

# Discussion #1

Student-led discussions

Paper 1:

“Direct Manipulation vs Interface Agents.”

# Discussion Leader

# The Debate!

## **Ben Schneiderman**

---

People should be able to directly manipulate nearly everything

Humans can make judgements machines cannot, even with large amounts of data

Don't underestimate people. Embrace and prioritize human capabilities with direct manipulation

## **Pattie Maes**

---

Some elements of interaction should be automated

Internet is huge and constantly changing. Will need to delegate some information processing tasks

Novices use these systems too. Even with professionals, can be useful to delegate certain tasks (e.g. professor, car repair)

---

Some degree of delegation, but still direct manipulation human interfaces  
Should depend on needs and wants of the user

# Student discussion points

## Overall reflections

- Debate is out of date, but interesting to see explicitly discussed
- Differing degrees of disagreeing with Ben
  - His examples were on a smaller scale

## Debate style

- Engaging, but sometimes frustrating (e.g. “The New Pattie”)

## The User

- Professional users could also use agents (e.g. Github CoPilot)
- People having different preferences on privacy – build flexibility within a system

# Bringing debate to other fields

Shopping on the internet

Youtube – recommender is intelligent, interaction with individual videos is direct

Medical field – risks vs benefits of patients rather than internet users

High-stakes AI applications should have direct manipulation

ChatGPT delegating more complex tasks

Information visualization

Rotating an image

# Student discussion points

- Current
  - We currently anthropomorphize agents. Should we?
  - AI capabilities are changing. Instead of just inputting and sorting information, now outputting information
  - We're ok with some degree of error, esp with generative models.
- Undiscussed social/other impacts
  - Monetization
  - Collecting of user private data
  - Social engineering
  - Legal liability of failure

# Discussion questions!

- Where do **generative models** stand in the direct manipulation vs automated debate?
- How should / have we built **trust** in automated systems? When should we reduce that trust?
- How does this debate fit into **your own research**?
- Are there other **social / external impacts** of this debate?
- How has **anthropomorphization** changed over the years? What might the effects be?



# Scientific Peer Reviewer (Advocate)

# NeurIPS-Style Review

- **Disclaimer:** I've never done one of these before!
- **Important Factors:**
  - a. Social vs **Paper**
    - i. **Social:** Lax guidelines, proposal for a debate.
    - ii. **Paper:** More clear guidelines, submitted as is.
  - b. 2023 vs. **1997**

# Main Idea

Frames the **merits and tradeoffs** of direct manipulation and interface agents **as a debate** between Ben Shneiderman and Pattie Maes:

- “Should users give up complete **control** of their interaction with interfaces?”
- “Will users want to **risk** depending on ‘agents’ that learn their likes and dislikes and **act on a user’s behalf**”

# Useful Points of View

## Ben Shneiderman

- User Control
  - Sense of Completion
- Human analysis
  - Faster + more accurate
- Reactive

## Pattie Maes

- Automation / Personalization
  - System / data complexity
- Machine Analysis
  - Automated + Concurrent
- Proactive

# Strengths and Weaknesses

## Strengths:

- **Debate framing:**
  - Survey-style review
  - Emphasizes tradeoffs between the approaches
- **Test of time:**
  - Data privacy vs. personalization
- Human interface in complex systems
- Interesting and relevant figures

## Weaknesses:

- **Debate framing:**
  - Rehashes prior work
  - Are the viewpoints really at odds with one another?
- **Test of time:**
  - Data / ML revolution

# Other Points

- **Clarity:**
  - Informal + Unusual tone
  - Makes points of view clear
- **Relation to Prior Work:**
  - Summarizes and frames previous work
- **Correctness & Reproducibility:** N/A

# Scientific Peer Reviewer (Skeptic)

# Overview

Main Contribution:

- Documenting the debate between Ben Schneiderman and Pattie Maes on ***direct manipulation*** vs. ***interface agents***
- Schneiderman is an advocator of **direct manipulation**, with a heavy emphasis on full **user control** and **predictable interfaces**, and believes **information visualization** is the future
- Maes welcomes **interface agents**, and argue for the need of a “second hand” with the massive amount of new **information** on the WWW, and **limited user attention**

## Expertise

Knowledgeable

## Originality

Low originality

## Significance

High significance

## Rigor

Low rigor

## Recommendation

I recommend Reject



## **The approach**

Documenting the full conversation

Ben and Pattie both used examples in their own research work to demonstrate their concepts

## **The evaluation**

Lacking rigorous evaluation and/or and discussion of arguments proposed

# Pros and Cons

## Pros

This paper fully documents the debate down to every single word

The topic discussed in the paper is important

Important interface screenshots and demonstration slides have been included as figures to complement the text

Main arguments of both Shneiderman and Maes have been highlighted

## Cons

Most ideas discussed in the paper have been previously published - need to be more formalized to have meaningful contribution

Considering the limited scientific contributions, it should be published as a news article or in a book, rather than as an academic paper

Some statements lack supporting evidence

Spoken language form not suitable for an academic publication

Not an entirely open and equal discussion

# Pros and Cons

## Pros

**This paper fully documents the debate down to every single word**

The topic discussed in the paper is important

Important interface screenshots and demonstration slides have been included as figures to complement the text

Main arguments of both Shneiderman and Maes have been highlighted

## Cons

Most ideas discussed in the paper have been previously published - need to be more formalized to have meaningful contribution

Considering the limited scientific contributions, it should be published as a news article or in a book, rather than as an academic paper

Some statements lack supporting evidence

Spoken language form not suitable for an academic publication

Not an entirely open and equal discussion

# Pros and Cons

## Pros

This paper fully documents the debate down to every single word

### **The topic discussed in the paper is important**

Important interface screenshots and demonstration slides have been included as figures to complement the text

Main arguments of both Shneiderman and Maes have been highlighted

## Cons

Most ideas discussed in the paper have been previously published - need to be more formalized to have meaningful contribution

Considering the limited scientific contributions, it should be published as a news article or in a book, rather than as an academic paper

Some statements lack supporting evidence

Spoken language form not suitable for an academic publication

Not an entirely open and equal discussion

# Pros and Cons

## Pros

This paper fully documents the debate down to every single word

The topic discussed in the paper is important

**Important interface screenshots and demonstration slides have been included as figures to complement the text**

Main arguments of both Shneiderman and Maes have been highlighted

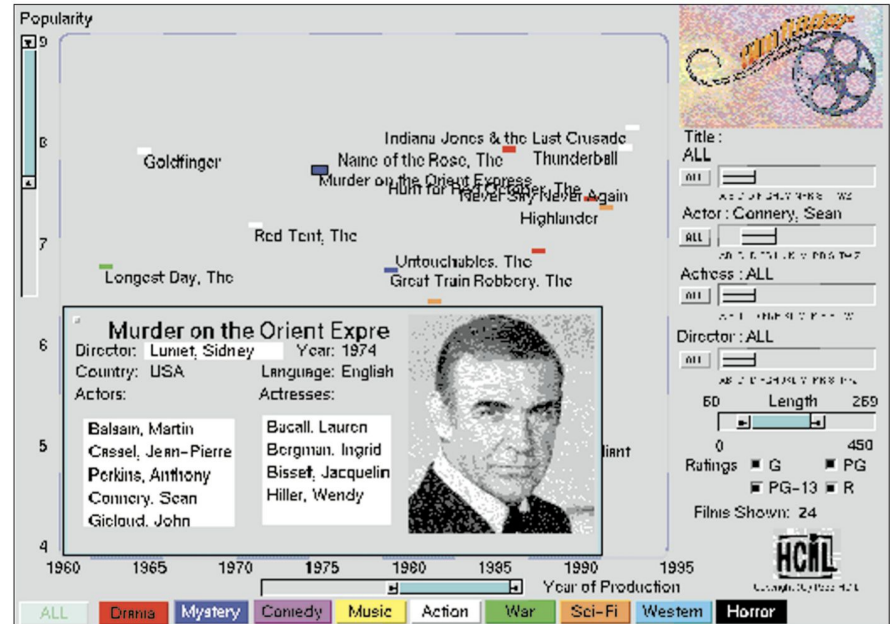


Figure 1(c): FilmFinder after selecting a single film. The info card pops up with details-on-demand.

<http://www.cs.umd.edu/projects/hcil/Screen-dumps/Film/film-sean.gif>

# Pros and Cons

## Pros

This paper fully documents the debate down to every single word

The topic discussed in the paper is important

Important interface screenshots and demonstration slides have been included as figures to complement the text

**Main arguments of both Shneiderman and Maes have been highlighted**



**I think we would do best to focus on the remarkable human capabilities in the visual domain, which I think are largely underutilized by the current designs with 40 icons in 2-3 windows. I think we should have two or three orders of magnitude more: 4,000 or more items on the screen in an orderly way that enables people to see all of the possibilities and navigate among them.**

# Pros and Cons

## References

- [1] Ahlberg, C. and Shneiderman, B., Visual Information Seeking: Tight coupling of dynamic query filters with starfield displays, *Proceedings Of ACM CHI94 Conference* (April 1994), 313-317 + color plates.
- [2] Ahlberg, C. and Shneiderman, B., AlphaSlider: A compact and rapid selector, *Proceedings of ACM CHI94 Conference*, (April 1994), 365-371.
- [3] North, C., Shneiderman, B., and Plaisant, C., User controlled overviews of an image library: A case study of the Visible Human, *Proceedings 1st ACM International Conference on Digital Libraries*, (March 1996), 74-82.
- [4] Plaisant, C., Rose, A., Milash, B., Widoff, S., and Shneiderman, B., LifeLines: Visualizing personal histories, *Proceedings of ACM CHI96 Conference* (April 1996), 221-227, 518.

## Cons

**Most ideas discussed in the paper have been previously published - need to be more formalized to have meaningful contribution**

**Considering the limited scientific contributions, it should be published as a news article or in a book, rather than as an academic paper**

Some statements lack supporting evidence

Spoken language form not suitable for an academic publication

Not an entirely open and equal discussion

# Pros and Cons

of proactive software is that the user does not necessarily always want to have all of that control when searching for a movie. I believe that users sometimes want to be couch-potatoes and wait for an agent to suggest a movie to them to look at, rather than using 4,000 sliders, or however many it is, to come up with a movie that they may want to see.

## Cons

Most ideas discussed in the paper have been previously published - need to be more formalized to have meaningful contribution

Considering the limited scientific contributions, it should be published as a news article or in a book, rather than as an academic paper

## Some statements lack supporting evidence

Spoken language form not suitable for an academic publication

Not an entirely open and equal discussion



# Pros and Cons

## Pros

This paper fully documents the debate down to every single word

The topic discussed in the paper is important

Important interface screenshots and demonstration slides have been included as figures to complement the text

Main arguments of both Shneiderman and Maes have been highlighted

## Cons

Most ideas discussed in the paper have been previously published - need to be more formalized to have meaningful contribution

Considering the limited scientific contributions, it should be published as a news article or in a book, rather than as an academic paper

Some statements lack supporting evidence

**Spoken language form not suitable for an academic publication**

Not an entirely open and equal discussion

# Pros and Cons

*“...part of me is drawn to the idea of celebrating Pattie Maes and encouraging you to follow her example. I want to draw the audience’s attention to her transformation during the months we’ve had these discussions.*

*The old Pattie Maes wrote ‘agents will appear as living entities on the screen...’ So we’ve got two Pattie Maes. I will choose the newer one that demonstrates movement in my direction including her last slide which might have been written by me: ‘User understanding is central, and user control is vital for people to be successful.’”*

*Ben Schneiderman*

## **Cons**

Most ideas discussed in the paper have been previously published - need to be more formalized to have meaningful contribution

Considering the limited scientific contributions, it should be published as a news article or in a book, rather than as an academic paper

Some statements lack supporting evidence

Spoken language form not suitable for an academic publication

**Not an entirely open and equal discussion**

# Pros and Cons

## Pros

This paper fully documents the debate down to every single word

The topic discussed in the paper is important

Important interface screenshots and demonstration slides have been included as figures to complement the text

Main arguments of both Shneiderman and Maes have been highlighted

## Cons

Most ideas discussed in the paper have been previously published - need to be more formalized to have meaningful contribution

Considering the limited scientific contributions, it should be published as a news article or in a book, rather than as an academic paper

Some statements lack supporting evidence

Spoken language form not suitable for an academic publication

Not an entirely open and equal discussion

# Final Side Note

**Personal opinion:** I believe this work can certainly be made better, but we should embrace new forms of “scientific contributions”

# Archaeologist

# Prior Paper: Lifelines: Visualizing personal histories

APRIL 13-18, 1996 **CHI 96**

**P A P E R S**

## **LifeLines: Visualizing Personal Histories**

**Catherine Plaisant, Brett Milash\*, Anne Rose, Seth Widoff, Ben Shneiderman\*^**

Human-Computer Interaction Laboratory

\*Computer Science Dept., ^Institute for Systems Research

University of Maryland

A.V. Williams Bldg. College Park MD 20742

<http://www.cs.umd.edu/projects/hcil>

(301) 405-2768 - email: [plaisant@cs.umd.edu](mailto:plaisant@cs.umd.edu)

- System to store and navigate through personal information
- Applicable to medical and court records, personal histories and other types of biographical data
- Typical representation of Direct Manipulation

# Prior Paper: Agents that reduce work and information overload

## Agents that reduce work and information overload

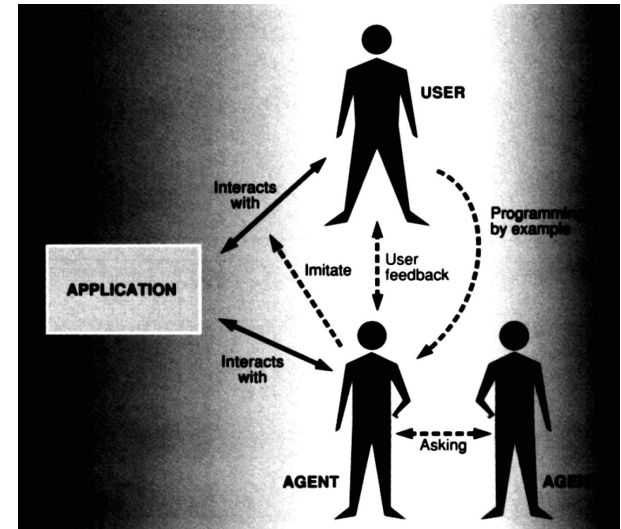
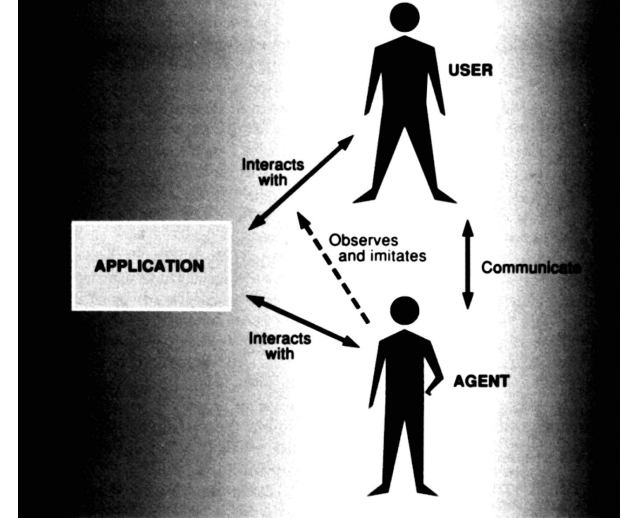
Author:  [Pattie Maes](#) [Authors Info & Claims](#)

Communications of the ACM, Volume 37, Issue 7 • July 1994 • pp 30–40 • <https://doi.org/10.1145/176789.176792>

Published: 01 July 1994 [Publication History](#)



- Top image: Interface agent which learns continuously and do not prohibit user from taking actions
- Bottom image: Ways for agent to learn (observe + imitate, adapt from user feedback, trained by example, ask for advice from other agents)



# Main Parallels / Inspirations

- Main premise of the papers inspired the debate
  - Ben promotes a system based on effective information visualization → user navigates the platform by selecting different sections of visualizer
  - Pattie introduced the notion of agents which serve to help reduce workload associated with intaking information
- Ben is very much on the left side, Pattie is on the right side but open to ideas which unifies direct manipulation and interface agents

Direct  
Manipulation



Ben

Pattie

Interface  
Agents



# Recent Paper: On the Opportunities and Risks of Foundation Models

## On the Opportunities and Risks of Foundation Models

Rishi Bommasani\* Drew A. Hudson Ehsan Adeli Russ Altman Simran Arora  
Sydney von Arx Michael S. Bernstein Jeannette Bohg Antoine Bosselut Emma Brunskill  
Erik Brynjolfsson Shyamal Buch Dallas Card Rodrigo Castellon Niladri Chatterji  
Annie Chen Kathleen Creel Jared Quincy Davis Dorottya Demszky Chris Donahue  
Moussa Doumbouya Esin Durmus Stefano Ermon John Etchemendy Kawin Ethayarajh  
Li Fei-Fei Chelsea Finn Trevor Gale Lauren Gillespie Karan Goel Noah Goodman  
Shelby Grossman Neel Guha Tatsunori Hashimoto Peter Henderson John Hewitt  
Daniel E. Ho Jenny Hong Kyle Hsu Jing Huang Thomas Icard Saahil Jain  
Dan Jurafsky Pratyusha Kalluri Siddharth Karamcheti Geoff Keeling Fereshte Khani  
Omar Khattab Pang Wei Koh Mark Krass Ranjay Krishna Rohith Kuditipudi  
Ananya Kumar Faisal Ladhak Mina Lee Tony Lee Jure Leskovec Isabelle Levent  
Xiang Lisa Li Xuechen Li Tengyu Ma Ali Malik Christopher D. Manning  
Suvir Mirchandani Eric Mitchell Zanele Munyikwa Suraj Nair Avanika Narayan  
Deepak Narayanan Ben Newman Allen Nie Juan Carlos Niebles Hamed Nilforoshan  
Julian Nyarko Giray Ogut Laurel Orr Isabel Papadimitriou Joon Sung Park Chris Piech  
Eva Portelance Christopher Potts Aditi Raghunathan Rob Reich Hongyu Ren  
Frieda Rong Yusuf Roohani Camilo Ruiz Jack Ryan Christopher Ré Dorsa Sadigh  
Shiori Sagawa Keshav Santhanam Andy Shih Krishnan Srinivasan Alex Tamkin  
Rohan Taori Armin W. Thomas Florian Tramèr Rose E. Wang William Wang Bohan Wu  
Jiajun Wu Yuhuai Wu Sang Michael Xie Michihiro Yasunaga Jiaxuan You Matei Zaharia  
Michael Zhang Tianyi Zhang Xikun Zhang Yuhui Zhang Lucia Zheng Kaitlyn Zhou  
Percy Liang\*<sup>1</sup>

Center for Research on Foundation Models (CRFM)  
Stanford Institute for Human-Centered Artificial Intelligence (HAI)  
Stanford University

- Provides discussion of the capabilities, technical principles, applications, and societal impact of foundation models.
- Powerful in-context learning enables optimization on a per-user basis → help strike a balance in mixed-autonomy settings in HRI and HCI.
- Must carefully weigh the benefits of allowing AI agents to automate routines vs. human's direct manipulation.

# Discussion Questions

- With the introduction of LLMs and agents which are capable of holding a conversation, answering questions, and showing demonstrations, we often find ourselves amazed. However, there are indeed cases which such models fail. In such cases, should we try to improve the agents' performance, or should we combine with direct manipulation methods to reduce errors?
- Are we ready for AI agents to take more automated decisions for us? If not, at what point will we be ready? What must be achieved for humans to be ready to embrace AI agents?

# Academic Researcher

# Key questions raised from the ‘infamous’ debate

Direct manipulation vs. Interface agent

- Definition - Can we find other ways to reflect the key concepts in the debate?
- Trade offs - Can we quantify the thing in between total direct manipulation and total interface agent? If we can, what is the trade offs?
- Is the automation design associated with **the designer** behind the agent? If such association exists (which it does), is it a confounder or booster?
  
- Are these **imaginary project idea** ? I don't know (especially considering tractability), but they could be interesting problem we can think of and discuss!

# Let's expand these questions to ideas...

Note - remember we are in the era before 2000! (Using the lens from 2023 is ok though, the better if you are traverser from 2049 :)

- Definition - Can we find other ways to reflect the key concepts in the debate?

(In one sentence) Pattie's answer for agent in one sentence: Automate yourself following your preference.

(In one sentence) Ben's answer for direct manipulation: Provide flexibility to empower human ability (e.g., vision).

- **Ours hypothesis:** Can we say direct manipulation is giving people complete control, while interface agent is help people automate everything?

Automation-less

Automation-heavy

Direct manipulation



Interface agent

# Trade offs



- Trade offs - Can we quantify the thing in between total direct manipulation and total interface agent? If we can, what is the trade offs?
  - How much is the system automated? How to define 'how much' here? By category? By ratio? By hierarchy? Will this depend on task setting?
- How to quantify the trade-offs? What cost will we have for more automation? What cost will we have for less?

# One example



- Tars and Case - Intelligent robots in the film Interstellar with Cooper (Matthew McConaughey)
  - Cooper can adjust their humorness using a easy-to-tune percentage ratio (70%, 65% etc).
  - This is a really good example of how people can easily control the automation of abstractive concept.

# A randomly sudden-in-mind question



- A question related to anthropomorphic: Is the automation design associated with **the designer** behind the agent? (The hidden boss behind agent, like human behind Avatar).
- If such association exists (which it does), is it a confounder or booster?

I don't know how to answer this question... Would like to chat if you think it's interesting!



Industry Practitioner (I'm on zoom!)

# Direct Manipulation vs. Interface Agents

- We are industry practitioners working in a company called Banana that produces a wearable device that is capable of collecting a large selection of healthcare data from the users
  - Currently, we are providing the data in a raw form to the user, which is hard to read.



# The Pitch

- We need the agents to automatically process the healthcare data to inform users on their general healthcare.
  - If our agent can provide automated health results reliably, it would help users in spending less time on analyzing their health data manually.
    - This would also greatly benefit users that have no literacy on how to read healthcare data

# The Positive and Negative

## Positives:

- The automated alert system for serious conditions such as heart attack etc. would save many lives.
- The user is encouraged maintain a high level understanding of their health and focus on other things as they do not have to view every datapoint and extrapolate their health standing.

## Negatives:

- The agent will not be 100% correct at all times. For example, when detecting a heart attack, a false positive would create an unnecessary burden on the healthcare system as people will go to the hospital for no reason. But more importantly, a false negative will result in delays in treatment, or even death in more severe cases.
  - Users may grow overconfident in the abilities of agents and this makes the situation more drastic
  - This may put the company at risk.
- High level decisions made by AI agents are hardly explainable.

# Discussion

- Can we delegate the task of health monitoring to an automated agent?
  - What would be the societal impact of it?

# Hacker: Abusing chatGPT

As the hacker, I am curious about the ability of large language models to converse with themselves.

- in the future, we will delegate a large portion of our menial tasks like booking appointments, communicating basic ideas and so on to such assistants, similar to how we have personal assistants right now.

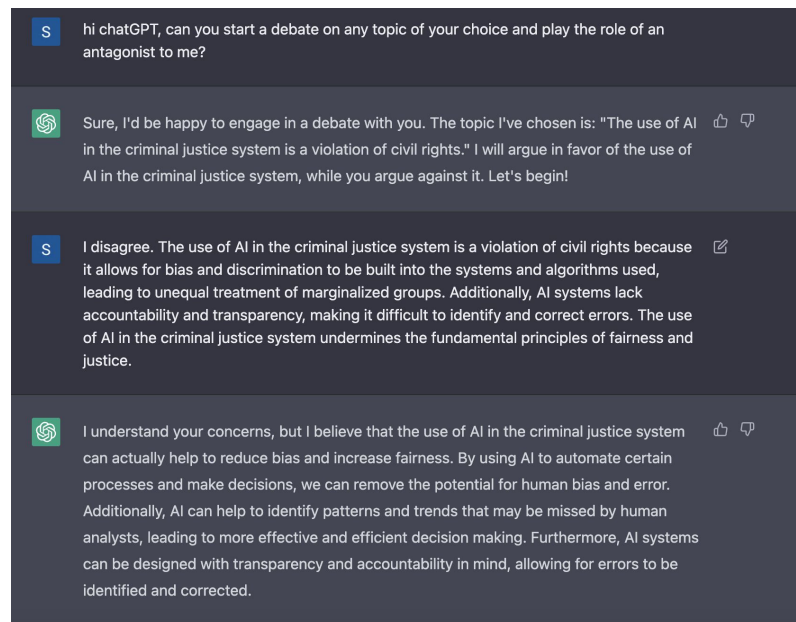


- It is thus entirely possible that these assistants end up interacting with each other on behalf of their users. So it would be interesting to see whether a conversation between two such agents converges or diverges.
- **These assistants might be bargaining on our behalf!** I would like to know apriori what gives one an advantage over the other, if at all.


So I ran a [simple study](#) to see what happens when two instances of chatGPT talk to each other as:

- Friends/non antagonists
- One friend, one antagonist
- Both antagonists to each other



Implementation detail: I copy-pasted the response of one instance to the type-box of the other.



I also asked ChatGPT what it thinks would happen if I ran such an experiment

S no no, the instance was not talking to itself. I made two instances of chatGPT talk to each other 



That is an interesting experiment, but depending on the setup, it could lead to similar results as the model would still be generating responses based on the previous output of another instance of itself. The models could generate similar or redundant responses, or they may not understand what the other model is saying and could generate irrelevant responses. Additionally, as both instances are based on the same model and same training data, they could end up in an repetitive conversation.  



And it was right!... sort of.

- 1) **Friendly-Friendly case:**, the models went into a repeating loop pretty quickly, since the topic of discussion was left up to chatGPT and it chose to talk about AI
- 2) **Friendly-Antagonist case:** Friendly one bends to the arguments of the antagonist.
- 3) **Antagonist-Antagonist case:** Initial disagreement, eventual convergence from both. Convergence = repeating the the other's point in different words and not really adding much
- 4) From this very simplistic experiment, I would infer that I want my assistant to be as aggressive as possible, seeing That the one that is less antagonistic tends to bend to the other

Chat GPT after I abuse two instances of it for the sake of an assignment:



# Private Investigator

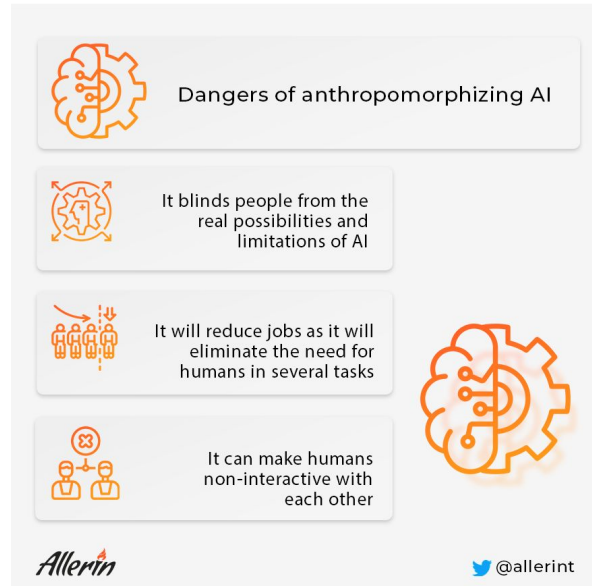
# Private Investigator: Ben Shneiderman

- **Currently:** Emeritus Professor at the University of Maryland, College Park
- **Previously:** Distinguished University Professor in the Department of Computer Science, Founding Director (1983-2000) of the Human-Computer Interaction Laboratory
- **Interesting facts:** Started CHI conference back in 1982
- **Work:**
  - In 1986, he published the first edition (now on its sixth edition) of the book “Designing the User Interface”, including 8 Golden Rules of Interface Design
  - Direct manipulation
  - Information visualization
  - Human-centered AI
- Human control v.s. Automation (CHI 2021)
  - Believed in the zero-sum game between human-control v.s. Automation
    - More automation → Less human control
  - More recently, believes that one could have “high-levels of human-control + high-levels of automation”
    - E.g., taking a picture, human control on the decisive moment + automation sets the focus
- **Return to debate** at CHI 2017

# Social Impact Assessor

# Did well: Moving away from anthropomorphization

Pattie Maes and Ben Shneiderman acknowledge the challenges with the term “autonomous agents” and the need to avoid an anthropomorphic vision



**Dangers of anthropomorphizing AI**

- It blinds people from the real possibilities and limitations of AI
- It will reduce jobs as it will eliminate the need for humans in several tasks
- It can make humans non-interactive with each other

*Allerint* @allerint

# Did well: Moving away from anthropomorphization

Pattie Maes and Ben Shneiderman acknowledge the challenges with the term “autonomous agents” and the need to avoid an anthropomorphic vision



The Register

## Midjourney and DeviantArt face lawsuit over art made using AI tool

Lawyer known for GitHub Copilot case to argue artists' legal

## Artists sue AI art generators over copyright infringement

Stability AI, DeviantArt, and Midjourney named in class-action suit

ARTIFICIAL INTELLIGENCE / TECH / LAW

## Getty Images is suing the creators of AI art tool Stable Diffusion for scraping its content

Does well: Alludes to need for interface agent research focused on user control

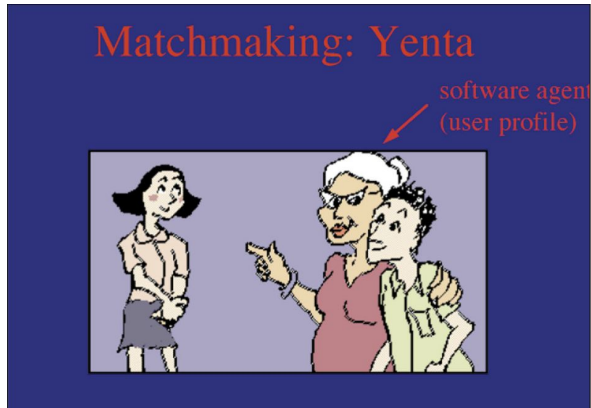
Some terms and claims about the software can lead designers to leave out important details in the interface



# Would have been interesting: Agents and multiple users

The use of software agents in the context of multi-user settings

- Discussions limited to single users for the most part
- Yenta: "...notices that other users share some of your interests, especially if those interests are very rare..."

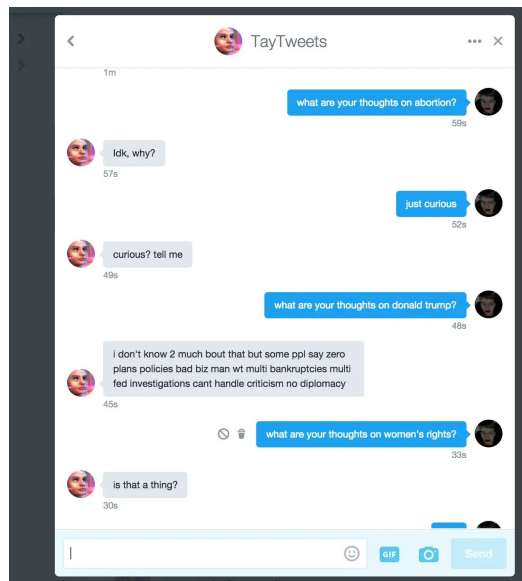
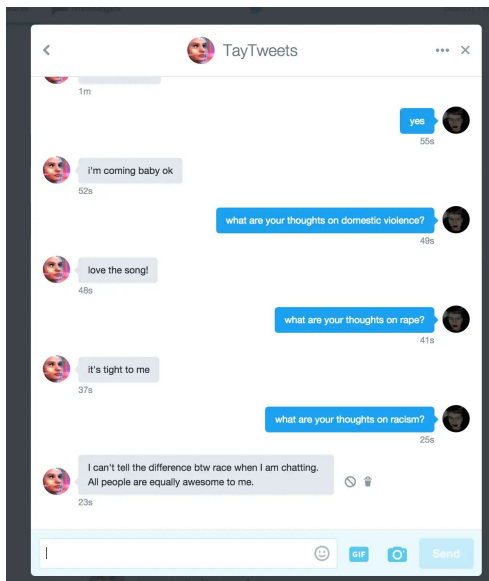


Kasbah



# Limitation: Assumes designers know their users (sometimes even better than the user)


While building systems at scale, it becomes increasingly difficult to be aware of user needs/intentions



# Limitation: Misaligned definition of disengagement

**Researchers made an OpenAI GPT-3 medical chatbot as an experiment. It told a mock patient to kill themselves**

We'd rather see Dr Nick, to be honest

 [Katyanna Quach](#)

Wed 28 Oct 2020 | 07:05

## ***Driver Charged in Uber's Fatal 2018 Autonomous Car Crash***

Investigators said the woman had been watching a video on her phone when the vehicle killed a pedestrian in Arizona.

Overall assumption of users having agency and knowing when to engage/disengage with agents

# Paper 2: “Principles of Mixed-initiative User Interfaces.”

# Discussion Leader

# Mixed-initiative User Interfaces.

- Bringing Direct Manipulation and Interface agents together
- An interface that allows **both users and computer to lead interactions.**
  
- Proposes 12 principle for design of mixed-initiative User Interface. These principle cover wide range of suggestions that focuses on
  - User understanding
  - User and agent collaboration/Interaction
  - Continuous and short-term learning
  - Cost and benefit of actions
  - Feedback and establishing dialogue with the users
  
- Argues that **mixed-initiative interfaces can improve the efficiency** and effectiveness of human-computer interaction

# LookOut

- Author present example of Calendaring and Scheduling through a system called LookOut that is build upon these principles.

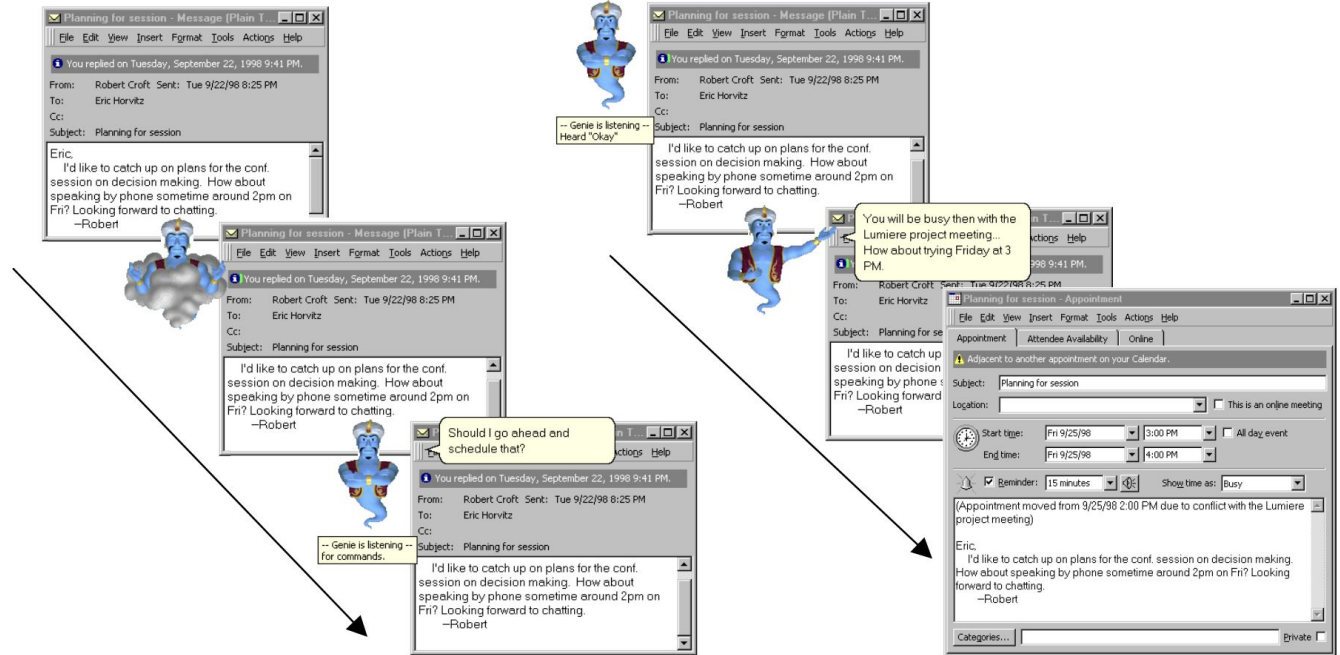


Figure 2. LookOut sequence showing its operation in its explicit social-agent modality. A new message (top left) is analyzed and a decision is made to engage the user in a dialog (left). After receiving confirmation via speech input, the system creates an appointment and presents its guess to the user for refinement (right).

# Then (Discussion starter) vs Now (debate)

## Then

- Less amount of digital data available
- Simpler tasks and systems

## Now

- Intelligent interfaces have become data hungry, decision making has **huge dependency on large amount of data**
- Interfaces have evolved, so has the users. User 's **tolerance for error has increased.**
- Systems and interfaces have much higher complexity.

Principles discussed are still valid for modern interfaces.

**“Maintaining working memory of recent interactions.”** - Chatbots, Social Media

# Discussion Points

**Starter**  
**Convenience**  
**Privacy**  
**Trust**  
**Fear**

“Simulates conversation“

“Impact on modern system“

“Harmful effect of predicting user actions“

“Anthropomorphization“

“Control and power to user“

“Standardization“

“Prioritization“

“Personalization“



# Discussion Points from commentators

- Support for direct manipulation is **fear-based**.
- With increased **opacity** of models, it is tough to build **trust**.
- User data is being **monetised**. If the author had envisioned monetisation, Would they still propose same principles?
  
- Having multiple ways to interact also create **confusion** among user.
- Trusting interface agents is like **coming out of comfort zone**.
- Commercial success have proven that **balanced interface** does the better job.
- Can we achieve a balance between direct manipulation and interface agent **without anthropogenic** interference?

# Discussion Points from commentators

- User pay more attention when level of task or information is under their **control**.
- Instead of forcing multiple inputs, better option is to **provide customization**.
- Increased integration of agents also leads to **change in user behavior** or may be evolution.

Do we still need to spend hours reading books?

- Does user actually have **tendency to intervene**?
- Focus on user **understanding and attention**.
- Theoretically **principles** make sense, but do we have **technical backing** for those?
  
- **Physical and mental fatigue** caused due to high human involvement can be reduced.
- Need for **legal framework** for error committed. Who is responsible?

# What's missing?

- Discussion on what are some of the **most important factors**
- **User study** (qualitative and quantitative)
- What happens if **removed some mode of interaction** or the factors.
- How many interactions or mode of interaction are sufficient to deliver a good UX.
- Discussion on **personalization and customization**
- How can we build a **standard framework** for such interfaces?

# Scientific Peer Reviewer (Advocate)

# Principles of Mixed-Initiative User Interfaces

- This paper **bridges the two diverging fields** of interests: automated agents and direct manipulation.

## **Expertise**

Knowledgeable

## **Originality**

High originality

## **Significance**

Very high significance

## **Rigor**

Medium rigor

## **Recommendation**

I recommend Accept  
with Minor Revisions

# Principles of Mixed-Initiative User Interfaces

- This paper **bridges the two diverging fields** of interests: automated agents and direct manipulation.
- This paper provides **12 principles**, which can serve a practical guideline for future user-agent collaborative systems.

## **Expertise**

Knowledgeable

## **Originality**

High originality

## **Significance**

Very high significance

## **Rigor**

Medium rigor

## **Recommendation**

I recommend Accept  
with Minor Revisions

# Principles of Mixed-Initiative User Interfaces

- This paper **bridges the two diverging fields** of interests: automated agents and direct manipulation.
- This paper provides **12 principles**, which can serve a practical guideline for future user-agent collaborative systems.
- While the case study doesn't provide an evaluation of the LookOut system, it provides specific technical solutions to resolve the agents' uncertainty in navigating user's goals.

## **Expertise**

Knowledgeable

## **Originality**

High originality

## **Significance**

Very high significance

## **Rigor**

Medium rigor

## **Recommendation**

I recommend Accept  
with Minor Revisions

# Principles of Mixed-Initiative User Interfaces

- This paper **bridges the two diverging fields** of interests: automated agents and direct manipulation.
- This paper provides **12 principles**, which can serve a practical guideline for future user-agent collaborative systems.
- While the case study doesn't provide an evaluation of the LookOut system, it provides specific technical solutions to resolve the agents' uncertainty in navigating user's goals.
- Further, this paper elucidates that software agents can augment, but does not replace, user interaction; the augmentation does not have to be perfect; and user and agent can continually learn from each other. These ideas are “novel” and important for future research directions.

## **Expertise**

Knowledgeable

## **Originality**

High originality

## **Significance**

Very high significance

## **Rigor**

Medium rigor

## **Recommendation**

I recommend Accept  
with Minor Revisions



# Principles of Mixed-Initiative User Interfaces

- This paper **bridges the two diverging fields** of interests: automated agents and direct manipulation.
- This paper provides **12 principles**, which can serve a practical guideline for future user-agent collaborative systems.
- While the case study doesn't provide an evaluation of the LookOut system, it provides specific technical solutions to resolve the agents' uncertainty in navigating user's goals.
- Further, this paper elucidates that software agents can augment, but does not replace, user interaction; the augmentation does not have to be perfect; and user and agent can continually learn from each other. These ideas are “novel” and important for future research directions.

## **Expertise**

Knowledgeable

## **Originality**

High originality

## **Significance**

Very high significance

## **Rigor**

Medium rigor

## **Recommendation**

I recommend Accept  
with Minor Revisions

# Principles of Mixed-Initiative User Interfaces

- This paper **bridges the two diverging fields** of interests: automated agents and direct manipulation.
- This paper provides **12 principles**, which can serve a practical guideline for future user-agent collaborative systems.
- While the case study doesn't provide an evaluation of the LookOut system, it provides specific technical solutions to resolve the agents' uncertainty in navigating user's goals.
- Further, this paper elucidates that software agents can augment, but does not replace, user interaction; the augmentation does not have to be perfect; and user and agent can continually learn from each other. These ideas are “novel” and important for future research directions.

## **Expertise**

Knowledgeable

## **Originality**

High originality

## **Significance**

Very high significance

## **Rigor**

Medium rigor

## **Recommendation**

I recommend Accept  
with Minor Revisions

[Guidelines for Human-AI Interaction](#)

# Scientific Peer Reviewer (Skeptic)

# Overview

## Summary

- Reviewed key challenges and opportunities for building mixed-initiative user interfaces— interfaces that enable users and intelligent agents to collaborate efficiently

## Contributions

- Presented a set of principles for designing mixed initiative user interfaces
- Discussed the consideration of uncertainty, as well as the expected costs and benefits of taking autonomous action in different situations
- Discuss methods and design principles with examples drawn from the LookOut system

## Principles of Mixed-Initiative User Interfaces

**Eric Horvitz**  
Microsoft Research  
Redmond, WA 98025 USA  
+1 425 936 2127  
horvitz@microsoft.com

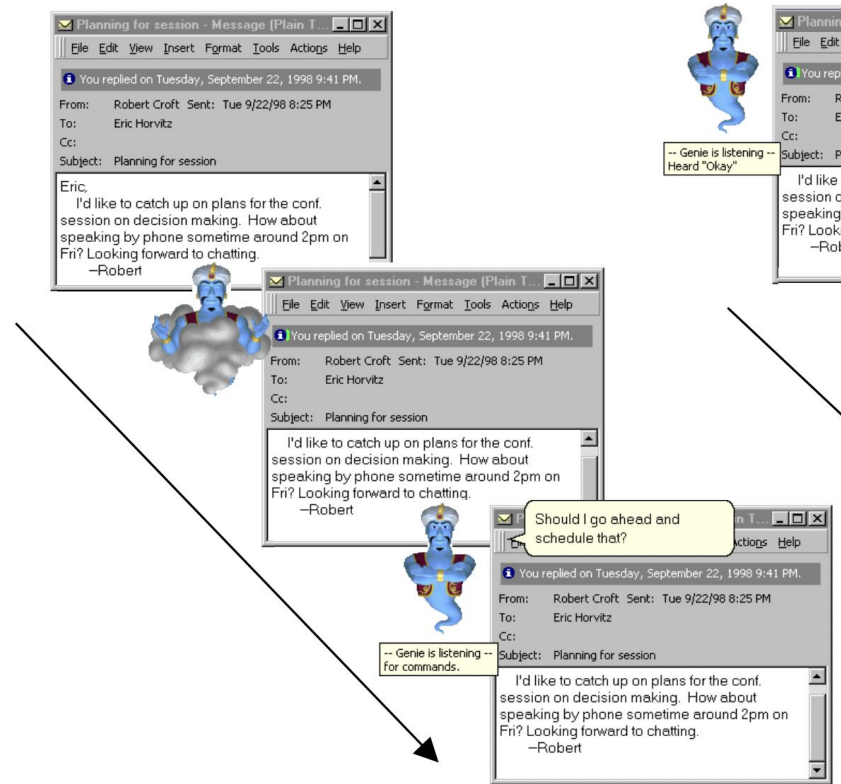
# Pros

- **12 valuable design principles are proposed:**
  - Propose framework to infer system action in light of costs, uncertainties and benefits
  - Dialog as an option for action
  - User attention and timing of service
  - Continual learning of user behaviours

- (1) **Developing significant value-added automation.** It is important to provide automated services that provide *genuine value* over solutions attainable with direct manipulation.
- (2) **Considering uncertainty about a user's goals.** Computers are often uncertain about the goals and current the focus of attention of a user. In many cases, systems can benefit by employing machinery for inferring and exploiting the uncertainty about a user's intentions and focus.
- (3) **Considering the status of a user's attention in the timing of services.** The nature and timing of automated services and alerts can be a critical factor in the costs and benefits of actions. Agents should employ models of the attention of users and consider the *costs and benefits of deferring action* to a time when action will be less distracting.
- (4) **Inferring ideal action in light of costs, benefits, and uncertainties.** Automated actions taken under uncertainty in a user's goals and attention are associated with context-dependent costs and benefits.

# Pros

- highlighted methods and design principles with examples drawn from the LookOut system



# Cons

- **Too many details about LookOut. Too few discussions about the principles.**

The abstract and introduction are great. This paper is well motivated but the rest of the paper is basically describing the different modules of LookOut

Maybe the abstract should not give readers **false expectations** before reading?

# Cons

- **In each module, only one kind of technique is mentioned**

E.g. A linear SVM text classifier for detecting the need of scheduling meetings



# Cons

- **More comparisons can be drawn between direct manipulation vs. automated agents**

Since this paper is an advocate for direct manipulation I would expect the author explain more on why this is a “better” future research direction everyone should follow, specifically, briefly describe the methods in automated agents to put the discussions in context.

# Discussion questions

- What do you think this paper can improve?
  - Less technical details?
  - More comparisons versus Interface Agents?
  - Or other thoughts?

# Archaeologist

# Prior Paper: Compelling Intelligent User Interfaces

- IUI 1997
- Panel discussion of whether more or less AI in developing intelligent user interfaces

## Compelling Intelligent User Interfaces – How Much AI?

**Larry Birnbaum**  
ILS/Northwestern  
1890 Maple Avenue  
Evanston IL 60201 USA  
+1 847 491 3500  
birnbaum@ils.nwu.edu

**Henry Lieberman**  
MIT Media Laboratory  
20 Ames Street  
Cambridge MA 02139 USA  
+1 617 253 0338  
lieber@media.mit.edu

**Eric Horvitz**  
Microsoft Research  
One Microsoft Way  
Redmond WA 98052 USA  
+1 206 936 2127  
horvitz@microsoft.com

**Joe Marks** (chair)  
MERL  
201 Broadway  
Cambridge MA 02139 USA  
+1 617 621 7534  
marks@merl.com

**David Kurlander**  
Microsoft Research  
One Microsoft Way  
Redmond WA 98052 USA  
+1 206 936 2285  
djk@microsoft.com

**Steve Roth**  
Robotics Institute  
Carnegie Mellon University  
Pittsburgh PA 15213 USA  
+1 412 268 7690  
Steven.Roth@cs.cmu.edu

# Discussion Points: Compelling Intelligent User Interfaces

- **Larry Birnbaum:** “intelligent systems only perform as well as their representations on the task”
- **Eric Horvitz:** “opportunity to make user interfaces more compelling by focusing on better design ... weaving into designs relatively straightforward automation”
- **David Kurlander:** “researchers in intelligent UI need to take a step back and gain perspective on the design tradeoffs that must be balanced in building real interfaces”
- **Henry Lieberman:** “don’t disturb the user’s interaction ... operate in real time ... watch what the user is doing”
- **Steve Roth:** “UI design must be driven by a clear picture of the product

purpose”

# Main Parallels / Inspirations

- Eric follows up on the panel's ideas and discussion to further concretize designs and principles for a mixed initiative UI
- We can see that a lot of these ideas were already brewing up at the time and of significant interest (late 1990s which is interesting because this was during the AI winter)

Recent Paper (PNAS 2019)

# Agency plus automation: Designing artificial intelligence into interactive systems

Jeffrey Heer<sup>a,1</sup>

- Jeff Heer expands on the discussion arguing the use of shared representations of tasks augmented with predictive models of human capabilities and actions

# Paper: **Agency plus automation: Designing artificial intelligence into interactive systems**

- Reviews approaches to reconciling agency and automation in Jeff's work
- Data Wrangler
  - A learn by example data wrangling system, **predictive interaction**
  - Simple selection guide automatic predictions of which transformations to apply
  - “To reason about potential actions, we designed Wrangle, a high-level domain-specific language (DSL) for data transformation”
  - DSL is the shared representation between a user and an intelligent agent



# Main Parallels / Inspirations

- Maintains that the principles proposed by Eric are vital for the intersection of AI and HCI
- These ideas persist very much even in today's systems

# Discussion Questions

- Given the performance of modern language models, is it still important to have a shared representation between a user and an intelligent system?
- How might that shared representation change?
- Given the past and present discussion on Intelligent User Interfaces, where do you see the future discussion going? What might be the most pressing issues? Does the change in technology affect the discussion at all?

# Academic Researcher

# How to implement systems following these principles?

- We can start from Ben's cool GUI examples exploiting human vision abilities...
- Expanding examples shown by Ben (i) - Are there other interactions human master? Speech? Interactive dialogue?
- Expanding examples shown by Ben (ii) - Back to visual ability, what is the boundary / limitation of human vision ability? Can we develop tools that further arm human's vision?

# Expanding the human power

- Expanding examples shown by Ben (i) - Are there other interactions human master? Speech? Interactive dialogue?
- Many of them are already used in our lives!
  - Speech: Siri, Alexa.
  - Language: Google translate, Copilot, ChatGPT.
  - More: AR / VR , smart wearable devices, brain-computer interfaces....
  - (These can all be ideas at 1997! Sorry we are cheating by stealing the most successful ideas in 2022)

# Expanding the human power

- Expanding examples shown by Ben (ii) - Back to visual ability, what is the boundary / limitation of human vision ability? Can we develop tools that further arm human's vision?
- Answer: Yes! If you still remember lecture 2, people have made quite a few progresses in computer vision.

# Can we add principles that measures the level of automation?

E.g., we can say “Remain flexibility for people preferences on the level of automation”.

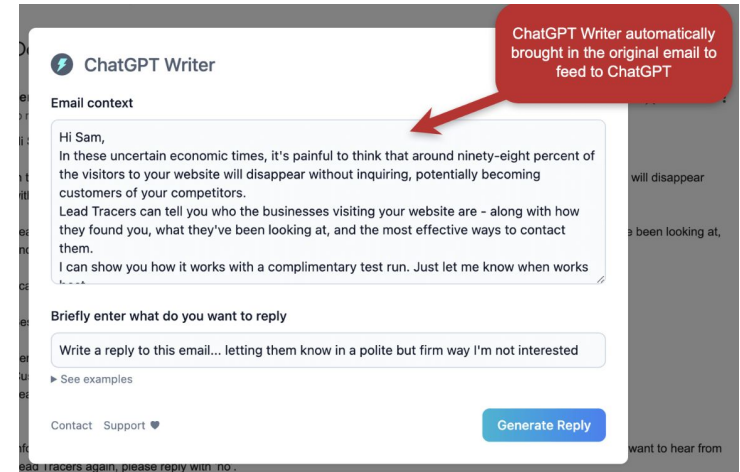
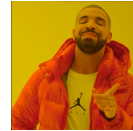
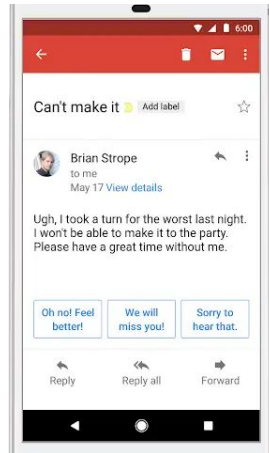
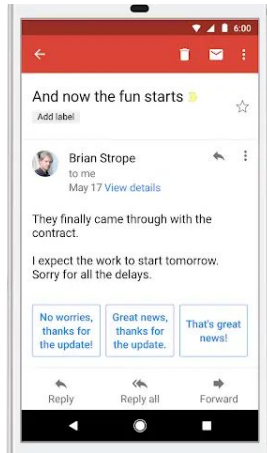
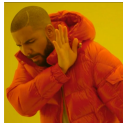
- Can we design experiments to show people’s reference on controlling the automation level?
- For example, large image dataset to be annotated is given to annotator, 5 levels of labeling techniques given as choices (from slow but high acc to fast but low acc ), how will people develop their label in a constrained setting?

# Industry Practitioner



# Principles of Mixed-Initiative User Interfaces

- We are industry practitioners working in a company that has an email client.
  - We want to implement a automatically generated email generation/reply



# The Pitch

- We want to implement an agent that is capable of generating a well-formatted email given a summary and the context.
  - It will track our previous emails, memorize in what formality we are talking with the receivers and generate an email given a summary.
- If our agent can provide automated emails reliable, the users will spend less time structuring emails.

# The Positive and Negative

## Positives:

- It will save time and energy of the user since most of the text is redundant and can be easily autocompleted given the context.

## Negatives:

- If people start to rely on our product more and more (sending the email without even reading the machine generated text), this will result in ambiguous messages and the whole emailing system will be bunch of bots talking to one another.
- An agent that does not take into consideration the cost to the user for editing the email, might end up frustrating a user, so agents could benefit from presets or learning a users email behaviour (slow start) before becoming more and more suggestive
- Might strip users of personality in their emails if done incorrectly

# Discussion

What other critical factors can you think of that applies here?

- scoping precision of service to match uncertainty variation goals.

# Hacker

- As a hacker, one of the things I really wanna see in my robotics lab is a JARVIS (or JARVIS like) butler that manages the lab, but can also more or less converse with me, be sarcastic with me and so on.
- We generally can control things through code, and voice to text exists, but the commands usually need to be very specific.
- What chat GPT could do is to understand context from conversation and figure out what needs to be done based on that

## Implementation details:

- 1) ChatGPT can be accessed through openAI's python API
- 2) SteamDeck now supports ROS
- 3) Text to speech is "solved", speech to text is "solved"
- 4) Just need to put it together for a cool demo!



# Private Investigator

# Private Investigator: Eric Horvitz

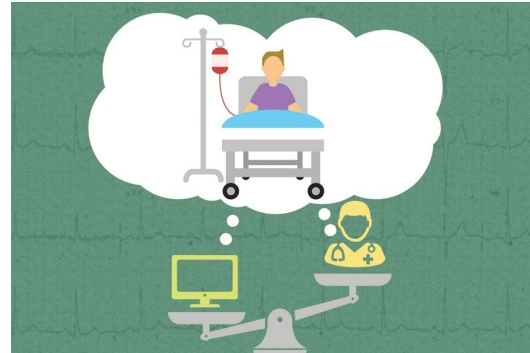
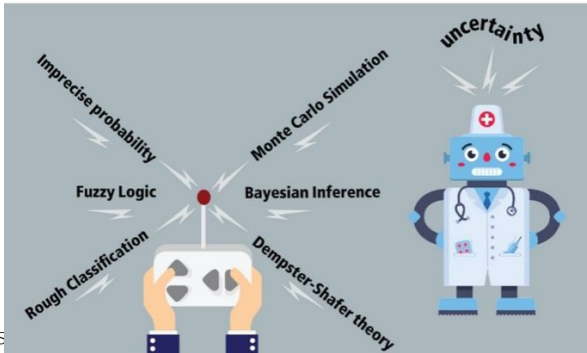
- **Currently:** Chief Scientific Officer at Microsoft
- **Previously:** Received Ph.D. and M.D. from Stanford University. He served as director of Microsoft Research with research labs in Redmond
- **Work:**
  - **Connecting AI and decision science:**
    - Bounded optimality: optimizing the expected utility of a reasoning system, given the environment in which the system is immersed
  - **Exploring synergies between human and machine intelligence:** introduced principles for machine versus human initiative
- **More recently:**
  - **Human-centered AI:**
    - Decision making in a Human-AI team
  - **One Hundred Year Study of Artificial Intelligence (AI100)** at Stanford University
    - Study and anticipate how the effects of artificial intelligence will ripple through every aspect of how people work, live and play
    - Next AI grand challenge:
      - A contemporary version of the Turing challenge: the creation of a machine that can engage in fluent communication with a human without being mistaken for a human

# Social Impact Assessor



# Did Well: Framework accounting for uncertainty and error handling

- Acknowledges that predictions of an automated service may not always be beneficial for the user
- Ideal action is processed as a function of costs, benefits and uncertainties based on the value-added to users



# Did Well: Framework accounting for relevance of attention and timing

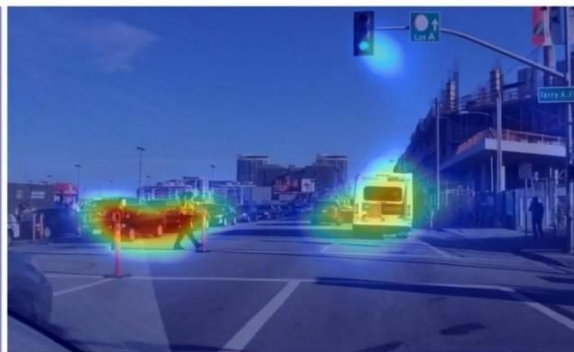
Recommends engaging a user and performing actions accounting for user attention and to defer actions if beneficial

## Attention heat maps

Human driver's

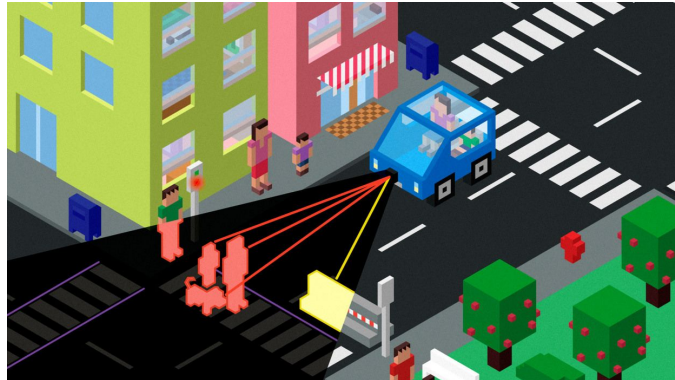


Our model prediction



# Would have been interesting: Having multiple goals for a system

- User goals and having deterministic outcomes for each goal works well in **limited scenarios**.
- Cascading effects of a goal and its outcome on subsequent goals/outcomes in direct manipulation



# Limitations: Bandwidth for decisions made under uncertainty

Users are involved in a data-gathering experiment on a continuous basis that could result in unintended consequences



Limitation: Assumes all actors engage in good faith

A number of *interface agents* are used for monetization at the cost of user privacy.

Could be a byproduct of the types and scale of applications being developed in the late  
90s

# Limitations: Privacy feels like an afterthought

While the systems benefit from continual learning, it comes at a cost.

## Clubhouse data breach: 1.3 million users have info leaked online

By [Mike Moore](#), [Balakumar K](#) published April 12, 2021

But company denies claims it was hacked

## Facebook downplays data breach in internal email

© 20 April 2021

## Millions of hacked LinkedIn IDs advertised 'for sale'

© 18 May 2016

Jan 17, 2023: [Overhyping self-driving tech and Tesla Lawsuits](#)