Agents and Interaction

Jim Warren
Professor of Health Informatics

Outline

- Interacting as an 'agent'
- User modelling
- Changing user behaviour
- Robots
- Information retrieval as an interaction paradigm
- The Social Network

Agents v. Direct Manipulation

- Point-and-click interfaces treat systems like tools
  - Like I would pick up a hammer and a nail
  - I don't have a conversation with the hammer – I directly manipulate it
- But another approach is to treat the system like an 'agent'
  - This is always present to some degree when using a command language, especially to the degree it resembles natural language
- It's just a 'stance' – a *mental model* – held by the user (and/or held and encouraged by the system designer)
  - But it can provide guidance as an interaction paradigm and lead to any of a host of compelling metaphors

Eliza

- Back in 1966, Weizenbaum made a simple dialogue agent that mimicked a Rogerian psychotherapist

HUMAN: You are like my father in some ways.
ELIZA: What resemblance do you see?
HUMAN: You are not very aggressive but I think you don't want me to notice that.
ELIZA: What makes you think I am not very aggressive?
HUMAN: You don't argue with me.
ELIZA: Why do you think I don't argue with you?
HUMAN: You are afraid of me.
ELIZA: Does it please you to believe I am afraid of you?
HUMAN: My father is afraid of everybody.
ELIZA: What else comes to mind when you think of your father?
HUMAN: Bullies.

My Eliza conversation

http://www.masswerk.at/elizabot/

ELIZA: Is something troubling you ?
YOU: I'm anxious
ELIZA: Is it because you are anxious that you came to me ?
YOU: Yes
ELIZA: I see.
YOU: I have some tough deadlines at work
ELIZA: That's quite interesting.
YOU: Yeah, but so much is required so soon
ELIZA: That is interesting. Please continue.
YOU: And I'm just concerned about getting it done in time
ELIZA: Do you know anyone else who is just concerned about getting it done in time ?
YOU: What?
ELIZA: Have you asked such questions before ?

Weizenbaum chose this domain because it was a particular case where no knowledge of the world was required. In his words, if I said…

> "I went for a long boat ride" and he responded "Tell me about boats", one would not assume that he knew nothing about boats, but that he had some purpose in so directing the subsequent conversation

User Model and 'Agency'

- What we need to make the dialogue better – more convincing, more helpful – is more explicit knowledge in the system
- Any knowledge the system stores about the user can be termed its *user model*
  - In particular, knowledge of user goals helps the system to function as an *agent*
- IBM definition:
  - *Intelligent agents* are software entities that carry out some set of operations on behalf of a user or another program with some degree of independence or autonomy, and in doing so, employ some knowledge or representations of the user's goals or desires
- Often a software component is only considered an 'agent' if it has a degree of autonomy in undertaking an action for the user

Trade-off for guessing

- From Eric Horvitz (1999)
  - Utility, u, of Action, A, or NOT Action A when the user does or does not have goal G, with probability of Goal given the Evidence, P(G|E)
    - Agent should do A when P(G|E) exceeds P*

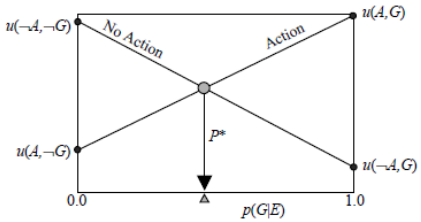

LookOut (Outlook variant, ha ha)

- Here the agent assumes Eric may well want to meet Mark (perhaps because Paul suggested it) and takes the Action to pop up the calendar

Anthropomorphic agent with speech recognition to play into agent metaphor

(Actually there's more going on in this example – Genie might've offered a specific appointment time, but it couldn't find one)

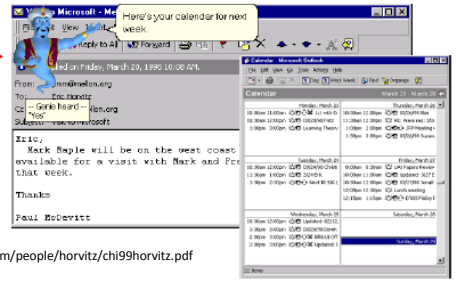

http://research.microsoft.com/en-us/um/people/horvitz/chi99horvitz.pdf

Uses of user models

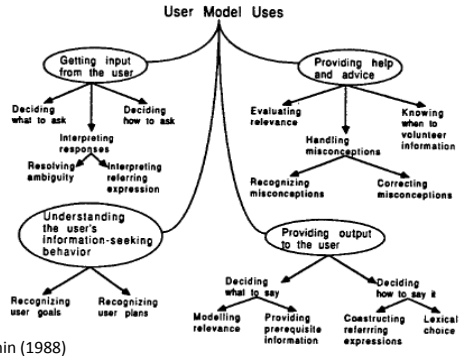

- Kass & Finin (1988)
  - http://www.csee.umbc.edu/~finin/papers/kass88.pdf

Dimensions of user models

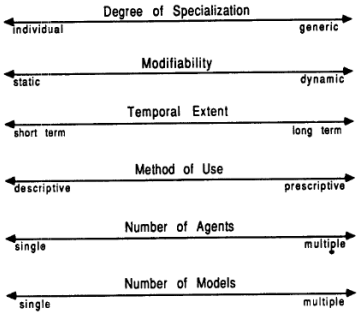

- Kass & Finin again

A less specialised user model might use 'stereotypes' such as 'novice' and 'expert' user

A short term model is needed just to track the immediate dialogue; longer-term models might track user characteristics between sessions

Changing user behaviour

- Rather than having strict 'agent' perspective (serving user goals), we might want to change the user's goals and behaviour
  - This could be our approach if we're implementing a component that will try to sell the user our product
  - Most of the user modelling and utility-of-action concepts apply just the same
- On the benevolent side, our goal might be a population health goal
  - Or an expressed long-term goal of a the user that they communicated by signing up to our agent's service (e.g., "I want to quit smoking – help me do it!"

Health Belief Model

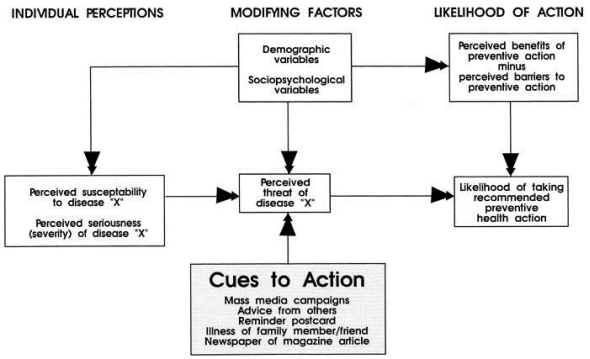

- http://www.nature.com/bdj/journal/v186/n9/fig_tab/4800135a_F1.html

Stages of Change (or 'Transtheoretical') Model

- We can estimate placement of the user on the model and adjust our actions (system outputs) accordingly to advance their progress

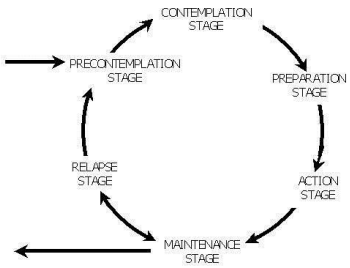

http://addictions.about.com/od/addictiontreatment/ss/The-Stages-Of-Change-Model-In-Addiction-Treatment.htm

STOMP (STop smoking Over Mobile Phone)

- Up to 480 customized text messages over the twenty-six week program duration

HME STOMP				
THE PROGRAM	STAGE	PERIOD	MESSAGE RATE	MESSAGE TYPE
	Pre-Quit	14 – 1 days prior to Quitting	1 – 2 per day	Cessation
	Quit Day	1 day	3 on day	Cessation
	Intensive	Quit Day – 4 wks	3 per day	Cessation
	Maintenance	Week 5 – End	1 every 3 days	Cessation
RELAPSE	Relapse Early or Late	4 weeks – After Quit Day only	3 per day	Relapse
CRAVE & SLIP UP	Anytime	50 Anytime	n/a	Crave Slip Up

STOMP (STop smoking Over Mobile Phone)

- Personalized Cessation Support – text message content tailored to the target participants
- Quit Tips – consistent and helpful text messages reminding the participant of the overall goal to quit smoking
- Culturally Relevant Messages – text messages tailored for specific cultural and language requirements
- Smoking Facts – general fact text messages that help reinforce smoking cessation
- Craving & Slip Up Support – responsive text message content for participants craving a cigarette or those who have smoked a cigarette
- Polling –participants can text their answers to questions posed by providers, and then view results.
- Message Blackouts –participants can designate one specific period per day during which STOMP will not send them messages
- Relapse Program – a 4-week intensive program which participants can enroll in if they started smoking again, but still wish to quit

So many user modelling aspects! http://www.hsaglobal.net/STOMP
http://www.quit.org.nz/39/help-to-quit/tools-to-help-you-quit-smoking#txt2quit
http://journal.nzma.org.nz/journal/118-1216/1494/

Telephone-Linked Care (TLC)

- A host of health promotion interventions have been developed under TLC from Boston University
  - Computer-managed phone calls
  - Uses a stored voice output read by an actor
  - Accepts simple voice input (Yes/No) or uses number pad ("Press 1 if Yes…")
- They map out the entire intervention
  - Identification of target demographic
  - Choice of psychological strategy
  - Logical flow of each call
  - Text options for each specific node

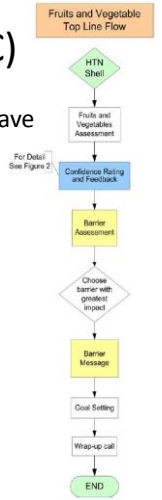

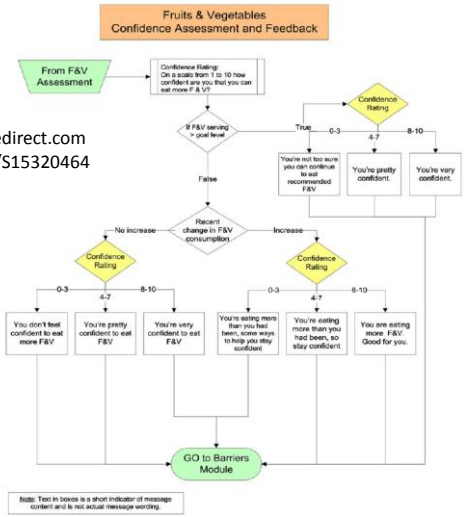

- http://www.sciencedirect.com/science/article/pii/S1532046406000347

How 'bout using a ROBOT?!

- TLC, and even STOMP, are actually very anthropomorphic
  - Txt'ing is something we do with real people
  - The TLC actor voice can engender engagement: guilt and even love
- But using a robot makes the anthropomorphic presence spatially tangible
- 'Cafero' waiter robot with clinical monitoring tools on the tray
  - Linux based navigation system on bottom
  - Windows touchscreen and voice interface up top

http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3243150/

Application / Study

- Elder care
  - Testing in a residential care facility (supported living: periodic caregiver visits, nurse on call)
  - Promoting quality use of medicine
    - Adherence to taking it (or knowing why not)
    - Physiological monitoring of effectiveness (and for safety)
    - Asking about side-effects
    - Providing education (and entertainment)
- Tested with morning medications of 12 residents
  (there's since been a long-term study, and then a larger trial, but this was an important iteration)
- Research ethics
  - Human research ethics committee approval of protocol; approval by aged care facility; telephone recruitment; individual signed informed consent of residents
  - Balance of risks and benefits: could cause people to double-dose, but there are a lot of medication administration errors in elderly with present workflows

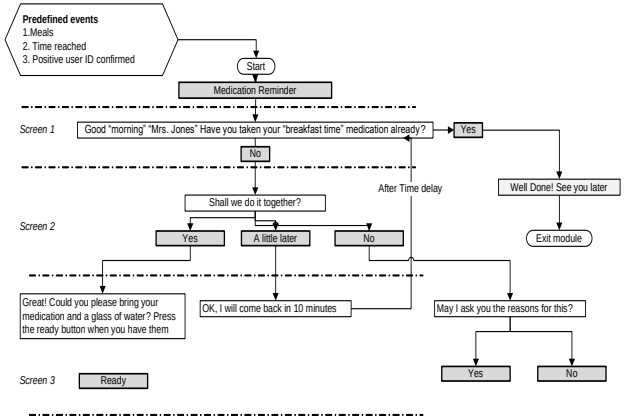

- Critical to design an *empowering* dialogue – not "You must do this", but "Shall we do this?" and with real options to say 'no' or 'not right now'
- Potential to learn a lot from the dialogue (e.g. patient refuses to take medication because it's meant to be taken with food, but they've been vomiting)

Measures / findings

- Video recorded
- Interviewed
  - Structured, open-ended
- Needed to tilt head lower!
- Patients like it and can use it well enough unless having significant dementia or macular degeneration
- Want features to video call and alert family

Information Retrieval (IR)

- IR is an interaction paradigm (and an extremely popular one!)
  - You ask a question and get an answer
  - On the Web this usually is done as you input a *search query* and get a list of Web sites the search engine feels will be relevant to your interest
  - The more the search engine tailors results to you, the more it's using its user model
    - E.g., it sees you're in .nz domain, so if you want a service like "pizza", it'll offer mostly .nz pizza sites (figuring you didn't want to import the pizza!)
- This fits the notion of agency
  - You trust the search engine (via what it's learned with its Web crawler, what it knows about you [and the world!], its algorithm to match your term to a page, and its algorithm to rank results) to present to you 'the answer'

Google PageRank

- Pages are ranked on the basis of the probability that a 'random surfer' will end up there

$$PR(A) = \frac{1-d}{N} + d\left(\frac{PR(B)}{L(B)} + \frac{PR(C)}{L(C)} + \frac{PR(D)}{L(D)} + \cdots\right)$$

- PR(A), PageRank of A = probability of surfer ending up on A, is sum of chance of surfer stopping on A (based on damping factor, d, and Web size, N) plus PageRank of all pages with link to A (B, C, D…) divided by the number of pages each page links to (L(B), L(C), L(D)…)

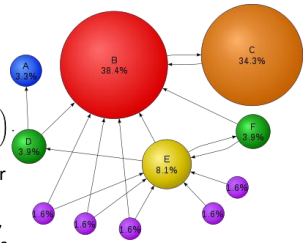

Page C doesn't have a lot of incoming links (well, just 1), but it's incoming link is the sole outgoing link of Page B, which is a very popular page

- Traditionally PageRank was crucial in Google ranking (order of presentation) of results with content matching your search query
- Of course there's been a lot of gaming; now mainly 'Panda' process based on similarity of page to other pages that human quality testers rated to be of high quality

IBM Watson – actually answers the question!

The Social Network

- Never mind having an AI answer my question
  - I want to know what actual people say the answer is!
- Functionally similar to IR over the Web
  - After all, people wrote the Web pages, so you were already searching for what people think…
  - But somewhat different as a user interface metaphor
    - A Google retrieval is *page* focused, the Social Network (FaceBook, Twitter, etc.) is *people* (or *user*) focused, or report-focused (Amazon customer reviews, Tripadvisor hotel reviews)
  - When there are faces next to the entries, you are emphasizing the Social Network metaphor

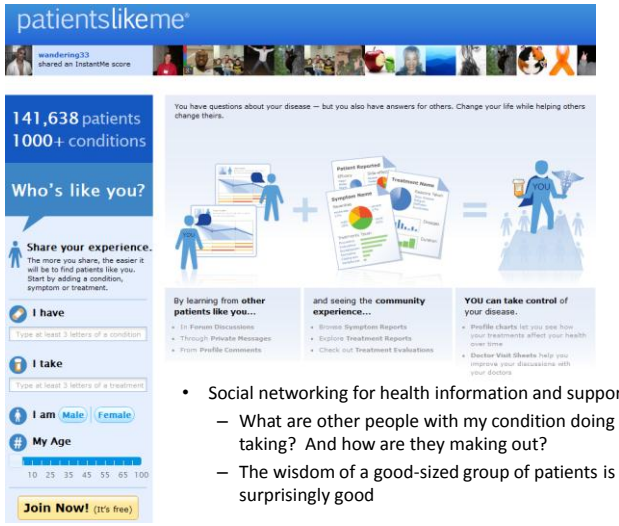

- Social networking for health information and support
  - What are other people with my condition doing / taking? And how are they making out?
  - The wisdom of a good-sized group of patients is surprisingly good

Healthy behaviour change based on groups and competing groups

Team members build relationships and provide support via team blog

Competition between teams to achieve best outcome (waka race)

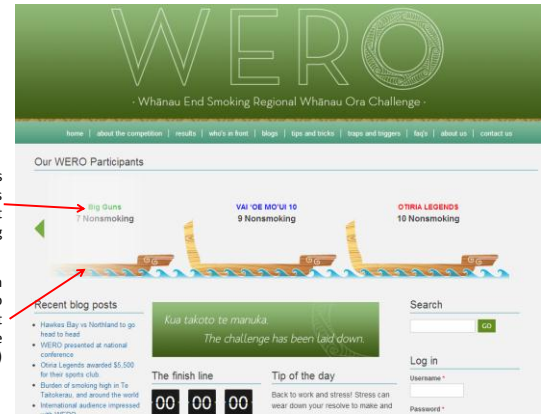

Conclusion

- Systems can be designed to interact with us as if they were a person (anthropomorphically, as 'agents')
- This doesn't have to be visual
  - Easier to make it work as txt or chat, or constrained stored voice messages and prompts
- User models are data used by the system to tailor its responses
  - These guide system's choice of action, and can be used to influence user behaviour
- Information retrieval (e.g. Google) has elements of agent based interaction and of Social Network (e.g. PageRank – lot's of people put in links to this page: it must be good)
- Social Network interaction is popular
  - Puts the system focus back on connecting us to other real people to search out their guidance (opinion, experience, emotional support)