

Welcome to cs303

Scott Klemmer, Phil Levis, Chris Manning
TA: Katherine Breeden

Wednesdays 3:15p-6:00p
<http://cs303.stanford.edu>

Rare feature: this course has three instructors. Scott's favorite course in graduate school was the grad OS course, b/c it was co-taught by Brewer & Hellerstein. Each asked the other good questions and offered another perspective on the material.

Today's three parts

- 1) An experiment -- you're the participants
- 2) We'll analyze the data, and introduce R along the way
- 3) We'll discuss our findings, and use that to introduce experimental design

3:15 sharp - 2 minutes: introduce high-level idea of course (phil, want to do this, as you're the OG instigator?); 2 -minutes: introduce staff (in essence, we each say "Hi my name is"); 1 minute: I'll explain goals for the day. (don't worry, we'll get more time to cover pragmatics later. and, as katherine points out, we can)

3:20 - Everyone gets a consent form and I explain the experiment

3:30 - Run through as many as possible.

4:00 Issues in Experimental Design.

4:30 break

4:45 Analyze the data. Introduce Fitts Law.

5:15 Phil introduces wireless

5:30 Explain assignment for next week & field questions about the class. (maybe anticipate a few and preemptively answer them)

5:45 done

EXAMPLE: POINTING DEVICES



Mouse. Engelbart and English

The Experiment

- Click on the blue circle
- Keep doing it till the system says done

Rotation

- If you're at an input device, use it until done
- If you're not, begin reading the paper.
- The readers are the queue.
- We'll rotate through.
- Do as many devices as you can until 4pm

Go!

Teaching staff can be readers.

Let's aggregate the data

- Via Sneakernet
- Katherine will put it into R, using the rbind function (which stacks all the data files)

What did we just do?

- Independent variables (aka Factors)
i.e., things we manipulated
- Dependent variables (aka Response Variables)
i.e., things we measured
- Control variables
i.e., things we held constant
- Random variables
i.e., things we allowed to vary

Kinds of validity

- Internal Validity: Manipulating independent variable causally explains change in dependent variable
- External/Ecological Validity: Generalizes to Real-World Situations
 - to people other than the participants
 - to “real” situations (is it contrived?)

What else might we have measured?

- Time on Task -- How long does it take people to complete basic tasks? (For example, find something to buy, create a new account, and order the item.)
- Accuracy -- How many mistakes did people make? (And were they fatal or recoverable with the right information?)
- Recall -- How much does the person remember afterwards or after periods of non-use?
- Emotional Response -- How does the person feel about the tasks completed? (Confident? Stressed? Would the user recommend this system to a friend?)

Why the repeated measures?

- Model the motor system reaction time...
- ...by factoring out the perception time

The importance of random assignment

- For example, say you wanted to find out whether people are faster at input in the morning or afternoon. You allowed people to come in whenever they wanted. What if people who have a preference for participating in the morning -- morning people -- are faster than people who have a preference for participating in the afternoon. You'll find that morning was faster. But the causal reason was the population difference, not the experimental manipulation. This confound is why a lot of economics is so hard -- it's computing correlations, but there's no manipulation. (That leads off to a longer story and more nuanced discussion, but that's not today's topic.) Random assignment is tool #1 in establishing causation.

Counterbalanced Assignment

- Say you worry that running speed will affect input speed (that fast runners are physically fast people, and they'll do better)
- You can assign participants so that running speed is balanced across conditions
- Do so with care: form matched pairs using a pre-test (work this out on the board)

A danger: regression

- Let's find heady coins
- First, let's flip all the coins (our pre-test)
- If they land heads more than half, we'll call them heady
- Now let's feed them a snack
- After a snack, do heady coins beat taily coins?
- Similarly, some of our runners will have had a relatively good run, and others a relatively bad one. They'll exhibit regression to the mean if re-tested.
- If the pre-test is used to counterbalance, and assignment is random, then the error goes away

Counterbalanced Ordering

- If one task is always first, learning is a confound.

Between vs. within subjects

- Within subjects
 - All participants try all conditions
 - + Can isolate effect of individual differences
 - + Requires fewer participants
 - - Ordering and fatigue effects
- Between subjects
 - Each participant tries one condition
 - + No ordering effects, less fatigue.
 - - Cannot isolate effects due to individual differences.
 - - Need more participants

Choosing Participants

- Representative of target users
 - job-specific vocab / knowledge
 - tasks
- Approximate if needed
 - system intended for doctors
 - get medical students
 - system intended for engineers
 - get engineering students
- Use incentives to get participants



A real potential pitfall



Source: PC World

Issues

- user sample
- statistical significance
- “newbie” effect / learning effects

If you read a bit more carefully into the study, you'll notice that the study is about initial adoption of the iPhone keyboard compared to users' current phones. Also, it isn't a survey, it was a study with one on one interviews where users typed and were timed.

The multitap (Non-QWERTY) users did the same or better with the iPhone than their current method, which suggests that multitappers may have an easier time adopting the iPhone's keyboard than QWERTY users. Which to me is interesting.

The study does not at any time attempt to say that QWERTY users will be twice as slow on the iPhone for as long as they use the iPhone, but it does say they may have more difficulty than multitap users initially. Which to me is interesting.

It would be interesting to see in a study some expert iPhone texters and have them switch to a QWERTY phone to see if there is a similar difference in typing efficiency.

The Hawthorne Effect



In the 1930's some studies were held at the Western Electric production facility outside Chicago in a place called Hawthorne. The intent of the study was simple enough: invite a handful of employees to participate in various working condition tests to determine which conditions were most conducive to increased production. Those conditions that "tested" best were then to be rolled out to the general production floor. One of things they tested was brighter lights. Production went up. Then they tested dimmer lights. Production went up. In fact, no matter what they tested, production went up!

“By singling out a small group of employees to participate in an exclusive trial, participants felt valued, special and important. The special attention they received gratified their ego and created a positive emotional bond with what they were trialing. The practical upshot was that the research trials effectively transformed the research participants into advocates for whatever it was they were trialing.”

Let's Open R

- Open a laptop, and follow along.

First, a very brief introduction to R

- `x <- c(1,2,4)`
- `x [1] 1 2 4`
- `x[3] [1] 4`
- `q <- c(x,x,8)`
- `mean(x) [1] 2.333333`
- `sd(x) [1] 1.527525`
- `y <- mean(x)`
- `y`
- `?datasets`
- `mean(Nile) [1] 919.35`
- `sd(Nile) [1] 169.2275`
- `hist(Nile)`
- `q()`

The web tutorials are great.

Let's Look at the Data

- `testdata<-read.csv("~/Dropbox/cs303/Intro/s42_2D_nomet__49844285.txt", header=T, as.is=T)`
- `names(testdata)`
- `plot(testdata[,13],testdata[,57])`
- `plot(testdata[, "A"],testdata[, "Duration"])`
- `color <- rep("black",length(testdata[, "A"]))`
- `color[testdata[, "A"]==256] <- "green"`
- `color[testdata[, "A"]==384] <- "blue"`
- `plot(testdata[, "W"],testdata[, "Duration"],col=color)`
- `plot(testdata[, "W"]/testdata[, "A"],testdata[, "Duration"],col=color)`
- `attach(testdata)`
- `plot(Duration ~ log(A/W), col=color)`
- `plot(Duration ~ log(A/W), col=color, xlab= "distance/size", ylab="milliseconds")`
- `abline(lm(Duration~log(A/W)))`

- Fitts' Law

- Time T_{pos} to move the hand to target size S which is distance D away is given by:

- $T_{pos} = a + b \log_2 (\text{Distance}/\text{Size} + 1)$
 - The log part is the “index of difficulty” of the target; it's units are bits

- summary

- time to move the hand depends only on the ***relative precision*** required


Try to hit a target without looking

- You can open your eyes after each step
- Then, try it for both a mac-style and windows-style menu bar

It models iterative, ballistic motion

- I have this slide in the input deck


What does Fitts' law really model?



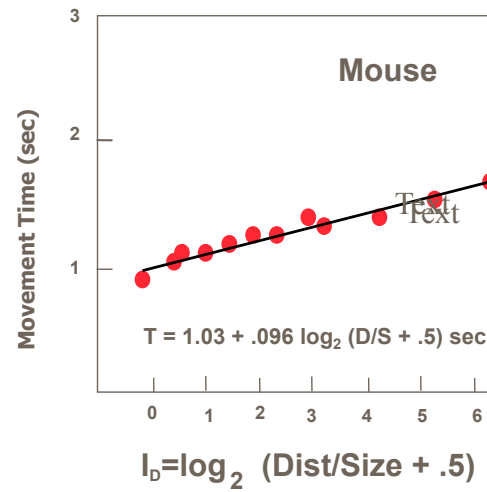
It was inspired by information theory

- It treats acquiring a target as specifying a number of bits
- i.e., in the Fitts' worldview, the human motor system is a noisy information channel
- Smaller target? More bits
- Further target? More bits

EXPERIMENT: MICE ARE FASTEST



WHY?



Why these results?

Time to position mouse
proportional to Fitts'
Index of Difficulty I_D .

Proportionality constant
= 10 bits/sec, same as
hand.

Therefore speed limit is
in the eye-hand system,
not the mouse.

Therefore, mouse is a
near optimal device.

50 years of data

Device	Study	IP (bits/s)
Hand	Fitts (1954)	10.6
Mouse	Card, English, & Burr (1978)	10.4
Joystick	Card, English, & Burr (1978)	5.0
Trackball	Epps (1986)	2.9
Touchpad	Epps (1986)	1.6
Eyetracker	Ware & Mikaelian (1987)	13.7

Reference:


Mackenzie, I. Fitts' Law as a research and design tool in human computer interaction. Human Computer Interaction, 1992, Vol. 7, pp. 91-139

Fitts' Law Example

Pop-up Linear Menu



Pop-up Pie Menu



- Which will be faster on average?
 - pie menu (bigger targets & less distance)


Source: Landay, James. "Human Abilities". CS160 UC Berkeley.

EXAMPLE: ALTERNATIVE DEVICES



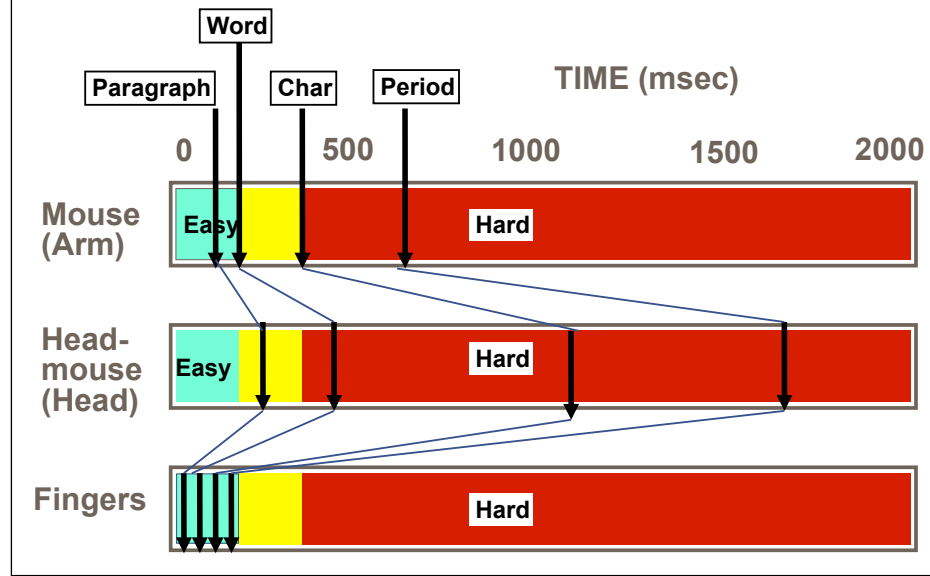
Headmouse: No chance to win

ATTACHING POINTING DEVICE




Use transducer on
high bandwidth
muscles

EXAMPLE: STRUCTURING THE TASK SPACE BY PROJECTING THE MODEL





EXAMPLE: BEATING THE MOUSE



Use transducer on
high bandwidth
muscles

