is, for small ϵ, do we find on average:

$$\left|\frac{f(x_n+\epsilon)-f(x_n)}{(x_n+\epsilon)-x_n}\right| = \frac{1}{\epsilon}|f(x_n+\epsilon)-f(x_n)| > 1$$

By contrast, two orbits converging to the same regular attractor will see their separation ratio tend to zero (or one).

Similar to how we derived a condition for stability of equilibria, this concerns the derivative of the map at each point along the orbit.

Lyapunov exponents

An n-dimensional map has a spectrum of n (not necessarily distinct) global Lyapunov exponents. The maximal (or characteristic) such exponent of the orbit $(\mathbf{x}_n)_n$ is the primary indicator of sensitivity.

Maximal global Lyapunov exponent

$$\lambda_1 = \lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} \ln\left(|\mathbf{f}'(\mathbf{x}_k)|\right)$$

If $\lambda_1 > 0$, the average derivative along the orbit is greater than 1, indicating that on average nearby orbits are moved further apart.

Provided our map is deterministic and bounded (as for $f(x) = 2x$, we have $\lambda_1 = \ln(2) \approx 0.69 > 0$), this can practically classify chaos.

Global Lyapunov exponent for the logistic map

For the logistic map, recall $f'(x) = r(1-2x)$. Thus:

$$\lambda(r, x_0) = \lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} \ln\left(|r(1-2x_k)|\right)$$

This is precisely zero at a bifurcation point. Note that, in principle, there may be multiple attractors and the initial state matters.

Observe $\lambda > 0$ for various values $r > 3.56995$ corresponding to chaos.

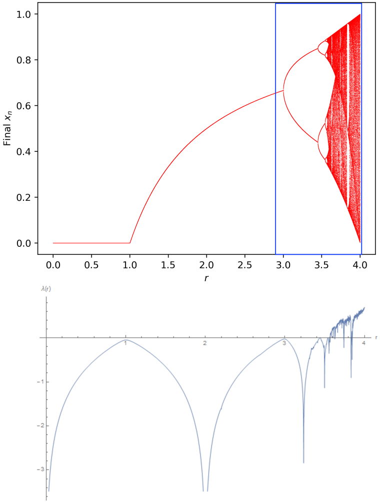

Lyapunov exponents in 2D systems

The maximal Lyapunov exponent can be calculate (practically, using eigenvalues of the Jacobian matrix) for higher-dimensional systems. Together with "is either population extinct", this is a useful tool for high-level classification of the system behaviour.

Consider this model:

$$\begin{aligned} x_{n+1} &= \max((1-p)rx_n(1-x_n) - cx_ny_n,\ 0) \\ y_{n+1} &= \max(prx_n(1-x_n) + \frac{r}{2}(1+x_n)y_n(1-y_n),\ 0) \end{aligned}$$

This predator-prey model features a cannibalistic predator subspecies which evolves directly from the prey species. It can survive independently but reproduces faster when it can feed upon x_n.

Mutant cannibal models

Fixing $p = 0.001$ and examining part of the (c, r)-parameter space:

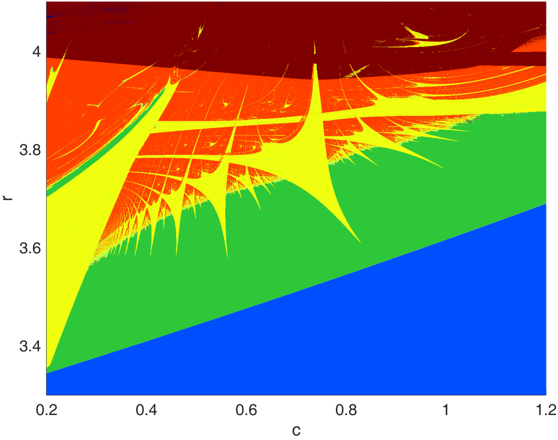

Red = y-only period 1
Orange = 2D chaos ($\lambda_1 > 0$).
Yellow = 2D periodicity.
Green = 2D quasiperiodicity.
Blue = coexistence period-1.

$x_0 = 0.1$, $y_0 = 0.0$.

10^5 transients; 10^6 iterations.

In 2D, two different routes to chaos. . .

Mutant cannibal models: strange attractors

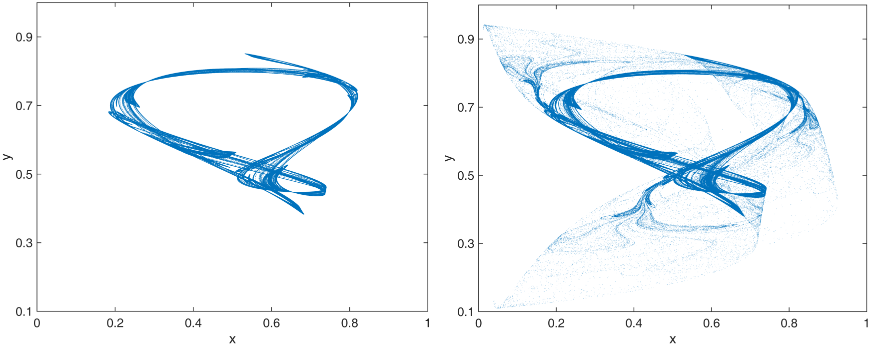

Generalising ecological models

Returning to the continuous-time description of our two-species standard model with general reproductive function f (e.g. logistic model) and predator functional response g (e.g. Lotka-Volterra):

Predator-prey model

Predator: $$\frac{\mathrm{d}y}{\mathrm{d}t} = \lambda g(x,y)y - dy$$

Prey: $$\frac{\mathrm{d}x}{\mathrm{d}t} = f(x)x - g(x,y)y$$

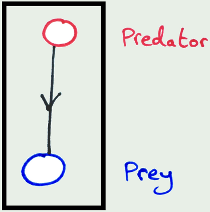

This can be extended in various ways with additional species, or sub-populations separated by space. . .

Generalising ecological models

Community models

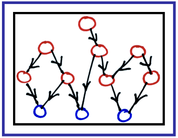

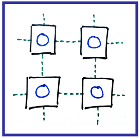

Meta-population models

Eco-evolutionary models

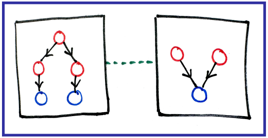

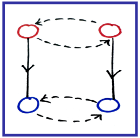

Meta-community models

Eco-Evolutionary Meta-Community Models

Eco-evolutionary meta-community models

Featuring:

- Mutation and evolution
- Multiple species
- Spatial structure

Conceptual challenges:

- Stability - how should we define stability of a whole ecosystem over longer timescales? (e.g. "community robustness")
- How to define species (one vs. multiple traits)?
- Dynamically determine feeding relationships (based on traits)?
- Population dynamics (reproduction, functional response)?
- Network structure of the spatial landscape?
- Dispersal mechanism (diffusion, probabilistic, adaptive)?

Eco-evolutionary community model: Example

Simulating evolutionary assembly of an ecological meta-community in space from first species.

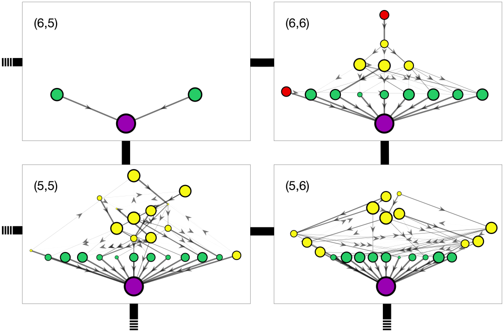

For each local population N:

$$\frac{\mathrm{d}N}{\mathrm{d}t} = F - M - P + \mu_i - \mu_e$$

where:

F = gains due to feeding;
M = natural mortality;
P = loss due to predation;
μ_e, μ_i = em/immigration.

Four local communities of a 6×6 meta-communities.

References

- Britton, N. F. (2003). *Essential mathematical biology*. Springer.
- Gleick, J. (2008). *Chaos: Making a new science.* Penguin.
- Alligood, K.T., Sauer, T.D., & Yorke, J.A. (1996). *Chaos - An Introduction to Dynamical Systems.* Springer.
- Sprott, J. C. (2003). *Chaos and time-series analysis.* Oxford university press.
- Abernethy, G. M., & McCartney, M. (2016). Analysis of a class of low-dimensional models of mutation and predation. *International Journal of Bifurcation and Chaos*, 26(11), 1630029.