

A Self-Help Guide for Autonomous Systems

Michael L. Anderson, Scott Fults, Darsana P. Josyula,
Tim Oates, Don Perlis, Matt Schmill,
Shomir Wilson, and Dean Wright

■ *Humans learn from their mistakes. When things go badly, we notice that something is amiss, figure out what went wrong and why, and attempt to repair the problem. Artificial systems depend on their human designers to program in responses to every eventuality and therefore typically don't even notice when things go wrong, following their programming over the proverbial, and in some cases literal, cliff. This article describes our past and current work on the metacognitive loop, a domain-general approach to giving artificial systems the ability to notice, assess, and repair problems. The goal is to make artificial systems more robust and less dependent on their human designers.*

Things almost never go exactly as we plan or hope; deviations from the ideal are a fact of life. As Heraclitus noted almost 2500 years ago, “only change endures.” Though Heraclitus was silent on whether the inevitable change will be good or bad, Murphy’s Law tells us to expect the worst. Given the pervasive and often perverse nature of the unexpected, it is not surprising that humans are adept at dealing with it—we would not have survived and flourished otherwise. This facility with the unexpected may account for our annoyance when our automated systems break in situations that we could deal with easily. It is then we need to remember that such failures account for much of the progress in AI, since—as this special issue attests—these annoying surprises goad us into insights and thus to improvements in the underlying algorithms.

But one can also view this phenomenon in another light. Instead of our having to rebuild the algorithms when—in an unanticipated situation—they make “stupid” mistakes that we could easily have dealt with, why not build them to do what we do: deal with the mistakes on their own? This is a tall order, and not a new one: it requires addressing the so-called brittleness problem (see Brachman [2006]). In this article, we present our own approach to this problem. First we describe the conceptual focus—which we call the *metacognitive loop* (MCL)—that is guiding our

work; then we present a number of domain-specific implementations of MCL, which we think provide significant evidence that MCL is a general-purpose methodology for building flexible (nonbrittle) programs in specific domains; and finally we discuss our current efforts toward a domain-independent implementation of MCL.

MCL Systems

Our starting point is to ask what mechanisms humans use when dealing with the unexpected. Our contention is that people learn what to expect, notice when an expectation fails, and then decide what to do about it.¹ More specifically humans tend to (and machines ought to):

N: Note when an anomaly occurs

A: Assess it

G: Guide a response into place

Looping through this NAG cycle—the basic form of MCL—leads to remarkably positive results when people do it, from noting that a successful behavior in one domain (for example, swinging a cricket bat) utterly fails in another (baseball) and is best jettisoned and a new one learned from scratch rather than tweaking the old, to noting and correcting a misunderstanding in conversation, to trying a new approach to a problem (for example, on the final exam in light of failures on the midterm). Moreover, asking for help or

giving up are in fact among the most important things in the human repertoire of repairs, for, as the saying goes, if you find yourself in a hole, the first thing to do is stop digging. Our work suggests that MCL can at the very least give machines the same simple wisdom that would have been quite useful to the DARPA Grand Challenge robot that continued to push against a fence it could not see, burning out its motors; or to the NASA satellite that, having obeyed instructions to turn away to point at another part of the sky, was thus rendered unable to receive further instructions and so stayed in that useless position until natural orbital forces reoriented it.

Noticing that there is a difference between what actually happened and what was expected to happen—that something has gone wrong—is the key step required to get MCL started. So, central to our approach is building systems that don't simply do things but *know* what they are supposed to be achieving. That is, they should have specific expectations for their own performance and the outcomes of their actions. Knowing may be half the battle, but it is only half; the next steps are also crucial. Rather than simply recognize problems, report them, and wait for human system designers to fix them, MCL systems are also self-diagnosing and self-repairing. Repairs can be as simple as trying a failed action again; or trying a new plan for the same goal, dynamically learning new operators; or using trial and error; or taking advice from a human; or embarking on self-training; and so on.

It is perhaps this aspect of our approach that is most distinct from the relevant prior work on which we are building, including case-based reasoning in general and introspective reasoning and learning from expectation failure in particular (Fox and Leake 1995, Gil 1994, Hammond 1986, Schank and Leake 1989). For instance, MCL systems don't just use expectation failures to trigger a single learning strategy, such as the refinement of planning operators, but can choose this option from among several; nor are MCL systems limited to selecting among multiple learning strategies (as with Cox and Ram [1999]), but can implement any of multiple disparate options for addressing anomalies. In addition, MCL systems don't just wait for a response plan to complete and then remeasure performance on the original task; rather they monitor the response just as they would any other action taken by the system and can intervene if the repair itself appears to be failing. Finally, the MCL systems we envision don't just work in a single domain with specialized knowledge but across many different domains, using ontologies capturing general coping strategies.

Imagining such systems is quite different from building them. To test the viability and domain generality of this approach, we have added MCL

components to a number of otherwise self-ignorant systems. This work has convinced us that MCL is indeed a powerful approach for endowing automated systems with the flexibility needed to deal with surprise. While these implementations are domain-dependent and differ significantly in various respects, they share some common features that serve to characterize the key aspects of MCL.²

We can briefly indicate some of these common features. Noting an anomaly amounts to noting a mismatch between an expectation $\text{Exp}(E)$ and an observed outcome $\text{Obs}(\neg E)$. This is the basis for the Note stage in all our implementations and domains.³ The Assess stage identifies the anomalous E and $\neg E$ (in some current context C) as being of some type T : $\text{Type}(E, \neg E, C, T)$. Here T might turn out to be highly domain specific (for example, sensor error, if the domain and context are sufficient to determine that) or very general (for example, simple logical contradiction). In addition, the Assess stage has, for each anomaly-type T , a prioritized list of possible responses. Finally, the Guide stage enacts responses from that list; if a given response fails, another is selected. Response failure can also lead to recursion if the NAG cycle is triggered by an anomaly in the attempt to guide a response. There is the danger here of nonterminating recursion, with repairs applied to failed repairs ad infinitum. However, as these repairs accumulate, the expected time to completion of the original task will move farther into the future, forcing MCL ultimately to fall back on a response such as asking for help (which may be costly but powerful) or simply giving up and moving on to something else. One final point: expectations themselves can be learned, as well as modified, by experience, and this is one of the most powerful aspects of MCL. Thus, not only is a particular instance of an anomaly dealt with by an MCL-endowed system, but the system's future expectations may change as a result. MCL can deal with learned expectations that are unreliable by either gathering more data to further refine them or by simply abandoning them altogether.

We have implemented these ideas in highly varied domains, including natural language human-computer dialogue, robot navigation, nonmonotonic reasoning, reinforcement learning, and a tank game. Three of these systems are described below, after which we will describe our efforts to implement domain-general MCL.

MCL-Enhanced Reinforcement Learning

Reinforcement learning (RL) is an established methodology that works very well in many settings, notably ones in which the reward structure is static or nearly static. But when that structure is changed suddenly and significantly, the performance of RL degrades severely and recovers excruciatingly slowly.

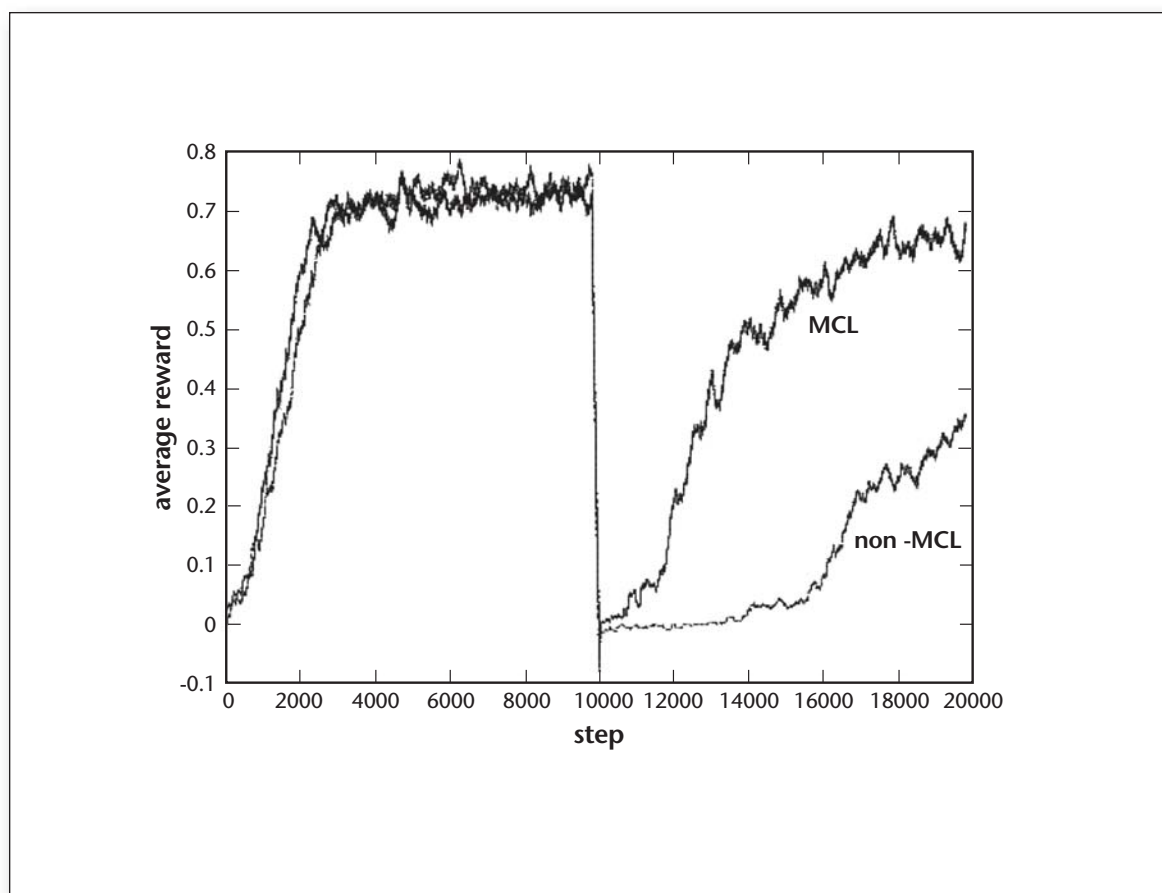


Figure 1: Learning Curves for a Standard Q-Learner Versus an MCL-Enhanced Q-Learner in the Face of Changes to the Reward Structure.

atingly slowly. In essence, RL algorithms need to “unlearn” what they have learned, step by step, since they have no way to recognize that the reward structure has changed, let alone assess what can be done about it. Yet it is clear that, given a drastic change that makes previous learning useless, the best policy is simply to throw it out and start over.

Using a variety of reinforcement learning algorithms (Q-learning, SARSA, and prioritized sweeping) we experimented with a simple 8 by 8 grid world with rewards in cells (1, 1) and (8, 8). The learner was trained for 10,000 steps, then the rewards were switched, and learning continued for another 10,000 steps. We compared the performance of standard RL algorithms to MCL-enhanced versions of the same algorithms. The MCL-enhanced RL algorithms maintained and monitored expectations about such things as average reward per step, value of future rewards, and average time to next reward. When these expectations were violated, the enhanced algorithms assessed

the nature of the violation and, using a simple decision tree, chose one of the available repairs (which included ignoring the problem, adjusting the learning parameter, or throwing out the current action policy and starting over).

A typical result is shown in figure 1. The horizontal axis is the step number, and the vertical axis is the average reward. Performance rises sharply and levels off until step 10,000 when the reward-switching occurs. At that point, performance falls dramatically and then begins to recover. However, the standard RL algorithms (in this case, Q-learning, seen as the lower curve) recover far more slowly and far less completely than the MCL-enhanced versions (the higher curve) of the same algorithms. Though given enough experience both variants will asymptote to the same point, in our experiments we found that the greater the degree of change in reward (such as swapping rewards for penalties, and vice versa), the greater and longer lasting were the transient benefits of MCL (Anderson et al. 2006).



Figure 2: A Bolo Tank Approaching a Neutral Pillbox beyond the Trees with a Refueling Station behind the Walls to the Left of the Tank.

MCL-Enhanced Human-Computer Dialogue

Another application area for MCL is natural language human-computer dialogue. Natural language is complex and ambiguous, and therefore, communication always contains an element of uncertainty. To manage this uncertainty, human dialogue partners continually monitor the conversation, their own comprehension, and the apparent comprehension of their interlocutor. Human partners elicit and provide feedback as the conversation continues and make conversational adjustments as necessary. We contend that the ability to engage in this metadialogue is the source of much of the flexibility displayed by humans when they engage in conversation (Perlis, Purang, and Andersen 1998). We have demonstrated that enhancing existing dialogue systems with a version of MCL that allows for metadialogic exchanges improves performance.

In one specific case tested, a user of TRAINS-96

(Allen et al. 1996), a simulation of a national train network that is controlled by natural language tells the system to “Send the Boston train to New York.” If there is more than one train in Boston, the system may well choose the wrong one to send—the user may have in mind the train that runs regularly between Boston and New York and so might respond: “No, send the *Boston* train to New York!” Whereas the original TRAINS-96 dialogue system responds to this apparently contradictory sequence of commands (Send, Don’t send, Send) by once again sending the very same train, our MCL-enhanced version of TRAINS notes the anomaly (that is, the contradiction in commands) and, by assessing the problem, identifies a possible explanation in its choice of referent for “the Boston train.” The enhanced system then chooses a different train the second time around, or if there are no other trains in Boston, it will ask the user to specify the train by name. The details of the implementation, as well as a specific account of the rea-

soning required for each of these steps, can be found in Traum et al. (1999).

More recently we have built another dialogue system, ALFRED,⁴ that uses the MCL approach to resolve a broader class of dialogue anomalies. The system establishes and monitors a set of dialogue expectations related to time, content, and feedback. For example, if the user says “Send the Metro to Boston,” ALFRED notices that it doesn’t know the word *Metro* (a failure of the expectation that it will find input words in its dictionary). Alfred’s first response is to try to determine what it can about the unknown word. Since Alfred knows the command “send” and its possible arguments, it is able to determine that *Metro* is a train. If it cannot determine from this which train the user is referring to, it will request specific help from the user, saying: “Which train is *Metro*?” Once the user tells the system that *Metro* is another word for *Metroliner*, it is able to correctly implement the user’s request (Josyula 2005).

A Bolo Player

Bolo is a multiplayer tank game that takes place in a world that contains various terrain types (roads, swamps, walls, and so on), refueling bases, and pillboxes. There are three types of pillboxes: neutral pillboxes fire on all tanks, friendly pillboxes fire only on other players’ tanks, and dead pillboxes pose no threat and can be captured to make them friendly. An important strategy in Bolo is to capture pillboxes, make them friendly, and then use them either offensively or defensively. Figure 2 shows a Bolo tank approaching a neutral pillbox.

Bolo can be played by humans, but it can also be played by programs. Such artificial Bolo players tend to play quite poorly and are easily fooled when unexpected complications arise (change of terrain, more dangerous pillboxes, and so on). Thus Bolo provides a good challenge domain in which to test MCL.

Our MCL-enhanced Bolo player is controlled by a simple hierarchical task network (HTN) planner with primitive actions that ground out in controllers. It maintains a variety of expectations, the primary one being that the tank it controls will not be destroyed.⁵ The initial HTN allowed the player to locate and capture dead pillboxes. However, the player did not have a plan to deal with hostile pillboxes, which fire on the tank, and so it was destroyed in its first such encounter. At this stage—when an expectation fails—MCL has a suite of actions to choose from, including means-ends analysis and operator refinement (Gil 1994, Wang 1995). In one scenario, our MCL-enhanced Bolo player was able to discover that firing on pillboxes offered a solution to the problem, even though it had no previous knowledge of the effect of that particular action. More precisely, the MCL compo-

nent searched through its past experience to try to locate salient differences in the conditions under which it succeeded in taking pillboxes, and those in which it failed. It found that only pillboxes with intact armor destroyed the tank, so the next step was to see if it had any actions that could reduce the armor of a pillbox. If it had known about an action that would do that, it would have tried the action immediately. In the case we tested, it had no such knowledge. Thus, it used a heuristic to rank all its actions according to how likely they were to have the desired effect, and then tested them until it found one that worked. Note how MCL turns failure into opportunity: in each case the system learned more about what effects its actions did and did not have, and in a way organized to support its ongoing mission.

Toward a Single Domain-Independent MCL Implementation

Humans are good at dealing with surprise (employing the NAG cycle) not only in a few specific domains, but generally. This, together with the successes we have seen with executing MCL in highly varied domains, suggests to us that a single, generalized MCL implementation might also be able to negotiate surprise across domains.

The idea of a domain-general MCL is this: there can be a general-purpose MCL-based implementation such that any given automated host system *S* that (as is typical) exhibits considerable brittleness can be interfaced with that implementation, allowing the latter to monitor and (when needed) guide *S* into significantly improved performance with respect to unanticipated situations, without the usual human recoding.

To achieve this goal, we are now building such a general-purpose MCL that encodes domain-general knowledge about anomalies in three ontologies, as shown in figure 3. There is one ontology for each phase of the NAG cycle.⁶

The Note phase of the NAG cycle uses an ontology of *indications*, which are sensory or contextual cues that something is wrong. Within this ontology, host-specific nodes at the lowest levels encode expectations about the behavior of sensor, state, and other values. Violations of these expectations activate domain-general nodes that characterize anomalies. For example, a reinforcement learning host system may maintain expectations about receiving rewards in certain states that, when violated, will trigger a domain-general indication such as MISSING DATUM.

Figure 4 shows fragments of the MCL *failure* (on the left) and *response* (on the right) ontologies, which are used by the Assess and Guide phases of the NAG cycle, respectively. In the former, dashed

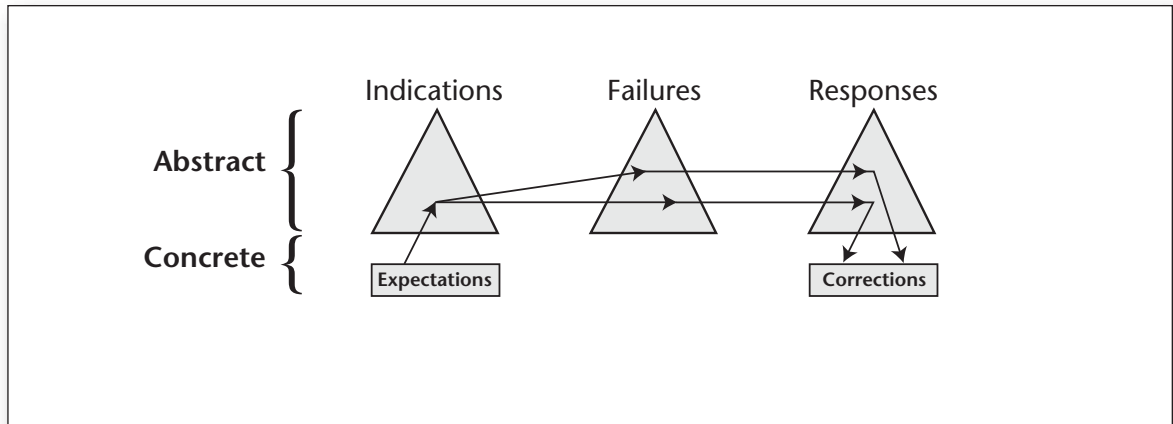


Figure 3: The Three MCL Ontologies: Indications, Failures, and Repairs.

arrows are interontological links coming from the indications ontology. That is, indications are linked to the types of failures they suggest. To continue our example of an RL host system, a MISSING DATUM could be due to a SENSOR FAILURE or a MODEL ERROR (that is, the expectation itself is wrong). These failure types can be generalized or specialized by moving up or down links within the failure ontology.

Failure nodes are linked to candidate courses of action in the response ontology, again denoted with dashed arrows. For example, if there is in fact a MODEL ERROR, and the model in question is a predictive model, one possible response is to MODIFY PREDICTIVE MODELS. This response is abstract, so it must be made concrete in terms of the host system by following downward links in the ontology. In the case of an RL host based on a Q-learning algorithm, a Q-table represents expectations about future rewards and can be adapted by increasing exploration, increasing the learning rate, or, if all else fails, resetting the Q-table and starting from scratch. The host-specific response node (the shaded nodes in the response ontology in the figure) with the highest expected utility is chosen and implemented.

Related Work

As has been noted already, the brittleness problem is not new, and autonomous systems have always needed mechanisms to deal with the unexpected. Systems as early as Shakey (Fikes 1971) were able to replan in light of failures or unmet preconditions. In recent years there has been a great deal of work in such areas as reactive computing, machine learning, planning and replanning, and fault detection, isolation, and recovery (FDIR); see Mitchell (1990); Williams and Nayak (1996); van

der Krogt, de Weerd, and Witteveen (2003); and Stroulia and Goel (1997) for just a brief sampling, and [2] for a general overview of metacognitive approaches to reasoning and learning.

One thing that is distinctive about our approach is its insistence on the general applicability of the NAG cycle, along with an architecture that encapsulates the metacognitive components, engaging them only when an expectation has been violated. This has advantages over some approaches—for instance, many reactive/adaptive systems (such as reinforcement learners) deal with change by continually training. This requires continual exploration, or deviation from the optimal action policy, whereas MCL systems can act optimally until they notice that something is wrong and then take remedial actions focused on the problem at hand.

For instance, although FDIR is clearly motivated by the same concerns as MCL—and has the same tripartite structure as the NAG cycle: (FD), (I), (R)—it typically takes an expert-systems or model-based approach to the issue that imposes significant limitations on the types and range of available diagnoses and repairs, often limiting the options to specific solutions known in advance to fix specific problems. By including the possibility of generating novel hypotheses about the causes of its performance and learning new models, operators, or action policies to deal with failures, MCL greatly increases the range of possible responses the system could implement. This is in line with our general preference (which we admit may not be appropriate for every domain) for increasing the agency and autonomy of our systems, and especially their freedom of action in responding to problems, rather than limiting it and hoping that circumstances do not stray from the anticipations of the system designer. Again, this increased freedom is highly focused, coming into play only when

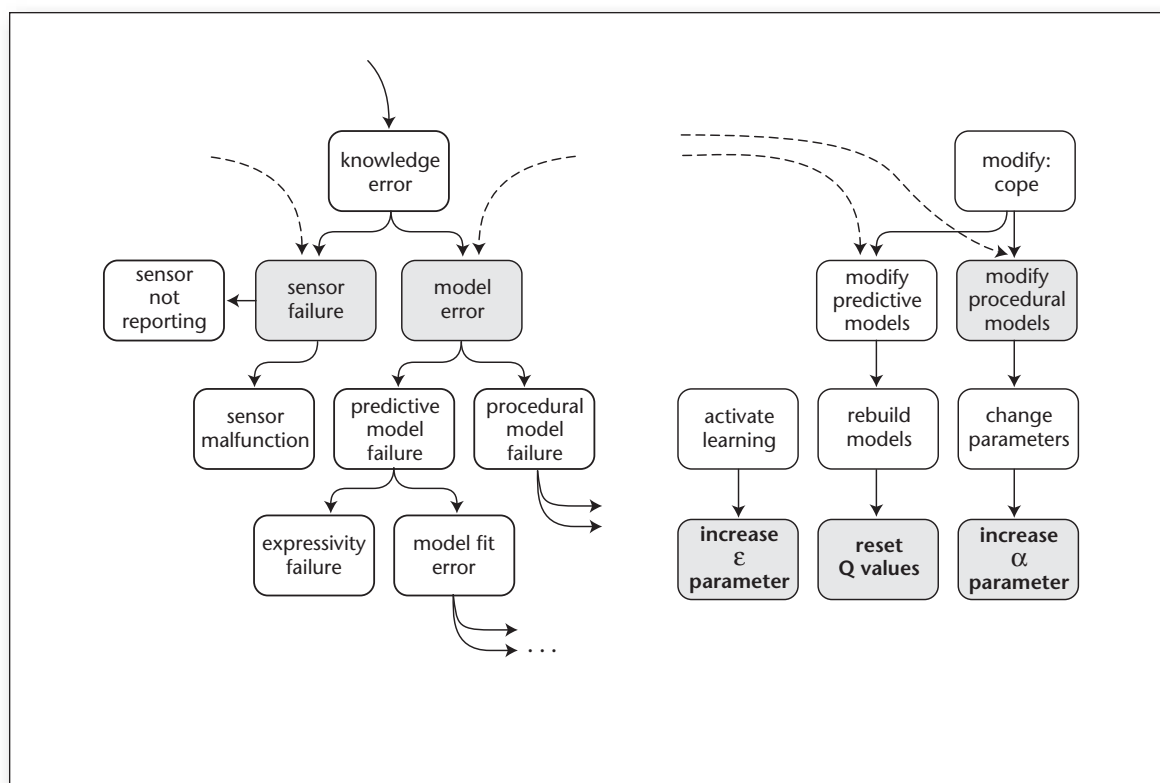


Figure 4: Fragments of MCL's Failure and Response Ontologies Used by the Assess and Guide Phases of the NAG Cycle, Respectively.

anomalies are noted that directly affect the performance of the system, about which expectations are maintained and monitored. MCL systems are self-aware, self-guided learners, able to decide whether, what, when, and how to learn; this combination of reasoning and learning, in which each can guide the other under the direction of the system itself, is one of the distinctive elements of the MCL approach.

Conclusion

Humans confront the unexpected every day, deal with it, and often learn from it. AI agents, on the other hand, are typically brittle—they tend to break down as soon as something happens that their creators didn't plan for. Brittleness may be the single most important and most difficult problem facing AI research. It is difficult because it is impossible to prepare an AI system for every possible contingency. It is important not only because brittle systems become unproductive when they fail, but also because if the goal is human-level intelligence, then systems must exhibit humanlike flexibility.

We believe that it is possible to replicate human-

like behavior by furnishing AI systems with one of the methods that humans use to deal with the unexpected, namely the Note-Assess-Guide cycle. In order to do this, we must enable AI systems to help themselves; they must establish expectations and monitor them, note failed expectations, assess their causes, and then choose appropriate responses.

We have, in fact, implemented this method in several distinct types of systems: a reinforcement learner, a human-computer dialogue agent, a tank game, a robot navigation system, and a common-sense (nonmonotonic) reasoner. In each case, the performance of the system was enhanced by MCL mechanisms. This, we believe, is promising enough to warrant our next step: domain-general MCL. Our vision of this single, domain-independent implementation of MCL is that it can be interfaced with any brittle automated system and thereby improve its performance in unexpected situations.

Notes

1. This intuitive characterization of human problem solving is also supported by work in cognitive psychology. See Nelson et al. (1994) and Kendler and Kendler (1962, 1969) for more information.

A Too-Clever Ranking Method

Tim Oates

I was a graduate student at a time when C4.5 ruled the machine-learning roost. I developed what I thought was an extremely clever method for detecting “bad” training instances. Each instance was scored, and those with the lowest scores could be removed before running C4.5 to build a decision tree with the remainder. I ran an experiment in which I removed the bottom 10 percent of the instances in a University of California, Irvine (UCI) data set. The resulting tree was smaller and more accurate (as measured by 10-fold CV) than the tree built on the full data set. Great! Then I removed the bottom 20 percent of the instances and got a tree that was smaller than the last one and just as accurate. At that point I had the feeling that this was going to make a great paper for the International Conference on Machine Learning (ICML).

So I kept going, removing an additional 10 percent of the instances at each step, getting smaller trees with no loss in accuracy. However, when I removed 80 percent of the instances, and was still getting the same result, I realized I had a problem. There was no way that 80 percent of the instances in any of the reversed UCI data sets were “bad.” After some time I realized I should run a control condition. What happens if I remove randomly selected training instances? Shockingly, I got the same results. The more randomly selected training instances I removed, the smaller was the resulting tree, with no loss in accuracy. My extremely clever ranking method was no better than a random number generator! After getting over the initial shock, I decided, with David Jensen, to pursue this more carefully with a larger sample of data sets. We found that this phenomenon was pervasive, both with respect to data sets and decision tree pruning mechanisms. We wound up writing papers on this topic that were published at the ICML, AAAI, and Knowledge Discovery in Datamining conferences, all because a surprising negative result made us look hard at what was going on.

Tim Oates is an associate professor in the Department of Computer Science and Electrical Engineering at the University of Maryland Baltimore County, and he is director of the university’s Cognition, Robotics, and Learning Laboratory. He received his Ph.D. from the University of Massachusetts Amherst in 2001 and spent a year as a postdoc in the MIT AI lab. His research interests include the sensorimotor origins of knowledge, language learning, grammar induction, automated development of representation.

2. For a more in-depth look at MCL and the motivation behind it, see Anderson and Perlis (2005).

3. In the case of logic-based domains, an anomaly often takes the form of a direct contradiction, E and $\neg E$. This is the case, for instance, not only in the nonmonotonic reasoning domain, but also in the natural language domain discussed in this article. For these, we employ active logic (Elgot-Drapkin and Perlis 2006), a time-sensitive infer-

ence engine specifically designed to allow an automated agent to reason in real time about its own ongoing reasoning, noting direct contradictions rather than inadvertently using them to derive all sentences.

4. Active logic for reason enhanced dialogue.

5. When the tank is destroyed, it reappears at a random location on the map.

6. Although the ontology-based view of MCL described here is still in the process of being implemented, it cleanly captures the more ad hoc implementations of the various MCL-enabled systems that we have developed.

References


- Allen, J. F.; Miller, B. W.; Ringger, E. K.; and Sikorski, S. 1996. Robust Understanding in a Dialogue System. In *Proceedings of the Thirty-Fourth Annual Meeting of the Association for Computational Linguistics*, 62–70. San Francisco: Morgan Kaufmann Publishers.
- Anderson, M. L., and Oates, T. 2007. A Review of Recent Research in Metareasoning and Metalearning. *AI Magazine* 28(1): 12–16.
- Anderson, M. L.; Oates, T.; Chong, W.; and Perlis, D. 2006. The Metacognitive Loop: Enhancing Reinforcement Learning with Metacognitive Monitoring and Control for Improved Perturbation Tolerance. *Journal of Experimental and Theoretical Artificial Intelligence* 18(3): 387–411.
- Anderson, M. L., and Perlis, D. R. 2005. Logic, Self-Awareness, and Self-Improvement: The Meta Cognitive Loop and the Problem of Brittleness. *Journal of Logic and Computation* 15(1): 21–40.
- Brachman, R. J. 2006. (AA)AI, More than the Sum of Its Parts. *AI Magazine* 27(4): 19–34.
- Cox, M. T., and Ram, A. 1999. Introspective Multistrategy Learning: On the Construction of Learning Strategies. *Artificial Intelligence* 112(1–2): 1–55.
- Elgot-Drapkin, J., and Perlis, D. 2006. Reasoning Situated in Time I: Basic Concepts. *Journal of Experimental and Theoretical Artificial Intelligence* 2(1): 75–98.
- Fikes, R. E. 1971. Monitored Execution of Robot Plans Produced by Strips. Technical Report 55, AI Center, SRI International, Menlo Park, CA 94025, April.
- Fox, S., and Leake, D. B. 1995. Using Introspective Reasoning to Refine Indexing. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 391–399. San Francisco: Morgan Kaufmann Publishers.
- Gil, Y. 1994. Learning by Experimentation: Incremental Refinement of Incomplete Planning Domains. In *Proceedings of the Eleventh International Conference on Machine Learning*, 87–95. San Francisco: Morgan Kaufmann Publishers.
- Hammond, K. J. 1986. Learning to Anticipate and Avoid Planning Problems through the Explanation of Failures. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, 556–560. Menlo Park, CA: AAAI Press.
- Josyula, D. P. 2005. A Unified Theory of Acting and Agency for a Universal Interfacing Agent. Ph.D dissertation, Department of Computer Science, University of Maryland, College Park, MD.
- Kendler, H. H., and Kendler, T. S. 1962. Vertical and Horizontal Processes in Problem Solving. *Psychological Review* 69(1): 1–16.
- Kendler, H. H., and Kendler, T. S. 1969. Reversal-Shift Behavior: Some Basic Issues. *Psychological Bulletin* 72(3): 229–232.
- Mitchell, T. M. 1990. Becoming Increasingly Reactive. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, 1051–1058. Menlo Park, CA: AAAI Press.
- Nelson, T. O.; Dunlosky, J.; Graf, A.; and Narens, L. 1994. Utilization of Metacognitive Judgments in the Allocation of Study during Multitrial Learning. *Psychological Science* 5(4): 207–213.
- Perlis, D.; Purang, K.; and Andersen, C. 1998. Conversational Adequacy: Mistakes Are the Essence. *International Journal of Human-Computer Studies* 48(5): 553–575.
- Schank, R. C., and Leake, D. B. 1989. Creativity and Learning in a Case-Based Explainer. *Artificial Intelligence* 40(1–3): 353–385.
- Stroulia, E., and Goel, A. K. 1997. Redesigning a Problem-Solver's Operations to Improve Solution Quality. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, 562–567. San Francisco: Morgan Kaufmann Publishers.
- Traum, D. R.; Andersen, C. F.; Chong, W.; Josyula, D.; Okamoto, Y.; Purang, K.; O'Donovan-Anderson, M.; and Perlis, D. 1999. Representations of Dialogue State for Domain and Task Independent Meta-Dialogue. *Electronic Transactions on Artificial Intelligence* 3(D): 125–152.
- Van Der Krogt, R.; De Weerd, M.; and Witteveen, C. 2003. A Resource Based Framework for Planning and Replanning. *Web Intelligence and Agent Systems* 1(3–4): 173–186.
- Wang, X. 1995. Learning by Observation and Practice: An Incremental Approach for Planning Operator Acquisition. In *Proceedings of the Twelfth International Conference on Machine Learning*. San Francisco: Morgan Kaufmann Publishers.
- Williams, B. C., and Nayak, P. P. 1996. A Model-Based Approach to Reactive Self-Configuring Systems. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, 971–978. Menlo Park, CA: AAAI Press.

Michael L. Anderson (michael.anderson@fandm.edu) is an assistant professor in the Department of Psychology at Franklin & Marshall College and a visiting assistant professor at the Institute for Advanced Computer Studies at the University of Maryland, College Park, where he is also a member of the graduate faculty in the Program in Neuroscience and Cognitive Science. He earned a B.S. in pre-medical studies from the University of Notre Dame, a Ph.D. in philosophy from Yale University, and did his postdoctoral training in computer science at the University of Maryland. Anderson is author or coauthor of over 50 publications in artificial intelligence, cognitive science, and philosophy of mind. His primary areas of research include an account of the evolution of the cortex through exaptation of existing neural circuitry (the “massive redeployment hypothesis”); the role of behavior, and of the brain’s motor-control areas, in supporting higher-order cognitive functions; and the role of self-monitoring and self-control in maintaining robust real-world agency.

Scott Fults graduated in 2006 from the University of Maryland with a Ph.D. in linguistics. His dissertation was concerned with how humans talk about and understand comparisons. He is currently a research scientist at the University of Maryland and works with Don Perlis on MCL and its application to natural language processing. He also does work in human semantic parsing.

Darsana P. Josyula is an assistant professor of computer

New Proceedings!



Available from AAAI Press (www.aaai.org/Press/)

science at Bowie State University, Bowie, MD, and is a member of the ALMECOM group at the Institute of Advanced Computer Studies, University of Maryland. She completed her Ph.D. in 2005 from the University of Maryland, College Park. Her research interests include logical reasoning, natural language understanding, metacognition, and BDI models.

Tim Oates is an associate professor in the Department of Computer Science and Electrical Engineering at the University of Maryland Baltimore County, and he is director of the university's Cognition, Robotics, and Learning Laboratory. He received his Ph.D. from the University of Massachusetts Amherst in 2001 and spent a year as a postdoc in the MIT AI lab. His research interests include the sensorimotor origins of knowledge, language learning, grammar induction, automated development of representation.

Don Perlis is a professor of computer science at the University of Maryland and a member of the artificial intelligence group there. He studies various aspects of commonsense reasoning, including the related areas of

cognitive modeling and philosophy of mind and language. An ongoing project of his research team is the use of time-situated metacognitive computation for enhanced flexibility and generality of reasoning.

Matthew Schmill is a researcher at the University of Maryland, Baltimore County. He received his PhD from the University of Massachusetts in 2004. His research interests include development, metacognition, building robust AI architectures, and computational finance.

Shomir Wilson is a Ph.D. student at the University of Maryland. He earned bachelor's degrees in computer science, mathematics, and philosophy from Virginia Tech in 2005. His research interests include cognitive modeling, commonsense reasoning, and natural language understanding and generation.

Dean Wright is a Ph.D. student at the University of Maryland Baltimore County. He holds an M.S. in computer science and an M.B.A. from Hood College.