ANPL | Autonomous Navigation and Perception Lab

Joint Incremental Inference & Belief Space Planning for Online Operations of Autonomous Systems

Elad I. Farhi

Advisor: Vadim Indelman

December 2020

TECHNION Israel Institute of Technology

TASP | TECHNION AUTONOMOUS SYSTEMS PROGRAM

Introduction

- Inference and Decision making under uncertainty impose a fundamental problem in Autonomous Systems (AS) and Artificial Intelligence (AI).

- At their core, Autonomous Systems require the following blocks

Inference & Perception

Obtain information from the environment (and\or other agents) and estimate state variables, using existing data

Planning

Plan next best action given current belief and objective function

Introduction

- The realistic problem is computationally intractable, hence usually approximated.
- Any reduction in computation time would pave the way to Online\ Realtime work.
- There are many variations of AS/ AI related problems

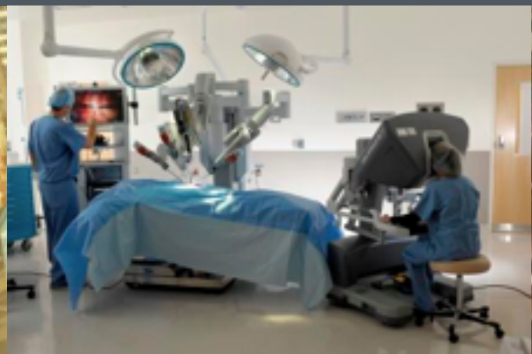

- Autonomous Navigation
- Search & Rescue scenarios
- Robot assisted
- Simultaneous Localisation & Mapping (SLAM)
- Business Decision Making
- Stock Market Prediction

Inference & Belief Space Planning (BSP) today

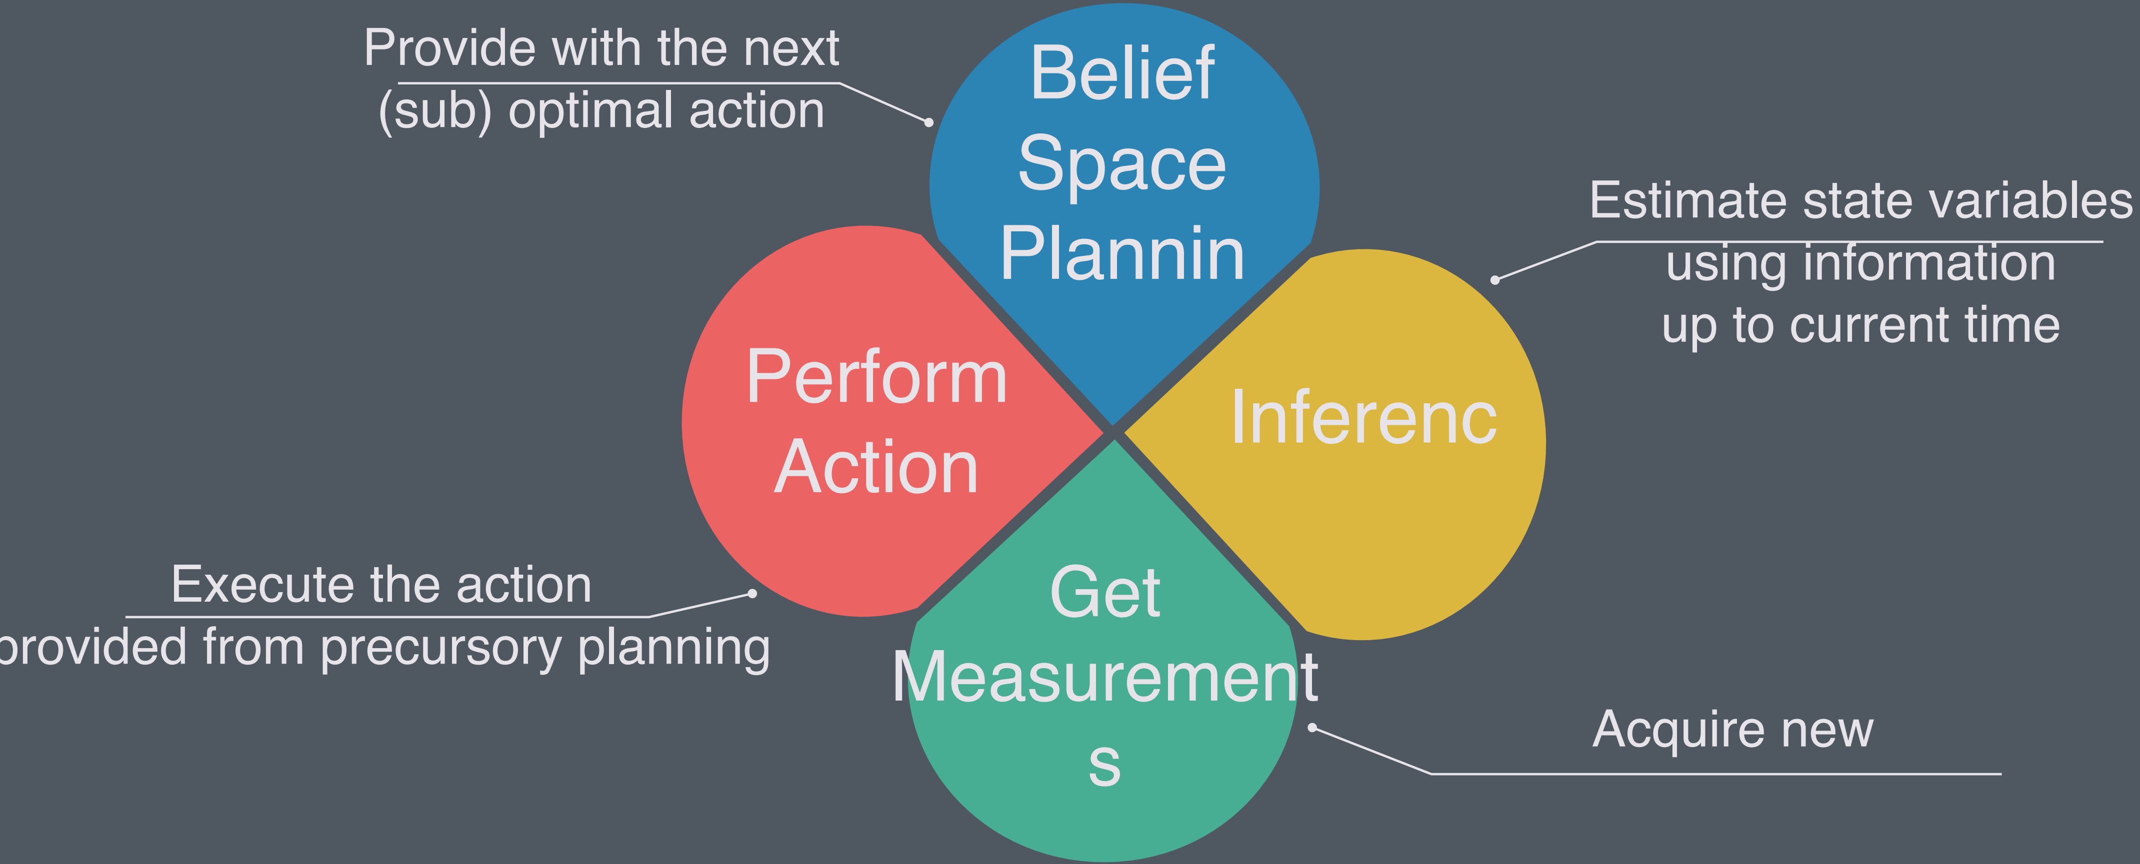

Inference & BSP today

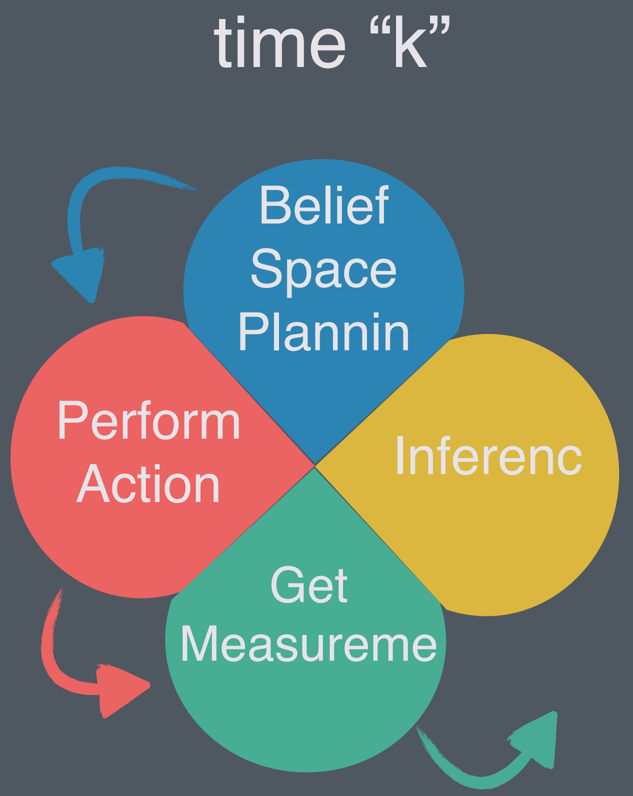

Inference & BSP today

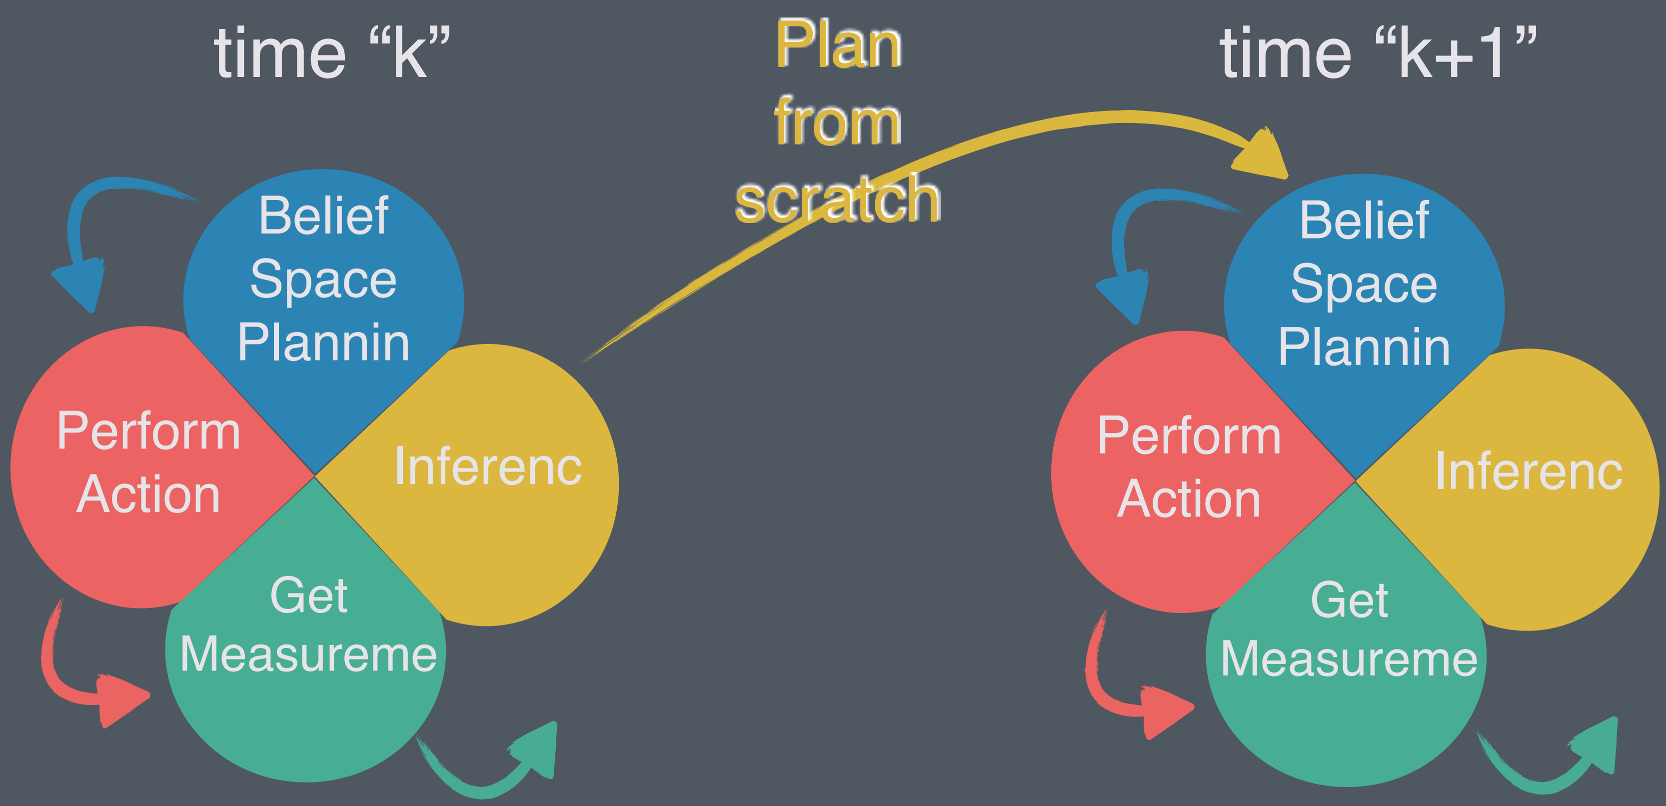

Inference & BSP today

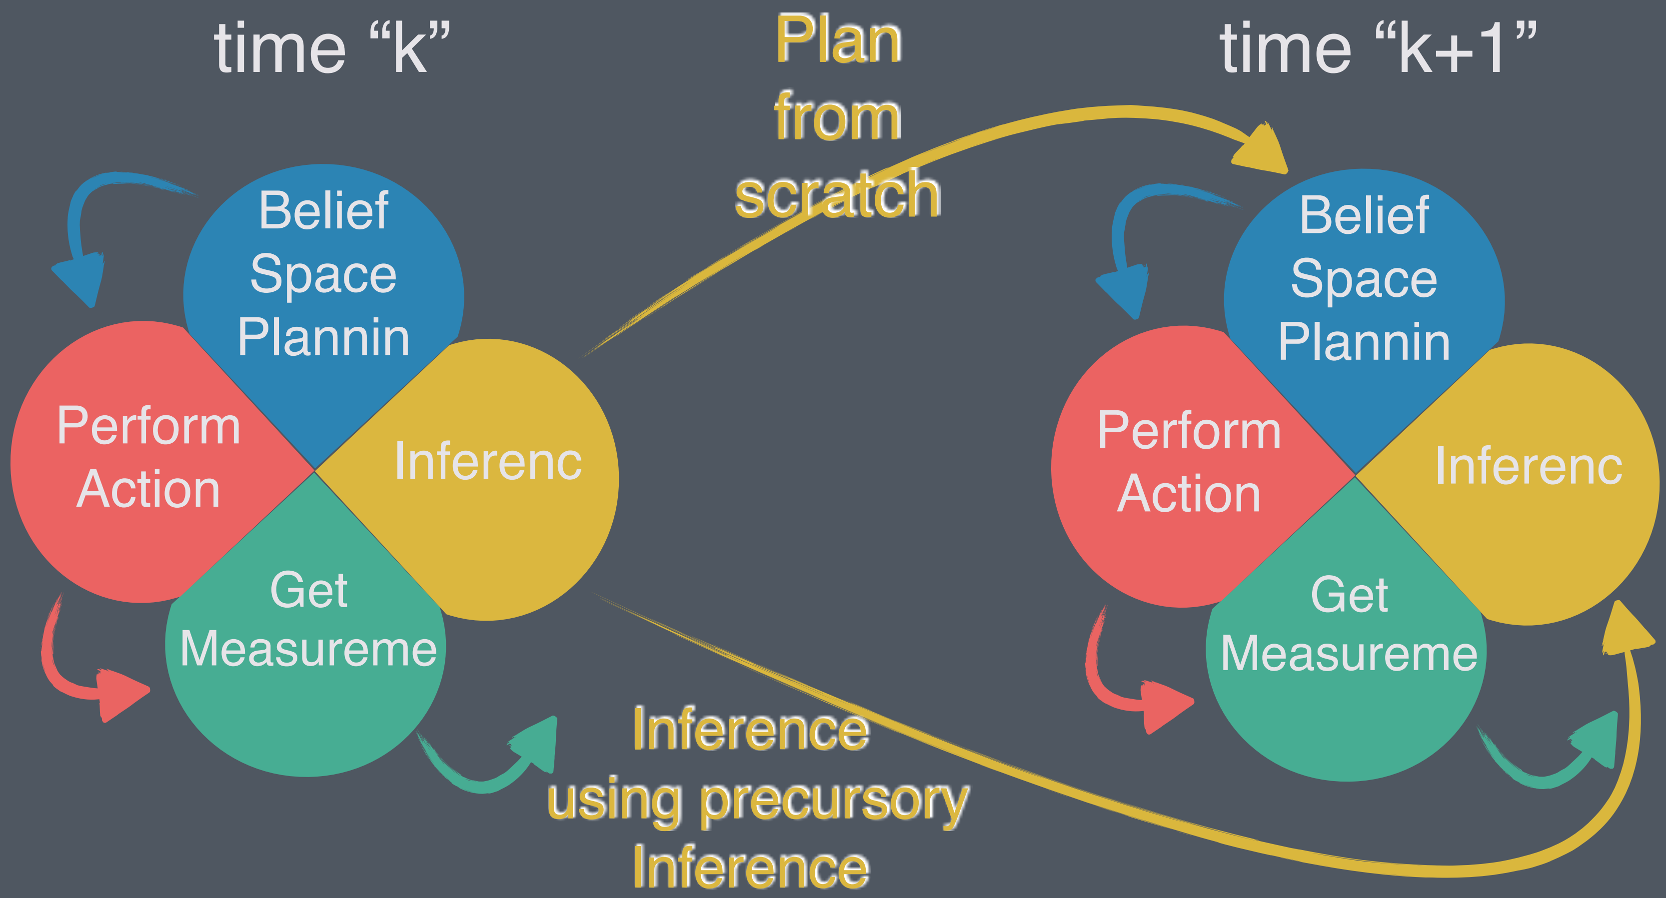

Inference & BSP today

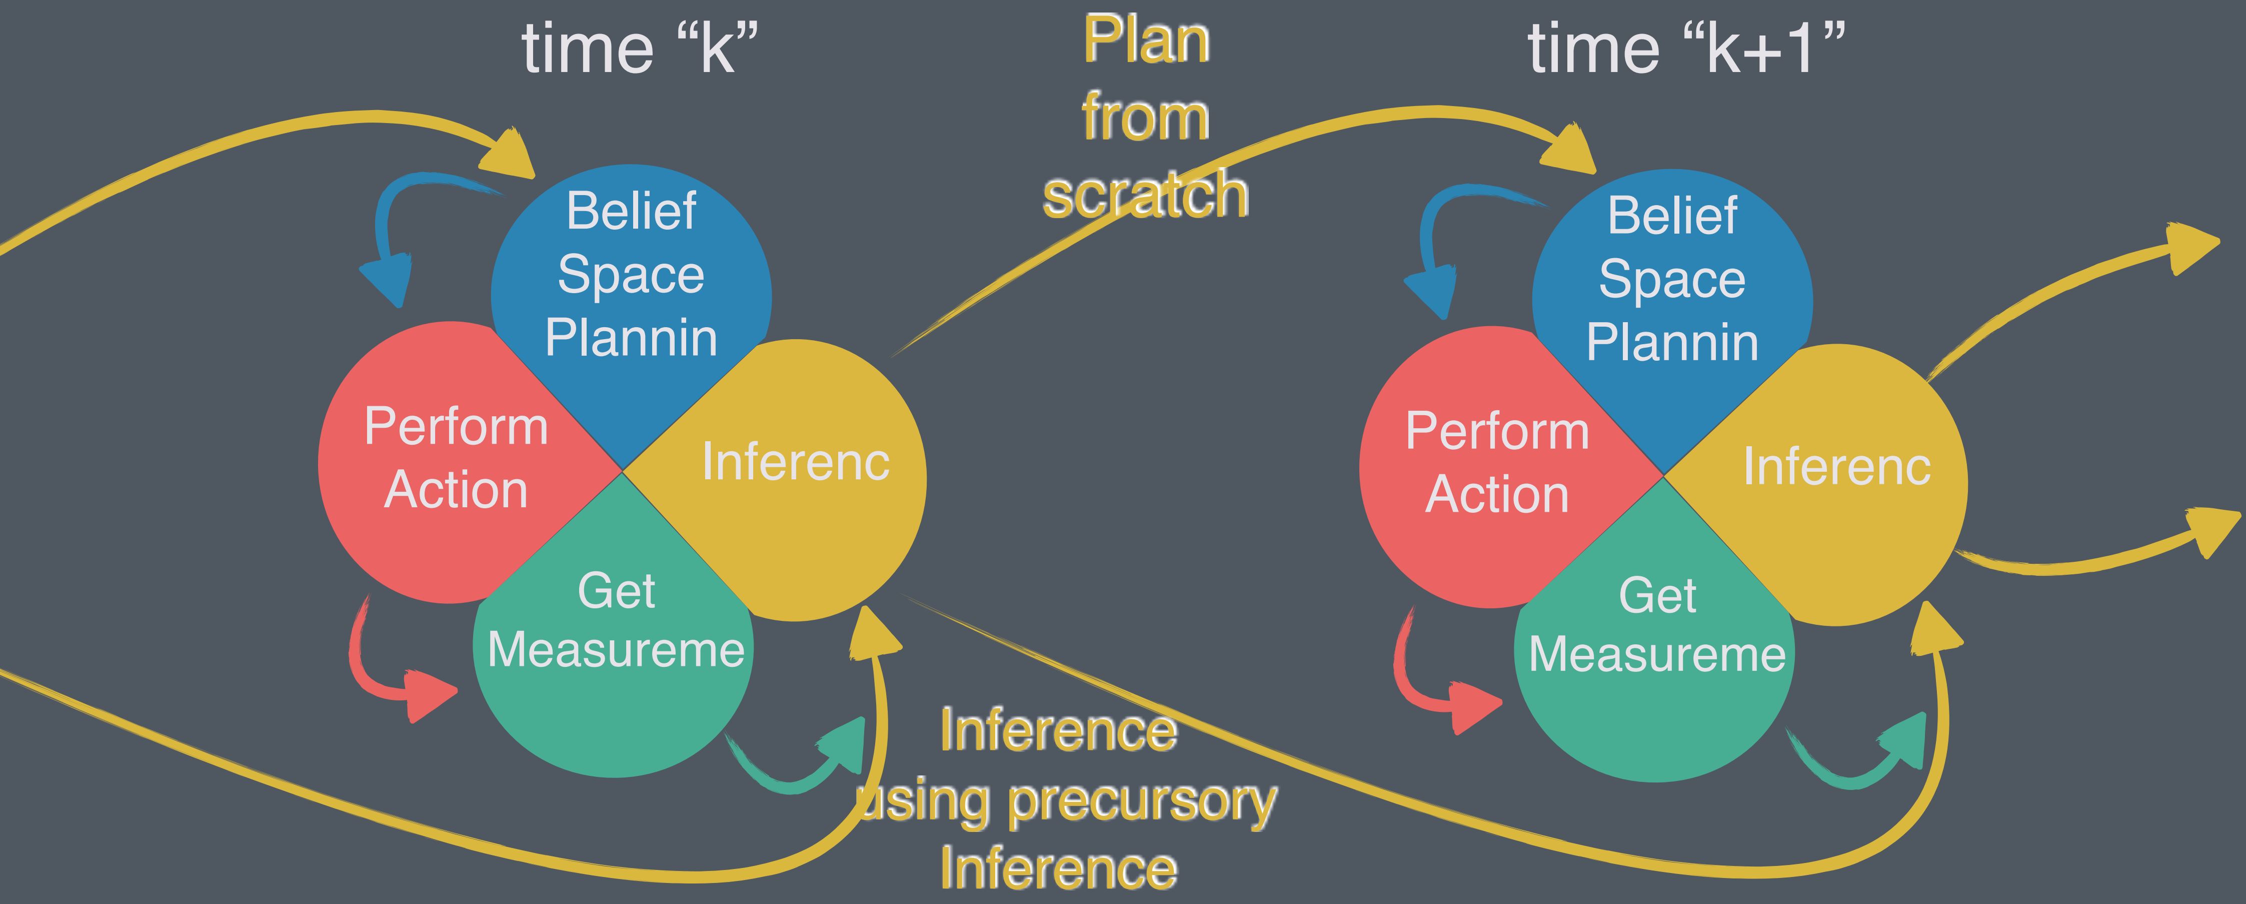

Inference & BSP today

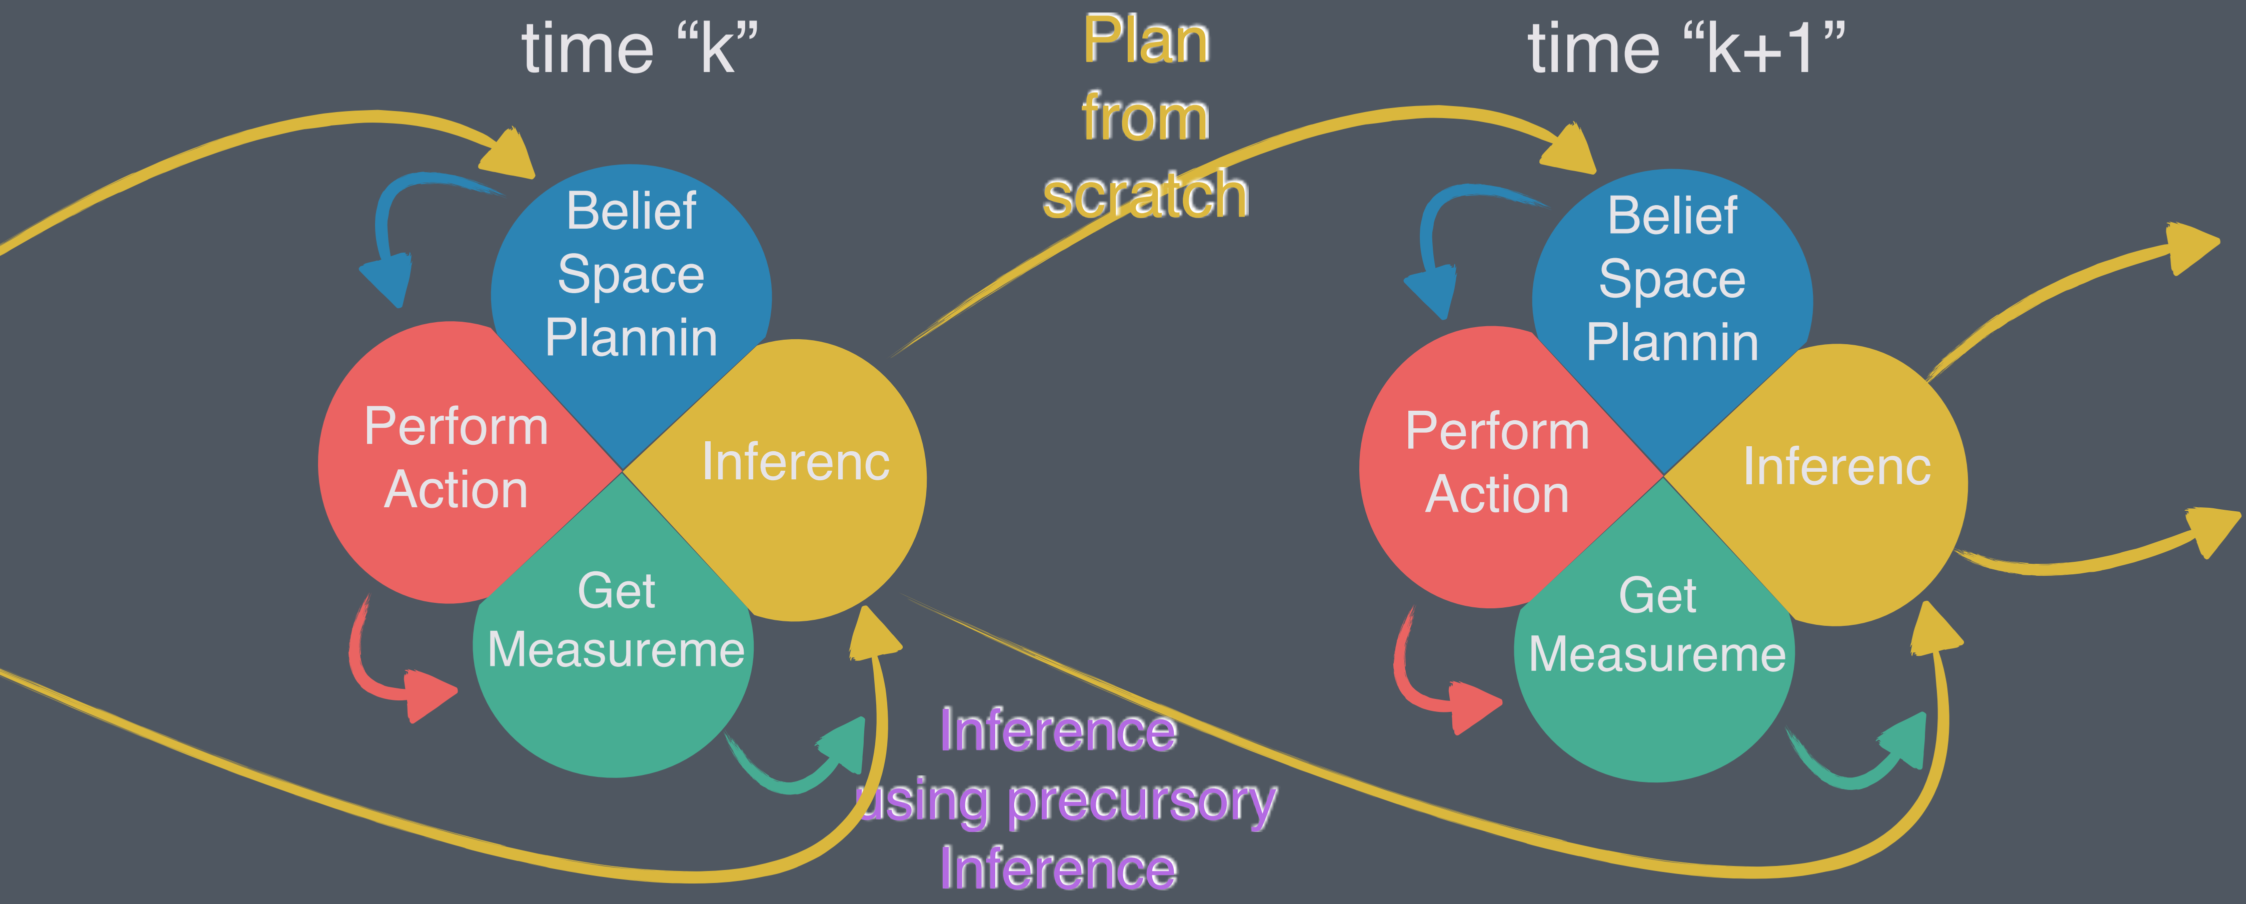

The Job Interview Example

- Let's say you have an interview for your dream job.
- You'll probably prepare yourself by going over all subjects you might be asked about.

The Job Interview Example

- The interview day has arrived, what would happen if you'll be asked on a subject,

Inconsistent Data Association
- you didn't cover ?
- close to what you have covered ?

Consistent Data Association
- identical to what you have covered ?

- Which would result in the quickest answer ?

Inference & BSP today

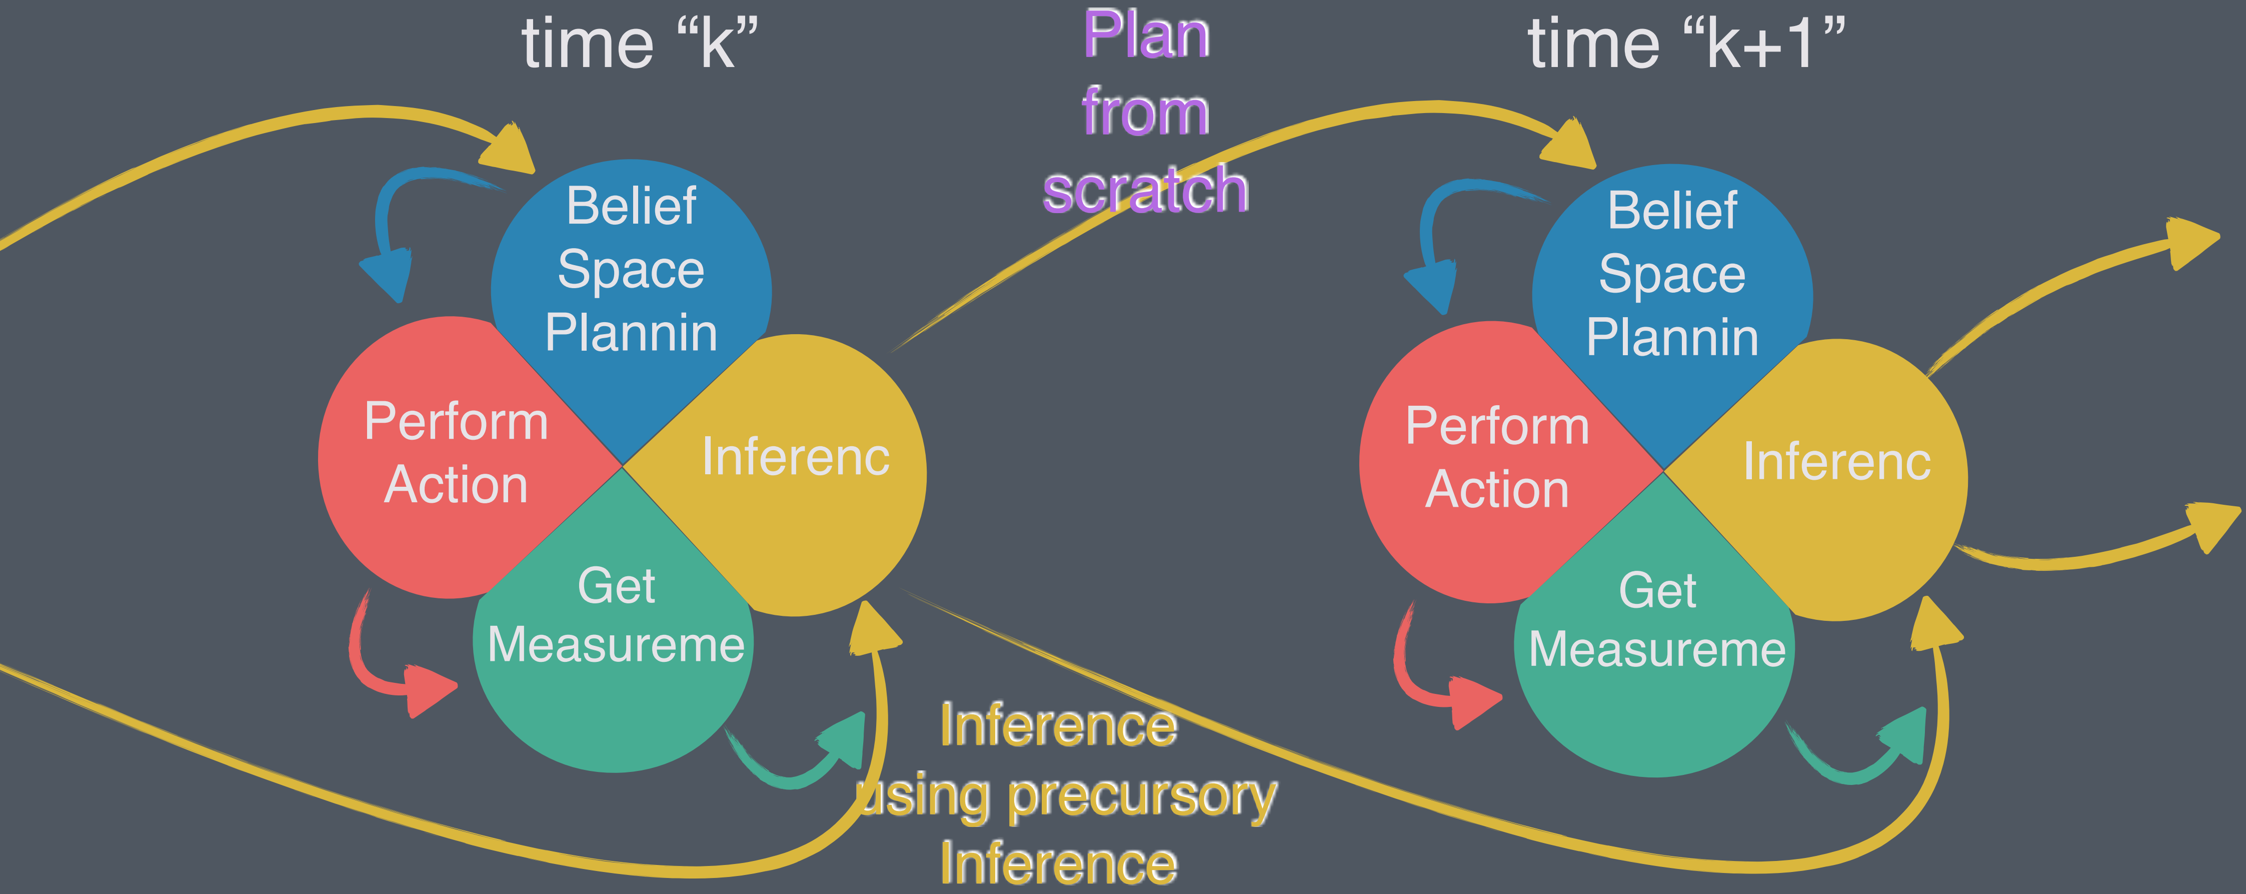

Home


Work
Home


Work
MALL
SALE
CINEMA
Home


Work
MALL
SALE
CINEMA
Home

The "Driving to Work" Example

- Will you plan everything from scratch?
- Or just update the appropriate segments of your original plan with this new information?
- While this toy example considers MDP setting with an observable and deterministic world

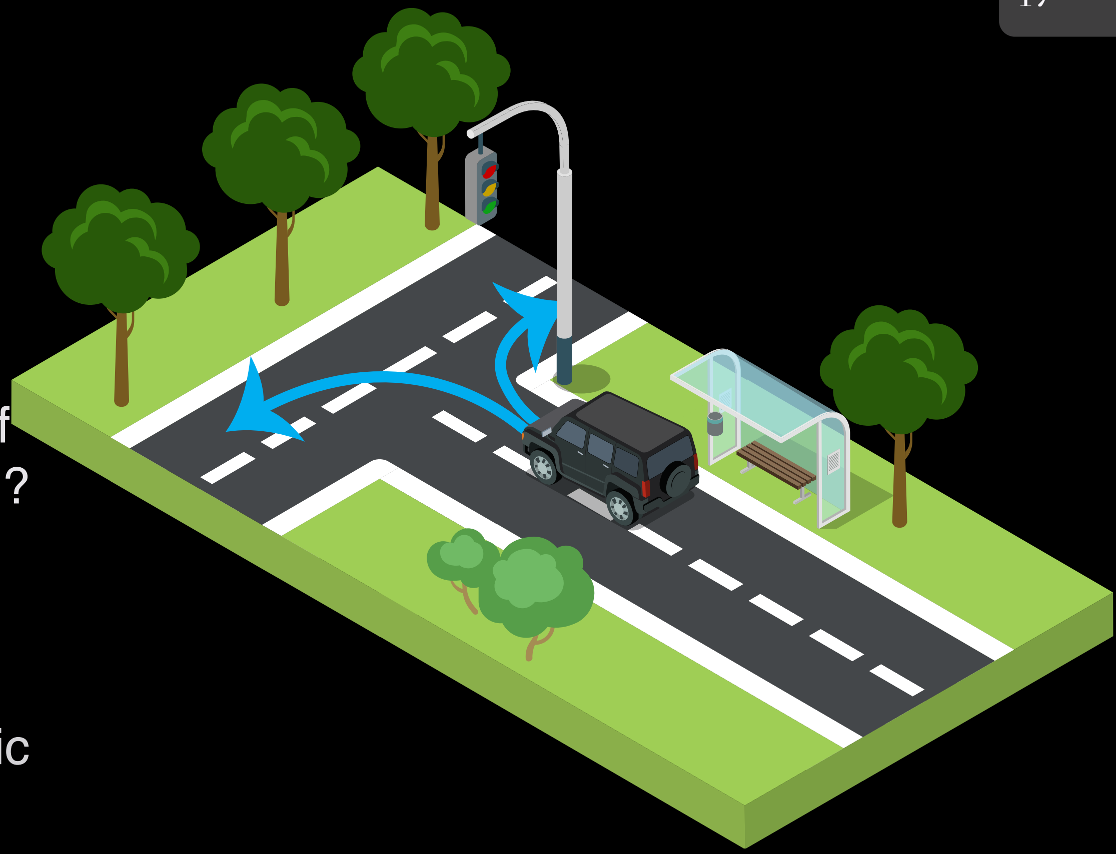

We consider the more general problem of a POMDP setting with an unknown world and a high dimensional state vector

Our Research Vision

Re-use prior calculations and information across inference and planning, for improved online autonomy, in particular in unknown/uncertain environments and high-dimensional state spaces.

Our Research Vision

Re-use prior calculations and information across inference and planning, for improved online autonomy, in particular in unknown/uncertain environments and high-dimensional state spaces.

Main Contributions

- Introducing Joint Inference & Planning - JIP as a novel paradigm shift from the common separation of inference and planning. (Farhi17icra) (Farhi19icra workshop)
- A novel approach for Re-Use BSP for efficient Inference update, named RUBI (Farhi17icra) (Farhi18ijrr conditionally accepted) (Farhi19icra workshop) (patent: WO2019171378A1)
- A novel approach for incremental eXpectation Belief Space Planning, named iX-BSP (Farhi19icra) (Farhi19icra workshop) (Farhi20journal to be submitted soon) (patent: US20200327358A1)

Our Novel Approach - Re-Use BSP for Inference

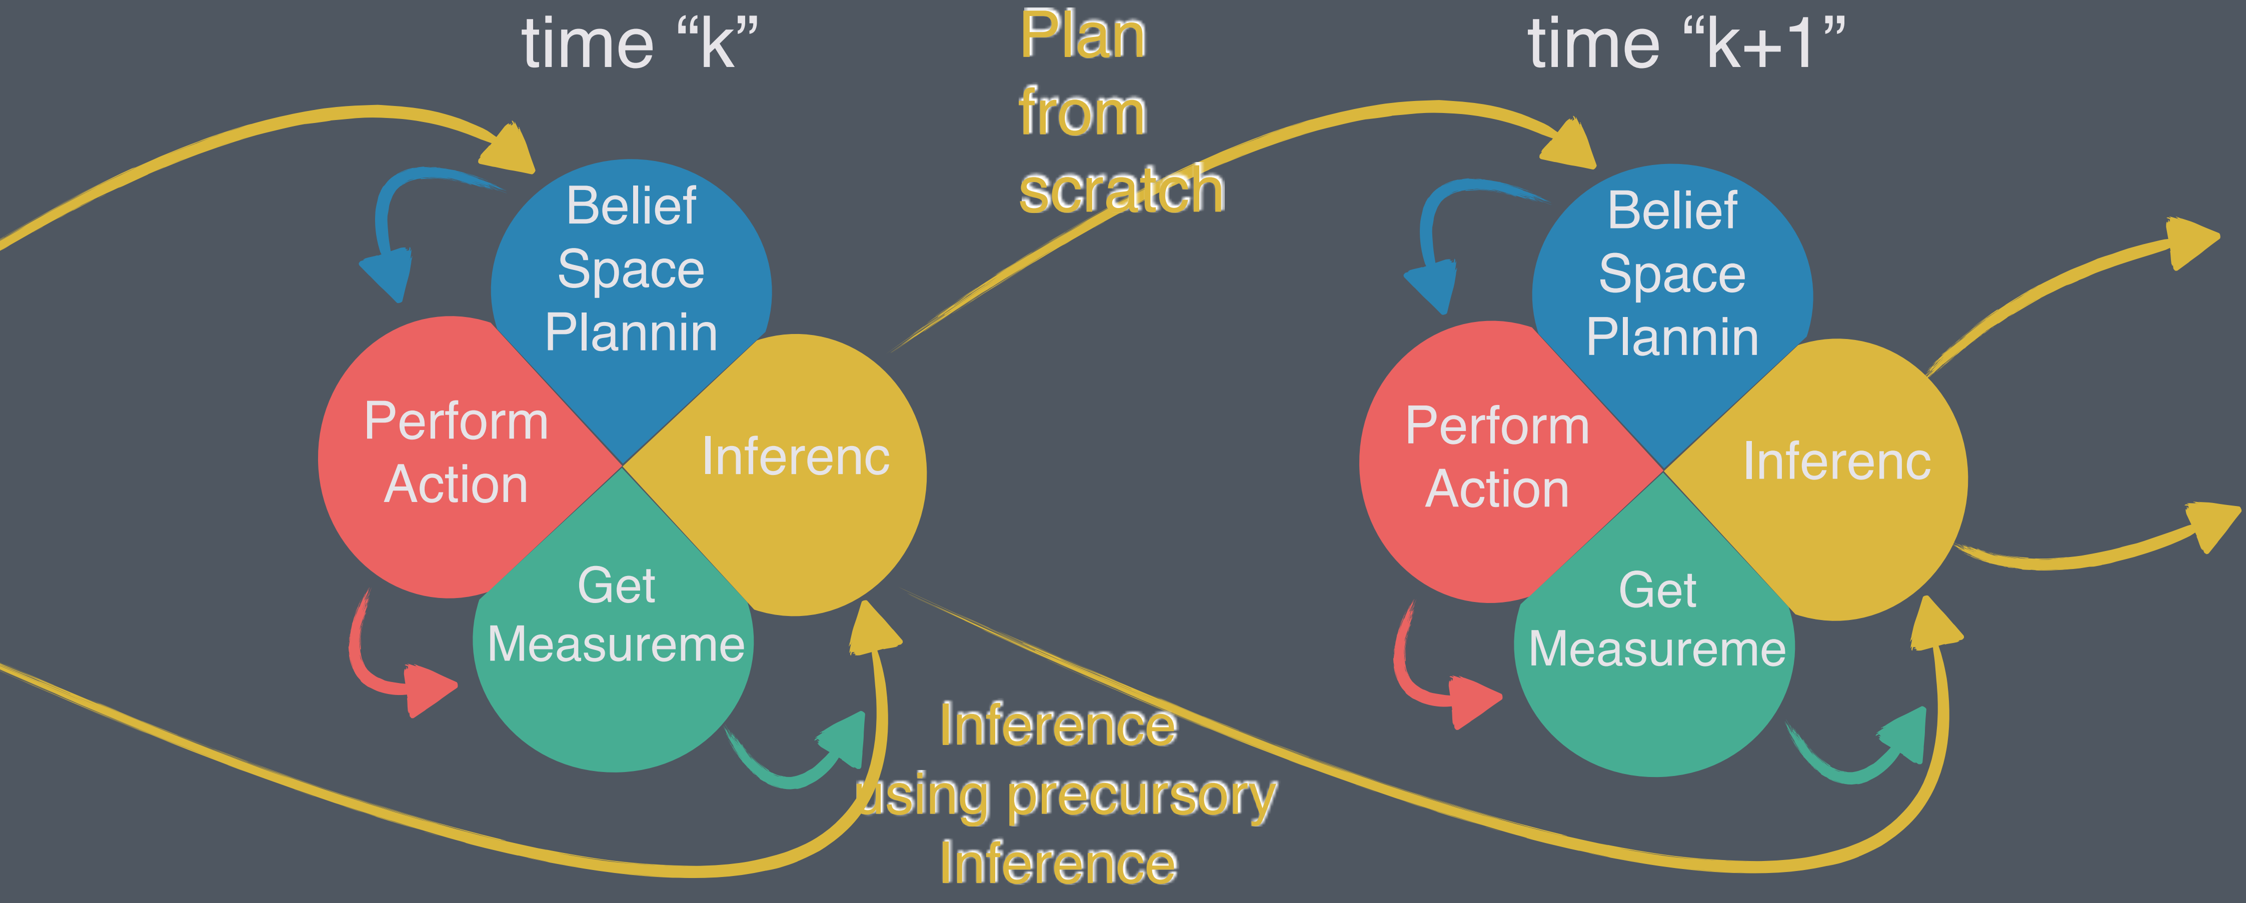

Our Novel Approach - Re-Use BSP for Inference

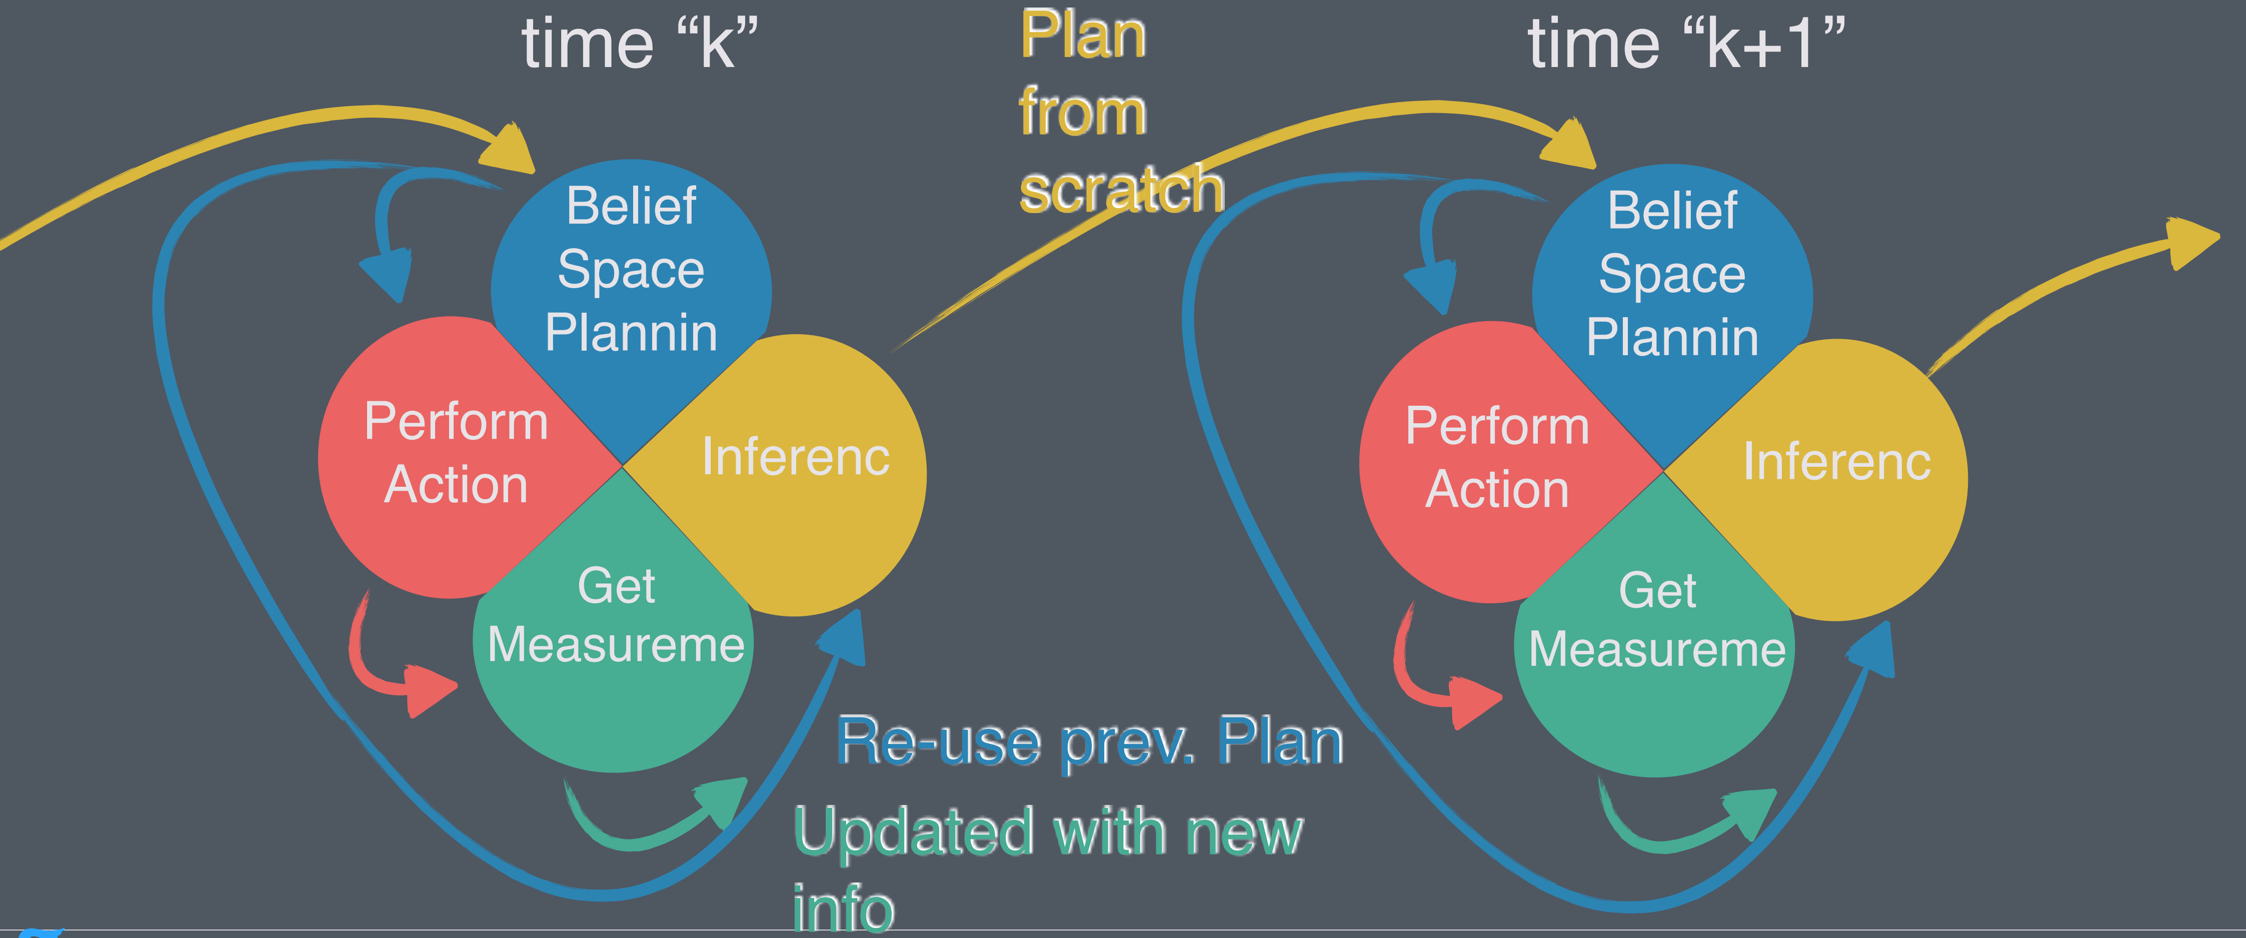

Our Novel Approach - Incremental eXpectation BSP

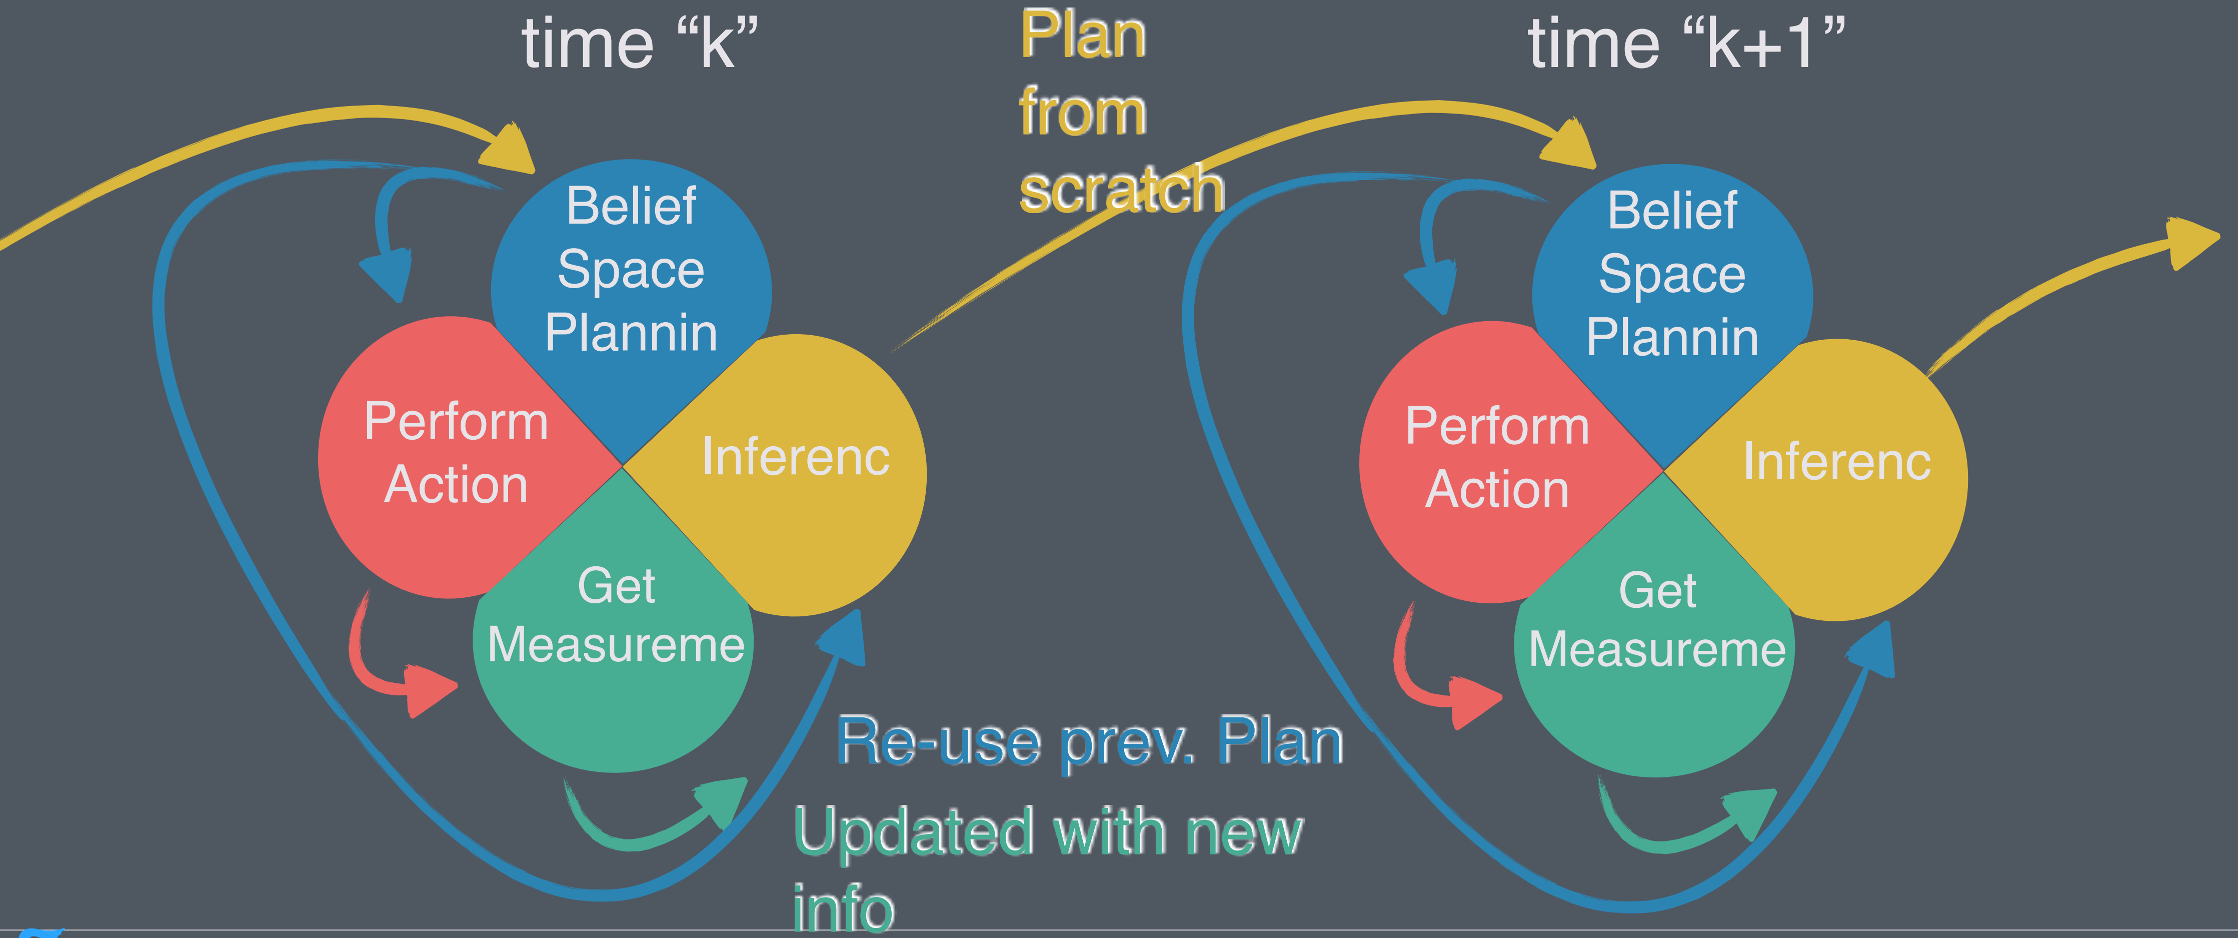

Our Novel Approach - Incremental eXpectation BSP

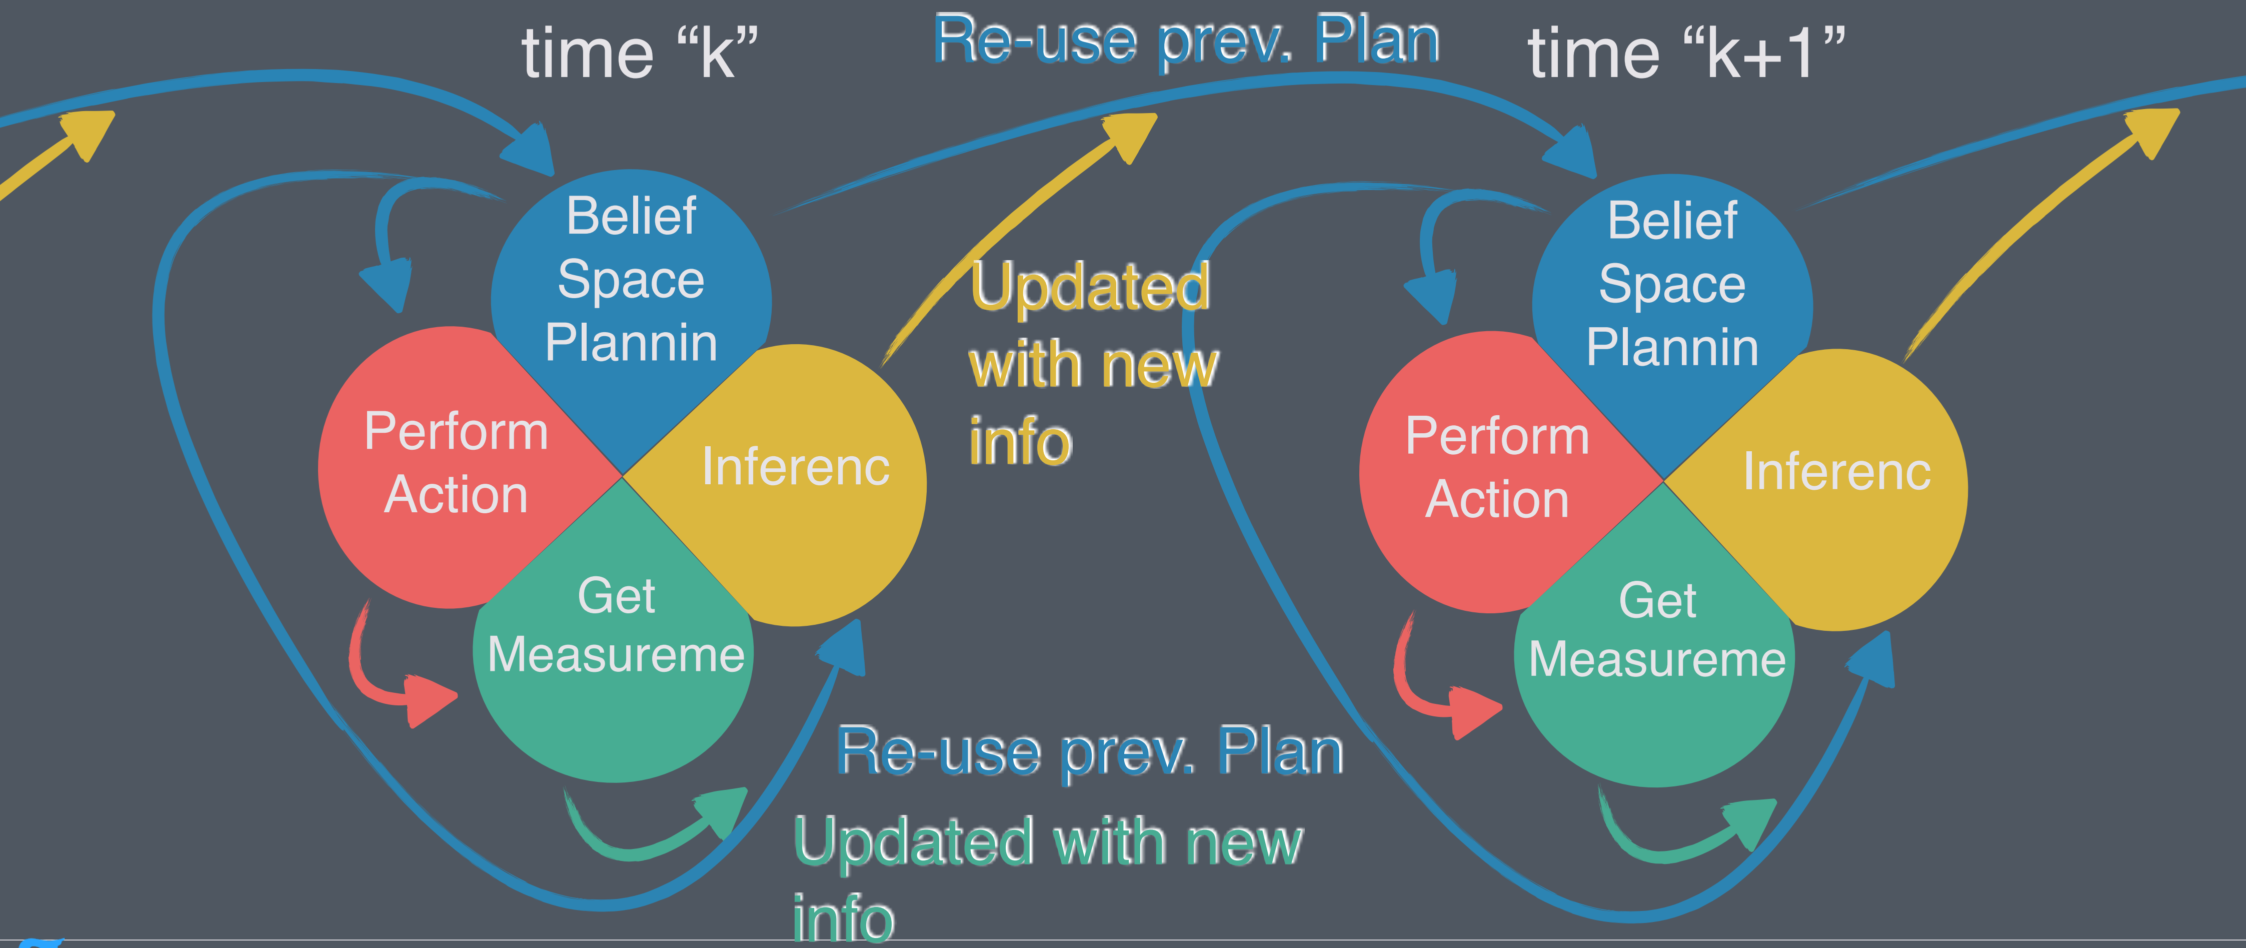

Notations & Formulation

- $\square_{t|k}$ - Referring to time t, while current time is k
- X_t - The joint state vector up to time t, i.e. smoothing problem (all robot poses and landmarks)
- $z_{1:t|k}$ - All measurements up to time t, while current time is k
- $u_{0:t-1|k}$ - All actions up to time $t-1$, while current time is k
- $b[X_{k|k}] = p(X_k | u_{0:k-1|k}, z_{1:k|k})$ belief at current time k
- $b[X_{k+i|k}] = p(X_{k+i} | u_{0:k+i-1|k}, z_{1:k+i|k})$ belief at planning horizon i
- $H_{k+i|k} \doteq \{z_{1:k+i|k}, u_{0:k+i-1|k}\}$ History at planning horizon i
- $H^-_{k+i|k} \doteq \{H_{k+i-1|k}, u_{k+i-1|k}\}$ Propagated history at planning horizon i

Research Outline

- Introducing Joint Inference & Planning
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

ANPL Autonomous Navigation and Perception Lab

Research Outline

- Introducing Joint Inference & Planning
  - Related Work
  - Unified Model for Inference & BSP
  - RUBI as part of JIP
  - iX-BSP as part of JIP
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

Related work on Inference & BSP Similarities

- Approximate solutions to the Markov Decision Process (MDP) case, for inference and planning, using inference optimization methods (Toussaint and Storkey 2006)
- Investigating the duality between inference and optimal control (Todorov 2008)
- Unified computational frameworks based on Dynamic Programing (Kobilarov 2015) and Factor Graph-FG (Ta 2014)
- Till this day, to the best of our knowledge,

  there is no Joint paradigm for inference and decision making under uncertainty

- Interestingly enough, inference and decision making under uncertainty in the human brain are tightly entwined, fact which provides motivation for AS & AI equivalent. (Schacter and Addis 2007) (Schacter and Addis 2009) (Race 2011)

Research Outline

- Introducing Joint Inference & Planning
  - Related Work
  - Unified Model for Inference & BSP
  - RUBI as part of JIP
  - iX-BSP as part of JIP
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

Research Outline

Introducing Joint Inference & Planning

- Related Work
- Unified Model for Inference & BSP
- RUBI as part of JIP
- iX-BSP as part of JIP

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

Unified Model for Inference & BSP

- Encapsulates both inference and planning separately.
- Enabling their "regular" functionality, as well as opening a gateway to new connections.

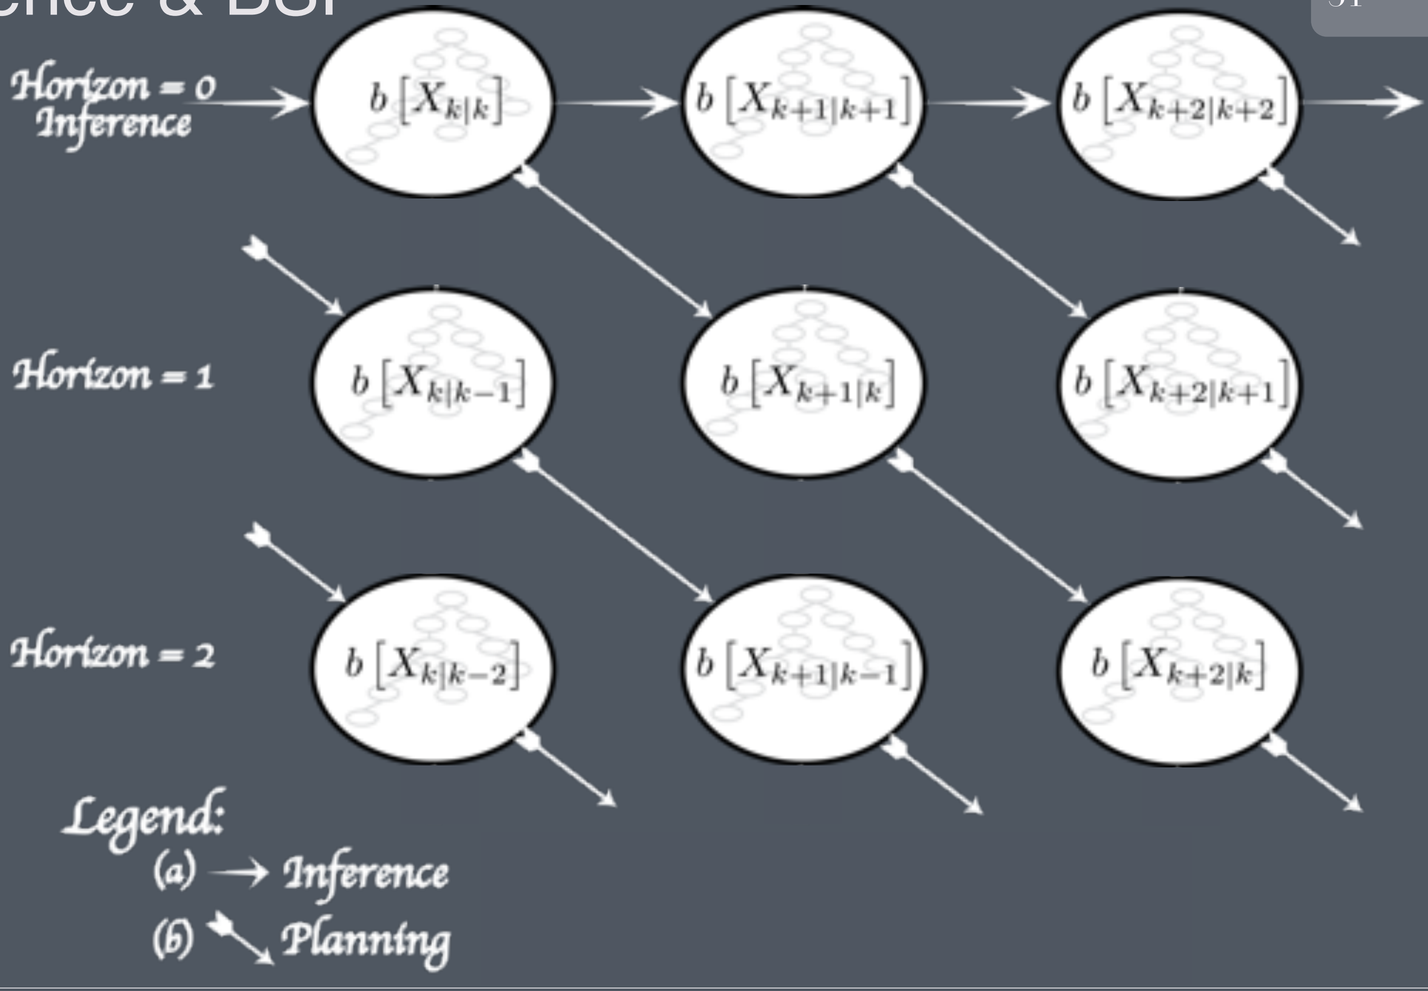

Research Outline

- Introducing Joint Inference & Planning
  - Related Work
  - Unified Model for Inference & BSP
  - RUBI as part of JIP
  - iX-BSP as part of JIP
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

Research Outline

- Introducing Joint Inference & Planning
  - Related Work
  - Unified Model for Inference & BSP
  - RUBI as part of JIP
  - iX-BSP as part of JIP
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

RUBI as part of JIP

- Conventional Bayesian inference - update inference using precursory inference
- We suggest a paradigm shift - update inference using precursory planning
- Saves valuable computation time without affecting estimation accuracy

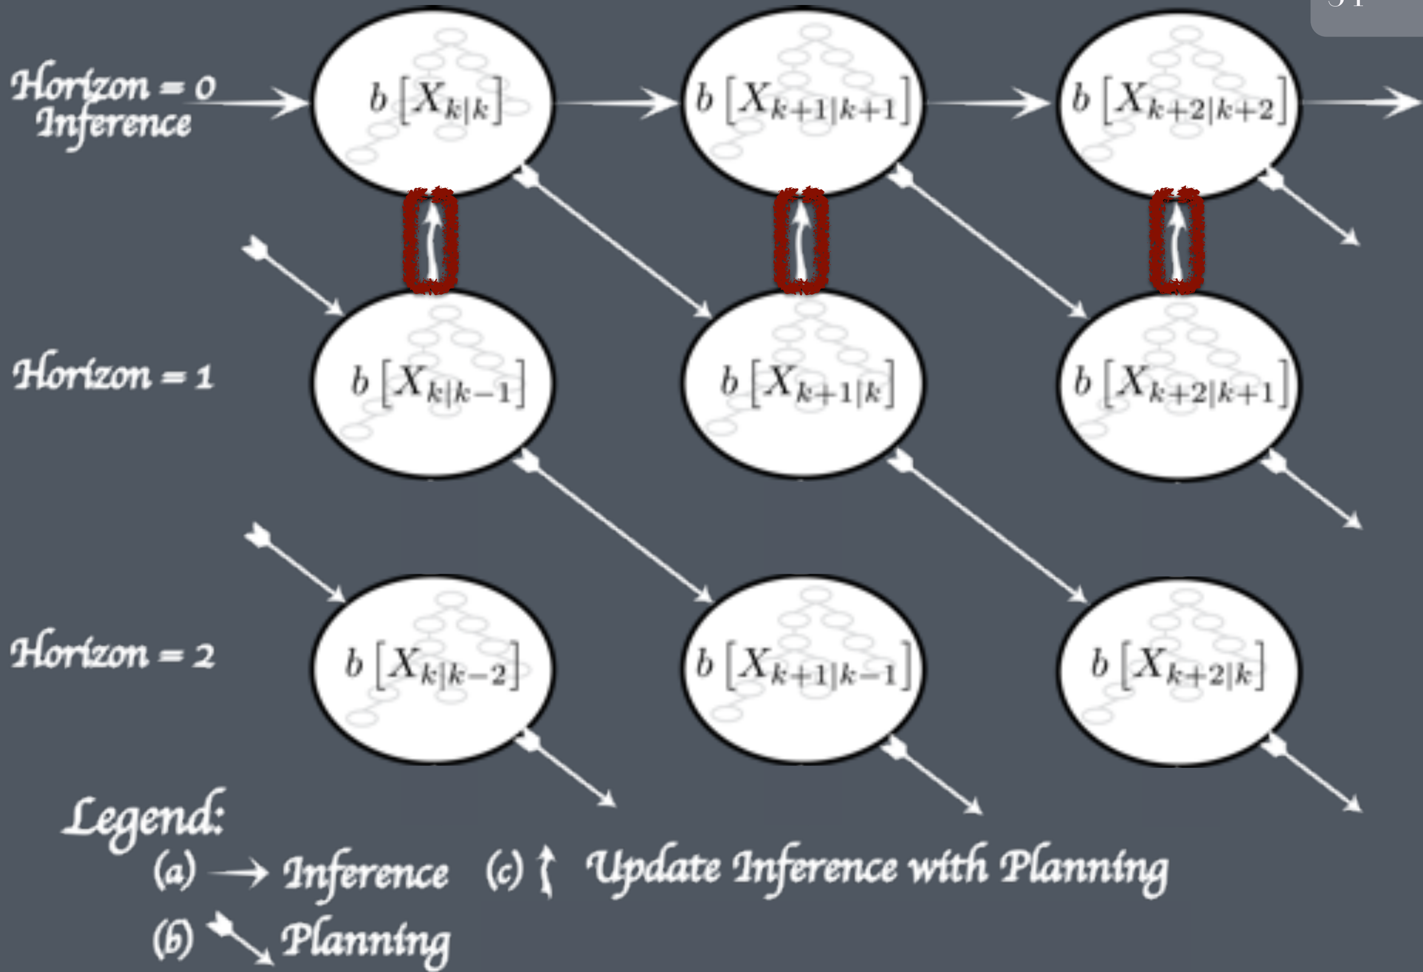

Research Outline

- Introducing Joint Inference & Planning
  - Related Work
  - Unified Model for Inference & BSP
  - RUBI as part of JIP
  - iX-BSP as part of JIP
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

Research Outline

- Introducing Joint Inference & Planning
  - Related Work
  - Unified Model for Inference & BSP
  - RUBI as part of JIP
  - iX-BSP as part of JIP
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

iX-BSP as part of JIP

- Uncertainty in the system and the environment forces re-planning in order to remain optimal
- Similarly to inference, planning can also benefits from re-using previous information
- Saves valuable computation time without affecting estimation accuracy

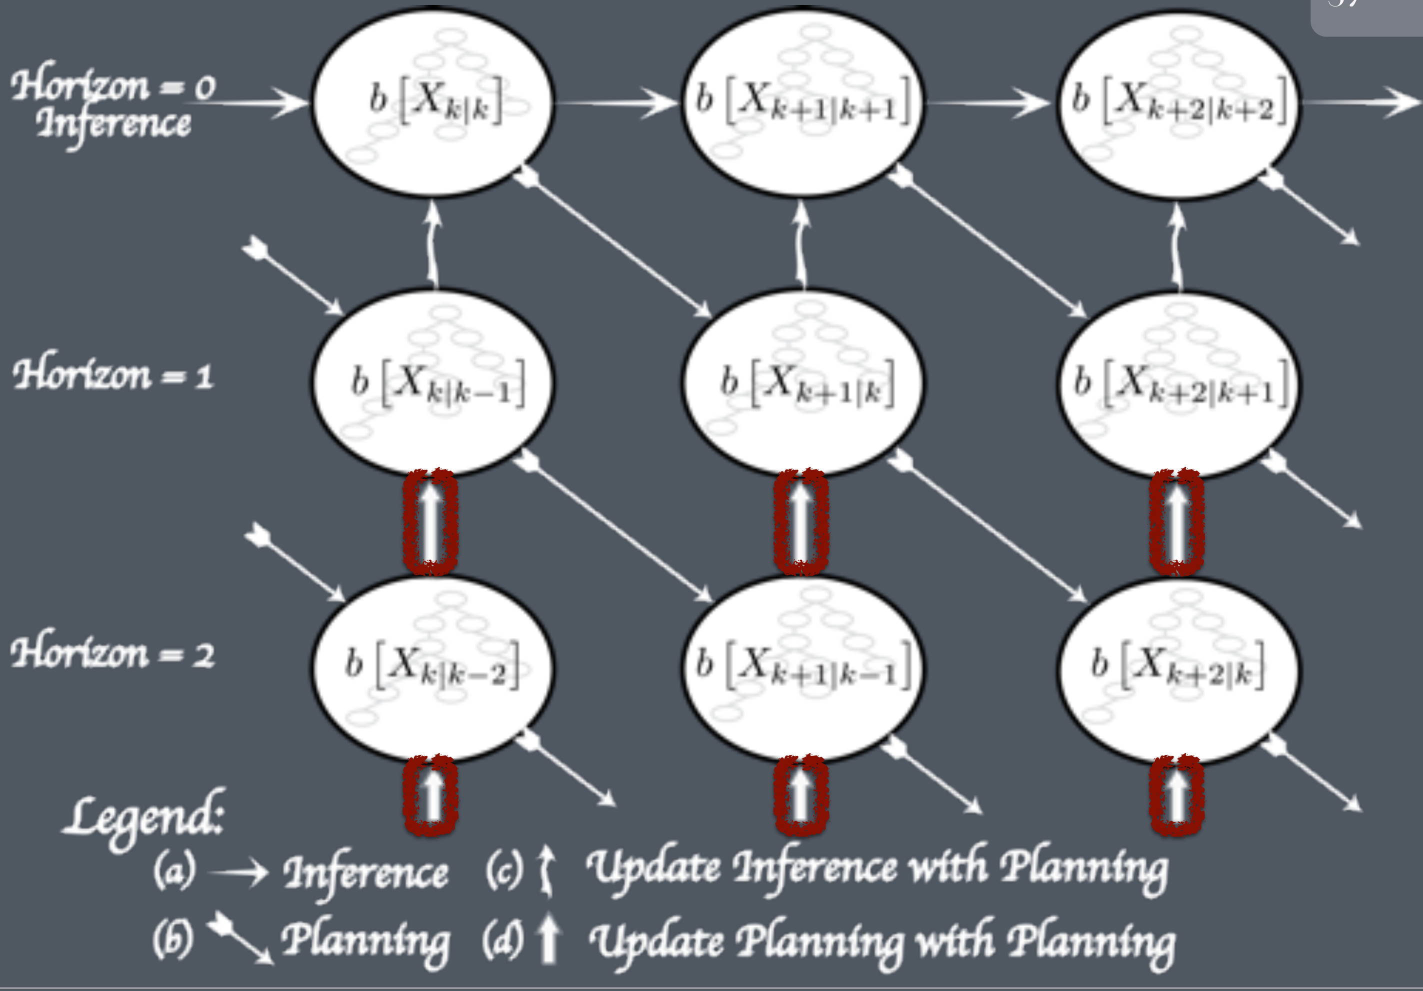

Research Outline

- Introducing Joint Inference & Planning
  - Related Work
  - Unified Model for Inference & BSP
  - RUBI as part of JIP
  - iX-BSP as part of JIP
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

Research Outline

- Introducing Joint Inference & Planning
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

- RUBI: Main Contributions
- Inference & BSP today
- Consistent DA assumption
  - Results
- RUBI
  - Results - simulation
  - Results - KITTI dataset

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

RUBI: Main Contributions

- A paradigm shift from standard Bayesian inference, inference update can be achieved more efficiently by updating precursory planning rather than precursory inference.
- Four exact methods for updating inference using precursory planning under the assumption of consistent data association and Gaussian models
- Paradigm for incrementally updating inconsistent data association
- Comparing RUBI to current state of the art in both simulative and real-world data, considering the problem of autonomous navigation in unknown environments.

(Farhi17icra) (Farhi18ijrr conditionally accepted) (Farhi19icra workshop) (patent: WO2019171378A1)

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

- RUBI: Main Contributions
- Inference & BSP today
- Consistent DA assumption
  - Results
- RUBI
  - Results - simulation
  - Results - KITTI dataset

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

- RUBI: Main Contributions
- Inference & BSP today
- Consistent DA assumption
  - Results
- RUBI
  - Results - simulation
  - Results - KITTI dataset

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

JIP - Joint Inference & BSP

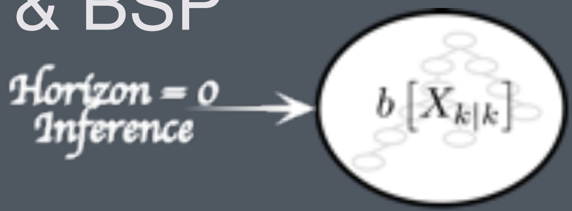

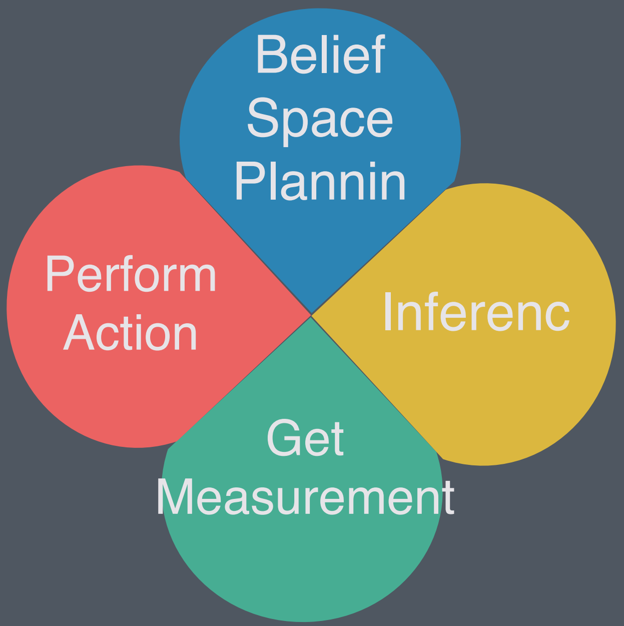

Legend:

(a) $\longrightarrow$ *Inference*

JIP - Joint Inference & BSP

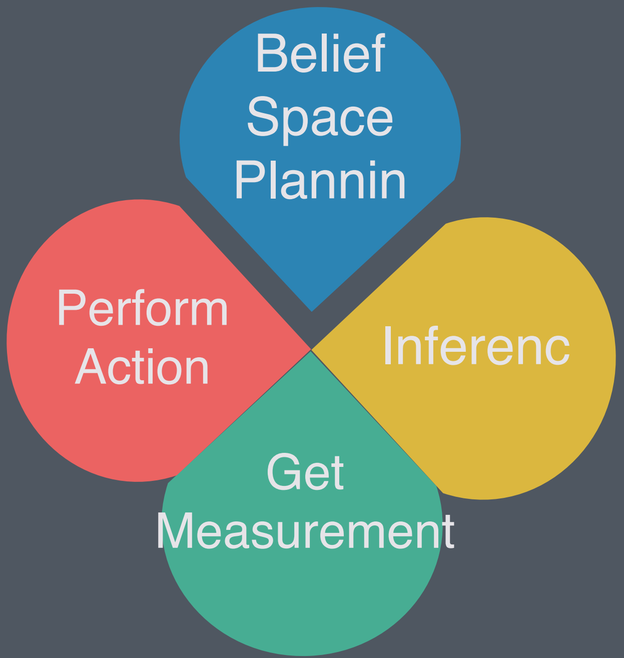

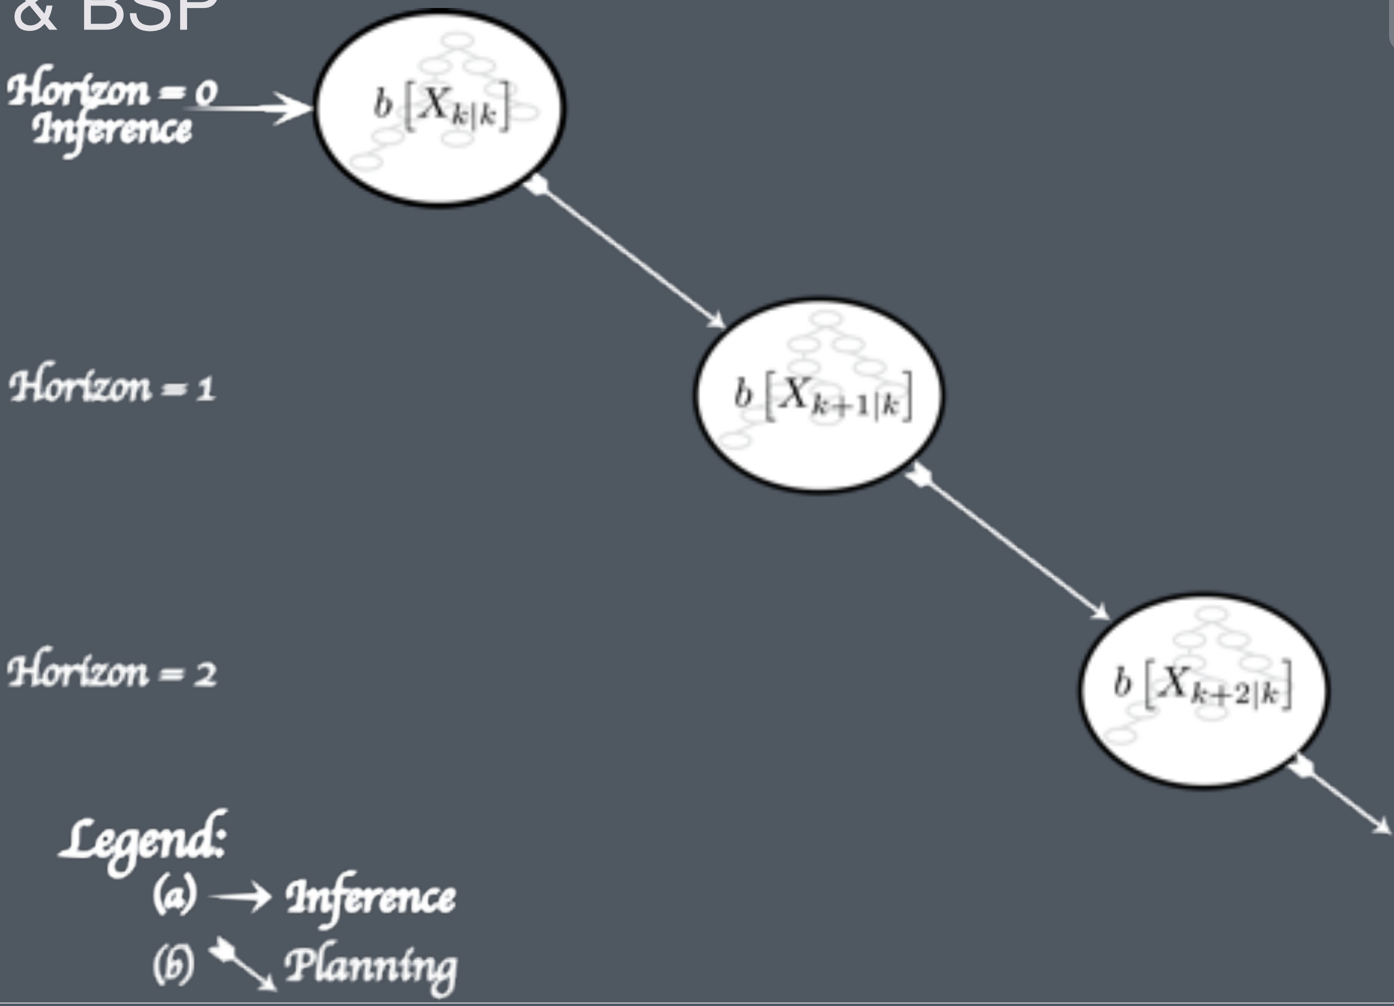

$\square_{t|k}$ - Referring to time t, while current time is k

Belief Space Planning Formulation Today

- BSP provides with the next (sub)optimal action(s), in reference to a Cost(Reward) function

$$u^{\star}_{k:k+L-1|k} = \underset{u_{k:k+L-1|k} \in \mathcal{U}_k}{argmin} \; J(u_{k:k+L-1|k})$$

$$J(u_{k:k+L-1|k})$$

Objective Value for horizon L

$\square_{t|k}$ - Referring to time t, while current time is k

Belief Space Planning Formulation Today

- BSP provides with the next (sub)optimal action(s), in reference to a Cost(Reward) function

$$u^{\star}_{k:k+L-1|k} = \underset{u_{k:k+L-1|k} \in \mathcal{U}_k}{argmin} \; J(u_{k:k+L-1|k})$$

$$J(u_{k:k+L-1|k}) \doteq \quad c_i\left(b[X_{i|k}], u_{i-1|k}\right)$$

Objective Value for horizon L

Future Belief

Future candidate action

Belief Space Planning Formulation Today

- BSP provides with the next (sub)optimal action(s), in reference to a Cost(Reward) function

$$u^{\star}_{k:k+L-1|k} = \underset{u_{k:k+L-1|k} \in \mathcal{U}_k}{argmin} \; J(u_{k:k+L-1|k})$$

$$J(u_{k:k+L-1|k}) \doteq \sum_{i=k+1}^{k+L} c_i \left(b[X_{i|k}], u_{i-1|k}\right)$$

Objective Value for horizon L

n

s

Future Belief

Future candidate action

$\square_{t|k}$ - Referring to time t, while current time is k

Belief Space Planning Formulation Today

- BSP provides with the next (sub)optimal action(s), in reference to a Cost(Reward) function

$$u^{\star}_{k:k+L-1|k} = \underset{u_{k:k+L-1|k} \in \mathcal{U}_k}{argmin} \; J(u_{k:k+L-1|k})$$

$$\underbrace{J(u_{k:k+L-1|k})}_{\text{Objective Value for horizon L}} \doteq \underbrace{\mathbb{E}_{z_{k+1:k+L|k}}}_{\text{Future measurements}} \left[\sum_{i=k+1}^{k+L} c_i \left(\underbrace{b[X_{i|k}]}_{\text{Future Belief}}, \underbrace{u_{i-1|k}}_{\text{Future candidate action}} \right) \right]$$

JIP - Joint Inference & BSP

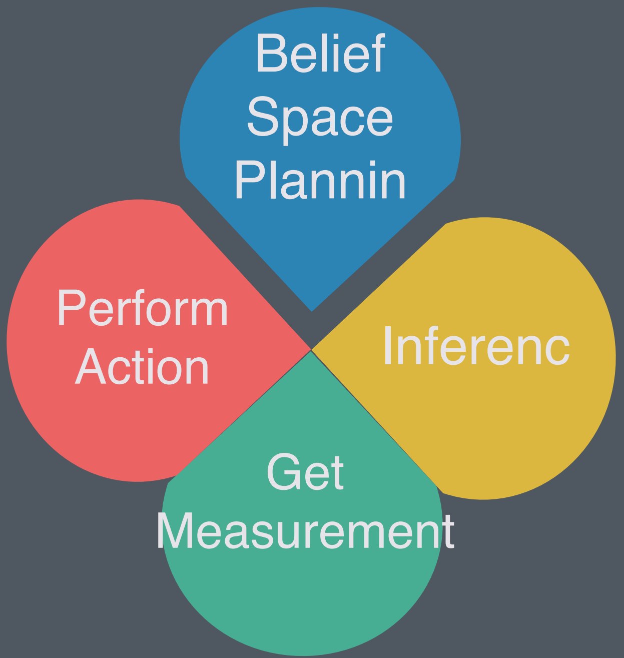

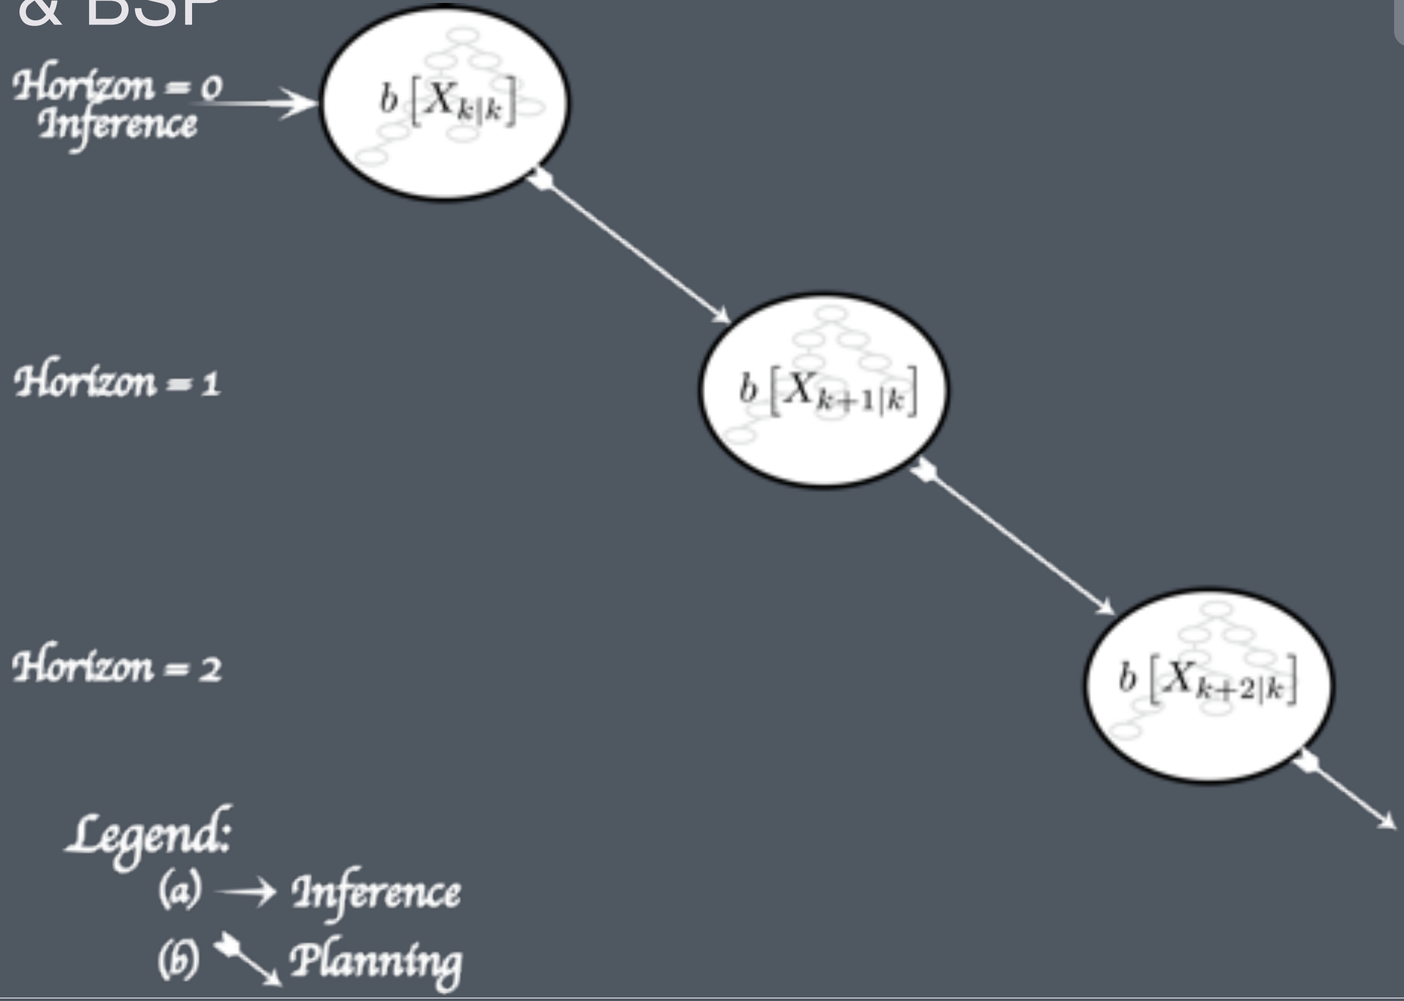

JIP - Joint Inference & BSP

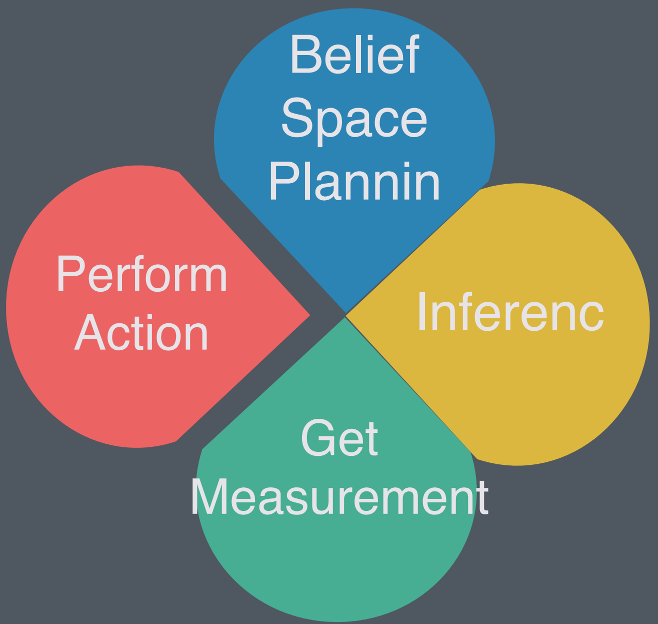

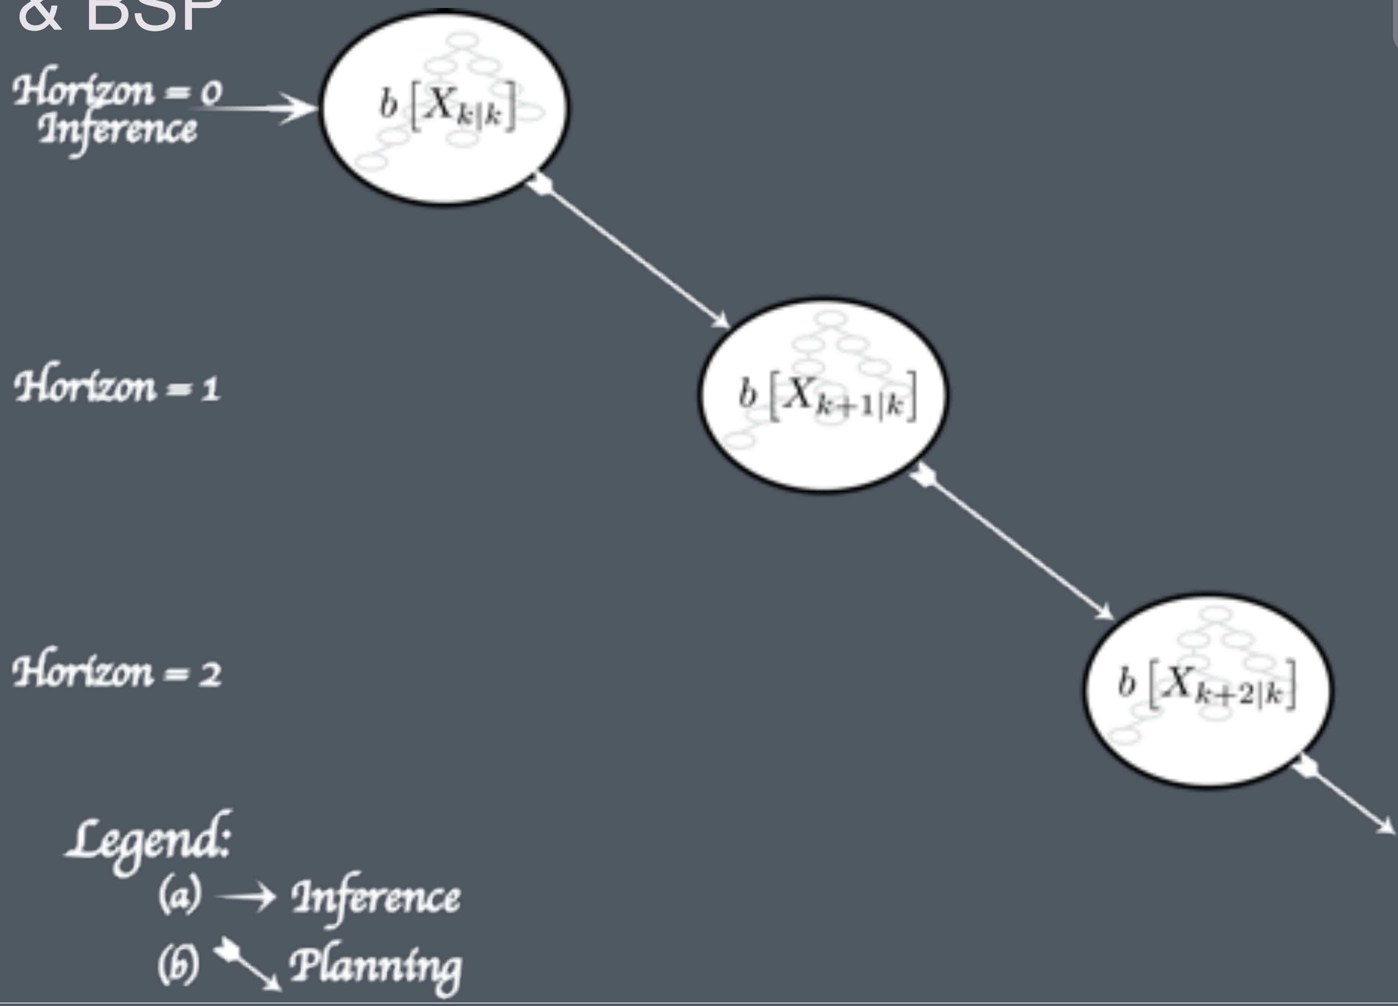

JIP - Joint Inference & BSP

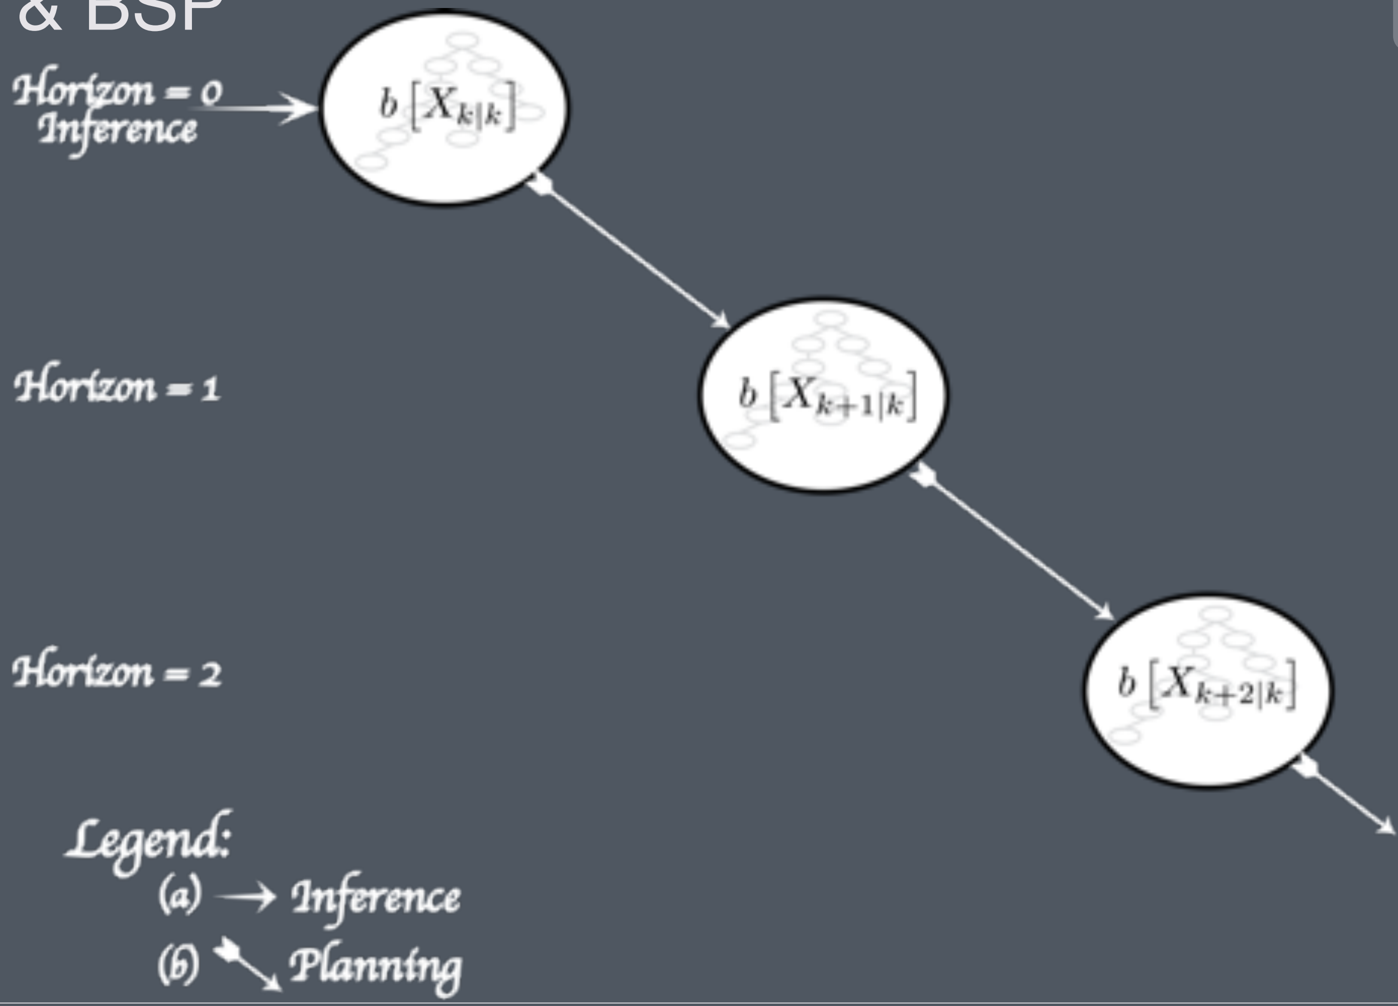

JIP - Joint Inference & BSP

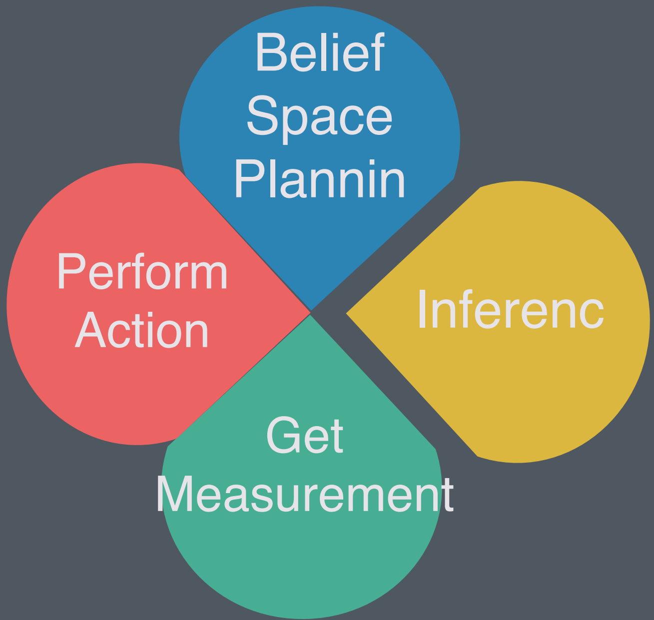

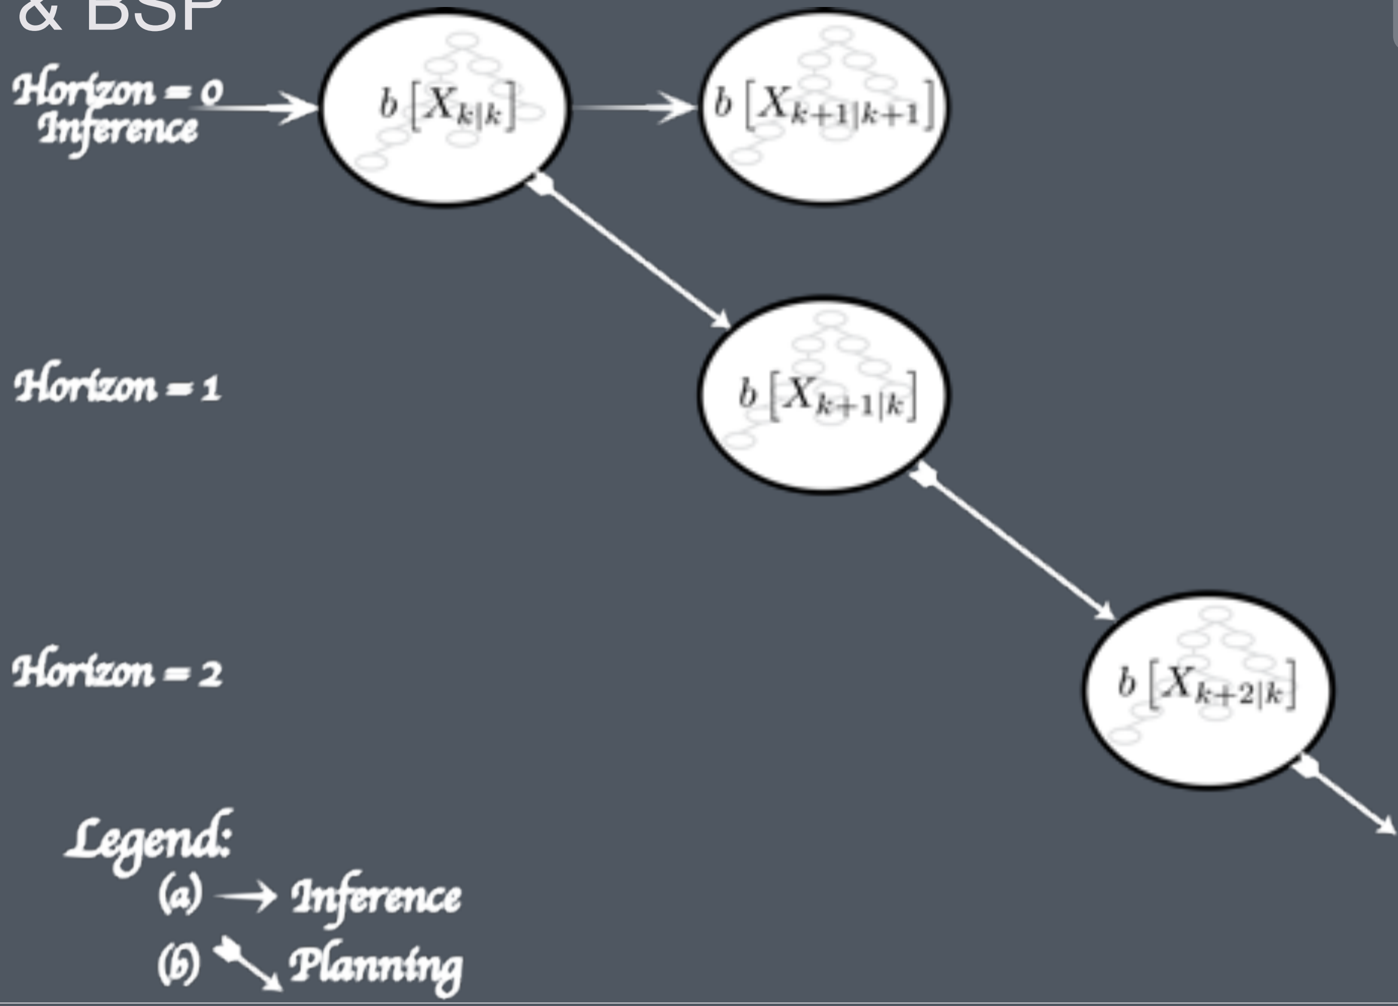

Inference Formulation Today

- Inference provides an estimation for the joint state

$$b[X_{k+1|k+1}] \propto \underbrace{p\left(X_0\right)}_{\text{Prior}} \prod_{i=1}^{k+1} \left[\underbrace{p\left(x_i | x_{i-1}, u_{i-1|k+1}\right)}_{\text{Motion Model}} \underbrace{\prod_{j \in \mathcal{M}_{i|k+1}}}_{\text{Data Association}} \underbrace{p\left(z^j_{i|k+1} | x_i, l_j\right)}_{\text{Measurement Model}} \right]$$

- For example, maximum a-posteriori (MAP) estimation

$$X^\star_{k+1|k+1} = \underset{X_{k+1}}{argmax}\ b[X_{k+1|k+1}]$$

$$\text{NLS} \Downarrow$$

$$A_{k+1|k+1} \cdot \triangle X_{k+1} = b_{k+1|k+1} \xRightarrow{Factorization} R_{k+1|k+1} \cdot \triangle X_{k+1} = d_{k+1|k+1}$$

$\mathcal{M}_{i|k+1}$

All landmark indices associated to measurements from time i, while current

Inference vs. Planning

- Inference and precursory planning (of the same action) differ in measurements and DA

$$\underbrace{b[X_{k+1|k}]}_{\text{Planning}} \propto \underbrace{p\left(X_0\right)}_{\text{Prior}} \prod_{i=1}^{k+1} \left[\underbrace{p\left(x_i | x_{i-1}, u_{i-1|k}\right)}_{\text{Motion Model}} \prod_{\underbrace{j \in \mathcal{M}_{i|k}}_{\text{Data Association}}} \underbrace{p\left(z_{i|k}^{j} | x_i, l_j \right)}_{\text{Measurement Model}} \right]$$

$$\underbrace{b[X_{k+1|k+1}]}_{\text{Inference}} \propto \underbrace{p\left(X_0\right)}_{\text{Prior}} \prod_{i=1}^{k+1} \left[\underbrace{p\left(x_i | x_{i-1}, u_{i-1|k+1}\right)}_{\text{Motion Model}} \prod_{\underbrace{j \in \mathcal{M}_{i|k+1}}_{\text{Data Association}}} \underbrace{p\left(z_{i|k+1}^{j} | x_i, l_j \right)}_{\text{Measurement Model}} \right]$$

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

RUBI: Main Contributions

Inference & BSP today

Consistent DA assumption

- Results

RUBI

- Results - simulation
- Results - KITTI dataset

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

RUBI: Main Contributions

Inference & BSP today

Consistent DA assumption

Results

RUBI

Results - simulation

Results - KITTI dataset

Consistent Data Association (DA) assumption

- For consistent DA

$$\mathcal{M}_{k+1|k} \equiv \mathcal{M}_{k+1|k+1}$$

DA from Planning

DA from Inference

$$\Downarrow$$

$$R_{k+1|k} \equiv R_{k+1|k+1}$$

- Hence in order to solve the inference problem (provide with a state estimation) we are left with updating the RHS vector

$$d_{k+1|k} \Rightarrow d_{k+1|k+1}$$

Inference Update

- We devised four different methods for updating the RHS vector

OTM	DU	OTM-OO	DU-OO
Orthogonal Transformation Matrix	Down-date Update	OTM - Only Observations	DU - Only Observations

The OTM & DU Methods

OTM

Orthogonal Transformation Matrix

DU

Down-date Update

$$d_{k+1|k+1} = Q^T_{k+1|k} \begin{bmatrix} d_{k|k} \\ \breve{b}_{k+1} \end{bmatrix}$$

$$d_{k+1|k} = (R^T_{k|k})^{-1}(R^T_{k+1|k} d_{k+1|k} - A^T_{meas} b_{meas})$$

$$R_{k+1|k+1} = R_{k+1|k}$$

$$R_{k+1|k+1} = R_{k+1|k}$$

The Only Observations (OO) Addition

OTM-OO

OTM - Only Observations

$$d_{k+1|k+1} = Q^T_{k+1|k} \begin{bmatrix} d^{Motion}_{k+1|k} \\ \breve{b}^{measurement}_{k+1} \end{bmatrix}$$

$$R_{k+1|k+1} = R_{k+1|k}$$

DU-OO

DU - Only Observations

$$R^{w/oT}_{k+1|k} R^{w/o}_{k+1|k} = R^T_{k+1|k} R^T_{k+1|k} - A^T_{meas} A_{meas}$$

$$d^{W/Omeasure}_{k+1|k} = (R^{w/oT}_{k+1|k})^{-1} (R^T_{k+1|k} d_{k+1|k} - A^T_{meas} b_{meas})$$

$$d_{k+1|k+1} = (R^T_{k+1|k})^{-1} (R^T_{k|k} d_{k|k} + A^T_{real-meas} b_{real-meas})$$

$$R_{k+1|k+1} = R_{k+1|k}$$

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

RUBI: Main Contributions

Inference & BSP today

Consistent DA assumption

- Results

RUBI

- Results - simulation
- Results - KITTI dataset

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

RUBI: Main Contributions

Inference & BSP today

Consistent DA assumption

Results

RUBI

Results - simulation

Results - KITTI dataset

Results - Putting JIP to the First Test

- We performed continuous BSP (POMDP case) in an unknown synthetic environment.
- Our four methods, coded in MATLAB, were compared to:
  - inference update using Standard batch approach - STD
  - inference update using iSAM2 efficient methodology (using C++ wrapper) - iSAM
- Robot was required to visit five targets whilst not crossing a covariance threshold.
- We considered known models with Gaussian additive noise and consistent DA

Our method produces an identical belief to the one received via iSAM,
hence only computation time would be compared and discussed.

The Map and Inference Update Total time

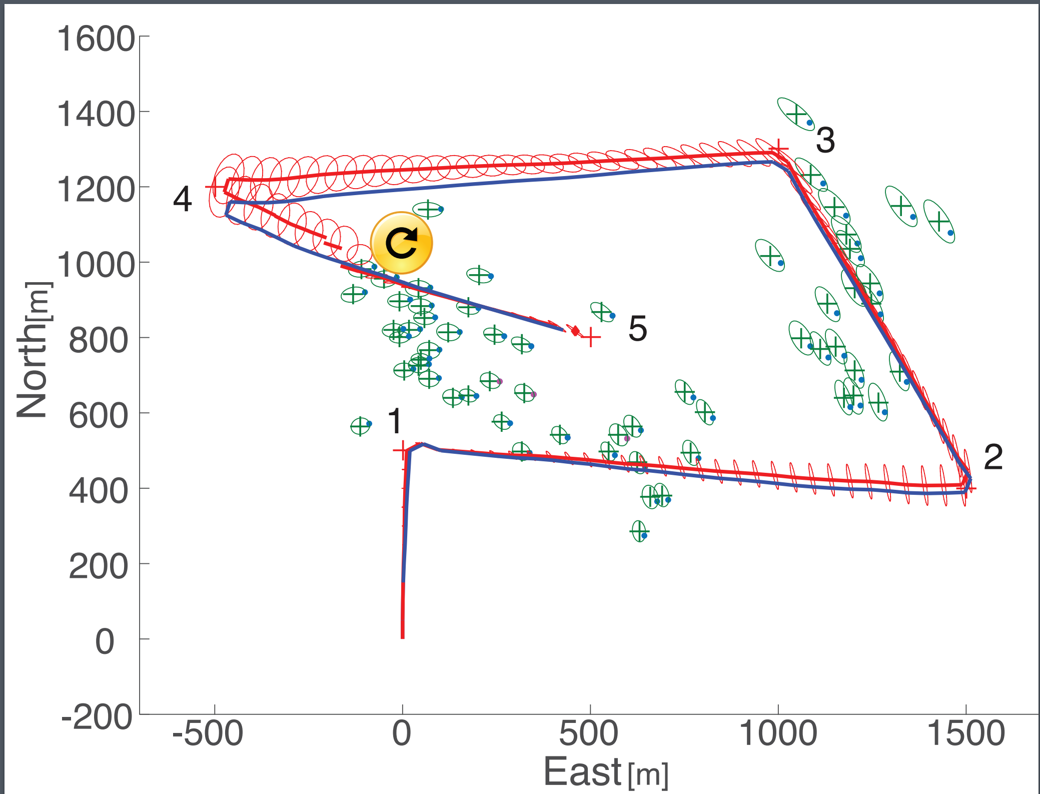

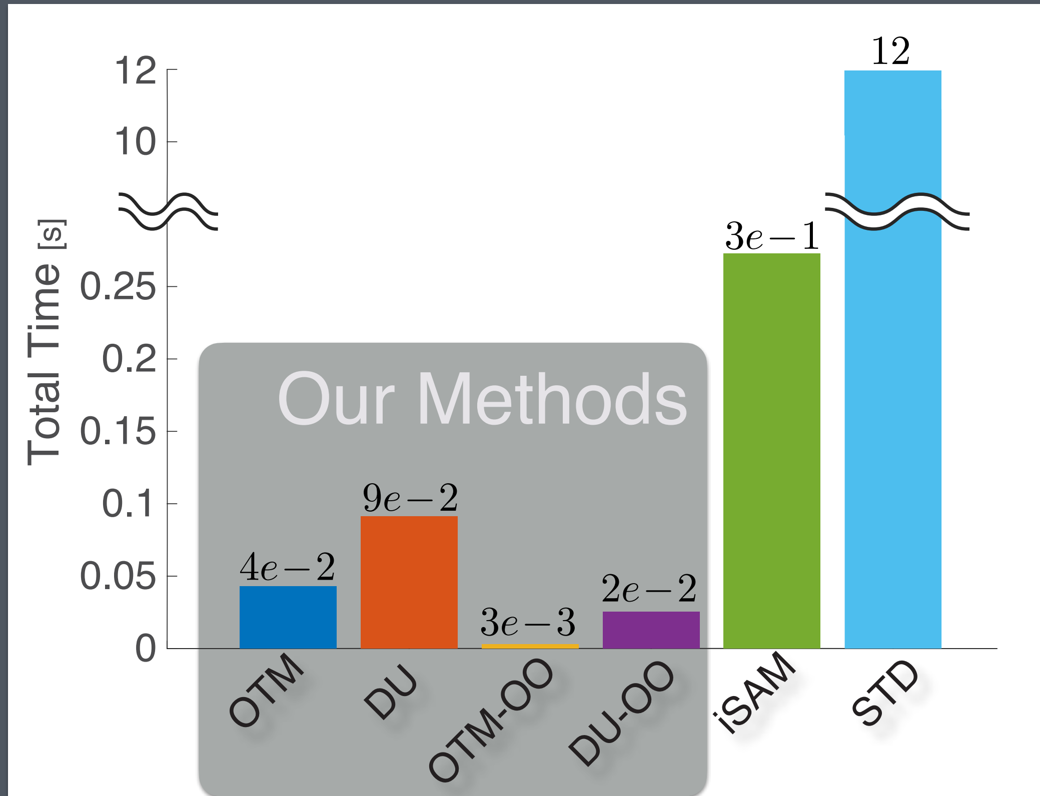

Performance Per-step

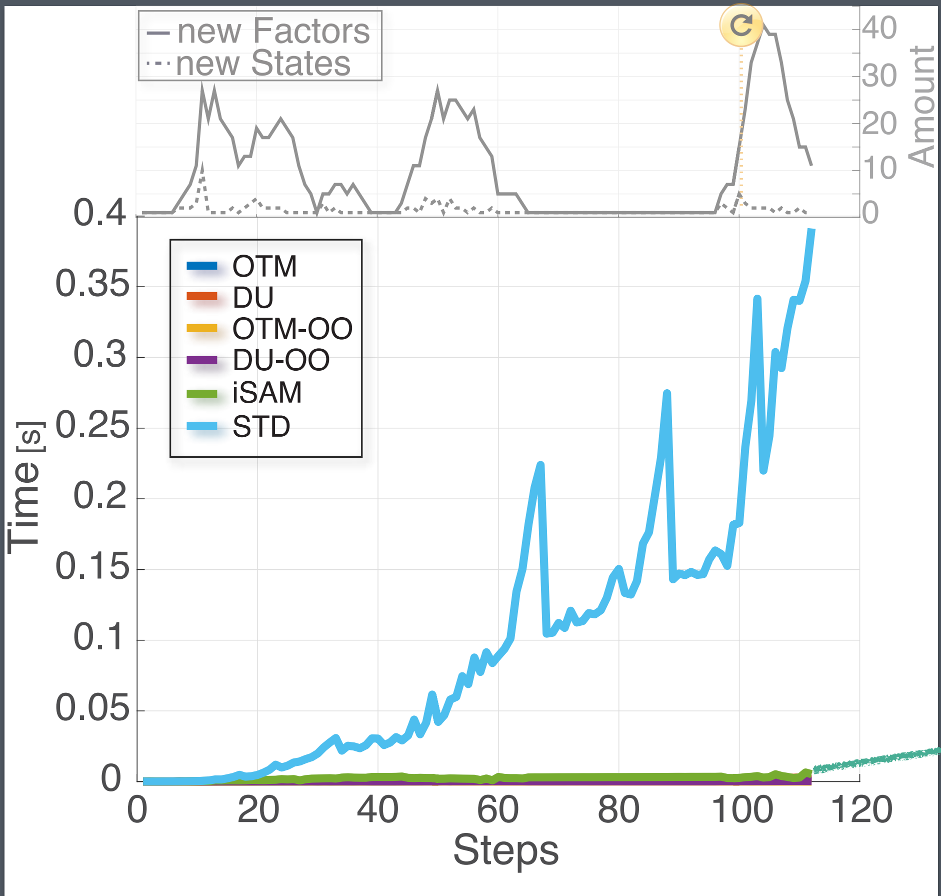

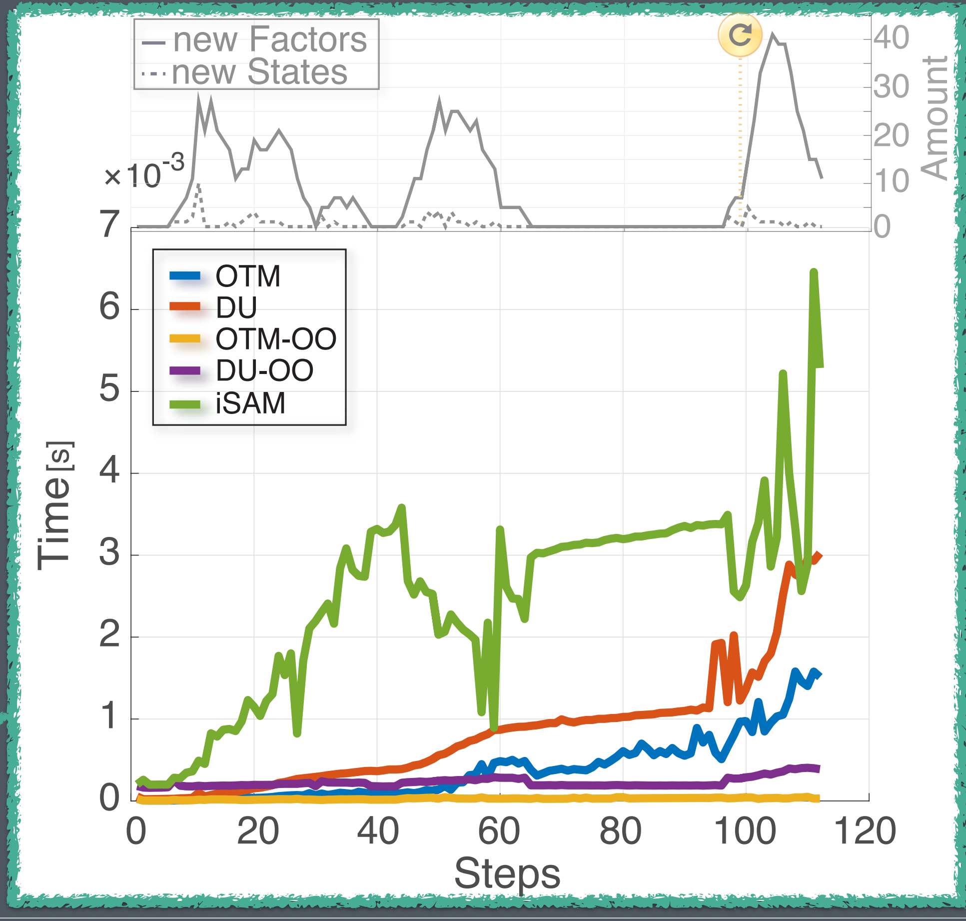

Robustness

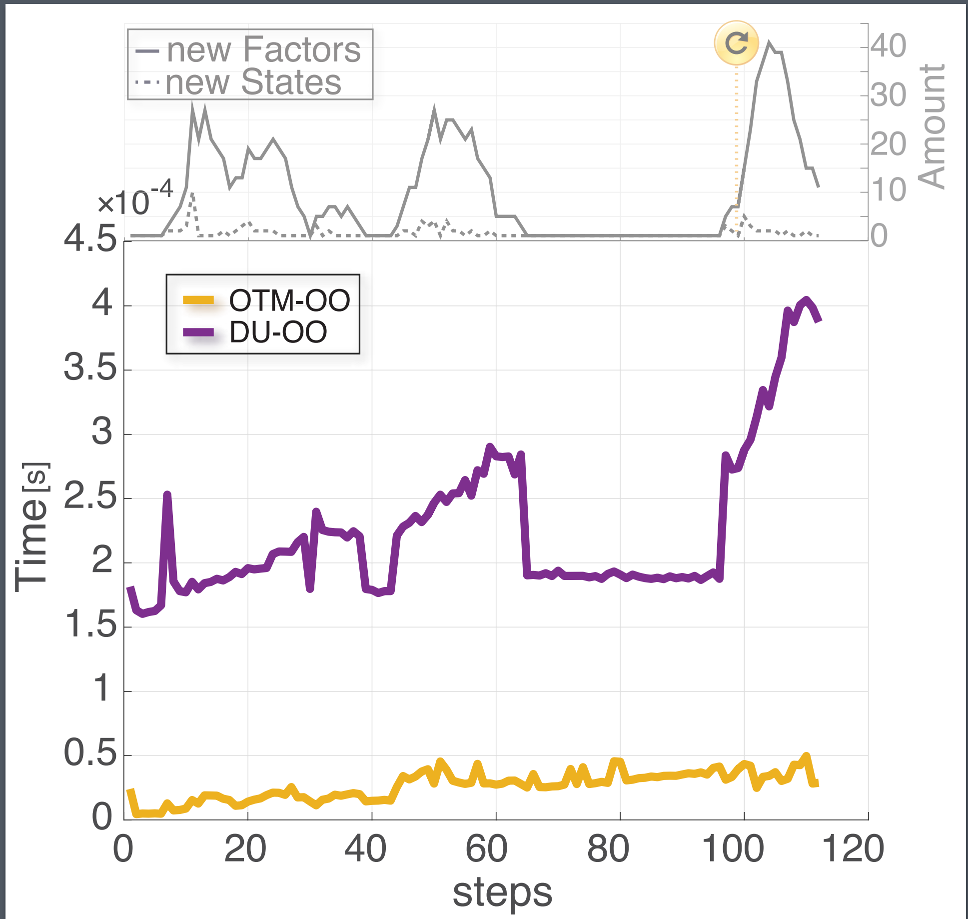
new Factors
new States
Amount
40
30
20
10
0
×10^-4
4.5
4
3.5
3
2.5
2
1.5
1
0.5
0
Time [s]
OTM-OO
DU-OO
0
20
40
60
80
100
120
steps

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

- RUBI: Main Contributions
- Inference & BSP today
- Consistent DA assumption
  - Results
- RUBI
  - Results - simulation
  - Results - KITTI dataset

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

RUBI: Main Contributions

Inference & BSP today

Consistent DA assumption

Results

RUBI

Results - simulation

Results - KITTI dataset

Relaxing the consistent DA assumption

- Accounting for data association inconsistency between inference and planning
- Once the DA inconsistency is dealt with, we revert to the previously presented solution - updating measurements
- The data association is corrected using QR update (existing equivalent graphical models)
- Thanks to QR update, not all variables are necessarily affected from correcting DA inconsistency

Inconsistent Data Association

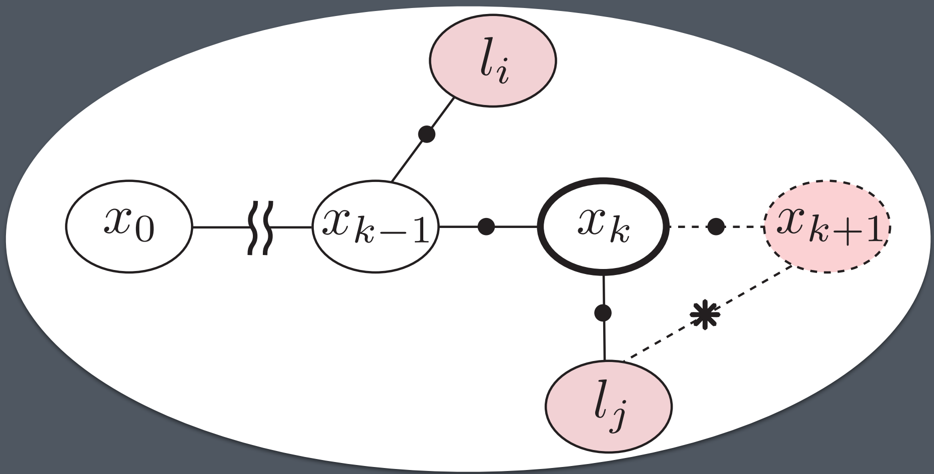

$b[X_{k+1|k}]$

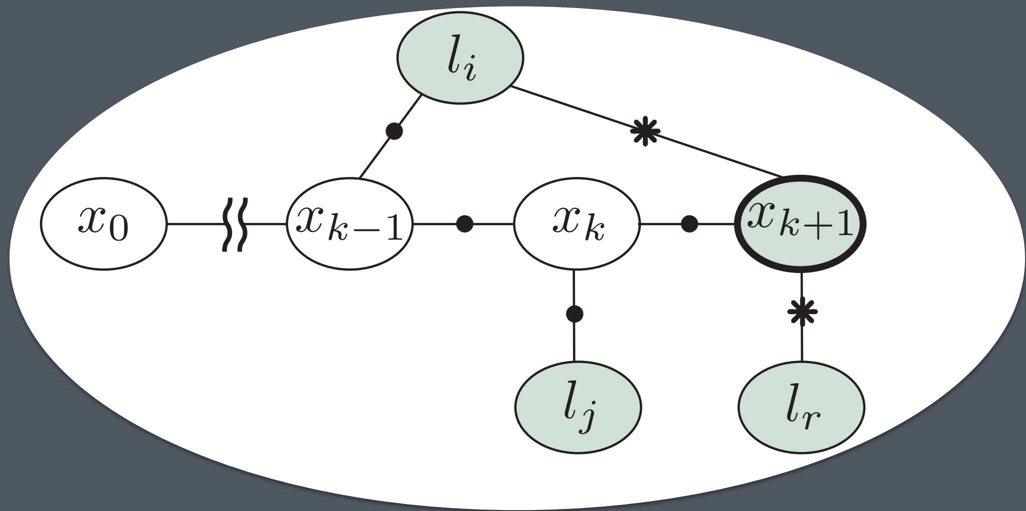

$b[X_{k+1|k+1}]$

Belief Graphical Representations

- (a) - Factor Graph
- (b) - Jacobian and RHS vector
- (c) - Bayes Tree
- (d) - Jacobian QR decomposition

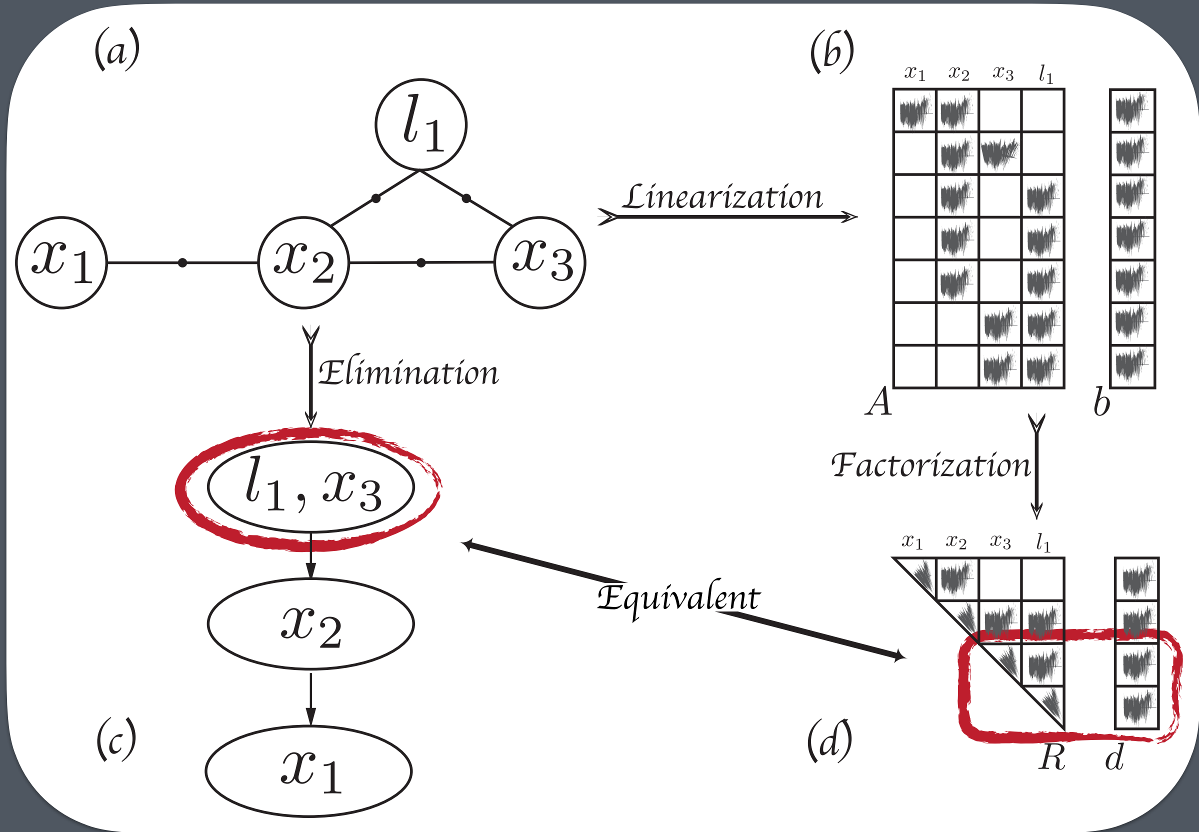

Correcting inconsistent DA

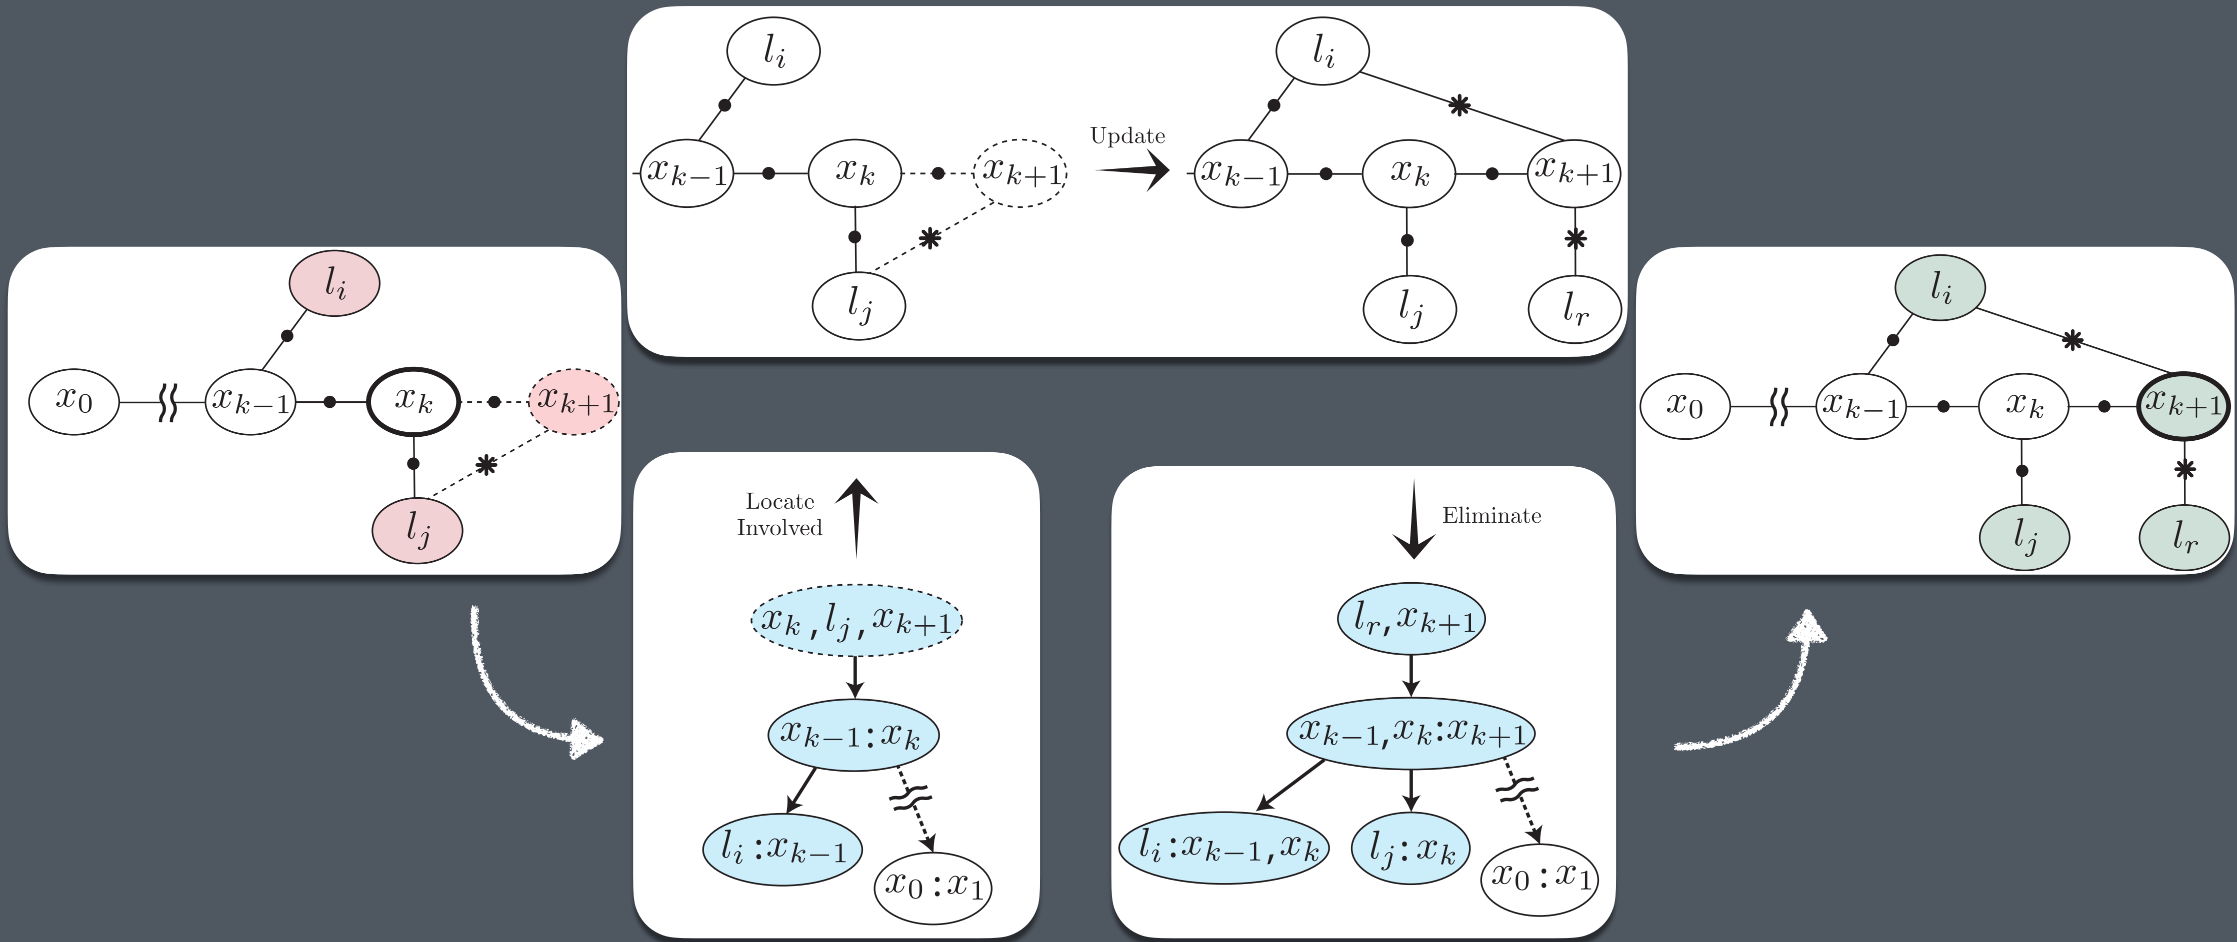

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

- RUBI: Main Contributions
- Inference & BSP today
- Consistent DA assumption
  - Results
- RUBI
  - Results - simulation
  - Results - KITTI dataset

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

RUBI: Main Contributions

Inference & BSP today

Consistent DA assumption

Results

RUBI

Results - simulation

Results - KITTI dataset

Results - Putting RUBI to the Test

- We performed continuous BSP (POMDP case) in an unknown synthetic environment.
- For inference update we use UD-OTM-OO, denoting a method which updates DA and update RHS vector using OTM-OO.
- inference update using iSAM2 efficient methodology (using C++ wrapper) - iSAM
- Robot was required to visit twelve targets whilst not crossing a covariance threshold.
- We considered known models with Gaussian additive noise

Our method produces an identical belief to the one received via iSAM,
hence only computation time would be compared and discussed.

The Map and Inference Update Total time

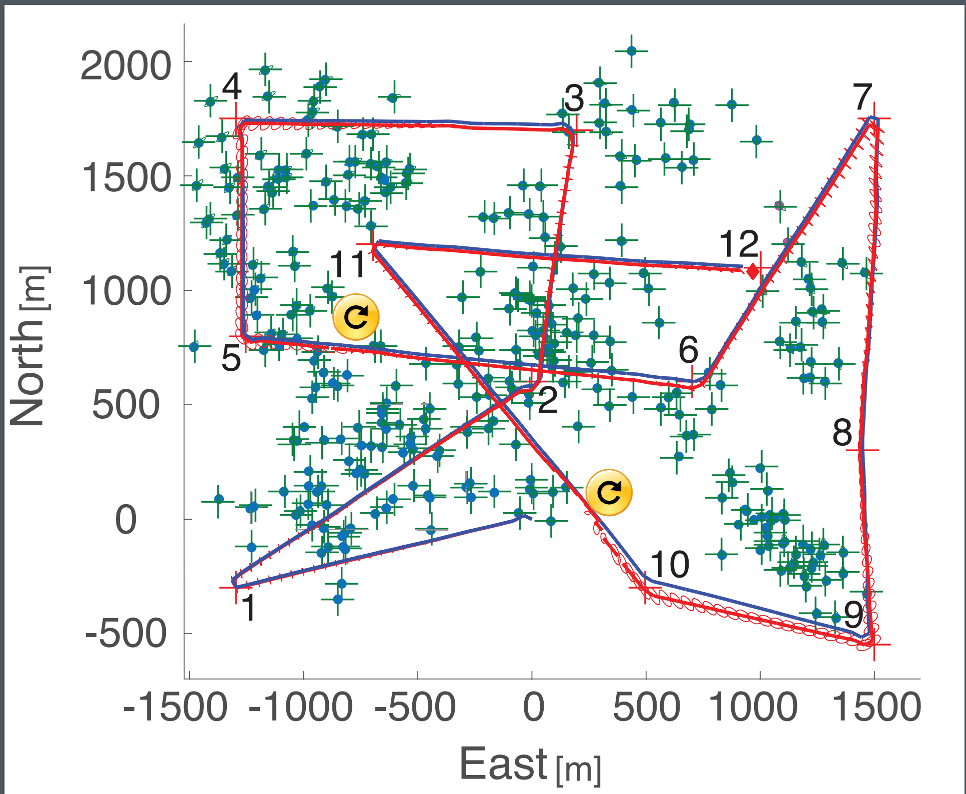

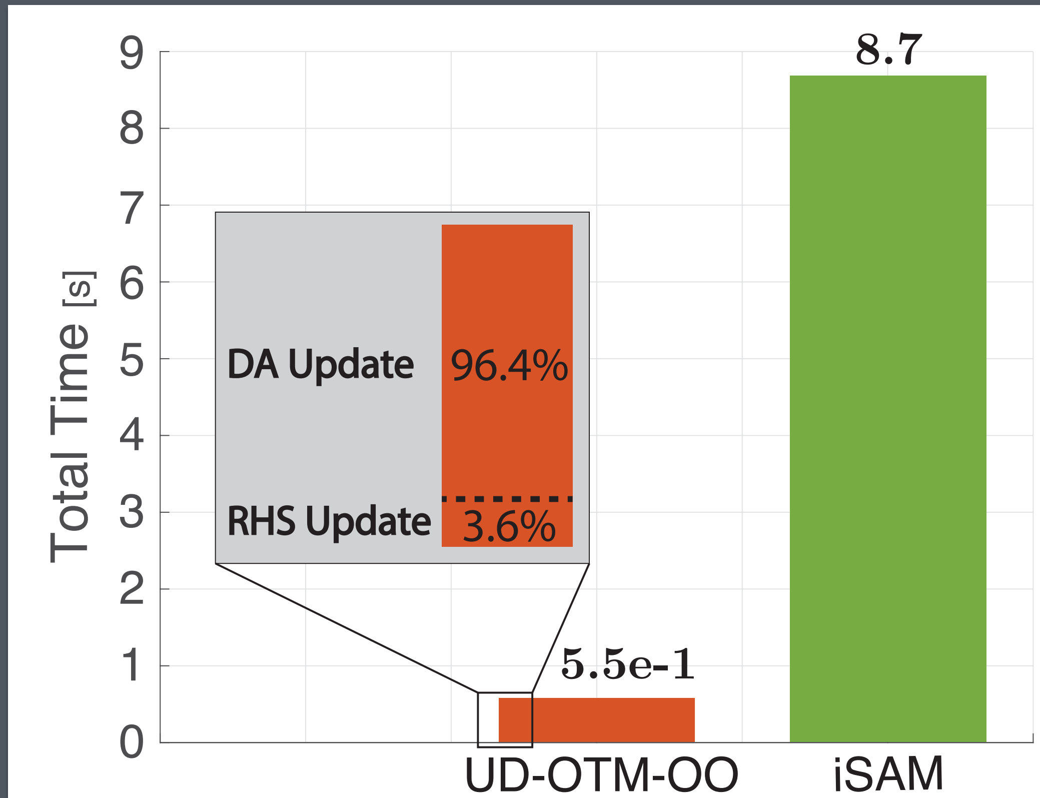

Performance Per-step - Inconsistent DA of 50%

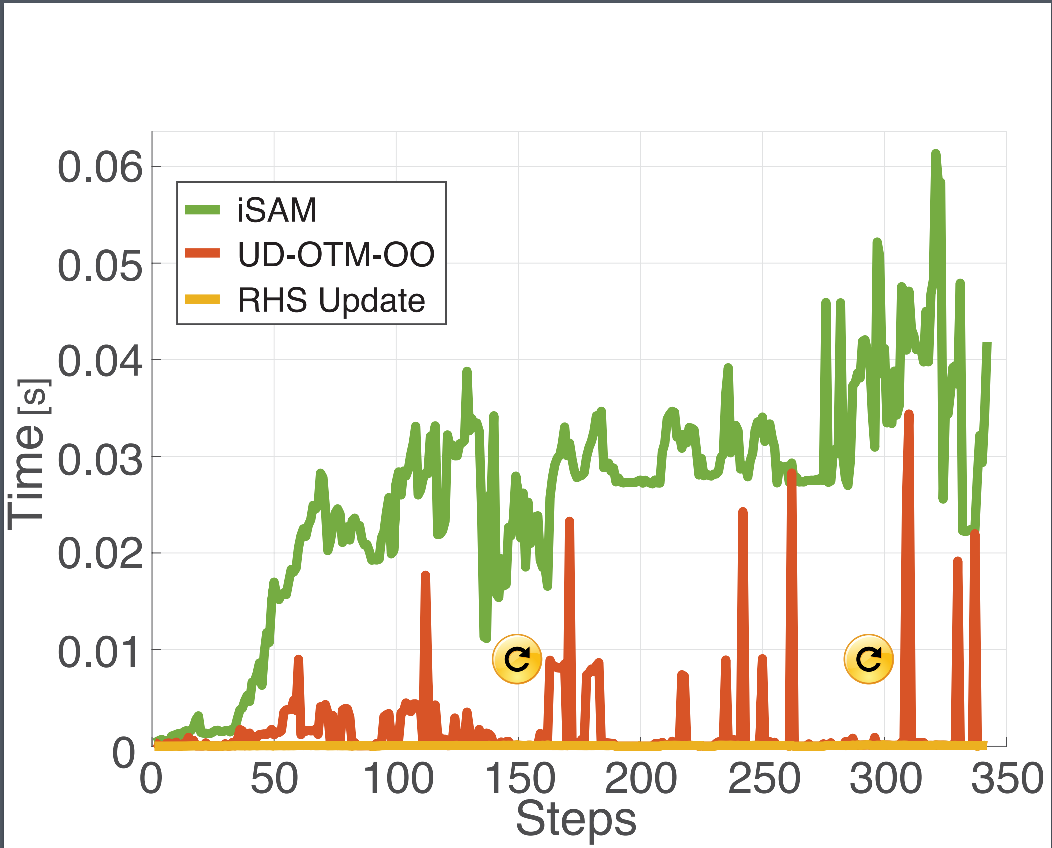

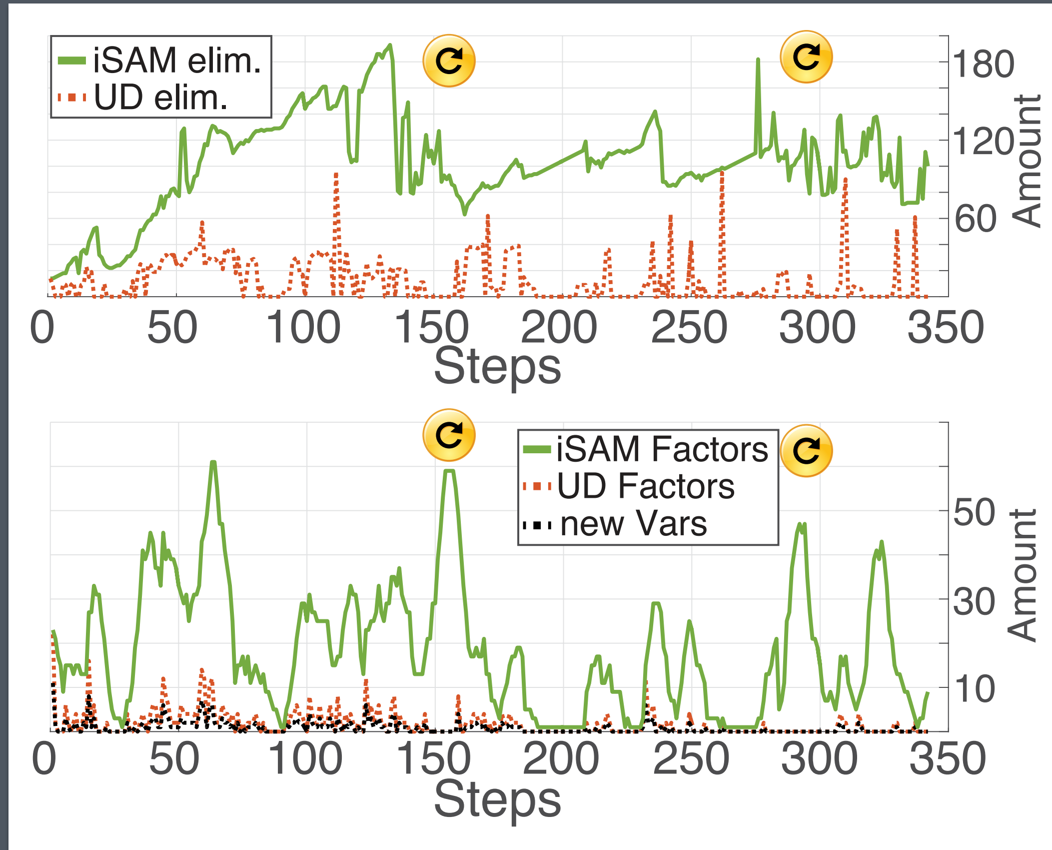

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

RUBI: Main Contributions

Inference & BSP today

Consistent DA assumption

Results

RUBI

Results - simulation

Results - KITTI dataset

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

- RUBI: Main Contributions
- Inference & BSP today
- Consistent DA assumption
  - Results
- RUBI
  - Results - simulation
  - Results - KITTI dataset

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

Results - Putting RUBI to a real world Test

- We used the well known KITTI dataset to compare UD-OTM-OO to iSAM
- KITTI is a passive SLAM dataset, so before each inference session we performed a planning session over the “optimal” action sequence.
- We used only the monocular stream as an input from the KITTI dataset.
- The robot started from an un-informative prior over its initial pose, with no prior knowledge over the environment.

Here, on real-world data, we would also compare the estimation difference between iSAM and UD-OTM-OO, although they are algebraically identical.

The Map and Inference Update Total time

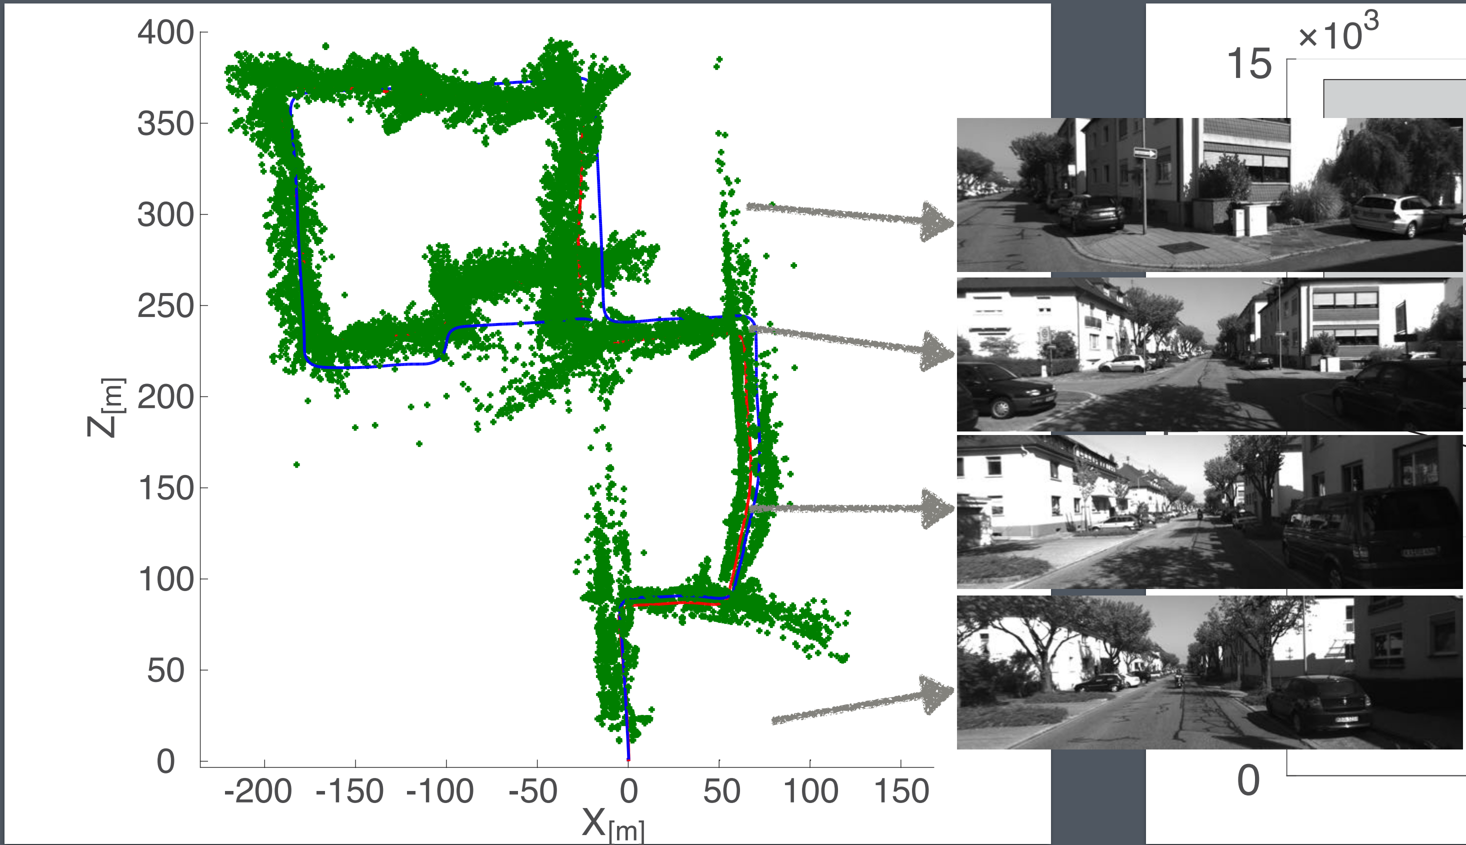

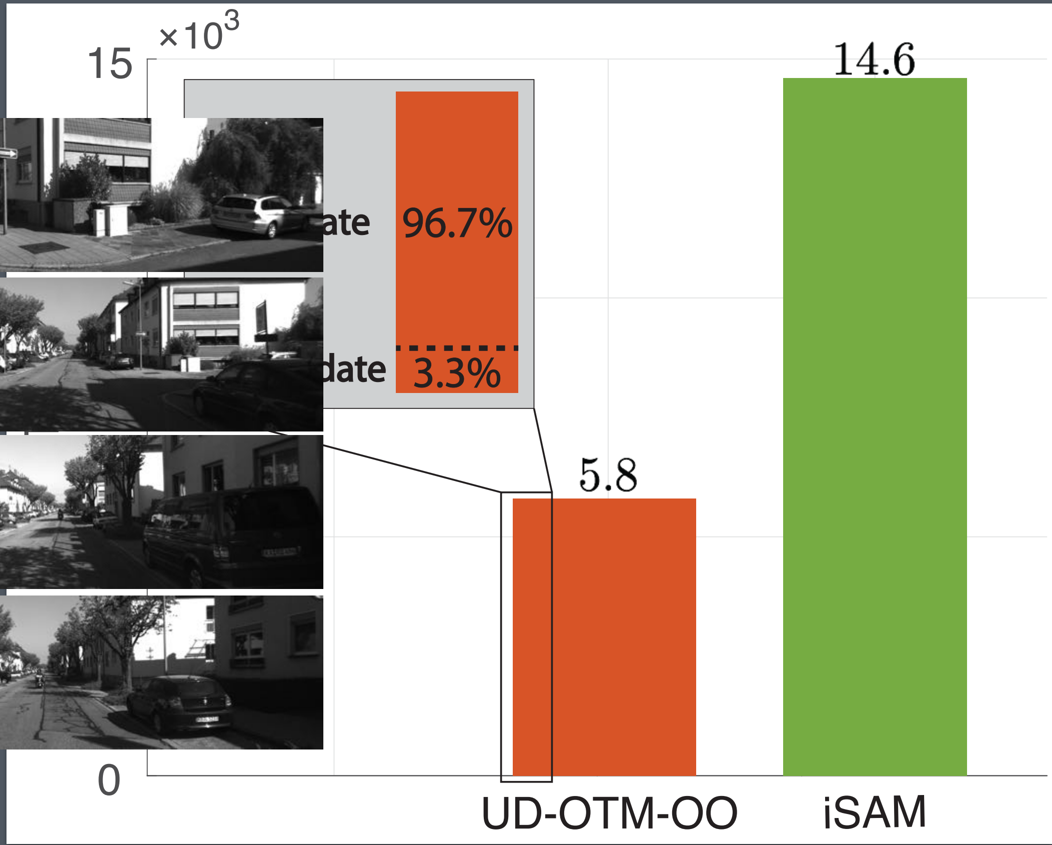

Performance Per-step

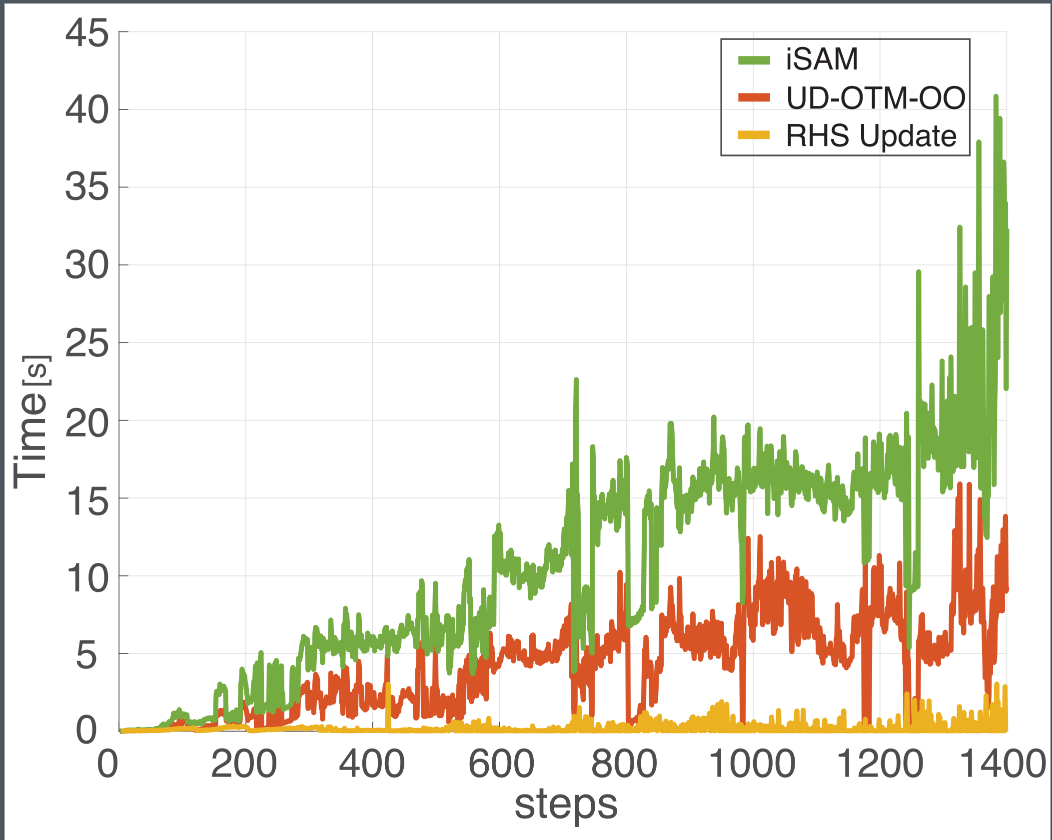

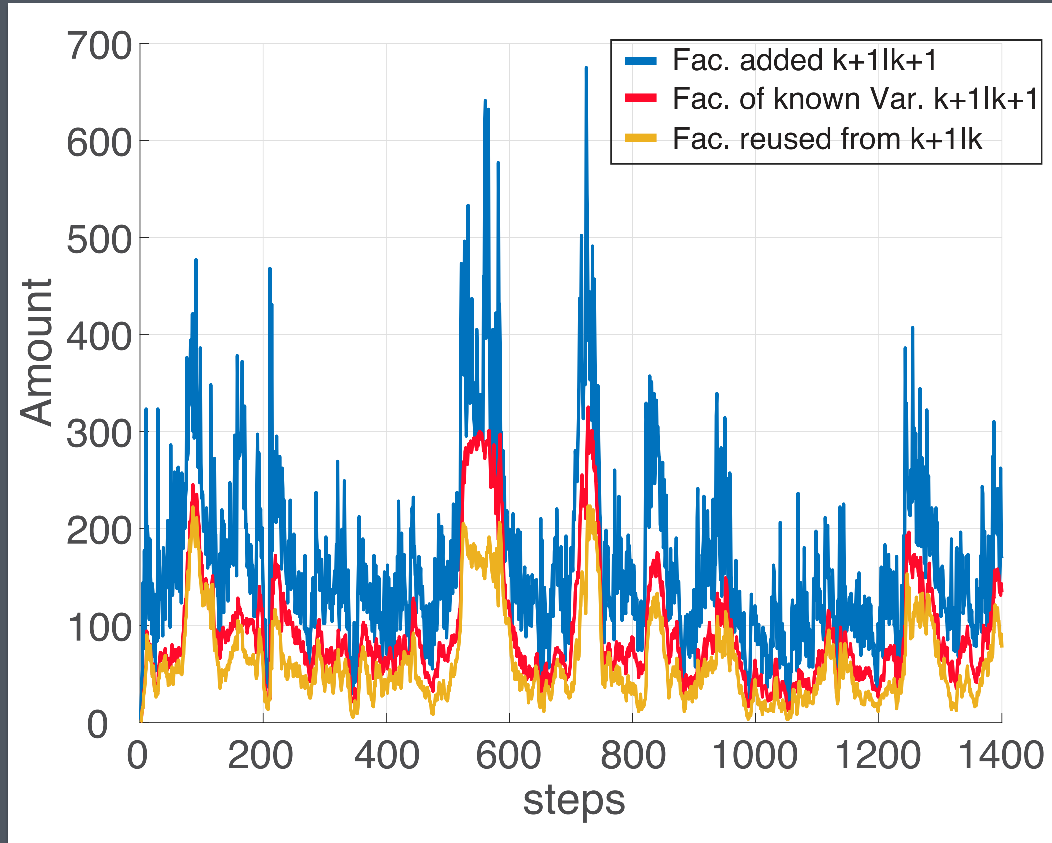

Translation estimation error - RUBI vs iSAM

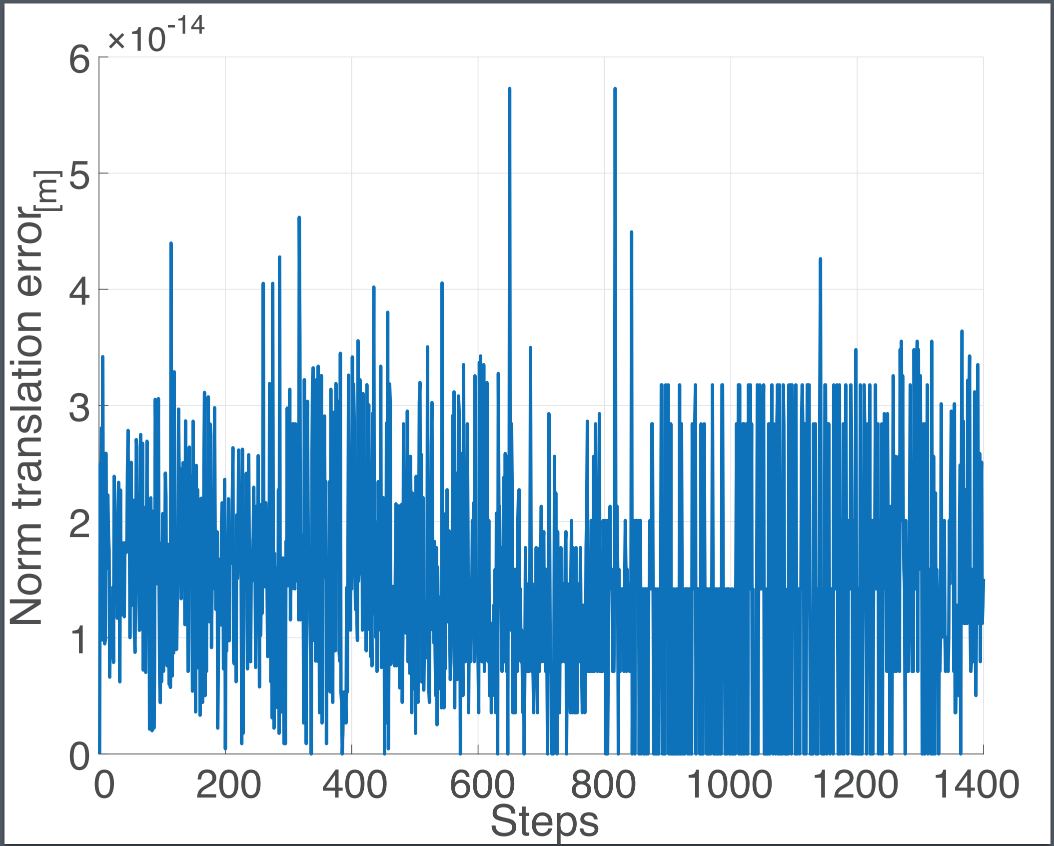

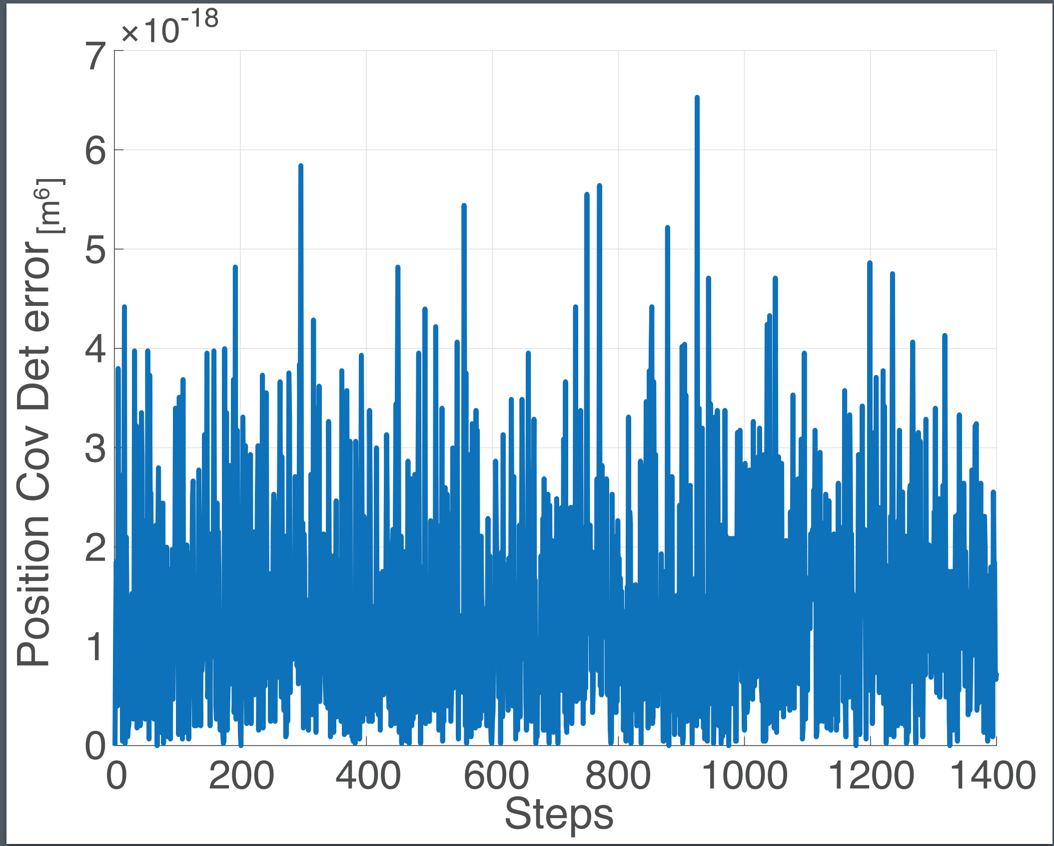

Orientation estimation error - RUBI vs iSAM

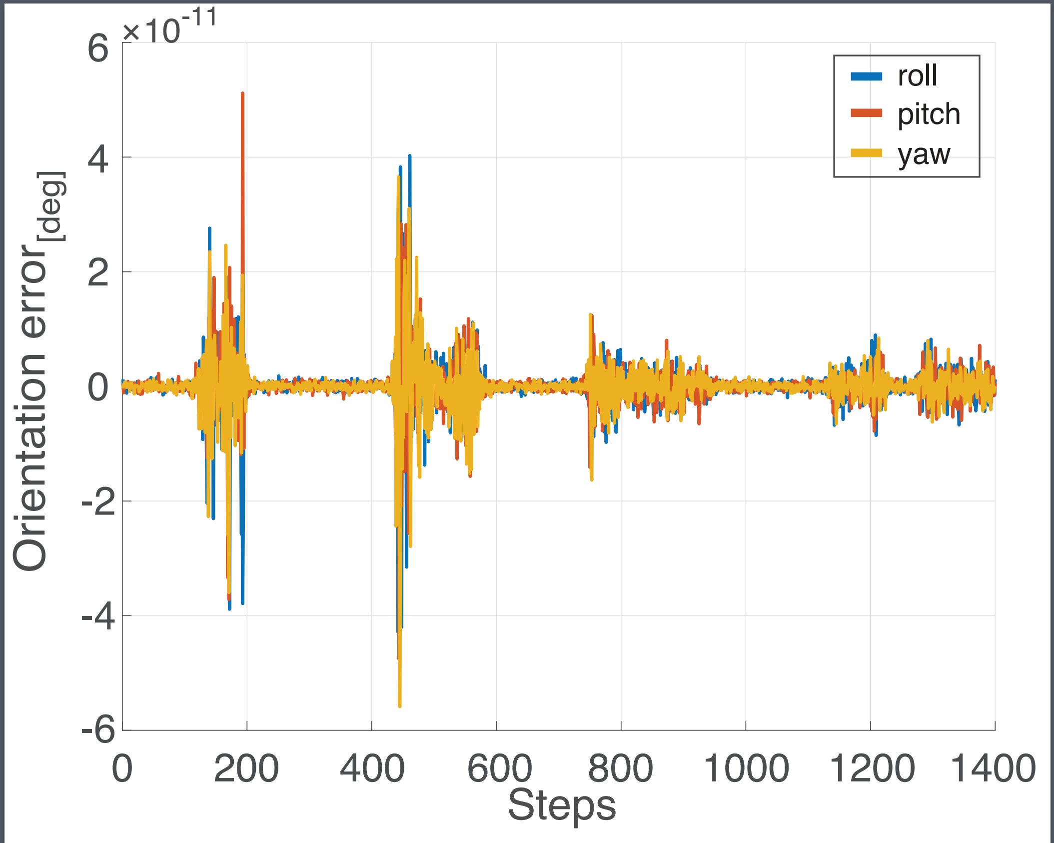

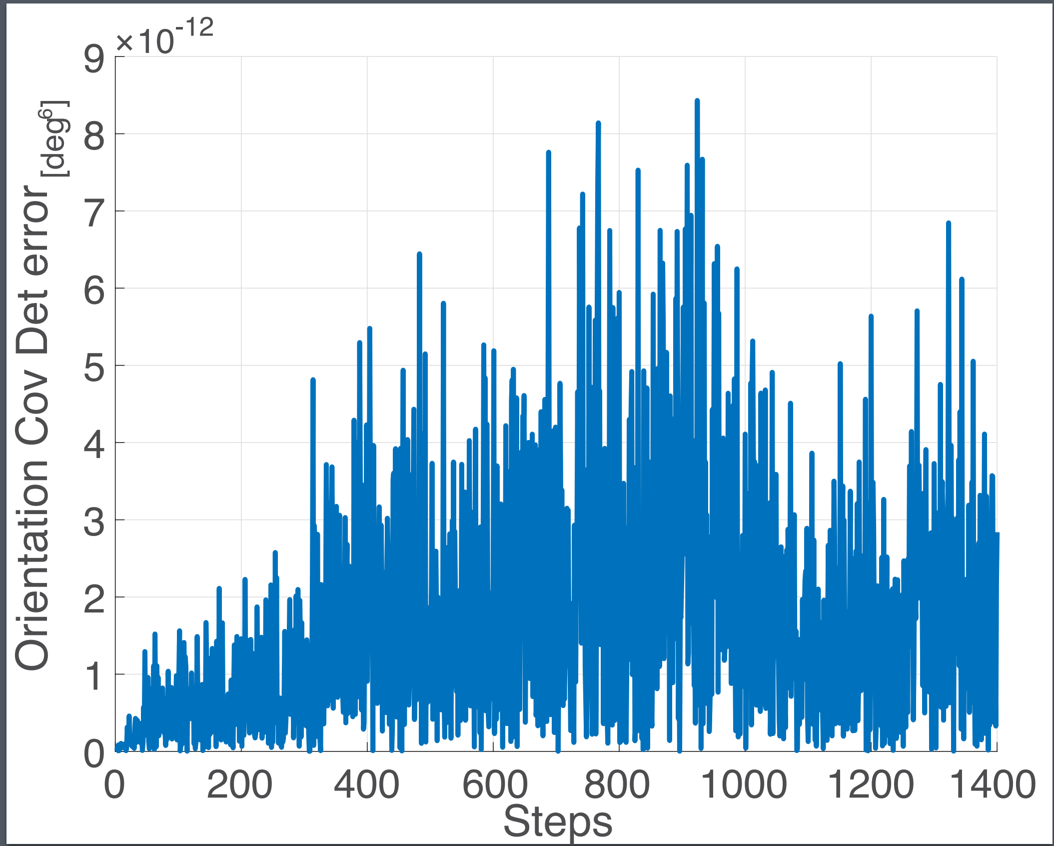

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

- RUBI: Main Contributions
- Inference & BSP today
- Consistent DA assumption
  - Results
- RUBI
  - Results - simulation
  - Results - KITTI dataset

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

- BSP formulation
- Related work
- iX-BSP
  - Results
- The Wildfire Assumption
  - Results
- iML-BSP
  - Results - simulation
  - Results - live

Belief Space Planning Formulation

- BSP provides with the next (sub)optimal action(s), in reference to a Cost(Reward) function

$$u^{\star}_{k:k+L-1|k} = \underset{u_{k:k+L-1|k} \in \mathcal{U}_k}{argmin} \; J(u_{k:k+L-1|k})$$

$$\underbrace{J(u_{k:k+L-1|k})}_{\text{Objective Value for horizon L}} \doteq \underbrace{\underset{z_{k+1:k+L|k}}{\mathbb{E}}}_{\text{Future measurements}} \left[\sum_{i=k+1}^{k+L} c_i \left(\underbrace{b[X_{i|k}]}_{\text{Future Belief}}, \underbrace{u_{i-1|k}}_{\text{Future candidate action}} \right) \right]$$

$\square_{t|k}$ - Referring to time t, while current time is k

Belief Space Planning Formulation

- BSP provides with the next (sub)optimal action(s), in reference to a Cost(Reward) function

$$u^\star_{k:k+L-1|k} = \underset{u_{k:k+L-1|k} \in \mathcal{U}_k}{argmin} \; J(u_{k:k+L-1|k})$$

Under Maximum Likelihood (ML) Assumption

$$\underbrace{J(u_{k:k+L-1|k})}_{\text{Objective Value for horizon L}} \doteq \sum_{i=k+1}^{k+L} c_i\left(\underbrace{b[X_{i|k}]}_{\text{Future Belief}}, \underbrace{u_{i-1|k}}_{\text{Future candidate action}}\right)$$

$\square_{t|k}$ - Referring to time t, while current time is k

Belief Space Planning Formulation

- BSP provides with the next (sub)optimal action(s), in reference to a Cost(Reward) function

$$u^{\star}_{k:k+L-1|k} = \underset{u_{k:k+L-1|k} \in \mathcal{U}_k}{argmin} \; J(u_{k:k+L-1|k})$$

$$\underbrace{J(u_{k:k+L-1|k})}_{\text{Objective Value for horizon L}} \doteq \underset{\underbrace{z_{k+1:k+L|k}}_{\text{Future measurements}}}{\mathbb{E}} \left[\sum_{i=k+1}^{k+L} c_i \left(\underbrace{b[X_{i|k}]}_{\text{Future Belief}}, \underbrace{u_{i-1|k}}_{\text{Future candidate action}} \right) \right]$$

$\square_{t|k}$ - Referring to time t, while current time is k

Belief Space Planning Formulation

$$J(u) = \int_{z_{k+1|k}} \mathbb{P}(z_{k+1|k}|H^{-}_{k+1|k}) \left[c_{k+1}(b[X_{k+1|k}], u_{k|k}) + \ldots \int_{z_{i|k}} \mathbb{P}(z_{i|k}|H^{-}_{i|k}) \left[c_i + \ldots \right] \ldots \right]$$

Future measurements

Measurement Weight

Future Belief

Future candidate action

$$\{z_{k+i|k}\}_1^n \sim \mathbb{P}(z_{k+i|k}|H^{-}_{k+i|k})$$

$$J(u) \approx \frac{1}{n} \sum_{\{z_{k+1|k}\}} \left[c_{k+1|k} + \frac{1}{n} \sum_{\{z_{k+2|k}\}} \left[c_{k+2|k} + \cdots \left[c_{k+L-1|k} + \frac{1}{n} \sum_{\{z_{k+L|k}\}} \left[c_{k+L|k} \right] \right] \right] \right]$$

$\square_{t|k}$ - Referring to time t, while current time is k

Belief Space Planning Formulation

$$J(u) = \int_{z_{k+1|k}} \mathbb{P}(z_{k+1|k}|H^-_{k+1|k}) \left[c_{k+1}(b[X_{k+1|k}], u_{k|k}) + \ldots \int_{z_{i|k}} \mathbb{P}(z_{i|k}|H^-_{i|k}) \left[c_i + \ldots \right] \ldots \right]$$

Future measurements | Measurement Weight | Future Belief | Future candidate action

Under Maximum Likelihood (ML) Assumption

$$z_{k+i|k} = \underset{z_{k+i|k}}{argmax}\ \mathbb{P}(z_{k+i|k}|H^-_{k+i|k})$$

$$J(u) \approx \quad c_{k+1|k} + \quad c_{k+2|k} + \quad c_{k+L-1|k} + \quad c_{k+L|k}$$

ML effect over estimation

- Gaussian prior on robot pose (mean at black square), two types of landmarks: high (1) and low (2) uncertainty.
- Robot considers two candidate actions: step left or step forward.
- Each colored pixel denotes a possible ground truth within the 1σ range, and the resulting cost value .
- Although “left” is statistically favorable, ML-BSP will choose forward

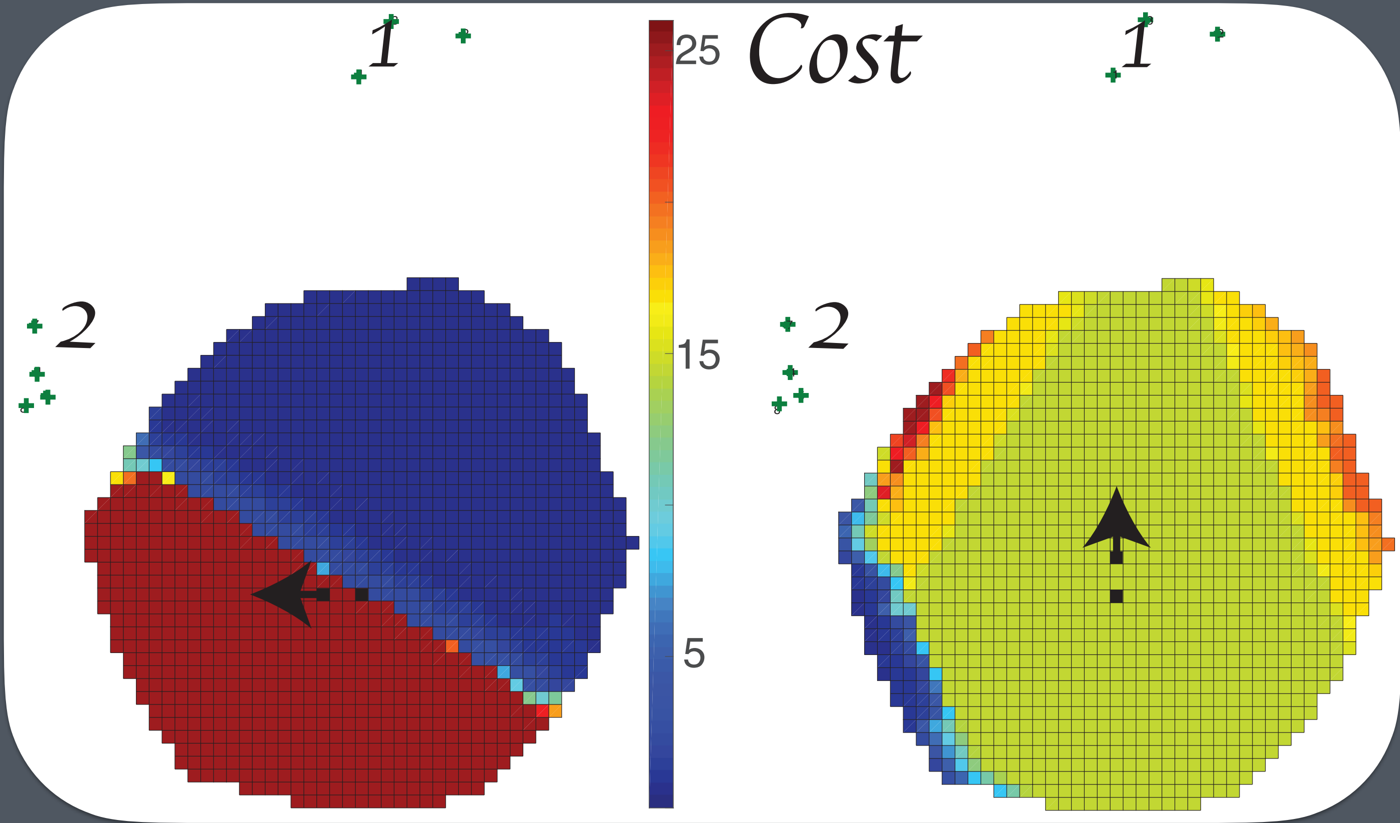

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

- BSP formulation
- Related work
- iX-BSP
  - Results
- The Wildfire Assumption
  - Results
- iML-BSP
  - Results - simulation
  - Results - live

Research Outline

- Introducing Joint Inference & Planning
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

- BSP formulation
- Related work
- iX-BSP
  - Results
- The Wildfire Assumption
  - Results
- iML-BSP
  - Results - simulation
  - Results - live

Related work on Incremental Decision Making Under Uncertainty

Research \ Characteristics	General Distribution?	Not using ML ?	Planning re-use ?
FIRM			
DESPOT			
is-DESPOT			
Platt11isrr			
Chaves16iros	Gaussian		
Kopitkov17ijrr	Gaussian		
ABT			
POMCP			
SARSOP			
Our iX-BSP			

Related work on Incremental Decision Making Under Uncertainty

- Building on POMCP, Adaptive Belief Tree (ABT) uses an offline calculated policy. When given as input the segments of the policy affected by posterior information, it freshly resample them (Kurniawati & Yadav 2016)
- While considering Gaussian belief under Maximum Likelihood (ML) assumption:
  - Utilizing a fixed shared location for all candidate actions for calculation re-use (Chaves & Eustice 2016)
  - Utilizing an augmented matrix determinant lemma to avert from belief propagation under information theoretic cost (Kopitkov & Indelman 2017)

Till this day, to the best of our knowledge,

Incrementally re-using decision-making under uncertainty has not been done for the general case.

Research Outline

- Introducing Joint Inference & Planning
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

- BSP formulation
- Related work
- iX-BSP
  - Results
- The Wildfire Assumption
  - Results
- iML-BSP
  - Results - simulation
  - Results - live

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

- BSP formulation
- Related work
- iX-BSP
  - Results
- The Wildfire Assumption
  - Results
- iML-BSP
  - Results - simulation
  - Results - live

iX-BSP: Main Contributions

- A novel paradigm for incremental expectation BSP, with selective re-sampling of future measurements.
- Identifying the problem of iX-BSP with selective re-sampling as a Multiple Importance Sampling problem, and provide the proper estimator using the balance heuristic
- Statistical comparison of iX-BSP to X-BSP (calculates expectation from scratch)
- Introduce the wildfire approximation to iX-BSP, which allows one to controllably trade accuracy for performance

(Farhi19icra) (Farhi19icra workshop) (Farhi20journal to be submitted soon) (patent: US20200327358A1)

iX-BSP: Main Contributions

- Supplying bounds and empirical results for the effect wildfire holds over the objective value
- Demonstrate how iX-BSP could also benefit approximations of X-BSP, e.g. iML-BSP
- Comparing iML-BSP to ML-BSP in both simulation and live experiments, considering the problem of autonomous navigation in unknown environments.

(Farhi19icra) (Farhi19icra workshop) (Farhi20journal to be submitted soon) (patent: US20200327358A1)

Comparing two planning sessions

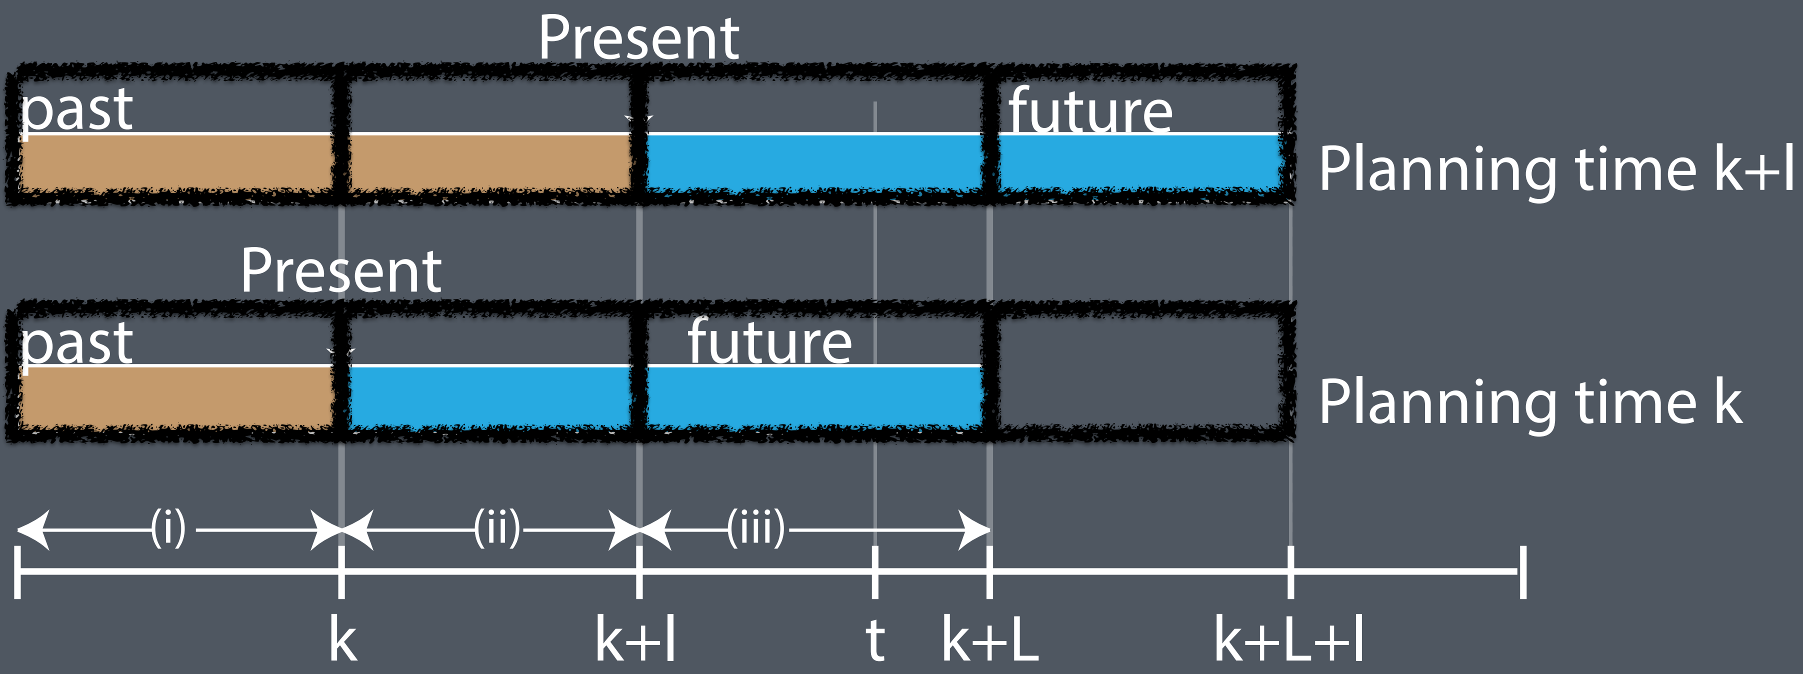

iX-BSP Governing Principle

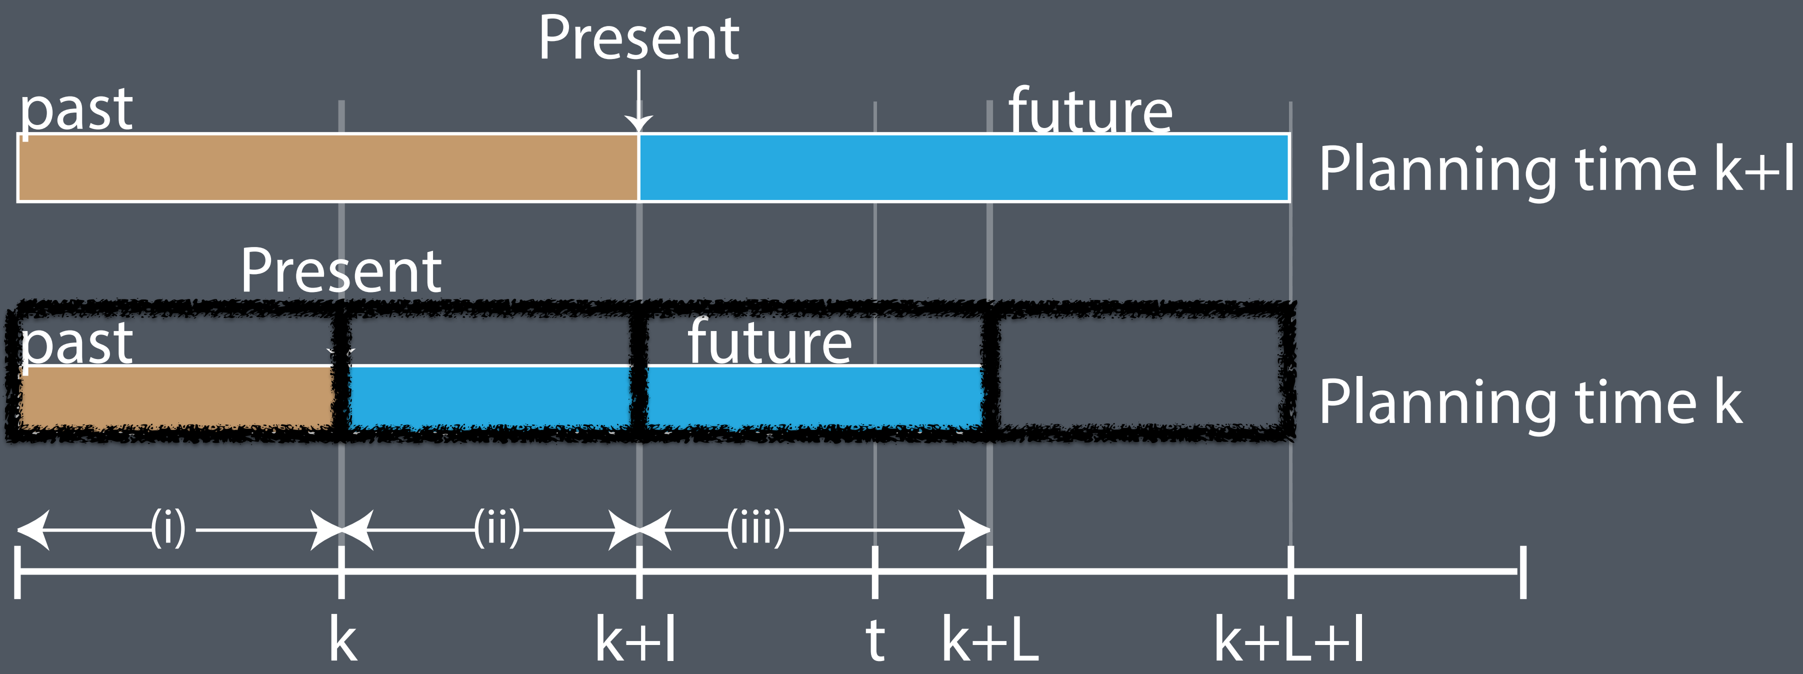

iX-BSP Illustration

- We illustrate full expectation-based BSP followed by our novel iX-BSP
- Instead of performing expectation from scratch, iX-BSP re-uses previous planning session(s
- In order to keep this illustration simple we assume the following:
  - a single candidate action
  - 2 samples per belief => $n_x = 2$, $n_z = 1$
  - Planning horizon of 3 steps

- Assume we completed inference for Current time t = 1
- $b[X_{1|1}]$ denotes belief at current time t = 1
- Belief uncertainty is illustrated by an ellipse
- next: Execute BSP to decide on next action

$b[X_{1|1}]$ 

X-BSP

Standard eXpectation BSP

- Consider $u_1 \Rightarrow u_2 \Rightarrow u_3$ sequence
- Propagate belief with candidate action u_1
- Obtain $b^-[X_{2|1}]$
- Sample measurements

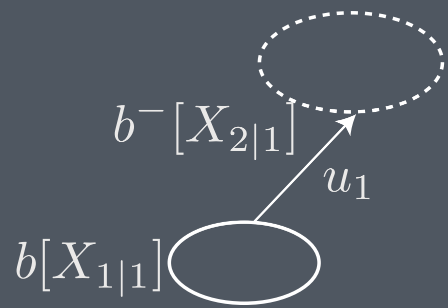

Sample measurements

- Since we do not have access to the measurement likelihood $\mathbb{P}(z_{k+i|k}|H^{-}_{k+i|k})$
- We sample states and given those states, measurements
- Based on the following equality

$$\underbrace{\mathbb{P}(z_{k+i|k}|H^{-}_{k+i|k})}_{\text{Measurement Likelihood}} = \int_{\underbrace{X_{k+i|k}}_{\text{Future state}}} \underbrace{\mathbb{P}(z_{k+i|k}|X_{k+i|k})}_{\text{Measurement Model}} \cdot \underbrace{\mathbb{P}(X_{k+i|k}|H^{-}_{k+i|k})}_{\text{Propagated belief}} dX$$

X-BSP

Standard eXpectation BSP

- Consider $u_1 \Rightarrow u_2 \Rightarrow u_3$ sequence
- Propagate belief with candidate action u_1
- Obtain $b^-[X_{2|1}]$
- Sample measurements

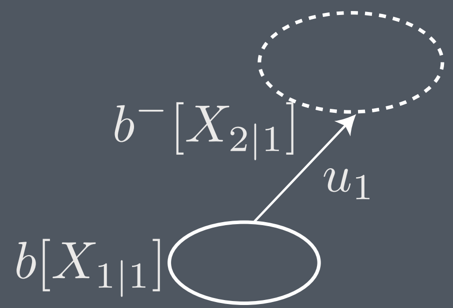

X-BSP

Standard eXpectation BSP

- Sampling two states, i.e. n_x=2
- and for each a single set of measurements, i.e. n_z=1
- Consider each of the sets and the propagated belief to obtain the posterior beliefs for future time t=2

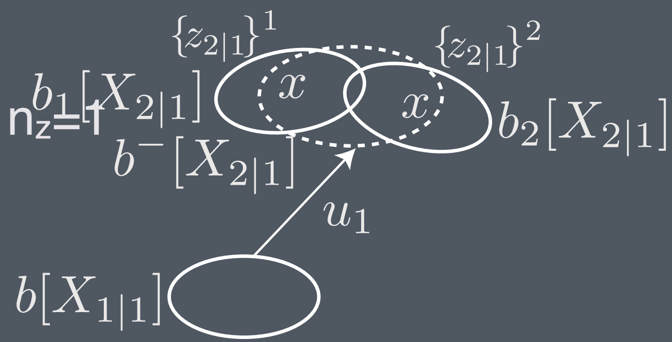

X-BSP

Standard eXpectation BSP

- And again for the second horizon step
- Propagate future beliefs
- Sample measurements
- Calculate future beliefs

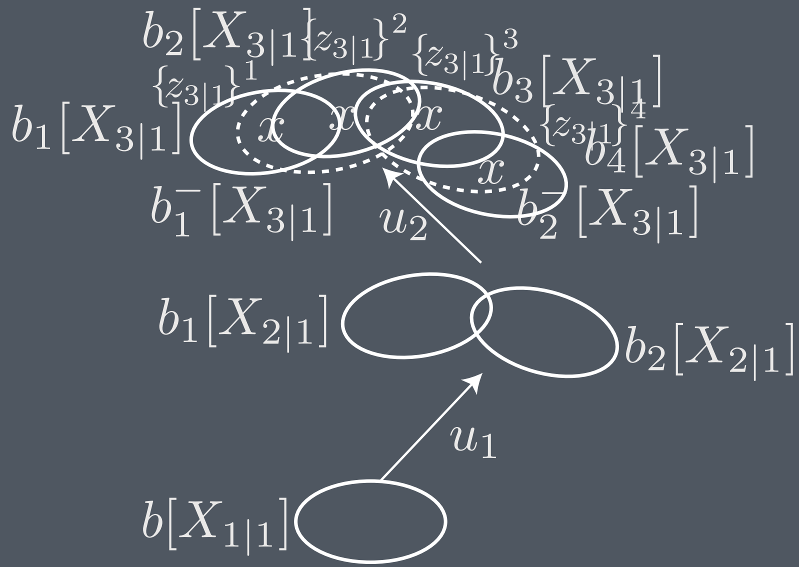

X-BSP

Standard eXpectation BSP

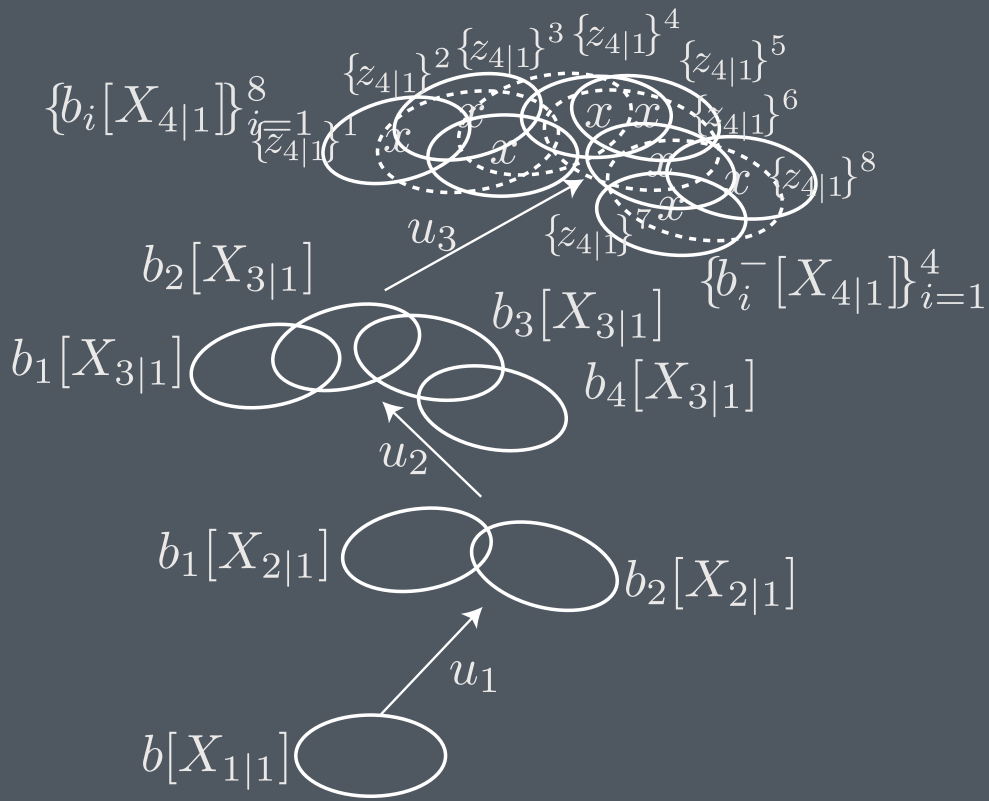

- And for the last horizon step
- Propagate beliefs
- Sample measurements
- Calculate future beliefs

~~Standard eXpectation~~ BSP

For Visual Reference: ML-BSP

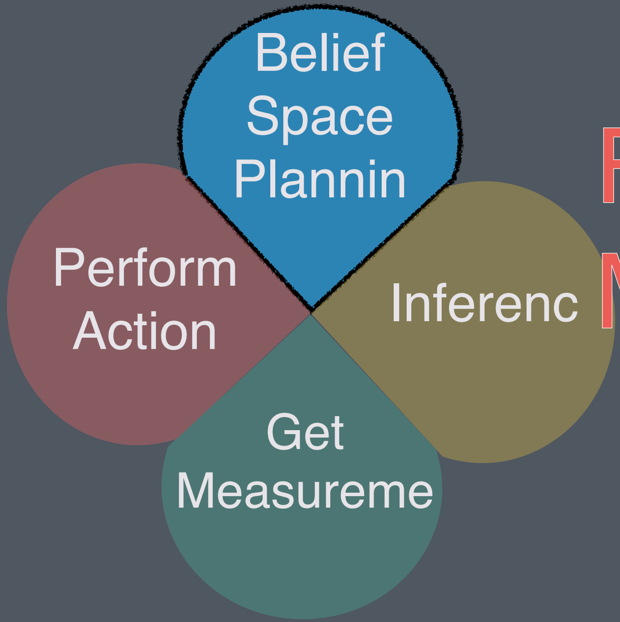

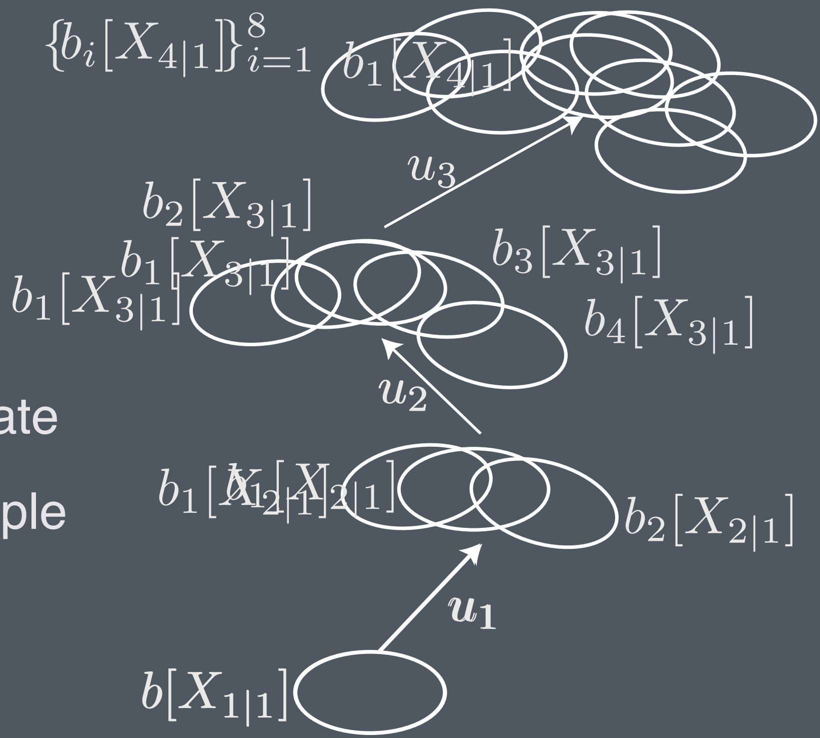

- At each horizon step we have one candidate belief
- Created using a single measurement sample
- More specifically, the most likely one

X-BSP

Standard eXpectation BSP

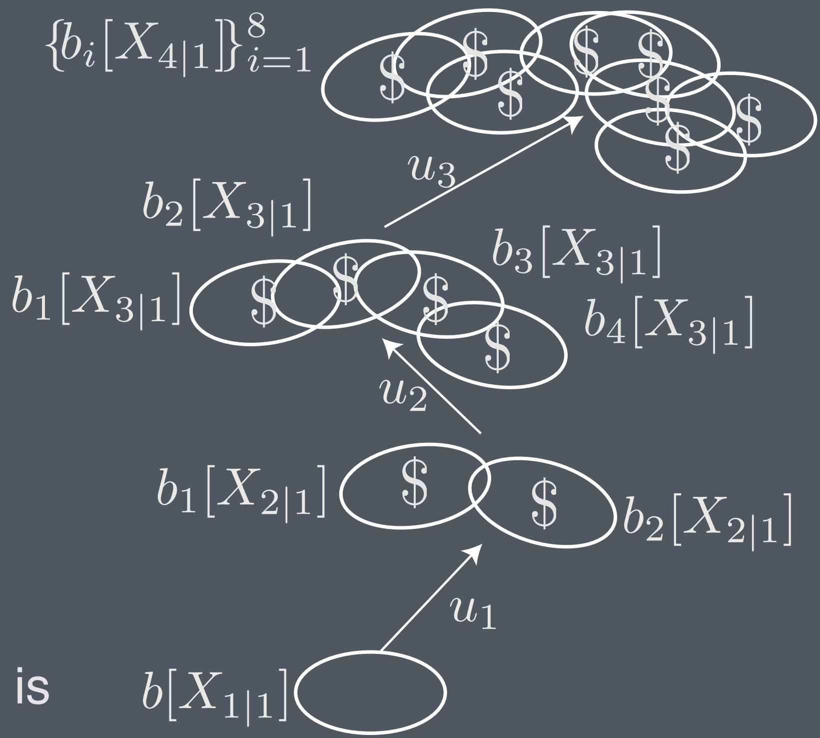

- For each belief we calc the reward(cost) value
- Rewards of the same action are averaged together
- The objective for each action sequence is calculated
- Action sequence with best objective value is chosen

- Execute action u_1

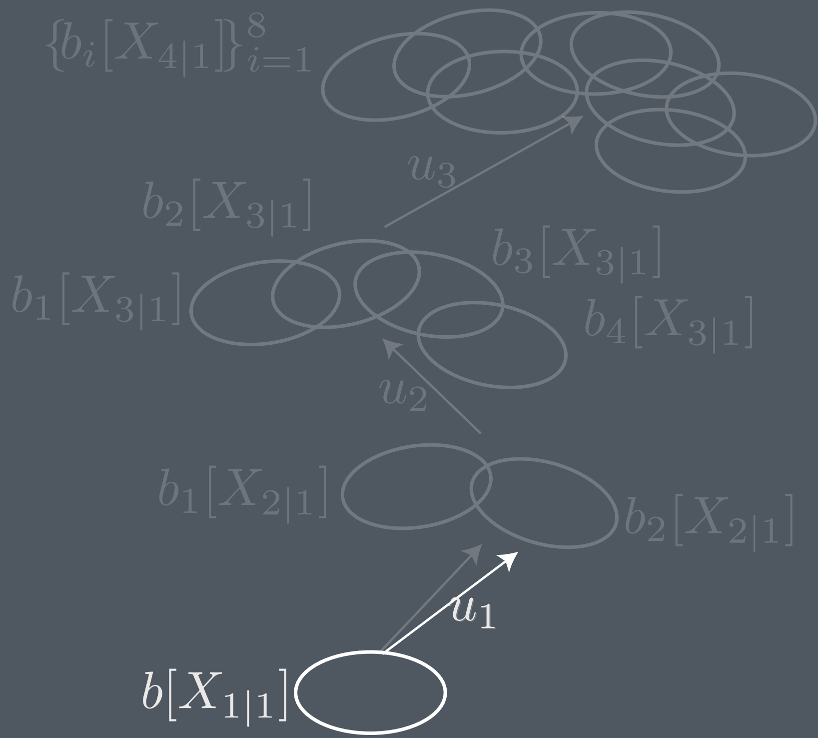

- Execute action u_1
- Get measurements for time t = 2

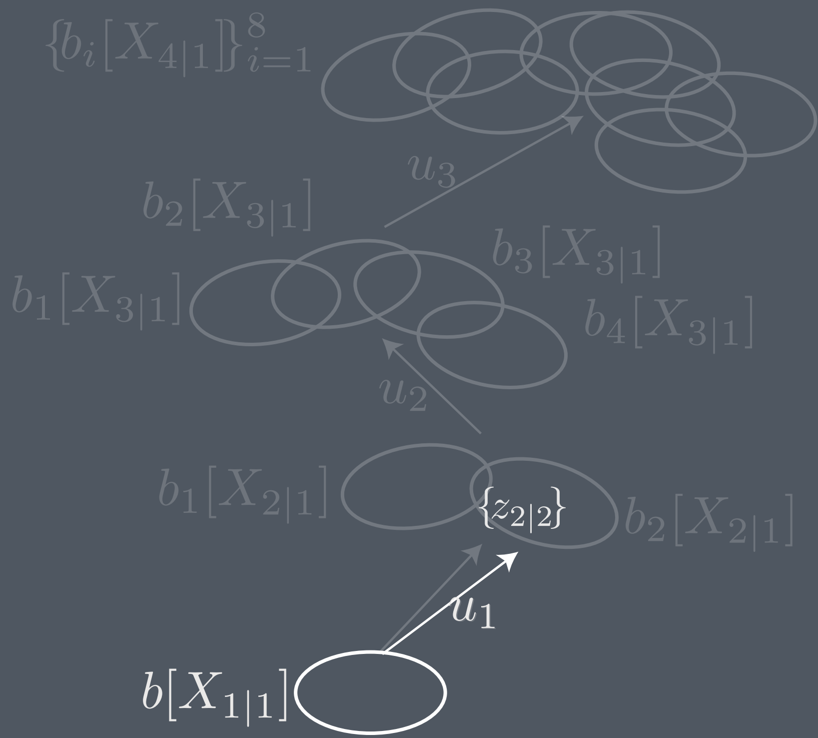

- Execute action u_1
- Get measurements for time t = 2
- Perform inference for time t = 2
- next: Execute iX-BSP to decide on next action

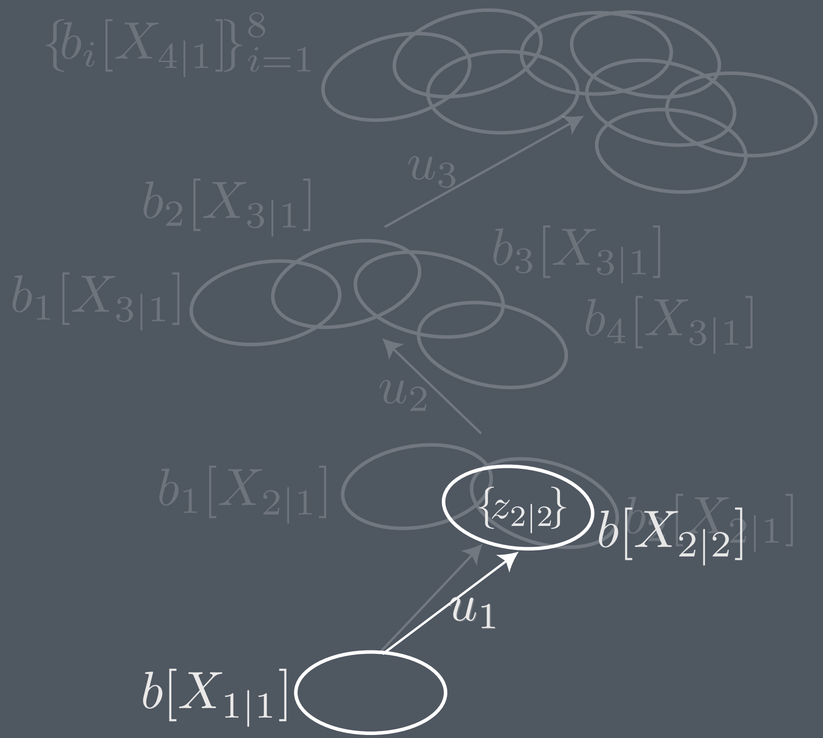

Incremental eXpectation BSP

iX-BSP

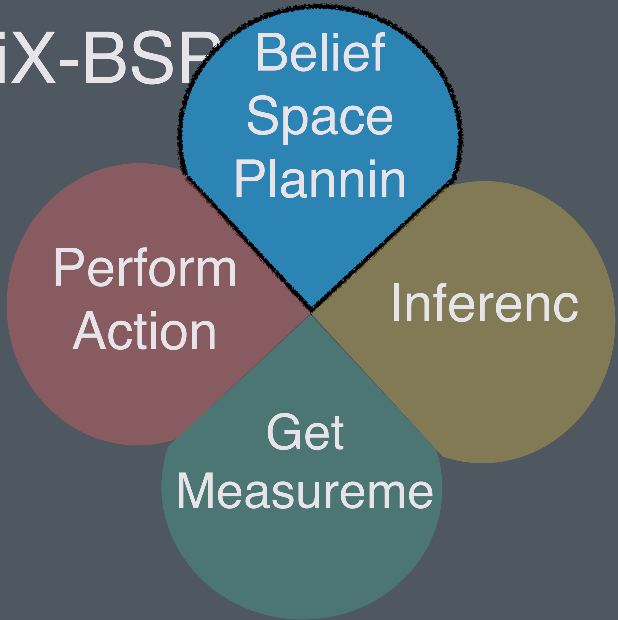

- Check which $b_i[X_{2|1}]$ is closest to $b[X_{2|2}]$

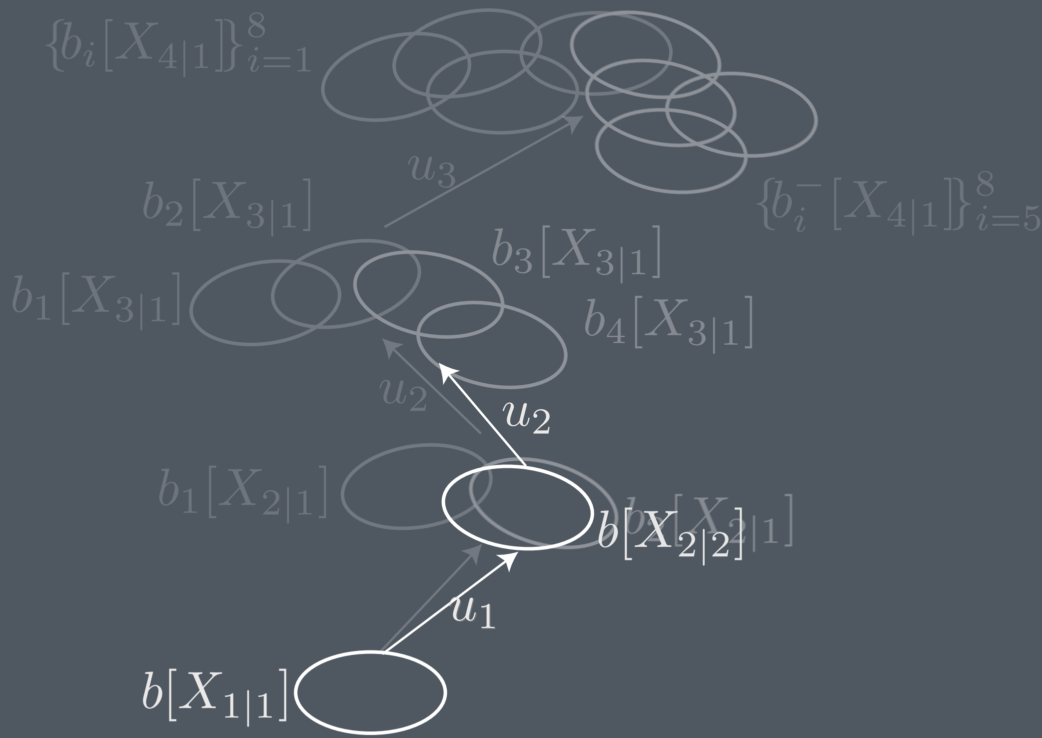

Belief Distance

Incremental eXpectation BSP

- When re-using a belief, iX-BSP updates it to match the posterior information.
- The more different it is, the more computation time is required to update it.
- For this reason we aspire to find the closest belief.
- After much consideration we chose to use the square root Jeffreys divergence - $\mathbb{D}_{\sqrt{J}}$

$$\mathbb{D}_{\sqrt{J}}(b, b') = \sqrt{\frac{1}{2}\mathbb{D}_J} = \sqrt{\frac{1}{2}\mathbb{D}_{KL}(b||b') + \frac{1}{2}\mathbb{D}_{KL}(b'||b)}$$

- We define some critical distance value - ε_c as the threshold for considering a belief as worth re-using.

Incremental eXpectation BSP

iX-BSP

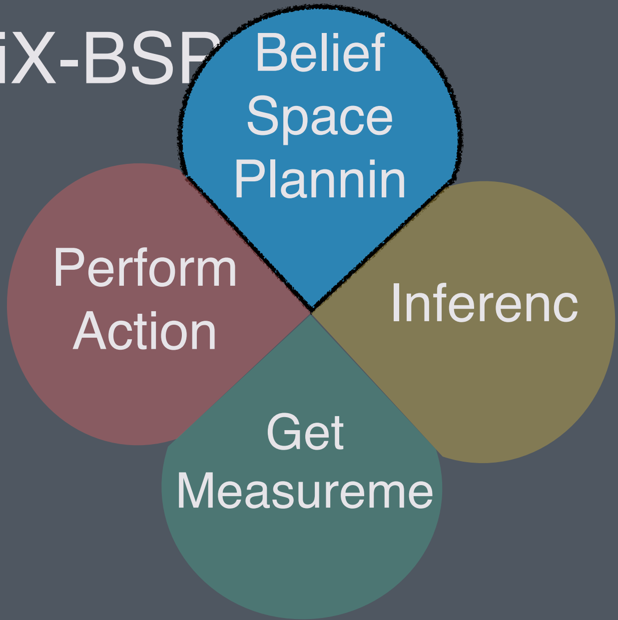

- Check which $b_i[X_{2|1}]$ is closest to $b[X_{2|2}]$
- Consider its children as candidates for re-use
- Consider $u_2 \Rightarrow u_3 \Rightarrow u_4$ sequence
- Propagate belief with candidate action u_2

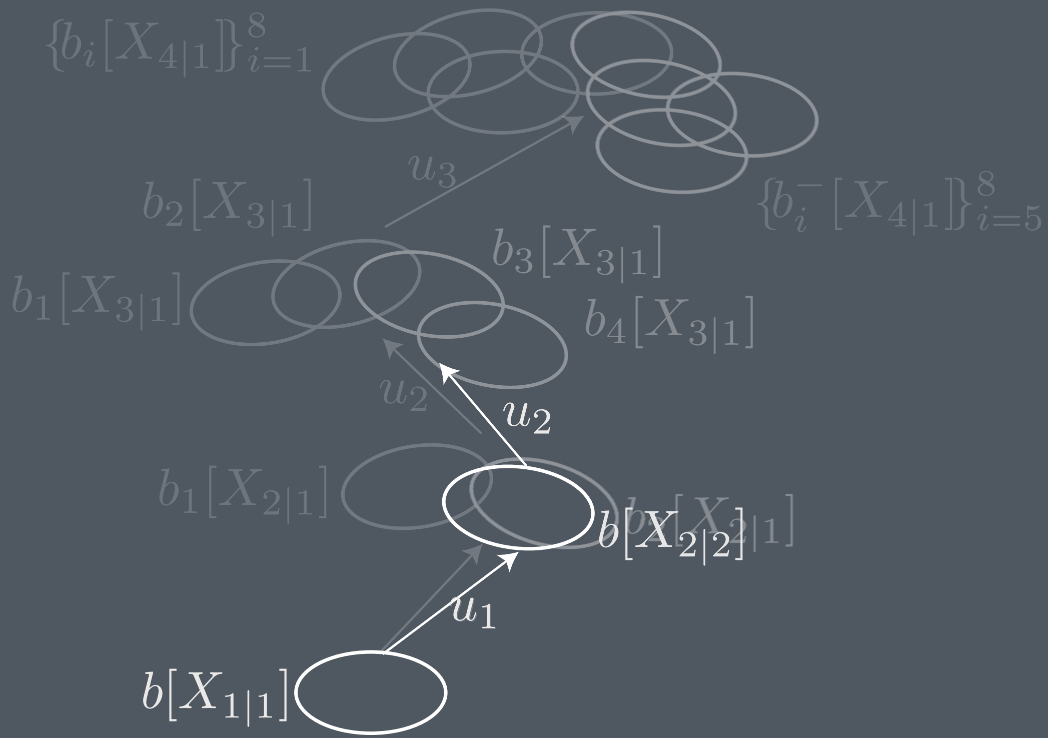

Incremental eXpectation BSP

iX-BSP

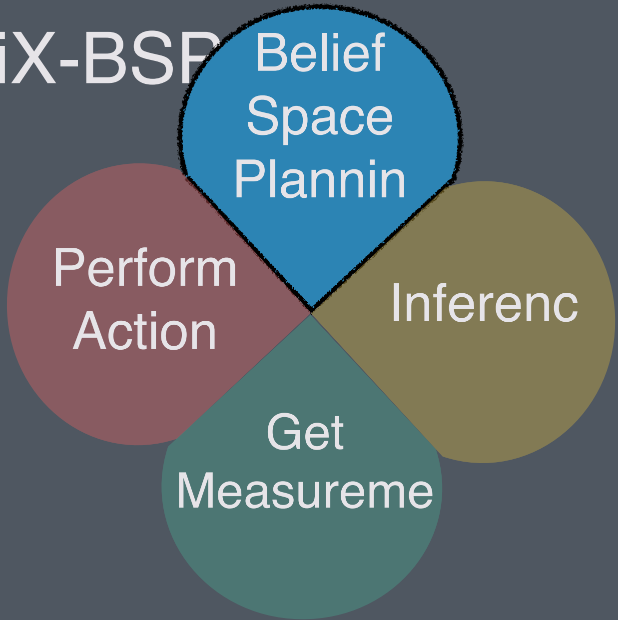

- Obtain $b^-[X_{3|2}]$
- Consider old samples
- Re-use representative samples (in blue)
- Re-sample the rest

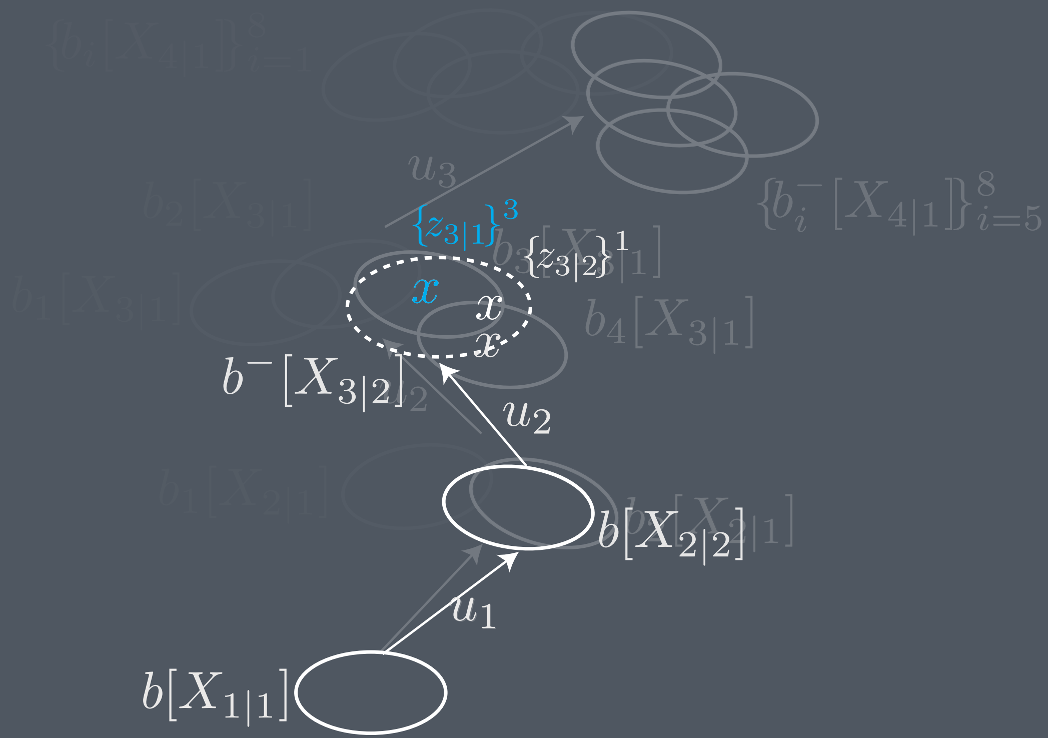

Representative Samples

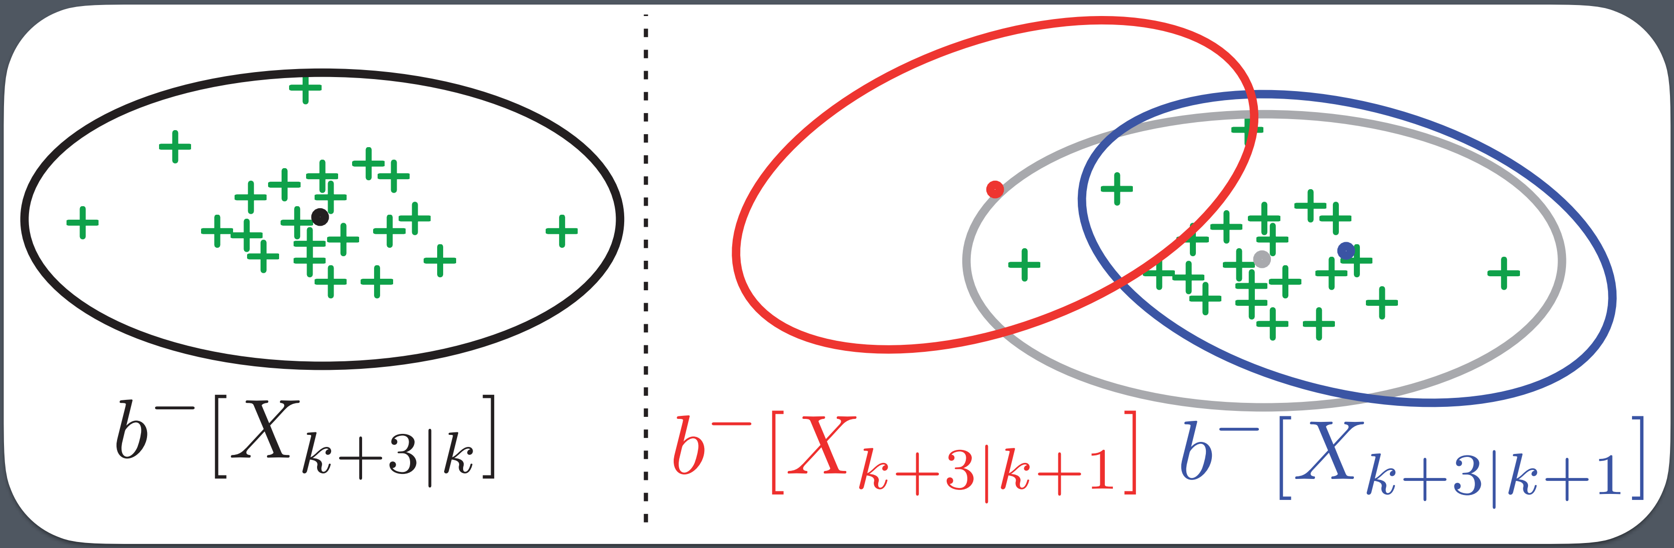

Incremental eXpectation BSP

iX-BSP

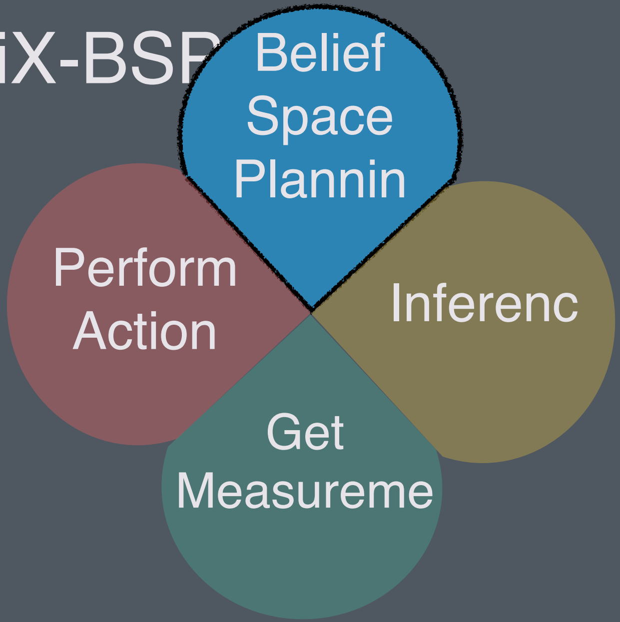

- For re-used samples, re-use beliefs
- Update these beliefs with info from t = 2
- For the rest of the samples - calc the beliefs
- Propagate future beliefs

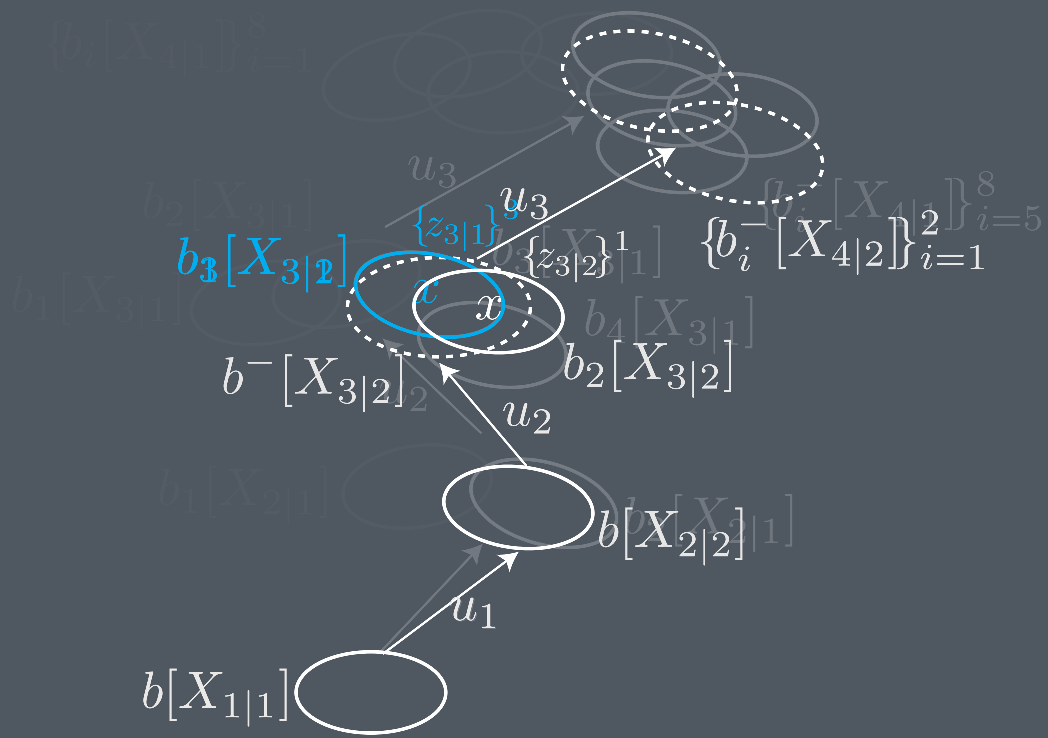

iX-BSP

Incremental eXpectation BSP

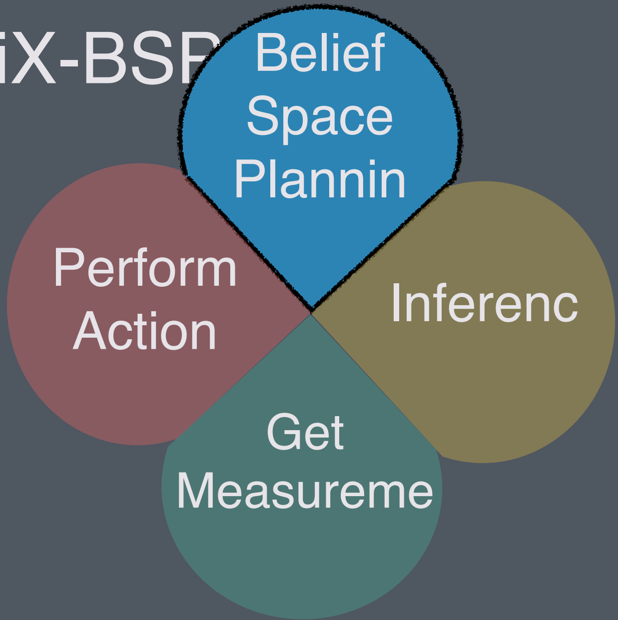

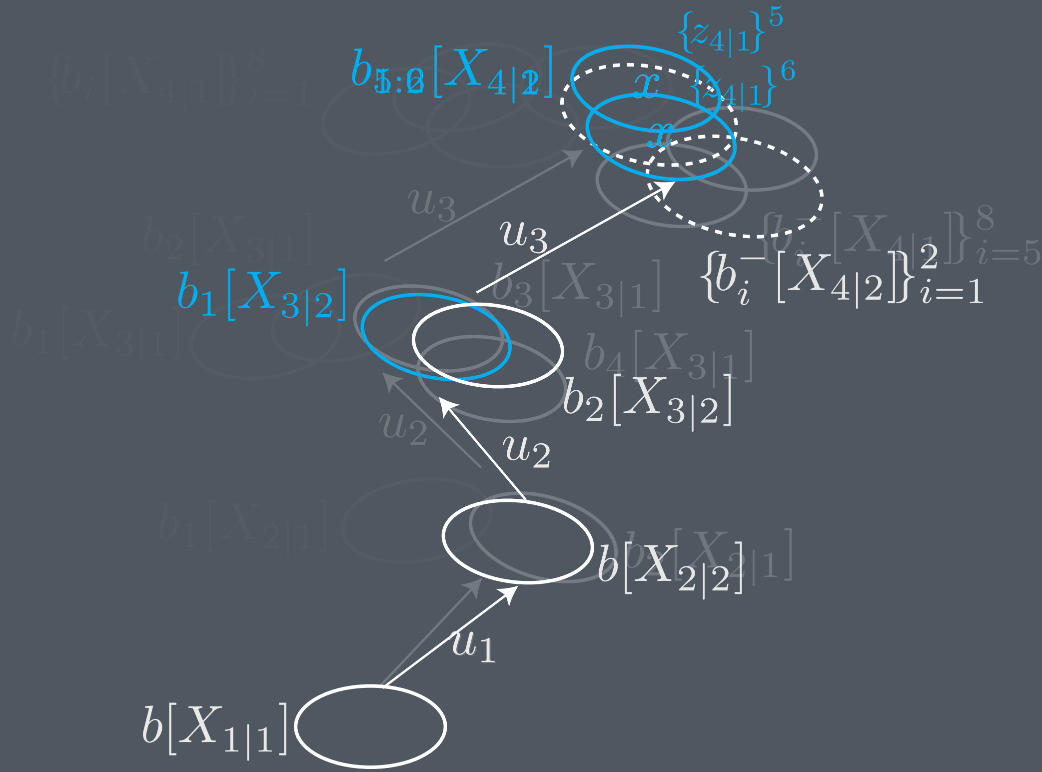

- Consider old samples
- Re-use representative samples (in blue)
- For re-used samples, re-use beliefs
- Update these beliefs with info from t = 2

iX-BSP

Incremental eXpectation BSP

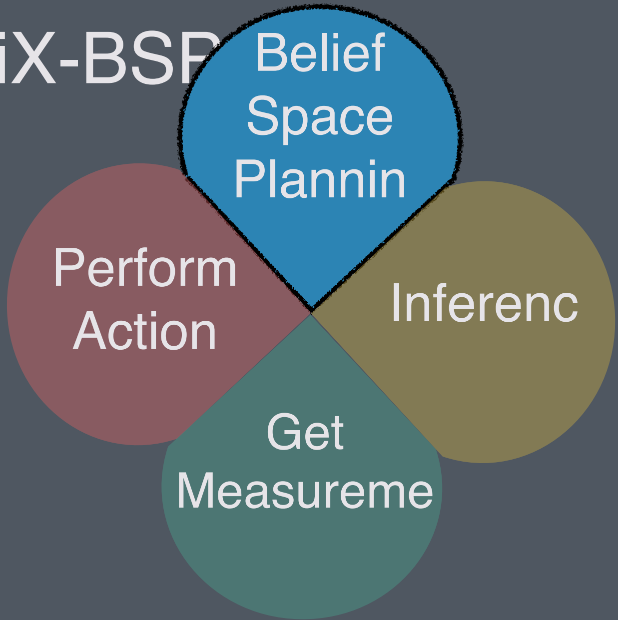

- Re-sample the rest of the measurements
- Calculate the rest of the beliefs
- Propagate future beliefs
- Last horizon step, i.e. use X-BSP

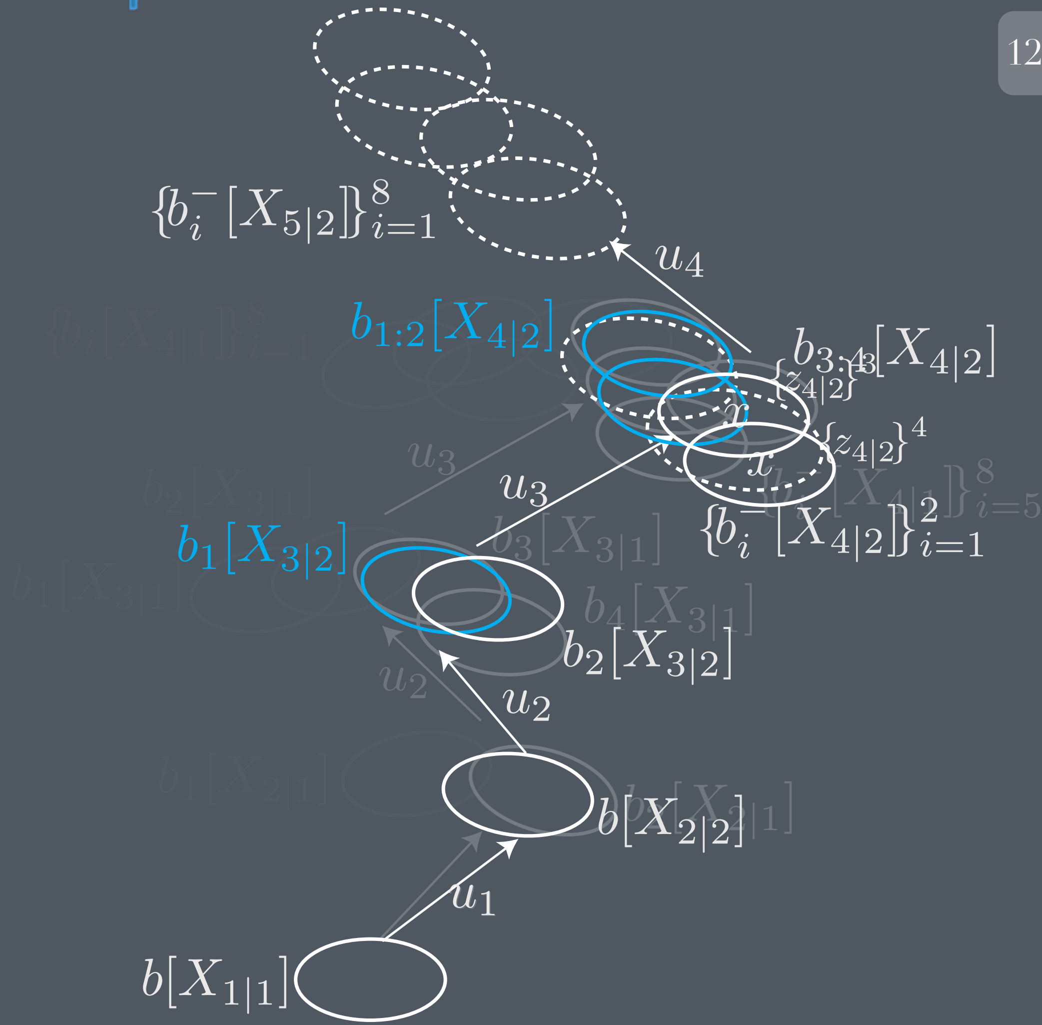

iX-BSP

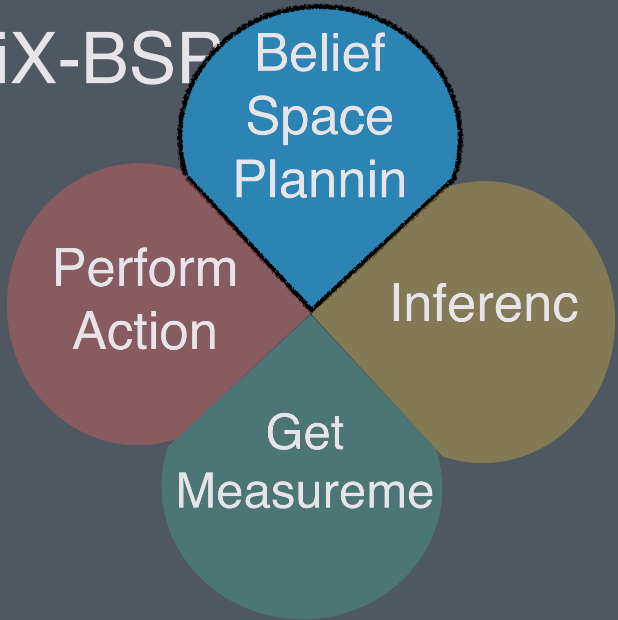

Incremental eXpectation BSP

Standard eXpectation BSP for last horizon step

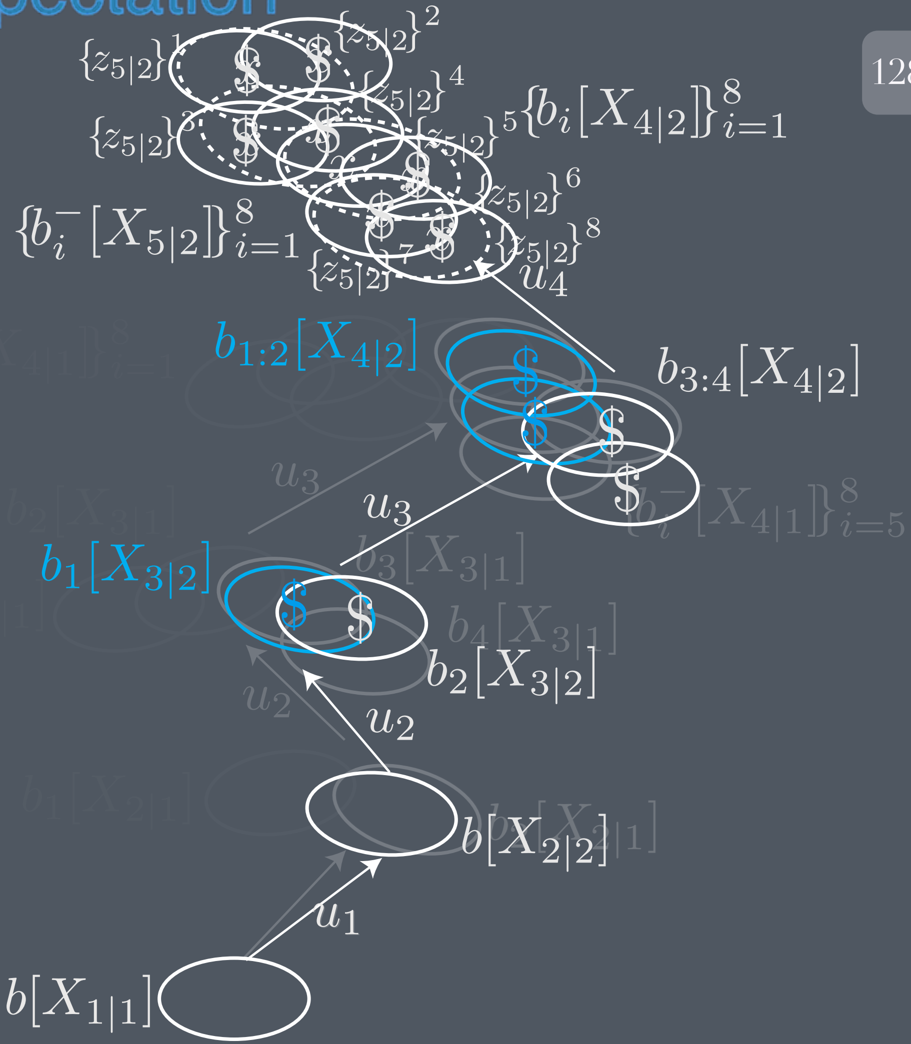

- Sample measurements
- Calculate the beliefs
- For each belief we calc/update the reward(cost) value
- Weighting rewards of the same action

Weighting rewards of the same action

- Based on color coding, per action, samples were taken from multiple measurement PDFs
  $\mathbb{P}(z_i | H_{k|k}, u_{k:i-1|k})$
- i.e. we are required to use importance sampling for proper weighting
- For this toy example we have a max of 2 different PDFs per step
- e.g. the weight corresponding to $b_1[X_{3|2}]$ is given by

$$w_3^1 = \frac{\mathbb{P}(z_3^1 | H_{2|2}, u_{2|2})}{\frac{1}{2}\mathbb{P}(z_3^1 | H_{1|1}, u_{1:2|1}) + \frac{1}{2}\mathbb{P}(z_3^1 | H_{2|2}, u_{2|2})}$$

- For the general case, each measurement can be sampled from a different measurement PDF

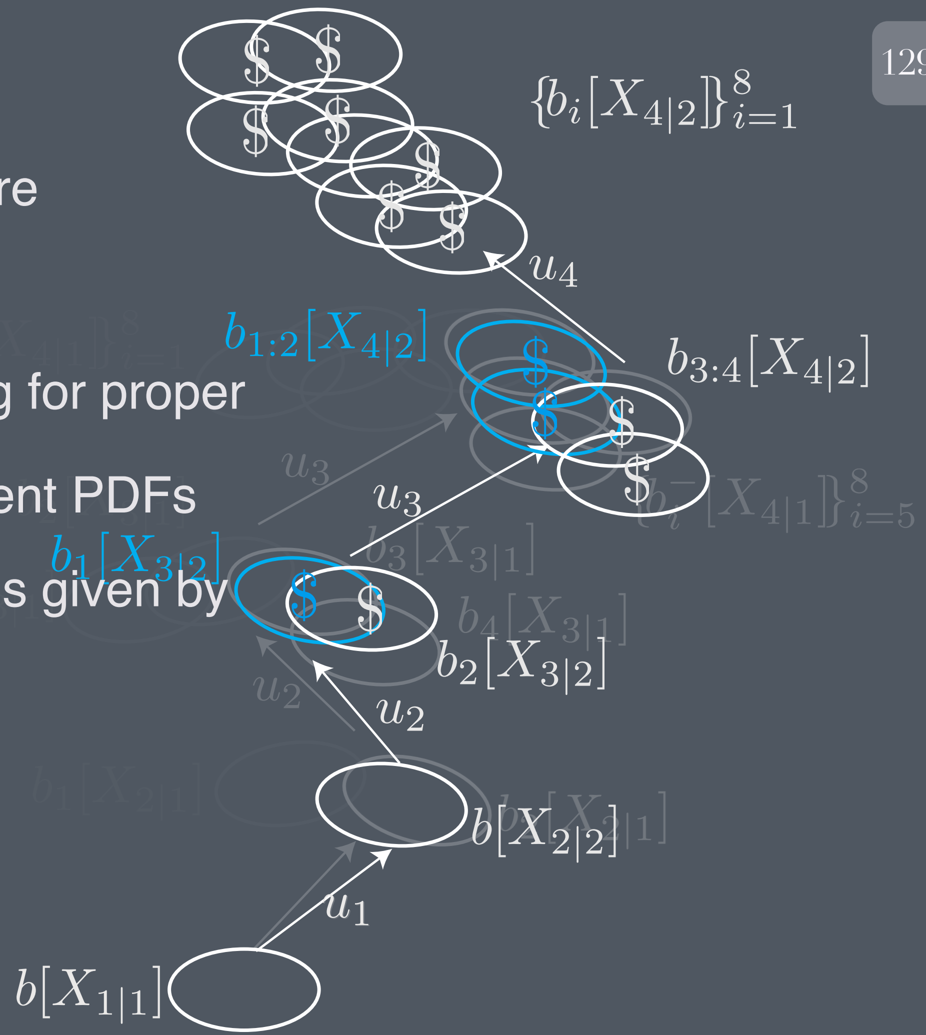

iX-BSP

Incremental eXpectation BSP

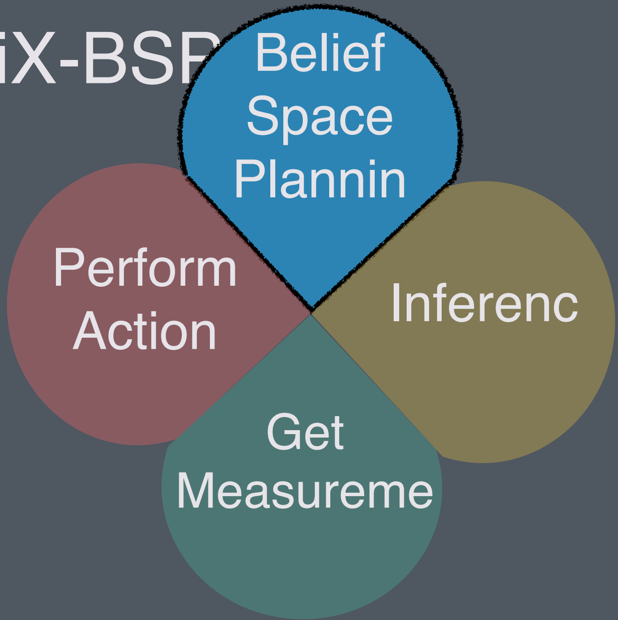

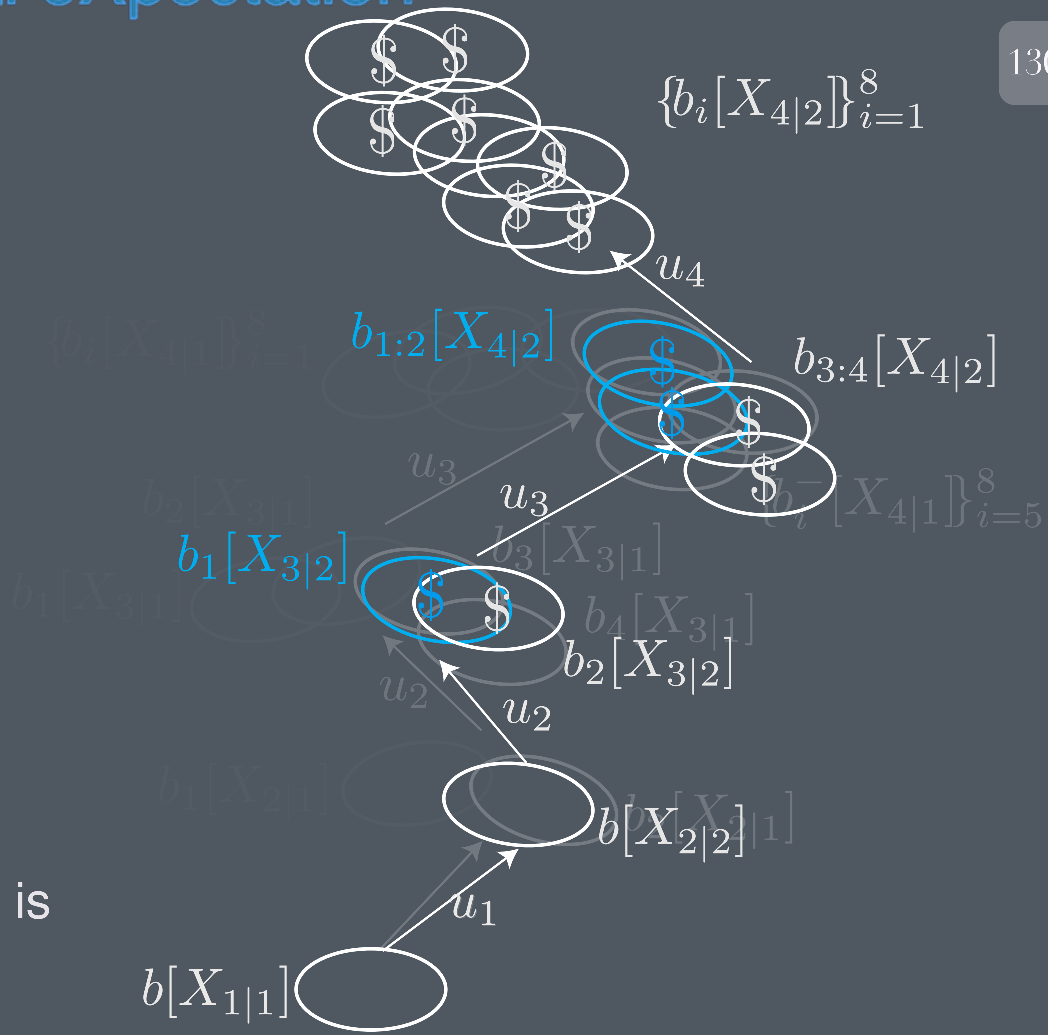

- Weighting rewards of the same action
- The objective for each action sequence is calculated
- Action sequence with best objective value is chosen

iX-BSP- Multiple Importance Sampling Objective Estimator

$$J(u') \approx \underbrace{\sum_{i=k+l+1}^{k+l+L}}_{\text{Horizon}} \left[\frac{1}{n_i} \underbrace{\sum_{m=1}^{M_i}}_{\text{Num of distrib.}} \underbrace{\sum_{g=1}^{n_m}}_{\text{Num of samples}} \underbrace{\omega_i(z^{m,g}_{k+l+1:i})}_{\text{Weight of } g_{th} \text{ sample of } m_{th} \text{ distribution}} \cdot c_i \left(b^{m,g}[X_{i|k+l}], u'_{i-1|k+l} \right) \right]$$

Using the Balance Heuristic:

$$\omega_i(z^{m,g}_{k+l+1:i}) = \frac{\mathbb{P}(z^{m,g}_{k+l+1:i} | H_{k+l|k+l}, u_{k+l:i-1|k+l})}{\Sigma^{M_i}_{\tilde{m}=1} \frac{n_{\tilde{m}}}{n_i} q_{\tilde{m}}(z^{m,g}_{k+l+1:i})}$$

Nominal distribution

All sampled distributions

Research Outline

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

BSP formulation

Related work

iX-BSP

Results

The Wildfire Assumption

Results

iML-BSP

Results - simulation

Results - live

ResultsiX-BSP

- We compare planning time of iX-BSP and standard BSP using expectation (X-BSP).
- We used 10 randomly generated maps, each with two goals.
- The robot is required to visit both goals with an objective that minimize Distance to Goal (D2G) and maximize information gain
- On each map we ran 20 rollouts (entire mission run), each with a different sampled initial ground truth position.
- The robot is equipped with a stereo camera and has no prior knowledge over the environment.
- We considered known models with Gaussian additive noise
- We compare planning computation time, excluding the last horizon step which is identical between the two

Randomly Generated Maps

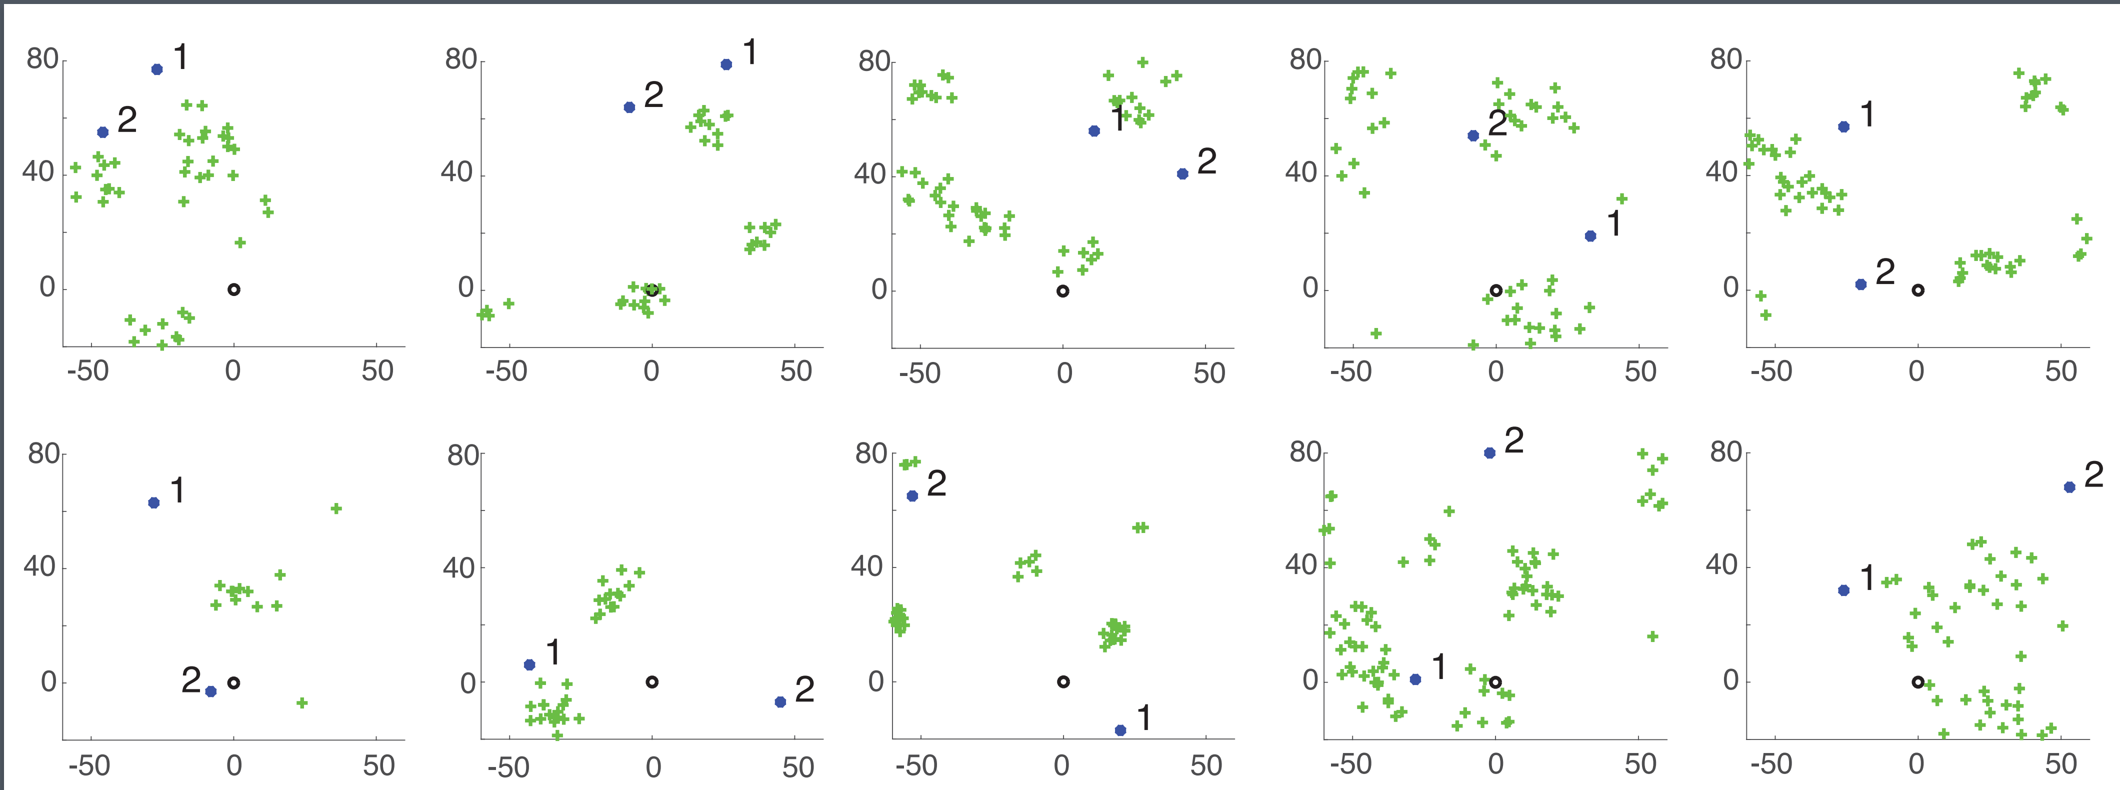

X-BSP vs. iX-BSP

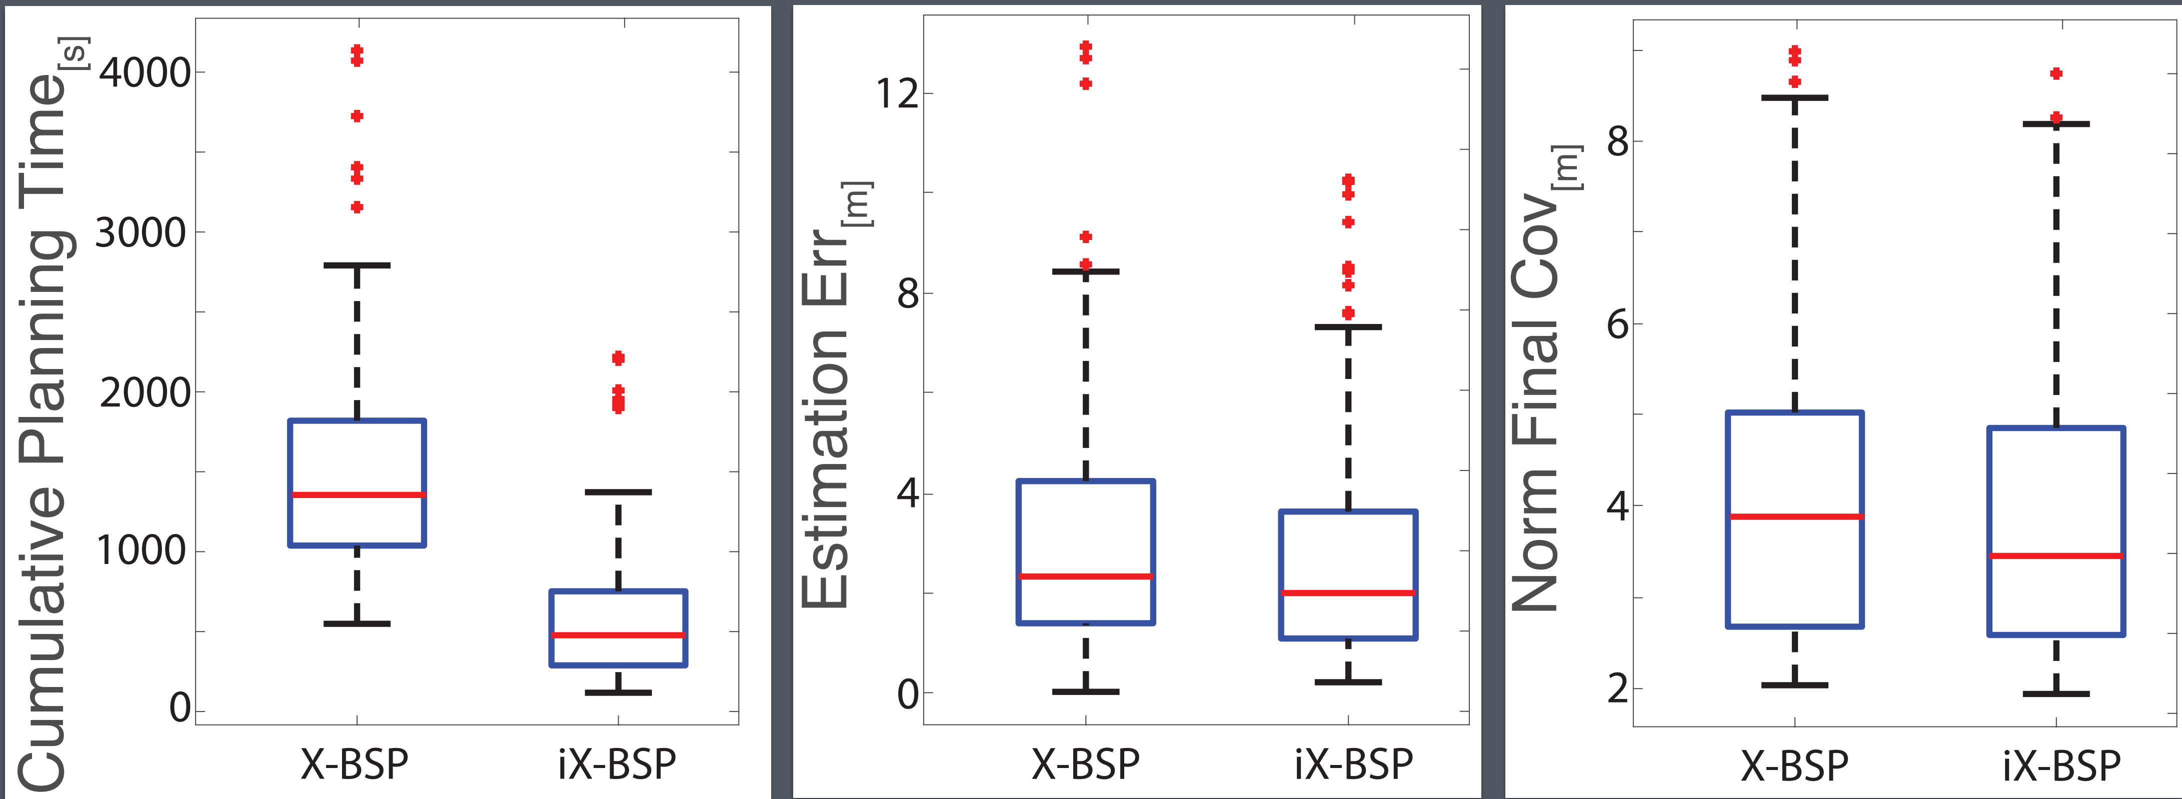

Research Outline

- Introducing Joint Inference & Planning
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

- BSP formulation
- Related work
- iX-BSP
    - Results
- The Wildfire Assumption
    - Results
- iML-BSP
    - Results - simulation
    - Results - live

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

- BSP formulation
- Related work
- iX-BSP
  - Results
- The Wildfire Assumption
  - Results
- iML-BSP
  - Results - simulation
  - Results - live

Introducing- the Wildfire approximation

- iX-BSP does not introduce approximations to the X-BSP solution, as it updates all posterior information, but sometimes a belief might be already "close enough"
- We introduce an approximation to iX-BSP called wildfire
- The wildfire threshold- ε_{wf} , sets an upper bound to consider beliefs as "close enough"

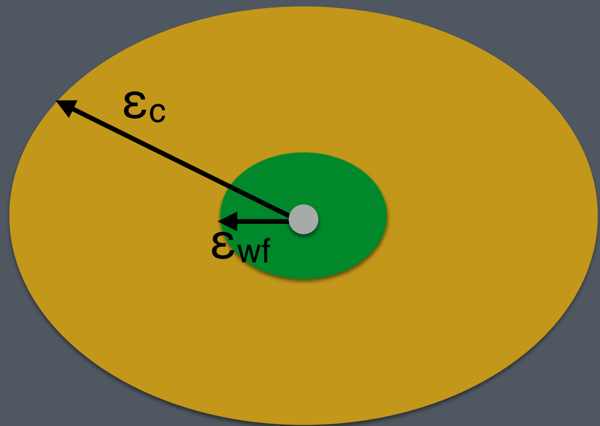

- Once a belief meets the wildfire condition, all its dependents are considered as wildfire as well (hence the name).

Wildfire bounds over objective value

- For $\varepsilon_{wf} = 0$, we consider only identical beliefs as close enough
- For $\varepsilon_{wf} = \infty$, we consider all beliefs as close enough and never update
- From these two edge-cases, we can deduce the choice of ε_{wf} would have a direct impact over the objective value
- Under an assumption of α-Holder reward function, we derived bounds for this impact

Wildfire bounds over objective value

$$\mid J_{k+l|k+l} - J_{k+l|k} \mid \leq \left(2\sqrt{ln2}\right)^{\alpha} \cdot \lambda_{\alpha} \cdot \left[L \cdot \epsilon_{wf}^{\alpha} + \sum_{i=k+l+1}^{k+l+L} \left(\sum_{j=k+l+1}^{i} \mathbb{E}\Delta_j \right)^{\frac{\alpha}{2}} \right]$$

Obj. error for using prev.

Holder param.

Wildfire threshold

Distance propagation along planning horizon

$$\Delta_i = \mathbb{D}^2_{\sqrt{J}}(b[X_{i|k+l}], b[X_{i|k}]) - \mathbb{D}^2_{\sqrt{J}}(b[X_{i-1|k+l}], b[X_{i-1|k}])$$

Distance propagation

Squared Distance between two beliefs at

Squared Distance between two beliefs at time i-1

Research Outline

- Introducing Joint Inference & Planning
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

- BSP formulation
- Related work
- iX-BSP
  - Results
- The Wildfire Assumption
  - Results
- iML-BSP
  - Results - simulation
  - Results - live

Research Outline

- Introducing Joint Inference & Planning
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

iX-BSP with wildfire

- We compare planning time of iX-BSP with and without the use of wildfire
- We used exactly the same scenario over the same 10 maps
- On each map we ran 20 rollouts (entire mission run), each with a different sampled initial ground truth position.

iX-BSP with wildfire

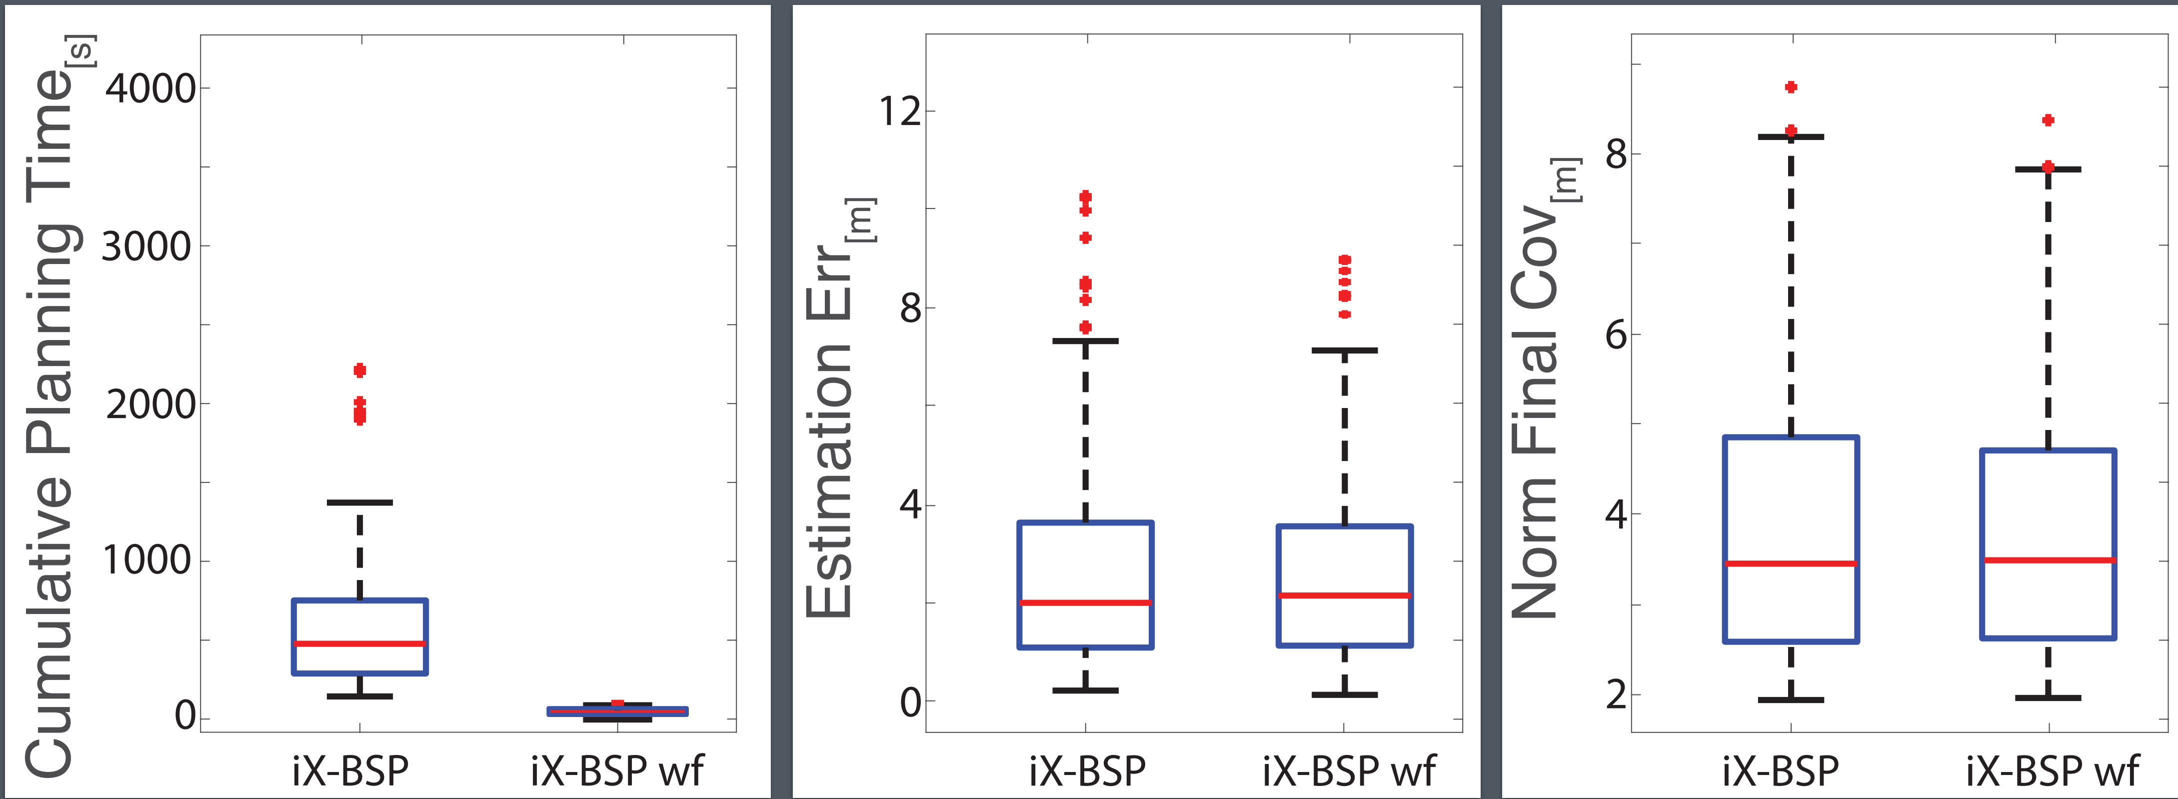

Empirical Objective bounds

- We would like to provide empirical results to the objective error
- We need to perform planning from two beliefs sharing a history with a specific distance between them
- We propagate a belief with predicted (right) and actual (left) measurements
- To control the distance between them we introduce specific noise to the actual action.

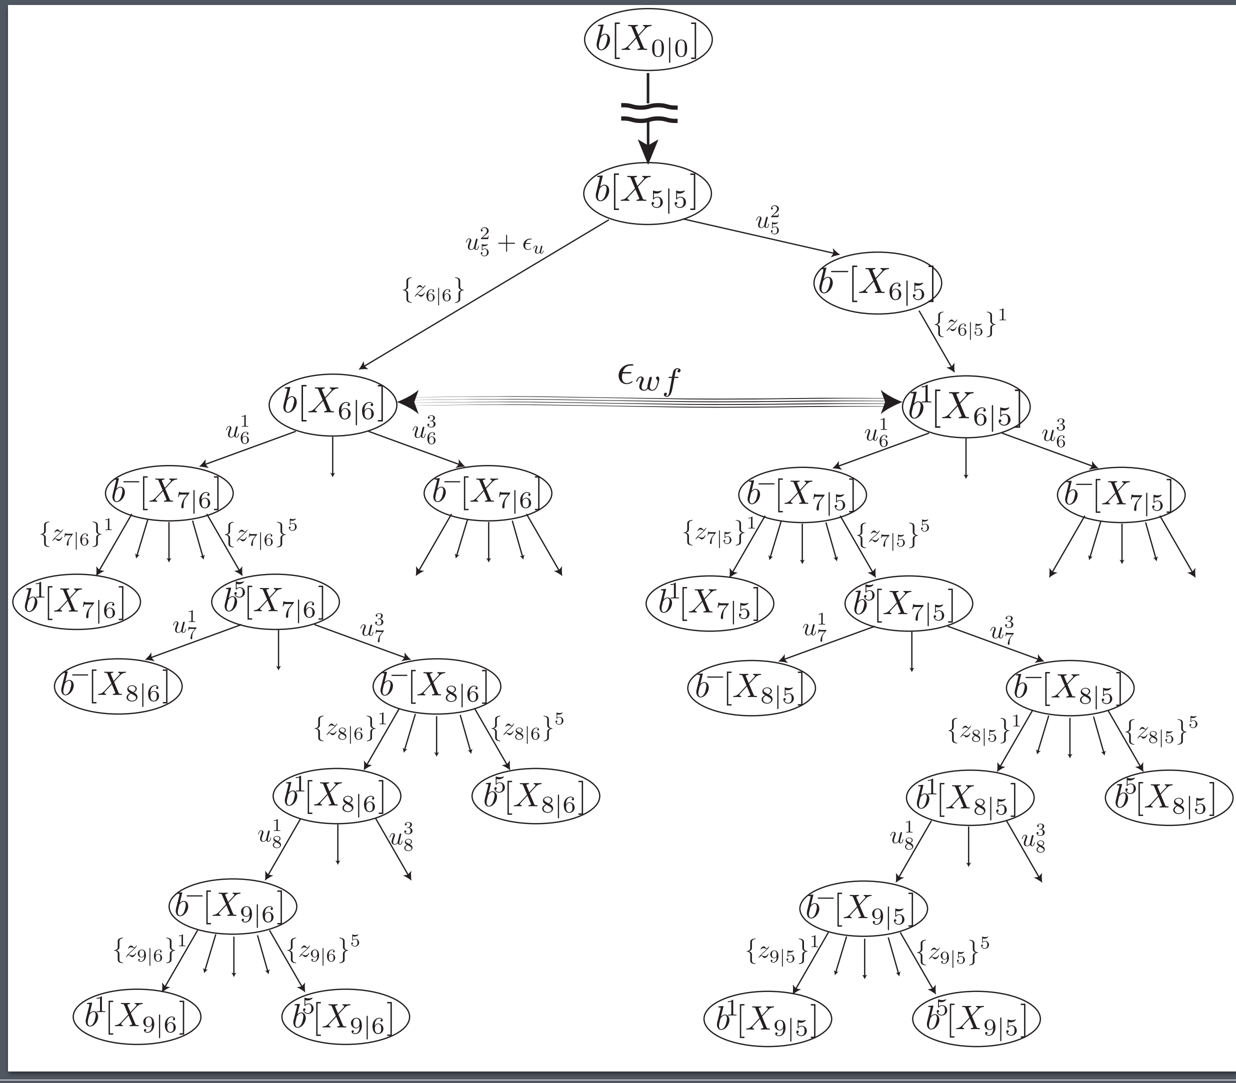

Objective error as a function of wildfire threshold

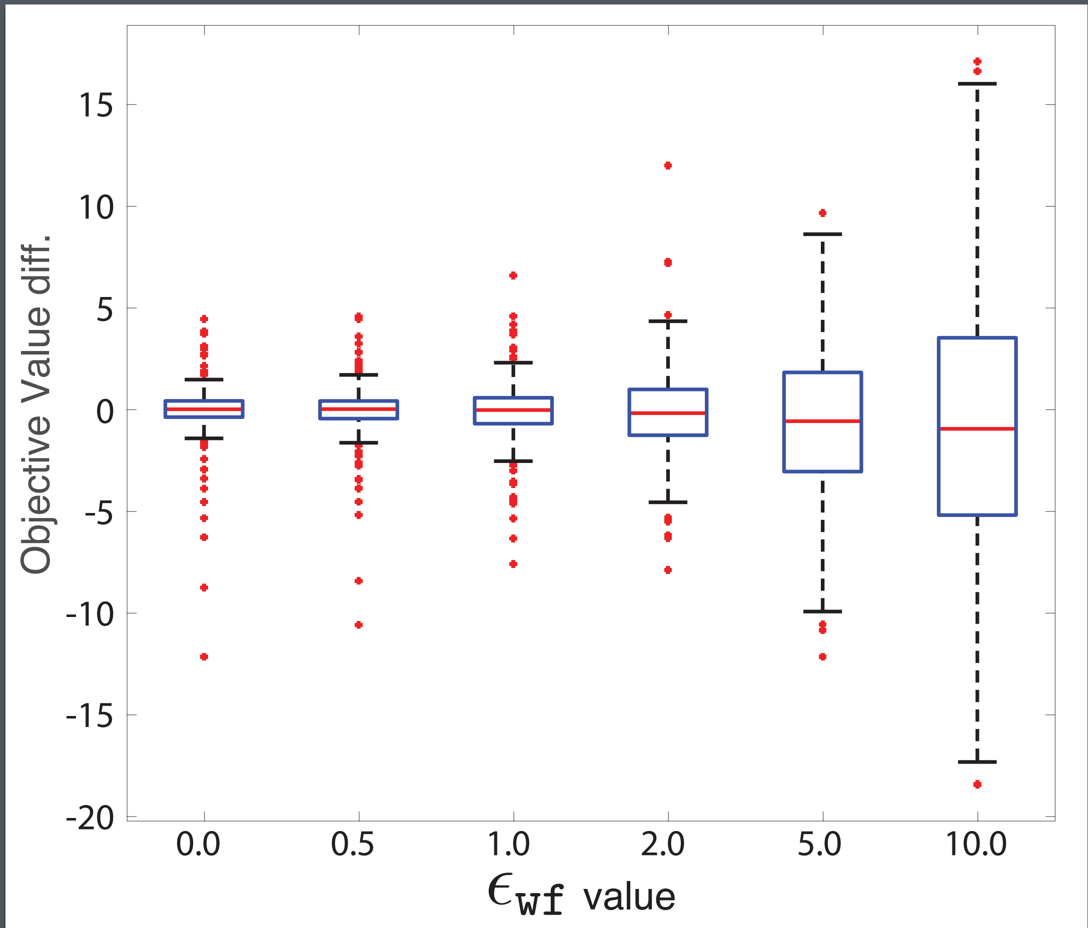

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

- BSP formulation
- Related work
- iX-BSP
  - Results
- The Wildfire Assumption
  - Results
- iML-BSP
  - Results - simulation
  - Results - live

Research Outline

- Introducing Joint Inference & Planning
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

- BSP formulation
- Related work
- iX-BSP
  - Results
- The Wildfire Assumption
  - Results
- iML-BSP
  - Results - simulation
  - Results - live

iML-BSP

- As iX-BSP was formulated over the original un-approximated problem of X-BSP, we believe it can also benefit all existing approximations of X-BSP
- To support this claim we introduce the ML approximation to iX-BSP, and denote the result as iML-BSP

$$J^{iML}(u') \approx \sum_{i=k+l+1}^{k+l+L} \left[w_i \cdot r_i \left(b[X_{i|k+l}], u'_{i-1|k+l} \right) \right]$$

$$w_i = \frac{\mathbb{P}(z_{k+l+1:i} | H_{k+l|k+l}, u_{k+l:i-1|k+l})}{q(z_{k+l+1:i})}$$

Nominal distribution

Sampled distribution

Research Outline

- Introducing Joint Inference & Planning
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

- BSP formulation
- Related work
- iX-BSP
  - Results
- The Wildfire Assumption
  - Results
- iML-BSP
  - Results - simulation
  - Results - live

Research Outline

- Introducing Joint Inference & Planning
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

- BSP formulation
- Related work
- iX-BSP
    - Results
- The Wildfire Assumption
    - Results
- iML-BSP
    - Results - simulation
    - Results - live

Results - simulation

- We compare planning time of iML-BSP and standard BSP using ML (ML-BSP).
- We used a randomly generated map, with two goals.
- The robot is required to visit both goals with an objective that minimize D2G and maximize information gain
- We ran 1000 rollouts (entire mission run), each with a different sampled initial ground truth position.
- The robot is equipped with a stereo camera and has no prior knowledge over the environment.

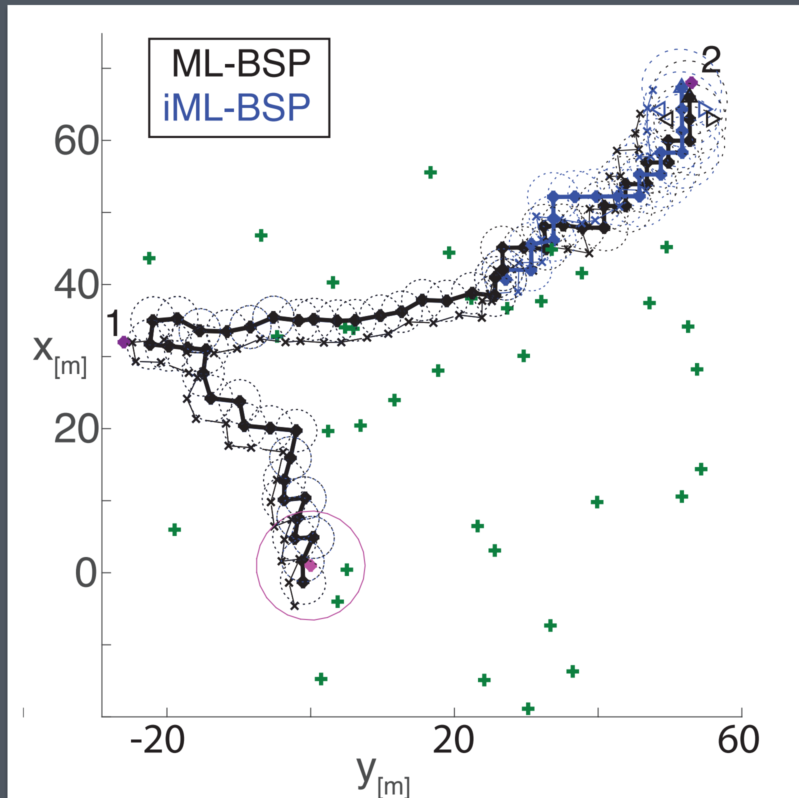

- We considered known models with Gaussian additive noise

Results - simulation

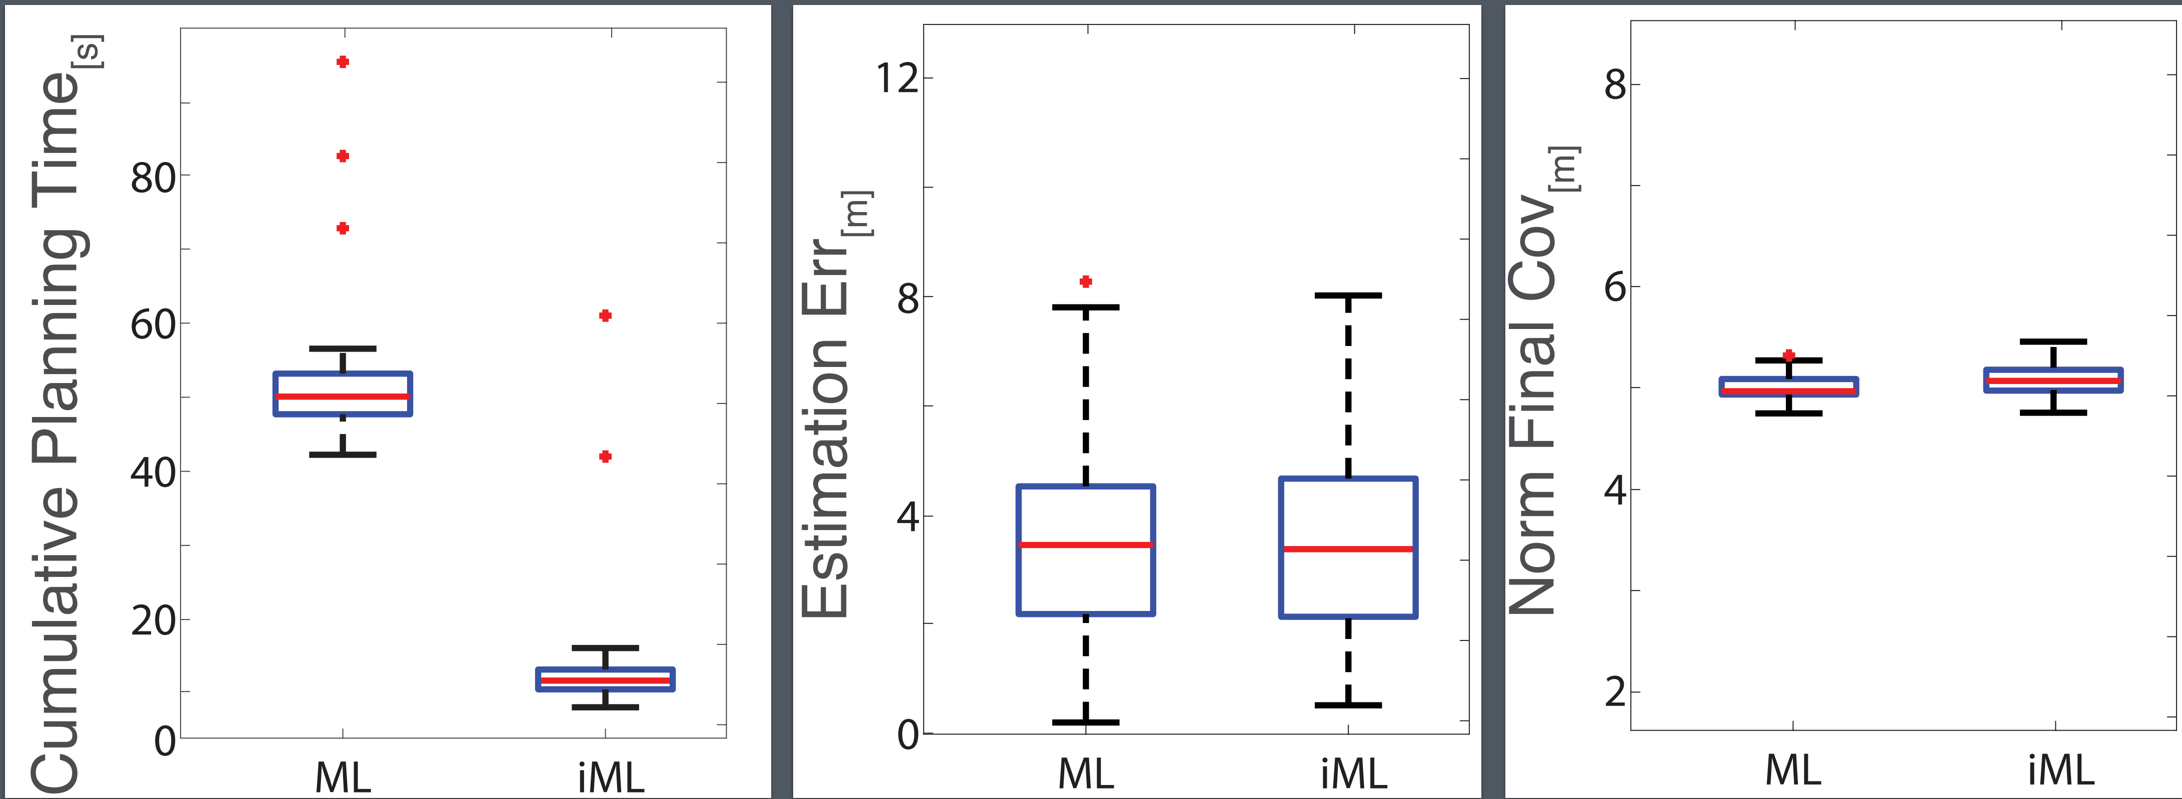

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

- BSP formulation
- Related work
- iX-BSP
  - Results
- The Wildfire Assumption
  - Results
- iML-BSP
  - Results - simulation
  - Results - live

Research Outline

- Introducing Joint Inference & Planning
- RUBI: Re-Use BSP for Inference update
- iX-BSP: incremental eXpectation BSP
- Concluding remarks
- Q&A

- BSP formulation
- Related work
- iX-BSP
  - Results
- The Wildfire Assumption
  - Results
- iML-BSP
  - Results - simulation
  - Results - live

iML-BSP live experiments

- We compare planning time of iML-BSP and ML-BSP
- We used the pioneer 3AT robot, equipped with ZED stereo camera and Hokuyo UTM-30LX Lidar
- The robot is required to visit set of goals with an objective that minimize D2G and maximize information gain
- We ran two experiments, 35_m and 148_m long.
- The robot has no prior knowledge over the environment, and no usage of offline calculations
- We considered known models with Gaussian additive noise

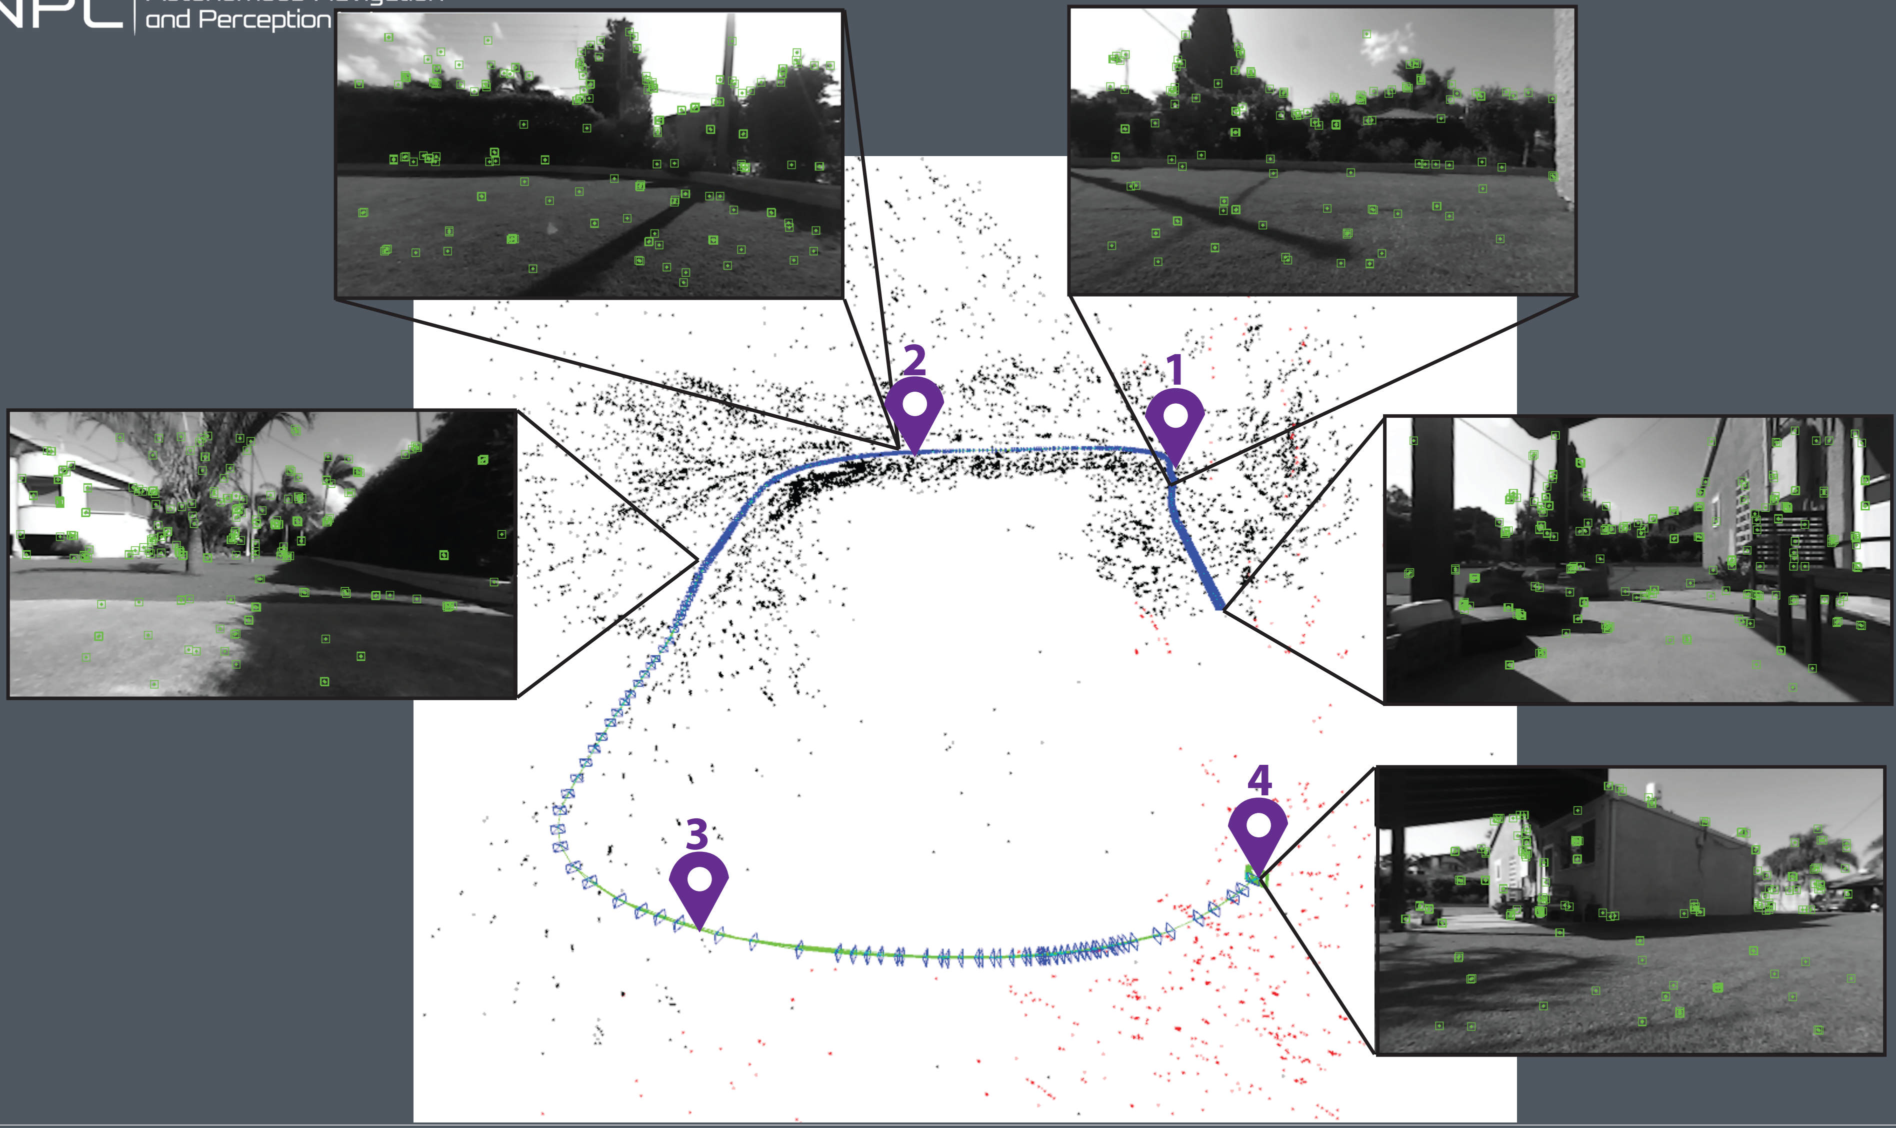
2
1
3
4

Planning computation time - 35_m run

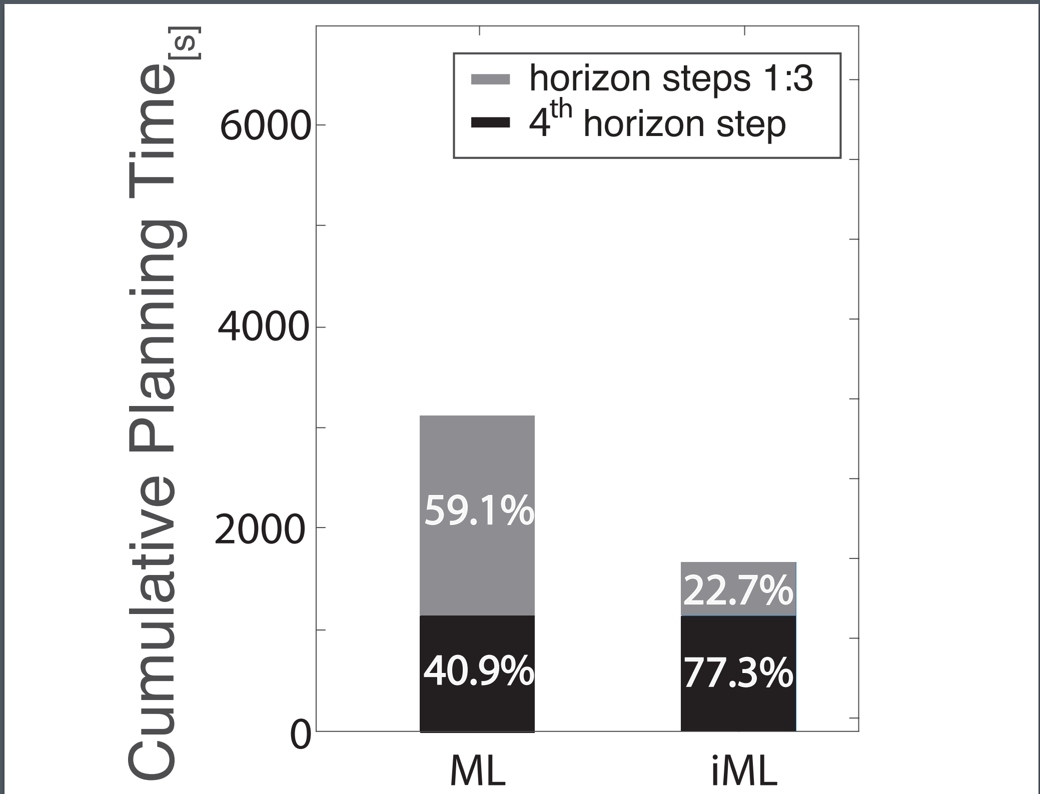

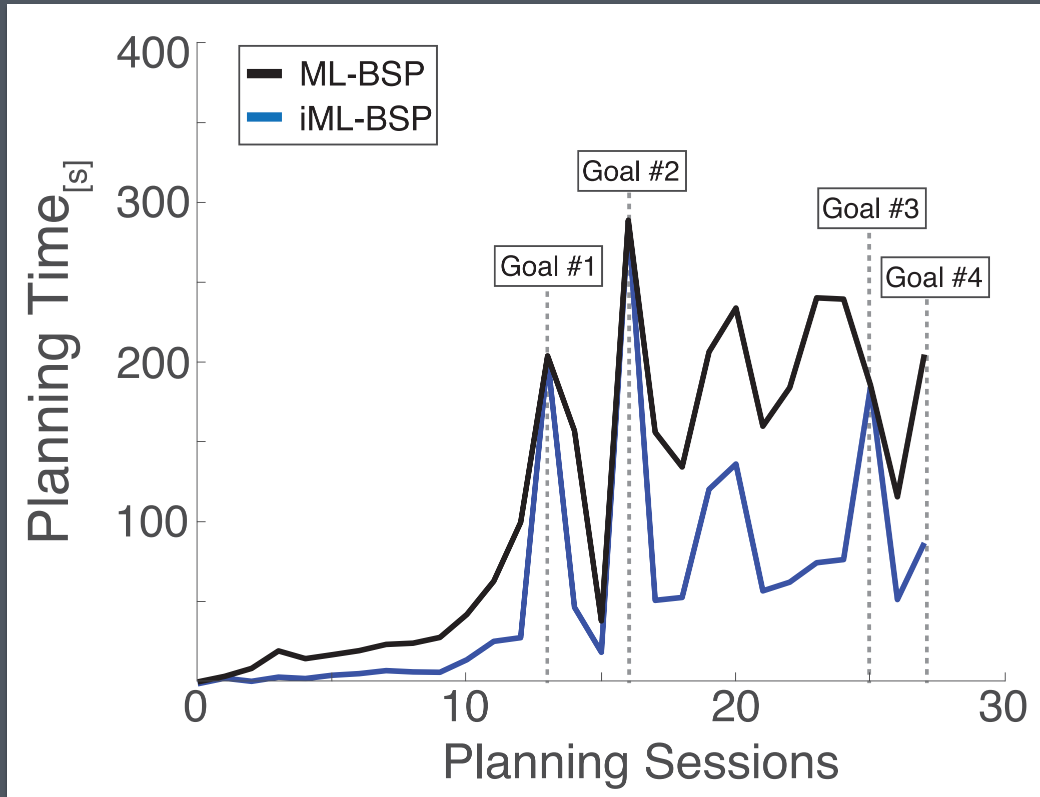

Involved Factors - 35_m run

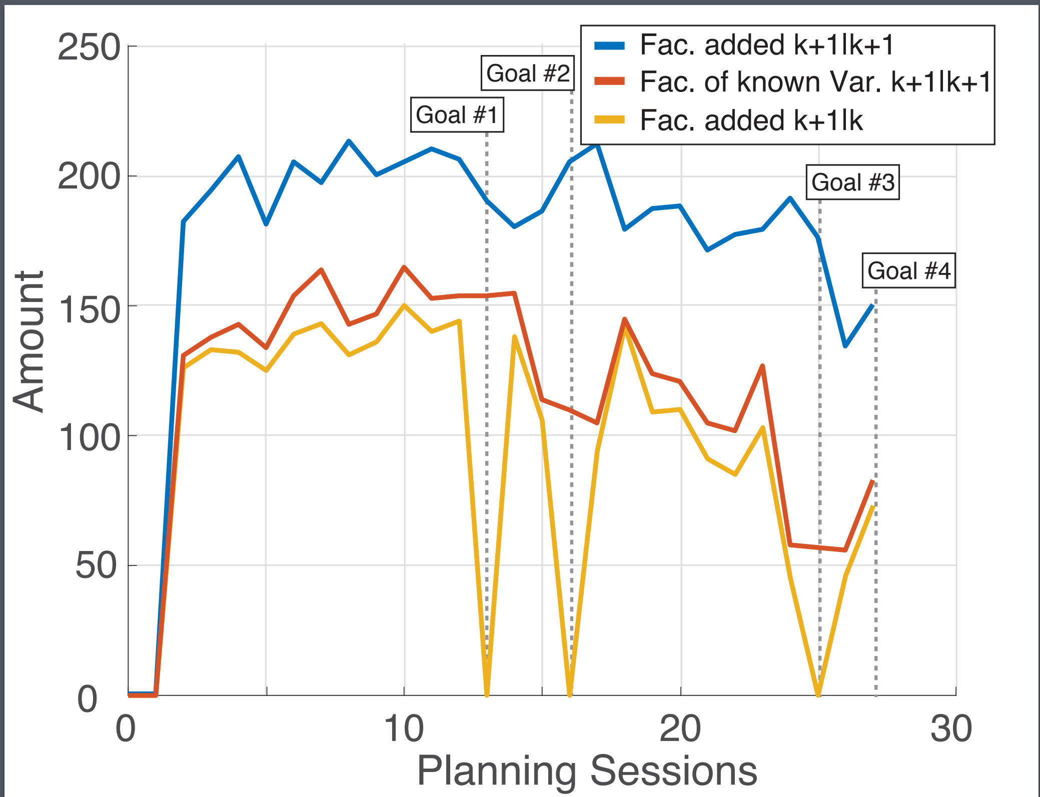

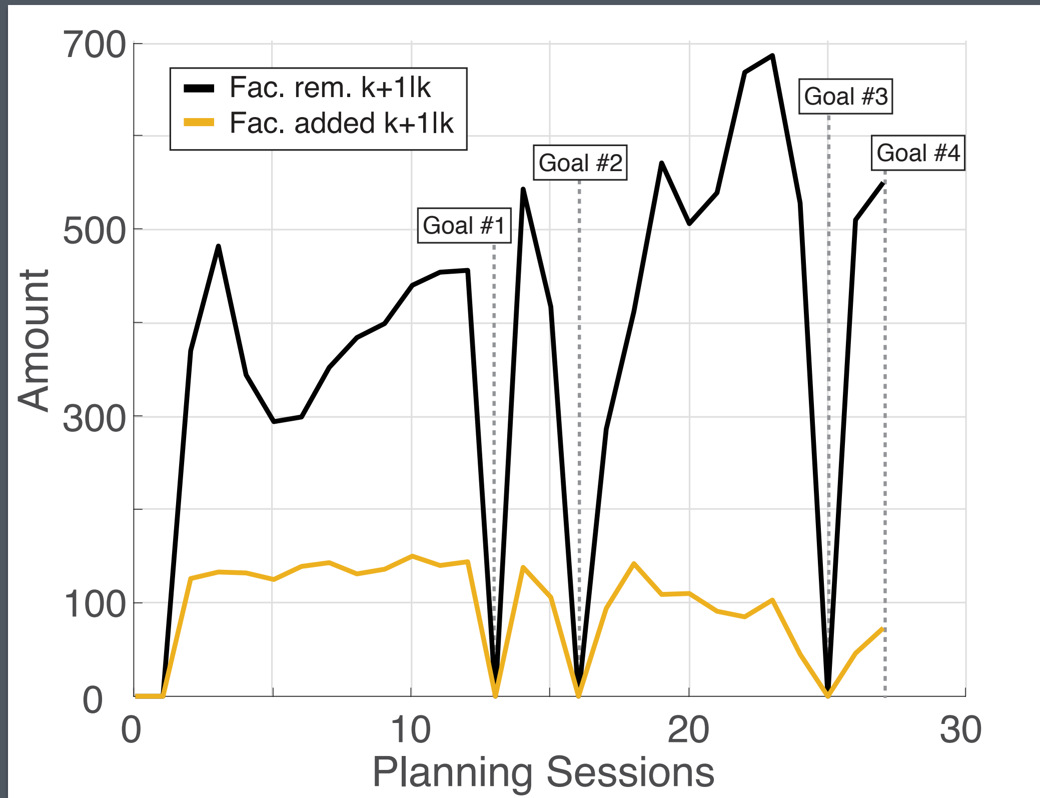

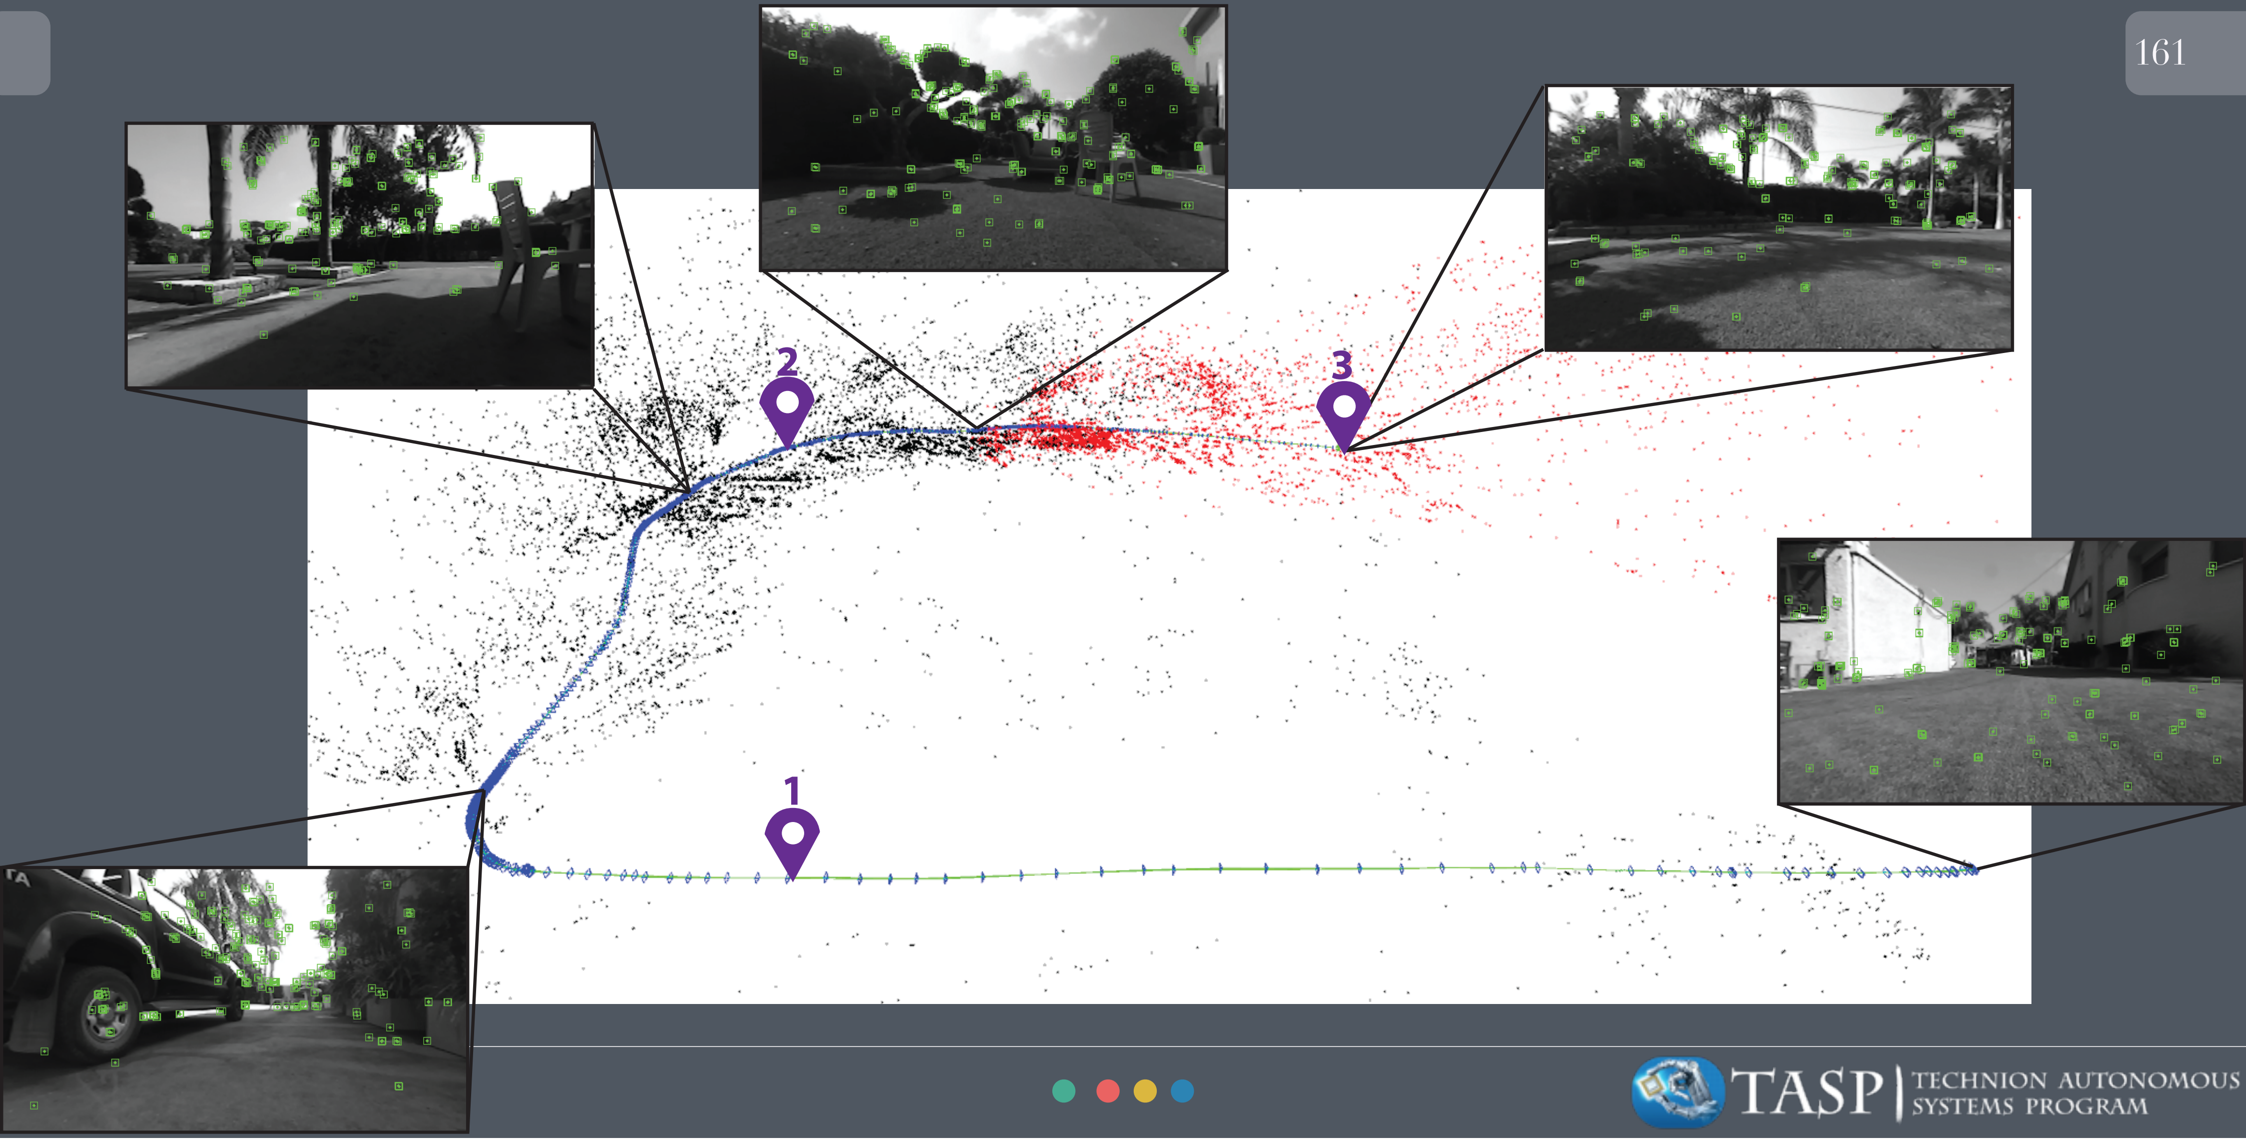
2
3
1

Planning computation time - 148_m run

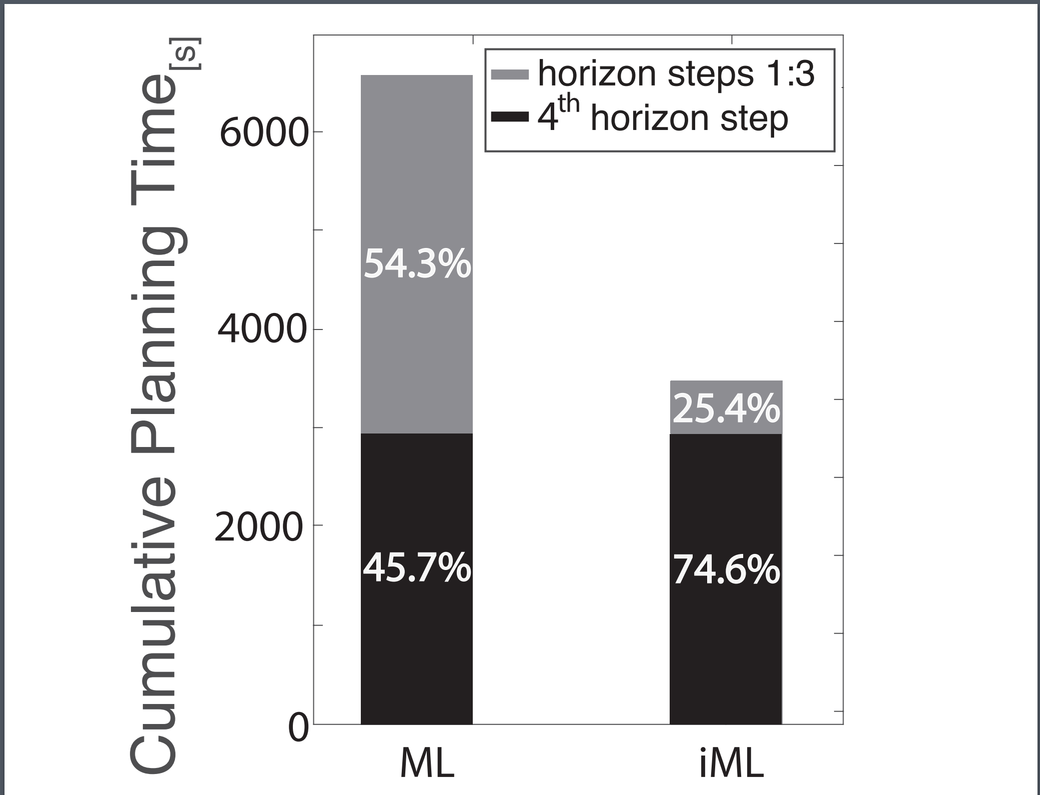

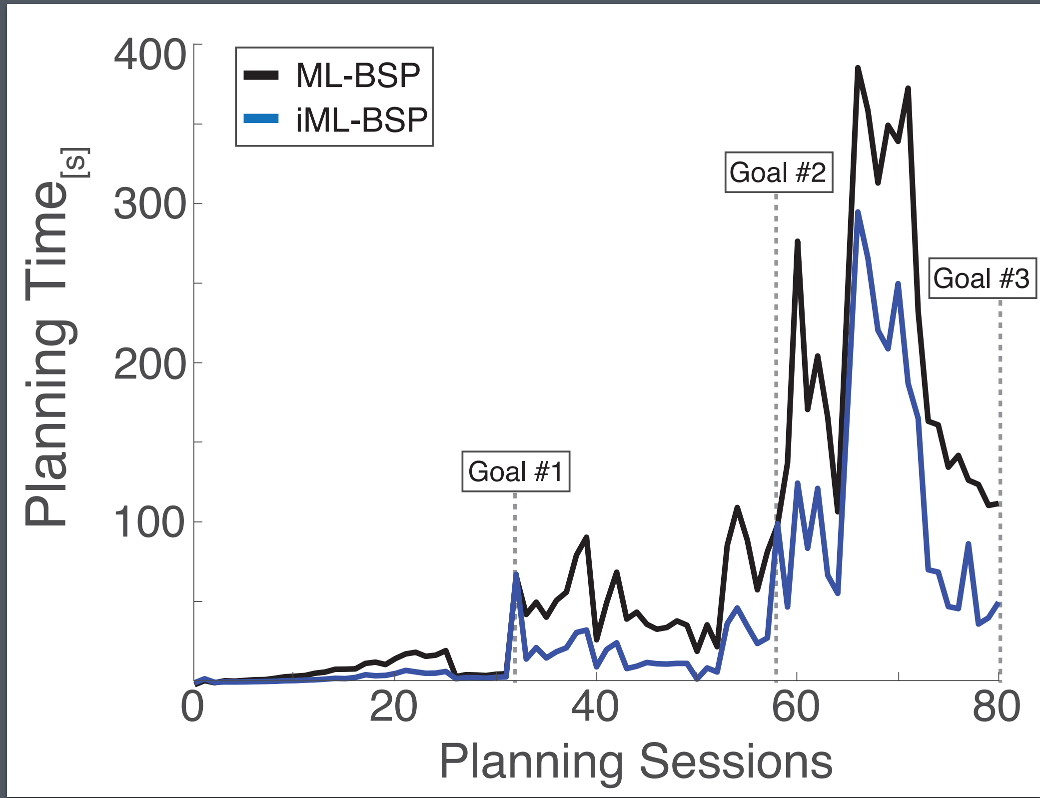

Involved Factors - 148_m run

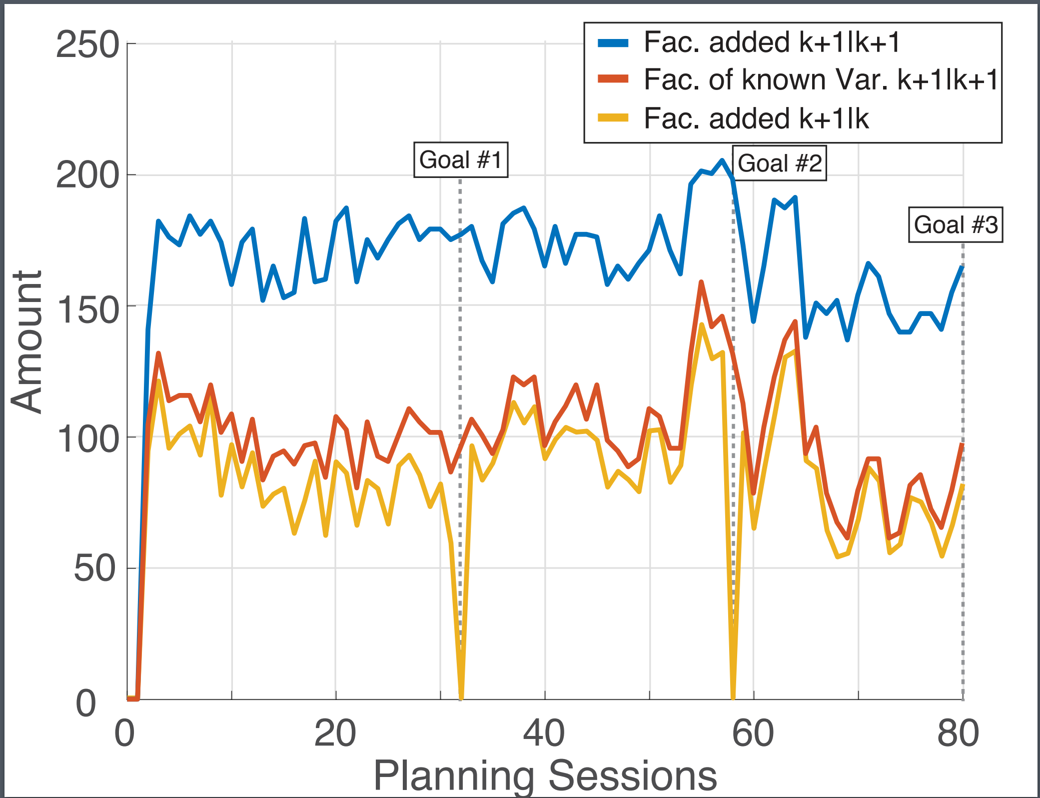

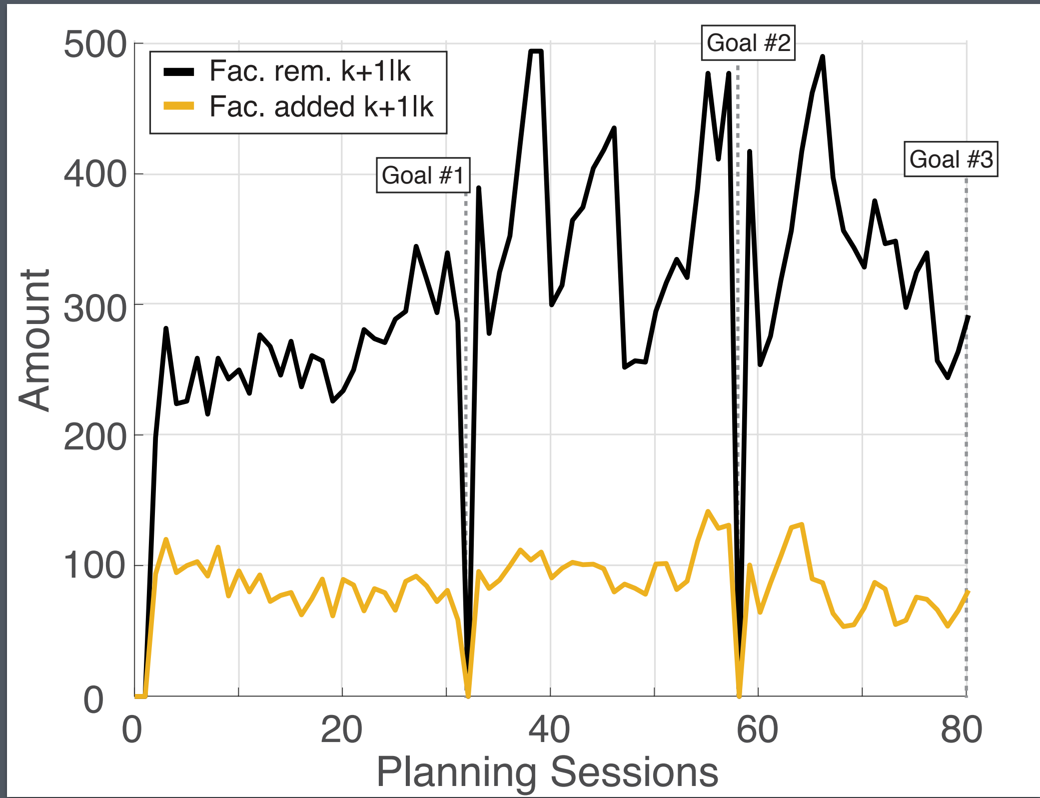

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

- BSP formulation
- Related work
- iX-BSP
  - Results
- The Wildfire Assumption
  - Results
- iML-BSP
  - Results - simulation
  - Results - live

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

Concluding remarks

- We introduced the novel concept of Joint Inference & Planning - JIP
- Inspired by JIP, we created two new approaches for inference update (RUBI) and BSP (iX-BSP)
  - RUBI provides efficient inference update using precursory planning session calculations.
  - iX-BSP provides efficient BSP by incrementally updating previous planning calculations.
- JIP, consisting of RUBI and iX-BSP, provides with an exact solution to the original standard plan-act-infer system, with a reduced computational effort
- This new approach of “symbiosis" might also pave the way into abilities we have yet to discov

Research Outline

Introducing Joint Inference & Planning

RUBI: Re-Use BSP for Inference update

iX-BSP: incremental eXpectation BSP

Concluding remarks

Q&A

Research Outline

Q & A Session

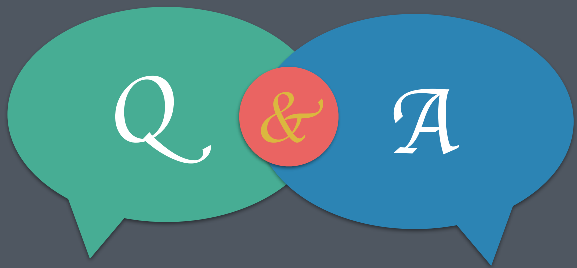

Thanks for Listening

We'll be answering Questions Now