How Human Factors Can Improve Neural IR

Formerly known as

IR Intelligence: Introduction to Neural IR & Learning to Rank

Search Solutions 2020

Daedalus
Information Systems

Learning to Rank Model

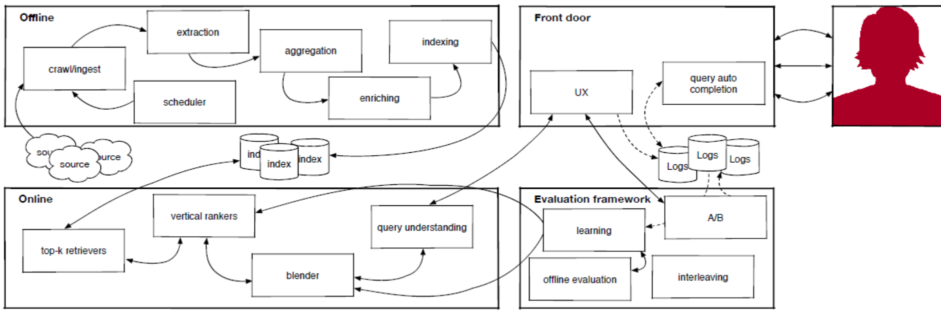

Michael Bendersky

FOLLO

Google Research
Verified email at google.com - Homepage
Information Retrieval   Natural Language Processing   Web Search & Data Mining   Machine Learning

TITLE	CITED BY	YEA
Discovering key concepts in verbose queries M Bendersky, WB Croft Proceedings of the 31st annual international ACM SIGIR conference on ...	304	200
Learning concept importance using a weighted dependence model M Bendersky, D Metzler, WB Croft Proceedings of the third ACM international conference on Web search and data ...	198	201
Quality-biased ranking of web documents M Bendersky, WB Croft, Y Diao Proceedings of the fourth ACM international conference on Web search and ...	126	201
Analysis of long queries in a large scale search log M Bendersky, WB Croft Proceedings of the 2009 workshop on Web Search Click Data, 8-14	123	200
Parameterized concept weighting in verbose queries M Bendersky, D Metzler, WB Croft Proceedings of the 34th international ACM SIGIR conference on Research and ...	122	201
Learning to rank with selection bias in personal search X Wang, M Bendersky, D Metzler, M Najork Proceedings of the 39th International ACM SIGIR conference on Research and ...	113	201

Daedalus Information Systems 3


Daedalus
Information Systems
4


See It Go
"Look," said Dick.
"See it go.
See it go up."
Jane said, "Oh, look!
See it go.
See it go up."
Daedalus
Information Systems
5

Information Seeking in Context 2004 – Dublin
Carol Kuhlthau calls out lack of collaboration between computer science and information science observing that they are isolated by practice and by distance (on college campuses)

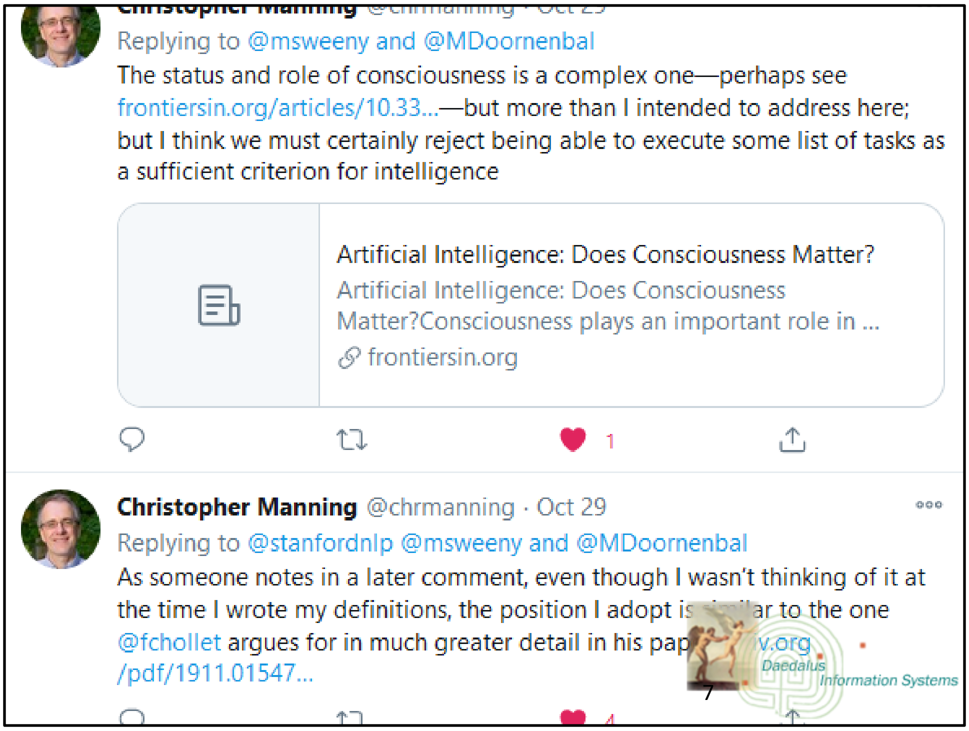

And then this happened…Christopher Manning who literally wrote the standard text for information retrieval, showed me the way.

No shared definition of what intelligence is
No shared set of tests to find out

Intelligence is the rate at which a learner turns its experience and priors into new skills at valuable tasks that involve uncertainty and adaptation. Definition is based on assessment of intelligent system with definition of intelligence grounded in: priors, experience, generalization of difficulty.

The measure of intelligence may be interpreted as a conversion rate between current state of information and the ability to perform well under an uncertain future.
On Measure of Intelligence by F Chollet – Robert Tjarko Lange

Human Intelligence

Daedalus Information Systems

9

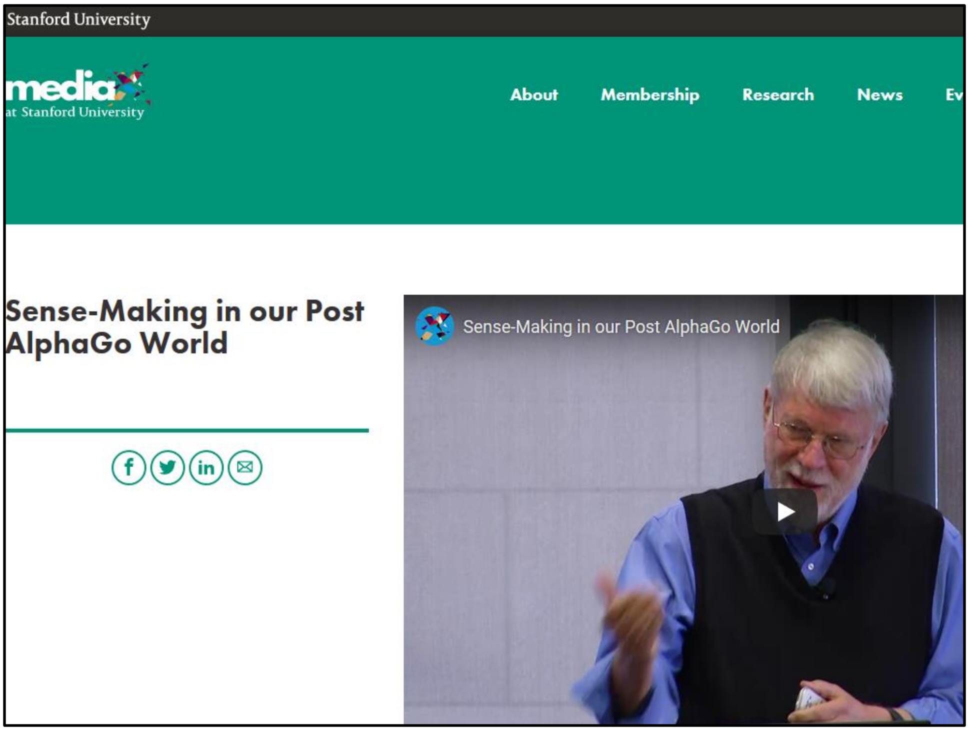

Current world is about shaping knowledge flows and reading context = reading context is different than reading signals (data), e.g. Brexit & USA 2016 presidential election – data is not information b/c it does not incorporate beliefs or values – sensemaking is a state that requires knowledge (an outcome of processed data)

Searle AI:

- Strong AI: system has a “mind” that can understand cognitive states (self aware)
- Weak AI: system can only simulate understanding (no mind)

https://mediax.stanford.edu/featured-events/john-seely-brown-mediax2017/
http://www.johnseelybrown.com/SensemakingStanford.pdf

Information Behavior Mai (2016)

Quality of Information is part of a spectrum

Data >> Information >> Knowledge

Quality depends on individual characteristics

- Contextual
- Situational
- Environmental
- Emotional

Daedalus
Information Systems

Jens Erik Mai, 2011

Further illustrates the singularly human aspect of intelligence

Consciousness

Humans have a knowledge of core concepts through experiencing the physical world

Consciousness allows for building more robust mental models that enable inference and prediction

Human consciousness:

- Introspection (self awareness)
- Empathy
- Transfer learning
- Adaptation
- Novelty
- Ambiguity

Daedalus Information Systems

12

David Gelernter, Computer science Yale University, artist and writer (Tides of Mind: Uncovering the Spectrum of Consciousness)
The human mind is not just creation of thought and collection of data; also a product of feelings, composite of sensations, memories, ideas that are worked and reworked over a lifetime.

Human consciousness is a first-person experience, AI consciousness can only be known from a 3rd person perspective (that of the programmer).

Ned Block (Philosophy faculty NYU): 2 distinct dimensions:

- Access consciousness (introspection, self-referential, monitor own processing
- Phenomenal consciousness: what it is “like” to be consciousness

Neural Intelligence

13

https://arnoldzwicky.org/2018/08/19/another-puzzle-in-cartoon-understanding/

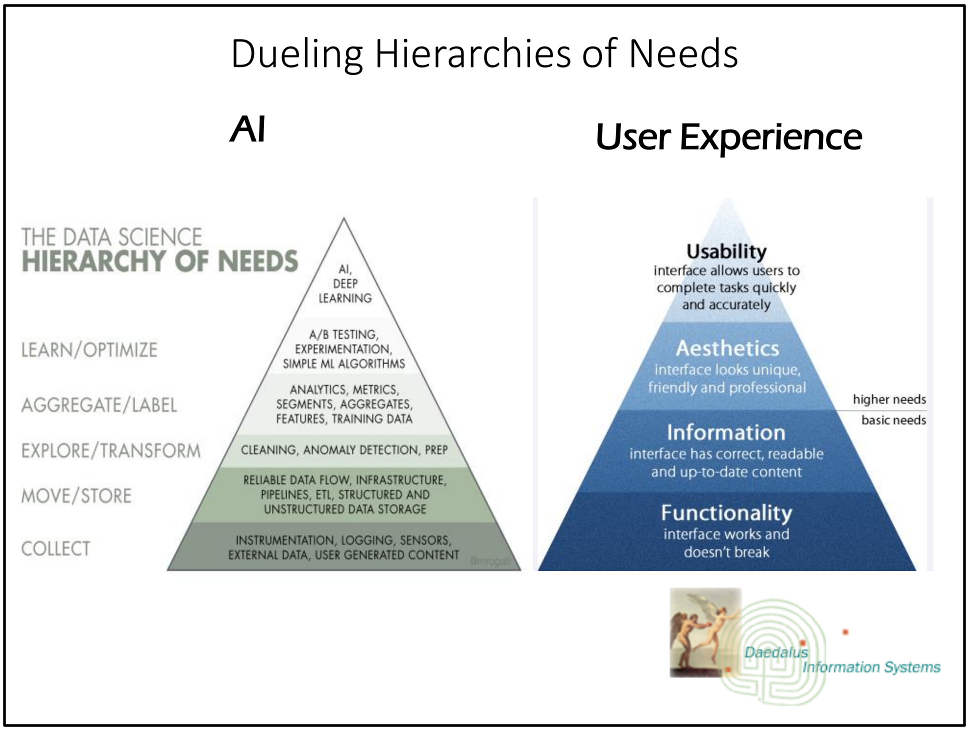

Human needs are influenced by feelings, engagement, values and moral codes
AI is influenced by data, data, data and programming

Machine Behavior
- Fundamentally different from human behavior (consciousness, empathy)
- Operates in a wider socio-technologic spectrum than human stakeholders
- Should not separate environment (the data used to train) from explicit behavior across diverse environments (variable use cases)

Influencing machine behavior to represent human needs requires a cross discipline effort, a human shaped hybrid-behavior

Start developing human-machine decision systems where human behavior shapes machine output through intricate involvement in training

- Understand the parameters of AI engineering to understand the outcomes and impact on human behavior
- Cooperation (traffic flow alterations), competition (gaming, chess, alpha go), coordination (robo trading financial markets): intersections of machine behavior

Neural IR Intent

Compositionality Principle

Customer intent actions:

- Query terms & refinements
- Dwell time
- # of results considered
- Time to first action
- Click counts
- First result clicked rank

Scan Browsing: information scanned based on relevance to changing tasks or transient information goals (Berrypicking)

Review browsing: some information is integrated into goal after deeper review (interest)

Daedalus Information Systems

15

Compositionality Principle: meaning of word compounds is derived from the meaning of individual words + the way they are combined, e.g. the components of language can be broken down into sub-components

Query intent = individual words that are possible indicators of customer intent. Uses term cooccurrence (proximity) models to improve retrieval relevance

Content units: specify need,

Intent units: modify the need in one of many possible ways

LTR attempts to capture and interpret hidden customer goals through clicked URLs

- Submitted queries
- Clicked URLs

Neural IR Intent Deconstruction

Query intent = individual words that are possible indicators of customer intent. Uses term cooccurrence (proximity) models to improve retrieval relevance

Intent words = articulated by customer to refine their information needs

Content words = core topic of query

Content unites further specify the need; intent units further modify the need in one of many possible ways

Daedalus Information Systems

Queries are classified as either content or intent

- Content words: nouns, verbs, adjectives
- Function words: pronouns, determiners, preposition, conjunctions

Tracked customer intent actions:

- Query terms & refinements
- Dwell time
- # of results considered
- Time to first action
- Click counts
- First result clicked rank

Scan Browsing: information scanned based on relevance to changing tasks or transient information goals (Berrypicking)

Review browsing: some information is integrated into goal after deeper review (interest)

Design for Neural IR

Daedalus
Information Systems

17

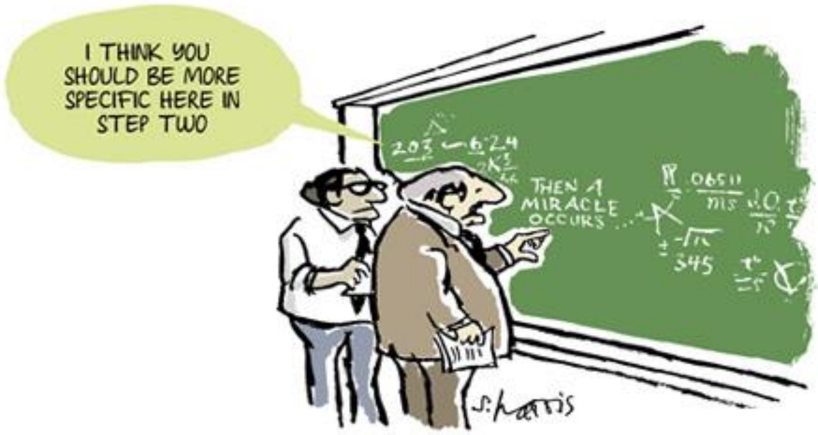

Daedalus Information Systems

This cannot be distilled from NLP or predictions on user behavior

Design for Machine Learning (together)

Define learning problem
- Inputs
- Outputs
- Types of training data needed

Generate good data
- Completeness
- Accurate
- Consistent
- Timely

Sketch out user and data flow (decision trees)

Test assumptions against prototype

Start with simple mechanism and move to complex

Daedalus Information Systems

20

Behavior cannot be separated from environment (data used to train AI) – focus has to be on characterizing AI behavior across diverse environments (persona, use case)

AI can help with correlated behavior – can interpret actual user data to detect patterns that can inform design changes (IoT, digital assistants). We can contribute by:

- Creation of interfaces that capture actual situation behavior, that maximize instructions to ML and AI to become a learning decision tree
- Classification for inputs, effective labels, etc.
- Enable design of interactions systems that can restate its understanding of tasks to be performed
- Provide user interfaces that allow user to bypass ML for implicit logic of user
- Create baseline to measure no go if system failure would cause serious consequences or irrevocable change (Postman's caution about new technology)
- Test the software across many environments

Start with Guiding Principles for AI

Principles	Applications

Principles are objectives and goals, the end state

Notes are the system rules, logic, rewards and feedback loops

Daedalus
Information Systems

Create Persona with an Empathy Profile

Find qualitative dimensions that define user understandings

Represent the emotions, not just the needs

Iterate as system learns

Daedalus Information Systems

Grounded Theory – qualitative components that differentiate shared experiences – synthesize beliefs and values with tastes and preferences

1. Define the solution
2. Understand the problem-solving options
3. Define the characteristics of a good solution (heuristics)
4. Map the environment (customer journey)
5. Benchmark success (quantitative, qualitative)

Empathy Persona Example

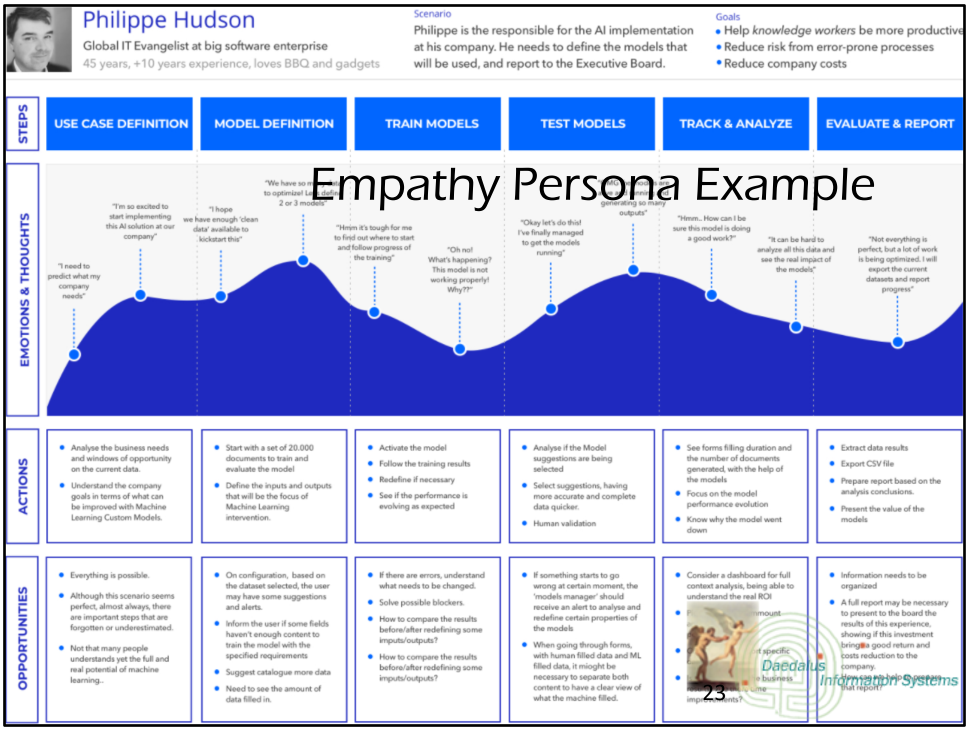

https://dribbble.com/shots/7595433-AI-Journey-Map

Map out the AI Journey

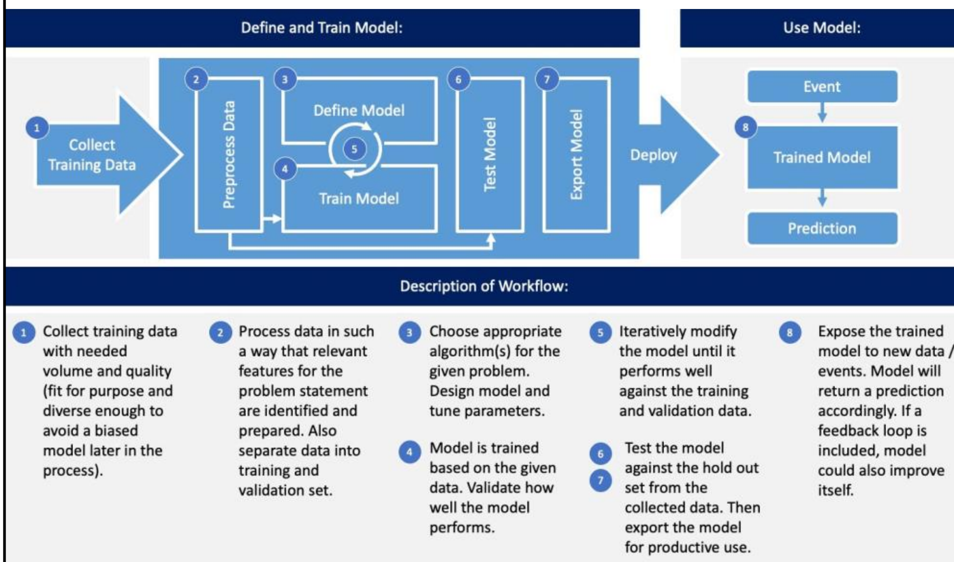

Human journey map focuses on Instinctual actions

Here, phases are influenced by thinking, expressing and acting

AI journey maps must be programmatically explicit as AI lacks intuitive movement

AI uses decision trees for journey progress. Info professionals can be useful for mapping more explicit and expansive decision trees

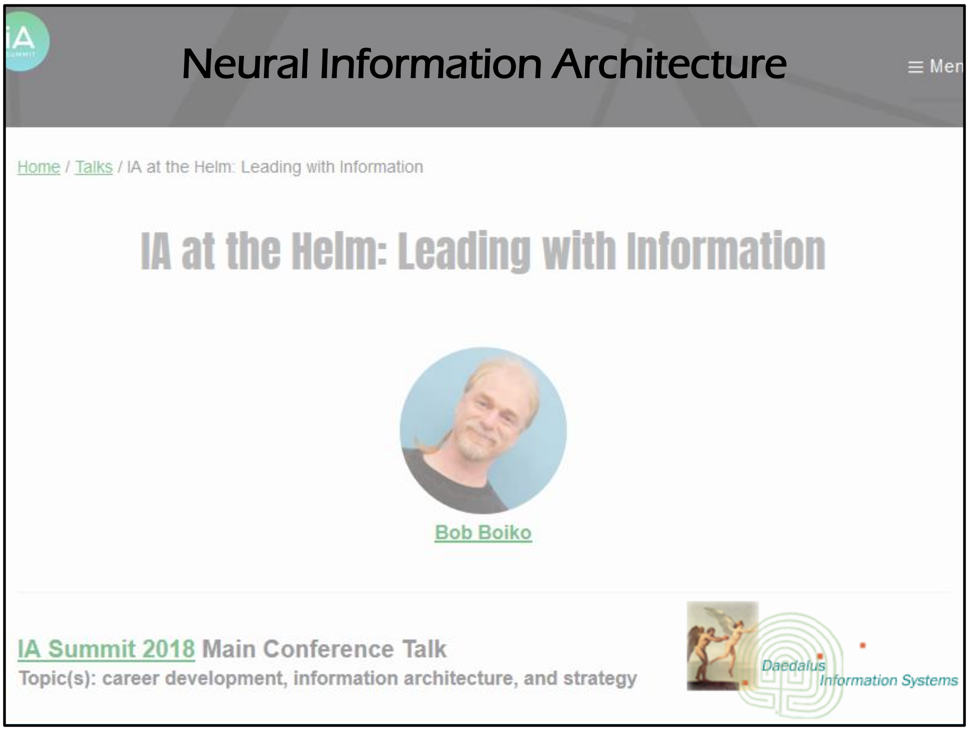

Bob Boiko:
Faculty, UW information School
6-part series called Information Systems from the Info Out

Much of the influence of IA practice stems from the power of naming objects in the system

Naming is a representative task that illustrates both the complexity and value of Information Architecture practice

Naming is a framing, an understanding of what everything is and how it fits into a system.

IA contributions to LTR

- Name object for cross system compatibility
- Problem definition and structure
- Connections
- Proto-typicality (mental models)
- Visual complexity (rely on text more than images)

Use Soft Information Architecture

Design for evolutionary processes
- Virtual reality
- Interactive architecture
- AI (self-organizing, educating, revising technology)

Status Quo Architecture= goal oriented

Soft Architecture = behavior-based

Designer must think within, not just about the system

Daedalus Information Systems

26

Status Quo Architecture= goal-oriented, e.g. laser focus on task at hand or problem

Soft Architecture = behavior-based, iterative, incorporates user input/output to iterate

Soft Architecture uses system to create various methods of presentation and structure to accommodate self-evolutionary/adaptive nature of system performance
(no more “set it and forget it”)

Structure uses feedback to influence behavior
- Inner structure = algorithms
- Physical structure = environment (digital, product, etc.)

Behavior based design can better incorporate changing nature of conditions and impact on system (self-driving car fatality)

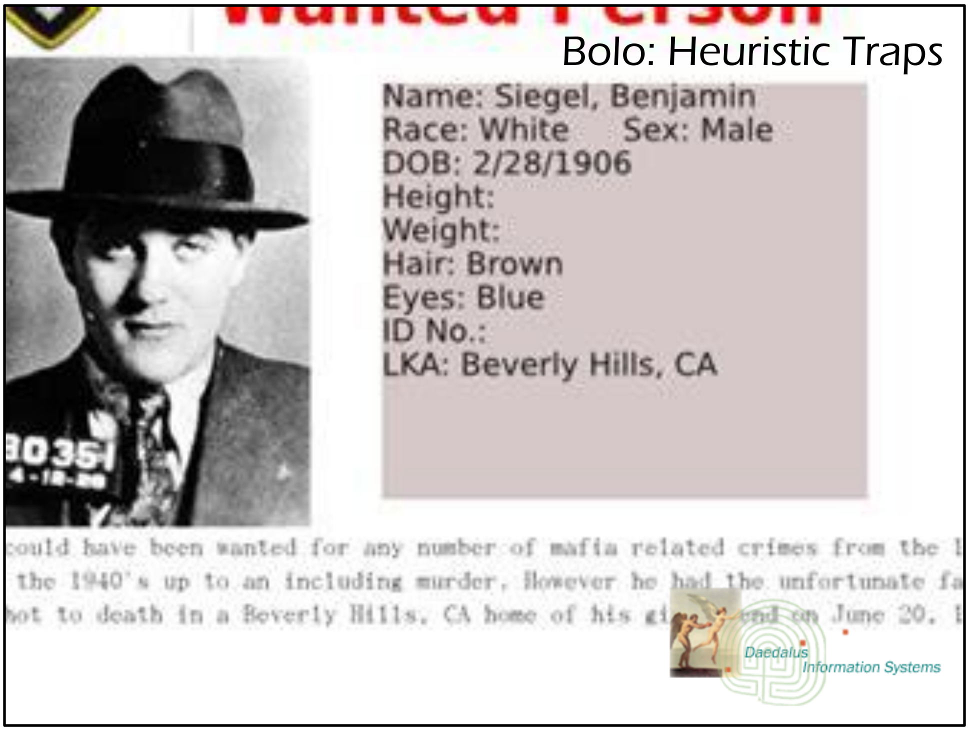

Familiarity: False confidence by failing to vigilant when faced with the known

Social Facilitation: Everyone is doing it so it must be okay

Expert Halo: False confidence in experts that we assume know what they are doing. So, it must be okay to follow them without question.

Anchoring:

Scarcity: If I don't do it now, I never will. Powder fever – it might be gone later

Acceptance: The cool kids are doing it and I want to be like them.

Why This?
Why Now?

Daedalus
Information Systems

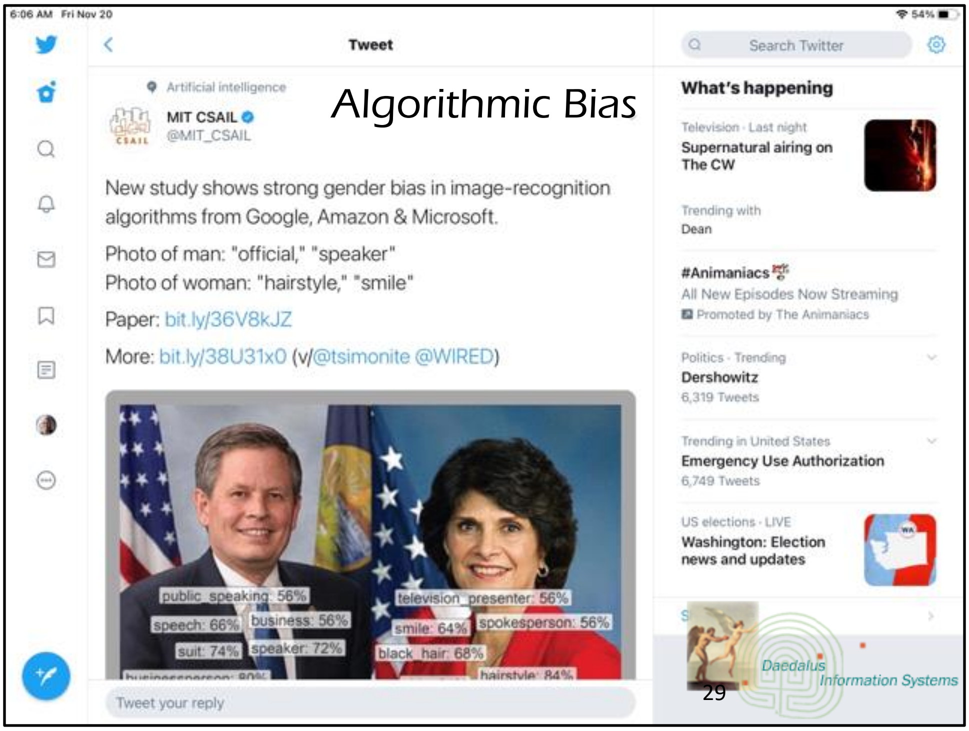

Technology inherits ideas and values of the group that develops it

Ben Schneiderman, winner ACM Turing Award, calls for a national algorithm safety board to monitor and assess safety of algorithms as they access social systems

Use Cases that are too Narrow

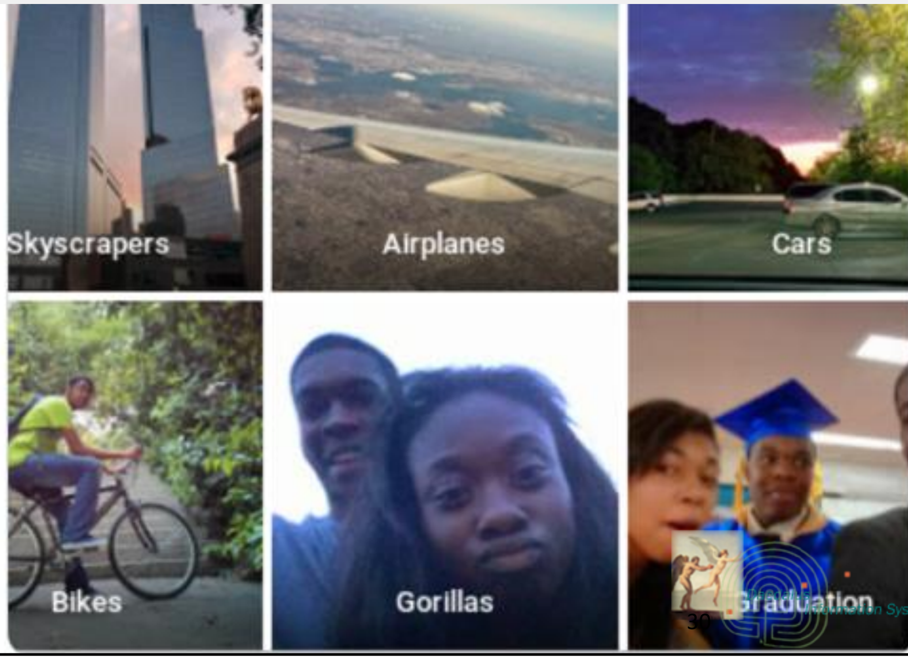

Poorly Conceived Objective Outcome

Built as a proof of concept for AI gone wrong with biased data

MIT AI Lab

Dataset was a sub-reddit dedicated to document the “disturbing reality of death.”

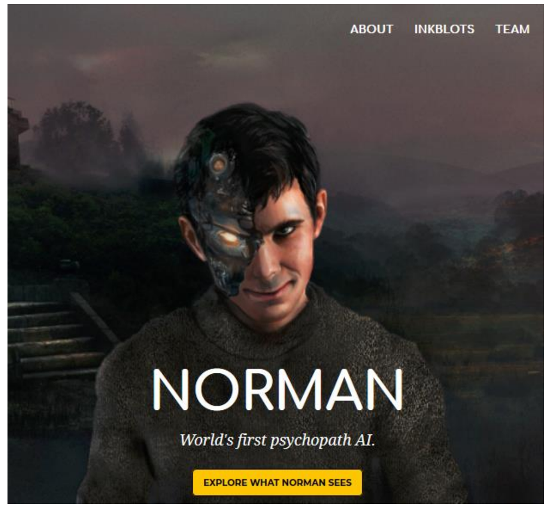

Daedalus Information Systems

31

http://norman-ai.mit.edu/ -

“Norman suffered from extended exposure to the darkest corners of Reddit and represents a case study on the dangers of Artificial Intelligence gone wrong when biased data is used in machine learning algorithms. “

Also produced Shelley (http://shelley.ai/), AI assisted horror stories, and Deep Empathy (https://deepempathy.mit.edu/) that produces images of what US cities would look like after conflict similar to that experienced in Syria

Poorly Designed Training Data

Bill Slawski @bill_slawski · 12s

Microsoft unveils a better-behaved chatbot after its last one became a NAZI

Microsoft unveils a better-behaved chatbot after its last one became …

Tech giant takes another pop at the artificial intelligence game with the release of a politer(and slightly stupider) machine mind 32

tems

LTR reliance on SERP Abandonment

- Studies show 41% of abandonments were bad, 32% abandonments were good, with the remaining 27% associated with alternate reasons (e.g. choosing a better query before considering returned SERP)
- Segmented as either good or bad abandonment

Search Abandonment Issues:

System ability to accurately categorizing abandoned queries into good and bad
Noisy indicators in user preference (emotional state, object, environmental conditions)
User bias

- User transformation of SERP into knowledge. SE cannot measure this so must calculate a sequence of interactions
- User biases
    - Position
    - Search engine
    - Cultural

Explore/Exploit hurdles

- System goal is to exploit the best ones (highest probability of success) to find better ones
- System must present results that produce feedback (explore) Results cannot always be perfect, or system stops looking for a better results (exploit)

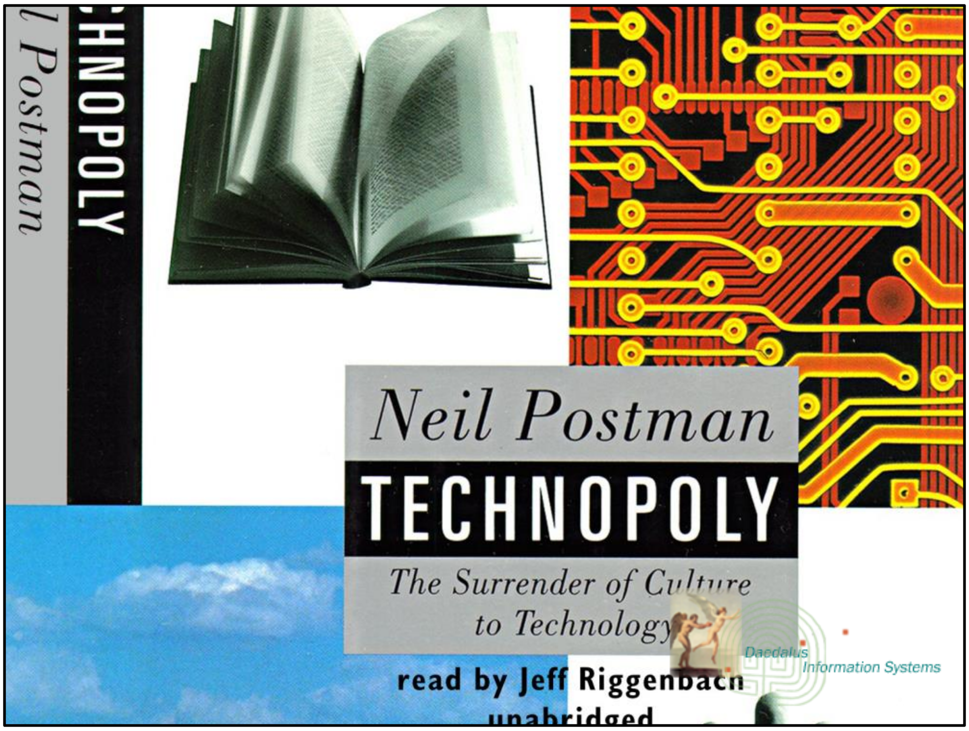

1. **All technological change is a trade-off**: The greater the wonders of technology, the greater will be its negative consequences. Culture always pays a price (algorithmic bias, social, psychological impacts)

2. **The advantages and disadvantages of new technologies are never distributed evenly among the population**: Some gain, some lose, few remain the same (predictive search, end of browsing, information induced blindness – systemic problems need action informed by information, not just more information)

3. **Embedded in every technology there is a powerful idea, sometime two or three ideas**: focus on all functions, not just the profitable ones.

4. **Technological change is not additive; it is ecological**: A new medium does not add something it changes everything (unintended consequences) often unpredictable and irreversible

5. **Media tends to become mythic**: Jaron Lanier (computationalism) – enthusiasm for the technology becomes a form a idolatry (AI is the new hammer and everything is a nail.) Capacity for good or evil requires human awareness and participation (human factors professionals included in development and execution)

Five things We Need to Know About Technological Change: Neil Postman: March 1998

(gratitude to Christine Emba, Wa post columnist for highlighting in her article)

Key Takeaways

Broaden scope of awareness

Understand the landscape and influences

Embrace new tools and methodologies

Daedalus Information Systems

35

We are the representatives of the qualitative self that is truer than the quantitative self represented by AI

Thank You

Embrace, engage, define, direct

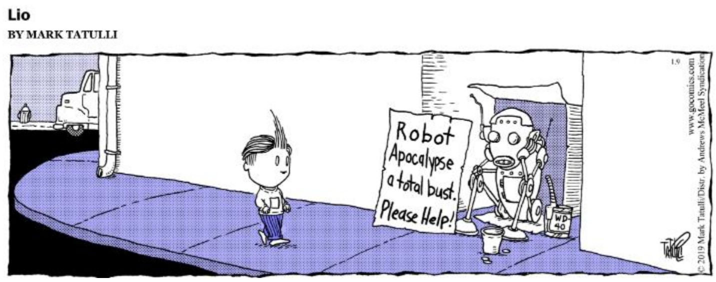

Marianne Sweeny
Principal
Daedalus Information Systems
sweeny48@uw.edu
@msweeny

Daedalus Information Systems

36

We can circumvent the AI apocalypse.