Contrastive Meta-Learning for Partially Observable Few-Shot Learning

Adam Jelley [1], Amos Storkey [1], Antreas Antoniou [1], Sam Devlin [2]

1 School of Informatics, University of Edinburgh

2 Microsoft Research Cambridge

Learning Representations from Partial Observations

- Many **contrastive and meta-learning** approaches learn representations by identifying **common features in multiple views**
- However, the formalism for these approaches generally **assumes features to be shared across views** to be captured coherently
- We consider the problem of learning a **unified representation from partial observations**, where useful features may be present in only some of the views

Partial Observation Experts Modelling (POEM)

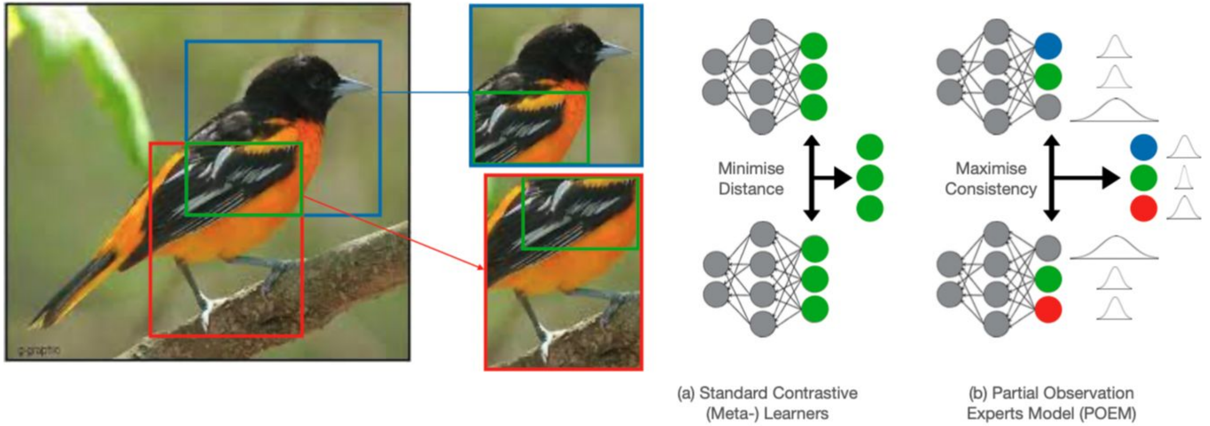

- Standard contrastive (meta-) learners **minimise relative distance** between representations to learn features that are consistent in views
- For partial observations, we instead want to **maximise consistency** by utilising representation uncertainty to capture which features of the entity are present in a view
- We achieve consistency **using a product of experts model** (Hinton, 2002) to combine and unify individual view representations that **weights features by their uncertainty**
- Since it is not obvious how to do so, we must **meta-learn** how to integrate views

POEM Theoretical Formalism

We consider a set S of V^m partial observations of M items where V^m can vary with m:
$S = \{X^m | m = 1, 2, \ldots, M\}$ where $X^m = \{x_1^m, x_2^m, \ldots, x_{V^m}^m\}$.

From a given support set X^m, we can compute a product of experts distribution z^m:

$$p(\mathbf{z}^m|\mathbf{X}^m) = \frac{p(\mathbf{z}^m)\prod_{v=1}^{V^m}\phi(\mathbf{z}^m|\mathbf{x}_v^m)}{\int d\mathbf{z}'\, p(\mathbf{z}')\prod_{v=1}^{V^m}\phi(\mathbf{z}'|\mathbf{x}_v^m)}$$

We can also compute the probability that a query view x^* would be generated from z^m:

$$p(\mathbf{x}^*|\mathbf{z}^m) = \frac{p(\mathbf{x}^*)\phi(\mathbf{z}^m|\mathbf{x}^*)}{p(\mathbf{z}^m)}$$

POEM Theoretical Formalism

These probabilities can be combined by marginalising over z^m to get the **marginal predictive**:

$$p(\mathbf{x}^*|\mathbf{X}^m) = \int d\mathbf{z}^m \;\; p(\mathbf{x}^*|\mathbf{z}^m) \;\; p(\mathbf{z}^m|\mathbf{X}^m)$$

$$= \int d\mathbf{z}^m \left(\frac{p(\mathbf{z}^m) \prod_{v=1}^{V^m} \phi(\mathbf{z}^m|\mathbf{x}_v^m)}{\int d\mathbf{z}' \; p(\mathbf{z}') \prod_{v=1}^{V^m} \phi(\mathbf{z}'|\mathbf{x}_v^m)} \right) \left(\frac{p(\mathbf{x}^*)\phi(\mathbf{z}^m|\mathbf{x}^*)}{p(\mathbf{z}^m)} \right)$$

$$= p(\mathbf{x}^*) \left(\frac{\int d\mathbf{z}^m \; \phi(\mathbf{z}^m|\mathbf{x}^*) \prod_{v=1}^{V^m} \phi(\mathbf{z}^m|\mathbf{x}_v^m)}{\int d\mathbf{z}' \; p(\mathbf{z}') \prod_{v=1}^{V^m} \phi(\mathbf{z}'|\mathbf{x}_v^m)} \right) = p(\mathbf{x}^*) \frac{\lambda(\mathbf{x}^*, \mathbf{X}^m)}{\lambda'(\mathbf{X}^m)}$$

where $\lambda(\mathbf{y}, \mathbf{X}) = \int d\mathbf{z} \, \phi(\mathbf{z}|\mathbf{y}) \prod_{v=1}^{V} \phi(\mathbf{z}|\mathbf{x}_v), \quad \text{and} \quad \lambda'(\mathbf{X}) = \int d\mathbf{z} \, p(\mathbf{z}) \prod_{v=1}^{V} \phi(\mathbf{z}|\mathbf{x}_v).$

POEM Theoretical Formalism

We can now **maximise the probability** that a given query view x^* comes from the correct corresponding item X^m, providing a **negative log likelihood objective** to be minimised:

$$\mathcal{L}(\{S_t\}, \{x_t^*\}) = -\sum_t \left[\log \frac{\lambda(\mathbf{x}^*, \mathbf{X}^{m^*})}{\lambda'(\mathbf{X}^{m^*})} - \log \sum_m \frac{\lambda(\mathbf{x}^*, \mathbf{X}^m)}{\lambda'(\mathbf{X}^m)} \right]$$

where $\lambda(\mathbf{y}, \mathbf{X}) = \int d\mathbf{z}\, \phi(\mathbf{z}|\mathbf{y}) \prod_{v=1}^{V} \phi(\mathbf{z}|\mathbf{x}_v), \quad \text{and} \quad \lambda'(\mathbf{X}) = \int d\mathbf{z}\, p(\mathbf{z}) \prod_{v=1}^{V} \phi(\mathbf{z}|\mathbf{x}_v).$

Interpretation of POEM Objective

- POEM **generalises Prototypical Networks** to a probabilistic latent space (with approximate equivalence when all uncertainties equal)
- This enables fair **comparison** of **partial** support and query views

Generalising Prototyping
(Weighting by Feature Uncertainty)

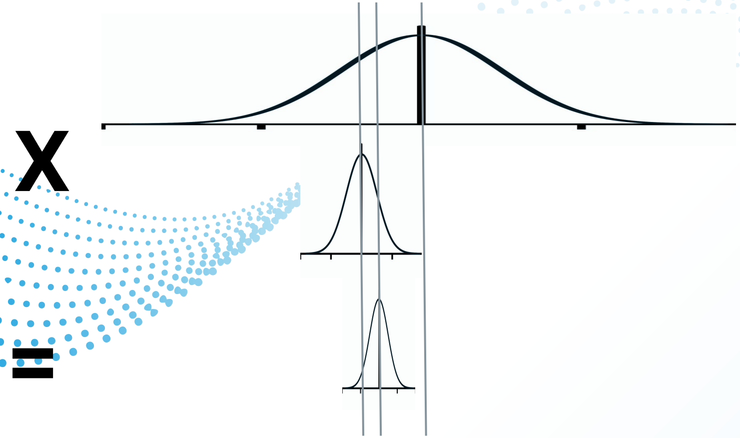

Generalising Querying
(Weighting by Support/Query Uncertainty)

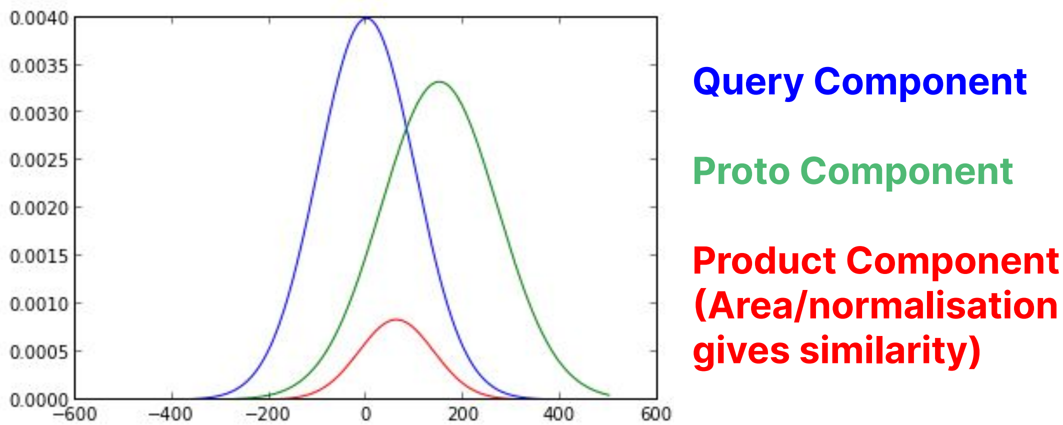

Partially Observable Meta-Dataset

We adapt the **Meta-Dataset** (Triantafillou et al., 2020) few-shot learning benchmark (dataset of few-shot datasets) to **incorporate partial observability** with strong cropping and standard contrastive augmentations

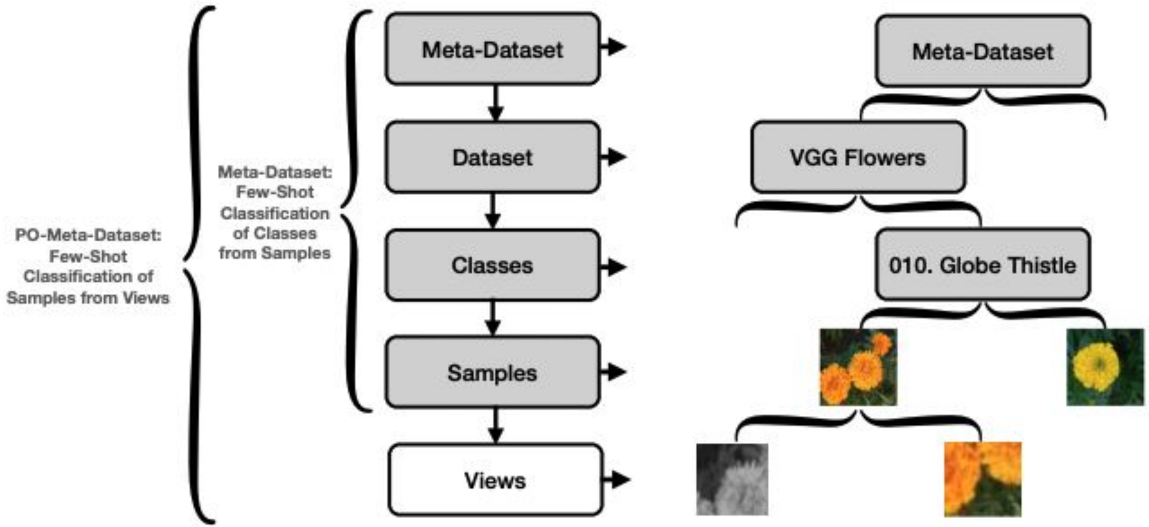

Results

POEM **outperforms** common meta-learning **baselines** at **few-shot learning** from **partial observations**, while matching state-of-the-art on fully-observable (standard) Meta-Dataset.

PO-Meta-Dataset:

Test Source	Finetune	ProtoNet	MAML	POEM
Aircraft	46.5 ± 0.6	48.5 ± 1.0	37.5 ± 0.3	$\mathbf{55.3 \pm 0.7}$
Birds	62.6 ± 0.7	67.4 ± 1.2	52.5 ± 0.6	$\mathbf{71.1 \pm 0.1}$
Flowers	48.5 ± 0.4	46.4 ± 0.7	33.5 ± 0.3	$\mathbf{49.2 \pm 1.5}$
Fungi	61.0 ± 0.2	61.4 ± 0.4	46.1 ± 0.4	$\mathbf{64.8 \pm 0.3}$
Omniglot	71.3 ± 0.1	87.8 ± 0.1	47.4 ± 1.0	$\mathbf{89.2 \pm 0.7}$
Textures	$\mathbf{83.2 \pm 0.4}$	76.7 ± 1.6	73.1 ± 0.4	81.4 ± 0.6

Meta-Dataset:

Test Source	Finetune	ProtoNet	MAML	POEM
Aircraft	56.2 ± 1.1	47.2 ± 1.2	35.9 ± 1.8	46.5 ± 1.5
Birds	52.6 ± 1.8	78.3 ± 0.5	65.2 ± 0.3	79.4 ± 0.3
Flowers	80.1 ± 2.0	84.2 ± 0.7	70.4 ± 0.4	83.6 ± 1.3
Fungi	33.6 ± 1.7	84.7 ± 0.2	18.9 ± 0.2	81.0 ± 0.1
Omniglot	89.6 ± 3.3	98.7 ± 0.1	94.7 ± 0.1	98.6 ± 0.1
Textures	60.4 ± 1.0	65.3 ± 1.2	56.1 ± 0.3	65.7 ± 0.8

All methods use a modernised, ImageNet-pretrained ResNet-18 backbone for fair comparison.

Application of POEM to Reinforcement Learning Agents in POMDPs

Many **reinforcement learning environments** are **partially observable** (POMDPs), and it may be helpful for an **agent** to **learn an environment representation** from its **partial observations**, e.g. in a maze:

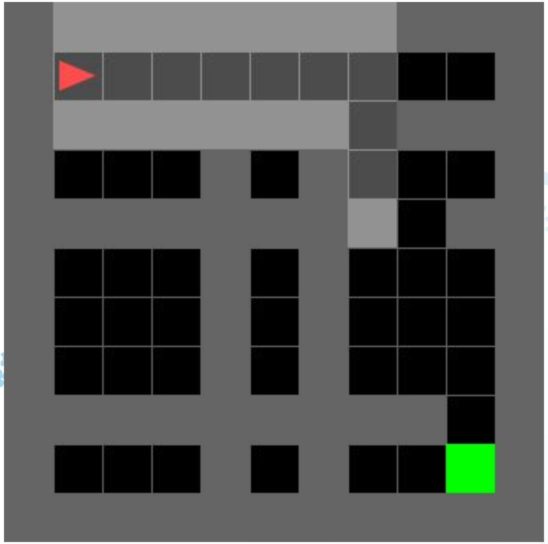

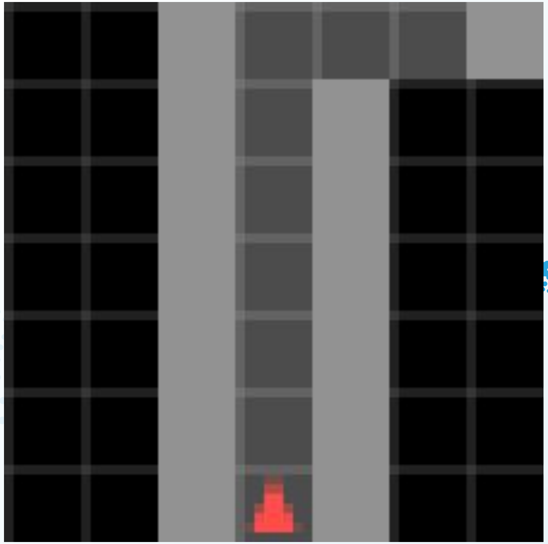

Training POEM on Agent Observations

- POEM can be used as an encoder, and can be trained **in an equivalent few-shot manner** (meta-learned) from offline data:
  - Given a **trajectory of partial observations** from a set of environments, and a new observation viewpoint, can you **determine which environment the observation is from**?

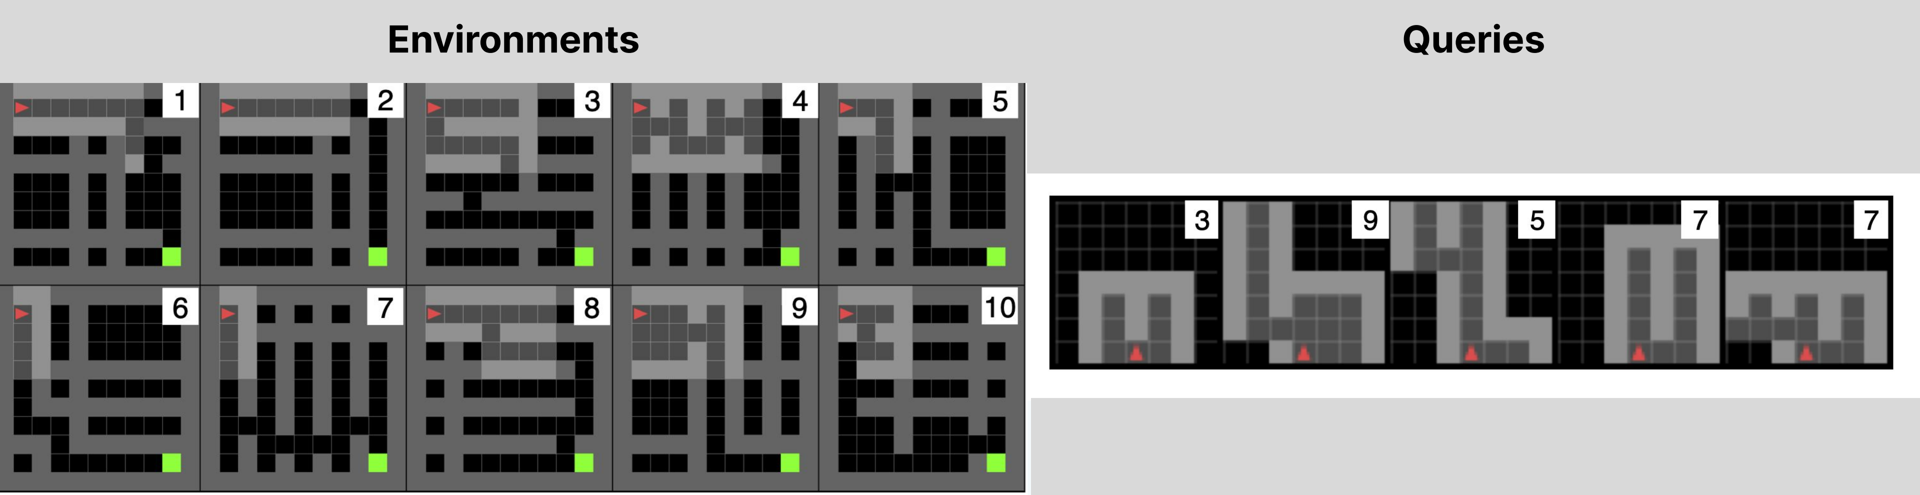

Environment Recognition Results

POEM demonstrates **benefits** over both meta-learning and recurrent network baselines at **environment observation recognition**

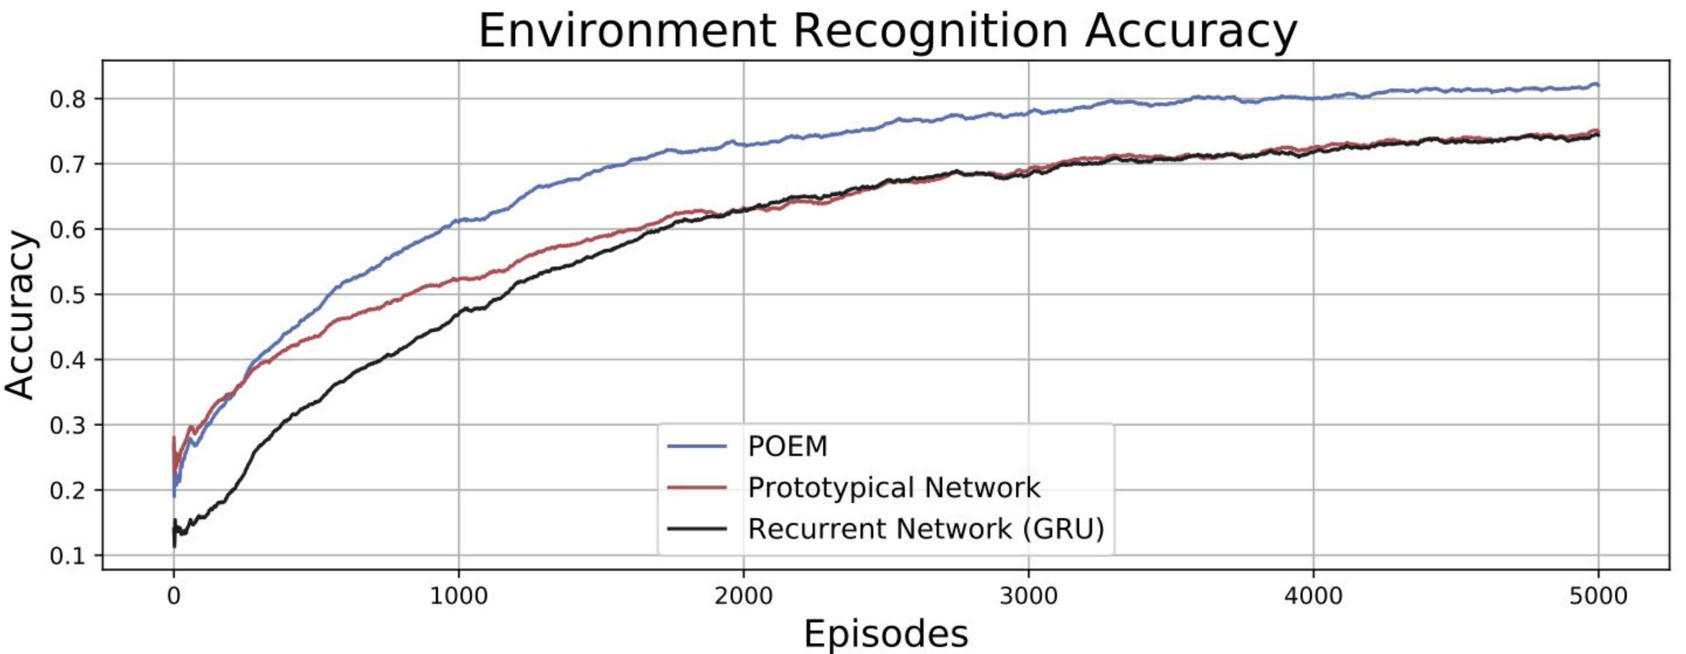

Investigating the Learned Representation

- **Decoder** can be **trained on learned representations** to **generate environment map**
- POEM encoder could now be used **online** to **integrate partial agent observations** into **updateable environment representation**, or decoded into online **environment map**

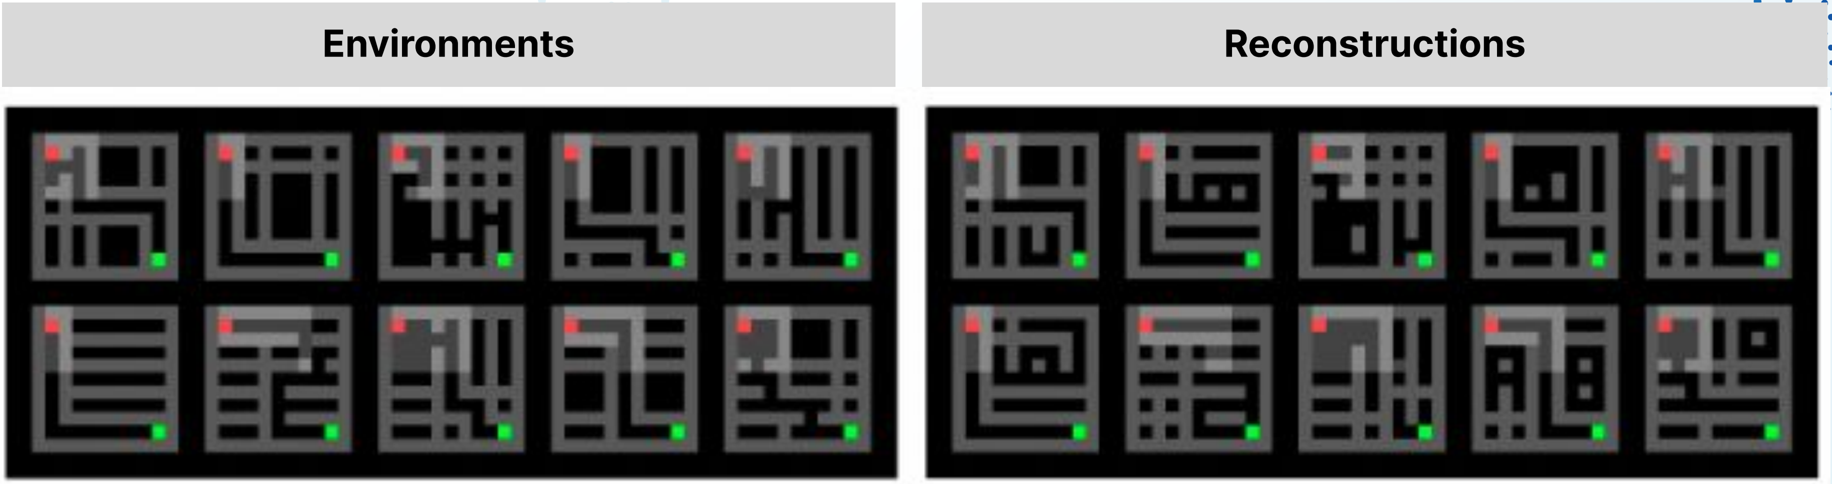

Summary

- Introduced **POEM**, a **contrastive meta-learning** approach for **few-shot learning** in **partially observable** settings

- **Probabilistic formalism** enables **consistent contrastive representation learning** from **partial observations**

- Potential **application** to **learning representations** of an **environment** from **agent-centric observations**