Learning from Design Experience in an Agent-Based Design System

Jarrod Moss, Kenneth Kotovsky
Department of Psychology, Carnegie Mellon University

Jonathan Cagan[1]
Department of Mechanical Engineering, Carnegie Mellon University

1. Motivation

A simple learning mechanism was added to an agent-based computational design system to see if it could then transfer knowledge across problems. An existing system, A-Design, was enhanced by giving it the ability to store useful design knowledge in a memory store so that this knowledge could be used in new design problems.

1. Background on A-Design

A-Design is a multi-agent computational system that automates the conceptual design process and is currently capable of solving several electromechanical design problems (Campbell, Cagan, & Kotovsky 1999, 2000). A-Design takes input and output constraints for a desired electromechanical device and produces an array of conceptual designs that satisfy the problem constraints. Along with the design problem, a number of evaluation criteria are also specified. For example, given the problem of designing a pressure gauge, the cost of the device should be kept to a minimum, the accuracy should be maximized, and the device should be as small as possible.

The generation of designs is accomplished through an iterative design process (see Figure 1). In each iteration of the design process, the various software agents interact to produce an array of candidate designs that satisfy the given problem constraints. These designs are all evaluated on the multiple objectives specified for the problem, and the designs are separated into good and bad designs based on a weighted evaluation function of these objectives. The good designs become part of the next iteration's population,

[1] Author of contact, Dept. of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, E-mail: cagan@cmu.edu

and they are allowed to "reproduce", in a manner described below, in the hope that some of their offspring will evaluate better than the original. The system will continue to develop and propagate designs to the next iteration until a set number of iterations has been reached.

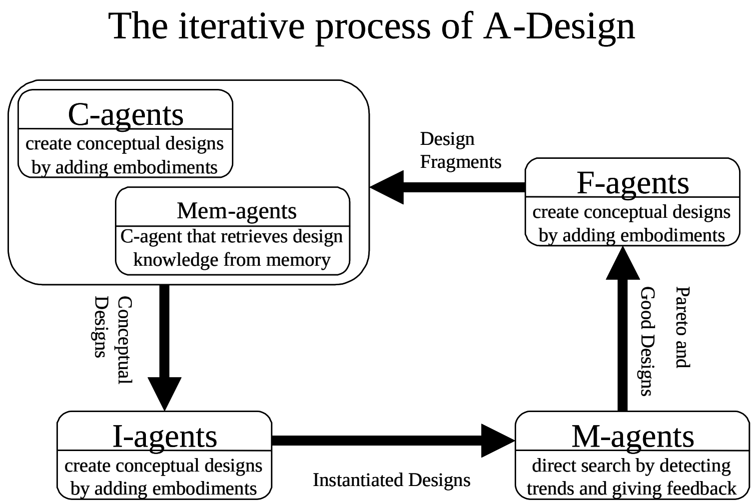

Figure 1. An iteration in A-Design is a complete cycle that starts with the C-agents

The actual construction of a candidate solution to the design problem begins with a group of creator agents (C-agents). C-agents select embodiments from A-Design's embodiment library and add these embodiments to an incomplete design. As components are added to an incomplete design, the design's properties are updated until the design satisfies the given input and output constraints of the design problem or until it is judged that the given design will probably never satisfy the design constraints. After a set of candidate designs has been constructed, the designs are passed to instantiation agents (I-agents) so that the conceptual design can be instantiated with real-world components. These agents select from a catalog of real components and instantiate a design by assigning these real components to the embodiments in a given design.

Manager agents (M-agents) guide the overall design process by choosing which C/I-agents to call at a given time and by extracting current trends from the design population. A trend is a particular set of agents or components that is present in a set of designs. Trends are looked for in the best and worst designs of a particular iteration. The trends from the best designs are passed back to the C/I-agents as a "todo" list, and the trends from the worst designs are passed back as a "taboo" list (Campbell, Cagan, & Kotovsky, 2001). Agents will try to avoid items on the taboo list while trying to produce items on the todo list. In this manner, the M-agents influence the design candidates that are generated in the next iteration.

In addition to keeping the good designs from the current iteration, copies of these good designs are passed to fragmentation agents (F-agents), who take out one or more components of the design. These fragmented designs are then reconstructed and become part of the next iteration's design population. This fragmentation process allows good designs to propagate similar designs to the next iteration with the hope that the changes made to the design will improve it.

2. Learning from Design Experience

In this study, A-Design was augmented to allow it to learn from its design experience. In order to do this, the existing trend extraction abilities of A-Design were utilized. Useful design knowledge was extracted after a problem had been solved by examining the best six designs produced in the final design iteration. Common subsystems were extracted from these designs, and these subsystems were added as chunks into a permanent memory store. These chunks could be a substantial part of the design or just a small set of commonly co-occurring components such as a rack and gear. Chunks were indexed in memory by the input and output constraints of the subsystem (see Figure 2). A-Design could then add these chunks to designs that it generated while solving new design problems.

Design Chunk

Rack-gear-chunk

Isa: design-chunk

Input-domain: translation

Input-interface: bolt

Output-domain: rotation

Output-interface: shaft-hole

Components: (rack gear)

Connectivity: port-2 of rack is connected to port-1 of gear

Figure 2. An example design chunk

A design chunk can be retrieved from memory based on its input constraints, its output constraints, or both. Three new C-agents (called Mem-agents, see Figure 1) were added to A-Design to embody these retrieval methods. As the C-agents are constructing designs, if one of the Mem-agents is called upon to add a component to an incomplete design, that agent retrieves a suitable chunk from memory and adds it to the design. The agents focus on a particular part of an incomplete design, and they try to retrieve a design chunk from memory based on the input/output constraints of this part of the design. These new agents along with the memory store are the core of A-Design's new learning and memory capabilities.

3. Results

Three electromechanical design problems were used to evaluate this new learning mechanism: a punch press, a pressure gauge, and a weighing machine (a scale). These specific problems were used to provide both problems that were similar to each other as well as problems that were significantly different. The pressure gauge and weighing machine are both measurement devices with a dial output, but the goal of the punch press is to amplify the small input force so that it is sufficient to drive a punch through some material. The weighing machine can be seen as similar to the pressure gauge but different from the punch press. A-Design was tested on these design problems to see if knowledge learned in one problem could be transferred successfully both within and across problems. Results indicate that A-Design applies learned knowledge very successfully in the same design problem where the knowledge was learned, and there was some successful transfer between problems as well. These results demonstrate the success of a relatively simple learning mechanism.

References

Cambell, M, Cagan J, and Kotovsky K: 1999, A-Design: An agent-based approach to conceptual design in a dynamic environment, *Research in Engineering Design* 11: 172-192.

Cambell, M, Cagan J, and Kotovsky K: 2000, Agent-Based synthesis of electromechanical design configurations, *Journal of Mechanical Design* 122: 61-69.

Cambell, M, Cagan J, and Kotovsky K: 2001, Learning from design experience: Todo/Taboo guidance, *Proceedings of the 2001 ASME Design Engineering Technical Conferences and Computers in Engineering Conference: Design Theory and Methodology Conference* DETC01/DTM-21687, Pittsburgh, PA.