FLAP for CAOS: Forward-Looking Active Perception for Clutter-Aware Object Search

Thorsten Gedicke, Martin Günther, Joachim Hertzberg

{tgedicke, martin.guenther, joachim.hertzberg}@uos.de

June 30, 2016

UNIVERSITÄT OSNABRÜCK

A mobile search and exploration system

"Explore your surroundings until you find the target,
and be quick about it!"

Application scenarios:

- Robot butler ("Bring me my favorite mug!")
- Search and rescue in disaster areas
- ...

Setup

- Willow Garage PR2
  - RGB-D sensor mounted on a pan/tilt-head
  - Telescoping torso
- Search regions on tabletops and inside shelves

Features

Our system...

- selects promising observation targets
- finds view poses that resolve occlusions
- is time-efficient (minimizes expected search time)
- uses available prior knowledge
- adapts to novel information gained during search using *continual planning*

State of the art

- Universal use of octrees for environment representation
- Exploitation of prior knowledge
  - Known map
  - Known distribution of target location probability
- Detection of promising sensing poses using ray casting
- Search strategies:
  - Greedy (always selects the currently "best" sensing pose)
  - Planning (using decomposition and/or pruning heuristics)

Observation

So far, all planning approaches become inefficient compared to a greedy strategy when replanning is involved. Let's change that!

Approach

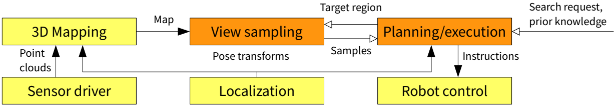

- Maintain an octree-based map (free/occupied/unknown) with defined search regions and target location probabilities
- Until the target is found:
  1. Find useful sensing poses using sampling and ray casting
  2. Predict execution times using a path planner
  3. Plan several observations ahead
  4. Navigate to the next view, observe, then back to (1)

Time-efficient continual planning

Goal

Construct a sequence of views that minimizes the expected time until the target is found (fast approximate solution).

- Depth-first branch-and-bound algorithm
- Parameterized limited horizon and pruning heuristics
  - Only branch up to a limited horizon of ψ steps.
  - Do not branch if a view's utility is worse than the greedy choice by a factor of ϕ.
  - Do not branch if an alternative exists that is both faster to reach and more likely to reveal the target (strict domination criterion ξ).
- Extreme cases: greedy vs. optimal

$$util(s_i|L) = \frac{p(s_i|L)}{t(L_n, s_i)}$$

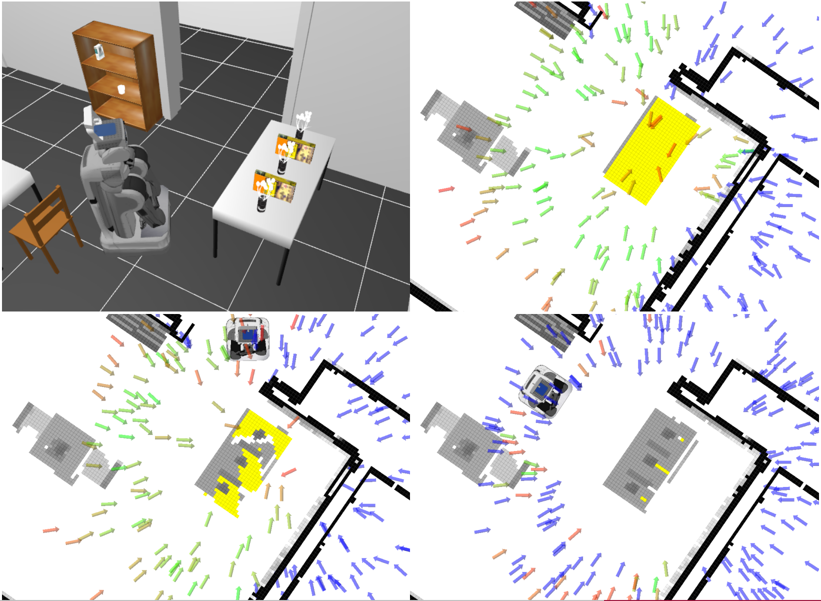

Evaluation of average search time

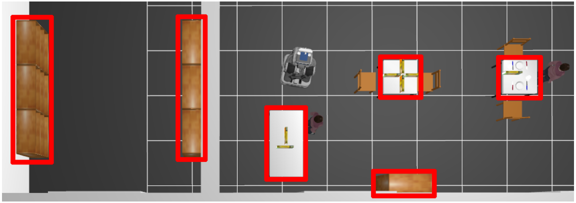

- Task: Find one object of target class within the search regions
- 20 test runs with and without a-priori knowledge of the map
- Random starting pose for each run
- Offline planning (with a known map) vs. continual planning (with and without a map)
- Simulation disregards issues of localization, navigation, and object recognition.

Result

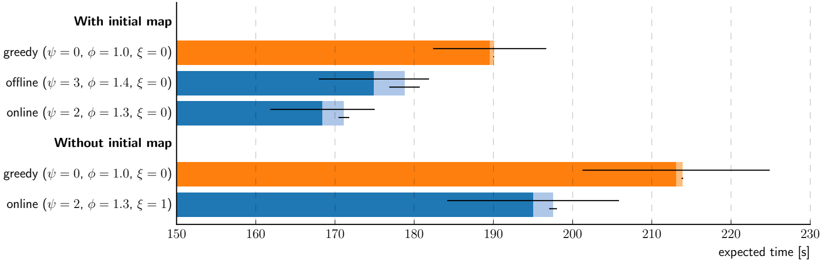

Continual planning achieves shorter average execution times

- 10 % if the map is known a-priori,
- 7 % else.

Conclusion

We presented a system for autonomous object search that uses continual planning to resolve unexpected occlusions on the go.

- Continual planning does pay off if the average planning time is kept short (below 5 s in our trials) *even if the environment is known from the start*.
- Higher payoff is possible by investing more planning time when cost functions are not purely time-based (e.g. energy).
- Open issues:
  - Integration with a probabilistic object location model
  - Modeling recognizability and spatial properties of objects
  - Examining the influence of different environments on optimal planner parameters

Visit us: `http://kos.informatik.uos.de/flap4caos/`