

Learning from Design Experience in an Agent-Based Design System

JARROD MOSS

*Department of Psychology
Carnegie Mellon University, Pittsburgh, PA 15213*

Corresponding author:

JONATHAN CAGAN

*Department of Mechanical Engineering
Carnegie Mellon University, Pittsburgh, PA 15213*

cagan@cmu.edu
412-268-3713
Fax: 412-268-3348

KENNETH KOTOVSKY

*Department of Psychology
Carnegie Mellon University, Pittsburgh, PA 15213*

Abstract. A learning mechanism based on findings in cognitive science was added to an agent-based computational design system to determine if this mechanism would enable the system to learn from its experiences and transfer useful design knowledge to new problems. Learning and transfer were evaluated by examining how well knowledge learned while solving one problem could be applied to solve the same problem more effectively as well as how well this knowledge could be transferred to new design problems. An existing system, A-Design, was enhanced by giving it the ability to learn and store useful design knowledge so that this knowledge can be used in new design problems. Three electromechanical design problems were used to evaluate this new learning mechanism, and results indicate that this basic cognitively-based learning mechanism is successful at transferring design knowledge to new problems with a few limitations. This knowledge transfer leads to a more effective design process.

Keywords: agents, synthesis, learning, cognitive science

The copyright for the final published version of this article is owned by Springer. <http://dx.doi.org/10.1007/s00163-003-0042-4>
The final publication is available at www.springerlink.com

1. Introduction

Human designers typically become better at designing devices in a domain as their experience in that domain increases. Computational design systems should also be able to learn from their design experiences so that they can learn to produce higher quality designs in a shorter amount of time. The work presented here is an attempt to enhance an existing design system, A-Design, to give it the ability to learn from its experiences. In particular, the learning mechanisms that were added to the system are based on ideas about learning derived from the cognitive science literature.

The process of acquiring expertise is gradual and has been documented in a number of other domains (Chase & Simon 1973; Larkin, McDermott, Simon, & Simon 1980; Reitman 1976; Richman, Staszewski, & Simon 1995). The reason for this improvement in performance with experience is due in part to the expert's ability to utilize previous experiences. One aspect of using previous design experiences is the ability of the designer to transfer knowledge learned in a past design problem to the problem that is currently being pursued. Incorporating previous knowledge into a new problem allows designers to explore different design possibilities and prevents them from having to solve the same subproblems multiple times. If a designer encounters a part of the current design problem that resembles a previously encountered design or part of such a design, the designer has the option to incorporate parts of that previous design into the current design. This process allows designers to reuse other designs and so prevents them from going through the effort of solving the same problem twice. Alternatively, if part of a design is performing poorly or is too costly, a designer may be able to substitute other parts which perform better through an analogy process based on other designs in memory. These are only a few of the ways in which the ability to transfer knowledge between problems is potentially useful in engineering design. However, transfer processes have only been implemented in a few computational design systems.

Case-based design systems such as CADET (Sycara, Chandra, Guttal, Koning, & Narasimhan 1991), ARCHIE-II (Domeshek & Kolodner 1991) , and Kritik2 (Goel, Bhatta, & Stroulia 1997) all have the capability of transferring the knowledge they have stored as cases to new problems. However, only some of the systems have the capability

of indexing new cases into memory so that design experience can accumulate over time (Goel et al. 1997). While these systems are able to implement a form of knowledge transfer, there are other forms of transfer in which the knowledge being transferred does not consist of an entire previously encountered design. There are some cases where transfer of more abstract knowledge may be beneficial. In these cases, another transfer process may be required such as the chunking process added to A-Design in this study. The chunks in A-Design are more abstract in the sense that they consist of embodiments and not actual instantiated components as described below. One type of abstract process that has received some attention in previous work is analogy.

Analogy is a powerful method that may be used to transfer knowledge from previous designs as well as from other knowledge and experiences from the designer's memory. Analogies could be used to increase understanding of a novel design by allowing the designer to map previous experience onto the new device. There are many other uses for analogy in design, and there has been some work on models of design by analogy (Bhatta & Goel 1996; Howe, Cohen, Dixon, & Simmons 1986; Huhns & Acosta 1988).

However, there have been fewer attempts to incorporate knowledge transfer into more search oriented design systems. These systems usually employ a form of search that takes advantage of computational power to search a large number of possible designs, and they have been based on traditional AI search techniques (Ulrich 1989; Welch & Dixon 1994), genetic algorithms (Brown & Hwang 1993), and simulated annealing (Szykman & Cagan 1995). This paper describes an attempt to modify such a design system in order to incorporate some knowledge transfer processes. This modification allows a powerful search based system to take advantage of learned design knowledge.

A-Design is a multi-agent design system based on an iterative stochastic algorithm which in many ways resembles a genetic algorithm (Campbell, Cagan, & Kotovsky 1999, 2000, 2003). The work presented here is an attempt to augment A-Design with a basic learning mechanism based on human cognition that allows it to transfer knowledge across design problems. The idea is to give the system the capability to extract knowledge as it solves design problems which can then be used to improve performance when solving novel problems. This process allows A-Design to work with

design knowledge at an abstract level so that the design knowledge is not necessarily tied to a specific design. In addition, using this knowledge does not require a complex analogical process. This work is motivated by findings from cognitive psychology which are described next. Following this description, an introduction to A-Design is presented followed by a description of the changes to the system that allow it to learn from its design experience. The results demonstrate that A-Design was able to transfer knowledge from one design problem to another, but there are some limitations to this process as it is currently implemented. These limitations and possibilities for overcoming them are discussed along with some of the reasons for studying the cognitive basis of engineering design.

2. Cognitive Basis

Experienced designers approach a design problem with a large body of potentially relevant experience and background knowledge, and this knowledge helps these designers to produce better designs than designers who have little or no design experience. However, the presence and amount of this experience is likely to be only one of many differences between expert and novice designers. In general, the differences between experts and novices in a domain are characterized by the amount of relevant knowledge in memory, the representation of this knowledge, and the organization of this knowledge. These are the types of expert/novice differences that have been found in domains such as chess, physics, electronics, Go, and medicine (Chase & Simon 1973; Egan & Schwartz 1979; Larkin et al. 1980; Patel & Groen 1991; Reitman 1976; van de Wiel, Boshuizen, & Schmidt 2000). Results of these studies indicate that the organization of knowledge in memory differs between experts and novices in a domain just as much as the amount of knowledge does. These findings can be used as initial guides when conducting psychological studies of expert/novice differences in the domain of engineering design, but they can also contribute ideas on how to incorporate cognitive principles into computational design systems.

One of the most common findings in the study of expertise is that experts have the ability to group large amounts of information into a single memory unit or chunk. This

chunking process has been a finding in a number of the domains mentioned above (Chase & Simon 1973; Egan & Schwartz 1979; Reitman 1976). For example, chess masters are able to recall the pieces on a mid-game chess board almost perfectly after only five seconds of exposure, but novice chess players can only recall five to seven pieces (Chase & Simon 1973). However, this finding is not due to extraordinary memory abilities, since both types of players recall only about seven pieces from randomly generated board positions. An analysis of the extraordinary recall ability of the chess master indicates that this recall is supported by the ability to group four to six memory pieces into a single unit of memory storage called a chunk. Novices appear to be unable to construct such chunks. This chunking ability is supported by a large knowledge base of common perceptual patterns of chess positions. Experts appear to be able to recognize such common patterns from past experience, and then the location of the pattern only has to be referred to in long-term memory when recalling the pieces. This chunking ability in a domain has been used to explain how experts are better able to perceive, analyze, and act on situations in their domain (Chase & Simon 1973).

In the study presented here, the main issue is the kinds of mechanisms that must be employed in order to allow a computational design system to demonstrate the ability to learn and utilize familiar design patterns. There is no attempt to actually model the particular cognitive processes that human designers employ in order to learn such information; the idea is rather to supply a design system with a simple chunking ability modeled on findings from studies of cognitive processes used in other domains of expertise. An agent-based design system called A-Design was modified to include new processes and agents that allow it to store chunks from past design experiences and utilize these chunks when solving new design problems. This new learning mechanism is simple in comparison to the cognitive processes utilized by humans in order to transfer knowledge to new problems, but its successes and limitations provide insight that should aid in future attempts to incorporate learning processes into computational design systems.

3. Background: A-Design

An introduction to A-Design is required in order to explain how the agents in the system produce designs, and how the system was modified in order to include the new learning processes. A detailed description of A-Design can be found in Campbell (2000) and Campbell et al. (1999, 2000).

A-Design is an agent-based design system that produces an array of conceptual design solutions in response to a set of input and output constraints for a design problem. A design problem is specified to A-Design by specifying the input and output constraints of the desired device in A-Design's representation. For example, a punch press could be specified to the system by indicating that the input is a downward force on a handle, and the output of the device is a much larger downward force that drives a punch into some material (Figure 1). Along with these general input and output specifications, a number of other input/output constraints are specified such as the desire to minimize handle displacement in the punch press (shown as a goal of zero meters of displacement in Figure 1). This problem also specifies that the punch should only be displaced by .25m. In addition to these constraints a number of other objectives can be specified such as minimizing the cost and weight of the device. Once a problem and all of the associated objectives have been specified, A-Design's iterative design process attempts to produce a design that optimizes these design goals by working at both the configuration and component levels. Components are taken from a user-defined catalog that contains information about actual parts.

Figure 1

3.1. Iterative Design Process

A-Design's iterative design process is similar to a genetic algorithm because in each iteration a number of new design candidates are produced by the system, and then the best of these designs are chosen to be the basis for the production of new designs in the next iteration. The basic structure of A-Design's iterative design process is shown in Figure 2 along with the set of agents associated with each part of the process.

The design process begins with a set of configuration agents (C-agents) who construct candidate designs using a library of embodiments. Each C-agent adds one embodiment to an incomplete design until either the design is complete or a maximum number of embodiments have been added to the design. A candidate design starts off as just a set of input and output constraints to which components can be added. These input and output constraints are represented in structures called functional parameters (FPs). An FP represents the characteristics of an interface between components, and a C-agent utilizes the qualitative information in an FP to determine which embodiments in the embodiment catalog can be connected to the incomplete design. The embodiment catalog contains information about types of components, but these components are not instantiated with actual parameter values such as length, resistance, etc. While this embodiment catalog is user defined, Chen and Brown (2002) have shown that it may be possible for A-Design to modify this catalog itself. There is also a second catalog, the component catalog, which contains a number of instantiated versions of each embodiment. This component catalog is used by the I-agents as described below. Once an embodiment is added to one of the FPs, the free ports of the embodiment are included as new FPs where future embodiments can be attached. A form of qualitative reasoning is used to update the constraints in the incomplete design as each embodiment is added. A candidate is complete once it has connected the input and output FPs and has qualitatively satisfied the input and output constraints. A C-agent chooses which embodiment to add to an incomplete design based on a set of preferences built into the agent, the current state of the incomplete design, and other influences originating from the manager agents in the system. For example, some C-agents may prefer hydraulic components while others prefer electrical ones or components connected in series over ones connected in parallel. The current state of the design can also influence which embodiment a C-agent selects. If a device should have a bounded displacement at the output such as in the punch press example in Figure 1, then an agent would prefer components which accomplished this goal when added to the system. The C-agents' choices are also influenced by feedback they receive from manager agents in the system as discussed below. Once the set of C-agents have constructed a set number of designs these designs are passed to a group of instantiation agents (I-agents).

Figure 2

I-agents take the configuration designs from C-agents and instantiate the parameters in the system with values obtained from a catalog of components. Each embodiment in the configuration has a set of parameters such as length, weight, or resistance. These agents also have a set of built in preferences. For example, some I-agents may prefer using components that are low cost while others prefer those that are low weight. Once all of the components in a design have been instantiated, equations that describe the behavior of the design can be extracted. These equations allow the design to be evaluated on how well it satisfies the constraints of the given problem. For example, the equations extracted from a punch press design are then used to evaluate how much the handle is displaced since that is one of the specified constraints. In this implementation, each device is evaluated along the dimensions specified in the problem, and these evaluations are combined into a linearly weighted sum. All of the devices are then sorted by this sum and Pareto optimal, good (based on weighting), and poor devices are separated based on this ordering. These design partitions are then passed to a set of manager agents (M-agents). For the purposes of the work in this paper Pareto and good designs are both lumped into one category of good designs. The A-Design system still distinguishes between the two, but the chunking and memory mechanisms described below treat good and Pareto designs in the same way.

M-agents take the current design population and produce feedback that controls how other agents in the system operate. First, the agents which contribute to good and poor designs are examined and the probabilities controlling how often those agents contribute to designs are adjusted. A C-agent that contributed to a number of good designs will be called more frequently than one who contributed to many poor designs. The M-agents keep track of the number of good designs that each C/I-agent contributes to, and the probability that a specific C/I-agent will be called upon in the future is a function of its past success. These statistics are used for all of the agents except for the M-agents. Managers also look for trends in the design population. Trends are groups of agents or connected embodiments that appear together in a number of designs (see example in Figure 3). Good trends are found by examining the best designs, and bad

trends are extracted from the worst designs. In this implementation, the six best designs and the six worst designs are used to extract trends. Good trends are placed on a “todo” list and bad trends on a “taboo” list. These lists work by encouraging agents to reproduce combinations of agents or embodiments that are on the todo list and discouraging agents from reproducing the groupings on the taboo list. In this manner, the todo/taboo lists allow the M-agent to influence the designs that are generated in the next iteration of the design process, but *they exert this influence only within a given run on a single problem.*

In addition to passing the good designs from the current iteration into the next iteration, all 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 each design will improve it. The design population now consists of the good designs plus the newly reconstructed designs. In the next iteration the C-agents produce the number of new designs necessary to bring the design population back to the original number of designs. This design population level is a parameter of the system. This iterative process is the basis of the A-Design system.

Figure 3

4. Learning Across Problems in A-Design

In this work, we extended A-Design so that it could learn knowledge from its design experiences for application to *future* designs. In order to do this A-Design needed to be able to extract knowledge from its design problem solving activities, and once it had this knowledge there had to be some way of applying it to new design problems. One type of information that A-Design already knew how to extract was the good trends in a set of designs produced in a single iteration. The sets of interconnected embodiments that appear on the todo list are subsystems that appear in a number of good designs, and so it is likely then that these subsystems perform well in the current problem and may be worth remembering for future use.

A-Design's design process is a series of iterations in which the current designs perform the same as or better than those in the previous iteration, and so the subsystems found on the todo list in the final iteration are the ones from the best designs produced. These subsystems were chosen as the ones to be remembered in the memory store. Each subsystem appearing on this list is extracted and placed into memory as a chunk of knowledge. A subsystem from the todo list usually has many open ports to which additional embodiments can be attached. It is also desirable to be able to call one of these open ports the input to the subsystem and to call another port the output. In order to identify input and output ports, the new system randomly selects one of the designs that the subsystem appeared in before it was placed on the todo list. The embodiment closest to the input of the original design is labeled as the input to the subsystem, and the same process is done to label one port the output. Each chunk is then placed in memory and indexed by its input and output constraints since this is the only information that is needed to determine if a component can be added to an incomplete design by the C-agents in the system.

An example chunk consisting of a belt and pulley connected together is shown in Figure 4. This chunk extraction process is only half of the knowledge transfer process. After a knowledge base of chunks has been constructed, a new set of C-agents is needed that have the ability to add chunks from memory to designs during the configuration part of the design process. These memory agents (Mem-agents) add components to incomplete designs just as the C-agents do except that they are adding chunks from memory instead of embodiments from the embodiment catalog. There is a three step process by which a chunk is added to a design (Figure 5). First, a memory agent is called to work on an incomplete design. The memory agent then examines the open FPs in the system to determine where chunks may be added. The agent then looks in memory for a chunk that is compatible with one or more of the FPs in the design and adds this chunk. As mentioned before, all chunks are indexed in memory by their input and output constraints which are essentially the input and output FPs for the chunk. Chunks can then be retrieved from memory either based on their input, their output, or both. Each of these retrieval methods is embodied in one Mem-agent. For example, the Mem-agent implementing the input retrieval strategy will search for chunks in memory whose input

constraints match one of the open FPs in the design, but the input-output Mem-agent would search for chunks in memory whose input/output constraints match two open FPs, one for input and one for output.

Figure 4

Figure 5

There are some cases in which multiple chunks in memory are retrieved as potential chunks to be added to a design. In this case, a Mem-agent needs a way of determining which chunk to add. The chunk could just be randomly selected from all possible chunks, or there could be some form of learning that takes place within the memory agent that allows it to pick the chunk most likely to produce a good design. The second option was implemented and is based on learning mechanisms found in the ACT-R model of human memory (Anderson & Lebiere 1998). Each memory agent is informed about the consequences of its actions by the manager agent, and the memory agent can then use this feedback to choose among multiple chunks when it has to. When a Mem-agent adds a chunk to a design, the design can either end up being evaluated as a good design or a poor design as discussed above. The proportion of good designs that result from using a particular chunk is calculated by dividing the number of good designs the chunk was used in by the number of times the chunk was used. This proportion is calculated based on design evaluation data from the previous five iterations. When multiple chunks are retrieved, the probability that any particular one is chosen is just the proportion of good designs for that chunk divided by the sum of the proportions for all of the chunks retrieved from memory. If the proportion of good designs for a particular chunk is below some minimum level then this minimum proportion is used instead of the actual value so that every chunk has at least a small probability of being chosen regardless of its past performance. These proportions are updated dynamically each iteration to include information about good and poor chunks from the previous five iterations. This mechanism is also similar to the strategy for move selection in the simulated annealing algorithm found in (Hustin & Sangiovanni-Vincentelli 1987). Each of the three Mem-agents maintains a separate set of these statistics because different

chunks in memory may be more suited to one of the three memory retrieval strategies outlined above.

These additions to A-Design give it the basic capabilities of noticing and storing common co-occurring elements in chunks, noticing when a current design can be augmented with information from memory, and the ability to learn which chunks in memory perform the best in a given design problem. In general, one of these chunks serves as a way from getting from a certain type of interface and energy to another type of interface and energy. This allows for a type of abstract knowledge to be gained through the chunking of patterns of components. These basic mechanisms should allow the system to accumulate some amount of expertise in design problem solving. Similar chunking mechanisms are well documented learning processes in human cognition, and the chunking described above is similar to the way in which certain types of knowledge are acquired in the process of becoming an expert in some domains such as chess and Go (Chase & Simon 1973; Reitman 1976). There are also some similarities between A-Design's chunking of common design elements and the way in which certain types of cognitive architectures chunk together problem solving steps (i.e. SOAR, Newell 1990).

5. Testing Learning

The chunk learning mechanism of A-Design was tested to see if it allowed the system to learn both within and between problems. Learning within the same problem is simply the case where A-Design works on a design problem, learns chunks, and then works on the same design problem again with the new chunks. Within problem learning should indicate whether the chunking mechanism is allowing the system to learn and apply useful knowledge about the design problem. Between problem learning occurs when A-Design applies chunks learned in one problem to a new design problem, and it is this type of knowledge transfer that the chunking mechanism was designed to accomplish.

Within problem learning was evaluated by running A-Design 20 times on one problem, and in each of these runs, a set of chunks was generated from the final design population. A-Design was then run again on the same problem 20 times, once with each

of the 20 sets of learned chunks. The number of iterations in each run was 60, and the design population was set to 120 designs. A graph of A-Design's performance on a design problem can be constructed by taking the evaluation score for the best design at each iteration and averaging this evaluation across all twenty runs of the problem. This leads to a graph similar to Figure 6 in which each line in the figure shows A-Design's average performance on one design problem. The performance difference between the initial problem and the second attempt at the problem was measured in three ways (see Figure 6). The evaluations of the initial designs, the evaluations of the final designs, and the number of iterations until a specified evaluation level is reached can be compared. The comparison of the number of iterations can be looked at as the number of iterations that were saved by the presence of chunks and it will be referred to as a measure of savings. The criterion evaluation level for each problem presented here is just the average evaluation score obtained in the final best design without using a set of stored chunks, i.e. just the performance of the A-Design system without the new learning mechanism. Any statistically significant differences in performance from the first attempt to the second can then be attributed to the chunks that were used in the second attempt.

This measure of savings was calculated by recording the iteration at which each run of A-Design reached the criterion level. If an individual run did not reach the criterion level by the end of 60 iterations, then the number of iterations to criterion was just taken to be 60 for that run. The average number of iterations for a given condition can then be calculated by averaging the number of iterations to criterion for each of the 20 runs in a condition. The average number of iterations to criterion can then be compared for any two conditions (i.e. for chunk and no chunk conditions). This yields a conservative estimate of the amount of savings since each run is capped at 60 iterations even if it would have taken longer to reach criterion.

Figure 6

Another potential metric that should be kept in mind is the actual running time for A-Design to finish a set of 60 iterations as outlined above. In general, the Mem-agents are computationally more expensive than normal C-agents. It is useful to determine if these extra costs significantly impact the running time. Run time results were examined

for all three problems, and it was found that the average time for the completion of one iteration did not differ significantly when Mem-agents were used. Therefore any benefits of the chunking processes do not come at the expense of significantly increased run time.

5.1. Design Problems

Three electromechanical design problems were used to assess the benefits of the chunking mechanism: a punch press, a pressure gauge, and a weighing machine. The punch press problem is the same as described above, a handle is pulled which forces a punch through some material at the output. Punch presses are evaluated based on their cost, the amount of input handle displacement, and how closely they conform to the specified output displacement and force. The pressure gauge problem has an input pressure source and the output is a dial display that reflects the amount of pressure coming from the pressure source, and this problem is evaluated on the cost, mass, efficiency, and dial accuracy of the gauge. The weighing machine as defined by (Campbell et al. 1999) takes a force input on a footpad and has a dial output, and it is evaluated on the cost, mass, dial accuracy, and input displacement of the device. An example of the linearly weighted evaluation function is presented in Figure 7. The evaluation function deals with real quantities like dollars, radians, and grams so the scales are difficult to interpret when combined with the weights. The relative weighting is the most important aspect of the evaluation function for our purposes. For example, the weights presented in Figure 7 emphasize dial accuracy while placing little emphasis on cost. In this equation, characteristics like cost can be read from the cost of the components, but to evaluate accuracy a behavioral equation must be extracted from the design so that accuracy can be calculated. Some example designs produced by A-Design for these problems are presented in Figure 8. These example designs are typical of the kinds of solutions A-Design produces for each problem. The component catalog used in this research is the same as the one that has been presented in previous A-Design work (Campbell 2000; Campbell et al. 1999, 2000). The only additions to the catalog were the input and output components for the new problems, such as the handle and punch for the punch press problem.

Figure 7**Figure 8**

These specific problems were utilized to provide both a problem that was similar to the weighing machine as well as a problem that was significantly different from the weighing machine. The pressure gauge is a measurement device with a dial output just like the weighing machine. However, 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. It was thought that this problem shares little with the weighing machine so it was used to evaluate how context dependent A-Design's new knowledge capabilities would be. For example, if A-Design learned design chunks from the weighing machine problem, then these chunks might be easier to apply in the pressure gauge problem since this problem is similar to the weighing machine. On the other hand, applying these same design chunks to the punch press might be more difficult and potentially less useful because the punch press does not have as much in common with the weighing machine. This means that the similarity between the devices probably has an impact on which learned chunks are transferred successfully.

The specific evaluation function used to evaluate a device does have an impact on the chunks learned from that design experience. For example, a pressure gauge design that minimizes cost above all else would be different from a design which primarily emphasizes the accuracy of the gauge. These different designs would lead to different chunks, and it may be that these chunks will transfer best to designs that have a similar emphasis in their evaluation functions. This aspect of chunking is not explored in the current work. Our primary goal was to evaluate the viability of the simple chunking mechanism described above. We tried to keep the emphasis in the evaluation functions as constant as possible. For example, both the weighing machine and the pressure gauge evaluations emphasize accuracy more than any other criteria.

Two sample chunks are shown in Figure 9. Figure 9a shows a chunk learned by A-Design while solving the weighing machine problem. This chunk was then used in the production of pressure gauge designs, and in particular, this chunk was used in the design

of the pressure gauge shown in Figure 8. Figure 9b is another example of a chunk learned in the weighing machine problem, and it was used in the production of the punch press design seen in Figure 8. The first chunk is larger as it contains 7 components as opposed to the smaller chunk composed of two components. Larger chunks were more frequently transferred between the weighing machine and pressure gauge problems than between the punch press and other problems. This finding was not unexpected and is due to the degree of similarity among the problems as noted above.

Figure 9

6. Results

6.1. Within problem results

Within problem results for the three design problems can be seen in Figures 10-12 and Tables 1-3. There appears to be within problem learning in each of the three problems (lower evaluation scores are better). A series of paired t-tests was run to assess the statistical significance of any differences in the three performance measures discussed above. In some cases the assumptions of the t-test did not hold (such as equal variances). In these cases nonparametric statistical tests were also run, but the results from these tests did not differ from those of the t-tests.

6.1.1. Weighing Machine

In the weighing machine problem only the comparison of the initial designs was not significantly different, $t(19) = 1.19$, $p = .12$, but both the comparison of the last designs, $t(19) = 2.47$, $p = .01$, and the amount of savings, $t(19) = 3.73$, $p < .001$, were significant. On average, the final design produced without chunks had an evaluation of 21.7 as compared to 19.3 with chunks, and the number of iterations required to reach an evaluation of 21.7 was 45.6 without chunks and 30.7 with chunks, 33% fewer iterations.

Figure 10

6.1.2. Pressure Gauge

In the pressure gauge problem the comparison of initial designs, $t(19) = 2.14$, $p = .02$, and savings, $t(19) = 1.82$, $p = .04$, were significant, but the comparison of final design evaluations was not, $t(19) = 1.19$, $p = .12$. On average, the initial design produced without any chunks was 50,302 as compared to 35,358 with chunks, and the number of iterations required to reach an evaluation of 3,258 was 44 without chunks and 32.7 with chunks, 26% fewer iterations.

Figure 11

6.1.3. Punch Press

In the punch press problem the comparison of initial designs, $t(19) = 3.32$, $p = .002$, final designs $t(19) = 3.34$, $p = .002$, and savings, $t(19) = 2.73$, $p = .007$, were significant. On average, the initial design evaluation produced without any chunks was 9,046 as compared to 6,675 with chunks, the final design produced without any chunks was 1,658 compared to 1,016 with chunks, and the number of iterations required to reach an evaluation of 1,658 was 38.9 without chunks and 23 with chunks, 41% fewer iterations.

Figure 12

6.2. Between problem results

The results for between problem transfer can be seen in Figures 13-15 and Tables 4-6. For each problem, a one way ANOVA was run for each comparison followed by a series of planned contrasts if the ANOVA indicated any significant differences. In some cases the assumptions of the ANOVA did not hold (such as equal variances). In these cases nonparametric statistical tests were also run, but the results from these tests did not differ from those of the ANOVAs.

6.2.1. Weighing Machine

Figure 13 indicates that the three conditions perform about the same, and there were no significant differences in initial evaluations, $F(2,57) = .74$, $p = .24$, final evaluations, $F(2,57) = .06$, $p = .47$, or savings, $F(2,57) = .804$, $p = .23$, for the weighing machine problem. Thus neither the pressure gauge nor the punch press chunks produced significant improvements in the weighing machine problem. In this case, the within problem transfer results outperformed all cases of between problem transfer.

Figure 13

6.2.2. Pressure Gauge

Figure 14 indicates that the two conditions with chunks perform better than solving the problem with no previous knowledge. There were significant differences in final evaluations, $F(2,57) = 2.94$, $p = .03$, and savings, $F(2,57) = 4.53$, $p = .008$, but not for the initial evaluations, $F(2,57) = .9$, $p = .21$. A series of contrasts reveal that chunks from the weighing machine problem allow the system to produce better designs than the no chunk condition, and both the weighing machine and punch press chunk conditions produce significant savings when compared to the no chunk condition. Final design evaluations in the weighing machine chunk condition had an average of 2,766 while the average in the no chunk condition was 3,258. Chunks from the weighing machine problem help the system to reach an evaluation of 3,258 in 28.8 iterations as compared to 44 iterations for the no chunk condition, and chunks from the punch press problem also allow the system to reach the criterion evaluation in 29 iterations, 34-35% fewer iterations in both cases. It can also be seen (Figures 14 and 11) that within problem transfer produced similar quality final designs and reached criterion in a similar number of iterations as did the two between problem transfer conditions. In fact, the two between problem transfer cases appear to outperform the within problem transfer, but these differences were not statistically significant.

Figure 14

6.2.3 Punch Press

Figure 15 indicates that the two conditions with chunks perform better than solving the problem with no previous knowledge. There were significant differences in final evaluations, $F(2,57) = 4.09$, $p = .01$, but not for the initial evaluations, $F(2,57) = .65$, $p = .26$, or savings, $F(2,57) = 1.37$, $p = .13$. A series of contrasts reveal that both of the chunk conditions produce significantly better final designs when compared to the no chunk condition. Final design evaluations in the weighing machine chunk condition had an average of 1,298 and in the pressure gauge chunk condition the average was 1,271, while the average in the no chunk condition was 1,658. The within problem transfer case (Figure 12) outperformed between problem transfer in this problem.

Figure 15

7. Discussion

The chunking mechanism did produce successful knowledge transfer in most cases. Transfer was found in most cases where it was expected, but interestingly there was no significant between problem transfer when working on the weighing machine problem. This was unexpected because of the similarity of the weighing machine and pressure gauge problems, and because there was transfer from the weighing machine problem to the pressure gauge problem. Performance on the punch press problem improved regardless of the types of chunks used which may indicate some unexpected similarities between the punch press and the other problems or the existence of some chunks that work regardless of the problem context in which they were learned. In general, within problem transfer was the most effective type of transfer as none of the between problem cases significantly outperformed the within problem case in any of the problems. Knowledge transfer within the same problem produced improvements in two to three of the performance measures in each problem. In the cases where significant between problem transfer did occur, the performance measures that improved were the final design evaluations and the savings measure. So while the knowledge from another

problem did not produce better initial designs in any of the three problems, it was able to improve performance throughout the iterative design process.

This eventual but not immediate improvement due to transfer could indicate that the chunks learned by the system could only be applied to a new problem at some late or intermediate stage of the design process. Another explanation is that only some of the chunks learned in another problem are beneficial to the current problem, and it takes many attempts to find which of the chunks are beneficial when applied to the current design problem. An analysis of the data produced by the system indicates that certain chunks did come to be preferred by the system in a certain problem, and this was mostly due to the success statistics that are accumulated by Mem-agents as discussed previously. It is also possible that some of the chunks had to be slightly modified in order to produce improvements in a particular problem, and this modification process did appear to operate as discussed below.

It was hypothesized that most of the between problem transfer would occur between the weighing machine and pressure gauge problems due to their similarity. In the pressure gauge problem, the within problem chunks lead to the most improvement followed by the chunks from the weighing machine problem and finally the punch press chunks. This graded transfer makes sense as it is easiest for the system to apply knowledge from the same problem and somewhat more difficult to apply knowledge from a similar problem. However, in the weighing machine problem there was no significant between problem transfer at all. This behavior appears to be caused by the embodiments that occur in the chunks learned from the pressure gauge problem. Every pressure gauge takes some pressure source as its input, and, given the catalog that was supplied, the only way A-Design has of transforming this pressure source into a translational motion is by using a hydraulic cylinder. So all pressure gauges have cylinders in them, and since chunks are found by extracting commonalities, a large portion of the chunks learned by A-Design in this problem have the cylinder embodiment in them. None of the best weighing machine designs produced had cylinders in them because of the high cost of this embodiment, and so these pressure gauge chunks do not appear to be very useful in the weighing machine problem. This seems like the most likely cause of the asymmetric transfer.

The results from the punch press problem were also a little surprising. It was originally thought that this problem had few similarities to either the weighing machine or pressure gauge problems, but in both cases transfer to the punch press improved results. One similarity that does exist between the pressure gauge and punch press problems is that good designs produced for both problems contain hydraulic cylinders, and it could be that this similarity made it easier to transfer chunks from one problem to the other. There were no such similarities noted for the weighing machine problem and the punch press problem though, and so this does indicate that at least some of the chunks learned in the weighing machine problem are useful even in dissimilar problems. These chunks were mostly small chunks like the one presented in Figure 9b, and these types of chunks are applicable in a wide number of designs.

It was also found that in a few cases A-Design's preexisting mechanisms interacted with the new Mem-agents to generalize learned chunks to new problems. In particular, if a chunk was useful enough that the design it was incorporated in was evaluated as a good design, then this design passed through the fragmentation process and was reconstructed. In this manner, F-agents were able to contribute to the utility of some chunks since parts of the chunk that did not contribute to improving the overall design quality had a chance to be removed while other parts of the chunk remained. This mechanism was an unexpected interaction between preexisting processes and the new chunking mechanisms, and it contributed to the success of the memorized chunks since they could be adapted to the new design problem. However, this mechanism was not able to overcome the presence of the cylinder in the pressure gauge chunks. One reason for this is that the cylinder is an expensive component, and its inclusion in a weighing machine will cause the design to evaluate poorly with respect to the many other less expensive weighing machine designs. This would mean that this chunk would never be in a good design, and so it would never reach the F-agents.

These results indicate that this chunk storage mechanism is a useful learning mechanism and has the potential to be an even more useful general learning mechanism with some further modifications. A-Design could adapt the learned chunks to new problems, and this learning mechanism could be combined with a more general analogy or reasoning mechanism to make the learned chunks even more beneficial. The necessity

for a more general mechanism can be seen by A-Design's inability to get around the predominance of hydraulic cylinders in the chunks learned in the pressure gauge problem in order to use them in the weighing machine problem. However, just this basic chunking mechanism was enough to enable the system to exhibit significant improvements on most problems in the number of iterations needed to produce designs of equal or better quality when compared to the system without chunking.

These results can be taken to reflect the widespread utility of chunking as it is found in cognitive structures in a variety of domains. As mentioned before, A-Design's chunking process is based on work in cognitive psychology which shows that building up an organized structure of domain relevant chunks is one component of the acquisition of expertise (Chase & Simon 1973; Egan & Schwartz 1979; Reitman 1976; Richman et al. 1995). In this manner experts begin to work with larger and more sophisticated units of knowledge as they gain experience. This network of knowledge can make solving common problems more efficient within the area of expertise (Larkin et al. 1980). A-Design currently has the ability to store chunks of knowledge, but these chunks are not highly organized in a hierarchical structure as they would be in an expert. Also, the existence of chunks only makes moderate changes to the way A-Design solves problems. While these limitations highlight A-Design's shortcomings as a cognitive model of engineering design, they also provide some indication of the potential advances that could be made through a better understanding of the cognitive processes involved in engineering design.

8. Conclusions

Overall, there was some success in producing knowledge transfer with a basic chunking mechanism. The mechanism takes advantage of A-Design's preexisting trend finding function to extract chunks, and all that was required to use these chunks was a new type of configuration agent that looks in a memory store of chunks instead of the embodiment library. This method of knowledge transfer still relies on the power of computational search in A-Design's iterative design process. This mechanism was able to produce impressive results when transferring chunks between problems. The system

operates on embodiment level similarities between problems, and, using only a basic approach, lacks a more powerful method of transfer. Such a transfer mechanism would probably rely on more knowledge based methods in addition to or instead of brute force computational search. For example, processes such as abstraction, qualitative reasoning, and analogy can produce knowledge transfer that operates across vastly different problems and domains. Even with the current limited processes, A-Design's transfer results are reminiscent of transfer results in cognitive psychology.

Research looking at knowledge transfer has shown that similarities between problems are one of the most important factors that determine the amount of transfer (Holyoak & Koh 1987; Kotovsky & Fallside 1989). Finding similarities between two different problems is also a key to using analogy (Falkenhainer, Forbus, & Gentner 1989; Holyoak & Thagard 1995). Analogy can be a powerful mechanism by which people can apply knowledge from one domain to another, and analogy may help an engineer map previous experiences onto novel design problems. However, all of these transfer processes operate on the contents of memory, and therefore a better understanding of the structure and content of the memories of expert designers will aid in understanding how such transfer processes operate. Cognitive science research has identified a number of general properties of expert memory in many other domains such as its hierarchical organization and chunking properties (Chase & Simon 1973; Egan & Schwartz 1979; Ericsson, Chase, & Faloon 1980; Patel & Groen 1991; Reitman 1976; Richman et al. 1995). We have also been extended this type of work to deal directly with expertise differences in engineering design (Moss, Kotovsky, & Cagan 2003). Our work highlights some of the differences in device representation that are associated with the early acquisition of expertise in engineering design. These results combined with new work on understanding the cognitive processes underlying design will yield a number of benefits including the construction of more knowledgeable design automation and assistance tools. A-Design has been the basis for our initial exploration into incorporating cognitive principles into design systems, and further research in this area should yield benefits for computational design systems and, we hope, lead to a better understanding of the cognitive processes underlying the design process.

Acknowledgments

The authors would like to acknowledge the assistance and support provided by Matthew Campbell while learning the intricacies of the A-Design system and while generating some initial ideas on how the system could be extended.

This research effort was partially sponsored by a National Defense Science and Engineering Graduate Fellowship and by the Air Force Office of Scientific Research, Air Force Material Command, USAF, under grant number F49620-01-1-0050. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of AFOSR or the U.S. Government.

References

- Anderson, JR, Lebiere, C (1998) *The atomic components of thought*. Lawrence Erlbaum, Mahwah, NJ
- Bhatta, SR, Goel, AK (1996) From design experiences to generic mechanisms: Model-based learning in analogical design. *Artificial Intelligence for Engineering Design Analysis and Manufacturing*, 10(2): 131-136
- Brown, DR, Hwang, KY (1993) Solving fixed configuration problems with genetic search. *Research in Engineering Design*, 5(2): 80-87
- Campbell, MI (2000) *The A-Design invention machine: A means of automating and investigating conceptual design*. Unpublished Ph.D. Dissertation, Carnegie Mellon University, Pittsburgh, PA
- Campbell, MI, Cagan, J, Kotovsky, K (1999) A-design: An agent-based approach to conceptual design in a dynamic environment. *Research in Engineering Design*, 11(3): 172-192
- Campbell, MI, Cagan, J, Kotovsky, K (2000) Agent-based synthesis of electromechanical design configurations. *Journal of Mechanical Design*, 122(1): 61-69
- Campbell, MI, Cagan, J, Kotovsky, K (2003) The A-Design approach to managing automated design synthesis. *Research in Engineering Design*, 14(1): 12-24
- Chase, WG, Simon, HA (1973) Perception in chess. *Cognitive Psychology*, 4(1): 55-81
- Chen, ZF, Brown, DR (2002) *Explorations of a two-layer A-Design system*. Paper presented at the Workshop on Agents in Design, Cambridge, MA.
- Domeshek, EA, Kolodner, JL (1991) Toward a case-based aid for conceptual design. *International Journal of Expert Systems*, 4(2): 201-220
- Egan, DE, Schwartz, BJ (1979) Chunking in recall of symbolic drawings. *Memory & Cognition*, 7(2): 149-158
- Ericsson, KA, Chase, WG, Faloon, S (1980) Acquisition of a memory skill. *Science*, 208(4448): 1181-1182
- Falkenhainer, B, Forbus, KD, Gentner, D (1989) The structure-mapping engine - algorithm and examples. *Artificial Intelligence*, 41(1): 1-63
- Goel, AK, Bhatta, SR, Stroulia, E. (1997) Kritik: An early case-based design system. In M. L. Maher & P. Pu (Eds.), *Issues and applications of case-based reasoning in design* (pp. 87-132). Mahwah, NJ: Lawrence Erlbaum
- Holyoak, KJ, Koh, K (1987) Surface and structural similarity in analogical transfer. *Memory & Cognition*, 15(4): 332-340
- Holyoak, KJ, Thagard, P (1995) *Mental leaps: Analogy in creative thought*. The MIT Press, Cambridge, MA
- Howe, AE, Cohen, PR, Dixon, JR, Simmons, MK (1986) Dominic: A domain-independent program for mechanical engineering design. *Artificial Intelligence in Engineering*, 1(1): 289-299

- Huhns, MN, Acosta, RD (1988) Argo - a system for design by analogy. *IEEE Expert-Intelligent Systems & Their Applications*, 3(3): 53-68
- Hustin, S, Sangiovanni-Vincentelli, A (1987, April) *Tim, a new standard and cell placement program based on the simulated annealing algorithm*. Paper presented at the IEEE Physical Design Workshop on Placement and Floorplanning, Hilton Head, SC.
- Kotovskiy, K, Fallside, D. (1989) Representation and transfer in problem solving. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: The impact of herbert a simon* (pp. 69-108). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Larkin, J, McDermott, J, Simon, DP, Simon, HA (1980) Expert and novice performance in solving physics problems. *Science*, 208(4450): 1335-1342
- Moss, J, Kotovsky, K, Cagan, J (2003) *Knowledge representation in engineering design: An initial investigation*. Paper presented at the Twenty-Fifth Annual Conference of the Cognitive Science Society, Boston, MA.
- Newell, A (1990) *Unified theories of cognition*. Harvard University Press, Cambridge, MA
- Patel, VL, Groen, GJ. (1991) The general and specific nature of medical expertise: A critical look. In K. A. Ericsson & J. Smith (Eds.), *Toward a general theory of expertise: Prospects and limits* (pp. 93-125). New York: Cambridge University Press
- Reitman, JS (1976) Skilled perception in Go: Deducing memory structures from inter-response times. *Cognitive Psychology*, 8(3): 336-356
- Richman, HB, Staszewski, JJ, Simon, HA (1995) Simulation of expert memory using epam iv. *Psychological Review*, 102(2): 305-330
- Sycara, K, Chandra, DN, Guttal, R, Koning, J, Narasimhan, S (1991) Cadet: A case-based synthesis tool for engineering design. *International journal of expert systems*, 4(2): 157-188
- Szykman, S, Cagan, J (1995) A simulated annealing-based approach to 3-dimensional component packing. *Journal of Mechanical Design*, 117(2): 308-314
- Ulrich, KT (1989) *Computation and pre-parametric design*. Unpublished doctoral dissertation, MIT, Cambridge, MA
- van de Wiel, MWJ, Boshuizen, HPA, Schmidt, HG (2000) Knowledge restructuring in expertise development: Evidence from pathophysiological representations of clinical cases by students and physicians. *European Journal of Cognitive Psychology*, 12(3): 323-355
- Welch, RV, Dixon, JR (1994) Guiding conceptual design through behavioral reasoning. *Research in Engineering Design*, 6: 169-188

Figure 1. The specification for the punch press problem.

Figure 2. A-Design's iterative design process (adapted from Campbell et al. 1999).

Figure 3. An example trend found by intersecting multiple designs with the overlapping components circled.

Figure 4. An example design chunk

Figure 5. Memory agents retrieve chunks from memory

Figure 6. An illustration of the three measures of performance

Figure 7. An example evaluation function used for the pressure gauge

Figure 8. Sample solutions produced by A-Design for each of the three design problems.

Figure 9. Sample chunks learned from the weighing machine problem and used in the pressure gauge and punch press problems in Figure 7.

Figure 10. Within problem results for the weighing machine problem

Figure 11. Within problem results for the pressure gauge problem.

Figure 12. Within problem results for the punch press problem.

Figure 13. Between problem results for the weighing machine problem

Figure 14. Between problem results for the pressure gauge problem.

Figure 15. Between problem results for the punch press problem.

TABLE 1. Within problem performance measures for the weighing machine.

	Initial evaluations		Final evaluations		Iterations to criterion	
	μ	σ	μ	σ	μ	σ
No chunks	773	1001	21.7	4.90	45.6	16.0
Using within problem chunks	465	632	19.3	3.29	30.7	20.3

TABLE 2. Within problem performance measures for the pressure gauge.

	Initial evaluations		Final evaluations		Iterations to criterion	
	μ	σ	μ	σ	μ	σ
No chunks	50,302	25,775	3,258	807	44.0	18.0
Using within problem chunks	35,358	26,275	3,005	643	32.7	23.5

TABLE 3. Within problem performance measures for the punch press.

	Initial evaluations		Final evaluations		Iterations to criterion	
	μ	σ	μ	σ	μ	σ
No chunks	9,046	1,327	1,658	524	38.9	21.4
Using within problem chunks	6,675	2,762	1,016	517	23.0	18.7

TABLE 4. Between problem performance measures for the weighing machine.

	Initial evaluations		Final evaluations		Iterations to criterion	
	μ	σ	μ	σ	μ	σ
No chunks	773	1001	21.7	4.90	45.6	16.0
Using pressure gauge chunks	700	1397	21.2	4.21	49.5	12.6
Using punch press chunks	405	352	21.4	4.73	43.2	18.4

TABLE 5. Between problem performance measures for the pressure gauge.

	Initial evaluations		Final evaluations		Iterations to criterion	
	μ	σ	μ	σ	μ	σ
No chunks	50,302	25,775	3,258	807	44.0	18.0
Using weigh. machine chunks	42,190	23,483	2,766	441	28.8	18.7
Using punch press chunks	52,670	2,8337	2,889	705	28.9	18.4

TABLE 6. Between problem performance measures for the punch press.

	Initial evaluations		Final evaluations		Iterations to criterion	
	μ	σ	μ	σ	μ	σ
No chunks	9046	1327	1658	524	38.9	21.4
Using weigh. machine chunks	8583	1493	1298	413	28.7	20.0
Using pressure gauge chunks	8472	2139	1271	491	30.7	20.9

Punch Press problem

Input-device: Handle

Output-device: Punch

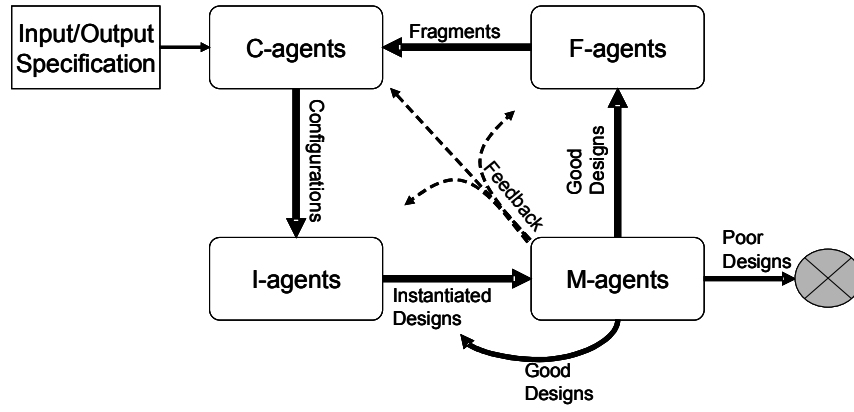
Input-force: 6 N

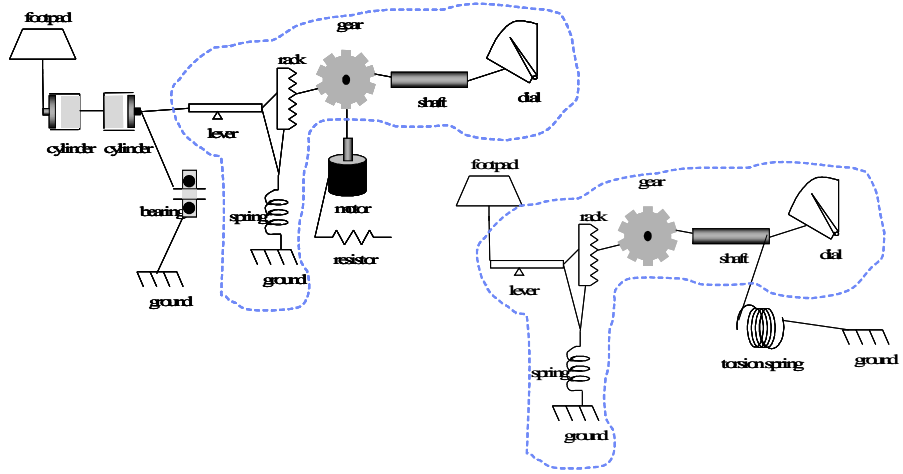
Output-force goal: 100 N

Input-displacement goal: 0 m

Output-displacement goal: .25 m

Objectives: minimize cost, minimize weight





Design Chunk

Belt-pulley-chunk

Isa: design-chunk

Input-domain: translation


Input-interface: bolt

Output-domain: rotation


Output-interface: shaft-hole

Components: (belt pulley)

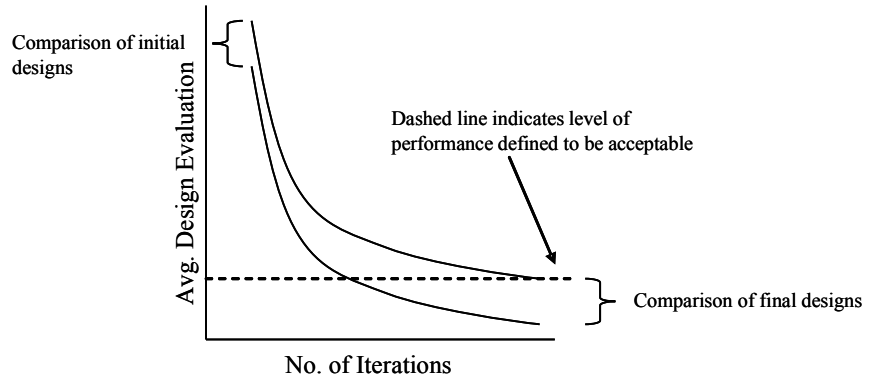
Connectivity: port-1 of component-1 is connected to port-1 of component-2



Incomplete Design




Memory agent retrieves chunk



$$\textit{Evaluation} = w_1 \textit{ cost} + w_2 \textit{ mass} + w_3 \textit{ efficiency} + w_4 \textit{ accuracy}$$

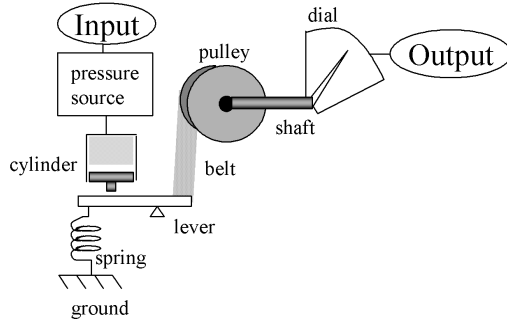
$$w_1 = 1, w_2 = 10, w_3 = 10, w_4 = 100$$

Weighing Machine




Design Objectives:
cost: \$14.61
mass: 0.14 kg
input dx: 5.3 cm
accuracy: .004 rad

Pressure Gauge




Design Objectives:
cost: \$532.76
mass: 1.7 kg
efficiency: 78%
accuracy: .02 rad


Punch Press




Design Objectives:
cost: \$741.74
input dx: 1.6 cm
output dx: 1.0 mm
output force: 101 N




Weighing Machine - Within problem




Pressure Gauge - Within Problem

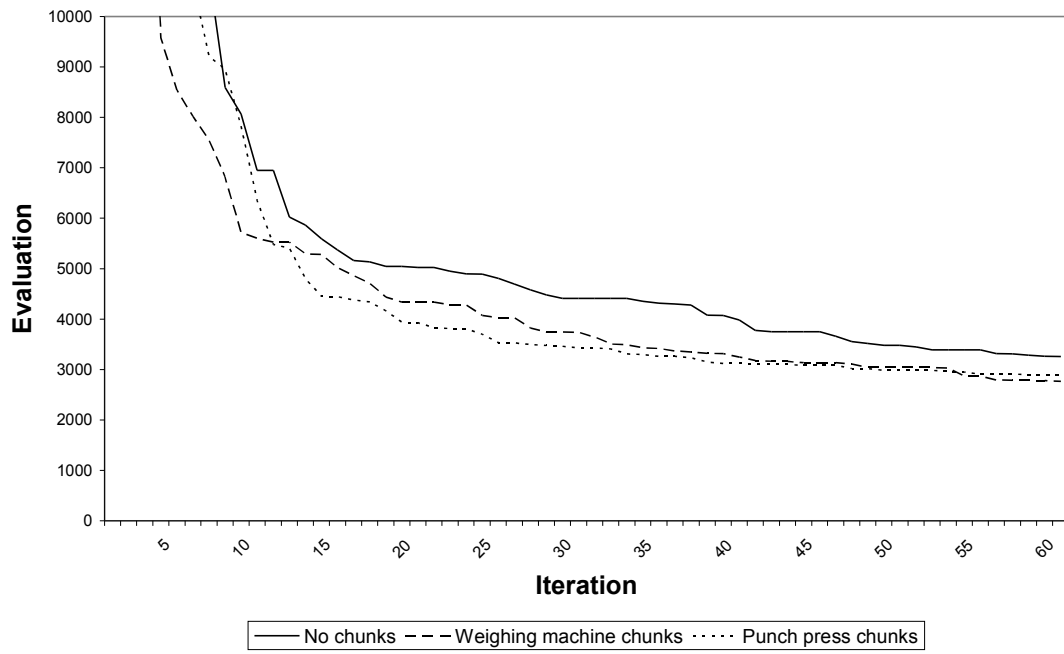


Punch Press - Within Problem

Weighing Machine - Between Problems



Pressure Gauge - Between Problems



Punch Press - Between Problems

