

An Abstract Computational Model of Learning Selective Sensing Skills

Pat Langley¹ (LANGLEY@CS.STANFORD.EDU)
Robotics Laboratory, Computer Science Dept.
Stanford University, Stanford, CA 94305

Abstract

In this paper we review the benefits of abstract computational models of cognition and present one such model of behavior in a flight-control domain. The model's central assumptions are that differences among subjects are due to differences in sensing skills, and that the main form of learning involves updating statistics to distinguish relevant from irrelevant features. We report an implementation of this abstract model of sensory learning, along with a system that searches the space of parameter settings in order to fit the model to observations. We compare the sensory-learning framework to an alternative based on the power law, finding that the latter fits the data slightly better but that it requires many more parameters.

Computational Models of Behavior

Computational models of human cognition date back to the 1950s, soon after researchers realized that computers had general symbol-processing capability. Early computer models like GPS (Newell, Shaw, & Simon, 1960) and EPAM (Feigenbaum, 1963) were implemented in basic list-processing languages like IPL-V and then in LISP. Later models of human behavior were cast in more theory-laden formalisms like production-system and schema languages (Newell, 1973; Norman & Rumelhart, 1975). Ensuing architectures such as ACT (Anderson, 1983) and SOAR (Newell, 1990) incorporated additional knowledge about the human information processor, forcing models stated within those frameworks to satisfy further theoretical constraints.

Many cognitive scientists view this progression as a positive development, leading toward what Newell (1990) has called *unified theories of cognition*. Nevertheless, computational models still require developers to introduce many assumptions, many not central to their theories, before they can produce behaviors and predictions. Moreover, features of models that developers do hold central are often not the source of their models' ability to explain psychological data. One example comes from Richman and Simon (1989), who argue that connectionist and discrimination-network explanations of word-recognition findings are due not to these models' core assumptions of parallel versus sequential processing, but from the way both models structure the task.

These observations suggest that detailed computer models of human behavior, though interesting from an AI perspective, may be misleading or at least unnecessary to explain many interesting phenomena. At first glance, mathematical models seem a natural alternative, in that they describe behavior at a much more abstract level. However, computational models were originally developed in response to perceived limitations of such mathematical methods, which were constrained to simple behaviors and often made restrictive assumptions of their own for the sake of analytical tractability.

Recently, Ohlsson and Jewett (1995; in press) have proposed a promising compromise between these two paradigms, which they refer to as *abstract models*. In this framework, the scientist still implements a running computer program that generates behavior, but the system omits details that are not essential to the phenomena one aims to explain. For example, to model learning in problem-solving domains, they suggest retaining the idea of search through a problem space, but removing details about the states and operators that define the space. Rather, one can describe the structure or connectivity of the space, and model the learning process using mechanisms that add connections or alter the probability of moving toward a goal state.

The idea of abstract computational models is not entirely new. For instance, Shrager, Hogg, and Huberman (1988) present an explanation very similar to Ohlsson and Jewett's for the power law of learning, which they coupled with a mathematical analysis. Rosenbloom and Newell (1987) present a different account of power-law learning, describing both a detailed computer model and an abstract model of this well-known phenomenon. Ohlsson and Jewett's contribution is the realization that neither the mathematical analysis nor the detailed model are necessary, and that researchers may often find it useful to work entirely at the level of abstract models.

However, work on abstract models remains rare, and Ohlsson and Jewett's research program has focused on cognitive tasks. In this paper, we adapt the approach to domains that have a significant sensory-motor component. Below we outline the PHOENIX domain, which involves control of a simulated airplane. After this, we briefly review ICARUS, a theory of the human cognitive architecture, and incorporate its core tenets into an abstract model of behavior on the PHOENIX task. Next we describe a variant model that addresses the influence

¹Also affiliated with the Institute for the Study of Learning and Expertise, 2164 Staunton Court, Palo Alto, CA 94306

of domain knowledge on sensing strategies, along with a simple account of sensory learning. Finally, we consider the model's fit to human behavior and compare its accuracy to that of an alternative account of learning, then discuss some broader issues that our approach raises.

The Sensory-Rich PHOENIX Domain

Goettl (1993, 1994) has described the PHOENIX domain, a simulated training environment that involves flying a simulated airplane through a series of rectangular gates that constitute a three-dimensional slalom course. The aim is to navigate the plane through these gates, preferably following as direct a route as possible. A cockpit window gives subjects information about the size, location, and orientation of the nearest gates as they would appear from an actual plane, along with a horizon line that reflects the plane's pitch and roll. The console display also gives numeric information on the flight speed, thrust, and altitude.

The gates are suspended in air, perpendicular to the ground and parallel to each other. The PHOENIX task begins with the plane facing and heading roughly in the direction of the nearest gate, but the subject must alter the plane's course to accomplish the task of flying through the gates in sequence. A joystick lets the subject affect the plane's pitch and roll, and thus its altitude and heading; additional controls can change thrust and thus flight speed, but this is less central to the basic task.

Goettl (1993) has analyzed the PHOENIX task into 19 separate component skills, which involve subtasks such as changing heading and changing altitude, which in turn break down into even more basic skills like altering the plane's pitch and roll. His experiments revealed a number of regularities in subjects' behavior on this task. For instance, he found that ability on most of the component skills identified during the task analysis were closely associated with ability on the overall slalom task. He also noted major differences in performance, especially between men and women, but also among subjects of the same sex. Determinants of task difficulty included the size of the gates and their distance apart.

Naturally, subjects improve their ability to fly the slalom course with practice. However, Goettl also found that part-task subjects (trained on the component skills) learned more slowly than those in the whole-task group (trained on the overall task), though the former did show positive transfer from practice on the component problems. In studies of a related task that involves shooting stationary targets, Goettl (1994) found that subjects trained on component tasks outperformed those trained on the whole task, provided they get interleaved practice on the components (i.e., one trial on each component per block), but not when they get segregated practice.

We will not attempt to explain all of the above phenomena here; for now we will focus on the basic fact of improvement with experience. However, the variety of results suggests the fertility of this domain for exploring behavior on complex sensory-motor tasks, which recommends it as a testbed for our ideas on abstract models.

A Model of Unskilled Sensing

Our approach to modeling human behavior in the PHOENIX domain builds on the ICARUS architecture (Langley, 1996), in which the basic unit of knowledge is the qualitative state. Each state S specifies a set of conditions that must hold for S to be active, along with optional information about actions to be performed during S , the effects of these actions, and likely successor states. The architecture operates in cycles, checking the conditions of the current state if one is active and selecting a new state from long-term memory otherwise. Constraints on perceptual attention limit the number of sensors updated on each cycle, with the system assuming that the values of unsensed features remain unchanged. When ICARUS detects that the activation conditions for the current state no longer hold, it checks to determine which successor state should become active or, if none hold, which other state seems most appropriate.

For this study, we assume that the agent has already mastered the basic skill of flying through a series of gates, which involves both knowledge of the component skills (states) and the order in which they should occur. Figure 1 shows one possible sequence of states involved in traversing a single gate, and the resulting flight path seen from above the plane. This sequence involves rolling the plane to the left, continuing the roll at the maximum allowed for some period, unrolling the plane right, and taking no action once the plane is aligned with the gate. This sequence assumes the plane is already aligned vertically; if the plane were below the gate, the sequence would also include states for altering the pitch to ascend followed by another state to level out. Alternative locations relative to the gate would produce similar paths based on analogous states, such as decreasing pitch and rolling right. We will not assume this precise decomposition of the slalom task, as other decompositions into states are possible, but we will posit a small number of states for each gate traversal.

Our model of behavior on the slalom task abstracts away from the details of ICARUS and the domain, and focuses on only a few essential parameters. In particular, we suppose that flying through each gate requires a sequence of s states and that each state has $r + i$ activation conditions that involve sensing, but that only r of these conditions actually differ between each state and its successor. This means that, in order to detect that the current state is no longer active, the agent need only sense one of these r relevant features. However, if the agent does not know which features to sense, its detection of state failure may be delayed, and thus it may continue carrying out the current actions longer than appropriate.

Our explanation of errors in this framework revolves around the idea that the agent must reach the final 'Fly Toward' state, in which the plane is aligned with the gate, before passing the gate's location. For a given location of the plane with respect to the gate at the outset of the state sequence, there will be a minimum number of time steps, ignoring time for sensing, for the agent to enter this final state. We will use t to represent the

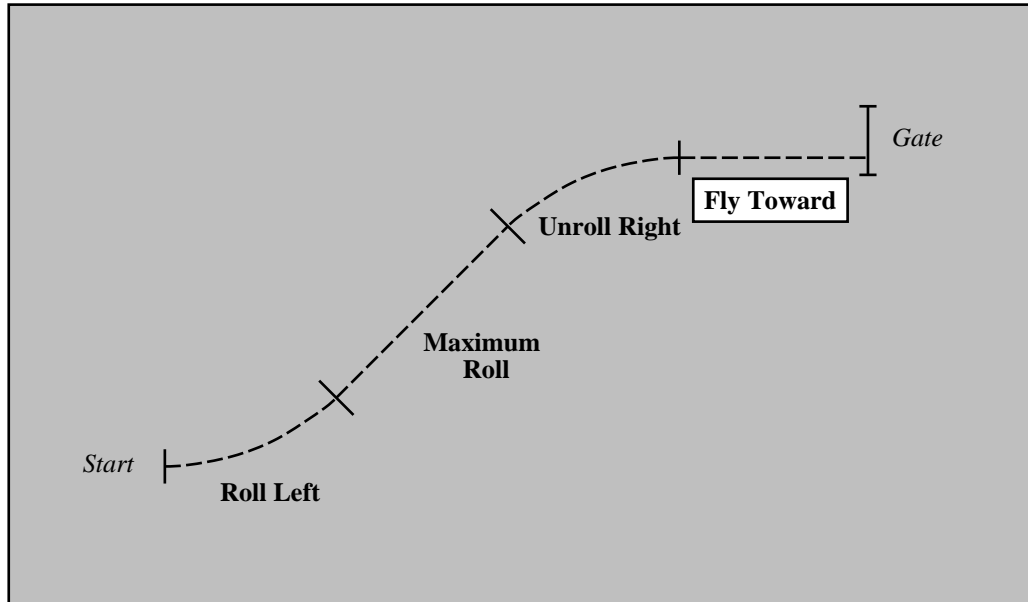


Figure 1: A four-state sequence that takes a plane through a gate when already vertically aligned. Each state continues for a number a time steps, until its activation conditions are no longer satisfied.

number of additional time steps available, beyond this minimum, before the plane passes the gate. Thus, the parameter t corresponds to the amount of ‘slack’ in a particular slalom task, with smaller values making the problem harder and larger ones making it easier.

According to this account, an agent that knows how to fly through a gate can still make errors because the number of time steps needed to detect a state shift may exceed the slack parameter t , causing the plane to miss the gate. We assume that the agent can sense only one feature on each time step, so that whether it notices a state shift depends on whether it senses relevant or irrelevant features. Lacking any knowledge of which features are relevant, we assume that the true novice has a probability

$$p = \frac{r}{r+i}$$

of selecting a relevant feature on each time step, and thus the same probability of noticing a state shift, once such a change occurs.

This model appears to have four parameters but actually has fewer. Note that the important factor is not the overall slack parameter t , but rather than amount of slack per state, $d = t/s$. Also, the actual number of relevant features r and irrelevant ones i matters less than p , the probability of detecting a state shift when one occurs. However, this quantity is determined not by r and i but by their ratio, $u = i/r$, which gives

$$p = \frac{1}{1+u}.$$

Taken together, the parameters d and u specify our abstract model of novice behavior on the PHOENIX slalom task, though it should apply equally well to other sensory-rich domains.

A Model of Skilled Sensing and Learning

The above model posits that the agent samples from among the $r+i$ state activation conditions from a uniform distribution, which produces the probability

$$p = \frac{r}{r+i} = \frac{1}{1+u}$$

of detecting a state shift on each time step after the shift occurs. However, if the agent has additional knowledge about the probability of each condition ceasing to hold, it can use a more selective strategy, based on nonuniform sensing, that produces a higher probability of detecting a state change when one occurs.

In order to model such skilled sensing behavior, we need some additional assumptions. The ICARUS architecture assumes that the agent associates a probability with each activating condition f of a state s , such that, when s is active and f is true, f will still hold on the next time step. Based on these estimates, ICARUS computes the probability that each activation condition (feature) of the current state has changed. Having limited attentional resources, the architecture must choose which features to sense. Here we assume that subjects use a *probability matching strategy*, which samples from among the available features in direct proportion to their estimated probability of changing when a state shift occurs. Probability matching has been implicated in a variety of decision tasks, making it a plausible candidate here.

We can model a subject’s knowledge about the relevance of features with one additional parameter, k , that represents the number of times the subject has observed a particular state transition in which the relevant features have changed and the irrelevant ones have not. We can incorporate this information into the probability of




Figure 2: The probability of detecting a state shift as a function of (a) parameters k and u in the novice model and (b) parameters n and c , when $u = 4$, in the skilled model.

selecting a relevant feature, which becomes

$$p = \frac{k \cdot r}{k \cdot r + i} = \frac{k}{k + u}.$$

This expression is equivalent to the novice quantity, $1/(1 + u)$, when $k = 1$, but the ratio approaches 1 as k goes to infinity. Figure 2 (a) shows the effect of k on p for different values of $i/r = u$.

Naturally, we do not claim that k remains constant, since subjects learn from their experience in the domain. Here we assume that the subject simply increments the value for k by 1 each time he observes a shift from one state to another, thus increasing the probability p of sensing a relevant feature. This suggests that we let $k = n$, where n is the number of times the subject has encountered the task. However, inspection of data for the slalom task reveals that some subjects start with much higher success rates than others. We can model these differences by introducing another parameter, c , that determines each subject's initial probability of sensing irrelevant features. In this revised model, we have $k = c + n$, so that

$$p = \frac{(c + n) \cdot r}{(c + n) \cdot r + i} = \frac{c + n}{c + n + u},$$

where the value for c partly determines the intercept in each subject's learning curve. Figure 2 (b) shows the effect of n on p for different values of c when $u = 4$.

Let us review the model and its structure. We have one parameter, $d = t/s$, that represents the difficulty of the task. We have a second parameter, $u = i/r$, that indicates the ratio of irrelevant to relevant features. Both d and u take on the same value for all subjects, since they are characteristics of the domain. However, we have a third parameter, c , that is specific to each subject, representing that person's initial bias toward sensing relevant features. The variable n also plays a role in the model, but we assume this represents the number of problems

the subject has solved.² Thus, given v subjects with w observations each, we must fit a model with $v + 2$ parameters to $v \cdot w$ data points. For the slalom task, we have 46 subjects and 8 measurements each, giving $46 \cdot 8 = 368$ values to constrain $46 + 2 = 48$ overall parameters.

Fitting the Model to Observations

In principle, we might derive a set of equations that follow from our model and use established statistical methods to determine the best-fitting values for each parameter. However, we have not found any closed-form solutions for the model, which rules out this approach. But it does not preclude us from incorporating the model's assumptions into an abstract computer program, using this program to predict results for given parameter settings, and searching the space of settings to find a good fit to the data.

We implemented the assumptions of the model in such a program, which we embedded in another program designed to search the space of parameter settings. The running model accepted the four variables described earlier – d , u , c , and n – as input and applied the strategy for selective sensing 1000 times to estimate the probability of successfully traversing a gate. The higher-level system computed the squared difference between the predicted and observed probability for each combination of subject and practice level. For the parameter d we told the system to consider only settings between 1 and 3; for u it examined settings from 1 to 20; and for c it considered values from 1 to 10.

The search program involves a number of iterative loops, the outermost devoted to finding the best d value and the next to finding the u setting. The three innermost levels iterate through the set of subjects, through

²Actually, each subject score is an average over 16 three-minute trials that involved separate passes through the slalom course, but these hold across subjects and thus are constant factors.

Table 1: (a) Sample parameter settings for the abstract sensory-learning model and the variance they explain (r^2) on data from the slalom task, along with (b) the parameters and r^2 for the power-law model. The best fit for the sensory-learning model ($d = 2$, $u = 19$) accounts for less variance than the power law but involves many fewer parameters.

(a) SENSORY-LEARNING MODEL			
DIFFICULTY d	RATIO u	INIT. BIAS c	r^2
1	4	[1-8]	0.111
1	12	[1-10]	0.651
2	12	[1-10]	0.580
2	19	[1-10]	0.680
3	12	[1-7]	0.317
3	20	[1-10]	0.652

(b) POWER-LAW MODEL		
SLOPE a	INTERCEPT b	r^2
[-1.17-0.07]	[-1.63-0.57]	0.827

values of n , and through settings for c . Inspection of the model's behavior over this parameter space suggested that, when only one parameter varies, the model's fit to the data follows a U-shaped curve. Thus, the system limited search somewhat by starting with a small parameter value and incrementing it only as long as this improved the fit, at which point it halted, having reached a local optimum given the values of other parameters.

Table 1 shows the variance explained (r^2) for a number of parameter settings, including the one that provides the best fit for Goettl's 46-subject data. The table includes a range of values for c , since this parameter varied across different subjects. Note that the best setting for u is 19, accounting for 68 percent of the variance, which implies that subjects considered 19 times as many irrelevant features as relevant ones. The PHOENIX flight simulator does have a complex display, so this value is not impossible, though it is higher than we expected.

One natural issue concerns how well our sensory-learning model compares to alternative explanations of the data. We plan to explore this question at length in future work, but we have already done some initial studies along these lines with a popular model that assumes learning obeys a negatively accelerated power law. Rosenbloom and Newell (1987) and Shrager et al. (1988) have shown that one can derive this law from assumptions about the task environment and learner, but both analyses deal with reaction times rather than error rates.

Here we simply assume that learning follows a power law of the form $E = bN^{-a}$, where E is the percent error after N training experiences, and where a and b are parameters specific to each subject. Taking the log of both sides gives the linear relation $\log(E) = \log(b) - a \log(N)$, which we can fit to the data using linear regression.

Table 1 also shows the parameter ranges and the r^2 that result from this process. The power law explains somewhat more variance (83 percent) than the sensory-learning model but includes nearly twice as many parameters; thus, we cannot claim that either is superior to the other on these data, and additional studies would appear necessary before we can draw any firm conclusions.

Preliminary analyses of results from another PHOENIX study, involving part-task training, suggest that rapid learners are less affected by the introduction of irrelevant features than slow learners (Goettl, personal communication, 1996). This appears consistent with our sensory-learning theory, but developing a detailed model for this experimental situation, and fitting it to the data, must await future work.

Discussion

Before closing, we should reexamine the theoretical status of our model and its relation to alternative frameworks. We have noted our debt to Ohlsson and Jewett for the notion of an abstract computational model, but our application of this idea differs somewhat from their own. We have used our abstract model, combined with a search engine, to fit data on particular subjects, whereas Ohlsson and Jewett instead explore how alternative models react to variations in parameter values, in order to determine whether their ability to cover phenomena depends on the underlying mechanism or on fortuitous parameter settings. These two approaches are not antithetical, but they do emphasize different issues.

Some readers will detect that our model of sensory learning has features in common with Estes' stimulus sampling theory, the basis for a wide variety of mathematical learning models. The two accounts both assume that subjects' decisions are probabilistic in nature, that they invoke a probability matching strategy, and that learning follows from simple changes to probability distributions. However, the details of the equations for performance and learning differ considerably, as do the underlying accounts that accompany the expressions.

Another issue concerns the degree to which our model, and others like it, explains the data or merely describes it. We hold that the model's processes and associated equations provide explanatory structure, whereas the parameter settings handle description within the structure. A more interesting question concerns the extent to which various model assumptions are necessary or merely sufficient to produce the data. A sufficient assumption can be replaced by another one that, with different parameter values, gives nearly the same results. In contrast, a necessary assumption seems required, in that no alternatives can fit the data, regardless of parameter settings. We have not yet attempted to analyze our account in this fashion, but abstract models seem well suited for such studies, as Ohlsson and Jewett have shown.

A final matter involves the generality of the abstract approach to modeling behavior. Our treatment has ignored many details of the PHOENIX task, such as particular sensory variables and component skills (states), and Ohlsson and Jewett have followed a similar line.

However, we might instead have developed an abstract model that included a separate parameter for each skill, provided data were available (e.g., from part-task studies) to estimate expertise on each. This approach to content-oriented abstract models might even let one distinguish between classes of knowledge, such as functional and structural (e.g., Stroulia & Goel, 1992), given these classes have different implications for subject's behavior.

Concluding Remarks

In this paper we reviewed an approach to cognitive simulation that Ohlsson and Jewett (1995) have called *abstract models*. We considered the advantages of this approach over traditional AI models of human behavior, which force one to specify a complete procedure that operates in the task domain even when the data provide insufficient constraints to justify such detail. We described a domain of this sort, studied by Goettl (1993, 1994), in which subjects must fly a simulated aircraft through a three-dimensional slalom course. Although we have implemented an AI system for this task, cast within the framework of a cognitive architecture, we found this system too complex for useful modeling of available data.

In response, we developed an abstract model of behavior on this task that incorporated parameters for task difficulty, the ratio of irrelevant to relevant features, and initial subject knowledge. The model's central assumptions are that skilled performance on this task involves selective sensing of relevant rather than irrelevant features, and that improvement comes from simple statistical learning about feature relevance. We implemented a program to search the space of parameter settings, and in this way found an instantiated form of the model which approached the fit for a power-law model that had twice as many parameters. These results do not prove that our sensory-learning account is the correct one, but they encourage us to continue exploring this class of models.

In future work, we plan to evaluate our abstract model on more detailed data that Goettl has collected for the PHOENIX domain, as well as compare it to other alternatives besides the power law. We also plan to draw on more sophisticated methods, some available in the statistical literature, for searching the space of parameter settings, and to produce more general tools that can be used with a broad class of abstract models. In the longer term, we hope to use the resulting system to develop and evaluate abstract models for a variety of learning tasks, in an effort to understand the potential of this approach to cognitive simulation.

Acknowledgements

This research was supported by Grant F49620-94-1-0118 from the Computer Science Division of the Air Force Office of Scientific Research. We owe thanks to Barry Goettl for useful discussions about the PHOENIX domain and for making his experimental data available, to Stellan Ohlsson, Wes Regian, and Jeff Shrager for their insights about abstract models, and to Stephanie Sage for fitting the power law to the flight-control data.

References

- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge: Harvard University Press.
- Feigenbaum, E. A. (1963). The simulation of verbal learning behavior. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought*. New York: McGraw-Hill.
- Goettl, B. P. (1993). Analysis of skill on a flight simulator: Implications for training. *Proceedings of the 37th Annual Meeting of the Human Factors and Ergonomics Society* (pp. 1257–1261). Santa Monica, CA: Human Factors and Ergonomics Society.
- Goettl, B. P. (1994). Contextual interference effects on acquisition and transfer of a complex motor task. *Proceedings of the 38th Annual Meeting of the Human Factors and Ergonomics Society* (pp. 1220–1224). Nashville: Human Factors and Ergonomics Society.
- Langley, P. (1996). *Learning sensory skills for control domains*. Unpublished manuscript, Robotics Laboratory, Computer Science Department, Stanford University, Stanford, CA.
- Newell, A. (1973). Production systems: Models of control structures. In W. G. Chase (Ed.), *Visual information processing*. New York: Academic Press.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Newell, A., Shaw, J. C., & Simon, H. A. (1960). Report on a general problem-solving program for a computer. *Information Processing: Proceedings of the International Conference on Information Processing* (pp. 256–264).
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Norman, D. A., & Rumelhart, D. E. (1975). (Eds.). *Explorations in cognition*. San Francisco: Freeman.
- Ohlsson, S., & Jewett, J. J. (1995). Abstract computer models: Towards a new method for theorizing about adaptive agents. *Proceedings of the Eighth European Conference on Machine Learning* (pp. 33–52). Heraklion, Crete: Springer Verlag.
- Ohlsson, S., & Jewett, J. J. (in press). Ideal adaptive agents and the learning curve. In J. Brzezinski, T. Maruszewski, & B. Krause (Eds.), *Idealization in psychology*. Amsterdam: Rodopi.
- Richman, H. B., & Simon, H. A. (1989). Context effects in letter perception: Comparison of two models. *Psychological Review*, *96*, 417–432.
- Rosenbloom, P. S., & Newell, A. (1987). Learning by chunking: A production system model of practice. In D. Klahr, P. Langley, & R. Neches (Eds.), *Production system models of learning and development*. Cambridge, MA: MIT Press.
- Shrager, J., Hogg, T., & Huberman, B. A. (1988). A graph-dynamic model of the power law of practice and the problem-solving fan effect. *Science*, *242*, 414–416.
- Stroulia, E., & Goel, A. K. (1992). Generic teleological mechanisms and their use in case adaptation. *Proceedings of the Fourteenth Conference of the Cognitive Science Society* (pp. 319–324). Bloomington, IN: Lawrence Erlbaum.