Editor: Dianne P. O'Leary, oleary@cs.umd.edu

Partial Solution to Last Issue's Homework Assignment

BEETLES, CANNIBALISM, AND CHAOS: ANALYZING A DYNAMICAL SYSTEM MODEL

By Dianne P. O'Leary

Last issue's installment of Your Homework Assignment featured the final problem in Dianne O'Leary's popular long-running department. In this issue, she offers a partial solution to it.

By using a *dynamical system* to model a flour beetle's life cycle, we can estimate the key parameters that describe its behavior. In the last issue, we presented the model and defined six parameters.

PROBLEM 1.

To get some experience with this model, plot the populations L, P, and A for 100 days for three sets of data: b = 11.6772, μ_L = 0.5129, c_{el} = 0.0093, c_{ea} = 0.0110, c_{pa} = 0.0178, $L(0)$ = 70, $P(0)$ = 30, $A(0)$ = 70, and μ_A = 0.1, 0.6, and 0.9. Describe the behavior of the populations in these three cases as if you were speaking to someone who isn't looking at the graphs.

Answer:

Figure 1 shows the results. When μ_A = 0.1, the solution eventually settles into a cycle, oscillating between two different values: 18.7 and 321.6 larvae, 156.7 and 9.1 pupae, and 110.1 and 121.2 adults. Thus the population at four-week intervals is constant. Note that the peak pupae population lags two weeks behind the peak larvae population, and that the adult population's oscillation is small compared to the larvae and pupae.

For μ_A = 0.6, the population eventually approaches a fixed point: 110.7 larvae, 54.0 pupae, and 42.3 adults.

In the third case, μ_A = 0.9, there is no regular pattern for the solution, so it's called *chaotic*. The number of larvae varies between 18 and 242, the number of pupae between 8 and 117, and the number of adults between 9 and 94.

PROBLEM 2.

Let μ_L = 0.5, μ_A = 0.5, c_{el} = 0.01, c_{ea} = 0.01, and c_{pa} = 0.01. Plot A_{fixed}, L_{fixed}, and P_{fixed} for b = 1.0, 1.5, 2.0, …, 20.0. To compute these values for each b, use `fsolve`, started from the solution with c_{el} = 0, to solve the equations $\hat{\mathbf{F}}(\mathbf{x}) = \mathbf{F}(\mathbf{x}) - \mathbf{x} = \mathbf{0}$. Provide `fsolve` with the Jacobian matrix for the function $\hat{\mathbf{F}}$; on your plot, mark the b values for stable equilibria with plus signs.

Answer:

Figure 2 shows the results. For the stable solutions, if we start with population values near A_{fixed}, L_{fixed}, and P_{fixed}, we'll converge to these equilibrium values.

PROBLEM 3.

(a) Let μ_L = 0.5128, c_{el} = 0.0, c_{ea} = 0.01, and c_{pa} = 0.09. For μ_A = 0.02, 0.04, …, 1.00, use the LPA relations to determine the population for 250 cycles. On a single graph, plot the last 100 values as a function of μ_A to produce the bifurcation diagram.

(b) Determine the largest of the values μ_A = 0.02, 0.04, …, 1 for which the constant solution is stable (that is, well-conditioned).

(c) Explain why the bifurcation diagram isn't just a plot of L_{fixed} versus μ_A when the system is unstable.

(d) Give an example of a value of μ_A for which nearby solutions cycle between two fixed values. Give an example of a value of μ_A for which nearby solutions are chaotic (or at least have a long cycle).

Answer:

Figure 3 shows the bifurcation diagram. The largest tested value of μ_A that gives a stable solution is 0.58. If we perform the computation in exact arithmetic, the graph would just be a plot of L_{fixed} versus μ_A. When the solution is stable, a rounding error in the computation produces a

 Copublished by the IEEE CS and the AIP 1521-9615/07/$25.00 © 2007 IEEE

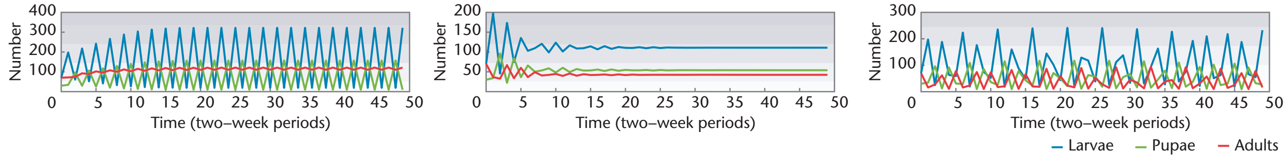

Figure 1. Results of the LPA model with three different choices of μ_A. Model predictions for $b = 11.6772$, $\mu_L = 0.5129$, $c_{el} = 0.0093$, $c_{ea} = 0.0110$, $c_{pa} = 0.0178$, $L(0) = 70$, $P(0) = 30$, $A(0) = 70$, and $\mu_A = 0.1$ (left), 0.6 (middle), and 0.9 (right). Number of larvae is in blue, pupae in green, and adults in red.

Table 1. Parameter estimates computed in Problem 4.

Colony	c_{el}	c_{ea}	c_{pa}	b	μ_L	μ_A	Residual
New: a	0.018664	0.008854	0.020690	5.58	0.144137	0.036097	5.04
Old: a	0.009800	0.017500	0.019800	23.36	0.472600	0.093400	17.19
New : b	0.004212	0.013351	0.028541	6.77	0.587314	0.000005	7.25
Old: b	0.010500	0.008700	0.017400	11.24	0.501400	0.093000	14.24
New : c	0.018904	0.006858	0.035082	6.47	0.288125	0.000062	4.37
Old: c	0.008000	0.004400	0.018000	5.34	0.508200	0.146800	4.66
New: d	0.017520	0.012798	0.023705	6.79	0.284414	0.005774	6.47
Old: d	0.008000	0.006800	0.016200	7.20	0.564600	0.109900	7.42

nearby point from which the iteration tends to return to the fixed point. When the solution is unstable, a rounding error in the computation can cause the computed solution to drift away. Sometimes it produces a solution that oscillates between two values (for example, when $\mu_A = 0.72$), and sometimes the solution becomes chaotic or at least has a long cycle (for example, when $\mu_A = 0.94$).

PROBLEM 4.

(a) Use `lsqnonlin` to solve the least-squares minimization problem, using each of the four sets of data in `beetledata.m`. In each case, determine the six parameters (μ_L, μ_A, c_{el}, c_{ea}, c_{pa}, and b). Set reasonable upper and lower bounds on the parameters and perhaps start the least-squares iteration with the guess $\mu_L = \mu_A = 0.5$, $c_{el} = c_{ea} = c_{pa} = 0.1$, and $b = 10$. Print the solution parameters and the corresponding residual norm.

(b) Compare your results with those that Brian Dennis and his colleagues computed (see `param_dl` in `beetledata.m`). Be sure to include a plot that compares the predicted values with the observed values.

Answer:

I used bounds of 0 and 1 for all parameters except b. For b, I used [0.1, 9.0]. The results are summarized in Tables 1 and 2 and contrast with the results of our model (new) with that of Dennis and his colleagues (old).

Figure 4 shows the predictions obtained from my parameters and from those of Dennis and his colleagues. Note that none of the models gives good predictions; we will see

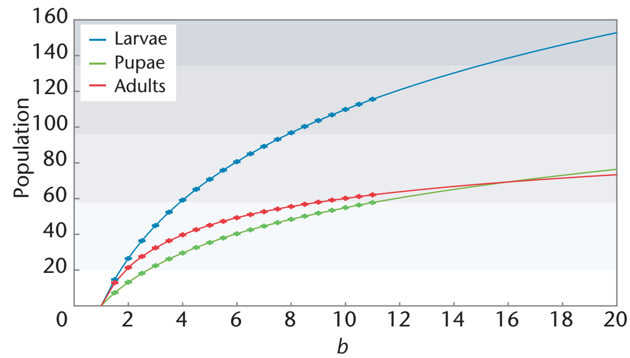

Figure 2. Equilibrium population as a function of b for $\mu_L = 0.5$, $\mu_A = 0.5$, $c_{el} = 0.01$, $c_{ea} = 0.01$, and $c_{pa} = 0.01$, $b = 1.0$, 1.5, 2.0, …, 20.0. Stable solutions are marked with pluses.

Table 2. Residual norms computed in Problem 4.

Colony	Norm of data vector	New residual	Old residual
Colony a	33.55	5.04	17.19
Colony b	33.70	7.25	14.24
Colony c	33.44	4.37	4.66
Colony d	33.68	6.47	7.42

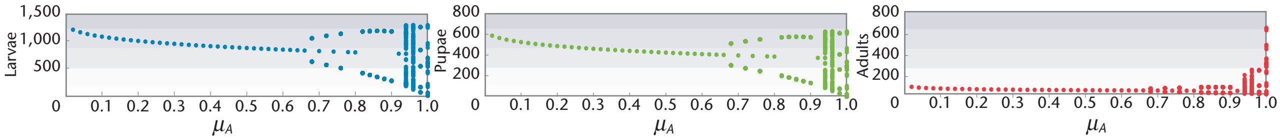

Figure 3. Bifurcation diagram for the data in Problem 3.

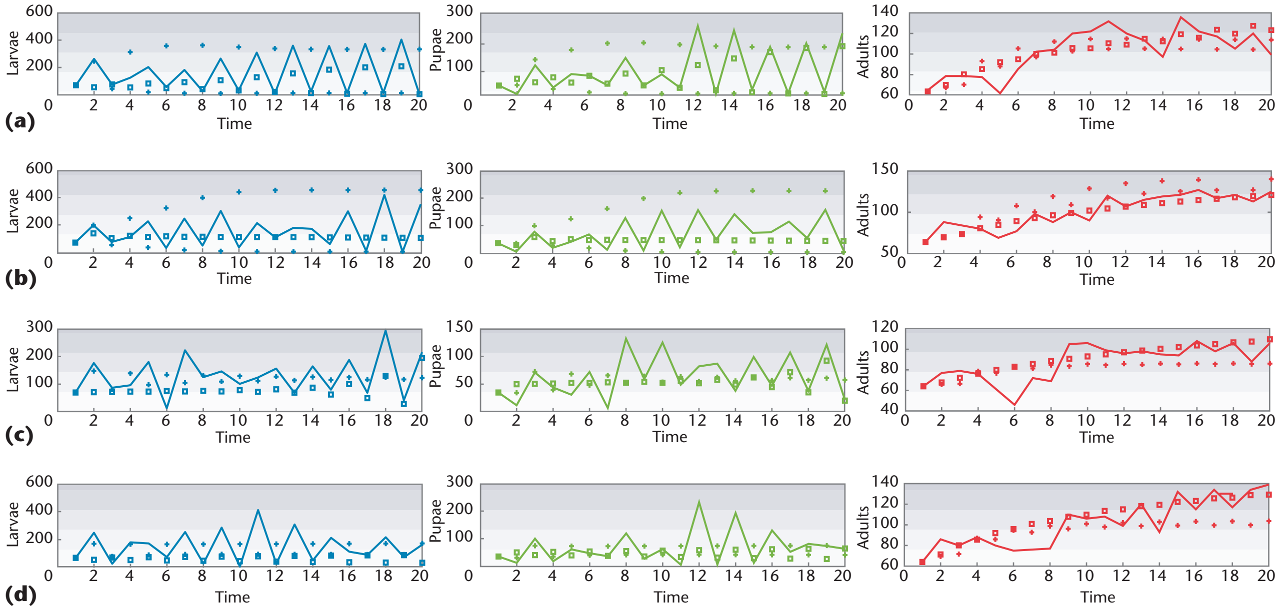

Figure 4. Model predictions for colonies (a) through (d). The solid line represents the data, the pluses are the predictions from Dennis and his colleagues, and the squares are our predictions.

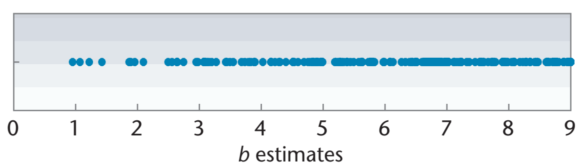

Figure 5. Values of *b* computed for colony b with 250 random perturbations of the log of the data, drawn from a normal distribution with mean 0 and standard deviation 1.

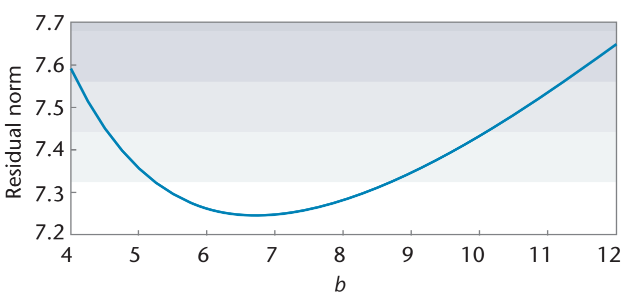

Figure 6. Changes in the residual as *b* is changed for colony b, leaving the other parameters fixed.

later that `lsqnonlin` finds only a locally optimal set of parameters—not necessarily the best choice overall. There is also the possibility of non-modeled errors in the data, and perhaps the beetles didn't respond well to the counting process.

PROBLEM 5

Consider the data for the second beetle colony. For each value $b = 0.5, 1.0, \ldots, 50.0$, minimize the least-squares function by using `lsqnonlin` to solve for the five remaining parameters. Plot the square root of the least-squares function versus b, and determine the best set of parameters. How sensitive is the function to small changes in b?

Perform further calculations to estimate the *forward error*—how sensitive the optimal parameters are to small changes in the data—and the *backward error*—how sensitive the function is to small changes in the parameters.

Answer:

When the data is randomly perturbed, the estimate of b for the second colony ranges from 4.73 to 6.83.

A larger upper bound for b tends to cause the minimizer to converge to a local solution with a much larger residual.

There are many ways to measure sensitivity:

- We might ask how large a change we see in b when the data is perturbed a bit. This is a *forward error* result.
- We might ask how large a change we see in the residual when the value of b is perturbed a bit. This is a *backward error* result.

To estimate the forward error, I repeated the fit after adding 50 samples of normally distributed error (mean 0, standard deviation 1) to the log of the counts. This is only an approximation to the error assumption usually made for counts—Poisson error—but by using the log function in their minimization, the authors are assuming that this is how the error behaves. Even so, the estimate shown in Figure 5 range from 1.00 to 9.00, quite a large change.

To estimate the backward error, I varied b, keeping the other parameters at their optimal values, and plotted the resulting residual versus b in Figure 6. We see that the residual isn't very sensitive to changes in b.

Then I minimized the residual as a function of the five parameters remaining after setting b to fixed values. From Figure 7, we conclude that for any value of b between 1 and 50, we can obtain a residual norm within 10 percent of the computed minimum over all choices of b. This model seems to give no insight into the true value of b.

But as a final attempt, I used a homotopy algorithm, repeating the computations from Figure 7, but starting each minimization from the optimal point found for the previous value of b. The resulting residuals, shown in Figure 8, are much smaller, and the b value is somewhat better determined—probably between 5 and 10. Even more interesting, the fitted model finally gives a reasonable approximation of most of the data (see Figure 9).

To check the reliability of these estimates, it would be a good idea to repeat the experiment for the data for the other three colonies and to repeat the least-squares calculations using a variety of initial guesses.

Dianne P. O'Leary is a professor of computer science and a faculty member in the Institute for Advanced Computer Studies and the Applied Mathematics Program at the University of Maryland. She has a BS in mathematics from Purdue University and a PhD in computer science from Stanford. O'Leary is a member of SIAM and AWM, and a fellow of the ACM. Contact her at oleary@cs.umd.edu; www.cs.umd.edu/users/oleary/.

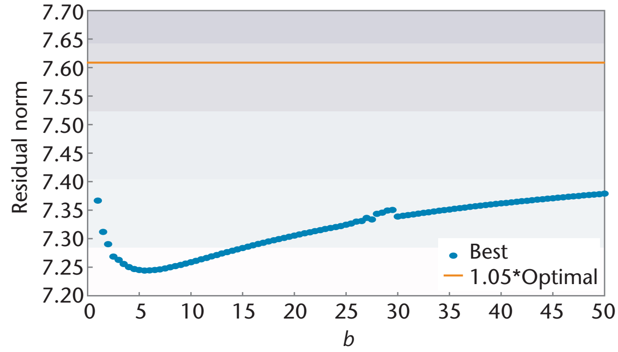

Figure 7. Best (smallest) residuals for colony b computed as a function of the parameter b (blue circles) compared with the red dotted line, indicating a 10 percent increase over the minimal computed residual.

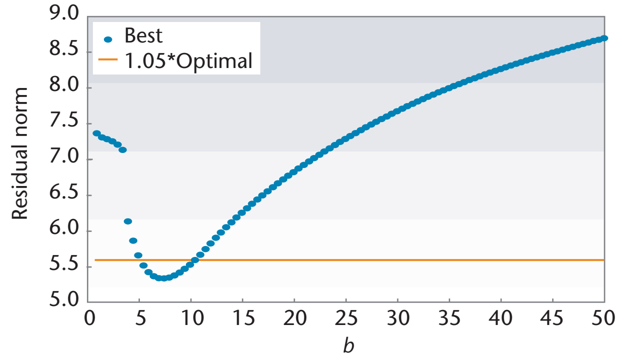

Figure 8. Best (smallest) residuals for colony b computed as a function of the parameter b (blue circles) compared with the red dotted line, indicating a 10 percent increase over the minimal computed residual, using homotopy.

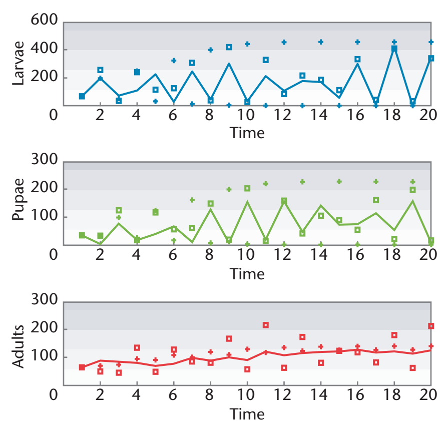

Figure 9. Revised model predictions for colony b, with parameters $c_{el} = 0.008930$, $c_{ea} = c_{pa} = 0$, $b = 7.5$, $\mu_L = 0.515596$, $\mu_A = 0.776820$. The solid line represents the data, the pluses are the predictions from Dennis and his colleagues, and the squares are our predictions.