Measurement Simplification in ρ-POMDP with Performance Guarantees

Presented by: **Tom Yotam**
Supervisor: **Assoc. Prof. Vadim Indelman**

July 25th 2023

ANPL
Autonomous Navigation and Perception Lab

Outline

- Background
- Approach
  - Reward Bounds
  - Observation Space Partitioning
  - Computational Complexity
  - Gaussian Beliefs
- Results
  - Bounds Analysis
  - Simulation
- Recap

ANPL
Autonomous Navigation and Perception Lab

Motivation

Decision Making Under Uncertainty

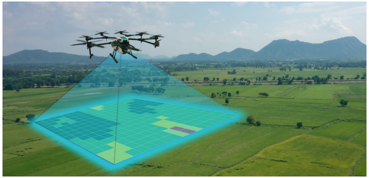

(a) Informative Planning

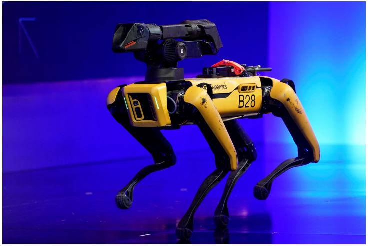

(b) Autonomous Agents

(c) Reinforcement Learning

Partially Observable Markov Decision Process

A POMDP formally: $(\mathcal{X}, \mathcal{A}, \mathcal{Z}, T, O, R)$

- state, action and observation spaces
- transition and observation models
- reward function

ANPL Autonomous Navigation and Perception Lab

Partially Observable Markov Decision Process

- Markovian transition model, i.e. $T(X, a, X') = \mathbb{P}(X'|X, a)$
- Each measurement is conditionally independent given the state, i.e. $O(X, z) = \mathbb{P}(z|X)$
- The reward is a function of the state

ANPL Autonomous Navigation and Perception Lab

Partially Observable Markov Decision Process

The true state is unknown

- The agent only observes the environment through noisy measurements
- It must maintain a probability distribution over the true state
- $b_k \triangleq b[X_k] = \mathbb{P}(X_k | z_{0:k}, a_{0:k-1}) \triangleq \mathbb{P}(X_k | h_k)$

ANPL
Autonomous Navigation and Perception Lab

POMDP - computational complexity

- Curse of dimensionality
- Curse of history

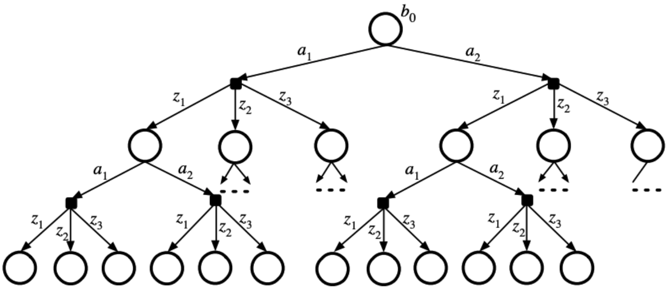

Figure: DESPOT Ye et al 2007

Plan-act-sense-infer

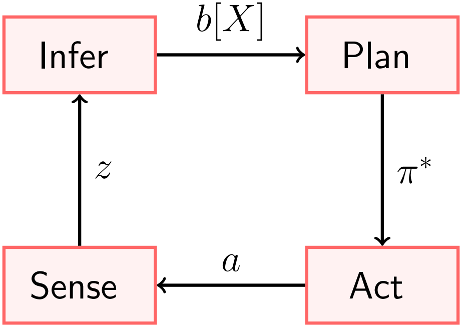

Plan-act-sense-infer

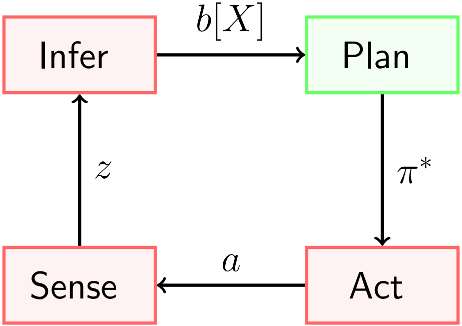

Plan

What is a policy?

- Maps belief states to actions, $\pi : \mathcal{B} \mapsto A$
- For some finite planning horizon ℓ, the *value* of a policy π:

$$V^{\pi}(b_k) = R(b_k, \pi_k(b_k)) + \mathop{\mathbb{E}}_{z_{k+1:k+\ell}} \left[\sum_{i=k+1}^{k+\ell} R(b_i, \pi_i(b_i)) \right]$$

ANPL Autonomous Navigation and Perception Lab

Plan

- Solving a POMDP is equivalent to finding the optimal policy π^* such that the value function is maximized.
- Can replace the optimal policy with the optimal action sequence (open loop)

ANPL Autonomous Navigation and Perception Lab

ρ-POMDP

- Reasoning about uncertainty is key for planning, AI, Machine Learning
- Quantifying uncertainty allows us to identify actions that reduce it

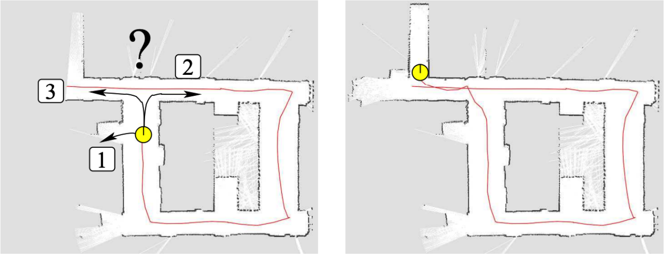

Stachniss et al. RSS'05

ρ-POMDP

- Extends the POMDP model to include belief dependent rewards
- $R(b, \pi(b)) \triangleq -\mathcal{H}(X) \equiv \mathbb{E}_{X\sim b}(\log b[X])$
- If both X, Z, are treated as random variables, the expected reward becomes the conditional entropy of these random variables

ANPL
Autonomous Navigation and Perception Lab

ρ-POMDP

- $\mathbb{E}_{Z}[R(b)] = -\mathcal{H}(X \mid Z) = -\mathbb{E}_{Z}[\mathcal{H}(X \mid Z = z)]$
- The expected reward at each ith look ahead step: $\mathbb{E}_{Z_{k+1:i}}[R(b_i, a_{i-1})] = -\mathcal{H}(X_i | Z_{k+1:i})$
- Future observations are drawn from the distribution $\mathbb{P}(Z_{k+1:i} \mid b_k, \pi)$ and $i \in [k+1, k+\ell]$

ANPL Autonomous Navigation and Perception Lab

Related Work

- POMDP online solvers
  - Sunberg and Kochenderfer - ICAPS'18
  - Ye et al. - JAIR'17
- Simplification in inference
  - Khosoussi et al. - WAFR'20
  - Zhang and Vela - CCVP'15
  - Carlevaris-Bianco, Kaess and Eustice - TRO'14
- Simplification in planning
  - Zhitnikov and Indelman - AI'22
  - Elimelech and Indelman - IJRR'22

ANPL Autonomous Navigation and Perception Lab

Contributions

- Novel observation space partitioning
- Analytical bounds on the expected reward, as function of partitioned space, that hold for all families of belief distributions.
- Partition tree that allows greater efficiency as we go down its hierarchy.
- Bounds that are adaptive and converge to the original solution.
- Hierarchy of efficient implementations for Gaussian beliefs

Outline

ANPL
Autonomous Navigation and Perception Lab

Approach - simplification

- To choose the optimal action from a pool of candidate actions - need to evaluate the reward function for each action.
- Instead, one can evaluate bounds on the expected reward function as a proxy

$$\mathcal{LB}_i \leq \mathop{\mathbb{E}}_{Z_{k:i}} \left(R(b_i)\right) \leq \mathcal{UB}_i$$

$$\sum_{i=k+1}^{k+\ell} \mathcal{LB}_i \leq J\left(b_k, a_{k:k+\ell-1}\right) \leq \sum_{i=k+1}^{k+\ell} \mathcal{UB}_i$$

ANPL Autonomous Navigation and Perception Lab

Approach - reward bounds

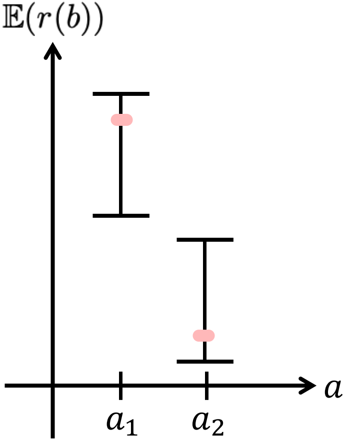

Approach - reward bounds overlap

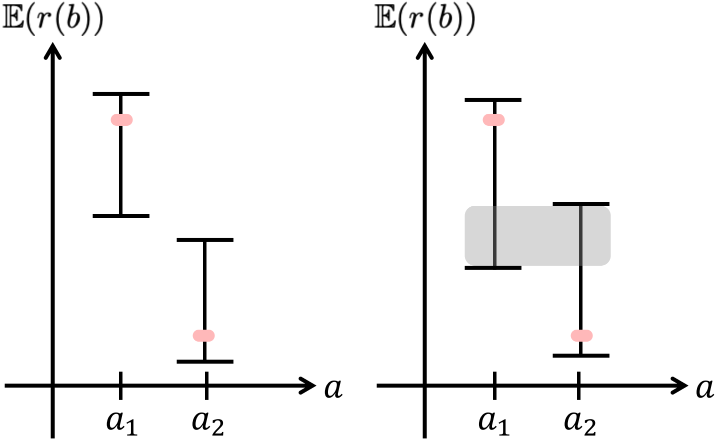

Multivariate Observation Space

- Consider a multivariate random variable $Z \in \mathcal{Z}$, that represents future observations:

$$Z = (Z^1, Z^2, \ldots, Z^m)$$

- Z^i is a random variable defined by a given sensing modality, and m is the number of such random variables

Multivariate Observation Space

For example, raw measurement of an image sensor

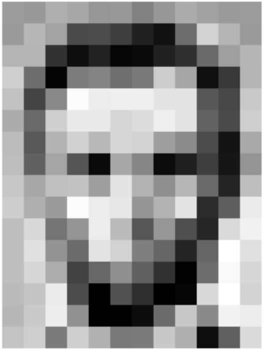

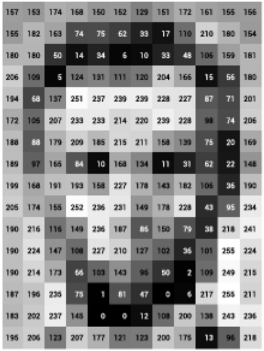

157	153	174	168	150	152	129	151	172	161	155	156
155	182	163	74	75	62	33	17	110	210	180	154
180	180	50	14	34	6	10	33	48	106	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	106	36	190
205	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	86	150	79	38	218	241
190	224	147	108	227	210	127	102	36	101	255	224
190	214	173	66	103	143	96	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	200	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

ANPL Autonomous Navigation and Perception Lab

Multivariate Observation Space

Or a factor graph

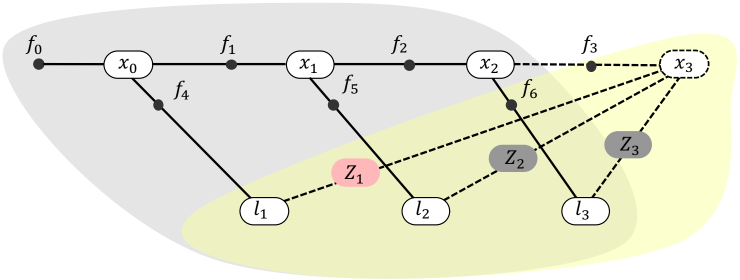

Partitioning of Multivariate Observation Space

- Consider the partitioning $Z^s \in \mathcal{Z}^s$ and $Z^{\bar{s}} \in \mathcal{Z}^{\bar{s}}$, such that:

$$Z^s = \{Z^1, Z^2, \ldots, Z^n\}$$

$$Z^{\bar{s}} = \{Z^{n+1}, Z^{n+2}, \ldots, Z^m\}$$

- $\mathcal{Z} = \mathcal{Z}^s \oplus \mathcal{Z}^{\bar{s}}$ (addition of subspaces)

ANPL
Autonomous Navigation and Perception Lab

Partitioning of Multivariate Observation Space

But why is this a good idea?

- Apply partitioning to a raw image measurement of size 20×20 binary pixels.
- Each pixel is represented by a random variable $Z^{x,y} \in \{0, 1\}$, and $Z \in \mathcal{Z} \subseteq (\mathbb{F}_2)^{400}$.

ANPL Autonomous Navigation and Perception Lab

Partitioning of Multivariate Observation Space

But why is this a good idea?

- Consider all of the different permutations for each pixel, 2^{400} in total, which defines $|\mathcal{Z}|$.
- If we partition $Z^s \triangleq \{Z^{x,y} \mid y \leq 10\}$ and $Z^{\bar{s}} \triangleq \{Z^{x,y} \mid y > 10\}$, we need only to consider 2^{200} permutations for each random variable.
- 2^{201} vs 2^{400}

ANPL
Autonomous Navigation and Perception Lab

The Big Picture

- Planning involves thinking about future observations (and actions), and evaluating a reward function
- This process is computationally expensive
- Partitioning the observation space makes this less expensive

Bounds on Expected Reward

$$\mathcal{LB} \leq \mathcal{H}(X|Z) \leq \mathcal{UB}$$

$$\mathcal{LB} \triangleq \mathcal{H}(Z^s \mid X) + \mathcal{H}(Z^{\bar{s}}|X) - \mathcal{H}(Z^s) - \mathcal{H}(Z^{\bar{s}}) + \mathcal{H}(X)$$

$$\mathcal{UB} \triangleq \mathcal{H}\left(Z^s|X\right) + \mathcal{H}\left(X\right) - \mathcal{H}\left(Z^s\right)$$

ANPL Autonomous Navigation and Perception Lab

Bounds on Expected Reward

Lemma 1

The conditional Entropy can be factorized as

$$\mathcal{H}\left(X|Z\right) = \mathcal{H}\left(Z|X\right) + \mathcal{H}\left(X\right) - \mathcal{H}\left(Z\right)$$

Bounds on Expected Reward

Theorem 1

The conditional Entropy can be bounded from above by

$$\mathcal{H}(X|Z) \leq \mathcal{UB} \triangleq \mathcal{H}\left(Z^s|X\right) + \mathcal{H}\left(X\right) - \mathcal{H}\left(Z^s\right)$$

Proof.

□

Bounds on Expected Reward

Theorem 1

The conditional Entropy can be bounded from above by

$$\mathcal{H}(X|Z) \leq \mathcal{UB} \triangleq \mathcal{H}\left(Z^s|X\right) + \mathcal{H}\left(X\right) - \mathcal{H}\left(Z^s\right)$$

Proof.

- $\mathcal{H}\left(X|Z^s\right) - \mathcal{H}\left(X|Z\right) = \mathcal{I}(X|Z^s; Z \setminus Z^s)$

□

Bounds on Expected Reward

Theorem 1

The conditional Entropy can be bounded from above by

$$\mathcal{H}(X|Z) \leq \mathcal{UB} \triangleq \mathcal{H}\left(Z^s|X\right) + \mathcal{H}\left(X\right) - \mathcal{H}\left(Z^s\right)$$

Proof.

- $\mathcal{H}\left(X|Z^s\right) - \mathcal{H}\left(X|Z\right) = \mathcal{I}(X|Z^s; Z \setminus Z^s)$
- $0 \leq \mathcal{I}(X|Z^s; Z \setminus Z^s) \rightarrow \mathcal{H}\left(X|Z\right) \leq \mathcal{H}\left(X|Z^s\right)$

$\square$

ANPL Autonomous Navigation and Perception Lab

Bounds on Expected Reward

Theorem 1

The conditional Entropy can be bounded from above by

$$\mathcal{H}(X|Z) \leq \mathcal{UB} \triangleq \mathcal{H}\left(Z^s|X\right) + \mathcal{H}\left(X\right) - \mathcal{H}\left(Z^s\right)$$

Proof.

- $\mathcal{H}\left(X|Z^s\right) - \mathcal{H}\left(X|Z\right) = \mathcal{I}(X|Z^s; Z \setminus Z^s)$
- $0 \leq \mathcal{I}(X|Z^s; Z \setminus Z^s) \rightarrow \mathcal{H}\left(X|Z\right) \leq \mathcal{H}\left(X|Z^s\right)$
- $\mathcal{H}\left(X|Z^s\right) = \mathcal{H}\left(Z^s|X\right) + \mathcal{H}\left(X\right) - \mathcal{H}\left(Z^s\right)$

□

ANPL Autonomous Navigation and Perception Lab

Bounds on Expected Reward

Lemma 2

Given two sets of expected measurements $(Z^s, Z^{\bar{s}})$*, the conditional Entropy can be factorized as*

$$\mathcal{H}(X|Z) = \mathcal{H}(Z^s|X) + \mathcal{H}(Z^{\bar{s}}|X) - \mathcal{H}(Z^s, Z^{\bar{s}}) + \mathcal{H}(X)$$

ANPL Autonomous Navigation and Perception Lab

Bounds on Expected Reward

$$\mathcal{H}(Z^s, Z^{\bar{s}}) = \mathcal{H}(Z^s) + \mathcal{H}(Z^{\bar{s}}) - \mathcal{I}(Z^s; Z^{\bar{s}})$$

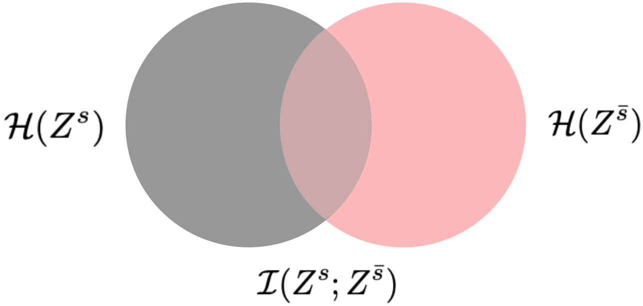

ANPL Autonomous Navigation and Perception Lab

Bounds on Expected Reward

Theorem 2

The conditional Entropy can be bounded from bellow by:

$$\mathcal{LB} \triangleq \mathcal{H}(Z^s \mid X) + \mathcal{H}(Z^{\bar{s}}|X) - \mathcal{H}(Z^s) - \mathcal{H}(Z^{\bar{s}}) + \mathcal{H}(X)$$

Proof.

□

Bounds on Expected Reward

Theorem 2

The conditional Entropy can be bounded from bellow by:

$$\mathcal{LB} \triangleq \mathcal{H}(Z^s \mid X) + \mathcal{H}(Z^{\bar{s}}|X) - \mathcal{H}(Z^s) - \mathcal{H}(Z^{\bar{s}}) + \mathcal{H}(X)$$

Proof.

- $\mathcal{H}(Z^s, Z^{\bar{s}}) = \mathcal{H}(Z^s) + \mathcal{H}(Z^{\bar{s}}) - \mathcal{I}(Z^s; Z^{\bar{s}})$

□

ANPL
Autonomous Navigation and Perception Lab

Bounds on Expected Reward

Theorem 2

The conditional Entropy can be bounded from bellow by:

$$\mathcal{LB} \triangleq \mathcal{H}(Z^s \mid X) + \mathcal{H}(Z^{\bar{s}}|X) - \mathcal{H}(Z^s) - \mathcal{H}(Z^{\bar{s}}) + \mathcal{H}(X)$$

Proof.

- $\mathcal{H}(Z^s, Z^{\bar{s}}) = \mathcal{H}(Z^s) + \mathcal{H}(Z^{\bar{s}}) - \mathcal{I}(Z^s; Z^{\bar{s}})$
- $\mathcal{I}(Z^s; Z^{\bar{s}}) \geq 0 \rightarrow \mathcal{H}(Z^s, Z^{\bar{s}}) \leq \mathcal{H}(Z^s) + \mathcal{H}(Z^{\bar{s}})$

□

Bounds on Expected Reward

Theorem 2

The conditional Entropy can be bounded from bellow by:

$$\mathcal{LB} \triangleq \mathcal{H}(Z^s \mid X) + \mathcal{H}(Z^{\bar{s}}|X) - \mathcal{H}(Z^s) - \mathcal{H}(Z^{\bar{s}}) + \mathcal{H}(X)$$

Proof.

- $\mathcal{H}(Z^s, Z^{\bar{s}}) = \mathcal{H}(Z^s) + \mathcal{H}(Z^{\bar{s}}) - \mathcal{I}(Z^s; Z^{\bar{s}})$
- $\mathcal{I}(Z^s; Z^{\bar{s}}) \geq 0 \rightarrow \mathcal{H}(Z^s, Z^{\bar{s}}) \leq \mathcal{H}(Z^s) + \mathcal{H}(Z^{\bar{s}})$
- Plug-in to lemma 3

□

ANPL Autonomous Navigation and Perception Lab

Bounds on Expected Reward

$$\mathcal{LB} \leq \mathcal{H}(X|Z) \leq \mathcal{UB}$$

$$\mathcal{LB} \triangleq \mathcal{H}(Z^s \mid X) + \mathcal{H}(Z^{\bar{s}}|X) - \mathcal{H}(Z^s) - \mathcal{H}(Z^{\bar{s}}) + \mathcal{H}(X)$$

$$\mathcal{UB} \triangleq \mathcal{H}\left(Z^s|X\right) + \mathcal{H}\left(X\right) - \mathcal{H}\left(Z^s\right)$$

ANPL Autonomous Navigation and Perception Lab

Hierarchical Partitioning

Can we simplify further?

- Unique encoding denoted as $Z^{n_i|m_j}$, n is the node number at the ith partitioning level, m is the node number at the parent partitioning level j

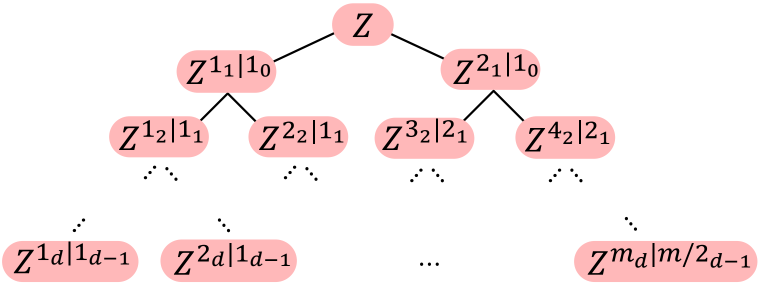

Bounds with Hierarchical Partitioning

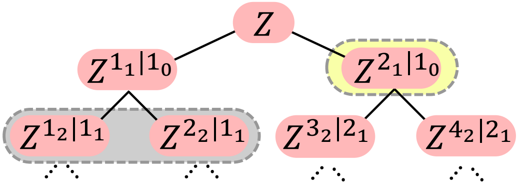

Computational Complexity

Lemma 2

Given two sets of expected measurements $(Z^s, Z^{\bar{s}})$, *the conditional Entropy can be factorized as*

$$\mathcal{H}(X|Z) = \mathcal{H}(Z^s|X) + \mathcal{H}(Z^{\bar{s}}|X) - \mathcal{H}(Z^s, Z^{\bar{s}}) + \mathcal{H}(X)$$

ANPL
Autonomous Navigation and Perception Lab

Computational Complexity

- Baseline Expected Reward

$$\mathcal{H}(X|Z)$$

Using lemma 2:

$$\mathcal{H}(Z^s|X), \mathcal{H}(Z^{\bar{s}}|X), \mathcal{H}(X), \mathcal{H}(Z^s, Z^{\bar{s}})$$

ANPL Autonomous Navigation and Perception Lab

Computational Complexity

- Baseline Expected Reward

$$\mathcal{H}(X|Z)$$

Using lemma 2:

$$\mathcal{H}(Z^s|X), \mathcal{H}(Z^{\bar{s}}|X), \mathcal{H}(X), \mathcal{H}(Z^s, Z^{\bar{s}})$$

- Expected Reward Bounds

$$\mathcal{H}(Z^s|X), \mathcal{H}(Z^{\bar{s}}|X), \mathcal{H}(Z^{\bar{s}}), \mathcal{H}(Z^s), \mathcal{H}(X)$$

ANPL
Autonomous Navigation and Perception Lab

Computational Complexity

- Baseline Expected Reward

$$\mathcal{H}(X|Z)$$

Using lemma 2:

$$\mathcal{H}(Z^s|X), \mathcal{H}(Z^{\bar{s}}|X), \mathcal{H}(X), \boldsymbol{\mathcal{H}(Z^s, Z^{\bar{s}})}$$

- Expected Reward Bounds

$$\mathcal{H}(Z^s|X), \mathcal{H}(Z^{\bar{s}}|X), \boldsymbol{\mathcal{H}(Z^{\bar{s}})}, \boldsymbol{\mathcal{H}(Z^s)}, \mathcal{H}(X)$$

ANPL Autonomous Navigation and Perception Lab

Computational Complexity

$$\boldsymbol{\mathcal{H}(Z^s, Z^{\bar{s}})} \text{ vs. } \boldsymbol{\mathcal{H}(Z^{\bar{s}}), \mathcal{H}(Z^s)}$$

- $\mathcal{H}(Z^s, Z^{\bar{s}}) = -\int\limits_{Z^s}\int\limits_{Z^{\bar{s}}} \mathbb{P}(Z^s, Z^{\bar{s}}) \log \mathbb{P}(Z^s, Z^{\bar{s}}) dZ^s dZ^{\bar{s}}$
- $\mathcal{H}(Z^s) = -\int\limits_{Z^s} \mathbb{P}(Z^s) \log \mathbb{P}(Z^s) dZ^s$
- $\mathcal{H}(Z^{\bar{s}}) = -\int\limits_{Z^{\bar{s}}} \mathbb{P}(Z^{\bar{s}}) \log \mathbb{P}(Z^{\bar{s}}) dZ^{\bar{s}}$

ANPL
Autonomous Navigation and Perception Lab

Computational Complexity

$$\boldsymbol{\mathcal{H}}(\boldsymbol{Z}^s, \boldsymbol{Z}^{\bar{s}}) \text{ vs. } \boldsymbol{\mathcal{H}}(\boldsymbol{Z}^{\bar{s}}), \boldsymbol{\mathcal{H}}(\boldsymbol{Z}^s)$$

- $O(|\mathcal{Z}^s||\mathcal{Z}^{\bar{s}}|)$
- $O(|\mathcal{Z}^s| + |\mathcal{Z}^{\bar{s}}|)$
- Same logic applies to the hierarchical partitions

ANPL
Autonomous Navigation and Perception Lab

Gaussian Belief - preliminaries

$b[X] \sim \mathcal{N}(\mu, \Sigma)$, with mean $\mu \in \mathbb{R}^N$ and covariance matrix $\Sigma \in \mathbb{R}^{N \times N}$

- $\Sigma_{k+1}^{-1} = \Lambda_{k+1}$
- $\Lambda_{k+i} \triangleq \Lambda_k^{\text{Aug}} + A_{k+1:k+i}^T \cdot A_{k+1:k+i}$

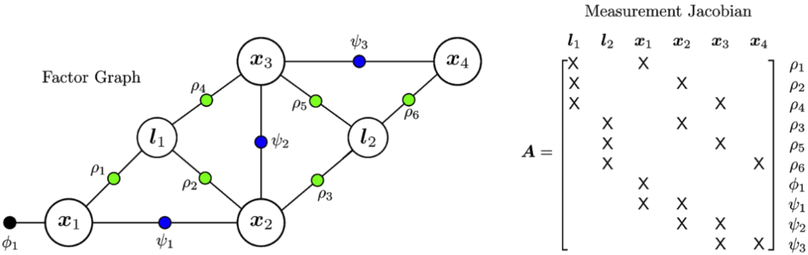

Figure: Caesar.jl'21

Gaussian Belief - preliminaries

$$b[X] \sim \mathcal{N}(\mu, \Sigma)$$

- $\mathcal{H}(X_{k+1}) = \frac{1}{2}(N \ln(2\pi e) - \ln|\Lambda_{k+1}|)$
- $\underset{Z_{k+1:i}}{\mathbb{E}}[\mathcal{H}(X_i|h_i)] = C - \frac{1}{2}\underset{Z_{k+1:i}}{\mathbb{E}}\left[\ln\left|\Lambda_k^{\text{Aug}} + A_i(z)^T A_i(z)\right|\right]$
- $C = N \ln(2\pi e)$

ANPL
Autonomous Navigation and Perception Lab

Partitioning of Gaussian Belief

- $\Lambda_i = \Lambda_k^{\text{Aug}} + A_i^T A_i$
- Observation partitioning corresponds to splitting the Jacobian into blocks

$$Z_i \quad \mapsto \quad (Z_i^s, Z_i^{\bar{s}})$$

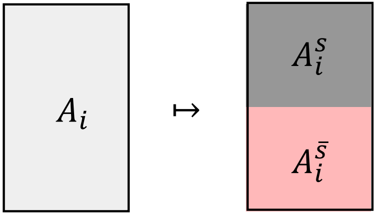

Gaussian Bounds

$$f(\Lambda, A) \triangleq |\Lambda + A^T A|$$

$$\mathcal{LB} = C - \frac{1}{2} \underset{Z_{k+1:i}}{\mathbb{E}} \left[\ln \frac{f\left(\Lambda_k^{\text{AUG}-}, A_i^s\right) \cdot f\left(\Lambda_k^{\text{AUG}-}, A_i^{\bar{s}}\right)}{|\Lambda_k^{\text{AUG}-}|}\right]$$

$$\mathcal{UB} = C - \frac{1}{2} \underset{Z_{k+1:i}}{\mathbb{E}} \left[\ln f\left(\Lambda_k^{\text{AUG}-}, A_i^s\right)\right]$$

Estimation of Expected Entropy

Assumption - the measurement Jacobian is not a function of Z

$$\mathcal{H}\left(X_i | Z_{k+1:i}\right) = C - \tfrac{1}{2}(\ln f\left(\Lambda_k^{\textsc{Aug}-}, A_i\right)) \rightarrow$$

$$\mathcal{LB} = C - \frac{1}{2}\ln \frac{f\left(\Lambda_k^{\textsc{Aug}-}, A_i^{s}\right) \cdot f\left(\Lambda_k^{\textsc{Aug}-}, A_i^{\bar{s}}\right)}{|\Lambda_k^{\textsc{Aug}-}|}$$

$$\mathcal{UB} = C - \frac{1}{2}\ln f\left(\Lambda_k^{\textsc{Aug}-}, A_i^{s}\right)$$

ANPL
Autonomous Navigation and Perception Lab

Methods for Determinant Calculation

Need to evaluate $f\left(\Lambda_k^{\text{AUG}-}, A_i^s\right), f\left(\Lambda_k^{\text{AUG}-}, A_i^{\bar{s}}\right)$

- Baseline:

$$|\Lambda + A^T A|$$

- Square root form:

$$|R^T R|$$

- rAMDL:

$$|\Lambda| \cdot |I_m + A\Lambda^{-1}A^T|$$

ANPL Autonomous Navigation and Perception Lab

rAMDL

- Apply Matrix Determinant lemma to posterior Information matrix (Kopitkov and Indelman - IJRR'17)
- Baseline is $O(N^3)$ while rAMDL is $O(m^3)$, Partitioning the measurements reduces m.

ANPL Autonomous Navigation and Perception Lab

rAMDL + Bounds

- Applying rAMDL to the determinants required for $\mathcal{LB}, \mathcal{UB}$:

$$|\Lambda_k^{\text{AUG}-} + (A_i^s)^T A_i^s|$$

$$|\Lambda_k^{\text{AUG}-} + (A_i^{\bar{s}})^T A_i^{\bar{s}}|$$

- When $s = \bar{s} \rightarrow |\Lambda_k^{\text{AUG}-} + (A_i^s)^T A_i^s| = O(\frac{m^3}{8})$

ANPL Autonomous Navigation and Perception Lab

Outline

ANPL
Autonomous Navigation and Perception Lab

Scenario

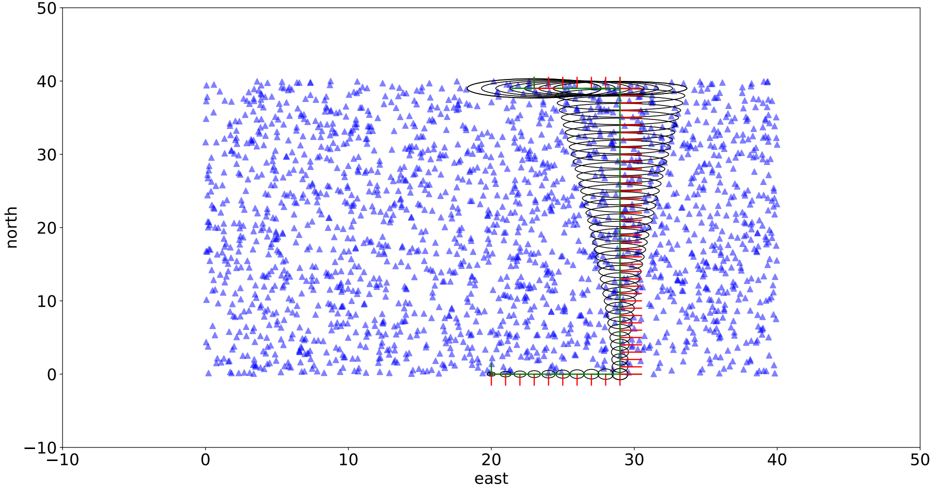

Scenario

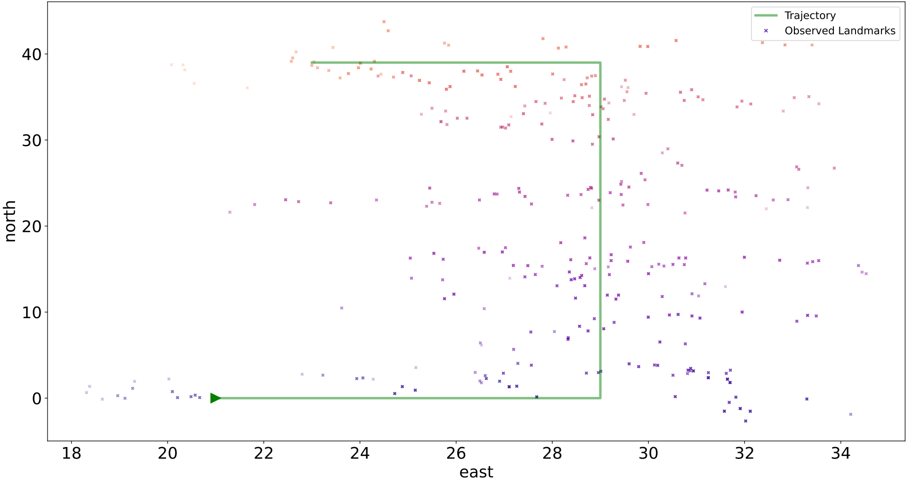

Trajectory
Observed Landmarks
north
east
40
30
20
10
0
18
20
22
24
26
28
30
32
34

Results - Bounds Analysis

For a specific action, $Z^s \rightarrow Z$ while $Z_{\bar{s}} \rightarrow \emptyset$

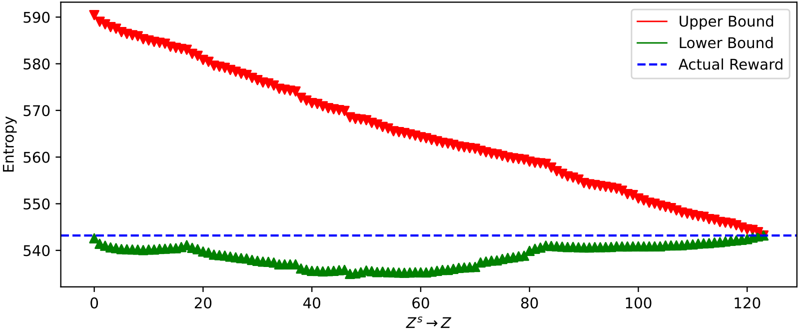

Results - Bounds Analysis

Going down the partition tree:

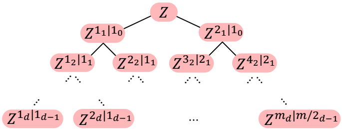

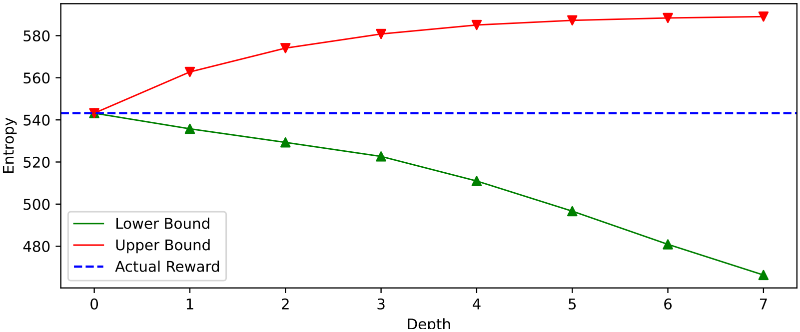

Results - Bounds Analysis

Myopic, random actions:

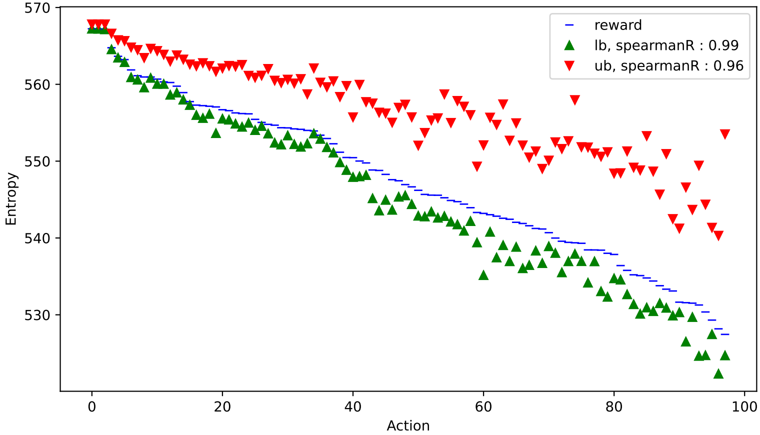

Results - Bounds Analysis

Two actions, different depths:

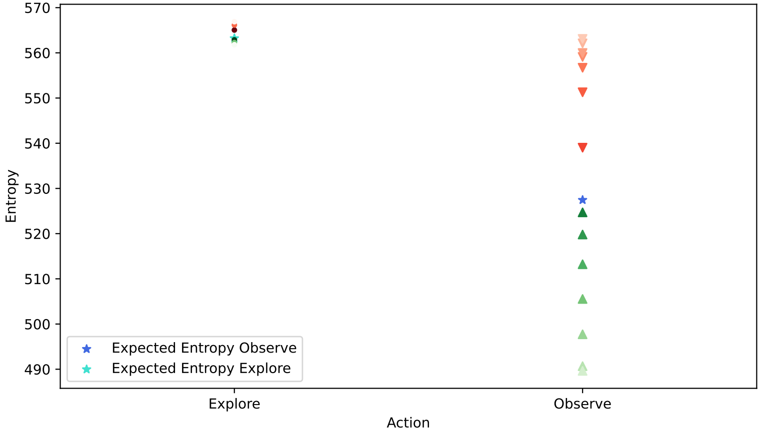

Results - Simulation

100 different paths, not including common terms

- Scenario 1 - 2956 factors without re-planning
- Scenario 2 - 2956 factors with re-planning
- Scenario 3 - 5904 factors with re-planning

Illustration - Scenario 1

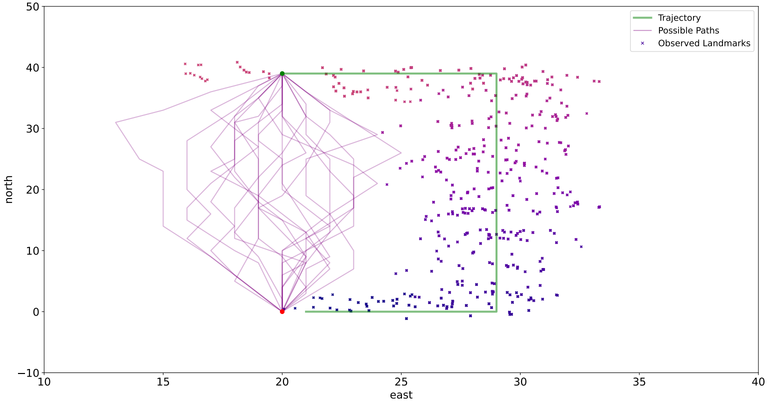

Bounds - Scenario 1

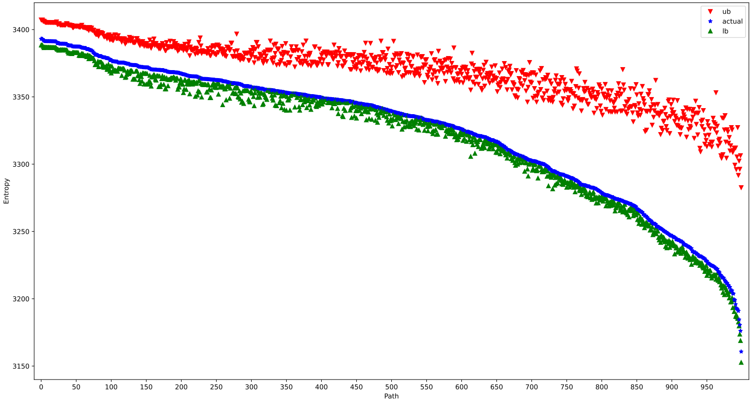

Illustration - Scenario 2

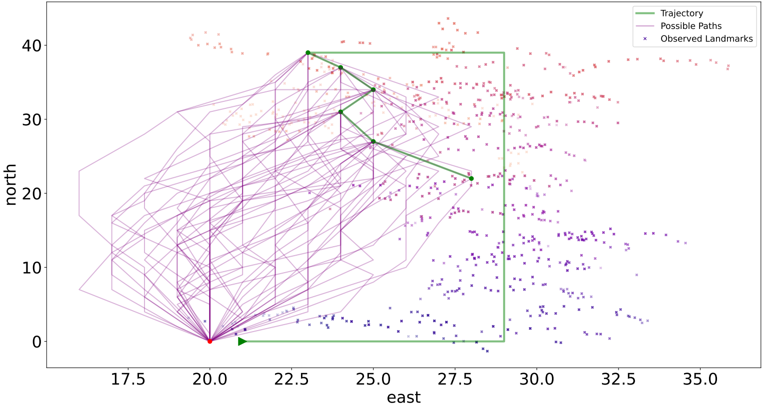

Results - run time

# Paths	# Factors	RP	rAMDL	MP (ours)
100	2956	No	11.521 ± 0.537	$\mathbf{6.888 \pm 0.155}$
100	2956	Yes	24.636 ± 1.381	$\mathbf{11.758 \pm 0.372}$
100	5904	Yes	84.376 ± 14.458	$\mathbf{32.069 \pm 4.913}$

Table: Total planning time in seconds (lower is better)

ANPL Autonomous Navigation and Perception Lab

Results - Visual Odometery

RoboMaster + ZED

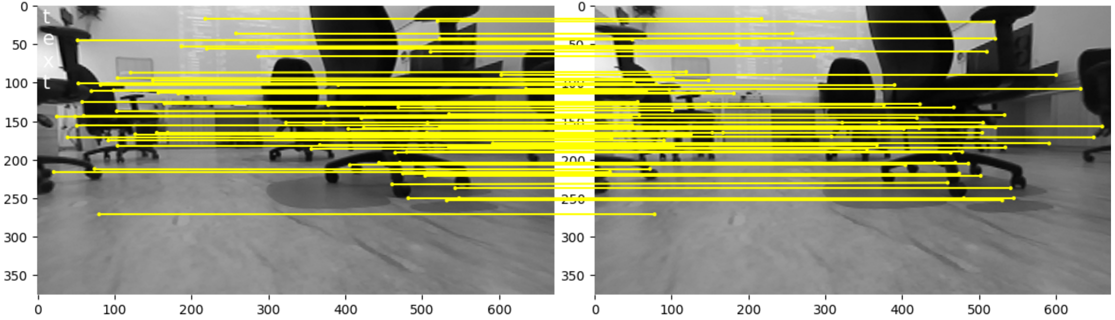

Illustration - Visual Odometery

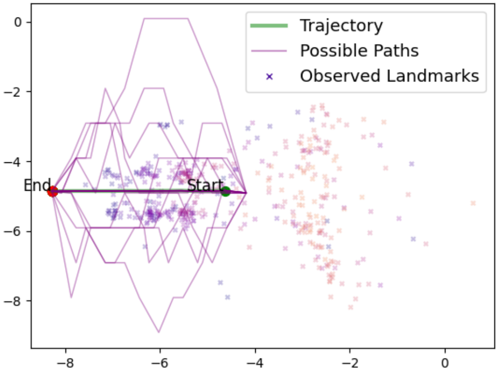

Bounds - Visual Odometery

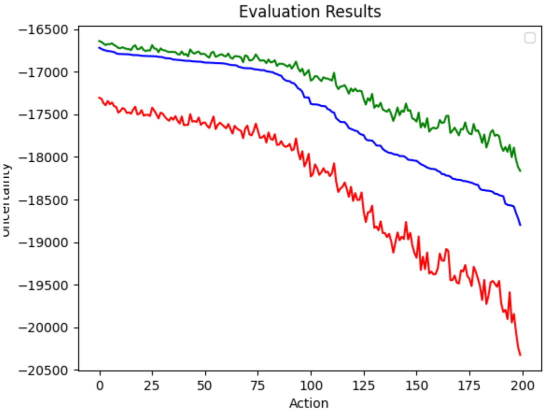

Results - run time

Method	time [sec]
MP (ours)	$\mathbf{585.507 \pm 27.153}$
rAMDL	802.545 ± 25.651
iSAM2	1764.835 ± 26.521

Table: Total planning time in seconds (lower is better)

worst-case [sec]	actual [sec]
166	97.6

Table: Covariance recovery time

ANPL Autonomous Navigation and Perception Lab

Conclusions

- We introduced a novel concept - observation space partitioning
- We proposed a simplified method to solving the POMDP planning problem using this concept
- We presented both theoretical and empirical studies of this method - both showing performance gains
- Multiple future research directions

ANPL Autonomous Navigation and Perception Lab

Questions