Autonomous Classification Under Uncertainty

PhD Seminar

Vladimir Tchuiev

Supervised by **Vadim Indelman**

Aerospace Engineering, Technion Israel Institute of Technology

Introduction: Object Classification

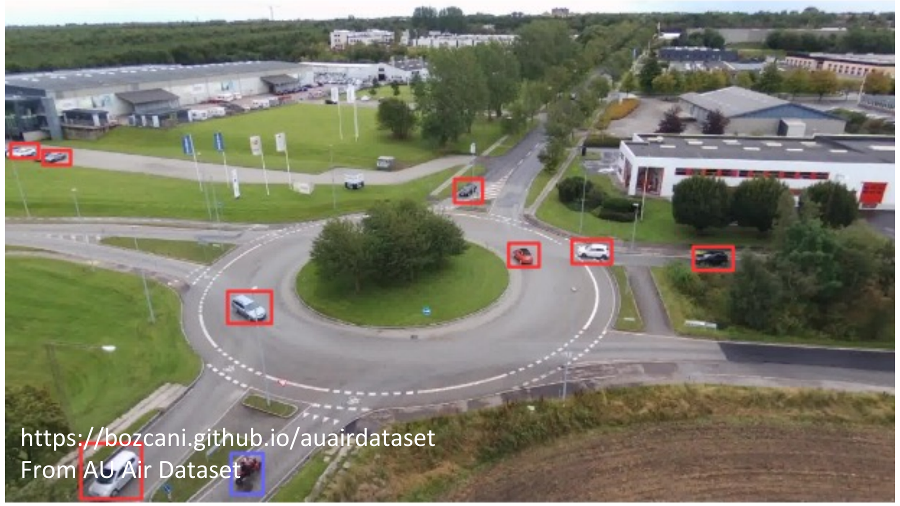

- **Object classification** is an important problem for autonomous vehicles and UAVs.
- Notable advancement in recent years with **deep learning and neural networks**.
- **Reliable** classification remains a challenge.

Introduction: Uncertainties in Object Classification

- Multiple factors affect classification accuracy:
  - Lighting
  - Occlusions
  - Resolution
  - **Viewpoint Dependency**
  - **Classifier epistemic uncertainty**
- **Viewpoint dependency:** certain relative viewpoints might introduce classification aliasing.
- **Epistemic uncertainty:** test data does not match the classifier's training data.

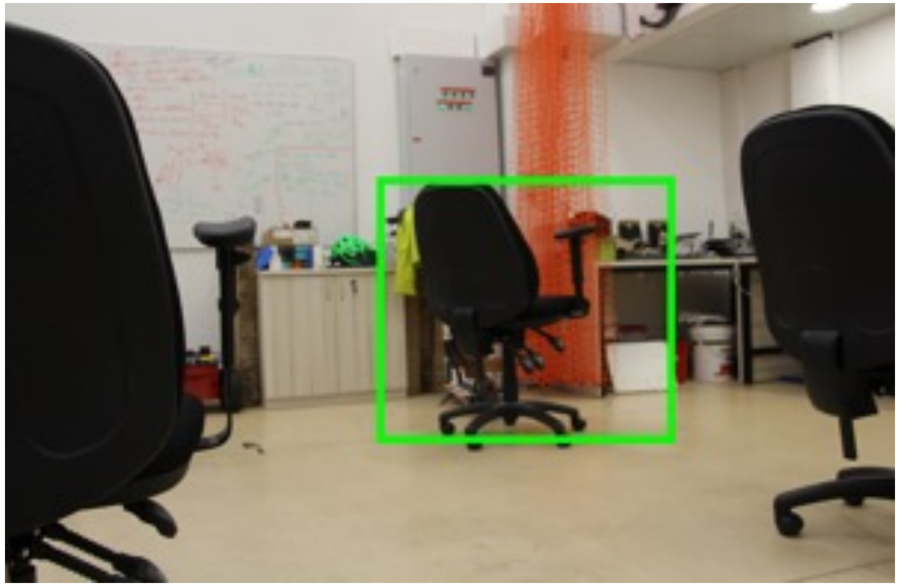

Introduction: Simultaneous localization and mapping (SLAM)

- Given measurements, construct a map of the environment and infer the robot's pose.
- Posterior Distribution:

$$\mathbb{P}(\mathbf{x}_{0:k}, \mathcal{X}^{o} | \mathbf{Z}_{1:k}, \mathbf{a}_{0:k-1})$$

Robot poses | Object poses | Measurements | Actions

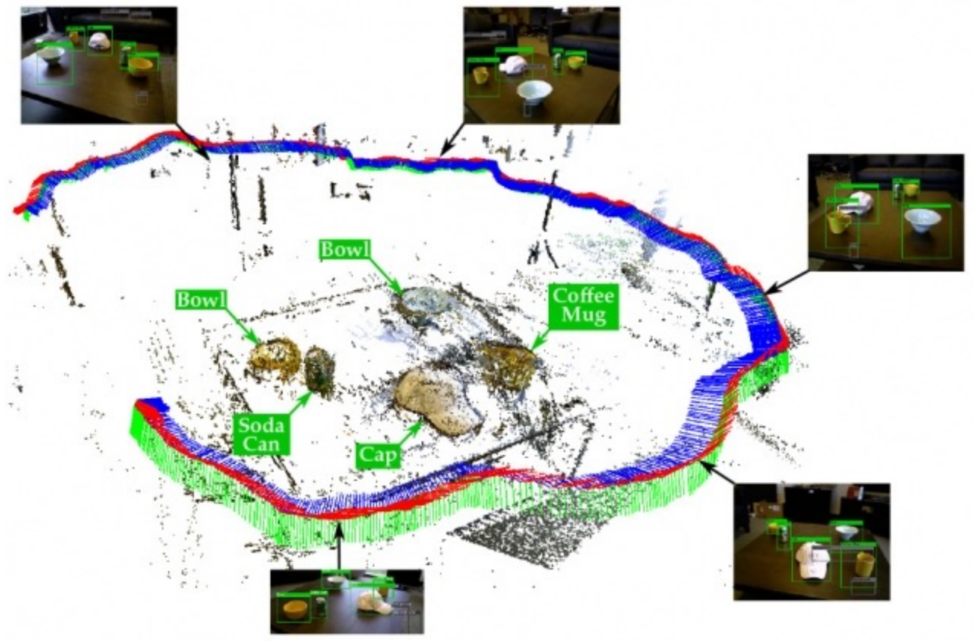

Pillai, Sudeep, and John Leonard. "Monocular slam supported object recognition." *arXiv preprint arXiv:1506.01732* (2015).

Introduction: SLAM

- Using Bayes rule and chain rule:

$$\mathbb{P}(x_{0:k}, \mathcal{X}^o | Z_{1:k}, a_{0:k-1}) \propto \mathbb{P}(x_0, \mathcal{X}^o) \prod_{t=1}^{k} \mathbb{P}(x_t | x_{t-1}, a_{t-1})\, \mathbb{P}(Z_t | x_t, \mathcal{X}^o)$$

- $\mathbb{P}(x_0, \mathcal{X}^o)$ - pose priors.
- $\mathbb{P}(x_t | x_{t-1}, a_{t-1})$ - motion model.
- $\mathbb{P}(Z_t | x_t, \mathcal{X}^o)$- measurement likelihood, where **data association (DA)** is important.
- **Data association:** assigning measurement to object/landmark.
- If **Gaussian**,$\mathbb{P}(x_{0:k}, \mathcal{X}^o | Z_{1:k}, a_{0:k-1})$ is computed via methods such as **iSAM2**.

Presentation Overview

- Data association aware semantic SLAM via viewpoint dependent classifier model (published in IROS 2019)
- Distributed semantic SLAM via viewpoint dependent classifier model (published in RAL/IROS 2020)
- Epistemic uncertainty aware sequential classification (published in RAL/IROS 2018)
- Posterior epistemic uncertainty aware inference and belief space planning (upcoming paper 2021)

Presentation Overview

- Data association aware semantic SLAM via viewpoint dependent classifier model (published in IROS 2019)
- Distributed semantic SLAM via viewpoint dependent classifier model (published in RAL/IROS 2020)
- Epistemic uncertainty aware sequential classification (published in RAL/IROS 2018)
- Posterior epistemic uncertainty aware inference and belief space planning (upcoming paper 2021)

DA Aware Semantic SLAM: Definitions and Problem formulation

- **Setting:** a robot observes objects within the environment, receiving:
  - Geometric measurements. E.g., range and bearing.
  - Semantic measurements of class probability vectors.
- **Key challenges:**
  - Classification aliasing.
  - DA aliasing.
- **We aim to maintain the hybrid belief:**

$$\mathbb{P}(\mathcal{X}_k, C, \beta_{1:k} | \mathcal{H}_k)$$

Robot and object poses | Object classes | DA realization | Measurements and actions

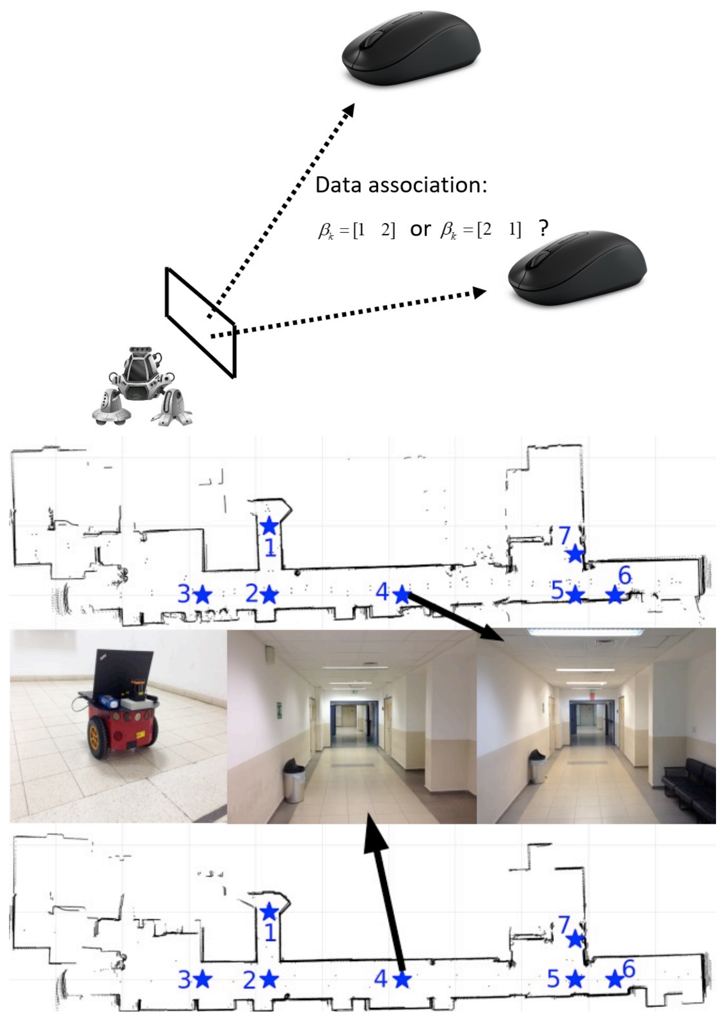

Pathak, Shashank, Antony Thomas, and Vadim Indelman. "A unified framework for data association aware robust belief space planning and perception." The International Journal of Robotics Research 37, no. 2-3 (2018): 287-315.

DA Aware Semantic SLAM: Contribution

We present an approach that:

- Maintains a ***hybrid belief*** over:
  - Robot and object poses.
  - Object classes.
  - DA hypotheses.
- Address **coupling** between classification and SLAM problem via a **viewpoint dependent classifier model**.

Leveraging the coupling between poses and classes to:

- Assist in **data association disambiguation**.
- Improve **classification** and **localization** performance.

Previous works:

- Consider **most likely class** semantic measurements.
- Utilize a viewpoint dependent classifier model with **solved data association**.

Published paper: Tchuiev, Vladimir, Yuri Feldman, and Vadim Indelman. "Data Association Aware Semantic Mapping and Localization via a Viewpoint-Dependent Classifier Model." In *IROS*, pp. 7742-7749. 2019.

DA Aware Semantic SLAM: Assumptions

- A single robot within a **static environment.**
- A **known** number of objects.
- Models: motion $\mathbb{P}(x_k|x_{k-1}, a_{k-1})$, geometric $\mathbb{P}(Z_k^{geo}|\mathcal{X}_k, \beta_k)$, and classifier $\mathbb{P}(Z_k^{sem}|\mathcal{X}_k, C, \beta_k)$, are **Gaussian.**
- The **object observation** model $\mathbb{P}(\beta_k|x^{rel})$ determines if DA realization is feasible given relative pose.

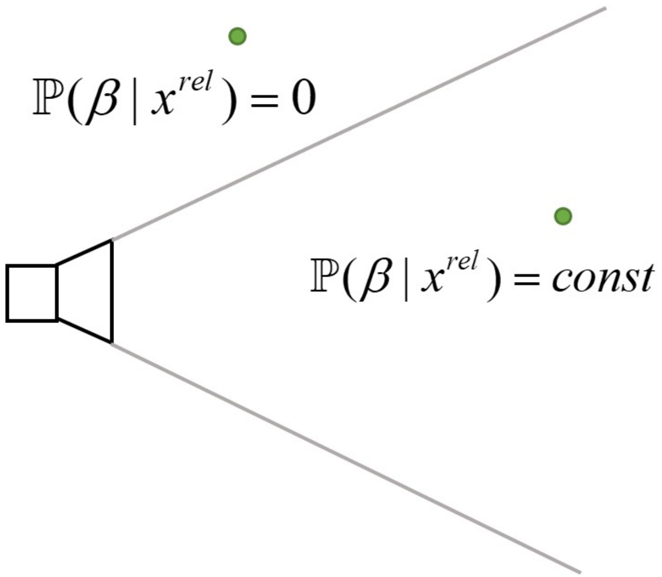

DA Aware Semantic SLAM:
The Classifier Model

- $z_k^{sem} \in \mathbb{R}^M$ is viewpoint dependent.
- The model is assumed Gaussian $\mathbb{P}\left(z_k^{sem} \middle| c, x^{rel}\right) = \mathcal{N}(h_c, \Sigma_c)$ where $h_c\left(x^{rel}\right)$ and $\Sigma_c\left(x^{rel}\right)$ depend on object class c and relative pose x^{rel}.

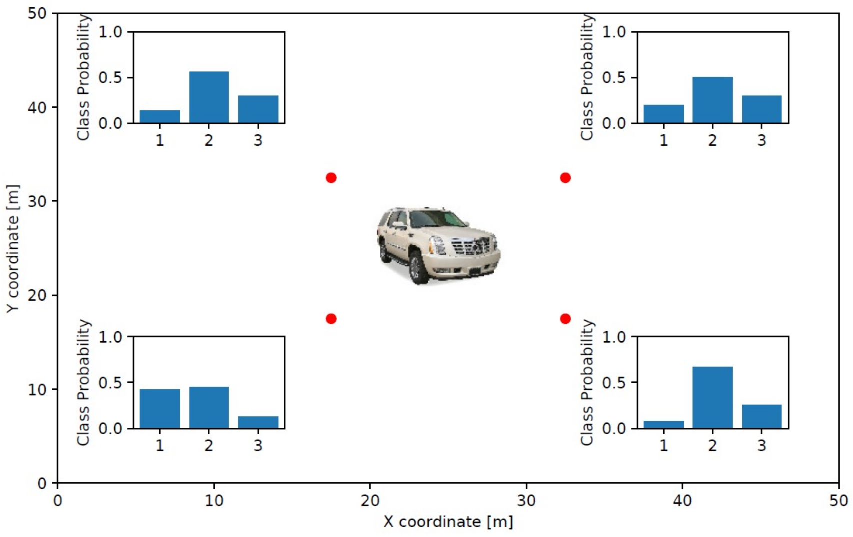

DA Aware Semantic SLAM: General Approach

- Split the **hybrid belief** to continuous and discrete parts:

$$\mathbb{P}(\mathcal{X}_k, C, \beta_{1:k} | \mathcal{H}_k) = \underbrace{\mathbb{P}(\mathcal{X}_k | C, \beta_{1:k}, \mathcal{H}_k)}_{b^C_{\beta_{1:k}}[\mathcal{X}_k]} \underbrace{\mathbb{P}(C, \beta_{1:k} | \mathcal{H}_k)}_{w^C_{\beta_{1:k}}}$$

- $b^C_{\beta_{1:k}}[\mathcal{X}_k]$ is the **continuous** belief given class and DA realization.
- $w^C_{\beta_{1:k}}$ is the **weight** of $b^c_{\beta_{1:k}}[\mathcal{X}_k]$, computed separately for each C and $\beta_{1:k}$.

DA Aware Semantic SLAM: Belief Update

- Continuous belief update:

$$b^{C}_{\beta_{1:k}}[\mathcal{X}_k] \propto b^{C}_{\beta_{1:k-1}}[\mathcal{X}_{k-1}] \cdot \mathbb{P}(x_k|x_{k-1}, a_{k-1}) \cdot \mathbb{P}(\mathcal{Z}_k|X_k, C, \beta_k)$$

- Weight update:

$$w^{C}_{\beta_{1:k}} \propto w^{C}_{\beta_{1:k-1}} \int_{\mathcal{X}_k} \mathbb{P}(\beta_k|\mathcal{X}_k) \cdot b^{C}_{\beta_{1:k}}[\mathcal{X}_k] d\mathcal{X}_k$$

- Small weights are **pruned** to keep the **number** of realizations **small**.
- **Viewpoint dependent classifier model** in $\mathbb{P}(\mathcal{Z}_k|X_k, C, \beta_k)$ assists in inference DA, and **reduces** the number of realizations when pruned.

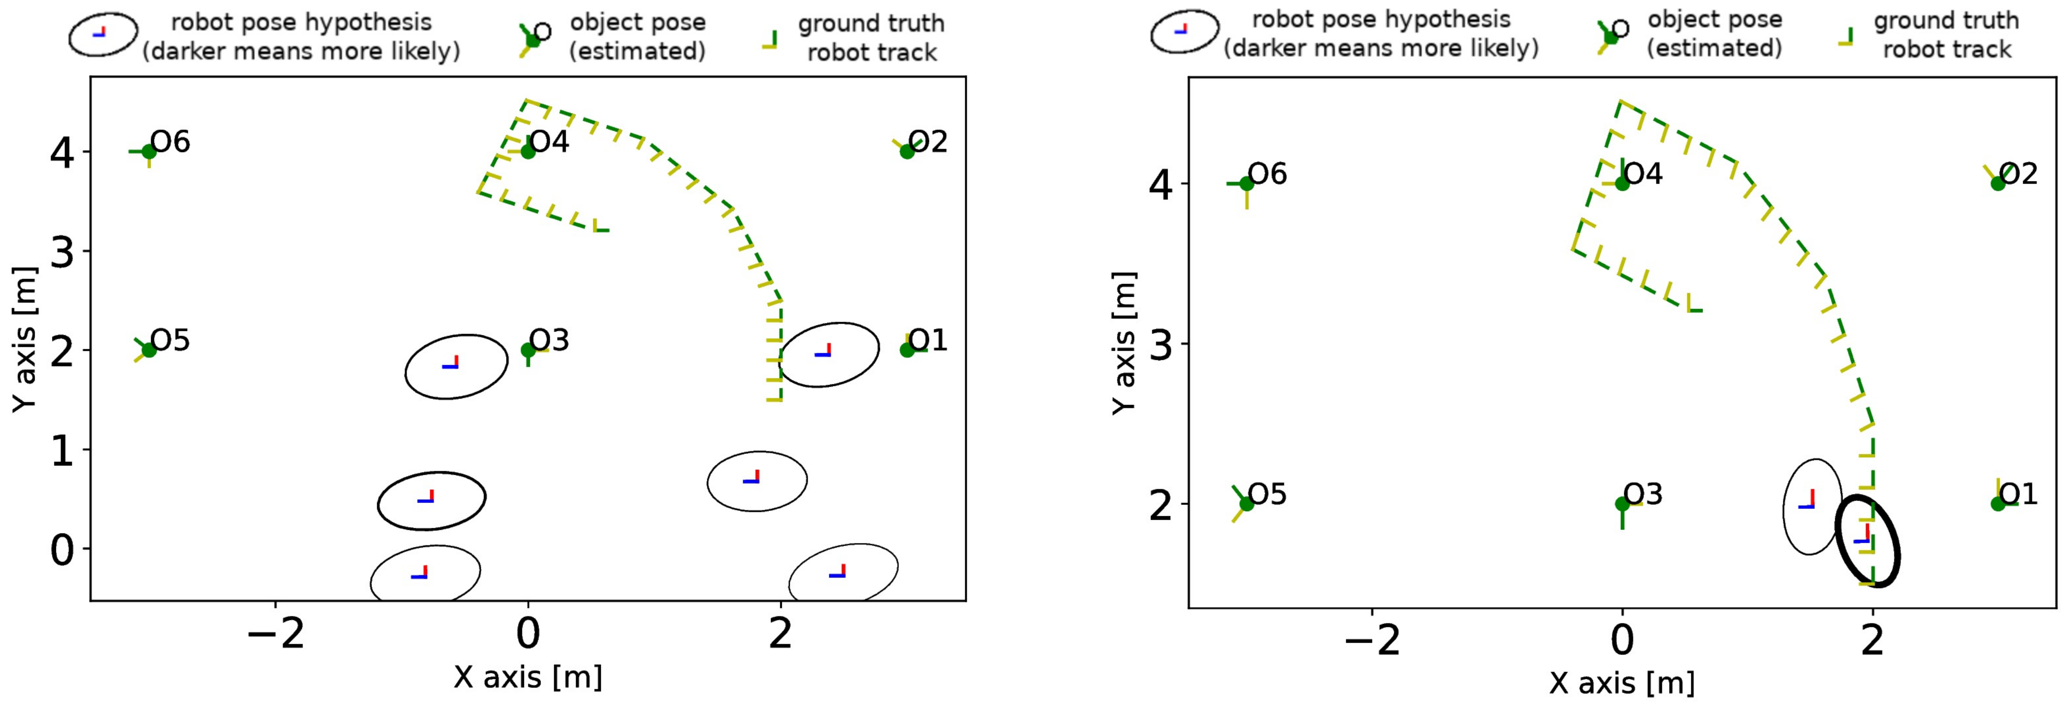

DA Aware Semantic SLAM: Simulation

- Comparison between **without** and **with** classifier model.
- **Highly aliased** scenario with 6 identical objects with different orientations.
- Uninformative prior on initial robot pose, causing **multiple probable hypotheses.**

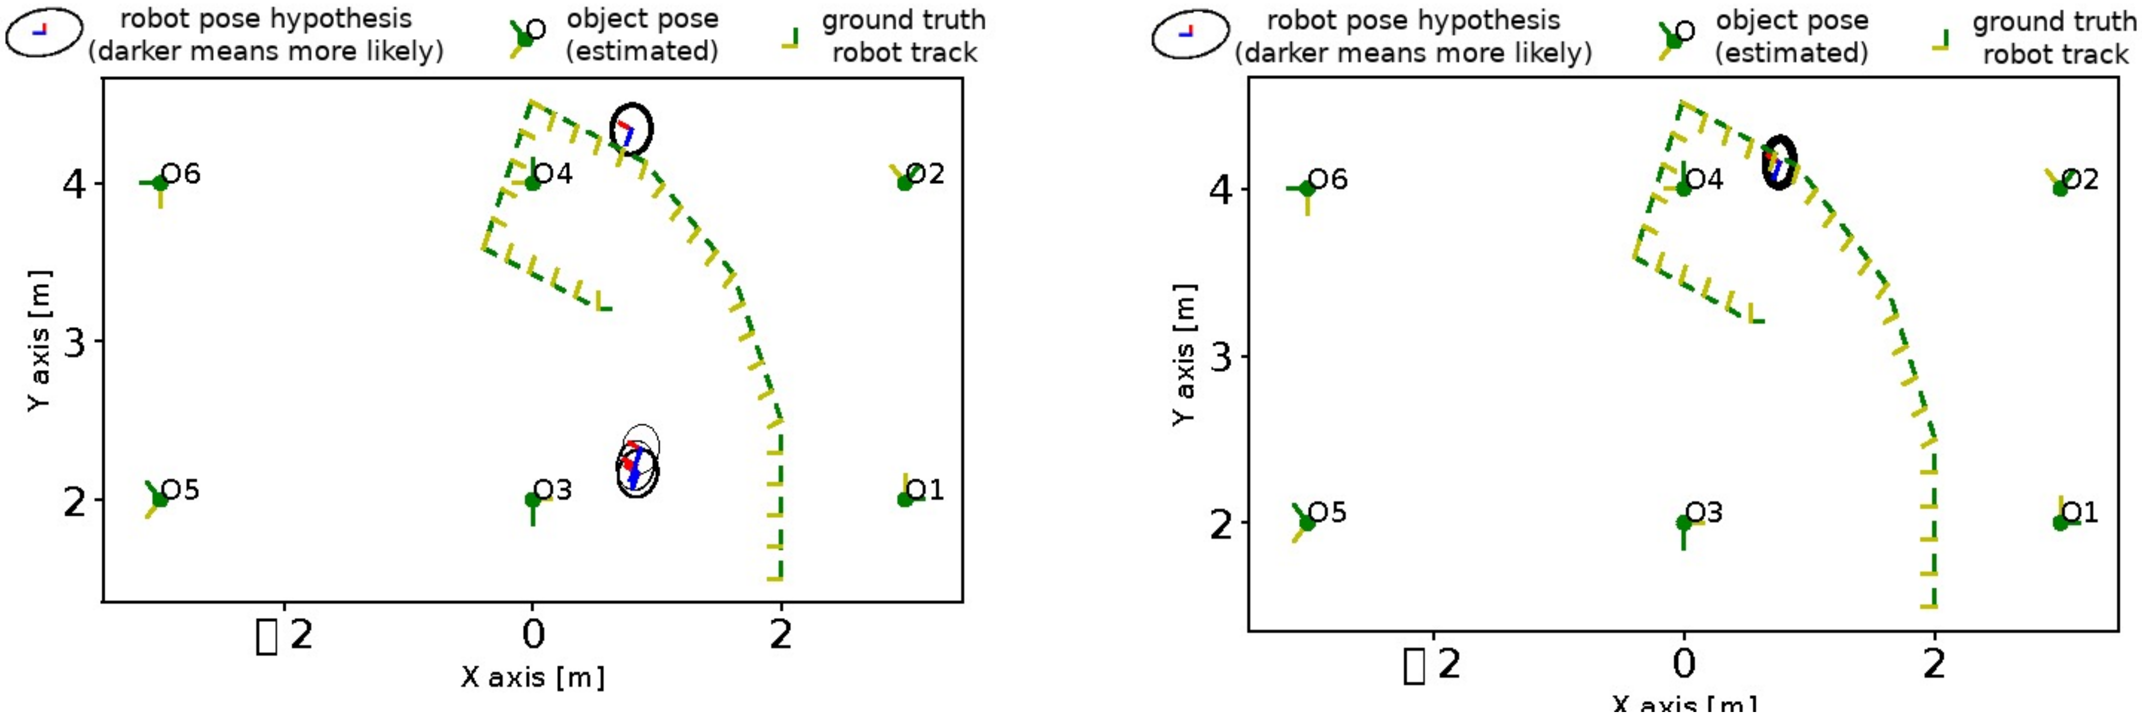

DA Aware Semantic SLAM: Simulation

❖ With classifier:
- ✓ **Fewer** belief components.
- ✓ **More accurate** localization.

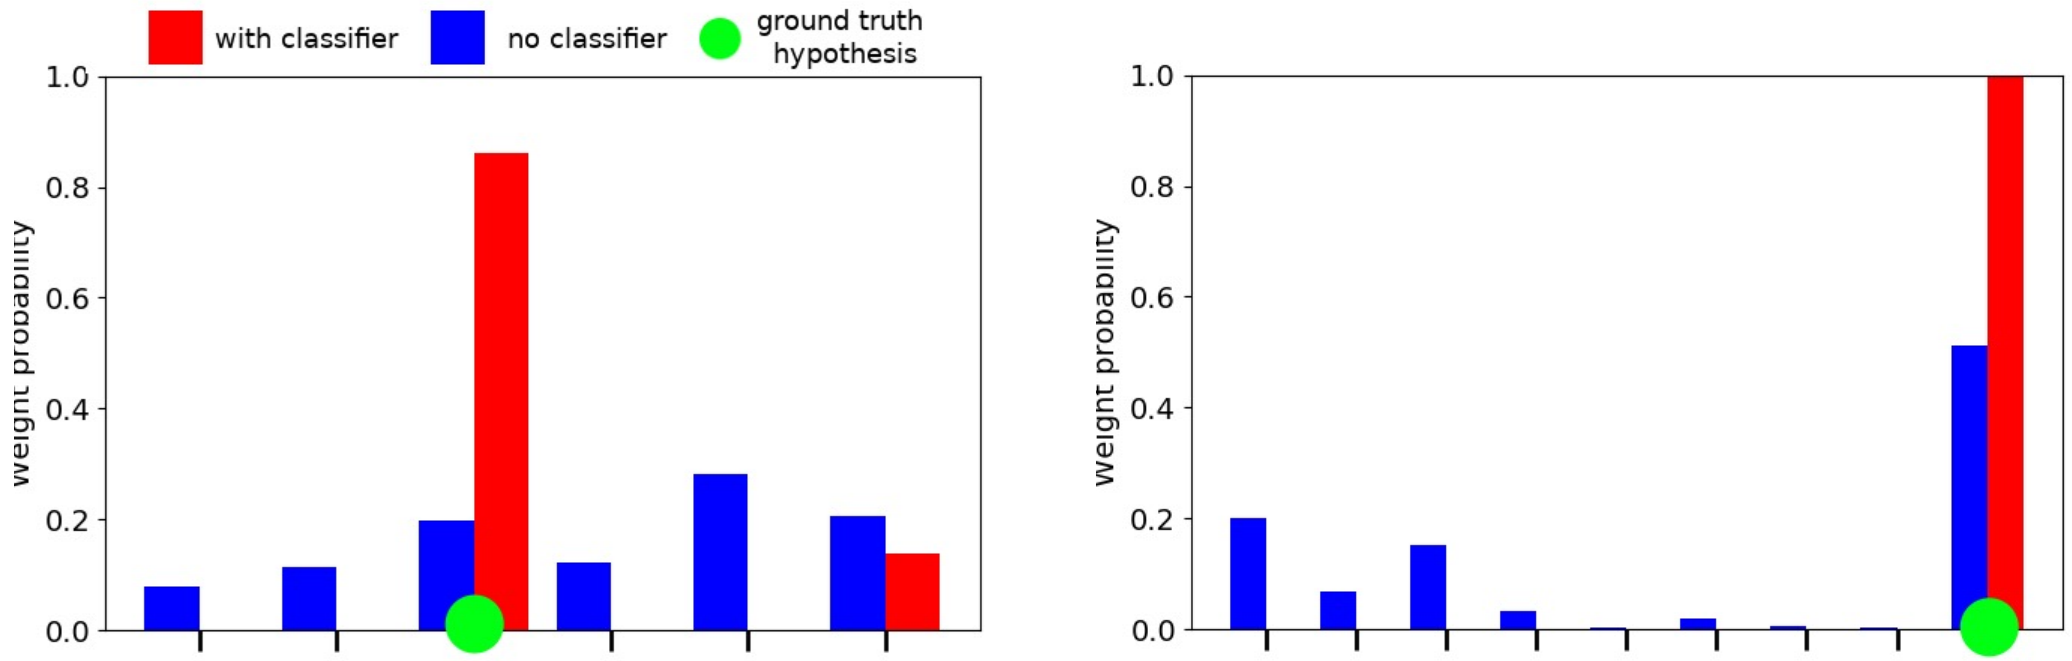

DA Aware Semantic SLAM: Simulation

- With classifier:
  - ✓ **Fewer** belief components.
  - ✓ **Stronger** disambiguation.

Presentation Overview

- Data association aware semantic SLAM via viewpoint dependent classifier model (published in IROS 2019)
- Distributed semantic SLAM via viewpoint dependent classifier model (published in RAL/IROS 2020)
- Epistemic uncertainty aware sequential classification (published in RAL/IROS 2018)
- Posterior epistemic uncertainty aware inference and belief space planning (upcoming paper 2021)

Distributed Semantic SLAM: Problem and Notations

- **Setting:** multiple robots observe objects within the environment, receiving:
  - Geometric measurements. E.g., range and bearing.
  - Semantic measurements of class probability vectors.
- **Key challenges:**
  - Classification aliasing.
  - Estimation consistency.
- DA is assumed solved.
- ***We aim to maintain joint hybrid belief:***

Robot Group R

$$\mathbb{P}(\mathcal{X}_k^R, C^R | \mathcal{H}_k^R)$$

Robot and object poses

Object classes

Measurements and actions

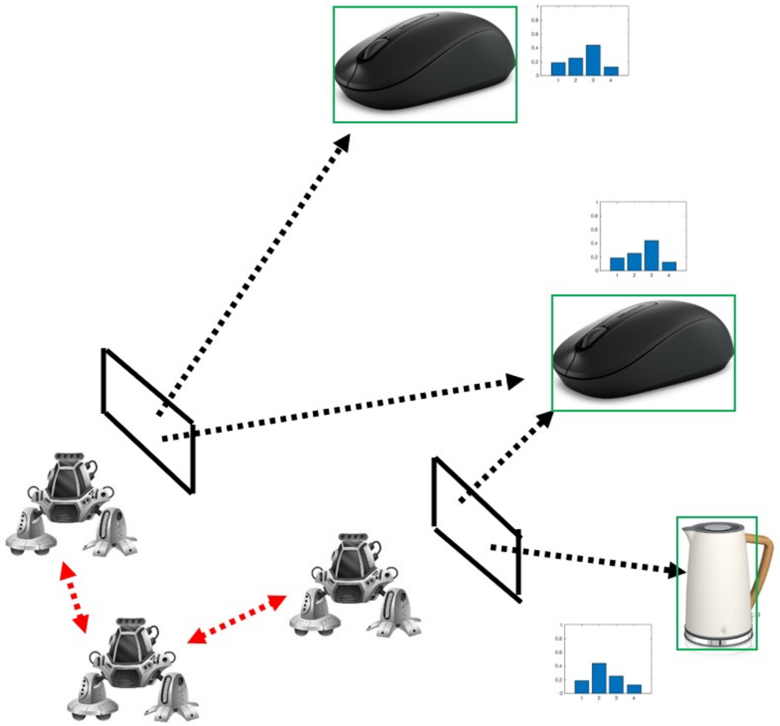

Distributed Semantic SLAM: Contribution

We present a *multi-robot* approach that:

- Maintains a *hybrid belief* over:
  - Robot and object poses.
  - Object classes.
- Address coupling between classification and SLAM problem via a viewpoint dependent classifier model.

We address estimation consistency:

- Continuous random variables.
- *Discrete* random variables.

Previous works:

- No semantic information in a multi-robot setting.
- Addressed double counting only for **continuous variables**.

Published paper: Tchuiev, Vladimir, and Vadim Indelman. "Distributed Consistent Multi-Robot Semantic Localization and Mapping." *IEEE Robotics and Automation Letters* 5, no. 3 (2020): 4649-4656.

Distributed Semantic SLAM: Double Counting

- In distributed systems, a measurement should be **counted** no more than **once**.
- **Relayed information** risks double counting.
- Double counting leads to **over-confident estimation.**
- **Example**: consider random variable c with data sets $Z_a = \{z_1, z_2\}$ and $Z_b = \{z_2, z_3\}$, the posterior is:
  - $\mathbb{P}(c|Z_a, Z_b) \propto \mathbb{P}(c)\frac{\mathbb{P}(c|z_1)\mathbb{P}(c|z_2)^2\mathbb{P}(c|z_3)}{\mathbb{P}(c|z_2)}$
  - Without the denominator $\mathbb{P}(c|z_2)$, this measurement is **double counted**.
- Double counting 'pushes' posterior to **extremes**.

Illustrated: effect of double counting on a 4 category variable with uninformative prior.

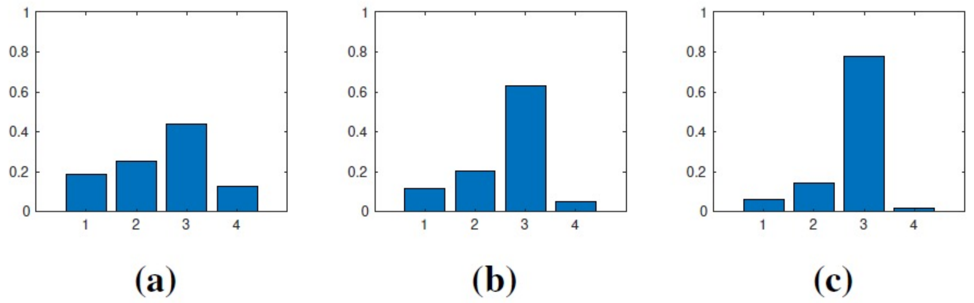

Distributed Semantic SLAM: General Approach

- Each robots maintains **two separate** hybrid beliefs:
  - Its **own** belief $\mathbb{P}(\mathcal{X}_k^r, C^r|\mathcal{H}_k^r) = \mathbb{P}(\mathcal{X}_k^r|C^r, \mathcal{H}_k^r)\mathbb{P}(C^r|\mathcal{H}_k^r)$
  - A **joint** belief $\mathbb{P}(\mathcal{X}_k^R, C^R|\mathcal{H}_k^R) = \mathbb{P}(\mathcal{X}_k^R|C^R, \mathcal{H}_k^R)\mathbb{P}(C^R|\mathcal{H}_k^R)$
- Each robot maintains a **stack** of **individual beliefs** of itself and from other robots.
- The robots **communicate the stacks** between them.
- After communication, the robots **update** the appropriate slot in the stack if the **received information is newer**.
- By **removing** the old information, the joint belief for every robot remains **consistent**.

ξ: Object pose marginals
ϕ: Object class marginals
k: Time stamp

	r_1 stack
r_1 info.	ξ, ϕ, k
r_2 info.	ξ, ϕ, k
r_3 info.	ξ, ϕ, k

	r_2 stack
r_1 info.	ξ, ϕ, k
r_2 info.	ξ, ϕ, k
r_3 info.	ξ, ϕ, k

	r_3 stack
r_1 info.	ξ, ϕ, k
r_2 info.	ξ, ϕ, k
r_3 info.	ξ, ϕ, k

Distributed Semantic SLAM: Experimental Setup

Ground Truth:

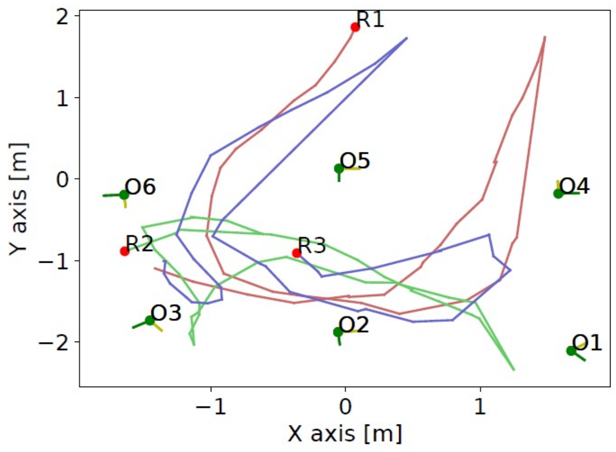

- **Scenario:** 3 robots communicating.
- 6 chairs at different orientations as objects.
- 3 candidate classes.
- Trained **classifier models**.
- *Comparing between 3 cases*:
  - Single robot.
  - Distributed.
  - With double counting.
- ***MSDE*** as classification benchmark:

$$MSDE \doteq \frac{1}{M}\sum_{i=1}^{M}\left(\mathbb{P}_{gt}(c=i) - \mathbb{P}(c=i|\mathcal{H}_k^R)\right)^2$$

Distributed Semantic SLAM: SLAM Graph Notations

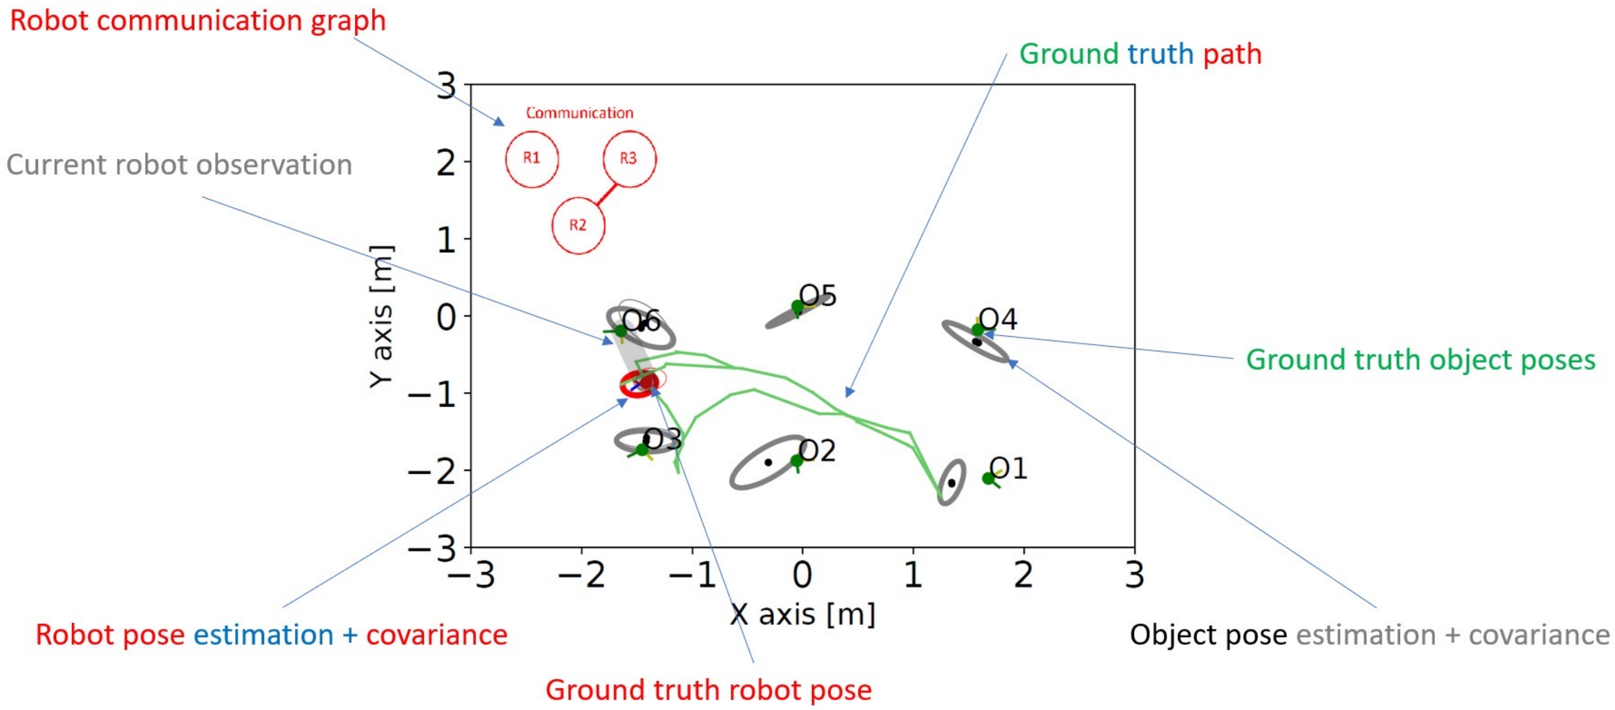

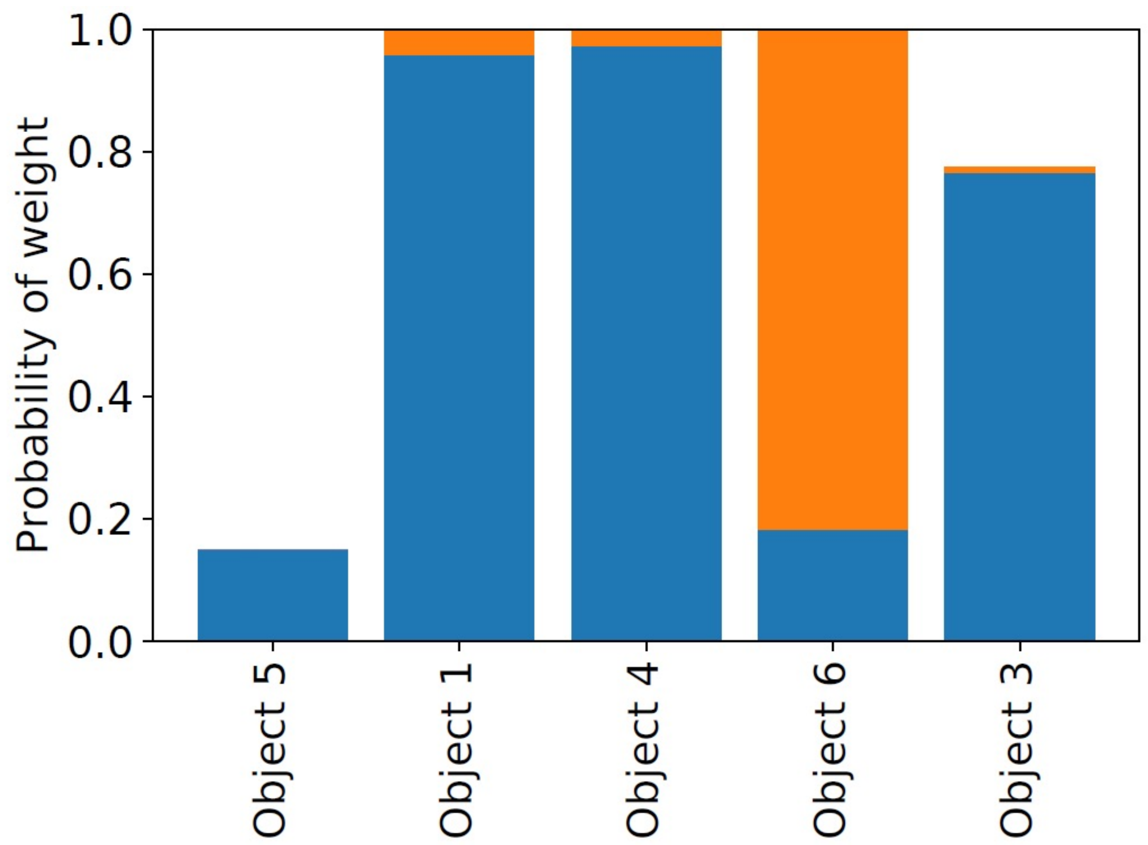

Distributed Semantic SLAM:
Class Probability Graph Notations

- Blue: class 1 probability.
- Orange: class 2 probability.
- White: class 3 probability.
- Class 1 is ground truth for all objects.

Distributed Consistent Multi-Robot Semantic Localization and Mapping

Vladimir Tchuiev and Vadim Indelman

Technion – Israel Institute of Technology

ANPL | Autonomous Navigation and Perception Lab

Technion Israel Institute of Technology

Summary Thus Far

- **An approach for semantic SLAM.**
- Maintain a **hybrid belief** over:
  - Robot and object poses.
  - Object classes.
- Leverage the coupling between poses and classes via a **viewpoint dependent classifier model.**
- The approach assists in **DA disambiguation**.
- The approach was expanded to a **distributed** setting.
- Avoids **double counting** for both continuous and discrete variables.

Presentation Overview

- Data association aware semantic SLAM via viewpoint dependent classifier model (published in IROS 2019)
- Distributed semantic SLAM via viewpoint dependent classifier model (published in RAL/IROS 2020)
- Epistemic uncertainty aware sequential classification (published in RAL/IROS 2018)
- Posterior epistemic uncertainty aware inference and belief space planning (upcoming paper 2021)

Introduction: Classifier Epistemic Uncertainty

- The classifier's **training set** is **limited**.
- During test time, when encountering data **outside the training set**, classification is **unreliable**.
- Results might be **catastrophic**.
- Can we reason about how **"certain"** a classification score is?

Introduction: Classifier Epistemic Uncertainty

- Class probability vector:

$$\gamma_k^i \doteq \mathbb{P}(c = i | I_k, w), \qquad \gamma_k \doteq [\gamma_k^1, \dots, \gamma_k^m]^T$$

- **Posterior** class probability vector:

$$\lambda_k^i \doteq \mathbb{P}(c = i | \gamma_{1:k}), \qquad \lambda_k = [\lambda_k^1, \dots, \lambda_k^m]^T$$

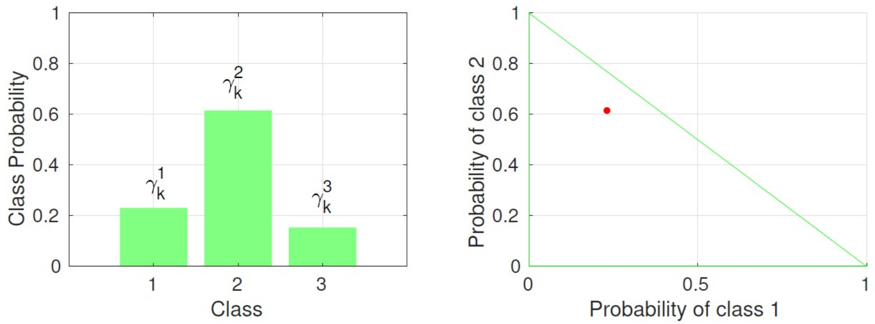

Introduction: Neural Networks

- We use a **Convolutional Neural Network (CNN)** classifier.
- The classifier parameters **(weights)** w are trained from labeled example dataset D.
- Given **fixed weights**, the classifier output is **deterministic** $\gamma_k = f_w(I_k)$.

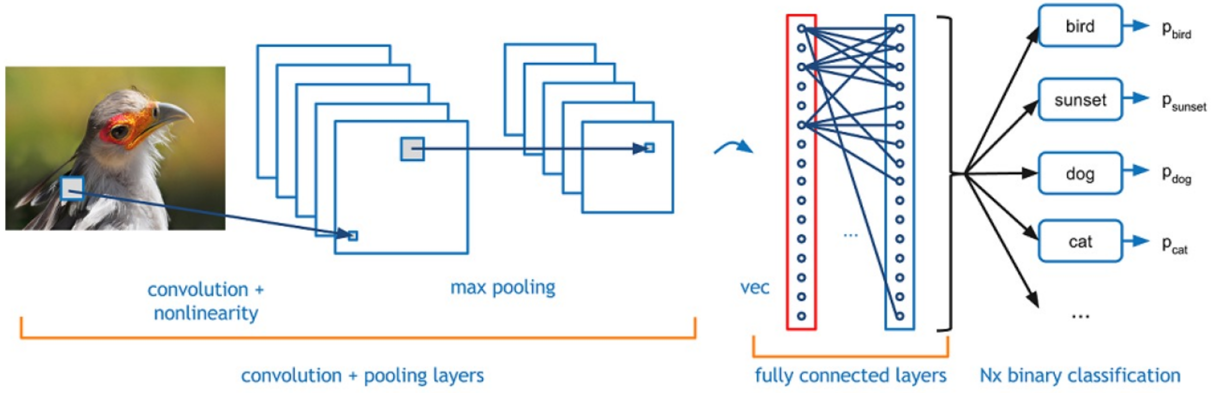

https://adeshpande3.github.io/A-Beginner%27s-Guide-To-Understanding-Convolutional-Neural-Networks/

Introduction: MC-Dropout

- **Dropout** modifies w by randomly **turning off neurons** and approximates $w \sim \mathbb{P}(w|D)$.
- We get **multiple** γ_k points corresponding to the weights: $\gamma_k \sim \mathbb{P}(\gamma_k|I_k, D)$.
- Epistemic uncertainty: **how close I_k is to the training set**?
- Although this work uses MC-dropout, it can utilize other epistemic-uncertainty-aware classifiers.

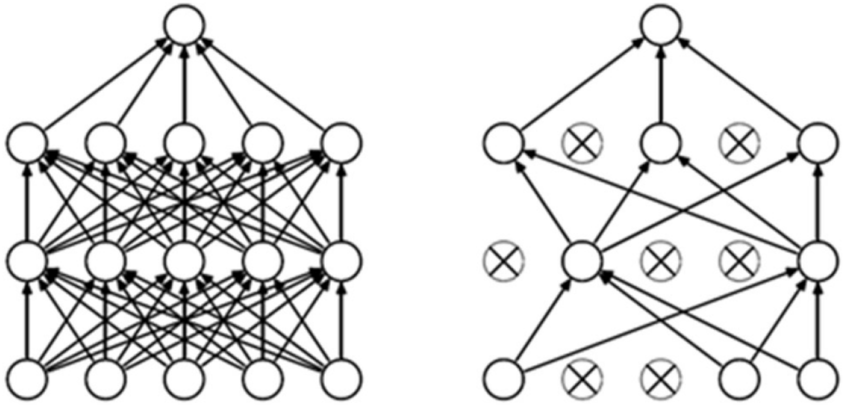

(a) Standard Neural Net (b) After applying dropout.

Srivastava, Nitish, et al. "Dropout: a simple way to prevent neural networks from overfitting", JMLR 2014

Introduction: Posterior Distribution Of Class Probability

- Eventually, we aim to infer $\mathbb{P}(\boldsymbol{\lambda}_k|\boldsymbol{I}_{1:k}, \boldsymbol{D})$.
- Because all γ are **random variables**, λ is **as well**.
- $\mathbb{P}(\boldsymbol{\lambda}_k|\boldsymbol{I}_{1:k}, \boldsymbol{D})$ may describe cases:
    - a) Out of distribution
    - b) High data uncertainty
    - c) **Confident prediction (Ideal scenario)**
    - d) Unconfident prediction

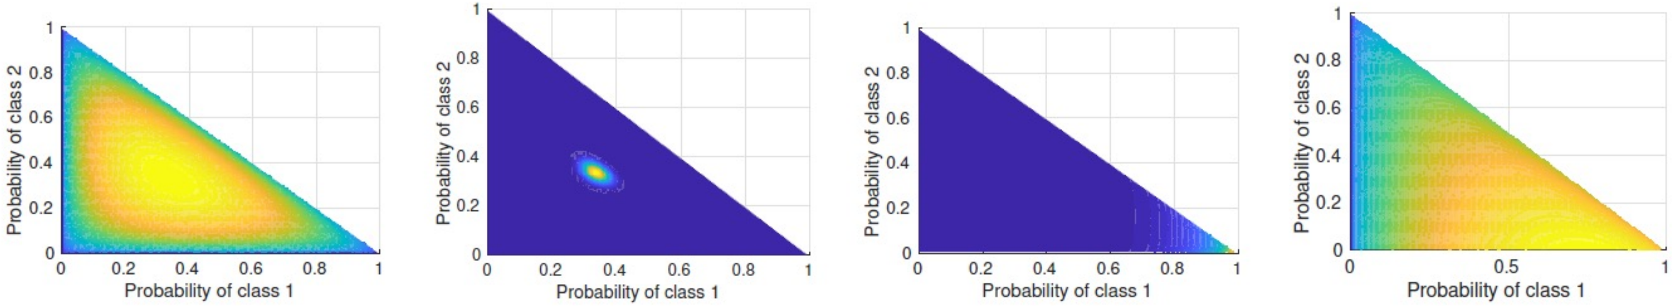

(a) (b) (c) (d)

Epistemic-Uncertainty-Aware Sequential Classification: Contribution

- We present **sequential classification** method for maintaining $\mathbb{P}(\boldsymbol{\lambda}_k|\boldsymbol{I}_{1:k}, \boldsymbol{D})$.
- We reason about the **posterior epistemic uncertainty** given the data thus far.
- ***Previous works:***
  - Sequential classification methods that reason about posterior $\mathbb{P}(c|\gamma_{1:k})$.
  - Infer epistemic uncertainty from classification from a single image only.
- Published paper: Tchuiev, Vladimir, and Vadim Indelman. "Inference over distribution of posterior class probabilities for reliable bayesian classification and object-level perception." *IEEE Robotics and Automation Letters* 3, no. 4 (2018): 4329-4336.

Epistemic-Uncertainty-Aware Sequential Classification: Assumptions

- A single object observed multiple times.
- **Classifier output** of $\{\boldsymbol{\gamma}_k\}$ that approximates $\mathbb{P}(\boldsymbol{\gamma}_k | \boldsymbol{I}_k, \boldsymbol{D})$.
- Uninformative prior for $P(c)$.
- Dirichlet distributed non-viewpoint dependent **classifier models**:

$$\mathcal{L}^i(\gamma_k) \doteq P(\gamma_k | c = i), \qquad \mathcal{L}(\gamma_k) = [\mathcal{L}^1(\gamma_k), \dots, \mathcal{L}^m(\gamma_k)]$$

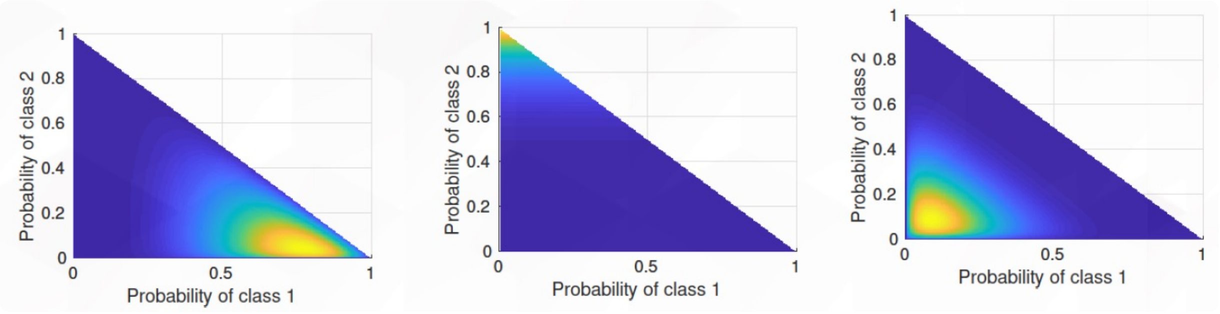

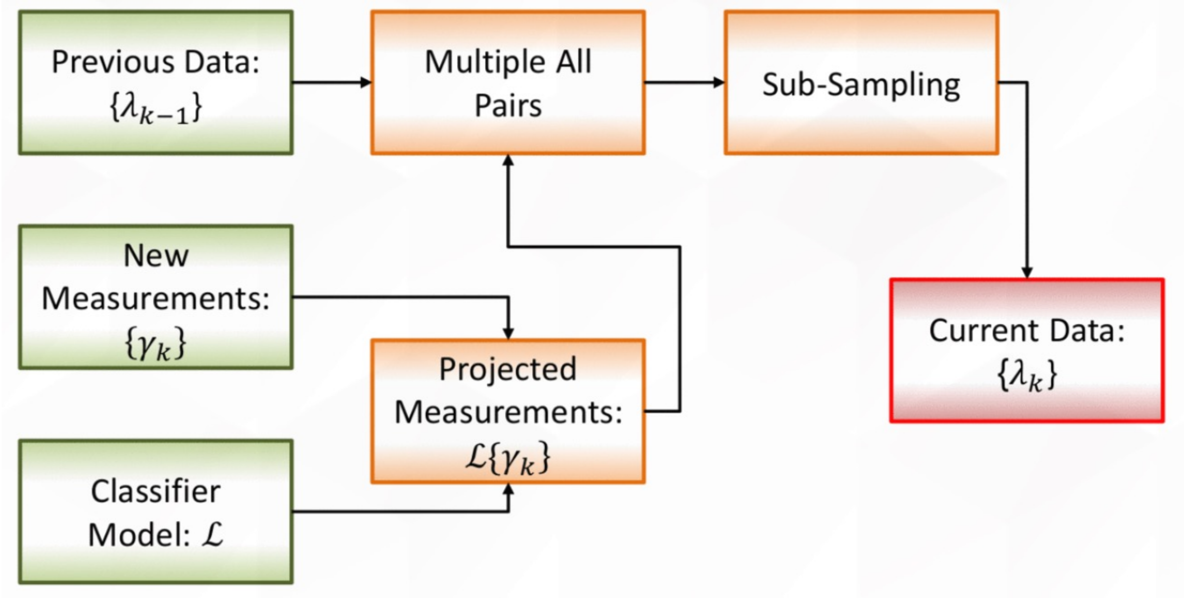

Epistemic-Uncertainty-Aware Sequential Classification: Approach

- Using Bayes rule: $\boldsymbol{\lambda_k^i \propto \lambda_{k-1}^i \mathcal{L}^i(\gamma_k)}$.
- Represent the distribution of each $\boldsymbol{\lambda}$ as a **point cloud** $\{\lambda\}$.
- Multiplying every γ_k and λ_{k-1} is **expensive**, we use sub-sampling to **reduce computation effort**.

Epistemic-Uncertainty-Aware Sequential Classification: Approach Illustration

❖Single step: posterior uncertainty **decreases**:

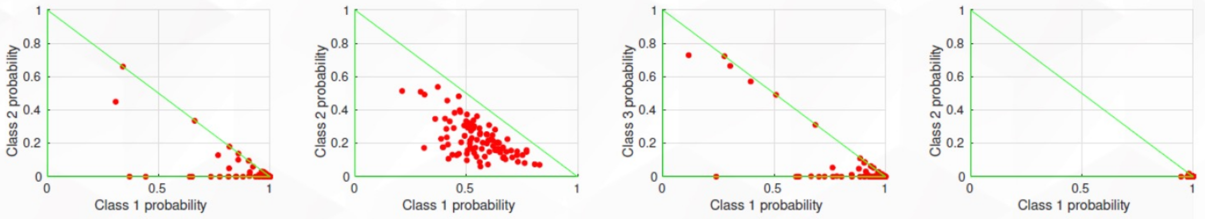

(a) $\{\lambda_{k-1}\}$ (b) $\{\gamma_k\}$ (c) $\{\mathcal{L}(\gamma_k)\}$ (d) $\{\lambda_k\}$

❖Single step: posterior uncertainty **increases**:

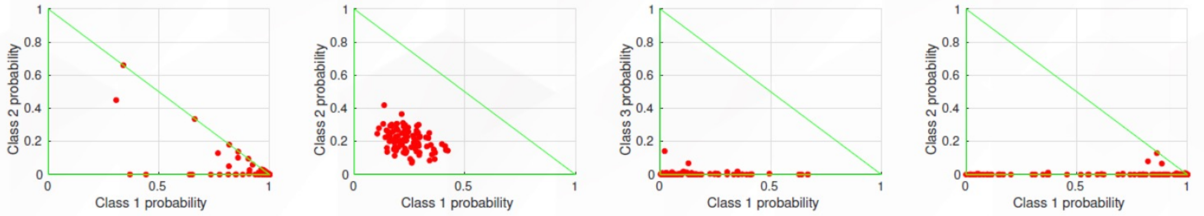

(a) $\{\lambda_{k-1}\}$ (b) $\{\gamma_k\}$ (c) $\{\mathcal{L}(\gamma_k)\}$ (d) $\{\lambda_k\}$

Epistemic-Uncertainty-Aware Sequential Classification: Experiment Setup

❖ Images of an object with **occlusion**, **blur**, and **different color filters**.

❖ 3 candidate classes, class 1 is correct.

❖ Compared between the following approaches:

- $\mathbb{P}(c|\gamma_{1:k})$, no classifier model.
- $\mathbb{P}(c|\gamma_{1:k})$, with classifier model.
- $\mathbb{P}(\lambda_k|I_{1:k}, D)$, all pairs considered.
- $\mathbb{P}(\lambda_k|I_{1:k}, D)$, with sub-sampling.

Epistemic-Uncertainty-Aware Sequential Classification: Experimental Results

❖ Our approach provides **superior classification** results.

❖ Provides access to **posterior epistemic uncertainty**.

❖ Sub sampling results are **close** to considering all pairs.

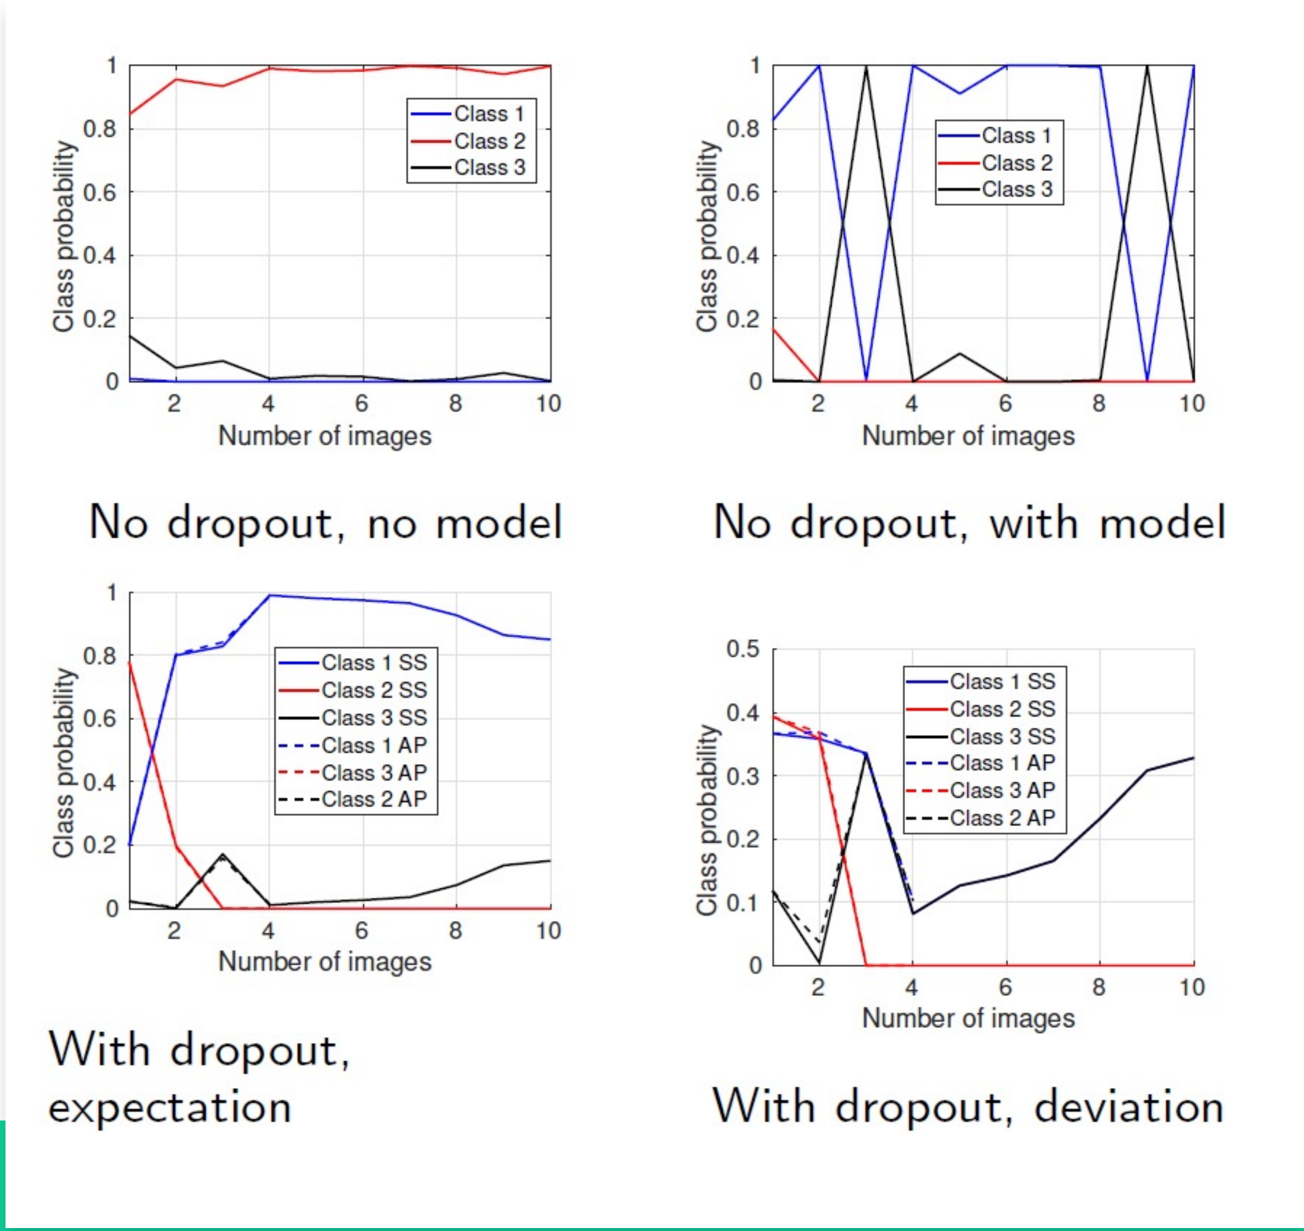

Summary Thus Far

We proposed maintaining the ***distribution over the posterior class probability*** for classification and extracting epistemic uncertainty.

We utilize a **cloud of class probability vectors** as a classier output.

To reduce computational effort, we proposed using a simple sub-sampling method.

We showed **superior results** to commonly used approaches for classification, as well as presenting ***epistemic uncertainty***.

Presentation Overview

- Data association aware semantic SLAM via viewpoint dependent classifier model (published in IROS 2019)
- Distributed semantic SLAM via viewpoint dependent classifier model (published in RAL/IROS 2020)
- Epistemic uncertainty aware sequential classification (published in RAL/IROS 2018)
- Posterior epistemic uncertainty aware inference and belief space planning (upcoming paper 2021)

Introduction: Active Classifier Epistemic-Uncertainty-Aware Inference and Planning

- **Up to this point** we presented methods for addressing:
  - Viewpoint dependency of classification scores.
  - Localization and mapping uncertainty.
  - Classifier epistemic uncertainty.
- Now we introduce two methods that **address both simultaneously** in inference:
  - Multi-Hybrid (MH)
  - Joint Lambda Pose (JLP)
- We extend the formulation of those two methods to **belief space planning**.

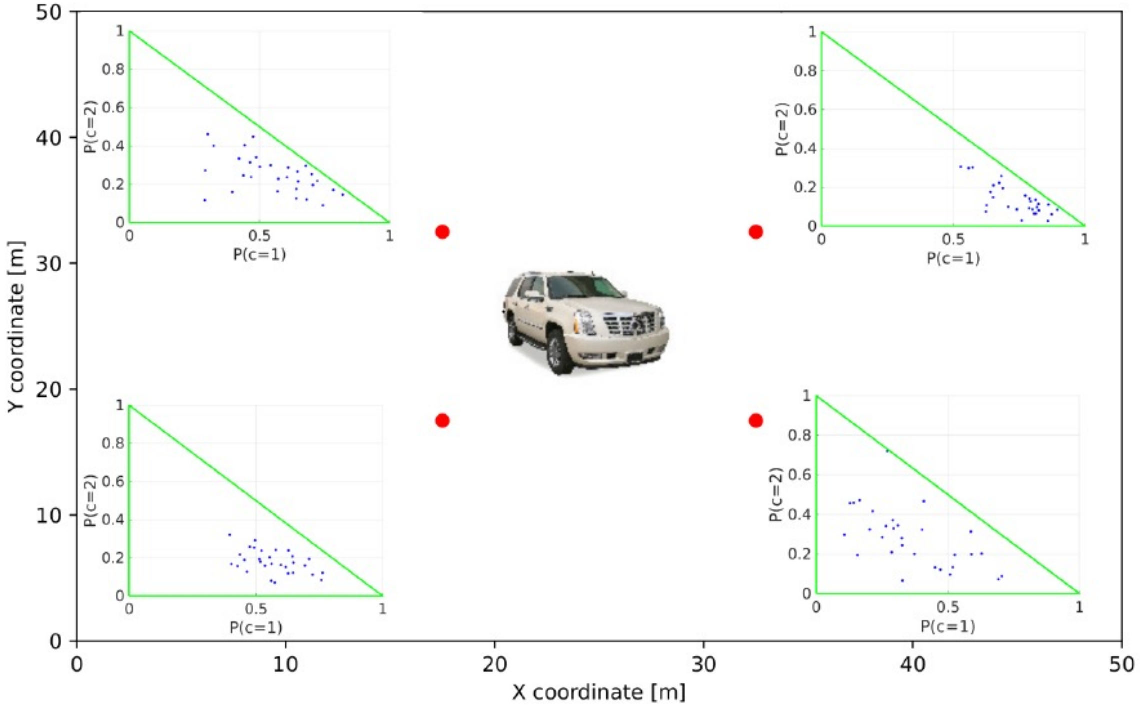

Multi-Hybrid (MH) and Joint Lambda Pose (JLP): Contributions

Maintain an epistemic uncertainty aware joint belief over poses and class probabilities:

- Multi-Hybrid **(MH).**
- Joint Lambda Pose **(JLP).**

Utilize a viewpoint dependent classifier uncertainty model to:

- Predicts **epistemic uncertainty** given **viewpoint**.
- Improve **classification performance** in inference.
- **Generate** predicted **measurements** for BSP.

Propose an information-theoretic reward over posterior epistemic uncertainty

Previous works:

- Don't consider classifier **epistemic uncertainty for BSP**.
- Epistemic uncertainty aware planning with **solved localization**.

Ongoing work for 2021 paper submission.

Introduction: Belief Space Planning (BSP)

❖A framework for **planning under uncertainty**.

❖**Objective Function**: given belief b_k, and an action sequence $a_{k:k+L}$:

$$J(b_k, a_{k:k+L}) = E_{Z_{k+1:k+L}}\left(\sum_{i=0}^{L} r(b_{k+i}, a_{k+i})\right)$$

- $r(\cdot)$ is the **reward function**.
- b_{k+i} is a function of observations Z_{k+i}

Introduction: Belief Space Planning (BSP)

- $J(b_k, a_{k:k+L})$ rewritten in a recursive form:

$$J(b_k, a_{k:k+L}) = \int_{Z_{k+1}} \mathbb{P}(Z_{k+1}|\mathcal{H}_k, a_k) \cdot J(b_{k+1}, a_{k+1:k+L}) dZ_{k+1}$$

- $\mathbb{P}(Z_{k+1}|\mathcal{H}_k, a_k)$: **measurement likelihood** term.

- The aim is finding an **optimal** action sequence:

$$a^*_{k:k+L} = \underset{a_{k:k+L}}{arg\ max}\ J(b_k, a_{k:k+L})$$

Introduction: Belief Space Planning (BSP)

❖ Key issue: generating predicted semantic measurements.

❖ ***Option 1***: generating **raw images**.

- High dimensional problem.
- Feasible only in specifically trained environments.

❖ ***Option 2***: generating **directly from classifier model**.

- Output dimension is much smaller.
- Can be generalized to more environments.

MH and JLP: Classifier Uncertainty Model

- Requirement: a viewpoint dependent model that fits both **inference** and **planning** (sampling).
- **Logit transformation** of a general probability vector $v \in \mathbb{R}^m$ to $lv \in \mathbb{R}^{m-1}$:

$$lv \doteq \left[\frac{\log v_1}{\log v_m}, \dots, \frac{\log v_{m-1}}{\log v_m}\right]^T$$

- γ_k is Logistical Gaussian distributed, therefore $l\gamma_k$ is Gaussian distributed:

$$\mathbb{P}(l\gamma | c, x^{rel}) = \mathcal{N}\left(h_c(x^{rel}), \Sigma_c(x^{rel})\right)$$

- Model's training set: $D_c \doteq \{x^{rel}, \{l\gamma\}\}$.
- Predicts **epistemic uncertainty**.

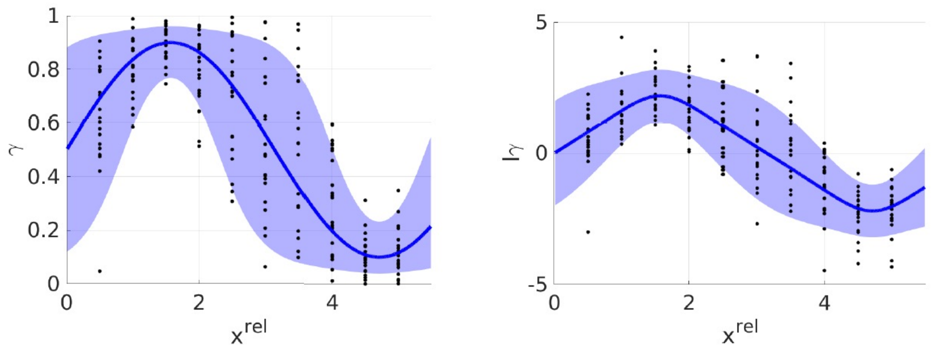

(a) γ space (b) $l\gamma$ space

MH Inference and Planning

- **We aim to infer the joint belief** $\mathbb{P}(\boldsymbol{\lambda}_k, \boldsymbol{\mathcal{X}}_k | \boldsymbol{\mathcal{H}}_k)$.
- We determine **fixed** weight realizations $w \in W$.
- Marginalizing over w:
$$\mathbb{P}(\lambda_k, \mathcal{X}_k | \mathcal{H}_k) = \sum_w \mathbb{P}(\mathcal{X}_k | \lambda_k, \mathcal{H}_k, w) \mathbb{P}(\lambda_k | \mathcal{H}_k, w)$$
  the R.H.S can be inferred via ***maintaining a hybrid belief per each w***.
- **In planning**, predicted measurements are generated via the classifier uncertainty model.
- MH is **computationally inefficient;** therefore, we propose **JLP.**

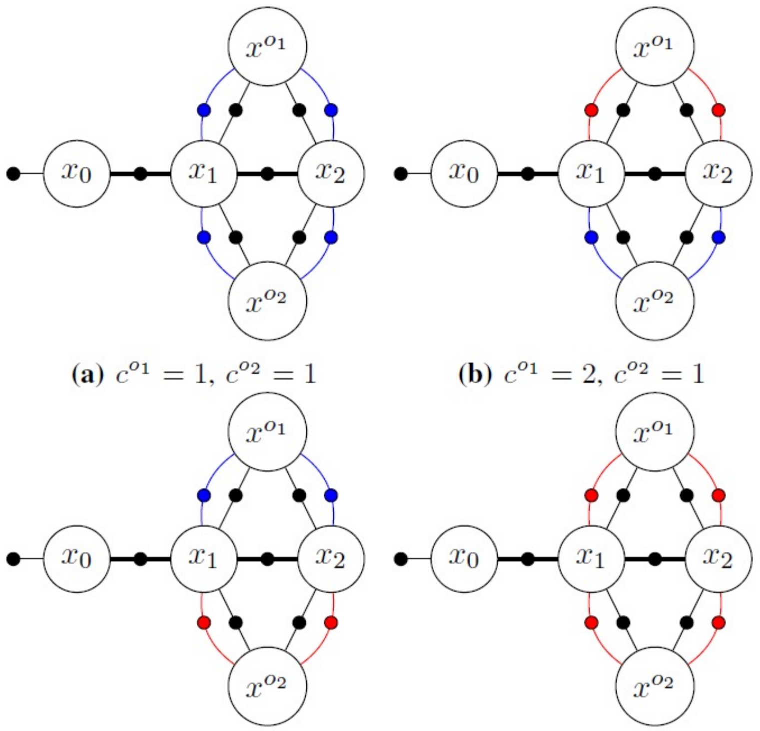

(a) $c^{o_1} = 1,\ c^{o_2} = 1$ **(b)** $c^{o_1} = 2,\ c^{o_2} = 1$
(c) $c^{o_1} = 1,\ c^{o_2} = 2$ **(d)** $c^{o_1} = 2,\ c^{o_2} = 2$

$\times|W|$

JLP Inference: Approach

- MH is **computationally expensive**; we propose a **more efficient alternative**.
  - MH maintains **multiple hybrid beliefs**.
  - JLP maintains a **single continuous belief**.
- **We aim to maintain the joint belief:**

$$\boldsymbol{b[l\lambda_k, \mathcal{X}_k] \doteq P(l\lambda_k, \mathcal{X}_k | \mathcal{H}_k, D)}$$

- Recursive formulation:

$$b[l\lambda_k, \mathcal{X}_k] \\ = \int_{l\lambda_{k-1}} \boldsymbol{P(l\lambda_k | l\lambda_{k-1}, \mathcal{H}_k, \mathcal{X}_k)} P(z_k^{geo} | \mathcal{X}_k) P(x_k | x_{k-1}, a_{k-1}) b[l\lambda_{k-1}, \mathcal{X}_{k-1}] dl\lambda_{k-1}$$

- Introducing the novel **JLP factor**.
- JLP is even more efficient than MH when considering multiple objects.

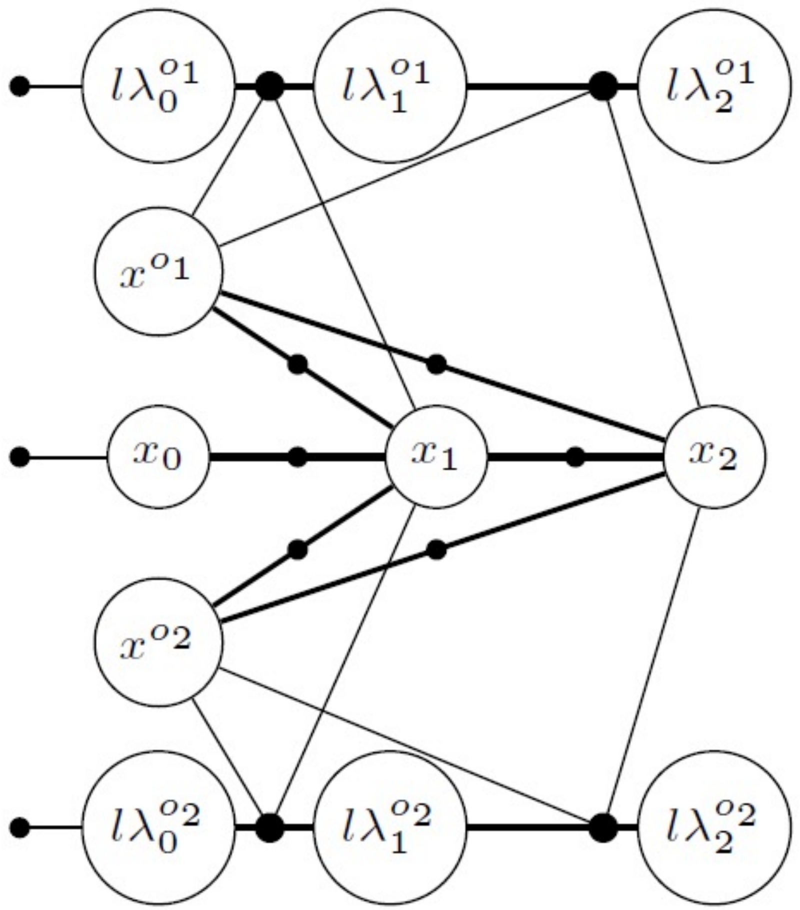

JLP inference: Approach

Under the condition below, the *JLP factor* is Gaussian and $l\lambda_k$ can be inferred by standard optimization methods.

❖ Recall the **classifier uncertainty model**:

$$\mathbb{P}\left(l\gamma \middle| c, x^{rel}\right) = \mathcal{N}(h_c, \Sigma_c)$$

❖ If $\Sigma_{c=i}(x^{rel}) = \Sigma_{c=j}(x^{rel})$ for all candidate classes, then the JLP factor is **Gaussian**.

❖ Even if the condition doesn't apply, the JLP factor **is approximately Gaussian** besides extreme cases.

JLP Planning: Measurement Generation

❖ Specifically for JLP, the objective function is:

$$J(b[l\lambda_k, \mathcal{X}_k], a_{k:k+L}) = E_{E(l\gamma_{k+1:k+L}),\Sigma(l\gamma_{k+1:k+L}),z^{geo}_{k+1:k+L}}\left(\sum_{i=0}^{L} r(b[l\lambda_{k+i}, \mathcal{X}_{k+i}], a_{k+i})\right)$$

❖ **Sampling of measurements:**

- **Geometric** from the measurement model.
- **Semantic** from the ***parameters*** ***of the classifier uncertainty model.***

❖ *Sampled measurements are used to infer predicted* $b[l\lambda_{k+i}, \mathcal{X}_{k+i}]$.

MH and JLP Planning: Reward Functions

- ***Maintaining*** $\boldsymbol{b[\lambda, \mathcal{X}]}$ ***opens access to a reward function of general type*** $\boldsymbol{r(b[\lambda, \mathcal{X}])}$ ***with possible variations:***
  - $r(\mathcal{X})$, e.g., distance to goal.
  - $r(b[\mathcal{X}])$, e.g., information-theoretic.
  - $r(E(\lambda))$, e.g., information entropy.
  - $r(b[\lambda])$, *a novel reward function type, planning over epistemic uncertainty.*
- The **posterior epistemic uncertainty** affects every reward.
- We use **negative of differential entropy** as reward:

$$r(b[\lambda]) = -H(\lambda) = \int_{\lambda} b[\lambda] \cdot log(b[\lambda])\, d\lambda$$

- $-H(\lambda)$ accounts for both $E(\lambda)$ (classification scores) and $\Sigma(\lambda)$ (epistemic uncertainty) **without hyperparameter tuning.**

MH and JLP Planning: Simulation Setup

- 9 objects in a 2D environment.
- 2 candidate classes.
- 5 motion primitives.
- **Two reward functions:**
  - $R_1 = \min(-\sum_{o \in O} H(\lambda), R_1^{max})$
  - $R_2 = -\sum_{o \in O} \sum_{c^o} E(\lambda^{c,o}) \cdot log(E(\lambda^{c,o}))$
- Compare between:
  - MH
  - JLP
  - Without Epistemic Uncertainty (WEU)
- *MSDE* as **classification benchmark**:

$$MSDE \doteq \frac{1}{m} \sum_{i=1}^{m} \left(\mathbb{P}_{gt}(c=i) - \mathbb{P}(c=i|\mathcal{H}_k^R) \right)^2$$

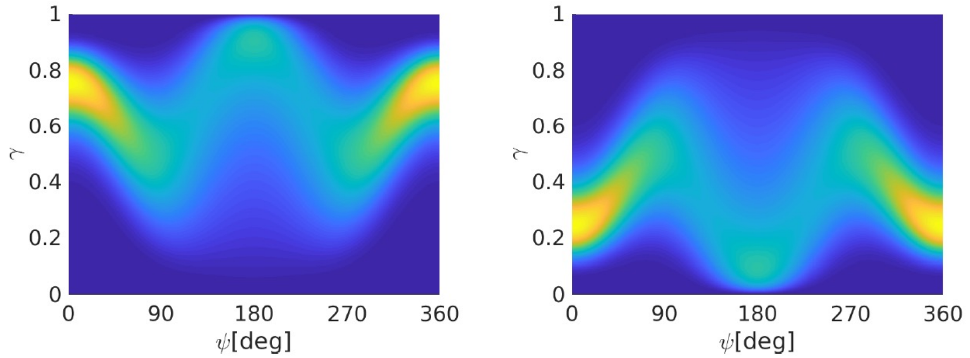

(a) $\mathbb{P}(\gamma^{c=1}|c=1,\psi)$ **(b)** $\mathbb{P}(\gamma^{c=1}|c=2,\psi)$

Ground Truth

Motion Primitives

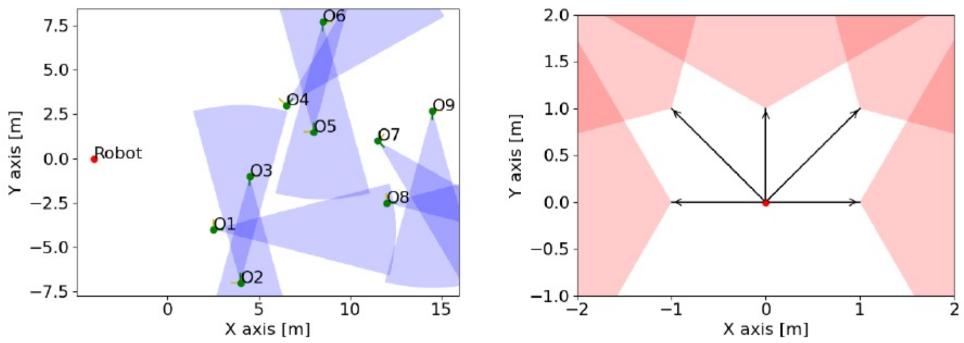

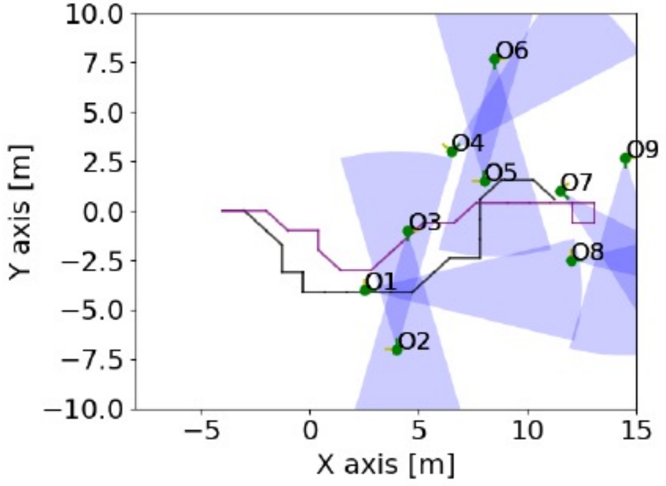

(a) R_1 (b) R_2 (c) WEU

MH and JLP Planning:
Simulation Results

- We show results for inference after actions already taken.
- Trajectories created by planning.

ANPL Autonomous Navigation and Perception Lab

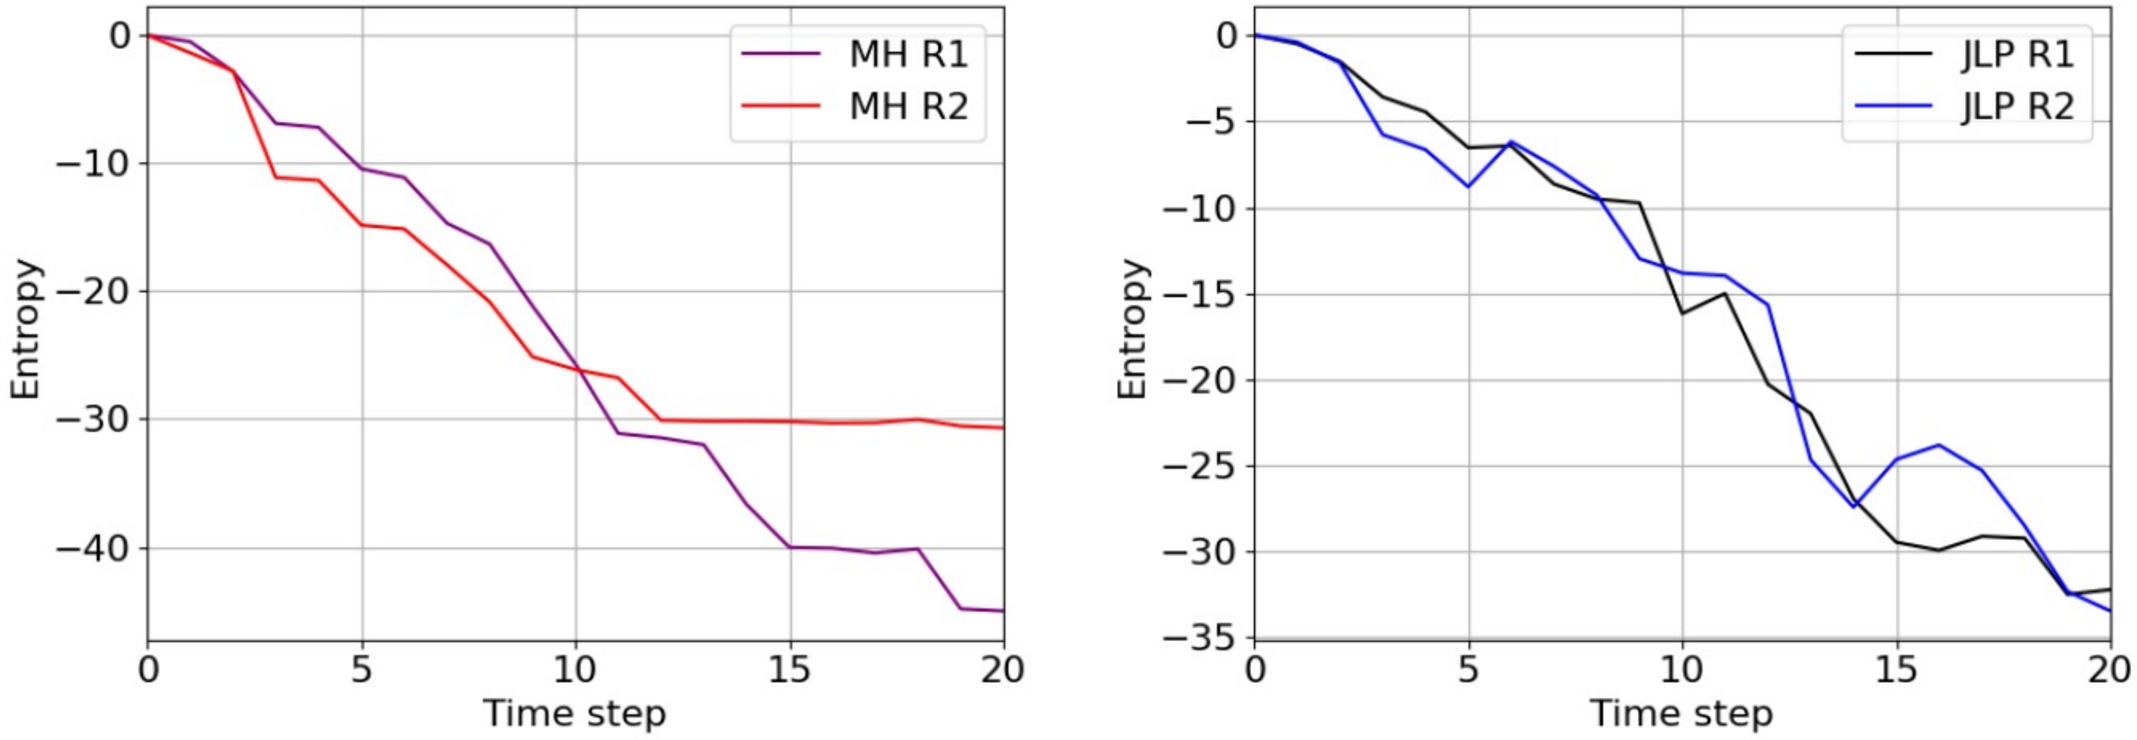

MH and JLP Planning: Simulation Results

- Entropy $\sum_{o \in O} H(\lambda^o)$ values as a function of time step.
- **Advantage** for using R_1 over R_2.

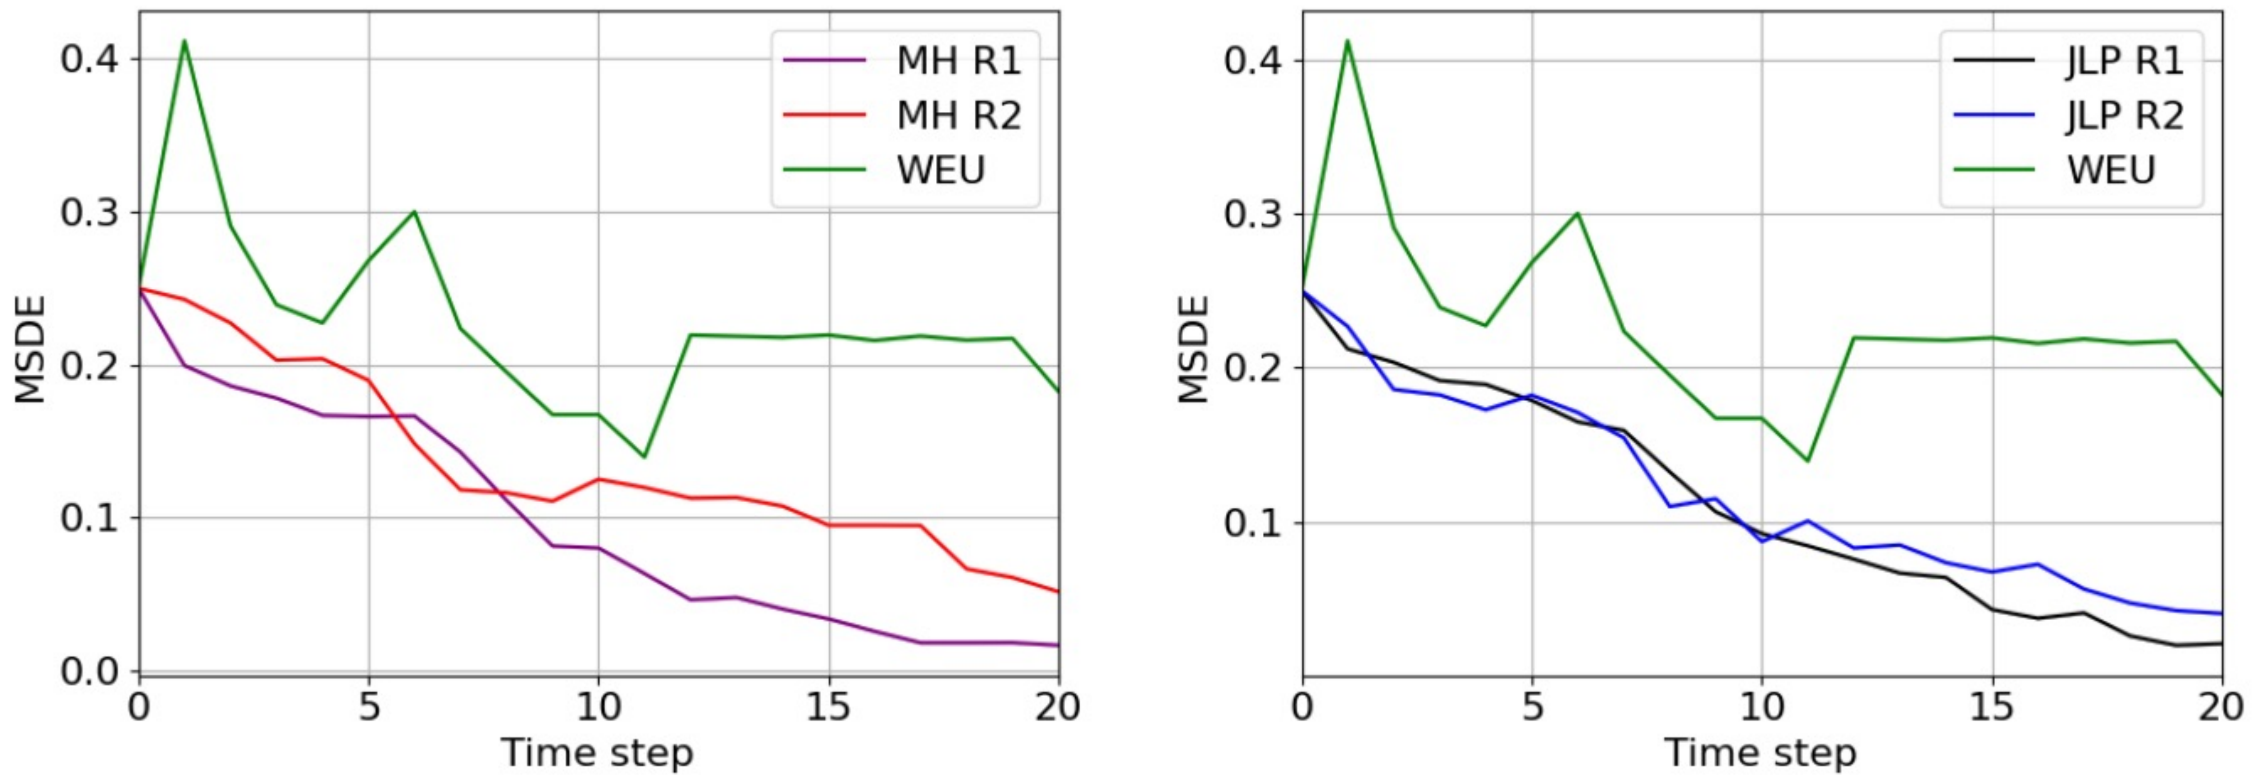

MH and JLP Planning:
Simulation Results

- ❖ MSDE results as a function of time step.
- ❖ **Advantage** for using R_1 over R_2, with both **outperforming WEU**.

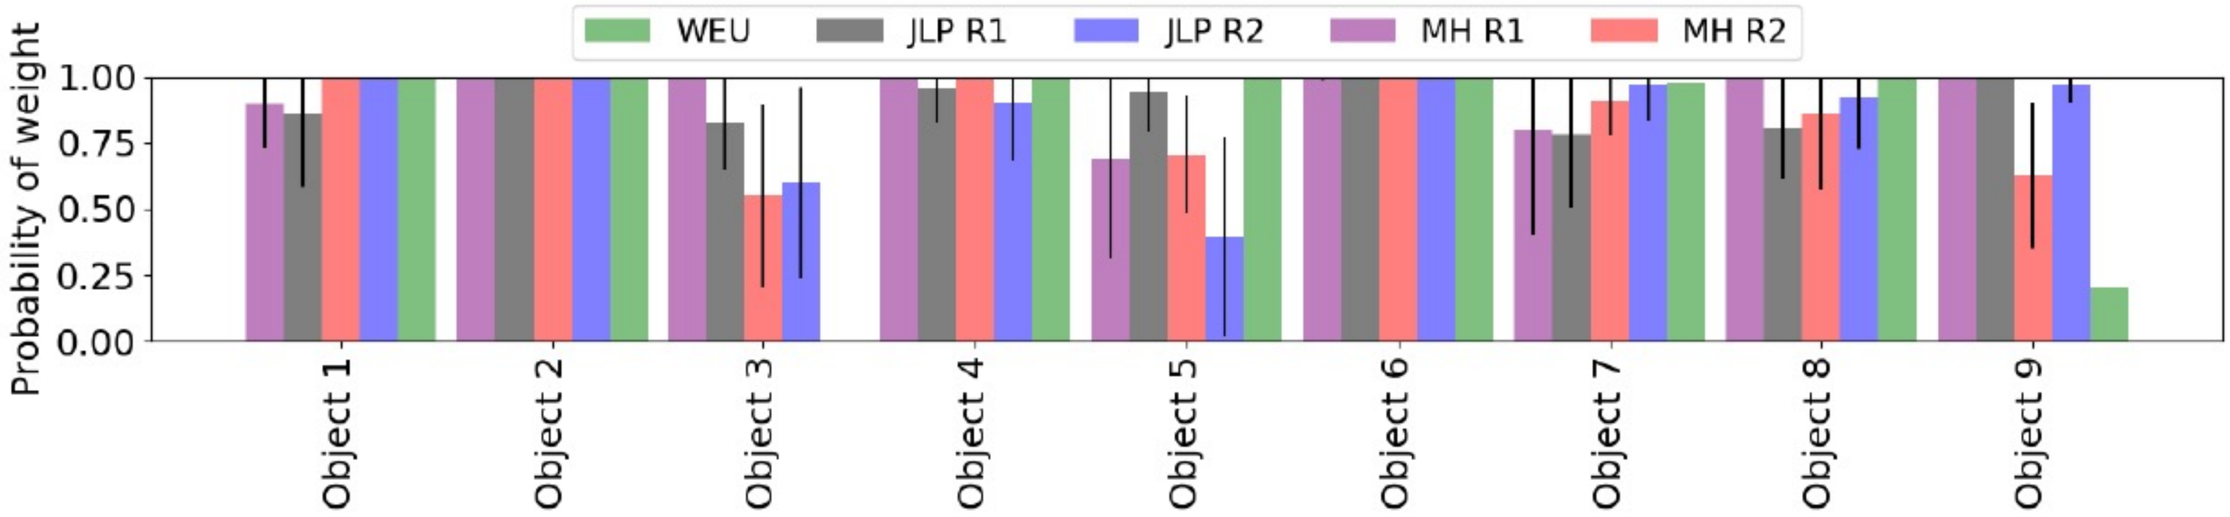

MH and JLP Planning: Simulation Results

- **Classification results** for the objects at $k = 20$: probability of the correct class.
- Black line represents the **posterior epistemic uncertainty**.
- **Advantage** for using R_1 over R_2. WEU tends to go to **extremes**.

MH and JLP Planning:
Simulation Results

- **Computation time** comparison between MH with 10 beliefs, JLP, and WEU.
- WEU is the fastest, JLP is comparable, while MH is the slowest.

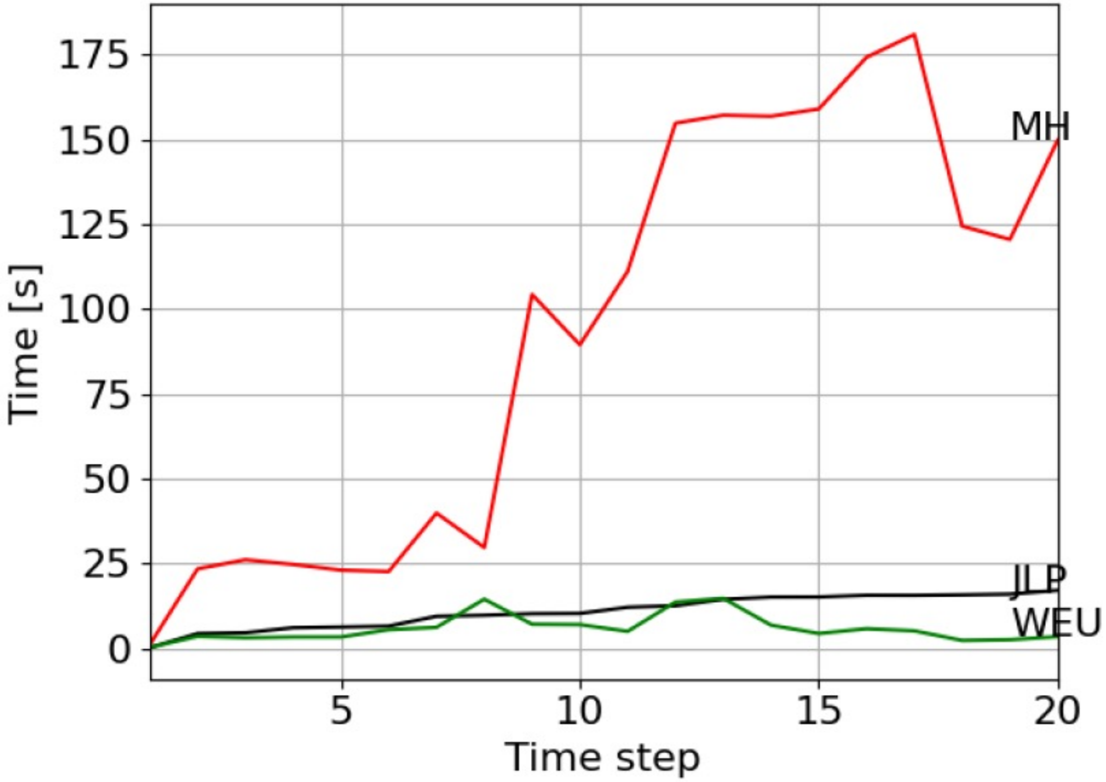

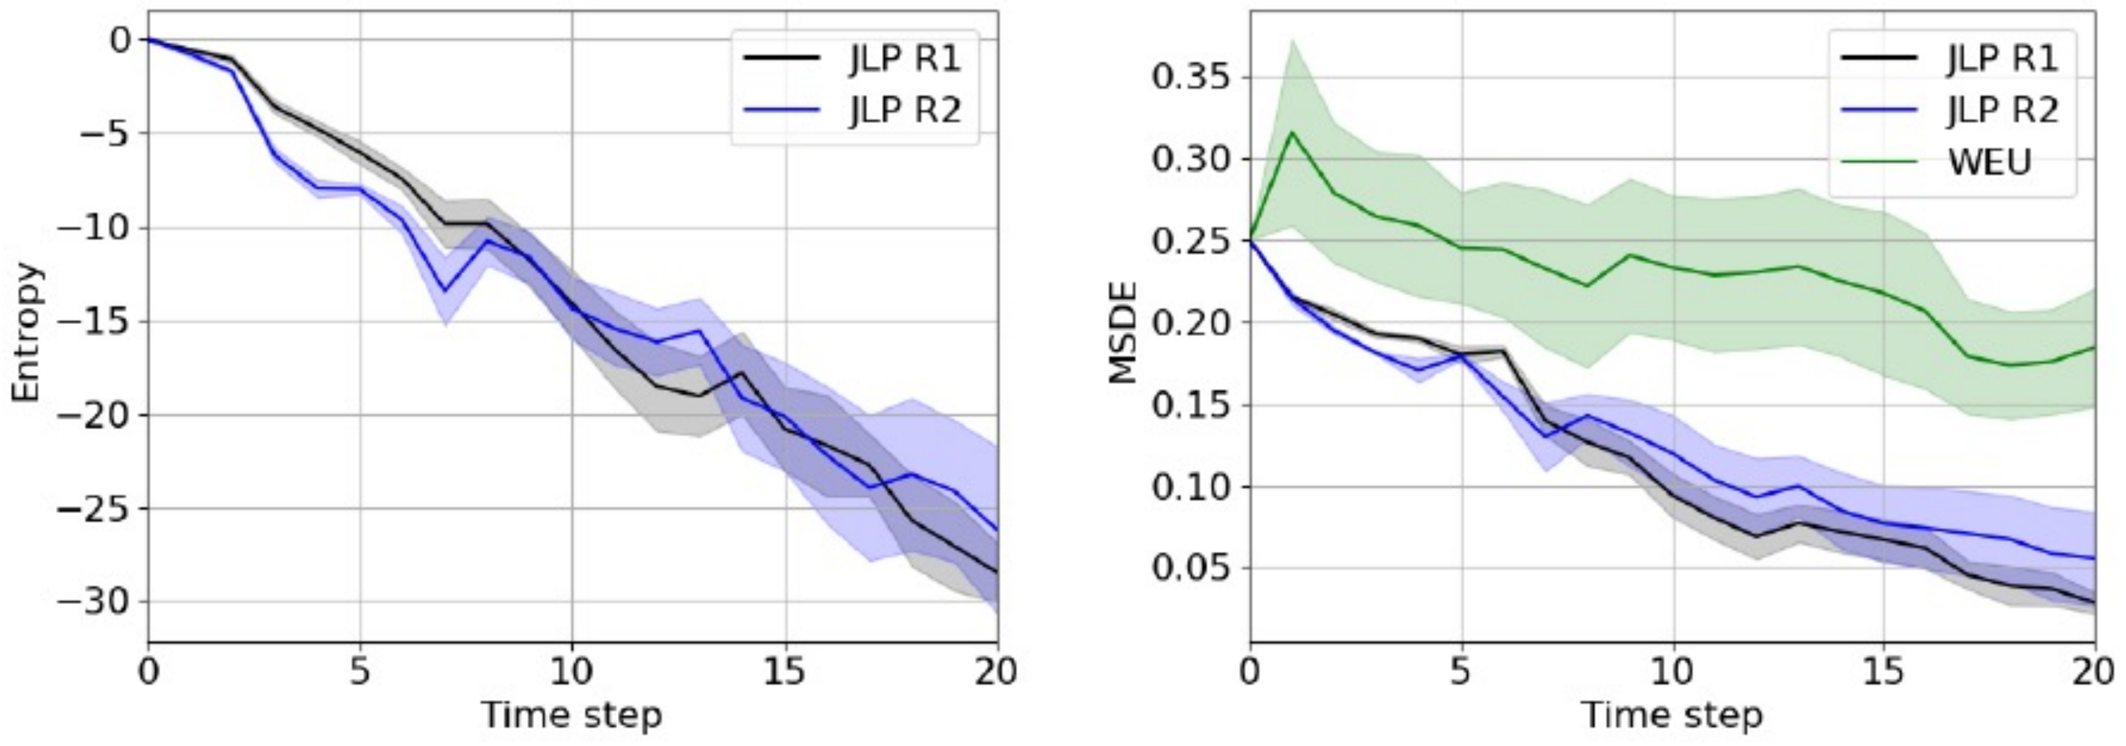

MH and JLP Planning: Simulation Results

- Statistical results for JLP with planning over R_1 and R_2 compared to WEU: entropy and MSDE.
- Colored area – one σ range.
- **Significant advantage vs WEU**, with R_1 having a small edge over R_2.

Summary

- **Uncertainties in object classification**
  - **Viewpoint dependency**.
    - A **semantic SLAM** approach that maintains a hybrid belief over poses and classes.
    - Expanding the approach to a **distributed** multi-robot setting.
    - Leveraging the coupling between poses and classes via a **viewpoint dependent classifier model**.
  - **Epistemic uncertainty.**
    - An approach that maintains the distribution of the **posterior class probability vector**.
    - MH and the faster JLP that reasons both about **viewpoint dependency** and **epistemic uncertainty**.
  - **Belief space planning**
    - Expand MH and JLP for **BSP**.
    - Use a **viewpoint dependent classifier uncertainty model** both for inference and BSP.
- Our approaches showed increased performance for **classification**, **localization**, and **data association disambiguation**.