Biologically Inspired Cognitive Architectures (2016) 16, 87–104

Available at www.sciencedirect.com

ScienceDirect

journal homepage: www.elsevier.com/locate/bica

RESEARCH ARTICLE

Towards real-world capable spatial memory in the LIDA cognitive architecture

Tamas Madl [a,b,*], Stan Franklin [c], Ke Chen [a], Daniela Montaldi [d], Robert Trappl [b]

[a] *School of Computer Science, University of Manchester, Manchester M13 9PL, UK*
[b] *Austrian Research Institute for Artificial Intelligence, Vienna A-1010, Austria*
[c] *Institute for Intelligent Systems, University of Memphis, Memphis, TN 38152, USA*
[d] *School of Psychological Sciences, University of Manchester, Manchester M13 9PL, UK*

Received 30 September 2015; received in revised form 11 February 2016; accepted 15 February 2016

KEYWORDS
Spatial memory;
LIDA;
Cognitive architecture;
Computational cognitive modeling

Abstract
The ability to represent and utilize spatial information relevant to their goals is vital for intelligent agents. Doing so in the real world presents significant challenges, which have so far mostly been addressed by robotics approaches neglecting cognitive plausibility; whereas existing cognitive models mostly implement spatial abilities in simplistic environments, neglecting uncertainty and complexity.
Here, we take a step towards computational software agents capable of forming spatial memories in realistic environments, based on the biologically inspired LIDA cognitive architecture. We identify and address challenges faced by agents operating with noisy sensors and actuators in a complex physical world, including near-optimal integration of spatial cues from different modalities for localization and mapping, correcting cognitive maps when revisiting locations, the structuring of complex maps for computational efficiency, and multi-goal route planning on hierarchical cognitive maps. We also describe computational mechanisms addressing these challenges based on LIDA, and demonstrate their functionality by replicating several psychological experiments.

* Corresponding author at: Austrian Research Institute for Artificial Intelligence, Vienna A-1010, Austria.
E-mail address: tamas.madl@gmail.com (T. Madl).

http://dx.doi.org/10.1016/j.bica.2016.02.001
2212-683X/

Introduction

Spatial representations are important for biological and artificial agents, to help them localize and navigate to important objects and places (such as food sources or shelters). Current computer models for learning spatial representations either neglect cognitive plausibility in favour of performance, such as simultaneous localization and mapping (SLAM) in robotics, or are incapable of running in large-scale, complex, uncertain environments perceived through noisy sensors.

Since biological cognition has been shaped by the structure, constraints, and challenges of the physical world, we argue that cognitive architectures should take these into account, as well. This argument is in accordance with the roadmap for the BICA Challenge, which also places importance on real-life capability (Samsonovich, 2012).

This paper describes an effort to take the LIDA (Learning Intelligent Distribution Agent) cognitive architecture (Franklin, Madl, D'Mello, & Snaider, 2014) closer to this goal. We hypothesize and implement approaches to tackle the sensory noise, uncertainty, and complexity of realistic environments. We also introduce a novel, conceptual and partially implemented, hierarchical spatial memory model, inspired by the neural basis of spatial cognition in brains, and provide a preliminary interface to realistic environments via the Robot Operating System (ROS) (Quigley et al., 2009). We demonstrate these extensions to LIDA in three-dimensional simulated environments that include simulated physics and high-quality graphics, based on the Player/Stage/Gazebo simulator.[1] This simulator presents the same interface to the agent as real devices, and an agent able to control a robot in Gazebo is also able to control the same robot in similar environments in the real world, without any changes to the control code (Rusu, Maldonado, Beetz, & Gerkey, 2007).

We build on and integrate our previous work investigating biologically and cognitively plausible implementations of Bayesian localization (Madl, Franklin, Chen, Montaldi, & Trappl, 2014), Bayesian nonparametric clustering for map structuring (Madl, Franklin, Chen, Trappl, & Montaldi, submitted for publication), and route planning based on activation gradients[2] (Madl, Franklin, Chen, & Trappl, 2013). The method for cognitive map correction (loop closing) is presented for the first time below. Although based on established mathematical tools from robotics, it is — to our knowledge — the first mechanism for large-scale cognitive map correction implementable in brains, and consistent with the replay phenomena observed in the rodent hippocampus (Carr, Jadhav, & Frank, 2011).

The present work is also (to our knowledge) the first to provide implementations of these mechanisms in a both cognitively and biologically plausible fashion (fitting behaviour data and implementable in brains), and integrated within the same cognitive architecture. Further contributions include concrete implementations of some features listed in the BICA Table (Samsonovich, 2010) which until now were only part of conceptual LIDA, including basic stereo colour vision, a cognitive map, spatial learning, and fusing information from multiple types of sensors and modalities via Bayesian update.

Related work

Apart from the complex perception problem, the most challenging problems for building spatial representations in realistic environments include localization and mapping under sensory noise, and correcting incorrect representations when revisiting known locations (loop closing). The robotics community has developed several solutions to these problems — see Bailey and Durrant-Whyte (2006), Durrant-Whyte and Bailey (2006), Thrun and Leonard (2008), and Williams et al. (2009). They have been designed to be accurate, not cognitively or biologically plausible, and rely on mechanisms that are difficult to implement in brains (e.g., many iterations performing operations on large matrices).

An exception is the partially connectionist RatSLAM system (Milford, Wyeth, & Rasser, 2004), which can learn robust maps in outdoor environments (Prasser, Milford, & Wyeth, 2006), and close large loops successfully if extended by a sophisticated data association method (Glover, Maddern, Milford, & Wyeth, 2010). Parts of it have been argued to be biologically plausible (Milford, Wiles, & Wyeth, 2010). However, RatSLAM has two disadvantages in the context of a cognitive model with long-term learning aiming for plausibility: (1) route planning only works along established routes (novel detours or shortcuts have not been demonstrated), (2) learned spatial information is mapped to a finite structure (attractor network) of fixed size which cannot be expanded.

On the other hand, models that emphasize plausibility — cognitive architectures and plausible spatial memory models — mostly focus on simplistic simulated environments, usually with no sensory noise and limited size/complexity. There are a few neurally inspired spatial memory models that can deal with a limited amount of uncertainty and noise (Barrera, Cáceres, Weitzenfeld, & Ramirez-Amaya, 2011; Burgess et al., 2000; Strösslin, Sheynikhovich, Chavarriaga, & Gerstner, 2005), but have only been tested in small indoor environments. See Madl, Chen, Montaldi, and Trappl (2015) for a review.

Spatial memory in brains

Spatial memory encodes, stores and recalls spatial information about the environment and the self-location of agents (biological or artificial), which they need to keep track of to navigate successfully. In most mammals, keeping track of position is achieved by path integration, which refers to updating the agent's position based on a fixed point and the estimated movement trajectory (based on information from proprioceptive and vestibular systems as well as sensory flow (Fortin, 2008; Mittelstaedt & Mittelstaedt, 1980)). It is a noisy process that accumulates large errors if uncorrected (Etienne, Maurer, & Sguinot, 1996).

[1] http://www.gazebosim.org/.

[2] Route planning in navigation space based on activation gradients has been proposed before (Burgess, Jackson, Hartley, & O'keefe, 2000; Schölkopf & Mallot, 1995), but not on a hierarchy — as it is in this work — which significantly improves its performance on multigoal problems.