Uncertainty Quantification of Optimal Control in Ecological Systems

Nina Chafee [1], Micah Chandler [2], Benjamin Lebdaoui [3], Meredith Orton [4]

[1] Haverford College, [2] University of North Georgia, [3] University of Montevallo, [4] LaGrange College

The Big Idea

Real-world decisions rely on models built from inherently noisy data. Noise propagates through the reconstruction (an inverse problem), making us less confident in the model and subsequent decision-making (optimal control). We propose a framework that propagates uncertainty from data through reconstruction into control, in order to quantify uncertainty in the outcome.

Defining the Problem

The local ecosystem partially consists of rabbits and wolves. Every few years, the rabbit population spikes, and our crops suffer. **How can we stabilize the rabbit population *without* starving the wolves?**

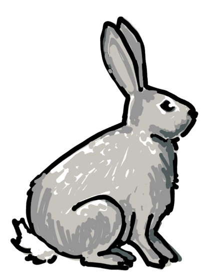

First, we can describe dynamics using a model F whose output y is time-dependent and relies on parameters θ.

The **Lotka-Volterra** model, $F(t;\theta) = [p(t), r(t)]$:

$$\begin{aligned} \dot{p} &= \theta_1 p - \theta_2 p\, r, & p(0) &= p_0 \\ \dot{r} &= -\theta_3 r + \theta_4 p\, r, & r(0) &= r_0 \end{aligned}$$

$$\begin{aligned} &\text{Output} && y = F(t;\theta) \\ &\text{Data} && d_i = F(t_i;\theta) + \epsilon,\ \epsilon \sim \mathcal{N} \end{aligned}$$

To predict the outcome of interventions, we need to accurately model the system.

INVERSE PROBLEM

GOAL: Solve the **inverse problem**; reconstruct the model that generated observed data.

PROBLEM: Noise in the data creates uncertainty in the model.

Incorporating Uncertainty

We use a Bayesian framework to quantify uncertainty, incorporating prior knowledge and observed data with Bayes' Rule: [4]

$$\Pi_{\text{post}}(\theta \mid d) = \frac{\Pi_{\text{like}}(d \mid \theta)\Pi_{\text{prior}}(\theta)}{\int \Pi_{\text{like}}(d \mid \theta)\Pi_{\text{prior}}(\theta) d\theta}$$

To create a representative sample of the **posterior distribution** of θ, we employ a Markov Chain Monte Carlo (**MCMC**) method: the Metropolis-Hastings algorithm.

Parameter Estimation

Estimates are based on twenty lynx and snowshoe hare population observations from 1900 to 1920. Predictions were made from 1920 to 1940, with a noise estimate calculated from residuals. [3]

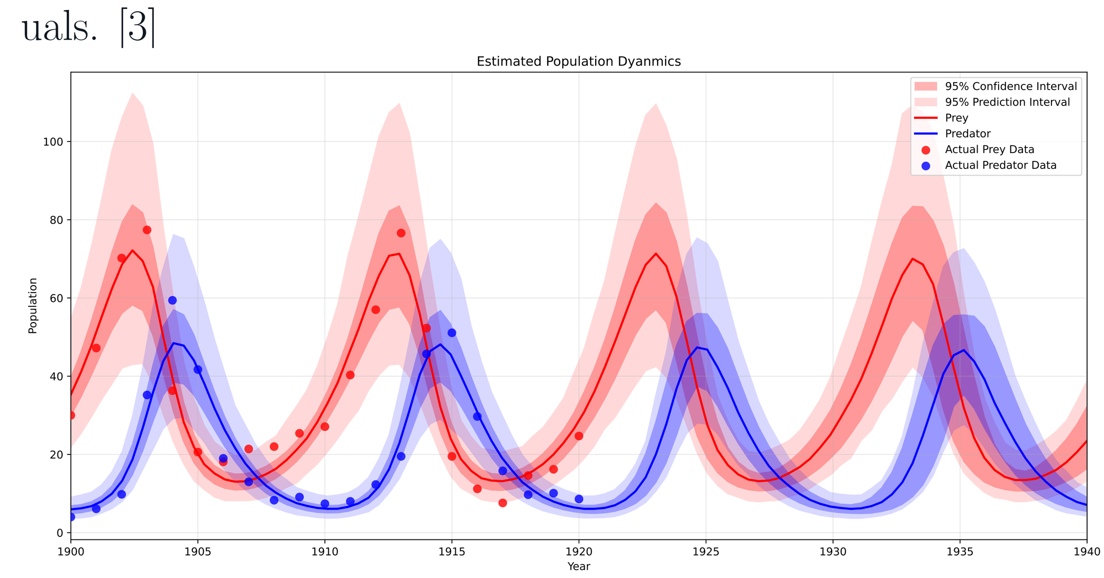

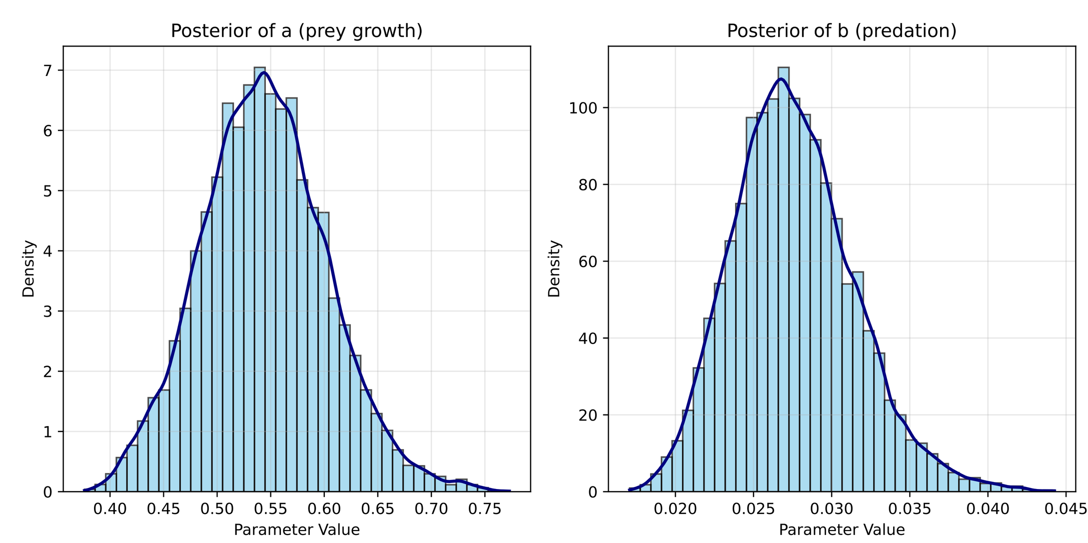

Example posterior distributions of θ_1 and θ_2 from MCMC

Guiding the System

OPTIMAL CONTROL

GOAL: Find the **optimal control** $\alpha(t)$ that efficiently generates desired behavior in the predator–prey system.

METHOD: Determine $\alpha(t)$ that minimizes a cost,

$$C(\alpha)_{x,t} = g(x(T)) + \int_0^T \mathcal{L}(x(s), \alpha(s), s) ds.$$

subject to dynamics and initial conditions of the model.

We want to guide the reconstructed model towards the steady state, $\left(\frac{\theta_3}{\theta_4}, \frac{\theta_1}{\theta_2}\right)$. Surprisingly, we can define $\alpha(t)$ as the optimal rate of reducing the wolf population.

The **controlled Lotka-Volterra** model:

$$\begin{aligned} \dot{p} &= \theta_1 p - \theta_2 p\, r, & p(0) &= p_0 \\ \dot{r} &= -\theta_3 r + \theta_4 p\, r - \alpha(t), & r(0) &= r_0 \end{aligned}$$

In this case,

$$g(x(T)) = \left(p(T) - \frac{\theta_1}{\theta_2}\right)^2 + \left(r(T) - \frac{\theta_3}{\theta_4}\right)^2$$

$$\int_0^T \mathcal{L}(x(s), \alpha(s), s) ds = \int_0^T \alpha(t)^2 dt$$

How can we solve this optimization problem?

Optimal Control Methods

PONTRYAGIN MAXIMAL PRINCIPLE (PMP) APPROACH

IDEA: Transform the optimization problem into a **system of differential equations** that satisfy boundary value conditions. [5]

A control $\alpha^*(t)$ is optimal if it minimizes the Hamiltonian:

$$\mathcal{H}(\alpha^*) \leq \mathcal{H}(\alpha) \quad \text{for all } \alpha \in \mathcal{A}.$$

The Hamiltonian combines system dynamics and costs:

$$\mathcal{H} = \mathcal{L}(p, r, \alpha, t) + \lambda_p \dot{p} + \lambda_r \dot{r}$$

where $\lambda_p(t)$, $\lambda_r(t)$ are costates corresponding to $p(t)$, $r(t)$.

- **State:** What the system is doing now
- **Costate:** How changes now affect future cost

The costate equations evolve backward in time:

$$\dot{\lambda}_p = -\frac{\partial \mathcal{H}}{\partial p}, \quad \dot{\lambda}_r = -\frac{\partial \mathcal{H}}{\partial r}$$

To find the optimal control, minimize $\mathcal{H}$:

$$\frac{\partial \mathcal{H}}{\partial \alpha} = 2\alpha - \lambda_r = 0 \quad \Rightarrow \quad \alpha^*(t) = \frac{\lambda_r(t)}{2}$$

Integrate the state equations forward and the costate equations backward. Iterate until the boundary conditions are satisfied:

$$\lambda_p(T) = 2\left(p(T) - \frac{\theta_3}{\theta_4}\right), \quad \lambda_r(T) = 2\left(r(T) - \frac{\theta_1}{\theta_2}\right)$$

This yields a unique optimal control $\alpha^*(t)$.

PSEUDO SPECTRAL APPROACH

IDEA: Change the controlled model (a system of ODEs) to an **algebraic system** with polynomial approximations. [2]

Use orthonormal Legendre polynomials

$$\phi(t) = \begin{bmatrix} P_1(t) & \dots & P_N(t) \end{bmatrix}$$

to approximate the control $\alpha(t)$ and state $x(t)$:

$$\alpha(t) \approx \sum_{j=1}^{N} a_j P_j(t) = a^T \phi(t)$$

$$x(t) \approx \begin{bmatrix} \sum_{j=1}^{N} x_{1,j} P_j(t) \\ \vdots \\ \sum_{j=1}^{N} x_{m,j} P_j(t) \end{bmatrix} = X\phi(t)$$

Enforce dynamics at points t_k:

$$\begin{aligned} &X\dot{\phi}(t_k) = f(X\phi(t_k), a^T\phi(t_k), t_k), \quad k = 1, \dots, M \\ &X\phi(0) = X_0 \end{aligned}$$

This yields an algebraic system we solve numerically, e.g. with GEKKO[1], for X and a. (MATH)

Comparing Approaches

UNCERTAINTY PROPAGATION

GOAL: Quantify uncertainty in optimal control

METHOD: Run optimal control methods on sample set of the posterior distribution of θ

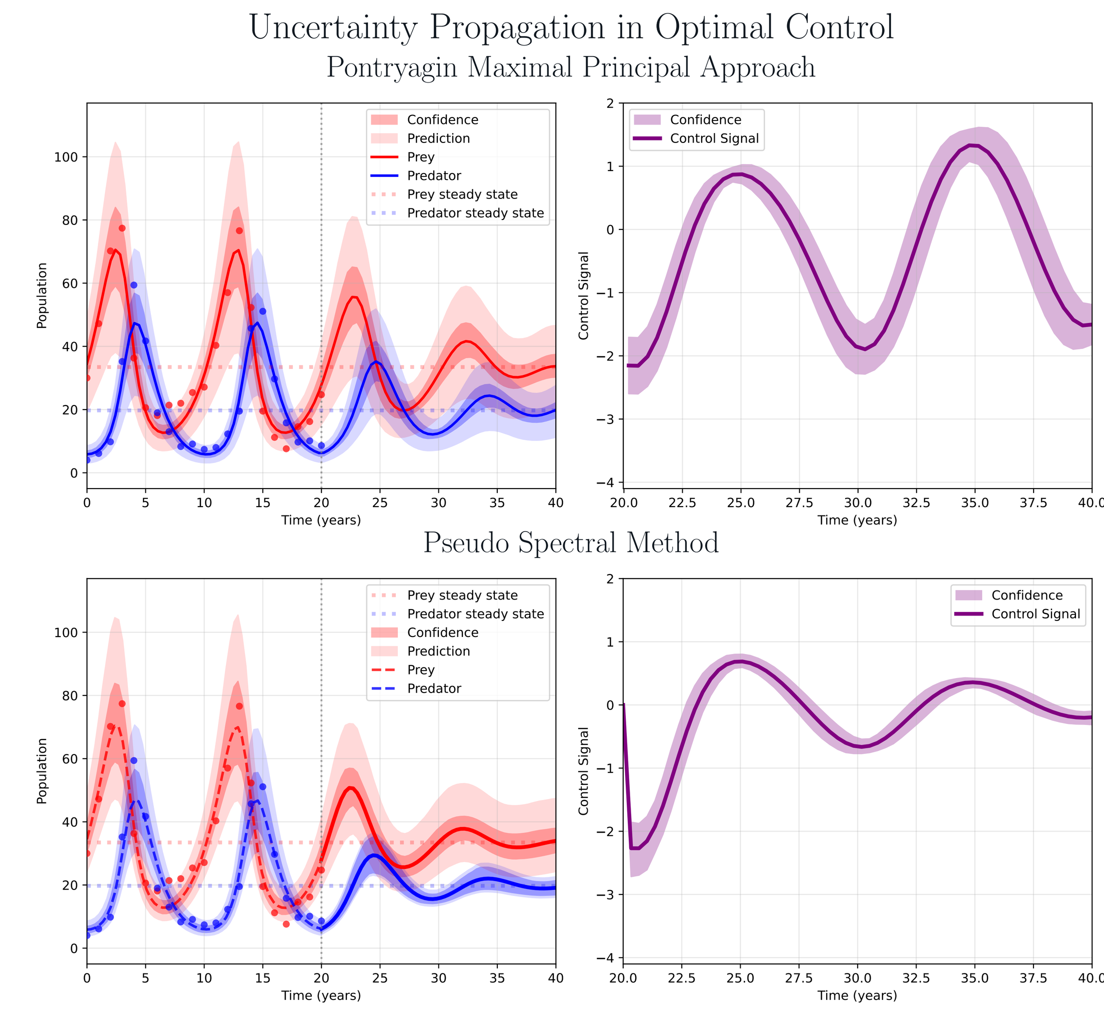

Population dynamics with 95% confidence and prediction intervals (left) and control signal with 95% confidence (right).

CONCLUSIONS:

- Estimated posterior distributions of parameters
- Successfully applied optimal control to guide Lotka-Volterra model towards stability
- Quantified uncertainty using posterior distributions

FUTURE WORK:

- Find **unique optimal control** based on entire posterior distribution rather than computing for each MCMC sample, quantify uncertainty from there
- Repeat pipeline with **higher dimensional ODEs**, e.g. SIR for disease transmission or incorporating eggs into Lotka-Volterra

Acknowledgments

This work is supported by the NSF award DMS-2349534. The authors would especially like to thank our advisor John Darges, as well as Emory University for facilitating the Models Meet Data REU 2025.

References

[1] Logan D. R. Beal, Daniel C. Hill, R. Abraham Martin, and John D. Hedengren. GEKKO Optimization Suite. *Processes*, 6(8):106, 2018. 10.3390/pr6080106

[2] David Benson. *A Gauss Pseudospectral Transcription for Optimal Control.* PhD thesis, Massachusetts Institute of Technology, 2005.

[3] Greg Brunkhorst. *ODE Lotka-Volterra With Bayesian Inference in Multiple Ways.* In: *PyMC Examples.* Edited by PyMC Team. DOI: 10.5281/zenodo.5654871

[4] Ralph C. Smith, *Uncertainty Quantification: Theory, Implementation, and Applications.* SIAM, 2024.

[5] Suzanne Lenhart and John T. Worman. *Optimal Control Applied to Biological Methods.* Taylor & Francis Group, 2007.