

An Architectural Account of Variation in Problem Solving and Execution

Pat Langley (patrick.w.langley@gmail.com)

Institute for the Study of Learning and Expertise
2164 Staunton Court, Palo Alto, CA 94306 USA

Background and Motivation

Theories of the cognitive architecture (Langley et al., 2009) aim to specify the constant features of the human mind. They make assumptions about the representation and organization of memory, as well as about performance and learning mechanisms that operate over its structures. They specifically do not include content that changes, whether rapidly through reasoning and problem solving or gradually through learning.

However, we also know that humans are highly adaptive and they can work on a given task in many different ways. This creates a tension between the desire to identify universals in human cognition and to explain observed variation. Some architectures, like Soar (Laird et al., 1987) and Epic (Kieras & Meyer, 1997), make few commitments except for knowledge representation and the basic cognitive cycle. In contrast, ICARUS (Langley et al., 2009) makes more specific claims about structures and mechanisms, committing to means-ends analysis for problem solving and to reactive control for execution.

In recent research, we have sought a middle ground, devising a new cognitive architecture that incorporates stronger assumptions than Soar and Epic, but weaker ones than ICARUS. In this paper, we present the phenomena it attempts to cover, its theoretical postulates, and planned extensions to the framework. We will see that the theory builds on classic ideas but combines them in novel ways that move beyond its predecessors.

Variation in Problem Solving

Humans exhibit the ability to solve novel problems, with an important case being generation of multi-step plans that achieve goals. The standard theory of problem solving, due to Newell and Simon (1972), states this involves search through a problem space that operates over states, goals, and operators encoded as symbol structures. Heuristics or rules of thumb guide this search process and make it tractable even with limited resources.

Despite these high-level regularities, people also show considerable variety in problem-solving behavior. For instance, they use means-ends analysis on puzzles like the Tower of Hanoi (Newell & Simon, 1972), but rely on forward search in games like chess (de Groot, 1978). Domain expertise accounts for some of these differences, but others appear due to generic strategies that determine search direction, operator selection, and other matters. Our research has led to an expanded theory of human problem solving that comprises five postulates:

- Plans are represented by hierarchical decomposition (AND) trees, in which each node is a problem and every child denotes a subproblem of its parent.
- Problem solving involves the recursive decomposition of problems into subproblems, with alternative candidates organized into a search (OR) tree.
- Problem solving operates in a cognitive cycle of five stages: problem selection, intention generation, subproblem creation, and failure/success checking.
- Strategic knowledge – encoded as domain-independent control elements – governs decisions at each stage to produce different problem-solving strategies.
- Domain expertise takes the form of generalized decompositions that specify how to break a problem into subproblems and that serve as higher-level operators.

We have incorporated these tenets into FPS, a problem solver that implements the theory (Langley et al., 2014). We have tested the system on multiple domains, including puzzles like the Tower of Hanoi and planning tasks involving logistics. We have demonstrated that FPS can solve problems in these domains, that its strategic knowledge reproduces familiar behaviors like means-ends analysis and forward search, and that domain knowledge reduces search and makes problem solving tractable.

Variation in Plan Execution

Humans also exhibit the ability to carry out complex sequential activities over time to achieve their goals (Miller et al., 1960). There are different accounts of such extended behavior, but they generally agree on major features. This topic has been studied mainly in the context of skilled behavior, but similar accounts apply to the execution of plans generated during problem solving, which we argue involve two distinct but linked processes.

Again, although human execution of skills and plans follows some high-level regularities, there remains considerable variation. The field of motor behavior saw a long debate about whether people utilize closed-loop (Adams, 1971) or open-loop (Schmidt, 1982) control, yet it is now clear this differs across individuals and situations. Our theory of plan execution includes four claims:

- Plans and skills are stored as hierarchical decomposition trees like those produced during problem solving.
- Plan execution operates by traversing these decomposition trees from top to bottom and from left to right, with physical actions occurring at terminal nodes.

- This process involves a cognitive cycle with five stages: intention selection, condition checking, intention enactment, perceptual inspection, and effects checking.
- Strategic knowledge – stated as domain-independent control elements – governs decisions at each stage of the cycle to produce different execution strategies.

We have incorporated these ideas into FPE, a flexible plan executor that embodies the theory. We have tested this module on the same domains as FPS, demonstrating that it can carry out hierarchical plans in simulated environments and that it supports a variety of execution strategies, such as open-loop and closed-loop control.

Interleaving Execution with Planning

Just as people exhibit different strategies for plan generation and plan execution, so can they interleave these two processes in distinct ways. This interaction has been less studied than either problem solving or execution in isolation, but some things are clear. In certain settings, a person can generate a complete plan before executing it in an open-loop manner. In others, the problem is so complex, as in difficult puzzles, or the environment is sufficiently unpredictable, as in playing chess, that they must alternate between extending and executing a plan.

Our theory of how humans interleave planning with execution adds two postulates to our earlier accounts:

- Domain-independent strategic control elements determine whether the system proceeds to the next stage of processing or transfers control to the other module.
- The primary loci of control reside in the fourth stage of problem solving – success checking – and the second stage of execution – condition checking.

We have extended FPS and FPE to incorporate these assumptions, letting them operate jointly as parts of an integrated architecture. This uses the problem solver to generate a plan, then shifts to execution and returns to problem solving when monitoring reveals the need.

Exactly when the shifts occur depends on the control elements. We have used different settings to reproduce execution only after forming a top-level plan, alternation between looking N steps ahead and taking a single action (as in game-playing systems), and execution upon solving each subproblem (as in ICARUS). Moreover, computational studies have shown that strategy effectiveness interacts with domain characteristics (Bai et al., 1015).

Strategy Variation and Adaptation

These accounts of variation in problem solving, plan execution, and interleaving make a clear contribution to our understanding of high-level cognition. However, they remain incomplete in that they do not clarify the conditions under which humans utilize different strategies. In ongoing work, we are extending the theory and its implementation to address this challenge.

The central idea is that strategy variation is the result of mental adaptation. The human cognitive architecture has access to each strategy handled by our theory, and it can switch among them as appropriate. This functionality requires access to meta-level information about the state of problem solving, such as branching factors in the forward and backward directions, and execution, such as the reliability of actions. This means that the strategic control elements must include conditions that match against such meta-level data.

We are currently extending the FPS and FPE modules to support such adaptive behavior. We intend to demonstrate that the revised architecture decides for itself which strategies to use during problem solving, plan execution, and interleaving of these activities. The result will be a more complete account of high-level cognition that is consistent with earlier theories but combines their ideas in novel ways to explain important phenomena.

Acknowledgements

This research was supported in part by Grant N00014-15-1-2517 from the Office of Naval Research. C. Pearce, Y. Bai, C. Worsfold, and M. Barley at the University of Auckland contributed to the results reported here.

References

- Adams, J. A. (1971). A closed-loop theory of motor learning. *Journal of Motor Behavior*, 3, 111–149.
- Bai, Y., Pearce, C., Langley, P., Barley, M., & Worsfold, C. (2015). An architecture for flexibly interleaving planning and execution. *Poster Collection: The Third Annual Conference on Advances in Cognitive Systems*.
- de Groot, A. D. (1978). *Thought and choice in chess* (2nd Ed.). The Hague: Mouton Publishers.
- Kieras, D., & Meyer, D. E. (1997). An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. *Human-Computer Interaction*, 12, 391–438.
- Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence*, 33, 1–64.
- Langley, P., Choi, D., & Rogers, S. (2009). Acquisition of hierarchical reactive skills in a unified cognitive architecture. *Cognitive Systems Research*, 10, 316–332.
- Langley, P., Pearce, C., Barley, M., & Emery, M. (2014). Bounded rationality in problem solving: Guiding search with domain-independent heuristics. *Mind and Society*, 13, 83–95.
- Miller, G. A., Galanter, E., & Pribram, K. A. (1960). *Plans and the structure of behavior*. New York: Holt, Rhinehart, & Winston.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Schmidt, R. A. (1982). More on motor programs. In J. A. S. Kelso (Ed.), *Human motor behavior: An introduction*. Hillsdale, NJ: Lawrence Erlbaum.