Learning to Navigate

Vladlen Koltun
Intel Labs

Vision, like other sensory functions, has its evolutionary rationale rooted in improved motor control. Although organisms can of course see when motionless or paralyzed, the visual system of the brain has the organization, computational profile, and architecture it has in order to facilitate the organism's thriving at the four Fs: feeding, fleeing, fighting, and reproduction.

- Churchland, Ramachandran, and Sejnowski

A critique of pure vision

Navigation via SLAM?

- SLAM addresses perception, not action
- SLAM does not tell us how to explore, navigate, accomplish goals
  - The planner will take care of it…
  - But planners assume that the map is accurate and sufficiently complete
  - When this assumption breaks, the whole system can fail

RealSense™ T260 Tracking Module

REALSENSE™ TRACKING MODULE T260

Intel® RealSense™ Tracking Module T260 Features:

Wide Field of View Stereo Fisheye Pair:
- 170 deg circular FOV on 800 pix (dia) monochrome sensor

MEMS IMU
- Bosch BMI055 @100-400Hz

Vision Processor
- Myriad 2

Complete 6DoF Solution
- Entire SLAM algorithm runs onboard; poses streamed out over USB

Low Power
- ~1 watt

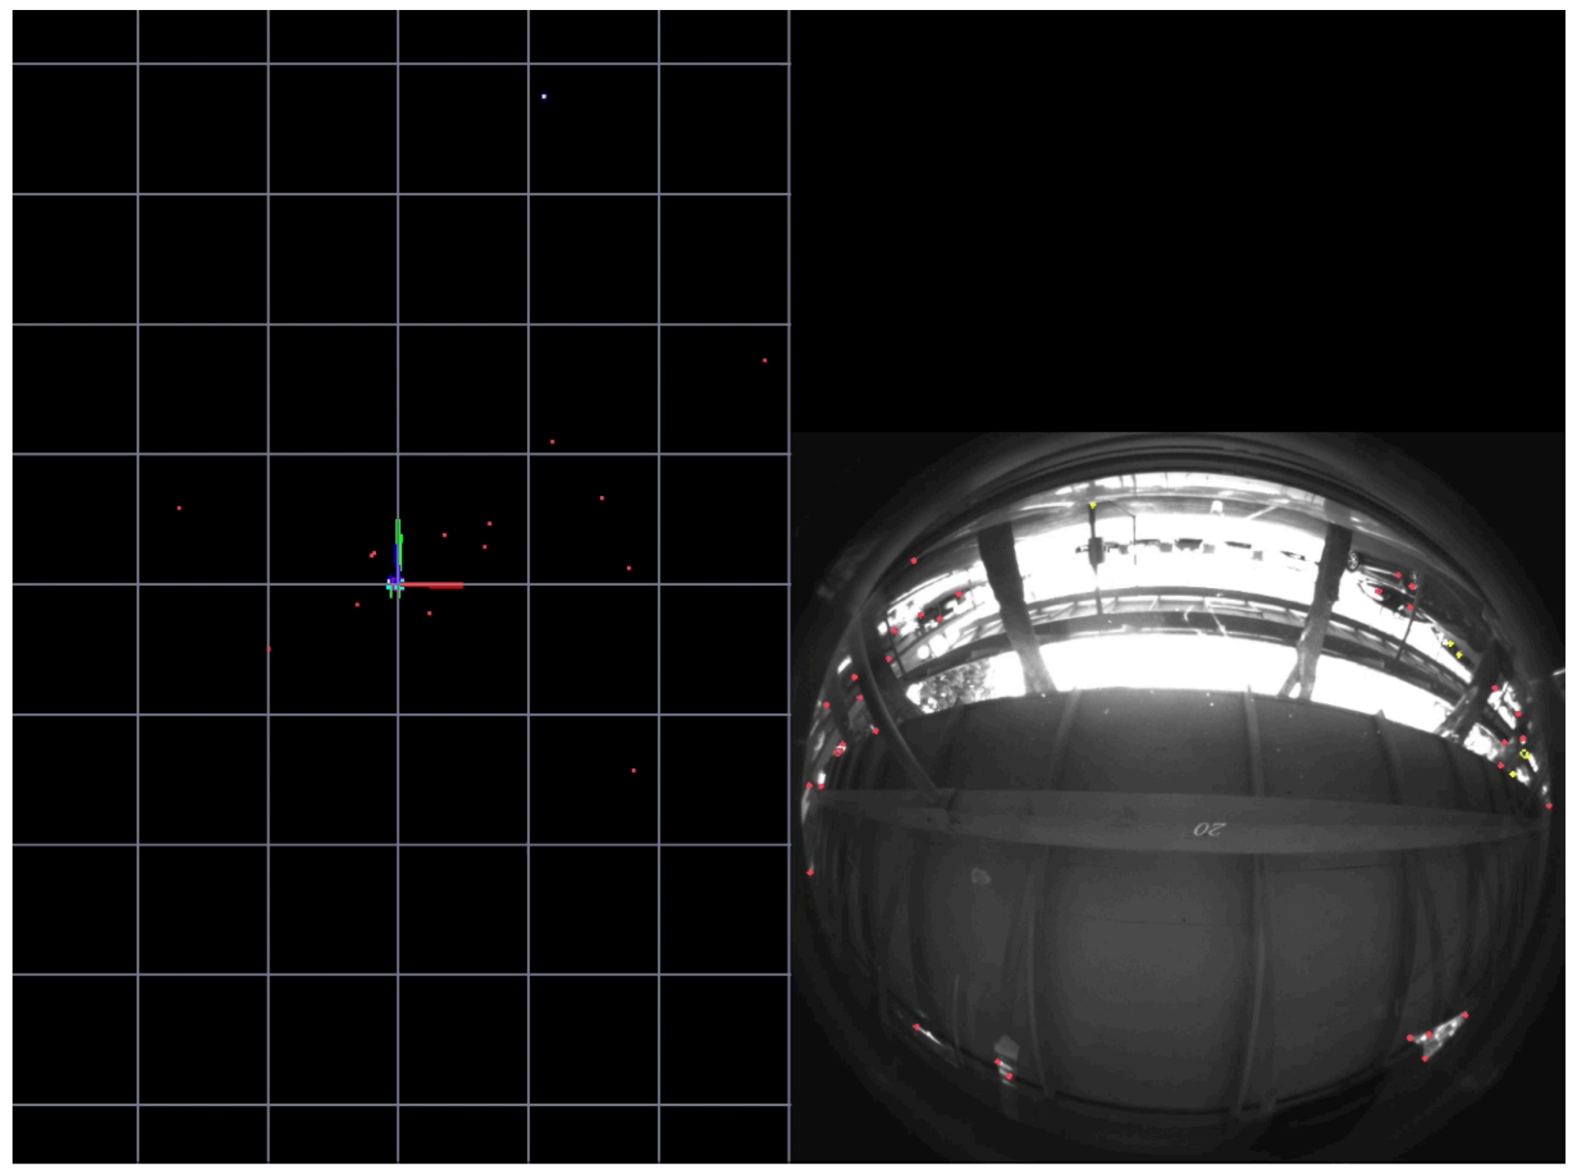

Slide content from Eagle Jones

Navigation via SLAM?

- High “dynamic range” (e.g., driving to work, flying to ECCV)
- Poor localization (e.g., perceptual aliasing, walking through a field)
- Environments that change over time
- Reflexive perception-action shortcuts (e.g., duck behind corner to avoid danger)
- Prior knowledge (navigate to a grocery store in a new city, navigate to a restroom in a restaurant)

Classic sensorimotor control pipeline

- Perception system constructs a map of the environment
- Planner generates waypoints
- Continuous controller (e.g., proportional-derivative) actuates the system towards the waypoints

Prize for place cells

Discoverers of brain's navigation system get physiology Nobel.

BY ALISON ABBOTT & EWEN CALLAWAY

Brain cells that make up the biological equivalent of a satellite-navigation system have garnered three scientists the 2014 Nobel Prize in Physiology or Medicine. The discovery of the cells sheds light on one of neuroscience's great mysteries — how we know where we are in space.

John O'Keefe of University College London won half of the prize for his discovery in 1971 of 'place' cells in the hippocampus, a part of the brain associated with memory. Edvard and May-Britt Moser, who are married and jointly run a lab at the Kavli Institute for Systems Neuroscience in Trondheim, Norway, share the other half for their 2005 discovery of 'grid' cells in an adjacent brain structure, the entorhinal cortex. Along with other navigation cells, grid and place cells allow animals to keep track of their position. Both cell types were discovered

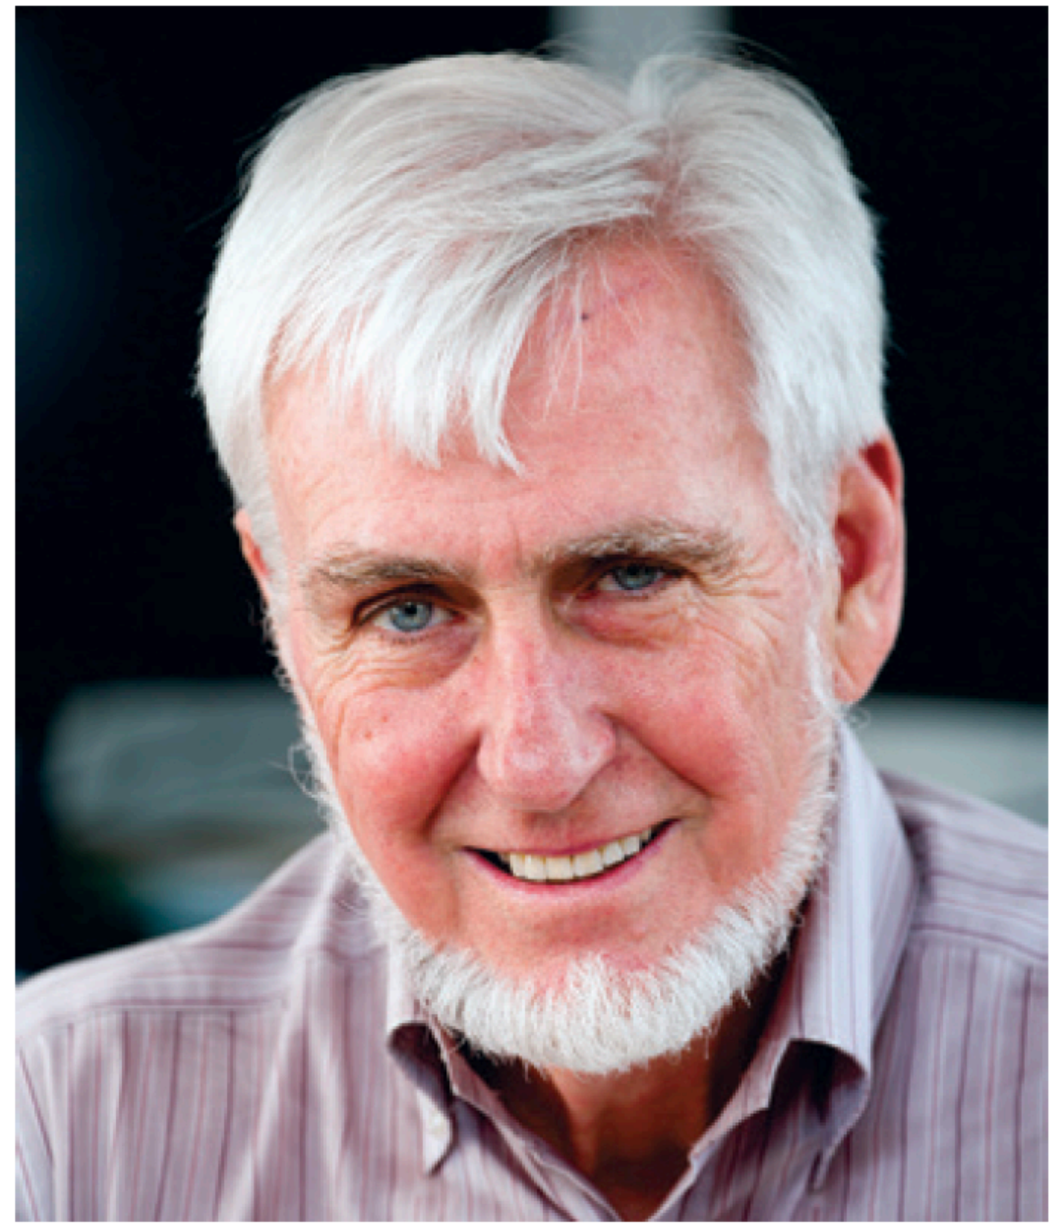

John O'Keefe, co-winner of a 2014 Nobel prize.

VOL. 55, NO. 4 JULY, 1948

THE PSYCHOLOGICAL REVIEW

COGNITIVE MAPS IN RATS AND MEN [1]

BY EDWARD C. TOLMAN

University of California

I shall devote the body of this paper to a description of experiments with rats. But I shall also attempt in a few words at the close to indicate the significance of these findings on rats for the clinical behavior of men. Most of the rat investigations, which I shall report, were carried out in the Berkeley laboratory. But I shall also include, occasionally, accounts of the behavior of non-Berkeley rats who obviously have misspent their lives in out-of-State laboratories. Furthermore, in reporting our Berkeley experiments I shall have to omit a very great many. The ones I *shall* talk about were carthe food box and eats. This is repeated (again in the typical experiment) one trial every 24 hours and the animal tends to make fewer and fewer errors (that is, blind-alley entrances) and to take less and less time between start and goal-box until finally he is entering no blinds at all and running in a very few seconds from start to goal. The results are usually presented in the form of average curves of blind-entrances, or of seconds from start to finish, for groups of rats.

All students agree as to the facts. They disagree, however, on theory and explanation.

DO ANIMALS HAVE COGNITIVE MAPS?

ANDREW T. D. BENNETT

Department of Pure Mathematics, University of Adelaide, Adelaide SA 5005, Australia, Department of Zoology South Parks Road, Oxford OX1 3PS, UK and **School of Biological Sciences, University of Bristol, Woodland Ro Bristol BS8 1UG, UK*

Summary

Drawing on studies of humans, rodents, birds and arthropods, I show that 'cognitive maps' have been used to describe a wide variety of spatial concepts. There are, however, two main definitions. One, *sensu* Tolman, O'Keefe and Nadel, is that a cognitive map is a powerful memory of landmarks which allows novel short-cutting to occur. The other, *sensu* Gallistel, is that a cognitive map is any representation of space held by an animal. Other definitions with quite different meanings are also summarised. I argue that no animal has been conclusively shown to have a cognitive map, *sensu* Tolman, O'Keefe and Nadel, because simpler explanations of the crucial novel short-cutting results are invariably possible. Owing to t repeated inability of experimenters to eliminate th simpler explanations over at least 15 years, and t confusion caused by the numerous contradicto definitions of a cognitive map, I argue that the cognit map is no longer a useful hypothesis for elucidating t spatial behaviour of animals and that use of the te should be avoided.

Key words: cognitive map, route, representation of space, spatial memory.

Do Humans Integrate Routes Into a Cognitive Map? Map- Versus Landmark-Based Navigation of Novel Shortcuts

Patrick Foo, William H. Warren, Andrew Duchon, and Michael J. Tarr
Brown University

Do humans integrate experience on specific routes into metric survey knowledge of the environment, or do they depend on a simpler strategy of landmark navigation? The authors tested this question using a novel shortcut paradigm during walking in a virtual environment. The authors find that participants could not take successful shortcuts in a desert world but could do so with dispersed landmarks in a forest. On catch trials, participants were drawn toward the displaced landmarks whether the landmarks were clustered near the target location or along the shortcut route. However, when landmarks appeared unreliable, participants fell back on coarse survey knowledge. Like honeybees (F. C. Dyer, 1991), humans do not appear to derive accurate cognitive maps from path integration to guide navigation but, instead, depend on landmarks when they are available.

Keywords: human navigation, virtual reality, path integration, visual landmarks

Do not ask whether they have a cognitive map, but how they find their way about

N.J. Mackintosh

University of Cambridge, UK

The publication of "The hippocampus as a cognitive map" (O'Keefe & Nadel, 1978) has had a remarkable impact, stimulating a huge amount of both behavioural and neurobiological research on spatial learning and memory, involving both laboratory and field studies, and employing a variety of novel techniques. The reviews of this general area of research provided by the previous contributors to this special issue attest to the progress that has been made since 1978. No one would now doubt that the hippocampus is implicated in the use of configurations of landmarks to locate a goal - in both mammals and birds, although the precise nature of that implication remains a matter of much speculation and debate. On balance, however, the behavioural evidence does not seem to have supported O'Keefe and Nadel's original hypothesis that true spatial or locale learning is quite distinct from simple associative learning and depends on the establishment of a cognitive map of the environment.

Human spatial representation: insights from animals

Ranxiao Frances Wang and Elizabeth S. Spelke

Human navigation is special: we use geographic maps to capture a world far beyond our unaided locomotion. In consequence, human navigation is widely thought to depend on internalized versions of these maps – enduring, geocentric 'cognitive maps' capturing diverse information about the environment. Contrary to this view, we argue that human navigation is best studied in relation to research on navigating animals as humble as ants. This research provides evidence that animals, including humans, navigate primarily by representations that are momentary rather than enduring, egocentric rather than geocentric, and limited in the environmental information that they capture. Uniquely human forms of navigation build on these representations.

SEMI-PARAMETRIC TOPOLOGICAL MEMORY FOR NAVIGATION

Nikolay Savinov*
ETH Zürich

Alexey Dosovitskiy*
Intel Labs

Vladlen Koltun
Intel Labs

ABSTRACT

We introduce a new memory architecture for navigation in previously unseen environments, inspired by landmark-based navigation in animals. The proposed semi-parametric topological memory (SPTM) consists of a (non-parametric) graph with nodes corresponding to locations in the environment and a (parametric) deep network capable of retrieving nodes from the graph based on observations. The graph stores no metric information, only connectivity of locations corresponding to the nodes. We use SPTM as a planning module in a navigation system. Given only 5 minutes of footage of a previously unseen maze, an SPTM-based navigation agent can build a topological map of the environment and use it to confidently navigate towards goals. The average success rate of the SPTM agent in goal-directed navigation across test environments is higher than the best-performing baseline by a factor of three.

Task: memory-based navigation

Walkthrough sequence (~5 min)

Current observation

Goal observation

Savinov et al., ICLR 2018

Topological memory for navigation

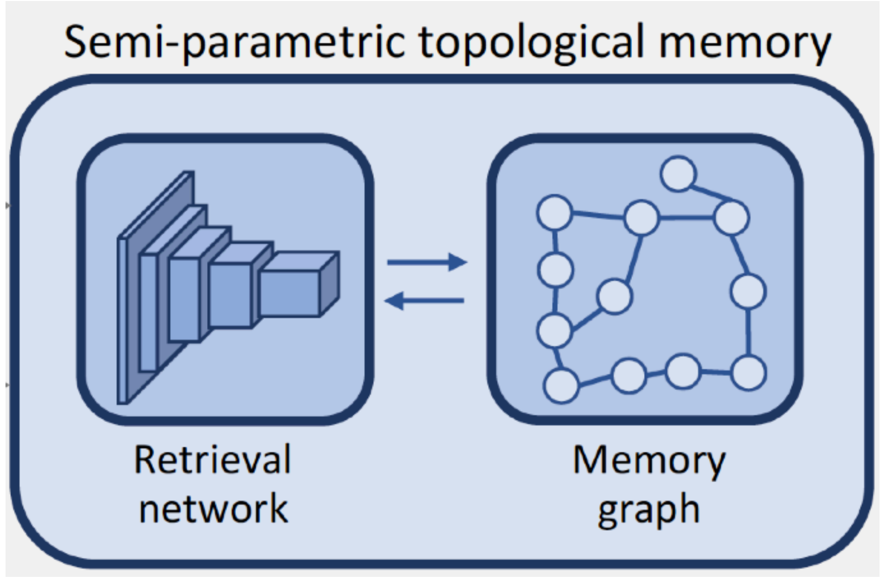

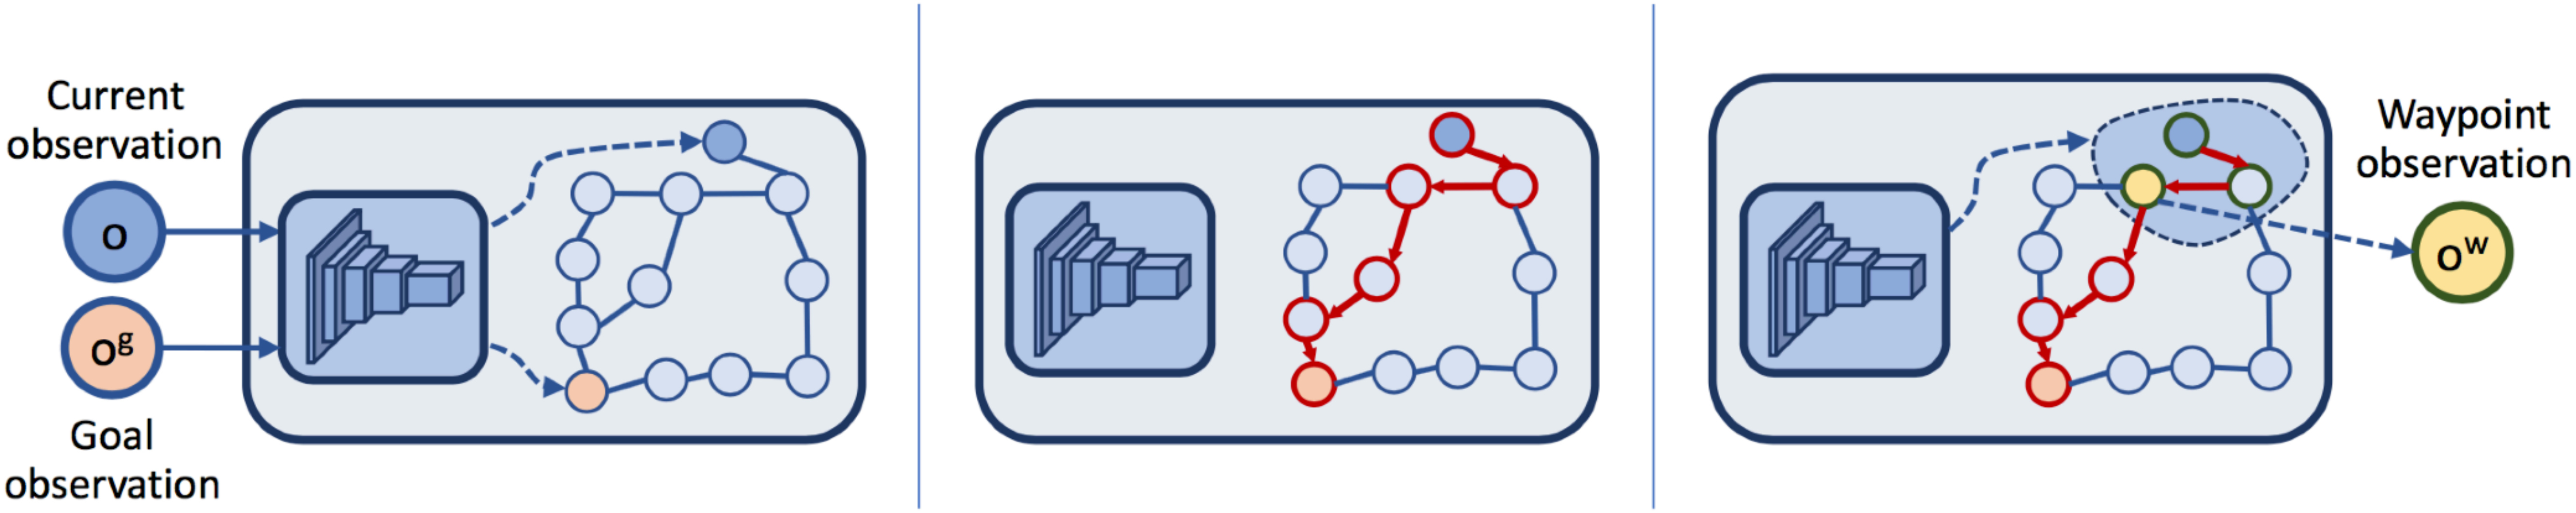

(a) Localization (b) Planning (c) Waypoint selection

Savinov et al., ICLR 2018

After being shown a 5-minute traversal video of
a previously unseen maze...

Savinov et al., ICLR 2018

On Evaluation of Embodied Navigation Agents

Peter Anderson
Australian National University

Angel Chang
Princeton University

Devendra Singh Chaplot
Carnegie Mellon University

Alexey Dosovitskiy
Intel Labs

Saurabh Gupta
UC Berkeley

Vladlen Koltun
Intel Labs

Jana Kosecka
George Mason University

Jitendra Malik
UC Berkeley and Facebook

Roozbeh Mottaghi
Allen Institute for AI

Manolis Savva
Princeton University

Amir R. Zamir
Stanford and UC Berkeley

Abstract

Skillful mobile operation in three-dimensional environments is a primary topic of study in Artificial Intelligence. The past two years have seen a surge of creative work on navigation. This creative output has produced a plethora of sometimes incompatible task definitions and evaluation protocols. To coordinate ongoing and future research in this area, we have convened a working group to study empirical methodology in navigation research. The present document summarizes the consensus recommendations of this working group. We discuss different problem statements and the role of generalization, present evaluation measures, and provide standard scenarios that can be used for benchmarking.

multaneous Localization and Mapping (SLAM) [6], which focuses on building a map of the environment and localizing the agent within the map. Traditional SLAM systems largely focus on constructing metric maps using geometric techniques. Navigation itself (i.e., what to do, where to go) is rarely considered. Representations constructed by SLAM systems are often not suitable for traditional motion planning methods and are prone to error when the environment changes over time.

In contrast, biological systems appear to use more flexible representations, can robustly navigate without precise localization or a metric map, and can bring substantial prior knowledge to bear when planning in previously unseen environments.

arXiv:1807.06757

Summary

- What matters is action
- SLAM is solving a problem that is both unnecessarily hard and incomplete
- Biological motivation suggests alternative approaches
- Semi-parametric topological memory (SPTM)
- Working group on empirical methodology in navigation research

Thank you