

Representation and Reasoning in Large Scale, Spatial-Temporal Planning Problems

Mark Crowley

Department of Computer Science
University of British Columbia, Vancouver, BC, Canada
crowley@cs.ubc.ca

My research investigates problems where decisions have to be made across a large landscape in the presence of spatial interrelations, uncertainty and change. I call these problems *large scale spatial-temporal planning problems (LST problems)*. LST problems exist all around us in important areas of society such as mining, forestry, agriculture, urban planning, infectious disease control and many others.

The decisions to be made in LST planning problems are which actions to take on small spatial areas of the landscape called *cells*. Examples of such actions are allowing residential development within a zoning region or harvesting a stand of trees in a forest. We assume that cells define a non-overlapping, covering set over the landscape and that they are atomic for decision making. Cells are the smallest division of the landscape that have actions associated with them. The problem must also include a utility model that defines which outcomes are preferred. The goal of LST planning is to find a good *policy* that describes an action for each cell in the landscape at each point in time.

The goal of my research is to discover efficient and general methods for planning in LST problems. My thesis is that we can improve upon current approaches to enable scalable, high quality decision making on LST problems by taking account of uncertainty and spatial relationships in the data.

The specific domain that will be used for most comparisons and for data will be forestry planning. The decision problem in forestry is to make decisions each year for every one of thousands of small areas of trees in the forest. The actions that can be taken on each of these cells include harvesting, treating the trees for pests or doing nothing at all. Plans generally need to forecast decades into the future to satisfy constraints imposed by resources and government regulations.

Forestry is an enormously important industry in Canada. Recently, forestry planning has become harder due to increasingly complex constraints on allowable policies based on changing ecological and recreational values for forests in society. The current unprecedented outbreak of Mountain Pine Beetle (MPB) in BC forests is another complicating factor. The spread of MPB across the forest every year lead to greater uncertainty in forestry models and greatly increases the spatial interdependence between cells. MPB can travel up to 20 km in one season so the conditions from the surrounding neighbourhood of a cell are now very relevant to decision making. Every action can influence the spread of MPB and thus it is important to use models that are as realistic as possible when computing policies. However, the increase in spatial interdependence between cells invalidates assumptions of some popular decision techniques in forestry planning that assume independent cells. There are other planning techniques in common use that can handle some spatial interdependence between cells but these do not deal well with uncertainty and there is a need for better tools[1].

Our approach will be to try to use the very realistic simulations available in forestry as the model and to perform planning without oversimplifying the problem by removing uncertainty and spatial interdependence. These are the defining features of LST planning problems. My research attempts to expand the size of problem that can be dealt with using existing decision theoretic techniques by taking advantage of the structure of LST problems. The decision theoretic model I intend to use are *Partially Observable Markov Decision Processes (POMDPs)*.

POMDPs provide a very general model for decision making and are an active area of research in AI that has seen dramatic improvements in solution techniques in recent years. Currently, I am carrying out tests using the state of art in POMDP point based solution methods such as Perseus[2] on moderate sized LST problems. A major limitation of these approaches is their scaling for problems with large state and action spaces. We want to take advantage of the structure of LST problems to improve this performance. One way we might do this is to use clustering to group similar cells together based on their properties and surroundings. A hierarchical abstraction model can be created that will adjust the number of cells and the number of clusters to be used based on the needs of performance and accuracy. We will then try to learn policies acting on these cell clusters which are far fewer in number than the number of actual cells. The intention would be to use learned policies from other levels of abstraction to inform learning the current policy and thus find a policy that is robust over different scales while always keeping the size of the problem down. These hierarchical policies can then be applied back to the concrete cells based on their specific context.

Another important goal of our LST framework is to use simulations from specific problem domains for the transition dynamics during planning. Non-model based planning algorithms such as Q-learning and Direct policy search (DPS) [3] provide promising alternatives to achieve this. DPS could be used to learn parameterized policies on the cell clusters by running these smaller states through external simulations and iteratively improving the policy. We intend to experiment with combinations of all these different techniques as each has advantages for different aspects of LST problems.

LST problems are widespread and are of huge importance to society but LST planning is a very challenging area of general decision making because of the enormous scale of most problems. However, they also provide a great opportunity by exhibiting a regular structure that we can take advantage of to solve larger decision problems than has been possible previously. Unfortunately, problems with an LST structure have not been directly addressed by current AI research. Meanwhile, researchers in numerous specific domains continue to come up with their own solutions based on technology familiar to them. I hope my research will raise awareness about this class of problems and encourage a dialogue that moves towards a general treatment of all LST problems.

References

1. J.P.Kimmins, et al.: Possible Forest Futures. Working Paper 2005-11. Pacific Forestry Centre, Canadian Forest Service, National Resources Canada. Victoria, Canada. (2005)
2. Spaan, M.T.J., Vlassis, N.: Perseus: Randomized point-based value iteration for POMDPs. *Journal of Artificial Intelligence Research* **24** (2005) 195–220
3. Ng, A.Y., Jordan, M.: PEGASUS: A policy search method for large MDPs and POMDPs. In: *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*. (2000)