

Linearization of Differential Equation Models

1 Motivation

We cannot solve most nonlinear models, so we often instead try to get an overall feel for the way the model behaves: we sometimes talk about looking at the **qualitative dynamics** of a system. Equilibrium points— steady states of the system— are an important feature that we look for. Many systems settle into a equilibrium state after some time, so they might tell us about the long-term behavior of the system.

Equilibrium points can be stable or unstable: put loosely, if you start near an equilibrium you might, over time, move closer (stable equilibrium) or away (unstable equilibrium) from the equilibrium. Physicists often draw pictures that look like hills and valleys: if you were to put a ball on a hill top and give it a push, it would roll down either side of the hill. If you were to put a ball at the bottom of a valley and push it, it would fall back to the bottom of the valley.




Figure 1: An example of stability: both A and B are equilibrium points, at the top of a hill and at the bottom of a valley. If the ball at point A is pushed in either direction, it will roll away down the hill. If the ball at point B is pushed a small amount in either direction, it will roll back to its initial point.

Mathematicians have many different definitions of ‘stability’, but we won’t worry too much about such details, except to say that we often distinguish between **local** and **global** stability. In the example above, point B is **locally stable** but not **globally stable**. If you only push the ball a short distance away, it will roll back to point B. If you push the ball far enough (i.e. beyond point A), it will not roll back to point B. More detailed information on stability can be found in books

on nonlinear differential equations or dynamical systems (for instance S. H. Strogatz's 'Nonlinear Dynamics and Chaos').

Linearization can be used to give important information about how the system behaves in the neighborhood of equilibrium points. Typically we learn whether the point is stable or unstable, as well as something about how the system approaches (or moves away from) the equilibrium point.

The basic idea is that (in most circumstances) one can approximate the nonlinear differential equations that govern the behavior of the system by linear differential equations. We can solve the resulting set of linear ODEs, whereas we cannot, in general, solve a set of nonlinear differential equations.

2 How to Linearize a Model

We shall illustrate the linearization process using the SIR model with births and deaths in a population of fixed size.

$$\dot{S} = \mu N - \beta SI/N - \mu S \quad (1)$$

$$\dot{I} = \beta SI/N - (\gamma + \mu)I \quad (2)$$

$$\dot{R} = \gamma I - \mu R, \quad (3)$$

with $S + I + R = N$. Since we assume that the population is closed, we can always calculate the value of R if we know S and I . Therefore we need only focus on the first two equations for S and I .

We denote an equilibrium point by attaching asterisks to the state variables: e.g. an equilibrium point of the SIR model may be written as (S^*, I^*) . Because an equilibrium point means that the values of S and I (and R) remain constant, this means that $dS/dt = dI/dt = 0$ when $(S, I) = (S^*, I^*)$.

If we imagine that both S and I are close to the equilibrium point, then the differences $S - S^*$, which we denote by x_1 , and $I - I^*$, which we denote by x_2 , will be small. An important point is that terms involving products of x_1 and x_2 (e.g. the quadratic terms x_1^2, x_2^2 or $x_1 x_2$) are much smaller still and so, to a very good approximation, can be ignored.

We can differentiate our expression for x_1 : $dx_1/dt = dS/dt - dS^*/dt$. Since S^* is a constant, we have $dx_1/dt = dS/dt$ and, using a similar argument, $dx_2/dt = dI/dt$. So we now have

$$\dot{x}_1 = \mu N - \beta SI/N - \mu S \quad (4)$$

$$\dot{x}_2 = \beta SI/N - (\gamma + \mu)I. \quad (5)$$

But these equations are in terms of the original variables, S and I .

There are two ways in which we can then obtain the linearization. One is a calculus-free method, the other uses the idea of Taylor series from calculus.

2.1 Non-calculus method:

We rewrite the previous equations in terms of x_1 and x_2 using $S = S^* + x_1$ and $I = I^* + x_2$.

$$\dot{x}_1 = \mu N - \beta(S^* + x_1)(I^* + x_2)/N - \mu(S^* + x_1) \quad (6)$$

$$\dot{x}_2 = \beta(S^* + x_1)(I^* + x_2)/N - (\gamma + \mu)(I^* + x_2). \quad (7)$$

If we multiply out the brackets that appear in the infection term, we have

$$\beta SI/N = (\beta/N)(S^* + x_1)(I^* + x_2) \quad (8)$$

$$= (\beta/N)(S^*I^* + S^*x_2 + I^*x_1 + x_1x_2) \quad (9)$$

$$\approx (\beta/N)(S^*I^* + S^*x_2 + I^*x_1). \quad (10)$$

Where, to reach the last step we made use of the fact that we noted above: the product x_1x_2 is very small indeed and so can be ignored.

This leads to the following:

$$\dot{x}_1 = \mu N - \beta S^* I^*/N - \mu S^* + \beta(S^*x_2 + I^*x_1)/N - \mu x_1 \quad (11)$$

$$\dot{x}_2 = \beta S^* I^*/N - (\gamma + \mu)I^* + \beta(S^*x_2 + I^*x_1)/N - (\gamma + \mu)x_2. \quad (12)$$

Since (S^*, I^*) is an equilibrium point, the original model equations tell us that $\mu N - \beta S^* I^*/N - \mu S^* = 0$ and $\beta S^* I^*/N - (\gamma + \mu)I^* = 0$. This allows us to cancel some terms in these equations and we are left with the following linearized equations:

$$\dot{x}_1 = (\beta/N)(S^*x_2 + I^*x_1) - \mu x_1 \quad (13)$$

$$\dot{x}_2 = (\beta/N)(S^*x_2 + I^*x_1) - (\gamma + \mu)x_2. \quad (14)$$

2.2 Calculus method:

By using a Taylor series expansion, we can arrive a little more quickly at the linearization.

As a shorthand, we write the right hand side of the dS/dt equation as $f(S, I)$ (e.g. $f(S, I) = \mu N - \beta SI/N - \mu S$) and the right hand side of the dI/dt equation as $g(S, I)$. We then expand about the point (S^*, I^*) to give

$$\frac{dS}{dt} = f(S^*, I^*) + (S - S^*)\frac{\partial f}{\partial S} + (I - I^*)\frac{\partial f}{\partial I} + (\text{higher order terms}) \quad (15)$$

$$\frac{dI}{dt} = g(S^*, I^*) + (S - S^*)\frac{\partial g}{\partial S} + (I - I^*)\frac{\partial g}{\partial I} + (\text{higher order terms}). \quad (16)$$

Here, both partial derivatives are evaluated at the point (S^*, I^*) . (In case you aren't familiar with partial derivatives: when you work out $\partial f/\partial S$, you imagine that I is a constant. For the SIR model, $\partial f/\partial S = \beta I/N - \mu$ and $\partial f/\partial I = \beta S/N$.)

Since (S^*, I^*) is an equilibrium point, we have that $f(S^*, I^*)$ equals zero, since $dS/dt = 0$ at this point and $g(S^*, I^*) = 0$ since $dI/dt = 0$. Remembering that $x_1 = S - S^*$ and $x_2 = I - I^*$ and that

$dx_1/dt = dS/dt$ and $dx_2/dt = dI/dt$ we have

$$\frac{dx_1}{dt} = x_1 \frac{\partial f}{\partial S} + x_2 \frac{\partial f}{\partial I} \quad (17)$$

$$\frac{dx_2}{dt} = x_1 \frac{\partial g}{\partial S} + x_2 \frac{\partial g}{\partial I}, \quad (18)$$

where we have again ignored the higher order terms since they are of much smaller size.

For the SIR model, this becomes

$$\dot{x}_1 = (\beta I^*/N - \mu)x_1 - (\beta S^*/N)x_2 \quad (19)$$

$$\dot{x}_2 = (\beta I^*/N)x_1 + (\beta S^*/N - \gamma - \mu)x_2. \quad (20)$$

These are the same equations that we had before. Once you are familiar with the process, it's very easy to obtain the linearized equations in this way.

2.3 Matrix Notation for the Linearization

We can write linearizations in matrix form:

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = \begin{pmatrix} \frac{\partial f}{\partial S} & \frac{\partial f}{\partial I} \\ \frac{\partial g}{\partial S} & \frac{\partial g}{\partial I} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}, \quad (21)$$

or in shorthand

$$\dot{\mathbf{x}} = J\mathbf{x}, \quad (22)$$

where J is the so-called Jacobian matrix, whose entries are the partial derivatives of the right hand sides of the differential equations describing the model, taken with respect to the different state variables of the model (e.g. S and I). We often write the entries of J as

$$J = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}. \quad (23)$$

We can do this linearization process for a model with any number of state variables: if there are n state variables, we get an n -dimensional set of coupled linear differential equations. (Bear in mind that for the SIR model there are three state variables, S , I and R , but our assumption that $S + I + R = N$ leads to our only having to consider a two dimensional system.)

3 What do we do with the linearization?

There is a well-developed theory for solving linear differential equations such as (22). We can only cover the briefest points here: for more information, find a book on differential equations or an introductory mathematical biology text. (You might start with chapter 5 of S. H. Strogatz's 'Nonlinear Dynamics and Chaos').

3.1 One Dimensional Case

It's perhaps simplest to start with the corresponding one-dimensional equation:

$$\dot{x} = \lambda x. \quad (24)$$

This equation has solution

$$x(t) = ce^{\lambda t}, \quad (25)$$

where c is the initial value of x (i.e. the value taken by x when $t = 0$). This equation describes exponential growth or decay.

If λ is greater than zero, then points move away from $x = 0$. Remembering that $x = 0$ corresponds to the equilibrium point, we see that non-zero points move away from the equilibrium as time passes: the equilibrium is unstable. If λ is less than zero, points move towards $x = 0$: the equilibrium is stable. If $\lambda = 0$, points neither move towards nor away from the equilibrium.

The sign of λ tells us about the stability of the equilibrium, and the size of λ tells us something about how quickly points move away from or towards the equilibrium. When we have a stable equilibrium, we sometimes talk about the **relaxation time**, which is defined to be $-1/\lambda$. This is the time taken for the distance between the point and the origin to decrease by a factor $1/e \approx 0.368$.

We can summarize the behaviors in the following figures, in which arrows denote the directions in which points move.



Figure 2: The left panel illustrates an unstable equilibrium, the right panel a stable equilibrium.

3.2 Two Dimensional Case

By analogy with the one dimensional case, we try a solution of the form

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} e^{\lambda t}, \quad (26)$$

where λ , v_1 and v_2 are constants. In shorthand notation we have

$$\mathbf{x}(t) = \mathbf{v}e^{\lambda t}. \quad (27)$$

If we differentiate $\mathbf{x}(t)$ directly, we see that its derivative is given by

$$\dot{\mathbf{x}}(t) = \lambda \mathbf{v}e^{\lambda t}. \quad (28)$$

But we know that $\mathbf{x}(t)$ has to satisfy $d\mathbf{x}/dt = J\mathbf{x}$, so we have

$$\dot{\mathbf{x}} = J(\mathbf{v}e^{\lambda t}). \quad (29)$$

$$= (J\mathbf{v})e^{\lambda t}. \quad (30)$$

Comparing (28) and (29) we see that \mathbf{v} and λ must satisfy

$$J\mathbf{v} = \lambda\mathbf{v}. \quad (31)$$

In linear algebra, such vectors \mathbf{v} are known as **eigenvectors** of the matrix J and the constants λ are known as **eigenvalues** of the matrix. Because these are important properties of matrices, there is quite a large theory devoted to calculating their values. (If you need to know more, a good place to start looking is a linear algebra textbook.) In the two dimensional case, the matrix J typically has two independent eigenvectors and two eigenvalues (although the two eigenvalues can be equal).

Because we have a two dimensional set of linear differential equations, the general solution is, in most cases ¹, given by the sum of two terms of the form (27): $\mathbf{x}(t) = c_1\mathbf{v}_1e^{\lambda_1 t} + c_2\mathbf{v}_2e^{\lambda_2 t}$. (The fact that one can add solutions together in this way is a fundamental property of linear systems, and one which makes them nice to work with.) The constants c_1 and c_2 are determined by the initial conditions. Notice that this form of the solution shows that the eigenvectors define two directions along which there may be exponential growth or decay. We will see this in action later on.

3.2.1 Finding the Eigenvalues and Eigenvectors in the 2D Case

We now need to find the eigenvalues and eigenvectors. Rearranging (31) gives us

$$(J - \lambda I)\mathbf{v} = 0, \quad (32)$$

where I is the two dimensional **identity matrix**

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}. \quad (33)$$

Theory from linear algebra shows that such an equation can only have a non-trivial solution (i.e. one in which \mathbf{v} is not equal to zero) if the **determinant** of the matrix $J - \lambda I$ is equal to zero.

For a 2×2 matrix with entries a_{ij} (e.g. as in equation (23)), the determinant of the matrix is given by the product of the two entries on the leading diagonal minus the product of the two entries on the other diagonal: $a_{11}a_{22} - a_{12}a_{21}$.

Using this definition, we see that in order for λ to be an eigenvalue of the matrix J , it must satisfy

$$(a_{11} - \lambda)(a_{22} - \lambda) - a_{12}a_{21} = 0. \quad (34)$$

(Notice that the ‘ $-\lambda$ ’ terms arise because we have $J - \lambda I$, and the matrix I has entries equal to one in the (1,1) and (2,2) positions.)

Equation (34) is a quadratic equation for λ : notice that this gives us (in most cases) two possible answers, λ_1 and λ_2 . If we had an n -dimensional system, we would get a polynomial of degree n . When n is anything other than a small number, such equations become difficult (or impossible) to solve. But in the 2D case we can get an explicit expression for the eigenvalues. At worst, we can always use the ‘quadratic formula’, but we can do a little better than this.

¹The general solution can be more complex if the matrix J does not have a complete set of linearly independent eigenvectors. In such cases, the general solution has terms of the form $p(t)\mathbf{v}e^{\lambda t}$, where $p(t)$ is a polynomial in t .

Multiplying out (34) gives

$$\lambda^2 - (a_{11} + a_{22})\lambda + (a_{11}a_{22} - a_{12}a_{21}) = 0. \quad (35)$$

We see that the linear term involves $a_{11} + a_{22}$, the sum of the entries on the leading diagonal of J . This sum is called the **trace** of the matrix J , $\text{tr } J$. The constant term of this quadratic is just the determinant of the matrix J , $\det J$. So we can rewrite (35) succinctly as

$$\lambda^2 - \text{tr } J\lambda + \det J = 0. \quad (36)$$

Since we know that the numbers λ_1 and λ_2 satisfy the quadratic equation for λ , we know that we can factor this quadratic in the following way

$$(\lambda - \lambda_1)(\lambda - \lambda_2) = 0. \quad (37)$$

Multiplying out this equation, we get

$$\lambda^2 - (\lambda_1 + \lambda_2)\lambda + \lambda_1\lambda_2 = 0. \quad (38)$$

Comparing (36) and (38) shows that

$$\lambda_1 + \lambda_2 = \text{tr } J \quad (39)$$

$$\lambda_1\lambda_2 = \det J. \quad (40)$$

We can immediately write down the sum of the eigenvalues and their product in terms of the trace and determinant of J : quantities that we can easily work out from the entries of J .

The formula for the solution of a quadratic equation, applied to (36) gives

$$\lambda = -\frac{1}{2} \text{tr } J \pm \frac{1}{2} \sqrt{(\text{tr } J)^2 - 4 \det J}. \quad (41)$$

Notice that an important difference between the 1D and 2D cases is that we can now have **complex** eigenvalues. This happens if the term inside the square root is negative, i.e. if $(\text{tr } J)^2 < 4 \det J$.

If we have a complex eigenvalue, which we can write in terms of its real and imaginary parts as $\lambda_1 = \rho + i\omega$, then it is fairly straightforward to show that $\lambda_2 = \rho - i\omega$ is also an eigenvalue. These eigenvalues have equal real parts but their imaginary parts have opposite signs: this is called a **complex conjugate pair**. There is a standard result that says if you have a polynomial whose coefficients are real, then complex eigenvalues always come in such complex conjugate pairs. (Notice that this result applies to our situation: the entries of the Jacobian matrix are real numbers.)

Notice that in the complex eigenvalue case, $\exp(\lambda t)$ is of the form $\exp\{(\rho \pm i\omega)t\}$, which can be multiplied out to give $\exp(\rho t)\exp(\pm i\omega t)$. The magnitude of $\exp(\lambda t)$ is given by the term $\exp(\rho t)$: the real part of the eigenvalue, ρ , describes whether one has growth or decay. In the decay case, the relaxation time is equal to $-1/\rho$. Since one can write

$$\exp(i\omega t) = \cos(\omega t) + i \sin(\omega t), \quad (42)$$

we see that the term $\exp(i\omega t)$ describes oscillatory behavior with constant amplitude, with (angular) frequency ω . The period of the oscillation is given by $2\pi/\omega$.

3.3 Complete Characterization of Behavior in 2D Case

We now make use of the above analysis of the eigenvalues in 2D to characterize all possible cases. We make use of the fact that $\det J = \lambda_1 \lambda_2$ and $\text{tr } J = \lambda_1 + \lambda_2$ to find the signs of the eigenvalues (or of their real parts in the complex case), as it is these quantities that determine stability. We will leave discussion of so-called ‘borderline’ cases until the end: these borderline cases involve ones in which either both eigenvalues are equal or one (or more) eigenvalues are equal (or have real part equal) to zero.

- **Both eigenvalues are real:** $(\text{tr } J)^2 > 4 \det J$

- $\text{tr } J > 0$

This means that the sum of the eigenvalues is positive.

There are two possible cases:

- $\det J > 0$

The product of the eigenvalues is also positive. This means that both λ_1 and λ_2 must be positive. In terms of the differential equation, we have exponential growth in the directions represented by both \mathbf{v}_1 and \mathbf{v}_2 . We call this situation an **unstable node**. If we order the eigenvalues so that $\lambda_1 > \lambda_2 > 0$, then we see that the growth is faster in the direction defined by \mathbf{v}_1 and so trajectories starting close to the equilibrium point tend to move away from the equilibrium in this direction.

- $\det J < 0$

The product of the eigenvalues is negative. This means that we must have one positive (call it λ_1) and one negative (λ_2) eigenvalue. We have exponential growth in the direction represented by λ_1 but exponential decay in the direction represented by λ_2 . This is called a saddle: points move towards the equilibrium in one direction, but away from the equilibrium in the other.




Figure 3: Left panel illustrates a stable node, with $\lambda_1 > \lambda_2 > 0$. Notice that trajectories leave the equilibrium in the direction of \mathbf{v}_1 . Right panel illustrates a saddle, with $\lambda_1 > 0 > \lambda_2$. Notice that trajectories approach the equilibrium in the direction of \mathbf{v}_2 , but leave in the direction of \mathbf{v}_1 .

- $\text{tr } J < 0$

This means that the sum of the eigenvalues is negative.

There are two possible cases:

- $\det J > 0$

The product of the eigenvalues is positive. This means that both λ_1 and λ_2 must be negative. In terms of the differential equation, we have exponential decay in the directions represented by both \mathbf{v}_1 and \mathbf{v}_2 . We call this situation a **stable node**. If we order the eigenvalues so that $\lambda_1 > \lambda_2$ (i.e. λ_2 is more negative), we see that the contraction in the \mathbf{v}_2 direction occurs more quickly than in the \mathbf{v}_1 direction. This means that trajectories approach the equilibrium along the line defined by the vector \mathbf{v}_1 .

- $\det J < 0$

The product of the eigenvalues is negative. This means that we must have one positive (call it λ_1) and one negative (λ_2) eigenvalue. As in the situation above, this corresponds to a **saddle**.




Figure 4: Left panel illustrates a stable node, with $0 > \lambda_1 > \lambda_2$. Notice that trajectories approach the equilibrium in the direction of \mathbf{v}_1 . Right panel illustrates a saddle.

- **Eigenvalues are Complex:** $(\text{tr } J)^2 < 4 \det J$

In this case $\text{tr } J$ equals the sum of the real parts of the two eigenvalues. But we know that these eigenvalues occur as a complex conjugate pair and so have equal real parts. The trace of J tells us whether the real parts of the eigenvalues are positive or negative.

- $\text{tr } J > 0$

The eigenvalues have positive real part, so we have exponential growth. Points move away from the equilibrium point, but do so in an oscillatory manner as discussed above. We call this an **unstable spiral**. Notice that the spiral may be elliptical in nature.

- $\text{tr } J < 0$

The eigenvalues have negative real part, so we have exponential decay. Points move towards the equilibrium point, but do so in an oscillatory manner as discussed above. We call this a **stable spiral**.




Figure 5: Stable (left panel) and unstable (right panel) spirals.

3.4 Borderline Cases

As mentioned above, borderline cases involve repeated eigenvalues (in which case we have $4 \det J = (\text{tr } J)^2$), situations in which one or more eigenvalues are zero ($\det J = 0$ and $\text{tr } J \leq 0$) or have zero real part ($\det J \geq 0$, $\text{tr } J = 0$).

If one or more eigenvalues is equal to zero, or have real parts equal to zero, then there is no motion in the direction(s) defined by the corresponding eigenvector(s). An important example is the **center**, when there is a pair of eigenvalues that are purely imaginary (recall that the signs of their imaginary parts must be opposite). The resulting motion involves points moving around on ellipses: there is no net motion towards or away from the equilibrium.

In the case when there is one zero eigenvalue and one real non-zero eigenvalue, points just move parallel to the direction defined by the eigenvector corresponding to the non-zero eigenvalue. They either move towards or away from the line described by the other eigenvector, depending on whether the sign of the eigenvalue is negative or positive. In this case, there is an entire line of equilibrium points.

The other borderline cases correspond to repeated eigenvalues, and can be viewed as limiting cases of nodes. If there are repeated eigenvalues, there may or may not be linearly independent eigenvectors. If there are, then the resulting equilibrium point is described as a **stable star** or an **unstable star** (depending on the sign of the eigenvalue). This is just like a node, but where the two rates of contraction or expansion are equal. The second possibility is that there is only one eigenvector. The resulting equilibrium is then known as a **type II node**.




Figure 6: Borderline cases: the center, the star (stable star shown) and type II node (unstable shown). These figures were taken from Appendix A of J.D. Murray's 'Mathematical Biology'.

4 Relationship Between Nonlinear and Linearized Behavior

Some important results tell us when we can expect the behavior of the linearization to give a qualitatively correct picture of the behavior of the full nonlinear system in the neighborhood of an equilibrium point.

Provided that we are not in the borderline cases mentioned above paragraph (i.e. the linearization describes a saddle, a node or a spiral) then the picture we get from the linearized model gives a very good description of the behavior of the nonlinear model, at least in the neighborhood of the equilibrium. In the borderline cases (e.g. center, star, type II node), the nonlinear terms will change the picture.

The linearization correctly characterizes the **stability** of different directions provided that the corresponding eigenvalue is not zero or has real part equal to zero. In particular, if the eigenvalue with largest real part is non-zero, our linearized analysis will correctly predict if the equilibrium is stable or is unstable. Our linearization is uninformative regarding stability if this eigenvalue (the **dominant eigenvalue**) has zero real part: in such cases, nonlinear terms determine stability.

Taken together, these statements explain why we didn't give so much attention to the borderline cases. For instance, a stable star in the linear picture will correspond to a stable equilibrium in the nonlinear model, but the nonlinearity will mean that it may exhibit node-like or spiral-like behavior. A center has neutral stability in the linear picture: nonlinear terms can lead to stable or unstable behavior.




Figure A.2. Summary diagram showing how $\text{tr } A$ and $\det A$, where A is the linearisation matrix given by (A.4), determine the type of phase plane singularity for (A.1). Here $\det A = f_x g_y - f_y g_x$, $\text{tr } A = f_x + g_y$, where the partial derivatives are evaluated at the singularities, the solutions of $f(x, y) = g(x, y) = 0$.

Figure 7: Summary of the two dimensional case: dynamical behavior in terms of the determinant and trace of the Jacobian matrix. This figure was taken from Appendix A of J.D. Murray's 'Mathematical Biology'.