COO-614-61

ICR QUARTERLY REPORT

Number 16
February 1, 1968

MAY 23 1968

MAY 23 1968

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

THE UNIVERSITY OF CHICAGO

THE INSTITUTE FOR COMPUTER RESEARCH

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency Thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

The Institute for Computer Research
The University of Chicago

MASTER

ICR Quarterly Report

Number 16
February 1, 1968

LEGAL NOTICE

This report was prepared as an account of Government sponsored work. Neither the United States, nor the Commission, nor any person acting on behalf of the Commission:

A. Makes any warranty or representation, expressed or implied, with respect to the accuracy, completeness, or usefulness of the information contained in this report, or that the use of any information, apparatus, method, or process disclosed in this report may not infringe privately owned rights; or

B. Assumes any liabilities with respect to the use of, or for damages resulting from the use of any information, apparatus, method, or process disclosed in this report.

As used in the above, "person acting on behalf of the Commission" includes any employee or contractor of the Commission, or employee of such contractor, to the extent that such employee or contractor of the Commission, or employee of such contractor prepares, disseminates, or provides access to, any information pursuant to his employment or contract with the Commission, or his employment with such contractor.

Prepared under the auspices of AEC Contract No. AT(11-1)-614.

Victor H. Yngve, Principal Investigator

PATENT REVIEW APPROVES RELEASE. PROCEDURES
GOVERNING PATENT REVIEW AND RELEASE ARE ON
FILE IN RECEIVING SECTION.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Staff of the Institute for Computer Research

Naomi Alton

Robert Ashenhurst

Jurgen Bounin

Dolores Brennan

William Brown

Rachel Chakrin

Minja Choe

Elaine Chow

Charles Fischer

Etta Franklin

William Georgas

Robert Graves

Richard Herwitt

Leonardo Herzenberg

Thea Hodge

Jane Jodeit

Paul Kosinski

Hirondo Kuki

George Kreisheimer

Vincent Kruskal

Glenn Manacher

Richard Miller

Lily Monheit

John Munn

Daniel Norman

Alex Orden

Walter Petryshyn

Charles Robinson

Clemens Roothaan

Lyle Settle

John Shepherd

Cheng-kiat Teng

Serena Torres

Nancy Trombetta

Cherie Weil

Michael Williams

Augustine Witt

Victor Yngve

TABLE OF CONTENTS

Section I

- A. Project Support Activity
- B. The Man-Machine Project (C.C.J. Roothaan)
- C. PDP-8 I/O Devices (J. Bounin)
- D. Phase Counter and Interface to the PDP-8 for an Interferometer
(J. Bounin & L. Herzenberg)

Section II

- A. The Linguistic Information Analysis and Retrieval System (S. Herbst)

Section III

- A. Essential Features of Purposive Movements (P.H. Greene)
- B. Cybernetic Problems of Sensorimotor Structure: Introductory
Remarks and Survey of a Study (P.H. Greene)
- C. Seeking Mathematical Models for Skilled Actions (P.H. Greene)
- D. An Aspect of Robot Control and Nervous Control of Skilled Movements:
Coordination of Two Effectors and Transfer of Adaptation

PROJECT SUPPORT ACTIVITY

On January 3 the Maniac III system ground to a halt due to an "impossible" overflow in the accounting file. This file is twenty-two tracks long with two words used per job giving a total of $\frac{320}{2} \times 22 = 3520$ jobs. This file is purged once every three months and the statistics are reported in the Quarterly Report. During the period from October 24 through January 3 (a period that includes the Christmas holidays), enough jobs were run on Maniac III to fill the file.

There were approximately 1300 student jobs submitted under the Inf Sci 310 identification, distributed among 33 students, or approximately 40 jobs per student or 10 per student assignment. The total time use by these students was 52 hours and 20 minutes of time. Other (unsupervised) student jobs consumed 84 hours for running 925 jobs (5.4 minutes/job). The increase in systems time was due to the development and introduction of a conversational link between Maniac III and the PDP-8.

MANIAC III USE

October 14, 1967 - January 3, 1968

Systems work	72 hours
Engineering computer time	50 hours
Computer testing of circuit boards	50 hours
Astronomical calculations	31 hours
Miscellaneous (including demonstrations)	<u>20 hours</u>
Total	223 hours

In the succeeding two months, Maniac III ran a total of 2196 jobs in 204.5 hours. The breakdown by use category follows.

MANIAC III USE

January 3, 1968 - February 26, 1968

Students	106 hours	1450 jobs
Systems staff	19 hours	300 jobs
Engineering testing	30 hours	75 jobs
Stockroom bookkeeping	2 hours	16 jobs
Astronomical calculations	21.5 hours	151 jobs
Circuit board testing program	12 hours	84 jobs
Miscellaneous	14 hours	120 jobs

THE MAN-MACHINE PROJECT*

C.C.J. Roothaan

Progress Report No. 12 (Final Report)

(for the period 1 October 1967 through 31 December 1967)

System/360 Mathematical Subroutines

Mr. Kuki has finished revising the System/360 Mathematical Function Library. The principal purpose of the revision was to adjust the library to pending floating point hardware changes. A secondary purpose was to accomplish improvements and to prevent loss of quality. These purposes have been achieved.

Improvements: 1. Intermediate underflow is completely avoided.

Terminal underflow is via floating point exception control.

2. The two complex square root subroutines and the two complex logarithm subroutines are now free of the possibility of both intermediate overflow and loss of accuracy due to partial underflow.

3. The use of so-called mini-max techniques in achieving "best" approximations has resulted in great improvements for the TANH, ERF, and GAMMA subroutines--and in modest improvements

for other subroutines. For example, the revised

*Work on this project, supported by a grant from the International Business Machines Corporation with C.C.J. Roothaan as principal investigator, has been carried on under the joint auspices of the Institute and the Computation Center.

version of the long form GAMMA subroutine is 28% faster than its predecessor for arguments between 1 and 2. It is also somewhat shorter.

4. Several refinements have been introduced. For example, the maximum relative errors have been cut in half for both of the square root subroutines by the introduction of virtual rounding. The long form exponential routine now uses an extra-precision reduction which assures virtual full precision answers regardless of argument size. CABS, LOG, ASCN, and SCNH all use special techniques to protect the full visible precision.

Members of the SHARE Numerical Analysis Project contributed advice leading to several improvements. In particular, the above mentioned refinements owe much to ingenious techniques suggested by Mr. William J. Cody of the Argonne National Laboratory and by Mr. Richard Turner of the Lewis Research Center.

Note: Some of the costs (such as machine time, keypunching, and communications) for revising the library were borne by IBM under their purchase order No. LV-B3445.

CHOP (Chicago Output Package)

7094 CHOP: Mr. Settle has written a Fortran II adapter, CHOPA, which provides complete linkage to CHOP for titling,

for writing supermatrices, for inserting BCD information into CHOP's page buffer(s), and for retrieving information already there. In addition, CHOPA intercepts each line generated by the WRITE OUTPUT TAPE 6 statement, treats it as a line segment, and puts it into CHOP's buffer. Where a given line segment is placed is determined solely by the CRANKERs which the user controls through simple calls. By interspersing WOT6 statements with CRANKER calls, an output stream can be easily rearranged into almost any conceivable format. (The printed format need bear no resemblance to the order in which WOT6 statements were executed.) The CRANKERs also determine whether lines are written horizontally or vertically. If the user does not use CRANKERs (and other options) his output will resemble standard FORTRAN output with the following exceptions:

- * Lines 1, 2, 3, 5, 64, 65, 66 (physical page) will be blank.
- * Line 4 will contain the date and page number.
- * Fortran output will be written on lines 6-63.
(A new page will begin on line 6.)
- * Only Fortran spacing characters 1, +, 0, -, and blank will be honored.

The user may provide OVERLAY subroutines which will be automatically called to add annotation to each page just before the page is sent to tape.

360 CHOP: Mr. Sevcik has changed CHOP to run under OS/360 and has made preliminary designs of macros for generating CHOP calling sequences. In addition to making it easier to call CHOP, these macros will detect several calling sequence errors.

CHAOS (CHicago Asynchronous Operations Scheduler) Procedure Library

Work is in progress on a disk-resident procedure library for CHAOS. The library is made possible by the CHAOS relocatable assembler and loader. The benefits of such a library include a substantial reduction in the size of "resident" CHAOS, and a greatly increased potential for the addition of new features.

CHAP (CHaos Assembly Program)

Miss Monheit has finished the loader. Both it and the assembler are operational.

CHAOS/7040

Mr. Kosinski has implemented a new memory allocator which allows memory blocks of arbitrary size to be obtained and released dynamically in the multiprogramming environment. A very fast self-searching mechanism is used to reduce the overhead for this service and recombination of the fragmented blocks is done only when needed.

This replaces the old fixed length memory block allocation scheme and plans are being made to take advantage of the flexibility provided by the availability of variable length blocks.

Mr. Kruskal has provided several typed-in commands so that operator communication typewriters may be "attached" and "detached" by the operators. If an operator typewriter is "detached", messages which normally would be sent to that typewriter are routed to an alternate typewriter according to a preference list (in case that alternate typewriter is also "detached"). Input is always accepted by an operator typewriter, however, so that a typewriter may be used to "attach" itself. Since there are three typewriters available to the operators in the Computation Center, the system is not crippled by one, or even two typewriters going 'down'.

Messrs. Kosinski and Kruskal have done the design work with regard to completely revising the disc input/output support. A new type of file, called a direct access file, is contemplated together with a revision of the sequential access file implementation. A direct access file would appear to be much like a core memory, that is, an ordered sequence of words from which any subsequence would be retrievable by address and length. This, coupled with a "lookbehind" buffering scheme, would take some of the memory space pressure off the CHAP assembler (which could use direct access files for its symbol tables).

Also being contemplated is a revision of the file cataloging scheme to 1) allow for the clean incorporation of direct access files; and 2) remove the necessity for keeping core resident so many tables concerning tasks or jobs.

PDP-8 I/O DEVICES

by J. Bounin

Several more devices have been recently added to the PDP-8: a memory multiplexer (DMO1, from DEC); an Interferometer Interface; a second tape transport (TU55, from DEC). The complexity of the current I/O arrangements on the PDP-8 seems to warrant a condensed write-up.

The PDP-8 itself, together with the tape-controller (TCO1); the two tape transports (TU55), and the memory multiplexer (DMO1) are mounted on a small truck. Thus this part of the equipment can be easily moved to other locations on or off campus for experimental runs.

All the other I/O devices are housed in standard relay racks, which provides at least some mobility. The need for mobility might arise in the case of the Analog-Digital Counter 139E/138E and it is anticipated for the interferometer interface.

The PDP-8 I/O bus, together with the I/O devices, is shown schematically in Fig. 1. All devices are connected logically in parallel, physically in series, onto this bus. The physical layout is illustrated in Fig. 2; one set of I/O cables enters every relay rack, one exits from it. Devices within one relay rack are internally jumpered on the I/O-bus. In the following notes a short description of the I/O devices is given and the set of I/O instructions that reference them described. More detailed information may be found as referenced.

Teletype model 33 ASR Keyboard/Reader

Function: 8-bit characters are received serially from the teletype into the teletype input buffer TTI, following which the keyboard flag is set.

Address: No. 03.

Interrupt: An interrupt is sent to the PDP-8 if the keyboard flag is set.

Instructions:

<u>name</u>	<u>mnemonic</u>	<u>octal code</u>	<u>operation</u>
Skip on Keyboard Flag	KSF	6031	If the keyboard flag is set, then the next instruction will be skipped.
Clear Keyboard Flag	KCC	6032	Clear AC and keyboard flag
Read Keyboard Buffer	KRS	6034	TTI is loaded into AC4...AC11.

Notes: a) These instructions can be combined.

b) For detailed information see DEC-F85, page 57.

Teletype model 33 ASR Teleprinter/Punch

Function: Receive an 8-bit character from AC4...AC11 into teletype output buffer TTO, and sent it serially to the teletype. The teleprinter flag is set whenever TTO is ready to receive the next character from the PDP8.

Address: No. 04.

Interrupt: An interrupt is sent to the PDP8 if the teleprinter flag is set.

Instructions:

<u>name</u>	<u>mnemonic</u>	<u>octal code</u>	<u>operation</u>
Skip on Teleprinter Flag	TSF	6041	If the teleprinter flag is set, then the next instruction is skipped.
Clear Teleprinter Flag	TCF	6042	The teleprinter flag is cleared to 0.
Load Teleprinter Buffer and Print	TPC	6044	TTO is loaded from AC4...AC11, and the character is sent serially to the teletype.

Notes: a) Instructions can be combined.

b) Detailed information may be found in DEC-F85, page 59.

Multiplexer DMO1

The multiplexer is not an active I/O device, so has no I/O address. Some of the timing signals on the I/O bus, however, are used for the priority selection of the data-break and during a data-break data is sent and received via the I/O bus.

DEC Tape Control Type TC 01DEC Tape Transport Type TU 55

Function: Every tape transport holds a 260 ft reel of magnetic tape. The transport is address-selected via a thumb wheel and a register in the control. The tape control can handle up to 8 tape transports, one at a time, and controls moving, recoding, reading and writing.

Capacity: 350 $\pm$ 55 bits/inch

849 036 lines/reel 3 bits/line

Standard format: 129 word/block

1470 blocks/tape = 45 core images/tape

Speed: 97 $\pm$ 14 in/sec

15000 characters/sec, 6 bits/character

Address: No. 76 and No. 77

Interrupt: An interrupt request is sent to the PDP-8 whenever the DECTape Control Flag is on, and either the error flag or the DECTape flag is set.

Instructions:

<u>name</u>	<u>mnemonic</u>	<u>octal code</u>	<u>operation</u>
Read Status Register A	DTRA	6761	Status register A is or'ed into AC.
Clear Status Register A	DCTA	6762	Clear status register A. Leave DECTape Control Flag and error flag undisturbed.
Load Status Register A	DTXA	6764	Exclusive -or AC ₀ ...AC ₉ into status register A. If AC10=1, clear error flags. If AC11=1, clear DECTape flag.

Skip on Flags	DTSF	6771	If error flags or DECtape is set, then skip.
Read Status Register B	DTRB	6772	Status register B is or'ed into AC.
Load Status Register B	DTXB	6774	AC6...AC8 is or'ed into the memory field position of Status register B.

Notes: a) Instructions within one address group can be combined.

b) For details see DEC-F85, page 108.

Maniac III/PDP-8 Interface

Function: The interface receives two 8-bit bytes from Maniac III, buffers them, and requests transfer of buffer content to the PDP-8. The PDP-8 can transfer a 12-bit word to Maniac III through the interface.

Address: No. 30

Interrupt: An interrupt will be caused if Maniac III sets the C-bit during the transmission of the second byte (F 33).

Instructions:

<u>name</u>	<u>menmonic</u>	<u>octal code</u>	<u>operation</u>
Transmit word to MIII	WXM3	6302	If the interface has received READY from Maniac III, then transmit AC to Maniac III, and skip.
Transmit last word to MIII	EOXM3	6306	Same as WXM3, but also send EOX.
Send MIII interrupt	INT M3	6304	Set SPARK toggle and PDP-8 toggle in Maniac III.
Read MIII data	RM3	6307	If two bytes have been received from Maniac III load buffer into AC and skip.
Test and load control data	TM3	6305	If INTERRUPT or EOF then load control data. If <u>INTERRUPT</u> then <u>reset</u> buffer. If INTERRUPT and EOF then skip.

Notes: a) These I/O instructions cannot be combined.

b) For detailed information see Quarterly Report No. 11, page F-C1.

Drawings are filed under DRWG. No. 388.

Analog - Digital Converter, Type 138EMultiplexer Control, Type 139E

Function: One of 64 analog channels is selected through the selector register, and the input level converted into a binary number of 6 to 12 bits in the buffer. The buffer can be read into the AC. The AD converter flag is set at the end of the conversion.

Addresses: No 53 and No. 54.

Interrupt: If the AD converter flag is set, an interrupt request is sent to the PDP-8.

Instructions:

<u>name</u>	<u>mnemonic</u>	<u>octal code</u>	<u>operation</u>
Skip on ADC Flag	ADSF	6531	If the ADC flag is set, skip the next instruction.
Convert	ADCV	6532	The ADC flag is cleared, and conversion is started.
Read ADC Buffer	ADRB	6534	The ADC flag is cleared, and the buffer loaded into the ACM
Clear Multiplexer Channel	ADCC	6541	The channel selector register is cleared.
Set Multiplexer Channel	ADSC	6542	The channel selector register is loaded from AC6...AC11.
Increment Multiplexer Channel	ADSC	6544	Increment content of channel selector register by one.

Notes: a). Instructions within one address group can be combined.

b) For detailed information see DEC-F85, page 68.

Memory Monitor

Function: The memory monitor looks at the address buses of the 16K Ampex memory, and then uses this address in truncated form to address the PDP-8 memory via data-break to increment the content of that location by one, thus creating a histogram of the Ampex memory activity.

Address: No. 33

Interrupt: An interrupt is caused a) if a cell in the histogram overflows,
b) if the data-break facility is not used.

Instructions:

<u>name</u>	<u>mnemonic</u>	<u>octal code</u>	<u>operation</u>
Monitor Skip	MCS	6331	Skip if no interrupt request; Clear AC.
Load Address	MLA	6332	Or the monitored address into AC; reset the OVERFLOW flip-flop flips that caused the interrupt.
Clear Monitor	MSR	6334	Set the RUN flip-flop, reset the OVERFLOW flip-flop.

Notes: a) These instructions can be combined.

b) The drawings for the memory monitor are filed under ICR
DRWG. No. 400.

c) A write-up on the memory monitor is planned for a future Quarterly
Report.

IC-Card-Tester

Function: A 96 bit buffer can be loaded from the PDP-8; the buffer provides input signals to the IC-Card under test. Outputs from the IC-card can be read via 48 lines, and transmitted to the PDP-8.

Addresses: No. 31 and 32.

Interrupts: No interrupts can occur.

Instructions:

<u>name</u>	<u>mnemonic</u>	<u>octal code</u>	<u>operation</u>
Count	CNT	6311	Increment address counter by one.
Load Output Buffer and Count	LBU	6312	Same as CNT, and load AC into the buffer section selected by the not yet incremented address counter.
Reset Counter and Buffer	RES	6314	Reset all 96 flip-flops in the output register. Reset address counter to 0.
Trigger	TRIG	6320	Provide a scope trigger pulse.
Load AD from IC-card	LAC	6332	The data on the 12 lines selected by the address counter is or'ed into the AC. Then the address is incremented.
Clear and Load AC	CLP	6323	Same as LAC, but clear AC first.
Reset Counter	RCO	6324	Reset the address counter to 0.

Notes: a) Instructions can be combined within one address group.

b) Drawings of the IC card tester are filed under ICR DRWG. No. 389.

Interferometer Interface

Function: This interface samples the signal from four photomultipliers, transmits the events to the PDP-8, derives count pulses for its 8 bit up/down counter from them, and detects error combinations in the input signals.

Addresses: No. 34 and No. 35.

Interrupt: An interrupt request to the PDP-8 can be caused by:

- a) any one of the 4 input signals;
- b) a strobe pulse from the 1 kc clock;
- c) counter underflow or overflow;
- d) any error.

Instructions:

<u>name</u>	<u>mnemonic</u>	<u>octal code</u>	<u>operation</u>
Load Counter and Errors into AC	ILC	6347	If an interrupt request is up because of b), c), or d), then the counter and error flip-flops are loaded into the AC. After an overflow the counter is preset to 200_8 .
Load Input Buffer	ILB	6357	If any buffer toggle is set, the buffer is or'ed into the AC, and the buffer cleared; otherwise, skip next instruction.

Notes: A detailed description of the interferometer interface can be found in this Quarterly Report, Section I-D. Drawings are kept under ICR DRWG. No. 402.

Fig.1

ALL CABLES: 92 Ω COAX.

9 COAX PER CABLE

6 CABLES IN I/O BUS

4 CABLES IN DATA BREAK BUS

Fig. 2

PHASE COUNTER AND INTERFACE TO THE PDP-8 FOR AN INTERFEROMETER

by J. Bounin and L. Herzenberg

Introduction

In the last Quarterly Report* R. Miller described two projects for investigating astronomical seeing and fringe visibility. Both these projects would use the same interferometer, essentially differing only in the kind of data collected during the experiment and its evaluation.

The rate of the data collection will be such as to require an automatic way of dealing with it. It was therefore decided to build some special equipment with counters, decoders and error-logic which would handle the input signals, and to interface this unit to a small computer, the PDP-8, for further sampling and storage of the (slightly) preprocessed data.

In the following article we describe the overall design and some particular details of the phase counter.

*

1) R. H. Miller, "Signal-to-noise Ratio in Fringe Visibility Measurements with a Michelson Stellar Interferometer," Quarterly Report No. 15, Section II A.

2) R. H. Miller, "A Scheme for Studying Astronomical 'Seeing'," Quarterly Report No. 15, Section III B.

Block Diagram

The four pulses coming from the photomultipliers of the interferometer, labeled $D_1 \dots D_4$, (see Fig. 2 of 2) are pulse-shaped and then integrated, as shown in Fig. 1.

One of the four flip-flops of the buffer register will be set, and a request for an I/O transfer of the buffer content to the PDP-8 issued.

The integrators feed two two-input level comparators, whose outputs are monitored by a transition detector. Depending on the state of the comparators an input-level change will cause an up-count pulse, a down-count pulse or an error pulse to be produced.

The count pulses control an 8-bit up/down counter capable of detecting over or underflow and thereby setting an overflow flip-flop.

Errors are caused through counter dead time and limited time resolution of the transition detector, and are recorded in error flip-flops.

Requests to the PDP-8 for readout of the error, overflow and counter information are triggered by either a 1 khz clock in the interface, or by setting of the overflow or error flip-flops.

Integrators

The signals coming from the photomultipliers, after passing through amplifiers and discriminators, are .5 μ sec long standard pulses. These pulses are fed into integrators producing outputs that correspond to the pulse repetition rate, and hence to the light variations of interest. The discharge RC time constant is 1 msec. The input pulses are used to gate a constant current source into the integrating capacitor. The charging rate is 50mv/ μ sec., or 25mv/pulse. The maximum voltage that can be reached is 5 volts. The discharge resistor is actually a potentiometer, so the voltage supplied to the comparator may be varied.

Transition Detector

The four outputs from the integrators are fed in pairs to two comparators which produce output signals $a(\bar{a})$ and $b(\bar{b})$, respectively:

$$D_1 > D_2 \rightarrow a,$$

$$D_3 > D_4 \rightarrow b.$$

Consider the comparator outputs as state of the machine. If sufficient input changes occur, the comparator will change state, and produce an output from the transition detector. State graph and flow table for this Mealy-type machine are shown in Fig. 2.

The transition detector is realized in the following manner (see Fig. 3): Changes of the inputs are detected by pulse amplifiers PA 1. This causes a strobe of the present state (which is remembered just long enough!) into delays Δ , and triggers timing delay ΔT . At the end of ΔT PA2 will produce an output-strobe which combines the new state with the memorized old state (from the delay Δ), and depending on the pair of states produce either an up-count pulse, a down-count pulse or an error pulse according to the state graph.

Up/Down Counter

The up/down pulses from the transition detector are fed to a binary up/down counter. This 8-bit counter was preset to 80_{16} , the center of the counting range. Whenever underflow or overflow is detected, the overflow toggle is set, and the error logic will request readout by the PDP-8, which in turn will preset the counter again.

Error Logic

Because of the finite time resolution of the logic elements dead times are encountered, and proper logic function during these periods cannot be guaranteed. However, the error-logic will provide an indication of any such high-frequency transitions which cannot properly be dealt with. First, slight feedback around the comparators introduces about 100mv of hysteresis, which assures that the com-

parators will not respond to individual pulses, thus their output cannot change more than twice in any $4\ \mu\text{sec}$ interval.

Next, the transition detector will give an error indication if an input changes state twice within $5\ \mu\text{sec}$ (2 MC rate) or both inputs (comparators) change state within that period.

The up/down counter itself is protected against input-rate exceeding 1.25 Mc, and any pulse sequence exceeding this rate will cause an error indication.

If a second error occurs before the first one is dealt with, an error-pile-up flip-flop will be set.

I/O Logic

The 1 kc clock of the interface or any error will cause an interrupt to the PDP-8, which will then request readout of the counter and the error flip-flops. The bit patterns of the data-word transmitted to the PDP-8 accumulator are shown in Fig. 4a. Furthermore, every input signal $D_1 \dots D_4$ will set a buffer flip-flop and cause a PDP-8 interrupt. The bit-pattern for the buffer read-out is shown in Fig. 4b.

Two PDP-8 I/O instructions "talk" to the interferometer interface, and serve to test interrupt requests, load data and reset the interface.

a) READ COUNTER AND ERRORS

mnemonic: ILC

octal code: 6347

address: No. 34

If the interrupt request to the PDP-8 was up because of the 1 kc clock, an error or an overflow, this instruction will select the interface (No. 34), load (or'ed) the counter and the error flip-flops into the AC, and then reset the error flip-flops and preset the counter after an overflow. If no interrupt was requested execution of this instruction will cause a skip, i. e. the next instruction is not executed.

b) LOAD INPUT BUFFER

mnemonic: ILB

octal code: 6357

address: No. 35

If the interrupt request was on because one or more of the input flip-flops were set (from input pulses $D_1 \dots D_4$), then ILB will load the buffer content into AC8...AC11, and then clear the buffer. If none of the buffer flip-flops was set, then the next instruction will be skipped.

The smallest program loop that will handle data from the interferometer interface will require 27μ sec for execution, and might look as follows:

ILB		
DCA	I	X10
ILC		
DCA	I	X11
ION		
ISZ		COUNT
JMP		*

with an additional JMP at location 0001 to instruction ILB. The auto-index registers X10 and X11 will contain the addresses for input data storage. COUNT will determine how often the loop will be repeated before exiting.

Remark.

The interferometer interface was built with DEC FLIP-CHIP modules. It is housed in a standard relay rack, together with amplifiers and power supplies for the photomultipliers.

Fig. 1

STATE — GRAPH

FLOW TABLE

	D1	D2	D3	D4	D1	D2	D3	D4
ab	ab	ab	ab	$a\bar{b}$	—	Down	—	UP
$a\bar{b}$	$\bar{a}\bar{b}$	$a\bar{b}$	ab	$a\bar{b}$	—	UP	Down	—
$\bar{a}b$	ab	$\bar{a}b$	$\bar{a}b$	$\bar{a}\bar{b}$	UP	—	—	Down
$\bar{a}\bar{b}$	$a\bar{b}$	$\bar{a}\bar{b}$	$\bar{a}b$	$\bar{a}\bar{b}$	Down	—	UP	—

	ab	$a\bar{b}$	$\bar{a}b$	$\bar{a}\bar{b}$
ab	—	UP	Down	E
$a\bar{b}$	Down	—	E	UP
$\bar{a}b$	UP	E	—	Down
$\bar{a}\bar{b}$	E	Down	UP	—

COMPARATORS:

$D1 > D2 \rightarrow a$

$D3 > D4 \rightarrow b$

Fig. 2

Fig.3

THE LINGUISTIC INFORMATION ANALYSIS AND RETRIEVAL SYSTEM

S. Herbst

(Mr. Herbst is a first year undergraduate student who undertook the project described in the paper on his own initiative using Maniac III computer time made available by the Institute.)

1. Overview

The Linguistic Information Analysis and Retrieval System (LIAR) is a series of real-time, linguistically oriented disk storage and retrieval routines in Maniac III Algol designed to respond to statements and questions in common English (or, compatibly, any other language using the Latin character set). LIAR V, the current and most refined version, is discussed here.

It is important in an information retrieval device of this kind that the input string be reduced as much as possible to only those conceptual units which adequately characterize the question presented, and characterize it rather uniquely. It is furthermore necessary that this breakdown be carried out in such a way that grammatically different forms of the same question be reduced to the same set of stimulus data, and interchangeably elicit the same response. The analysis of input strings by LIAR, basically, depends on reference to a stored list of sentence patterns, the rules for which will be described presently. The stimulus information which results from this breakdown has two components: BASE, which is the main subject of the question, and MOD, pertinent modifying information which reflects on the main subject, and whose purpose is to distinguish different questions

concerning the same main subject. In the sentence, "What is the capital of Italy", for example, BASE is "Italy" and MOD is "capital". The reason for the twofold stimulus is to provide a dual level data structure, the address for random storage being derived from BASE (see sec. 3) and therefore, all information concerned with a particular main subject being located together. Although the simulation of human cognitive structures is not the goal of LIAR, the hierarchical framework facilitates random retrieval and speeds up search operations. Storage and retrieval procedure is dealt with in sec. 3; syntactical breakdown of input is discussed in sec. 2.

Provisions are being made to enable LIAR V to perform in French as well as English. At first, this will consist of purely independent operation in either language, with no cross-reference ability. As synonym replacement procedures are refined, however, (a synonym-synonym pair being stored as nothing more than a specific type of question-answer pair) interlingual mobility will allow information stored in one language to be retrieved in the other. Present synonym handling can replace English words with English synonyms, allowing for increased latitude in the form of the input question.

Examples of LIAR operation are given at the end of this report; the final section is a brief discussion of intended future developments of the LIAR routine and an idea of possible applications for the system. The following sections contain somewhat detailed discussion of the internal operation of LIAR V, and brief description of certain accessory facilities.

2. Breakdown of sentences

A. Control words. LIAR's performance does not rely on thorough grammatical analysis of input sentences. Rather, the structure which LIAR attributes to a statement or question is governed by the presence or absence, and position, of various "control words" such as "of," "is," and "because." Each of these words is assigned a fixed numerical value. Standard sets of words which are commonly used to begin a question ("who is," "how many," etc.) are separately treated as "operators." Those operators which imply simple associations and do not significantly modify the contents of the question ("what is," "who is") are ignored; all others are assigned the arbitrary value of 10. The words "a," "an," and "the" are ignored whenever encountered. All words not specifically assigned are given the value 1. Control words and their numerical equivalents are stored in a small sequential file. Operators and associated output information used to generate complete output statements from questions utilizing them are stored in the main random-access memory file.

B. Sentence structures. The sequence of numerical values for words in a sentence makes up the "structure" of that sentence. For "What is the capital of Italy" the structure is:

```

1 - capital0000000000000000
15 - of00000000000000000000
1 - Italy000000000000000000

```

The numerical structures are stored sequentially in the latter part of the file in which the control words reside. Information stored after each structure dictates how the words associated with its components are to be arranged to form BASE and MOD. In the above example, this information takes the form "3/1", implying that the third component ("Italy") comprises BASE and that the first ("capital") comprises MOD.

The information "5/3/9999/1" stored after an appropriate sentence structure indicates that an association is to be made with the fifth component in BASE and the third in MOD, that the resultant answer is to be stored in BASE, and a second association made with the first component in MOD. In this way the question, "What are the chief exports of the capital of Italy" is reduced ultimately to two associations: that of "italy" and "capital" with "Rome", and subsequently that of "Rome" and "chief exports" with the appropriate answer. The question "What are Italy's capital's chief exports", broken down according to a different sentence structure, would yield the same pair of associations as the question above it. In general, grammatically different forms of the same basic question or statement, analyzed according to different sentence structures, yield the same response.

Furthermore, information associated with structures pertaining to statements not only set up BASE and MOD, but also dictate the manner in which the answers to the inherent questions are to be extracted from the statements. The statement "The color of grass is green" should

generate the same BASE and MOD as the corresponding question, for example, and store the answer "green" on the disk. All answer material is supplied to LIAR either in the form of direct statements or as the answers to questions which have been asked and to which no response is known.

3. Storage and retrieval

The main memory of LIAR is, as has been mentioned, randomly accessed; the storage address is calculated from the numerical values of a sampling of the characters in BASE. The coded form of each data field on the disk is shown in the diagram below.

ADDRESS--

CHECK WORD 1

CHECK WORD 2

CODED OUTPUT

The address and the two check words are derived arithmetically from a predetermined sample of the characters in the two stimulus fields. The first check word is calculated from BASE, the second from MOD. The three numbers determine rather uniquely the particular BASE and MOD input, although none alone is unique to the string from which it was calculated. This triple check insures adequately that the information stored can be retrieved only by the same combination of BASE and MOD. Output information is stored in successive words, five characters packed decimally per word. Any number of words can be used for output characters,

the last labelled as such by adding $10^{**}10$ to the ten-digit decimal number. Information thus coded is unpacked and typed out as the answer to the question.

Since, as mentioned previously, the random address is not unique for a particular BASE, let alone a specific BASE and MOD pair, it is necessary to provide a certain degree of latitude as to the exact location of the data field. In storage, if the word addressed is occupied, the data is stored in the next available location. In retrieval, then, LIAR searches 50 or more words beginning with the one addressed until either it finds the answer to its question or has scanned a specified range, in which event it admits, "I do not know." Furthermore, although under normal conditions the storage of a new answer to a previously stored question replaces the old answer, it is possible to place the answer instead in a free location further on, allowing more than one answer to a particular question. This concept is useful primarily in storage of synonyms and foreign word equivalents, where more than one possibility exists. Eventually, much used words will be associated with synonyms in their own language and with equivalents in others, implying multiple associations with each word, that is, with each pair of BASE and MOD.

LIAR V's main data file is 80000 words in length. Assuming a minimum of 3 words per association, or an average of about 4, LIAR's theoretical capacity is approximately 20000 pieces of information. Under practical conditions, of course, the realized limit will be somewhat less. To assure reasonable conservation of space, LIAR V will

in the near future be able to store data fields containing long output strings in a separate sequential file. The file will then be utilized in storing pieces of text and other long strings which would impair the random allocation of space were they put in the main file.

4. Form of output

The output data stored in LIAR's memory generally comprises only a part of the answer supplied to a given question. In order to return the information in the form of a complete sentence, often part of the input question is repeated with the addition of various accessory words. The form of this generated prefixtural material is governed by the form of the operator. "What is", for example, initiates the output of the portion of the question which follows it, followed by the word "is". The retrieved answer then follows, completing a sentence. The question "Why is a raven like a writing desk", for instance, would have an answer of the form: "A raven like a writing desk because"

All input and output is by means of the Maniac III console typewriter, although card input may be used. Since Maniac III Algol has no provision for alphabetic character input, a small subroutine in Assembler III is used to receive the input question or statement.

5. Additional features

A. Pronoun. When a solitary pronoun or possessive pronoun (stored as control words) is encountered in an input sentence it is replaced with the contents of BASE for the last sentence processed.

After the question, "What is the capital of Italy", for example, the question "What is its population" will be interpreted as meaning "What is the population of Italy". Such replacement, although not always appropriate, can provide limited additional versatility in handling users' input without detailed grammatical and contextual analysis.

B. Replacement of MOD. It is often necessary for LIAR to employ a MOD in storage which is not physically present in the input statement. The sentence "The grass is green" implies, of course, "The color of the grass is green". It is important in such a case to recognize that "green" refers to a color, and not an equivalent to "grass". LIAR V, accordingly, is capable of storing the intended MOD for a particular word, specifically labelled as such, so that when scanning the words in a sentence it can load MOD with the appropriate stored string. Having been fed "The grass is green", therefore, LIAR V is able to answer the question, "What is the grass's color", whereas it will not simply associate "grass" by itself with the answer "green".

C. Why and because. When a question beginning with "Why..." is presented to LIAR V, the word "why" is loaded into MOD. If, however, an input statement is of the form, "A because B", it is necessary that provision be made to manually load "why" into MOD before storage. Sentence structures corresponding to various forms of this statement provide indication that "why" is to be stored, and the question, "Why is A" would elicit the response "A because B".

D. Subroutines. It is possible with LIAR V, instead of responding to a question with an alphanumeric answer, to execute a subprogram whose name is stored as an answer. Such subprograms could infinitely expand LIAR's potential; enabling it to perform mathematical manipulations or play music. The extension is an obvious one, but apparently significant.

6. Future applications

At present, LIAR is in the experimental stage--the immediate goal is to refine techniques used in the routines and to increase speed and efficiency in operation with small quantities of data. Eventually, it is hoped, LIAR will provide practical, rapid access to large quantities of verbal information. Either the data stored will be concentrated around a particular subject area, comprising a relatively complete record of that area, or it will be composed of scattered pieces of information supplied by individual users. Applications in the realm of the first case are, it would seem, apparent. The second case would apply in the context of a time-sharing system, in which case each user will be able to keep at his fingertips large amounts of personal records, and could at the same time have the benefit, in whole or in part, of information supplied by the entire user body. LIAR could, then, provide a link between participants in a time-sharing system, and would store answers to such questions as, "What is the latest word on Dr.'s research project" or "What was that idea I had last night". It is expected that, to this end, improved

synonym replacement operations will provide for increased latitude in the form of the input questions. All records would, unlike written notes or notices, be permanent unless specifically changed, retrievable if needed after long periods of time. The virtue of a retrieval system such as LIAR is, then, that unrelated verbal information is difficult to organize and store physically, it being much more efficient to be able to ask a simple question and have it answered, have tables of data supplied, etc.

It is intended, of course, that LIAR eventually will be equipped with protection devices to insure limited accessibility to certain "private" data, and also that it be tied in with other systems programs to provide verbal access to them.

At the moment, LIAR is in Algol; it would be difficult to make the improvements which are constantly being made were the program in Maniac III Assembler language. However, it will become necessary eventually to translate LIAR into Scheme M to run on the new computer when Maniac III is replaced. The transition shall be made--LIAR shall not die.

APPENDIX I : EXAMPLES OF OPERATION

→ The mass of a Deuterium atom is 2.01410 a.m.u.

STRUCTURE: 1 - mass0000000000000000	STORED INFO:
15 - of00000000000000000000	
1 - deuterium atom000000	10000000003 (statement)
20 - is00000000000000000000	1
1 - 2.01410 a.m.u.000000	5 (answer)
	0
BASE: deuterium atom00000000	
MOD: mass00000000000000000000	(ANSWER STORED)

→ What is the mass of a Deuterium atom

STRUCTURE: 1 - mass0000000000000000	STORED INFO:
15 - of00000000000000000000	
1 - deuterium atom000000	3
	1
BASE: deuterium atom0000000000	0
MOD: mass00000000000000000000	0

ANSWER: 2.01410 a.m.u.

→ What is a Deuterium atom's mass

STRUCTURE: 5 - deuterium atom000000	STORED INFO:
1 - mass00000000000000000000	
	1
BASE: deuterium atom0000000000	3
MOD: mass00000000000000000000	0
	0

ANSWER: 2.01410 a.m.u.

→ How many days in June

STRUCTURE: 10 - how0000000000000000	STORED INFO:
1 - days00000000000000000000	
15 - in00000000000000000000	1
1 - june00000000000000000000	12
	0
BASE: june00000000000000000000	0
MOD: how days000000000000000000	

ANSWER: 30

APPENDIX II: OUTPUT EXAMPLES (INPUT TERMINATED BY "<")

ask away.

What is the capital of Italy<

the capital of Italy is Rome

ask away.

The mass of a Deuterium atom is 2.01410 a.m.u.

thank you

ask away.

What is the mass of a Deuterium atom<

the mass of a deuterium atom is 2.01410 a.m.u.

ask away.

What is a Deuterium atom's mass<

a deuterium atom's mass is 2.01410 a.m.u.

ask away.

What is a Deuterium atom<

a deuterium atom is I do not know.

an isotope of hydrogen

thank you

ask away.

What is a Deuterium atom<

a deuterium atom is an isotope of hydrogen

ask away.

What is an atom of Deuterium<

an atom of deuterium is I do not know.

A Deuterium atom

thank you

ask away.

What is the mass of an atom of Deuterium<

the mass of an atom of deuterium is 2.01410 a.m.u.

ask away.

Plants are green because of chlorophyll<

thank you

ask away.

Why are plants green<

plants are green because of chlorophyll

ask away.

When is the 4th of July<

the 4th of July is annually

ask away.

LIAR V is alive<

thank you

ask away.

What is LIAR V<

liar v is alive
ask away.
What is alive<
alive is liar v
ask away:
thank you<
you're welcome.

ESSENTIAL FEATURES OF PURPOSEIVE MOVEMENTS

P.H. Greene

I. Introduction to a Problem Area

This paper introduces and motivates a series of studies on cybernetic problems of sensorimotor structure in brain, artificial control system, and robot. I hope to show that many of these problems arise, not only in guiding movements, but in purely perceptual tasks, notably those of object recognition by machines, that are commonly studied in isolation from problems of movement. The general statements of this paper will be frequently cited in later papers to show how examples fit into the overall aim of the investigation.

Brains made animals move long before they thought or perceived much, and we might seek good brain models by looking at characteristic features of this control. Besides, merely tracing sensory coding back further and further into the brain is sometimes like putting coded television pictures further and further inside a machine, which would not be equivalent to perception by the machine. What are some characteristic features of movement?

1. Write by holding your pen still and moving the paper instead with your other hand. You can do it, although you have never done it before. You are somehow doing the "same thing" with different muscles that move differently and require different signals from your brain. Your brain could never store recipes for performing every variant of

every action by every different set of muscles. So how does an action "get into" a particular set of muscles? How could a robot guide complex linkages to perform skilled actions?

2. When a cat turns its head to follow a mouse, the act of turning will tune movement control centers in such a way that, at any time, the brain has only to decide "Jump!" and the jump will be in the right direction, because all has been prepared beforehand by reflexes.

3. The cat's brain is like a general who commands, "Take Hill 7!". A general does not have to know the details, because officers at lower levels adjust the plan of attack in accordance with local conditions. Likewise, the same command "Jump" makes the cat do different things under different conditions, because of adjustments in lower centers that monitor things like the position of the head.

4. Although a battle plan requires simultaneous actions by many soldiers, the general does not have to regulate that many independent actions: provided with officers who know how to modify these plans to fit particular conditions, and when to switch from one plan to another, the general has only to select from a small set of overall tactics.

5. Many activities of animals and of the most interesting computer programs and machines follow the above scheme. A group in the USSR Academy of Sciences has been systematically studying this aspect of the organization of movements,¹ but there has unfortunately been too little study in the USA. The neglect surely follows from the subtlety of organization needed for this kind of control, in which the top controller

does not know exactly what is being done below. This is because a low level plan might, at any time, be replaced by another, if the two happen to coincide throughout some range of conditions surrounding the existing conditions, or be replaced by a third plan that differs from either, but can be used because somewhere there is an adjustment that compensates for the difference. Since even the basic plans differ from time to time, muscle to muscle, and person to person, you may wonder how to search for any uniform scheme to organize this fluid situation, to get the "same action" into different muscle groups, and to recombine the same movements into different acts.

6. Customary mathematical studies of control systems fail to address these aspects of control. In subsequent papers, I shall point out some directions that I believe are open to rigorous, and often mathematical, theory. To prepare the way, I shall have to review areas that have traditionally been associated with rigorous analysis, and in which significant progress is continually being made; and next, compare the results of this analysis with problems that we should all like to understand, but that have not seemed amenable to rigorous treatment. The remainder of the study will attempt such treatment.

7. In so doing, I shall suggest that if I already had as much of a cat, or of a robot, as in Paragraphs 1-5, I should have so much patterned spatial information embodied in its various adjustment schemes, that to extract it would be to achieve cleverer perceptual pattern recognition than I should expect from most models that have been designed specifically for perceptual tasks.

II. How Will Knowledge of These Things Help Us?

Analysis of the above examples, and others, suggested the importance of studying systems that have a repertoire of basic actions they can perform and transformations they can apply to these actions. These systems synthesize their outputs by first selecting basic actions that are something like what they want to do. They apply appropriate transformations to the basic actions in order to fit them together smoothly and to shape them into the exact actions that are desired. When these generalities are given specific content through examples, common mathematical problems of organization are seen to arise that are applicable to several areas of research:

(a) Nervous control of movements. For example, it is commonly held that a function of the cerebellum is to modify individual pieces of actions that have been initiated by other systems, in order to match them to each other and to initial conditions, and to smooth the transitions from one to another. It may also regulate posture, and thus determine some of the initial conditions, parameters, and constraints of the actions.

We see how many of these variables must be controlled when we look further into the example of handwriting in the preceding section. The "same" action can be performed in many different ways--for example, a person can write large or small, with his arm held high or low or loaded with a weight, or with a longer or shorter pencil, or even with a

pencil held in his teeth, and although the muscles move differently in each case, the same type of handwriting always results.

The point of the hierarchical control of the preceding section which fits together structures already present at lower levels, is to avoid having to regulate and correlate all these variables individually.

(b) Control of robots. A robot with many degrees of freedom of movement would, like a person, have an infinity of motions that could lead to a single result. How can these vast classes of possible movements and their interrelations be recorded in the memory so that appropriate movements become available when needed for performing useful tasks that have been specified in terms of their effects upon the environment (like picking up a tool), rather than in terms of movements of particular joints and motors? In computer terminology, one must construct something like a compiler for actions.

In order to control a device, such as the robot, in which many degrees of freedom are combined to produce skilled movements, or in order to control an industrial process having many independent variables, it seems expedient to use control systems that need only a few degrees of freedom to select from repertoires of actions organized at lower levels, and then to shape the chosen actions into more approximate versions of the desired actions, by means of a repertoire of transformations that can be used like subroutines in many different contexts.

(c) Recognition of objects and pictures. If the process of recognition consists, in large part, of resynthesis of the image from elements produced

by the recognizer (either elements equivalent to geometrical figures, or else abstract symbols representing them)--and this is true of the recognition systems that do realistic tasks--then the problem of associating parts of pictures with appropriate modifications of some meaningful figural elements raises some of the same structural questions as does control of movements.

(d) Hybrid computation and automatic synthesis of nonlinear circuits.

These problems also arise if one wishes to design a system that could synthesize a network for performing a specified function, given a collection of available components and ways of interconnecting them and of shaping the response. For example, the responses of the nonlinear components might be approximated by piecewise-linear curves, and thus, within some neighborhood of a point in the configuration space of the network, its behavior may be approximated by that of an equivalent linear circuit. The desired responses are pieced together by choosing from equivalence classes of elements whose characteristics coincide over some neighborhood of the current configuration. Circuit elements must be introduced to maintain inputs to components within particular linear segments of their response curves, or to perform the appropriate switchings and adjustments when these segments are abandoned and members of other equivalence classes take over.

Similar problems of piecing together arise in certain methods of solution of equations or optimization of control by hybrid computers.

For example, an accurate but slow digital computer may produce a precisely known drift- and noise-free approximate solution that can be corrected by fast analogue computing elements, in such a way that analogue errors affect only a small part of the result. The digital computation may itself be an improvement of a crude analogue computation.

III. General Motivation

These remarks have introduced a program to formulate and study some cybernetic problems that I think we must face in order to learn how the nervous system regulates skilled actions, and how to build machines that can perform skilled acts of perception and movement.

Since I believe that even the best mathematical studies of these problems have concerned components of cybernetic systems, more than the desired functioning of systems comprised by these components, and since in most systems, the relevant mathematical methods differ at different system levels, I shall need to demonstrate the existence of some unfamiliar questions and some ways of dealing with them. This requires patient consideration of specific tasks for brains and robots, in order to perceive these structural problems, and to gain enough practice thinking about them, so that one may be in a position to draw even elementary conclusions.

Although any clever idea about regulating actions might presumably be applied to building useful machines, the study of what actually happens in the nervous system is more demanding. Theoretical investigation alone cannot, therefore, tell us how the nervous system works, but at most, help us learn to be a bit smarter about knowing just what to look for experimentally, in order to learn how it works. The biologist J.B.S. Haldane said, "The world is not only stranger than we know; it is stranger than we can know." The first aim of this study is to train ourselves so that we can know a bit more.

REFERENCE

1. Numerous articles in Biophysics, the English translation of the Russian journal Biofizika, in the section entitled "Biophysics of complex systems. Mathematical models."

CYBERNETIC PROBLEMS OF SENSORIMOTOR STRUCTURE: INTRODUCTORY REMARKS AND SURVEY OF A STUDY

P.H. Greene

This note sketches the purpose of a series of reports, because of recurring questions people ask, because parts of the discussion appear elsewhere, out of sequence, and because it saves space to tell at the start how each paper fits into the general program, rather than repeating this each time. The remarks and outline will be used for reference in forthcoming reports.

I. Why this type of study?

This study treats problems of understanding how we perform coordinated movements and acts of perception. The present conceptualization of models of perception emphasizes aspects that have been slighted in most attempts at rigorous investigation. Some of the same problems arise in designing a machine to perform skilled movements and acts, as in creating the sensorimotor organization required to deal, not with meaningless jumbles of sensory impressions, but with objects having permanent identities and moving in predictable trajectories, and the coordination required to manipulate these objects and to locomote about a varied environment.

Unfortunately, it is not enough to look at the physiological mechanisms of the nervous system without simultaneously taking steps to increase our understanding of the structure of the information handling

and information creating functions these structures perform. A specific instance will appear in the simple machine introduced in another paper ("An Aspect of Robot Control and Nervous Control of Skilled Movements: Coordination of Two Effectors and Transfer of Adaptation," Section III-D of this Quarterly Report). Even if we were magically to receive a preprint from the future telling us just which molecules or neurons embodied those functions, we would have no idea what to look for in these structures, because we have no idea what must be embodied, or even what it might mean to embody whatever it is. If Aristotle appeared today, and we trained him to use voltmeters and oscilloscopes, and we sat him down in front of a computer solving equations of quantum field theory, and we asked him to make measurements and to tell us "what the computer was doing," he would make less progress initially by probing into the computer and measuring things, than by going off by himself and training himself to understand modern concepts of mathematics and the physical world. Then he could return to the computer. The present study will, fortunately, not require so much preliminary self-improvement. We shall start directly with an examination of interesting biological and artificial systems, proceed to a slightly abstract interpretation of the features that seemed important, circle back for a closer look at biological systems, making use of the concepts and vocabulary previously introduced; look once more at concepts suggested by studying aspects of artificial systems, and return to consideration of biological experiments, hoping each cycle to come closer to isolating manageable chunks of what we can at present only point out as

global problem areas.

These remarks may appear obvious to many. I have emphasized them because of attempts by so many investigators to restrict their attention to components--either physical, like neural circuits, or mathematical, like methods for optimization, synthesis of feedback systems, or partitioning spaces with hyperplanes. Why can we know so much about the anatomy and physiology of the cerebellum, and yet not know "what it is doing" to aid the coordination and smooth execution of movements? We shall see examples of immeasurably simpler systems that exhibit problems of control theory that it takes practice even to perceive and think about before one can know what to study experimentally. The initial stage of this study, comprising several reports, will attempt to give a bit of this practice.

II. Survey of the study

We may begin with the observation that to select, from all the things one could measure in the nervous system, just those things that are most relevant to purposive activity, is almost a guess or an act of faith.

*A childishy simple, though artificial, analogy: The fact that a system is busily using trigonometric functions might provide an additional clue we need in order to surmise what the system is doing, and to know where to probe further. If Aristotle had the wit to divine the existence of programs in his computer, and to decode a subroutine that is setting $x(t)$ equal to $1.999x(t-1) - x(t-2)$, would he, by means of this "physiological" analysis of "summation" and "inhibition of x " and "synaptic delays," know what the subroutine "was doing," or that it was doing the same thing as a superficially different subroutine, unless he also knew such fanciful things as the meaning of difference equations and the functional properties of $\sin$ and $\cos$?

If we are guessing, we shall be no worse off to guess something that sounds like characteristic and interesting things the nervous system does, rather than something that does not, or that seems to be a detail, or an individual component of unknown usefulness, or a question of efficiency that might be interesting to ask after we knew how the system worked in the first place.

To help us find a good guess, a forthcoming report presents a thumbnail characterization of some customary mathematical approaches to problems of perception and actions, involving transformation and partitioning of input spaces. They will not be treated in detail, as this report is intended merely for later reference and as a setting for my claim that other types of problem exist as well, when we try to understand how these customary basic mechanisms might be coordinated into more comprehensive structures. A brief summary of parts of this report appears in another paper ("Seeking Mathematical Models for Skilled Actions", Section III-C, I.2; this Quarterly Report) which follows.

After this glance at mathematical models, we are in a position to survey interesting biological and artificial systems that do useful things, to see whether the mathematical models describe their essential features. This survey¹ is briefly summarized in another paper ("Seeking Mathematical Models for Skilled Actions", Section III-C, I.3, this Quarterly Report), and shows, at the level of general discussion, how notions of fitting together structures produced by a system pervades theories of a wide variety of successful systems. The procedures of the previous paragraph serve as

mechanisms, but the characteristic overall motor and perceptual function may be described, less as a transformation of the stimulus, than as a construction of actions, or a reconstruction of images, by forming combinations of patterns inherent in the system (whether innate or recently learned), that are selected, triggered, and modified, in part, by the stimuli. This raises a question: if we decide to study a few models first, should we not include these ideas of fitting together, rather than restricting ourselves to the "stimulus-response" and transformation approaches that have been the major concern of mathematical investigators to date? So far, the ideas have been introduced in a general non-mathematical way: the aim is merely to survey some interesting systems that suggest an attitude regarding fundamental problems to study. It is useful to know which direction to take before setting out.

Before we can formulate mathematical analyses of these procedures for fitting things together, we must move closer to the problems that have been identified, and attempt to isolate manageable subproblems. In Part II of Section III-C of this Quarterly Report, experiments on sensory adaptation to prismatically displaced visual images, and general observations on the coordination of movements in a simple robot, are used to pose eight general types of structural problem suggested by the present attitude. This argument maintains that a theory should contain at least some of these eight ingredients, if it pretends to deal with the interesting and characteristic aspects of purposive sensorimotor behavior that have been introduced.

This section is essential to the mathematical formulations that follow, because it explains the purpose of each mathematical consideration that will be introduced later, and it will be frequently cited in succeeding papers.

Section III-D, "An Aspect of Robot Control and Nervous Control of Skilled Movements," of this Quarterly Report, shows how some of these general problems arise in the simple instances of coordination in a robot. These are the first specific examples in which the mathematical problems will later be shown to arise, and they will be further elaborated in subsequent papers. Studies in preparation show (1) how an observable that turns out to be important for the robot is exactly the sort of stimulus that elicits the biological orienting reflexes; (2) how the robot's use of this information turns out to be a concrete realization of a previously known philosophical study of an abstract perceptual process for labeling equivalence classes of movements oriented with respect to positions in external space; and (3) how these mechanisms can help us think about the fine adjustments imposed upon spinal motor neurons and interneurons by higher centers in the nervous system.

All this preparation will then give us a perspective from which to survey a large number of biological studies of movement. Our concepts will enable us to make a selection from the vast body of biological information of typical studies that define our subject area.

Finally, all of the previous information will be used in developing a mathematical formulation with the aim of being able to say something

rigorous about what at present are attractive, but unanalyzable, global problems.

REFERENCE

1. Greene, P.H. "Models for Perception and Action," Proceedings of the First Annual Princeton Conference on Information Sciences and Systems, March 30-31, 1967, Department of Electrical Engineering, Princeton University, Princeton, New Jersey, 1967, pp. 245-253.

SEEKING MATHEMATICAL MODELS FOR SKILLED ACTIONS*

Peter H. Greene

Part I. Some Models that Present Problems1. The Need for New Concepts

We wish to understand perception and skilled action in man or machine. The investigation, of which this paper summarizes a preliminary part, reveals some new mathematical problems that may be important in studying the same systems that other persons are already studying quite competently from a different aspect. This paper continues an explanation of the area in which these problems lie, and the reasons why they seem important.

We should like to be able to say something revealing about the remarkable adaptiveness and flexibility of the movements we perform. One can hold a pencil between the teeth and write one's name, and the writing, allowing for clumsiness, comes out something like one's own handwriting. We surely do not have some mapping from all forty or so degrees of freedom in our shoulder, hand, and arm, to the dozens of degrees of freedom in our trunk, shoulders, and neck. It would take a table larger than the size of the earth to hold all that information, and even if the information could all be stored away, it could probably never

*This is an expanded version of a lecture presented at the Biomechanics Symposium, Rock Island Arsenal, Illinois, April 5-6, 1967. This research has been supported by the Information Systems Branch, Office of Naval Research, under Contract No. Nonr 2121(17) NR 049-148. Reproduction, in whole or in part, is permitted for any purpose of the United States Government.

be retrieved. Instead, there is likely some crude way in which one can trace the movements of either hand or head back to some general purposes, such as moving forwards, or upwards, and perhaps one has some way of building up a complex action from these general purposes, in some hierarchical fashion--or so people say. The trouble is, what does that mean? What should we look for to understand this? We could measure what the neurons are doing, but what are we supposed to measure? If we went to the State Department to learn about foreign policy, and weighed everybody, we would get precise numerical results, but they would not be very significant in respect to foreign policy. We should, instead, have to find out what are the meaningful things to observe, and this is of course true in any scientific investigation, including the present one. Just as physicists had to find out that certain combinations of mass and velocity, mv and $\frac{1}{2}mv^2$, were more meaningful than others, say mv^3 , we must find meaningful elements to observe in order to study coordinated actions and perceptions.

We wonder, for example, how we perceive meaningful objects moving in definable trajectories, instead of seeing jumbles of impressions, in which we mix up this page and a little branch from some tree, and the window blinds. Or, we watch a baby's abilities grow. He learns to turn his head, get the nipple between his lips, and suck. He learns to turn his head and eyes and look at something. He learns to move his hand and grasp something. It is easy to feel that there is a little bit in common among these things--he does not have to learn each of these things independently. That is, we feel, the reason why he becomes able to grasp anything, whatever

it is, using whatever means happen to be most convenient. Instead of merely coming to associate independent movements, the baby may build them all from common elements in such a way that each is built out of some ingredients of the others. Therefore, some people talk intuitively about things like a "schema of grasping," and we wonder whether we can say something precise about these appealing, but vague ideas.

Whenever attempts at precision have been the most mathematical, especially in psychology, the theories seem to ignore many of these problems of developing structure. They tell how fast or how strong things are, or how firmly they are connected, but not what they are. Their rigorous concepts do not grasp the elusive structures mentioned above. Anatol Rapoport¹ once described the investigations of a man from Mars who, chancing to arrive on earth in the midst of a chess match, and being impressed by this activity, created a mathematical theory of chess that predicts the average rate at which the pieces disappear from the board. It predicts this average rate with great accuracy, and it is replicable: if you take the first 10,000 and the next 10,000 games, the rate of depletion will be exactly the same. The Martian even developed equations that told him the distribution of the length of the game. But these are not the aspects of chess that we care about. Some highly developed mathematical formulations of perception, learning, pattern recognition by computers, and the like, seem to resemble the Martian's theory of chess. The purpose of this paper is to list some ingredients that should be contained in any theory of sensorimotor organization

of man or machine that aspires to be more relevant than the Martian's theory of chess.

2. Customary Formulations

Section 1 has asked the following questions: How does the nervous system, or how could a machine, regulate actions that must be organized in terms of a purpose, rather than in terms of particular muscle movements or movements of particular mechanical linkages? How can a person or a machine perceive things, and not just fragmentary and ever-changing patterns of excitation of sense organs? Attempts to answer these questions are widespread. Ignoring their specific methods of execution, we may characterize the common conceptual basis for most such attempts as a search for ways to partition an abstract space, representing input configurations, into regions corresponding to input types, or for ways to transform the inputs in a fashion that renders them more recognizable, or that exhibits invariants by which they may be classified. We shall see some examples in this section. In essence, the control system is regarded as a chain of mappings from input elements, through transforming elements, to output elements. In certain models, these are designated "sensory," "association," and "motor" elements, respectively. This schema is generally not conceived as a particular choice among many, hopefully, the best choice, but is implicitly taken to be the unique mode of description. An aim of this study is to explore a wider choice of models, emphasizing those in which the "sensory" organizations may themselves

be composed of motor patterns or their derivatives. To be an element of sensory experience, in such models, is to be part of an organization of these motor derivatives. An item of perception, in this view, necessarily contains elements of other perceptual items in order to be an element of perceptual experience at all: one does not merely aim to introduce suitably strong links between previously independent "sensations." Of course, an effective system would be likely to employ as components the useful transformations studied by other investigators. It will therefore be useful to mention briefly some of the customary approaches. It should be clear that these are not being criticized. They are useful components, but in any system, the problems of organization of the overall system differ from the problems of organization of the components, and the present study is directed toward the organization of the systems in which the components function.

It is easy to name examples of useful transformations of inputs. For instance, some transformations help us extract a signal from noise. These are essential to the task of interpreting weak signals that have come a long way, such as radar signals bounced off Venus. Statisticians have ways of transforming variables so as to separate the variability that distinguishes different classes of objects from the chance variability within classes of the same kind of object.

Another typical useful transformation is a transformation applied to an optical image in order to emphasize some feature, such as a boundary between light and dark regions. Local transformations are helpful--for

instance, many visual systems employ lateral inhibition for this purpose, the effect being to diminish the excitation at a point by some function of the average excitation in a neighborhood of the point. It should of course be realized that blind application of local transformations is inadequate--for example, the above contrast enhancement, together with the finest brightness discriminations, would result in perceiving every surface as broken up into myriad meaningless subregions. One really has to combine crude input transformations with more intelligent processing that then informs the sensory system when and where it is appropriate to inspect more closely.

If we receive a telegram that says HAPPP BIRTHDAY, we know what message was intended, and the reader, in fact, may not have even noticed the mistake. When error correction like this is done systematically, we may regard it as being equivalent to a function that takes a signal and transforms it into a meaningful signal. It is a function, that, when applied to HAPPP, yields HAPPY. Alternatively, we may think of the abstract space representing all possible signals, and partition it into regions, each containing one meaningful signal (e.g., HAPPY) and a number of meaningless signals (e.g., HAPPP) that resemble the meaningful one. Then any garbled message is replaced by the message obtained when the meaningful signal in a class is substituted for any other signal from that class that occurs in the message.

One of the useful ways a statistician transforms his data to separate classes is to construct a discriminant function. Suppose he has a collection of skulls that can be classified into two races. Suppose that each skull can

be characterized by twenty measurements. It may be that some combination of these twenty measurements will serve as a good indicator of the race. This simplifies the classification of a new skull, for it allows twenty numbers to be replaced by one number. If the statistician is lucky, this will be possible, and if he is even luckier, a linear combination, i.e., a weighted sum, of the twenty measurements may suffice. If this is the case, the statistician has rules that tell him how to compute the weights. Suppose he has found a linear combination that is positive for all skulls of race A and negative for all skulls of race B. Then we may describe what the statistician has done in two mathematically equivalent ways. First, we may say that he has found a function that transforms the twenty input variables into a more useful output variable. Second, we may observe that a boundary case separating races A and B is any skull for which the linear combination equals zero. The set of all measurements for which this is the case forms a surface, in this case a hyperplane, because of the linearity, in the twenty-dimensional space of the measurements. Thus, we have described the situation in two equivalent ways: either as the application of a suitable transformation, or as the partitioning of an input space into regions representing different classes of object. If linear combinations do not separate the classes, the statistician might try more complicated functions, corresponding to curved decision boundaries that might manage to squeeze in between the right points in the input space, and if this is insufficient, he might try to separate the points by using several decision boundaries, but the idea is the same. It is clear that these

procedures are useful, but equally clear that there are problems for which they would miss the point. For example, the classes might depend upon some property of the eigenvalues of a matrix composed of the observed numbers that could be easily ascertained by looking at the eigenvalues, but that has no clear relation to the observed distribution of the numbers. In general, it may be better to apply some sequence of deductions to the observations than to try to make a decision based on locations in the raw data space. Whether this is the case depends upon the particular problem, and no one supposes that there is a universal answer.

Suppose that a linear combination of some functions of the data values will serve to distinguish two types, as above. Instead of computing this combination, as above, suppose that the statistician picks weights at random, and if his sum comes out negative when it should be positive, he pushes some of the weights upwards, and if the sum comes out positive when it ought to be negative, he pushes some weights downwards. It can be proved, under our supposition that the linear discriminant exists, that if he does this long enough, he will eventually arrive at a suitable discriminant, and he can stop pushing the weights. If a machine automatically pushes the weights up and down, it is called a perceptron. Perceptrons can solve many problems that involve iterative adjustment of weights^{2,3}, but, as in the case of the statistical procedure they automate, there are problems for which they would be inappropriate, or would at most function as components of a larger process, to be plugged

in wherever a suitably weighted combination of things would be useful. Perceptrons might try to classify visual patterns by forming weighted combinations of the outputs of property detectors, although in realistic recognition tasks,⁴ involving objects seen together with other objects, against various backgrounds, it has been found that such transformations must be part of a more complicated multilevel procedure involving numerous deductions. Perceptrons have been used to perform the whole task of recognition from beginning to end, without deductions at higher levels of processing, and in so doing, have often been applied, not to the outputs of property detectors, but to the light intensities at individual points of the pictures to be recognized, which seems odd, since they do not consider any of the structural features of the picture that people use for recognition. However, they work in simple instances.

The point is, that the methods described above are very useful in some situations, are useful in other situations as elements of complex methods described in terms of questions of a different nature (e.g., deductive methods, or properties of eigenvalues), and are irrelevant to other situations. They may be representative of certain processes performed by the brain, but not of other brain processes, and it might be fruitful to look more closely at what the brain is doing before deeming it worthwhile to elaborate these models.

The same general comments apply to learning theory, which has dealt in an elaborate way with the consequences of postulates regarding the rate of strengthening of numerically characterized connections between...

observable units of gross behavior. A closer look at the logical and biological problems of behaving adaptively--without preconceptions as to the form or the mathematics of these problems--might reveal significant subproblems and meaningful chunks of these subproblems (analogous to the physicist's mv and $\frac{1}{2}mv^2$) that follow simpler laws than the chunks currently studied by learning theorists, and that might help decide among competing models of learning. Might the customary chunks be derived from preconceptions that miss many of the relevant aspects of sensorimotor problems, like the Martian's theory of chess? They might and they might not. This paper is part of an attempt to find out.

3. A Common Feature of Some Successful Systems

With these comments as background, it is possible to explain the general course of the study. A previous paper⁵ attempted to delineate, in a general way, a broad area in which interesting problems lie. That paper considered a number of interesting systems that do something practical: model-reference adaptive control systems for stabilizing and guiding systems such as airplanes that must rise quickly from the ground to high altitudes, in which internal models of system functioning are used in order to provide stability and simplicity of transfer function, despite rapid changes in conditions; physiological systems for stabilizing vision during movements of the body, and physical systems for determining the location of a target detected by a turning radar antenna, by means of

of one or another version of comparing image locations with models of intended movements or expected appearances; the Newell-Simon-Shaw chess playing computer program, in which moves are selected by analysis of the productions of a handful of single-purpose move generators, in order to avoid a huge tree search; the method of ethology, the science of instinctive behavior and its relation to learning, in which an organization of actions that looks like the rough diagram of the chess player provides a simple way of storing the capacity for instant production of a specific adaptive action under a wide variety of circumstances that the animal might not have experienced before; computer programs for picture recognition that can be said to reconstruct a model of the picture out of picture elements contained in the program; the model of the psychologist Piaget for early sensorimotor development of infants, in which a system that could be discussed in terms of the chess player and the picture recognizer provides an organization suitable for achieving a growing hierarchical organization of skilled actions; the theory of thought proposed by Freud in the last chapter of The Interpretation of Dreams, in which a model, again resembling the chess-player-ethological-Piaget organization, seems a plausible one for explaining certain structural properties of thought; and, finally, certain computer-controlled systems for graphic display and transformation of drawings, and recognition and manipulation of objects, that treat visual images not as collections of independent points, but as highly structured configurations of meaningful figural elements processed in ways that could be described in some of the same terms used by Gestalt

psychologists, but with perfect precision.

Discussion of these systems revealed a common thread linking diverse models for diverse aspects of behavior, even though the proponents of some of these models were indifferent or antagonistic to one another's aims. This common fundamental aspect is the synthesis of an output or of a recognized image from structures actively produced within the system, and not merely from the transformation of inputs or subdivision of an input space. A basic task in all cases is the selection, modification, and fitting together of these structures, and the usual mathematical treatments say nothing about some of the most significant problems revealed in this task, as I hope to show in this paper, and in a mathematical outgrowth of the arguments of this paper, to be presented elsewhere. The usual models are not wrong--they probably function as components of the overall synthetic task. I repeat the remark in Section 1 that the present study assumes that the components work successfully, and deals with the organization of the systems in which they function.

The examples all suggested that the system as a whole does not just have something going in and something coming out. Rather, the process one sees the most is the system itself producing a collection of structures that are fitted together. In the examples, the structures were often elicited by causing excitable systems to turn on and produce some pattern of activity that could sometimes be produced spontaneously, in the absence of stimuli. Instances of such patterns were described

in the previous paper. These patterns were such that they could be fitted together into a useful response. The attitude adopted was that the response could be better described, less as a transformation of the stimulus, than as a combination of patterns inherent in the system, that are selected and triggered by the stimulus. Then they are modified to match one another and the requirements of the particular situation.

All this discussion has reviewed purely verbal consideration of extremely general features. The present paper will look more closely at specific examples of these features in an attempt to discern interesting questions to ask about them, and interesting structural features that can be studied mathematically. Although this discussion will not be mathematical--except in the most qualitative way--it will place us in a position to discriminate a theory that at least tries to face these questions from a theory, however highly developed, that, like the Martian's theory of chess, ignores the features most characteristic of the problems.

Part II. Introduction to New Problems: What a Theory Should Contain

1. On Not Avoiding the Problems

The preceding section presented a common feature that served a wide variety of purposes in a number of natural and artificial systems. If this feature is really so important, then rushing to work out the details of an easily formulable mathematical model that misses the most common problems introduced by this feature may be in part an instance of looking for the lost ring under the street light, where it is brighter than

at the place the ring was dropped.

This common feature, the construction of motor actions and perceptual structures by fitting together patterns produced by the system, when stated abstractly, does not necessarily differ from the more common formulation in terms of transformation of inputs. The two formulations, however, are expressions of different points of view that lead to different selections of structural problems to study. The structural problems studied in learning theory, decision theory, and the theories of machines designed to adjust weights and to discriminate among patterns by partitioning input spaces are well known. This section begins the task of showing that the other point of view, while presupposing the customary mechanisms as components, at the same time really leads to new structural problems that are now ready to be studied mathematically.

This can only be done by presenting a series of specific examples of problems that arise in fitting together the productions of specific systems in specific situations. How could such productions be fitted together to produce the integrated behavior patterns mentioned in Part I, for instance, the things described vaguely as "hierarchical organization," and a "schema of grasping" so that in some way each of a number of grasping actions "is built out of some ingredients of the others?" What can we say to get our hands on these important things so as to study them rigorously?

This part of the paper will, through examples, introduce eight general types of structural problem that seem characteristic of sensorimotor behavior.

If these problems are really as characteristic as claimed, then any theory, however "rigorous," that omits all these ingredients must surely resemble the Martian's theory of chess. Although this part of the paper points out several general directions for study without showing how to follow them, it will already provide a point from which to look back at existing mathematical models to observe whether they address themselves to the problems. In addition, each of the examples is needed to explain the meaning of the definitions in a mathematical treatment in preparation, and to provide labels for problems that will be referred to in discussing physiological results.

2. Mappings of Spaces

A well known type of sensorimotor experiment⁶ is a good vehicle for introducing several of the desirable theoretical ingredients. If a subject wears spectacles bearing prisms, the apparent positions of objects are displaced, and when the subject first reaches for them, he reaches wrongly. In the laboratory, the subject reaches below a table toward the apparent image of a target reflected by a mirror placed so that the image appears to be beneath the table. In this way, the subject sees the object, but does not see his hand as he reaches, and cannot correct his error. If, however, the subject practices looking at his hand through the spectacles, while he moves it about, he somehow adapts to the prismatic displacement, so that when subsequently tested, he reaches more accurately towards the object under the table. With enough practice looking

at his hand through the prisms, he may adapt completely, and reach towards the true position of the object. If the spectacles are now removed, the subject reaches wrongly again. It is as if objects appear, because of the prisms, to be located at a certain distance, say, to the right of where they actually are. The subject, therefore, comes, through practice watching his hand, to expect things to be a certain distance to the left of where they appear to be. He is thus able to compensate for the prismatic displacement. Now, when the spectacles are removed, he continues to reach to the left of where he sees the target, even though there is no longer any prismatic displacement to compensate, so that he reaches wrongly until he readapts to the normal state. This naive description very likely confuses questions concerning the changes occurring in the subject, and the interpretation of the underlying mechanism given here is not to be taken seriously. It was just a quick way to describe how the experiment looks. The true nature of the underlying mechanism is the subject of lively controversy among persons who do do experiments like this. There is, for example, disagreement regarding the locus of the adaptation within the subject: is there a recorrelation of what is seen and the movement chosen in order to reach it; or does the change lie entirely within the proprioceptive system, so that the subject feels his hand to be at the apparent location of the target when it is actually at the objective target location, and he therefore continues to choose the "same" movement, which turns out differently; or is some other interpretation more appropriate?

We may use the initial error upon removal of the spectacles as a measure of the change that has occurred during training. If the subject has been exposed to a negligible duration of watching his hand through the spectacles, then when he removes the spectacles, there will be no initial error, while if he had a lengthy exposure to the displaced image of his hand, there would be a large initial reaching error upon removal of the spectacles. We may plot the initial error E upon removal against the training time T in order to see the progress of the change in the subject. Figure 1 schematically illustrates the results obtained by Hay and Pick⁷ in three versions of an experiment in which subjects adapted by wearing prisms during normal activities. These three experiments involve the pairwise correlations of three sensorimotor systems: the visual system, the auditory system, and the arm system. (Once more, the descriptive interpretations may not be ideal, but will suit the present purpose.)

The first version is the situation that has already been described: the subject is required to point, without visual feedback, to a target that he sees. The initial error, as a function of training duration T (the length of time he moved about while wearing prisms), is given by curve 1 in Figure 1. In other words, for relatively short prism exposure, initial error upon removal of the prisms is roughly proportional to exposure duration T , while for longer exposure, the initial error upon removal levels off at an amount of change slightly less than the deviation Δ induced by the prisms. This test has measured a change in the relation between the visual system and the control system of the arm.

Figure 1. Changes in coordination between eye, ear, and hand during six days exposure to optical displacement. (Redrawn schematically from Hay and Pick.⁷)

The second test measures a change in the relation between the auditory and visual systems, induced by the discrepancy between the direction of sounds and the prismatically displaced appearances of the things that produced the sounds. In the second test, the subject hears a sound that comes from behind one of a number of labelled targets that he can see, and he must face the sound and name the target behind which he believes the sound source to be located. As shown by curve 2, the initial error upon removal of the prisms rises more gradually than curve 1 with increasing exposure duration, and it levels off at somewhat lesser amount than the eye-hand error. The third test measures a change in the relation between the

auditory system and the system controlling the arm. In this test, the blindfolded subject faces a sound source, and is required to point toward the source. Curve 3 shows the experimental result. The pointing error first increases with increasing exposure duration, and then decreases, so that with long enough prism exposure, the pointing error is again zero, as it would have been if there were no exposure.

Hay and Pick also tested eye-arm, ear-eye, and ear-hand relations at various times after prism exposure of long duration, and they tested relations between other pairs of systems at the instant of removal of the prisms after varying exposure duration. These tests will not be described here, because the results described so far are adequate to make the point of this section.

Hay and Pick interpret adaptation to involve a proprioceptive change in the arm system, so that the subject thinks he is pointing where he sees the target, while he is actually pointing closer to where the target really is. This is in accord with the view of Harris.⁸ Hay and Pick also observe that the function represented by curve 1 (eye-hand) looks like the sum of the function represented by curve 2 (ear-eye) and the function represented by curve 3 (ear-hand), and suggest that the eye-hand change is the sum of a slowly arising, but lasting, non-proprioceptive ear-eye change, and a quicker, transitory proprioceptive ear-hand change. It is not clear that all these views are consistent, or that there is any good reason why one kind of change is lasting and the other transitory.

All the curves, including those not described here, are adequately accounted for by a hypothesis of Held.⁹ The point of this section is to use Held's interpretation as an example of a hypothesis that raises many interesting and complicated questions of a certain kind. To study the mechanisms of adaptation, if Held is right, or to find good ways to prove Held wrong, one should be able to consider these questions deeply. These questions typify the first kind of problem to which a sensorimotor theory should address itself, if it is not to be like the Martian theory of chess, and we must now look at Held's interpretation to see what the questions are. It will be clear that briefer examples could have introduced the same questions, but the present example was used, because it will be used again and again to introduce several different problems.

Held believes that eye-arm adaptation consists in a changed correlation between something in the visual system representing visual direction and something in the arm system representing arm direction, rather than a proprioceptive change that would affect tasks involving the arm alone. (He has evidence against changes in the latter tasks. Also, his interpretation does just as well as Harris' interpretation of the observations described by Harris.⁸ However, the present purpose is not to support Held's view, but to show that they raise questions that are worthy of study.)

Thus, for Held, the eye-arm change represented by curve 1 is a primary effect of prism exposure, induced by the discrepancy between where the subject sees the target and where he should reach for it. Similarly, the ear-eye change represented by curve 2 is a primary effect induced by the

discrepancy between where the subject hears a target and where he should look for it. The strangely transitory ear-arm change represented by curve 3 is, for Held, an immediate consequence of a very interesting hypothesis.

Held supposes that the correlations between direction representatives of two non-visual systems, such as ear and arm, are not established directly, but proceed via the visual system, or some system of variables in one-to-one correspondence with the visual direction representatives. Then the change in the ear-arm relation would be the resultant of the change in the ear-eye relation and the change in the eye-arm relation, rather than being a primary adaptive change, as supposed by Hay and Pick. The shape of curve 3 follows from this hypothesis in a way that will be spelled out in detail, because it will be used to illustrate several points in the sequel.

Once more describing the situation perhaps naively, in order to be brief, suppose that a person sees a visual target at direction x , where x is an angle (or distance) to the left or right of straight ahead, positive to the right, negative to the left. In normal circumstances, he reaches for it in direction x in arm coordinates, and he expects to hear it in direction x in ear coordinates. That is, the normal environment has induced a certain correlation between eye and hand direction representatives and a correlation between eye and ear direction representatives, and we may label the correlated representatives with the same number x , the objective direction of the object that is seen, reached for, and heard.

If the subject wears prisms that displace images by the angle Δ to the right, then the same object, actually lying in direction x , appears to lie in direction $x + \Delta$. If the subject is to grasp this object, as he is repeatedly required to do during the period of training, in which he pursues customary activities, he must, of course, reach in direction x , where the object lies, rather in direction $x + \Delta$, where he sees the object. This discrepancy, according to Held, results in a remapping of the arm coordinates onto the visual coordinates, whereby a visual direction v comes to be correlated with a hand direction that comes closer to $v - \Delta$ as training proceeds. Let us represent this gradually adapted hand direction as a function of training time T in the form $v - f_1(T)\Delta$, where $f_1(T)$ represents the fraction of full adaptation to the prismatic displacement Δ . It is easy to see the form of f_1 . When the prisms are first removed, an object actually at x will appear to be at $x + \Delta$, and the subject, reaching as though he still wore prisms, will reach to $x + f_1(T)\Delta$. Hence, $f_1(T)\Delta$ will be the subject's initial error, shown by curve 1 of Figure 1, and $f_1(T)$ represents the initial error as a fraction of the error Δ that would result if there were complete adaptation. Therefore, f_1 rises from 0 to a maximum somewhat less than 1, and its graph has the shape of curve 1.

In similar fashion, the discrepancy between an object seen at v and heard at $v - \Delta$ leads, according to Held, to a new correlation between eye and ear coordinates. Since the adaptation is gradual, the

subject comes to expect, after training of duration T , that something seen at v will be heard at $v - f_2(T)\Delta$, where f_2 represents the adaptation as a fraction of full adaptation to the displacement Δ . Curve 2 in Figure 1 shows the form of f_2 , which rises more gradually than f_1 to a somewhat lower maximum.

The last two paragraphs merely introduce symbolism that, with slightly different interpretations, could represent the views of Hay and Pick. The significant part of Held's hypothesis, as stated above, is that the changed correlation between the ear coordinates and the hand coordinates is not a primary effect of the training, but simply a consequence of the ear-eye and eye-hand adaptation. This is because Held supposes that the correlation between the location of a sound, and the direction of the arm in pointing to the sound, is achieved by first mapping the ear coordinates into the corresponding visual coordinates, and then mapping the latter into corresponding arm coordinates. It might seem odd that the ear and hand do not communicate directly, but one might speculate that something associated with the visual system provides the only scale calibrated well enough to be used for remembering and measuring positions and angles, so that other systems make use of this scale, much as we might use a ruler to compare the sizes of two objects, rather than bringing the objects themselves together, especially if they must be observed at different times.

If, for whatever reason, we adopt the hypothesis of a visual intermediary, we easily predict the subject's behavior when he points

to a sound. Suppose the sound really comes from direction s . Then it also has the direction s in ear coordinates. If the subject has worn prisms for a time T , his eye coordinate exceeds his ear coordinate by an amount $f_2(T)\Delta$, so that the auditory direction s is represented by the visual direction $s + f_2(T)\Delta$. This visual intermediary must be transformed to an arm direction. Since we know that the visual coordinate exceeds the arm coordinate by an amount $f_1(T)\Delta$, the resulting arm direction is $s + f_2(T)\Delta - f_1(T)\Delta = s - [f_1(T) - f_2(T)]\Delta$. That is to say, the subject points $[f_1(T) - f_2(T)]\Delta$ to the left of the actual direction of the sound, and this, as shown by curve 3, is just what Hay and Pick observed, since curve 3 represents, approximately, the difference between the errors represented by curves 1 and 2.

If one simply writes down the stimulus coordinate in each of the Hay and Pick experiments described here, and transforms in the indicated way to the corresponding coordinate in the response system, one obtains the same results as Hay and Pick in all but one detail (a gradual increase in ear-hand error after the initial growth and decline, in their experiment I), for which Hay and Pick offer no explanation either.

Obviously, any number of much simpler examples could have been used to illustrate the fact that the nervous system establishes mappings between variables ("mapping" simply being a synonym for "function"). The point of using this example will gradually emerge. Some of the reasons are:

(a) This is a complicated example that obviously requires much more theoretical and experimental investigation in order to be well understood. Hence, this example is a good illustration of the point of this section, that a good sensorimotor theory must tell us what to ask in order to study the establishment of mappings between systems.

(b) The intersensory mappings are really more complicated than the way they were described above. They are actually likely to be composed of many coordinated partial mappings, and this aspect of the present example will be used to illustrate the points of some of the succeeding sections of this paper.

(c) The difference between two ways of mapping the same system onto another will become significant in succeeding sections, and in mathematical formulations that will appear elsewhere, and the difference $f_1(T) - f_2(T)$ will provide a simple example.

(d) Information might be mapped into a particular system because that system has subroutines for working on the information. We shall be interested in ways in which we might map information into some system, work on it, and then put it back where it came from.

One ingredient, then, that ought to be present in a sensorimotor theory that is better than the Martian chess theory, is some information about the nature, establishment, and modification of mappings between sensorimotor systems. Very little is known, for example, even about how to identify the features that represent movements and positions and are correlated with one another. Although the moral of this section is

expressed in terms of generalities, it indicates a direction that will be explored in greater detail, and, together with the conclusions of succeeding sections will, even in the present general form, help provide a perspective for viewing existing mathematical theories of sensorimotor organization.

3. Sets of Allowable Transformations

Continuing our examination of the prism experiments, we may make the general observation that the subject has some mechanism, such that, once this mechanism has adapted to the displacement, it no longer matters that the spectacles are being worn. Somehow, displacement of visual images is a change for which the nervous system can compensate. Therefore, if two sets of visual experiences differ by this displacement, then for some purposes this difference does not count, because it will be compensated for. This general situation seems important in many of our activities. For instance, we can move things, and it is unimportant how heavy they are, and we can push things with a stick, and it is unimportant just how long the stick is--we always make the right compensations. We shall therefore be interested in some set of allowable transformations (which will, in the sequel, sometimes be called "fudges" for short), so that if we really ought to have a particular mapping, but instead of getting it, we get something that is "in the same ball park"--it only differs from the desired mapping by one of this allowable class of transformations--then what we get will be considered as good as the mapping we wanted, because

there is a mechanism that can automatically convert it to what we wanted.

The second desirable aspect of a sensorimotor theory, then, is that it should consider equivalence classes of things that differ by one of a class of allowable differences, as usual, this will take on meaning only when we come to specific contexts.

4. Equivalence Classes

Consider the relation between training and testing in the experiment in which moving one's hand while watching it through prismatic spectacles led to some sort of adaptation? What is adapted? During exposure to the displaced image of his hand, the subject typically moved his forearm back and forth, while his upper arm remained relatively fixed. During testing he often kept his elbow relatively fixed while he used his whole arm as a unit to point. Thus, at least two degrees of freedom of the arm are concerned, two kinds of rotation through angles that will be called R_1 and R_2 . There are many different combinations of R_1 and R_2 that all cause the hand to point to the same place. We may represent these combinations schematically by Figure 2, in which a few of the equivalent combinations are represented by curves labelled according to the number of degrees to the right of center to which the hand points. During training, the subject may only change R_1 , so that in this schematic representation, he is moving back and forth along the horizontal broken line. Adaptation

consists in learning that the combination B on this line that used to result in apparant position 0° now results in, say, 1° , the combination C that used to be 1° is now 2° , the combination A that used to be -1° is now 0° , etc., in case the prisms shift everything 1° to the right. When you test this adapted subject, he may indicate directions by maintaining a fixed R_1 while varying R_2 , so that he is now moving

Figure 2. The curves represent combinations of two degrees of freedom of the arm that result in the same hand positions.

along the vertical broken line. However, he immediately knows, for instance, that the configuration at F corresponds to 1° , that G corresponds to 2° , that H corresponds to 3° , etc., even though he had no training with these configurations. He knows that F now corresponds to 1° because of his experience relabeling configurations like B, and similarly for the other configurations. The relabeling

of configurations on the horizontal line somehow extends to configurations on the vertical line. In other words, the subject has not merely relabeled individual configurations: he has relabeled the equivalence classes represented by the curves. We shall often be interested in such equivalence classes, in which configurations are classed together on the basis of their effects on the environment, rather than merely in terms of their internal degrees of freedom. The mappings we consider will often have to relate sets of such equivalence classes rather than sets of individual configurations. Another obvious instance of the use of equivalence classes determined by effects relating to the environment is involved in seeing an object in a fixed position while moving the trunk, the head, and the eyes in the head. A range of configurations of retinal image location, and trunk, head, and eye positions are regarded as equivalent with respect to perceived location of the viewed object. Mathematical studies in progress show an example in which individual elements are related only in a probabilistic way, according to some traditional kind of learning theory, so that examination of the raw data would only yield a theory about conditional probabilities, while consideration of the proper equivalence classes, would reveal units that mapped deterministically. We must therefore learn what combination of things to look at, and this general point, to be given specific content in the sequel, is the third desirable ingredient of a sensorimotor theory.

5. Fitting Pieces Together

Experiments on adaptation of reaching behavior to changes in visual images will still provide examples. Suppose that a normal subject is asked to point to a visible target. The perceived location of this target, in inches to the right or left of center, is represented by the abscissa T in Figure 3.

Figure 3. Reaching for a target. Subject reaches R inches to right of center when he perceives target to be T inches to right of center. C_1 , normal. C_2 , displaced image. M_1 , magnified image. M_2 , magnified and displaced image. T_1 , target position during training. T_2 , target position during testing.

Where does the subject reach? Since he is a normal subject, who has, of course, had extensive experience reaching for things, if he sees the target 1 inch to the right, he reaches 1 inch to the right, and if he sees it 2 inches to the right, he reaches 2 inches to the right, and so on:

he reaches where he sees the target. If the place he reaches, in inches to the right or left, is represented by the ordinate R , then this relation $R = T$ for the normal subject is represented by curve C_1 . Now let us introduce optical displacement, as in the prism experiments. Now, when the subject sees the target 1 inch to the right, he really ought to reach (say) 2 inches to the right, and when he sees it 2 inches to the right, he ought to reach 3 inches to the right, in case the prism has displaced the images 1 inch to the left. Therefore, the broken line C_2 represents the new relation, $R = T + 1$ that the subject ought to learn. As in the initial description of the prism experiment in 5.2, the description in terms of perceived target location would need to be thought out more carefully, if the experiment were really being described adequately. At present, a crude description of the experiment is being presented merely to introduce an idea that should be explored.

One way of performing the experiment is as follows. The subject practices reaching many times to a target that always remains in one particular perceived location, T_1 . Therefore, what the subject learns is represented by just one point P_1 , having coordinates $T = T_1$, $R = T_1 + 1$. After the subject has learned this, he is tested, using a target that remains in a single perceived location T_2 , different from T_1 . If the subject behaved according to a type of generalization curve often considered in learning theories, one might expect the learning to have resulted in a new reaching v. apparent target position curve that departed from C_1 only in the vicinity of the training abscissa T_1 . In that case, the subject would reach for the

test target seen at T_2 in a way represented by point Q. However, he actually reaches with the same amount of shift that he learned when reaching for the training target at T_1 , as indicated by point P_2 .¹⁰ In general, the subject reaches to various targets in a way represented by the broken line C_2 . In actuality, we cannot perform the experiment with such reproducible precision¹¹, so that the present description is, in reality, an idealization of the observed behavior.

We might say that the subject has learned to introduce a certain shift into all his responses. However, a different way of speaking about the same situation will prove useful for unifying the discussion of examples that will follow. We could say that all the subject needs is the one point P_1 , and he can "plug in" the whole curve C_2 . Now recall the basic model for this whole study: the significant thing to study is a collection of excitable "boxes." If a "box" is turned on, it can produce pieces of structures that are fitted together to produce adaptive behavior. We may think of the whole curve C_1 as one of the pieces that can be produced by the system that governs reaching. The variously displaced curves that could have been produced, given various amounts of image displacement, constitute a repertoire of pieces that could have been produced by this system. All this system needs is one parameter to specify which piece from the repertoire to use, and the single point P_1 supplies the equivalent of this information. This contrasts with the behavior that would have resulted if the subject had compiled a table listing apparent target positions, together with their associated reaching positions, in which case, the experience represented by P_1 .

might have led him merely to revise one entry or a few neighboring entries in the table. The ability to apply the shift learned at target position T_1 to other target positions like T_2 is part of our ability to deal with spatial relationships among objects in space, rather than our being limited to learning individual muscular movements, and we should like to be able to ask better questions about how the nervous system accomplishes this.

Next, suppose that the subject puts on magnifying lenses, and practices reaching toward targets long enough for him to learn the appropriate relation between apparent target location and reaching movement just as thoroughly as an untrained subject knows the normal relation. Then the trained subject's reaching relation is represented by something like line M_1 in Figure 3, which has a lower slope than C_1 , because, owing to the magnification, changes in apparent position now correspond to smaller changes in reaching.

Finally, suppose that a displacement is introduced, as before, in addition to the magnification. Suppose that the subject, already behaving according to relation M_1 for magnification alone, is now trained, with displacement added, in the task of reaching to a target at the fixed position T_1 , and suppose that the displacement is such that P_1 again represents the proper behavior. This time, if the subject is tested using an apparent target position T_2 , will he reach in the way represented by point P_2 , as if he plugged in the curve C_2 from his original repertoire,

or will he reach as represented by P_3 , as if he utilized a new repertoire of potentially available curves, selected or created by his training with magnification alone? (Adequate treatment of this type of experiment would require more careful consideration of which changes were functions of apparent target position, and which were functions of real target position, as ascertained by reaching. Also, the notion of apparent target position, is, as before, unclear. However, in this discussion, the experiment is used only as a hypothetical illustration of kinds of questions that could be asked. Actually, there is too much variability to perform this experiment accurately enough to obtain useful information. However, it might be possible to perform such experiments, using tasks that require motor behavior alone.)

This discussion suggests that if one were building a machine that is required to reach for things and to adapt to displacements in its sensory information (or if it has to reach to one side of where it sees an object, because it has to touch the object, not with its fingers, but with the end of a screwdriver), one might consider having a repertoire of relationships, like the straight lines above, that one could select and plug in by picking appropriate values of a parameter. One might then fit pieces of several of these curves into a more complicated curve.

For instance, the relation between the movement of a handle moved by a person's hand, and the movement of a visible object that is guided by the handle, might be some nonlinear function, shown in Figure 4, in which the abscissa is the desired position of the controlled object, and the ordinate

is the position of the hand or the handle, as will be explained.

Figure 4. Fitting pieces together.

The curve C represents the relation between the target position x of the guided object (abscissa) and the position h of the handle (ordinate), and hence, the desired position of the hand, that will guide the object to this target position. The hand system is required to approximate this relation, using the repertoire of relations at its disposal. Two of these hypothetical relations are represented in Figure 4 by the curves R_1 and R_2 , in which the ordinate h now denotes the hand position resulting from a relation in the repertoire, when the device that produces this relation is stimulated by an object at position x .

If the system has learned a way to choose, for each value of x , that action from its repertoire that best approximates the desired action, in this case, the action represented by curve C , that will enable the hand to guide the handle in the desired way, then the system has found a way to fit together its repertoire of movements into a very simple action that is adapted to one type of effect on the environment. Of course, these remarks are not restricted to the particular interpretation of the curves as positions of a hand and a handle. Curve C could be any desired input-output relation, and curves R_1 and R_2 elements from some repertoire that are selected to approximate C .

The examples in this section have illustrated the fourth general kind of thing a sensorimotor theory should contain, that is, a study of piecewise construction of action structures from a repertoire of simpler structures. This study would accord with the general point of view, summarized in Section 3, that many interesting systems synthesize their outputs by fitting together elements from repertoires of partial actions. Of course, all complicated systems fit things together. The real content of these remarks will come in the discussion of specific examples of fitting things together, and in the particular aspects of fitting things together that will turn out to be important. These examples, as well as a discussion of how they apply to physiological experiments on motor coordination, will be presented elsewhere.

6. Shaping and Fitting by Stimulus-Dependent Transformations

In the example illustrated by Figure 4, suppose that the system has already learned how to choose, for each value of x , that action from its repertoire that best approximates the desired action, as discussed in Section 5. Obviously, performance would be improved if the system had ways of modifying the responses R_1 and R_2 from its repertoire, so as to make them more closely fit the desired response curve C . In the hypothetical example of Figure 4, R_1 should be lowered more and more as x varies from the center of interval U_1 toward either end of U_1 , and the same rule of transformation, this time over the interval U_2 , will benefit R_2 , provided that the minimum amount of "lowering" is taken as negative, i.e., as raising. Corrections like this could be made if, for example, the system possessed a device that could lower the output of any other device (in neurological terms, inhibit any other subsystem), provided that it was given parameters specifying the location and extent of the region representing the appropriate values of the input to the other device, and specifying the minimum amount of lowering and where it is to be applied, and how the amount of lowering varies with stimulus distance from the point of minimum application. Given these parameters at the time it is called into action, this device could then apply an appropriate correction, as a function of the stimulus, to make an element of the original repertoire of actions conform more closely to the desired action.

Another way in which the above hypothetical transformations would be useful may be seen by inspecting curves R_1 and R_2 . Suppose that the

hand is controlling the handle so as to track a moving object that starts in region U_1 and moves into a part of U_2 that does not overlap U_1 . Then, unless the hand switches from the mode operation corresponding to R_1 to the mode corresponding to R_2 precisely at the point where the curves R_1 and R_2 cross, there will be a sudden jerk in the movement of the system tracking the object. Since a given action might require switching at some other point, it would be useful to have a central control device that could examine a general specification of the performance of an action in terms of the modes of operation to be used, and constraints on the switching points, and then provide the information needed to call for the appropriate transformations to smooth the response at the point of switching.

Another useful ability would be the ability to interpolate between primary modes of operation by specifying the weights in weighted combinations, of some sort, of the primary elements of the repertoire. Again, a control system would have to supply weights that are functions of the stimulus and of parameters of the operation.

One general advantage of these procedures is that a system might start out with an easily produced repertoire of actions that would enable it to perform vital functions--that would get it through the day, so to speak--until it had learned to fit these together into coordinated actions. Examples of fitting together a repertoire of actions produced by the system have been cited in Section I.3.

A more subtle advantage of constructing actions by choosing preformed partial actions and transforming these to fit the purpose is that such a multilevel procedure allows a relatively simple central control system to govern actions that must be performed by effectors having a large number of degrees of freedom. Roughly speaking, even if the effector system has ten degrees of freedom, crude versions of desired actions might be performed by choosing from a one-dimensional list of modes such as R_1 and R_2 (the coordination of whose internal ten degrees of freedom does not worry the central unit, since this coordination has been wired into the units that do R_1 and R_2). For example, it is well known that the position of a cat's head when it is watching a mouse presets enough preformed response mechanisms so that a relatively simple decision to jump will be adequate to bring the cat into a stable attitude at the desired location. A number of physiological experiments concerning this general strategy of constructing coordinated movements will be reviewed elsewhere, in the context of an attempt to formulate enough theoretical notions to tell what to ask about such a strategy.

The fifth general ingredient that should be present in an adequate sensorimotor theory is, then, an understanding of the problems involved in applying a repertoire of standard types of stimulus- and parameter-dependent transformations to a repertoire of standard partial acts, in order to fit these acts smoothly together, and to shape them into the desired action.

7. Excitable Boxes as Mappings

Continuing with the example of the last two sections, observe in Figure 4 that function R_1 is defined on some subset U_1 of the stimulus space X , and R_2 is defined on an overlapping subset U_2 . One may think of subset U_i ($i = 1$ or 2 , in this case) as the set of stimuli (or whatever variables and parameters might be represented by x in some other case) that are adequate to excite "box" i , which then selects from its repertoire some function (depicted by curve R_i) that is defined on U_i . At this stage of the discussion, the above is merely a fancy way to restate what has been said more straightforwardly, but this description will turn out to be useful in motivating the mathematical formulation that follows from the general treatment of this paper, a more detailed consideration of hypothetical examples of robot control, and a review of physiological experiments on motor coordination. These remarks, and all the others of Part II, are collected together to define a general attitude with which to start investigating sensorimotor problems, and to avoid repeatedly presenting fragments of these remarks, with the resultant need to redefine the same examples each time.

8. Atlases: Discrepancies at Joints

Continuing with the preceding example, in which the relations R_1 and R_2 are fitted together to construct an approximation to the target-handle relation C , we may imagine how a control system might coordinate the two relations. Suppose there are two boxes B_1 and B_2 that include

relations R_1 and R_2 in their respective repertoires. A central control system, given the knowledge of the task and of the target position x , can, in some way not considered here, decide which box to activate, and which mode of action to select from the repertoire of that box.

The active box uses this information to compute a signal s , the actual physical signal, e.g., a voltage, that is required to cause the particular physical mechanism of the selected mode to transport the hand to the position $h(x)$. Figure 5 lumps together, into a system called S , the central decision mechanism and the initial stages of the boxes B_1 and B_2 that compute the signal s , and shows the resulting s being sent to the remaining stages of either box B_1 and B_2 that produce R_1 and R_2 .

Figure 5. Switching from one mode to another.

This way of grouping the functions of the system may be summarized by saying that system S examines the stimulus and decides whether to activate R_1 or R_2 . Think of R_1 and R_2 as the actions of two different systems of motors and linkages. The input to a motor is a signal s

(think of a voltage). Since the output h of either motor must approximate some prescribed function of x , and the two motor systems work differently, there must be a change in the value of s as the system S switches from R_1 to R_2 somewhere within the overlap of their domains of activity.

Now, suppose that the target position x changes from one for which R_1 is adequate towards those for which R_2 is preferable.

Then, as x approaches the value for which R_1 and R_2 are equally good or bad, or for which switching should occur for any other reason, the system S must get ready to switch over to R_2 . It might be valuable for S to anticipate this switch, in the way that the cat anticipates his mode of jumping (Sec. 6), by pretuning the effectors, so that a simple decision to switch can be made at a moment's notice. To do this, S might store a table, or employ a function generator, giving, as a function of x , the change Δs that must be added to s as the system switches. In addition to its value in performing the switchover, such a table might allow some other system to know what change in s to expect in order to regulate something that is influenced by x , and it might allow the selection of a transformation to smooth the switchover. For one reason or another, then, it may be useful to know the difference between the two ways of supplying s to the motor systems (i.e., the B_1 way or the B_2 way) as a function of x throughout their common domain.

In what is at present a rather nonspecific fashion, the above may be described as a problem of fitting together two of the pieces out of

which the total action is constructed, and recording how the pieces differ at the joints, and what transformations it takes to smooth away the discrepancies between the pieces in order to make them fit. In a specific application, it might be useful if these transformations happened to be included in a repertoire of transformations that the system found easy to perform as subroutines wherever needed, this repertoire being an example of the collection of allowable transformations, compensations, or "fudgings" that was described in Sec. 3 as second on the list of ingredients that should go into a good theory.

In examples that will be used elsewhere in the mathematical formulation of sensorimotor problems, there will be a number of specific instances of fitting things together and comparing them at the joints. In one example that has been considered, dealing with adaptation to change in parameters of a task, the things that differ are different mappings of a space representing states of a manual control system when the parameter has one value, onto a space representing states when the parameter has another value. In situations like this, mathematicians use a vocabulary derived from customary terminology about maps and atlases. Suppose we wish to make flat maps of the world that we can bind into a book. We customarily consider the world to be divided into a convenient collection of overlapping regions, and we have a collection of charts, each representing

one of the regions. A particular city, say New York, might appear on two charts, perhaps those corresponding to "Northeastern United States" and "Middle Atlantic States." For brevity, let us refer to these two regions as U_1 and U_2 , just as in the description of the two overlapping domains of mappings in the hand-handle example, Figure 4, and let us call their intersection U_{12} . The subscripts correspond to the individual charts: in a conventional atlas, one might look up "Middle Atlantic States" in the index, and be referred to chart 2, which is our map of U_2 . Looking up New York in the index, one might find the references "Chart 1, C-4; Chart 2, H9." This, of course, means that the regions U_1 and U_2 of the world have been mapped into the Euclidean plane, with coordinates given by letter-number pairs, and that a city, like New York, that appears on two maps, has two different coordinate pairs in the plane, corresponding to the two different ways of mapping the region U_{12} where U_1 and U_2 overlap. Now, the two different images of U_{12} that appear on charts 1 and 2 of the atlas will probably not look exactly the same, because, for example, mapmakers' projections usually introduce different distortions of distances in different parts of the charts. If two images looked so different that only some exceedingly involved transformation, or lengthy list of coordinates, could reveal their point for point correspondence, then the atlas would be inconvenient to use. In an ordinary atlas, however, the correspondence is easily recognizable upon inspection, and is usually approximated quite well, over the small regions involved, by a very simple function--perhaps just a change of coordinate origin and direction. Suppose, in our example,

that the maps of U_1 and U_2 face the same way, and are drawn to the same scale. We know that New York is represented by "Chart 1, C-4; Chart 2, H-9." Suppose that Washington, D.C. also lies within U_{12} . Then, if we look up Washington in the index, and read "Chart 1, B-2;...", we might expect this listing to continue "...Chart 2, G-7." All this information is shown in Figure 6.

Figure 6. Two charts of an atlas.

In this figure, the two odd shapes on the left schematically represent the regions U_1 and U_2 (or accurately represent odd shaped areas within U_1 and U_2), and the grid on the right represents the coordinate plane, together with the images of the various regions and points under the two mappings. U_{12} and its two images are hatched. Of course, for convenience, a geographical atlas does not superimpose both images on the same copy of the coordinate plane, but represents them individually on two copies of the coordinate plane on different pages.

The dashed arrows indicate the discrepancy between the two mappings of U_{12} onto the coordinate plane. As drawn in the figure, they show the fates of two points in that mapping of the second image of U_{12} onto the first image of U_{12} , given by the inverse of mapping 2 followed by mapping 1. It is this composite function that one expects to be simple enough, according to some criterion that depends upon one's particular purposes. A mathematical atlas is defined in just this way: it is a collection of mappings of overlapping regions of a given space into a coordinate space in such a way that the discrepancies between any two mappings of the same region belong to some allowable class of transformations. Two different geographical atlases might be regarded as equivalent, for practical purposes, even though the regions and mappings may be different, if the maps are easily relatable to one another, in the sense that has been described. This may be expressed, in the mathematical context, by saying that the union of all the maps in both

atlases itself constitutes an atlas as defined above.

All this may be summarized by saying that in our theory, we shall sometimes be interested in studying how a structure may be built up by fitting smaller structures together, taking account of the transformations that may be necessary to make them fit together, and of discrepancies between different ways of fitting them together. This is the seventh item on the list of what a theory of sensorimotor structure should contain. It has been given specific content in work in preparation.

9. Hierarchies

Psychologists and others who have studied sensorimotor coordination speak of building up hierarchies of actions, and they have pointed out phenomena that seem to be most naturally described in such terms. The things one can say about such organization are, however, quite limited at the present time, and it is important to develop ways of knowing what to ask and observe in order to understand these organizations. Specific examples will once more suggest additions to the list of things a theory should treat, if it is to be better than the Martian's.

Consider the action of moving my right hand, while I am standing upright, from left to right along a horizontal line in a frontal plane that passes through a point about one foot in front of my chest. The successive positions of my upper arm and forearm are shown in the simplified diagram of Figure 7, which also shows the directions of change of shoulder and elbow angles.

Figure 7. Moving the right hand in a straight line from left to right. In this simplified drawing, the arrows on the upper arm and forearm indicate the direction of change of shoulder and elbow angles. The former reverses its sense at the end of the movement, and the latter near the middle of the movement.

I do not know how my nervous system really organizes this action.

I am merely using this as an example of how one might imagine programming a robot built somewhat like me to do this task. Now I might have a recipe in my head telling just how to move in that particular straight line. That means I would need to have another recipe for how to move in a third line, a still different recipe for how to move in a slightly curved path, and so on. Moreover, I would need different versions of each recipe if my trunk were to be held in a slightly different position, or if I were holding a small weight. Starting from the intention of moving in some path, I would be confronted with all these independent recipes, and I would not know any relations among them.

There are two obvious drawbacks to this scheme. First of all, there is the problem of storing every recipe for every combination of all the parameters that define the actions. The second drawback is that I cannot

profit from any similarities between actions. After all, the organization involved in moving along one straight line has a lot in common with that involved in moving along a nearby straight line, or along a similarly located path that curves a bit. I may do almost identical things in all three cases, except that I might, say, have to pull my forearm in a bit more or less. Therefore, I might feel, intuitively, that I could build up the actions in stages that could be used as subroutines wherever they fit. This would at once show relations between actions and reduce the number of independent things to be stored.

Suppose we imagine a way of breaking the left to right straight line movement into stages. First of all, I (or the robot) might have a reflex that causes me to reach out in front of me whenever I see an object, or whenever I am hungry. Such a reflex might serve nicely as one of the easily built-in mechanisms that, even before there is any time for learning, help me get through the day in a suitable habitat, because it is fairly good for reaching for food or support. At this stage, the exact distance and direction of reaching might depend less upon the location of the stimulus than on the degree of hunger or excitability. All that now has to occur in a second stage is a fairly uniform clockwise rotation at the shoulder, and we already have a fair approximation to the desired movement. Now, all that must be added to obtain a perfectly straight line is a decelerated flexion of the elbow, followed by an accelerated extension, with a deceleration and slight reversal of the shoulder rotation at the very end. We could say that the stages consist

of (1) a fixed reaching, i.e., specification of arm variables that are constant functions of the stimulus; (2) a constant velocity, or uniformly changing angle, of the shoulder rotation; (3) a constant decrease in velocity of flexing the elbow, from positive to negative values, together with a constant decrease in velocity of the shoulder movement of stage 2, so that this velocity passes through zero just before the end of the action.

This organization has several virtues. First, it starts out, not with random combinations of atomic movements, but with a collection of easily pre-programmed movements that might provide for immediate necessities before there is time for learning.

Second, this initial repertoire may constitute a good vocabulary for building up more complicated actions, because they occupy a strategic position between the "machine language" of the atomic motor movements and the complicated "task-oriented" organization of skilled actions. A discussion of this point, in relation to the psychologist Piaget's views on infant sensorimotor development was given in a previous paper.¹²

Third, this organization reveals relations between actions by showing what stages they have in common. For example, when hungry, I could reach or I could suck. If I reach, I could reach with my left hand or with my right. If I have chosen to use my right hand, I could add the second stage above, which introduces a particular type of shoulder rotation, or I could move my shoulder a different way. Having progressed to the particular second stage chosen, I could introduce a

third stage in various ways. I have, in this way, constructed a family tree of actions, and I can trace the tree backwards from two actions to learn what ancestors they have in common. The actual organization might be more complicated than such a simple tree, but the point is that any such hierarchical organization tells me which actions are cousins; they are no longer given by isolated recipes. Two movements that are performed by different effector systems might, by being traced backwards, both be recognizable as instances of reaching out and to the left. Information like this must surely be part of what is needed to understand the question posed at the beginning of this paper: how can one perform the "same" type of writing by holding the pencil in one's teeth?

The fourth virtue of the hierarchical organization is that the individual stages might be very simple, as in the example, so that it would be easy to build in the capability for their performance, using the raw materials that people think should be available in brains and robots.

The fifth virtue is that a small collection of these simple procedures may occur so often in a wide variety of tasks that a great lessening of storage requirements may be achieved by using the procedures as subroutines that are performed by specialized systems, and that may be called upon whenever necessary. Looking once more at the present decomposition of straight line movement, we see the following rough description of subroutines: (a) a repertoire of constant actions (modes of

of reaching out that are constant functions of the stimulus location, depending for their selection upon internal factors like energy level); and (b) a "more and more" or "less and less" subroutine--e.g., flex the elbow more and more and then less and less as a function of shoulder angle, or cause the flexion velocity of the elbow to become less and less. These are examples of the stimulus-dependent transformations of Section 6. They are subroutines that might be used for many different purposes, so that one might imagine the value of having, for example, a procedure that, when furnished with rate parameters and the names of two variables, would cause the first variable to decrease at the specified rate, as the second variable increased. Whenever people try to guess the functions of parts of the brain that regulate actions, they describe simple kinds of regulations like these, and it would be good to know if organized actions could really be compounded from such ingredients.

An example of the need for this kind of theory is given by the fact that more is known about the microscopic anatomy and electrical circuitry of the cerebellum than of any other part of the brain, yet no one really knows "what it does." It is involved somehow in the programming of skilled movements, and people believe that an important function of the cerebellum is to modify general strategies for action that have been initiated elsewhere, by adjusting timing parameters and gains for smooth and swift execution of the strategy. Guesses as to the role of the known cerebellar circuit elements in performing this function tend to involve only such crude smoothing and modulating activities as resemble the subroutines described

above.¹³ Our "more and more", as a function of the stimulus, might just be one of the things the cerebellum could easily do. One clearly needs more theory just in order to know what it is that one should look for in trying to get beyond these vague statements, and to know what aspects of the limitless set of measurable things in the brain are relevant to these concerns. A mathematical formulation in progress studies how one could build such simple things as adding, averaging, smoothing, and interpolating, into complex actions.

We are interested, then, in clarifying what can be meant by the hierarchical organization often supposed to exist in sensorimotor coordination, and we are interested in tracing things back to reveal their common ingredients. Knowing non-trivial things to say about these problems is a prerequisite to understanding the development, for example, of the "schema of grasping" mentioned at the beginning of this paper. It was there pointed out that many differently performed grasping movements have something in common, and that one feels that the child is somehow learning to take advantage of these common elements. One feels vaguely that a particular highly organized grasping act--say, grasping with the hands--is not merely associated as a whole with some simpler grasping act--like turning the head, seizing the nipple with the lips, stabilizing the head, and sucking,--but that, in some unclear sense, the complex act seems to be built out of ingredients of the simpler one. Once more, theory is needed before one can even know what to ask about such questions, and the purpose of the study growing out of this paper is to

provide some of that theory.

10. Summary of Things a Sensorimotor Theory Should Contain.

Part II has reviewed general sorts of things a theory of sensorimotor behavior should contain:

1. It should identify significant structures, such as abstract mathematical spaces that represent particular aspects of actions (specific examples in preparation), and should study the kinds of functions that map one into another. We saw an example that suggested some kind of representation of auditory space and manual space in visual space, and an example (in preparation) of adaptation of a manual task to a change in parameters will provide an example of an abstract mathematical structure that represents a particular aspect of this manual task.

2. It should take note of the existence of particular transformations of these structures that the nervous system is good at compensating (as in adaptation to optical displacement of images, change in size, weight, or dynamics of objects to be manipulated), or good at producing (to convert one action into a slightly different one). For some purposes, two mappings or two structures will be considered equivalent if they differ by one of these allowable transformations.

3. Actions and their modifications are best described with respect to their effects on the environment, rather than in terms of the configuration of all the degrees of freedom that produce them. Similarly, what we see is best described in terms of external objects, rather than the particular collection of retinal cells that happen to be stimulated at a particular time.

This has the consequence that, as in the example of optical displacements, the mappings and adjustments often action large sets of configurations as a whole, and we must try to understand what these sets are, and how they come to be treated as units. In a more technical example under investigation, it appears that certain equivalence classes are mapped in an orderly way, discernable if one knows enough to look for the classes, while the raw individual data are mapped haphazardly.

4. A theory should deal with the ways in which structures can be pieced together to make more complex structures. Psychologists usually talk rigorously about only one kind of fitting together--the learned association of preexisting structures, an association characterized by individual numerical variables that represent strengths or conditional probabilities. One may see the existence of other problems that one can study about fitting things together; they will be necessary in order to make sense out of the vague notions of building "schemata" of actions incorporating elements of other actions.

5. In order to fit basic actions together into a variety of complex actions, there must be some coordination of the collections of structures representing the basic actions and of the collections of transformations and adjustments that may be applied to the basic actions. The mathematical formulation under investigation treats things that can be said about ways in which spaces of transformations operate on other spaces.

For instance, in the hypothetical description of reaching given in Section 9, there might be a stimulus space in which the stimulus is at

position x . I might have a mapping into an effector space, given by a function associating each stimulus position x with a variable that specifies the way I reach. Next, I might have a space of applicable transformations: instead of reaching in one particular way, I might be told by some other system to reach a little higher when I am starting from a given position; and perhaps I plug in my "more and more" routine, so that I reach to some level when my arm is in one place, a little higher when it is in another place, still higher in a third, and so on. I am therefore interested in situations in which an input space is mapped into an output space, where sometimes the output space is a space of movements, and sometimes a space of transformations that may be applied to some other space of movements or transformations. Just as in the traditional approaches described in Section I.2, it will also be of interest to study input-output functions. It will be the types of things said about them that will be different from the usual.

6. An example of pointing at a target was described as a collection of mappings of regions of a stimulus space into some other spaces, and this description suggested that it might be worth regarding these mappings as some of the pieces that are to be fitted together.

7. Fitting mappings together, as suggested above, was described in the terminology of an atlas of charts. The organization of a model for controlling a device that pointed to targets suggested that one should be interested in the discrepancies between the pieces where they are fitted together, and the transformations that are required to make them

fit together. This will be given concrete content in examples of coordination in a robot, and will be more fully discussed in mathematical formulations in progress.

8. For many reasons, students of sensorimotor behavior often speak of a hierarchical organization of actions, but they cannot say much that is significant about this, and have only begun to know a few things to look for experimentally in order to learn more about it.

An example of a hypothetical hand movement suggested that some kind of hierarchical organization would have the virtues of providing a good "vocabulary" of basic programs for building up more complex actions in a form that might be more easily built into a system than the complex actions; of revealing which actions are "cousins" because of common elements; and of decomposing complex actions into a few stages that might recur as subroutines in many different actions, and might be similar to the things that persons suppose neural circuits to be capable of doing. Any good theory should certainly tell us what non-trivial things we can ask about these hierarchical organizations--these "schemata" of actions--in which one action is built up out of ingredients of others. Since some observers of sensorimotor behavior seem to be studying ways in which one movement might be composed of the same elements as another, only fastened together differently, or else different elements, only fastened together the same way (in some sense), a theoretical formulation should attempt to say something sensible about such possibilities.

11. Value of These General Considerations

The preceding sections presented a collection of vague and general suggestions. Nonetheless, they show a certain area of inquiry that should be entered by any investigator who wishes to study real problems of sensorimotor organization. Many traditions of research have grown up, whose followers assert that all one needs to know is how association strengths develop, or how conditional probabilities change with time, or how conductivities change in networks so as to perform better transformations of inputs. This may be true in one sense: the mechanisms may be composed entirely of conditioned associations, probabilistic learning devices, adaptive networks, and the like. Nonetheless, if the problems listed in this paper exist and are truly important, then the organization of the components, such as associations and networks, must be such as to take the problems into account and deal with them, and we must understand the problems in order to understand the relevant things that the components are doing. This is no more mysterious than the obvious fact that to understand how a television set works, one needs to understand both how the transistors and capacitors work and how oscillator and sweep circuits and feedback systems work; and the concepts and the mathematics are different at the two levels. This is obvious, because there does not happen to be a tradition of asserting that if we would only learn more about how the components work, we would know all there is to know about all possible circuits and devices.

Similarly, given a picture in the form of a mosaic of tiles, every thing in sight is constructed by glueing in tiles, yet there are important picture elements, such as subject, perspective, composition, etc., that cannot even be stated in terms of adjacencies of tiles (the analogue of associations in psychology, or "synaptic" conductivities in network models). That two horizontal lines in the picture must meet at a vanishing point, can be translated into a statement about tiles only by referring, in a particular mosaic representation of the picture, to collections of tiles defined only by the circumstance that they constitute what are intended to depict horizontal lines. One must understand geometrical structures like lines and their properties, in addition to adjacencies and adhesive properties of tiles. Again, two different mosaic realizations of the same picture can be described as similar only in regard to their higher level properties. I should be ashamed to belabor the obvious, were there not this tradition of insisting that we may not know the answers now, but if we only learn more about this or that component process, we surely shall. The present study is devoted to structural problems faced in attempting to synthesize skilled behavior from the elements customarily proposed; this study does not criticize the choice of elements, but attempts to learn more about what it is that these elements must do.

If these arguments carry weight, then any suggestions whatsoever are useful for understanding behavior, if they help an investigator get pointed in the right direction to see significant aspects of behavior; and if one ignores these aspects, one's theory, like the Martian's,

is simply not dealing with the characteristic problems of sensorimotor behavior.

12. Research in the Suggested Directions

The preceding sections might, in a general way, suggest the kinds of programs that should be built into robots, although the real content of the suggestions will come only from consideration of specific examples of tasks. The discussion has told nothing about how the nervous system really works; if successful, it is an exercise for the reader, that will help make him a better person, who, when he looks at the nervous system, may see more. To be successful in this way, each of the topics in this paper will need to be supplemented by examples showing how these ingredients occur in sensorimotor tasks, and by theory that is more than mere formalism or exotic vocabulary, and that tells something practical one can do, now that one understands the examples. This section will briefly indicate some ways in which this program is being followed.

The first task is to obtain examples. Just as this paper looked in more detail at examples illustrating the general attitude developed in the study reviewed in Section I.3, so will it be important to look still more closely at specific instances of the general types of sensorimotor coordination discussed in this paper. This is being done, and some results have been achieved. The first such detailed study is devoted to an example of sensorimotor coordination in an artificial system. This begins by describing relevant features of a particular control system that learns to

provide the optimal value of some parameter required by a given system. An example, chosen arbitrarily from the literature illustrates a simple form of associative learning that serves as a starting point for thinking about the problems studied in this paper. In a generalization of this example, one has a control system that is required to do something that is a function of input variables and a parameter. If the current value of the parameter is slightly different from the value of the parameter when the task was first learned, the system might remember what it used to do when it had its present input, and then introduce a correction to take into account the changed value of the parameter. This procedure is considered from the point of view of the hierarchical organization of basic actions and transformations discussed in the present paper. Finally, a case is considered in which a system has two different forms of feedback of the results of its actions, so that it can either remember what it used to do when one of the monitored variables had its present value, or what it used to do when the other monitored variable had its present value. The corrections for the changed parameter will, in general, be different in the two cases, and a consideration of additional features of hierarchical organization discussed in this paper shows situations in which it would be advantageous to be able to generate a second order correction to transform one of the two first order corrections into the other. An example of this situation has been considered, in which a robot must learn how to coordinate its right hand with its left in order to transfer an object from the left hand to the right hand at any point that both hands can reach. The parameter

value defining standard conditions under which initial learning took place can be the size or weight of the object, or the positions of the supporting members of the body, such as the trunk. The two forms of feedback might be visual monitoring and monitoring by means of feelings in receptors in the joints. The situation that had been discussed abstractly arises in some very natural tasks of becoming more adept at the hand-to-hand transfer by learning to do it by feel instead of by visual feedback. If one assumes that the robot has already learned to adapt to changes of the parameter when he uses visual monitoring, then a close analysis of what the robot needs to do in order to transfer this adaptation to monitoring by feel reveals specific instances of several of the problems reviewed in this paper. These include problems of comparing two different mappings of one space on another--which can now be discussed in specific terms, because the spaces discussed in generalities in this paper can be explicitly defined. Analysis of the task of transfer of adaptation reveals that there are two ways to do this, if one restricts himself to simple conditioning models not requiring extensive memory and computation, and that one way is much easier than the other. The easier way turns out to illustrate several points in the present paper, and, in addition, turns out to be a concrete realization of a previously known abstract study that can be interpreted in terms of the problem of labeling equivalence classes of movements, rather than individual movements (Section II.4).

A discussion of the organization of a mechanism for producing the motions of the arms in the above example reveals specific instances of

the problem of fitting pieces together (Section II.5). A simple example of a system that accomplishes this illustrates the organization of Figure 5 in Section II.8, and the decomposition of movements into a combination of simple subroutines applied to reaching movements. (Section II.9; especially the fifth virtue discussed therein).

An interesting feature of all the structural considerations that arise in these examples is that they cannot be defined solely in terms of the physical structure of the machines, but come into existence only when one considers specific tasks these machines must perform. This is not because the tasks have some mystical Gestalt quality. They are actually quite easy to specify with complete precision. Natural examples of these tasks turn out to exhibit the structural features discussed in this paper, while all that is visible when looking at the machine without considering what it does is a collection of input-output functions of the types considered in many customary theories. I hope to use these examples and more like them to show just how specific tasks generate problems that cannot even be defined in treatments like automata theory and state-variable approaches to systems, that attempt to work in complete generality. These problems do not exist "in general", although the discussion of this paper and the examples that follow from it suggest that they have many common features that might allow a general theoretical treatment.

Work currently in a tentative stage attempts to provide a general mathematical framework for some of these common features. It deals

with ways of fastening together pieces of spaces of actions that are themselves operated upon by spaces of transformation.

All these examples, together with a detailed study of behavioral and physiological evidence regarding problems introduced here, are in the process of further development.¹⁴ They are introduced here in order to show that the points made in this paper form a coherent introduction to real problems that arise when one studies practical problems of sensorimotor coordination, rather than mathematical models, however highly developed, that do not address themselves to elements of structure that are demonstrably present.

REFERENCES

1. Rapoport, A., Verbal discussion at University of Illinois Symposium on Self Organization, June, 1961, in von Foerster, H. and G. W. Zopf, Jr. (Eds.), Principles of Self-Organization, New York, Pergamon Press, 1962, pp. 251-252.
2. Rosenblatt, F., Principles of Neurodynamics, Washington, D.C., Spartan Books, 1962.
3. Tzypkin, Ya. Z., "Optimization, Adaptation, and Learning in Automatic Systems," in Tou. J. T. (Ed.) Computer and Information Sciences--II, New York: Academic Press, 1967.
4. For example, the studies on recognition and manipulation of objects done in Project MAC, Massachusetts Institute of Technology. See, e.g., Guzman-Arenas, A., Some Aspects of Pattern Recognition by Computer MAC-TR-37 (Thesis), Feb., 1967, and Artificial Intelligence Memoranda by M. Minsky and S. Papert on their robot.
5. Greene, P.H., "Models for Perception and Action," Proceedings of the First Annual Princeton Conference on Information Sciences and Systems, March 30-31, 1967, Department of Electrical Engineering, Princeton University, Princeton, New Jersey, 1967, pp. 245-253.
6. Held, R., "Plasticity in Sensory-Motor Systems," Sci. Amer., 213 (1965), pp. 84-94.
7. Hay, J.C. and Pick, H. L., Jr., "Visual and Proprioceptive Adaptation to Optical Displacement of the Visual Stimulus," J. Exp. Psychol., 71 (1966), pp. 150-158.

8. Harris, C.S., "Perceptual Adaptation to Inverted, Reversed, and Displaced Vision," Psychol. Rev., 72 (1965), pp. 419-444.
9. Held, R., Verbal presentation at Conference on Adaptation to Sensory Rearrangement, Department of Psychology, Massachusetts Institute of Technology, November, 1964.
10. Harris, C.S., Adaptation to Displaced Vision: A Proprioceptive Change. Doctoral dissertation, Harvard University. Ann Arbor, Michigan: University Microfilms, 1963, No. 63-8162.
11. Sekuler, R. W. and Bauer, J. A., Jr., "Adaptation to Prismatic Displacements: Hand Position and Target Location," J. Exp. Psychol., 72 (1966), pp. 207-212.
12. Greene, P.H. and Ruggles, T., "CHILD and SPOCK (Computer Having Intelligent Learning and Development; Simulated Procedure for Obtaining Common Knowledge)", IEEE Trans. on Military Electronics, Vol. MIL-7 (1963), pp. 156-159.
13. Arbib, M. A., Franklin, G.F., and Nilsson, N., "Some Ideas on Information Processing in the Cerebellum," Proceedings of the Summer School on Mathematical Models of Neuronic Networks held at Ravello, Italy, June 4-19, 1967. (To appear).
14. Greene, P.H., Quarterly Reports, Institute for Computer Research, University of Chicago, (to appear).

AN ASPECT OF ROBOT CONTROL AND NERVOUS CONTROL OF
SKILLED MOVEMENTS: COORDINATION OF TWO EFFECTORS
AND TRANSFER OF ADAPTATION.

P.H. Greene

1. Introduction

A person can perform the same action in many different ways-- for example, he can write with his arm held high or low or loaded with a weight, or even with a pencil held in his teeth, and although the muscles move differently in each case, the same handwriting always results. A robot with many degrees of freedom of movement would likewise have an infinity of motions that could lead to a single result. How can these vast classes of possible movements and their interrelations be recorded in the memory so that appropriate movements become available when needed for performing useful tasks that have been specified in terms of their effects upon the environment (like picking up a glass), rather than in terms of movements of particular joints and muscles?

Surely, in the above example, a person does not store mappings within his brain between temporal sequences of the fifty-dimensional configurations of all his arm muscles, for each position of the arm in relation to the rest of the body, and for each length and weight of the writing instrument, and temporal sequences of all the dozens of degrees of freedom of his head and neck muscles under a similar variety of circumstances.

A number of studies, including a sustained program in an institute of the USSR Academy of Sciences in Moscow¹, systematically explore a long known principle of guidance of such systems having many degrees of freedom. This is accomplished by a hierarchical organization that makes use of lower dimensional submanifolds in the configuration space of the effectors--i. e., combinations of movements constrained by organization at a lower level to have a small number of independent degrees of freedom. The highest control center selects an appropriate combination that almost fits what it wants to do. Transformations at lower levels shape this combination into the desired action. Not all movement combinations can be performed in this way--people cannot move their arms and legs in arbitrary combinations--but most of the useful combinations can be performed. Were it not for the restrictions on what can be done, nothing would be done at all--if we had to control all the degrees of freedom of all the molecules of a typewriter, this would, in theory, allow the movements of typing, but we would never get any typing done.

There is no adequate theory of such organization. Those mathematical models that have been extensively developed do not address themselves to problems of organization that are demonstrably present. It is important to develop mathematical formulations that are truly descriptive of the tasks of interest. This paper describes a limited aspect of these problems that has some interest in itself, and will prove useful in introducing more general formulations.

2. A Learning Task

The Introduction asked how a control system could store data allowing performance of an action under a variety of conditions, using a variety of sensorimotor systems. Interesting problems arise when we assume the existence of learning systems similar to those in the literature, and require them to perform a task of this nature. It is not a priori clear to everyone that learning theory can deal with the structural complexity of sensorimotor learning, and it is interesting to examine ways in which adaptive behavior can be related to commonly proposed learning mechanisms. The same remarks apply to the presently tenuous relation between the richness of purposive sensorimotor behavior and the electrical interactions observed in neural circuits.

As an example, we may consider an instance of coordination of two effectors, when one must take up where the other leaves off. This gives a concrete example of the problem of fitting structures together. If the task is to transfer an object from the left hand to the right hand, the association of each possible configuration of the left arm system at the time of transfer with a chosen configuration of the right hand system that allows the transfer of the object establishes a mapping between the left-hand and the right-hand spaces. In this example, the mapping is supposed to depend upon a parameter p , and we suppose that the robot wishes to save storage space by learning the association for a standard value $p = p_0$, and correcting for small changes in p . Assuming that he has already learned

to do this, using visual guidance, can he transfer this learning to performance monitored according to feel, e. g., proprioceptive signals from his joints? Moreover, can he learn a second-order correction to transform the visually based correction into the proprioceptively based correction so that he can switch from one form of monitoring to the other during performance? This would also be useful if the second order correction happened to be a simple kind of function of the arm configurations, so that its generation is easier than learning and storing a whole new set of first order correction.

Suppose we have a robot that must move an object it is holding in its left hand, in some specified way, and then must transfer the object to its right hand. To perform this particular movement, the robot must provide its left arm motor with a signal s_L that is some specified function of time t and of some parameter p . This parameter might represent the weight of the object, or the way in which the object is held, or the position of some body member, such as the trunk, that supports the arm, or anything else that influences arm position as a function of the input signal for the motor. The figure shows examples of s_L plotted against t for a few values of p . These curves constitute the particular action, assumed known by the robot, who can produce the variations of s_L that they represent.

At a variable time, chosen by the robot for reasons that do not concern us here, the right hand interrupts the action of the left hand described above by grasping the object, which is released by the left hand,

and by moving the object in some specified way determined by a signal s_R , which is some function, supposed given, of time and various parameters and variables.

One of the variables entering into the definition of s_R will be the initial value of s_R at the time that the hand to hand transfer is made. This will depend, at least, upon the position reached by the left hand at the time of transfer. Call this initial value z . Then to a particular moment of transfer, for a particular value of the parameter p , will correspond a point, such as P in the diagram, and to this point P there corresponds, for a particular right hand movement, an appropriate $z = z(P)$.

Now z is a function of the point P where the left hand motion leaves off. How is P to be recognized by the system that must specify z ? There may be a number of variables that could be used to parametrize the curves and so specify P .

In general, the z system might depend upon sensors that recognized P by some two variables that we shall call u and v . Then the curve corresponding to any given value of p may be parametrized by either of u and v , and points on different curves can be associated in the memory if they have the same u - or v -component of the stimulus to the sensors. Then, in the diagram, P and R , and S and Q are associated according to v , while P and Q , and T and R are associated according to u . Nothing is to imply that one or both v and u might not be identical with one or more of t and s_L , in which case the appropriate broken curves of equal v or u would simply be vertical or horizontal straight lines.

For picturesqueness, v will be thought of as visually sensed position of the hand (i.e., the sensed direction of the eye while looking at the hand), and u will be thought of as kinesthetically sensed position. They will be called the "look" and the "feel" for short. Then, under different loads, or different positions of the supporting members, etc., the same time or the same s_L might correspond to a different position of the arm, hence a different look. This is natural, because if the load is different, the same signal to the motor will cause a different speed, while if the supporting members such as the trunk are facing in a different direction, the hand will arrive in a different place. Likewise, a different feel will result, either because the load is different, or because with a constant load the feeling of force will depend upon the angle the arm makes with the vertical, and this is determined in part by the position of the supporting members. According to this interpretation, P and R look the same, and so do S and Q , while P and Q feel the same, and so do T and R . Although these interpretations will be maintained for vividness and simplicity, and for possible application, it should be kept in mind that u and v need not have these meanings, but may be any two parametrizing variables.

Returning to the idea of storing information about actions performed under standard conditions, together with ways to compensate for small departures from these conditions, let us suppose that the curve for $p = p_0$ represents the standard action, and that $p = p_0 + \Delta p$ represents a small departure. Suppose that the robot already possesses the following learning

and learning capabilities: we suppose that for $p = p_0 + \Delta p$, the z correction mechanism is already conditioned to v , so that the right hand does what it used to do (for $p = p_0$) when it looked the same ($\dot{v} = v(P)$), and adds the correction $\Delta z(R \rightarrow p) = \Delta z)_{v=v(P)}$. This does not necessarily mean that the robot performs the actions and corrections sequentially, i.e., although he might do so. A more interesting possibility is that at the same time the original command is transmitted to the effectors, another signal to the effectors changes their state in such a way that the original command produces a slightly different result, which only for purposes of discussion is regarded as being equal to the original result plus a correction. The decomposition into movements, corrections, and touchup corrections need exist only in the mathematical analysis.

Now, however, the z system must do what it did when it felt the same ($u = u(P)$), add the old correction $\Delta z)_{v=v(P)}$ that was accurate only when the system was doing what it did when it looked the same, and finally, add a touchup correction to allow for the fact that the appropriate total correction is not $z(R \rightarrow P) = \Delta z)_{v=v(P)}$, but $Z(Q \rightarrow P) = \Delta z)_{u=u(P)}$. We assume the robot can associate simultaneous stimuli and responses, and has thus learned an association $v = \phi_p(u)$ for $p = p_0$, this being similar to the visual-kinesthetic mappings studied by Held, etc. Assume that corrections are approximately linear, and hence additive, within the tolerances allowed by the z -system.

3. Motivation for this procedure.

The reasons for achieving the correction in this roundabout way are: (a) One purpose of the example is to study what happens when ordinary associative learning leads the system to do what it used to do when some stimulus component was the same, and one of the things it used to do was to apply the v-based correction. (b) The v-based correction table plus an easily generated touchup correction rule might occupy less memory space, and take less time to learn, than a v-based table plus an equally large u-based table. One reason that the correction rule might be simple is that the difference between the values of a function on two nearby curves is likely to vary much more slowly with position than does the value of the function on one of the curves. Thus, even though the function may have to be represented by a table giving its values at closely spaced points, the differences might vary so slowly as to be representable by a very small number of values at widely spaced points, with at most a simple interpolation rule. (c) Even if all these tables and rules are present, the possibility of mapping v-based actions onto u-based actions is a step toward allowing the synthesis of actions accomplishing the same purpose from whatever sensorimotor elements happen to be most conveniently available. (d) There is experimental reason for believing that changes in the parameters of actions to be performed by animals are reflected by changes in state of spinal motor systems in such a way that the same command coming from higher levels causes an appropriately changed performance of the action.

Such organization allows the higher levels to function like a general who gives an order "Take hill 75!" without having to know all about local conditions, because officers at lower levels will cause this act to be realized in specific ways appropriate to these conditions. Analysis of examples helps us know how to begin looking for the organization of information needed to perform this tuning of lower levels.

The basic motivation of the above is to present an example of a task that a robot might typically have to do, a task that makes use of the very same learning capabilities that everyone assumes, and to show that the performance of this task involves some of the structural problems that are the subject of this investigation, in addition to the laws of the individual mechanisms of learning, etc. The present example presupposes some conventional learning mechanism that learns to pick the correct z for slight variations in the parameter p .

Replacing visual control of movements by proprioceptive control appears to be a commonly occurring feature of the development of skilled actions, and this is an immediate reason for choosing this example. There are also longer range goals concerning the development of correlations between different sensory systems, important in the development of the ability to perceive objects and not just fleeting impressions.

4. Achieving intersensory transfer of adaptation.

In this simple example, it is very easy to look at the diagram to answer the question posed above, namely, what sources of information are available in order to utilize the previous learning in the new learning in the simplest way. This will be a simple example of constructing a coherent picture of all the ways of performing an action under all possible circumstances; and then looking at the "anatomy" of information pathways and causal relationships in this abstract picture, in much the same spirit as an anatomist looks for neuronal pathways in the brain.

Additivity means that, if we write $\Delta z(X \rightarrow Y)$ for the correction produced by the present model, in case it has formed a z table in which the first point, X , corresponds to standard conditions and the second point, Y , to non-standard conditions, then, for any three points X , Y , Z , we have $\Delta z(X \rightarrow Z) = \Delta z(X \rightarrow Y) + \Delta z(Y \rightarrow Z)$, within allowable error. The corrections to be considered here are nothing more than the linear terms one has from differential calculus: if z is a function of a variable x , then Δz , to a linear approximation, is $(dz/dx)\Delta x$.

Our assumption of some standard conditioning model says that the system will do what it did when some stimulus component was the same, and somehow correct for the differences in the other components. This means that all such corrections may be thought of as comparisons of z along paths of constant v , u , or p , for these are the detectable stimulus components, or else as composites of such elementary corrections, provided that the composites are not too complicated.

The last paragraph tells us that the desired correction $\Delta z(Q \rightarrow P)$ must be expressed in the form of a constant- v correction plus a touchup correction either as $\Delta z(Q \rightarrow S) + \Delta z(S \rightarrow P)$ or else as $\Delta z(R \rightarrow P) + \Delta z(Q \rightarrow R)$, since these have the required form, and all other decompositions, such as via U and S , are more complicated, or else, like the path RTP , require just that type of information, e.g., $\Delta z(R \rightarrow T)$ that is not yet learned and that it is the robot's task to learn. Once again, it is important to note that these sequential decompositions are purely conceptual; the robot may do them in a different order, or not at all, in case he realizes their equivalent through tuning of lower levels, as explained above.

All we need do is look at the four terms in these two decompositions to see which way is simpler, and to interpret the result in a form that will be useful for later purposes.

Suppose, first, that the path QSP is attempted. What information is required to be known, and how is it to be used? The first term is $\Delta z(Q \rightarrow S)$. This is a v -based correction, which is in the memory, but is accessible to the correction mechanism only if the control system specifies $v(Q)$, since this correction is $\Delta z_{v=v(Q)}$. Now, $v(Q)$ is not currently being observed by the v -sensor, but, according to our assumption, concerning associative memory, $v(Q)$ has become associated with $u(Q)$, which is being observed by the u -sensor, since $u(Q) = u(P)$, the current value of u . The first term, therefore, requires the following chain of information:

Observe current $u = u(P)$

$$= u(Q) \xrightarrow{\text{(apply } \phi_P)} v(Q) \xrightarrow{\text{(look up } \Delta z)} \Delta z_{v=v(Q)}$$

This requires the composition of two different transformations, together with the programming (which had to be developed somehow) to make the robot perform just this composition. This all seems moderately complicated, not mathematically, but as a task for a simple robot.

The second term in the decomposition QSP is $\Delta z(S \rightarrow P)$: How is this to be expressed in terms of things the robot could notice or remember? Analysis of this is longer, and shows that our simple robot could not find the term. (See Appendix.)

To summarize, it is impractical for the robot to approximate $\Delta z(Q \rightarrow P)$ by means of the path QSP, because the first term, $\Delta z_{v=v(Q)}$ is hard to remember, and the second term is impossible.

Suppose, now, that the path QRP is attempted, i. e., $\Delta z(Q \rightarrow P)$ is approximated by $\Delta z(R \rightarrow P) + \Delta z(Q \rightarrow R)$. Will this be easier than QSP?

The first term, $\Delta z(R \rightarrow P)$ is easily accessible:

Observe current $v = v(P)$

$$= v(R) \xrightarrow{\text{(look up } \Delta z)} \Delta z_{v=v(R)}$$

I. e., the system merely performs the v-correction it used to perform when things looked the way they do. If we insist that the v-correction, as well as the touchup correction, must be entirely u-based, then we do not allow

the robot to look at the current v , but require him to develop associations between u and v along the $p_0 + \Delta p$ curve:

$$\begin{aligned} \text{Observe current } u = u(P) & \xrightarrow{\phi_{p_0 + \Delta p}} \Delta v(P) \\ & = v(R) \xrightarrow{\text{(look up } \Delta z)} \Delta z_{v=v(R)} \end{aligned}$$

This procedure does require $\phi(-, p_0 + \Delta p)$, which was also a difficulty along path SP, but this time the procedure is easier on two counts. First, the argument is $u(P)$, which is currently being observed, rather than the unknown $u(S)$ or $v(S)$. Second, in the previous attempt, the transformation $\phi_{p_0 + \Delta p}^{-1}$ at S had to be applied to the result of an application of ϕ_{p_0} at Q , while in the present case, the transformation $\phi_{p_0 + \Delta p}$ is applied to direct observation. It is easy to imagine ways in which juggling both ϕ and ϕ^{-1} at different points--especially if the same association mechanisms must work for both--would be harder than performing ϕ and using the result to look up Δz .

For the second term, $\Delta z(Q \rightarrow R)$, we have either:

$$\begin{aligned} \Delta z(Q \rightarrow R) &= (dz/du)[u(R) - u(Q)] \\ &= (dz/du)[u(R) - u(P)] \\ &= (dz/du)[\phi_{p_0}^{-1}(v(R)) - u(P)] \\ &= (dz/du)[\phi_{p_0}^{-1}(v(P)) - u(P)] \\ &= (dz/du)[\phi_{p_0}^{-1}(\phi_{p_0 + \Delta p}(u(P)) - u(P)], \end{aligned}$$

or else,

$$\Delta z(Q \rightarrow R) = (dz/dv)[v(R) - v(Q)]$$

$$\begin{aligned}
&= (dz/dv)[v(P)-v(Q)] \\
&= (dz/dv)[v(P) - \phi_{p_0}^{-1}(u(Q))] \\
&= (dz/dv)[v(P) - \phi_{p_0}^{-1}(u(P))] \\
&= (dz/dv)[\phi_{p_0 + \Delta p}^{-1}(u(P)) - \phi_{p_0}^{-1}(u(P))].
\end{aligned}$$

Consider the first possibility. This requires realization of the composite function $\phi_{p_0}^{-1} \phi_{p_0 + \Delta p}$, which we imagine to be hard to achieve, although this time, the argument is actually being observed. In order to avoid this complication, the robot would have to use the expression in the fourth line, which involves the v -sensors, rather than being entirely u -based. For possible application to some other task, for which this might be acceptable, the fourth expression may be given an interpretation in the visual and proprioceptive terms that were used for concreteness when introducing the model. The factor dz/du is merely a gain that need not be calculated, because it can be set at random and raised or lowered if results are consistently too small or too large. This factor must be stored in a table. The term in brackets may be interpreted as the difference between what is currently being felt ($u(P)$) and the way it used to feel, under standard conditions $p=p_0$, when it looked the way it currently looks ($v(P)$). In other words, it is the discrepancy between the observed feeling and the remembered feeling that has come to be associated with the observed appearance. This is one of the discrepancies between mappings found to be significant in mathematical analysis of more general situations.

Now consider the second possibility. The manipulations shown here begin in the v-system, and, at first glance, it appears that their result depends partly upon the v-system, because of the factor dz/dv . If, however, the robot can set this factor at random, and adjust it when necessary, as described above, then it does not depend upon v any more than does dz/du , and there is no problem. The expression in brackets this time does not require composing ϕ with ϕ^{-1} for different values of p, but requires two separate applications entirely u-based, while the other requires the v-sensor as well. In our concrete terminology, the expression represents the discrepancy between what is actually being seen (fourth line), or else the look currently associated with what it actually being felt (fifth line), and the remembered way it used to look, under standard conditions, when it felt the way it now feels.

Reviewing the discussion of the "anatomy" of the causal pathways in this model, we see that if the robot is to use conventional learning mechanisms in the task of realizing the u-based compensation for Δp in the form of the v-based correction plus a touchup correction, then it had better use one of the methods associated with path QRP, rather than QSP, and probably would find that method simplest that starts by translating all observed and remembered feelings into visual coordinates. With all this practice spelled out in all boring detail, it is now easy to see these things directly in the system diagram, but that is the virtue of such practice. General analysis in preparation will reveal the value of having the details explicitly set down.

Observe that all the "anatomy" depended upon specification of a task as well as the underlying system. Different tasks lead to different abstract structures to be "anatomized.*" This shows that there are theoretical problems entirely different from those of understanding the constituent mechanisms of a machine or brain.

The fact that the structures arise in specific tasks shows that there are problems that cannot be handled by general mathematical theories of systems, unless they concern themselves with the specific tasks the systems are required to perform, and the environments in which the system operates. This is not to say that general mathematical features of these problems do not exist.

Examples like the present one are currently being used in a more abstract study of ways of linking systems. To give a taste of the kind of problem considered, we might require the robot above to switch from visual to proprioceptive control in the middle of an action. Will he get the same result if he gets to proprioceptive control via visual plus correction as if he went directly to proprioceptive control? Imagine a much more complicated robot who had more elaborate sets of corrections and transformations of actions, and whose actions were composites of many smaller actions fitted together. Imagine that the visual to proprioceptive

* A pure "physiologist", studying this robot solely at the component level would merely report things like: " s_L is a function $s_L(t, p)$, s_R (initial) is a function of s_L (final), etc." The problems we consider will not occur to him, and he will not have any special reason to look for a system that measures $u(R) - u(Q)$.

transformation is not unique, but that one specific transformation happened to be learned through some chance associations. With all this complication, when could we be sure that the two methods of arriving at proprioception matched, at least well enough that discrepancies could be patched up by some of the transformations in the robot's repertoire? In some mathematical generalizations that appear to capture many of the features of this messy, but realistic situation, the question can be answered in case the spaces that are mapped upon each other are built up of pieces of a specifiable type.

APPENDIX: EXPRESSING $\Delta z(S \rightarrow P)$, THE SECOND TERM
IN THE DECOMPOSITION QSP, IN TERMS OF THINGS
THE ROBOT COULD NOTICE OR REMEMBER.

Any decomposition of this increment by means of a path through a point on the p_0 curve must involve the very unbiased correction the robot needs; hence the correction must be stored, if at all, as a function of variables at P and at S. Then, since $\Delta z(S \rightarrow P)$ may be approximated either by $(dz/dv)(v(P) - v(S))$ or by $(dz/du)(u(P) - u(S))$, and both $v(P)$ and $u(P)$ are currently being observed, we must consider the derivatives and the values at the future position S.

Suppose that the control system already had the second factor in either expression. Then, the derivative, which should really be evaluated at some undetermined point between S and Q, but which according to our assumptions, is roughly constant over the region we are considering, is just some numerical coefficient representing a sensitivity factor or a gain. This coefficient need not be calculated at all, but can be set to a random initial value, and raised a bit whenever Δz turns out to be consistently too small, or lowered a bit whenever Δz turns out to be consistently too large. Thus, this factor is not much of a problem, but it does require storage.

The difficulty lies in predicting $v(S)$ or $u(S)$, for which task we may inspect the following possibilities:

Observe current $v = v(P)$
 $= v(R) \dots\dots$

(This takes us in the wrong direction. Remembering what it did when things looked the same will not help the robot predict the stimuli at S. Moreover, it was assumed that the robot already knew how to produce Δz while tracking the movement with the v-sensors, so that these elaborate methods would be unnecessary. The task is, rather, to track the movement with the u-sensors.)

Observe current $u = u(P)$
 $= u(Q) \xrightarrow{\text{(apply } \phi_p)} v(Q)$
 $= v(S) \xrightarrow{\phi_{p_0 + \Delta p}^{-1}} u(S)$

(This is hopeless, since not only does it involve the composition of two different transformations, but the second of these is not available in the memory, for the memory only contains ϕ for $p = p_0$, not $p = p_0 + \Delta p$, and moreover, the dash stands for either $v(S)$ or $u(S)$, the very things the robot is trying to predict, so that even if we assume $\phi_{p_0 + \Delta p}$ already learned by association the procedure seems impractical.

Note that the difficulty has nothing to do with the fact that finding the values at S constitutes a prediction. Every statement would hold verbatim if the curves of constant v and u happened to be slanted differently, or if the positive direction of the time axis were reversed, so that finding these values constituted a recollection rather than a prediction. In the diagram, as drawn, an analogous case of recollection would occur if the system had the opposite task of transferring u-based adaptation to v-based adaptation, and tried to decompose the correction $\Delta z(R \rightarrow P)$ by using the path RTP. The same difficulties would obtain.

REFERENCES

1. Numerous articles in Biophysics, the English translation of the Russian Biofizika, in the section entitled "Biophysics of complex systems. Mathematical models."
2. Greene, P.H., "Models for Perception and Action," Proceedings of the First Annual Princeton Conference on Information Sciences and Systems, March 30-31, 1967, Department of Electrical Engineering, Princeton University, Princeton, N. J., 1967, pp. 245-253.
3. Greene, P.H., "Seeking Mathematical Models of Skilled Actions," Rock Island Arsenal Biomechanics Symposium, April 5-6, 1967, Expanded version, IOR Quarterly Report No. 16 (part C).

Figure 1. Ways to correct for the change in p .