

- SARSOP (Geoff Hollinger)
 - Review of POMDPs
 - one solution approach, sampling belief space
 - Key Idea
 - method for sampling belief space to solve POMDP
 - don't need to solve whole POMDP, only from initial belief state
 - take belief samples from the optimally reachable belief space
 - Covering Number
 - running time is linear in covering number
 - hard part is generating the 'cover'
 - SARSOP Algorithm
 - significant contribution of paper in 'sample' and 'prune' functions
 - upper bound using Sawtooth
 - Backup operator
 - Sampling near R^*
 - on board: chart of actions and upper/lower bounded reward region
 - choose action with highest upper bound, can never get rid of it
 - Gap termination
 - Sampling in SARSOP (Algorithm 3)
 - Pruning
 - APPL software implements SARSOP
 - good for comparison to other algorithms as they are provably within factor of optimal
- SARSOP Applications
 - 10x10x10 grid world
 - states is BIG number (2^{1000})
 - belief state exponential in states
 - underwater robot, only localized in certain areas
 - grabbing the tetris block
 - integrated exploration
 - Performance comparison
 - faster (sometimes much) than HSVI2
 - comparable (but worse) performance
 - Applications are all toy, mostly not useful on real problems
 - discussion: simple useful problems?
 - video of robot chasing dilbert
- Anytime RRTs
 - regular Rapidly-exploring Random Trees
 - Earlier Improvements
 - goal bias - similar to best first search

- nearest neighbor could become k-nearest neighbor
 - An idea from ARA*
 - start with a suboptimal solution with epsilon bound, then make improvements
 - Anytime RRT Algorithm
 - Repeat RRT search, but improve solution bound every time with previous best, don't expand solution that pass our bound
 - discussion:
 - as the algorithm gets close to optimal it may have difficulty finishing a solution and could become stuck. Does it need random restarts?
 - can the RRT algorithm be more parallel by growing more trees?
 - taken to a large number becomes a PRM
 - Anytime RRT Planning
 - Bias the point sampling to choose points that may lead to an improved solution
 - extending the tree: choose cheapest of possible extensions
 - Performance vs RRT
 - random performance of RRT does not approach Anytime RRT over time
 - One and Three robot path planning with cost obstacles
- The Belief Roadmap
 - Motivation
 - want to minimize uncertainty without doing a POMDP
 - Kalman filter review
 - Factoring the Covariance
 - can 'star' the operations from start to T to get the final transform matrix
 - The Algorithm
 - first build the belief roadmap
 - 'Experimental' Validation
 - beacons are contrived to be off of the shortest path
 - this alg follows the beacons
 - navigation across MIT campus
 - walls not taken into account
 - Related discussion
 - if beacon positions are unknown at beginning becomes Active SLAM, open problem
 - Paper on trace or determinant of covariance or information matrix