AAAI 2009 Fall Symposium Series

Arlington, Virginia – November 5-7, 2009

Panel Discussion:

Comparative Repository of Architectures, Models, Tasks and Data

Chair: Christian Lebiere

Agenda

4:00 – 4:20	Owen Holland: Keynote Presentation
4:20 – 4:30	Christian Lebiere: Topic Introduction
4:30 – 5:20	Statements: (alphabetical order – 3 minutes each) Bernard Baars, Nick Cassimatis, Antonio Chella, Ben Goertzel, Steve Grossberg, Owen Holland, John Laird, Frank Ritter, Stuart Shapiro, Robert St. Amant, Andrea Stocco, Ron Sun, Kristinn Thorisson, Pei Wang
5:20 – 5:45	Panel Discussion with Audience Q&A

Objective

To identify the necessary means to achieve greater rates of convergence and incremental progress in cognitive modeling through the use of a shared repository of computational cognitive architectures, models, tasks and data.

Why do we need a repository?

1. **To facilitate direct comparison of different architectures.** The development and study of cognitive architectures has continued for many decades, yet it is difficult to see how different architectures map onto each other, what features or components are missing, and what progress has been made. Different communities speak different languages and largely ignore each other. Detailed models and data are seldom available in a comparable format, making direct model comparisons partial at best and tendentious at worst. Therefore, we need to
    a. Develop a common language for description of architectures and paradigms.
    b. Use common tasks and data as benchmarks for subfield.
    c. Develop shared, widely accepted comparison metrics.

    As a first step toward these goals, we propose to create a comparative table of main cognitive architectures using collective efforts of many researchers involved in their design,

study and usage. The first prototype of the table (intended as a simplest, top-level crystallization center) is posted at http://members.cox.net/bica2009/cogarch/ (please watch for updates). Reaching a cross-community agreement on the structure and the content of the table at this level will be a step allowing us to move down to details and forward – to further goals. We are planning to accomplish this first step during the discussion panel. Given this general agreement as a pre-requisite, the following goals are outlined below.

2. **To provide a centralized resource, that modelers, students, and teachers can access when they want to start a modeling research project.** The repository should facilitate:
    a. Finding all the available architectures for one's needs and purposes. This includes source code, executable program, documentation, papers and support community.
    b. Finding all the existing models for a given task or problem. This can be difficult since while some tasks are well identified, others can arise in many different formats.
    c. Finding relevant behavioral (or neuroscience) data for a given task. This means raw data rather than the aggregate analyses provided in publications to provide additional constraints for the developing of increasingly refined models.
    d. Finding an implemented version of experimental tasks. Too much time (as much as half by some estimates) is spent in modeling projects on (re)implementing and connecting to task environments. Often different modelers abstract away from a common task and thus prevent models from being directly comparable.

3. **To have an immediate and organized way to access an overview of relevant information.**
    a. Key findings of different subfields. Models are often developed in response to experiment results and validated by experimental findings. It would be helpful to have all relevant results available to provide broad and integrated rather than partial and selective accounts.
    b. Consistent and comparable record of activity for the various architectures. This would provide an archival record of the range of coverage of each architecture. It would highlight the core business of each architecture and would encourage keeping models updated to keep credit for current version of architecture.

4. **To enable the reuse of models.**
    a. Consistency in architectural parameters across models. Providing complete access to all model parameters would discourage excessive (i.e., post hoc) parameter fitting and encourage the adoption of consensus values.
    b. Ontological consistency in domain representation. Availability of standard ways of encoding knowledge for specific domains would enable the development of more complex, comprehensive models validated over broader range of findings.

5. **To encourage the development of modeling tools and standards.**
    a. Developing modeling tools (e.g., for model editing or parameter search) is a rather esoteric niche with little benefits. Making them available to a broad community would benefit the community through improved productivity.

b. Promoting the development of standards, such as for the integration of architectures and models and tasks environments. This would raise productivity as well as provide additional constraints on models.

How are we going to spread it?

1. Giving proper academic credit for uploading materials. Submissions might be subject to peer review. Full descriptions of models can be references with a DOI system or counted as online publication as in Scholarpedia (which includes a smart revision system).
2. Making models and data available in repository should be made a condition of publication and/or funding (as it happens in other fields). This can easily be done for specialized conferences (e.g., NIPS, ICCM).
3. Making behavioral data available for a given task will establish it as a de facto benchmark for its subfield. This will lead to a convergence towards a standardized set of tasks that will keep expanding rather than remain static and thus subject to be gamed.
4. Tying the repository into another system would allow users to make use of that system with no extra investment in effort. Examples of such external systems include simulation systems (e.g., Unreal Tournament), model running and parameter optimization system (e.g., MindModeling) and experiment system (e.g., Eprime). This would also enforce some code-compliance and standard policies.

How are we going to make it work?

Uploading tasks and code as currently existing is *not* enough. The following issues should be considered.

1. A standard API between cognitive architectures and task simulation environments is needed to assure portability across tasks and models. Tasks and models could only be included when they are compatible, ensuring interoperability. If both tasks and models comply with the API, both scientific (principled model comparison, separation between task and model) and technical (reusability, productivity) goals will be enhanced and the exponential growth associated with systems embracing common standards (Internet, PC) will then be possible. The primary scientific obstacle to such an API is to agree upon a common level of description across models and architectures. The alternative is to adopt a multi-level model approach that integrates models across multiple grain scales.
2. Models need to be updated and kept current. Developers will have incentives to maintain their code up-to-date to claim cumulative credit from models developed under previous versions of their architecture. The main issue is how much standardization should be required (e.g., fixed parameters, common knowledge representation) for an architecture to claim an integrated account across the models that it supports.
3. Infrastructure funding should be provided by some source, including DoD research agencies, DARPA, NSF, even foundations. The alternative is incremental funding through individual projects contributing ancillary development to the repository.
4. Before proceeding with the implementation, some informal polls or surveys should be taken to study the modelers' habits and needs, like Elsevier does with new products.

What are we going to discuss at the panel?

We need to share ideas and develop consensus. Below is a non-exclusive list of topics elaborating on the points above that could be divided among speakers and discussed in the Q&A. Given time limits, every panelist should (a) briefly address the first question, and then (b) focus on one key idea in their statement.

1. Do we agree on the structure and the content of the architecture comparison table? If not, then what do we need most in order to reach a general pre-requisite for a dialogue?
2. How to give academic credit for contributions to the repository?
3. How to foster the use of common tasks and the reuse of models?
4. What kind of data should be used to constrain architectures and models?
5. How to update models as architectures evolve?
6. How to foster integration between architectures?
7. How to converge on a common API for tasks and models?
8. How to reconcile architectures at different levels of abstraction?
9. How to enforce usage (lessons from fields with experience, incl. Biology, Physics)?