

Ronald Vance Bjarnason Jr.

2048 KEC, Oregon State University • Corvallis, OR 97331-5501 • 209-648-8979 • ronny@eeecs.oregonstate.edu

My research strengths are analysis and optimization of real world problems.

I have a strong background in algorithms and have twice taught the undergrad algorithms class, emphasizing application to real world problems. My PhD research has involved problems defined over large data sets with computationally intensive exact solutions. Implementation of accurate and computationally efficient approximation algorithms has been key to our success.

Earlier years of my PhD program were funded by a relational reinforcement learning grant with Prasad Tadepalli. Our chosen domain, Klondike Solitaire, is rich with complex relational dependencies. In time our work transitioned to the fully observable version of the game, Thoughtful Solitaire. At the time, the most successful algorithm was a nested adaptation of rollouts using prioritized action rules, reported at NIPS'04 – winning 70% of Thoughtful games while reporting just over 1% of games as dead ends. Our research evaluated the relationship between features of the root nodes and features of the leaf nodes following this complex rollout search.

In addition to our stated research objectives, I optimized the rollout algorithm. A feature-based search heuristic was implemented to evaluate nodes directly. The search tree was compressed and pruned by eliminating redundant actions and solving a relaxed search problem. The resulting heuristic search algorithm solved over 82% of games in less than one third the time of the previous algorithm. Additionally, the pruning method identified over 8.5% of games as dead ends, reducing uncounted games from 29% to less than 9.5%. We described our algorithm in a 2007 ICGA Journal article that reported these results. In the past year, my official research has turned to a different domain, but I am preparing an upcoming conference paper that will extend this work to standard Klondike.

For the past year, we have been working closely with our local fire department on improving response times of fire and emergency response units. The policy space of this domain is severely restricted due to requirements of real-life implementation. These restrictions prevent us from implementing popular optimization techniques. We have already achieved significant results with a simple two-phase hill climbing algorithm that alternates optimizing vehicle/personnel station assignments and incident response policies. Our simulator approximates emergency response for 24,740 actual 9-1-1 requests received by local dispatchers from Oct 2003 to Sept 2008. We evaluated our learned policy using published standardized metrics. Our results improved over a variety of baseline policies and actual performance of our local department. We also simulated projected expansion of our city population and the subsequent optimized response policies. Our findings were quite surprising to our collaborators in the fire department, as our results recommended reapportioning resources and closing a fully-staffed station. These results have been detailed in a recently submitted conference paper. Our work in this domain continues with other machine learning solutions. We are currently adapting semi-supervised and reinforcement learning techniques in both the restricted and unrestricted policy spaces, as we extend our work to a general class of multi-scale optimization domains.

As my research career moves forward, I will continue to develop novel and efficient solutions to interesting problems. My background with reinforcement learning, heuristic search and approximation techniques will allow me to objectively analyze a broad spectrum of difficult problems. My experience with algorithm analysis will enable me to efficiently implement these solutions.