

Routine Evolution as the Microgenetic Basis of Skill Acquisition

Philip E. Agre
Department of Computer Science
University of Chicago

Jeff Shrager
System Sciences Laboratory
Xerox Palo Alto Research Center

Recent work in several areas has begun to demonstrate the value of microgenetic studies of the evolution of human activities. By paying close attention to the circumstantial detail that chronometrically based statistical methods suppress, microgenetic methods redirect our theoretical attention from putative mental mechanisms to the inherent logic of activities *in situ*. We have been studying in detail the actions of a single subject engaged in a routine photocopying job. Although her actions speed up according to a classical power law, the changes in her performance derive primarily from a constant and quite complex evolution of her patterns of interaction with her environment. Microgenetic analysis of the evolution of three aspects of the subject's activity demonstrates the dialectical relationship between the person-in-activity and the setting in which the activity takes place. Our goal is to build a theory of skill acquisition that does justice to the detailed organization and situational embedding of activity.

1 Introduction

If you do something over and over, you get better at it. A substantial literature understands this commonplace observation in terms of something called a *practice effect*. In this literature, "better" generally means "faster" and what gets faster is hypothesized to be mental processing. When this formulation is operationalized experimentally, one obtains a curve recording how long an experimental task has taken as a function of how much practice the subject has had with it. It is among the most robust results of experimental psychology that this curve, when averaged across subjects, approximates a power function, so that the subjects can be viewed as getting, say, thirty percent better on each iteration relative to some asymptote. This is called the "power law of practice."¹

Many cognitive accounts of the power law and other practice effects have been proposed, including the shifting of procedural subunits from controlled to automatic processing (Shiffrin & Schneider, 1977), the substitution of procedural subunits by remembered results (Newell & Rosenbloom, 1981; Siegler & Shrager, 1984); and the strengthening and associated speedup or collapsing of procedural subunits (Anderson, 1983). The principal concern within this literature has been to isolate a speedup effect independent of complex changes in the activity. This has led most investigators to study practice effects by combining data from several subjects, thus suppressing the sharp changes in performance time that result from qualitative modifications of each individual subject's behavior.

A number of authors have pointed out that performance may improve as a result of changes in the way the subject carries out the task rather than by acceleration of mental processing. Crossman (1959) derived general laws of practice by statistical combination of different possible procedures for accomplishing a task. Anzai and Simon (1979) proposed a theory of skill acquisition incorporating a number of mechanisms including chunking, avoiding bad moves, and generalization. Cheng (1985) warned of the possibility of general "restructuring," a process she likens to changes in the brain's software. Siegler (1987) demonstrates the dangers of statistical combination of performance times in inadvertently hiding important individual differences in the procedures that compose skill, or even worse, of giving credit to a procedure that does not exist at all. Finally, Siegler and Jenkins (1989) observed the discovery of a new strategy in one child's arithmetic performance.

All of these projects begin from the assumption that speed-up is a unitary phenomenon. This assumption leads to a search for general mental speedup mechanisms, such as those based on strengthening, chunking, or restructuring. This strategy, we feel, limits these research programs in three ways. First, it leads some

¹ We recommend Newell and Rosenbloom (1981) for a careful history and analysis of the power law.

of them to suppress the qualitative structure of activities so as to focus on aggregate properties such as elapsed time. Second, even those authors who attend to the details of individual performance claim domain independence for their theories; as a consequence, the nature of “domains” is never clarified or given a substantive role in the theory. And third, the search for mental speedup mechanisms leads *all* of these researchers to treat activities and changes as wholly mental even though, as we will argue, they are not. All of the accounts hypothesize that what changes is some structure located primarily in the subject’s mind, such as a plan, algorithm, or memory configuration.

To overcome these limitations, we propose a different research strategy for investigating the evolution of skills through practice. We propose taking the evolving organization of an activity to be a central phenomenon, not a marginal one. Specifically, we propose viewing individuals’ activities as taking place through interactions with the world and not just in their heads. Activities derive their organization not only from mental processes and structures but from interactions with a potentially changing environment. The topic of study, then, is not “speed-up” or “performance improvement” as such but rather the qualitative processes through which routine activities evolve over time. We refer to this phenomenon as “routine evolution.”

This strategy is consistent with a broad movement in computational and social theories toward viewing human activities as essentially situated in particular concrete circumstances (Agre, in preparation; Lave, 1988; Suchman, 1987). This is not to deny the existence of mental processing, but it is to give interactions with the environment a central role in organizing activity and in determining its course in particular cases. Lave (1988) has recommended adopting as the units of analysis the *person-in-activity* and the *setting* in which the activity takes place. These two entities participate in a dialectical relationship with one another. The term “dialectical” has three implications: that the person-in-activity and setting interact with one another; that they reciprocally influence one another; and that they cannot sensibly be defined except with reference to one another. On this view, routine evolution is best viewed as reflecting the evolution of the person/setting dyad and not simply of mental contents or physical environments on their own. Thus the challenge for research is how to study this unfolding dialectic of interaction as a natural, qualitative, self-organizing process.

2 Routine Evolution in Photocopying

In pursuit of a theory of routine evolution, we are conducting detailed analyses of the evolution of particular routine activities performed by particular individuals. In this paper we describe one of our early studies. We chose to observe routine evolution in the activity of photocopying because it is a very common office task, making it easy to find non-expert subjects with some experience. Indeed, it is common enough that one might not even expect to find practice effects at all.

We employed a deceptive experimental paradigm so that the subjects’ awareness of being observed would not lead them to try self-consciously to optimize their performance. We advertised at Stanford University for a temporary office worker to help out in preparations for an upcoming conference. No mention was made of a psychology experiment. When the subject, whom we shall call Sarah, arrived she was shown to a room containing a Xerox 1075 copier (Figure 1). We had mounted a camera in the ceiling of the room, overlooking the entire copier, and had hidden a microphone near the copier. The camera was not easily visible, and in the debriefing Sarah reported having no awareness that she was being observed. Upon completion of the task, Sarah received \$20, after which the experimental cover was lifted and she was debriefed. She was also given the option of having the videotape erased immediately and on the spot without any questions asked, and without pressure from the experimenter, to the extent that this was possible.

There were four sorts of copying task, but only one of these concerns us here. In this task, three papers were selected from a collection of readings from a Stanford psychology course, bound by a 0.5 inch black “Cerlox” binder, with front and back covers of heavy green paper. A post-it on the book’s cover read: “Need 3 copies of the indicated papers—any style but stapled would be nice.” The first of the three papers was 17 pages in length. The copy of this paper in the spiral-bound readings collection begins on page 67 (a right




Figure 1: The Xerox 1075 photocopier: 1. Recirculating Document Handler (RDH), 2. Control Panel, 3. Copier Body, 4. Output Trays. The START button is the leftmost on the control panel, and the copier glass is on the copier body centered under the RDH.

hand page) of the 296 page book. Sarah began to copy this paper after about one hour and six minutes of work, and finished about four minutes later.

We begin our analysis by looking at the chronometric aspects of Sarah's performance. As we will see in a moment, Sarah's activity in this task is quite complex and changes constantly. Is she actually improving in terms of speed? Figure 2 shows the time in hundredths of seconds between odd-numbered presses of the copier's START button. We depict the time for successive pairs of pages because left-page-right-page is the natural cyclical unit in copying a bound document. We can see that although she begins fairly slowly, she rapidly increases in speed. In fact, this plot is very close to a power law. We hypothesize that if we average such data over a number of subjects, we will be able to produce an excellent classical power law of practice.




Figure 2: The elapsed time for each successive pair of copies.

3 Microgenetic Analysis

Having demonstrated that there is indeed speedup in this simple activity, and that it appears to be compatible with the traditional practice effect findings, we now want to locate the source of this effect. Chronometric methods cannot register the fine qualitative structure in the evolution of Sarah's routine for the copying task. Instead, we follow Vygotsky's advice (1978: 58–75; Scribner, 1985) to study the subject's activity as a historical phenomenon, finding its essence in its processes of change rather than in the "fossilized" outward structure that it might exhibit if allowed the time to become wholly mechanical. Specifically, we employ microgenetic methods such as employed by Siegler and Jenkins (1989). These methods involve the detailed documentation and coding of particular stretches of activity, a sort of micro-longitudinal tracking of the routine evolution process across a long period. The analysis we discuss in this paper covers the four minutes of Sarah's copying activity, observed at video-frame granularity of 1/30th of a second.

We will develop our argument by recounting a series of moments in the evolution of Sarah's photocopying routine. In doing this, we have two goals. First, we wish to demonstrate the utility of microgenetic analysis in making sense of the intricacies of the evolution of this activity. And second, we wish to demonstrate that even when activity is routine, it is continually self-organizing through interaction between the person-in-activity and the setting. A prerequisite of routine activity is the regularity of the relevant aspects of the environment, and this regularity is itself a product of the routine activity.

Sarah's actions while copying the first few pages are markedly circumspect. She settles the book carefully into place, readjusts the controls, looks around her a great deal, and checks the results. None of these aspects of her activity is apparent once her routine has settled down in the later part of the task. It is worth distinguishing this generic circumspection from actions which are present in the beginning but which later prove inessential. Some actions drop out; for instance, she does not close the RDH² cover after the third page. This change saves her about 6.3 seconds per cycle.³ In addition, aspects of the activity slide on top of one another in time and begin to merge together. For example, her pressing of the START button—initially a clearly separate step—merges with other actions such as walking to the right of the copier or continuing to position the book with her left hand.

Sarah uses aspects of the setting to coordinate her activity, but the relationship between these changes over time in the task. For example, because Sarah is making three copies in this task, the copier flashes very brightly three times (at 0.8 second intervals) after the START button has been pressed and before the book can be set up for the next page. Over the course of the task these flashes grow into a pivotal resource in the organization of her activity. This development has two related aspects. First, a period of dead time of about 2.4 seconds intervenes between pressing START and the point at which the book can be moved. Second, the precise moment when manipulation of the book can resume can be identified by means of the third flash. Near the beginning of the task a moment passes between the third flash and Sarah's turning toward and setting up the next page. As time goes on, however, she begins setting up *even before* the third flash (that is, after the first or second flash) so that she is ready to go as soon as the third flash arrives. In the course of the task, this change reduces the time per cycle by about 5.7 seconds.

Having surveyed these general points let us focus on the evolution of a particular segment of each cycle, namely the period after each odd-numbered page, when Sarah turns to the next page. At this point in the cycle the right hand page is face down on the copier glass and the left hand page is draped over the right edge of the glass. She needs to turn the book face up, put it down on the copier, turn the page, and replace the book on the glass so that the next (left hand, even-numbered) page is face down on the glass. In each case, Sarah proceeds by a rapid and fluent sequence of hand-over-hand moves that takes some patience to record accurately. The specific way in which she accomplishes these moves provides a clear example of the

² "Recirculating Document Handler" – see Figure 1.

³ We provide these approximate measurements purely as part of the analysis of the source the power law shown in Figure 2. It is always *extremely* difficult to measure in a principled way the chronometric contribution of these aspects of the activity. In this case, for example, closing the RDH leaves Sarah's right hand closer to the START button than when she is touching the book and does not close the RDH. Therefore, although some time is lost in RDH manipulation, her hand spends less time traveling to the START button. Any chronometric operationalization would have to find a systematic way of correcting for such effects.

dialectical character of routine evolution.

Table 1 catalogs which hand Sarah uses to accomplish each of eight functions in the page turning activity on each of the eight cycles of the copying task (R is right, L is left). It is important to understand that these "functions" are gross categories that we, the theorists, are picking out from Sarah's behavior. We do not know if she would describe the activity in these terms, and of course much more is going on in the page turning activity than this table captures.

With that in mind, we can describe the evolution of Sarah's page turning routine in three parts, as depicted in Table 1. Part A consists of the three functions: lifting the book to turn it face up, flipping it face up, and something we will call "assistance" that generally consists of grasping the free-hanging side of the book and pulling it around in the appropriate direction. Part B consists of lifting the page to be turned and flipping the page to the opposite side. Part C consists of lifting the right side of the now face up book, and flipping it face down. This operation too requires assistance with the free-hanging side of the book. Note that at the end of parts A and C, both of Sarah's hands are grasping the book and she uses both hands to settle the book in place, either face up as in part A (to turn the page) or face down as in part C (to continue copying), on the copier glass.

Last Page Copied

	1	3	5	7	9	11	13	15		
Functions	lift to flip face up	R	R	R	R	R	L	R	L	part A
	flip book face up	L	L	L	L	L	L	L	L	
	flip assist	R	R	R	R	R	R	R	R	
	lift page	L	R	L	R	R	R	R	R	part B
	turn page	L	R	L	L	L	L	L	L	
	lift to flip face down	R	R	R	R	R	R	R	R	part C
	flip book face down	L	L	L	L	L	L	L	L	
	flip assist	R	R	R	R	R	R	R	R	

Table 1: Which hand was used in each of eight functions in the page turning activity over the session.

At this level of analysis, Sarah uses the same hands in part C throughout the task. Her manipulation takes place in the sequence: right hand (lift), left hand (flip), right hand (assist). We will abbreviate this as "RLR". The fact that we do not see change in this sequence at this level of analysis should not be interpreted as absolute stability of this part of the activity. Although Sarah uses the same hands for these functions in the same sequence, there is considerable adjustment involved, for instance, in repositioning her hands and the book in order to accomplish smooth grasps.

Contrast the stability of part C with the relative instability of part A, which goes from a fairly stable RLR sequence to LLR in the page turns following pages 11 and 15. Finally, notice that part B begins inconsistently in the first three cycles (the left part of Table 1), moving from LL to RR to LL, but that in the last five cycles (the right part of Table 1) it has settled down into an RL pattern.

Our task, then, is to explain three contrasting phenomena: the unsystematic change in part A, the systematic change toward stability in part B, and the lack of change in part C. Let us consider these one at a time and then attempt to synthesize what we have learned.

Part A, turning the book face up, begins after the copier has finished copying a given odd-numbered page. As we have mentioned, Sarah has an evolving relationship with the third flash of light, which marks this point. Over a few cycles she learns to anticipate this flash, so that her hands are on the book ready to go when the flash arrives. Later, though, the dead time during the copier's flashing is sometimes filled by other activities. During the flashes for pages 5, 7, and 9, she walks to the right (off the bottom of the video

image), perhaps to check the copies in the side output tray. During pages 11 and 15, perhaps satisfied with the state of the output tray, she stands with her left hand resting on the spine of the book during the flashes. And during the flashes for page 13 she appears to take a moment to relax: she steps back from the copier, looks about, touches the copies sitting on the top of the copier, pulls her sleeve up slightly, and walks back toward the copier, arriving at the precise moment of the third flash. The variation in this segment of the activity is at least partially systematic, in her increasingly accurate anticipation of the third flash, but it is also influenced by a variety of other matters, such as her concerns (whatever they might have been) about the copies in the output tray. As a result of these developments, her physical relationship to the copier at the beginning of part A changes from cycle to cycle. The change is small at first but then grows larger as more elaborate activity comes to fill the dead time. She picks up the book at the beginning of step A with (approximately speaking) whatever hand is closer to its right edge. Which hand is closer, in turn, depends on her lateral (i.e., left or right) position in front of the copier, which in turn depends on what has been going on before to affect that position. Thus, in Table 1, we see that she starts out using RLR hand sequences to turn over the book but that this pattern breaks down later on. The alternation between RLR and LLR does not appear to be systematic, though she might have settled into an approximately stable LLR pattern if the document had more pages.

The evolution of part B is conditioned by several other factors. As part B begins, she has just finished turning the book face up (part A). In particular, she has just finished laying the book down, so that it is now in a highly standardized relationship to her body and especially to her hands. Her way of turning over the page evolves in a systematic way. As she is finishing with part A, her right hand is located at the top right corner of the book and her left hand is located along the bottom edge of the left half of the book. Her task at this point, of course, is to flip over the right-hand page. On the first such page, she begins lifting the top right corner of the page with her right hand but then also begins lifting the bottom of the page with her left hand. She then finishes turning the page with her left hand. On page 3, she begins by moving her right hand to the bottom right hand corner of the book and her left hand to the bottom of the left page of the book, as if smoothing the left page. She turns the page using her right hand, with token guidance from her left hand. On page 5, she moves her left hand to the bottom of the right page and turns the page entirely with her left hand. Meanwhile, her right hand remains on the upper right hand corner of the right page and appears to make a vestigial attempt to turn up the corner just before her left hand turns the page. On page 7 she settles into the emerging pattern of moving both hands to the bottom of the right hand page, lifting the page at the lower right hand corner with her right hand and then turning it over with her left hand. As her left hand turns the page, her right hand pauses for 0.4 seconds and then proceeds to the right edge of the book to prepare to grasp it for part C. Page 9 is very similar to page 7, with the exception that her right hand does not pause between handing off the turned-up page and proceeding to the right edge of the book. All of the remaining page turns look identical to this one, though in some of them it takes her an extra moment to get the page turned up, presumably due to pages sticking together. What we see here is a highly logical evolution of her page-turning routine. Each step changes incrementally from the previous one until the routine settles into a pattern that is stable so far as the representation in Table 1 is concerned. It is worth noting that none of the previous tasks in the experiment had called for her to turn pages in a bound book. She has presumably turned many pages in books before, but the particular way she has evolved to turn pages here appears well-adapted to several aspects of the physical setting: her standing posture, the height of the machine, the lateral placement of the book relative to her body, the mechanical properties of the Cerlox binding, and the affordance offered by the right edge of the book to the grasping operation which will commence part C.

In part C, Sarah uses the same hands for the same functions on every page. Parts A and C involve highly analogous tasks of flipping over the book, yet part A changes and part C does not. The reason for this discrepancy is plain in looking at the tape and from the analysis above. Whereas the relationship of Sarah's body to the setting changes before part A due to different activity during the flashes, the process of flipping the book over (part B) repositions the book with respect to her body. The processes of flipping the book over and turning the page have left her, again, in a highly standardized physical relationship to the book. Table 1 does not register the single exception, which occurs on page 13, when the flipping-over operation

in part A has left the book lying entirely on the copier glass, so that its right edge does not readily afford grasping. In this case, she quickly pushes the book to the right with her left hand before grabbing it with her right and proceeding with the regular RLR pattern.

Before drawing conclusions about the evolution of routine activity, let us gather some observations. Part A of the activity differed from parts B and C in that its initial conditions were subject to variation as a side-effect of previous activity (waiting on the copier). Parts B and C enjoyed very stable initial conditions due to the regularities in outcome of the previous operations (namely those of parts A and B, respectively). Part B differed from part C in that it underwent a steady, incremental process of evolution. Part C, by contrast, did not evolve. But as part B evolved, the “joint” between parts B and C evolved as well, as Sarah’s hand began heading toward the right edge of the book immediately after handing the page off to her left hand. In general, the various parts of the evolving activity stay knitted together despite their various wanderings.

We have spoken of activity in terms of a dialectical relationship between a person-in-activity and the setting. Let us now substantiate this characterization in the present case in terms of the three aspects of a dialectical relationship that we listed before. Sarah and the copier plainly interact with one another: Sarah’s actions are contingent in fine detail on the physical situation, upon which Sarah acts in turn. Moreover, Sarah and the copier exert a mutual influence on one another, as evidenced in the copier’s evolving physical state and Sarah’s evolving skill. And as this mutual influence proceeds, it becomes impossible to characterize either the copier’s state or Sarah’s skill except in terms of their reciprocal relationship. As part of this process Sarah’s actions exhibit a growing fluency and a continually increasing adaptation to the specific details of the setting.

The conditions for stability of Sarah’s activity are naturally analyzed in terms of this interaction. Parts B and C of this activity, unlike part A, could stabilize because their initial conditions were stable. One possible explanation of why part B evolved whereas part C did not is that part C, unlike part B, corresponds to a task that she has had to perform throughout the copying session, so that it is likely to have become more highly evolved. On the other hand, the difference may simply reflect the analytic filtering imposed by the representation in Table 1. In any event, the stability we observe is the stability of the person/setting dyad; changes are changes in this dyad. Even changes to Sarah’s skill, which might be construed as mental changes, cannot readily be defined except in terms of the shifting patterns of interaction in which they participate. The subject’s actions exhibit an increasing fluency and coherence, but these are emergent properties of a more complex process whose logic is the logic of photocopying.

4 Concluding Remarks

Practice is for us not a single “effect,” or ten effects, but the concrete unfolding of a given activity in accord with its own logic. Such a view carries with it a severe methodological challenge: to find a way of studying human activity that both treats episodes of activity as singular events and facilitates the construction of accounts that have broad application. The microgenetic method is substantially more useful than the chronometric method for these purposes because it makes the qualitative organization of the learning process much more evident.

The microgenetic method thus enables us to overcome the three limitations of mentalistic accounts that we discussed at the outset. First, the detailed twists and turns of practice curves which are often treated as deviations from a general trend are, for us, the principal phenomena. The general trend that shows up as the smooth curve one obtains by averaging many individual performance is just that, a general trend, and not a reflection of a single underlying factor. Second, our analysis requires unpacking the detailed organization of the activity and the logic of its change. It thus calls for an analysis of the positive structure of the domain. Indeed, the notion of a domain of activity requires further analysis in terms of its broader social structuration, as suggested by Lave (1988) and the activity theorists (Davydov & Radzikhovskii, 1985). And third, in contrast to theories that locate skill and its acquisition in mental processes, we would locate activity and its evolution in the dialectical relationship between persons-in-activity and settings. Our data provide

clear motivation for such a view. We have not simply added a setting to the mind; we have enriched the analysis of each by taking account of its reciprocal involvement with the other.

This alternative understanding of skill acquisition makes less natural the reification of particular aggregate properties of the activity such as “speed” and “performance” or objective global principles such as a law of minimum effort. But then an important new question arises. These reified concepts had functioned to provide an account of the macroscopic coherence of the development of skills. Given that the subject’s activity could have evolved along many possible paths, why does it evolve in the way it does? In other words, what kind of organizing principle *does* give skill acquisition its overall coherence? Routine evolution is the unfolding of the logic of a particular activity, but this logic itself is mediated by cultural constructions of the activity and of the considerations that might bear on it (piece-work versus hourly wages, precision, comfort, safety, aesthetics, relationships with other people, and so forth). A theory of practice effects will ultimately entail an understanding of the cultural meanings of practice and of skill.

References

- Agre, P. E. (in preparation). *The dynamic structure of everyday life*. Cambridge University Press.
- Anderson, J. R. (1983). *The architecture of cognition*. Harvard University Press.
- Anzai, Y., & Simon, H. A. (1979). The theory of learning by doing. *Psychological Review*, 86(2), 124-140.
- Cheng, P. W. (1985). Restructuring versus automaticity: Alternative accounts of skill acquisition. *Psychological Review*, 92(3), 414-423.
- Crossman, E. R. F. W. (1959). A theory of the acquisition of skill-speed. *Ergonomics Journal*, 153-166.
- Davydov, V. V., & Radzikhovskii, L. A. (1985). Vygotsky’s theory and the activity-oriented approach in psychology. In J. V. Wertsch (Ed.), *Culture, communication, and cognition: Vygotskian perspectives*. Cambridge University Press.
- Laird, J., Rosenbloom, P., & Newell, A. (1986). *Universal subgoalting and chunking: The automatic generation and learning of goal hierarchies*. Boston: Kluwer Academic Publishers.
- Lave, J. (1988). *Cognition in practice: Mind, mathematics, and culture in everyday life*. Cambridge University Press.
- Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Scribner, S. (1985). Vygotsky’s uses of history. In J. V. Wertsch (Ed.), *Culture, communication, and cognition: Vygotskian perspectives*. Cambridge University Press.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84(2), 127-190.
- Siegler, R. S. (1987). The perils of averaging data over strategies: An example from children’s addition. *Journal of Experimental Psychology: General*, 116(3), 250-264.
- Siegler, R. S., & Jenkins, E. (1989). *How children discover new strategies*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Siegler, R. S., & Shrager, J. (1984). Strategy choices in addition and subtraction: How do children know what to do? In C. Sophian (Ed.), *Origins of cognitive skills*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Suchman, L., (1987). *Plans and situated actions*. Cambridge University Press.
- Vygotsky, L., (1978). *Mind in society: The development of higher psychological processes* (M. Cole, V. John-Steiner, S. Scribner, and E. Souberman, Eds.). Cambridge, MA: Harvard University Press.

Philip E. Agre, Department of Computer Science, University of Chicago, 1100 East 58th St., Chicago, IL, 60637, (312)702-6209, agre@gargoyle.uchicago.edu; Jeff Shrager, Xerox Palo Alto Research Center, 3333 Coyote Hill Rd., Palo Alto, CA, 94304, (415)494-4338, shrager@xerox.com