A Scalable Framework for Representation and Reasoning in Large Scale, Spatial-Temporal Planning Problems

PhD Proposal

Mark Crowley
http://www.cs.ubc.ca/∼crowley

Department of Computer Science
University of British Columbia

8th May 2008

A Scalable Framework for Representation and Reasoning in LST Planning Problems

PhD Proposal

Mark Crowley
http://www.cs.ubc.ca/~crowley

Department of Computer Science
University of British Columbia

8th May 2008

Forestry Planning as an LST Problem

Landscape (Forest)

Cell (Stand)

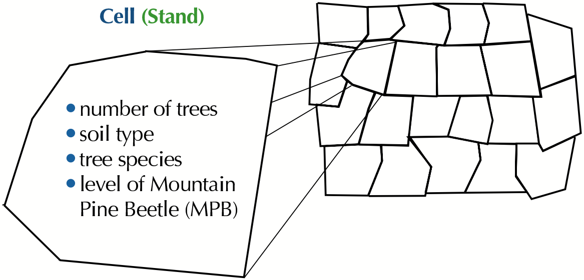

Properties

- number of trees
- soil type
- tree species
- level of Mountain Pine Beetle (MPB)

Utilities

- economic: value of lumber
- social : recreational use
- ecosystem health: habitats

Actions

- clearcut all trees
- cut some trees (thin)
- treat for pests
- salvage dead trees
- do nothing

Constraints

- bounds: annual cut limit
- spatial : migration corridors
  viewsheds

Forestry Planning is Even Harder Now

Mountain Pine Beetle (MPB)

Dendrochtonus ponderosae

Current Approaches in Forestry Planning

Forestry planning is carried out at many different scales:

strategic – tactical – operational

Linear Strata - treat stands as independent, group by attributes. Use linear programming to find optimal policy.

Stochastic Local Search - use tabu search, genetic algorithms and simulated annealing

- these methods can account for some spatial relations
- still do not deal well with uncertainty

The Big Picture

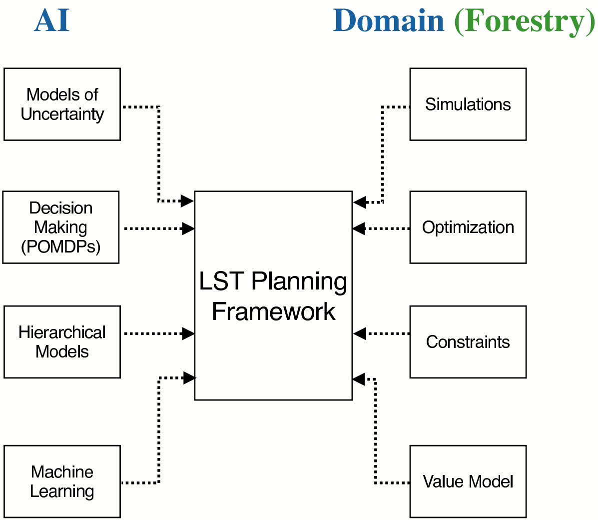

Modeling Uncertainty - Bayesian Network

A Bayesian network (BN) (Pearl, 1988) models the conditional independence relationships between a finite set of random variables. A BN is a directed acyclic graph where each node is independent of its non-descendants given its parents.

a cell

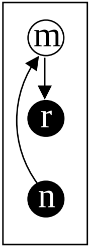

Legend

m - *MPB infestation level*
r - *red tree count*
n - *number of trees*
black nodes = observed, white = unobserved

Modeling Uncertainty and Dynamics - Dynamic BNs

A dynamic Bayesian network (DBN)(Dean and Kanazawa, 1989) models a dynamic process using a two step BN. Arcs are added to model dependency between variables across timesteps.

Legend

m - *MPB infestation level*
r - *red tree count*
n - *number of trees*
a_t^i - *action on cell i at time t*
black nodes = observed, white = unobserved

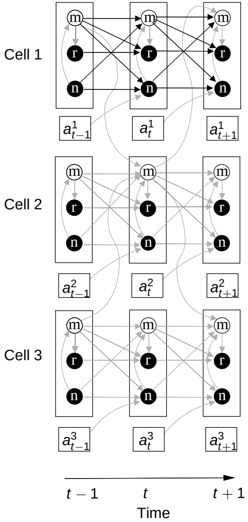

Modeling Uncertainty and Dynamics - Dynamic BNs

A dynamic Bayesian network (DBN)(Dean and Kanazawa, 1989) models a dynamic process using a two step BN. Arcs are added to model dependency between variables across timesteps.

Legend

m - *MPB infestation level*
r - *red tree count*
n - *number of trees*
a_t^i - *action on cell i at time t*
black nodes = observed, white = unobserved

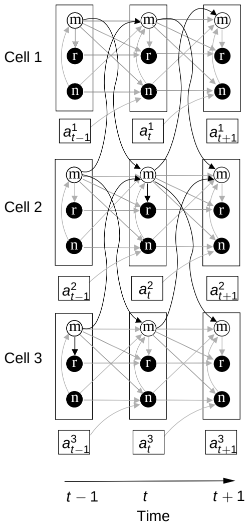

Now spread that DBN across entire landscape

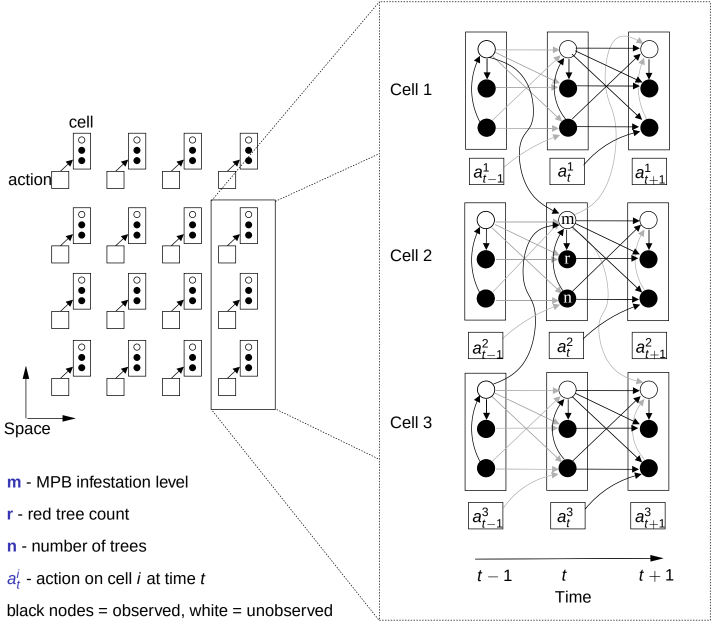

m - MPB infestation level

r - red tree count

n - number of trees

a^i_t - action on cell i at time t

black nodes = observed, white = unobserved

Decision Making

A Markov Decision Process (MDP) is model for representing decision problems. An MDP is a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T})$, where:

- $\mathcal{S}$ a finite set of states
- $\mathcal{A}$ a finite set of possible actions that can be taken at any time-step
- $\mathcal{R}(s, a)$ gives the expected reward for being in state s after taking action a
- $\mathcal{T}(s'|s, a)$ gives the probability of ending up in the next state s' if action a is performed in the current state s

Decision Making

A Partially Observable MDP (POMDP) adds noisy observations of state to the MDP definition:

$\mathcal{Z}$	a finite set of possible observations of the stat
$\mathcal{O}(z\|a, s')$	gives the probability of observing z when action a is taken and results in state s'

Some Example POMDP Solution Methods

Model Based

Some guarantees, very active, getting better all the time.

- Reinforcement Learning - value iteration, policy iteration
- Point Based Belief State Methods - (Perseus[8])
- Factored States - Decision Diagrams or linear value functions often used compress states [7]
- Relational MDPs - learn general policies on relational model, generalize to larger domains [4]

Model-Free

Good if the dynamics come from external simulations

- Reinforcement Learning - Q-learning
- Direct Policy Search - optimize a parameterized policy (PEGASUS[5])
- Hierarchichal RL - combine policies over many timescales (MAXQ[3])

Hierarchical Abstraction of State Space

Perform planning on smaller, abstract states arranged in a hierarchy

Prototypes - use clustering to define a set of prototypical cells that can represent the state. Vary the number of clusters and properties used.

Advantages:

- adjustable state size
- real world planning is fundamentally hierarchical

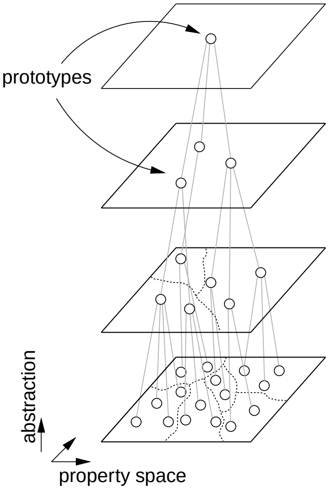

Abstract states $\longrightarrow$ abstract actions/policies

Atomic Actions

- actions initially defined on cells
- actions now need to apply to clusters of cells

Parameterized Actions

- *"cut 5% of the trees"*
- *"cut trees nearest to roads up to an area of x ha"*
- *"cut x ha beginning with stands near roads"*

The Future is not the same as the present

Modeling the state at full detail far into the future has problems:

1. **We can't afford to** - long planning horizon and huge state space
2. **We don't know how** - uncertainty in model, change over time
3. **We don't really want to** - real plans are updated continuously, plan *may* be followed for a one or two years

So, allow abstraction to increase over time

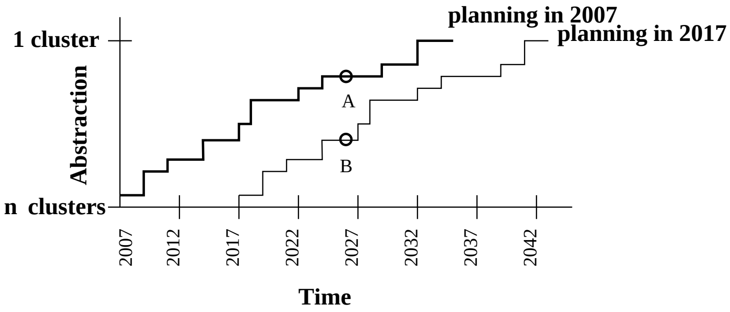

An abstraction schedule will be followed as future timesteps are considered. The schedule could be determined:

- as needed, to balance performance and accuracy
- ahead of time, to indicate level of interest at points in the future

Evaluation is Difficult

Simulation - Use forestry simulations to see results

Comparison - Compare our results to forestry planning solutions

- can we achieve higher utility?
- can we deal with larger problems?
- can we deal with more complex problems?

Qualitative Evaluation - Manual qualitative evaluation by decision making experts in the field

Conclusion

- LST Planning covers a class of decision making problems that are highly relevant to society and challenging to solve
- AI has all the tools we need to make progress on this class of problems
- There is data and simulations to use from domains such as Forestry
- This research will begin building an LST planning framework that :
  - integrates uncertainty
  - integrates spatial relations
  - uses hierarchical abstraction to perform planning efficiently on subproblems

Thanks for Listening

any questions?

Bibliography

Xavier Boyen and Daphne Koller.
Exploiting the architecture of dynamic systems.
In *AAAI '99*, 1999.

T. Dean and K. Kanazawa.
A model for reasoning about persistence and causation.
Computational Intelligence, 5:142–150, 1989.

Thomas G. Dietterich.
An overview of MAXQ hierarchical reinforcement learning.
Lecture Notes in Computer Science, 1864/2000, 2000.

C. Guestrin, D. Koller, C. Gearhart, and N. Kanodia.
Generalizing plans to new environments in relational mdps.
In *The International Joint Conference on Artificial Intelligence (IJCAI-03)*, 2003.

Andrew Y. Ng and Michael Jordan.
PEGASUS: A policy search method for large MDPs and POMDPs.

Variables Become Entangled Over Time

A major challenge using DBNs is that independent nodes become dependent over time.

The BK Algorithm

(Boyen and Koller, 1999)

1. project state to an approximate belief state by breaking links between weakly related clusters of nodes
2. run dynamics forward
3. repeat

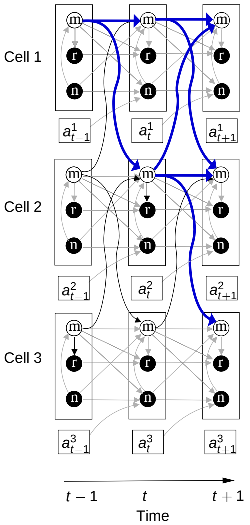

Variables Become Entangled Over Time

A major challenge using DBNs is that independent nodes become dependent over time.

The BK Algorithm

(Boyen and Koller, 1999)

1. project state to an approximate belief state by breaking links between weakly related clusters of nodes
2. run dynamics forward
3. repeat

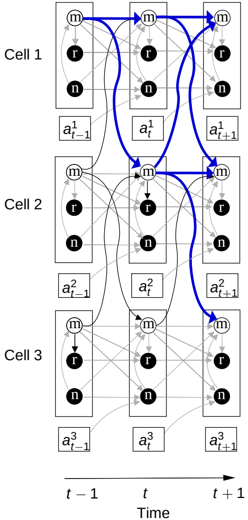

Spatial planning with prototypes

Neighbourhood Sampling

- most spatial relationships have an upper bound on distance (e.g. MPB spread 0-20km)
- use prototypes to create sample neighbourhoods that are *big enough* to model spatial relationships but *much smaller* than the total landscape

A	A	B	A	C
B	C	B	C	C
B	A	A	B	B
A	C	B	C	B
B	A	B	B	A

C	B	C
A	A	B
B	C	B

B	C	C
A	B	B
B	A	B

C	B	C
B	A	B
B	B	A