

Preface

Cognitive architectures form an integral part of robots and agents. Architectures structure and organize the knowledge and algorithms used by the agents to select actions in dynamic environments, plan and solve problems, learn, and coordinate with others. Architectures enable intelligent behavior by agents, and serve to integrate general capabilities expected of an intelligent agent (for example, planning and learning), to implement and test theories about natural or synthetic agent cognition, and to explore domain-independent mechanisms for intelligence.

As AI research has improved in formal and empirical rigor, traditional evaluation methodologies for architectures have sometimes proved insufficient. On the formal side, rigorous analysis has often proved elusive; we seem to be missing the notation required for formally proving properties of architectures. On the empirical side, experiments which demonstrate generality are notoriously expensive to perform, and are not sufficiently informative. And at a high-level, evaluation is difficult because the criteria are not well defined: Is it generality? Ease of programmability? Compatibility with data from biology and psychology? Applicability in real systems?

Recognizing that scientific progress depends on the ability to conduct informative evaluation (by experiment or formal analysis), this workshop addresses the methodologies needed for evaluating architectures. The focus is on evaluation methodology, rather than specific architectures. There are many researchers investigating architectures, but surprisingly little published work on evaluation methodology. Thus the workshop's immediate goal is to generate discussion of a wide spectrum of evaluation challenges and methods for addressing them. The next step is to harness such discussions to propose guidelines for evaluation of architectures that would be acceptable to the AI community, and allow researchers to both evaluate their own work, and the progress of others. We believe such guidelines will facilitate the collection of objective and reproducible evidence of the depth and breadth of an architecture's support for cognition, and its relationship to human or other natural cognition.

Based on the contributions to this workshop the following key questions will be raised to motivate the workshop discussion, with the goal of providing answers:

- What are the underlying research hypotheses one explores with architectures?
- Which functions/characteristics turn an architecture into an architecture supporting intelligence?
- How are architectures to be compared in an informative manner?
- What evaluation methods are needed for different types of cognitive architectures?
- What are the criteria and scales of evaluation?
- How should we validate the design of a cognitive architecture?
- Are there any relevant formal methods? Can we prove properties of architectures?
- Can we develop a common ontology for describing architectures and/or the various sets of requirements against which they can be evaluated?
- How can data-sets and benchmarks (standardized tasks) be used to evaluate architectures?
- Are there useful case-studies?
- How can we determine what architectures to use for different tasks or environments? Are there any trade-offs involved?

We hope that this workshop serves as a stepping stone in defining metrics and evaluation criteria for intelligent systems; thus leading to the development of a set of practicable guidelines for researchers in the AI community.

– Gal A. Kaminka (*Bar Ilan University, Israel*)

– Catherina R. Burghart (*University of Karlsruhe, Germany*)