

# Human-Centred AI

---

Dr Charles Martin

## Announcements

- all tutorial marks have been entered.
- congrats to those (107) legends who **maxed** out 100% on tutorials!
- SELT should be open this Week (??) — fill in your SELT!
- Your SELT is a gift to future students, and gives me confidence to advocate for good practice!

## Final Project

- brings all skills from the course together
- idea is that these are mini HCI-research papers
- prepare you for further research, exploration in HCI and assess all skills in the course
- **one week left for your final project**
- let's look at the forking numbers for the assignment repo

# Plan for the class

- a few notes wrapping up the course
- human-centred AI — research and design challenge for the future (Shneiderman, 2022)
  - basically covering the content of this book, available for free via library and online.
- revision on key ideas



**Figure 1:** Today we will (try) to cover *all* of Shneiderman (2022)

# Course stats

- Max of 330 students (down to ~295) — **biggest ever HCI class at ANU.**
- Git Commits: >370
- Slides: 565 slides, 38185 words, 386 images
- Tutorials: 10 new tutorials, 11283 words
- Forum:
  - 3K threads, 363 answers, 858 comments.
  - 185K views, biggest viewer viewed 7035 threads (???) (I only viewed 4095!)
  - Charles answered 180+ questions

Takeaway: we wrote the equiv of a PhD thesis (~50K words) this semester—yikes! 😞

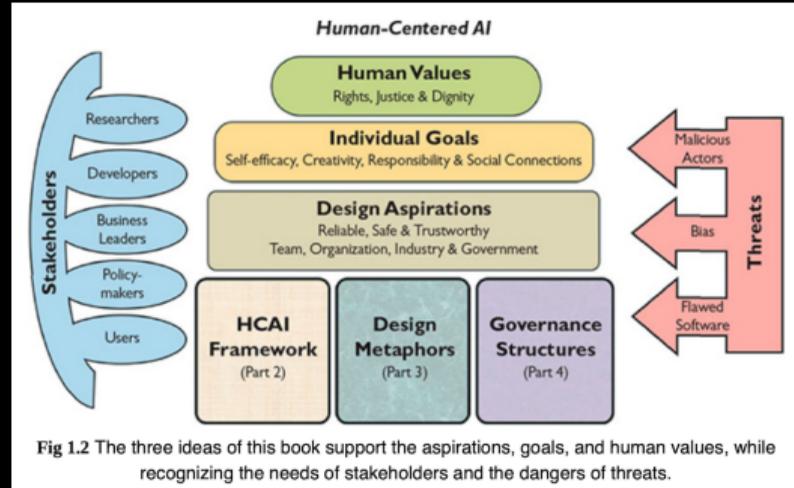


Figure 2: Github Insights on HCI this year...

# What is Human-Centered AI (HCAI)?

---

# What is Human-Centered AI (HCAI)?



**Figure 3:** Map of Shneiderman (2022) view of Human-Centered AI

Shneiderman (2022) proposes combining human-centred thinking with AI-based algorithms to create HCAI

- increase chance of AI empowering rather than replacing people
- give equal attention to human users and other stakeholders when developing AI systems
- HCAI developers would value meaningful human control
- serve human values: rights, justice, dignity
- goals: self-efficacy, creativity, social connection

## Why is this important?

- AI is the transformational technology of our time (*Charles, 2025*)
- Applying AI well is a *massive* HCI challenge, probably bigger than any before
- We don't have good enough frameworks to understand HCAI

### Evidence:

- lots of AI papers involving humans wrongly sent to CHI conference, frustrating reviewers
- AI collaboration technology currently at the level of CLI (chatGPT) or confusing magical tools (agentic coding, e.g., copilot in VSCode)
- *So many* research questions in this area.
- news preoccupation with risks of AI taking over, not capacity of actual products

## How is HCAI different to AI?

Shneiderman (2022) points to two differences between AI and HCAI:

1. **process:** HCAI builds on user experience design methods (e.g., data gathering, stakeholder engagements, etc.) in use of systems that employ AI and ML
2. **product:** HCAI systems supposed to be *supertools* that amplify, augment, empower and enhance human performance

Critical points here:

- Research methods for AI/ML, primarily statistical/logical modelling, very quant, much different than design methods.
- HCAI *supertools* emphasise human control and the enhancement of the human's capability and experience, not the AI/ML algorithms capability.
- Examples of *supertools*: digital cameras, navigation systems

## White Robot vs Supertool

AI reporting frequently applies simplistic tropes for AI systems.

- “white robot” (technoutopian)
- “terminator” (technodystopian)

These tropes mirror people, not computer systems. White robots are naive but noble people (can be turned to evil), terminators are evil people (can be redeemed)

### Supertool

- Shneiderman’s “supertool” enables us to consider how AI fits into existing tool use.
- Not a complete solution, still lots of questions and unknowns!



**Figure 4:** No white robots (Photo by Maximalfocus on Unsplash)

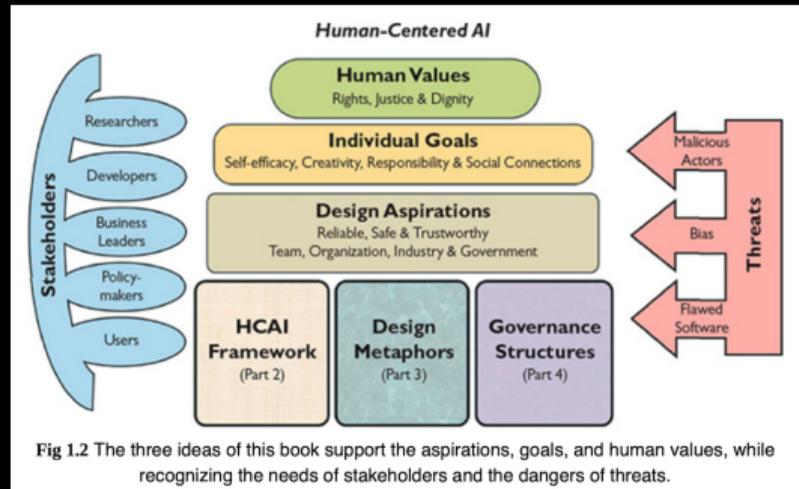
# Human-Centred AI Structure

Three big ideas:

- HCAI framework: guide human-centric thinking for creative design
- design metaphors: practical ideas for designing HCAI systems (*not white robots*)
- governance structures: practice steps to realise ethical principles

Support:

- design aspirations
- individual goals
- human values (most important!)



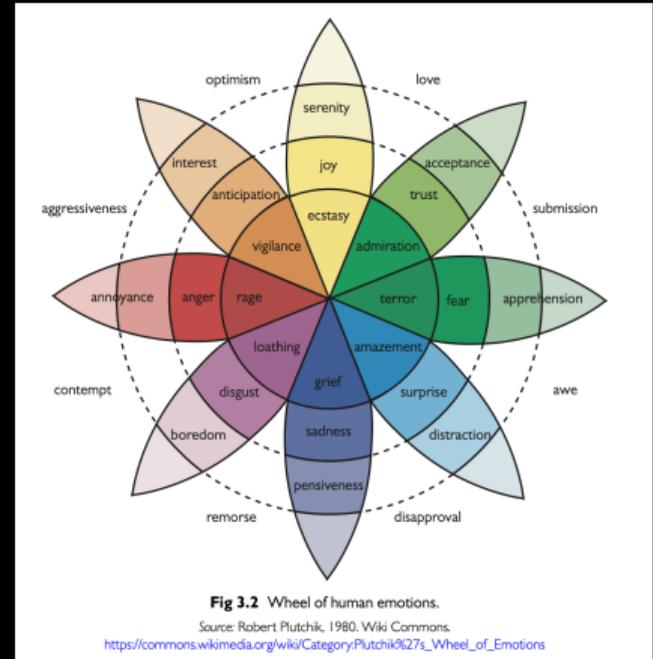
**Figure 5:** Map of Shneiderman (2022) view of Human-Centred AI

# People and Computers

Shneiderman (2022) argues against AI goals that equate computers with people.

- people are already good at lots of things (celebrate that)
- argues to change from simulating human behaviour to enhancing it.
- As an example:
  - people are good at emotions, and sensitive to them.
  - people can be disturbed by emotional simulations in robots

Rather than anthropomorphic systems: argues for Creativity Support Tools (CSTs) (see *Expressive Interaction*), direct manipulation interfaces (see *Interfaces*)



**Figure 6:** The *wheel* of emotions.

# HCAI frameworks

---

# HCAI frameworks



**Figure 7:** All systems need a framework to hold it up...

## Why start with a framework?

Interaction designers love good frameworks.

Two good reasons to investigate a framework for HCAI.

1. Understand different capabilities of present and near future systems.
2. See opportunities (gaps) for new design ideas, or different versions of the same idea.

## Rising above levels of automation

- Shneiderman (2022) critiques standard models of autonomy/automation by computers.
- e.g., Parasuraman et al. (2000) model of levels of autonomy.
- typically, autonomy is seen as 1D spectrum
- so if the computer is autonomous, the human does nothing.

What if both computer and human can have autonomy?

---

level

---

10

the computer...

---

(High) decides everything and acts autonomously, ignoring the human informs the human only if the computer decides to informs the human only if asked executes

7

automatically, then necessarily informs

## Reverse Centaurs and Bad Kinds of Automation

Cory Doctorow recently writes about “reverse centaurs” as a model for bad kinds of automation.

The idea is:

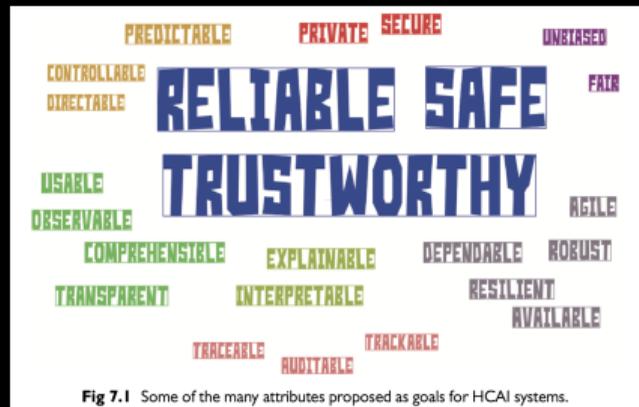
- a **centaur**: a powerful steed (the horse) with the head of an (intelligent and knowledgeable) person.
- a **reverse centaur**: a person’s (slow and weak) body with the head of a (much lower intelligence) animal.

Examples:

- freelance writer put under ridiculous deadlines acting as an “accountability sink” for chatGPT.
- vibe-coding a whole app: “checking” for bugs is slow and hard and puts responsibility on the vibe-coder to find problems

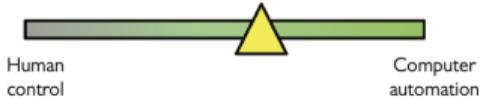
# Defining Reliable, Safe, Trustworthy Systems

- **reliable:** produce expected responses when needed
  - comes from good *engineering*, verification, validation
  - technical audit trails when things go wrong
- **safe:** social/cultural commitment
  - reporting failures and near misses
  - review of problems and solutions
- **trustworthy:** what *people* want from a system
  - independent oversight (certification, regulation, insurance)
  - (communication and understandable context)



**Figure 8:** Attributes of reliable, safe, and trustworthy systems.

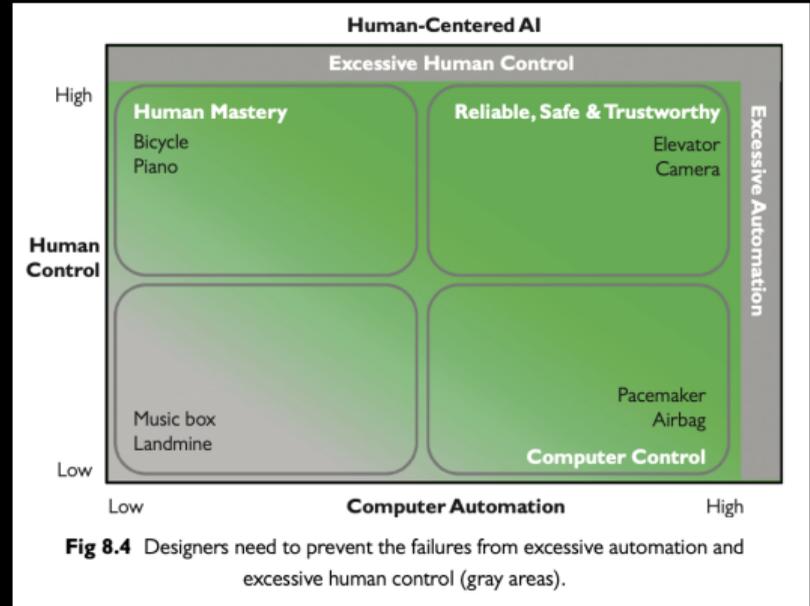
# Two Dimensional HCAI Framework



**Fig 8.1** Misleading one-dimensional thinking suggests that designers must choose between human control and computer automation.

## Figure 9: The 1D spectrum.

- existing 1D framework poses an unhelpful zero-sum tradeoff in automation
- **problem:** full automation seems like the only goal
- **2D framework:** computer and human control can both be useful in different contexts
- helps understand existing systems,



## Figure 10: Mapping systems and locating dangers

## Eight Design Guidelines for HCAI

Shneiderman, famous for the *Eight Golden Rules for Design* proposes a new set!

1. Overview first, zoom and filter, then details-on-demand
2. Preview first, select and initiate, then manage execution
3. Steer by way of interactive control panels
4. Capture history and audit trails from powerful sensors
5. People thrive on human-to-human communication
6. Be cautious when outcomes are consequential
7. Prevent adversarial attacks
8. Incident reporting websites accelerate refinement

# Processes

- **Overview first, zoom and filter, then details-on-demand:**
  - give users agency over information abstraction, make sure the details are available!
  - Relates to visualisation in particular, but many other kinds of information.
- **Preview first, select and initiate, then manage execution:**
  - for automated processes, users should have a preview of all steps, ways to select a plan, and to manage the execution. - works for navigation, parking, cameras (how about MS Copilot?)
- **Steer by way of interactive control panels:**
  - give users ways to steer processes. “control panels” can be joysticks, sliders, gestural controls.
  - this may be a good place for *expressive interaction* knowledge to help map simple controls to complex outcomes.
  - envisages AR/VR as good way to assist here.

# History, Sharing, and Hazards

- **Capture history and audit trails from powerful sensors:**
  - Cars and aeroplanes capture lots of data about usage to help diagnose issues when things go wrong, so can apps.
  - Problem: could this lead to excessive data collection and privacy issues.
- **People thrive on human-to-human communication:**
  - people like to share content and experiences
  - collaborative work makes us better.
  - consider: buzz about chatGPT relates to gurus/power-users sharing tips.
- **Be cautious when outcomes are consequential:**
  - humility is a good attribute for designers
  - thorough evaluation and continuous monitoring more important with high risk
  - (how do LLMs manage this? poorly so far!)

# Responsibility and Feedback

- **Prevent adversarial attacks:**
  - failures from vandals or baddies disrupt normal use and cause badness in the world.
  - how can we prevent these issues? Security and governance but also design.
- **Incident reporting websites accelerate refinement:**
  - openness to feedback leads to technology improvement
  - something like “bug bounties” for HCAI systems

# Design Metaphors

---

# Design Metaphors

Metaphors useful for clarifying design ideas.

- Steve Jobs famously called computers “bicycles for the mind”
- Shneiderman (2022) looks for metaphors to dig further into the two dimensions of the HCAI framework
- starting point are *science* and *innovation*
  - **science:** understand principles that make intelligent behaviour possible in a computer
  - **innovation:** create computers that amplify human abilities

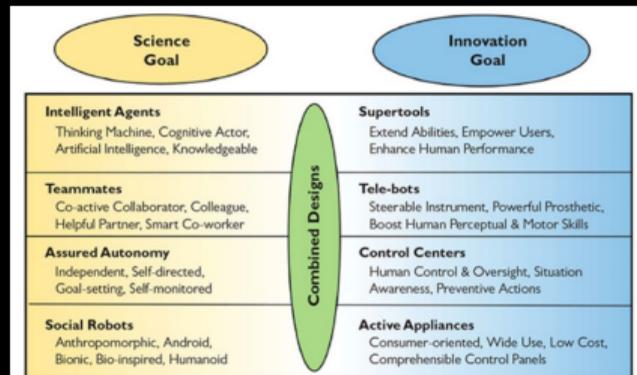


Fig. 11.1 Terminology and design metaphors for the science goal and innovation goal, with the possibility of combined designs that take the useful features from each goal.

**Figure 11:** combining science and innovation goals to find design inspiration

## Science vs Innovation Goals

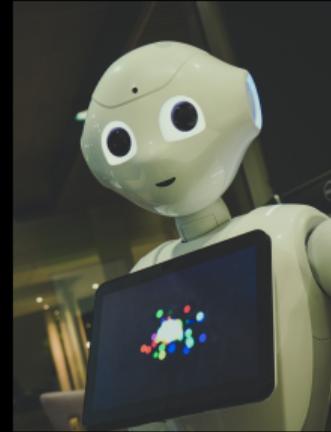
Science and innovation goals can come into conflict. Significant amounts of design sometimes needed to turn a scientific output into a product.

### **Autonomous social robots**

- Science goals: a general purpose robot elders, parcel delivery, etc.
- Innovation goals: tune solution for each context of use.

### **Online meeting services**

- Science: devices / software support collaboration.
- Innovation: Microsoft Teams, Zoom, Google meet.



**Figure 12:** Pepper is good for demos, but what context of use does it support? Photo by Owen Beard on Unsplash.

# Intelligent Agents and Supertools

- early computers were labeled “thinking machines” and “electronic brains”.
- Martin (1993) argues that such myths slowed workplace acceptance and created unrealistic expectations
- Turing (1950) asked “Can Machines Think?”
  - Envisioned machines competing with humans in intellectual fields
- debate: a deceptive metaphor rather than a meaningful test (Natale, 2021)
- pop culture influence of powerful thinking machines and robots



**Figure 13:** HAL 9000 in movie “2001: A Space Odyssey.” Image Source

## Towards supertool designs

*combine intelligent agents with human-controlled tools to ensure trust, usability, and control*

- human-computer symbiosis explored from 60s: humans make decisions; computers handle routine tasks (Licklider, 1960)
- human-AI design debate: automation vs direct user control (Shneiderman & Maes, 1997).
- Maes: Advocated for proactive software agents that anticipate user needs.

AI in conferences & applications:

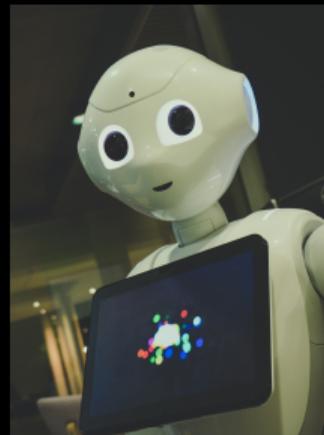
- science Goal (AI researchers): automated performance (e.g., self-driving cars, reading x-rays).
- innovation Goal (HCAI researchers): promote human-supervised tools (e.g., SIGCHI, UXPA, Augmented Humans conferences)

## Teammates and Tele-bots

- **Teammates:** modelling human-robot interaction after human-human interaction
- **Telebots:** remote-controlled systems with high degrees of automation but ultimately human control

Teammate approach has many challenges (Klien et al., 2004):

- Unmet expectation to disappointment.
- False beliefs about robot autonomy/responsibility.
- Emotional attachment can mislead usage.



**Figure 14:** Is Pepper useful as a teammate?

## Designs for Tele-bots

- Human operators remain accountable (e.g., remote surgery)
- Leverage what machines do better, don't copy humans
- Responsive controls (e.g., DaVinci surgical system)
- Enable users to fix, personalise, provide feedback for future design iterations
- incorporating human creativity
  - tools should support human innovation, not replace it.
  - enable users to fix, personalise, provide feedback for future design iterations.



Fig 14.1 DaVinci Surgical System from Intuitive Surgical

Source: <http://www.intusurg.com>

**Figure 15:** The DaVinci surgical robot is a successful telebot application

**Question:** is agentic coding AI (e.g., CoPilot) a telebot or teammate?

# Assured Autonomy and Control Centers

**Autonomy:** delegation to an authorised entity to take action within specific boundaries (David & Nielsen, 2016)

- autonomous systems are hard and interesting to design
- but when autonomy falls short, frustration is enhanced
- problems can include: reduction in human attention and awareness
- ironically, workload can **increase**, more vigilance is required

## Assured Autonomy

- “assured” with HCAI, formal methods for proving correctness, testing and



**Figure 16:** Tesla Autopilot Crash, 2018. Image Source.

# Control centres and supervised autonomy

- Control panels and remote control centers.
- Visual monitoring, audit trails, and feedback.
- Retrospective analysis of failure data.

## Examples:

- Aviation: pilots, co-pilots, TRACON, ARTCC, FAA certification, NASA Rovers.
- Healthcare: ICU monitoring.
- Social media / e-commerce: alerts, feedback, interlocks.

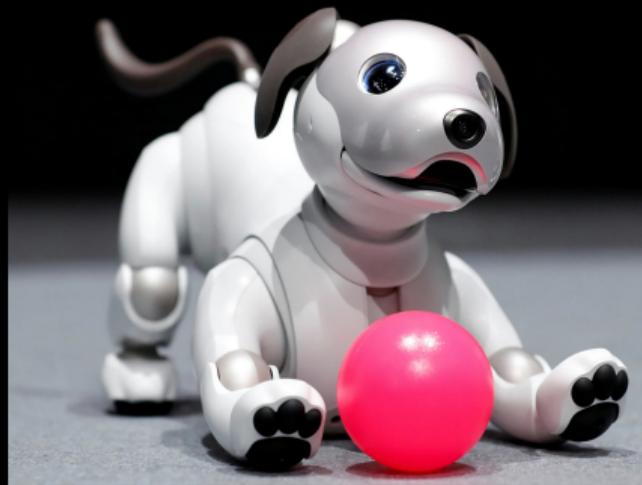


**Figure 17:** Supervised autonomous systems should have some kind of control, with lots of feedback and advanced interfaces

## Social Robots and Active Appliances

What kind of AI systems should we live with?

- Human-like figures from stories and myths (e.g., golem, Frankenstein's creation)
- Repeated failures: gimmicky, unrealistic addressing real needs.
- Animal robots (e.g., Sony AIBO) offer emotional connection without promising intelligence
- Simpler, clearer use cases (e.g., therapy, companionship).
- Avoids ethical concerns tied to humanoid robots.



**Figure 18:** AIBO. Image Source.

# Practical Robots in Daily Life

Good example of a successful household robot:  
**washing machine.**

- dishwashers, pool cleaners, security systems
- integration of sensors, automation, and machine learning.
- improved user interfaces (e.g., smartphone apps, touchscreens)
- increasingly robot-like without anthropomorphic form

## Challenges:

- lack of standardised controls across devices.
- interface consistency, especially for:



**Figure 19:** The promise of active appliances

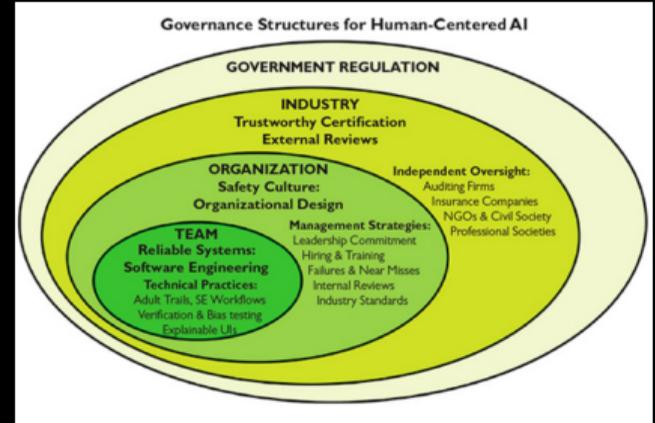
# Governing Structure

---

# Governing Structure

*What we need to govern is the human application of technology, and what we need to oversee are the human processes of development, testing, operation, and monitoring. (Bryson, 2020)*

- Reliable systems through software engineering teams
- Safety culture in organisations
- Trustworthy certification via industry oversight
- Regulation by government agencies



**Figure 20:** Four levels of governance structures for human-centered AI.

## Traditional AI vs. Human-Centered

---

### Traditional AI

Build AI algorithms and systems, stressing the autonomy of machines

Research human behaviour study and emulation

Benchmarks on algorithm performance

Less engagement with end users in early design phases

Success defined by technical functionality

Priority on machine intelligence and control

### Human-Centered AI (HCAI)

Emphasises human autonomy through well-designed user interfaces

Build on AI foundations to amplify, augment, and enhance human performance

Elevates human performance and satisfaction

Advocates user-centered participatory design by engaging diverse stakeholders

Success defined by how well systems support human goals, activities, and values

Ensures meaningful human control and values customer and consumer needs

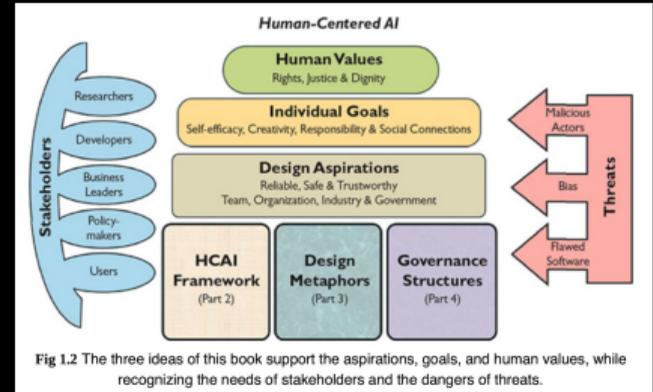
# Takeaways from Human-Centred AI

- tensions between the scientific goals of AI/ML and innovation goals of HCI community
- possible to bring together scientific and innovation goals to develop new designs
- autonomy does not have to be a zero-sum 1D spectrum
- pop-culture reflects us, not always a good predictor for best ideas

Alan Kay (designed first windowed UI at Xerox PARC) said:

*the best way to predict the future is to invent it*

so go do that!



**Figure 21:** Map of Shneiderman (2022) view of Human-Centered AI

# Course Wrap-up and Revision

---

# Course Wrap-up and Revision

What was this course about again?

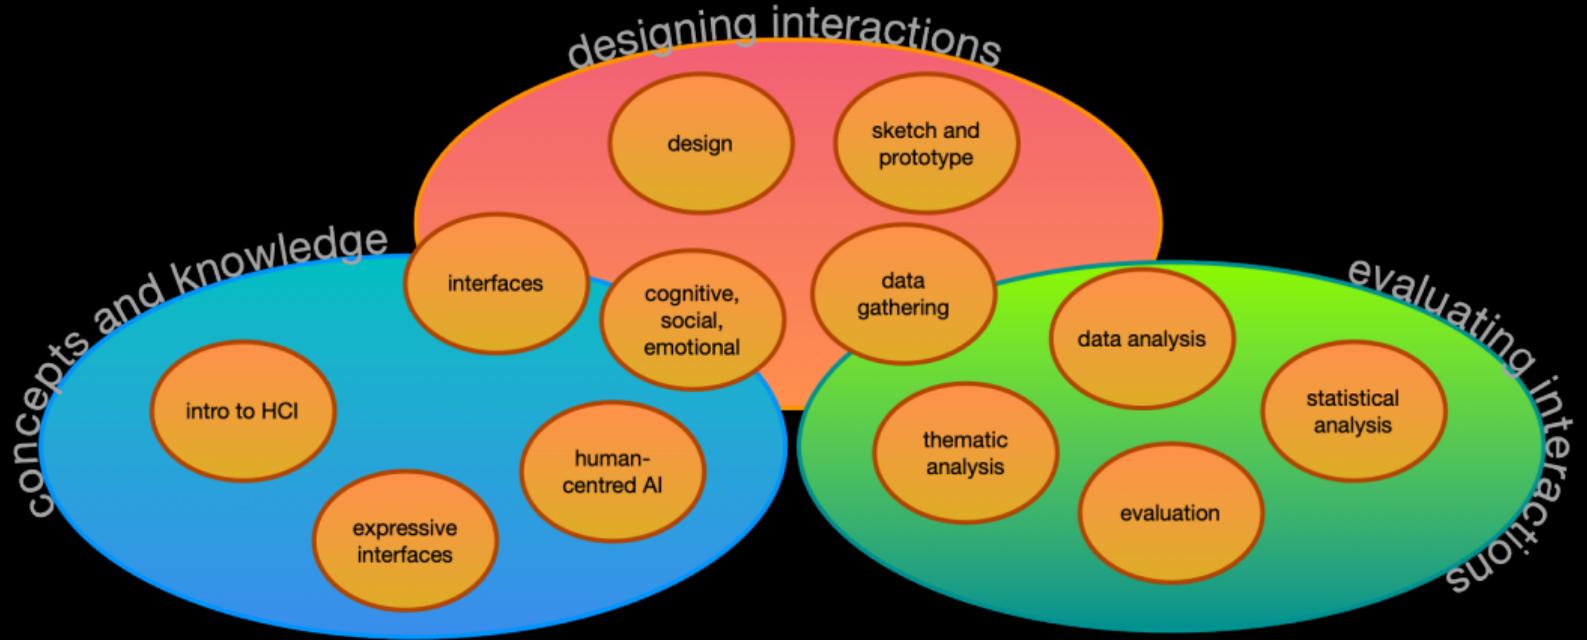


**Figure 22:** from this...



**Figure 23:** to this...

# core concepts, design, and evaluation



**Figure 24:** A map of concepts in our HCI course with loose-ish connections.

## stakeholders and perspectives

	designer	programmer	business manager	researcher
<b>concepts and knowledge</b>	"what concepts are the basis for my decisions?"	"which parts of software are exposed to user?"	"what heuristics will ensure success?"	"can established knowledge be challenged?"
<b>designing interactions</b>	"what processes should I follow?"	"how will requirements be established?"	"how can I iterate or pivot to success?"	"how can I find solutions to my research problem?"
<b>evaluating interactions</b>	"how do I know the	"how can I test user	"what signals tell me if	"how to balance

## HCI history: The “wave” theory

- First wave: Human Factors in Computing (1980-1992) (Chignell et al., 2023)
  - Optimising/measuring efficiency, cognitive psychology approach, studying individual users. Lab setting.
- Second wave: Cognitive revolution — mind and computer coupled (1992-2006) (Kaptelinin et al., 2003)
  - Optimise interactions, hypothesis testing, affordances, activity theory, user-centred design. Work environments.
- Third wave: Situated perspectives (Bødker, 2015) (2006-)
  - Consumer tech, participation and sharing, pervasive computing, AR, tangible interaction, home environments,
- Fourth wave: Entanglement HCI (Frauenberger, 2019) (2019-)
  - Computer and humans entangled in society: focus on values, accessibility, diversity, policy, law, ethics, individuals' and society's responsibilities

# Usability Goals (Interaction Design, Beyond HCI)

- Effective to use (effectiveness)
- Efficient to use (efficiency)
- Safe to use (safety)
- Having good utility (utility)
- Easy to learn (learnability)
- Easy to remember how to use (memorability)

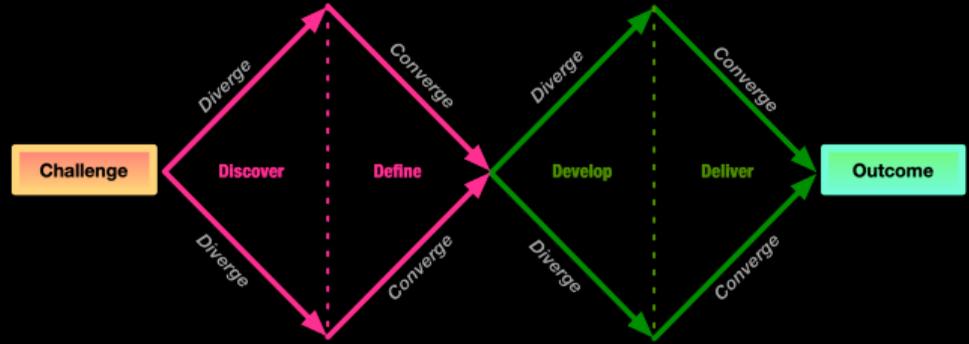
This version from: (Rogers et al., 2023)



**Figure 25:** Raffaele et al. (2016), illustrating Rogers et al. (2023)

# Design Stages

1. *Discover*: understand the problem and the people affected
2. *Define*: define the problem clearly so that it can be addressed
3. *Develop*: create ideas, prototypes, sketches, etc, that might address the problem
4. *Deliver*: test potential solutions to find promising directions, and iterate

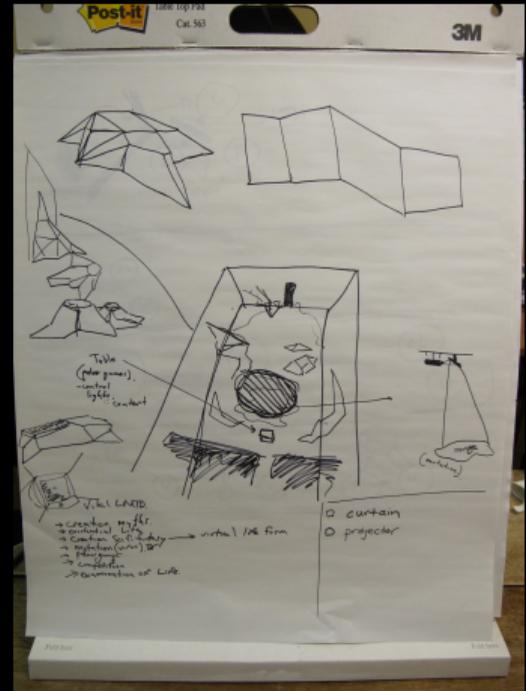


**Figure 26:** The double diamond model of design (adapted from Design Council, 2025)

# What is a sketch?

- quick
- timely
- inexpensive
- disposable
- plentiful
- clear vocabulary
- distinct gesture
- minimal detail
- appropriate degree of refinement
- suggest and explore rather than confirm

(Buxton (2007), p.111-113)



**Figure 27:** Sketching a stage in 2010

# What is a prototype

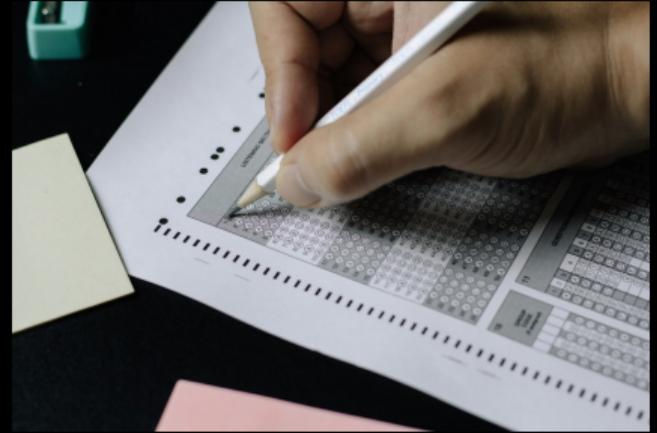
- “*primitive form*”
- the form that comes before... something.
- in this context:
  - a *testable* form
  - a form we can *experience*
- enables evaluation and iteration
- *primitive*: should be somehow rough or limited



**Figure 28:** A prototype AI musical instrument.

## Data Gathering: Interviews and Questionnaires

- **Interview techniques:** structured, semi-structured, open
- **Questionnaires:** closed, open, rating scale questions
- **Established questionnaires:** Software Usability Survey (SUS), NASA Task Load Index (TLX), Creativity Support Index (CSI)
- **DIY questionnaires** can be tricky to do well!
- All useful, but need to be justified
- Require different types of analysis, both can involve quantitative and qualitative.



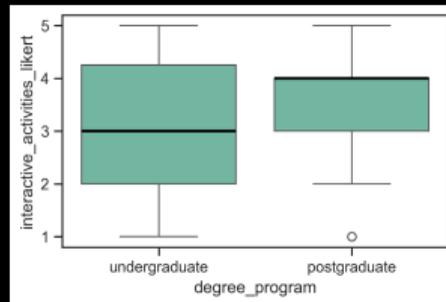
**Figure 29:** Photo by Nguyen Dang Hoang Nhu on Unsplash

# Basic Quantitative Analysis

- **descriptive statistics**
  - minimum, maximum
  - lower and upper quartile
  - median and mean
  - number of data points (count)
- **plot distribution**
  - **scatter plot:** see all the data! good for checking outliers and comparing aspects of data
  - **boxplot:** useful to compare distributions clearly **charles approved plot!**

These approaches may be enough to make clear research findings!

stat	interactive activities
min	1
25%	2
50%	3
75%	4
max	5

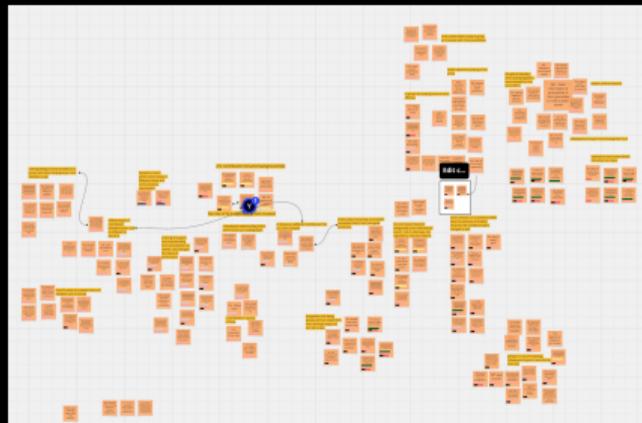


# Basic Qualitative Analysis: Thematic Analysis

Lots of qualitative techniques but our focus is (Reflexive) Thematic Analysis (RTA) (Braun & Clarke, 2022), a well-known and accessible methodology.

1. familiarise with the data
2. coding (short labels, multiple rounds)
3. generating initial themes (higher level than codes)
4. developing, reviewing, and refining themes

Your themes become the findings of your qualitative analysis.



**Figure 31:** A Miro board from Yichen Wang's thematic analysis (2025)

## Thin vs thick themes

There are different types of themes, and a common distinction:

- Themes that categorise groups of codes: bucket themes, semantic themes, *thin* themes
- Themes that interpret the codes, revealing hidden information: latent themes, *thick* themes

Charles (2025; i.e., these slides!) suggests that **4** is a key heuristic for assessing theme thickness. (Disclaimer: may be revised in future!)

Number of words heuristic:  
*If your theme is <4 words, it might be a bit thin.*

Number of themes heuristic:  
*If you are proposing >4 themes, they might be a bit thin.*

Source: Charles, 2025. 🤖

## 45 years of interface types!

- Command Line
- Graphical
- Multimedia
- Virtual reality
- Web
- Mobile
- Appliance
- Voice
- Pen
- Touch
- Touchless
- Haptic
- Multimodal
- Shareable
- Tangible
- Augmented reality
- Wearables
- Robots and drones
- Brain-computer
- Smart
- Shape-changing
- Holographic

# Cognitive, Social, Emotional Interaction

Cognitive Processes (Eysenck & Brysbaert, 2023):

1. Attention
2. Perception
3. Memory
4. Learning
5. Reading, speaking, listening
6. Problem solving, planning, reasoning, decision making

Social and Emotional aspects

- conversation (face-to-face vs remote)
- co-presence
- Emotions and behaviour relate
- Models of emotional design
- Affective Computing and Emotional AI
- Persuasive Technologies
- Anthropomorphism

# Developing an evaluation plan

- Evaluation Goal/Aims
- Participants
- Setting
- Data to collect
- Methods
- Ethical Considerations and Consent
- Data capture, recording, storage
- Analysis method
- Output(s) of evaluation process



**Figure 32:** How to evaluate this app?

# Statistical Analysis and Significance Testing?

- going beyond descriptive statistics...
- **significance testing:** quantifying differences in mean
  - $t$ -tests: for comparing two means
  - ANOVA: for comparing 3+ means, incl. repeated measures
  - non-parametric alternatives: Mann-Whitney  $U$ , Wilcoxon signed ranks
  - $\chi^2$  test: comparing categorical data
- correlation analysis
- regression

One-way ANOVA:

```
from scipy.stats import f_oneway
import statsmodels.api as sm
from statsmodels.formula.api import ols

# group by 'independent' column and compare
groups = [group['dependent'].values for group in data.groupby('independent')]
f_stat, p_value = f_oneway(*groups)
```

```
# create a Model from a formula and data
model = ols('dependent ~ C(independent)', data).fit()
anova_table = sm.stats.anova_lm(model, 1)
```

Factorial ANOVA:

```
# factorial anova: example effects 50f
```

## What is an expressive interaction?

Mapping sensed **gestures** to an expressive output that is fed back to the user.

- **gestures:** the use of motions by the limbs or body as a means of expression
- can be unintentional, control, or ancillary gestures
- from non-human actors (e.g., the movement of a leaves on a branch of a tree)
- *“any sort of motion, that may be understood as an expression of something”*

The interaction itself is expressive, and the output is an expression as well. We consult *Composing Interactions* (Baalman, 2022) as a



**Figure 33:** Sensing movement and touch to create music, Atau Tanaka performing in 2010 (Tanaka, 2010)

## Key challenges of HCI research

1. **problem finding:** computers are ubiquitous and everybody<sup>1</sup> seems to get along with them ok, so where are the problems? Where can we make impact to help people?
2. **design processes:** need them to lead to new contributions, not clones of existing designs. If you are making something new, how do you know what the requirements are? At any scale, design and prototyping is expensive.
3. **evaluation:** working with people is time-consuming and uncertain. Data analysis can easily take months for a research paper.
4. **communication:** after doing *all this work* (see above), there can be a lot of explain in a paper! Clear communication is a huge challenge. Reviewers love to poke holes in any and all aspect of a design and evaluation process.

Conclusion: HCI research is kinda hard!

<sup>1</sup>everybody to some value of people we know about

# So long, and thanks for the all the fish!

We got through a lot this semester!  
Thanks for coming on this journey with me!  
Good luck with your final projects and your  
other assessments this semester!



**Figure 34:** Meet you *at the bar* for questions. 🍸🥤🍵☕  
Unfortunately no drinks served!



# References

---

## References i

- Baalman, M. (2022). *Composing interactions: An artist's guide to building expressive interactive systems*. V2\_Publishing.
- Bødker, S. (2015). Third-wave HCI, 10 years later—participation and sharing. *Interactions*, 22(5), 24–31. <https://doi.org/10.1145/2804405>
- Braun, V., & Clarke, V. (2022). *Thematic analysis: A practical guide*. Sage Publications.
- Bryson, J. j. (2020). The artificial intelligence of the ethics of artificial intelligence: An introductory overview for law. *The Oxford Handbook of the Ethics of Artificial Intelligence*.

## References ii

- Buxton, B. (2007). *Sketching user experiences: Getting the design right and the right design*. Morgan Kaufmann Publishers Inc.
- Chignell, M., Wang, L., Zare, A., & Li, J. (2023). The evolution of HCI and human factors: Integrating human and artificial intelligence. *ACM Trans. Comput.-Hum. Interact.*, 30(2). <https://doi.org/10.1145/3557891>
- David, R. A., & Nielsen, P. (2016). *Defense science board summer study on autonomy*.
- Design Council. (2025). *The double diamond: A universally accepted depiction of the design process*. <https://www.designcouncil.org.uk/our-resources/the-double-diamond/>

## References iii

- Eysenck, M. W., & Brysbaert, M. (2023). *Fundamentals of cognition (4th ed.)*. Routledge. <https://doi.org/10.4324/9781003384694>
- Frauenberger, C. (2019). Entanglement HCI the next wave? *ACM Trans. Comput.-Hum. Interact.*, 27(1). <https://doi.org/10.1145/3364998>
- Kaptelinin, V., Nardi, B., Bødker, S., Carroll, J., Hollan, J., Hutchins, E., & Winograd, T. (2003). Post-cognitivist HCI: Second-wave theories. *CHI '03 Extended Abstracts on Human Factors in Computing Systems*, 692–693. <https://doi.org/10.1145/765891.765933>

## References iv

- Klien, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., & Feltovich, P. J. (2004). Ten challenges for making automation a "team player" in joint human-agent activity. *IEEE Intelligent Systems*, 19(6), 91–95. <https://doi.org/10.1109/MIS.2004.74>
- Licklider, J. C. (1960). Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics*, 1, 4–11.
- Martin, C. D. (1993). The myth of the awesome thinking machine. *Commun. ACM*, 36(4), 120–133. <https://doi.org/10.1145/255950.153587>
- Natale, S. (2021). *Deceitful media: Artificial intelligence and social life after the turing test*. Oxford University Press.

## References v

- Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 30(3), 286–297.  
<https://doi.org/10.1109/3468.844354>
- Raffaele, R., Carvalho, B., Lins, A., Marques, L., & Soares, M. M. (2016). Digital game for teaching and learning: An analysis of usability and experience of educational games. In A. Marcus (Ed.), *Design, user experience, and usability: Novel user experiences* (pp. 303–310). Springer International Publishing.
- Rogers, Y., Sharp, H., & Preece, J. (2023). *Interaction design: Beyond human-computer interaction, 6th edition*. John Wiley & Sons, Inc.  
<https://quicklink.anu.edu.au/kv9b>

## References vi

- Shneiderman, B. (2022). *Human-centered AI*. Oxford University Press.  
<https://doi.org/10.1093/oso/9780192845290.001.0001>
- Shneiderman, B., & Maes, P. (1997). Direct manipulation vs. Interface agents. *Interactions*, 4(6), 42–61. <https://doi.org/10.1145/267505.267514>
- Tanaka, A. (2010). Mapping out instruments, affordances, and mobiles. *Proceedings of the International Conference on New Interfaces for Musical Expression*, 88–93. <https://doi.org/10.5281/zenodo.1177903>
- Turing, A. M. (1950). Computing machinery and intelligence (1950). *Mind*, 49, 433–460.