Gardenförs, P. (2004). Cooperation and the evolution of symbolic communication. In K. Oller & U. Griebel (Eds.), *Evolution of communication systems: A comparative approach* (pp. 237–256). Cambridge, MA: MIT Press.
Gott, S. P., & Lesgold, A. M. (2000). Competence in the workplace: How cognitive performance models and situated instruction can accelerate skill acquisition. In R. Glaser (Ed.), *Advances in instructional psychology* (pp. 239–327). Mahwah, NJ: Lawrence Erlbaum Associates.
Katz, S., O'Donnell, G., & Kay, H. (2000). An approach to analyzing the role and structure of reflective dialogue. *International Journal of Artificial Intelligence in Education, 11(3),* 320–343.
Lave, J., & Wenger, E. (1990). *Situated learning: Legitimate peripheral participation.* Cambridge, England: Cambridge University Press.
Lesgold, A. (in press). Assessment to steer the course of learning: Dither in testing. In E. L. Baker, J. Dickieson, W. Wulfeck, & H. O'Neil (Eds.), *Assessment of problem solving using simulations.* Mahwah, NJ: Lawrence Erlbaum Associates.
Lesgold, A. (1996). Quality control for educating a smart work force. In L. B. Resnick, & J. Wirt, (Eds.), *Linking school and work: Roles for standards and assessment* (pp. 147–191). San Francisco: Jossey-Bass.
Lesgold, A., & Nahemow, M. (2001). Tools to assist learning by doing: Achieving and assessing efficient technology for learning. In D. Klahr & S. Carver (Eds.), Cognition and instruction: Twenty-five years of progress (pp. 3070–346). Mahwah, NJ: Lawrence Erlbaum Associates.
Levy, F., & Murnane, R. J. (2005). *The new division of labor: How computers are creating the next job market.* Princeton, NJ: Princeton University Press.
Rasmussen, J. (1986). *Information processing and human–machine interaction: An approach to cognitive engineering.* New York: North-Holland.
Schön, D. (1987). Educating the reflective practitioner. San Francisco: Jossey-Bass.
Thorndike, E. L., & Woodworth, R. S. (1901). The influence of improvement in one mental function upon the efficiency of other functions (I). *Psychological Review, 8,* 247–261.
Willingham, W. W. (1985). *Success in college: The role of personal qualities and academic ability.* New York: College Entrance Examination Board.

CHAPTER
6

The Algebraic Brain

John R. Anderson
Carnegie Mellon University

My enduring intellectual interests were all formed in the 4 years of graduate education with Gordon Bower. I basically imprinted on the things he was interested in during the period 1968 to 1972. This included an interest in mathematical psychology, artificial intelligence, learning and memory, and applications of psychology to improving classroom performance. Whereas Gordon's interests have ranged far and wide over the decades, I have remained stuck on these topics. This chapter describes current work in my laboratory bringing these various threads together. I am going to describe a formal information-processing model of how children learn to solve linear equations and test predictions of this model for activation patterns in five brain regions.

THE EXPERIMENT AND THE ACT–R MODEL

In the experiment modeled in detail (Qin, Anderson, Silk, Stenger, & Carter, 2004), 10 students aged 11 to 14 spent 6 days practicing solving simple linear equations. The first day (Day 0) they were given private tutoring on how to solve a restricted set of equations and practiced paper-and-pencil solutions of such problems with a private human tutor. On the remaining 5 days, they practiced on a computer the solution of three classes of equations:

0-step: e.g., $1x + 0 = 4$

1-step: e.g., $3x + 0 = 12$, $1x + 8 = 12$

2-step: e.g., $7x + 1 = 29$

Each day they went through 10 computer-administered blocks of such equations. Each block consisted of 16 trials with four instances of the four possible types of equations (there are two subtypes for the one-step equations). Figure 6–1 presents their latency and the predictions of a model implemented in the ACT–R architecture (Anderson et al., 2004).

The ACT–R Architecture

According to the ACT–R theory, cognition emerges through the interaction of a number of independent modules. Figure 6–2 illustrates the modules relevant to algebra equation solving:

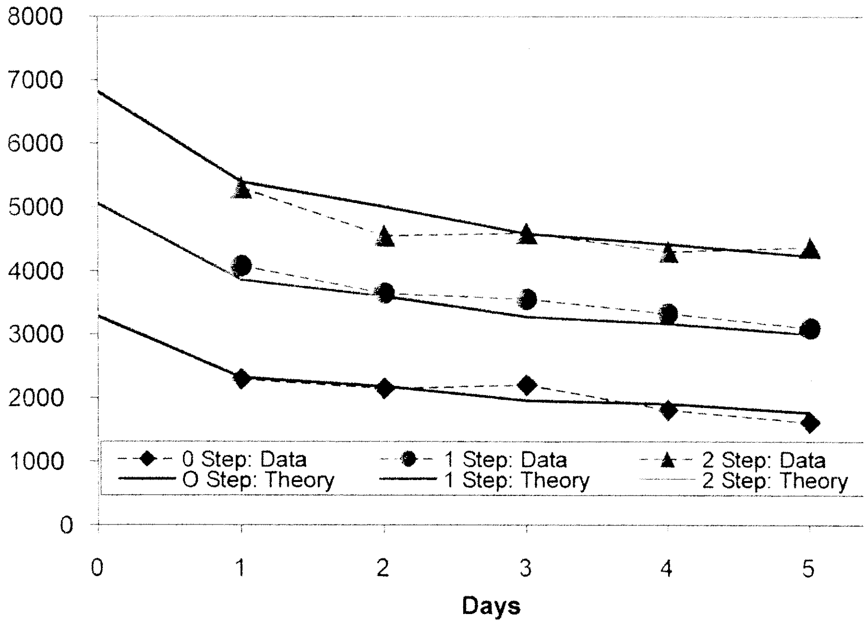

Figure 6–1. Mean solution times (and predictions of the ACT–R model) for the three types of equations as a function of delay. Although the data were not collected, the predicted times are presented for the practice session of the experiment (Day 0).

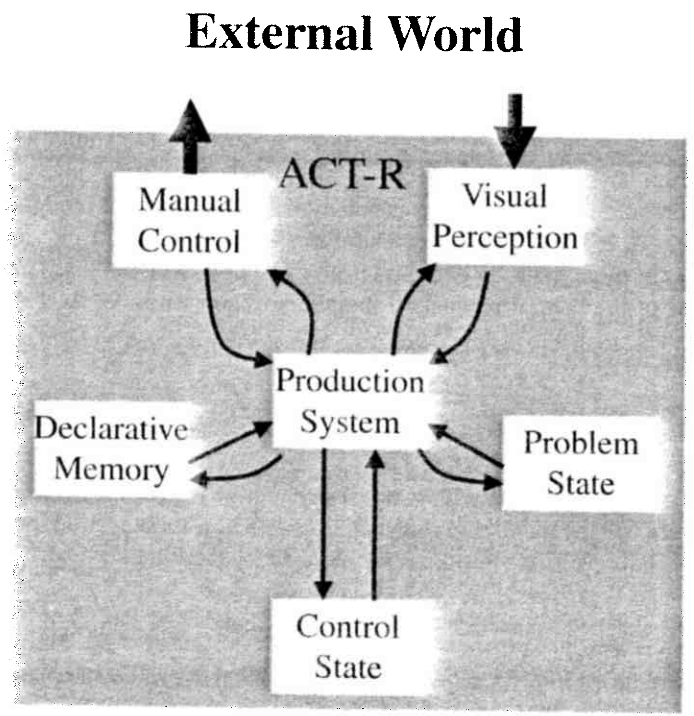

Figure 6–2. The interconnections among modules in ACT–R 5.0.

1. A visual module that might hold the representation of an equation such as "$3x - 5 = 7$."
2. A problem state module (sometimes called an imaginal module) that holds a current mental representation of the problem. For instance, the student might have converted the original equation into "$3x = 12$."
3. A control module (sometimes called a goal module) that keeps track of one's current intentions in solving the problem—for instance, one might be trying to apply the unwind strategy described later.
4. A declarative module that retrieves critical information from declarative memory such as that $7 + 5 = 12$.
5. A manual module that programs the output such as "$x = 4$."

Each of these modules is capable of massively parallel computation to achieve its objectives. For instance, the visual module is processing the entire visual field and the declarative module searches through large databases. However, each of these modules suffers a serial bottleneck such that only a little information can be put into a buffer associated with the module—a single object is perceived, a single problem state represented, a single control state maintained, a single fact retrieved, or a single program for hand movement executed. Formally, each buffer can only hold what is called a "chunk" in ACT–R, which is structured unit bundling a small amount of information. ACT–R does not have a formal concept of a working memory, but the current state of the buffers constitutes an effective working memory. Indeed, there is considerable similarity between these buffers and Baddeley's (1986) working memory "slave" systems. Communication among these modules is achieved via a procedural module (production system in Fig. 6–2). The procedural module can respond to information in the buffers of other modules and put information into these buffers. The response tendencies of the central procedural module are represented in ACT–R by production rules.

The ACT–R Model

The ACT–R model begins with a set of declarative instructions, given in Table 6–1, that encode the unwind strategy. To illustrate how these instructions apply to example equations, first consider a simple 0-step equation like:

$$1 * x + 0 = 2.$$

These instructions imply a sequence of operations that can be summarized:

Instruction 1a: Create image " = 2."

Instruction 2b: Unwind-right "1 * x + 0."

Instruction 3a: Focus on "1 * x" and unwind it.

Instruction 2c: Unwind-left "1 * x."

Instruction 4a: Focus on "x" and unwind it.

Instruction 2a: The answer is 2.

Whereas for a two-step equation like

$$7 * x + 3 = 38$$

they imply a sequence of operations that can be summarized:

Instruction 1a: Create image " = 38."

Instruction 2b: Unwind-right "7 * x + 3."

Table 6–1
English Rendition of Instructions Given to ACT-R Model for Equation Solving

1) To solve an equation, encode it and
   - a). If the right side is a number then imagine that number as the result and focus on the left side and unwind it.
   - b). If the left side is a number then imagine that number as the result and focus on the right side and unwind it.
2) To unwind an expression
   - a). If the expression is the variable then the result is the answer.
   - b). If a number is on the right unwind-right
   - c). If a number is on the left unwind-left
3) To unwind-right, encode the expression (of the form "subexpression operator number") and
   - a). If the operator is + or − and the number is 0 then focus on the subexpression and unwind it.
   - b). Otherwise invert the operator, imagine it as the operator in the result, imagine the number of the expression as the second argument in the result, evaluate the result, and then focus on the subexpression and unwind it.
4) To unwind-left encode the expression (of the form "number operator subexpression") and
   - a). If the operator is * and number 1 then focus on the subexpression and unwind it.
   - b). Otherwise check that the operator is symmetric, invert the operator, imagine it as the operator in the result, imagine the number as the second argument in the result, evaluate the result, and then focus on the subexpression and unwind it.

Instruction 3b: Change image to " = 38 − 3"; this to " = 35"; and focus on "7 * x" and unwind it.

Instruction 2c: Unwind-left "7 * x."

Instruction 4b: Change image to " = 35/7"; this to " = 5"; and focus on x and unwind it.

Instruction 2a: The answer is 5.

Figure 6–1 shows that ACT–R is able to reproduce the speed-up seen in the participants. The key to understanding this speed-up in the ACT–R model is to understand how the preceding instructions were interpreted. These instructions are encoded as declarative structures and ACT–R has general interpretative productions for converting these instructions to behavior. For instance, there is a production rule that retrieves the next step of an instruction:

> IF one has retrieved an instruction for achieving a goal
> THEN retrieve the first step of that instruction.

There are also productions for retrieving particular arithmetic facts such as

> IF one is evaluating the expression "*a* operator *b*"
> THEN try to retrieve a fact of the form "*a* operator b = ?"

Using such general instruction-following productions is laborious and accounts for the slow initial performance of the task.

Though multiple types of learning are occurring in this experiment, it is mainly production compilation that is accounting for the speed-up (see Taatgen, 2005; Taatgen & Anderson, 2002). This is a process by which new production rules are learned that collapse what was originally done by multiple production rules. In this situation, the initial instruction-following productions are compiled over time to produce productions to embody procedures that efficiently solve equations. For instance, the following production rule is acquired:

> IF the goal is to unwind an expression
> and the expression is of the form "subexpression = 0"
> THEN focus on the subexpression.

Figure 6–3a illustrates a typical trial at the beginning of the Day 1 and Figure 6–3b illustrates a typical trial at the end of the Day 5. In both cases, the model is solving the two-step equation $7 * x + 3 = 38$. The figure illustrates when the various modules were active during the solution of the equation and what they were doing. The Day 1 trial (Fig. 6–3a) takes 6.1 seconds and the Day 5 trial (Fig. 6–3b) takes 4.1 seconds. However, these do not reflect the extremes of the learning curve according to ACT–R. The very first trial on Day 0 takes 8.4 seconds in the model. With an infinite amount of practice, the model would take 1.7 seconds during which it would only read the equation and type the answer, having compiled the answer into production rules for that problem. Still, the contrast between parts a and b of Figure 6–3 gives a sense for what is happening over the course of learning. It is worth emphasizing a number of general features of the activity in the figure before discussing the detail of what is happening in individual buffers:

1. Multiple modules can be active simultaneously. For instance, early on in the figure there is a point where the goal module is noting that it is implementing

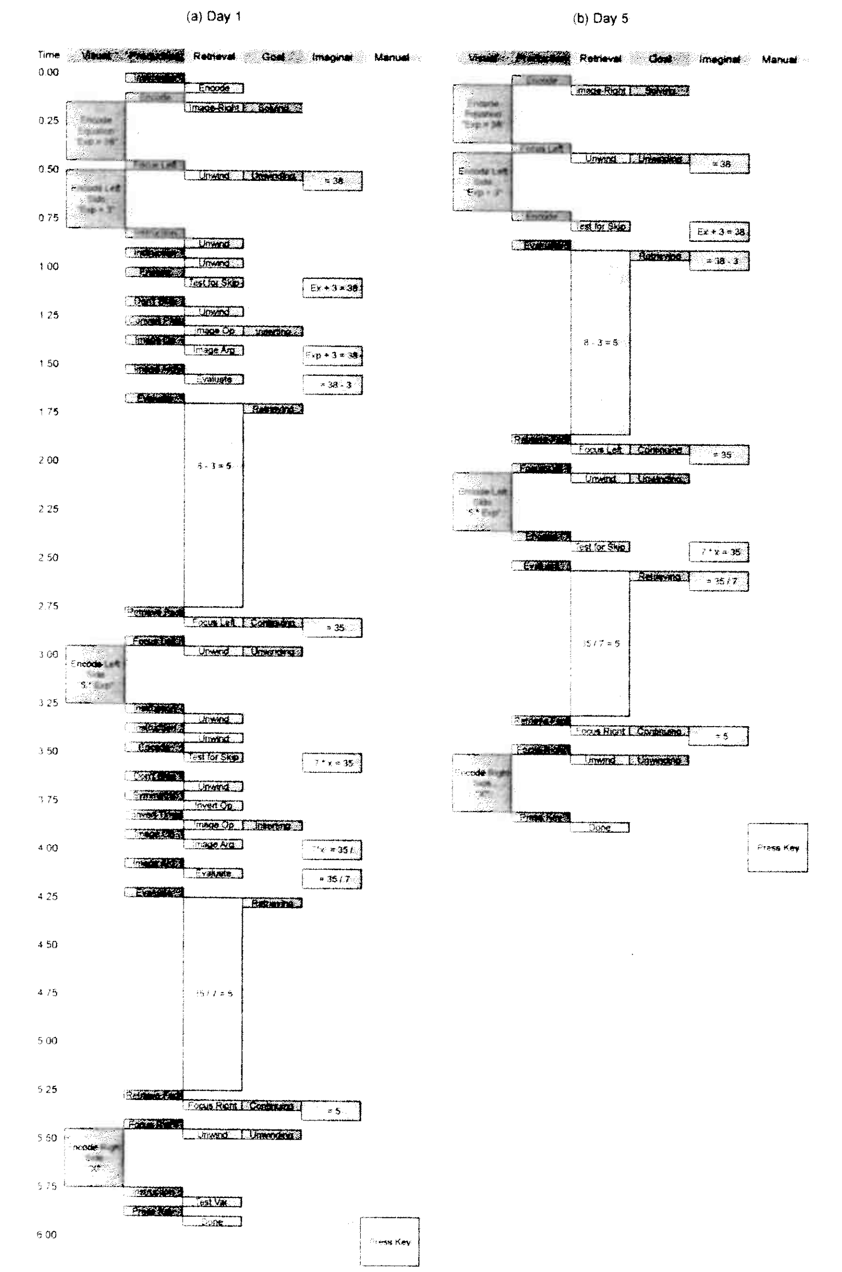

Figure 6–3. Comparison of the module activity in ACT–R during the solution of a two-step equation on Day 1 (part a) with a two-step equation on Day 5 (part b). In both cases the equation being solved is $7 * x + 3 = 38$.

the unwind strategy, while an image of the right-hand side of the equation ("= 38") is being encoded in the imaginal buffer, while the next step in the unwind strategy is being retrieved, and while the visual system is encoding the left-hand side of the equation. Certain of these activities tend to be on the critical path because they are taking longer than the other processes and further processing has to wait for them to complete. In these cases, the times of the other operations have no effect on total time. For instance, often the visual encoding of the equation is holding up other operations and the durations of these other operations do not matter.

2. Much of the speed-up in processing is driven by collapsing multiple steps into single steps. A particularly dramatic instance of this is noted in Figure 6–3 where five production firings, five retrievals, two control settings of the goal, and two imaginal transformations are compressed into one production, one retrieval, one control setting, and one imaginal transformation.[1] Production compilation can compress these internal operations without limit. What it cannot collapse are the external operations such as visual encodings or manual operations. These external operations define the bounds of the compilation. Whereas the example in Figure 6–3 shows multiple productions being collapsed, the actual learning proceeds slowly in ACT–R and takes all 5 days to achieve the transformation in Figure 6–3. Given enough practice, ACT–R would collapse all equation solving simply into a series of visual encodings and manual operations and there would be no effect of equation complexity (nor any real thought occurring). However, to do so ACT–R would have to essentially build into production rules the capacity to recognize each possible equation and produce its solution. The combinatorics of this are so overwhelming (so many different possible equations) that it would never happen in the normal course of learning to solve equations.
3. A second, lesser source of speed-up is the reduction of retrieval times. This reflects an increase in the base-level activation of the facts used in this experiment and as such it is an example of subsymbolic activation learning. This subsymbolic learning is a relatively minor contributor to the learning in Figure 6–1 for two reasons. First, the basic instructions get used over and over again and are already strongly encoded during Day 0 and there is not that much room for further speed-up. Second, the arithmetic facts do not repeat very often over the course of the experiment and are getting little practice over their baseline. In other situations, subsymbolic activation processes can have a major effect on performance. However, over the period of time studied in this experiment, the major learning is happening at the symbolic level in terms of creating new production rules.

[1]Although this is a particularly dramatic example of production compilation, there are many other instances in Figure 6–3 that I have not noted to avoid overly cluttering the figure.

Comments on the Model

The model has some considerable virtues. It actually interacts with the same software as the participants and so really does the task—there are no vague, unspecified bridging assumptions in its predictions. The model is not handcrafted to do the task but learns from instruction. Moreover, it is has the same model of instruction following that has been used in a number of other efforts in our lab (Anderson et al., 2004, Anderson, Taatgen, & Byrne, 2005). Only two parameters were estimated to fit the data. One was a scale factor that determines the length of the retrieval episodes and the other was the time for encoding a fragment of an equation from the screen (300 ms.). Although the fit to Figure 6–1 is pretty good, the reader might well harbor some doubt about whether this really justifies all the detail in the model in Figure 6–3. There are many unseen steps of processing. We have been engaged in a program of brain imaging to try to bring some converging data to bear on these assumptions.

USING fMRI TO TEST ACT–R MODELS

We have now completed a large number of functional magnetic resonance imaging (fMRI) studies of many aspects of higher level cognition (Anderson, Qin, Sohn, Stenger, & Carter, 2003; Anderson, Qin, Stenger, & Carter, 2004; Qin et al., 2003; Sohn, Goode, Stenger, Carter, & Anderson, 2003; Sohn, Goode, Stenger, Jung, Carter, & Anderson, 2005) and based on the patterns over these experiments we have made the following associations between a number of brain regions and modules in ACT–R. In this chapter, we are concerned with five brain regions and their ACT–R associations:

1. **Caudate (Procedural):** Centered at Talairach coordinates $x = -5$, $y = 9$, $z = 2$. This is part of a set of subcorticial structures called the basal ganglia that we associate with the procedural system.
2. **Prefrontal (Retrieval):** Centered at $x = -40$, $y = 21$, $z = 21$. This includes parts of Brodmann Areas 45 and 46 around the inferior frontal sulcus.
3. **Anterior Cingulate (Goal):** Centered at $x = -5$, $y = 10$, $z = 38$. This includes parts of Brodmann Areas 24 and 32.
4. **Parietal (Problem State or Imaginal):** Centered at $x = -23$, $y = -64$, $z = 34$. This includes parts of Brodmann Areas 39 and 40 at the border of the intraparietal sulcus.

5. **Motor (Manual):** Centered at $x = -37$, $y = -25$, $z = 47$. This includes parts of Brodmann Areas 3 and 4 at the central sulcus.

Predicting the BOLD Response

We have developed a methodology for relating the profile of activity in modules like those in Figure 6–3 to Blood Oxygen Level Dependent (BOLD) responses from the brain regions that correspond to these modules. Figure 6–4 illustrates the proportion of time that a particular module was active at various points during a

(a)

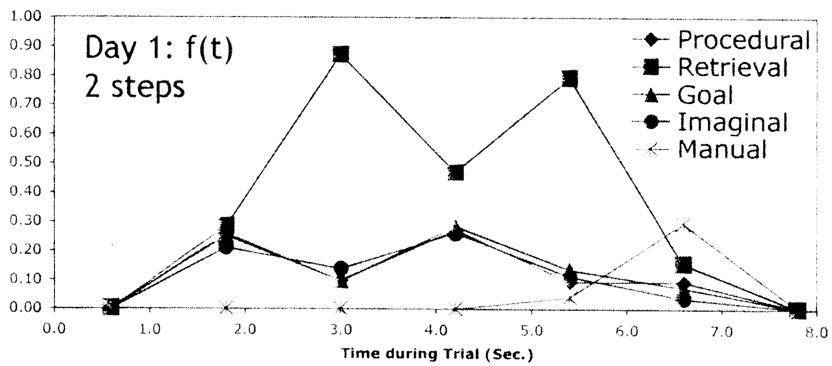

(b)

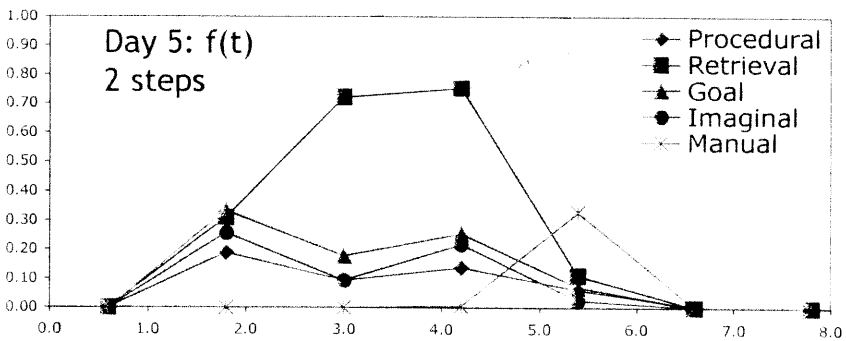

Figure 6–4. The degree of engagement of the various modules during a trial on Day 1 (part a) and Day 5 (part b).

trial on Day 1 (Part a) and Day 5 (Part b) for the two-step equations. These numbers would be directly obtainable from Figure 6–3, except that Figure 6–4 reflects the average engagement over the whole day not just at the beginning of Day 1 (Fig. 6–3a) and the end of Day 5 (Fig. 6–3b). The basic model we have developed of the BOLD response claims that the while a module is engaged it is producing a hemodynamic response in the corresponding region. We have adopted the standard gamma function that other researchers (e.g., Boyton, Engle, Glover, & Heeger, 1996; Cohen, 1997; Dale & Buckner, 1997; Glover, 1999) have used for the BOLD response. If the module is engaged it will produce a BOLD response t time units later according to the function:

$$B(t) = m\left(\frac{t}{s}\right)^a e^{-(t/s)}$$

where m governs the magnitude, s scales the time, and the exponent a determines the shape of the BOLD response such that with larger a the function rises and falls more steeply. The time to peak for the BOLD response is $a * s$ and the magnitude area under the curve is $m * s * \Gamma(a)$ where Γ is the gamma function ($\Gamma(a) = (a - 1)!$) The BOLD response accumulates whenever the region is engaged. Thus, if $f(t)$ is an engagement function giving the probability that the region is engaged at time t, then the cumulative BOLD response can be obtained by convolving the two functions together:

$$\mathrm{CB}(t) = \int_0^t f(x)\, B(t - x) dx$$

This is the observed BOLD response. Its area is proportional to the total time that the region is engaged. Thus, if a module is active for T seconds, then the area under the BOLD response is $T{*}m{*}s{*}\Gamma(a)$.

In summary, a model for the time course (Fig. 6–3) of this task yields engagement functions $f(t)$ like those in Figure 6–4. By convolving the engagement functions with the BOLD function, one can obtain predictions for the BOLD response in the regions associated with the modules. Most of the parameters of this model are set according to prior values established for ACT–R, but fitting the latency in Figure 6–1 did require estimating parameters for the time to encode the equation and the duration of the retrievals. Having now committed to the time course of each module, predictions immediately follow for the time course of the BOLD response. The exact height and shape of the BOLD response depends on the magnitude (m), the scale (s), and the exponent (a) for the region that corresponds to that module. However, the strong parameter-free prediction is that the relative areas under the BOLD responses in two conditions for a region will reflect the relative amounts of time this region is engaged in these two conditions. Thus, the BOLD response provides a direct check on assumptions about the amount of time various modules are engaged in doing a task.

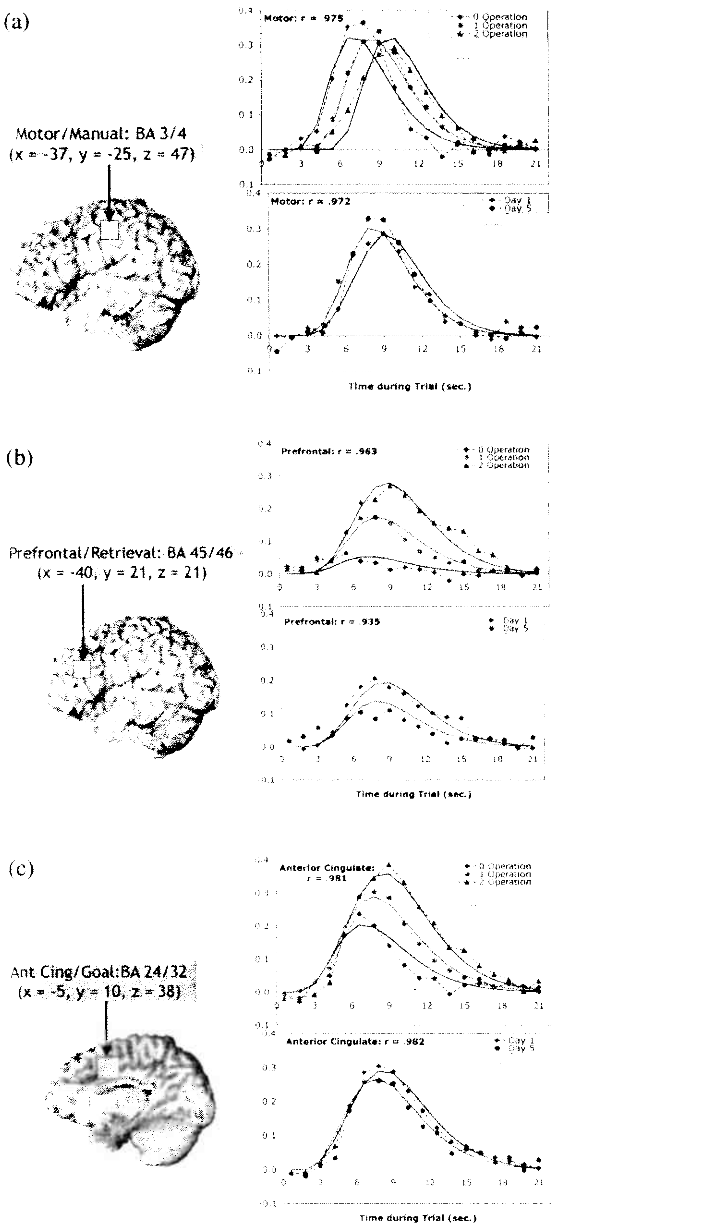

Figure 6–5. Use of module behavior to predict BOLD response in various regions: (a) Manual module predicts motor region; (b) Retrieval Module predicts prefrontal region; (c) Control/Goal module predicts anterior cingulate region

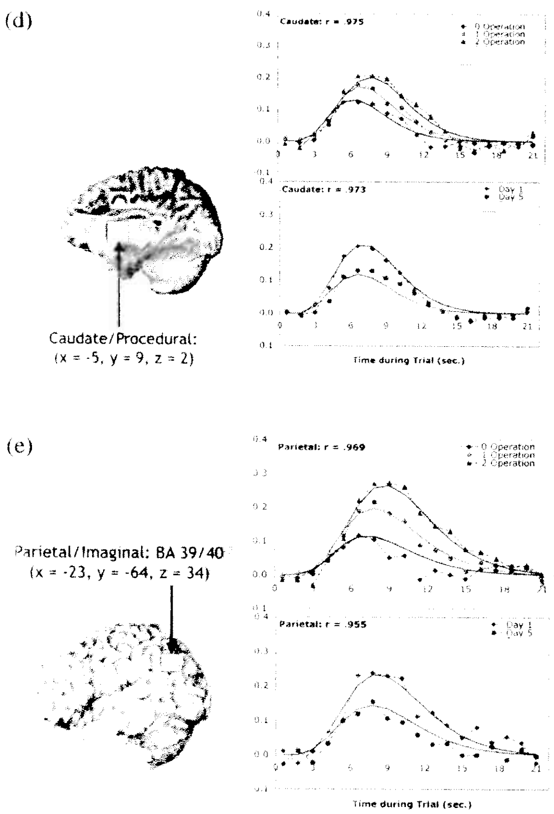

Figure 6–5. *(continued)* (d) Imaginal/Problem State module predicts parietal region; (e) Procedural module predicts caudate region.

Table 6–2 gives the estimated parameters for the BOLD response and Figure 6–5 shows how well this model predicts the BOLD responses in the six conditions achieved by crossing day and number of steps of transformation for each of the five associated regions. To simplify matters and to make the functions more comparable, the exponent of the BOLD response was set to 3 for all regions. To keep the data presentation readable and get better estimates, Figure 6–5 averages either over days or over conditions.[2]

[2]In fact, none of the regions showed a significant interaction between practice and number of steps or between practice, number of steps, and scan.

Table 6–2
Parameters Estimated and Fits to the Bold Response

	Motor/ Manual	Prefrontal/ Retrieval	Parietal/ Imaginal	Cingulate/ Goal	Caudate/ Procedural
Magn(m)	0.531	0.073	0.231	0.258	0.207
Exponent(a)	3	3	3	3	3
Scale(s)	1.241	1.545	1.645	1.590	1.230
Correlation	0.975	0.963	0.969	0.981	0.975
Chi-square (105 *df*)	88.93	82.60	95.21	123.27	81.03

Characterizing the Differences Among the Brain Regions

The first impression one probably gets from Figure 6–5 is that the BOLD responses for the five regions look a lot alike. All show a characteristic hemodynamic response that goes up and comes down with the trial structure. Furthermore, most regions show a stronger response for more transformations and a stronger response on Day 1. This is quite characteristic of imaging results where disparate regions of the brain give quite similar responses to the material. Without a strong theory to guide one's expectations, one is in danger of missing the differences and concluding that the whole brain (or at least those regions that respond—not all regions in the brain respond to the task structure in this experiment) is reflecting a global response to the task. However, if one knows where to look, there are characteristic differences. Although this one experiment does not reveal all the differences in the behavior of all five regions, it does reflect many of the important differences that we have identified over our experiments. These are enumerated next.

First, and most apparent, in Figure 6–5a the motor region is giving basically the same hemodynamic response in all conditions. The effect of the slower conditions is to delay when that hemodynamic response occurs. This is what would be expected given a relatively strong understanding of what regions of the brain control the hand. Although the motor region is transparently giving a different response than the other regions (on both theoretical and observational grounds), its correlation with the BOLD responses in other regions averages .66. Thus even it might be confused with the other regions unless one had a theory to tell one where to look to find the relevant differences.

Two of the other regions have distinct signatures. The prefrontal region (Fig. 6–5b) is distinguished by the very weak response it generates in the case of 0 steps. According to the model, this case involves some brief retrievals of instructions but no retrieval of number facts. We have often modeled this condition by assuming no retrieval and predicted a flat function but a slight rise can be discerned. The striking feature of the anterior cingulate (Fig. 6–5c) is that there is almost no effect of learning whereas there is a robust effect of number of steps on magnitude of the response. The goal component in ACT–R is engaged in maintaining state at points where the system is engaged in a retrieval of an arithmetic fact (this is because the retrieval buffer cannot be used to hold the next step of instruction). Every time it engages in retrieval of an arithmetic task it must note this so that it will wait for the fact before going on. Once the fact is retrieved it must reset the state so that it can proceed with solving the equation. Thus, the number of retrieval operations is one factor influencing the number of state-setting operations in the goal buffer. The number of arithmetic retrievals changes in this experiment with the number of steps in solving the equation since each step requires retrieval of a fact. However, there is little reduction in these retrievals with practice. In principle, with enough practice they would eventually drop out but there are so many individual facts that they just do not repeat enough in equation solving.

The other two regions (the parietal in Fig. 6–5d and the caudate in Fig. 6–5e) can be distinguished from the other three regions because they lack the features that identify the other three. However, there is little difference in the response that we see in these two regions. They approximately reflect the average response of all the areas, showing substantial effects of both number of steps of transformations and days of practice. The caudate is fit according to the number of rules that fire, which naturally increases with steps and decreases with days. The parietal is fit according to number of mental re-representations of the equation, which also increases with steps and decreases with days. There is a subtle difference between the two with the caudate showing a relatively larger effect of days and the parietal showing a relatively larger effect of steps. We find differences between these two regions in experiments that vary modality of presentation from visual to aural with the parietal responding less to auditory presentation than visual and the caudate responding more (Sohn et al., 2005). Note that in the comparisons of Figures 6–5d and 6–5e, the caudate gives a relatively weak response and has a poorer signal-to-noise ratio. This is unfortunate because according to the theory it should be the one region that is involved in all cognitive tasks, reflecting number of production rules fired.

Assessing Goodness of Fit

The figures contain measures of correlations between the predictions and observed behavior. These are averaged over either days or operations but Table 6–2 gives correlations among all 108 points for each region. Although this is a conventional measure of quality of fit, it has a number of problems. For instance, correlation is only sensitive to whether the shapes match up and not to whether the actual predicted numbers match up.

The quantitative correspondence can be assessed by the chi-square statistics in the table, which measure the degree of mismatch against the noise in the data. They are calculated as

$$\chi^2 = \frac{\sum_i (\hat{X}_i - \overline{X}_i)^2}{s^2_{\overline{X}}}$$

where the denominator is estimated from the interaction between conditions and participants. This has 105 degrees of freedom, calculated as 108 minus the 3 parameters estimated for the BOLD function. By this measure, all of the areas are being modeled as well as can be expected because they all yield nonsignificant chi-squares (it would have to be 130 or greater to be significant at the .05 level).

Table 6–3 reports the outcome of trying to fit each module to each region's activation profile and calculating a chi-square measure of misfit. With 105 degrees of freedom, the 5-percentile tails for the chi-square distribution are at 82 and 130. As we noted with respect to Table 6–2, all the modules give acceptable fits (less than 130) to their ascribed regions. A few other modules give acceptable fits to other regions although not as good. In particular, the modules other than the manual module all give approximately equal fits to the parietal and caudate regions. As noted, these regions approximately show the average response of all the regions.

Table 6–3
Chi-Square Measure of Fits between Regions and Modules

	Motor	Prefrontal	Cingulate	Parietal	Caudate
Manual	88.93	452.05	724.66	426.40	333.89
Retrieval	493.22	82.60	350.32	101.88	133.13
Goal	255.91	194.94	123.27	171.74	111.01
Imaginal	384.66	125.66	210.47	95.21	101.82
Procedural	347.05	163.76	286.28	114.93	81.03

CONCLUSIONS

We should note that there is no reason why such data and methodology should be limited to testing the ACT–R theory. Many other information-processing theories could be tested. The basic idea is that the BOLD response reflects the duration for which various cognitive modules are active. The typical additive-factors information-processing methodology has studied how manipulations of various cognitive components affect a single aggregate behavioral measure like total time. If we can assign these different components to different brain regions, we have essentially a separate dependent measure to track each component. Therefore, this methodology promises to offer strong guidance in the development of any information-processing theory.

ACKNOWLEDGMENTS

This research was supported by the NSF Grant ROLE: REC—0087396 and ONR Grant N00014–96–1–0491.

REFERENCES

Anderson, J. R. (in press). Using brain imaging to guide the development of a cognitive architecture. In W. E. Gray (Ed.), *Integrated models of cognitive systems.* New York: Oxford University Press.

Anderson, J. R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., & Qin, Y. (2004) An integrated theory of mind. *Psychological Review, 111,* 1036–1060.

Anderson, J. R., Qin, Y., Sohn, M-H., Stenger, V. A., & Carter, C. S. (2003) An information-processing model of the BOLD response in symbol manipulation tasks. *Psychonomic Bulletin & Review, 10,* 241–261.

Anderson, J. R., Qin, Y., Stenger, V. A., & Carter, C. S. (2004). The relationship of three cortical regions to an information-processing model. *Cognitive Neuroscience, 16,* 637–653.

Anderson, J. R., Taatgen, N. A., & Byrne, M. D. (2005). Learning to achieve perfect time sharing: Architectural implications of Hazeltine, Teague, & Ivry. *Journal of Experimental Psychology: Human Perception and Performance, 31,* 749–761.

Baddeley, A. D. (1986). *Working memory.* Oxford, England: Oxford University Press.

Boyton, G. M., Engel, S. A., Glover, G. H., & Heeger, D. J. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. *Journal of Neuroscience, 16,* 4207–4221.

Dale, A. M., & Buckner, R. L. (1997) Selective averaging of rapidly presented individual trials using fMRI. *Human Brain Mapping, 5,* 329–340.

Glover, G. H. (1999). Deconvolution of impulse response in event-related BOLD fMRI. *NeuroImage, 9,* 416–429.

Qin, Y., Anderson, J. R., Silk, E., Stenger, V. A., & Carter, C. S. (2004). The change of the brain activation patterns along with the children's practice in algebra equation solving. *Proceedings of National Academy of Sciences, 101,* 5686–5691.

Qin, Y., Sohn, M.-H., Anderson, J. R., Stenger, V. A., Fissell, K., Goode, A., & Carter, C. S. (2003). Predicting the practice effects on the blood oxygenation level-dependent (BOLD) function of fMRI in a symbolic manipulation task. *Proceedings of the National Academy of Sciences of the United States of America, 100*(8), 4951–4956.

Sohn, M.-H., Goode, A., Stenger, V. A, Carter, C. S., & Anderson, J. R. (2003). Competition and representation during memory retrieval: Roles of the prefrontal cortex and the posterior parietal cortex. *Proceedings of National Academy of Sciences, 100,* 7412–7417.

Sohn, M.-H., Goode, A., Stenger, V. A, Jung, K.-J., Carter, C. S., & Anderson, J. R. (2005). An information-processing model of three cortical regions: Evidence in episodic memory retrieval. *NeuroImage, 25,* 21–33.
Taatgen, N. A. (2005) Modeling parallelization and speed improvement in the skill acquisition: from dual tasks to complex dynamic skills. *Cognitive Science, 29,* 421–456.
Taatgen, N. A., & Anderson, J. R. (2002). Why do children learn to say "Broke"? A model of learning the past tense without feedback. *Cognition, 86,* 123–155.

CHAPTER
7

Remembering Images

Stephen M. Kosslyn
Harvard University

Gordon H. Bower had a profound influence on my eventual academic fate, well before I actually met him. When I arrived at Stanford in 1970, Gordon was on sabbatical (teaching in Austria, as I recall). One of the other faculty members gave me a preprint of an article he wrote, "Mental Imagery and Associative Learning," later to appear in the volume edited by Gregg (1972). In that paper, a single line struck me right between the eyes (I vividly remember reading it, late at night in the original Stanford coffee house). He said something to the effect that "If images are like pictures, and can be scanned and otherwise inspected ... ." This throwaway thought from Gordon (which he valued so little that he deleted it from the final version) instantly led me to have two ideas. First, I realized that—exactly analogous to the now-famous "mental rotation" experiments of Lynn Cooper, Jackie Metzler, and Roger Shepard (see Shepard & Cooper, 1982)—I could measure how much time people required to respond when they had to scan different distances across a drawing they were visualizing. If Gordon were correct, more time should be required to scan greater distances across an imaged object. Second, I also realized that I could use such response times more generally, as a kind of "mental tape measure," to assess structural properties of the underlying representation. The thought was that if image representations (the short-term memory representations, which somehow give rise to the experience of "seeing with the mind's eye") were in some sense pictorial, then space in the representation should embody actual space. If so, I conjectured, then the time to