
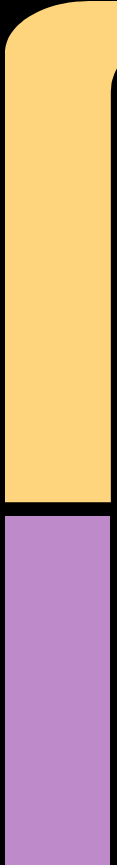

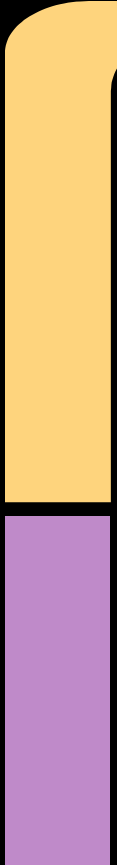


Machine Learning and Wildfire

Using simulated landscapes to learn
fire suppression policies

- 
- ◆ Background
 - ◆ My simulator
 - ◆ How to “learn” a Policy
 - ◆ Preliminary Results
 - ◆ Next Steps
- 

- 
- ◆ Background
 - My simulator
 - How to “learn” a Policy
 - Preliminary Results
 - Next Steps
- 



Why Wildfire Policies?





What is Machine Learning?





◆ Background

◆ My simulator

How to “learn” a Policy

Preliminary Results

Next Steps







The Simulator

Featuring:

- Tree Growth
- Timber Harvest
- Fuel Build-up
- Firespread

Important Notes:

- Spatially connected
- Abstract Models
- Abstract Units

- 
- ◆ Background
 - ◆ My simulator
 - ◆ How to “learn” a Policy
- Preliminary Results
- Next Steps
- 



Policies

What we'd like to know about a new fire:

- An idea of how likely this will be a “bad” fire in the long run
- A suggestion of whether to suppress this fire or let it burn
- An idea of why the computers ultimately suggest what they do

Policies - When to Suppress?



Policies - When to Suppress?


Wind Speed

- Increases fire-spread
- Increases the likelihood we'll suppress
- **Positive Weight**

Moisture Content


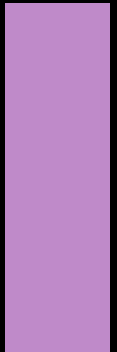
- Decreases fire-spread
- Decreases the likelihood we'll suppress
- **Negative Weight**

Policies - When to Suppress?

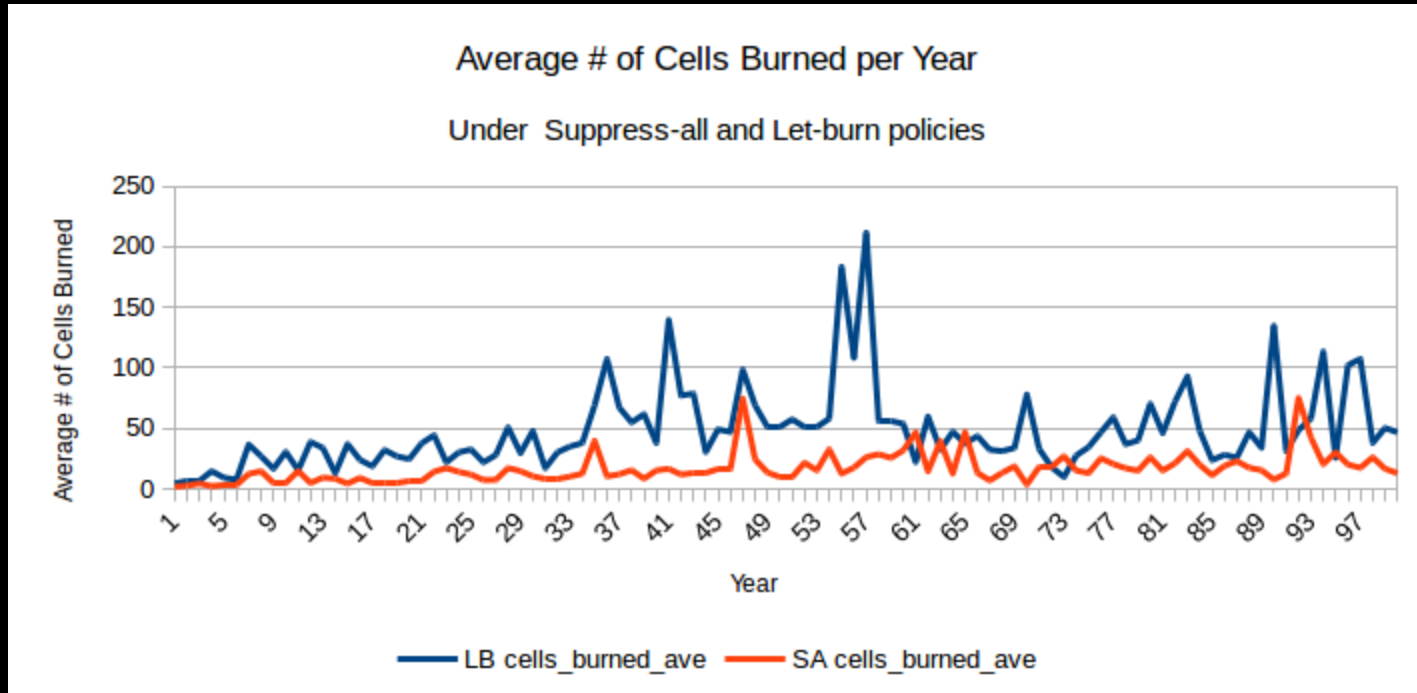


Generating a Policy


1. Simulate a bunch of landscapes for many years each.
2. Wait...
3. Wait...
4. Wait...
5. Policy = [1.5, -4.2, 3.3, -0.02]
[wind, rain, fuels, etc...]

- 
- ◆ Background
 - ◆ My simulator
 - ◆ How to “learn” a Policy
 - ◆ Preliminary Results
- Next Steps
- 


Outcomes of Different Policies



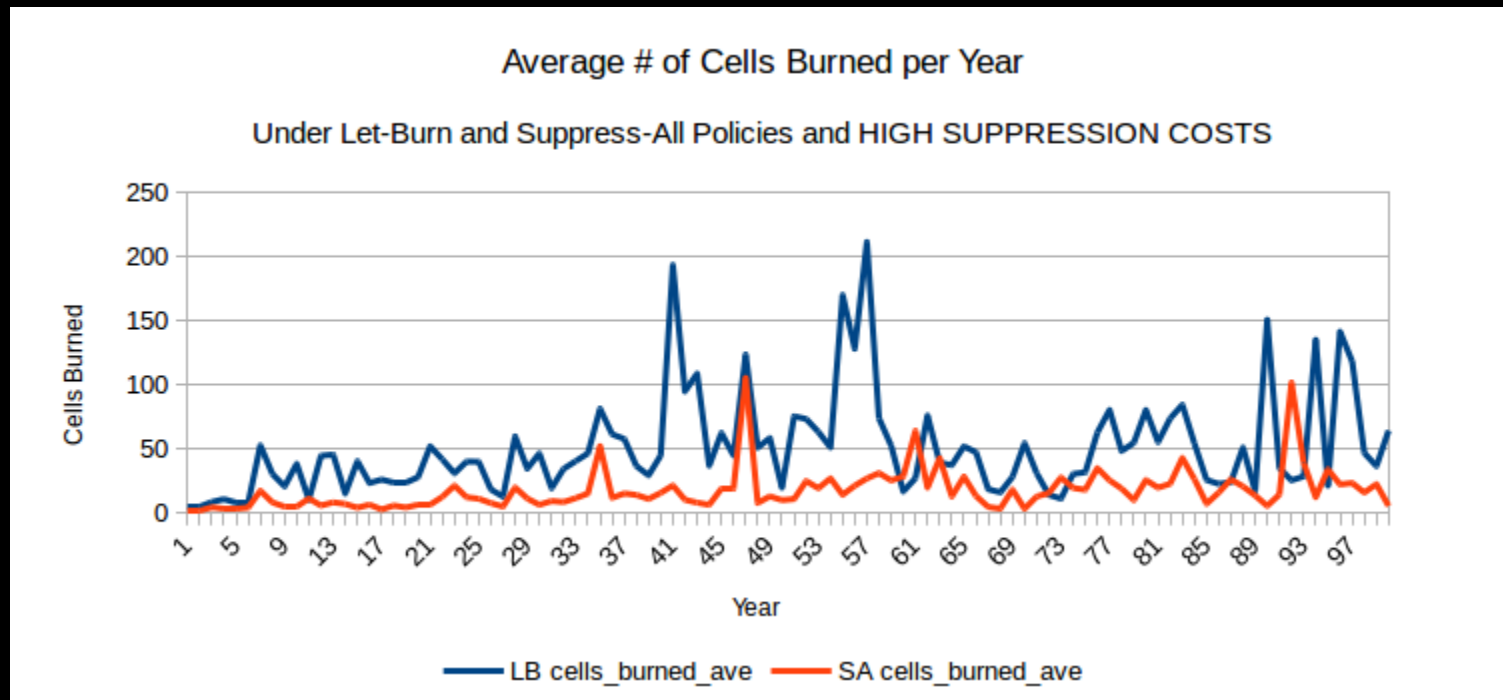
Outcomes of Different Policies




Outcomes of Different Policies


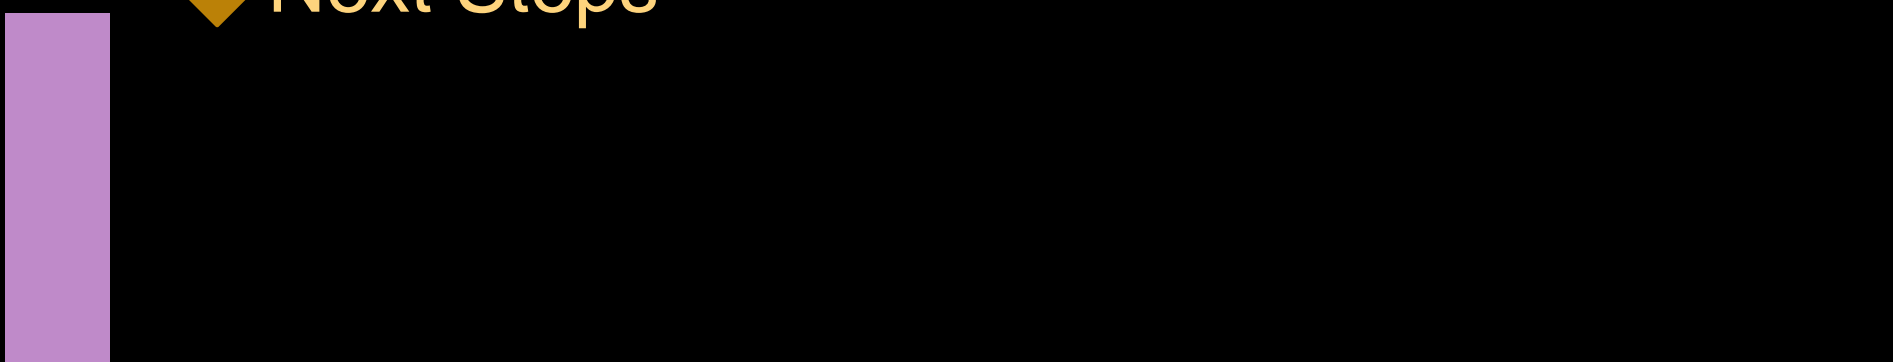


Outcomes of Different Policies



Outcomes of Different Policies



- 
- ◆ Background
 - ◆ My simulator
 - ◆ How to “learn” a Policy
 - ◆ Preliminary Results
 - ◆ Next Steps
- 



Next Steps

- Continue balancing the simulator
- Improve mathematics for the machine learning components
- Find out if the machine learning algorithms can find a balanced policy - suppressing some fires, letting others burn

Thank You!


To my Collaborators

Rachel Houtman
Sean McGregor
Claire Montgomery
Tom Dietterich
Ron Metoyer
Chris Lauer

Our Funding Agencies

NSF
USFS
ARCS Foundation
OSU Provost


Outcomes of Different Policies



Outcomes of Different Policies

Cells burned in 100 Landscapes with a Suppress-All Policy

Average, Min, and Max Cells Burned per year





Example Fire Spread Simulations

