

Imagery in cognitive architecture: Representation and control at multiple levels of abstraction

Action editor: Ron Sun

Samuel Wintermute^{*,1}

Department of Electrical Engineering and Computer Science, University of Michigan, 2260 Hayward Street, Ann Arbor, MI 48109, USA

Received 22 October 2011; accepted 12 December 2011

Available online 16 February 2012

Abstract

In a cognitive architecture, intelligent behavior is contingent upon the use of an appropriate abstract representation of the task. When designing a general-purpose cognitive architecture, two basic challenges related to abstraction arise, which are introduced and examined in this article. The perceptual abstraction problem results from the difficulty of creating a single perception system able to induce appropriate abstract representations in any task the agent might encounter, and the irreducibility problem arises because some tasks are resistant to being abstracted at all. The first contribution of this paper is identifying these problems, and the second contribution is showing a means to address them. This is accomplished through the use of mental imagery.

To support imagery, a concrete (highly detailed) representation of the spatial state of the problem is maintained as an intermediate between the external world and an abstract representation. Actions can be simulated (imagined) in terms of this concrete representation, and the agent can derive abstract information by applying perceptual processes to the resulting concrete state. Imagery works to mitigate the perceptual abstraction problem by allowing a given perception system to work in a wider variety of tasks, since perception can be dynamically combined with imagery, and works to mitigate the irreducibility problem by allowing internal simulation of low-level control processes.

To demonstrate these benefits, an implementation is described, which is an extension of the Soar architecture. An agent in this architecture that uses reinforcement learning and imagery to play an arcade game and an agent that performs sampling-based motion planning for a car-like vehicle are described, demonstrating the perceptual abstraction and irreducibility problems and the associated use of imagery to mitigate those problems in complex AI tasks.

Previous AI systems have incorporated imagery-like processes, however, the functional benefit of imagery in those systems has typically been characterized as the ability to perform more efficient inference through the use of a specialized representation. The use of imagery here shows further benefits related to the perceptual abstraction and irreducibility problems, enriching the broader understanding of the role of imagery in cognitive systems.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

People are confronted with a range of situations in their everyday lives that are characterized by a need for precise interaction with the spatial aspects of their surroundings. As a few extreme examples, consider catching a ball, solving

a jigsaw puzzle, or parallel parking. To catch a ball, a person must position their hand in a place where the ball will arrive; whether or not a given position meets this criterion depends upon the exact velocity of the ball and the influence of gravity. To solve a puzzle, a person must find which pieces fit together, which is a property that depends on the precise details of the shapes of both pieces. And to parallel park a car, the complex relationship between the controls of the car and its position on the street determines whether or not a given action sequence will result in successful parking.

* Tel.: +1 734 887 7624.

E-mail address: sam.wintermute@soartech.com

¹ Present address: Soar Technology, Inc., 3600 Green Court, Suite 600, Ann Arbor, MI 48105, USA

To create AI systems able to achieve this range of behavior, or cognitive models capable of accounting for it in humans, we pursue the development of a cognitive architecture. A cognitive architecture is a theory of the fixed processes underlying intelligent behavior (Langley, Laird, & Rogers, 2009). While it may be possible to create special-purpose models for each particular task, working in the context of an architecture causes a strict decomposition of the system into task-independent aspects (architecture) and task-dependent aspects (knowledge). Adherence to this decomposition allows the resulting system to be more easily extended to new tasks, as the architecture can be re-used, and leads toward the development of a unified theory of cognition, rather than many fragmented theories (Newell, 1990).

Here, we explore architectural structures supporting intelligence in complex spatially-oriented tasks. An important dimension through which these issues can be viewed is that of *abstraction*. In complex spatial tasks, an agent can benefit from using an abstract internal representation of the structure of the task. Abstraction removes unnecessary detail, making learning and knowledge representation more tractable. However, the need for abstraction seems to conflict with the need to account for the precise detail of the spatial world. Since abstraction is a process supported by the perception and action systems of a cognitive architecture, it is a critical issue to consider when designing one.

The basic conflict between abstraction and precision is analyzed in this article. As will be explained in the next section, there are two problems inherent in designing a general-purpose cognitive architecture capable of abstract representation while still maintaining the spatial precision necessary for intelligent behavior. First, the diversity of tasks an intelligent agent must address is large, and it is difficult to create a single perception system to create appropriate abstract representations in all such tasks. This difficulty is the *perceptual abstraction* problem. Second, some tasks are resistant to being abstracted at all, as is the case when the appropriate action varies continuously as a function of the details of the environment: this is the *irreducibility* problem.

Aspects of these challenges have manifested (and have been addressed) in research in several subfields of AI, including robotic motion planning, qualitative reasoning, and reinforcement learning. However, the root perceptual abstraction and irreducibility problems have not previously been identified and studied in a general manner. Problems of abstraction have received little emphasis in AI research compared to other perceptual challenges such as dealing with sensor noise and partial observability. The development of a task-independent cognitive architecture presents a context where abstraction challenges come to the forefront, though.

In this article, the perceptual abstraction and irreducibility problems are introduced, and a theory of basic architectural mechanisms that can work to mitigate these problems

is proposed. The crucial aspects of the theory include the use of both abstract and concrete (highly detailed) representations of the state of world, continuous action controllers which access the concrete representation, and imagery capability, where the concrete representation is internally manipulated, with the results feeding back to the abstract representation. In our terminology, “imagery” refers exclusively to this sort of manipulation of concrete information, rather than to the generic process of internally representing (imagining) alternate states of the world, which may be involve an abstract representation alone. The theory has been implemented by augmenting the Soar cognitive architecture (Laird, 2008; Newell, 1990) with general-purpose memories and processes for handling spatial information. Agents instantiated in the architecture provide demonstrations of both the operation of the architecture itself and the benefits of the underlying theory.

The theory is inspired by psychological research in mental imagery. This research has provided strong evidence that people maintain and manipulate visual and spatial information at a level close to that of perception, reusing the same systems that process perceptual data to process internally generated (imaginary) data (Kosslyn, Thompson, & Ganis, 2006). This work in this article builds on existing work on computational imagery systems, particularly that of Lathrop (2008), who created a pilot implementation of a mental imagery extension for Soar, but also drawing on other theories and systems (e.g., Barsalou, 1999; Glasgow & Papadias, 1992; Grush, 2004; Huffman & Laird, 1992; Kurup & Chandrasekaran, 2006; Ullman, 1984).

Imagery capability has been proposed as an important cognitive architectural component (e.g., Chandrasekaran, 2006; Lathrop, 2008). Much of the motivation for this inclusion has been drawn from psychological research. Arguments about the utility of imagery outside of psychological concerns—functional arguments for imagery—have also been made (e.g., Lathrop, Wintermute, & Laird, 2011). Typically these arguments are based on research examining benefits in terms of increased inference efficiency afforded by imagery-like processing (Glasgow & Papadias, 1992; Huffman & Laird, 1992; Larkin & Simon, 1987), or through demonstrations of particular domains where imagery use is beneficial, such as planning to coordinate a team of military scout robots (Lathrop & Laird, 2009). The use of imagery presented here, as a means of mitigating the perceptual abstraction and irreducibility problems, shows different functional benefits of imagery beyond inference efficiency. In that way, this work contributes towards a more thorough understanding of the role of imagery in cognitive architecture.

To elaborate on this, prior functionality-based examinations of imagery have assumed that, since abstract propositional representations and concrete perceptual representations can in principle encode the same information (Anderson, 1978), the primary functional role for imagery is to allow more efficient inference. However, the

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات