Teaching Agents how to Map: Spatial Reasoning for Multi-Object Navigation




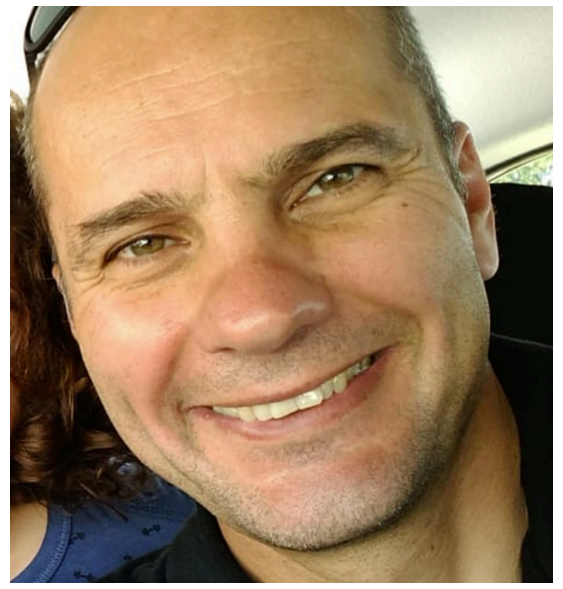

Pierre Marza
Laetitia Matignon
Olivier Simonin
Christian Wolf

pierre.marza@insa-lyon.fr

Multi-Object Navigation

Target objects

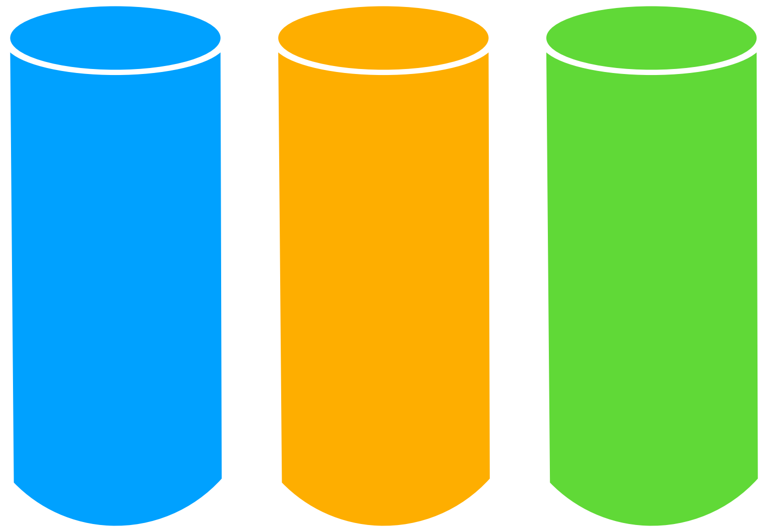

Matterport 3D [2]

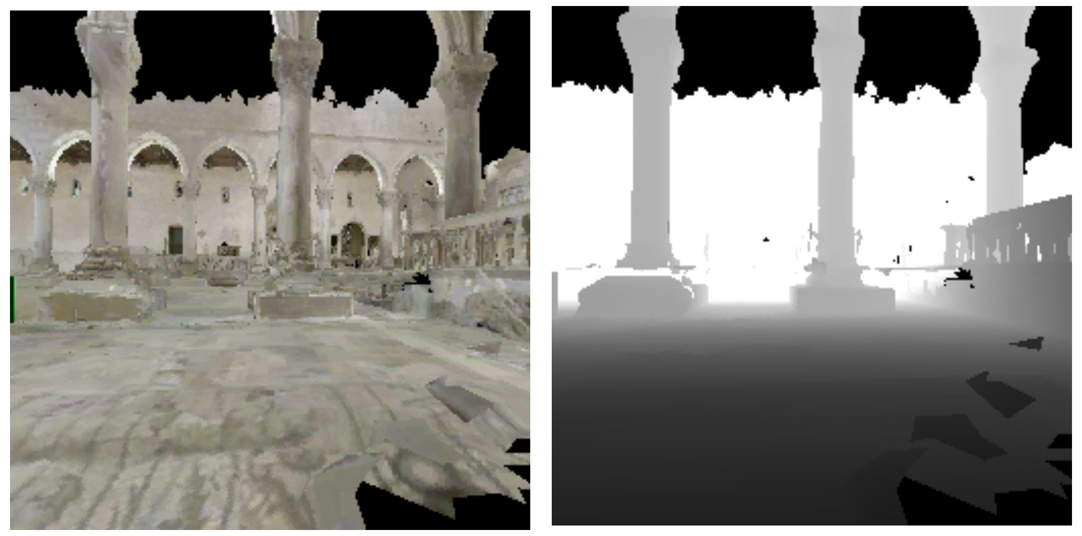

Important abilities

- Building a useful representation of the environment
- Taking advantage of such representation to plan and navigate efficiently

[1] Wani et al. MultiON: Benchmarking Semantic Map Memory using Multi-Object Navigation, NeurIPS 2020
[2] Chang et al. Matterport 3D: Learning from RGB-D Data in Indoor Environments, 3DV 2017

Multi-Object Navigation

Why is it interesting to benchmark mapping capabilities ?

- **Sequential** task
    - Remembering previously encountered objects
    - Mapping the environment

- **External objects** as objectives
    - Agent can't rely on knowledge about indoor layouts
    - Focus on memory

[1] Wani et al. MultiON: Benchmarking Semantic Map Memory using Multi-Object Navigation, NeurIPS 2020
[2] Chang et al. Matterport 3D: Learning from RGB-D Data in Indoor Environments, 3DV 2017

Baseline Architectures

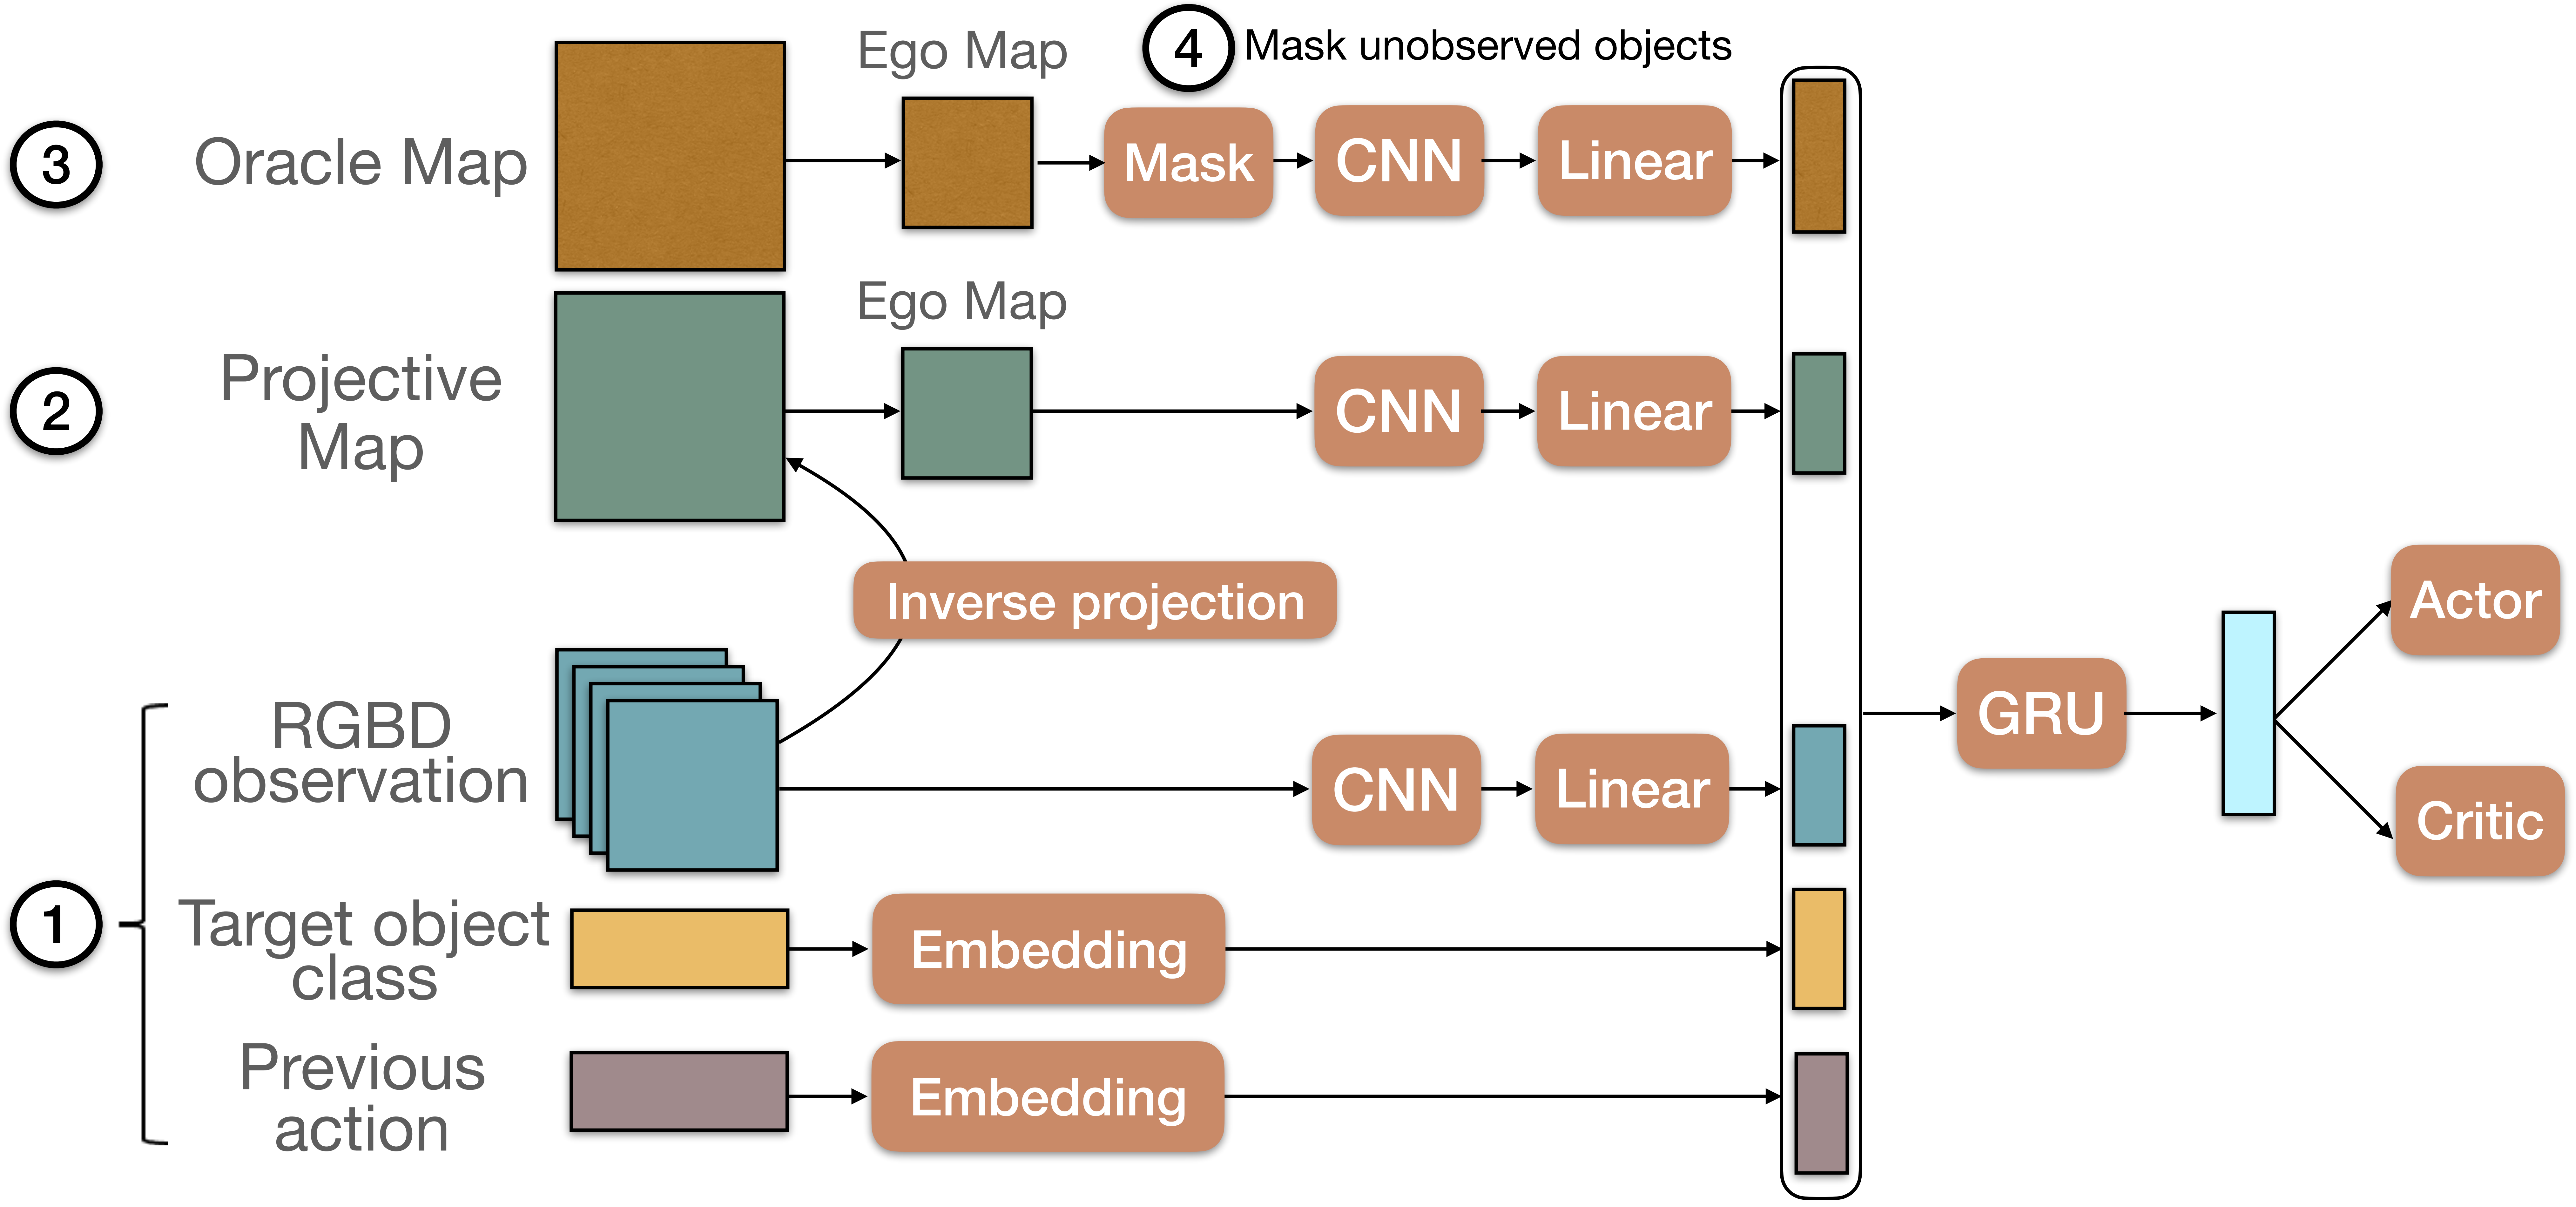

[1] Wani et al. MultiON: Benchmarking Semantic Map Memory using Multi-Object Navigation, NeurIPS 2020

Baseline Architectures

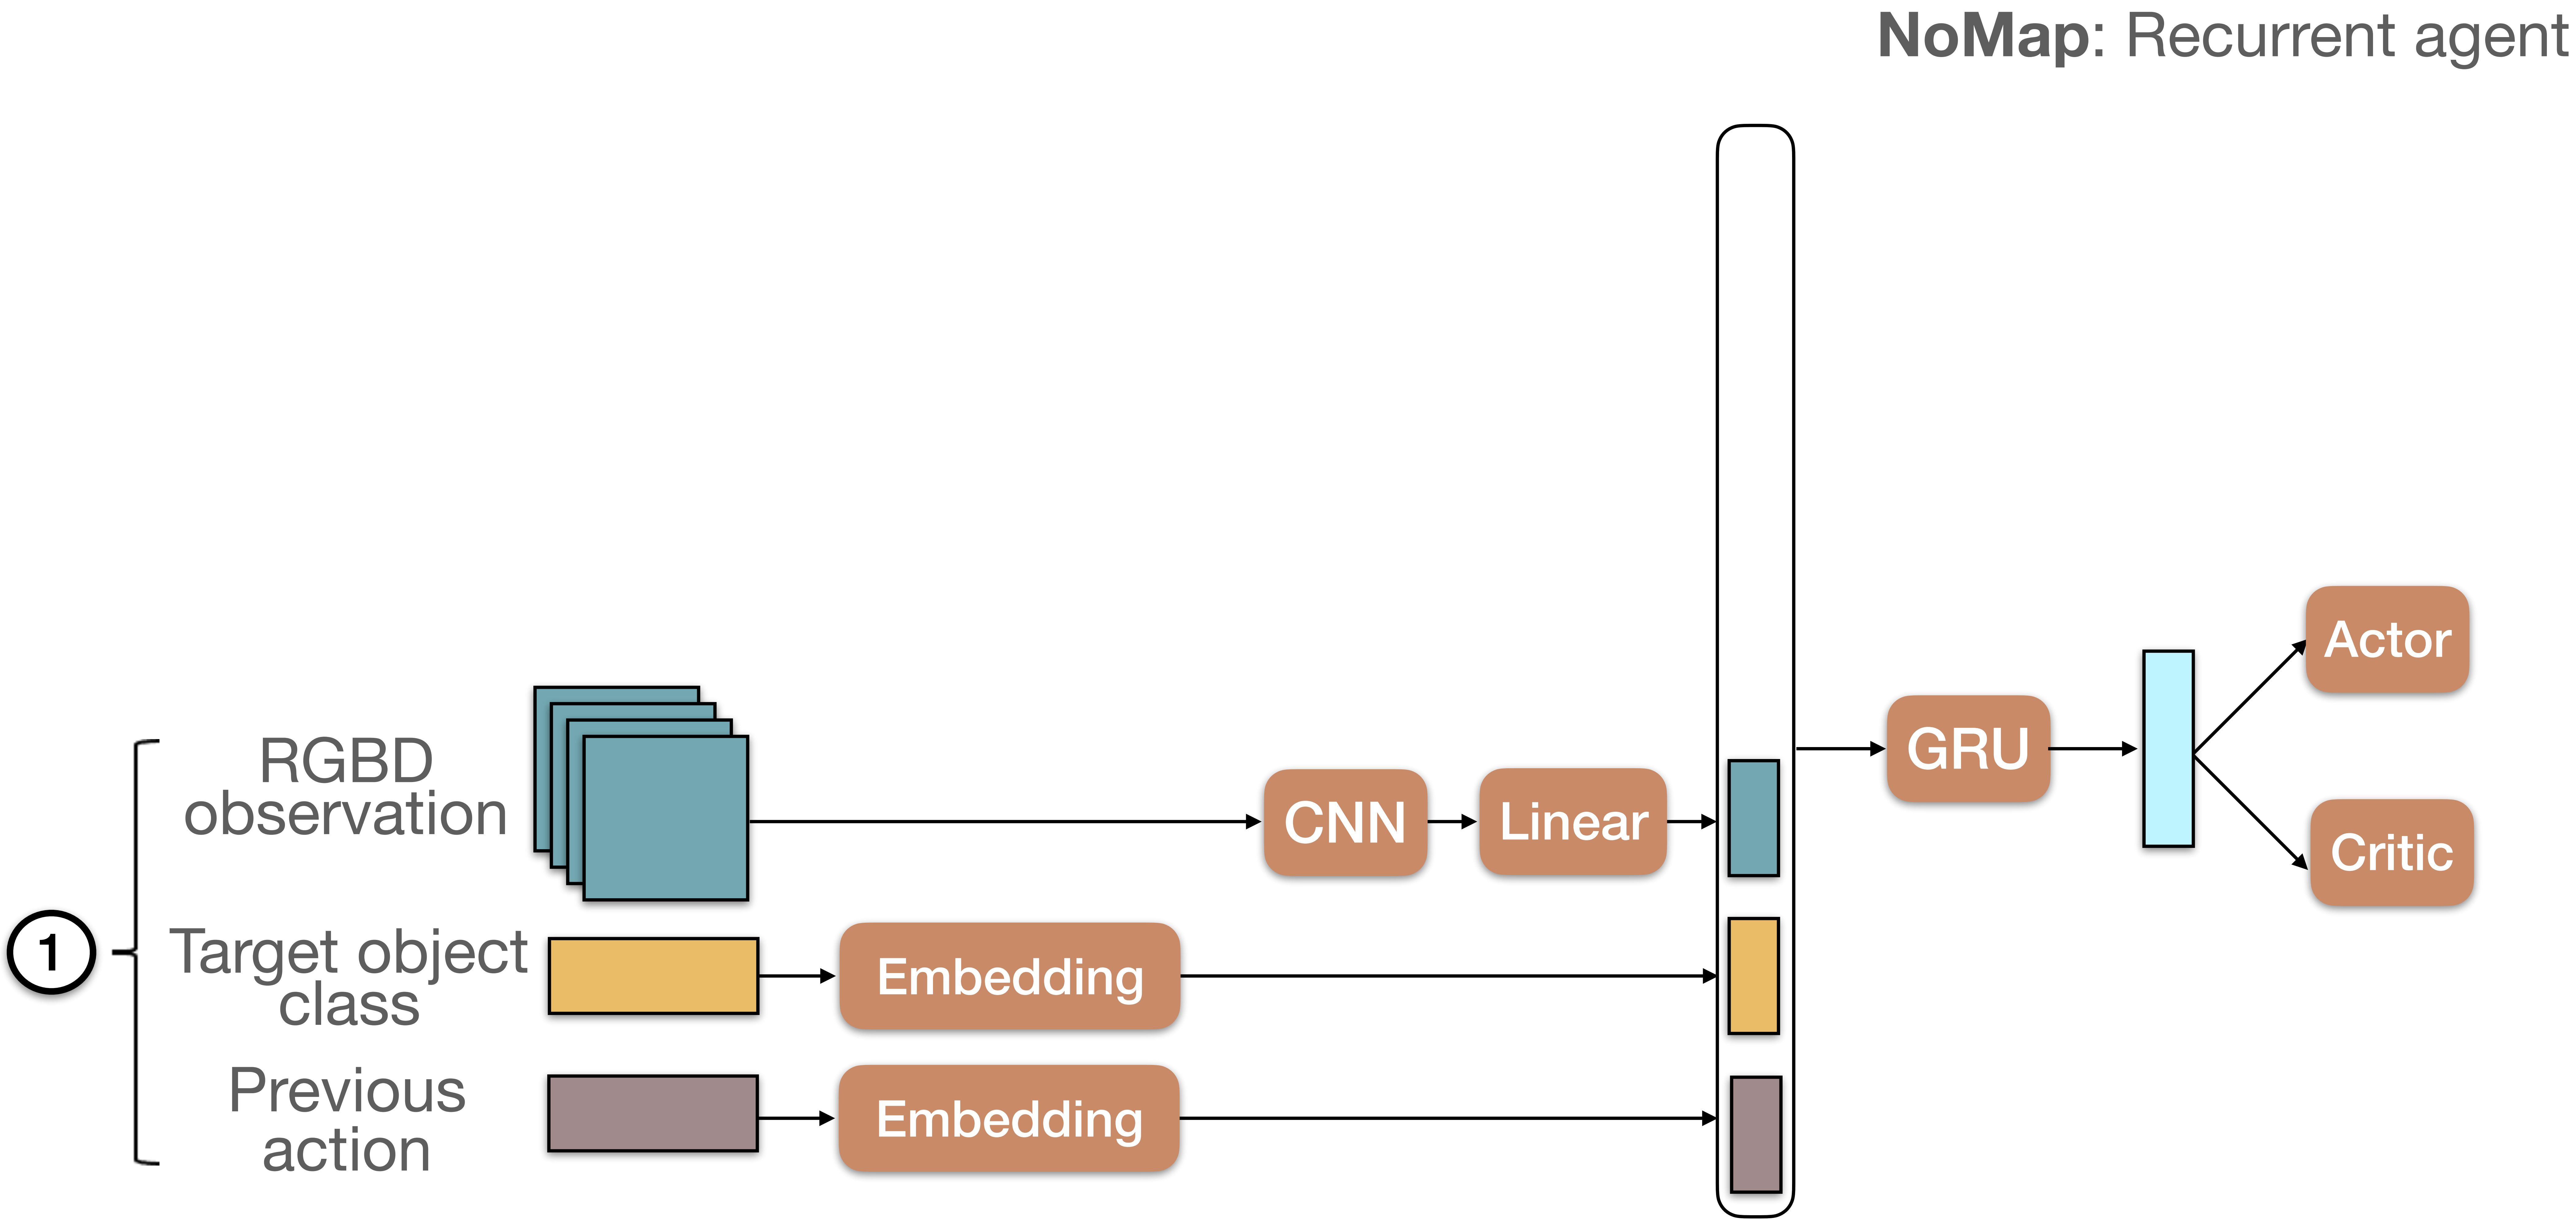

[1] Wani et al. MultiON: Benchmarking Semantic Map Memory using Multi-Object Navigation, NeurIPS 2020

Baseline Architectures

ProjNeuralMap: Projects extracted image features using depth [3]

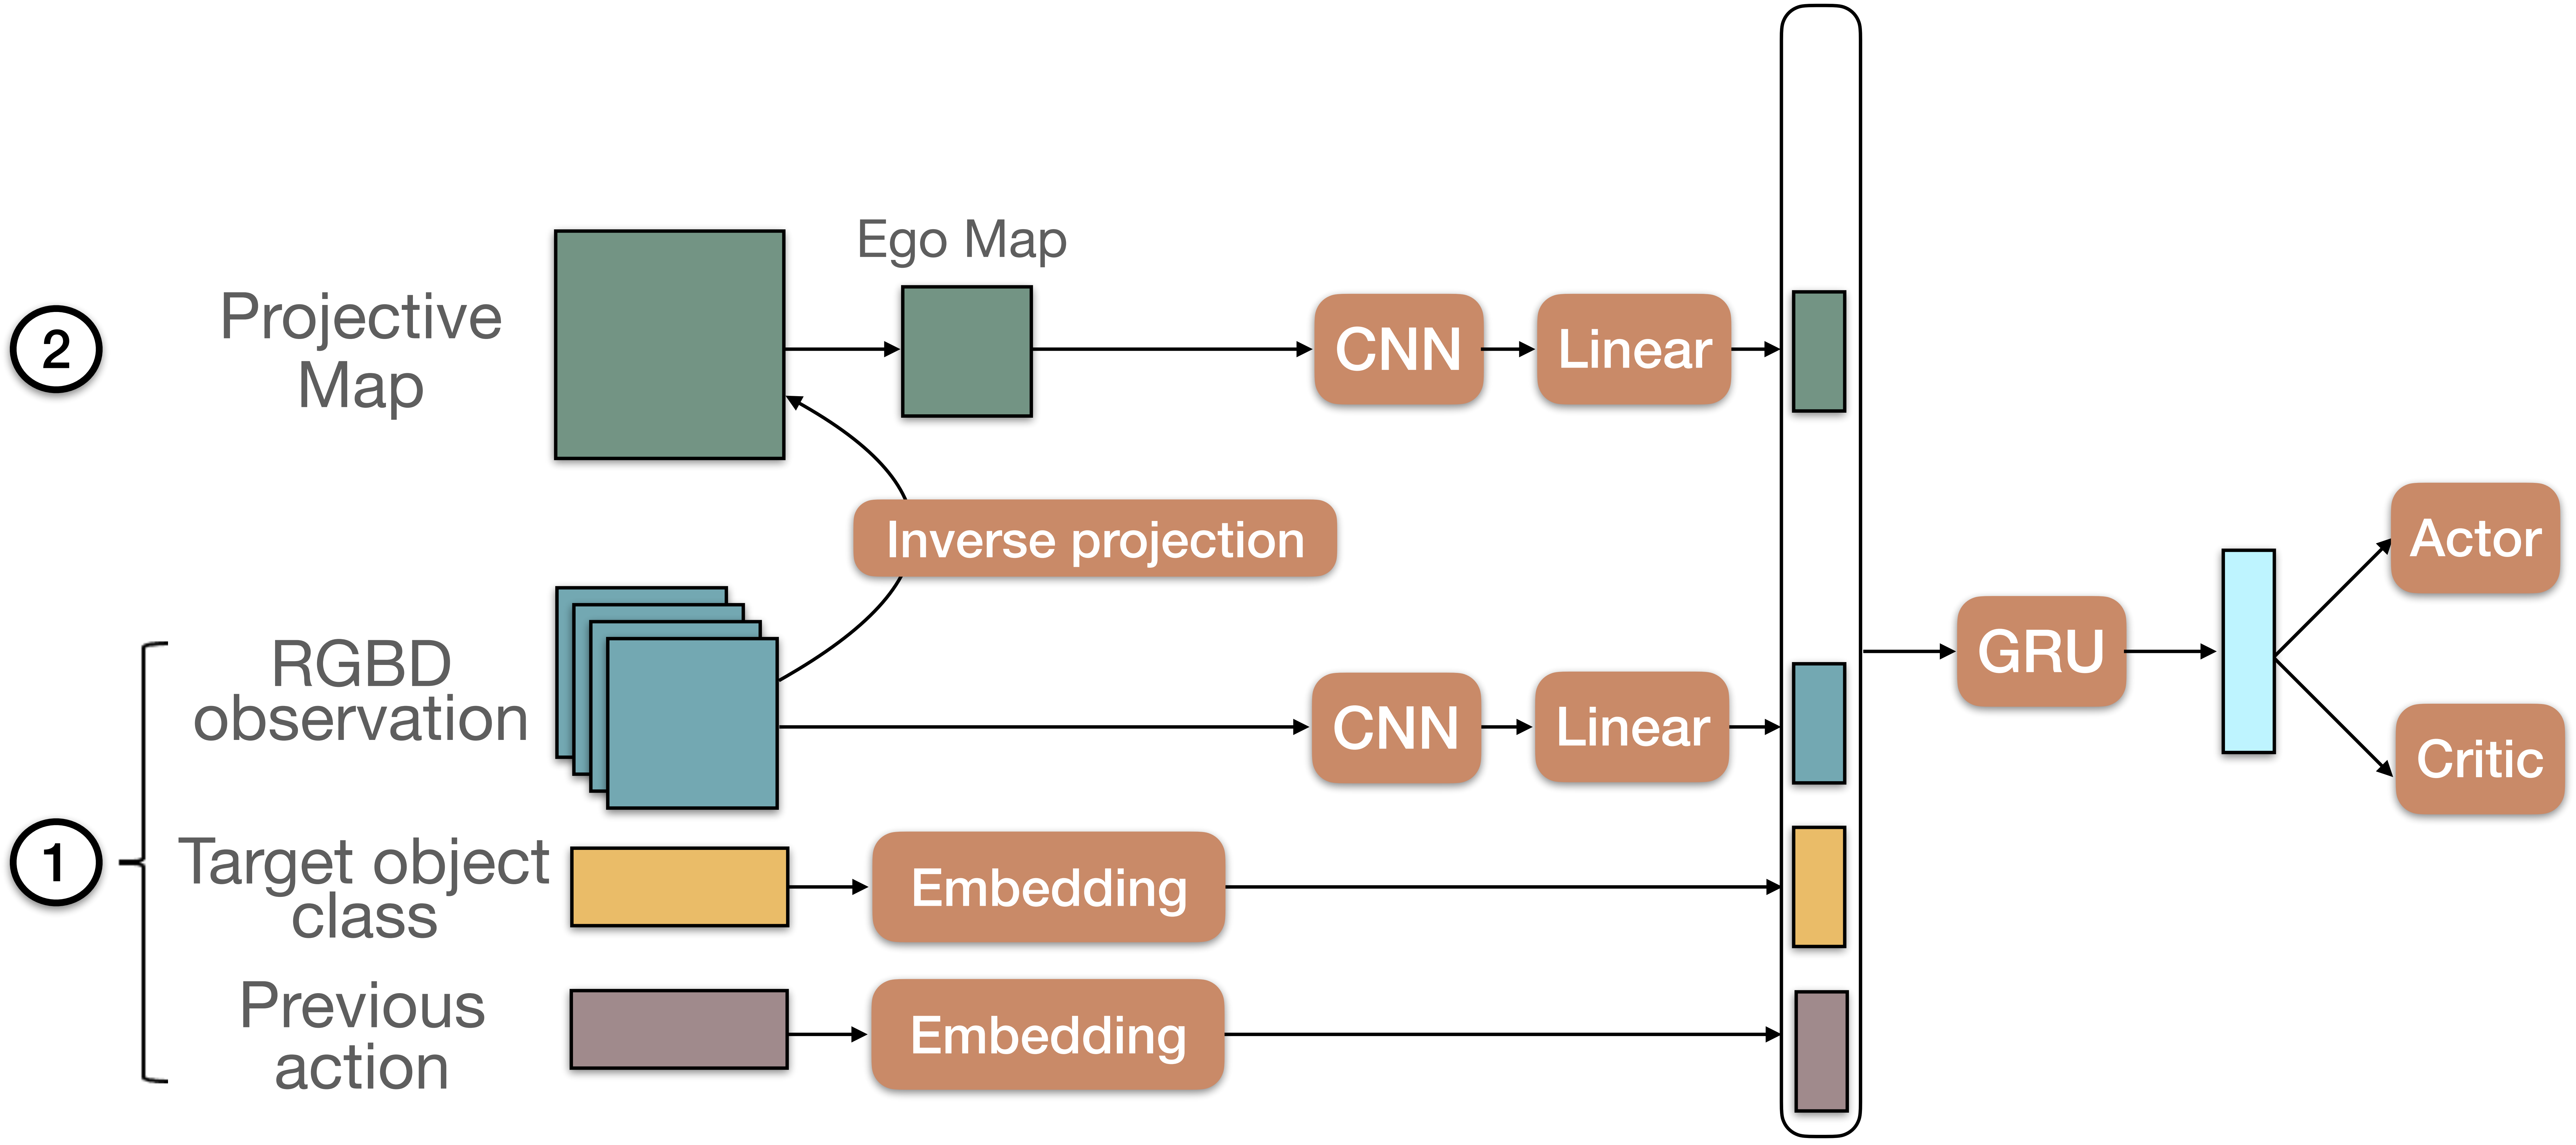

[1] Wani et al. MultiON: Benchmarking Semantic Map Memory using Multi-Object Navigation, NeurIPS 2020
[3] Henriques et al. Mapnet: An allocentric spatial memory for mapping environments, CVPR 2018
[4] Beeching et al. EgoMap: Projective mapping and structured egocentric memory for Deep RL, ECML-PKDD 2020

Baseline Architectures

OracleMap: GT occupancy map and object locations

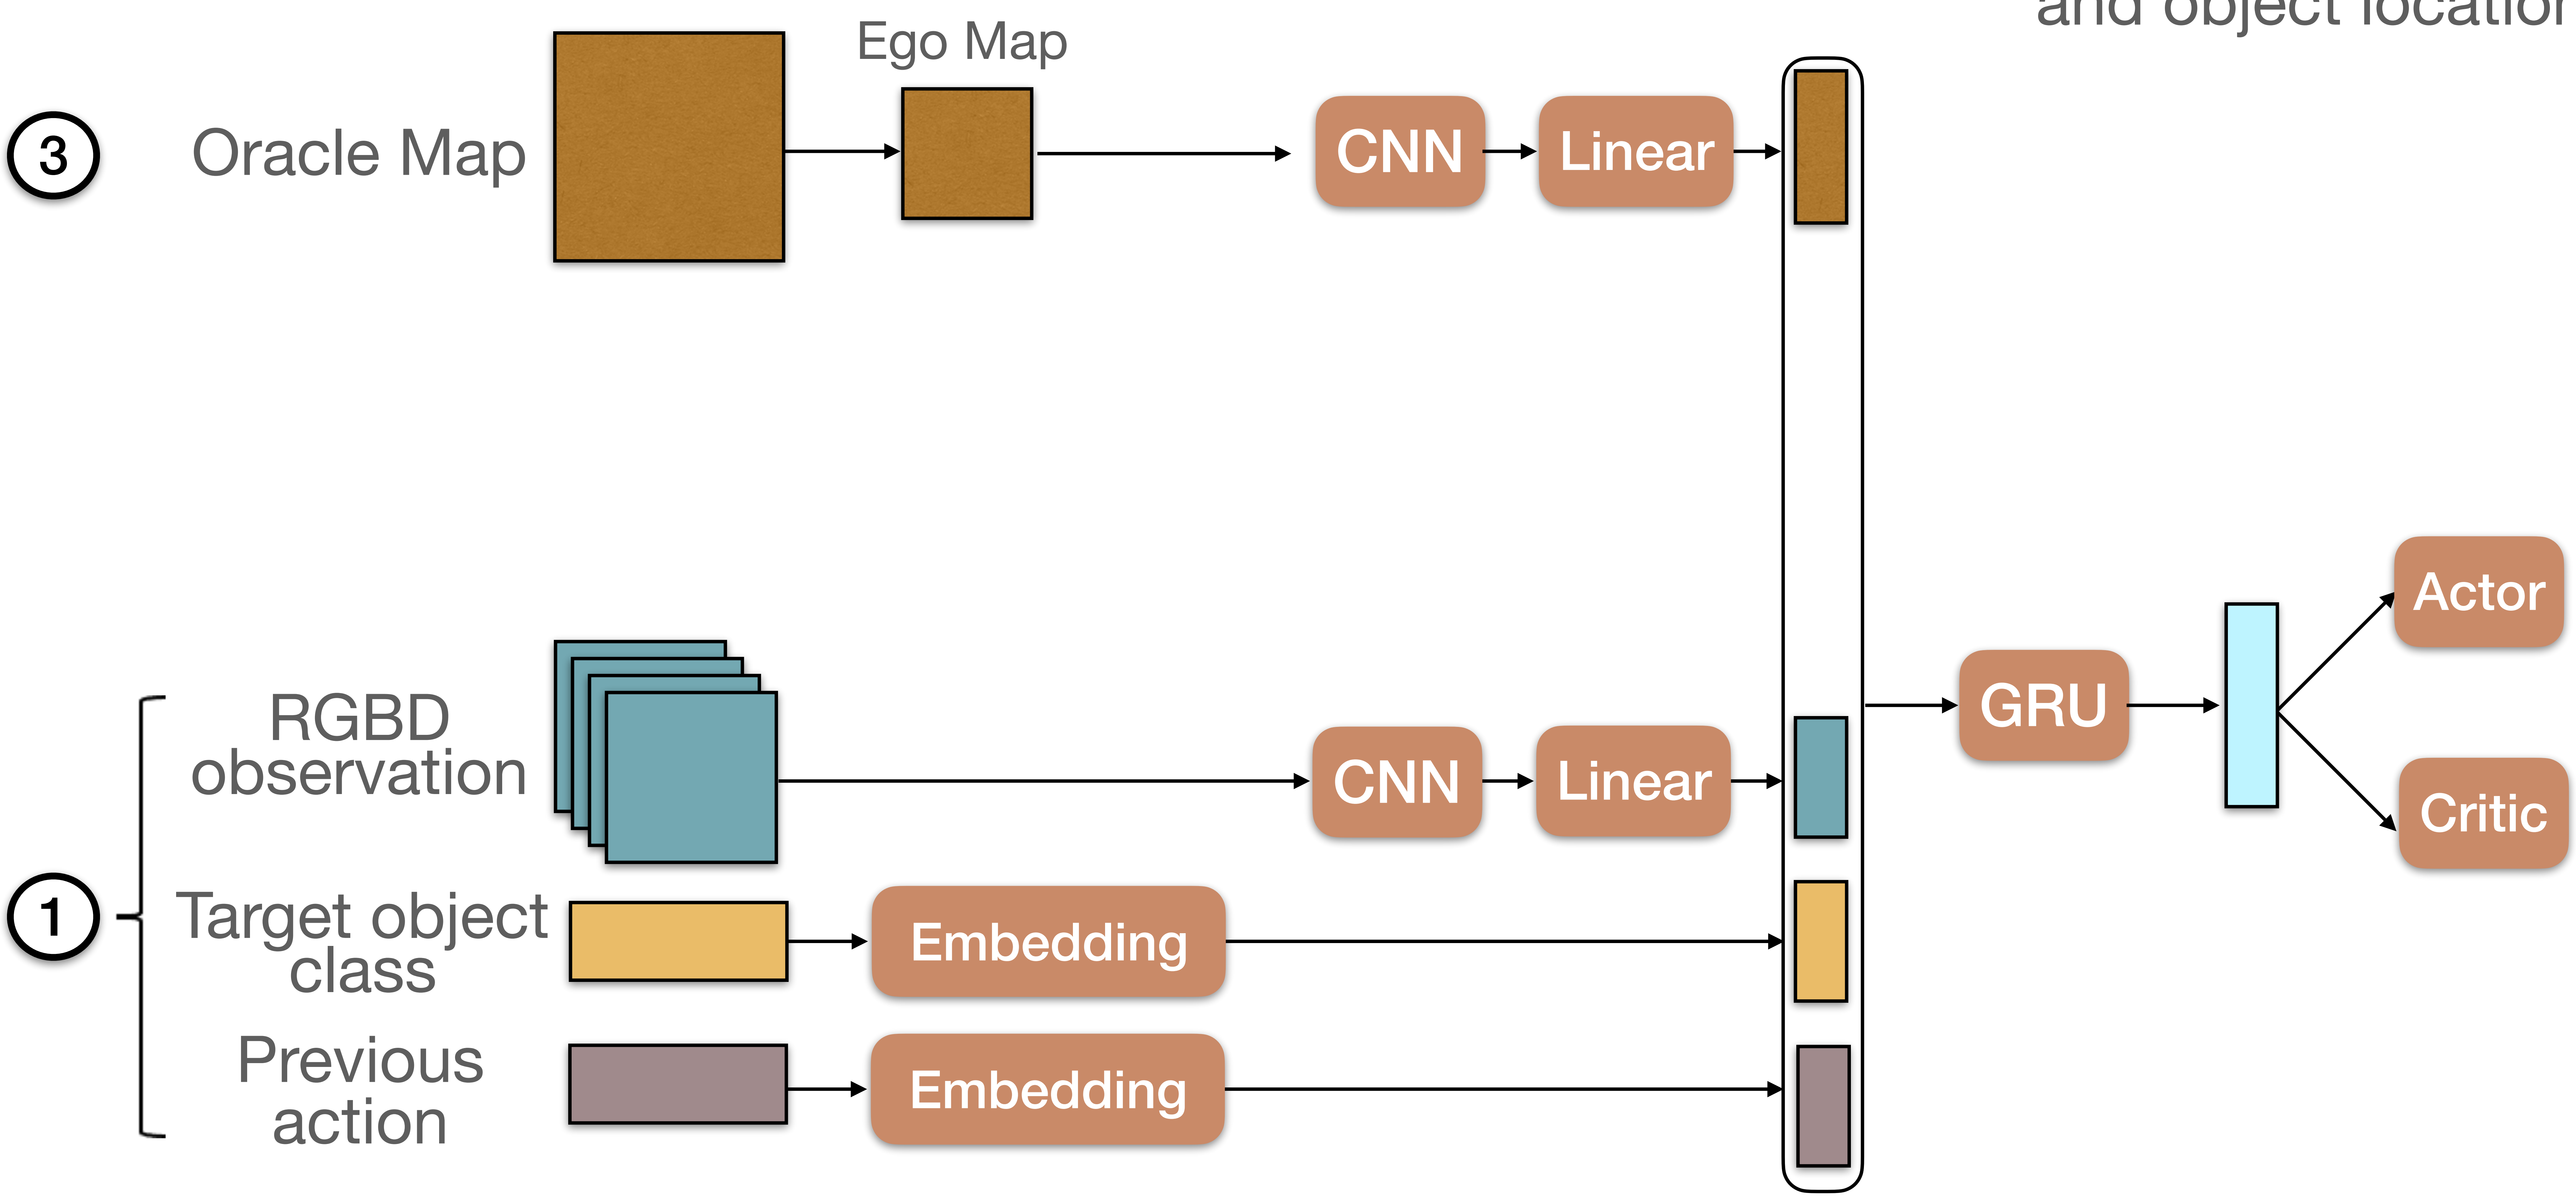

[1] Wani et al. MultiON: Benchmarking Semantic Map Memory using Multi-Object Navigation, NeurIPS 2020
[3] Henriques et al. Mapnet: An allocentric spatial memory for mapping environments, CVPR 2018
[4] Beeching et al. EgoMap: Projective mapping and structured egocentric memory for Deep RL, ECML-PKDD 2020

Baseline Architectures

OracleEgoMap: GT map revealed in regions already seen

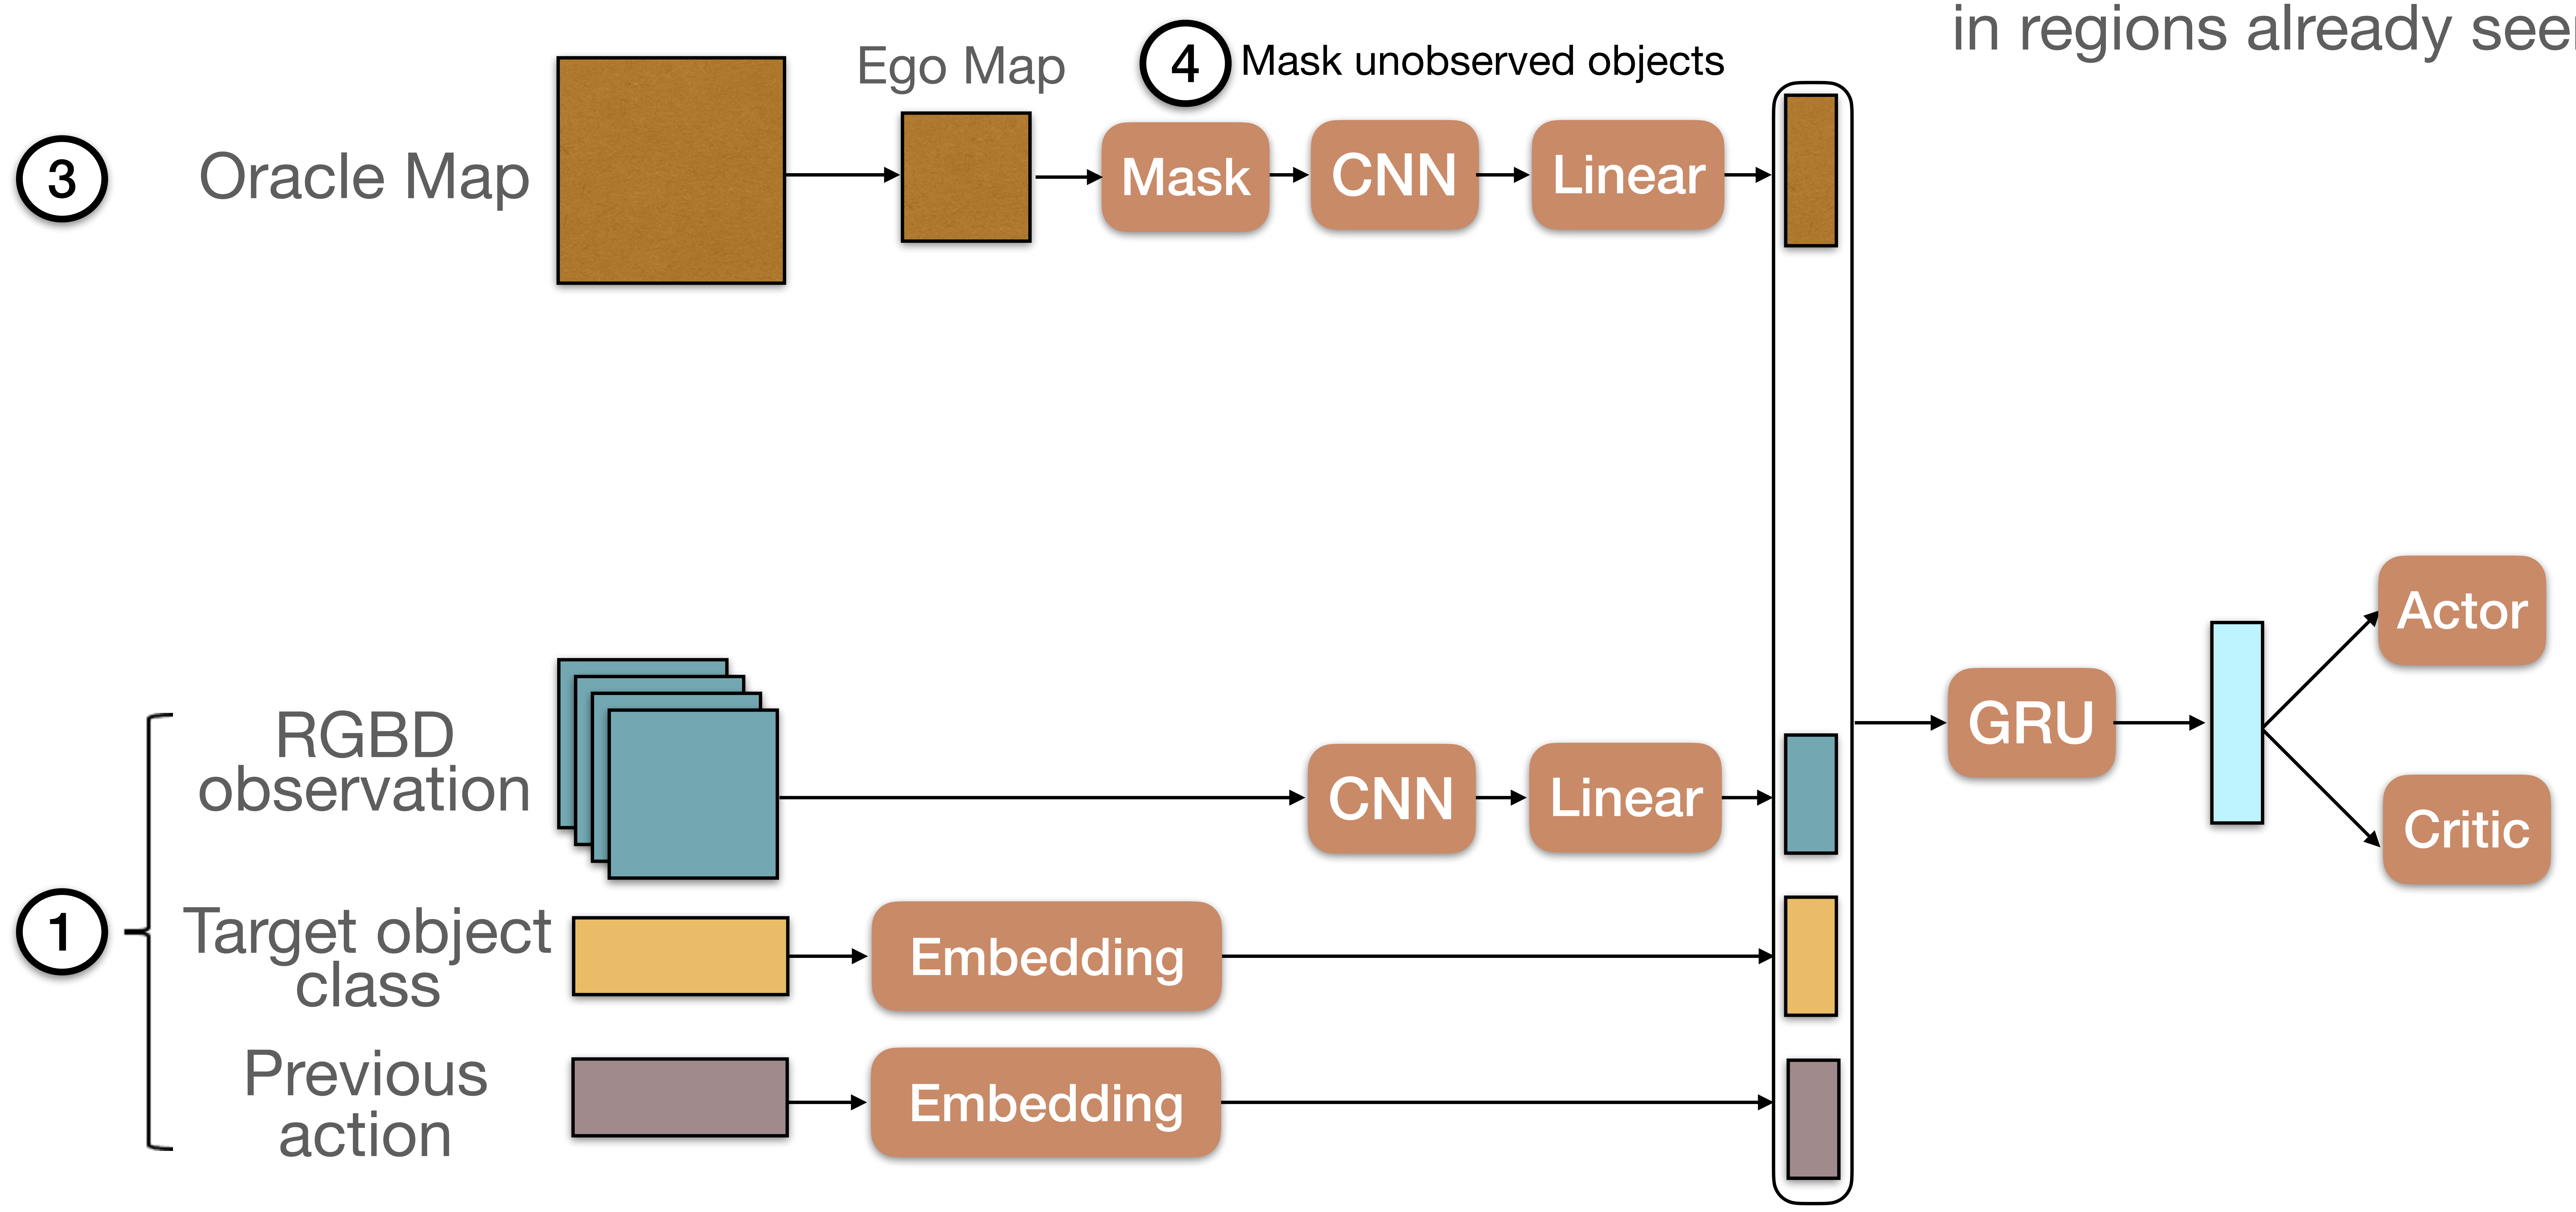

[1] Wani et al. MultiON: Benchmarking Semantic Map Memory using Multi-Object Navigation, NeurIPS 2020
[3] Henriques et al. Mapnet: An allocentric spatial memory for mapping environments, CVPR 2018
[4] Beeching et al. EgoMap: Projective mapping and structured egocentric memory for Deep RL, ECML-PKDD 2020

Baseline Architectures

Trained with Proximal Policy Optimization (PPO) [5]

$$R_t = 1_{[\text{reached-goal}]} \cdot R_{\text{goal}} + R_{\text{closer}} + R_{\text{time-penalty}}$$

[1] Wani et al. MultiON: Benchmarking Semantic Map Memory using Multi-Object Navigation, NeurIPS 2020
[5] Schulman et al. Proximal policy optimization algorithms, arXiv preprint, 2017

Inspiration

Behavioral Studies of Human Spatial Navigation

- **Sense of direction**
  - scene- and orientation- dependent pointing (SOP)
- **Judgment of relative distance**
  - Compare the relative distance to several goals

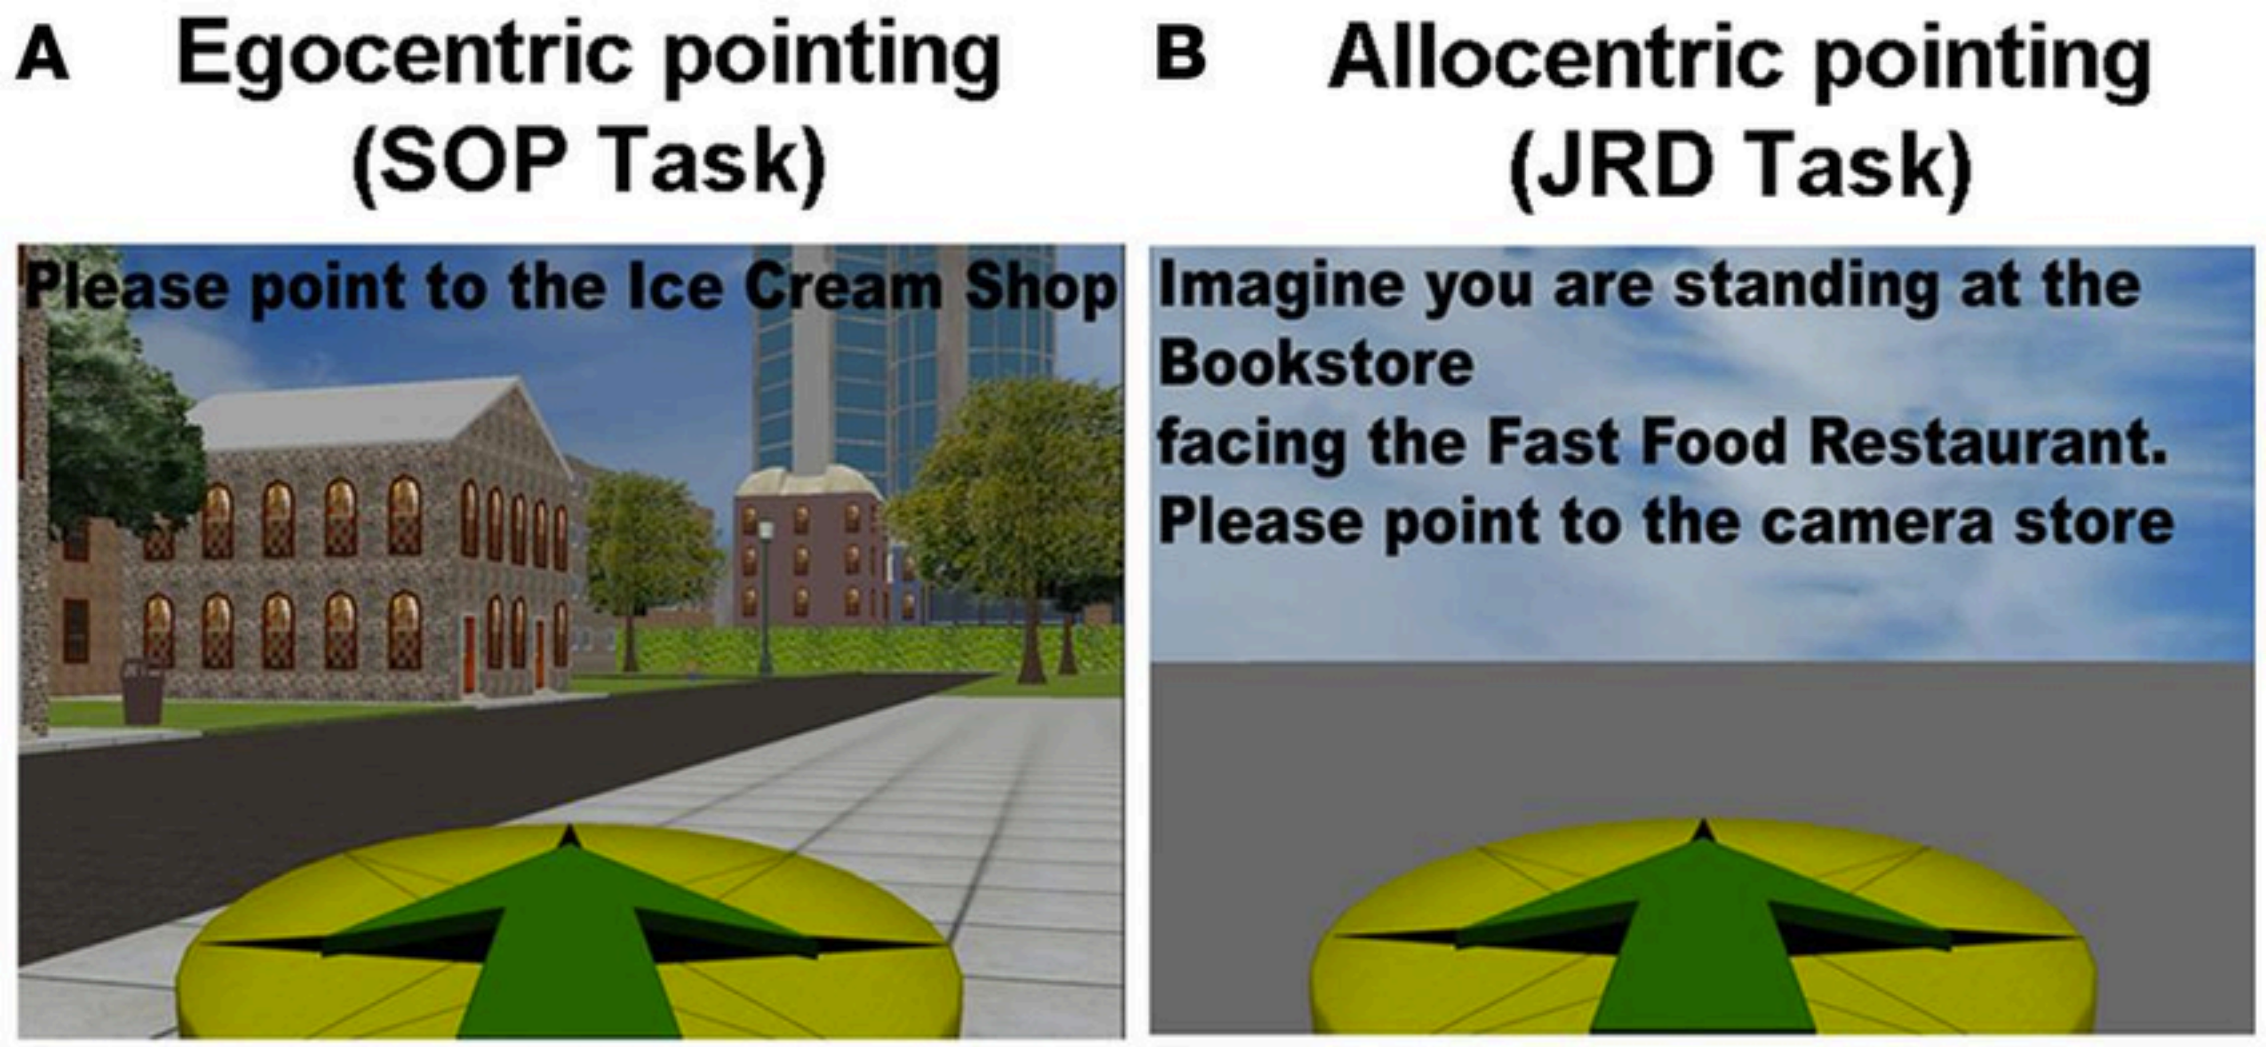

Reproduced from [6]

[6] Ekstrom et al. A critical review of the allocentric spatial representation and its neural underpinnings: Toward a network-based perspective, Frontiers in Human Neuroscience 2014

Auxiliary tasks

Direction prediction

- **Classification** problem
- Angles in the range [0, 360] divided into bins

$$\mathscr{L}_\phi = \frac{1}{|\mathscr{U}_k|\, T} \sum_{\tau \in \mathscr{U}_k} \sum_{t=0}^{T-1} \left[-1_t \sum_{c=1}^{K} \phi^*_{t,c} \log p(\hat{\phi}_{t,c}) \right]$$

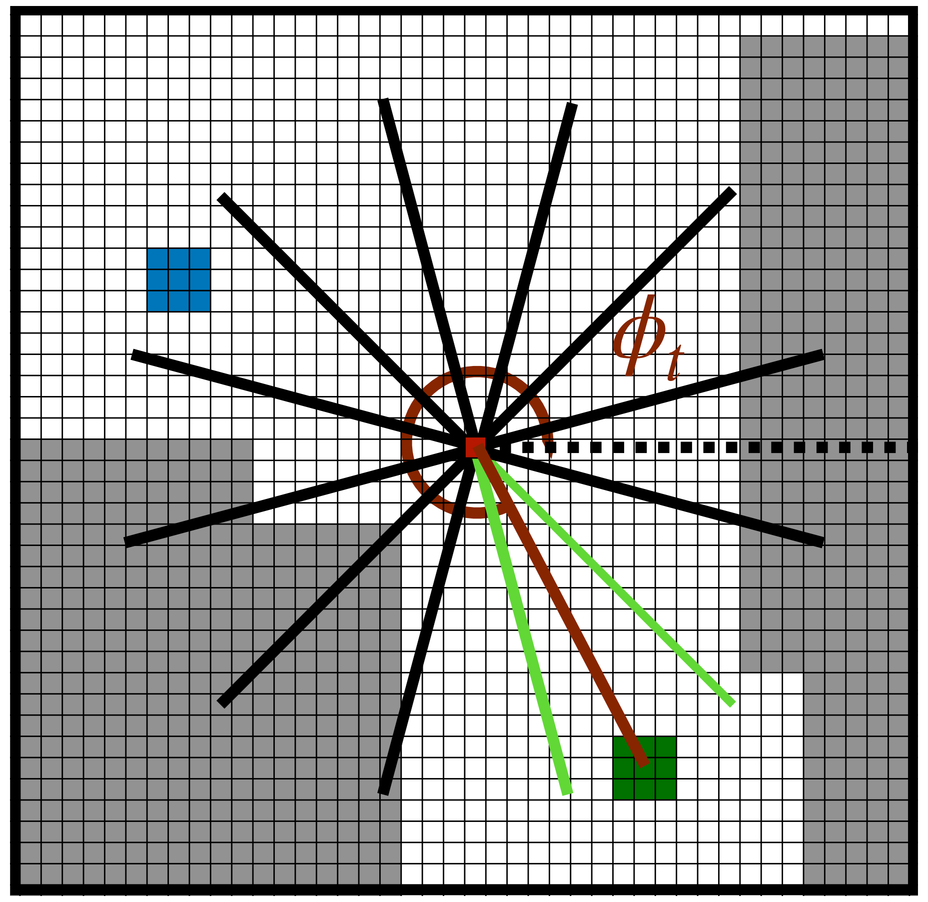

Distance prediction

- **Classification** problem
- Euclidian distances on the grid egocentric map divided into bins

$$\mathscr{L}_d = \frac{1}{|\mathscr{U}_k|\, T} \sum_{\tau \in \mathscr{U}_k} \sum_{t=0}^{T-1} \left[-1_t \sum_{c=1}^{L} d^*_{t,c} \log p(\hat{d}_{t,c}) \right]$$

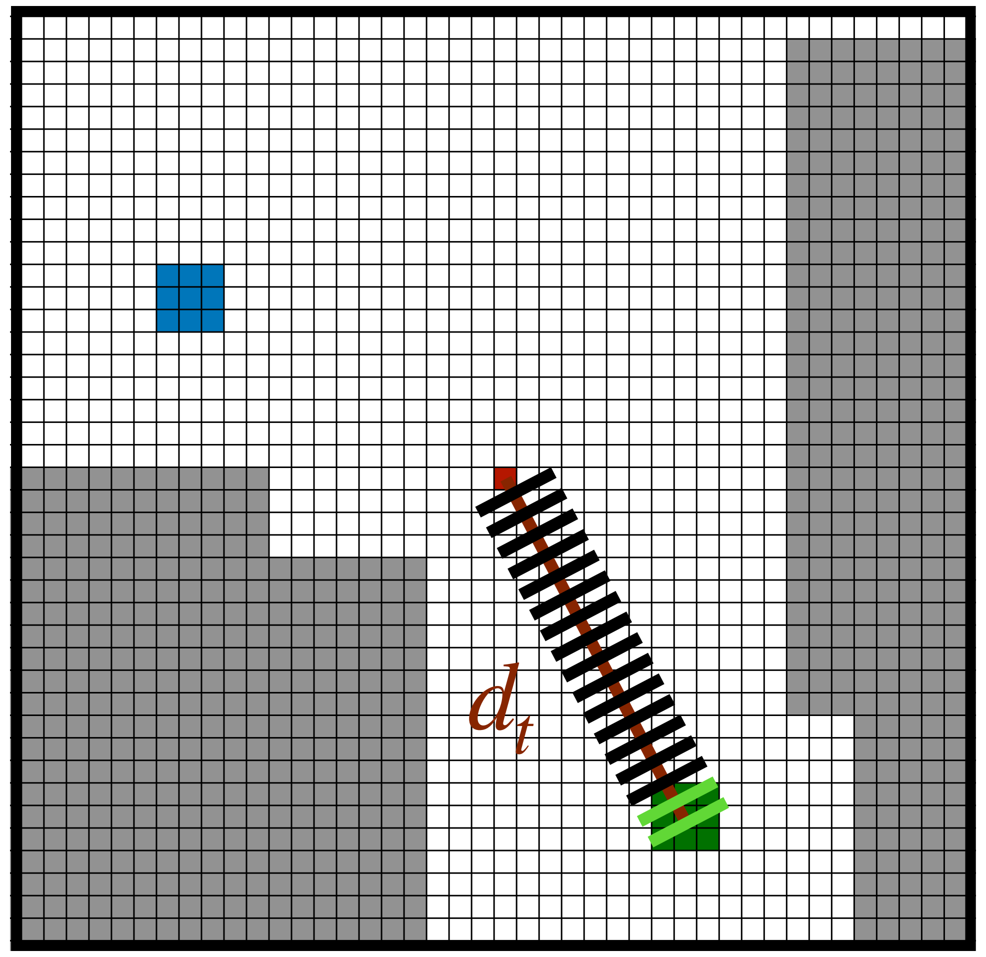

Only target objects that have already been seen

Auxiliary tasks

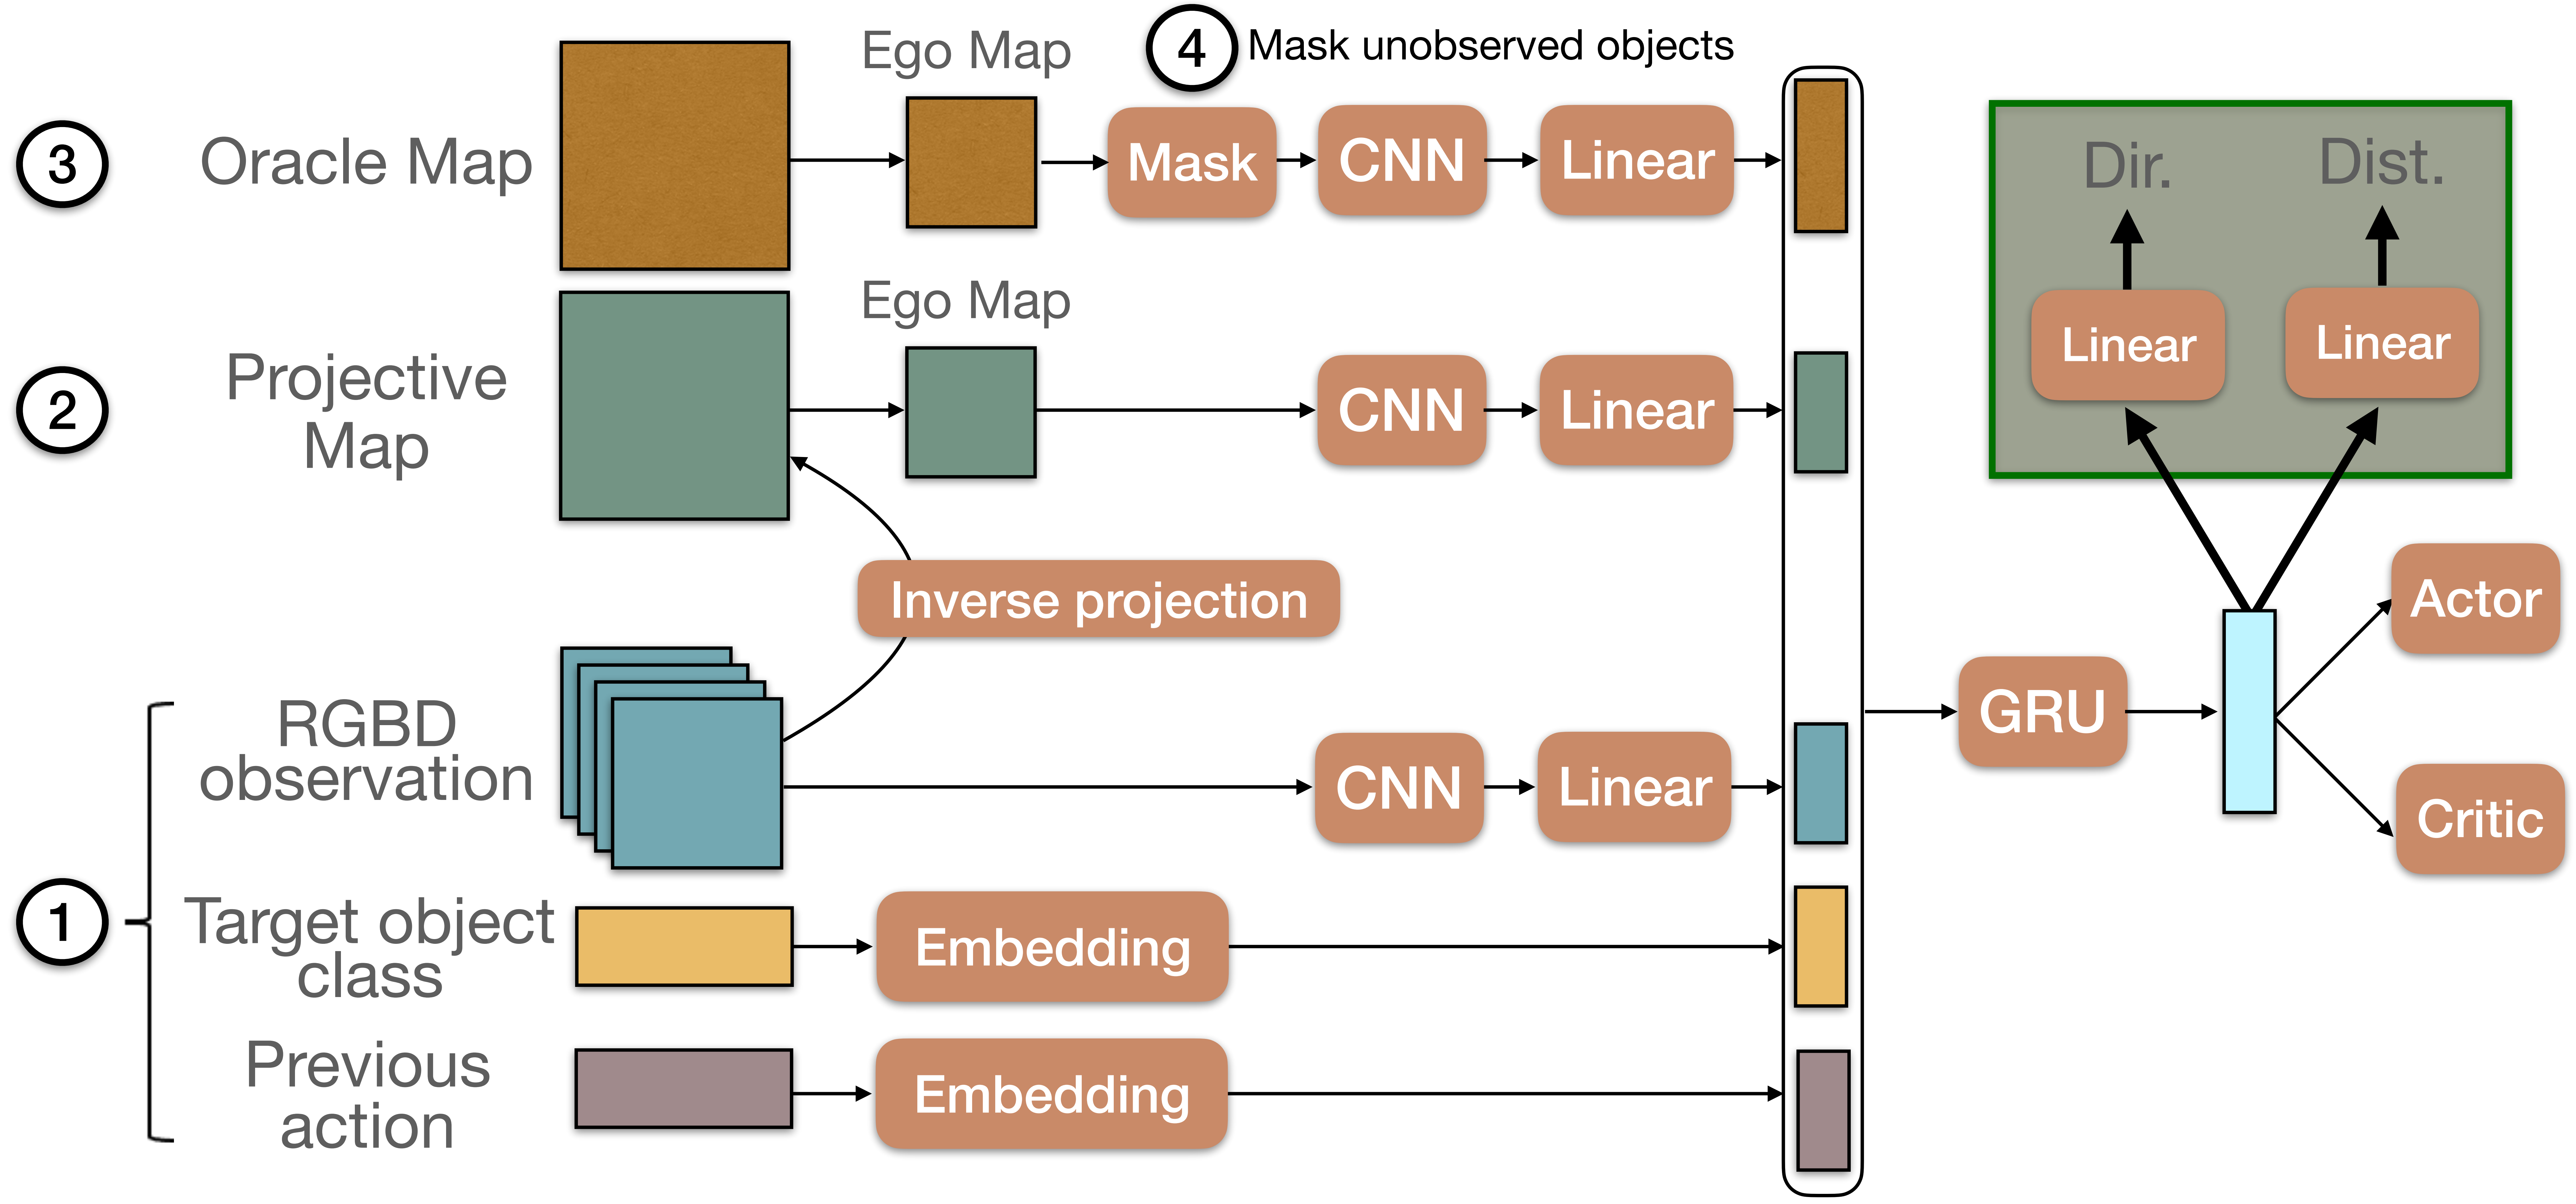

[1] Wani et al. MultiON: Benchmarking Semantic Map Memory using Multi-Object Navigation, NeurIPS 2020
[2] Chang et al. Matterport 3D: Learning from RGB-D Data in Indoor Environments, 3DV 2017
[3] Henriques et al. Mapnet: An allocentric spatial memory for mapping environments, CVPR 2018
[4] Beeching et al. EgoMap: Projective mapping and structured egocentric memory for Deep RL, ECML-PKDD 2020

Experimental results

Actions

- **FORWARD**: moves foward 0.25m
- **LEFT**: turns left 30°
- **RIGHT**: turns right 30°
- **FOUND**: signals the agents thinks it has reached the target

Metrics

- **Success**: Percentage of successful episodes
- **Progress**: Percentage of objects found in an episode

- **SPL** (Success weighted by Path Length): $SPL = s \cdot d / \max(p, d)$
  - s is the success binary indicator
  - p is the distance travelled by the agent
  - d is the total shortest path

- **PPL** (Progress weighted by Path Length): $\mathrm{PPL} = \bar{s} \cdot \bar{d} / \max(p, \bar{d})$
  - $\bar{s}$ is the progress
  - $\bar{d}$ is the shortest path to reach all found objects

Experimental results

Ablation Study - Impact of each loss on validation performance

Agent	Dir.	Dist.	Success	Progress	SPL	PPL	Comparable
OracleMap	–	–	51.4± 2.0	61.2± 0.8	41.3± 1.5	49.0± 0.7	–
OracleEgoMap	–	–	37.1± 1.0	51.8± 0.9	29.7± 0.7	41.2± 1.1	–
ProjNeuralMap	–	–	27.3± 3.5	44.8± 2.6	19.5± 0.9	32.5± 0.3	✓
	✓	–					✓
	✓	✓					✓

Experimental results

Ablation Study - Impact of each loss on validation performance

Agent	Dir.	Dist.	Success	Progress	SPL	PPL	Comparable
OracleMap	–	–	51.4± 2.0	61.2± 0.8	41.3± 1.5	49.0± 0.7	–
OracleEgoMap	–	–	37.1± 1.0	51.8± 0.9	29.7± 0.7	41.2± 1.1	–
ProjNeuralMap	–	–	27.3± 3.5	44.8± 2.6	19.5± 0.9	32.5± 0.3	✓
	✓	–	43.0± 5.7	58.9± 4.6	30.7± 4.9	42.1± 4.0	✓
	✓	✓					✓

Experimental results

Ablation Study - Impact of each loss on validation performance

Agent	Dir.	Dist.	Success	Progress	SPL	PPL	Comparable
OracleMap	–	–	51.4± 2.0	61.2± 0.8	41.3± 1.5	49.0± 0.7	–
OracleEgoMap	–	–	37.1± 1.0	51.8± 0.9	29.7± 0.7	41.2± 1.1	–
ProjNeuralMap	–	–	27.3± 3.5	44.8± 2.6	19.5± 0.9	32.5± 0.3	✓
	✓	–	43.0± 5.7	58.9± 4.6	30.7± 4.9	42.1± 4.0	✓
	✓	✓	**54.2 ± 3.5**	**67.4 ± 2.3**	**37.8 ± 0.8**	**47.4 ± 0.4**	✓

Both auxiliary tasks have a positive impact and are complementary

Experimental results

Test performance - Do the auxiliary tasks improve the downstream objective ?

Agent	**Aux. Sup.**	**Success**	**Progress**	**SPL**	**PPL**	**Comparable**
OracleMap	–	41.0± 1.8	50.3± 0.9	32.2± 0.9	39.4± 0.4	–
OracleEgoMap	–	25.8± 1.1	41.0± 1.0	19.7± 0.7	30.7± 1.3	–
ProjNeuralMap	–	18.0± 1.3	34.4± 1.7	12.3± 0.4	24.1± 0.1	✓
	✓	**38.0 ± 2.4**	**52.6 ± 2.0**	**25.7 ± 0.2**	**36.2 ± 1.1**	✓
NoMap	–	7.4± 0.2	21.7± 0.2	6.0± 0.1	17.3± 0.4	✓
	✓					✓

Experimental results

Test performance - Can an unstructured recurrent agent learn to map ?

Agent	**Aux. Sup.**	**Success**	**Progress**	**SPL**	**PPL**	**Comparable**
OracleMap	–	41.0± 1.8	50.3± 0.9	32.2± 0.9	39.4± 0.4	–
OracleEgoMap	–	25.8± 1.1	41.0± 1.0	19.7± 0.7	30.7± 1.3	–
ProjNeuralMap	–	18.0± 1.3	34.4± 1.7	12.3± 0.4	24.1± 0.1	✓
	✓	**38.0** ± **2.4**	**52.6** ± **2.0**	**25.7** ± **0.2**	**36.2** ± **1.1**	✓
NoMap	–	7.4± 0.2	21.7± 0.2	6.0± 0.1	17.3± 0.4	✓
	✓	22.4± 2.0	38.2± 2.0	15.2± 2.2	26.4± 2.3	✓

Experimental results

Winning entry of the MultiON Challenge, CVPR 2021 Embodied AI Workshop

Agent/Method	— Test Challenge —				— Test Standard —			
	Success	Progress	SPL	PPL	Success	Progress	SPL	PPL
Ours (Auxiliary losses)	**55**	**67**	**35**	**44**	57	70	36	45
Team 2	52	64	32	38	62	71	34	39
Team 3	41	57	26	36	43	57	27	36
ProjNeuralMap (Challenge baseline)	–	–	–	–	12	29	6	16
NoMap (Challenge baseline)	–	–	–	–	5	19	3	13

CVPR 2021 Embodied AI Workshop: https://embodied-ai.org/

MultiON Challenge: http://multion-challenge.cs.sfu.ca/

MultiON Challenge video : https://www.youtube.com/watch?v=ghX5UDWD1HU

Video presenting our method : https://www.youtube.com/watch?v=boDaAORoKho

UNIVERSITÉ DE LYON Lyon 1
INSA INSTITUT NATIONAL DES SCIENCES APPLIQUÉES LYON
LIRIS
citi lab

Teaching Agents how to Map: Spatial Reasoning for Multi-Object Navigation


Pierre Marza


Laetitia Matignon


Olivier Simonin

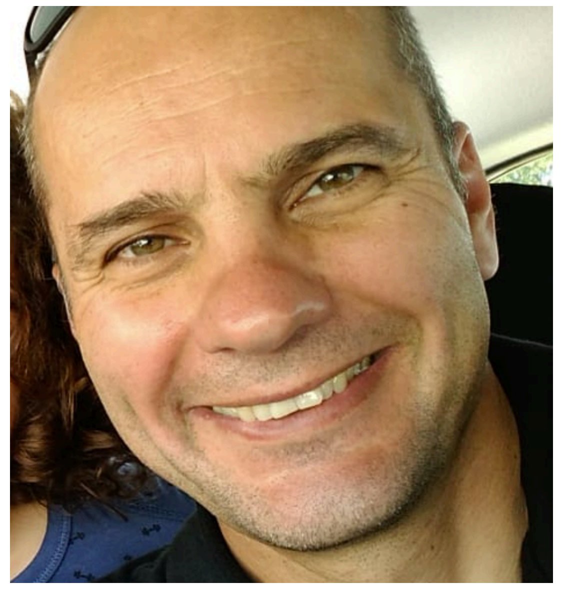
Christian Wolf

pierre.marza@insa-lyon.fr