Experience-based Reasoning as the Basis of a General Artificial Intelligence Architecture

Jim Davies

Institute of Cognitive Science
2201 Dunton Tower, Carleton University
1125 Colonel By Drive, Ottawa, Ontario, K1S 5B6, Canada
613-620-2888
jim@jimdavies.org

Abstract

I propose a challenge to the experience-based reasoning community, including analogy and case-based reasoning researchers, to create a general artificial intelligence architecture based on the principles of experience-based reasoning. This ambitious proposal would push the core ideas of the field to their limits, and if successful raise awareness of the field. I list some major cognitive tasks that such an architecture should be able to handle, and briefly sketch how an experience-based reasoning architecture might accommodate them.

Introduction

The structure of the field of artificial intelligence (AI) is organized by its methodologies. Major players include Bayesian reasoning, artificial neural networks, logic, and probability theory. Most AI research consists of applying one of these methodologies to a problem, with researchers, sensibly, applying the most promising methodology to the problem in question. But there is a growing interest in modeling general intelligence (Voss, Goertzel, & Pennachin, 2006; McCarthy, Minsky, Sloman, Gong, Lau, Morgenstern, Meuller, Riecken, Singh, & Singh, 2002). It has been argued that a diversity of methods will be required to achieve this (McCarthy et al., 2002). But without pushing existing methods to their limits we cannot know exactly where each will be useful and where each will fail. The policy of always applying the most promising method fails to demonstrate the boundaries of each method's abilities. My grand challenge is to test the limits of an AI methodology, experience-based reasoning (EBR), by pushing it outside of its comfort zone.

Artificial intelligence has seen a good deal of success with using EBR methods. Experiential reasoning (including analogical and case-based reasoning) is considered one of the major paradigms in the field. However, there have been few efforts to use EBR to implement *all* of the major problems of AI in within a single architecture.

The Grand Challenge. I propose that the EBR research communities create AI architectures based as much as possible on experience-based reasoning. If successful, then the AI community will have a promising new architecture. The places where the project fails will show the limitations of experience-based reasoning with a greater specificity than has yet been achieved. In either case, we contribute to the goal of designing artificial general intelligence.

In order to specify the goals of this project, it is important to describe the list of the things that a general AI architecture should be able to do. Chapter headings from some AI textbooks include methods such as search, genetic algorithms, logic, and neural networks. Since these are methodologies competing with EBR, it is unreasonable to ask that an architecture be able to do them. Indeed, the Soar architecture (Laird, Rosenbloom, and Newell, 1987) primarily uses a production system methodology. Instead, these are methodologies that implement cognitive *tasks* that intelligent creatures can do. We should look to an AI architecture to implement these higher-level tasks, such as classification, perception, problem solving, reasoning, and natural language understanding. For example, heuristic search can be used to implement problem-solving, so an architecture based on EBR should be able to implement problem-solving as well, but with its unique approach.

There is no agreed-upon exhaustive list of cognitive tasks, but most lists would contain some versions of the following items: attention, theory creation, conceptual

Davies, J. (2009). Experience-based reasoning as the basis of a general artificial intelligence architecture. *IJCAI workshop on Grand Challenges For Reasoning From Experiences*. July 11, Pasadena, California.

change, inference and reasoning, categorization and classification, decision making, emotional response, hypothetical thinking, language processing, memory retrieval and storage, problem-solving and planning, perception, and learning.

In this paper I will sample some of these cognitive tasks, and describe sketches of how they could be implemented with an EBR approach.

Experience-based reasoning. The goal of an EBR architecture is to use analogy and case-based reasoning methods to implement the major tasks involved in intelligent thought. In this section I will review what EBR methods are, in the broadest sense.

The core idea is to apply knowledge from past experience to inform new experience. Researchers have identified several steps that are typically involved (Leake, 1999), although there is not universal agreement on what the steps are, nor how many there are. Not all of these sub-processes will prove necessary for the implementation of a general AI architecture. In the next section I describe how the tasks might be implemented, and will refer to these steps.

The first step is *situation assessment*, which is preparing the current experience so that it works with the indexing scheme used for the past experiences. This facilitates retrieval of appropriate experiences.

The second step is *retrieval*, which is the choosing from memory which experiences to consider and use for reasoning. The selected experience is referred to in the case-based reasoning (CBR) and analogy literature as the "source" or "base."

The *mapping* step is the selection of correspondences between the elements of the source and the current situation (the target). For example, in an analogy between a car and a horse, the wheels of the car might map to the legs of the horse.

Similarity assessment is determining the relevant similarities and differences between the source and target. When referring to a man as a bear, we know to infer that the individual in question is big, and also know to not infer that he slaps salmon out of rivers.

Transfer and adaptation is the process of applying the knowledge of the source to the target (transfer), which might involve changing the nature of the knowledge (adaptation).

Evaluation is the determination of how well the transferred knowledge applies to the target. If the transferred knowledge is a solution to a problem, the evaluation determines if the solution is adequate.

Finally, *storage* is the process of encoding the new, completed experience in memory so it can be used as a source by reasoning processes.

EBR uses past experiences to guide actions in new ones. When a reasoner is first starting out, with no experiences in memory, there is nothing to retrieve. We can call this the boostrapping problem.

In general this can be overcome in two ways. First, the reasoner can act randomly, and remember the resulting experience in terms of the extent to which the reasoner's goals were achieved. These random acts and consequences form the initial memories that EBR uses. This is a kind of learning by doing. For many low-level AI tasks, I will suggest some version of this "random acts" boostrapping method.

Alternately, a reasoner can hear about or observe other agents acting (this is not only a bootstrapping process, but also another means to obtain memories in a working system). In many CBR systems, the initial cases are provided by the user from some database, perhaps created by knowledge engineers. In this context I classify this ability to obtain virtual experiences as "learning by observation" rather than learning by doing.

Now I will describe several core cognitive tasks, and sketch how an EBR system might implement them.

How EBR Could Implement Several Core Cognitive Tasks

There are some areas of cognition that have received a good deal of attention in the EBR literature, so I will not elaborate on them here, such as planning (Veloso & Carbonell, 1993; Hammond, 1989), language processing (Burke, 1998)[1] and problem solving, which is so common in CBR research that it a part of CBR's very definition.[2]

Attention. Attention is the process a reasoner uses to decide which internal representations or aspects of the environment to concentrate on. In this paper "concentrating on" something means to spend resources processing it. The EBR architecture will handle attention shifts from one thing to another based on analogies with previous attentional shifts. The process will be bootstrapped with random shifts of attention. As the reasoner is rewarded and punished for effective and ineffective attention shift decisions, those shifts are stored as cases to be retrieved and applied to new situations.

As the reasoner learns abstract symbols, it will be able to transfer attentional shifts from one thing to another based on more abstract similarities between cases.

To make this more concrete, I will focus on the shifting of visual attention. Visual attentional shifts take two forms: motion of the eyes, in which the reasoner moves the eye or sensor to be pointed at what they want to process, and motion of the "attentional window" (Kosslyn, 1994) in

[1] For a bibliography see http://www.lrdc.pitt.edu/Ashley/TCBR_BIB.html
[2] http://en.wikipedia.org/wiki/Case-based_reasoning retrieved March 13, 2009 from the Wikipedia.

which the reasoner makes changes in what part of the input is to be processed. That means that while the eyes (or cameras) remain steady, the reasoner can attend to different parts of the visual signal. What gets transferred are eye/camera movements as well as attentional window movements associated with memories similar to the current sensory state.

These shifts get associated with success or failure depending on feedback regarding what *should* have been attended to. For example, if the reasoner finds out later that there was something important that should have been attended to, the architecture can punish the attentional shift that was used as a source, making it less likely to be retrieved in the future.

Decision Making. Previous decision-making episodes can be used to inform new decisions. Similarly to the categorization example below, the decision-making strategies can be encoded in the cases. Since the term "decision-making" can potentially apply to many other psychological processes (such as classification and planning), I will restrict the meaning in this paper to the selection of a course of action. When viewed this way, it is easy to conceive of cases in which there is a situation description and an action taken, rather like a production rule. In addition to the methods in the CBR literature, the cognitive modeling architecture ACT-R (Anderson & Lebiere, 1998) has a means for selecting productions based on success rates. Once selected, CBR methods can be used to transfer and adapt the cases to the new situation.

Theory and Explanation Generation. This includes both scientific and more everyday generation of explanations. Theory creation is a complex process with many subtasks. I will address a few of them.

Often the creation of a new theory involves the creation of a mental model of some system in the world to serve as an explanation. To do this with EBR, the reasoner would transfer a mental model of how something works from memory to the novel target phenomenon to be understood. For example, a reasoner might try to understand how a film set functions based on its experience with a theater company.

A good deal of work has been done in this area already. Model creation has been explored in Davies, Nersessian, & Goel (2001) for the case of James Clerk Maxwell's creation of the theory of electromagnetism, based on an analogy with gear systems.

Sometimes a new situation requiring explanation will have similarities to other situations in memory that have explanations associated with them. Those similar cases are retrieved, and the explanations transferred to the target situation. For example, if a reasoner observes violence, the reasoner can transfer an explanation of why (e.g., anger on the part of the one engaging in the violence) from a previous experience. The Interactive KRITIK system (Goel, Gomez de Silva Garza, Grue, Murdock, & Recker, 1997) and PHINEAS (Falkenhainer, 1990) generated explanations using EBR for physical devices, and the SWALE system (Leake, 1992) provided creative explanations for anomalous situations, and (HaCohen-Kerner, 1995) built a system for finding explanations in Chess.

The creation of equations is particularly important for scientific theory creation. When a theory is turning from a mental model into an equation, analogy can help determine the parts of the equation. For example, if a number has to be positive, there could be cases in the reasoner's mind in which a value is squared, which always results in a positive number. This could be transferred to the target situation.

Memory. The cognitive tasks involving memory are regarded as some of the most fundamental in cognitive psychology. Memories are stored, transferred to long-term memory from the various short-term memories, altered, and retrieved.

Case-based reasoning already has theories of memory, in its own way: memories are cases, indexed so that they can be found efficiently and effectively (e.g., Kolodner, 1983; Hammond, 1989; Birnbaum, 1985; Schank, 1999). This has proven effective for the AI community, and might be enough for an AI architecture. If EBR is to be a contender as a *cognitive* architecture, however, the theory is not at the level that most psychologists tend to be interested in. Rather than focusing on the use and structure of complex memories, psychologists tend to focus on things such as memorization, the relation of memory to attention, forgetting, and the different kinds of memory storage (e.g., working, long-term, the visuo-spatial sketchpad, etc.). Their data are based on millisecond response times and accuracy measures, rather than broad questions of knowledge transfer that CBR is typically concerned with. So what would an EBR theory of memory look like at this lower-level?

I will discuss one of the many memory-related tasks: memory retrieval. There is wide agreement that memory retrieval occurs based on some notion of semantic similarity and something approximating spreading activation. The process could bootstrap by retrieving based on random features of the query stimuli. If appropriate memories are returned, then that method is more likely to be used when future memory-retrieval episodes occur. For example, if one sees a bear on a sunny day, retrieving other memories of sunny days will prove less useful than retrieving other instances of seeing bears. What the reasoner would be doing is retrieval and adaptation of past memory retrievals to inform new ones.

There are many ways that spreading activation can occur, and many decisions a modeler must make when implementing it. For example, what is the multiplier used when one activated node spreads to an adjacent? Are certain kinds of edges in the network more likely to spread in certain conditions? All of these decisions could be tried, stored as cases, and retrieved with an EBR system trying to spread activation.

EBR theory was not intended to model behavior at this low of a level. Perhaps using EBR to determine the details of spreading activation is pushing the idea too far. We won't know for sure until we try.

Conceptual Change. We can imagine two parts of conceptual change: the creation of a new conception, and the transfer of preference from the old to the new conception. While the reasoner holds the previous conception, a new conception is being learned (or created) and slowly understood. At some point the new conception might be judged as more sensible than the old, and a conceptual change occurs, transferring preference from the old to the new. Even when conceptual change occurs, the old theory is still more or less intact in the reasoner's mind.

The learning of a new conception involves understanding a mental model that someone else has already figured out for themselves. It is trying to understand someone else's communicated mental model (Norman, 1983). To some extent, when we learn something complicated we must reinvent it for ourselves. To this extent the process is identical to that described in the theory generation section. The rest is creating a mental model out of the communicated symbol structures. When told that, for example, an electron is like a planet in a solar system, the properties of the solar system are transferred to the electron target (Gentner & Schumacher, 1986).

The transfer of acceptance from the old to the new conception can be viewed as a decision-making task as described above. The reasoner evaluates the two choices: to transfer acceptance now, or not. As described in the decision making section, the decision is made based on the results of similarly-made decisions in memory. In this case, the important aspects of the predicted outcomes are used as queries into memory. Previously made decisions are retrieved, and, based on how they turned out, one is selected to be used to determine the outcome of this decision.

Inference and Reasoning. Many case-based reasoning systems do some kind of inference in the form of knowledge transfer. However more classical forms of reasoning strategies can also be implemented. Reasoners experience cases of others' reasoning, through reading and listening. Also, they can create their own reasoning methods, with varying degrees of success. Individual reasoning cases (I'll call them "arguments") are retrieved, modified, and re-used.

For example, a reasoner might hear an example of Modus Ponens: "If I have the money I will buy an ice cream. I have the money, so I'm going to buy an ice cream." One of the various means discovered to abstract cases might generate an abstract case resembling "if P then Q, P, therefore Q." Then, when encountering a new case of the form "if P then Q, P" the reasoner can retrieve and use the previous case to conclude Q. Other argument forms can also be implemented in this way.

Categorization and Classification. In this paper I will refer to "classification" as the placement of something into a category, and "categorization" as the creation of categories.

A memory of classification events can be built up through observation of others' and one's own experiences. For example, a reasoner might observe someone classify a brown, furry, barking thing as a dog. Alternatively, a reasoner might (randomly or through CBR) classify a cow as a dog. Subsequent action or inference can provide the feedback for these cases. When a child mis-classifies a cow as "doggie," an adult might correct her, making that particular case of classification less likely to be retrieved in the future.

Categorization is more complicated. How can a reasoner use previous examples of creating categories to inform the creation of new ones? Categories are useful in that through classification we can predict unobserved properties and attribute values. Certain computational processes such as factor analysis and principle components analysis find categories automatically according to correlation of attribute values. It is tempting to make such a classification strategy "innate" to the architecture, rather than having it learned, but if we're pushing experience-based reasoning to its limits we need to try to make the strategy of "using correlation as a basis for categorization" a part of cases that can be retrieved.

There are other ways to categorize. For example, when people freely sort novel stimuli into categories of their own choosing, they often select a single attribute (or dimension) and sort according to the values (this is called a 1D sort in the psychology literature. See Ahn & Medin, 1992). This is a robust finding and it has proven rather difficult to show that people attend to correlational structure at all in category creation (Billman & Davies, 2005).

These are two of the potentially many strategies that could be used to create categories. The success of the inferences made as a result of these strategies would provide the reasoner's feedback. Then, when encountering a new situation where category creation is required, cases of the uses of different strategies can be evaluated and retrieved.

Emotional Response. Similarly to decision making, emotional responses can be viewed as actions in response to certain stimuli. Since emotions tend to ready other processes for action, these "emotional productions" could be evaluated as decisions as well.

Right now the modeling of emotions is more important for cognitive models of humans than for general AI architectures. In human psychology, emotions are evolutionarily old and most of them are run in the System 1 part of the brain (contextual, innate, fast, modular, etc., see Stanovich & West, 2003), and we might not want to claim that System 1 works by experience-based reasoning.

Nevertheless, representing emotional response this way might be a useful AI strategy.

Conclusion

The idea of EBR as the basis of cognition is not new. Hofstadter (2001) suggested that analogy is the core of cognition. However his paper focuses on high-level cognition, to which analogy has traditionally been successfully applied. I am suggesting taking on the difficult task of modeling lower-level cognitive tasks, the architecture-level tasks, with EBR as well.

Forbus and Gentner (1997) have suggested that the breadth of human commonsense reasoning and learning takes the form of analogical reasoning and learning from experience. The Companion Cognitive System project (Forbus & Hinrichs, 2004; Forbus, Klenk, & Hinrichs, 2008) is also working with the hypothesis that most learning and reasoning can be handled with analogical reasoning. I applaud their efforts; theirs is the most sophisticated EBR architecture in existence.

However even the Companion architecture breaks from EBR for some tasks, such as logical reasoning. Their papers do not describe exactly why EBR cannot be used for logical reasoning. If it cannot, this should be expressed as an explicit AI finding. Architectures such as the Companion architecture could produce findings about newly-found limits and capabilities of EBR in the context of a general intelligence.

Nobody has yet tried to consistently apply EBR to cognition at lower levels, and few have even tried to apply it consistently at higher levels, that is, with the same analogical or case-based mechanism. Attempting to do this is important for two reasons, the first scientific and the second practical. First, though EBR has been shown to be important in many areas, its usefulness has not been systematically pushed to the breaking point. There might be task-related breaking points, by which I mean there might be classes of tasks it cannot account for, such as, perhaps, memory retrieval. There also might be level-related breaking points. For example, when EBR is the basis of memory retrieval, it might break down.

Second, such architectures have the potential to raise EBR's profile. Most of cognitive psychology is the study of low-level tasks like memory and attention. EBR as a mechanism will not interest most cognitive psychologists unless it addresses the phenomena they find important. Indeed, even in case-based reasoning's home sub-discipline of artificial intelligence it receives little attention. Some contemporary AI textbooks mention it only in passing.[3]

Experience-based reasoning has the potential to be the basis of general artificial intelligence architectures. Creating such architectures will force us to attempt to model cognitive tasks not traditionally implemented with EBR. These efforts will prove to be valuable contributions to AI, whether EBR succeeds or not. Given the growing interest in artificial general intelligence, the exploration of experience-based AI architectures is our grandest challenge.

References

Ahn, W. and Medin, D. L. 1992. A two-stage model of category construction. *Cognitive Science*, 16, 81--121.

Anderson, J. R., and Lebiere, C. 1998. *The Atomic Components of Thought*. Lawrence Erlbaum Associates, Mahwah, New Jersey.

Billman, D. O. and Davies, J. 2005. Consistent contrast and correlation in free sorting. *American Journal of Psychology*. 118 (3), pp. 353--383.

Burke, R. 1998. Surveying Opportunities for Textual CBR. *Papers from the AAAI-98 Workshop on Textual Case-Based Reasoning*. p. 13.

Birnbaum, L. 1985. A Short Note on Opportunistic Planning and Memory in Arguments. *Proceedings of the International Joint Conference on Artificial Intelligence*, 281-283.

Davies, J. R., Nersessian, N. J. and Goel, A. K. 2001. The role of visual analogy in scientific discovery. *Model-Based Reasoning: Scientific Discovery, Technological Innovation, Values*. Pavia, Italy.

Falkenhainer, B. 1990. A unified approach to explanation and theory formation. In J. Shrager and P. Langley (Eds.), *Computational Models of Scientific Discovery and Theory Formation* pp157-196: San Mateo, CA: Kaufmann.

Forbus, K. and Gentner, D. 1997. Qualitative mental models: Simulations or memories? *Proceedings of QR-97*, Cortona, Italy.

Forbus, K. and Hinrichs, T. (2004). Companion Cognitive Systems: A step towards human-level AI. *AAAI Fall Symposium on Achieving Human-level Intelligence through Integrated Systems and Research*, October, Washington, DC.

[3] Russel and Norvig, who wrote the currently most popular AI textbook (2003), mention it on a single page (708) in a "historical notes" section. Negnetitsky (2005) doesn't mention it at all.

Forbus, K., Klenk, M. and Hinrichs, T. *2008*. Companion Cognitive Systems: Design Goals and Some Lessons Learned. *AAAI Fall Symposium on Naturally Inspired Artificial Intelligence*.

Gentner, D., and Schumacher, R. M. 1986. Use of structure-mapping theory for complex systems. Presented at the *Panel on Mental Models and Complex Systems, IEEE International Conference on Systems, Man and Cybernetics*.

Goel, A. K., Gomez de Silva Garza, A., Grue, N., Murdock, J. W., and Recker, M. M. 1997, Functional Explanations in Design. *Program of the IJCAI'97 Workshop on Modeling and Reasoning about Function*. Nagoya, Japan, August 25.

Hammond, K. J. 1989. *Case-based planning: viewing planning as a memory task*: Academic Press Professional, Inc.

Hofstadter, D. R. 2001. Analogy as the Core of Cognition. in *The Analogical Mind: Perspectives from Cognitive Science*, edited by Dedre Gentner, Keith J. Holyoak, and Boicho N. Kokinov:MIT Press.

Kolodner, J. L. 1983. Constructive memory: A computer model. *Cognitive Science*, 7:4, 281-328.

Kosslyn, S. M., 1994. *Image and Brain*: MIT Press.

Laird, J., Rosenbloom, P. and Newell, A. 1987. Soar: An Architecture for General Intelligence. *Artificial Intelligence*, 33: 1-64.

Leake, D. B. 1992. *Evaluating Explanations: A Content Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Leake, D. B., 1999. Case-based Reasoning. in W. Bechtel and G. Graham (Eds.) *A Companion to Cognitive Science*: Blackwell.

McCarthy, J., Minsky, M., Sloman, A., Gong, L., Lau, T., Morgenstern, L., Meuller, E. T., Riecken, D., Singh, M. & Singh, P. 2002. An architecture of diversity for commonsense reasoning. *IBM Systems Journal*, 41(3):530-539.

Negnevitsky, M. 2005. *Artificial Intelligence: A Guide to Intelligent Systems:* Addison Wesley.

Norman, D. A. 1983. Some Observations on mental models. In D. Gentner and A. Stevens (Eds.), *Mental Models*: L. Erlbaum Associates.

Russel, S. and Norvig, P. 2003. *Artificial Intelligence: A Modern Approach, Second Edition:* Prentice Hall.

Schank, R. C. 1999. *Dynamic Memory Revisited*: Cambridge University Press, 2nd Edition.

Stanovich, K. E., and West, R. F. 2003. The rationality debate as a progressive research program. *Behavioral and Brain Sciences*, *26*, 531-534.

Veloso, M.M. and Carbonell, J.G. 1993. Derivational Analogy in PRODIGY: Automating Case Acquisition, Storage, and Utilization. *Machine Learning,* 10(3):249-278.

Voss, P. , Goertzel, B., and Pennachin, C., Eds., 2006. *Essentials of general intelligence Artificial General Intelligence*:Springer

HaCohen-Kerner, Y. 1995. Learning Strategies for Explanation Patterns: Basic Game Patterns with Application to Chess. *Proceedings of the International Conference on Case-Based Reasoning.* 491-500