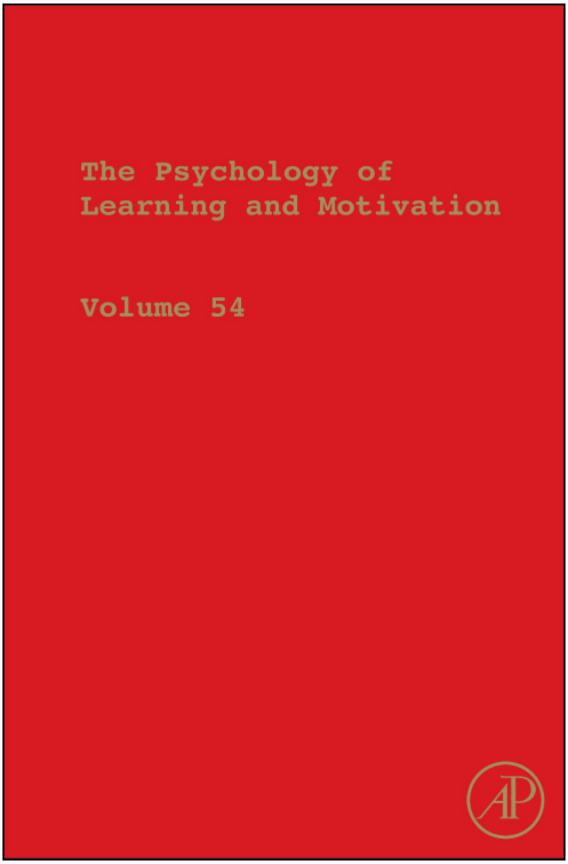

From: Gordon D. Logan and Matthew J. C. Crump, Hierarchical Control of Cognitive Processes: The Case for Skilled Typewriting. In Brian Ross, editor: *The Psychology of Learning and Motivation,* Vol. 54, Burlington: Academic Press, 2011, pp. 1-27.
ISBN: 978-0-12-385527-5

CHAPTER ONE

Hierarchical Control of Cognitive Processes: The Case for Skilled Typewriting

Gordon D. Logan *and* Matthew J. C. Crump

Contents

Psychology of Learning and Motivation, Volume 54
ISSN 0079-7421, DOI: 10.1016/B978-0-12-385527-5.00001-2

Abstract

The idea that cognition is controlled hierarchically is appealing to many but is difficult to demonstrate empirically. Often, nonhierarchical theories can account for the data as well as hierarchical ones do. The purpose of this chapter is to document the case for hierarchical control in skilled typing and present it as an example of a strategy for demonstrating hierarchical control in other cognitive acts. We propose that typing is controlled by two nested feedback loops that can be distinguished in terms of the factors that affect them, that communicate through intermediate representations (words), that know little about how each other work, and rely on different kinds of feedback. We discuss hierarchical control in other skills; the relation between hierarchical control and familiar concepts like automaticity, procedural memory, and implicit knowledge; and the development of hierarchical skills. We end with speculations about the role of hierarchical control in everyday cognition and the search for a meaningful life.

1. Introduction

The idea that cognition is controlled hierarchically is ubiquitous but enigmatic. Hierarchical control was a critical issue in the cognitive revolution against the behaviorists in the 1950s (Lashley, 1951; Miller, Galanter, & Pribram, 1960) and it remains a common feature in modern theories of executive control in cognitive science and cognitive neuroscience (Badre, 2008; Cooper & Shallice, 2000; Logan & Gordon, 2001; Miller & Cohen, 2001; Norman & Shallice, 1986). Nevertheless, the case for hierarchical control is weaker than it ought to be, with much of cognitive psychology focused on simple tasks that can be explained readily in terms of the stimulus–response associations that the behaviorists favored (Logan, 1988; Shiffrin & Schneider, 1977). Many of the more complex tasks can be explained by nonhierarchical control as well as by hierarchical control (see Botvinick & Plaut, 2004, 2006 vs. Cooper & Shallice, 2000, 2006a, 2006b). In this chapter, we present the case for hierarchical control in skilled typewriting. We present a theory that claims there are two hierarchically nested feedback loops in skilled typewriting and we present evidence that the two loops can be distinguished by factors that selectively influence them, that words are the interface between the two loops, that the two loops share little knowledge beyond words, and that the two loops rely on different feedback. We offer this analysis as a method for making the case for hierarchical control in other instances of cognitive control, and we draw implications from our analysis of typewriting to other issues and topics in cognitive control of thought and action.

2. What is Hierarchical Control?

2.1. Hierarchy

The case for hierarchical control must begin with definitions of hierarchy and control so we know what we are looking for when we examine typewriting. A hierarchy is a representation that has at least two levels with a one-to-many mapping of elements in the higher level to elements in the lower level (Markman, 1999; Novick & Hurley, 2001). The texts that typists type satisfy this definition: Texts are made of paragraphs, paragraphs are made of sentences, sentences are made of words, and words are made of letters. Typists' psychological representations of the texts must reflect this structure, so typing is driven by hierarchical representations. However, the debate about hierarchical control is not about hierarchical representation, but rather, about whether the processes that operate on the representations are also hierarchical. Hierarchical representation does not imply hierarchical processing. The same processes could operate at different levels of a hierarchical representation. For example, Schneider and Logan (2006, 2007) had subjects perform a sequence of tasks and argued that the sequence and the tasks were represented at different levels of a hierarchy but suggested that the same memory retrieval processes operated on both levels (also see Botvinick & Plaut, 2004). The case for hierarchical control requires more than demonstrating hierarchical representations.

2.2. Control

A process is controlled if it is willfully directed toward the fulfillment of a goal (Logan, 1988; Miller et al., 1960). "Willfully" is a difficult component of this definition. Some argue that a process is willful if it is chosen voluntarily (Arrington & Logan, 2004). Others argue that a process is willful if it can be interrupted on demand (Logan, 1982; Logan & Cowan, 1984). However, there is no consensus on conceptual or operational definitions of will (Wegner, 2002). "Goal directed" is an easier component and we will focus on it in this chapter. We assume that typing is willful. People do not type by accident.

Miller et al. (1960) expressed this definition of control in a generic feedback loop called a TOTE unit. *TOTE* stands for "*t*est, *o*perate, *t*est, *e*xit," which involves a comparison of the current state with the goal state (test), followed by the execution of an operation intended to reduce the difference between the current state and the goal state (operate), followed by another comparison of the current state with the goal state (test). If the current state matches the goal state, the task is completed, and the system goes on to pursue other goals (exit). If the current state does not match the goal state, the task is not yet complete, and the operate phase is engaged again, iterating until the goal state is attained. For example, the finger movements in typing the letter "T" can be described by a TOTE that compares the goal state ("is the index finger above the T key?") with the current state of the finger ("the index finger is above the F key"), resulting in successive operations (movements toward the T key) until the goal state is satisfied (the index finger is above the T key and ready to strike it). TOTE units embody control because the operations are directed toward the attainment of a goal.

2.3. Hierarchical Control

Miller et al. (1960) offered TOTE units as a cognitive or cybernetic alternative to the generic stimulus–response bond in behaviorist psychology. TOTE theory was a conceptual advance over stimulus–response bonds because TOTEs have more structure (see Chomsky, 1959). They can be concatenated to create extended chains of complex behavior, which Miller et al. called plans. For example, the plan for typing the letter "T" can be described by a sequence of two TOTEs, in which the first moves the finger from its current location to the location above the T key, as described in the earlier example, and the second depresses the key (the test is "is the key depressed?" and the operation is "push the index finger down"). More importantly for our purposes, TOTEs can be nested hierarchically. The operate phase of one TOTE can be replaced by another TOTE or a series of TOTEs that describe the details of the operation. Thus, the two-TOTE plan for moving the finger and depressing the key can be viewed as the operate phase of a superordinate TOTE that types the letter.

The operations in the subordinate TOTEs need not resemble the operations in the superordinate TOTE, so hierarchical processes (nested operations) can be distinguished from hierarchical representations (nested goals). Moreover, TOTEs can be modular and informationally encapsulated (Fodor, 1983): a superordinate TOTE need not be aware of the operations in its subordinate TOTEs. The superordinate TOTE only needs to know that the subordinate TOTEs moved the current state toward the superordinate goal state. These features are important in typewriting, so we will adopt the idea of nested TOTEs—nested feedback loops—as our generic definition of hierarchical control.

2.4. The Case for Hierarchical Control, 0.0

The case for hierarchical control in any domain requires several demonstrations. First, there must be at least two levels of processing that can be distinguished by manipulations of experimental factors. Selective influence by itself is not strong evidence for hierarchical processing because it is also found for separate processes that are not arranged hierarchically (e.g., Sternberg, 1969). Nevertheless, processes that cannot be distinguished from one another cannot be said to be organized hierarchically, so this is a necessary but not sufficient step in developing the case. To make this part of the case, we review studies of selective influence in typewriting.

Second, the two levels of processing must operate on different kinds of information that exist at different levels of an informational hierarchy, with the higher level operating on higher-level information than the lower level. One way to make this part of the case is to demonstrate an informational hierarchy and show that information at some intermediate level serves as the interface between levels of processing. The higher level deals with units at this intermediate level and higher; the lower level deals with units at this intermediate level and lower. To make this part of the case, we review studies that show the role of words as the interface between language processes and motor control in typewriting.

Third, the two levels of processing must divide the intellectual labor required to perform the task in a way that is consistent with hierarchical processing. The higher level must deal with larger structures and broader goals than the lower level. Moreover, the higher level should not know much about what the lower level is doing. The higher level should issue commands and determine whether they are executed without knowing the details of how the lower level executes the commands. The lower level should be informationally encapsulated, so the details of its processing are not available to the higher level (Fodor, 1983). To make this part of the case, we review studies that show that skilled typists know little about how they type.

Fourth, the two levels of processing must utilize different kinds of feedback, appropriate to the goals they address. The higher level should process feedback about higher-level goals and the lower level should process feedback about lower-level goals. Demonstrating sensitivity to different kinds of feedback requires identifying the goals that drive each level of processing and identifying the states of the cognitive system or states of the environment that signal progress toward those goals. To make this part of the case, we review studies that identify different levels of feedback that are utilized in skilled typewriting.

Finally, the levels of processing must be integrated in an overarching theory of the computations involved in successful performance. The first four steps mostly address whether two or more levels can be distinguished. The computational analysis specifies the relation between the levels, situating them in a larger system whose parts are individually necessary and jointly sufficient to perform the task at hand. This part of the case has already been made in typewriting: several computational theories have been proposed and tested (John, 1996; Rumelhart & Norman, 1982; Salthouse, 1986;

Wu & Liu, 2008). Our analysis of typewriting abstracts two hierarchically nested levels of processing from these previous theories. Our efforts at establishing the first four steps in the case are grounded in these more complex theories, which specify the hierarchical relation between the two levels in computational terms.

3. The Two-Loop Theory of Typewriting

Typewriting is a recent skill in human history. The QWERTY keyboard on which English speakers type was patented on July 14, 1868 by Christopher Latham Scholes. In 1876, the first book was created on a typewriter (*Tom Sawyer* by Mark Twain). By 1900, typewriters were in common use and Scholes was heralded as an emancipator for bringing women into the workplace as typists. From 1900 to 1980, typists were a small group of trained professionals who spent most of their time transcribing text—copy typing. The proliferation of personal computers in the 1980s brought typing to the masses. Now, typing is ubiquitous. Most American homes, businesses, and schools have computers, and most college-age people have strong typing skills. In 800 college-age typists we tested, mean typing speed was 68 words per minute (SD = 18; range = 21–126), similar to yesterday's professional typists. In 246 typists we surveyed, the average age at which they started typing was 10 years. They had typed for 10.8 years and currently spent 4.4 h per day on their computers. Seventy-eight percent had formal training in typing, averaging 44 weeks in duration. The nature of their typing was different from yesterday's professional typists: Only 9% of their typing involved transcription, 51% involved composition (e-mail, essays), and 40% involved other activities (instant messages, search engines, user ID, etc.). In addition, they averaged 40 text messages per day (range = 0–500).

Compared to most tasks studied in cognitive science and cognitive neuroscience, typewriting is a complex activity. It involves many successive responses that are strongly constrained. The input is constrained by language processing and the output is constrained by the QWERTY keyboard. The act of typing itself is constrained by the need for speed and accuracy. All of the characters in each word have to be typed as quickly as possible in the correct order (Lashley, 1951). These constraints engage a wide variety of processes, ranging from language comprehension and generation to hand and finger movements (Johns, 1996; Rumelhart & Norman, 1982; Salthouse, 1986; Shaffer, 1976; Wu & Liu, 2008).

Our research has been guided by a simple model of typewriting that divides the many processes into two nested feedback loops or TOTEs: an *outer loop* that begins with language comprehension or generation and ends with a series of words to be typed, and an *inner loop* that begins with a word to be typed and ends with a series of keystrokes (Logan & Crump, 2009). It is said that science carves nature at the joints. Our outer-loop inner-loop theory is intended to carve typewriting at a major joint, like severing an arm at the shoulder. Other theories address smaller joints, distinguishing more components and focusing on finer details of performance such as the timing of keystrokes (Gentner, 1987; Soechting & Flanders, 1992; Sternberg, Monsell, Knoll, & Wright, 1978; Terzuolo & Viviani, 1980; Viviani & Laissard, 1996) and the kinematics of finger movements (Flanders & Soechting, 1992; Gordon, Casabona, & Soechting, 1994; Soechting & Flanders, 1992). We do not deny the importance of these smaller joints. Instead, we argue that the distinction between the outer loop and the inner loop is important and general. It encompasses theories that address finer details and it leads to testable hypotheses and useful extensions to other instances of cognitive control.

In recent years, we have amassed evidence that makes a strong case for hierarchical control in typewriting. We have shown that the outer loop and inner loop are affected by different factors, that they communicate at the level of words rather than sentences or keystrokes, that the outer loop knows little about the workings of the inner loop, and that the two loops rely on different kinds of feedback. In the remaining pages, we describe the experiments that support these claims and suggest how our experimental procedures might be used to generalize these claims to other instances of cognitive control.

4. Distinguishing the Outer Loop and the Inner Loop

A straightforward way to distinguish the outer loop and the inner loop is to examine response times and interkeystroke intervals in discrete typewriting tasks, in which a single word is presented and subjects are instructed to type it as quickly as possible. Response time is the interval between the onset of the word and the registration of the first keystroke. In our theory, response time measures the duration of outer-loop processes that identify the word on the screen and pass it to the inner loop. Response time must also include the duration of inner-loop processes that prepare and execute the first keystroke. Thus, response time measures the duration of both loops.

Interkeystroke interval is the interval between successive keystrokes. It is tempting to think of interkeystroke interval as a response time, but we should not yield to that temptation (Lashley, 1951). High-speed films of typists typewriting show that finger movements often occur in parallel, and the finger movement for one keystroke often begins before the finger movement for the preceding keystroke ends (Flanders & Soechting, 1992). Thus, it is better to think of interkeystroke interval as measuring the finishing times of concurrent, temporally overlapping processes. In our theory, interkeystroke interval measures processing in the inner loop. The inner loop prepares and executes successive keystrokes, and interkeystroke interval measures the rate at which this occurs.

4.1. Distinguishing the Loops by Selective Influence

In our theory, outer and inner loops can be distinguished by experimental factors that selectively influence response time and interkeystroke interval (Sternberg, 1969). For example, Logan and Zbrodoff (1998) ran a Stroop (1935) task with typewritten responses and found that the congruency of the color and the word affected response time but not interkeystroke interval. This suggests that congruency affects the choice of the word to type, which is the business of the outer loop, but not the execution of the keystrokes in the word, which is the business of the inner loop. Similarly, Logan (2003) had typists type words presented on the left or right side of a central fixation point and found Simon-type interference for words typed entirely with the left or right hand. The congruency of stimulus and hand locations affected response time but not interkeystroke interval, suggesting that congruency affects the choice of the word to type but not the execution of keystrokes.

Some of the factors that are important in typewriting do not selectively influence a single stage of processing, and this limits the utility of selective influence in distinguishing the outer loop from the inner loop. For example, Crump and Logan (2010a) examined repetition priming in typewriting, presenting words several times and comparing performance on repeated words with performance on new words that had not been presented in the experiment. We found that repetition reduced both response time and interkeystroke interval, suggesting that repetition affects only inner-loop processing, which contributes to both measures. However, repetition could affect both inner- and outer-loop processing. Other data suggest that repetition priming facilitates perceptual and conceptual processing (Logan, 1990), which is part of the outer loop, so we can reject the hypothesis that repetition affects only inner-loop processing, but the data themselves cannot distinguish between the hypotheses. The problem is that repetition priming does not selectively influence outer- and inner-loop processing.

More generally, factors that do not selectively influence the outer and inner loops cannot be used to distinguish between them (Sternberg, 1969). This limits the utility of the contrast between response times and interkeystroke intervals as a way to distinguish between outer and inner loops. Even more generally, the problem of selective influence shows the limits of defining processes in terms of the factors that affect them (Garner, Hake, & Eriksen, 1956; Logan, Coles, & Kramer, 1996; Sternberg, 1969).

We prefer to define processes in terms of the computations they require. Finally, the method of selective influence works only with discrete typing tasks in which response time can be defined meaningfully. It does not work with continuous typing tasks, which only provide information about inter-keystroke intervals.

4.2. The Case for Hierarchical Control, 1.0

Studies of selective influence suggest that the outer and inner loops are affected by different factors, but they also indicate that the two loops are sometimes affected by the same factors. Thus, studies of selective influence suggest that two separate processes underlie typewriting but do not provide strong evidence for hierarchical control.

5. Words as the Interface Between Outer and Inner Loops

There is abundant evidence that words are important units in typewriting. Manipulations of units larger than the word have little effect on typewriting: scrambled sentences are typed as quickly as intact ones (Fendrick, 1937; Gentner, Larochelle, & Grudin, 1988; Shaffer & Hardwick, 1968). However, manipulations of units smaller than the word have a strong effect: scrambled words and random letter strings are typed more slowly than intact words (Fendrick, 1937; Gentner et al., 1988; Shaffer & Hardwick, 1968). Studies that manipulated preview of the text to be typed found that increasing preview from 1 to 8 characters (approximately one word) increased typing speed, but further increases in preview up to 40 characters produced no further increase (Hershman & Hillix, 1965; Shaffer, 1973). More generally, reading (250–350 words/min) and speaking (120–200 words/min) are much faster than typing (50–100 words per minute; Rayner & Clifton, 2009), so the effects of units larger than the word may be absorbed while the outer loop waits for the inner loop to finish the current word.

We interpret these effects as evidence for the proposition that words are the interface between the outer and inner loops. The outer loop generates a series of words to be typed, through language generation or comprehension, and passes them one by one to the inner loop. The inner loop takes each word and translates it into a series of letters to be typed, translates the letters into a series of keystrokes, and executes them one by one on the keyboard. The hierarchical relationship between words and letters is mirrored in the hierarchical relationship between the outer and inner loops: One word in the outer loop corresponds to several letters in the inner loop.

This perspective predicts that words will activate their constituent letters and the corresponding keystrokes in parallel. We tested this prediction in three ways.

5.1. Words Prime Constituent Letters in Parallel

Crump and Logan (2010b) developed a priming technique, in which typists were given five- or seven-letter words as primes. On some trials, the prime was followed by another copy of itself, and typists typed the word. On other trials, the prime was followed by a single-letter probe, and typists typed the letter. The letter was either the first, middle, or last letter in the prime word or a randomly chosen letter that did not appear in the prime. We found that typists typed the single-letter probes faster if they appeared in the prime word than if they did not, suggesting that the prime activated all of its constituent letters. Priming was greater for the first letter than for the middle and last letters, but there was no difference in priming for the middle and last letters. This suggests that there is an advantage to the first letter, perhaps because it must be typed first, but the middle and last letters were activated similarly, which is consistent with the hypothesis that all letters in the word are activated in parallel.

The priming words were presented visually, so it is possible that the priming effects were perceptual rather than motoric: seeing the letter in a word may have sped up perceptual processing of single-letter probes. We addressed this possibility in two ways. First, we presented auditory primes rather than visual ones and found the same effects: response time to single-letter probes was faster when the letters were part of the prime word than when they were not. Priming was greater for the first letter than for middle and last letters, but middle and last letters were primed equally. This shows that visual presentation of the primes is not necessary to produce within-word priming effects. Second, we presented strings of consonants as primes. The strings of consonants should be represented as several units in the outer loop, and we assumed that only the first unit would be passed to the inner loop. Consistent with this assumption, we found priming for the first letter of the string but no priming for the middle or last letter. This shows that visual presentation of the primes is not sufficient to produce within-word priming effects.

5.2. Words Activate Spatial Locations of Constituent Letters

Logan (2003) had typists type single words presented to the left or right of a central fixation point. There were three types of words: *LEFT/right* words were constructed such that all letters were typed entirely in the left or right hand (e.g., *rest*), *LEght/riFT* words were constructed such that the first two letters were typed with one hand and the remaining letters were typed with the other (e.g., *swim*), and *Light/rEFT* words were constructed such that the first letter was typed in one hand and the remaining letters were typed in the other hand (e.g., *dump*). If the words activated their constituent keystrokes in parallel, then there should be a Simon-type effect (Simon & Small, 1969), in which words presented on the same side of the screen as their constituent keystrokes should be typed faster than words presented on the opposite side of the screen. Importantly, the Simon effect should be stronger when more letters were typed in one hand; distributing letters across the hands should weaken the effect. Consistent with this prediction, the Simon effect was stronger with LEFT/right and LEght/riFT words than with Light/rEFT words. The effect of letters beyond the first suggests that the spatial locations of all the letters were activated in parallel.

5.3. Words Activate Motor Representations of Constituent Letters

Logan, Miller, and Strayer (2011) presented LEFT/right, LEght/riFT, and Light/rEFT words centrally and measured the lateralized readiness potential in the electroencephalogram while typists typed them. The lateralized readiness potential is the difference in electrical potential between electrodes located over the left and right motor cortex (C3 and C4 in the international 10–20 system; Jasper, 1958). It reflects the process of response selection, measuring the growth in activation of motor representations of responses that are about to be executed (Coles, 1989). Logan et al. focused on the early part of the lateralized readiness potential time-locked to the first keystroke to measure activation of motor representations while the first response was selected. If the constituent keystrokes of a word are activated in parallel, then the amplitude of the lateralized readiness potential for the first keystroke should decrease systematically as progressively more keystrokes are activated in the opposite hand. Thus, the lateralized readiness potential should be greater for LEFT/right words than for LEght/riFT words, and greater for LEght/riFT words than for Light/rEFT words. However, if the constituent keystrokes are activated in series, then there should be no effect of subsequent letters on the lateralized readiness potential for the first letters of LEFT/right, LEght/riFT, and Light/rEFT words. The results were consistent with parallel activation: the amplitude of the lateralized readiness potential for the first keystroke decreased monotonically from LEFT/right to LEght/riFT to Light/rEFT words.

5.4. The Case for Hierarchical Control, 2.0

The idea that words are the interface between the outer loop and the inner loop provides strong evidence for hierarchical control. Outer-loop processes operate on structures larger than the word in comprehending and generating language, and that processing results in a series of words that are passed one at a time to the inner loop. Inner-loop processes translate the word into letters, motor plans addressed to keyboard locations, and ultimately, keystrokes. Three lines of evidence provide strong support for the notion that outer-loop words translate to inner-loop motor plans in parallel, reflecting the one-to-many mapping that is characteristic of a hierarchy. Together with the evidence that the two loops may be influenced selectively by experimental factors, this evidence strengthens the case for hierarchical control of typewriting.

6. The Inner Loop is Informationally Encapsulated

Our outer-loop inner-loop theory assumes that the intelligence required to type text is divided between the loops. The outer loop is concerned with language generation and comprehension, and its job is to produce a string of words to be typed. The inner loop is concerned with translating words into letters, finger movements, and keystrokes, and its job is to produce a series of keystrokes on the keyboard. With this division of labor, the outer loop does not need to know what the inner loop does. It only needs to provide the inner loop with words to be typed, one at a time. We believe that the outer loop does not know what it does not need to know. The inner loop is informationally encapsulated (Fodor, 1983), so the outer loop does not know how the inner loop does what it does. In this respect, typing is like other skills that exhibit paradoxical dissociations between explicit knowledge and implicit knowledge: practitioners of high-level skills know how to perform very well but do not know much about how they do it (Beilock & Carr, 2001; Beilock, Carr, MacMahon, & Starkes, 2002; Terzuolo & Viviani, 1980).

6.1. The Outer Loop Does Not Know Which Hand Types Which Letter

The outer loop knows which words must be typed and it is able to spell the words, but it usually does not break words down into letters before passing them to the inner loop. The previous section summarized the evidence suggesting that the outer loop passes whole words to the inner loop, and the inner loop breaks the words down into letters and assigns the letters to particular hands and particular keyboard locations. This division of labor suggests that the outer loop does not know which hand types which letter but the inner loop does (it must because it types letters correctly).

Logan and Crump (2009) showed that the outer loop does not know which hand types which letter by having typists type only the letters assigned to one hand. This was very disruptive, as you can confirm for yourself by typing only the right-hand letters in this sentence. In one experiment, we had typists type whole paragraphs and told them to type only the left-hand letters (or only the right-hand letters). With these instructions, typists typed 14 words per minute and made errors on 33% of the words. When the same typists typed the same texts under instructions to type normally (i.e., to type all letters), their typing speed was 80 words per minute and they made errors on only 6% of the words. Similar results were found when typists typed single words preceded by a cue that told them whether to type the words in one hand (LEFT or RIGHT) or to type all the letters (WHOLE). The requirement to type only the letters in one hand increased the response time by 454 ms (a 54% increase), interkeystroke interval by 153 ms/keystroke (a 104% increase), and error rate by 16% (a 304% increase). A control experiment in which the letters to be typed were cued by color ("type only the red letters") produced no disruption, suggesting that the difficulty lay in discovering which hand typed which letter.

These gargantuan disruptions are paradoxical: They suggest that skilled typists do not know which hand types which letters, yet they choose the correct hand 5–6 times/s in normal typing. Our two-loop hypothesis resolves the paradox by proposing that the inner loop is informationally encapsulated. In order to discover which hand types which letter, the outer loop must observe the inner loop's output. To type letters from only one hand, the outer loop must slow the inner-loop's cycle time so that it has time to observe which hand was selected and inhibit the keystroke if necessary. Other investigators have shown similar disruptions from drawing attention to the details of performance in other skills (Beilock & Carr, 2001; Beilock et al., 2002). Our research provides one explanation for the disruption: performance must slow down so the outer loop can observe the details, and that disrupts timing and the fluency of performance.

6.2. The Outer Loop Does Not Know Where Letters Are on the Keyboard

Another paradox in typing skill concerns knowledge of where the letters are located on the keyboard. Our intuitions as typists tell us we have little explicit knowledge of letter location, yet our fingers find the correct locations five to six times per second. We suggest that this is also a consequence of the division of labor between outer and inner loops and another example of encapsulated inner-loop knowledge. In this case, there may be stronger motivation for encapsulating knowledge about letter location in the inner loop: the locations of letters in words rarely correspond to the locations of letters on the keyboard (e.g., the letters in *pout* are in opposite left-to-right order on the screen and on the keyboard). Encapsulating knowledge about letter location and communicating information about letters through the intermediary representations of words may reduce the costs of stimulus–response incompatibility.

Liu, Crump, and Logan (2010) had skilled typists make explicit judgments of the relative locations of keys on the keyboard, using standard procedures from the literature on spatial memory (McNamara, 1986; Stevens & Coupe, 1978). We asked typists to imagine that they were standing on a key on the keyboard (e.g., F) facing a particular direction (e.g., the space bar) and then point to the location of another letter (e.g., W), indicating the direction with a mouse. When typists judged the relative direction with no keyboard in the testing room and so could only rely on explicit (outer loop) knowledge of keyboard locations, the absolute angular error (the unsigned difference between the actual angle and the judged angle) was 47°. A control group who made the same judgments when it could see the keyboard in the room had an absolute angular error of 28°. Another control group who made the same judgments after typing the letters on a keyboard covered by a box to prevent it from seeing the keys had an absolute angular error of 23°.

In a second experiment, Liu et al. (2010) had skilled typists use a mouse to drag a depiction of a key to its location relative to another key. The absolute angular error was smaller overall, but it was still greater for typists who had to imagine the keyboard (29°) than for typists who could look at the keyboard (17°) or type the letters on a keyboard covered by a box to block vision (14°). Absolute distance error—the unsigned distance between the correct location and the location to which they dragged the key—was 81 mm for typists who imagined the keyboard, 58 mm for typists who saw the keyboard, and 54 mm for typists who typed the letters on a keyboard covered by a box.

This experiment allowed us to compare the relative precision of explicit and implicit knowledge of letters on the keyboard. In the explicit judgments, the standard deviation of the distance between the correct and judged location (signed distance error) was 28 mm. To estimate the standard deviation of the distance between correct and judged location in implicit judgments, we assumed that location was represented implicitly as a bivariate normal distribution centered on the key and that the percentage of correct responses reflected the proportion of the distribution that fell within the boundaries of the key. Typists who imagined the keyboard typed 93% of the keystrokes correctly in a typing test, so we assumed that 93% of the distribution fell within the boundaries of the key. We estimated a z score for the radius of the bivariate normal distribution by taking the square root of the 93rd quantile of a chi-square distribution with 2 degrees of freedom. The quantile was 5.2. The square root corresponds to a z score of 2.28, which corresponds to a standard deviation of 4.2 mm for implicit knowledge of key location. Thus, explicit knowledge of keyboard location is (28/4.2=) 6.7 times less precise than implicit knowledge of keyboard location. This analysis underestimates the difference in precision because it assumes that all typing errors are misplaced keystrokes, and that is not the case. Insertions, deletions, and transpositions are more common (Lessenberry, 1928; Logan, 1999). The standard deviation of implicit knowledge of key location may be much smaller than 4.2 mm.

6.3. The Case for Hierarchical Control, 3.0

Two lines of research provide strong evidence that the outer loop does not know what the inner loop is doing. This informational encapsulation is characteristic of the division of labor in hierarchical control, in which the higher level issues commands and notes that they are executed but does not know the details of how the commands are executed (Beilock & Carr, 2001; Beilock et al., 2002; Fodor, 1983). Together with the evidence for selective influence and limited word-level communication between loops, the evidence for informational encapsulation makes the case for hierarchical control even stronger.

7. The Outer Loop and the Inner Loop Rely on Different Feedback

Feedback loops are defined in terms of the goals they are intended to achieve, the operations they carry out in order to achieve them, and the feedback they evaluate to determine whether or not the operations were successful. TOTE units evaluate feedback in the test phase by comparing the goal state and the current state. Thus, feedback can be identified by discovering the goal state and the current states—mental or physical—to which the TOTE is sensitive. Different TOTEs should be sensitive to different feedback. Nested TOTEs should be sensitive to a finer grain of feedback than the superordinate TOTEs in which they are nested. This suggests that the outer and inner loops should be sensitive to different kinds of feedback, and the feedback for the inner loop should be finer-grained than the feedback for the outer loop. We have two lines of evidence supporting this proposition.

7.1. The Inner Loop Relies on the Feel of the Keyboard

Crump and Logan (2010c) asked whether the inner loop relied on different feedback than the outer loop, assessing the role of the "feel" of the keyboard in supporting skilled typing (also see Gordon & Soechting, 1995; Terzuolo

& Viviani, 1980). Our research was motivated in part by dueling intuitions from our own experience as typists about the role of the keyboard in skilled typing. On the one hand, we find it very difficult to type in the air or on a tabletop without a keyboard to support our typing. This suggests that the feel of the keyboard is essential. On the other hand, we believe that typing is a general skill that can be transferred readily to new keyboards. Otherwise, we would be reluctant to buy new computers or switch between keyboards on desktops and laptops. However, commercial keyboards are very similar, with keys of similar sizes at similar distances in similar layouts. Transfer outside of these familiar parameters may be difficult. Ultimately, the question is empirical, so we designed an experiment to test it.

Crump and Logan (2010c) had typists type words on a regular keyboard and on "deconstructed" keyboards that successively removed familiar tactual and proprioceptive cues. First, we removed the keys from a regular keyboard and had typists type on the rubber buttons underneath them. This removes the usual tactual feedback while still providing the resistance or "give" of the regular keyboard, which the rubber buttons provide. Relative to the regular keyboard, the response time slowed by 144 ms (21%), interkeystroke interval slowed by 117 ms/keystroke (75%), and error probability increased by 0.08 (62%). Then, we removed the buttons and had typists type on the flat plastic panel underneath the buttons, in which the circuitry is embedded. This removes the resistance of the keyboard as well as the feel of the keys. Relative to the regular keyboard, the response time slowed by 296 ms (43%), interkeystroke interval slowed by 321 ms/keystroke (207%), and error probability increased by 0.23 (175%). Finally, we tested the typists on a commercially available laser projection keyboard that projected a life-size image of the keyboard on a tabletop. Like the flat keyboard, the laser keyboard removes the feel of the keys and the resistance of the buttons. Relative to the regular keyboard, the response time increased by 323 ms (47%), interkeystroke interval increased by 160 ms/keystroke (103%), and error probability increased by 0.28 (207%). Similar effects were found in typing paragraphs. Typing speeds were 76, 52, 30, and 43 words per minute on the regular, button, flat, and laser keyboards, respectively. These large disruptions suggest that the feel of the keyboard is an important source of feedback that supports inner-loop performance.

As a further test of the importance of feedback from the keyboard, we had 61 typists place their fingers on a blank piece of paper as they would if they were resting on the home row, and we traced the outline of their fingertips. The outline was curved, following the natural contour of the fingertips, and not straight, as it would be if the fingers were resting on the keyboard. The mean discrepancy from a straight line was 12.5 mm, which is about two thirds of the distance between the home row and the top or bottom row. This suggests that the feel of the keyboard is important in maintaining the proper alignment of the fingers on the keys.

7.2. The Outer Loop Relies on the Appearance of the Screen

An important function of a feedback loop is to detect errors in performance. The processes that implement the operate phase of a TOTE may not always reduce the discrepancy between the current state and the goal state. Sometimes they increase it. In tasks like typewriting, operations that increase the discrepancy between the current state and the goal state produce errors such as typing the wrong letter, omitting a letter that ought to be typed, or typing the right letters in the wrong order (Lessenberry, 1928; F. A. Logan, 1999). Typists must detect these errors and correct them (Long, 1976; Rabbitt, 1978).

The two-loop theory of typewriting claims that the outer loop and inner loop process different kinds of feedback, which implies that the two loops detect errors in different ways. The outer loop generates a series of words to be typed and so should monitor the accuracy with which words are typed. We suggest that it monitors the computer screen for the appearance of the intended word. If the intended word appears as it should, the outer loop assumes that it was typed correctly and moves on to the next word. If the intended word does not appear as it should, the outer loop assumes it was an error and asks the inner loop to make the screen look right. The inner loop generates a series of keystrokes and monitors proprioceptive and kinesthetic feedback to ensure that the right fingers moved to the right locations and struck the right keys. If the movements match intentions, typing should remain fast and fluent. If there is a mismatch, typing should slow down or stop.

To test these claims, Logan and Crump (2010) had typists type single words and created mismatches between what typists actually typed and what appeared on the screen. We corrected errors that typists made, so the screen matched their intentions but their motor behavior did not, and we inserted errors that typists did not make, so their motor behavior matched their intentions but the screen did not. To measure outer-loop error detection, we had typists report whether or not they typed each word correctly. We assumed that the outer loop monitors the appearance of the screen and so would report corrected errors as correct responses and inserted errors as actual errors, showing cognitive illusions of authorship. Typists confirmed this prediction in an experiment with two alternative posterror responses (correct, error), calling corrected errors "correct" on more than 80% of the trials and inserted errors "error" on more than 70% of the trials, claiming authorship for the appearance of the screen even though it contradicted their motor behavior. In another experiment, we told typists we would correct some errors and insert some errors, and we gave them four alternative posterror responses (correct, error, corrected error, inserted error). We found cognitive illusions of authorship for corrected errors: typists were as likely to call them correct responses as corrected errors. We found no such illusion for inserted errors: typists called inserted errors "inserted errors" as often as they called actual errors "errors."

The inner loop was not susceptible to these cognitive illusions of authorship. We measured inner-loop error detection by assessing posterror slowing. People often slow down on trials after an error in choice response time tasks (Laming, 1968; Rabbitt, 1966) and typists slow down after erroneous keystrokes (Gordon & Soechting, 1995). We found posterror slowing for actual errors and corrected errors but no posterror slowing after inserted errors. Thus, the inner loop knew the truth behind the illusion of authorship.

The contrast between explicit error reports and posterror slowing suggests a dissociation between outer-loop and inner-loop error detection. Explicit reports of correct responses occurred both with actual correct responses, which produced no posterror slowing (since there were no errors), and with corrected errors, which produced posterror slowing. Explicit reports of erroneous responses occurred both with actual errors that exhibited posterror slowing and with inserted errors that exhibited no posterror slowing. This dissociation provides further support for our distinction between the outer loop and the inner loop.

7.3. The Case for Hierarchical Control, 4.0

Two lines of research provide strong evidence that the outer loop and inner loop rely on different kinds of feedback. Reliance on different feedback is strong evidence that the two loops engage in different computations, which supports the claim that they are different processes. Together with the evidence for selective influence, limited word-level communication, and informational encapsulation, the evidence for reliance on different feedback makes a strong case for hierarchical control in typewriting. Coupled with more detailed analyses of the computations involved inside the two loops (John, 1996; Rumelhart & Norman, 1982; Salthouse, 1986; Wu & Liu, 2008), our case for hierarchical control of typewriting is complete and compelling.

8. Beyond Typewriting

8.1. Hierarchical Control in Other Skills

Can the case for hierarchical control in typewriting be generalized to other skills? The answer depends on how unique typewriting is in the range of skills humans are capable of performing. Typing is like other skills in that proficiency is attained only after extensive practice. Our typists had 11 years of practice, logging in the 10,000 h necessary for truly expert performance (Ericsson, Krampe, & Tesch-Römer, 1993). Large amounts of practice may be necessary to develop the autonomy and modularity we see in inner-loop processing. Hierarchical control may be seen in other skills that involve similar amounts of practice.

Typewriting may be different from other skills in that it is grafted onto preexisting skills that are already well developed. Children learn to speak around 2 years of age and learn to read around 5 or 6 years of age. Their reading skills are grafted onto well-developed language skills, providing a new input modality. Our survey of typists suggests that children learn to type around 10 years of age, when language skills are quite sophisticated and reading skills are well developed. Typing skill is grafted onto these preexisting skills, providing a new output modality. We believe that this developmental history invites the development of modular typing skill, grafting a new inner loop onto a preexisting outer loop. Hierarchical control may be seen in other skills that graft new inputs or outputs onto preexisting skills. Grafting new skills onto old ones may be sufficient to develop hierarchical control but it may not be necessary. Speech production is controlled hierarchically, although it develops at the same time as language comprehension (Dell, 1986; Levelt et al., 1991).

Typewriting may be different from other skills in that it is the result of performance that matters and not the performance itself. The end product of typing is an external text that conveys an intended meaning. The effort involved in producing the product does not matter much as long as the product looks as it should. Thus, the outer-loop processes that generate the intended meaning need not be concerned with the inner-loop processes that translate it into keystrokes. Skill at playing music is different from typewriting in that the performance itself matters more than the plan that generates it. The expressive aspect of music results directly from the nuances of the physical interaction of the musician's effectors with the instrument. Guitar players evoke emotion with timing, vibrato, and bending and sliding notes (Juslin, Karlsson, Lindström, Friberg, & Schoonderwaldt, 2006). Piano players evoke emotion by varying timing and striking the keys gently or robustly (Repp & Knoblich, 2004; Shaffer, Clarke, & Todd, 1985). The outer loop is directly concerned with inner-loop processes to be sure they convey the intended emotion. Thus, musicians may be more aware of what their fingers are doing than typists are. A strong case can be made for hierarchical control in skilled musical performance (Palmer, 1997; Shaffer, 1982), but the inner loop may not be as modular and informationally encapsulated as it is in typewriting.

8.2. Hierarchical Control, Automaticity, Procedural Memory, and Implicit Knowledge

For better or for worse, a large amount of psychology is built around binary distinctions (Newell, 1973; Platt, 1964). Indeed, our distinction between the outer loop and the inner loop offers another one. Historically, three binary distinctions have been important in the psychology of skill: automatic and controlled processing (Shiffrin & Schneider, 1977), declarative and procedural memory (Cohen & Squire, 1980), and explicit and implicit knowledge (Roediger, 1990). How do these ideas relate to the idea of hierarchical control?

It is tempting to say that the outer loop is controlled and the inner loop is automatic. Indeed, the outer loop controls the inner loop and the inner loop is relatively autonomous. However, we believe there are automatic and controlled components in both loops. Our theory shifts the emphasis from whether processes are controlled to how processes are controlled. Both loops are controlled because they are willfully directed toward goals, but they control different things (relying on different feedback; Crump & Logan, 2010c; Logan & Crump, 2010). Both loops have automatic processes (e.g., lexical activation: Levelt, 1989; finger movements: Gordon et al., 1994), but the processes are controlled in that they serve larger goals (Bargh & Ferguson, 2000) and can be interrupted easily (Logan, 1982; Long, 1976; Rabbitt, 1978; Salthouse & Saults, 1987).

Similar considerations apply to the distinctions between declarative and procedural memory, and between explicit and implicit knowledge. The inner loop is a paradigm case of procedural memory, but the language comprehension and generation processes in the outer loop also involve procedural memory (Levelt, 1989). Much of the knowledge in the inner loop is implicit, but it can be made explicit easily (although typists may have to slow down to do so; Logan & Crump, 2009), and much of the knowledge in the outer loop is implicit as well. Indeed, skilled performers often have more explicit knowledge about their skill than novices do (Beilock & Carr, 2001), although that knowledge may not be used directly to control performance. As with controlled and automatic processing, the key question in skilled performance is not whether memory is declarative or procedural or whether knowledge is explicit or implicit, but rather, how declarative and procedural memory, and explicit and implicit knowledge support performance. All the three binary distinctions were developed to address simpler tasks than typewriting, which typically involved single responses to single stimuli. We should not expect distinctions developed for simple tasks to generalize transparently to complex tasks like typewriting. The distinctions may apply to simple components of complex skills, but the components must be organized and coordinated to produce complex behavior, and that may require new concepts and new distinctions like TOTE theory (Miller et al., 1960), two-loop theory (Logan & Crump, 2009), and computational theories of typewriting (John, 1996; Rumelhart & Norman, 1982; Salthouse, 1986; Wu & Liu, 2008). Complex tasks may require more complex explanations.

8.3. The Development of Hierarchical Control

Skills like typewriting are acquired, so the hierarchical control they entail must develop during skill acquisition. It is not clear how this happens. Bryan and Harter (1899) documented plateaus in learning skill at sending and receiving Morse code in telegraphers and argued that other complex skills should show similar plateaus (see LaBerge & Samuels, 1974). Improvements in the speed of sending and receiving messages occurred in steps as telegraphers learned letters, then words, and then phrases. Bryan and Harter interpreted the plateaus as indicating a hierarchy of habits, which implies hierarchical control, so we might see similar plateaus when other hierarchical skills are acquired. However, Keller (1958) reviewed skill acquisition studies that had been published in the years following Bryan and Harter and found little objective evidence for plateaus in the learning curves. Plateaus may be a subjective phenomenon.

Newell and Rosenbloom (1981) proposed a theory of skill acquisition for hierarchically structured tasks. They argued that lower-level chunks were smaller than higher-level chunks and would repeat more often in the task environment. Consequently, lower-level chunks should be learned faster than higher-level chunks. MacKay (1982) proposed a theory of skill acquisition based on strengthening connections that made the opposite prediction. He argued that learning was proportional to the difference between current strength and the maximum strength, and that this difference was greater for higher-level structures than for lower-level structures. Thus, there should be greater learning at higher levels of the hierarchy.

It is not clear what these theories would predict for typewriting. As we noted earlier, typing skill is grafted onto preexisting language skills, so the higher-level chunks and structures may already be learned before lower-level hand and finger skills are acquired. In learning one's first musical instrument, higher-level structures involving musical phrases, scales, and keys may develop at the same time as lower-level skills that connect one's effectors to the instrument.

Anderson (1982, 1987) proposed a theory of skill acquisition based on the idea of collapsing a series of steps into a single step, which he called "composition." In his theory, composition can operate at a single level, in which the steps in the series simply drop out (e.g., instead of counting to determine that $2 + 3 = 5$, we simply remember the answer). It can also describe hierarchical skills, in which control of a sequence of actions is passed down from the cognitive system to the motor system. Anderson offers the example of dialing a phone number. As it becomes familiar, we may no longer think of the individual numbers; instead, we think of calling the person and let our fingers take care of the numbers. This proposal is unsatisfactory for understanding skills like typewriting, in which the important thing to explain is the way the motor system manages to execute a sequence of responses.

At present, we know for sure that hierarchical skills can be acquired with practice, we are reasonably sure that extensive practice is necessary to develop them, and we have several theories of skill acquisition that may be helpful in understanding how hierarchical control develops. However, we do not yet have a satisfactory explanation. The acquisition of typing skill may be harder to understand nowadays because people begin to acquire typing skill when they are children, so changes in skill are confounded with large changes in development. One strategy is to degrade skill by degrading the input (having skilled typists type nonwords; Crump & Logan, 2010b) or by degrading the output (having skilled typists type on flat keyboards; Crump & Logan, 2010a, 2010c). These degradations may push skilled typists back toward the beginning of the learning curve to a point at which their typing is no longer controlled hierarchically. If so, then training with degraded inputs and outputs may allow us to observe the development of hierarchical control in adult subjects (for a similar strategy for studying the acquisition of skill at arithmetic, see Zbrodoff, 1995).

8.4. Nested Control Loops in Everyday Cognition

Nested control loops may be at work in many kinds of cognition. The inner loops take care of immediate goals, while the outer loops ensure that broader goals are satisfied. Our analysis has focused on skills like typewriting that require years of practice to attain proficiency, but nested control loops may be created ad hoc when they are required. A central tenet in theories of cognitive control is that executive processes are flexible, allowing us to address new situations coherently and solve novel problems expeditiously (Logan, 1985; Miller & Cohen, 2001; Monsell, 1996; Shiffrin & Schneider, 1977). The ability to create ad hoc control loops fits well with this tenet.

Our use of language provides examples of ad hoc nested control loops. In a conversation, the inner loop may ensure that sentences express the intended meaning, while the outer loop ensures that the point of the conversation gets across. When meeting a new person, the inner loop lets us talk coherently, while the outer loop lets us make a good impression. In negotiation, the inner loop may raise issues and debate points made by the other party, while the outer loop addresses a hidden agenda. The specific goals may be ones we have never addressed before (Chomsky, 1959), yet we accomplish them with some degree of satisfaction.

An interesting possibility is that we create outer loops as we need them. Most of the time, we may be "middle management," carrying out tasks without much regard for higher goals. But when something unusual occurs, like an error, an interruption, or a crisis, we may step back and consider how our activity relates to higher goals (Vallacher & Wegner, 1987). We may not think much about the route home, but an accident on the road ahead may prompt us to create a new plan. We may not think much about survival while writing a chapter, but a fire alarm may revise our priorities. The outer loop may be dormant until we need it, or it may not even exist until we create it on the fly.

The idea of nested control loops may even provide some insight into the quest for the meaning of life. We may feel our lives are meaningful when we are pursuing intermediate goals so actively that we do not have time to consider whether we are fulfilling higher-level goals. We may be happier as middle managers (Frankl, 1959). When we reflect on higher-level goals, we may find nothing of value ("Daddy, I don't want to go to Europe." "Shut up and keep swimming.") or we may invent ad hoc goals to justify our existence. For now, we are happy to work toward the goal of finishing this chapter. What then?!!

ACKNOWLEDGMENTS

This research was supported by grants BCS 0646588 and BCS 0957074 from the National Science Foundation.

REFERENCES

Anderson, J. R. (1982). Acquisition of cognitive skill. *Psychological Review*, *89*, 369–406.
Anderson, J. R. (1987). Skill acquisition: Compilation of weak-method problem solutions. *Psychological Review*, *94*, 192–210.
Arrington, C. M., & Logan, G. D. (2004). The cost of a voluntary task switch. *Psychological Science*, *15*, 610–615.
Badre, D. (2008). Cognitive control, hierarchy, and the rostro-caudal organization of prefrontal cortex. *Trends in Cognitive Sciences*, *12*, 193–200.
Bargh, J. A., & Ferguson, M. J. (2000). Beyond behaviorism: On the automaticity of higher mental processes. *Psychological Bulletin*, *126*, 925–945.
Beilock, S. L., & Carr, T. H. (2001). On the fragility of skilled performance: What governs choking under pressure? *Journal of Experimental Psychology: General*, *130*, 701–725.
Beilock, S. L., Carr, T. H., MacMahon, C., & Starkes, J. L. (2002). When paying attention becomes counterproductive: Impact of divided versus skill-focused attention on novice and experienced performance of sensorimotor skills. *Journal of Experimental Psychology: Applied*, *8*, 6–16.
Botvinick, M., & Plaut, D. C. (2004). Doing without schema hierarchies: A recurrent connectionist approach to normal and routine sequential action. *Psychological Review*, *111*, 395–429.
Botvinick, M., & Plaut, D. C. (2006). Such stuff as habits are made on: A reply to Cooper and Shallice (2006). *Psychological Review*, *113*, 917–927.
Bryan, W. L., & Harter, N. (1899). Studies of the telegraphic language. The acquisition of a hierarchy of habits. *Psychological Review*, *6*, 345–375.
Chomsky, N. (1959). Review of Skinner's *Verbal Behavior*. *Language*, *35*, 26–58.
Cohen, N. J., & Squire, L. R. (1980). Preserved learning and retention of pattern-analyzing skill in amnesia: Dissociation of knowing how and knowing that. *Science*, *210*, 207–210.
Coles, M. G. H. (1989). Modern mind reading: Psychophysiology, physiology, and cognition. *Psychophysiology*, *26*, 251–269.
Cooper, R., & Shallice, T. (2000). Contention scheduling and the control of routine activities. *Cognitive Neuropsychology*, *17*, 297–338.

Cooper, R. P., & Shallice, T. (2006a). Hierarchical schemas and goals in the control of sequential behavior. *Psychological Review*, *113*, 887–916.

Cooper, R. P., & Shallice, T. (2006b). Structured representations in the control of behavior cannot be so easily dismissed: A reply to Botvinick and Plaut (2006). *Psychological Review*, *113*, 929–931.

Crump, M. J. C., & Logan, G. D. (2010a). Episodic contributions to sequential control: Learning from a typist's touch. *Journal of Experimental Psychology: Human Perception and Performance*, *36*, 662–672.

Crump, M. J. C., & Logan, G. D. (2010b). Hierarchical control and skilled typing: Evidence for word level control over the execution of individual keystrokes. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *36*, 1369–1380.

Crump, M. J. C., & Logan, G. D. (2010c). Warning, this keyboard will de-construct: The role of the keyboard in skilled typewriting. *Psychonomic Bulletin & Review*, *17*, 394–399.

Dell, G. S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review*, *93*, 283–321.

Ericsson, K. A., Krampe, R. Th., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, *100*, 363–406.

Fendrick, P. (1937). Hierarchical skills in typewriting. *Journal of Educational Psychology*, *28*, 609–620.

Flanders, M., & Soechting, J. F. (1992). Kinematics of typing: Parallel control of the two hands. *Journal of Neurophysiology*, *67*, 1264–1274.

Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Frankl, V. E. (1959). *Man's search for meaning: An introduction to logotherapy*. Boston, MA: Beacon Press.

Garner, W. R., Hake, H., & Eriksen, C. W. (1956). Operationism and the concept of perception. *Psychological Review*, *63*, 149–159.

Gentner, D. R. (1987). Timing of skilled motor performance: Tests of the proportional duration model. *Psychological Review*, *94*, 255–276.

Gentner, D. R., Larochelle, S., & Grudin, J. (1988). Lexical, sublexical, and peripheral effects in skilled typewriting. *Cognitive Psychology*, *20*, 524–548.

Gordon, A. M., Casabona, A., & Soechting, J. F. (1994). The learning of novel finger movement sequences. *Journal of Neurophysiology*, *72*, 1596–1610.

Gordon, A. M., & Soechting, J. F. (1995). Use of tactile afferent information in sequential finger movements. *Experimental Brain Research*, *107*, 281–292.

Hershman, R. L., & Hillix, W. A. (1965). Data processing in typing. *Human Factors*, *7*, 483–492.

Jasper, H. H. (1958). The ten-twenty electrode system of the International Federation. *Electroencephalography and Clinical Neurophysiology*, *10*, 371–375.

John, B. (1996). TYPIST: A theory of performance in skilled typing. *Human-Computer Interaction*, *11*, 321–355.

Juslin, P. N., Karlsson, J., Lindström, E., Friberg, A., & Schoonderwaldt, E. (2006). Play it again with feeling: Computer feedback in musical communication of emotions. *Journal of Experimental Psychology: Applied*, *12*, 79–95.

Keller, F. S. (1958). The phantom plateau. *Journal of the Experimental Analysis of Behavior*, *1*, 1–13.

LaBerge, D., & Samuels, S. J. (1974). Toward a theory of automatic information processing in reading. *Cognitive Psychology*, *6*, 293–323.

Laming, D. (1968). *Information theory of choice-reaction times*. London: Academic Press.

Lashley, K. S. (1951). The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior* (pp. 112–136). New York: Wiley.

Lessenberry, D. D. (1928). *Analysis of errors*. Syracuse, NY: Smith and Corona Typewriters, School Department.

Levelt, W. J. M. (1989). *Speaking: From intention to articulation*. Cambridge, MA: MIT Press.
Levelt, W. J. M., Schriefers, H., Vorberg, D., Meyer, A. S., Pechmann, T., & Havinga, J. (1991). The time course of lexical access in speech production: A study of picture naming. *Psychological Review*, *98*, 122–142.
Liu, X., Crump, M. J. C., & Logan, G. D. (2010). Do you know where your fingers have been? Explicit knowledge of the spatial layout of the keyboard in skilled typists. *Memory & Cognition*, *38*, 474–484, 138.
Logan, G. D. (1982). On the ability to inhibit complex actions: A stop-signal study of typewriting. *Journal of Experimental Psychology: Human Perception and Performance*, *8*, 778–792.
Logan, G. D. (1985). Executive control of thought and action. *Acta Psychologica*, *60*, 193–210.
Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*, 492–527.
Logan, G. D. (1990). Repetition priming and automaticity: Common underlying mechanisms? *Cognitive Psychology*, *22*, 1–35.
Logan, F. A. (1999). Errors in copy typewriting. *Journal of Experimental Psychology: Human Perception and Performance*, *25*, 1760–1773.
Logan, G. D. (2003). Simon-type effects: Chronometric evidence for keypress schemata in typewriting. *Journal of Experimental Psychology: Human Perception and Performance*, *29*, 741–757.
Logan, G. D., Coles, M. G. H., & Kramer, A. F. (1996). Converging operations in the study of visual selective attention: Introduction. In A. F. Kramer, M. G. H. Coles, & G. D. Logan (Eds.), *Converging operations in the study of visual selective attention* (pp. xv–xxv). Washington, DC: American Psychological Association.
Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, *91*, 295–327.
Logan, G. D., & Crump, M. J. C. (2009). The left hand doesn't know what the right hand is doing: The disruptive effects of attention to the hands in skilled typewriting. *Psychological Science*, *10*, 1296–1300.
Logan, G. D., & Crump, M. J. C. (2010). Cognitive illusions of authorship reveal hierarchical error detection in skilled typists. *Science*, *330*, 683–686.
Logan, G. D., & Gordon, R. D. (2001). Executive control of visual attention in dual-task situations. *Psychological Review*, *108*, 393–434.
Logan, G. D., Miller, A. E., & Strayer, D. L. (2011). Electrophysiological evidence for parallel response selection in skilled typists. *Psychological Science*, *22*, 54–56.
Logan, G. D., & Zbrodoff, N. J. (1998). Stroop type interference: Congruity effects in color naming with typewritten responses. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 978–992.
Long, J. (1976). Visual feedback and skilled keying: Differential effects of masking the printed copy and the keyboard. *Ergonomics*, *19*, 93–110.
MacKay, D. G. (1982). The problem of flexibility, fluency, and speed-accuracy tradeoff in skilled behavior. *Psychological Review*, *89*, 483–506.
Markman, A. B. (1999). *Knowledge representation*. Mahwah, NJ: Lawrence Erlbaum Associates.
McNamara, T. P. (1986). Mental representations of spatial relations. *Cognitive Psychology*, *18*, 87–121.
Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167–202.
Miller, G. A., Galanter, E., & Pribram, K. (1960). *Plans and the structure of behavior*. New York: Holt.

Monsell, S. (1996). Control of mental processes. In V. Bruce (Ed.), *Unsolved mysteries of the mind* (pp. 93–148). Hove, England: Erlbaum.

Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), *Visual information processing* (pp. 283–308). New York: Academic 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* (pp. 1–55). Hillsdale, NJ: Erlbaum.

Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation: Advances in research and theory* (pp. 1–18). New York: Plenum.

Novick, L. R., & Hurley, S. M. (2001). To matrix, network, or hierarchy: That is the question. *Cognitive Psychology*, *42*, 158–216.

Palmer, C. (1997). Music performance. *Annual Review of Psychology*, *48*, 115–138.

Platt, J. R. (1964). Strong inference. *Science*, *146*, 347–353.

Rabbitt, P. M. A. (1966). Errors and error correction in choice-response tasks. *Journal of Experimental Psychology*, *71*, 264–272.

Rabbitt, P. M. A. (1978). Detection of errors by skilled typists. *Ergonomics*, *21*, 945–958.

Rayner, K., & Clifton, C. (2009). Language processing in reading and speech perception is fast and incremental: Implications for event-related potential research. *Biological Psychology*, *80*, 4–9.

Repp, B. H., & Knoblich, G. (2004). Perceiving action identity: How pianists recognize their own performances. *Psychological Science*, *15*, 604–609.

Roediger, H. L. (1990). Implicit memory: Retention without remembering. *American Psychologist*, *45*, 1043–1056.

Rumelhart, D. E., & Norman, D. A. (1982). Simulating a skilled typist: A study of skilled cognitive-motor performance. *Cognitive Science*, *6*, 1–36.

Salthouse, T. A. (1986). Perceptual, cognitive, and motoric aspects of transcription typing. *Psychological Bulletin*, *99*, 303–319.

Salthouse, T. A., & Saults, J. S. (1987). Multiple spans in transcription typing. *Journal of Applied Psychology*, *72*, 187–196.

Schneider, D. W., & Logan, G. D. (2006). Hierarchical control of cognitive processes: Switching tasks in sequences. *Journal of Experimental Psychology: General*, *135*, 623–640.

Schneider, D. W., & Logan, G. D. (2007). Retrieving information from a hierarchical plan. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *33*, 1076–1091.

Shaffer, L. H. (1973). Latency mechanisms in transcription. In S. Kornblum (Ed.), *Attention and performance IV* (pp. 435–448). New York: Academic Press.

Shaffer, L. H. (1976). Intention and performance. *Psychological Review*, *83*, 375–393.

Shaffer, L. H. (1982). Rhythm and timing in skill. *Psychological Review*, *89*, 109–122.

Shaffer, L. H., Clarke, E. F., & Todd, N. P. (1985). Metre and rhythm in piano playing. *Cognition*, *20*, 61–77.

Shaffer, L. H., & Hardwick, J. (1968). Typing performance as a function of text. *Quarterly Journal of Experimental Psychology*, *20*, 360–369.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, *84*, 127–190.

Simon, J. R., & Small, A. M., Jr. (1969). Processing auditory information: Interference from an irrelevant cue. *Journal of Applied Psychology*, *53*, 433–435.

Soechting, J. F., & Flanders, M. (1992). Organization of sequential typing movements. *Journal of Neurophysiology*, *67*, 1275–1290.

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' methods. In W. G. Koster (Ed.), *Attention and performance II* (pp. 276–315). Amsterdam: North Holland.

Sternberg, S., Monsell, S., Knoll, R. L., & Wright, C. E. (1978). The latency and duration of speech and typewriting. In G. E. Stelmach (Ed.), *Information processing in motor control and learning*. New York: Academic Press.

Stevens, A., & Coupe, P. (1978). Distortions in judged spatial relations. *Cognitive Psychology*, *10*, 422–437.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, *18*, 643–662.

Terzuolo, C. A., & Viviani, P. (1980). Determinants and characteristics of motor patterns used for typing. *Neuroscience*, *5*, 1085–1103.

Vallacher, R. R., & Wegner, D. M. (1987). What do people think they're doing? Action identification and human behavior. *Psychological Review*, *94*, 3–15.

Viviani, P., & Laissard, G. (1996). Motor templates in typing. *Journal of Experimental Psychology: Human Perception and Performance*, *22*, 417–445.

Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.

Wu, C., & Liu, Y. (2008). Cuing network modeling of transcription typing. *ACM Transactions on Computer-Human Interaction*, *15*, 1–45, (Article 6).

Zbrodoff, N. J. (1995). Why is 9 + 7 harder than 2 + 3? Strength and interference as explanations of the problem-size effect. *Memory & Cognition*, *23*, 689–700.