reprinted from

HUMAN–COMPUTER INTERACTION

INTERACT '90

Proceedings of the IFIP TC 13 Third International Conference on
Human–Computer Interaction
Cambridge, U.K., 27–31 August, 1990

Edited by

D. DIAPER
Department of Computer Science
University of Liverpool
Liverpool, U.K.

D. GILMORE
Department of Psychology
University of Nottingham
Nottingham, U.K.

G. COCKTON
Department of Computing Science
University of Glasgow
Glasgow, U.K.

B. SHACKEL
HUSAT Research Institute
and
Department of Human Sciences
Loughborough University of Technology
Loughborough, U.K.

1990

NORTH-HOLLAND
AMSTERDAM · NEW YORK · OXFORD · TOKYO

Human–Computer Interaction – INTERACT '90
D. Diaper et al. (Editors)
Elsevier Science Publishers B.V. (North-Holland)

Force-to-Motion Functions for Pointing

Joseph D. Rutledge Ted Selker
IBM T.J.Watson Research Center, Yorktown N.Y. 10598
SELKER@ibm.com

A pointing device which can be operated from typing position avoids time loss and distraction. We have built and investigated force-sensitive devices for this purpose. The critical link is the force-to-motion mapping. We have found principles which enable a force joystick to match the function and approach the performance of a mouse in pure pointing tasks, and to best it in mixed tasks, such as editing. Examples take into account task, user strategy and perceptual-motor limitations.

1. INTRODUCTION

Various workers over the past two decades have investigated and compared a variety of analogue devices for use in computer interface pointing tasks [1, 4]. The usual conclusion has been that the mouse has the advantage over alternatives, and the current commercial fashion seems to agree.

We have been intrigued with the 1.5 [2] or so seconds required to make an excursion from the keyboard to the mouse and return; in applications which intermix pointing and typing, this can be significant. Also, the mouse has other inherent disadvantages, especially in environments which provide restricted space or where dangling wires or loose bits of equipment are a hazard.

Our thesis is that it is possible to point efficiently without moving the hands from the normal touch typing home position. This requires locating the pointing device either in the immediate vicinity of the J or F keys (the index finger being rather clearly the finger of choice), or below the space bar, convenient to the thumbs. We first investigated the use of the J or F keys themselves, to serve for both pointing and typing. This requires that the user tell the computer which use is intended. A number of mode switch possibilities are available, but after preliminary experiments we concluded that the cognitive load of making the switch was serious, and shifted attention to a miniature joy-stick, located between the G and H keys in "no-hands land" where it does not interfere with normal typing. This POINTING STICK is the subject of the studies reported here.

The constraints of space in the keyboard eliminate the kind of position-to-position mapping used for the mouse - hence an isometric or force joystick. We could map force applied to the joystick to the velocity of the cursor, to its position, or perhaps to some combination. We report here on the first choice, the conventional rate joystick. The function relating force to velocity is critical to the performance of the Pointing Stick, and leads to the principle results reported here.

The force joystick has a long history of investigation and use [2]. It has been found that pointing times could be expected to be perhaps 20% slower than for a mouse performing the same tasks. Another concern is the "feel" - the subjective impression of exact control of the cursor, and that its movements are the "natural" response to actions.

Many people find pointing with the position of a mouse natural. Can pointing with a rate joystick also feel natural? The rate joystick appears to have an immediate disadvantage here, since the most natural response to a hand motion (for many people) is a movement of proportional magnitude, independent of duration. An analogous discordance will be recalled by anyone who has taken the controls of a light aircraft for the first time - the aircraft responds to a control offset with a rate of change, not with a direct change. As in that case, we find that users very quickly become accustomed to the rate mode of response, and find it natural.

The less tangible aspect of "feel" is the positive control; here the force to motion function is critical. Good "feel" seems to correlate, up to a point, with the more easily measured speed of pointing tasks, especially with small targets.

This paper reports the result of an investigation of a class of force-to-motion functions (***transfer functions***) and their effect on the speed of several experimental pointing tasks for our in-keyboard pointing device, the Pointing Stick.

2. TRANSFER FUNCTIONS

Our exploration of the space of transfer functions began with three families of mathematically simple mappings of force to cursor velocity - linear, parabolic, and a sigmoid parabolic, obtained by reflecting the initial part of the parabola in the point 1/2,1/2 [$(v = f)$, $(v = f^2)$, and $(v = 2 \times f^2, 0 \leq f \leq 1/2;$ $v = 2 \times (1/2 - (1 - f)^2, 1/2 \leq f \leq 1; v = 1, f > 1)$]. Force f and velocity v have scale factors (coefficients), making each of these a 2-parameter family of functions. From experience with these functions, we arrived at the following conjectures:

1. A 'solid' feel, that a point can be held, requires a 'dead band' near zero force, in which the cursor does not move, even if the finger is not perfectly steady.

2. Pointing at small targets requires accurate control of low speed motion - one pixel at a time must be possible. This needs to be done without excess strain in fine motor control, hence the slope of the function at low speed should be low.

3. For long-distance cursor movements, high speed is required. However, we found that when eye-tracking became inaccurate, overall speed was reduced. A high-speed dash off the screen, or to somewhere distant from the target, is counter-productive. In less extreme form, one has the impression of playing golf - a long-distance, partially controlled 'drive', followed by "now where is it - oh, there", then perhaps another, shorter shot, recovery, and finally a low-speed 'putt'. This suggests that a limitation of maximum speed to the eye-tracking limit will be desirable.

4. As a final touch, users like to feel that they can make the cursor dash across the screen almost instantly, and there may be occasions when one wants to reach the opposite edge and start again from there. To accommodate this, we add a steep rise near the top of the force scale. This probably adds little if anything to speed of performance, but it does no damage, and seems to increase acceptance.

Of the simple functions, the sigmoid parabolic seems the most promising, according to the conjectures. This was borne out in informal experiments. However, its behavior near zero was less than 'solid'. The addition of a 'dead slow' plateau suggested itself, following a true dead band. This gives no motion at all for very low force, followed by a region of predictably slow motion somewhat independent of force, then followed by a rapid but smooth acceleration. Similarly, in the upper range, we would like to be able to easily 'cruise' just below the eye-hand-tracking limit, without danger of exceeding it. An upper plateau provides this, reached smoothly from the acceleration regime (Figure 1).

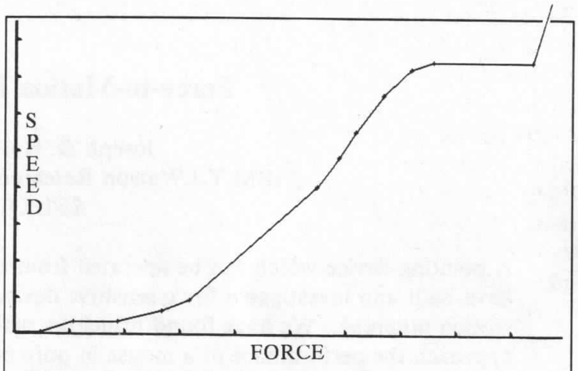

Figure 1. Transfer Function 2Plateau

The ordinate of this graph is force, the abscissa is cursor velocity, in percent of the corresponding scale factors. The velocity scale factor (multiplier of v in the above formulas) is 1500 pixels/second, or on our screen, 66 cm/second. The force scale factor (multiplier of f) was fixed for these experiments at a comfortable value of 225 grams; all sensitivity adjustments were done with the velocity scale.

3. APPARATUS

The Pointing Stick, as used in these experiments, is a steel rod of 2 mm diameter and 2 cm length, mounted on an acrylic base. A section near the base has orthogonal flats to which miniature semiconductor strain gages are bonded (Figure 2).

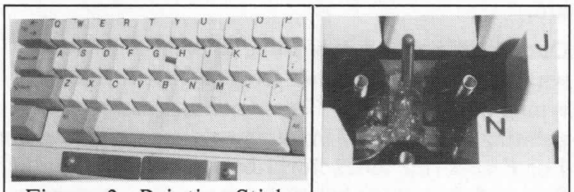

Figure 2. Pointing Stick

The base is glued on the sub-key surface of an IBM PS/2 keyboard, so that the stick protrudes approximately 4 mm above the surface of the keys in their rest position, between the G and H keycaps, which are relieved at their bases to allow space for it. The top is rounded to provide a comfortable fingertip grip. To provide mouse button signals, two microswitches and operating buttons are mounted nearly flush just below the space bar, convenient to the thumbs.

The keyboard was placed about 6 cm from the edge of the desk, allowing subjects to use it as a rest for the heel of the hand. The keyboard retains its normal function as the keyboard of a PS/2 Model 80 computer, which presented and recorded the experiments.

The strain gage outputs of the Pointing Stick are conditioned by an IBM PC/Portable computer, equipped with a LabMaster A/D, D/A, and clock board. The computer makes resistance measurements on the pointing stick gages at 10 millisecond intervals, and emits a set of four pulse trains simulating standard Hawley Mouse signals, for speeds from 2 to about 10,000 pixels/second. Either these signals or signals from a standard mouse feed the PS/2 via an interface box (supplied by Microsoft during 1988-89) converting to serial PS/2 format. The experimental display is an IBM Type 8514 PS/2 color display, displaying 640 pixels horizontal and 350 vertical. Parameters are specified and results given in a coordinate system with 0,0 at screen center and $-1000 \leq X \leq 1000$, $-750 \leq Y \leq 750$, or approximately 0.14 mm per unit. Software in the PC/Portable allows full generality in generating, modifying, and applying transfer functions. The mouse is a Microsoft InPort(tm) Mouse purchased during 1989.

4. EXPERIMENTAL PROCEDURE

Two related experimental procedures were used. In both, Subject is seated before the computer display and keyboard in normal typing position, hands on the keys. Either the Pointing Stick or the mouse may be used; if the mouse, it is located adjacent to the keyboard on the preferred side, on a foam pad at about the level of the top of the keyboard. After signing in and entering experimental parameters, Subject initiates a trial by pressing a key ("t"). At the end of the trial, a score is presented, and the experimenter may choose to commence another trial, present an average score for the most recent group of trials, change experimental parameters, or terminate the experiment. The content of the 'trial' depends on the particular experiment.

1. Target Shooting. Subject selects targets presented as circles of random size and position on the screen. The situation being abstracted here is that of a user engaged in a typing task interspersed with single pointing actions; a pointing action begins and ends with the hands in typing home position. The 'trial' consists of 10 repetitions of the following: a blank screen is presented, with the mouse cursor (arrow) somewhere on it. Subject presses the J key (F if left-handed). The arrow appears at screen center, and a target outline appears at a random position on the screen. Subject moves to the pointing device, brings the arrow to point within the target, and presses a 'mouse-button' (on the mouse if a mouse is in use, the button below the space bar if the Pointing Stick). A hit (splash) or miss (beep) is signaled by the computer. For a hit, the target and splash symbol remain on the screen until Subject returns to the keyboard and presses the J or F key again; for a miss, the screen blanks, ready for the next shot. For each shot, six items are recorded: target position (X,Y), target size, and three times: the time from initial keypress to first pointer movement, to 'hit', and to keyboard return. Misses are generally excluded from the data in analysis. Subject identification, experimental parameters, transfer function in use, date and time, and any other relevant conditions are also recorded in the same file.

The targets are circles of diameter randomly chosen from a uniform distribution between limits specified as an experimental parameter (usually 20 and 100 screen units, corresponding to the range from one character to a representative icon). Targets which extend beyond the screen edges, or are within one diameter of the center, are excluded.

2. Maze Running. A field of targets is presented which requires a sequence of pointings of varying directions and distances. Immediately upon the initiation of a trial, the screen is blanked and a field of numerals is presented, with the arrow in screen center. The object is to select the numerals in numerical order. Initially "1" is highlighted; as soon as it is selected by pressing the appropriate 'mouse' button with the arrow within the highlight, the highlight moves on to "2", and so on. For two-digit numerals, only the first digit is highlighted. Misses (inappropriate button presses) are disrewarded with a brief low-pitched sound, and counted. Each numeral must be successfully selected before the subject can proceed. An *event* begins with one successful selection (or the beginning of the trial) and ends with the next. The duration of each event is recorded. When the last numeral has been selected, the trial ends and the total elapsed time and number of errors are reported.

The same maze is used for a series of runs, so that in place of the random pointings of the other experiment, the maze presents a fixed sequence of pointings which is quickly learned. The targets are of fixed size, and, most important for mouse - Pointing Stick comparisons, the keyboard is not involved at all - this is a pure pointing task.

5. SUBJECTS

Subjects were 6 men ages 22-30 employed as co-op students at the T.J. Watson Research Center. All were experienced and proficient mouse users, but, aside from video game experience, naive to the Pointing Stick or any similar device. Subjects performed the experiments in random order, until scores had settled (no significant difference between first and last 10 trials of a series of at least 40 trials (10 shots or 16 maze events per trial). Subjects reached different levels of proficiency, and comparisons are first within subject; those considered significant are consistent across subjects.

6. ANALYSIS

We can compare performance in our experiments in several ways. The simplest is to average measurements over enough trials for the target distance and size to average out. To see the effect of size or distance we can take a measurement as a function of the parameter in question, averaging out the other. Another option, used by previous authors [1] is to use the Fitts Difficulty Index, $DI = \log_2(D/S + .5)$, to collapse distance and size into a single parameter. We can then fit a line to the resulting point set, either before or after averaging points with similar DI, to obtain a two-parameter characterization and visualization of the data set, with a correlation coefficient to characterize the adequacy of the fit. This gave nice results, with correlation coefficient in the neighborhood of .98 for our larger data sets with averaging over intervals of 0.25 in DI.

For the maze experiment, total times are directly comparable between trials, and can be used as a sensitive measure of performance. To preserve the momentum of a sequence of pointing tasks, errors were tolerated in the maze experiment. In the target shooting experiment, to make all events directly comparable, we followed earlier workers [1] in dropping pointings in which errors occurred. One might question the effect of the different treatment of errors in the two experiments. When events in which errors occurred are eliminated from the maze data, the effect on the overall results is to increase the speed by perhaps 5%, without any qualitative change. Subjects were in part motivated by the scores which they saw at the end of each trial. In the maze the penalty for an error was loss of time, but more time might be lost in waiting to be sure of a hit before pressing the button. In the target experiment, errors did not directly affect the score, and it might be advantageous to deliberately miss a difficult target; we saw no suggestion that this occurred. The error rate was considerably higher in the maze experiment.

7. RESULTS

The velocity scale must be in a reasonable range - a control with a low top speed, or one which jumps uncontrollably at the slightest touch, is clearly unsatisfactory. The exact setting is less obvious. We repeatedly found that our intuition led to excessive sensitivity. The more interesting questions concern the shape of the transfer function, once the scaling is optimized.

In preliminary experiments we selected the following transfer functions for more careful characterization:

- Three linear functions with velocity scale factors respectively 1.5, .75 and .375. These are LIN1a, LIN1b and LIN1c.
- Two parabolic functions with velocity scale factors 1 and 2, called PAR1 and PAR2.
- Our current favorite shown above, 2Plateau. Its velocity scale factor of 1.5 puts the upper plateau of 2Plateau at 1120 pixels or about 50 cm per second.

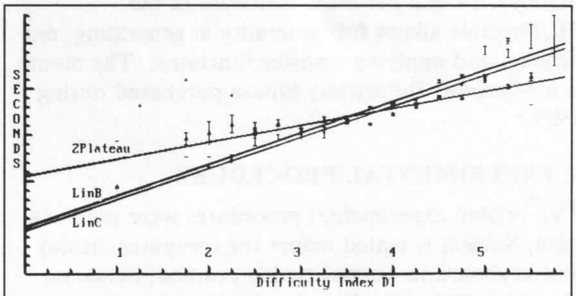

Figure 3. Comparison of 2Plateau with Linear Transfer Functions

Figure 3 is a plot of time against DI for 'target shooting' with 2Plateau, LIN1b, and LIN1c. The linear functions are faster at low difficulty (mainly distance - the range of target sizes in this experiment was 20 to 50 units). The simple numerical average times from keyboard to hit, for example, were 1.61, 1.71, and 1.65 seconds (average distances 645 ±1, sizes 35.5 ±0.5 for all three runs). Excluding points representing targets of size < 35 left the time against DI regression lines for 2Plateau and LIN1c essentially unchanged, but reduced the slope of the LIN1b line from .33 to .23. It appears that despite the small range of target sizes, the effect of size is significant.

The 'maze running' experiment gave a clearer distinction. Average run times and standard deviations in a sequence of runs, for one subject, were:

function	average time	S.D.	trials	slope t/DI
2Plateau	23.9	2.3	20	.30
LIN1b	27.9	2.4	30	.34
LIN1c	29.5	2.9	20	.51
LIN1a	27.8	2.4	30	.31
2Plateau	23.6	1.7	40	.30

LIN1a and LIN1b are not distinguished, but LIN1c differs from them at about 1 sigma, and all from 2Plateau at 2 sigma.

PAR1 and PAR2 gave performance similar to 2Plateau. Subjects reported objectionable fatigue using PAR1 and PAR2. The lower sensitivity could be compensated, but at the cost of physical effort - see discussion below. More sensitive parabolic functions were rejected in early screening as inadequately controllable for fine pointing.

Comparisons with sigmoid parabolic functions gave similar results - no significant differences in speed in either experiment, but noticeable differences in 'feel' and in fatigue effects.

8. POINTING STICK VERSUS MOUSE

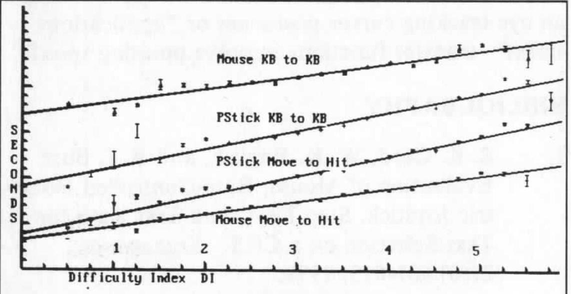

Figure 4. Mouse vs. Pointing Stick

Figure 4 shows the general result. The lower line fits the pointing time for the mouse, in the 'target shooting' experiment, taken from first movement (after 'homing') to selection of the target (hit). The upper line is the same run of the experiment, but timed from keyboard to keyboard (homing times included). The middle pair of lines gives the same information for the pointing stick. The averaged measurements for these runs are as follows:

	mouse	SD	Point	SD
Keyboard to first move	.64	.11	.39	.08
First move to hit	.76	.19	1.18	.35
Hit to keyboard	.72	.12	.09	.13
Keyboard to keyboard	2.12	.26	1.66	.39

Note that the time to reach the Pointing Stick is higher than expected, nearly 2/3 that for the mouse, despite the much shorter distance. The return time for the mouse is much longer than for the Pointing Stick.

The 'maze running' experiment, as a (nearly) pure pointing task, gives results very similar to the central part of the above experiment. The respective time against *DI* regression lines lie close to those for 'first move to hit' for both the mouse and the Pointing Stick. For most (but not all) subjects there was a significant delay between the hit on target n and the first move toward target n + 1, of the order of 0.1 second for the mouse and approaching twice this for the Pointing Stick. Best average times observed for the traversal of a sixteen point maze, starting at screen center, were 15.7 seconds, S.D. 1.8, 60 consecutive runs, for the mouse, and 20.0 seconds, S.D. 1.3, 120 consecutive runs, for the Pointing Stick.

9. DISCUSSION

In comparisons of mouse with Pointing Stick, it must be kept in mind that the subjects were highly experienced mouse users, but novices with the Pointing Stick. Therefore the comparisons can be used only as upper bounds on the differences to be expected in practice. Even so, for an isolated pointing action the Pointing Stick still has an advantage.

We have no firm explanation of the time from keyboard to first movement with the Pointing Stick, or of the difference in hit-to-first-move times in the maze between mouse and Pointing Stick. It is tempting to speculate that about 0.2 seconds is occupied in mental preparation for the move, that this is overlapped with the reaching action in the case of the mouse, and that the relative unfamiliarity of the Pointing Stick accounts for the longer time observed in the maze. The subject who exhibited very short hit-to-move times in the maze was using the LIN1 transfer functions, with slow cursor movement, and was observed to be 'shooting on the fly', never apparently stopping at a target; this strategy was not otherwise observed.

The relatively long return-to-keyboard time for the mouse is consistent with the fact that a key is a smaller target than the mouse.

The comparison of Pointing Stick transfer functions shows a wide range of subject adaptability in using strategies appropriate to the case in hand. For high sensitivity functions they automatically used intermittent contact with the stick, for low sensitivity they maintained contact and (in one case) adopted 'shoot-on-the-fly'. There may in fact be individual differences in optimum transfer function, although we have not observed this. In addition to the observed speed differences between linear and non-linear functions, differences of 'feel' and fatigue were observed, supporting our conjectures that at least two stable speeds, with an appropriate ratio between them, are desirable. The lower plateau of 2Plateau, at 1.5 cm/second, is appropriate for character-sized targets, but a bit fast for pixel targets, which would be needed for a drawing application. While subjects could perform at speed with PAR1 and PAR2, the force required for long fast movements was too much to sustain for more than a few minutes of operation, while more sensitive functions made fine pointing too difficult.

We observed time/*DI* regression line slopes in the range of 0.12 (for the mouse) to .20±.03 for the Pointing Stick with optimal transfer function, and considerably higher with other functions. These contrast with apparently corresponding slopes of about 0.10 found previously [1, 3]. The latter effect is expected, for functions with low maximum speed - time increases linearly with distance, not logarithmically. For other functions, the explanation is presumably deeper, and requires further investigation.

10. CONCLUSIONS

We have been exploring alternative analogue pointing devices for computer interfaces. Laptop computers have no space for a mouse, and space is a problem in many office and other settings as well. The distraction and time of reaching for and returning from a mouse concerns us. We first considered adding sensors to a key under the index finger in a normal keyboard; signaling the use of the key for pointing or typing was distracting. We have placed joysticks in several keyboards and find the Pointing Stick between the G and H keys very useable. In experimenting with analogue pointing devices we have found the Pointing Stick can best the mouse in many situations.

For intermixed pointing and keyboard tasks the Pointing Stick is faster than the mouse. When three or more consecutive pointings occur the mouse can be up to 25% faster than the Pointing Stick. We note also that our Pointing Stick users' pointing speed continues to improve.

Our experience has been that users consistently over estimate their ability to control a fast pointing device. Reducing the rate of change for low speeds as in the parabolic, sigmoid parabolic and 2Plateau (Figure 1) functions increases subjects' speed for selecting small objects. The presence of two plateaus, with the proper ratio between them, makes precise control possible at relatively high sensitivity, greatly improving comfort and reducing fatigue. Adding the high speed tail of 2Plateau made users more comfortable with the Pointing Stick. Before this was added, two users literally bent the Pointing Stick (probably pressing over 5 pounds with their index fingers).

11. FUTURE DIRECTIONS

Following [5, 6] we have informally modeled pointing as a feedback control process, attempting to maintain what we think of as critical damping, which we find to yield the highest speed. A more critical treatment of this area should yield improvements in ease of use and in speed.

Other classes of force-to-motion functions are possible, in particular some degree of force-to-position mapping. Pure force-to-position mapping seems infeasible, but some mixed strategy, perhaps force-to-position locally with force-to-velocity at greater distances, should be worth investigating.

We informally measured how fast a subject could run our maze with his eyes; this was about 12 seconds or 3 seconds faster than the fastest pointing measured. Could this 25% speed difference be bridged? Could an eye tracking cursor positioner or "applications smart" transfer functions improve pointing speed?

BIBLIOGRAPHY

1. S. K. Card, W. K. English, and B. J. Burr. Evaluation of Mouse, Rate-Controlled Isometric Joystick, Step Keys, and Test Keys for Text Selection on a CRT. *Ergonomics*, 21(8):601-613, 1978..

2. Stuart K. Card, Thomas P. Moran, and Allen Newell. *Psychology of Human Computer Interaction*. Lawrence Erlbaum Associates, 1983.

3. P. M. Fitts. The Information Capacity of the Human Motor System In Controlling The Amplitude of Movement.. *Journal of Experimental Psychology*, 47:381-391, 1954..

4. P. M. Fitts and J. R. Peterson. Information Capacity of Discrete Motor Responses. *Journal of Experimental Psychology*, 67(2):103-112, 1964..

5. T. O. Kvalseth. Information Capacity of Two-Dimensional Human Motor Responses. *Ergonomics*, 24(7):573-575, 1981..

6. D. E. Meyer, S. Kornblum, R. A. Abrams, C. E. Wright, and J. E. K. Smith. Optimality in Human Motor-Performance - Ideal Control of Rapid Aimed Movements. *Psychological Review*, 95(3):340-370, 1988..