Multi-Object Navigation with dynamically learned neural implicit representations

Pierre Marza[1], Laëtitia Matignon[2], Olivier Simonin[1], Christian Wolf[3]

INSA Lyon[1] UCBL[2] Naver Labs Europe[3]

ICCV23 PARIS

Project Page

Introduction

Problem

- Can we learn implicit representations of the structure and semantics of a 3D scene in real time, and train an agent with Reinforcement Learning (RL) to use them to solve a visual navigation task?
- How to efficiently query an implicit representation at navigation time to follow shorter paths to goals?

Considered task: Multi-Object Navigation (MultiON) [?]

- Requires an agent to navigate in a photo-realistic 3D environment from RGB-D observations and reach a sequence of target objects (colored cylinders) in a particular order.
- Two sub-tasks: exploration and semantic mapping.

Semantic Finder

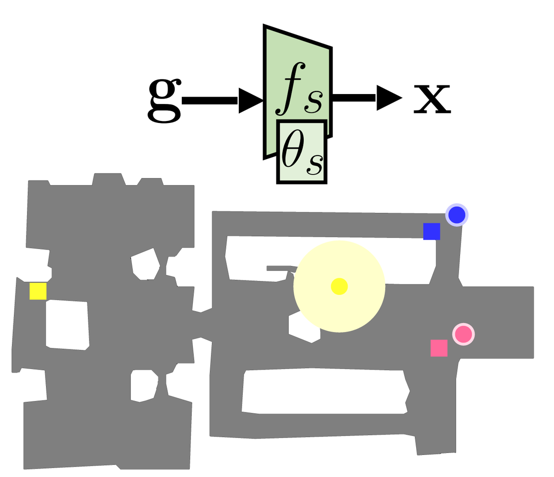

- Predicts the position of an object (colored circle) specified through an input query (colored squares show GT locations).
- Querying the location of an object is a single forward pass.
- Prediction uncertainty (shaded area) is also computed.

Occupancy and Exploration Implicit Representation

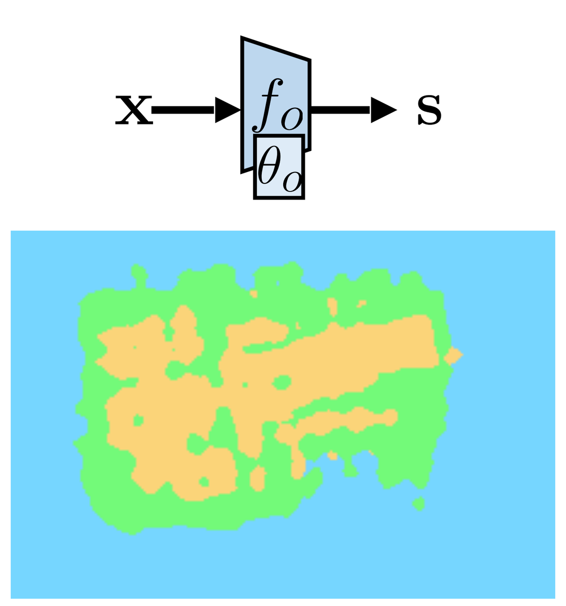

- Continuous representation of free navigable space and obstacles.
- Predicts occupancy **s** as a classification problem with three classes *{Obstacle (green), Navigable (orange), Unexplored (blue)}*

Global Reader

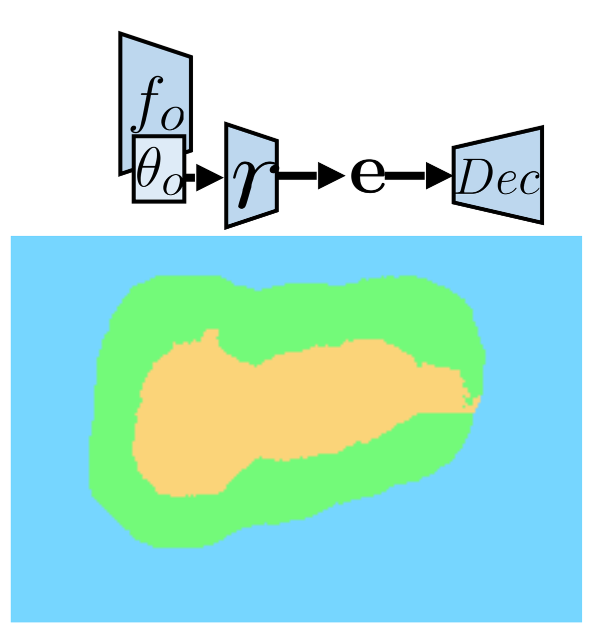

- f_o can be queried for a position, but reading out information over a large area this way requires multiple reads.
- Global reading mechanism summarizing the known state of the environment in one forward pass.
- r is a Tansformer model predicting an embedding **e** directly from the weights of f_o (θ_o).
- A decoder is only used to train r.

Navigating with implicit representations

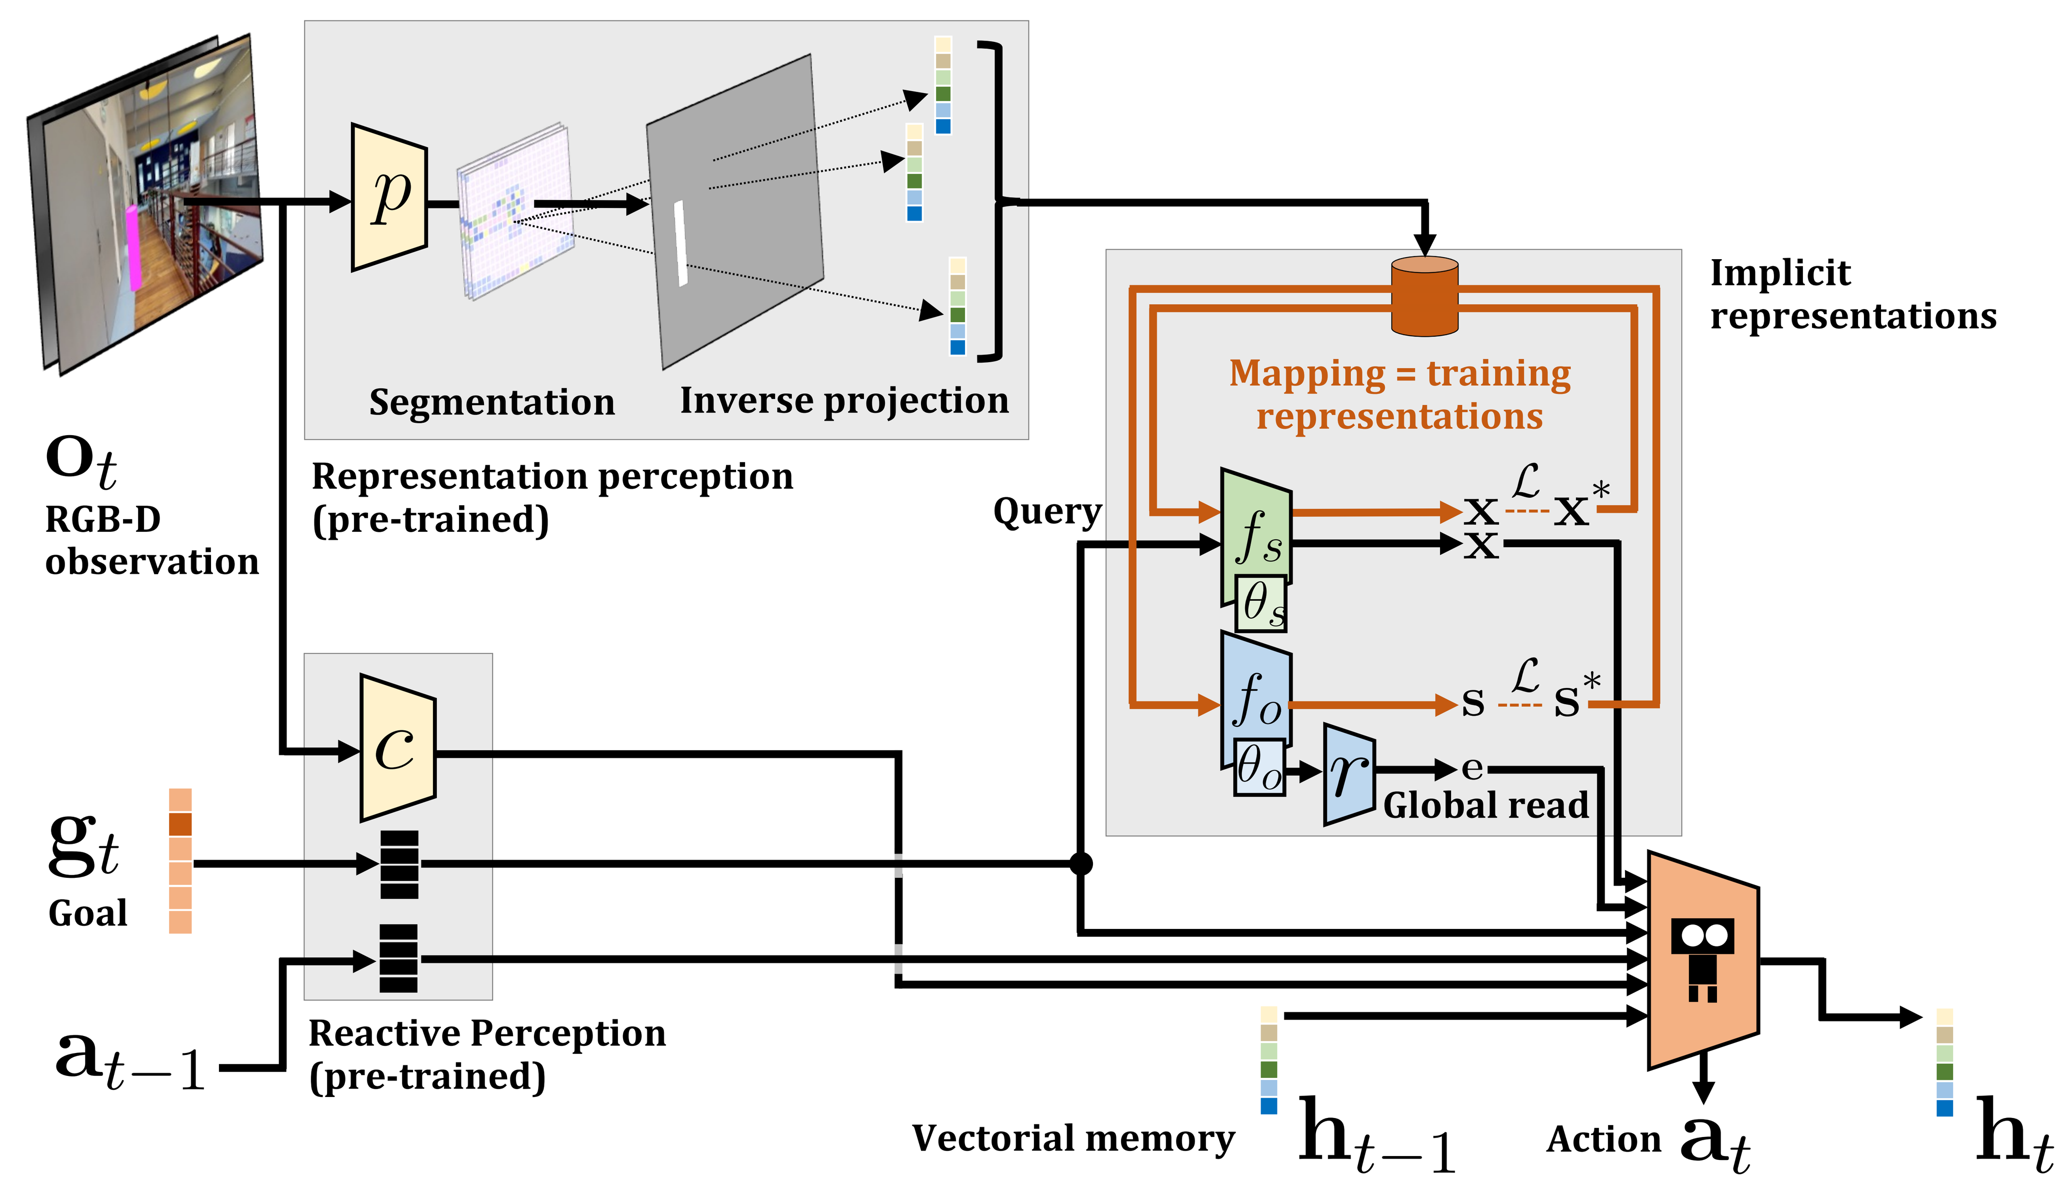

Mapping means training! f_s and f_o maintain a compact and actionable representation of the observed scene, and as such need to be updated at each time step from the current observation. Given their implicit nature and implementation as neural networks, updates are gradient based and done with SGD. The implicit representations are therefore *trained from scratch at each episode even after deployment*.

Agent training The agent is trained with RL (Proximal Policy Optimization [?]). The inner training loops of the implicit representations are supervised (red arrows) and occur at each time step in the forward pass, whereas the RL-based outer training loop of the agent occur after N acting steps (black arrows).

Quantitative results

Baselines

- *NoMap* [?]: Recurrent agent.
- *ProjNMap* [?]: Recurrent agent with a 2D top-down map of CNN-extracted features.
- *OracleMap/OracleEgoMap* [?]: *ProjNMap* with oracle map (goal locations, ✓ in ORC).
- *ProjNMap + AUX* [?]: Augmenting the training supervision of *ProjNMap* with mapping-related auxiliary losses (✓ in AUX).

Training details

- ρ: pre-training of input encoders from [?] (*w/* or *w/o pre-train*).
- α: finetuning of input encoders with RL.
- *w/ curriculum*: 3-steps agent RL training to reduce compute requirements: 0–$30M$ frames with no impl. repr., $30M$–$50M$ frames with f_s only and $50M$–$70M$ frames with f_s, f_o and r.
- γ: both implicit representations are accessible to the agent since the beginning of RL training (*w/o curriculum*).

Metrics

- *Success*: Percentage of successful episodes (all targets found in order).
- *Progress*: Percentage of found targets.
- *SPL*: Success and path efficiency.
- *PPL*: Progress and path efficiency.

	Agent	ρ	α	γ	**Success**	**Progress**	**SPL**	**PPL**	AUX	ORC
(a)	OracleMap† [?]	–	✓		50.4± 3.5	60.5± 3.1	40.7± 2.2	48.8± 1.9	–	✓
(b)	OracleEgoMap† [?]	–	✓		32.8± 5.2	47.7± 5.2	26.1± 4.5	37.6± 4.7	–	✓
(c)	NoMap† [?]	–	✓		16.7± 3.6	33.7± 3.3	13.1± 2.4	26.0± 1.7	–	–
(d)	ProjNMap† [?]	–	✓		25.9± 1.1	43.4± 1.0	18.3± 0.6	30.9± 0.7	–	–
(e)	NoMap	✓	–		42.3± 1.5	56.7± 0.9	28.1± 1.0	37.8± 1.8	–	–
(f)	ProjNMap [?]	✓	–		39.7± 2.3	55.4± 1.4	28.7± 1.1	40.1± 1.9	–	–
(g)	Implicit (Ours) *w/ curriculum w/ pre-train*	✓	–	–	46.7± 3.0	60.1± 3.1	35.1± 1.4	44.8± 1.0	–	–
(h)	ProjNMap + *AUX* [?]	N/A	✓	N/A	57.7 ± 3.7	70.2 ± 2.7	37.5 ± 2.0	45.9 ± 1.9	✓	–
(i)	Implicit (Ours) *w/o curriculum w/ pre-train + AUX*	✓	✓	✓	58.3± 0.8	69.4± 1.1	**43.8**± 1.0	**52.1**± 1.6	✓	–
(j)	Implicit (Ours) *w/o curriculum w/o pre-train*	–	✓	✓	54.8± 3.6	68.0± 3.4	41.7± 1.9	51.3± 1.6	–	–
(k)	Implicit (Ours) *w/o curriculum w/o pre-train + AUX*	–	✓	✓	57.9± 2.0	69.5± 0.6	**43.3**± 2.2	**51.9**± 3.7	✓	–

Training of the Global Reader

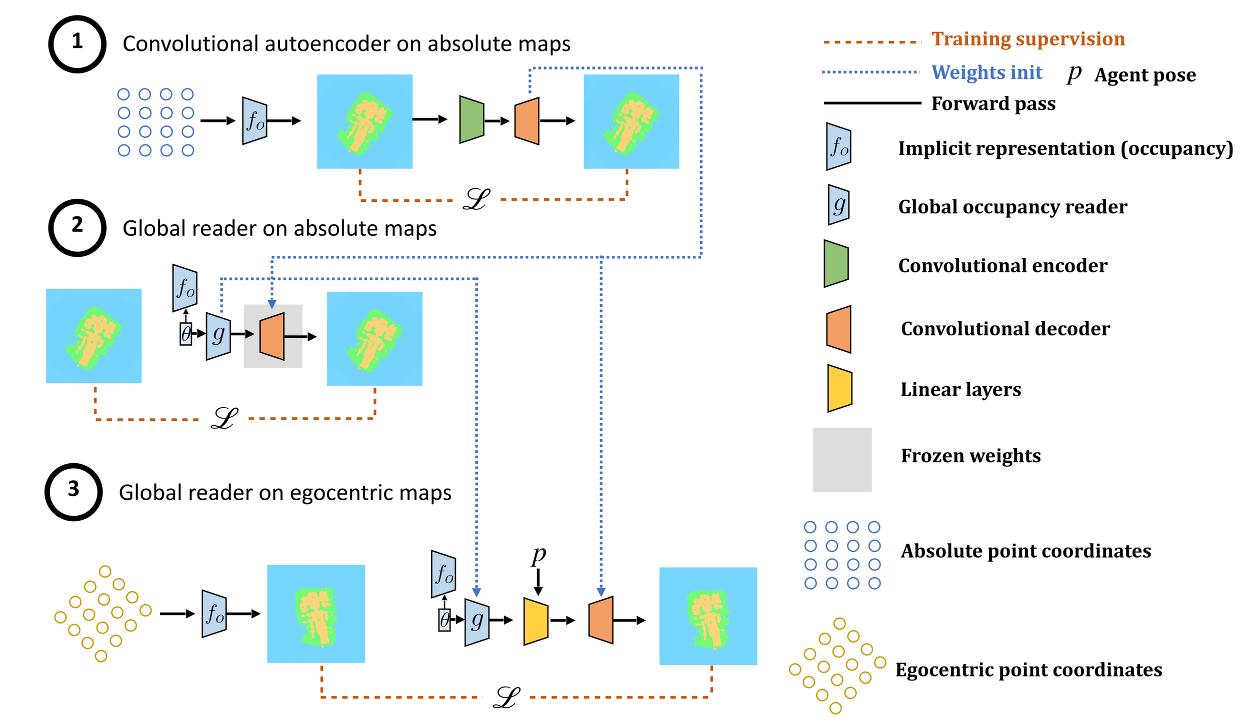

- Trained from a dataset of pairs: MLP weights θ_o parametrizing an implicit representation f_o and an egocentric map obtained by discretizing f_o.
- Trained to decode the egocentric map from the set of weights as input.
- 3-steps training works better than single-step training from scratch.
- After this pre-training, r is not adapted during the RL training.

Lifelong learning

- Evaluating the capacity of f_s to hold the learned information while being continuously trained.
- The mean error in distance for the predicted position of seen targets quickly decreases and stabilizes.

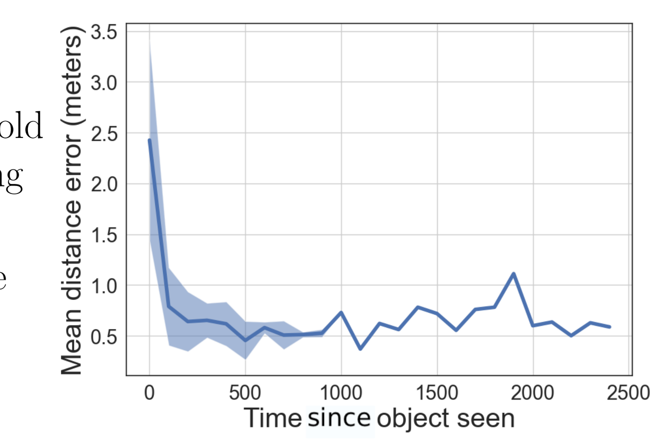

Capacity of the semantic representation

- Study performed independently of the *MultiON* benchmark: synthetic dataset to evaluate the capacity of f_s to store large numbers of objects.
- f_s trained from dummy queries: (a) one-hot queries with same dimension as nb. objects; (b) random query with dimension 9; (c) random query with same dimension as nb. objects. Mean distance prediction error is reported.

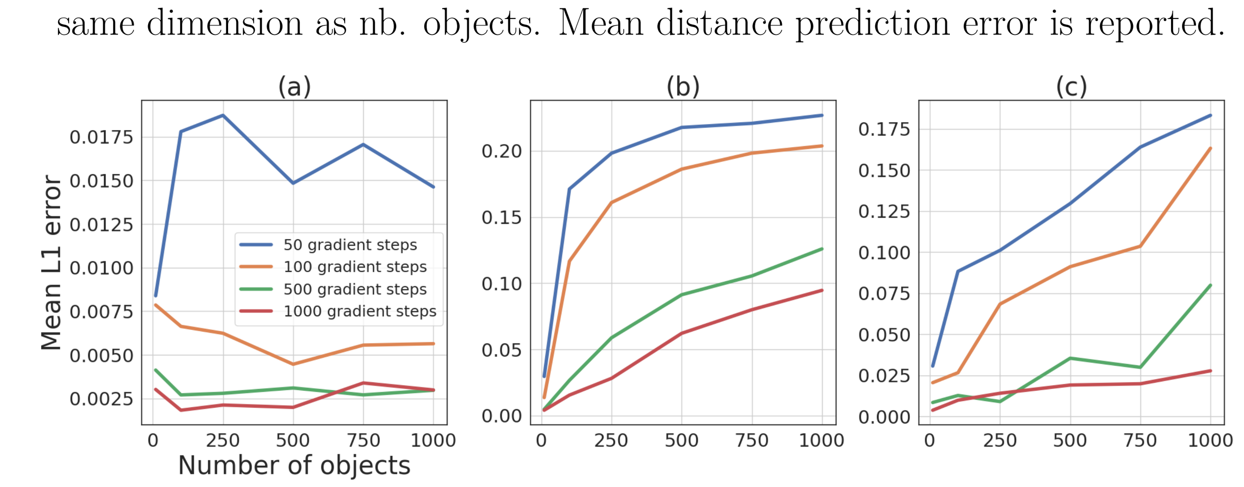

References

[1] João F. Henriques and Andrea Vedaldi. Mapnet: An allocentric spatial memory for mapping environments. In *CVPR*, 2018.

[2] Pierre Marza, Laetitia Matignon, Olivier Simonin, and Christian Wolf. Teaching agents how to map: Spatial reasoning for multi-object navigation. In *IROS*, 2022.

[3] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint*, 2017.

[4] Saim Wani, Shivansh Patel, Unnat Jain, Angel X. Chang, and Manolis Savva. Multion: Benchmarking semantic map memory using multi-object navigation. In *NeurIPS*, 2020.