A Control Framework for Immunology: Threat Detection, Learning, and Stability

Matthew M. Peet*, Peter S. Kim and Peter P. Lee
*Illinois Institute of Technology

IEEE Conference on Decision and Control
Orlando, FL

December 13, 2011

IEEE

How To Recognize a Threat?

The Innate Immune Response

Threats: viruses, bacteria, parasites

Detection: Pattern Recognition Receptors (PRRs) identify Pathogen-Associated Molecular Patterns (PAMPs).

- TLR3 recognizes double-stranded RNA (viruses)
- TLR4 recognizes polysaccharides (bacteria)
- TLR5 recognizes bacterial flagellin
- TLR9 recognizes unmethylated CpG-containing DNA (common in viruses and bacteria)

Response: Macrophages, Dendritic Cells attack pathogens, amplify immune response, and recruit monocytes.

- Activation (Phagocytosis, Lysis)
- Cytokine signaling attracts monocytes (yield more DCs and MΦs).
- Cytokine signaling causes inflamation.
- Antigen presentation

Problems with Innate Response

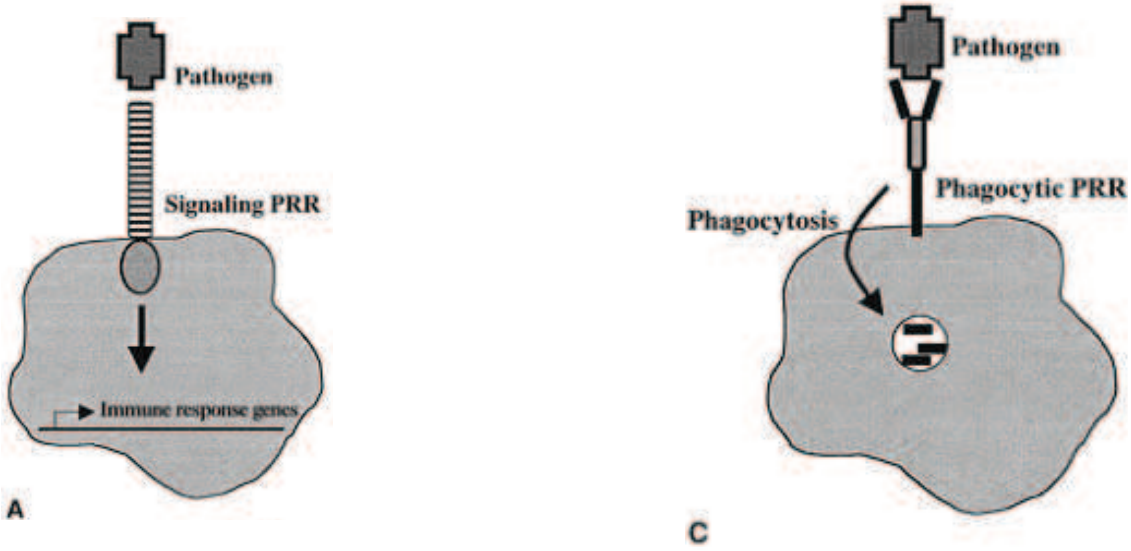

Paul, Fundamental Immunology

Problems with innate immunity:

- Slow
- No immunity
- Not robust
- No response to cancer

The Adaptive Immune System?

A Secondary System

Adaptive Immunity is new.

- Not present in plants

Several Functions

- Respond quickly to known threats - Immunity
- Identify threats missed by PRRs

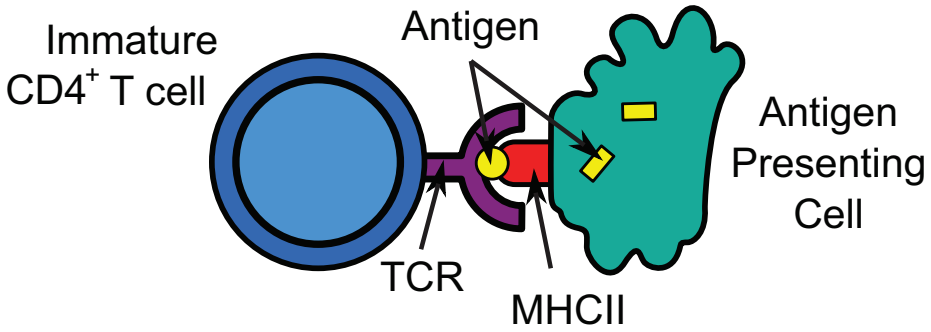

Figure: T Cell Receptors are only bind with one antigen (peptide)

The key to adaptive immunity is that it is *antigen-specific*.

- The adaptive response targets a single biological marker (antigen).
- In contrast to PRR defense, which targets entire classes of cells.

The Adaptive Immune System

How does it work?

Antigen-Presenting Cells (APCs) sweep up antigens

- Macrophages, Dendritic Cells, B-cells
- Antigens are presented to T cells

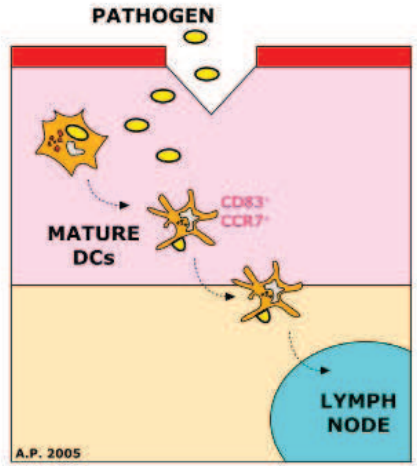

Response: T cells train B cells and killer T cells

- B cells produce antibodies which bind to a single type of antigen.
- Killer T cells induce apoptosis in infected cells.

In this talk, we focus on the T cell dynamics.

The Adaptive Immune System

The Decision-Making Process

- Should a presented antigen be targeted?

Congressional Committee: Decision-makers congregate in Lymph nodes.

- Helper T cells vote to amplify immune response.
- Regulatory cells vote to suppress immune response.
- Memory cells of both types can override decisions.

Constraints

- All antigens look the same (more or less).

Consequences

- Targeting of self-antigens results in auto-immune disease.
  - Type-I diabetes; graph vs. host; allergies; septic shock.
- Tolerance of hostile pathogen results in chronic disease.
  - Cancer, HIV, parasites.

The Adaptive Immune System?

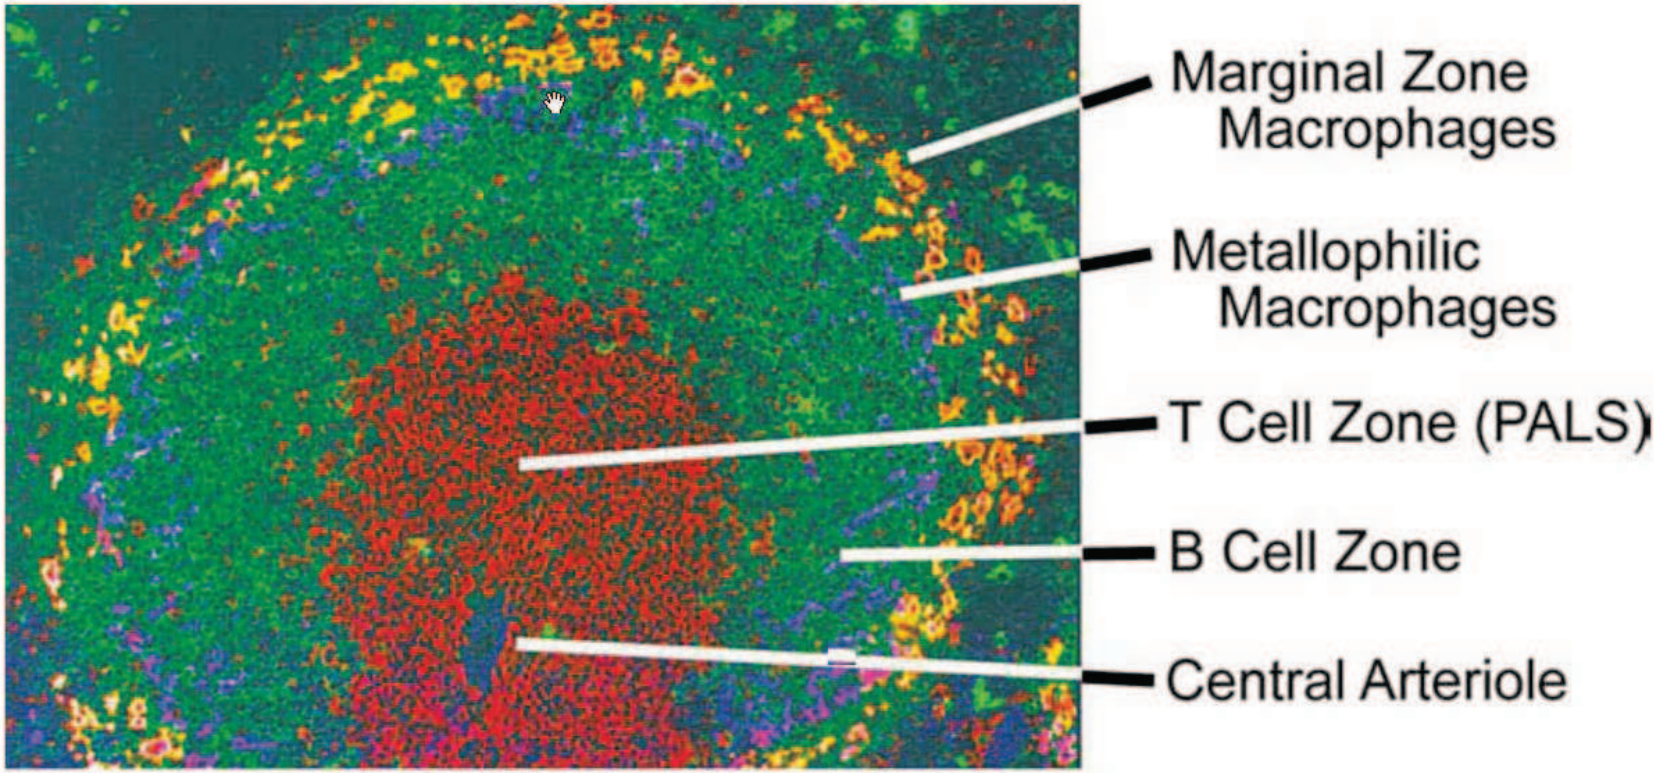

Figure: Decision-Making in the Lymph Nodes (`C. Zindle`)

Outline of Our Model

Direct Modeling of the immune system is impossible/useless.

- An emerging field with lots of uncertainty.
- Time-series data not available.
- Too much complexity.
  - Nonlinear with thousands of possible states

We will pick our fights carefully

- Self-nonself discrimination.
- Threat communication and triggering.
- Maintain stability of response.

Basics of the Control System

What are we looking for?

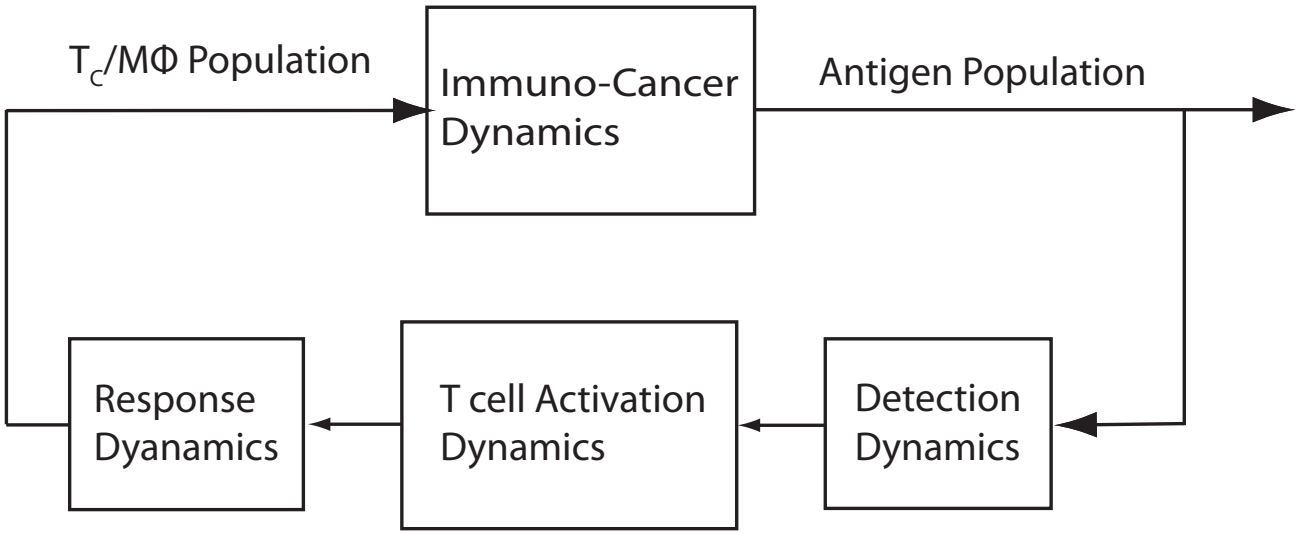

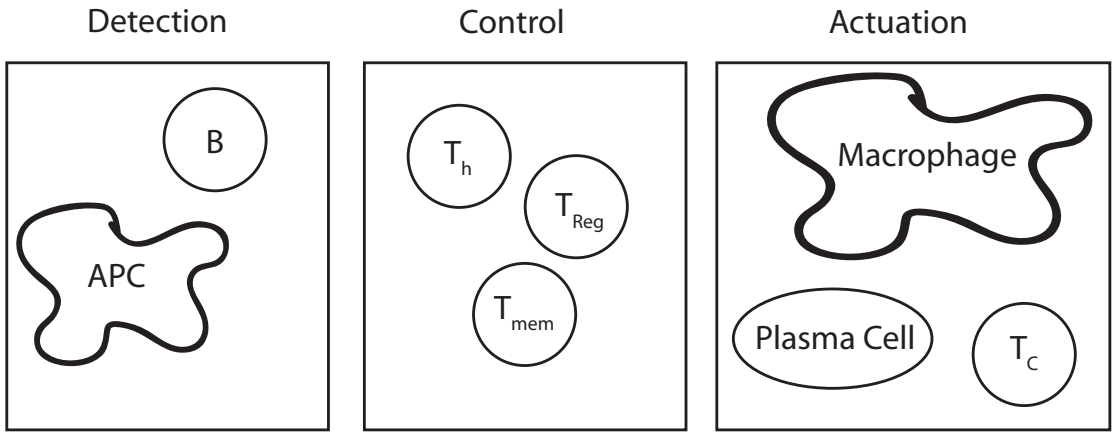

A Basic Model

Proportional Response: Sensor

The first step is a common model of proportional response.

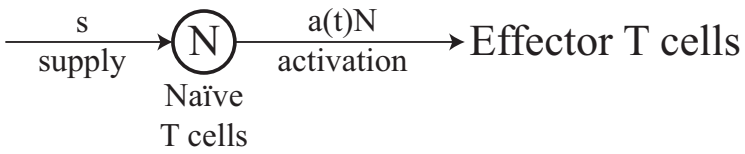

Hypothesis: A stabilized reservoir of naïve T cells is available.
Sensor: Helper Cell Dynamics

$$\frac{dE(t)}{dt} = R_{Ea}Na(t) - d_E E(t),$$

N is the size of the pool of Naïve T cells. R_{Ea} is a reaction rate. d_E is death/loss rate. $a(t)$ is antigen concentration. System at steady-state has

$$E(t) = \frac{NR_{Ea}}{d_E}a(t)$$

Threat Detection

Derivative Control

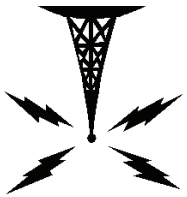

Friendly Objects Don't Move

Consider first-order differential approximation

- Trigger an alarm if:
  - $\dot{x}(t) \cong \frac{x(t)-x(t-\tau)}{\tau} \neq 0$

More generally: Define threat based on behavior

- We consider rate of change in antigen concentration.

Threat Detection: Derivative Response

First Order Approximation

Observation: The T_{reg} response is delayed.

Assume T_{reg} and T_h populations both in steady state.

$$E(t) = K_E a(t), \qquad R(t) = K_R a(t-\tau)$$

Regulator cells de-activate helper cells.

$$\begin{aligned}\frac{dE(t)}{dt} &= r_{Ea} a(t) E(t) - r_{RE} R(t) E(t) \\ &= \left(r_{Ea} a(t) - K_{RE} a(t-\tau)\right) E(t)\end{aligned}$$

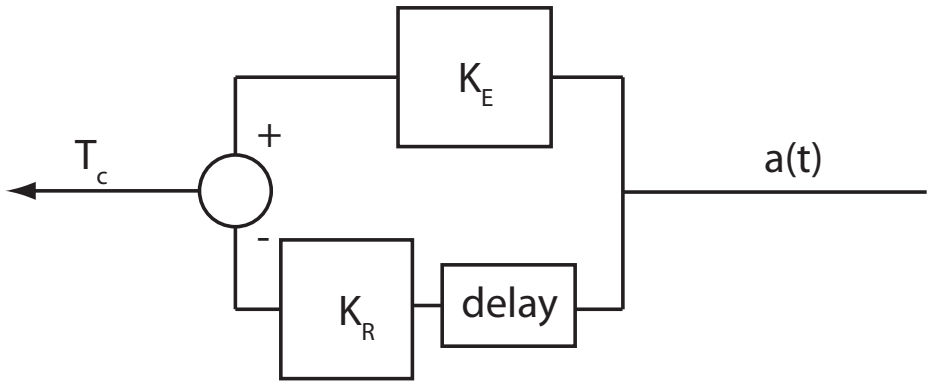

Now, include the steady-state actuator dynamics

The Activation Dynamics: Derivative Gain

Actuator Dynamics

Proportional-Differential Control

$$\frac{dE(t)}{dt} = K_1 a(t) E(t) + K_2 \frac{(a(t) - a(t-\tau))}{\tau} E(t)$$
$$\cong (K_1 a(t) + K_2 \dot{a}(t))\, E(t)$$

where

- $K_1 = (r_{Ea} - K_{RE})$ and $K_2 = \tau K_{RE}$.

If the system is in balance:

- If $r_{Ea} \cong K_{RE}$, there is no proportional response.
- Further, if a threat is persistent, $a(t) = a(t-\tau)$, then $\dot{E}(t) = 0$, so **the threat is ignored.**

Conclusion

- No cell is able to determine threat level.
- Threat is determined by overall balance of T_{reg}/T_{eff} populations.

Return to Motion Detection

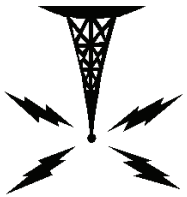

Problem: The signal $x(t) - x(t-\tau)$ is not strong or persistent.

Solution

- Use $x(t) - x(t-\tau)$ as a trigger:

The Activation Dynamics: Trigger Mechanism

A Switching Model

Observation: T_h cell proliferation is driven by cytokine IL-2.

$$\frac{dp(t)}{dt} = r_p E(t) - d_p p(t).$$

- p is concentration of IL-2.
- Assume dynamics are fast.

$$p(t) = \frac{r_p E(t)}{d_p}.$$

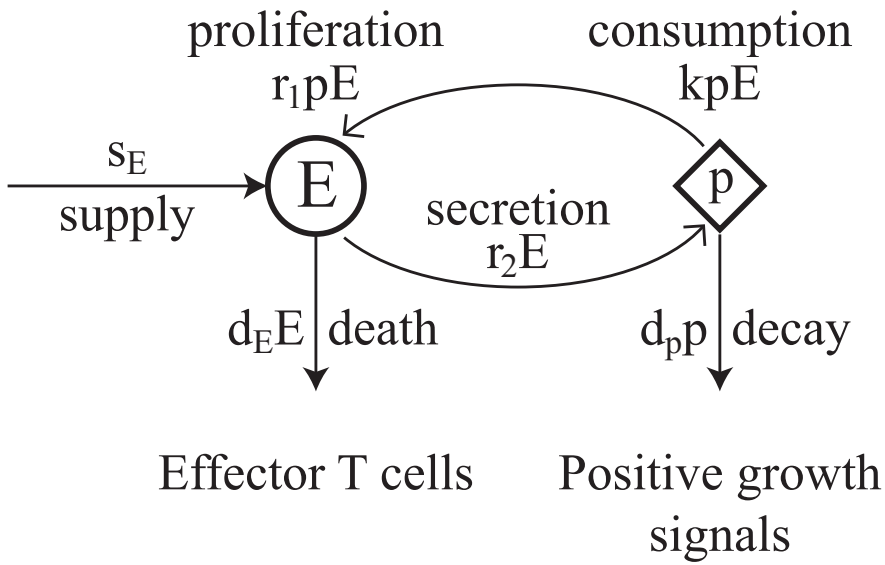

Figure: Release and Absorption of Growth Signals

Effector Cell Dynamics become

$$\frac{dE(t)}{dt} = -d_E E(t) + r_E E(t)^2 \frac{r_p}{d_p} + u(t)$$

The Activation Dynamics: Trigger Mechanism

Stability Threshold

The one-dimensional Effector Dynamics: $\dot{E}(t) = f(E(t)) + u(t)$

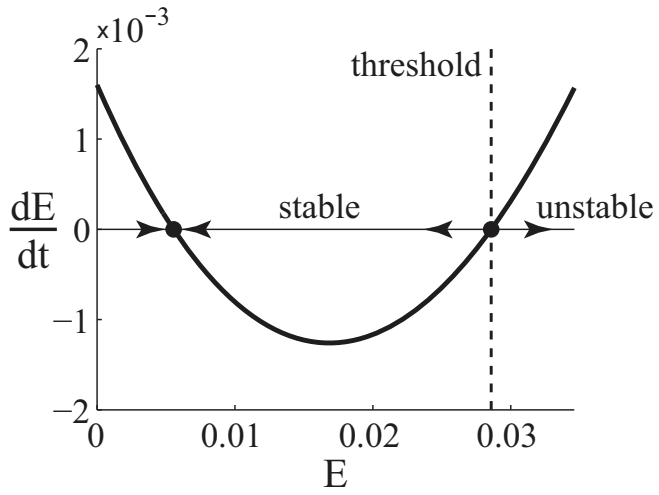

When $u(t) < u_{trig}$:

- Two Equilibria : one stable, one unstable.
- $u_{trig} = d_E^2 \frac{d_p}{4 r_p r_E}$

When $u(t) > u_{trig}$:

- No equilibria, exponential growth.
- If $u(t)$ returns to 0, growth continues anyway.

The Activation Dynamics: Containment

Integral Control

Unbounded (unstable) exponential growth is unrealistic.

- We model contraction using a long-lived iT_{reg} population which emerges from the helper T cell population.

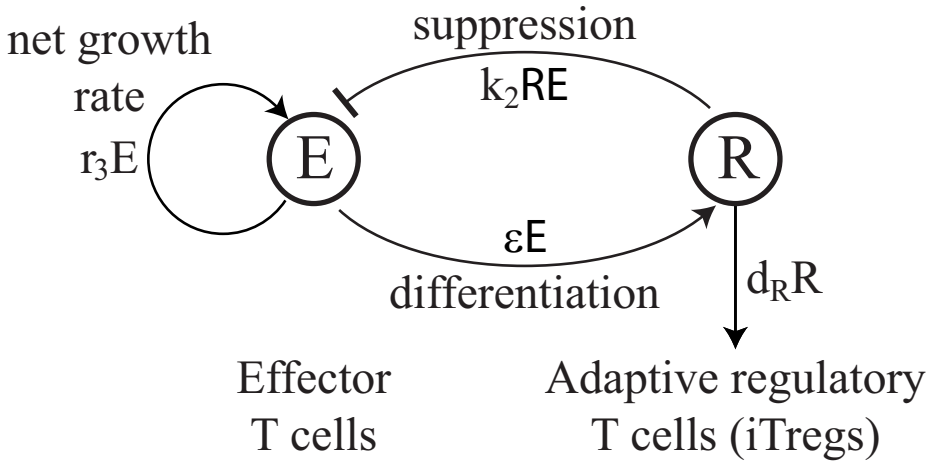

$$\frac{dR_i(t)}{dt} = \nu_R p(t) E(t) - d_{Ri} R_i(t).$$

- ν_R is the emergence rate via cytokines.

The Activation Dynamics: Containment

Integral Control

If we assume the death rate d_R is relatively small. Then we have

$$R_i(t) \cong K_i \int_0^t E(s)ds$$

Question: Is this enough to overcome the positive feedback loop?
To answer this we use **Sum-of-Squares** Optimization

- An approach to optimization over the cone of positive polynomials
- Find a Lyapunov function $V(x) \geq \epsilon \|x\|^2$
- With Negative Derivative:

$$\nabla V(x)^T f(x) \leq -\alpha \|x\|^2$$

Regions of Stability

Lyapunov Stability Analysis

- We find a degree 6 Lyapunov function.
- Use nominal values of the parameters.

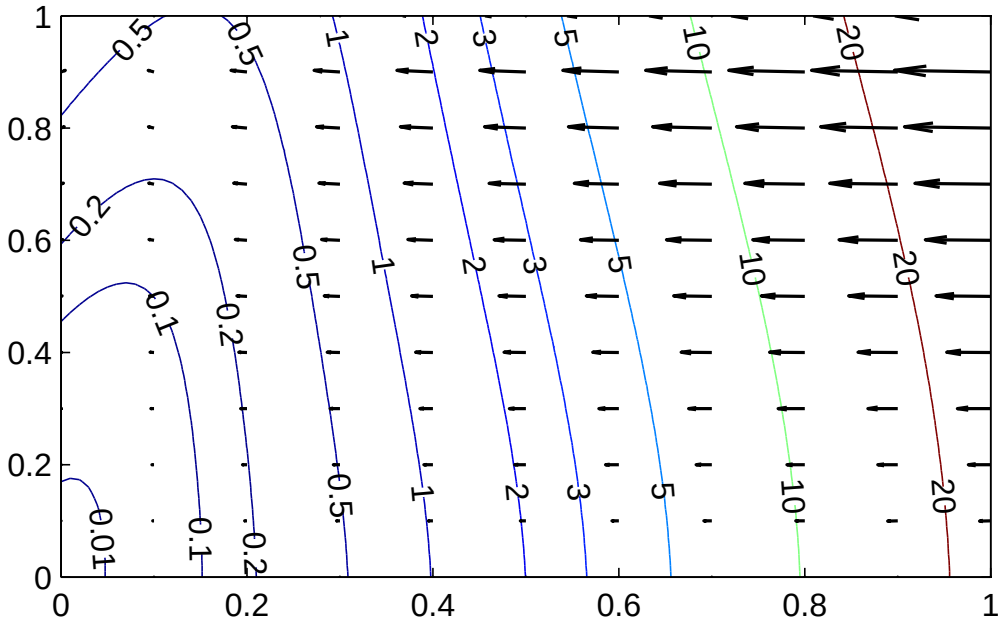

Figure: Lyapunov Level Sets and Vector Field: Helper vs. Regulatory Cell Concentration

Regions of Stability

We can automate the search over the parameter space.

- ν_R is the differentiation rate of iT_{reg} cells
- r_{RiE} is the suppression rate of helper cells by iT_{reg} cells

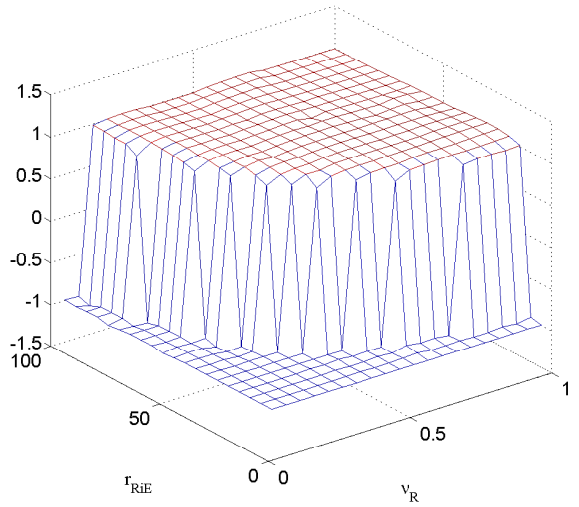

Figure: Stability for ν_R vs. r_{RiE}. Generated from SeDuMi on a grid. 1 implies stability. -1 means indeterminate

Parameter Region of Stability:

$$\nu_R \cdot r_{RiE} > 12.$$

Why is the Control Perspective important?

Consider the idle system on an automobile

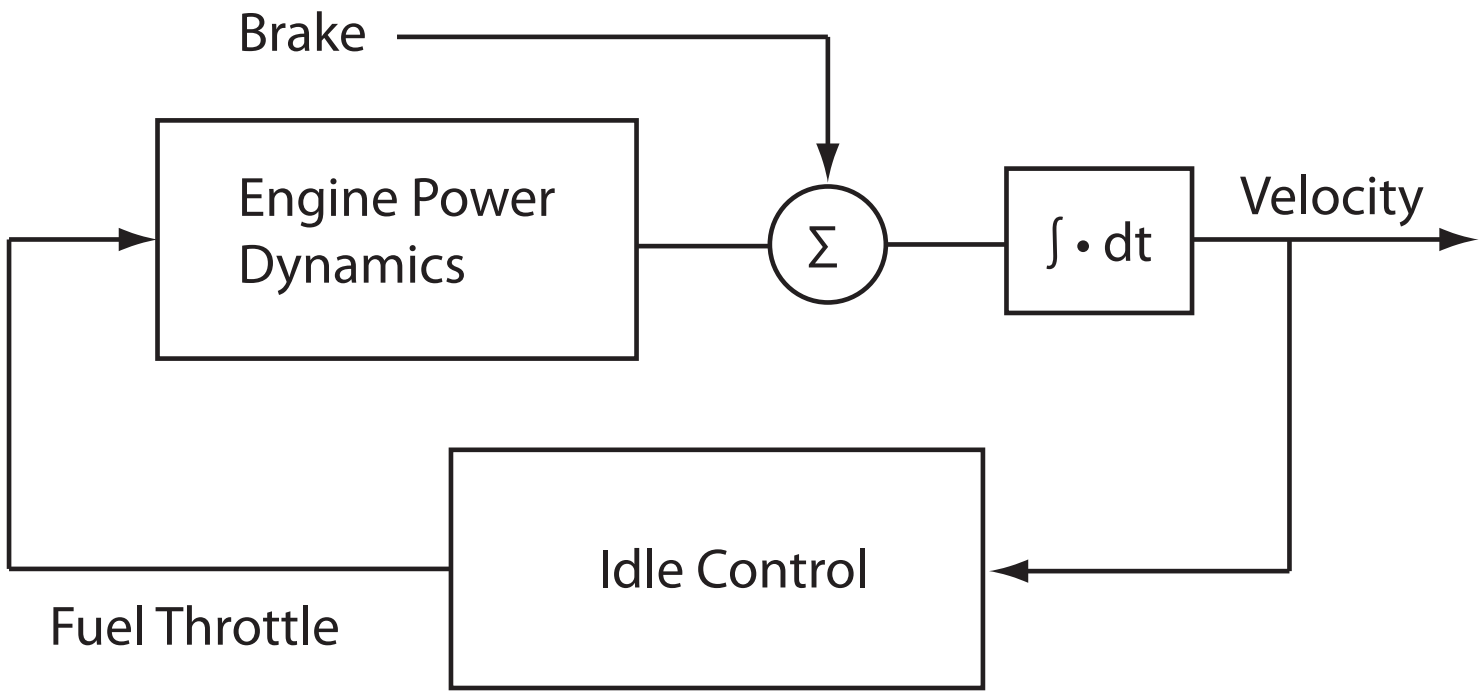

Figure: Illustration of the automotive idle control system

For a malfunctioning automotive idle: What is the better solution -

- Apply the Brakes?
- Re-calibrate the fuel sensor?

Conclusion

Modeling Immune Response as a Control System

The System Responds to Behavior

- Optimal dosing strategies may induce tolerance
  - Reduce rejection in transplantation
- Experimental tests in preparation

Ongoing Work:

- Modeling Memory.
- Optimal Control theory - Modeling Evolution.

Web Site:

`http://mmae.iit.edu/~mpeet`