an Article from **SCIENTIFIC AMERICAN**

The Skill of Typing

How can skilled typists type as fast as they do? A century of study has not produced a definitive answer. A leading hypothesis is that a typist learning to type well learns to make mental processes overlap
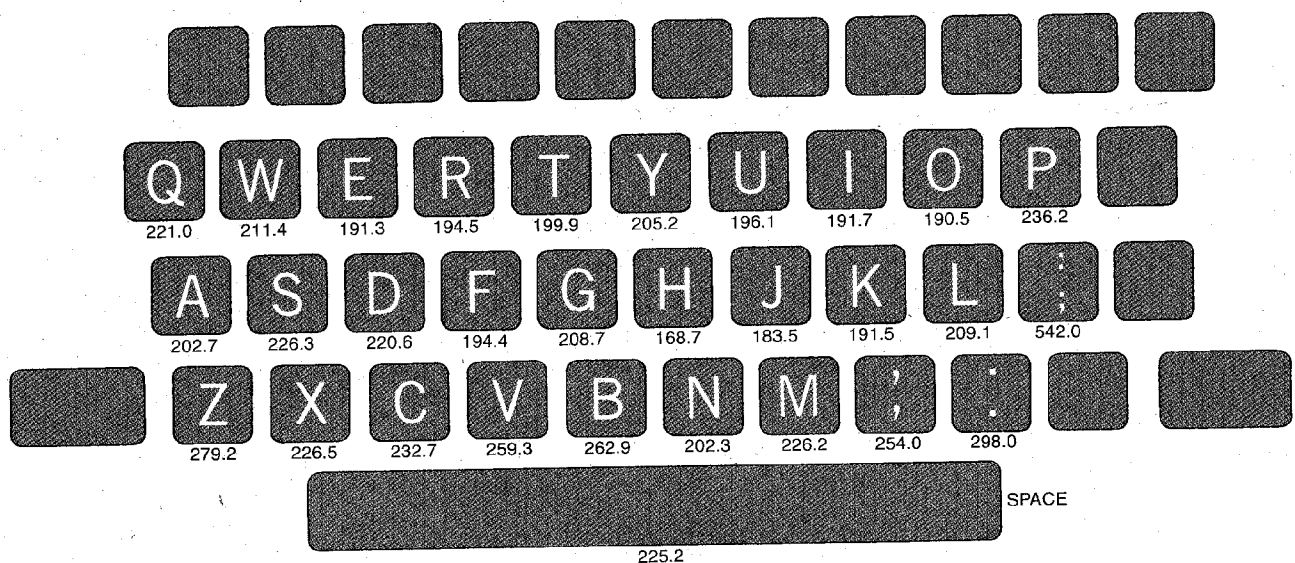
by Timothy A. Salthouse

From almost the time the first typewriter came into use, the skill of typing has intrigued experimental psychologists, and it intrigues them even today, because the rate at which typists (even average ones) perform exceeds by far the rate that laboratory tests quite common in psychology would lead a psychologist to predict. Consider a typical secretarial task, the retyping of a document. In essence it is the transcription of a sequence of symbols (letters, numbers and so on) that come under the typist's gaze into a sequence of motor acts: the pressing of typewriter keys. Thus it closely· resembles a series of choice-reaction-time tasks, in which a subject is presented with a single visual stimulus from a set of two or more stimuli after being instructed to rapidly press a particular button for each of the possible stimuli. Under optimal conditions

(highly practiced people and a minimal number of stimulus-response alternatives) the average latency, or delay between the presentation of the stimulus and the pressing of a button, is approximately 250 milliseconds. The paradox of typing is that a latency of 250 milliseconds yields a typing rate of 48 words per minute (assuming five keystrokes per word). Yet speeds of twice that rate are fairly common.

How is it possible? What have skilled typists learned that enables them to overcome what appear to be fundamental limitations? Everyone has a minimum reaction time, which should set limits on the maximum typing rate, and yet skilled typists have developed a means of overcoming their perceptual and motor restrictions. Their achievement may have significance far beyond the skill of typing. After all, the detailed

understanding of any skill, in the form of knowing how skilled people differ from less skilled people, should have implications for the selection of students and for their training in skills quite generally. Understanding the skill of typing may even have significance for rehabilitation therapy.

Typewriters came into use in the U.S. in the latter part of the 19th century, a period when experimental psychology was also becoming established. Perhaps that is why a number of early psychological studies were done on various aspects of typing. The research uncovered most of the major phenomena of typing: that it goes faster when keystrokes alternate between the hands, when successive keystrokes represent letter pairs with a high frequency of occurrence in the language, when the material being

(keyboard diagram with timing values)

Key	Time
Q	221.0
W	211.4
E	191.3
R	194.5
T	199.9
Y	205.2
U	196.1
I	191.7
O	190.5
P	236.2
A	202.7
S	226.3
D	220.6
F	194.4
G	208.7
H	168.7
J	183.5
K	191.5
L	209.1
;	542.0
Z	279.2
X	226.5
C	232.7
V	259.3
B	262.9
N	202.3
M	226.2
,	254.0
.	298.0
SPACE	225.2

TYPING SPEED exceeds the speed that psychological tests would lead one to predict. The number below each key on this typewriter keyboard is the average time required for that keystroke by 10 skilled typists the author studied. (Times are in milliseconds.) Most of the times are notably fast: less than 250 milliseconds (*colored numbers*). The same typists were asked to rest each index finger on a certain key and press it whenever the letter *l* or *r* (for "left" or "right") appeared on a screen. Response times were then 500 to 600 milliseconds.

typed is meaningful to the typist and so on. Thus the research provided some of the first scientific evidence pertaining to man-machine interaction. Indeed, the research probed what is sometimes considered one of the most maladaptive man-machine interfaces ever invented: the standard, or QWERTY, typewriter keyboard.

The QWERTY keyboard (named, of course, for the sequence of letters across its top row) had been devised in the 1870's without empirical research. It was a response to a defect of the earlier machines, namely that moderate speeds of typing tended to make keys jam together. QWERTY has been criticized for requiring disproportionate effort by the weakest fingers of each hand, and several alternatives have been proposed. Unfortunately QWERTY is now so thoroughly established that the benefits of a better arrangement are unlikely to offset the cost of retraining present-day typists and scrapping present-day typewriters and other keyboard machines. Clearly specialists in man-machine interaction should participate in the design of virtually all the equipment people use.

One of the earliest studies of typing was reported by William F. Book in 1908. Book had been much influenced by a classic study of telegraphers published 11 years earlier by William Bryan, an experimental psychologist, and Noble Harter, a former telegrapher who was doing graduate work in psychology. These earlier investigators had concluded from interviews, systematic observation and records of performance of telegraphers that a telegrapher in the process of becoming expert focuses attention on successively larger units, beginning with dots and dashes, the fundamental elements of telegraphy, and progressing to characters, syllables, words and even phrases. There is now considerable evidence that many skills are organized in this hierarchical fashion. Moreover, it is clear that skills such as reading would be almost impossibly difficult if the reader always analyzed written language in units of isolated letters. Book was therefore on firm theoretical ground in suggesting that skilled typing was achieved by expanding the size of verbal units from single characters to words and perhaps even phrases.

Fifteen years after Book's proposal the importance of large units of analysis to the skill of typing was empirically demonstrated by John E. Coover. He phrased his finding succinctly: "If copy is presented one letter at a time, so that as soon as the letter is typed another automatically appears, the expert's performance is reduced to a series of reaction times to the letters, and his rate is greatly reduced." This observation is readily verified with the aid of a computer. In the procedure I have implemented on a computer in my laboratory at the University of Missouri at Columbia typists see written material in a "preview display" that can be varied from one character to 39 characters. That is, the material shown to them on a video terminal can range from zero to 38 characters ahead of the character they are typing. They type at a standard keyboard, and the time (in milliseconds) between their successive keystrokes is automatically recorded and stored, along with the character typed.

Many typists have now been tested with this procedure. The results are uniform in that every subject tested typed very slowly when the preview was most limited and typed faster when the preview size increased. The preview size at which the typing rate first equaled its normal rate (its rate with unlimited preview and with a record of the typing that was visible to the typist) ranged from three to seven characters. (The 10 typists in the study varied in age from 20 to 40 and had gross typing rates, that is, rates uncorrected for typing errors, that varied from 50 to 75 words per minute.) If the typing rate is affected by the visual availability of characters up to seven spaces in advance of the character being typed, one must conclude that the typist attends to these characters. Thus there must be a gap or span between the characters receiving the attention of the eyes and the character whose key is being pressed. Similar estimates of such a gap have been reported in analyses of eye movements recorded in typists.

How does preview confer an advantage, so that the limitations of reaction time are circumvented by skilled typists? The reason is still in dispute. The possibility Book would favor is that typists developing their skill move from a mode in which they analyze character by character to one involving larger units such as words and phrases. Thus
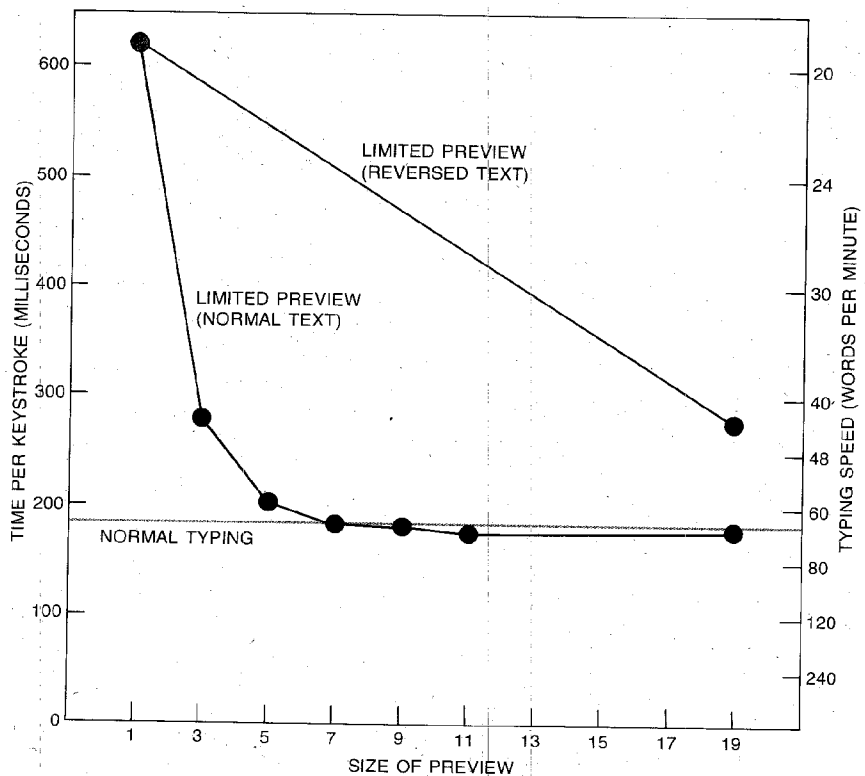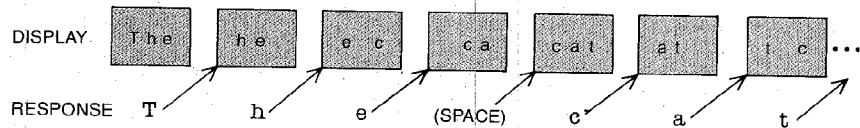

LIMITED PREVIEW diminished the typists' speed. Each typist faced a screen on which a text always remained a certain number of characters ahead of the character the typist was typing. (In the example at the top the text is "The cat climbed . . ." and the preview size is three.) The typists' speed approached normal speed only when the preview size approached seven. A similar effect was seen even when the text was randomized by reversing each word letter by letter.

the skilled typist is assumed to impose groupings on both the input (the text to be typed) and the output (the sequence of keystrokes). The input groupings amount to perceptual "chunks." Probably they consist of syllables or common words, evidently having three to seven characters. The output groupings are similar to what contemporary theorists call motor programs and consist of integral sequences of essentially automatic keystrokes. In effect the typing of a given output grouping is taken to be ballistic: it requires only a start signal, after which nothing affects the execution of the movement. According to this hypothesis, maximum typing speed requires a preview of three to seven characters because a smaller preview cripples the ability to group.

Shortly after Book made his suggestion a number of people proposed the competing idea that typing speeds circumventing the choice reaction time are accomplished by an overlapping of "processing operations." August Dvorak, a strong advocate of this position, was also a critic of QWERTY and the inventor of an alternative keyboard. Dvorak, like Book, asked why typing strokes are faster than isolated choice-reaction-time responses. "The answer," he wrote in 1936, "is that even with a two-letter sequence, while the first finger is stroking, the second finger is starting its play for position and overlapping its stroke with the first."

The processing operations can be given modern labels remarkably similar to the ones incorporated in recent "stage theories" of choice-reaction-time tasks. They might begin with encoding: the perception of each character. Then would come the categorization of the stimulus (say into the class of the left- or right-handed keystrokes), the fixed decision as to which finger must type and finally the preparation and execution of the actual typing movements. On this hypothesis maximum typing speed requires a certain preview because the visual availability of text characters allows the sequence of processing operations to be executed in parallel for a number of characters, the late stages of one sequence accompanying the earlier stages of others instead of the strict succession dictated by the availability of only one character at a time.

Some support for the "chunking hypothesis" (the one that Book would favor) is available from an analysis of the time intervals between the keystrokes made by a typist who repeatedly types the same text with unlimited preview. Under these circumstances a typist is generally quite consistent in the temporal pattern of keystrokes composing a given word, as though the individual keystrokes were under the control of a single motor program. The consistency cannot simply be attributed to the differ-

INPUT: The cat climbed . . .

OUTPUT   T   h   e   c   a   t   c   l   i   m   b   e   d

TIME——>
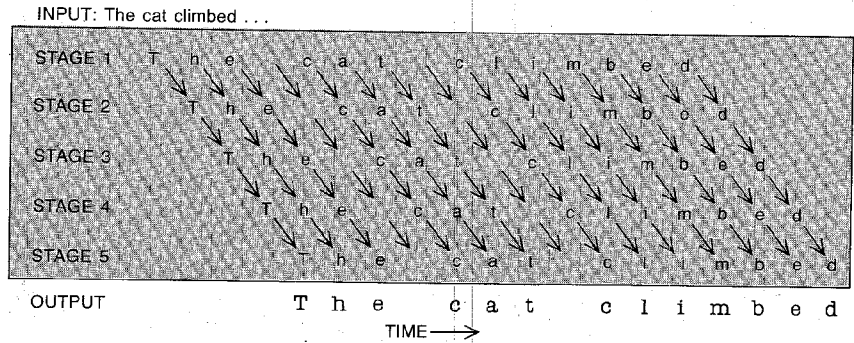
CHUNKING HYPOTHESIS is an attempt to account for both the preview effect and the remarkable speed of skilled typists. It proposes that the material to be typed gets mentally "chunked" into multicharacter units (here words) and that the typing of each unit is governed by an autonomous "motor program." If the preview is too limited, chunking is counteracted.

INPUT: The cat climbed . . .

OUTPUT   T   h   e   c   a   t   c   l   i   m   b   e   d

TIME——>
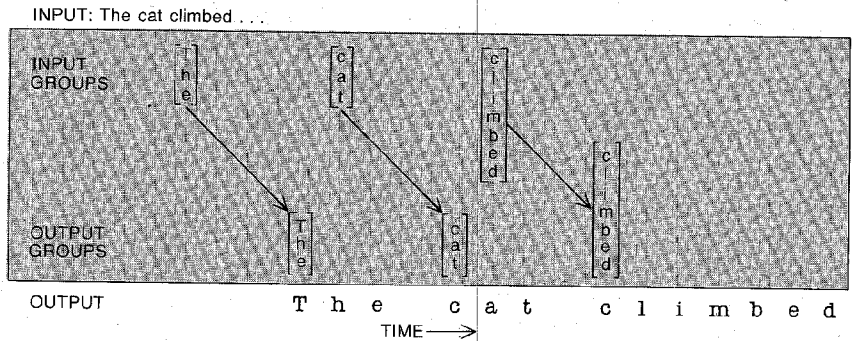
OVERLAPPING OPERATIONS is a competing hypothesis. It proposes that typing of each character proceeds in steps (here five) such as perception of a character, its assignment to a hand and so on, and that the steps overlap. If preview is too limited, overlapping is counteracted.

The quick brown fox jumps over the lazy dog
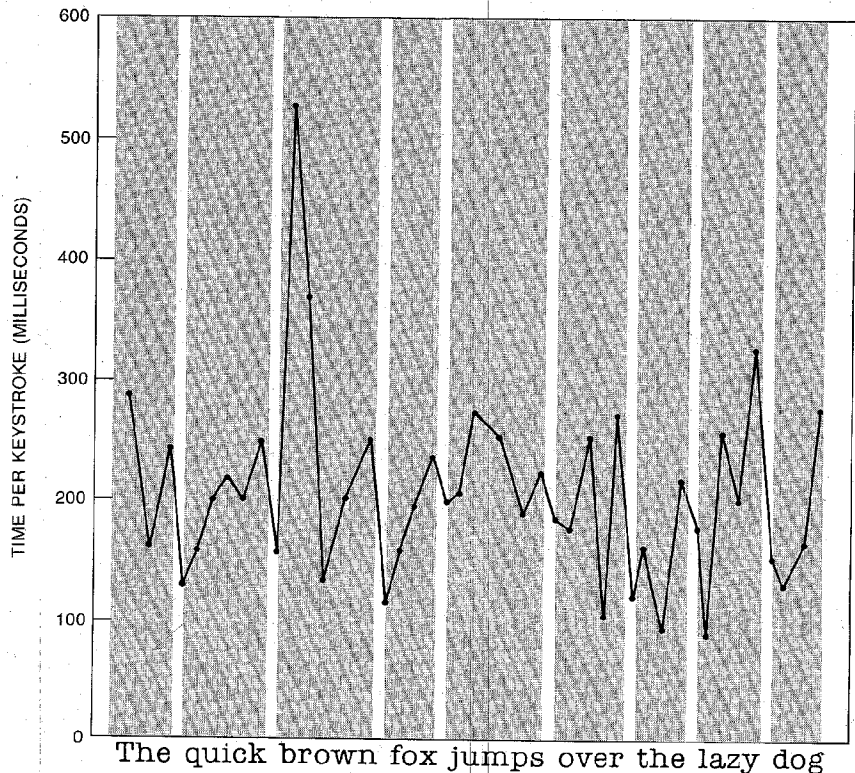
TEMPORAL PATTERN OF KEYSTROKES over 10 typings of "The quick brown fox..." by one typist seems at first to favor the chunking hypothesis. A given letter (such as o) is typed at different speeds in different words, as if it were under the control of different motor programs.

131

ing accessibility of various keys or to variations in the strength or dexterity of individual fingers, because the same character turns out to be typed with different latencies in different contexts. An example is the letter *o* in the familiar sentence "The quick brown fox jumps over the lazy dog." The typists I studied typed it with an average latency of 370 milliseconds in the word *brown,* 160 milliseconds in *fox,* 185 milliseconds in *over* and 130 milliseconds in *dog.*

Further results, however, suggest important qualifications of the chunking hypothesis. If the hypothesis is interpreted literally and the characters of a word are organized into a strictly ballistic motor program, one might predict that errors made at any point in its exe-

cution would go undetected until at least the end of the program. The prediction can be examined by an analysis of the errors committed in normal typing. It has been reported (in fact, the first reports came soon after Book's pioneering work) that the keystrokes accompanying the first intimation on the part of the typists that an error has been made are partially inhibited, in that they are weaker or delayed with respect to normal typing. These phenomena might serve, then, as a marker of the instants when errors are detected. The markers do have a drawback: it is difficult to predict a typist's groupings. A preferable strategy is to assume that groupings do not exist. Errors could then be detected immediately. A failure to find evidence of immediate error detection

would therefore tend to support the notion that typing is controlled in multicharacter groups.

I determined the median latencies for the keystrokes before and after each of four types of errors: substitution errors (typing *caet* instead of *cart,* for example), intrusion errors (*carrt* instead of *cart*), omission errors (*cat* instead of *cart*) and transposition errors (*crat* instead of *cart*). The interpretation of the omission results is difficult: the latency of the omitted keystroke might be incorporated into that of the next keystroke. In addition some errors are not detected unless the typist proofreads the typed copy. No error-detection phenomena could reasonably be expected from undetected errors. Nevertheless, my results suggest that a large proportion of errors are detected immediately, and not at the end of multicharacter groups. This is not to say output groupings cannot exist. It does, however, indicate that the control of typing is not completely relinquished to multicharacter motor programs.

Additional evidence for the conclusion that keystrokes are controlled in units smaller than words or syllables has recently been provided by Gordon D. Logan of the University of Toronto. Logan employed a technique in which typists were instructed to stop typing whenever they heard a tone. He reasoned that if word-length motor programs controlled typists' keystrokes, the typists would type to the end of a word regardless of both the "position" of the stop signal and the number of letters remaining to be typed. Logan found that regardless of word length the typists hearing the tone actually stopped after an average of only one or two letters.

A third piece of evidence that weakens the chunking hypothesis emerges from a comparison of the typing of normal and random texts. The chunking hypothesis leads to two predictions about the typing of random texts. First, the process of chunking calls for the matching of grouped input material to preestablished motor programs, and so the process should be disrupted by material that does not have the familiarity of normal English. Hence the rate of typing should be greatly reduced. Second, a random text is not amenable to chunking even with unlimited preview. Hence the rate of typing should not be affected by reducing the size of the preview.

These predictions were examined by having typists type sentences where each word had been reversed letter by letter. The resulting text is not genuinely random, but it does have the advantage of disrupting the familiar letter sequences and word organizations of English while preserving the letter frequencies and the pattern of interword spacing. Preview sizes of 19 characters and one character
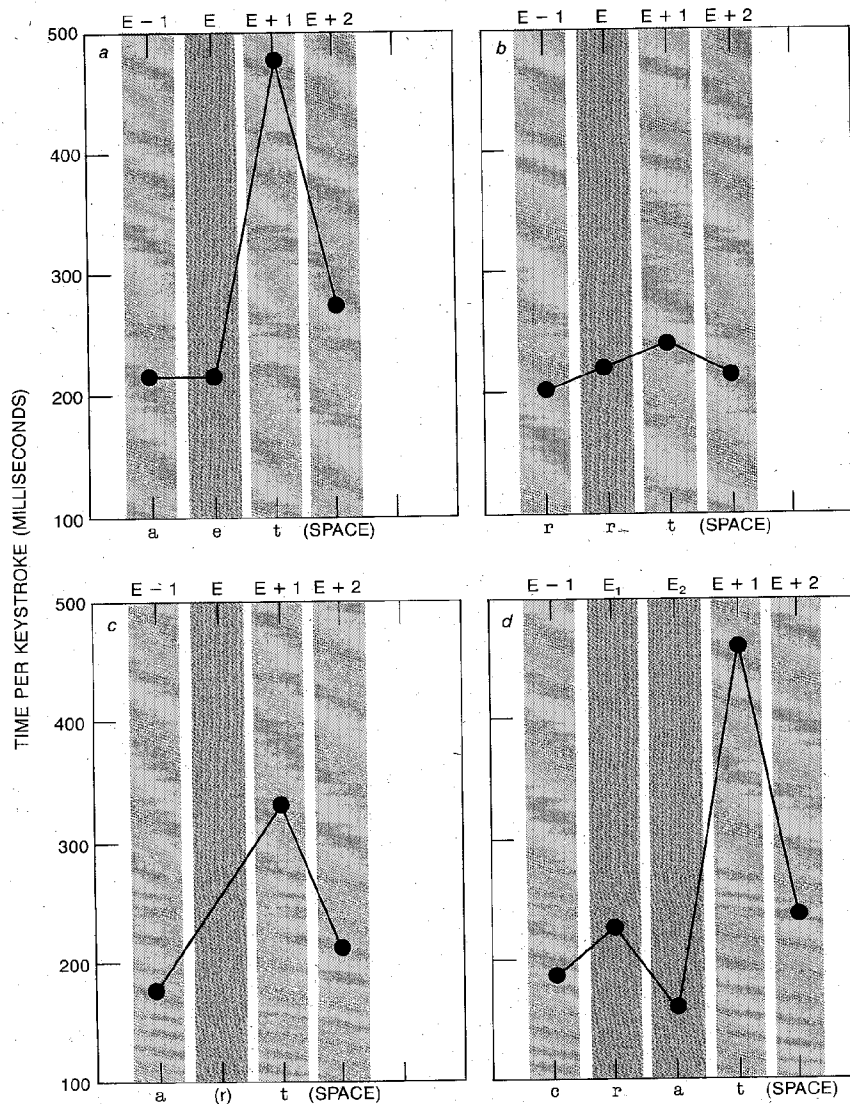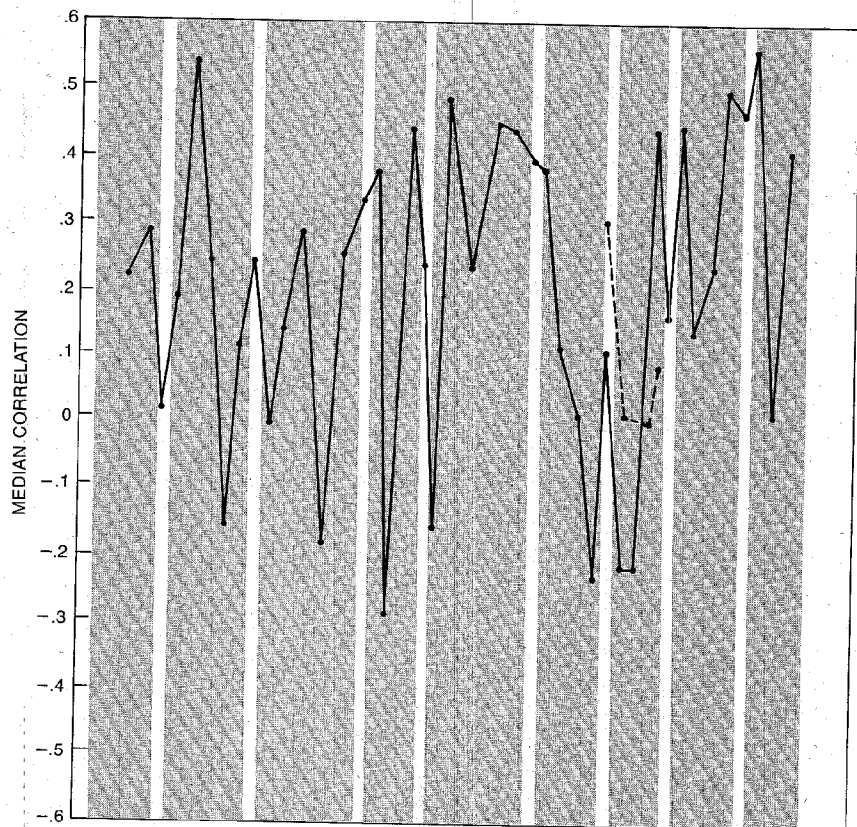

TYPING ERRORS weaken the chunking hypothesis. Four errors are shown: substitution (*a*), such as *caet* instead of *cart;* intrusion (*b*), such as *carrt* instead of *cart;* omission (*c*), such as *cat* instead of *cart,* and transposition (*d*), such as *crat* instead of *cart.* In each case typists slowed one keystroke after the error (*dark color*) was made. The slowing is taken to signal awareness that something is wrong. The chunking hypothesis implies errors cannot be detected that soon.

were presented, and typists performed under each condition twice. The first prediction proved to be true: typing speed was impaired. The second prediction proved to be mistaken: the viewing of 19 characters resulted in markedly faster typing than the viewing of one character at a time. Thus the preview effect is pronounced. Indeed, increasing the number of visible characters provided nearly as much relative advantage with the meaningless material as it would with normal text. It seems to follow that the benefits of presenting more than one character at a time are not simply attributable to the input chunking and output chunking of meaningful patterns.

The chunking hypothesis is weakened still further by the finding that the latencies of successive keystrokes have a relatively low correlation. That is, the time it takes a typist to make a given keystroke proves to have little relation to the time it takes the typist to make the next keystroke. If the successive keystrokes were within an output grouping, one might expect to find such a relation. If the first keystroke takes longer, the second one should too. After all, they would both be under the control of a single higher-order process.

Finally, the chunking hypothesis is weakened by the finding that there is no relation between the length of a word and either the latency of the first keystroke in the word or the median interval between successive keystrokes in the word. The reason a word-length effect might have been expected is that the time needed for a motor program to be "retrieved" and "prepared for execution" (to borrow the terminology of computer science) should be related to the number of units (separate keystrokes) governed by the program.

No single item among these lines of evidence is conclusive in itself. Taken together, however, the findings strongly imply that typing speeds circumventing choice-reaction-time rates are unlikely to be accountable solely to a switch from single-character to syllable- or word-length units of analysis. Why, then, does the chunking hypothesis seem so compelling? One explanation might be the strong subjective impression among most skilled typists that they need only think of a word and it seems to be typed automatically. This impression may be due more to the typist's input than to any output processes. In particular, the lifelong habits of reading probably include the perception of written language in multicharacter input units. Thus a manuscript in the process of being typed may get perceptually chunked by syllable or word. After that the preparation and execution of the discrete motor acts that constitute typing

The quick brown fox jumps over the lazy dog

LACK OF SYSTEMATIC CORRELATION between the speeds of keystrokes made in succession by people typing "The quick brown fox..." also weakens the chunking hypothesis. Each dot shows the correlation between the time required for a given keystroke and the time required for the preceding keystroke. The scale ranges from 1 through 0 to −1. A correlation of 1 means that an increase in one of the variables is always accompanied by a proportionate increase in the other. A correlation of −1 means that an increase in one variable is always accompanied by a proportionate decrease in the other. A correlation of 0 means that the variables have no relation. The broad scatter of points suggests that the characters in each word cannot be part of a higher-order chunk. Broken line in the chart shows correlations among keystrokes obtained from the typing of *the* in nonrepetitious material. Again no relation emerged.
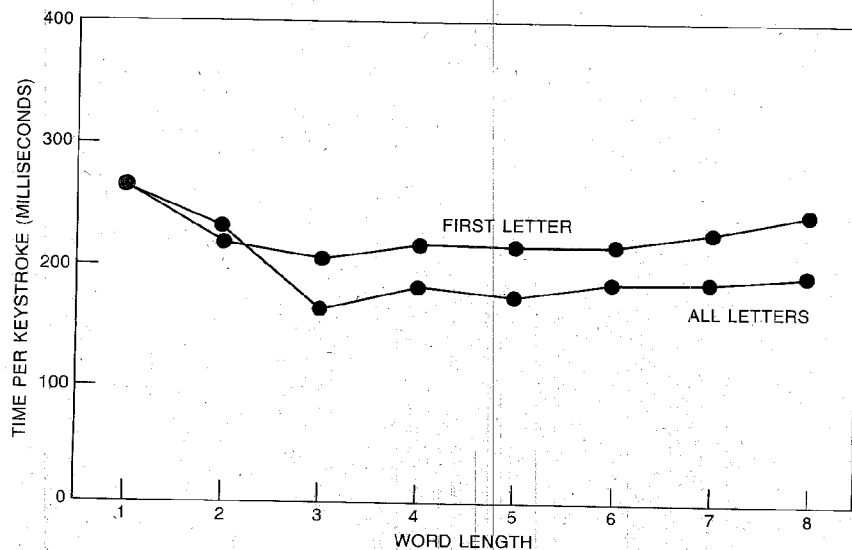

LACK OF WORD-LENGTH EFFECT weakens the chunking hypothesis still further. The hypothesis predicts that motor programs for long words should take longer to "retrieve" and "execute." Yet the median time per keystroke needed to type a word (*black curve*) and the time needed to type the first letter in it (*colored curve*) are independent of the length of the word.

133

may proceed in ways utterly inaccessible to conscious awareness.

How can the evidence that seems to favor the chunking hypothesis (for example the consistency in the temporal patterns of keystrokes making up the words in "The quick brown fox...") be explained if not by a motor program? Specifically, how can the evidence be explained by the competing hypothesis that skilled typists make their processing operations overlap? One possibility is that the time required for a keystroke depends on factors such as the position of the preceding keys and subsequent keys and the direction in which specific fingers must move. The importance of these factors is documented by the finding that the time needed for a keystroke is less when the preceding keystroke was made by a finger of the opposite hand. My own studies of normal typing show that opposite-hand sequences of two keystrokes are executed an average of 40 milliseconds faster than same-hand sequences. Similar effects have been reported by earlier workers investigating sequences involving the same finger or different fingers (the latter are faster) and the same row or different rows of the keyboard. (The results are a bit complicated; for example, the index finger moves inward on a row faster than it moves to a different row.)

One might also expect that the relative frequency with which particular sequences of letters have been typed in the past is important, so that commonly typed sequences are executed faster than less-practiced ones. This possibility can be examined by calculating the correlation between two variables: on the one hand the frequency of occurrence of various two-letter sequences in English and on the other the latency between the two keystrokes in each sequence as the sequences arise in the course of normal typing. The range of occurrences is great: from zero percent for sequences such as *qz* to 3.68 percent for the sequence *th*. Therefore I worked with the logarithms of the frequencies of occurrence. The correlation turned out to be low: only some 4 percent of the variability in the latency between keystrokes could be attributed to the variability in the logarithm of the frequency of occurrence of two-letter combinations. Nevertheless, the average latency decreased by about 52 milliseconds for each log-unit increase in frequency. Thus a frequency effect emerges. I have not done a similar analysis for longer sequences, but it is reasonable to guess that three-letter combinations might also show a frequency effect. It has been estimated that only 104 three-letter sequences (out of 26³, or 17,576) account for nearly half of the occurrences of three-letter sequences in English.

So far I have emphasized the similar-
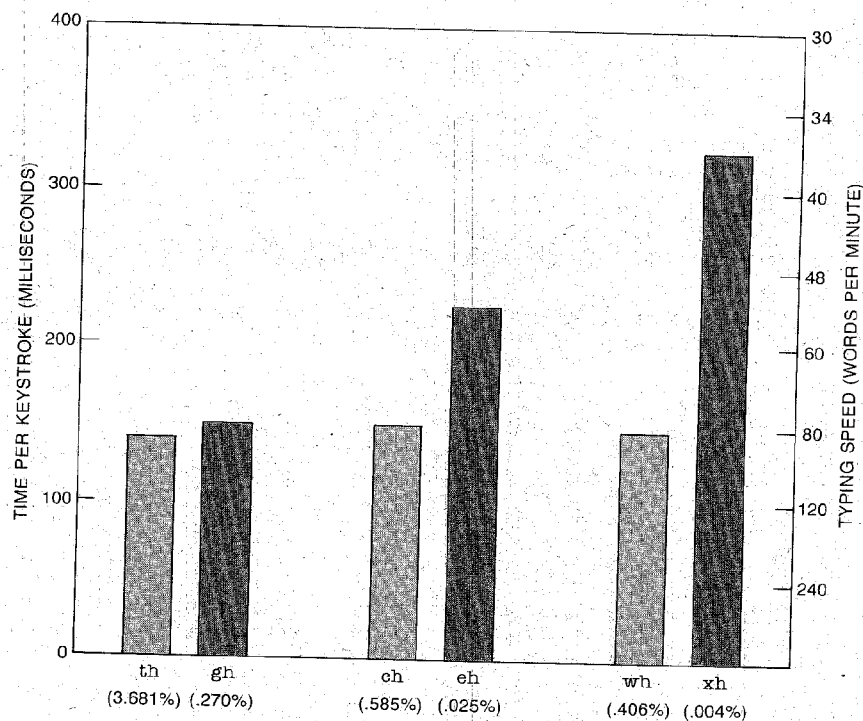
TYPING OF TWO-LETTER COMBINATIONS suggests one reason typing speeds for a given letter tend to vary from word to word. In each case the letter *h*, typed by the right index finger, was preceded by a letter typed with the opposite hand. In *th* and *gh* (*left*) the *t* and the *g* are typed by the left index finger; *th* is commoner in English. (Frequencies of occurrence are in parentheses.) In *ch* and *eh* (*middle*) the *c* and the *e* are typed by the left middle finger; *ch* is the commoner. In *wh* and *xh* (*right*) *w* and *x* are typed by the left ring finger; *wh* is the commoner. The time required to type *h* is greater for the *h* in each of the less common combinations.

ities among typists. I should also mention the differences. One of these is the extent to which typists comprehend what they type. In the course of my studies some typists reported they had virtually no idea of the content of what they typed. Others said they could comprehend the material at least as well as if they had merely read it. Their subjective reports are supported by their scores on posttyping comprehension tests, which ranged from 12.5 percent correct responses to 75 percent.

Another way in which typists differed was in the frequency with which they made errors of various types. The least frequent type of error for all typists was transposition, but the most frequent type was either intrusion, omission or substitution. The duration of the keystrokes associated with errors also varied. Some typists made extremely fast keystrokes by pressing two fingers almost simultaneously. The first keystroke would come with a normal latency, the second would come perhaps 10 milliseconds later. The error (often an intrusion error) might be either one of the strokes.

All things considered, it appears that the limits of reaction time are circumvented primarily because skilled typists have learned to overlap the perform-

ance of many of the operations involved in making keystrokes. This overlapping is impossible when the material to be typed is displayed only a few characters at a time. What does this conclusion imply for the understanding of skill? First, the results from the study of typing are consistent with results from the study of other perceptual-motor activities in suggesting that a cognitive component is present in nearly all skilled behavior. Attention to characters ahead of the one being typed may be due to the development of multicharacter units of analysis and execution or it may be due to the overlapping of processing operations. Either way, the mechanisms are highly rational adaptations to the limits of processing character by character.

Second, the processes mediating between the input and the output in typing are exquisitely coordinated: the skilled typist does not often hesitate. The intensive practice necessary to achieve high levels of skill appears, then, to result in the elimination of unnecessary operations, in the ability to execute more than one operation at a time and perhaps in a reduction in the attention demanded of the typist by certain operations. These characteristics may be common to a wide variety of activities and so may serve as some of the goals toward which any kind of training should be oriented.