

Announcements

Presentation sign-ups due today (4/7).

Remember that signing up means listing a paper.

Tools for Thought

CIS 7000-001

Andrew Head & Danaé Metaxa

What is a tool for thought?

Explicit, fully committal definitions are rare in the HCI literature.

"Things that help us think."

Think: *define problems, solve them, understand the world, make decisions, do logic, formulate and review argument, or more generally*
produce new or refined knowledge from old.

Things: *objects, representations, notations, languages, interfaces, custom, methods [Appleton 2022]*

Didn't we already have this lecture?

If our definition is "things that help us think" then...

- **Graphs** are a tool for thought [viz lecture]
- **Logo** is a tool for thought [ed tech lecture]
- **Twitter** is a tool for thought [social media lecture]
- **Figma** is a tool for thought [content creation lecture]
- **The mouse** is a tool for thought [input/output lecture]

Yes!

So why do we have another lecture talking about it?

Tools for thought as a way of framing HCI outcomes

With a "tools for thought" framing, we aspire to increase cognitive capacity to **achieve new outcomes not possible before.**

This involves helping people take on **complex problems**, **comprehend complex situations**, and come up with **solutions** previously beyond their reach [Engelbart 1962].

How do we do this?

Today

Extended cognition

Undefined, non-linear problem solving

Reverse centaurs

Where are we today?

Extended Cognition


Augmenting human intellect

[Engelbart 1968]

In 1968, Douglas Engelbart performs the "Mother of All Demos" in Silicon Valley.

The demo showed off the fruits of the labor of the Augmentation Research Center (ARC): precursors to many computer tools for thought: documents, hyperlinks, diagramming, remote collaboration, etc.

The Mouse



Augmenting human intellect

[Engelbart 1962]

To augment human intellect, we alter and connect a number of components:

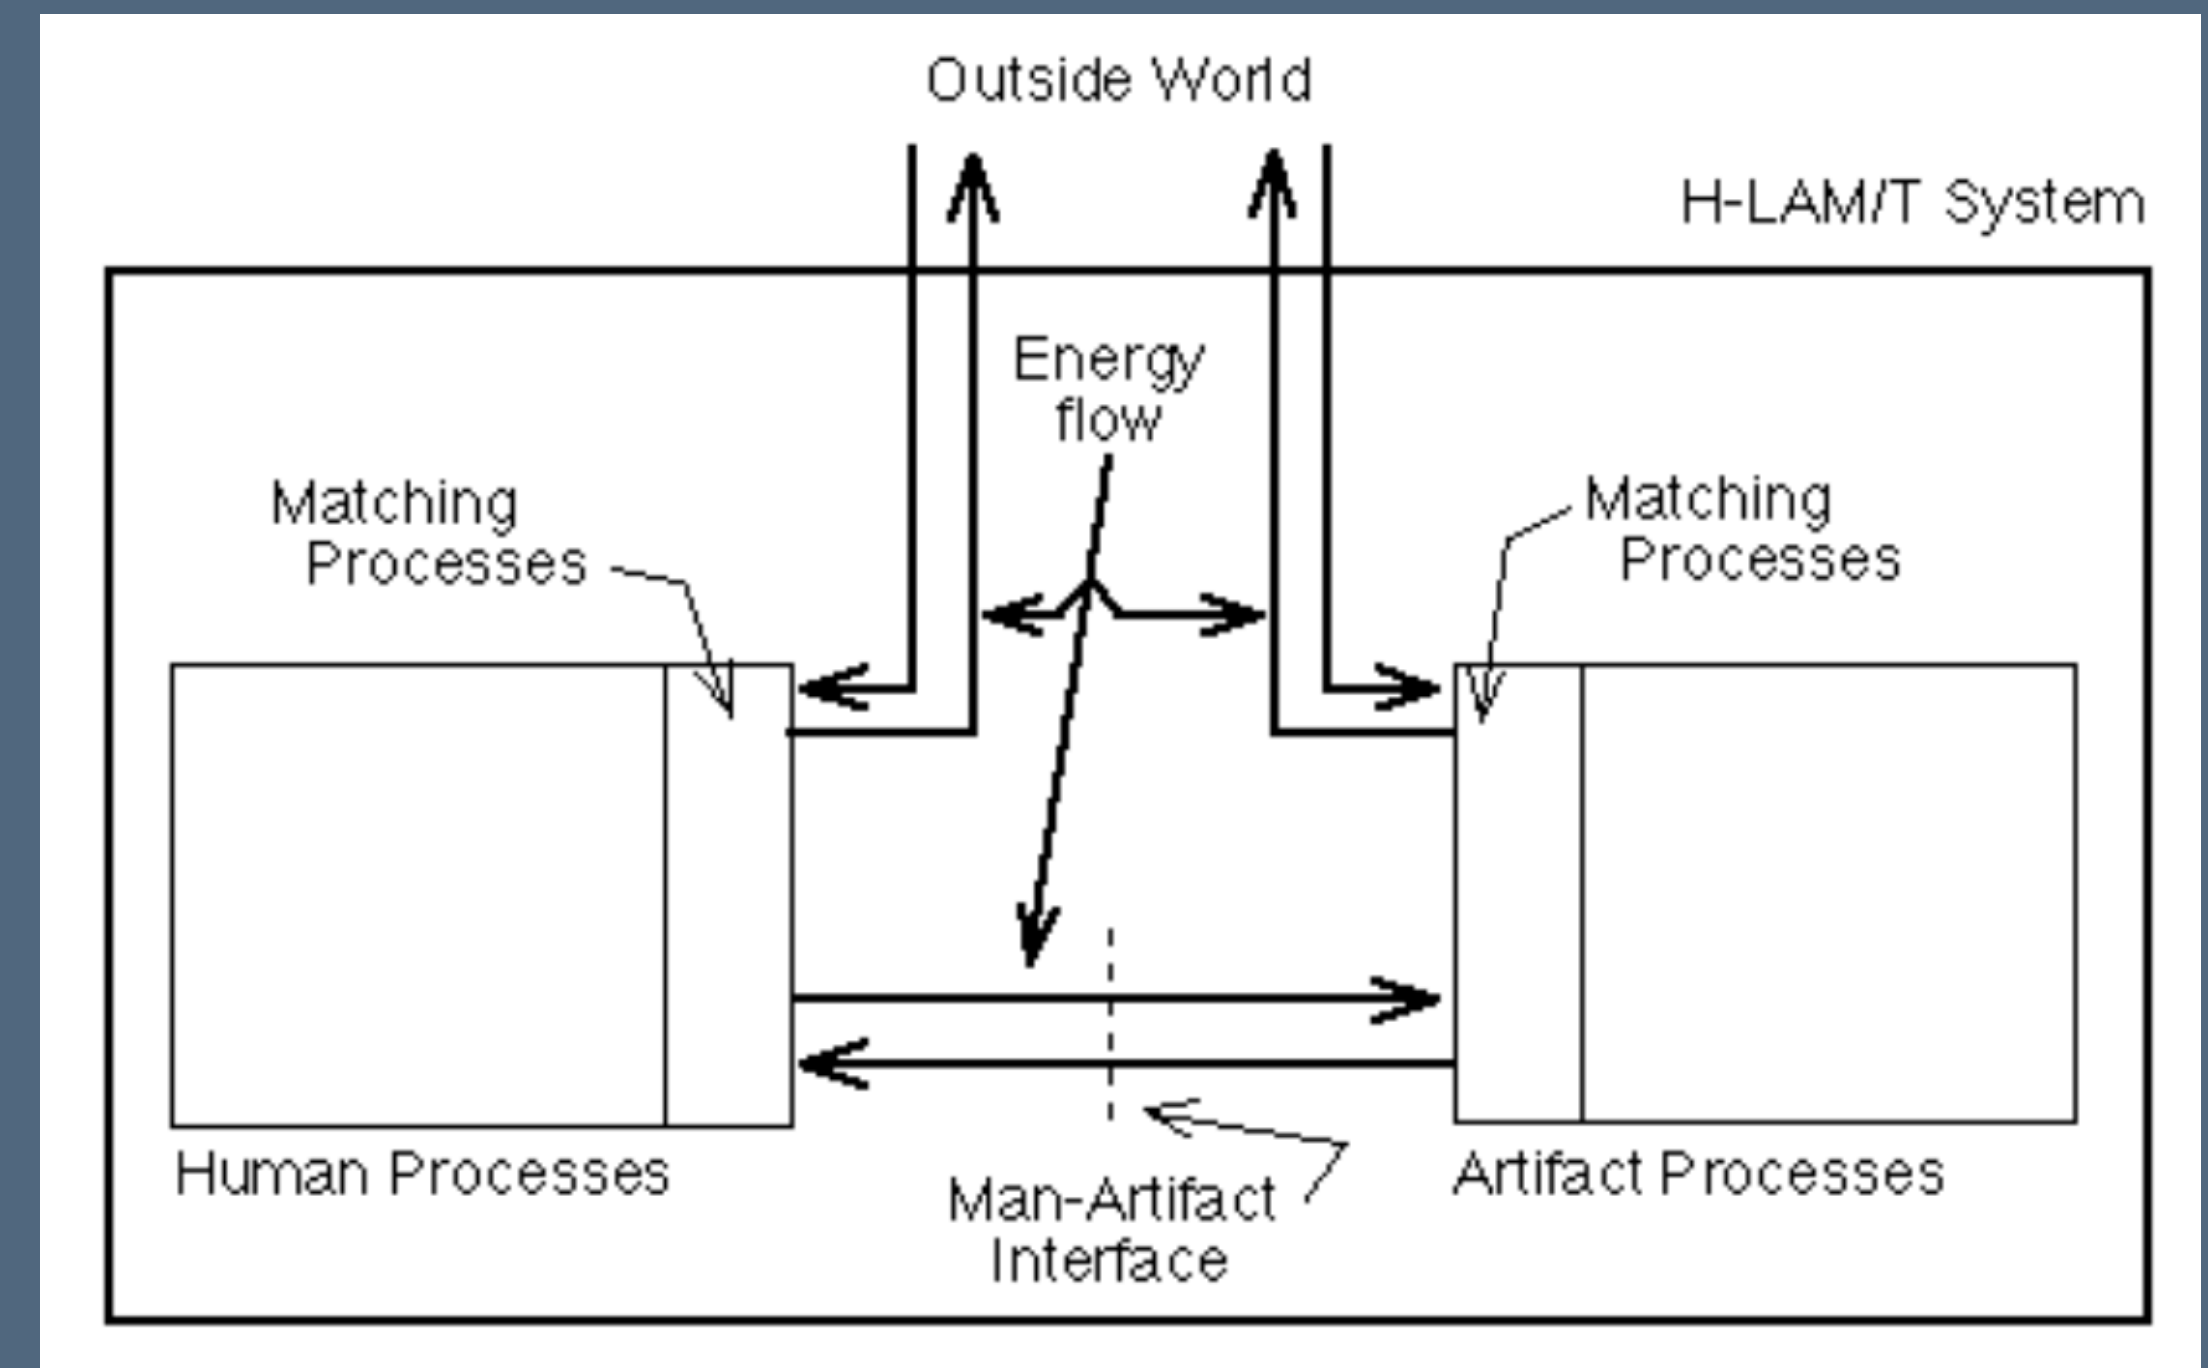
Human (H): self-explanatory

Artifacts (A): physical objects for manipulation of things, materials, and symbols

Language (L): how information travels between the world and the mind

Methodology (M): organization of activity

Training (T): conditioning for a human to develop skill with A, L, M.



What makes a good representation?

[Scaife and Rogers 1996]

Compute the following product:

$$\text{LXVII} \times \text{X} = ?$$

Try again with this one:

$$68 \times 10 = ?$$

Recall from before that picking the right representations can make cognitively demanding tasks trivial.

External cognition

[Scaife and Rogers 1996]

Scaife and Rogers name a few factors that make our graphical representations more effective:

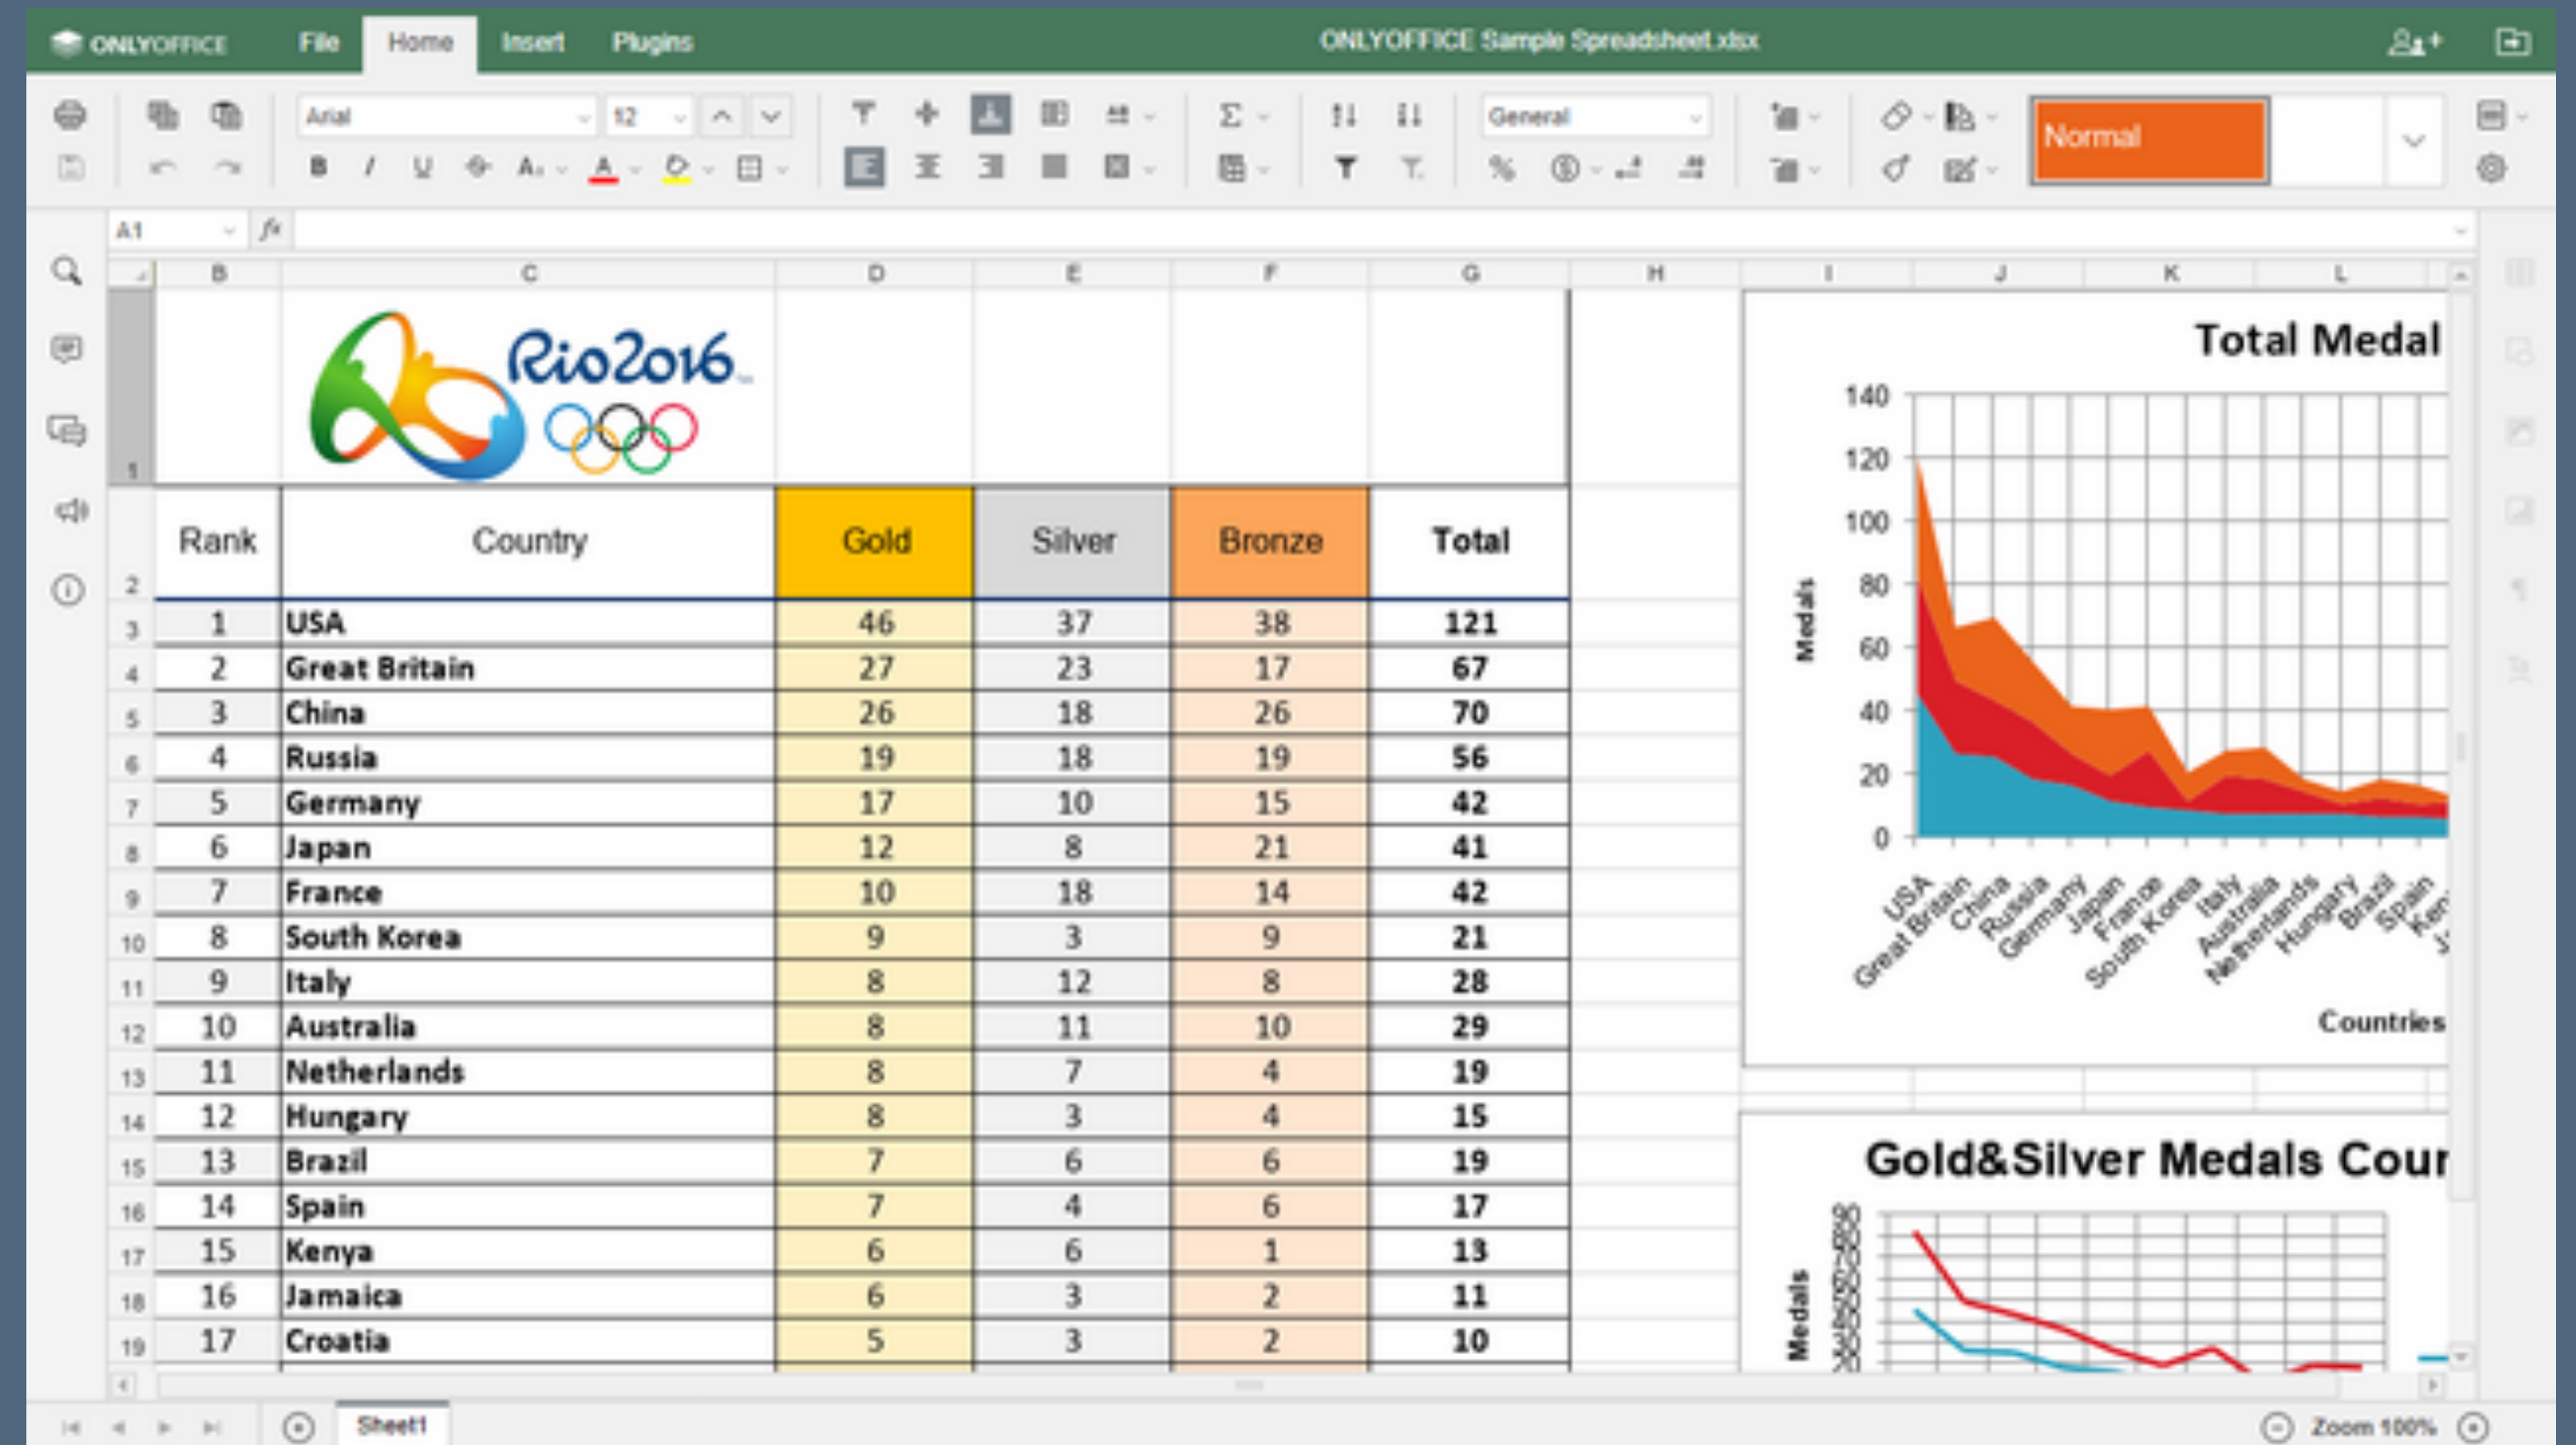
- **Computational offloading:** reducing the effort to solve a problem without affecting its underlying information content.
- **Re-representation:** casting a problem into a new representation that makes problem solving easier or harder.
- **Graphical constraining:** when visual aspects constrain the kinds of inferences that can be drawn about the represented world.

Computational offloading

[Scaife and Rogers 1996]

Reducing the effort to solve a problem without affecting its underlying information content.

Example: spreadsheet, maintains many computations that would otherwise need to be done individually.



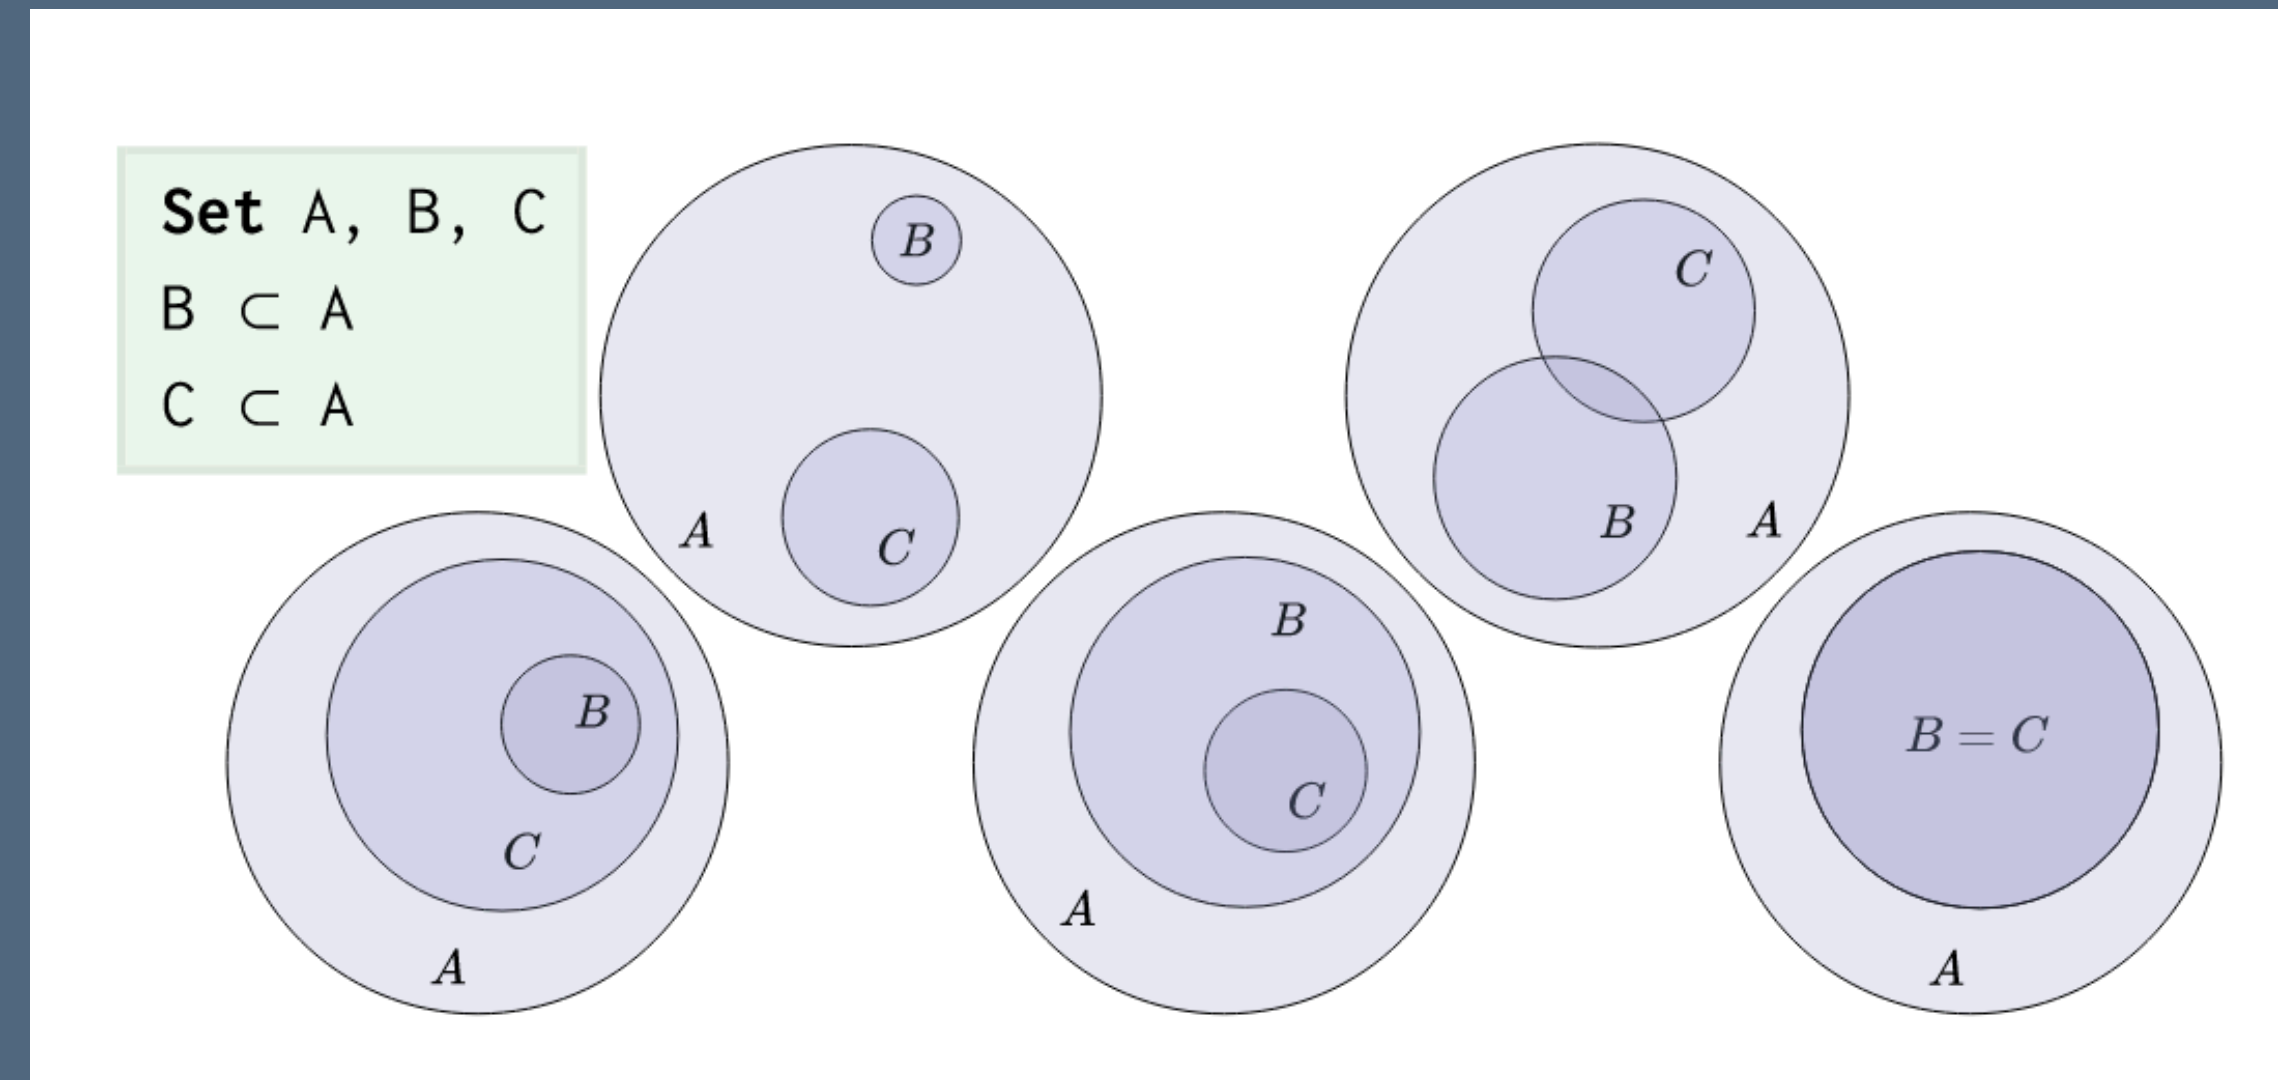
https://commons.wikimedia.org/wiki/File:OnlyOffice_spreadsheet_editor.png

Re-representation

[Scaife and Rogers 1996]

Casting a problem into a new representation that makes problem solving easier or harder.

Example: generating mathematical diagrams to convey relations between variables [Ye et al. 2020].




Re-representation

[Scaife and Rogers 1996]

Casting a problem into a new representation that makes problem solving easier or harder.

Example: visualizing hidden work of an LLM to invite inquiry into that work [Xie et al. 2024].

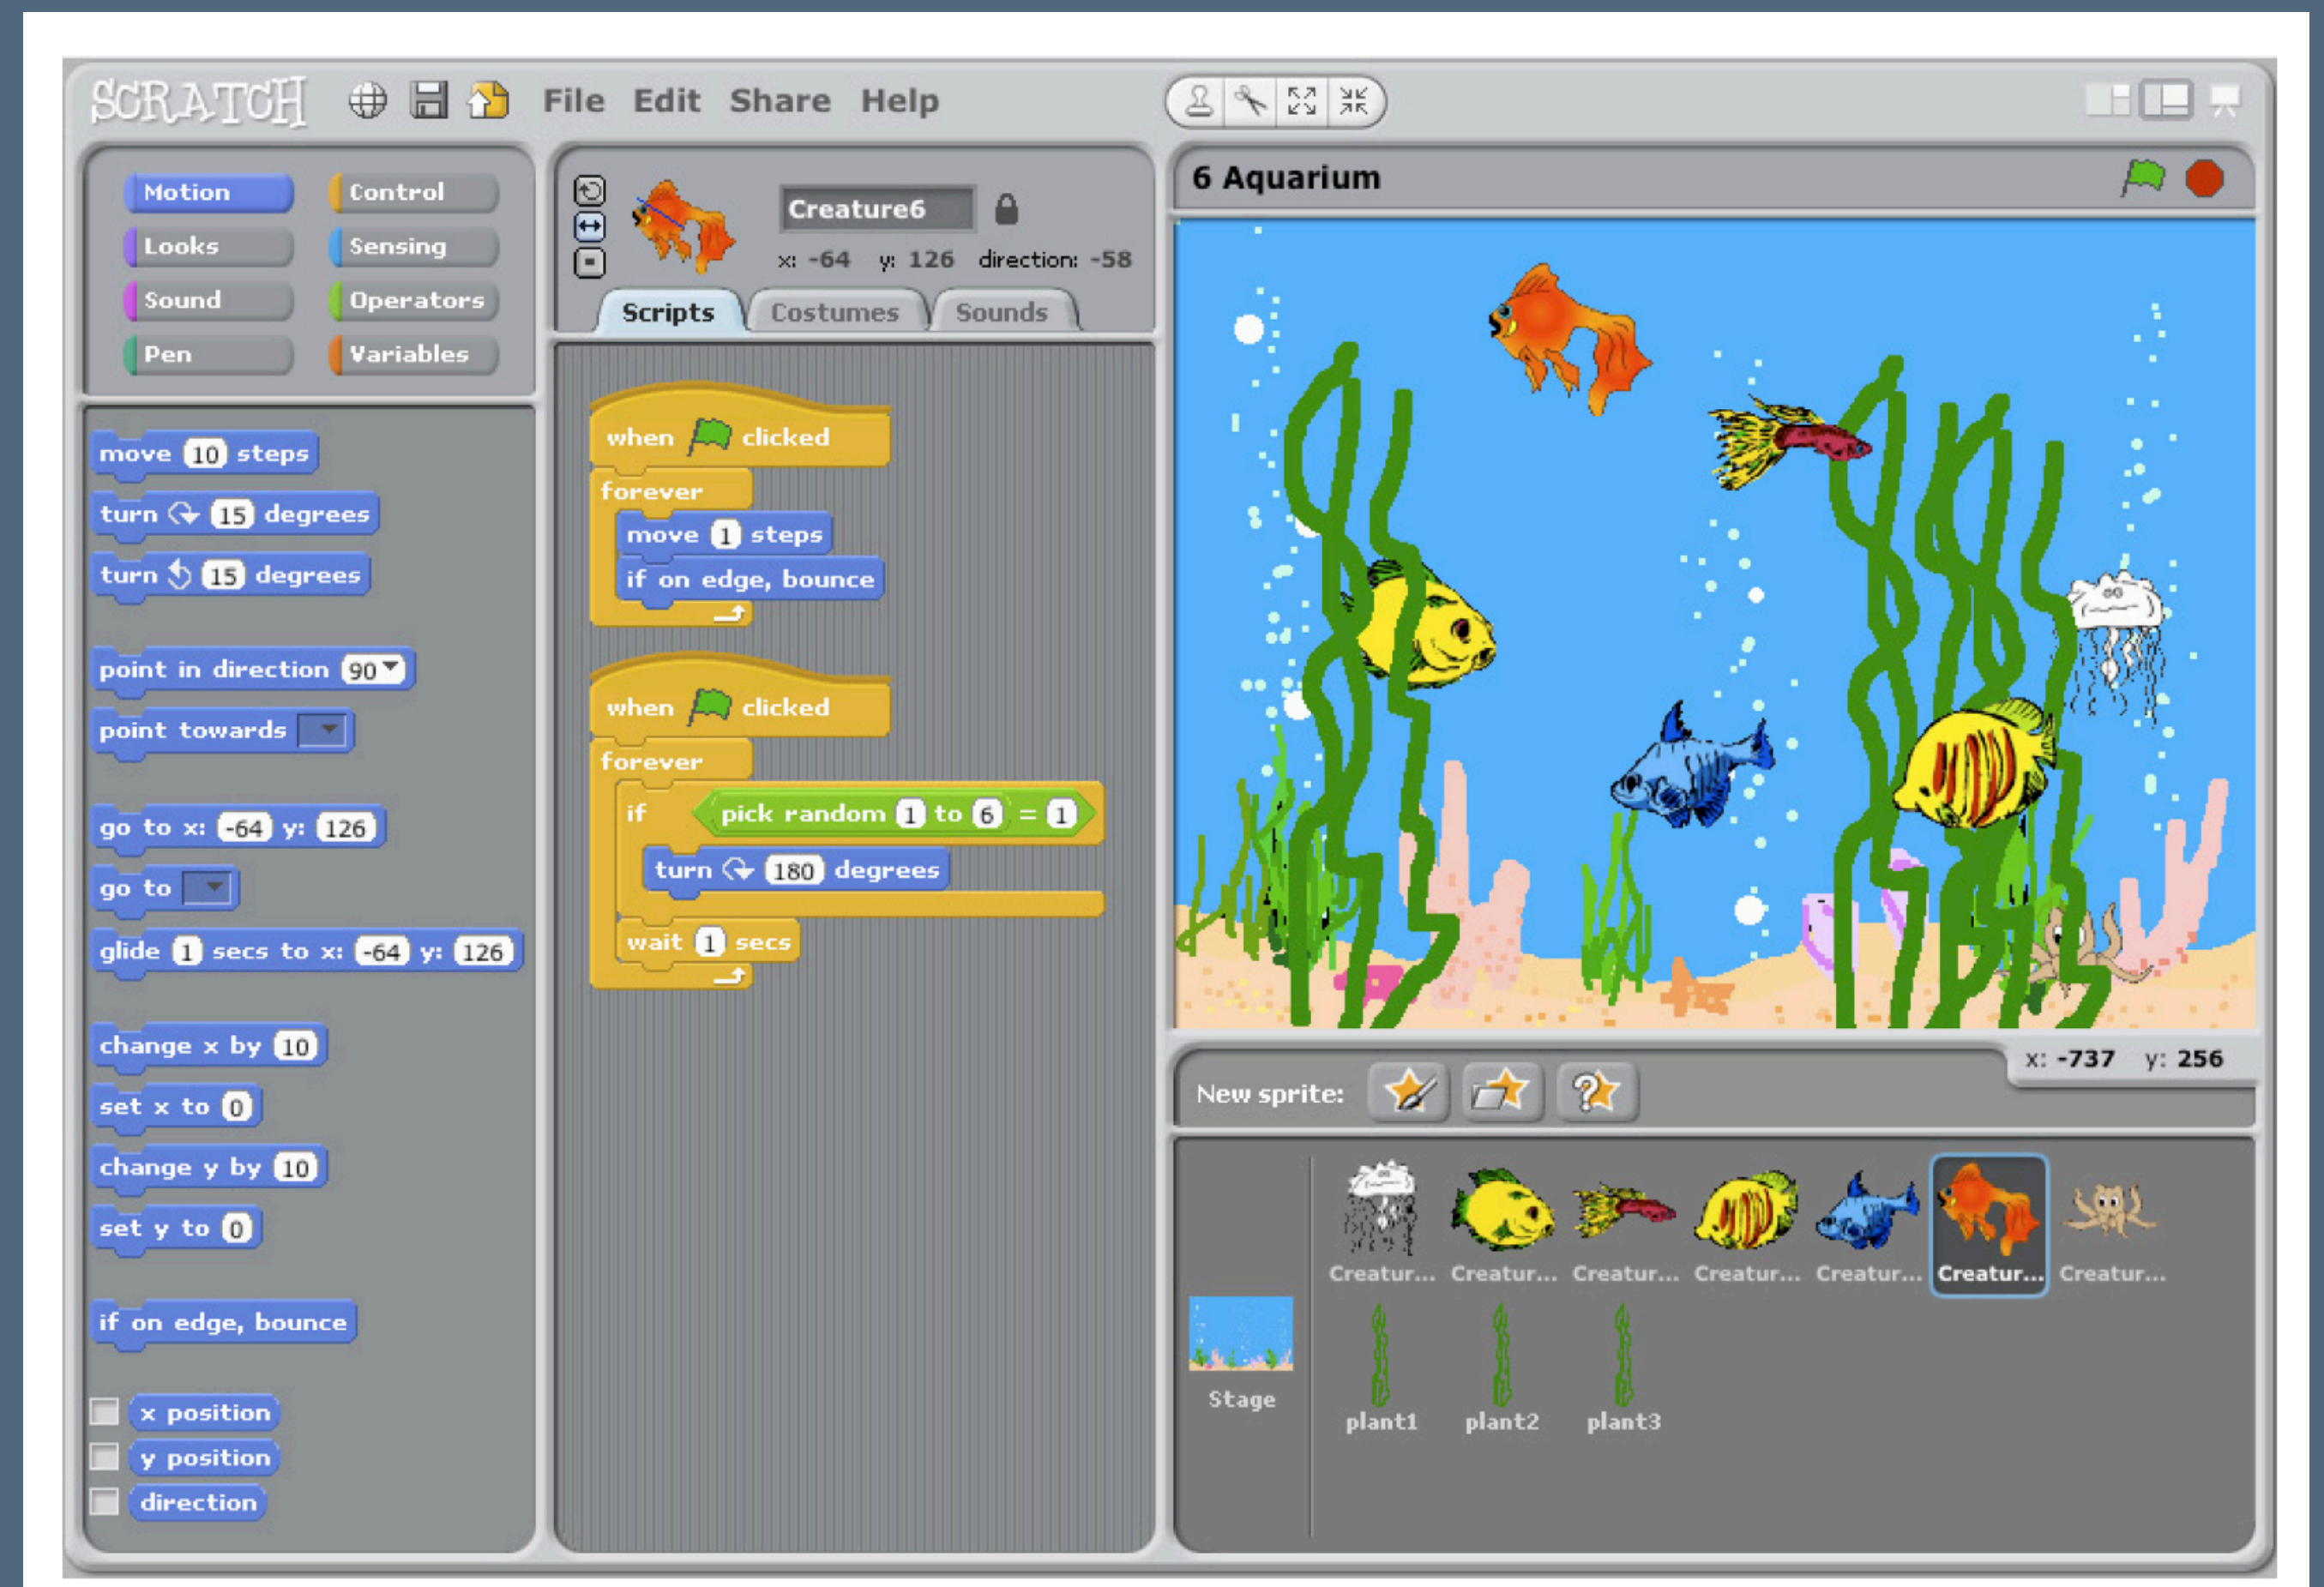


Graphical constraining

[Scaife and Rogers 1996]

When visual aspects constrain the kinds of inferences that can be drawn about the represented world.

Example: block-based programming languages constraining how a program can be constructed [Maloney et al. 2010].



Graphical constraining

[Scaife and Rogers 1996]

When visual aspects constrain the kinds of inferences that can be drawn about the represented world.

Example: AI highlights that tell you what *can't* be skipped when skimming
[Gu et al. 2024]

The recent recognition of a link between increasing rates of deforestation and increasing global climatic warming has focused new attention on the ecological role of forests. Deforestation threatens the continued existence of forests, and their loss would lead to an immediate, irreversible destabilization of the climate because the destruction of forests contributes to increased atmospheric concentrations of such heat-trapping gases as carbon dioxide and therefore to the acceleration of global warming.

The world is at present accumulating carbon dioxide in the atmosphere from two well-known sources: the combustion of fossil fuels and deforestation. Deforestation results in higher levels of carbon dioxide in the atmosphere because the carbon stored in plants and trees is released when trees decay or are burned. A third source, the warming-enhanced decay of organic matter in forests and soils, especially in the middle and higher latitudes, is now being recognized as potentially significant. Evidence is accumulating that carbon from this source is beginning to have global effects. Thus, two of the three sources of carbon dioxide in the atmosphere are directly related to the survival and health of forests.

In summary, good representations...

[Scaife and Rogers 1996]

- Do the computation for us
- Cast into fluid, perceivable, manipulable representations
- Constrain the space of solutions

What makes a representation become a natural part of our thinking?

Extended mind hypothesis

[Clark and Chalmers 1998]

When the tools we use are heavily integrated into the way we process information, they might as well be part of our "mind."

For instance, consider a person with Alzheimer's, Otto, who relies on a notebook to look up where places are. Can we consider this notebook part of his cognition?

It is a **constant** in Otto's life

It is **directly available**
without difficulty

He automatically **endorses**
information from it

Extended mind hypothesis

[Clark and Chalmers 1998]

The extended mind hypothesis is an interesting provocation suggesting that if the coupling between a person and their artifacts is reliable, then perhaps it should be viewed as part of the same mind, working towards the same goal.

Sensemaking

Example sensemaking situation

You are planning a trip for you and your family.

Intended end result: a schedule of places and things to do

Starting point: a destination and set of vague constraints

How might you get to the intended result?

Come up with ideas

Collect information

Organize it

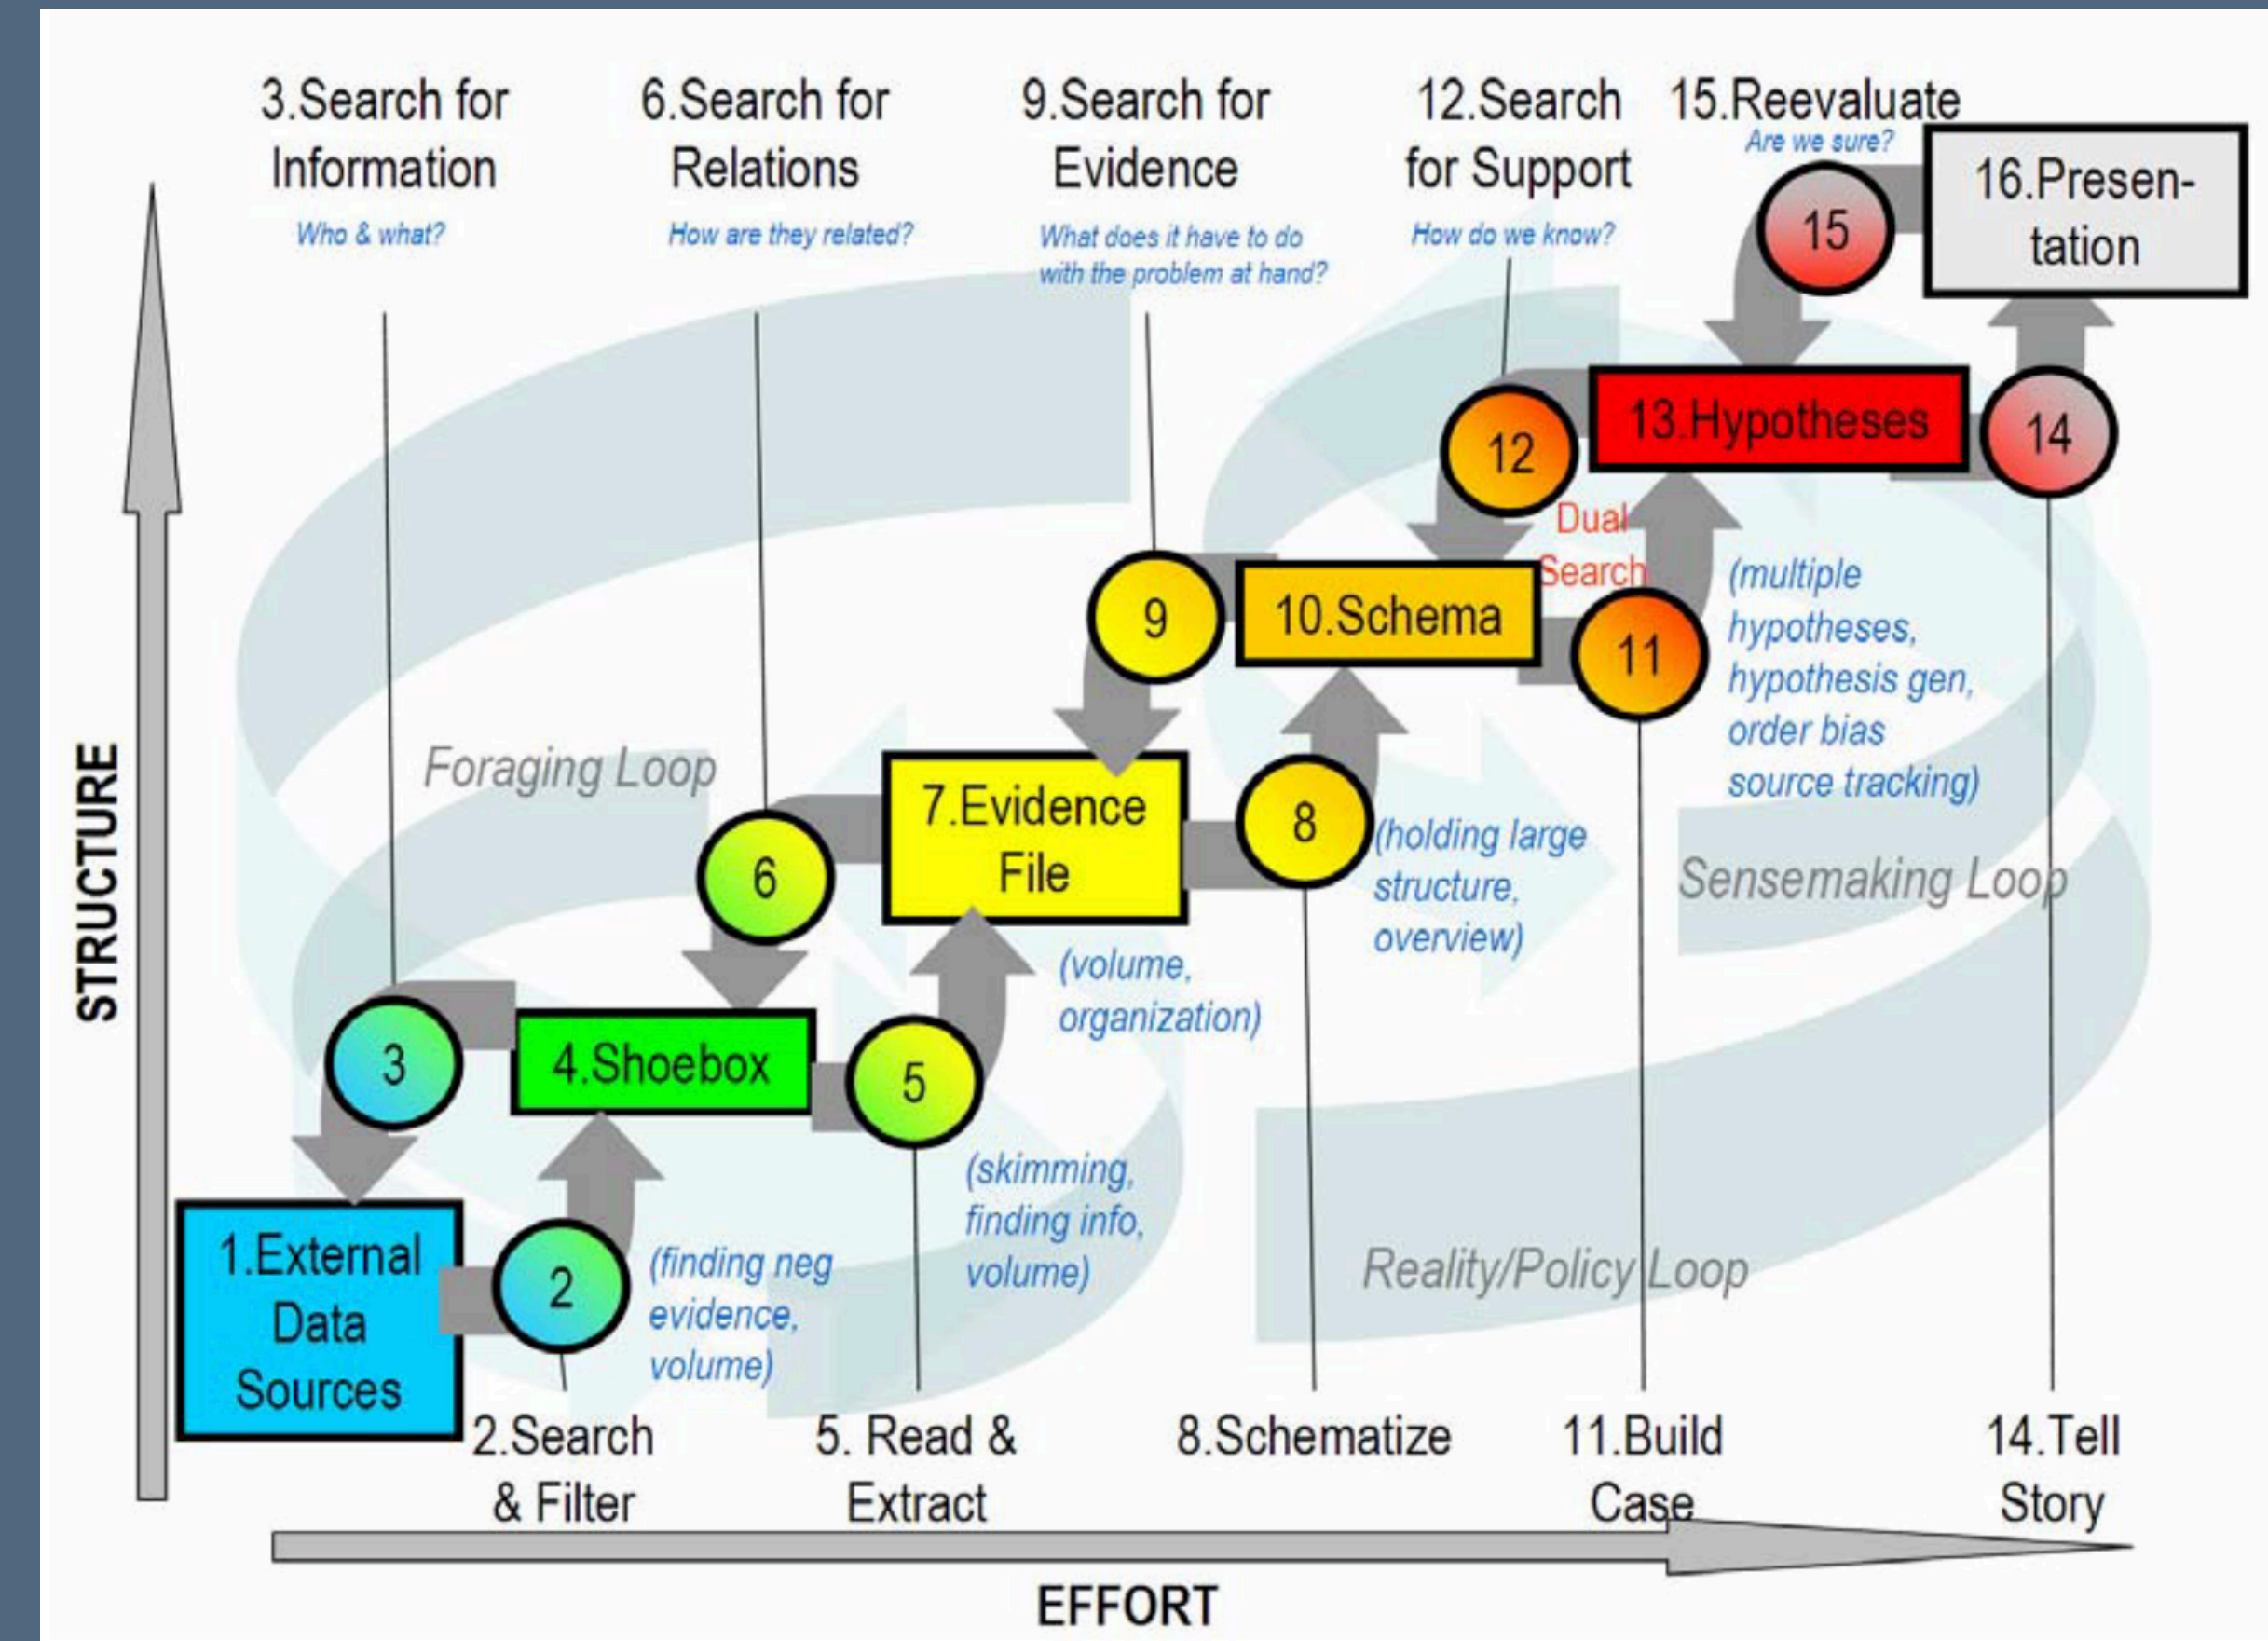
Reflect and iterate

Share outcome

The sensemaking process

[Pirolli and Card 2005]

Derived from a cognitive task analysis of intelligence analysts. Still, this framework has stood the test of time. It has appeared in derivative form in countless HCI papers and talks.




The sensemaking process

[Pirolli and Card 2005]

At its coarsest, described as:

1. Information
2. Schema
3. Insight
4. Product

Schema: set of learned templates for how best to interpret and process information



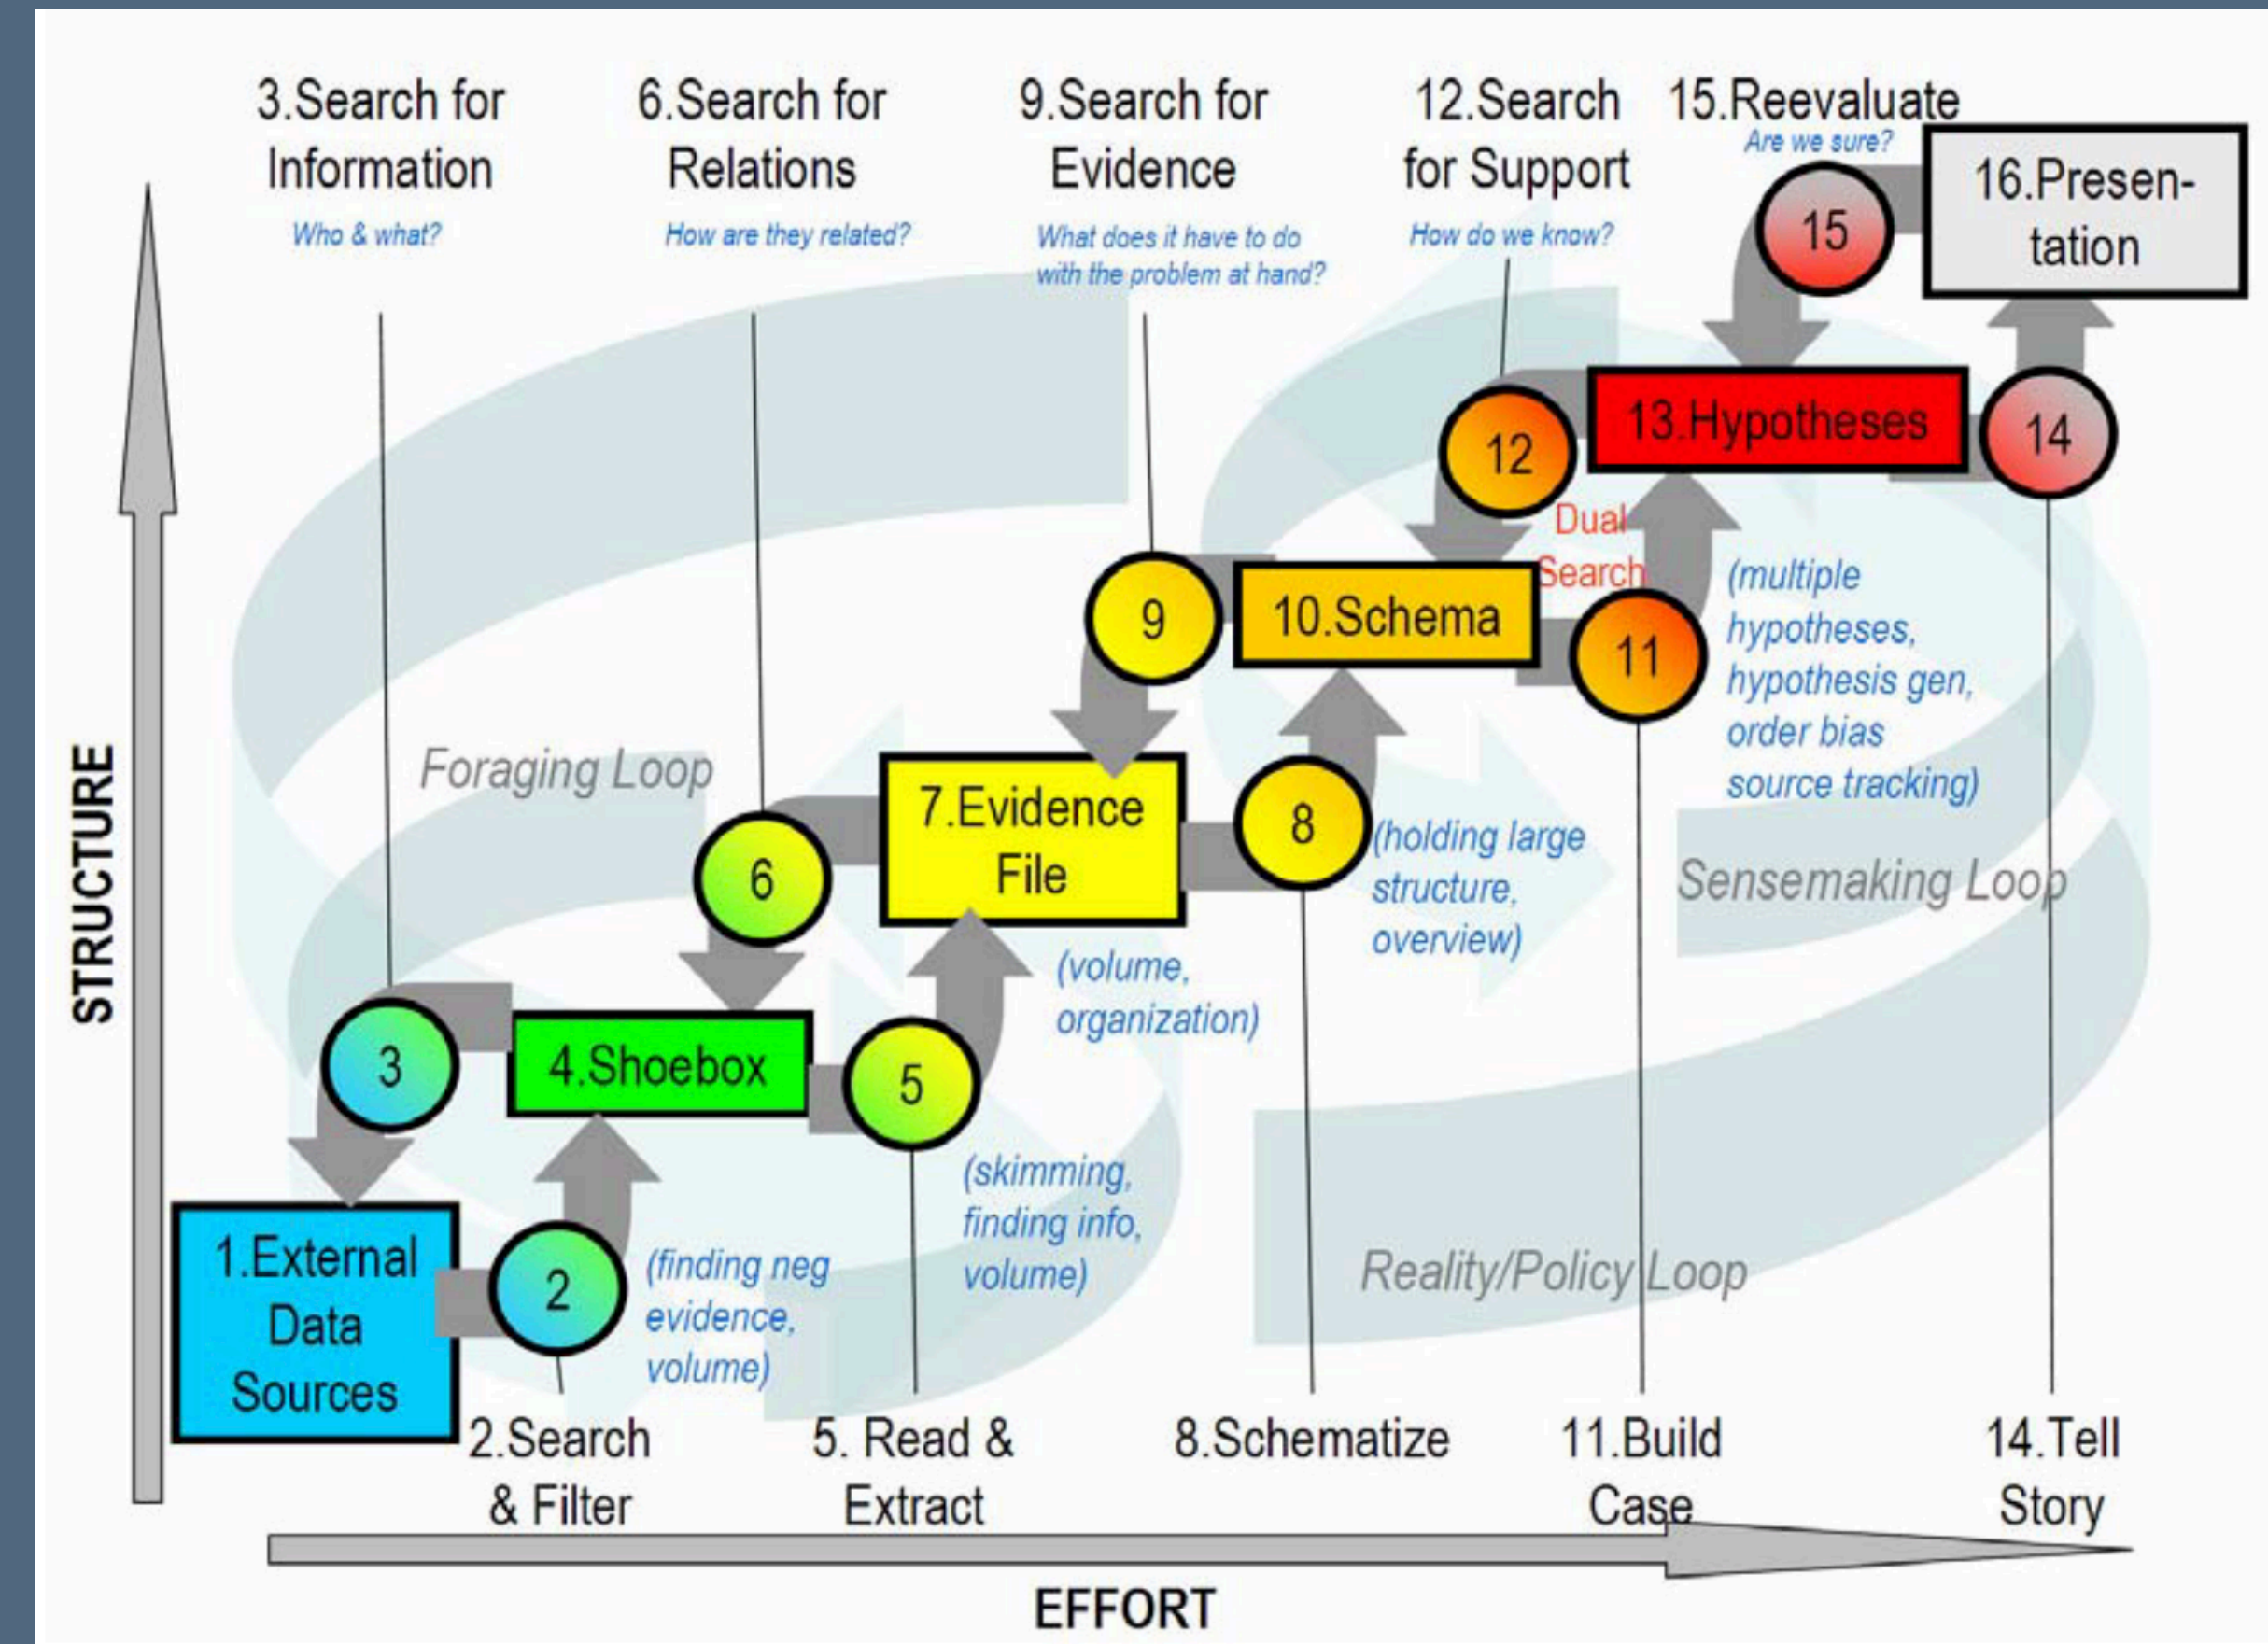
The sensemaking process

[Pirolli and Card 2005]

Loop 1. Foraging loop. Seek, search, filter, read, extract information, possibly into a schema.

Loop 2. Sense making loop. Iterative development of an understanding that fits the evidence.


What interfaces do you use for sensemaking?



The sensemaking process

[Pirolli and Card 2005]

Sensemaking tools address some steps or transitions in the sensemaking process, lowering the barriers to foraging or synthesis.




Search and Filter

[Chang et al. 2016]

Reduce cost of capturing information into the shoebox
[Chang et al. 2016].

Support capturing on mobile phones using more easily executable input actions (long press with context window determined by pressure).

Reduce temptation to overinvest in precision at early stage (prohibit exact text selection).



Search for Relations

[Chang et al. 2023]

Put encountered information in context.

Augment the citation with information about who else cited this work and how else you have come across it in the past.

(At the same time, assist in decisions around searching for more citations).

Assieme [33], and Libra [56] improve existing general-purpose...
Assieme: finding and leveraging implicit references in a web search interface for programmers
2007. R. Hoffmann, J. Fogarty and Daniel S. Weld
TLDR In a study of programmers performing searches related to common programming tasks, it is shown that programmers obtain better solutions, using fewer queries, in the same amount of time spent using a general Web search interface.(show more)
Cited by 1 of your recent readings...
83% SCROLLED
Crystalline: Lowering the Cost for Developers to Collect and Orga...
2022. Michael Xieyang Liu, A. Kittur and B. Myers
Opened 11 minutes ago (remove from history)
hide citation context
Here, developers not only need to find information pertinent to their problem [8, 59, 97, 115], which is the first step in such complex sensemaking tasks [106, 123], but also collect and synthesize relevant information into structured knowledge so that they can make progress towards fully understanding the decision space [53, 71, 72, 81].

could build on current systems that already recommend contextual help to debug programming errors (like Unakite [56]) or present example code of that data are search systems idea galleries (li in exploration c convert video, by a search syst these limitation from these find
7 CONCL
This paper intro mation present query suggesti by detecting p of issued queri domain-genera
notes-based query suggestions helped people explore broad multi-
[13] Robert Capra, Jaime A ferences in the use of Proceedings of the 38th
READ 18% SCROLLED
Unakite: Scaffolding Developers' Decision-Making Using the Web
2019. Michael Xieyang Liu, Jane Hsieh, Nathan Hahn, Angelina Zhou, Emily Deng, Shaun Burley, C. Taylor, A. Kittur and B. Myers
TLDR The evaluation results show that Unakite reduces the cost of capturing tradeoff-related information by 45%, and that the resulting comparison table speeds up a subsequent developer's ability to understand the trade-offs by about a factor of three.(show more)
Saved 3 minutes ago
Saved from: Crystalline: Lowering the Cost for Developers to Collect an...
In addition, though existing tools provide users with the flexibility and agency to synthesize the collected information into useful representations, such as comparison tables [15, 81] or knowledge maps [87], developers still need to perform these organizing operations manually.
[21] Allan Cooper, Mira D

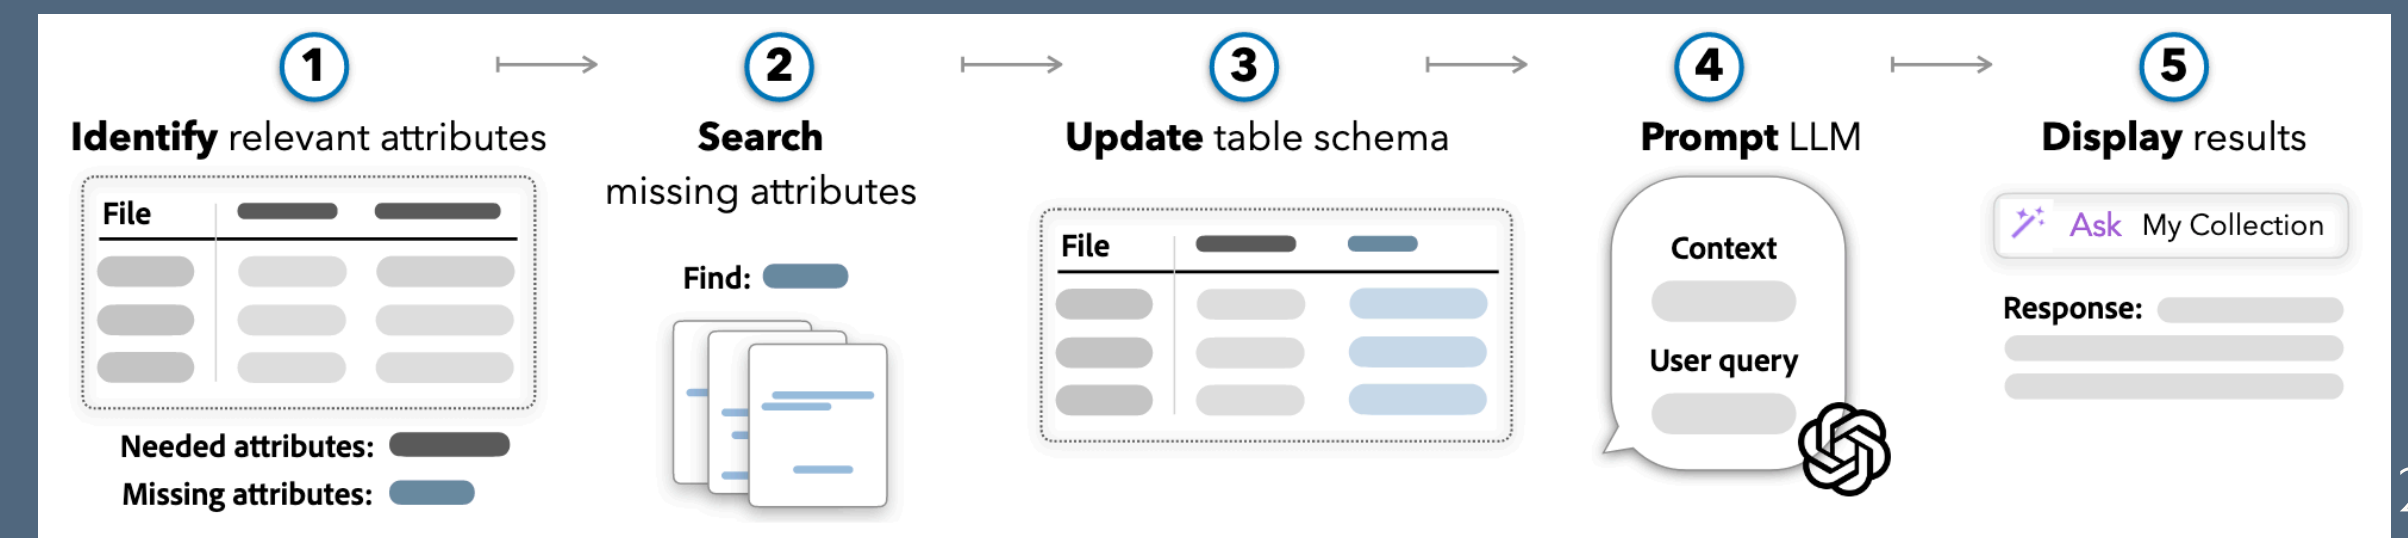
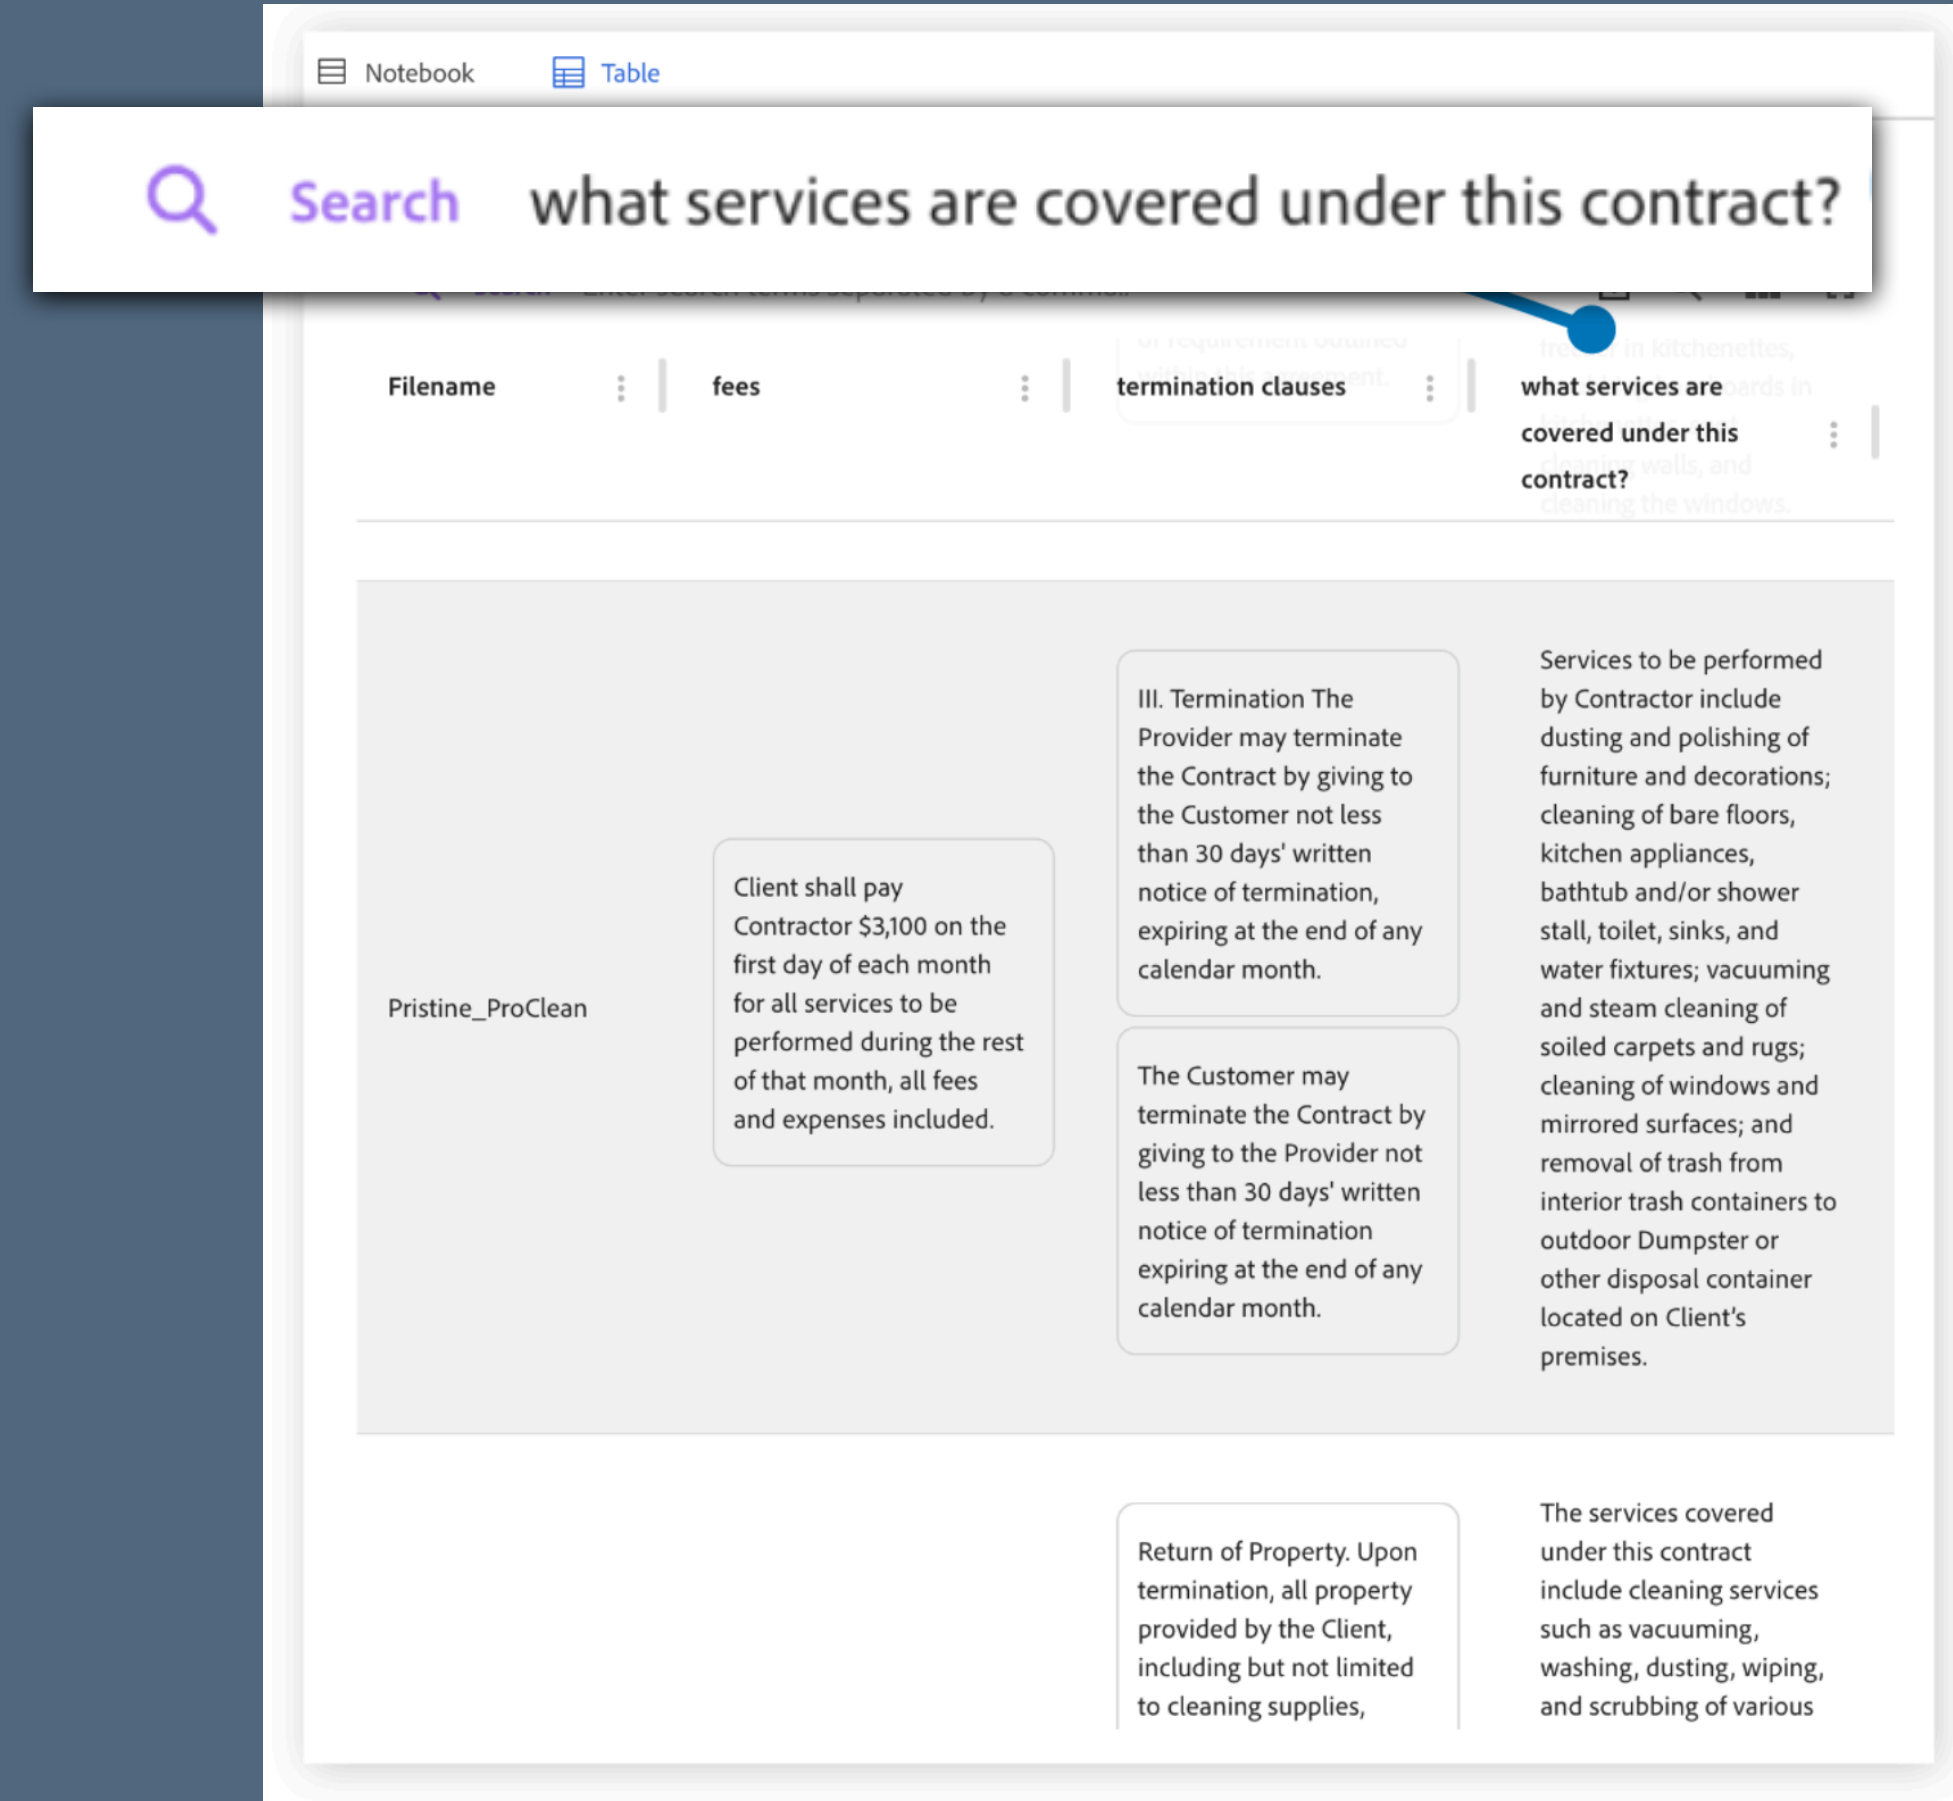
Schematize

[Fok et al. 2024]

Fit found information into a meaningful structure.

Given a user query, extract information answering it from business documents. Place them in a table that supports more fluid comparison and synthesis.

Bonus: hooks back to the foraging step (can click on extractions to open them up in the original doc).

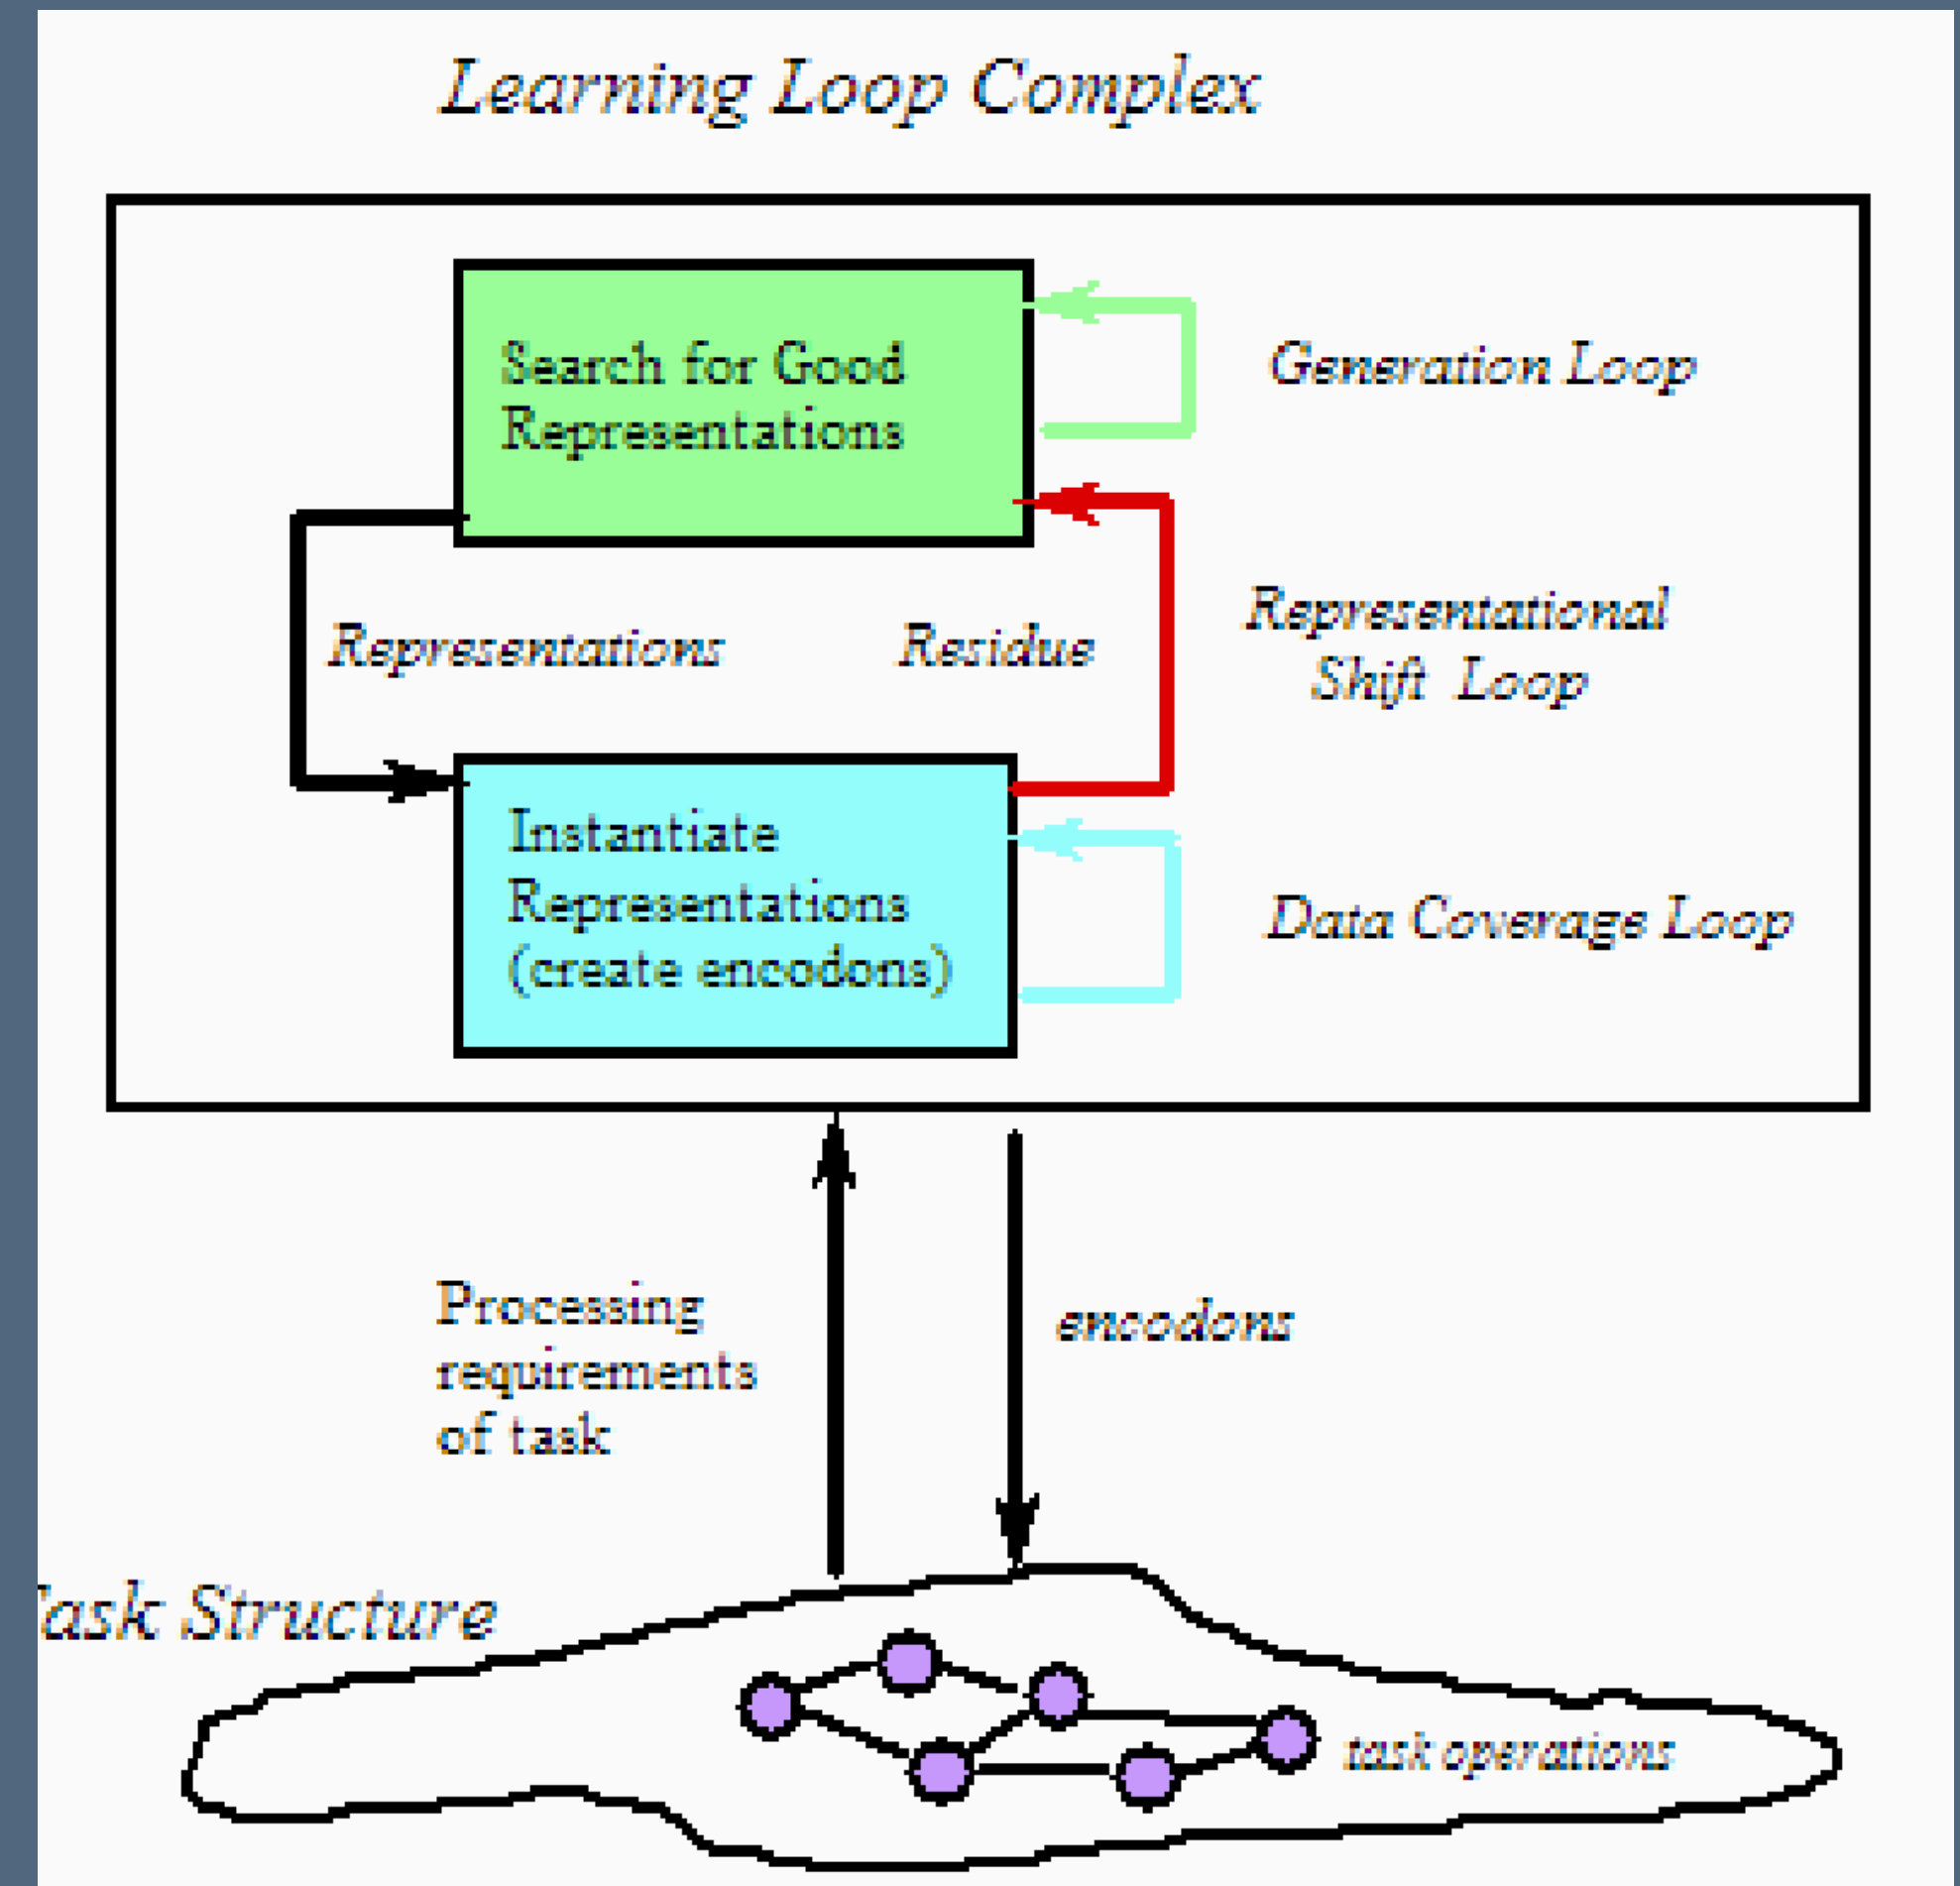


Surprise in sensemaking

[Russell et al. 1993 by way of Pirolli and Card 2005]

If sensemaking involves finding the right schema to fit information into, then there are problems if that schema is wrong.

When important information does not fit into our schema, we have **residue**. In some situations, it is important that we help people incorporate that residue.



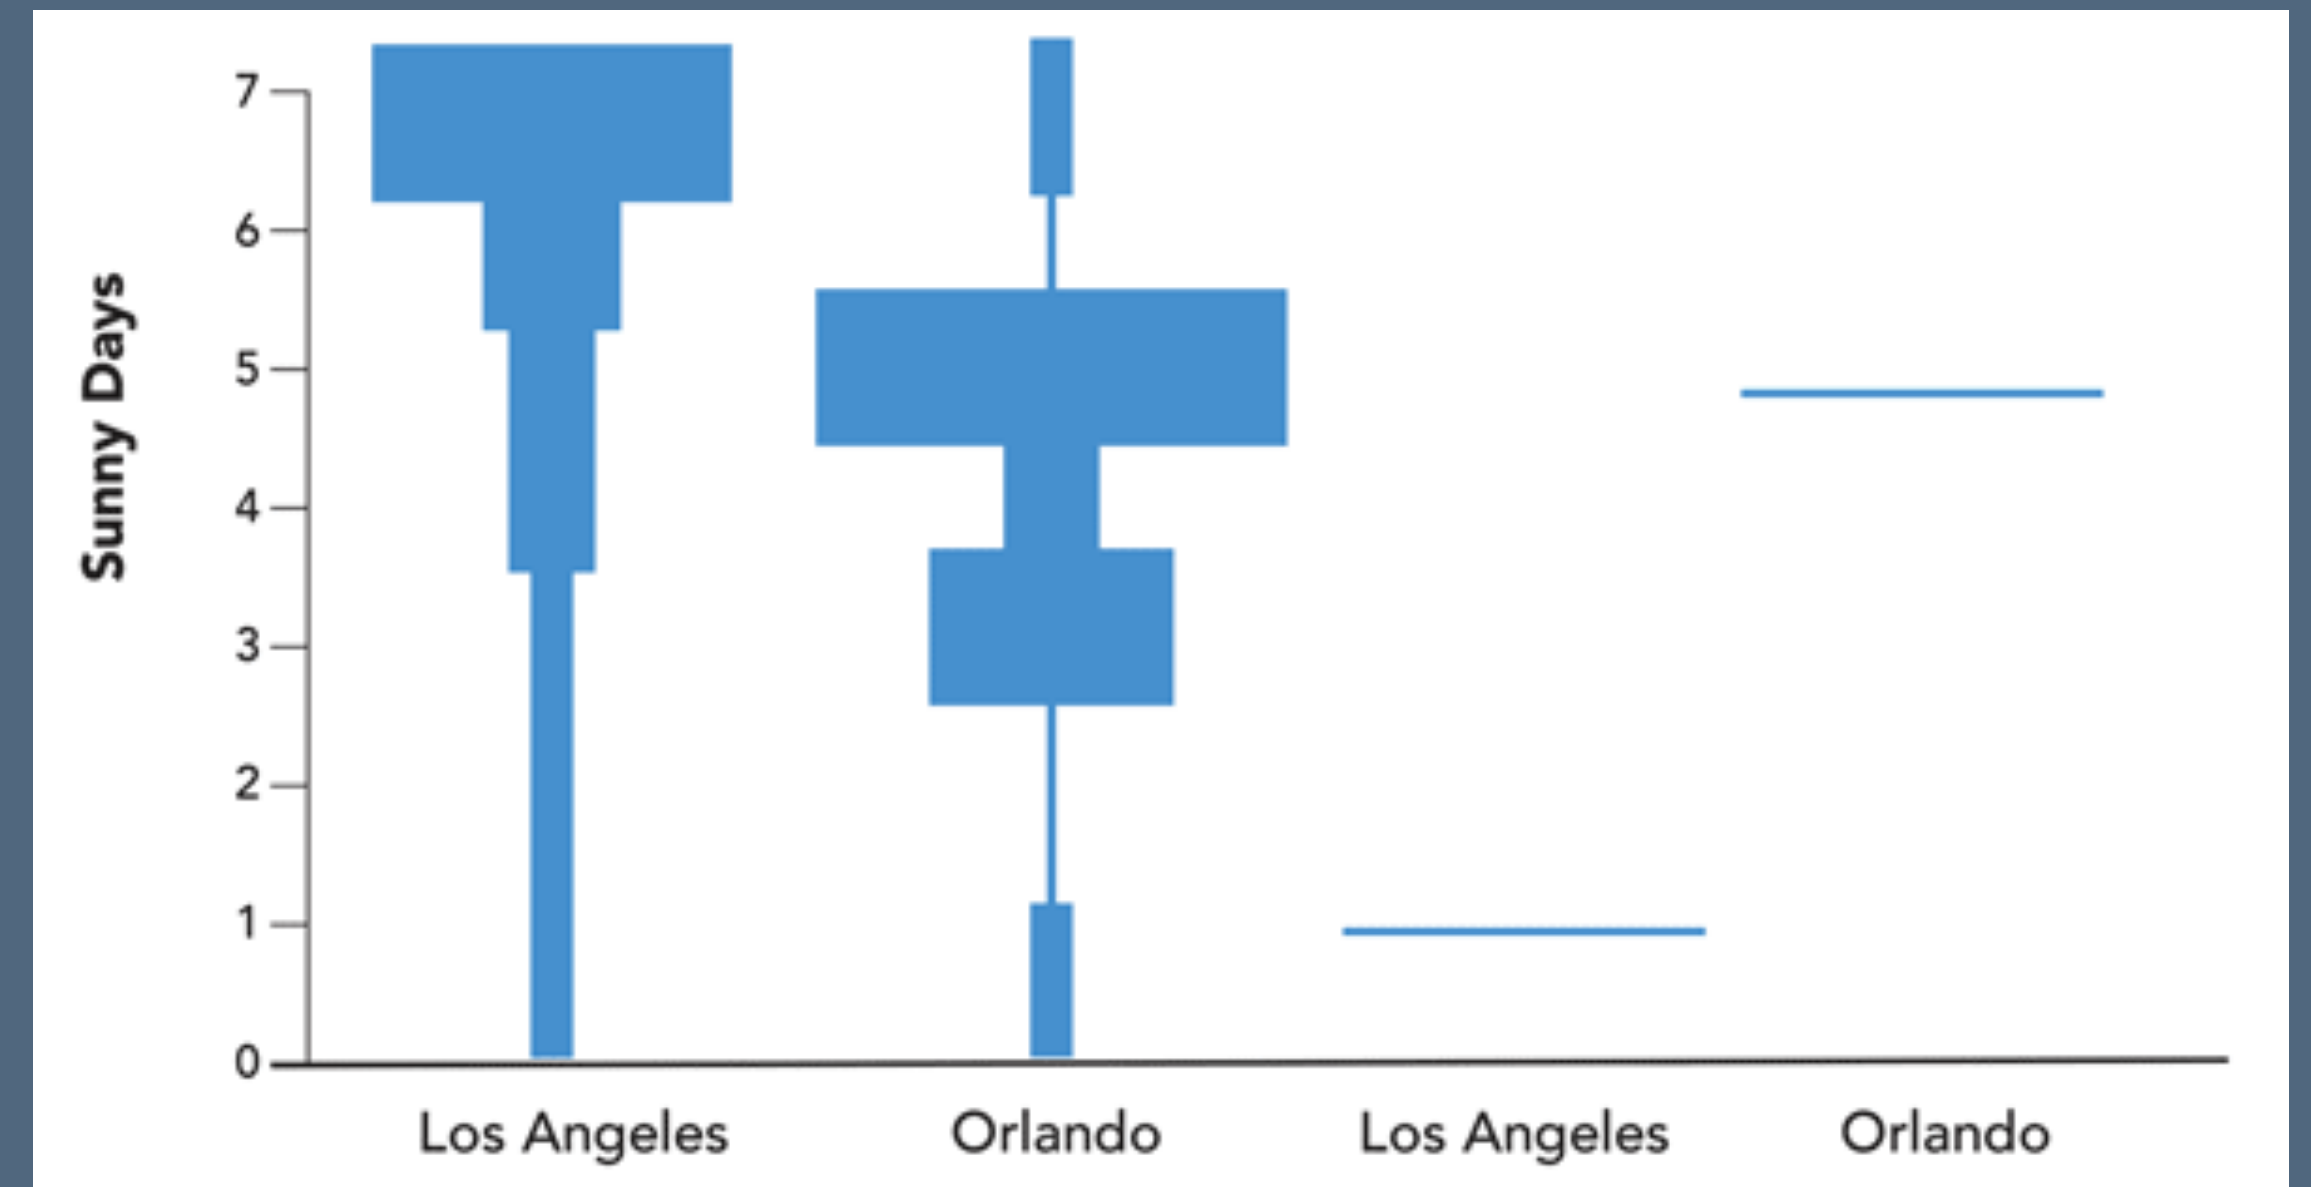
Showing the residue

[Kale et al. 2019]

Trying to keep someone from misreading the statistical model?

Show them examples in the form of hypothetical outcome plots.

Can improve readers' ability to infer the actual underlying trend in the data in low-evidence situations.



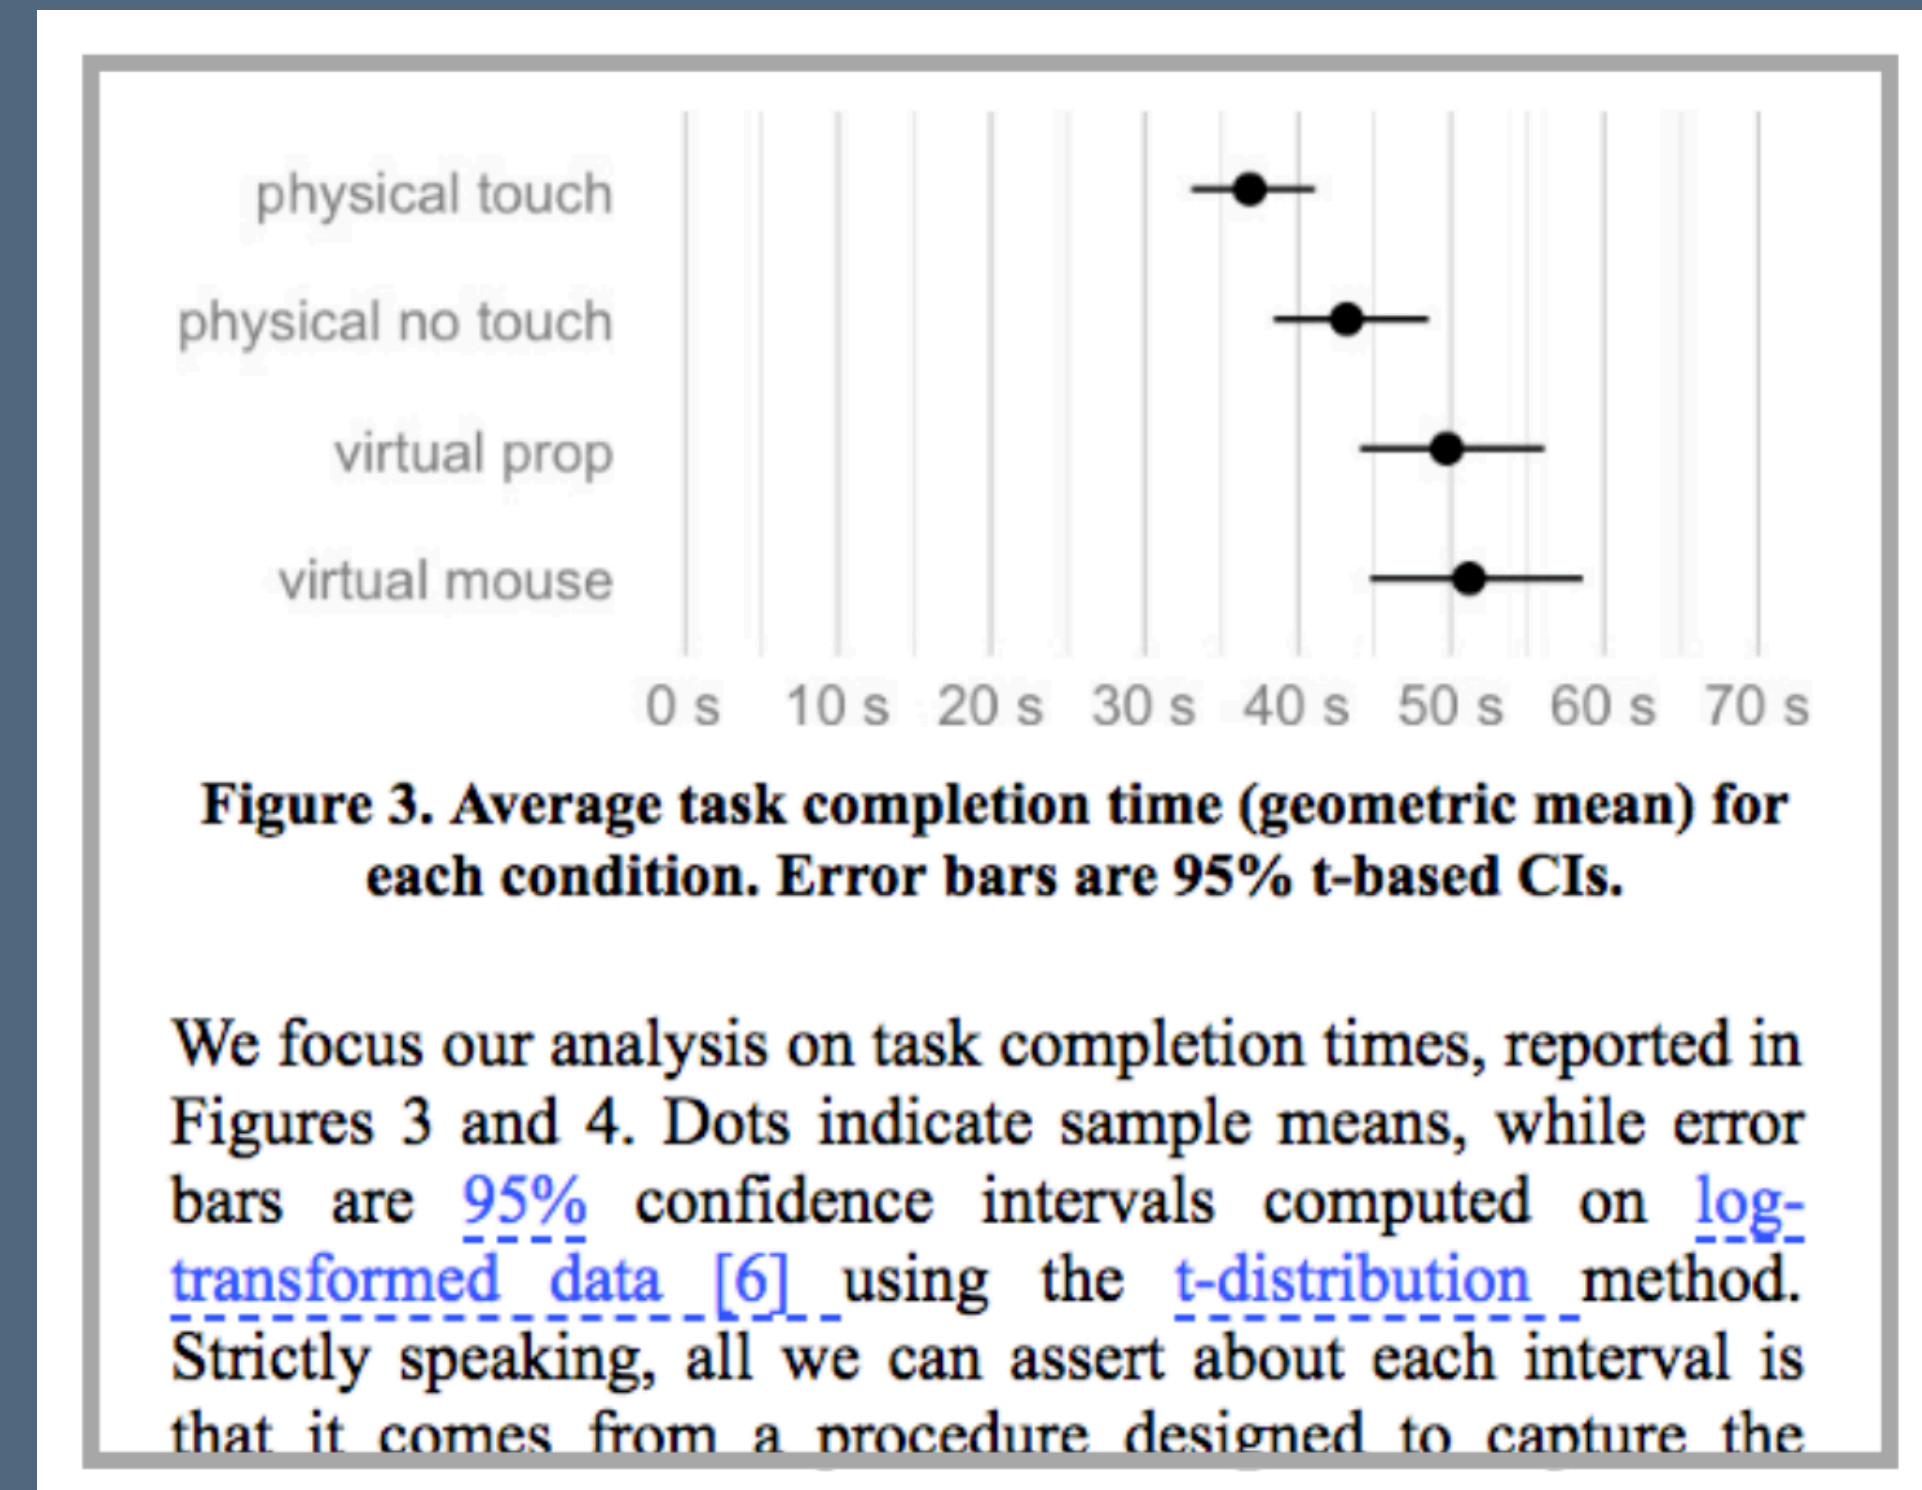
Showing the residue

[Dragicevic et al. 2019]

Trying to keep someone from misreading the statistical model?

Allow people to play around with the parameters and see how it affects the predictions.

Create a new medium of explorable multiverse reports, where parameters are baked into the document to permit this kind of information seeking.

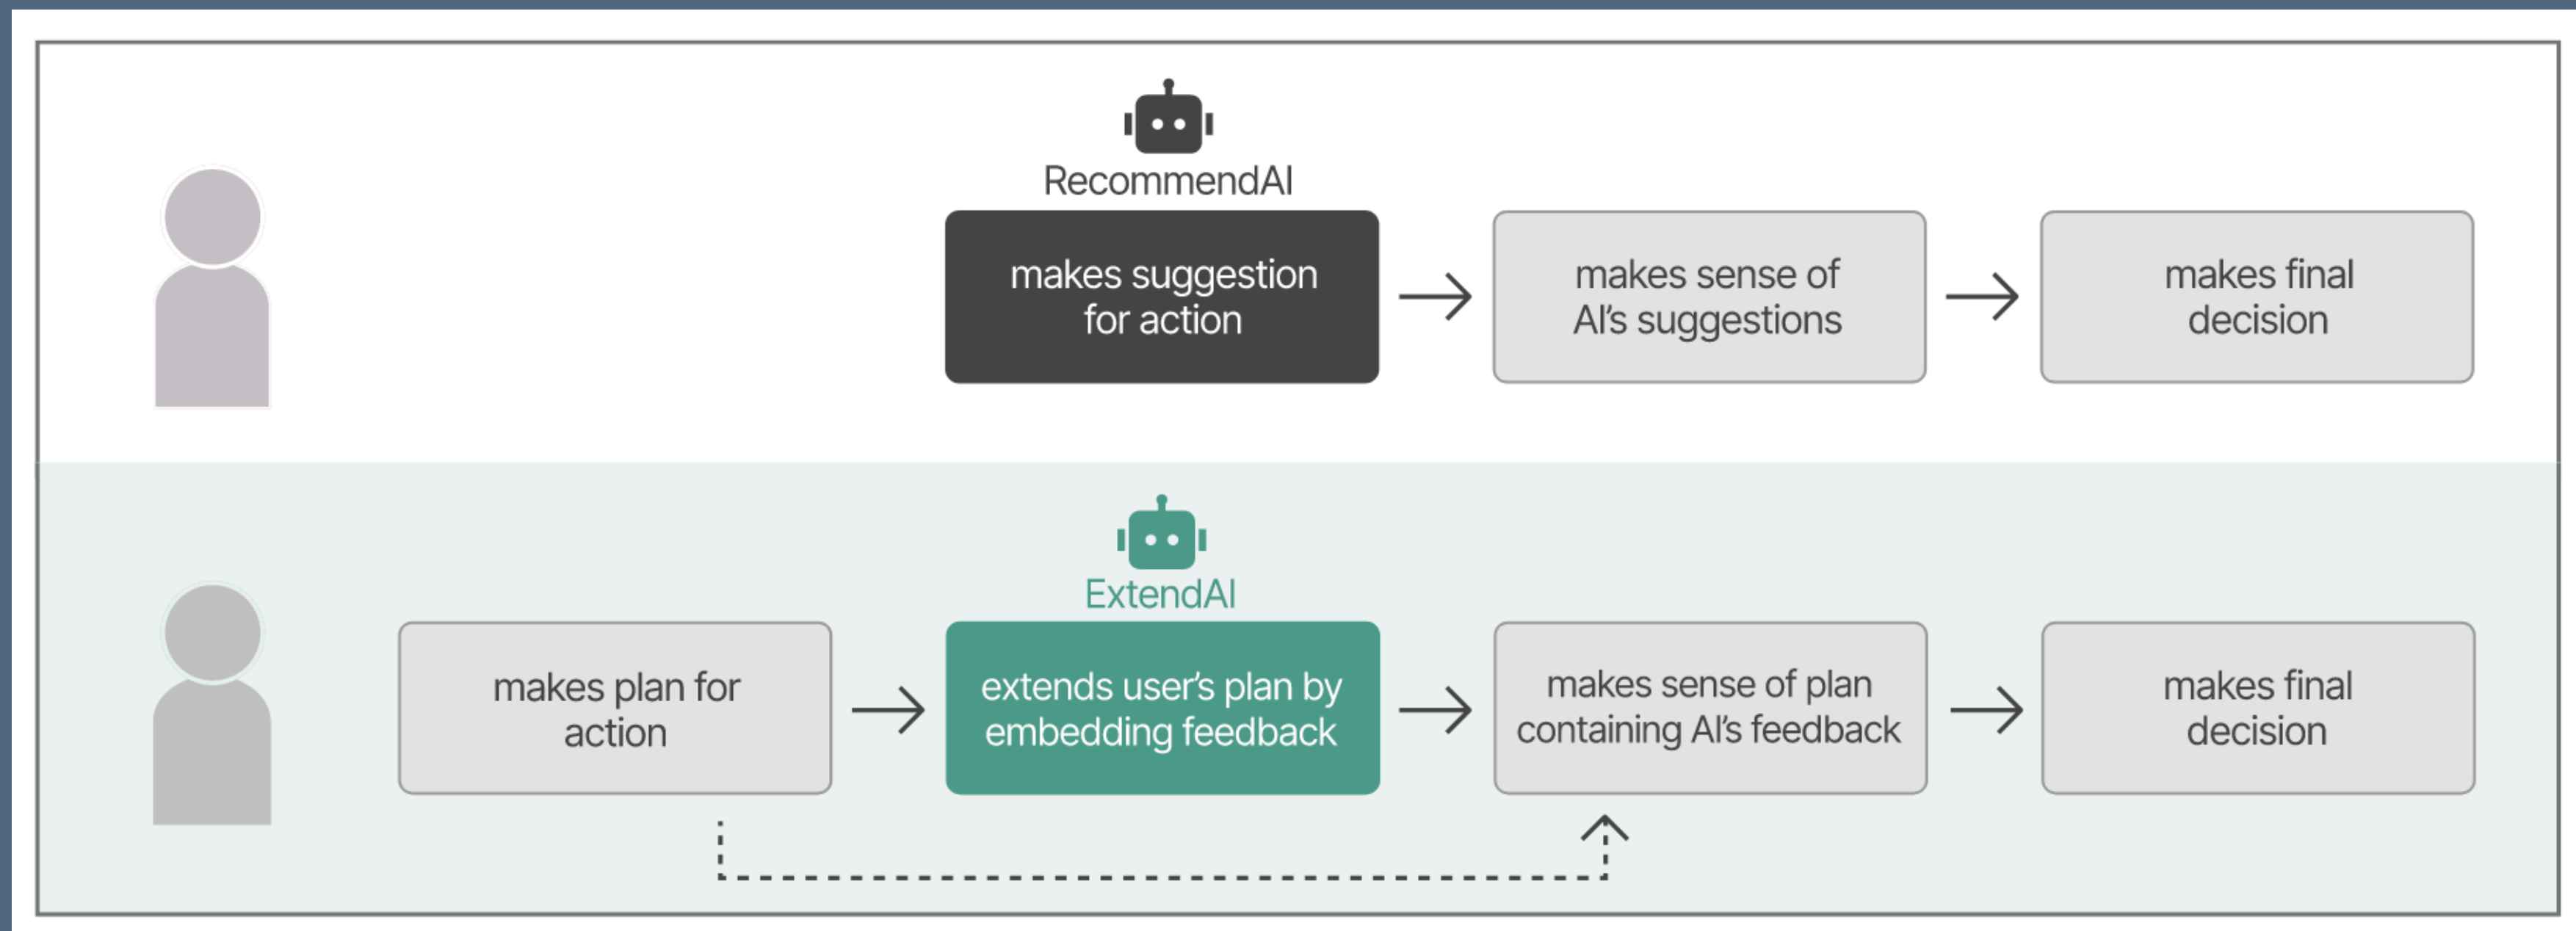


Sensemaking with an AI

[Reichert et al. 2025]

Using an AI as a partner in making sense of information with an intent to make a high-stakes decision.

What model of interaction is better? *RecommendAI* or *ExtendAI*?



Reverse Centaurs


Centaur

[adapted from Doctorow 2025]

Mind of a human, propulsion of a horse.

Used to describe human-AI interaction where the person **works towards their goals with the power of AI behind them.**

Example: While writing a new article, I used AI to produce a transcript of a bunch of MP3s I have of podcasts, and was able to search them for the perfect quote.



[https://commons.wikimedia.org/wiki/File:Centaur_\(PSF\).jpg](https://commons.wikimedia.org/wiki/File:Centaur_(PSF).jpg)

Reverse centaurs

[adapted from Doctorow 2025]

Mind of a horse, propulsion of a human.

Used to describe human-AI interaction where the **person's efforts are put to work in service of an AI.**

Example. Amazon warehouse workers fetching packages under the prescribed scheduling of the AI. A content creator who does not write themselves but rather fact-checks AI-generated drafts.

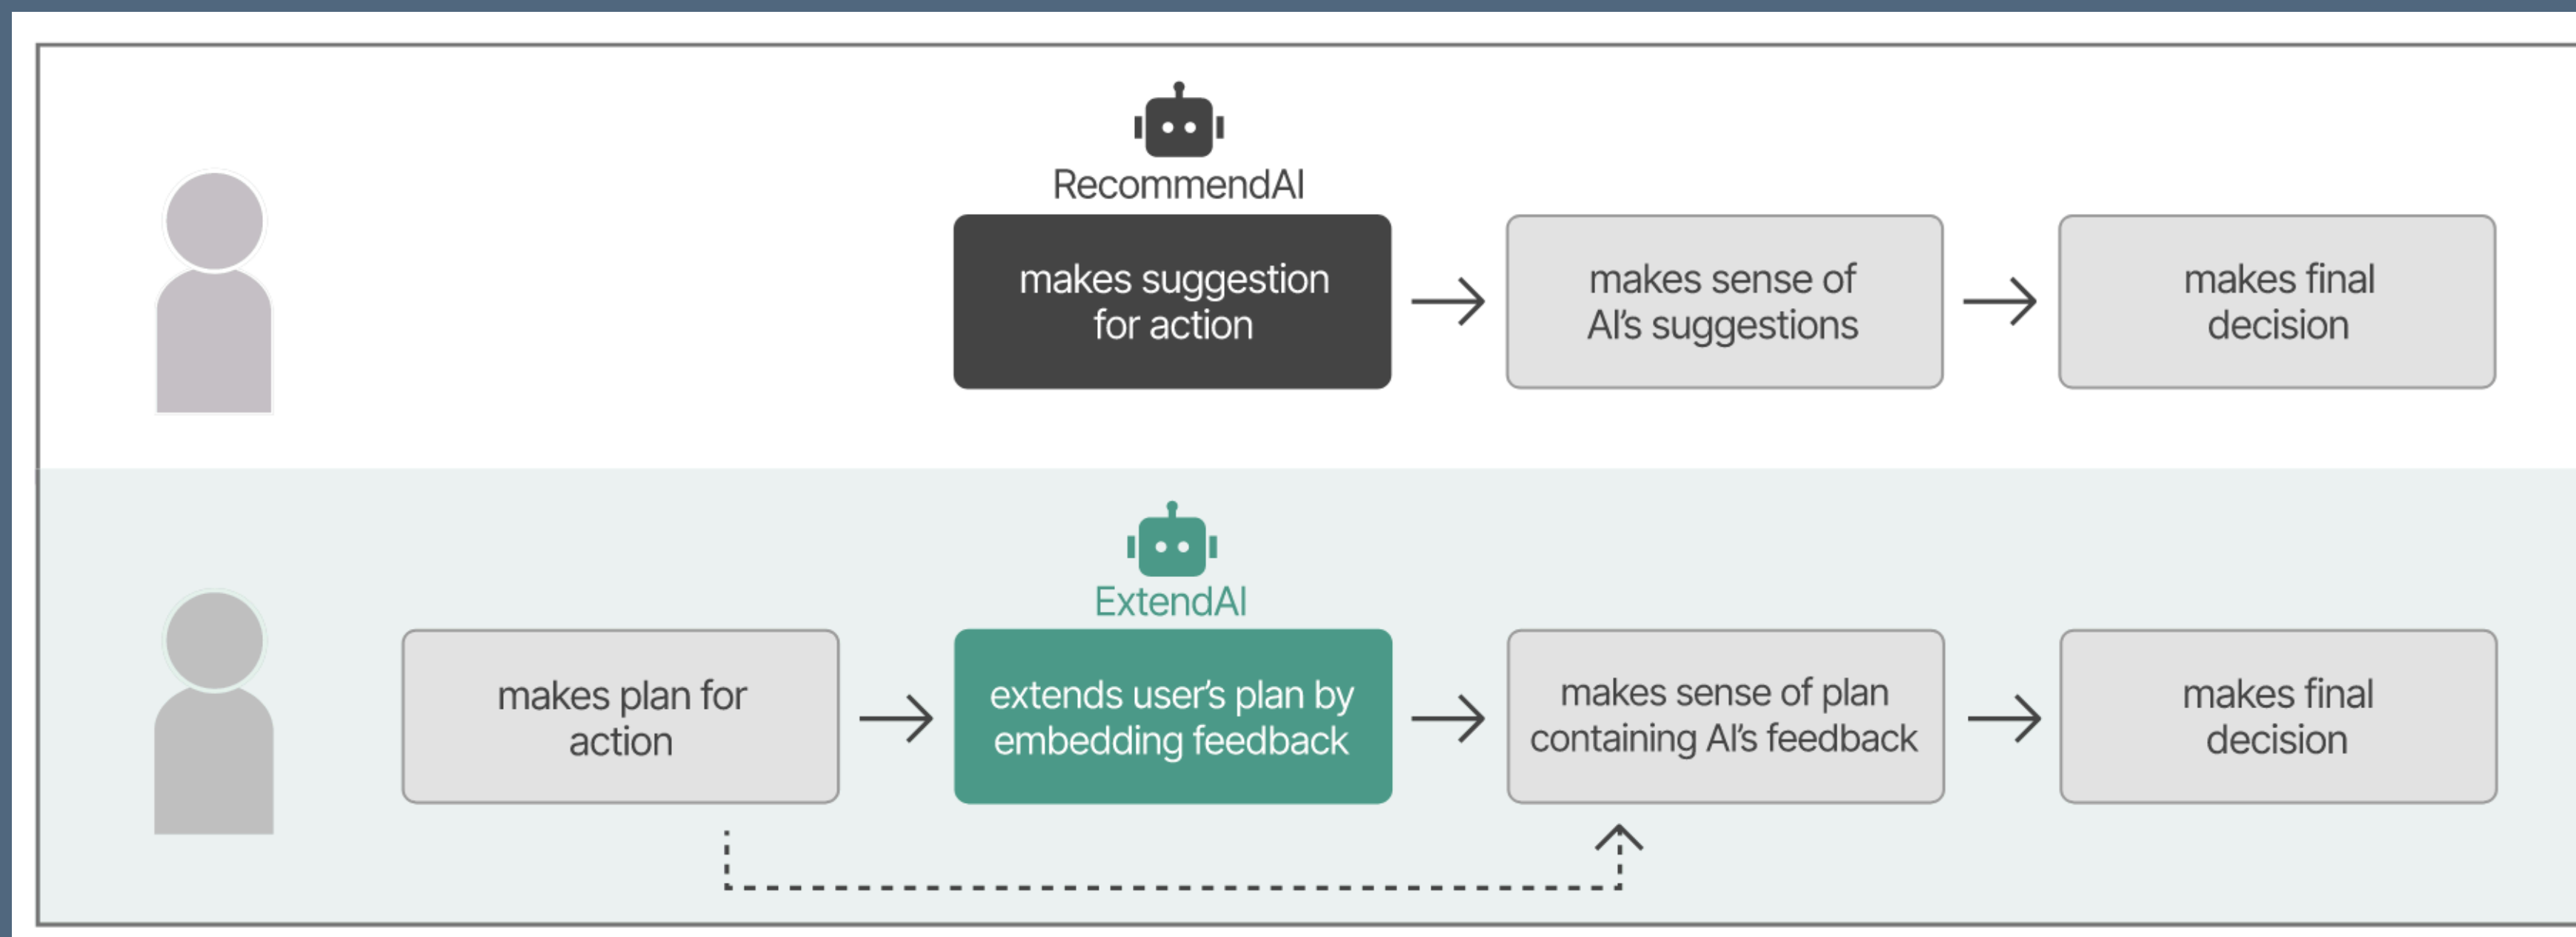


Sensemaking with an AI

[Reichert et al. 2025]

RecommendAI. Centaur or reverse centaur?

ExtendAI. Centaur or reverse centaur?



After "Man-computer symbiosis"

[Licklider 1960]

"Man-computer symbiosis is probably not the ultimate paradigm for complex technological systems. It seems entirely possible that in due course, electronic or chemical machines will outdo the human brain in most of the functions we now consider exclusively within its province." - Licklider

This sounds like anticipating the reverse centaur in 1960. Not exactly a stirring vision of where this is all going.

What role should/must humans play in thought?

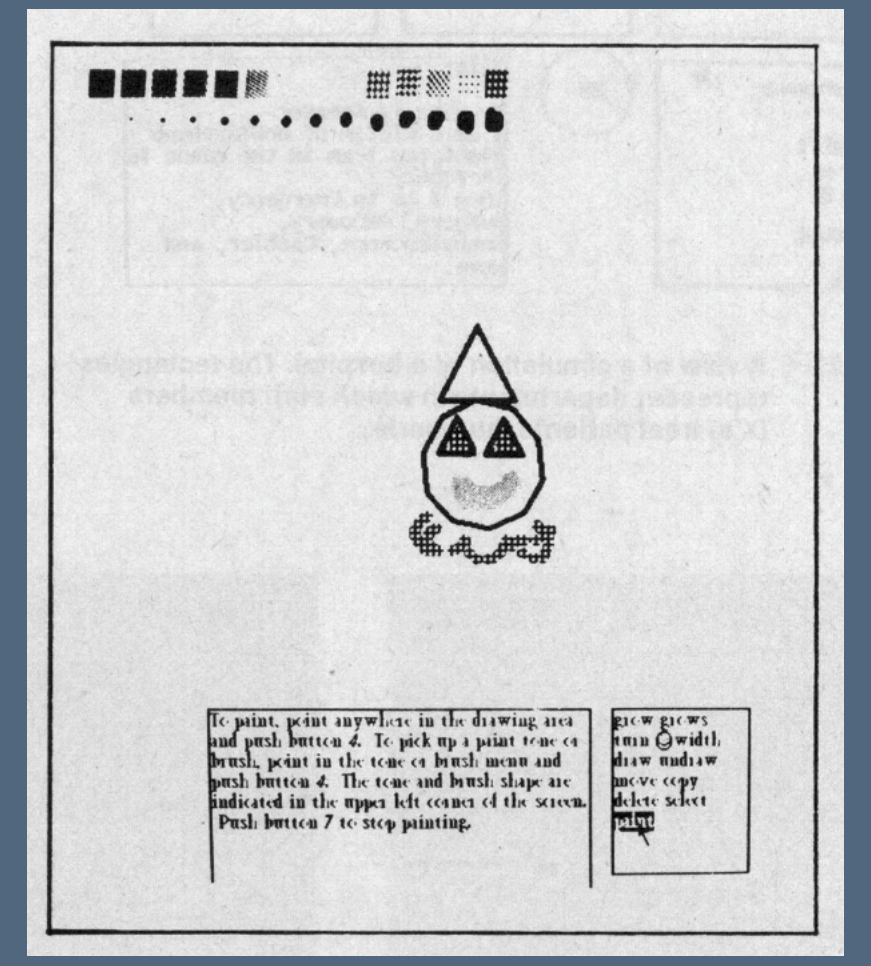



How far have we come?

One more blast from the past

[Kay and Goldberg 1977]

"Although thinking goes on in one's head, external media serve to material thoughts and, through feedback, to augment the actual paths the thinking follows."

Was Kay and Goldberg's vision transformative or not? Are there any parts we haven't achieved?



Have we missed the goal?

Kay famously said that the computer revolution hasn't happened yet.

At the time, many of his visions around more malleable programming of the machines and adaptable experiences had, and they still haven't.

When asked, Engelbart apparently told one of his collaborators in 2006 that only "2.8 percent" of his vision had been achieved [Landau 2018].

Some challenges for tools for thought

[Matuschak and Nielsen 2019]

Developing tools for thought is something of a **public goods problem**. Software firms do not have much incentive to develop them, because it is costly to do well. Discovered software patterns can be easily copied.

And many good tools for thought are public goods. Think about language, writing, math notation, the spreadsheet design. They are universal, belong to no one, and thus can scaffold coordination between us all.

There is also an **expertise problem**. Matuschak and Nielsen posit that good tools for thought require domain experts working on real, cutting-edge problems, where tools for thought fall out of the experience as an aid to help them do their work.

Summary

Tools for thought are a wide category of tools where the common feature is their support for **production of new or refined knowledge from old knowledge**.

One kind of tool for thought is **representations**. These offload human computation, support problem-solving through manipulation, and helpfully constrain the space of solutions.

Sensemaking tools support thought as it evolves in messy, nonlinear, iterative projects. They help people arrive at appropriate schemas while foraging and making sense of information against those schemas.

Tools can either promote the human or AI as lead thinker, resulting in centaurs or **reverse centaurs**. A recent conversation for our community is how to advocate for centaurs where it seems more work is aimed towards reverse centaurs.

References

Appleton, Maggie. "Tools for Thought as Cultural Practices, Not Computational Objects." 2023. May 13. <https://maggieappleton.com/tools-for-thought>.

Chang, Joseph Chee, Nathan Hahn, and Aniket Kittur. 2016. "Supporting Mobile Sensemaking Through Intentionally Uncertain Highlighting." Proceedings of the 29th Annual Symposium on User Interface Software and Technology (Tokyo, Japan), UIST '16, October 16, 61–68. <https://doi.org/10.1145/2984511.2984538>.

Chang, Joseph Chee, Amy X. Zhang, Jonathan Bragg, et al. 2023. "CiteSee: Augmenting Citations in Scientific Papers with Persistent and Personalized Historical Context." ACM Conference on Human Factors in Computing Systems, April 19, 1–15. <https://doi.org/10.1145/3544548.3580847>.

Clark, Andy, and David Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 7–19.

Doctorow, Cory. 2025. "Commentary: Cory Doctorow: Reverse Centaurs." Locus Online, September 1. <https://locusmag.com/feature/commentary-cory-doctorow-reverse-centaurs/>.

Dragicevic, Pierre, Yvonne Jansen, Abhraneel Sarma, Matthew Kay, and Fanny Chevalier. 2019. "Increasing the Transparency of Research Papers with Explorable Multiverse Analyses." Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (New York, NY, USA), CHI '19, May 2, 1–15. <https://doi.org/10.1145/3290605.3300295>.

Engelbart, Douglas C. 1962. AUGMENTING HUMAN INTELLECT: A CONCEPTUAL FRAMEWORK.

Engelbart, Douglas C. and others. 1968. The Mother of All Demos. <https://dougengelbart.org/content/view/209/>

Gu, Ziwei, Ian Arawjo, Kenneth Li, Jonathan K. Kummerfeld, and Elena L. Glassman. 2024. "An AI-Resilient Text Rendering Technique for Reading and Skimming Documents." Proceedings of the CHI Conference on Human Factors in Computing Systems, May 11, 1–22. <https://doi.org/10.1145/3613904.3642699>.

Kale, Alex, Francis Nguyen, Matthew Kay, and Jessica Hullman. 2019. "Hypothetical Outcome Plots Help Untrained Observers Judge Trends in Ambiguous Data." *IEEE Transactions on Visualization and Computer Graphics* 25 (1): 892–902. <https://doi.org/10.1109/TVCG.2018.2864909>.

Kay, A., and A. Goldberg. 1977. "Personal Dynamic Media." *Computer* 10 (3): 31–41. <https://doi.org/10.1109/C-M.1977.217672>.

References

- Landau, Valerie. n.d. "How Douglas Engelbart Invented the Future." *Smithsonian Magazine*. Accessed April 6, 2026. <https://www.smithsonianmag.com/innovation/douglas-engelbart-invented-future-180967498/>.
- Licklider, Joseph CR. 1960. "Man-Computer Symbiosis." *IRE Transactions on Human Factors in Electronics*, no. 1: 4–11.
- Matuschak, Andy, Michael Nielsen, Andy Matuschak, and Michael Nielsen. 2019. *How Can We Develop Transformative Tools for Thought?* <https://numinousproductions.ttft>.
- Pirolli, Peter, and Stuart Card. 2005. *The Sensemaking Process and Leverage Points for Analyst Technology as Identified Through Cognitive Task Analysis*.
- Reichert, Leon, Zelun Tony Zhang, Elisabeth Von Oswald, Yuanting Liu, Yvonne Rogers, and Mariam Hassib. 2025. "AI, Help Me Think—but for Myself: Assisting People in Complex Decision-Making by Providing Different Kinds of Cognitive Support." *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, April 26, 1–19. <https://doi.org/10.1145/3706598.3713295>.
- Scaife, Mike, and Yvonne Rogers. 1996. "External Cognition: How Do Graphical Representations Work?" *International Journal of Human-Computer Studies* 45 (2): 185–213. <https://doi.org/10.1006/ijhc.1996.0048>.
- Xie, Liwenhan, <https://orcid.org/0000-0002-2601-6313>, View Profile, et al. 2024. "WaitGPT: Monitoring and Steering Conversational LLM Agent in Data Analysis with On-the-Fly Code Visualization." In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. ACM Other Conferences. <https://doi.org/10.1145/3654777.3676374>.
- Ye, Katherine, Wode Ni, Max Krieger, et al. 2020. "Penrose: From Mathematical Notation to Beautiful Diagrams." *ACM Transactions on Graphics* 39 (4). <https://doi.org/10.1145/3386569.3392375>.

References

- Hou, Xinying, Zihan Wu, Xu Wang, and Barbara J. Ericson. 2024. "CodeTailor: LLM-Powered Personalized Parsons Puzzles for Engaging Support While Learning Programming." Proceedings of the Eleventh ACM Conference on Learning @ Scale, July 9, 51–62. <https://doi.org/10.1145/3657604.3662032>.
- Kulik, James A., and J. D. Fletcher. 2016. "Effectiveness of Intelligent Tutoring Systems: A Meta-Analytic Review." Review of Educational Research 86 (1): 42–78. <https://doi.org/10.3102/0034654315581420>.
- Maloney, John, Mitchel Resnick, Natalie Rusk, Brian Silverman, and Evelyn Eastmond. 2010. "The Scratch Programming Language and Environment." ACM Trans. Comput. Educ. 10 (4): 16:1–16:15. <https://doi.org/10.1145/1868358.1868363>.
- Papert, Seymour. 1980. Mindstorms: Children, Computers, and Powerful Ideas. Basic Books, Inc.
- Ren, Haocheng, Muzhe Wu, Gregory Thomas Croisdale, Anhong Guo, and Xu Wang. 2025. "Rubikon: Intelligent Tutoring for Rubik's Cube Learning Through AR-Enabled Physical Task Reconfiguration." Proceedings of the 2025 ACM Designing Interactive Systems Conference (New York, NY, USA), DIS '25, July 4, 3549–62. <https://doi.org/10.1145/3715336.3735788>.
- Soloway, Elliot, Mark Guzdial, and Kenneth E. Hay. 1994. "Learner-Centered Design: The Challenge for HCI in the 21st Century." Interactions 1 (2): 36–48. <https://doi.org/10.1145/174809.174813>.
- VanLehn, Kurt. 2006. "The Behavior of Tutoring Systems." Int. J. Artif. Intell. Ed. 16 (3): 227–65.
- VanLehn, Kurt. 2011. "The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems." Educational Psychologist 46 (4): 197–221. <https://doi.org/10.1080/00461520.2011.611369>.
- Weitekamp, Daniel, Erik Harpstead, and Ken R. Koedinger. 2020. "An Interaction Design for Machine Teaching to Develop AI Tutors." Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, April 21, 1–11. <https://doi.org/10.1145/3313831.3376226>.