


AD-A100 359

CALIFORNIA UNIV SAN DIEGO LA JOLLA CENTER FOR HUMAN --ETC F/6 5/10
SIMULATING A SKILLED TYPIST: A STUDY OF SKILLED COGNITIVE-MOTOR--ETC(U)
MAY 81 D E HUMELHART, D A NORMAN NU0014-79-C-0323
CHIP-102 NL

UNCLASSIFIED

1 of 1
4L
300MB



END
DATE
SERIAL
7 8!
DTIC

LEVEL II



SIMULATING A SKILLED TYPIST:


A STUDY OF SKILLED COGNITIVE-MOTOR PERFORMANCE

AD A100359

David E. Rumelhart and Donald A. Norman

DTIC
ELECTIC

JUN 10 1981



CENTER FOR HUMAN INFORMATION PROCESSING

Research support was provided by the Personnel and Training Research Programs, Psychological Sciences Division, Office of Naval Research, under Contract Number N00014-79-C-0323, Contract Authority Identification Number, NR 157-437. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Office of Naval Research. Approved for public release; distribution unlimited. Reproduction in whole or in part is permitted for any purpose of the United States Government.

UNIVERSITY OF CALIFORNIA, SAN DIEGO
LA JOLLA, CALIFORNIA 92093

CHIP 102
MAY 1981

WING FILE COPY

CHIP Report 102

May 1981

Simulating a Skilled Typist:
A Study of Skilled Cognitive-Motor Performance

David E. Rumelhart

Donald A. Norman

Department of Psychology

and

Program in Cognitive Science

University of California, San Diego

Copyright © 1981 David E. Rumelhart and Donald A. Norman

Approved for public release; distribution unlimited.

Program in Cognitive Science, C-015
Center for Human Information Processing
University of California, San Diego
La Jolla, California 92093

Accession For	
FILE	X
REF	
USE	
BY	
DATE	
TIME	
INITIALS	
A	

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

DD FORM 1 JAN 73 1473

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Abstract

We review the major phenomena of skilled typing and propose a model for the control of the hands and fingers during typing. The model is based upon an Activation-Trigger-Schema system in which a hierarchical structure of schemata directs the selection of the letters to be typed and, then, controls the hand and finger movements by a cooperative, relaxation algorithm. The interactions of the patterns of activation and inhibition among the schemata determine the temporal ordering for launching the keystrokes. To account for the phenomena of doubling errors the model has only "type" schemata -- no "token" schemata -- with only a weak binding between the special schema that signals a doubling and its argument. The model exists as a working computer simulation and produces an output display of the hands and fingers moving over the keyboard. It reproduces some of the major phenomena of typing, including the interkeypress latency times, the pattern of transposition errors found in skilled typists, and doubling errors. Although the model is clearly inadequate or wrong in some of its features and assumptions, it serves as a useful first approximation for the understanding of skilled typing.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

Simulating a Skilled Typist:
A Study of Skilled Cognitive-Motor Performance

David E. Rumelhart
Donald A. Norman
University of California, San Diego

The surprising thing about human control of the limbs is the contrast between the ease with which it is done and the complexity of the system. Piano playing rates often reach 30 notes per second for sustained periods. Typing champions reach close to 200 words per minute (17 letters per sec). The hand is an intricate mechanical system, with complex control problems. Just how the fingers should be moved to reach some target depends upon the positions of the hand, the lower and upper arms, the shoulder, and the body position and angle. The computation of the proper movement is difficult, for the number of degrees of freedom is large and multiple solutions are possible (see Saltzman, 1979). The ten fingers each have three joints, each bone controlled in its up and down motion (extension and flexion) by two tendons, each tendon controlled by a muscle located in the lower hand. Additional muscles within the hand control the side-to-side motion of the fingers (abduction, adduction, flexion, and rotation) and the configuration of the hand. In total, there are about 50 separate movements of the parts, not counting the relevant movements of wrist, lower arm, upper arm, and body.

Typing is not a trivial task. A surprisingly large percentage of the population cannot type. Among those who are expert, considerable training has been required to reach that status. In the early days of typewriting, it took a rare act of courage and belief to attempt to type

We thank Eileen Conway, Donald Gentner, Jonathan Grudin, Geoffroy Hinton, Paul Rosenbloom, and Craig Will for their assistance and many discussions with us on the nature of the typing data, their work in collecting and interpreting typing errors, and their discussions on the underlying response mechanisms. Gentner has provided a large set of keypress reaction time data, and Grudin, Gentner, and Conway have taken high-speed films and video tapes of skilled typists. This ongoing research and the several large corpora of data on typing performance have been of considerable assistance to us in the preparation of this paper. (The several studies will be published as separate research reports.) Stu Card, Tom Moran, and Bill Verplank have provided useful discussions and analyses, as well as data on typing performance and keyboard design.

Research support was provided by the Office of Naval Research and the Naval Air Development Center under contract N00014-79-C-0323.

Requests for reprints should be sent to David E. Rumelhart, Program in Cognitive Science C-015; University of California, San Diego; La Jolla, California, 92093, USA.

without looking at the keyboard. In those days, the arguments were over whether it was best to use 2, 3, or 4 fingers, and the concept that someone could touch-type, without looking at the keyboard, using all 10 fingers, was not taken seriously (see the discussion in Beeching, 1974, pp. 39 ff.). Matters are not made easier by the fact that almost all contemporary typewriters use the Sholes (qwerty) keyboard, designed in 1873 to minimize jamming of the keys by maximizing the distance between frequently typed pairs of keys, without regard for ease of learning or typing.

Our goal is the study of human skilled performance. In this paper we concentrate on the problem of control of skilled motor movements. In particular, we study typing, for this is a highly skilled motor activity with enough experts available that we can study performance with some ease. Moreover, it is fairly easy to monitor keypress latencies, although more difficult to monitor the actual trajectories of the fingers, or the times at which the fingers begin their movements.

THE BASIC PHENOMENA OF TYPING

In this section we review briefly some of the empirical results that have been found in the study of typing. We especially emphasize those that we find most suggestive of the underlying mechanisms. Following this section we present the model for the simulation of these phenomena, and then, in the last part of the paper, we return to examine the phenomena in some detail, contrasting the model performance with that of the data from the literature and from our own observations.

The fundamental phenomena fall into three categories: those involving timing of keystrokes, those involving errors, and those involving the general organization of the typing process. We discuss these in turn, using the data recorded as one of us (DER) typed a 90,000 keystroke manuscript from dictation. Unless otherwise specified, the data discussed in the paper come from our corpus. Figure 1 shows the Sholes keyboard and the standard American mapping of fingers to keys.

The Timing of Keystrokes

People can type very quickly. World champion typists can type at rates up to 200 words per minute. This involves a stroke every 60 milliseconds, close to the neural transmission time between the spinal cord and the periphery. There cannot be much feedback between strokes being performed so rapidly. Even relatively ordinary typists can routinely generate strokes at rates almost as rapid as this. For example, of the 1656 times th was typed by our subject, 414 times the interval was less than 63 msec. The th interstroke interval was less than 75 msec half of the time.

Speed, however, is the simplest of the phenomena that need to be accounted for. Overall, there are five different sets of timing phenomena that provide strong constraints on the structure of a possible model of typing:

STANDARD QWERTY KEYBOARD





Figure 1. The standard American keyboard (the "Sholes" or "qwerty" keyboard) and the mapping of fingers to keys.

1. people can type very quickly;
2. cross hand interstroke intervals are shorter than those within hands;
3. within hand interstroke intervals appear to be a function of the reach from one key to the next;
4. the time for a particular interstroke interval can depend on the context in which it occurs;
5. there is a negative correlation between the intervals on successive strokes -- especially when the alternate strokes occur on alternate hands.

We discuss these in detail at the end of the paper.

Pattern of Errors

Errors are of special importance, for some of them give strong clues as to the underlying mechanisms. Thus, the existence of transposition, doubling, and alternation errors has played a major role in determining the structure of the model.

Transposition errors. One of the most common and most interesting categories of errors is transposition, the reversal of two adjacent letters. The large majority of these errors occur across hands. Shaffer (1976) reports that of his subject's transposition errors, about 90 percent were cross hand transpositions. For our subject, the figure is 75 percent. Examples of transpositions from our subject include:

because -> becuase
which -> whihc

Of the within-hand errors, half involved adjacent keys (such as e and r and o and p), as in

supremely -> supermely

One interesting example was reported by Shaffer (1976):

went down -> wne todnw

the four strokes on the right hand (the n, space, o, and n) have all been displaced with respect the five left hand strokes.

Doubling errors. When a word contains a doubled letter, the wrong letter is sometimes doubled. Thus, look can become lokk. This error was pointed out by Lashley (1951) and by Shaffer (1976) as being diagnostic as to the nature of motor control. Our corpus of transcription typing included only one example of a doubling error of this sort:

school -> scholl

but we have collected many doubling errors from our samples of composition typing (while using the laboratory computer). For example:

gibbs -> giibs
Screen -> Serren

Alternation reversal errors. This is akin to the doubling error, but with an alternating sequence. Thus in the word these the ese is an alternation. Samples observed during composition typing include:

these -> thses
there -> threr
were -> wrer

Other errors. In addition to the errors of transposition, doubling, and alternation, a number of other forms of errors occur. However, these are not so critical in determining the form of the typing model, and so discussion of these is deferred until later. (For example, we believe that some of these errors come from factors outside the scope of the present model.) These other errors include:

homologous errors;

capture errors;

omission errors;

misstroke errors.

These errors will be discussed at the end of the paper.

The General Organization of Typing

Finally, there are two other observations that we have used as strong constraints: the overlapping of hand movements and the unit of organization of the strings to be typed.

Skilled typists move their hands toward the keys in parallel. In our laboratory, Gentner, Grudin, and Conway (1980) have carried out a photographic analysis of a skilled typist using a high speed motion camera. The results of this study show the fingers of the hand in almost constant motion, with fingers starting to move toward their destination before the several preceding characters have been typed. A serial model of typing in which each finger in turn makes its stroke is incorrect. Rather, there seems to be a coordinated structure that allows the control of several fingers simultaneously.

The units of typing seem to be largely at the word level or smaller. In Shaffer's (1973) study of the units of typing for one skilled typist, he had his subject type normal prose, random words, random letters, and foreign words. He found almost no difference between typing prose and random words (the mean keystroke latency was 107 msec for prose and 104 msec for random words). Random letters were typed much more slowly than normal text (192 msec), and German text was typed at an intermediate rate (149 msec: the typist didn't know German). Shaffer also found that when look-ahead was limited, the typist needed to see at least eight letters ahead of where she was typing in order to maintain her normal typing rate. (The rule of thumb that has emerged from this and similar studies is that a typist looks ahead about one

second's worth of text: slower typists look ahead fewer letters, faster typists more.)

A MODEL OF TYPING

Control of the fingers poses a number of complexities, and the cognitive specification of the actions to be performed must be compatible with both the existing knowledge of mental structures and of the phenomena of typing, especially the factors discussed in the previous section. Our analyses of these issues lead us toward a model that has the following properties:

1. control of action sequences by means of schemata;
2. selection of appropriate motor schemata through a combination of activation value and triggering condition;
3. the representation of letter typing by means of a pure type theory (i.e., one with no type-token distinction).
4. the need for distributed (local) rather than concentrated (central) control of movement.

We start with the assumption that motor control of learned movements is represented by means of motor schemata. A motor schema is an organized unit of knowledge, differing from the form of knowledge widely studied in the literature on memory, language, and thought only in that it has as its output the control of body movements. This is not a new concept. Actually, the term "schema" was originally introduced into psychology for the use in skilled motor control by Head (1926) and is still used for that purpose (cf. Schmidt, 1976).

We propose that one of the functions of schemata is to act as motor programs. The term "motor program" is to be understood by analogy with the the term "computer program." We believe there has been some confusion in the literature on skills in this regard, with critics of the notion of motor programming acting as if a program were a fixed action sequence, specified in complete detail before the actual movements. According to our view, motor programs are flexible, interactive control structures, capable of calling upon sub-programs, passing parameters to be bound to program variables, and making local decisions as a result of current conditions (which might include information from feedback channels, from perception, or other sources of knowledge). A motor program is not a fixed action pattern of movements. It is a set of specifications or control statements that govern the actions that are to be performed, with considerable flexibility in the specification of the actions. A program specifies the rules that are to be followed in the action, not the actual motions.

The ATS Formalism

The basic framework that we follow is called an Activation Triggered Schema system (ATS). The model consists of a set of schemata, each with activation values. A schema has an activation value that reflects the total amount of excitation that it has received. The normal, resting value for a schema is zero. It can increase when the schema is "activated" or decrease when the schema is "inhibited." Schemata interact with one another, and the activation value reflects this interaction, as well as the effects of decay and other sources of activation and inhibition. When appropriate conditions have been satisfied, a schema may be "triggered," at which time its procedures become operative and control whatever operations they specify.

Different schemata are often interconnected. Moreover, one schema may call upon other schemata to perform specific tasks, much as a computer program calls upon subroutines or coroutines. When one schema calls upon another, the initiating schema is called the "parent schema" and the called schema is the "child schema." Each schema can serve in any or all of three ways: as a program in control of operations, as a parent schema that initiates the operation of other schemata, or as a child schema, invoked by a parent.

A particular schema might be invoked by a parent schema, set into motion some operations, and then itself serve as a parent to its child schemata. Usually, but not necessarily, when a child schema has completed its operations, control returns to the parent schema. Thus, the schema for typing the word the might be initiated by the triggering of its parent schema, which then controls the activation and triggering of the child schemata for the letters t, h, and e, which in turn activate the child schemata that control the actual finger, hand, and arm movements.

The Simulation Model

Figure 2 illustrates the basic structure of the model. The model incorporates the principles discussed in the previous sections, plus specific control mechanisms for the activations and selection of particular hand and finger movements. The input of the model is a string of characters that constitute the text to be typed. The output is a sequence of finger movements, either displayed on a visual computer-controlled display as the movement of the hands and fingers over a typewriter keyboard, or as a series of coordinate locations for the relevant body parts.

The Hand

The basic hands employed in the model can be seen in Figures 1 and 2. Each hand consists of a palm and the five fingers: thumb, index, middle, ring, and little. The model hands are a simple schematic approximation to the real hands. We let the palm and each of the five fingers have two degrees of movement: up and down the keyboard and inward and outward from the center of the keyboard. (Fingers actually move "up" and "down" by extension and flexion in the third dimension, but in the model, we view the hands from above in a two-dimensional projection. Thus, the extension and flexion of the fingers are represented by lengthening and shortening.) "Upward" is defined to mean toward the top row of the keyboard and "downward" to mean toward the bottom row (the space bar). Sideways motion is relative to the hands, so that "inward" means toward the center of the keyboard (toward the thumbs) and outward means towards the ends of the keyboard (toward the little finger). We assume that the palm can move any distance in any of these directions. Each finger is assumed to have a region of additional movement relative to the location of the palm. The exact amounts of movement of each finger relative to the palm are parameters of the model and depend on the finger in question. We selected values that roughly correspond to the comfortable reach of our subject (DER). Thus, in the simulation, the index finger can move 1.0 key spaces inward, only 0.05 key spaces outward, and 0.8 rows upward and downward. Because the little finger is at the end of the hand, its horizontal range of movement is asymmetrical: it can move 1.0 key space outward, 0.05 key spaces inward, and 0.3 rows upward and downward. For simplicity we assume that the thumb is rigidly attached to the palm and can only move horizontally (inward and outward). Table 1 gives the assumed amount of movement in each direction for each of the fingers.

It is important to note that there is a tradeoff between palm movement and finger movement. Thus, each final position can be reached by infinite combinations of palm and finger positions. Taking advantage of




Figure 2. The information processing system involved in typing: The perceptual processes output their identification of the input to a buffer which maintains the information while the appropriate response schemata are activated and take control of the actions. The input of the model is words; the output is the movement of the hands and fingers across the keyboard.

Table 1

Amount of Finger Movement Permitted for Each Finger in the Simulation Model

	Up	Down	In	Out
Thumb	0.0	0.0	2.00	2.00
Index	0.8	0.8	1.00	0.05
Middle	1.0	1.0	0.10	0.10
Ring	0.7	0.7	0.10	0.10
Little	0.3	0.3	0.05	1.00

this tradeoff is part of the process of efficient typing.

The Keyboard

The keyboard is modeled after the conventional Sholes or "qwerty" keyboard, restricted to the letter keys, semicolon, comma, period, and RETURN. We do not model the use of the SHIFT key or the number keys. The position of the RETURN key is taken to be just to the right of the semicolon.

Finger assignments are those used in the standard American method for the teaching of touch typing (shown in Figure 1). The "home" position of the hands is taken to be that with the fingers of the left hand resting on the a, s, d, and f keys, and those of the right hand on j, k, l, and ;. The palms are located so that the fingers can be on the home keys with the fingers in neutral position relative to the palms.

The Perceptual Processes and the Parser

There are a number of interesting issues involved in determining an appropriate synchronization between the operation of the perceptual processes and the rest of the typing model. For example, the perceptual processes have to read the words of the manuscript at just the proper range of rates so that there are always character strings to be typed available to the parser in the buffer, but not so many strings that the buffer overflows. Although these are of real importance to the typing process, they are not included in the simulation model and so we do not address them here. We do assume that the output of the perceptual system is words, not letters. The operation of the typing model begins where the perceptual processes leave off: with strings of words in the buffer of the parser.

The job of the parser is to transform the word strings in the buffer to appropriate activations of the key press schemata. The basic task of the parser is to divide the words into the individual characters, then activate the appropriate key press schema for that character.

The Activation Process

Figure 3 illustrates the basic assumptions of the activation process using the word very as an example. First, the schema for the word is activated by the perceptual system and parser. This, in turn, activates each of the child schemata for keypresses. Each keypress schema specifies the target position, with position encoded in terms of a keyboard centered coordinate system; upward-downward (upward is positive) and inward-outward (inward is positive). These target positions are sent to the response system which then must configure the palm and finger positions properly. Each keypress schema inhibits the schemata that follow it. This means that proper temporal ordering of the keypress schemata is given by the ordering of the activation values. In addition, the activation values are noisy, which leads to occasional errors.




Figure 3. The interaction of activations when the word very is to be typed. The very schema is activated by the perceptual system and parser, activating each of the schema for the component letters. Each letter schema specifies the target finger position, specified in a keyboard coordinate system. L and R stand for the left and right hands, and I and M for the index and middle fingers. In the coordinate system, the first variable is upward/downward (up is positive) and the second variable is inward/outward (inward is positive). The keypress schemata receive information about the current finger position from the response system. Each letter schema inhibits the activation of all letter schemata that follow it; inhibition is shown by the lines with solid circles at their termination.

The response system feeds back information to the keypress system about the current location of the fingers. This is essential for the triggering conditions of the keypress schemata. Whenever the current finger position is within some criterion distance of its target position, and the relevant schema is the one most highly activated, then the triggering conditions are satisfied and the actual keystroke is launched. Shortly after the launch (the amount of time being a parameter of the model), the keypress schema deactivates itself, resulting in a release of inhibition for all the succeeding keypress schemata.

Repeated letters imply that there are no token schemata. The existence of doubling and alternation errors pose special problems. Consider the word book. According to the arguments we have just presented, the word would be represented by schemata for each of the letters: b o o k. It is easy to see how such a representation could lead to transposition errors (such as boko) but not to doubling errors. It would be easy to make up a schema for a doubled letter (so that the word would be represented by the schemata b double-o k), this would not lead to the doubling errors either.

The doubling error turns out to have two major implications. First, it implies that there are special schemata that signal the existence of doubled letters, and that occasionally these schemata get applied to the wrong letters. In computational terms, this means that the binding between the arguments of the special schemata for doubling occasionally get made improperly. Second, the need for a special schema to mark doubled letters implies a difficulty in having the regular letter schema signal the double. Why isn't the word book represented by the schemata b o o k? The reason would seem to be that this would require two instances (tokens) of the schema for o; the existence of the doubling error implies that such repeated tokens of a schema is not possible.

Thus, the existence of doubling errors forces us to a pure "type" model, in which each letter can only have a single keypress schema; the keypress schemata exist only as "types," with no "token" schemata. There must be a special schema that signals the presence of a doubled letter. Moreover, there must be a weak binding between the special schema and the arguments upon which it operates. In the model, the arguments are not bound to the schemata, but are established via activation values. The most highly activated keystroke schema is triggered when it is within a criterion distance of its target. After triggering (and the resulting launch of the keystroke), this keypress schema can become "bound" to a doubling schema if one exists with a higher activation value than its own. Because activation values are noisy, occasionally this leads to errors in the linking of keypress schemata to a doubling schema.

The existence of alternation errors leads to the same conclusion; there must be a special schema that signals the presence of alternating letters, with a weak binding between the schema and its arguments, and the mechanism proposed for alternations is similar to that for doubling.

Here is an example of the typing of a word with a doubled letter. The word book is represented by the activation of four schemata: b double o and k. Each schema inhibits all that follow it, in the regular fashion. The operation now is much as we illustrated before except that after a keystroke schema has been triggered, it checks for the existence of a double schema whose activation value exceeds its own value. The b will initially be the schema most highly activated, and when the finger gets within a criterion distance from the key, the b keystroke will be launched. Now the double schema will have the highest activation level. However, the double schema does not command any motor response, and it allows control to be passed to the schema with the next highest activation value. The next schema is the one for o. It proceeds normally. As the keystroke is launched, the o schema notes that there is a double schema whose activation value is greater than its own. Whenever this condition occurs, the keypress schema deactivates the double schema and, after the keystroke, does not deactivate itself. As a result, at the completion of the keystroke, the schema is again triggered, launching itself a second time. At the launching of this second keystroke the doubling schema is no longer present, so that typing of the rest of the word can continue. Noise in the activation levels occasionally causes this mechanism to go awry so that the double schema gets associated with the wrong keypress schema, causing the wrong character to be doubled.

We suspect that there is a similar mechanism for alternating keystrokes. However, when we added such a mechanism to a version of the model, it led to types of alternation errors that were never observed. Let the symbol "-" stand for space. Words such as -a- that had but one letter were spelled a-a rather than -a-. It is possible that the problem here lies with the parsing mechanism, not with the alternation schemata. Perhaps it is wrong to treat a word such as a or I as an alternation of space letter space but rather, each word should either be followed or preceded by a space (so that only a single space is ever tied in with a word). Because we did not attempt to model the perceptual and parsing process, we did not pursue this possible explanation.

The assumption that there are no "token" nodes, but only "type" nodes causes special problems for any word that contains repeated instances of the same letter that is not part of a double or an alternation (e.g., the e and p of perception). As a result, if a word contains a repeated instance of a letter, processing of that word is blocked at the repeated letter until the keystroke for the previous instance has been completed.

Movement

One of the issues in the control of movement is the degrees of freedom problem; there are many ways to perform any given task, yet the motor programming system must determine one of the many possibilities. But this can be a virtue, for it allows any given movement to serve more than a single goal. The extra degrees of freedom can be used to optimize movement toward many goals at the same time. In the case of typing,

this means that the set of arm-hand-finger configurations can be chosen so as to optimize the striking of as many keys in a sequence as possible, without requiring an overwhelming computation.

In the model, each active schema pushes its relevant hand and finger toward its desired key at the same time, and the final overall configuration is determined by the competition among these forces. Each schema pushes with a force proportional to its activation level. As a result, the forces are weighted so as to cause the letter schema that is next in line to be typed to approach its key most quickly. The actual location of each finger is determined by the sum of the extensions of the finger and the hand. To type a particular typewriter key, it is only necessary that the end position be correct. Each keypress schema specifies its desired endpoint by specifying the total extension of the hand plus the relevant finger. The endpoint configuration is reached through an iterative relaxation process that only involves local computation. That is, the desired target extension is sent to the muscle system, and each finger and hand combination moves a very small distance toward the goal. The algebraic sum of all of the competing movements toward the various keys represented by the activated schemata determines the total movement toward the goal. Then the process is repeated, with a small movement resulting from each iteration of the process. Because of the unequal weighting of activations, the process will eventually cause the most highly activated schema to move its finger-palm configuration to within a criterion distance from its target key, satisfying the trigger conditions and launching the keystroke.

The positions of the fingers and palms are represented in the system as a sum of muscle lengths (i.e., a total extension). Thus, a position specification feeds into a comparator which uses the specification as a reference signal. This comparator also receives inputs from a "length detector" for each muscle. These lengths are summed, leading to a measure of the location of the end of the finger. If this positional feedback is greater than the reference signal, the relevant agonist muscles are inhibited and the relevant antagonist muscles activated. If the positional feedback is less than the reference signal, the rate of activation of the muscles would be unchanged.

Each schema activates its relevant muscle system proportional to its degree of activation. The activation of all currently active typing schemata send activation to the same muscle system, their rates summing algebraically. The movement is determined by the summed result from all of the response schemata driving the muscles. The actual velocities and hand configurations are determined by the interaction among the competing forces driving the various fingers toward their ultimate goals.

At some point the finger gets into its appropriate position and its triggering conditions are satisfied. This initiates the actual keystroke movement, launching a ballistic keystroke movement. Upon launching, a number of changes occur: the schema begins to deactivate itself and reduces its activation level. In addition, the launching of other fingers on the same hand is inhibited for a period after the launch has

been initiated. There is no feedback for the keystroke once it has been launched; the movement is ballistic, and so, after a brief delay, the schema fully deactivates itself. Aside from this, the system operates normally during the launch, and other keypress schemata may have their triggering conditions met and launch their own keystrokes before the earlier ones have been completed.

This movement scheme has the possibility of a "deadlock" in which the algebraically summed pressures on a hand balance so that the hand freezes and never reaches its goal. In order to forestall such a possibility, there is an additional inhibition among those keystroke schemata commanding the same hand: a progressively increasing inhibition, increases as the time between keystrokes increases. This scheme ultimately resolves all deadlocks.

An Example

Consider how the model might type the string very well. The perceptual processes interpret the input and present the words to the parser. Then the word and keypress schemata are activated, with each activated schema inhibiting the keypress schemata for those keypresses that follow. Figure 4 shows the patterns of activation of the schemata and Figure 5 shows the resulting hand and finger positions photographed from the actual model display approximately every 4 time units throughout the typing of very well. Figure 4 also shows the times when each keystroke is launched, and the times that each key is actually pressed.

At time unit 0 in Figures 4 and 5, the typing of the string has just begun. The hands are still in their home position and the v schema is most highly activated, followed by the schemata for e r y space and w. Activation stops at this point because the next schema needed is that for the e of well, but this keypress schema is already active. No further activation of schemata can occur until after the e in very has been typed.

For the first few time units, the activation values remain constant (except for noise) while the hands adjust their positions. By time unit 3, the left index finger is clearly reaching down toward the v key, while the left middle and ring fingers are extending toward the e and w keys. Meanwhile, the right index finger is moving up toward the y key. Note that since all of the pressures on the right hand pull the hand up, while some on the left hand pull it up and others pull it down, the right hand is further up than the left.

By time unit 7, the left index finger has already been launched. Meanwhile, the left middle and ring fingers have extended even more while the right index finger is still closer to the y key. In Figure 4 we see that the launching of the v schema reduced its activation level slightly, thus increasing the activation levels of the others. At time unit 11, the v key is pressed, and the e, y and w fingers are already nearly in place. The v schema totally deactivates itself, further




Figure 4. Patterns of activation for the various keypress schemata during the typing of the string very well. D stands for double schema. See the text for details.




Figure 5. The hand and finger configurations photographed from the display of the working model during the typing of the string very well. Each frame of the figure displays the patterns approximately 4 time units after the previous frame. The * indicates an actual press of a key.

elevating the remaining activated schemata.

The remaining parts of the figures illustrate the completion of the process. Note that by time unit 22, both the r and the y have been launched, with the r just slightly ahead of the y. Also by this time, the e in very has already been struck, allowing the activation of the schemata for the remaining letters of well (e, D --for double--and l). Because the activation levels are noisy, on occasion a schema other than the next one in line is most highly activated. When this happens there is the possibility for a transposition error. In this particular run no errors occurred even though the wrong schema was most active on several occasions. This is because the triggering conditions were never met for the wrong schema when it was also accidentally the most highly active.

APPRAISAL

In this section, we return to the fundamental phenomena of typing and show how they are produced by the model. In order to evaluate the model, we gave it a text of slightly over 2,000 words to type. The pattern of keystrokes and times were collected from the simulation and analyzed in exactly the same fashion as we had analyzed the data from our human subject.

Timing

People Can Type Very Quickly

We believe that the overall structure of the model explains the speed with which people can normally type. According to the model, people type quickly because they are carrying out many actions at once. The very fast 60 msec interstroke intervals don't represent responses initiated and completed in 60 milliseconds. Rather, since responses are made in parallel, their termination points differ by 60 milliseconds, but the responses themselves may take much longer than that to complete (see Gentner, Grudin, and Conway, 1980). Moreover, although our model requires feedback as to the location of the fingers, this feedback is required only up until the launch of the stroke. At that point, the response is assumed to be ballistic. In this way, very quick feedback from the execution of a response is not required, instead, all that is required is continual feedback of location from the fingers.

Cross Hand Interstroke Intervals Are Shorter than Those Within Hands

In the pattern of interstroke intervals, the first and most obvious observation to be made is that when successive strokes are made on different hands, the interstroke intervals are shorter than when both strokes are made with the same hand. This result has been reported often (cf. Shaffer, 1976). Part A of Figure 6 shows the distributions of keystrokes within and between hands from our data. The difference between the median interstroke intervals for the two kinds of strokes is about 54 msec. The difference is even more striking when we compare interstroke intervals on different hands to those which require the same




Figure 6. Part A shows the overall distributions of keystrokes within and between hands from our data. Part B compares the distributions of times to type a y when in the word by and when in the word my. Part C compares the distributions of times to type an e when in the word me and when in the word be.

finger on both strokes. Figure 6B shows the almost non-overlapping distributions of times to strike the y in the word by and in the word my. The b and the y are struck by different hands whereas the m and y are struck by the same finger. Figure 6C shows the distribution of times for the e in the words me and be. Here the difference is somewhat smaller than in Figure 6B since the b and e do not require the same finger.

In the model, as in skilled typists, cross hand interstroke intervals are faster than those from consecutive strokes on the same hand. Table 2 compares the average times of the model for within and between hand interstroke intervals and compares those times with the comparable medians for our subject. Clearly, the same general difference is observed in each -- within hand strokes are substantially slower. According to the model there are two reasons for this result. First, and most important, when both letters are typed with the same hand, the typing of one key is likely to pull the finger from the next key it is supposed to strike. In contrast, if the two keys are on different hands, the activation level for the second keypress is likely to have the highest value on that hand. Therefore, the finger for the second key is very likely to be near its target key at the time the first keystroke is launched. In this way, the cross hand stroke will, on the average, occur sooner than the within hand stroke. Second, when a keystroke is launched, the launching of subsequent strokes on the same hand is inhibited, but launching of subsequent strokes on the other hand is not. Thus, overlapped launches are allowed between hands, but not within hands. This too contributes to the advantage of between hand strokes in speed.

Within Hand Interstroke Intervals Appear to Be a Function of the Reach from One Key to the Next

Keystroke times are affected by the distance between keys. To see this, first consider the time required to type an e when the preceding letter uses the same finger. (The e is located on the top row and is typed with the middle finger of the left hand.) It takes 165 msec to repeat an e (to move from the e to the e), 201 msec to go from the home row key of d to e and 215 msec to go from the bottom row key of c to e. Part A of Figure 7 illustrates these results. Now let us consider the times when different fingers are used. Again, consider how long it takes to type an e, but this time, when the preceding key is typed with the index finger of the left hand. These times and the relative keyboard positions are shown in Figure 7B. To go from the top row key of r to the e (also on the top row) only takes 145 msec. To go from the t to the e takes 159 msec. The times to go from the middle row keys of f and g to the e are 168 and 178 msec, respectively. And the times from the bottom row keys of v and b are 178 and 195 msec, respectively. To a first order approximation, the longer the reach the longer the interstroke intervals. Note, that we should not be surprised that the re and te times are less than the ee times since the ee strokes cannot overlap at all, whereas those from t to e and from r to e can.

Table 2
Interstroke Intervals

	Within Hand	Between Hand
Simulation (mean)	7.60	4.71
Data (median)	167.67	111.37

ACTUAL DATA




Figure 7. Part A shows a comparison of times to type a key when the preceding letter was on the same hand and the same finger: e to e, d to e, and c to e. Part B shows a similar comparison, except with the preceding letter on the same hand going from a letter typed with the index finger to an e (typed with the middle finger). The arrangement of the letters in the figure is the same as the arrangement of the keys on the keyboard. Times are shown in milliseconds.

The model generates a pattern of within hand interstroke intervals as a function of the reach between consecutive strokes that is similar to that of skilled typists. Part A of Figure 8 illustrates the simulated interstroke times for within finger strokes. If anything, the model shows a stronger dependence on distance here than do human subjects. This may be because the model assumes that each finger attempts to move at constant velocity. It may well be that long excursions like that between the c and the e keys, actually involve higher velocities than for shorter excursions, such as from d to e.

Part B of the figure shows the average simulation results for the times between each of the keys struck by the left index finger to the striking of the e key. The model shows the same general pattern of results as the data. The longest reach involving the b key is the longest time. The shortest reach keys, from r and t, show the shortest times. The ge time is slightly fast relative to fe. One oddity of the simulation results is that in the simulation, the te time is faster than the re time. We do not fully understand this. It could be that with the sample text typed by the model, the contexts within which re occur tend to give less time for the e finger to get near the e key while the r is being struck than while the t is being struck.

Correlations between the model and subjects. Although these demonstrations would seem to indicate that the simulation results do show about the right pattern of interstroke intervals, it would be useful to get a more general measure. For this reason, we chose the 66 most common bigrams from our data, and correlated the simulation data with the average of the data from our subject (DER) combined with the interstroke data from five other subjects, each of whom did transcription typing on a corpus of about 50,000 strokes. (These data were collected by Donald Gentner and kindly made available to us for this purpose.) Figure 9 shows the scatter plot for the mean interstroke interval time over these six subjects compared with the times generated by the simulation. The overall correlation between the model and the averaged data is about 0.86. Thus, although the fit is not bad, the model clearly does not account for all that is happening.

Just how well does the model account for the data? One problem with assessing the fit of the model is our lack of knowledge of how much different subjects differ from one another. Thus, the model may actually capture much of the processes that underlie typing, on the average, without actually fitting any single subject's performance in detail. One way of checking the relationship among the model and that of the several different typists is to compare the fit offered by the model for the data with that offered by another subject. That is, how well do the data from one subject compare with the data of the other subjects? How well any individual subject can describe the rest of the subjects is a measure of the amount of common variance. The results are shown in Table 3. The model accounts for the individual subjects about as well as each of the subjects accounts for the others. This means that the deviations of the model are well within the variability of the subjects. Thus, any further refinement of the timing aspects of the model must be focused on

SIMULATED RESULTS




Figure 8. Interkeystroke times from the simulation. Part A shows times for within finger strokes. Part B shows time between each of the keys struck by the left index finger to the striking of the e key. The arrangement of the letters in this figure is the same as the arrangement of the keys on the keyboard. Times are means of arbitrary model units from the typing of a sample text.




Figure 9. The scatter plot for the mean interstroke interval time for six subjects compared with the times generated by the simulation. Data for five subjects were collected by Donald Gentner and the sixth (DER) is our subject.

Table 3

Correlation Coefficients for six typists and the typing model

SUBJECT:	Average	Simulation	DER	S1	S2	S3	S4	S5
Average	1.00							
Simulation	.86	1.00						
DER	.83	.71	1.00					
S1	.45	.55	.34	1.00				
S2	.85	.78	.69	.45	1.00			
S3	.90	.77	.88	.43	.79	1.00		
S4	.86	.69	.81	.45	.78	.79	1.00	
S5	.85	.88	.74	.41	.85	.81	.76	1.00

the individual differences among typists.

The Time for a Particular Interstroke Interval Can Depend on the Context in Which It Occurs

Shaffer (1978) has shown that the time between successive keystrokes depends on the larger context in which these strokes appear. For example, he collected repeated typings of the words wintry and wink and found that even when the preceding word was controlled, the wi transition was longer in the word wink than in the word wintry. Similarly, the in transition was found to be longer in the word wintry than in the word wink. (Shaffer's data are shown in Table 4.) Thus, a particular interstroke interval depends on the preceding as well as the following letters. These results are reminiscent of the "co-articulation" effects common in speech production (Kent & Minifie, 1977; Liberman, Cooper, Shankweiler, & Studdert-Kennedy, 1972).

The model also predicts that the time for a particular bigram depends on the larger context in which the bigram occurs. This dependency occurs because the surrounding strokes are pushing and pulling the fingers toward other keys, thus slowing or speeding the progress between keys. In order to see if the context dependencies are generally of the sort reported by Shaffer (1978), we ran our model on the words wink and wintry to compare the pattern with that reported by Shaffer. The results are given in Table 4. The time to the first letter is somewhat longer for wink than wintry, because the w in wintry is speeded along by the t and r which are both pulling the hand upward. The time from the w to the i is also a bit longer for the word wink, presumably because the k is holding the i back a bit, whereas the y in wintry is pulling toward the upper row. The i to n time is faster in the word wink. There are three possible reasons for this. First, because it took longer for the i to be struck, there was more time for the index finger to move toward the n. Second, when the i is released both the n and the k work together to bring the hand down. Third, the y in wintry is trying to hold the hand and index finger up and slowing the time to the n.

However, although the context effects in the model are similar to those found by Shaffer, the timing pattern is otherwise quite different. Thus, in the simulation, the time to strike the i key is very short, whereas it is long in Shaffer's data. There are other anomalies as well. Thus, in the model, the time to type the n and the k in wink are about the same, whereas in the data, they differ substantially. In wintry, the model predicts that t should be typed much more rapidly than the data indicate, and the data show the n and the r to take about the same time, which is not true in the model. So, overall, the model handles some of the phenomena, but misses others.

There Is a Negative Correlation between the Intervals on Successive Strokes -- Especially When the Alternate Strokes Occur on Alternate Hands

Table 4

Mean Response Latencies as a Function of Context

	w	i	n	k
Simulation (arbitrary units)	7.9	3.0	6.8	6.0
Shaffer (msec)	100	140	90	150

	w	i	n	t	r	y
Simulation (arbitrary units)	6.9	2.0	10.4	1.8	6.5	4.8
Shaffer (msec)	90	110	140	90	140	100

It is possible to treat the successive intervals between keystrokes in a corpus of typing data as a time series and look for dependencies between successive strokes. If one interval is longer than usual, is the following interval also longer than usual, is it shorter than usual, or is it unaffected? The autocorrelation of successive intervals is a measure of this dependency.

Shaffer (1978) set out to study the nature of this relationship. He had to control for a number of potential problems. First, as we have just noted, the time between successive strokes on the same hand is largely determined by the reach between those strokes. This factor would tend to reduce the autocorrelation to zero. In order to get around this problem, Shaffer noted that the variance between strokes on different hands was less than the variance between strokes on the same hand. Moreover, these times are not subject to the problems of differential reach on the two hands. He produced texts consisting entirely of words in which successive strokes are made with alternate hands (e.g., with and such). Another problem involves long term variations in the typist's rate of typing. If a typist slows up for some period and then speeds up, this produces a positive autocorrelation. To avoid this spurious autocorrelation, Shaffer decided to look over the relatively short span of 100 keystrokes (10-15 seconds of typing), find the autocorrelation within each block of 100 strokes, and then take the average over blocks, thus minimizing the effects of overall drift in typing speed. When these precautions were taken, Shaffer found the results illustrated in Table 5; the autocorrelation with a lag of 1 (between strokes i and $i+1$) was negative -- that is, long intervals tended to be followed by shorter intervals and vice versa. However, the autocorrelation with a lag of 2 (between strokes i and $i+2$) was positive.

Figure 10 shows the distribution of times between successive strokes in the word with from our data. Note that, although the keystrokes alternate between the hands, the wi time is fast, the it time slow, and the th time fast. As Shaffer points out, such behavior is to be expected from a metronome model of typing in which the typist initiates a stroke regularly to some sort of internal clock beat, but then there is some variance in the actual execution of the response. In this case a short interval will cause the next one to be long. As we show, this behavior is also to be expected from models not involving such an internal clock.

All other things being equal, the model predicts a negative correlation between adjacent strokes. In general, the prediction follows from the fact that a slow response gives the following finger more time to get into position, thereby speeding up the following response. If there happens to be a fast response, the following stroke has less time to get into position and so the second interstroke interval is somewhat longer. In order to quantify this effect in our simulation model, we extracted all strings of keypresses which involved either a right-left-right-left or a left-right-left-right sequence of strokes (excluding spaces and returns). We then computed the correlation between the three

Table 5

Autocorrelations of Response Latency

Autocorrelation with Lag	Simulation	Shaffer (1978)
1	-0.58	-0.28
2	+0.44	+0.10

"WITH" Latency Distributions




Figure 10. The distribution of keystroke times in typing the word with.

interstroke intervals thus defined. The results of this computation are shown in Table 5. Successive strokes are always negatively correlated. The correlation between the adjacent times (corresponding to an autocorrelation with a lag of one) is -0.58 . The correlation between the first and third intervals (corresponding to an autocorrelation with a lag of 2) is 0.44 . The pattern has the same direction as the data reported by Shaffer, but somewhat more extreme.

Errors

Transposition Errors

This class of errors was discussed earlier in the paper. The proportion of transposition errors in the model is determined by the amount of noise in the activation levels. In our data we observed transpositions at about a rate of 1 for every 1800 keystrokes. In the simulation, we examined 12,000 keystrokes, so if the error rate had been the same as that of our subject, we would only have seen about 7 errors. We therefore adjusted the noise level to yield errors at a rate of about 1 for every 30 keystrokes. Despite the large difference in rate, the basic pattern of errors is similar. For example, a large majority (76 percent) of the transpositions in the simulation occur across hands. Although this is a somewhat smaller percentage than Shaffer's subject shows, it is about the same as our subject and comparable to values reported for others. Examples of the cross hand errors generated by the model include:

special -> speical
course -> coruse
dramatic -> dramaite

We also found examples of multiple transpositions in the simulation results (although none were as impressive as the example from Shaffer that we reported earlier). We did find

dismal -> dsiaml

(in which the right hand keys appear to be delayed relative to the left)
and

vitamins -> ivtmaisn

(in which, with the exception of the final "s", the left hand strokes appear to be delayed with respect to the right).

We also found within-hand transpositions in the simulation data. Examples of these include:

master -> masetr
result -> ersult

The reason that the model gives this pattern of transposition errors should be clear enough. A transposition error can only occur if the wrong schema has its trigger conditions satisfied prior to the correct schema. This means that the wrong finger must both be in place to strike its key and its schema must momentarily have the highest level of activation. This is more likely to occur across hands than within hands because the next finger on the other hand is less constrained by the motions of other fingers: the next finger to type on the opposite

hand always has a speed advantage. The longer a finger is in position, the greater the chance that noise will momentarily cause its schema to become most active and thus be triggered. Within hands, the transpositions should involve keys that are near one another, so that the movement of the palm helps rather than hinders movement toward the keys by the next finger that is to launch a keystroke.

Doubling Error

Doubling errors played an important role in the development of the theory behind the model so it is interesting to see how well the model did at producing them. At the level of noise employed in the simulation, 17 doubling errors were generated. Examples included:

sheer -> sherr
following -> foolowign
little -> liittle

Thus, as with actual observations of errors, we can find doubling errors in which the double is assimilated to either the preceding or following letter.

Alternation Reversal Errors

The version of the simulation model that we ran had no provision for the alternation encoding and thus, none of these errors were generated. Inclusion of errors of this type requires the addition of a special function whose sole purpose would be to encode alternations and would thereby generate the errors. As mentioned earlier, an attempt to devise such a mechanism led to kinds of errors that are never observed.

Homologous Errors

In this class of errors, the wrong hand is selected, but then within the hand, the correct finger and the correct finger movement is performed. Thus, the erroneous stroke is anatomically homologous to the correct one. An analysis of typing errors from the collection of Lessenberry (1928) indicates that homologous errors occur quite frequently. Lessenberry's corpus contains 60,000 typing errors from people of all skills. However, in the various corpora of errors that we have examined that come from skilled typists, there are few homologous errors. We suspect that errors of this type occur more frequently the less the skill of the typists, but we do not have firm evidence for this belief.

The model does not generate homologous errors. We believe that these errors are probably generated by errors in the developing response schemata. The mechanism would have to involve a weak binding between the keypress schemata and the argument that specifies the hand to be controlled, perhaps much in the same way that we provided a weak binding between a schema and its arguments in the case of doubling errors; the

The analysis of the error matrix was performed by Jonathan Grudin and Craig Will.

right finger and place will be commanded, but on the wrong hand.

Capture Errors

This category of error occurs when one intends to type one sequence but gets "captured" by another that has a similar beginning (Norman, 1981). Examples include:

efficiency -> efficient
incredibly -> incredible
normal -> norman

There is no provision in the model for capture errors. We suspect that these occur at either the perceptual or encoding stage of the system, neither of which are dealt with in the model.

Omission Errors

These errors occur when a letter in a sequence is omitted. Examples include:

amount -> amont
education -> educatin
lunches -> luches

Shaffer (1976) has argued on the basis of his timing studies that omissions typically result from aborted strokes, with the timing the same as if the letter had been typed. It is as if the stroke simply involves a depression of the key insufficient to result in the letter being printed. The single most common omission that we observed involves the o in ion suggesting that the ring finger is pulled away too soon as the index finger reaches for the n key.

The model had no provision for omission errors. Presumably, a mechanism could be constructed in which assumed a certain percentage of launches were defective, with the deactivation of the schema occurring before the actual conclusion of the launch.

Misstrokes

We classified about 10 percent of the errors in our collection as misstrokes, defined to occur whenever a key nearby the intended key is inappropriately struck. Examples include:

awareness -> awareneww
because -> becajse
believe -> beoieve
the -> tje

The simulation model did produce a number of misstrokes. Misstrokes occurred when the launch time was longer than usual so that the keystroke schema was already self inhibited before the launch was complete. This allowed the hand to pull away and when the stroke completed it struck a nearby key. Examples of such misstrokes are

distances -> disrances
entire body -> entird body
however -> howevdr

In each case, a nearby key is struck as the hand and/or finger moves toward its next appointed task.

The General Organization of the Typing Process.

Skilled Typists Move Their Hands Toward the Keys in Parallel

The parallel overlapping response patterns observable in skilled typists is clearly apparent in the structure of our model. In the model the hands are constantly in motion, each finger always pushing toward its next key. If anything, our model overplays this aspect. It is unlikely that expert typists always move their fingers as soon or as far as they can, as our model does. Thus, our model might be considered an extreme case of the kind of behavior actually engaged in by expert typists.

The Units of Typing Seem to Be Largely at the Word Level or Smaller

The units of typing in our model are words and letters. No larger units are employed. We assume that words are initially encoded in terms of words which, in turn, activate the relevant keypress schemata. Thus, our model operates as well on random word strings as on prose. More problematic, however, is the fact that our model, as it actually works, would perform as well on random letter strings as on words. As we have already noted, human typists are much poorer on random letter strings. There are a number of reasons why our model may behave wrongly here. In the first place, we have not equipped our model with a lexicon of words, so it simply assumes that there are word schemata for every string of letters bounded by spaces. We also know that random letter strings are poorly perceived and poorly remembered. Thus, we can expect that a certain amount of the slowdown in typing random letter strings comes from the perceptual and memory processes. In the simulation, the effects of perception and memory limitations are not considered.

Sequences Involving Cross Hand Strokes Seem to Take Longer to Program than Those Involving Only Within Hand Strokes

Sternberg, Monsell, Knoll, and Wright (1978) used a discrete trial method in which subjects were first shown a string to be typed and then, on a signal, typed the string. They found that the preparation time (the time to the first stroke) increased linearly with the length of the string. They also found that the preparation time was longer for those strings that involved strokes on both hands than for those involving only a single hand. The typing time increased non-linearly with the length of the string, and strings involving two hands were faster to type than those involving one.

We made no attempt to deal with the preparation time phenomena. Our efforts have concentrated on the study of continuous typing and have offered no analysis of the discrete trial situation analyzed by Sternberg *et al.* In continuous typing, it is not clear what the analog to preparation time would be, except, possibly, the time between words. Of

course, our model does show that the duration of typing words involving both hands is less than for words involving a single hand.

For the moment, the preparation time results lie outside of the range of the model. It is possible that the planning time has to do with the set-up of the response schemata, and that the setting-up of the inhibition patterns across two hands is more complex and takes more time than the patterns within a hand. Within the context of continuous typing, such set-up times would not be noted. (Note too that the typing speeds observed by Sternberg et al. for their discrete trials were much faster than the normal typing speeds for these same subjects when typing in a continuous task.)

CONCLUSION

In this paper we have presented an analysis of many of the phenomena of typing and we have presented a computer simulation model that captures the appropriate spirit of the phenomena. The result has both positive and negative features.

Clearly the model does not yet offer a complete account of the typing process. The studies now underway in our laboratory point out some of the problems. There are striking individual differences among typists, so that two people who type at the same (high) average rate may differ substantially in the manner in which they type. In our model, we have made no provision for the substantial individual differences that can be observed in response patterns, no provision for several of the forms of errors that have been observed, nor no satisfactory account of the input and encoding processes involved in transcription typing.

We have presented strong arguments for the form of the mechanism that can account for both doubling and alternation errors. However, despite the success of the mechanism for doubling errors, we were not able to get essentially the same mechanism to work satisfactorily for the case of alternations. One possible reason for the failure may lie in the parsing and encoding stage. The problems we faced dealt with spaces when there were single letter words. It is possible that this case is handled by the stages of parsing and encoding, not by the typing mechanisms covered by the model.

There are some interesting asymmetries in typing behavior of our subject that we have not captured. Thus, the bigrams er and re are both typed rapidly, but er is typed considerably faster than is re. The model does not predict this. This result probably comes from the physical characteristics of the hand -- such as the differences in agility and length of the index and middle fingers -- perhaps coupled with the peculiar way in which the letters are arranged in diagonal rows on the keyboard and the resulting angles that the palms make with the keyboard. Our model of the hand is too simple, neglecting the differences in strengths and speeds of the fingers, and treating each finger independently of the others. In fact, some of the fingers are coupled together, with the tendons for some of the fingers bound within the same

tendon sheath. (There was a time around the turn of the century when virtuoso piano players had these sheaths surgically cut in the hopes of gaining more independent control of the fingers.)

In spite of these problems the model does capture many of the essential aspects of typing. We have shown how a very simple set of ideas can account for a wide range of phenomena when combined in a reasonable simulation model of the task. Note that despite the lack of an internal clock or metronome for timing, the model provides a reasonably good account of the timing patterns observed among skilled typists, including the prediction of negative correlations among successive keystrokes, a characteristic of metronome models. In similar fashion, there are no specific context dependencies built into the model and yet the time that it takes to strike keys depends upon the context in which they occur. We have no specific stored timing patterns for specific words, yet the model predicts that words have characteristic time profiles. We have no specific mechanism for transposition errors, yet our model generates the correct types of transposition errors. Moreover, the coordinative structure assumed within the model yields a qualitative emulation of the pattern of overlapping movements shown in a high speed film of a typist.

A number of conclusions can be drawn from our studies. First, the existence of doubling errors strongly implies the existence of a pure "type" representation of the keyboard schemata, with their arguments only loosely bound. In the model, the arguments are not attached to the schemata, but rather are picked up when they are needed by an evaluation of activation values. Second, the nature of the skill requires simultaneous, parallel control of the fingers and hands, and this requires some form of negotiation process to turn the potentially competitive movements into cooperative ones. The degrees of freedom problem is turned into a degrees of freedom virtue. Third, the model must incorporate the entire environment within which the typist operates, from the reading of the text, to the cognitive and motor control systems, to the shapes and mechanical characteristics of the hands, finger, and keyboard. Indeed, some of the limitations of the model may really result from limitations of how well we dealt with the environment surrounding the control processes. Perhaps the central conclusion to be drawn from our analysis of typing deals with the nature of skilled motor coordination. We propose that the motor control system carries out its computations relatively locally and in parallel. We presume that such a conclusion will be proven for all skills involving high speed performance.

References

- Beeching, W. A. Century of the typewriter. New York: St. Martin's Press, 1974.
- Gentner, D. R., Grudin, J., & Conway, E. Finger movements in transcription typing. San Diego, California: Center for Human Information Processing, University of California, San Diego. Technical Report 8001. May, 1980.
- Head, H. Aphasia and kindred disorders of speech. Vol. 1. Cambridge, England: Cambridge University Press, 1926. (pp. 448 ff.)
- Kent, R.D., & Minifie, F. D., Coarticulation in recent speech production models. Journal of Phonetics, 1977, 5, 115-133.
- Lashley, K. S. The problem of serial order in behavior. In L. A. Jeffress (Ed.), Cerebral mechanisms in behavior. New York: Wiley, 1951.
- Lessenberry, D.D. Analysis of errors. Syracuse, New York: L.C. Smith and Corona Typewriters, Inc., School Department, 1928. [Reprinted in Dvorak, A. Typewriting Behavior; Psychology Applied to Teaching and Learning Typewriting. New York: American Book Company, 1936.]
- Liberman, A. M., Cooper, F. S., Shankweiler, D. P. & Studdert-Kennedy, M., Perception of the speech code. In E. E. David, Jr. & P. B. Denes (Eds.), Human communication: A unified view. New York: McGraw-Hill, 1972.
- Norman, D. A. Categorization of action slips. Psychological Review, 1981, 88, 1-15.
- Saltzman, E. Levels of sensorimotor representation. Journal of Mathematical Psychology, 1979, 20, 91-163.
- Schmidt, R. A. The schema as a solution to some persistent problems in motor learning theory. In G. E. Stelmach (Ed.), Motor control: Issues and trends. New York: Academic Press, 1976.
- Shaffer, L. H. Latency mechanisms in transcription. In S. Kornblum (Ed.), Attention and performance IV. New York: Academic Press, 1973.
- Shaffer, L. H. Intention and performance. Psychological Review, 1976, 83, 375-393.
- Shaffer, L. H. Timing in the motor programming of typing. Quarterly Journal of Experimental Psychology, 1978, 30, 333-345.

Rumelhart & Norman
May 7, 1981

Simulating a skilled typist
41

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

CHIP Technical Report List

1. David M. Green and William J. McGill. On the equivalence of detection probabilities and well known statistical quantities. October, 1969.
2. Donald A. Norman. Comments on the information structure of memory. October, 1969.
3. Norman H. Anderson. Functional measurement and psychophysical judgment. October, 1969.
4. James C. Shanteau. An additive decision-making model for sequential estimation and inference judgments. October, 1969.
5. Norman H. Anderson. Averaging model applied to the size-weight illusion. October, 1969.
6. Norman H. Anderson and James C. Shanteau. Information integration in risky decision making. November, 1969.
7. George Mandler, Richard H. Meltzer, Zena Pearlstone. The structure of recognition. Effects of list tags and of acoustic and semantic confusion. November, 1969.
8. Dominic W. Massaro. Perceptual processes and forgetting in memory tasks. January, 1970.
9. Daniel Graboi. Searching for targets: The effects of specific practice. February, 1970.
10. James H. Patterson and David M. Green. Discrimination of transient signals having identical energy spectra. February, 1970.
11. Donald A. Norman. Remembrance of things past. June, 1970.
12. Norman H. Anderson. Integration theory and attitude change. August, 1970.
13. A.D. Baddeley and J.R. Ecob. Reaction time and short-term memory: A trace strength alternative to the high-speed exhaustive scanning hypothesis. November, 1970.
14. A.D. Baddeley. Retrieval rules and semantic coding in short-term memory. December, 1970.
15. Roy D. Patterson. Residue pitch as a function of the number and relative phase of the component sinusoids. March, 1971.
16. George Mandler and Marilyn A. Borges. Effects of list differentiation, category membership and prior recall on recognition. May, 1971.
17. David E. Rumelhart, Peter H. Lindsay, and Donald A. Norman. A process model for long-term memory. May, 1971.
18. David E. Rumelhart and Adele A. Abrahamson. Toward a theory of analogical reasoning. July, 1971.
19. Martin F. Kaplan. How response dispositions integrate with stimulus information. August, 1971.
20. Martin F. Kaplan and Norman H. Anderson. Comparison of information integration and reinforcement models for interpersonal attraction. August, 1971.
21. David M. Green and R. Duncan Luce. Speed-accuracy tradeoff in auditory detection. September, 1971.
22. David E. Rumelhart. A multicomponent theory of confusion among briefly exposed alphabetic characters. November, 1971.
23. Norman H. Anderson and Arthur J. Farkas. New light on order effects in attitude change. March, 1972.
24. Norman H. Anderson. Information integration theory: A brief survey. April, 1972.
25. Donald A. Norman. Memory, knowledge, and the answering of questions. May, 1972.
26. David J. Weiss. Averaging: An empirical validity criterion for magnitude estimation.
Norman H. Anderson. Cross-task validation of functional measurement. June, 1972.
27. David E. Rumelhart and Patricia Siple. The process of recognizing tachistoscopically presented words. August, 1972.
28. Ebbe B. Ebbesen and Richard Bowers. The effects of proportion of risky to conservative arguments in a group discussion on risky shift. September, 1972.
29. Ebbe B. Ebbesen and Michael Haney. Flirting with death: Variables affecting risk taking on our nation's highways. September, 1972.
30. Norman H. Anderson. Algebraic models in perception. November, 1972.
31. Norman H. Anderson. Cognitive algebra: Information integration applied to social attribution. December, 1972.
32. Jean M. Mandler and Nancy L. Stein. Recall recognition of pictures by children as a function of organization and of distractor similarity. January, 1973.
33. David E. Rumelhart and Donald A. Norman. Active semantic networks as a model of human memory.
Marc Eisenstadt and Yaakov Kareev. Towards a model of human game playing. June, 1973.
34. George Mandler. Memory storage and retrieval: Some limits on the reach of attention and consciousness. July, 1973.
35. Kent L. Norman. A method of maximum likelihood estimation for stimulus integration theory. August, 1973.
36. Yaakov Kareev. A model of human game playing. August, 1973.
37. Donald A. Norman. Cognitive organization and learning. August, 1973.
38. The Center for Human Information Processing: A Five Year Report — 1968-73.

39. Larry D. Rosen and J. Edward Russo. Binary processing in multi-alternative choice. October, 1973.
40. Samuel Himmelfarb and Norman H. Anderson. Integration theory analysis of opinion attribution. December, 1973.
41. George Mandler. Consciousness: Respectable, useful, and probably necessary. March, 1974.
42. Norman H. Anderson. The problem of change-of-meaning. June, 1974.
43. Norman H. Anderson. Methods for studying information integration. June, 1974.
44. Norman H. Anderson. Basic experiments in person perception. June, 1974.
45. Norman H. Anderson. Algebraic models for information integration. June, 1974.
46. Ebbe B. Ebbesen and Vladimir J. Konečni. Cognitive algebra in legal decision making. September, 1974.
47. Norman H. Anderson. Equity judgments as information integration.
Arthur J. Farkas and Norman H. Anderson. Input summation and equity summation in multi-cue equity judgments. December, 1974.
48. George Mandler and Arthur Graesser II. Dimensional analysis and the locus of organization. January, 1975.
49. James L. McClelland. Preliminary letter identification in the perception of words and nonwords. April, 1975.
50. Donald A. Norman and Daniel G. Bobrow. On the role of active memory processes in perception and cognition. May, 1975.
51. J. Edward Russo. The value of unit price information. An information processing analysis of point-of-purchase decisions. June, 1975.
52. Elissa L. Newport. Motherese: The speech of mothers to young children. August, 1975.
53. Norman H. Anderson and Cheryl C. Graesser. An information integration analysis of attitude change in group discussion. September, 1975.
54. Lynn A. Cooper. Demonstration of a mental analog of an external rotation.
Lynn A. Cooper and Peter Podgorny. Mental transformations and visual comparison processes: Effects of complexity and similarity. October, 1975.
55. David E. Rumelhart and Andrew Ortony. The representation of knowledge in memory. January, 1976.
56. David E. Rumelhart. Toward an interactive model of reading. March, 1976.
57. Jean M. Mandler, Nancy S. Johnson, and Marsha DeForest. A structural analysis of stories and their recall: From "Once upon a time" to "Happily ever after". March, 1976.
58. David E. Rumelhart. Understanding and summarizing brief stories. April, 1976.
59. Lynn A. Cooper and Roger N. Shepard. Transformations on representations of objects in space. April, 1976.
60. George Mandler. Some attempts to study the rotation and reversal of integrated motor patterns. May, 1976.
61. Norman H. Anderson. Problems in using analysis of variance in balance theory. June, 1976.
62. Norman H. Anderson. Social perception and cognition. July, 1976.
63. David E. Rumelhart and Donald A. Norman. Accretion, tuning and restructuring: Three modes of learning. August, 1976.
64. George Mandler. Memory research reconsidered: A critical view of traditional methods and distinctions. September, 1976.
65. Norman H. Anderson and Michael D. Klitzner. Measurement of motivation.
Michael D. Klitzner and Norman H. Anderson. Motivation x expectancy x value: A functional measurement approach. November, 1976.
66. Vladimir J. Konečni. Some social, emotional, and cognitive determinants of aesthetic preference for melodies in complexity. December, 1976.
67. Hugh Mehan, Courtney B. Cazden, LaDonna Coles, Sue Fisher, Nick Maroules. The social organization of classroom lessons. December, 1976.
- 67a. Hugh Mehan, Courtney B. Cazden, LaDonna Coles, Sue Fisher, Nick Maroules. Appendices to the social organization of classroom lessons. December, 1976.
68. Norman H. Anderson. Integration theory applied to cognitive responses and attitudes. December, 1976.
69. Norman H. Anderson and Diane O. Cuneo. The height + width rule in children's judgments of quantity. June, 1977.
Norman H. Anderson and Clifford H. Butzin. Children's judgments of equity. June, 1977.
70. Donald R. Gentner and Donald A. Norman. The FLOW tutor: Schemas for tutoring. June, 1977.
71. George Mandler. Organization and repetition: An extension of organizational principles with special reference to rote learning. May, 1977.
72. Manuel Leon. Coordination of intent and consequence information in children's moral judgements. August, 1977.
73. Ted Supalla and Elissa L. Newport. How many seats in a chair? The derivation of nouns and verbs in American Sign Language. November, 1977.
74. Donald A. Norman and Daniel G. Bobrow. Descriptions: A basis for memory acquisition and retrieval. November, 1977.

75. Michael D. Williams. The process of retrieval from very long term memory. September, 1978.
76. Jean M. Mandler. Categorical and schematic organization in memory. October, 1978.
77. James L. McClelland. On time relations of mental processes: A framework for analyzing processes in cascade. October, 1978.
78. Jean M. Mandler and Marsha DeForest. Developmental invariance in story recall. November, 1978.
Jean M. Mandler, Sylvia Scribner, Michael Cole, and Marsha DeForest. Cross-cultural invariance in story recall. November, 1978.
79. David E. Rumelhart. Schemata: The building blocks of cognition. December, 1978.
80. Nancy S. Johnson and Jean M. Mandler. A tale of two structures: Underlying and surface forms in stories. January, 1979.
81. David E. Rumelhart. Analogical processes and procedural representations. February, 1979.
82. Ross A. Bott. A study of complex learning: Theory and methodologies. March, 1979.
83. Laboratory of Comparative Human Cognition. Toward a unified approach to problems of culture and cognition. May, 1979.
84. George Mandler and Lawrence W. Barsalou. Steady state memory: What does the one-shot experiment assess? May, 1979.
85. Norman H. Anderson. Introduction to cognitive algebra. June, 1979.
86. Edited by Michael Cole, Edwin Hutchins, James Levin and Naomi Miyake. Naturalistic problem solving and microcomputers. Report of a Conference. June, 1979.
87. Donald A. Norman. Twelve issues for cognitive science. October, 1979.
88. Donald A. Norman. Slips of the mind and an outline for a theory of action. November, 1979.
89. The Center for Human Information Processing: A Description and a Five-Year Report (1974-1979). November, 1979.
90. Michael Cole and Peg Griffin. Cultural amplifiers reconsidered. December, 1979.
91. James L. McClelland and David E. Rumelhart. An interactive activation model of the effect of context in perception. Part I. April, 1980.
92. James L. McClelland and J.K. O'Regan. The role of expectations in the use of peripheral visual information in reading. February, 1980.
93. Edwin Hutchins. Conceptual structures of Caroline Island navigation. May, 1980.
94. Friedrich Wilkening and Norman H. Anderson. Comparison of two rule assessment methodologies for studying cognitive development. June, 1980.
95. David E. Rumelhart and James L. McClelland. An interactive activation model of the effect of context in perception. Part II. August, 1980.
96. Jean M. Mandler. Structural invariants in development. September, 1980.
97. David E. Rumelhart and Donald A. Norman. Analogical processes in learning. October, 1980.
98. James A. Levin and Yaakov Kareev. Personal computers and education: The challenge to schools. November, 1980.
99. Donald A. Norman and Tim Shallice. Attention to action: Willed and automatic control of behavior. December, 1980.
100. David E. Rumelhart. Understanding understanding. January, 1981.
101. George Mandler. The structure of value: Accounting for taste. May, 1981.
102. David E. Rumelhart and Donald A. Norman. Simulating a skilled typist: A study of skilled cognitive-motor performance. May, 1981.

Note: Requests for CHIP reports should be addressed to the author. Reports are also available through the Library Loan Service of the University of California, San Diego, La Jolla, California 92093.