Towards real-world navigation with deep differentiable planners

Shu Ishida, João F. Henriques {ishida, joao}@robots.ox.ac.uk Visual Geometry Group, University of Oxford

EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines & Systems

Engineering and Physical Sciences Research Council

Lattice PointNet (LPN)

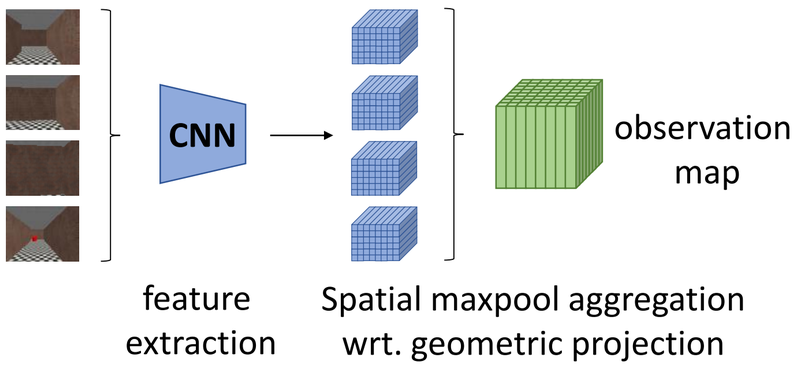

Collision Avoidance Long-term Value Iteration Network (CALVIN)

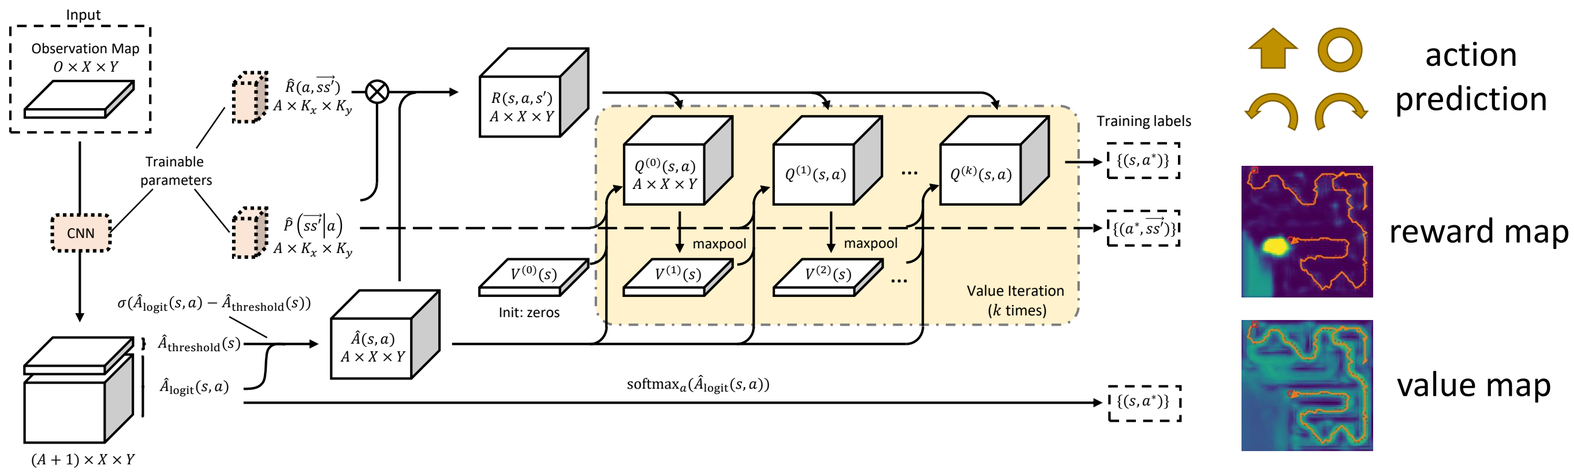

Problem

Given expert demonstrations (trajectories of poses, RGB-D images and chosen actions), learn a navigation strategy that finds the target while avoiding obstacles in novel environments. The nature of obstacles and targets are learnt rather than given.

Method

CALVIN improves upon VIN[1] in four ways:

- Decomposes transition model
- Penalises invalid transitions better
- Better sampling strategies
- Robust navigation in partially known environments

$$P\left(\overrightarrow{ss'}\middle|a\right) = \hat{P}\left(\overrightarrow{ss'}\middle|a\right) \times \hat{A}(s,a,o)$$

shared motion model

transition availability model (predicted from observation)

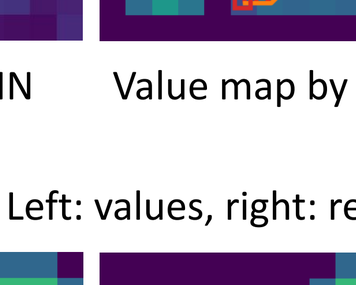

Grid Maze Environments

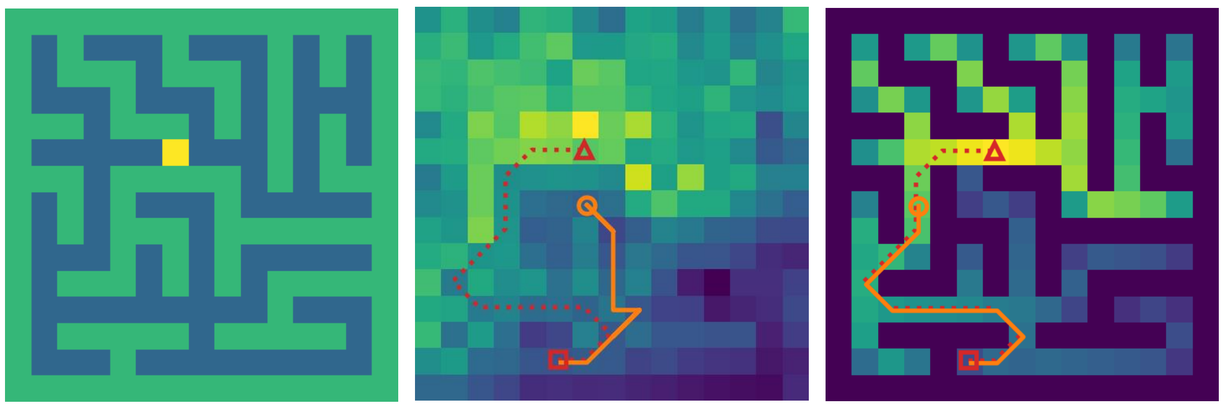

Maze | Value map by VIN | Value map by CALVIN

Embodied CALVIN

Left: values, right: rewards

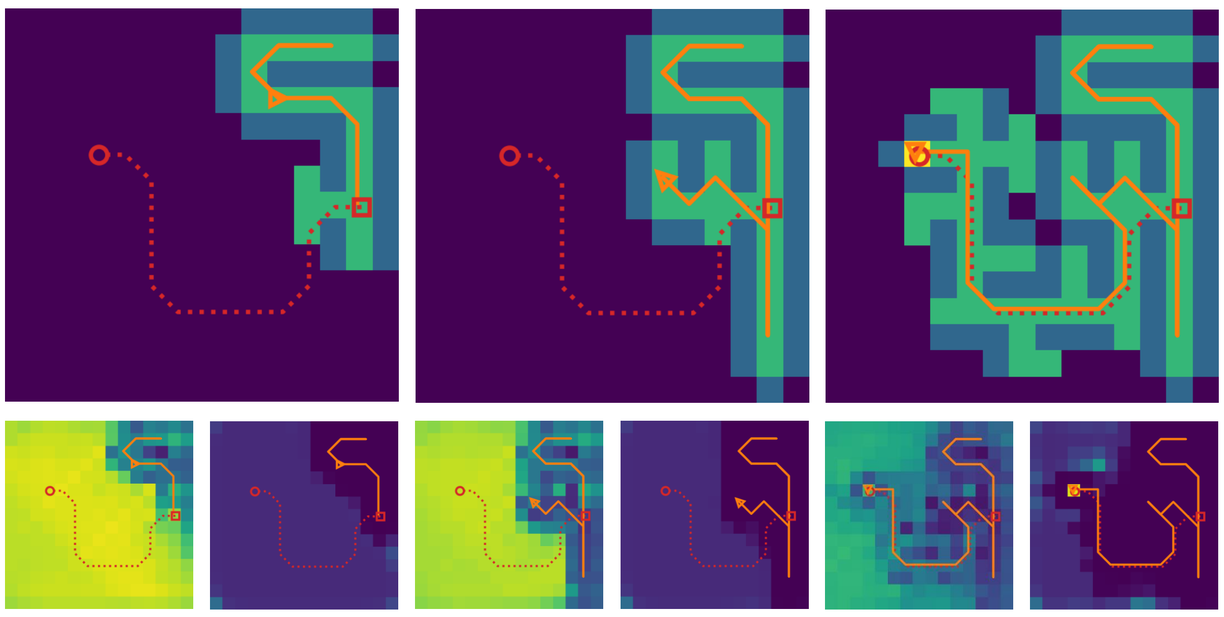

Miniworld[2] Environments

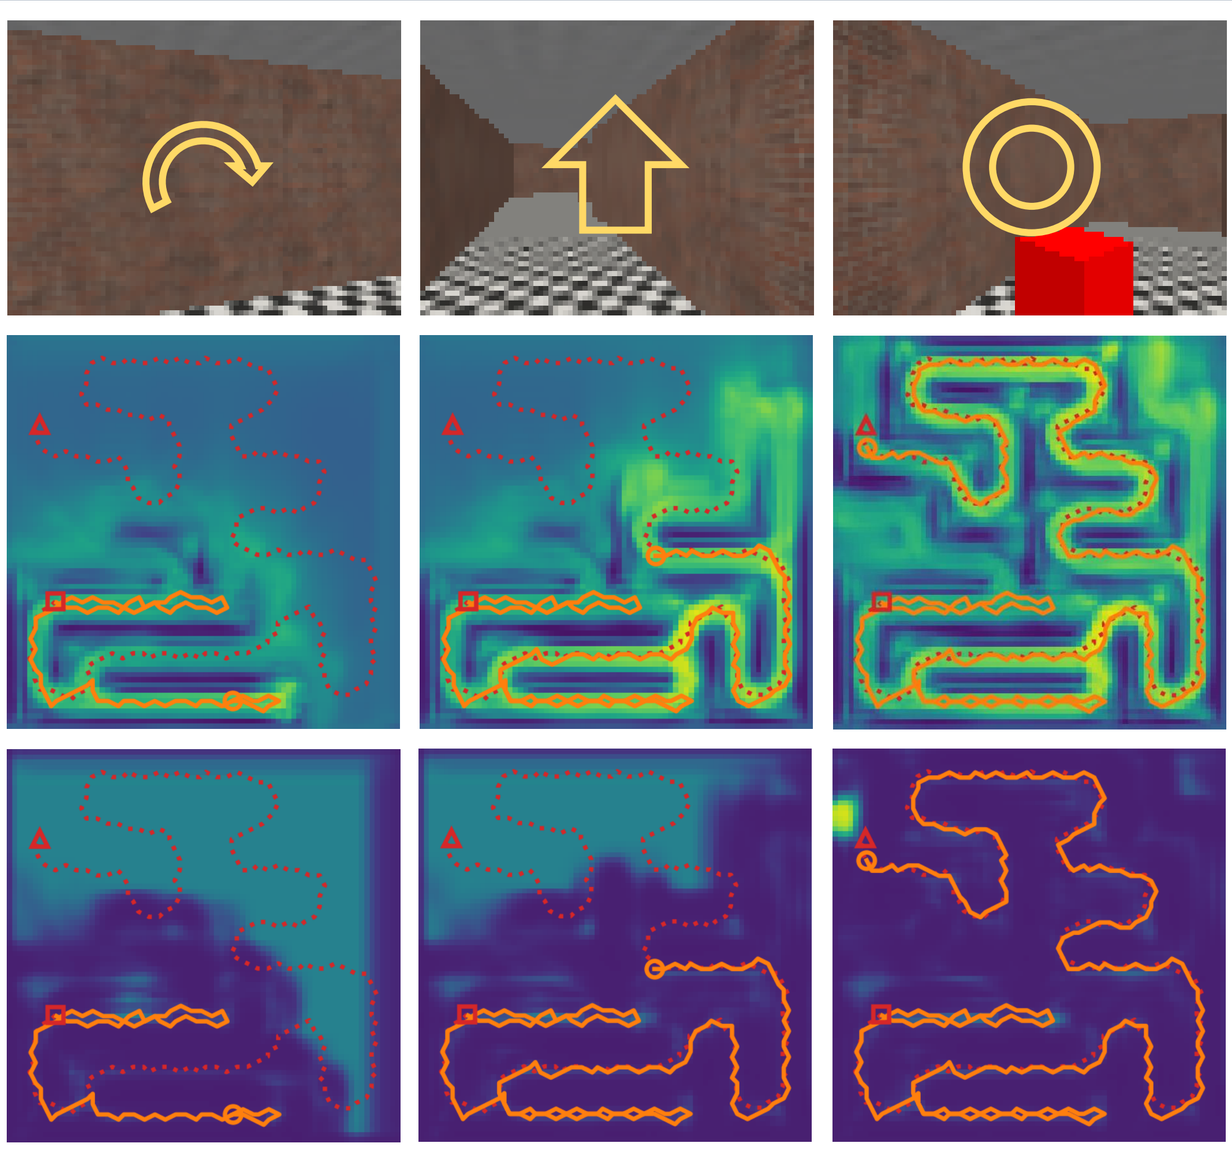

Learnt agent motion model

Active Vision Dataset[3]

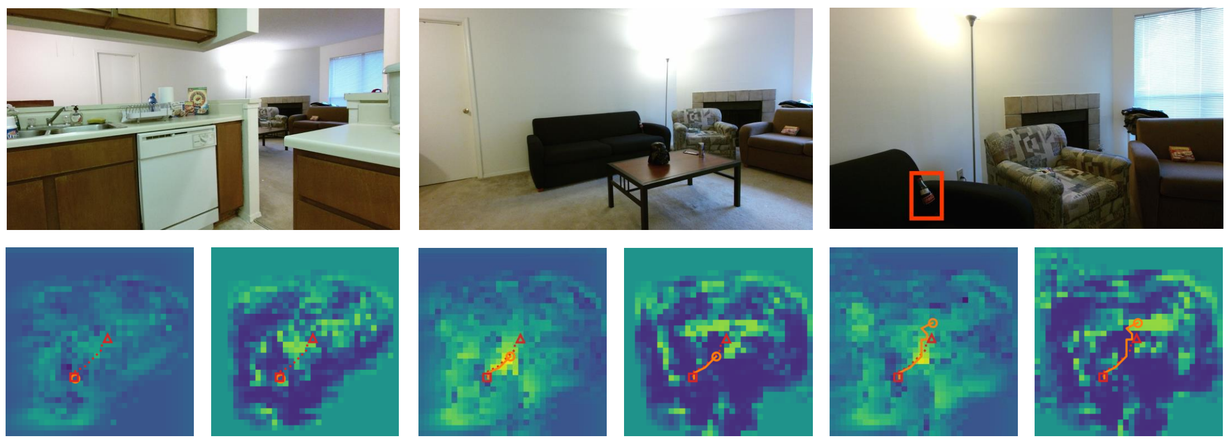

[1] Tamar et al., “Value Iteration Networks”, NeurIPS 2016.
[2] M. Chevalier-Boisvert, https://github.com/maximecb/gym-miniworld, 2018.
[3] Ammirato et al. “A datasetfor developing and benchmarking active vision”, ICRA 2017.