Machine Learning and Wildfire

Using simulated landscapes to learn fire suppression policies

Optimal Wildfire Suppression

- Good Fires vs Bad Fires
- Saving catastrophe for the future
- Managing for Risk

Wildfire Suppression Policies

- A landscape changes after every fire
  - whether we suppress the fire or not
- The effect of each choice now is based on the interaction of prior choices with the evolving landscape
- A small change in a choice-making policy can lead to very different states, and affects what the policy will have to work with in the future.

Searching for “Better” Policies

- “Self-interacting” decisions
  - Cannot look at single, independent state transitions because the result is very short-sighted
  - There may be hidden future rewards
- State Stability
  - Especially dynamic stability
  - Otherwise the time-horizon we choose determines the apparent value of a policy

Mathematics

- Expected Value of using a policy...
  - We cannot visit every possible state
  - Even Monte Carlo simulations take too long

A policy as a probability

- Probability of taking action X
- Following this type of policy produces a series of probabilistic steps
  - Joint Probability of taking this “pathway”
- Changing the policy will change the suggested joint probability of making the decisions we already made.

Logistic Function Policy

State values = [*a*, *b*, *c*]

Policy parameters = [*x*, *y*, *z*]

decision probability = logistic(a*x + b*y + c*z)

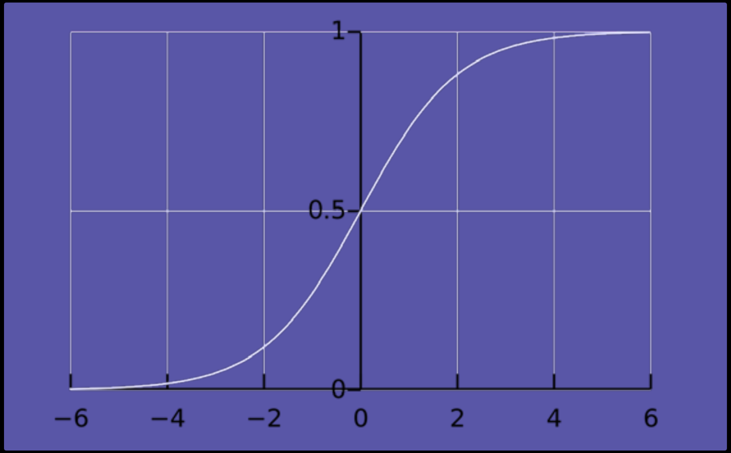

Objective Function

Average over all pathways:

$$\frac{\textit{joint probability under "this" policy}}{\textit{joint probability under original policy}} \times \textit{net present value of the pathway}$$

- The actions we make in a pathway don't change…
  - Only the policy's suggested probability of *taking* those actions
- The "weighted" net present value can be adjusted up or down by changes to the numerator policy

Optimization

- The objective function derivative is calculable
  - Gradient Descent
- Once a local optima is found, new pathways can then be generated using the policy associated with that optima
  - rinse and repeat...

Our Simulator

- 129x129 grid of forest cells
  - Timber age and associated value
  - Dead fuels on the ground which accumulate
- Weather
- Logging
- Firespread
- Economic Values:

  Timber cut - Suppression Costs

Current Results

- Extremely sensitive to the original policy
- Very low joint probabilities are problematic
- Overfitting

BUT

- Responsive to large changes in suppression cost.

Continued Development

- Move up to a larger “real-world” simulator
- Improve mathematics for the machine learning components
- Find out of the machine learning algorithms can find a balanced policy - suppressing some fires, letting others burn

High Suppression Cost

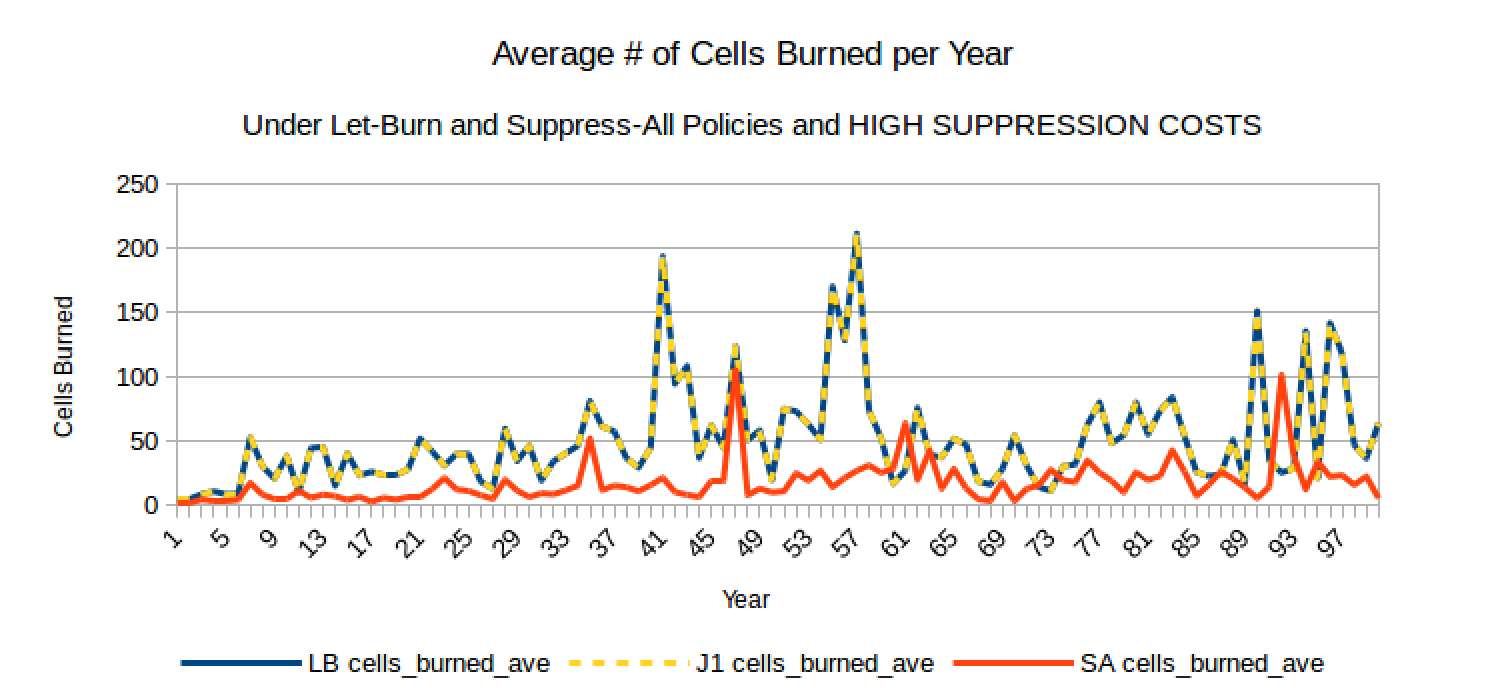

Zero Suppression Cost

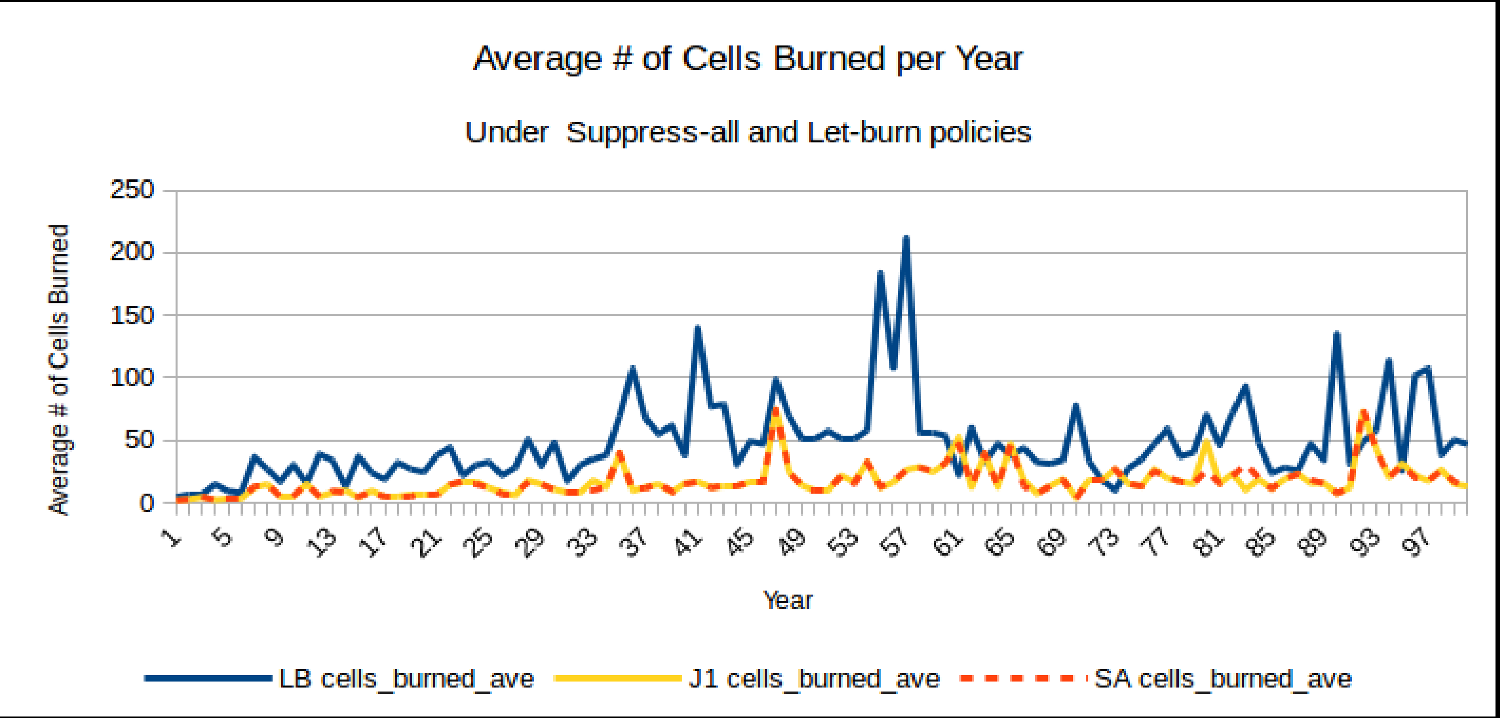

Thanks!

To my Collaborators

Rachel Houtman

Sean McGregor

Claire Montgomery

Tom Dietterich

Ron Metoyer

Chris Lauer

Our Funding Agencies

NSF

USFS

ARCS Foundation

OSU Provost