Dynamics

Any physical system, such as neurons or muscles, will not respond instantaneously in time but will have a time-varying response termed the *dynamics*. The dynamics of neurons are an inevitable constraint of the times taken to charge cell bodies by ionic currents. Dynamics limit the ultimate speed of computation, but they also can be helpful. One way is to allow compact specifications for the generation of actions. Figure 1 shows an example of a trajectory used to write the letter A generated by the dynamics of a network of neurons. To generate the trajectory for the A, points on the letter are specified only during a training phase. After that, generating the trajectory shown in the figure requires only the starting point.

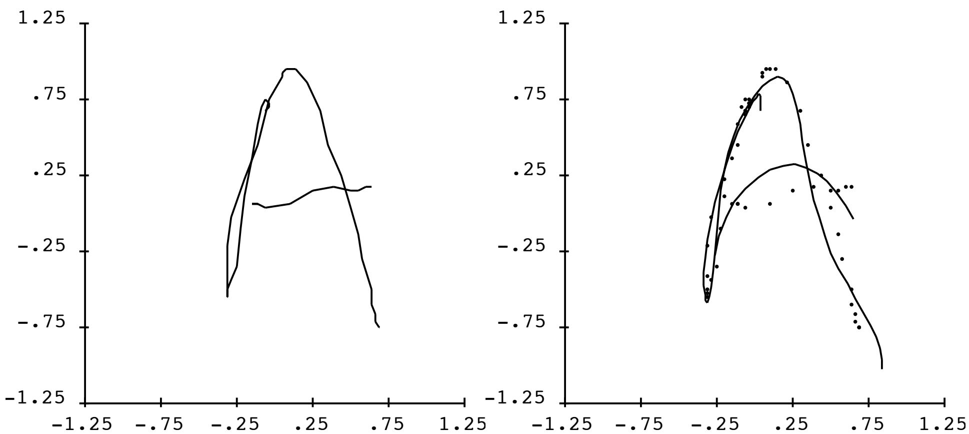

Figure 1: Dynamics used to generate an exemplar of the letter A. To generate this letter, a neural network is trained by being given examples of points along a trajectory, as shown in the upper part of the figure. The trajectory is actually a smooth curve conjoined with a "pen up/pen down" component of the state vector that signals when to write. The result of learning is that when given a static signal denoting the letter A, the natural time constants of the network's weights enable it to dynamically generate an exemplar. (From Simard, 1991.)

Dynamical systems can be summarized as a single vector equation

$$\dot{\boldsymbol{x}} = \boldsymbol{F}(\boldsymbol{x})$$

where $\dot{\boldsymbol{x}}$, read as "x dot," denotes temporal derivative $\frac{d\boldsymbol{x}}{dt}$ and, in vector notation,

$$\boldsymbol{x} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} \text{ and } \dot{\boldsymbol{x}} = \begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \vdots \\ \dot{x}_n \end{pmatrix}$$

and

$$\boldsymbol{F}(\boldsymbol{x}) = \begin{bmatrix} f_1(x_1, x_2, \ldots, x_n) \\ \vdots \\ f_n(x_1, x_2, \ldots, x_n) \end{bmatrix} \tag{1}$$

where n is the *order* of the system. The dynamical system vector equation has a nice graphical interpretation. Where $\boldsymbol{x}$ is the position of the state vector, then $\dot{\boldsymbol{x}}$ is its instantaneous velocity. And since $\dot{\boldsymbol{x}} = \boldsymbol{F}(\boldsymbol{x})$, then, for any point $\boldsymbol{x}$, $F(\boldsymbol{x})$ describes the change in state of the system when started at $\boldsymbol{x}$, as shown in Figure 2.

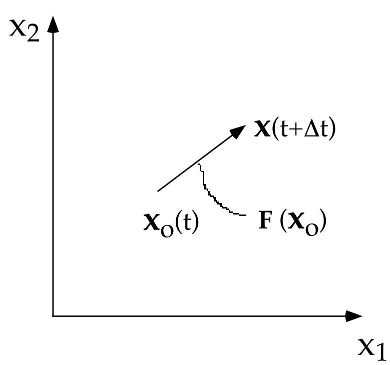

Figure 2: The dynamic function has a straightforward interpretation as the instantaneous direction of motion of the system in the state space.

Linearizing a NL dynamical system

Although the nonlinear system $\dot{\boldsymbol{x}} = \boldsymbol{F}(\boldsymbol{x})$ is difficult to solve in general, its behavior about *equilibrium points* can be very informative. Such points in the state space can be *stable*, wherein the state vector moves back after a perturbation, or *unstable*, wherein the state vector moves away after a perturbation.

Equilibrium point analysis can be formalized in terms of a linear approximation to the system at the point in question. The way to linearize a nonlinear system is to let $\boldsymbol{x} = \boldsymbol{x}_0 + \Delta\boldsymbol{x}$ where $\boldsymbol{x}_0$ is an equilibrium point. Then this equation can be expanded in a Taylor series (see appendix). Writing the terms up to first order explicitly and denoting the rest as higher-order terms (HOTs),

$$\dot{\boldsymbol{x}}_0 + \Delta\dot{\boldsymbol{x}} = \boldsymbol{F}(\boldsymbol{x}_0) + \boldsymbol{F}'(\boldsymbol{x}_0)\Delta\boldsymbol{x} + \text{HOTs}$$

so that, since $\dot{\boldsymbol{x}}_0 = \boldsymbol{F}(\boldsymbol{x}_0)$, and it is assumed that the HOTs can be neglected,

$$\Delta\dot{\boldsymbol{x}} = \boldsymbol{F}'(\boldsymbol{x}_0)\Delta\boldsymbol{x} \tag{2}$$

This system is linear, since the matrix $\boldsymbol{F}'(\boldsymbol{x}_0)$ is constant, being evaluated at the equilibrium point. This matrix $\boldsymbol{F}'(\boldsymbol{x})$is a matrix of *partial derivatives*, and is called the Jacobian, as shown in Eq 3:

$$\boldsymbol{F}'(\boldsymbol{x}) = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \frac{\partial f_2}{\partial x_1} & & & \\ \vdots & & & \\ \frac{\partial f_n}{\partial x_1} & \cdots & \cdots & \frac{\partial f_n}{\partial x_n} \end{bmatrix} \tag{3}$$

Example: Linearizing a Dynamical System For a second-order system let

$$\boldsymbol{F}(\boldsymbol{x}) = \begin{pmatrix} x_1 + x_2^2 \\ x_1 x_2 + x_2^2 \end{pmatrix}$$

Then the Jacobian is given by

$$\boldsymbol{F}'(\boldsymbol{x}) = \begin{bmatrix} 1 & 2x_2 \\ x_2 & x_1 + 2x_2 \end{bmatrix}$$

At equilibrium, $\boldsymbol{F}(\boldsymbol{x}) = \mathbf{0}$, so that the point $\boldsymbol{x}_0 = \begin{pmatrix} -1 \\ 1 \end{pmatrix}$ is an equilibrium point. At this point the Jacobian $\boldsymbol{F}'(\boldsymbol{x}_0)$ simplifies to

$$\boldsymbol{F}'(\boldsymbol{x}_0) = \begin{bmatrix} 1 & 2 \\ 1 & 1 \end{bmatrix}$$

Equation 2 is crucial, since it turns out that the eigenvalues of the Jacobian matrix govern the dynamic behavior of the system. Where the approximation is exact, the eigenvalues describe the behavior for the entire state space. In the more general case of a nonlinear system, the approximation is only local to equilibrium points, but the eigenvalues are still extremely useful in predicting local behavior.

Models for Neuron Dynamics

A system of n neurons can be modeled by an equation of the form

$$\dot{\boldsymbol{x}} = \boldsymbol{g}(W\boldsymbol{x})$$

where W is a matrix denoting synaptic connections and $\boldsymbol{g}$ is a function that prevents the value of $\boldsymbol{x}$ from becoming too large or too small. The linear form of the model is given by $\dot{\boldsymbol{x}} = W\boldsymbol{x}$. Our ultimate objective is to describe ways of changing W to attain the dynamical behavior that achieves some goal, but before we can achieve it, we must understand how a given W governs dynamic behavior.

Another way of achieving dynamical trajectories is to "control" or "force" the system with a specific time-varying signal, perhaps generated by additional neurons. In this case the model of the system is extended to denote the *control* vector $\boldsymbol{u}$:

$$\dot{\boldsymbol{x}} = \boldsymbol{F}(\boldsymbol{x}, \boldsymbol{u})$$

As with the unforced case, insights into this system can be gained by studying its linear form,

$$\dot{\boldsymbol{x}} = A\boldsymbol{x} + B\boldsymbol{u}$$

The distinction between achieving behaviors by changing W or changing $\boldsymbol{u}$ is that W is assumed to vary much more slowly than $\boldsymbol{u}$ whereas the latter is assumed to vary on a timescale with the state vector $\boldsymbol{x}$.

Solving Linear Systems

From the foregoing discussion you can appreciate that linear differential systems are important as an approximation to nonlinear systems. The reason for this importance is that they are easy to analyze. The solution to the general case can be obtained as the composition of the solutions to first- and second-order systems. Thus we start by solving these key examples.

Radioactive Decay The first example models the process of radioactive decay, as radioactive particles lose their radioactivity by emitting subatomic particles. The rate of this process is governed by the current amount of the substance.

Beginning with M_0 of an element, and letting $x(t)$ denote the amount at any time t,

$$\dot{x} = -ax$$

To solve this expression, note that the equation is very close to setting the value of x equal to its derivative. This observation motivates trying e^t, since it is its own derivative. In fact, trying e^{ct},

$$ce^{ct} = -ae^{ct}$$

or

$$c = -a$$

so that e^{-at} is a solution.

Note that $x(t) = Ae^{-at}$ works for any A. This fact is fortunate because we have to be able to match the initial condition $x(0) = M_0$, since $e^0 = 1, A = M_0$, and $x = M_0e^{-at}$.

Undamped Harmonic Oscillator Figure 3 shows the classical second-order system of a spring and mass. Motion is constrained in the vertical direction. In the horizontal direction the only force is the spring force, which is a linear function of displacement, that is, $F = -kx$.

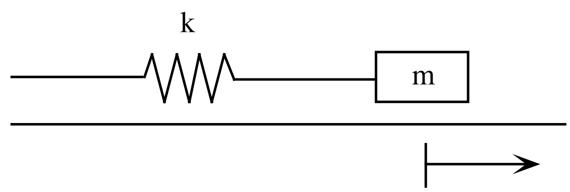

Figure 3: The classical second-order system: an undamped harmonic oscillator consisting of a mass connected to a spring. Of interest is the behavior of the system when the mass is displaced from the equilibrium point.

To derive the dynamic equation, start with Newton's law

$$F = ma$$

and substitute for the spring force. Acceleration is simply the second temporal derivative (in the notation "x double dot"):

$$-kx = m\ddot{x}$$

Rearranging,

$$\ddot{x} + \frac{k}{m}x = 0$$

For the solution, as in the first-order case, try $x(t) = e^{at}$ so that

$$\ddot{x} = a^2 e^{at}$$

Substituting in the dynamic equation leads to

$$a^2 e^{at} + \frac{k}{m}e^{at} = 0$$

or, in other words,

$$a^2 = -\frac{k}{m}$$

Thus in this case the solution has imaginary roots; that is, where imaginary $i = \sqrt{-1}$,

$$a = \pm\sqrt{\frac{k}{m}}i$$

So the solution has the form

$$x(t) = Ae^{iwt} + Be^{-iwt}$$

where $w = \sqrt{\frac{k}{m}}$ and where A and B are constants that can be determined by the initial conditions. Assume the spring starts out with zero velocity and is stretched by d, then from the initial condition on position,

$$x(0) = d = Ae^0 + Be^0 = A + B$$

and next from the initial condition on velocity,

$$\dot{x}(0) = 0 = iwAe^0 - iwBe^0$$

which implies that

$$A - B = 0$$

Combining these two equations in A and B,

$$A = B = \frac{d}{2}$$

Thus, finally, the solution is given by

$$x(t) = \frac{d}{2}e^{iwt} + \frac{d}{2}e^{-iwt}$$

This solution can also be written as a cosine. Recall Euler's formula,

$$e^{i\theta} = \cos\theta + i\sin\theta$$

which allows the cosine and sine to be expressed as

$$\cos\theta = \frac{e^{i\theta} + e^{-i\theta}}{2} \text{ and } \sin\theta = \frac{e^{i\theta} - e^{-i\theta}}{2i}$$

so that $x(t) = d\cos wt$. What this result means is that you can forget about sine and cosine and consider only exponentials. An imaginary part of the exponential indicates an oscillation in the solution.

The General Case

Although the previous example is a second-order linear differential equation (LDE) with constant coefficients, it can be expressed equivalently as two first-order LDEs. To see this possibility, define

$$x_1 = x \text{ and } x_2 = \dot{x}$$

Then

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = \begin{bmatrix} 0 & 1 \\ -\frac{k}{m} & 0 \end{bmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

The general form of this transformation, where $x_{n+1} = \frac{d^n x}{dt^n}$, works for LDEs with constant coefficients of *any* order. In other words, any n^{th}-order LDE can be expressed as the sum of n first-order LDEs. The effect of this transformation is that the solutions that you studied for the first- and second-order case extend to the higher-order cases as well. For an n-dimensional vector, each component will either grow or decay exponentially, and it may oscillate as well if its exponent turns out to have an imaginary part. With this fact in mind, let us solve the general case, which is specified by

$$\dot{\boldsymbol{x}} = A\boldsymbol{x}$$

where A is a matrix of constant coefficients.

How do we solve this system? Since $x = x_0 e^{at}$ worked before, try

$$\boldsymbol{x}(t) = \boldsymbol{x}_0 e^{At} \tag{4}$$

where $\boldsymbol{x}_0$ is a vector of initial conditions. What about e^{At}? Since

$$e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \cdots$$

define e^A as

$$e^A \equiv I + A + \frac{A^2}{2} + \cdots$$

and substitute into both sides of $\dot{\boldsymbol{x}} = A\boldsymbol{x}$:

$$\begin{aligned} A\boldsymbol{x} &= A[I + At + \frac{(At)^2}{2!} + \frac{(At)^3}{3!} + \cdots]\boldsymbol{x}_0 \\ \dot{\boldsymbol{x}} &= [A + \frac{2A(At)}{2!} + \frac{3A(At)^2}{3!} + \cdots]\boldsymbol{x}_0 \end{aligned}$$

Since these two are equal, as you can quickly verify, the solution defined by Eq 4 actually works. It's not practical to try and find it this way, as the series expansion of e^A is too cumbersome to work with, but suppose that the matrix A happened to be diagonal; that is,

$$A = \Lambda = \begin{bmatrix} \lambda_1 & & 0 & \\ & \lambda_2 & & \\ & & \ddots & \\ & 0 & & \lambda_n \end{bmatrix}$$

In this case, all the equations are decoupled and

$$e^A = e^\wedge = \begin{bmatrix} e^{\lambda_1} & & & \\ & e^{\lambda_2} & & \\ & & \ddots & \\ & & & e^{\lambda_n} \end{bmatrix}$$

so that solving this system reduces to solving n versions of the one-dimensional case. That is,

$$\boldsymbol{x}(t) = e^{At}\boldsymbol{x}(0) = e^{\wedge t}\boldsymbol{x}(0) = \begin{bmatrix} x_1(0)e^{\lambda_1 t} \\ x_2(0)e^{\lambda_2 t} \\ \vdots \\ x_n(0)e^{\lambda_n t} \end{bmatrix}$$

Now remember that the matrix A can be diagonalized by a similarity transformation:

$$A^* = T^{-1}AT = \begin{bmatrix} \lambda_1 & & 0 & \\ & \lambda_2 & & \\ & & \ddots & \\ & 0 & & \lambda_n \end{bmatrix}$$

where T is the matrix of eigenvectors. This result motivates transforming $\boldsymbol{x}$ into a new coordinate system by

$$\boldsymbol{x}^* = T^{-1}\boldsymbol{x}$$

using

$$\dot{\boldsymbol{x}} = A\boldsymbol{x}$$

Write

$$T^{-1}\dot{\boldsymbol{x}} = T^{-1}ATT^{-1}\boldsymbol{x}$$

that is,

$$\dot{\boldsymbol{x}}^* = A^*\boldsymbol{x}^*$$

Since A^* is diagonal, $\boldsymbol{x}^*$ has the simple solution format already described, that is,

$$\boldsymbol{x}^*(t) = \begin{bmatrix} x_1^*(0)e^{\lambda_1 t} \\ x_2^*(0)e^{\lambda_2 t} \\ \vdots \\ x_n^*(0)e^{\lambda_n t} \end{bmatrix}$$

To discover $\boldsymbol{x}(t)$, use the fact that $\boldsymbol{x} = T\boldsymbol{x}^*$.

This is the main conclusion: that the dynamics of an n^{th}-order LDE with constant coeffiecients is completely characterized by its eigenvalues and eigenvectors.

Example Let the matrix A be given by

$$A = \begin{bmatrix} 1 & 2 \\ 1 & 3 \end{bmatrix}$$

Then

$$|A - \lambda I| = \lambda^2 - 4\lambda + 1 = 0$$

has the solution $\lambda = 2 \pm \sqrt{3}$. Thus T, which has the eigenvectors as columns, is given by

$$T = \begin{bmatrix} 1 & 1 \\ \frac{3+\sqrt{3}}{2} & \frac{3-\sqrt{3}}{2} \end{bmatrix}$$

and the solution is given by

$$\begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix} = \begin{bmatrix} 1 & 1 \\ \frac{3+\sqrt{3}}{2} & \frac{3-\sqrt{3}}{2} \end{bmatrix} \begin{pmatrix} x_1^*(0)e^{(2+\sqrt{3})t} \\ x_2^*(0)e^{(2-\sqrt{3})t} \end{pmatrix}$$

Of course the initial conditions will be expressed in the original coordinate system and must be transformed into the * coordinate system by T^{-1}.

Intuitive Meaning of Eigenvalues and Eigenvectors

Recall the graphical interpretation of $\dot{\boldsymbol{x}} = \boldsymbol{F}(\boldsymbol{x})$, shown in Figure 4. For a general nonlinear system there will be many equilibrium points, as implied by the diagram, but for a linear system $\dot{\boldsymbol{x}} = A\boldsymbol{x}$ there is a single equilibrium point at the origin.

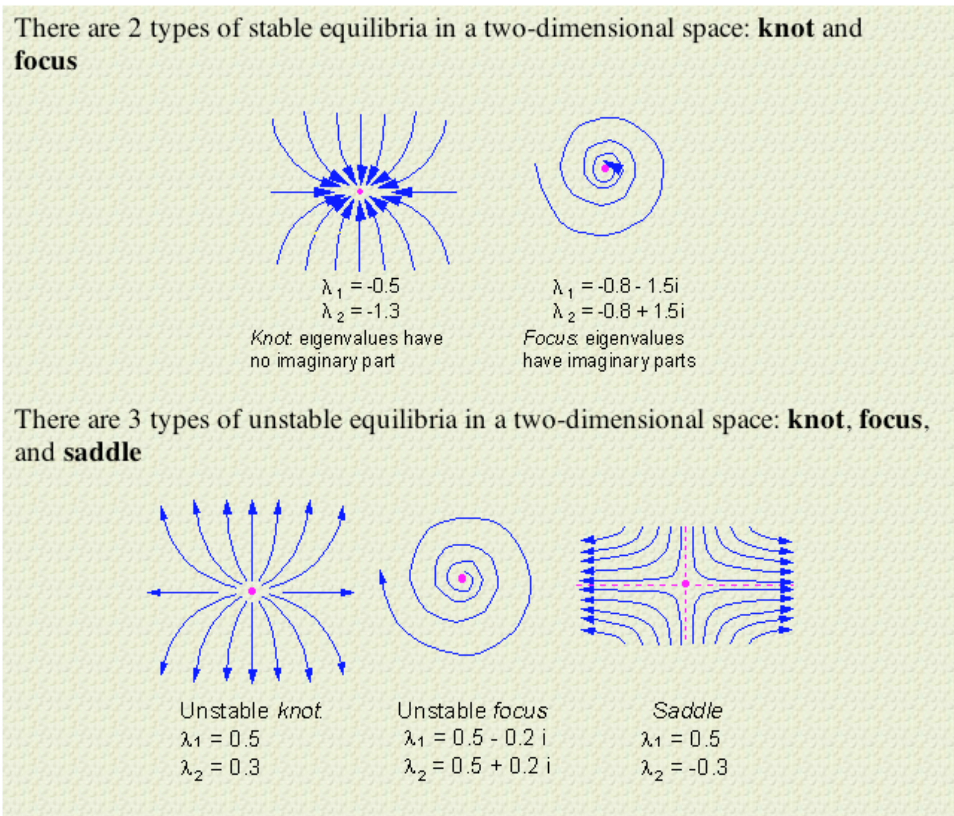

Figure 4: From: (http://home.comcast.net/ sharov/PopEcol/lec9/equilib.html) A nonlinear system may be understood by characterizing the behavior of trajectories of the system linearized near equilibrium points, as in this second-order system. This behavior is entirely determined by the eigenvalues and eigenvectors of the linearized system.

Nonlinear Systems

Linear systems can be analyzed because the behavior over the entire state space can be determined. But nonlinear systems are more complicated, and their behavior can only be analyzed in the neighborhood of equilibrium points. Remember that these are the points for which

$$F(\boldsymbol{x}) = 0$$

Linear systems are simple to characterize because they have a single equilibrium point. The behavior of a system is determined entirely by the values of the eigenvalues and directions of the eigenvectors of the system. One of the main reasons that nonlinear systems, by far the general case, are not so easily handled is that they may contain multiple equilibrium points. How should they be analyzed?

The important realization is that characterizing the global behavior is too difficult and the best one can hope for is to characterize the behavior local to equilibrium points. The key issue is stability. One wants to know what will happen to the state vector when it is displaced from an equilibrium point. Broadly speaking, there are three kinds of behavior.

- Asymptotic stability: A displaced state vector will eventually return to the equilibrium point.
- Instability: A displaced state vector will continue to move away from the equilibrium point.
- Marginal stability: A displaced state vector will oscillate near the equilibrium point.

Example: Linearizing a Nonlinear System

Consider the case of two coexisting populations, one of deer and the other of wolves. The population of deer will grow unabated in the absence of a predator, but will decline with the growth of the wolf population. Similarly, the population of wolves has a tendency to grow and needs the deer for food, so it will increase as a function of the deer population. These effects can be captured by the nonlinear differential equations

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = \begin{pmatrix} ax_1 - bx_1x_2 \\ -cx_2 + dx_1x_2 \end{pmatrix} \tag{5}$$

where a, b, c, and d are all assumed to be positive coefficients. For an equilibrium point, $\dot{x_1} = \dot{x_2} = 0$, so it is easy to see that this reasoning leads to an equilibrium point:

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} \frac{c}{d} \\ \frac{a}{b} \end{pmatrix}$$

Now use Equation 2 to linearize Equation 5 about this equilibrium point. First compute the Jacobian,

$$F'(x) = \begin{bmatrix} a - bx_2 & -bx_1 \\ dx_2 & -c + dx_1 \end{bmatrix}$$

so the system at the equilibrium point is given by

$$\begin{pmatrix} \Delta\dot{x}_1 \\ \Delta\dot{x}_2 \end{pmatrix} = \begin{bmatrix} 0 & \frac{-bc}{d} \\ \frac{ad}{b} & 0 \end{bmatrix} \begin{pmatrix} \Delta x_1 \\ \Delta x_2 \end{pmatrix}$$

From this the characteristic equation is

$$\lambda^2 = -ac$$

so that the two roots are

$$\lambda_{1,2} = \pm\sqrt{ac}\; i$$

From the definitions of stability given earlier it is clear that this equilibrium point is marginally stable. When displaced from the equilibrium point, it will oscillate forever about that point without ever converging or becoming unstable. Thus a small increase in the deer population will result in a subsequent increase in the wolf population, which in turn will result in a subsequent decrease in the deer population, and so on.