

Researching Responsible Natural Language Processing

Week 4 – 12th October 2026

What makes a ~~design~~ HCI paper?

John Vines

Professor of Design Informatics

john.vines@ed.ac.uk



THE UNIVERSITY *of* EDINBURGH

design
informatics

**Designing
Responsible**
Natural
Language
Processing

Outline of session

- Brief introduction to HCI as a design and research field
- Hear from you all about what papers you found
- Introduction to value sensitive design

What is a design paper?

~~**What makes a design paper?**~~

What makes a HCI paper?

HCI as a “design” field

“Human–computer interaction is a discipline concerned with the **design**, evaluation and implementation of interactive computing systems for human use and with the study of major phenomena surrounding them.”

Hewett, et al. 1992. *ACM SIGCHI curricula for human-computer interaction*.

HCI as a “design” field

HCI as a term sometimes gets used to describe the phenomena of humans interacting with computers – it’s a thing that people do

But it’s also an interdisciplinary field, that involves “design” in a broadest possible sense....

- HCI researchers study groups of people to understand how technology might support them better, drawing out implications or directions for technology design
- HCI researchers study groups of people to understand how technology is currently used by them, how they struggle with it, what works well and what doesn’t – drawing out implications for (re)designing technology, and developing good or bad practices
- HCI researchers experiment with designing new ways of interacting with or using technologies – building novel prototypes with new interaction modes, and giving these to people to test them out
- HCI researchers come up with methods and techniques that enable “non-designers” to input into design processes - something we call co-design, or participatory design..

HCI as a field of research



Started off as the intermingling of computer science and cognitive science and cognitive psychology

... then started to draw more on social psychology, social science and sociology ...

... and now connects many different disciplines together

Credit for diagram: Ruth Stalker-Firth, 2018.

Interdisciplinarity in HCI



John M. Carroll. 2026. Interdisciplinarity in HCI: What Happened in Our Big Tent? interactions 33, 2 (March - April 2026), 32–37.

<https://doi.org/10.1145/3796981>

The “Waves” and “Paradigms” of HCI

First Wave

1970s – 1990s

Expert users in workplaces

Focused on one-to-one interactions
between a “user” and a “machine”

Heavily influenced by cognitive
science and cognitive ergonomics

Lab-based studies

Modelling human behaviour

The “Waves” and “Paradigms” of HCI

First Wave

1970s – 1990s

Expert users in workplaces

Focused on one-to-one interactions
between a “user” and a “machine”

Heavily influenced by cognitive
science and cognitive ergonomics

Lab-based studies

Modelling human behaviour

Second Wave

1990s – 2000s

Many users in workplaces

Focused on groups of people
working with a range of applications
to get jobs done

Still a lot of cognitive science, but
also influence of social psychology
and ethnography

Human studies in contextual
settings and understood as
practices

The “Waves” and “Paradigms” of HCI

First Wave 1970s – 1990s	Second Wave 1990s – 2000s	Third Wave 2000s -
Expert users in workplaces	Many users in workplaces	Technology in homes, in culture, in leisure, in everyday settings
Focused on one-to-one interactions between a “user” and a “machine”	Focused on groups of people working with a range of applications to get jobs done	Technology no longer just desktop computers – different modes of interaction, moving between contexts
Heavily influenced by cognitive science and cognitive ergonomics	Still a lot of cognitive science, but also influence of social psychology and ethnography	Focus on understanding experience and emotion rather than just usability and tasks
Lab-based studies		
Modelling human behaviour	Human studies in contextual settings and understood as practices	More influence of ethnography, cultural theories, design research

Where to find HCI research - SIGCHI

[About](#)[People](#)[Conferences](#)[Resources](#)[News](#)[Join Us](#) ➔

Special Interest Group *On* Computer-Human Interaction

ACM SIGCHI is the leading international community of students and professionals interested in research, education, and practical applications of Human Computer Interaction.

[Learn more](#) ➔

*Committed to growing
our global and local
presence*

Our *Twenty Eight* Conferences

SIGCHI sponsors/co-sponsors 28 Human-Computer Interaction conferences annually, including our flagship conference, CHI (Human Factors in Computing Systems), which has been organized every year since 1983.

<https://sigchi.org/>

Where to find HCI research – SIGCHI Conferences



CHI

ACM Conference on
Human Factors in
Computing Systems



CSCW

ACM Conference On
Computer-
Supported
Cooperative Work
And Social
Computing



CUI

ACM Conversational
User Interfaces



DIS

ACM Designing
Interactive Systems



IUI

ACM Conference on
Intelligent User
Interfaces



UIST

ACM Symposium on
User Interface
Software and
Technology



UMAP

ACM Conference on
User Modeling,
Adaptation and
Personalization



CI


ACM Collective
Intelligence
Conference



COMPASS

ACM SIGCAS/
SIGCHI Conference
on Computing and
Sustainable
Societies

Where to find HCI research – SIGCHI on ACM DL



The screenshot shows the ACM Digital Library website with the SIGCHI (Special Interest Group on Computer-Human Interaction) page. The top navigation bar includes links for Journals, Magazines, Proceedings, Books, SIGs, Conferences, and People. A search bar is present with the text "Search ACM Digital Library". Below this, a secondary navigation bar for SIGCHI includes links for SIG Home, Publications, Upcoming Events, Authors, Affiliations, and Award Winners. The main header for SIGCHI is displayed in a large, stylized font, followed by the text "Special Interest Group on Computer-Human Interaction". A search bar within the SIGCHI section is also visible. The breadcrumb trail indicates the path: Home > SIGs > SIGCHI. The main content area features a paragraph describing the SIGCHI group, its interdisciplinary nature, and its goals. To the right of the text are two blue buttons: "Join SIG" and "Recommend ACM DL". Further right is a circular SIGCHI logo and a link to the "SIGCHI Website".

ACM DIGITAL LIBRARY

Association for Computing Machinery

University of Edinburgh

Browse About Sign in Register

Journals Magazines Proceedings Books SIGs Conferences People

Search ACM Digital Library

Advanced Search

SIG Home Publications Upcoming Events Authors Affiliations Award Winners

SIGCHI
Special Interest Group on Computer-Human Interaction

Search within SIGCHI

Home > SIGs > SIGCHI

SIGCHI • Special Interest Group on Computer-Human Interaction

The ACM Special Interest Group on Computer-Human Interaction is the world's largest association of professionals who work in the research and practice of computer-human interaction. We are an interdisciplinary group of computer scientists, software engineers, psychologists, interaction designers, graphic designers, sociologists, multi-media designers, and anthropologists, just to name some of the domains whose special expertise come to bear in this area. What brings us together is a shared understanding that designing useful and usable technology is an interdisciplinary process, and when done properly it has the power to transform persons' lives. SIGCHI has close to forty active local chapters across five continents, to promote local support networks for HCI professionals. SIGCHI members receive registration discounts to SIGCHI-sponsored and co-sponsored conferences, a subscription to "Interactions" magazine, and the monthly electronic ACM SIGCHI TechNews. SIGCHI membership also grants access to all SIGCHI publications in the ACM Portal.

Join SIG


Recommend ACM DL

SIGCHI

SIGCHI Website

<https://dl.acm.org/sig/sigchi>

Where to find HCI research beyond SIGCHI!



Categories > Engineering & Computer Science > Human Computer Interaction ▾

	Publication	h5-index	h5-median
1.	Computer Human Interaction (CHI)	139	185
2.	Proceedings of the ACM on Human-Computer Interaction	88	138
3.	International Journal of Human-Computer Interaction	83	124
4.	Behaviour & Information Technology	73	116
5.	IEEE Transactions on Affective Computing	72	115
6.	International Journal of Human-Computer Studies	64	104
7.	Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies	64	95
8.	Virtual Reality	62	115
9.	International Journal of Interactive Mobile Technologies	60	86
10.	ACM Transactions on Computer-Human Interaction	55	90
11.	ACM Designing Interactive Systems Conference	54	71
12.	ACM Symposium on User Interface Software and Technology	53	79
13.	ACM/IEEE International Conference on Human Robot Interaction	52	78
14.	Frontiers in Virtual Reality	52	76
15.	IEEE Virtual Reality Conference	52	66
16.	International Conference on Intelligent User Interfaces (IUI)	50	90
17.	Universal Access in the Information Society	48	74
18.	IEEE Transactions on Human-Machine Systems	45	71
19.	HCI International	45	64
20.	Multimodal Technologies and Interaction	42	62

https://scholar.google.co.uk/citations?view_op=top_venues&hl=en&vq=eng_humancomputerinteraction

Over to you all 2 mins each!

Tell us...

- How did you find exploring the CHI 2026 proceedings?
- Was it challenging, or easy, or ... ? Finding a paper for you and your team mate?
- Tell us about the papers you found for you and your team mate – why did you choose them?

Joshua!

For me...

RESEARCH-ARTICLE | **OPEN ACCESS** | 



Framing Responsible Design of AI for Mental Well-Being: AI as Primary Care, Nutritional Supplement, or Yoga Instructor?

For Coco...

RESEARCH-ARTICLE | **OPEN ACCESS** |  



Bloom: Designing for LLM-Augmented Behavior Change Interactions

Coco!

For me...

RESEARCH-ARTICLE | **OPEN ACCESS** |  

The AI Memory Gap: Users Misremember What They Created With AI or Without

For Joshua...

RESEARCH-ARTICLE | **OPEN ACCESS** |  

"I think this is fair": Uncovering the Complexities of Stakeholder Decision-Making in AI Fairness Assessment

Jenny!

For me...

RESEARCH-ARTICLE | **OPEN ACCESS** | 

Filtering the Invisible: A Feminist HCI Perspective on Informal Infra-structuring in Gig Labor

For Meryem...

RESEARCH-ARTICLE | **OPEN ACCESS** | 

Reimagining Sign Language Technologies: Analyzing Translation Work of Chinese Deaf Online Content Creators

Meryem!

For me...

RESEARCH-ARTICLE | **OPEN ACCESS** |   

Does My Chatbot Have an Agenda? Understanding Human and AI Agency in Human-Human-like Chatbot Interaction

For Jenny...

RESEARCH-ARTICLE | **OPEN ACCESS** |   

Writing with AI Can Reduce Gender Bias in Hiring Evaluations

Clodnah!

For me...

RESEARCH-ARTICLE | **OPEN ACCESS** |   | 

“I Don’t Think RAI Applies to My Model” – Engaging Non-champions with Sticky Stories for Responsible AI Work

For Adam...

RESEARCH-ARTICLE | **OPEN ACCESS** |  

Participatory AI Justice in HCI: A Scoping Review

Adam!

For me...

RESEARCH-ARTICLE | **OPEN ACCESS** |  

Cocoa: Co-Planning and Co-Execution with AI Agents

For Tianrui...

RESEARCH-ARTICLE | **OPEN ACCESS** | 

Making Multimodal LLMs Reliable Chart Data Extractors: A Benchmark and Training Framework

Tianrui!

For me...

RESEARCH-ARTICLE | **OPEN ACCESS** |  

PaperTok: Exploring the Use of Generative AI for Creating Short-form Videos for Research Communication

For Evelyn...

RESEARCH-ARTICLE | **OPEN ACCESS** |  

Lost in Transcription: Subtitle Errors in Automatic Speech Recognition Reduce Speaker and Content Evaluations

Evelyn!

For me...

RESEARCH-ARTICLE | **OPEN ACCESS** |  

A Shared Look: Detecting Deepfakes with Inter-Subject Neural Synchrony

For Clodnah...

RESEARCH-ARTICLE | **OPEN ACCESS** |  

Queer, Nonbinary, or Ambiguous? Rethinking Voice Labels through Queer Theory in HCI

Jamie!

For me...

RESEARCH-ARTICLE | **OPEN ACCESS** |  



The Privacy Paradox of LLMs: User Perceptions and the Reality of PII Leakage

For Termeh...

RESEARCH-ARTICLE | **OPEN ACCESS** |  



Orality: A Semantic Canvas for Externalizing and Clarifying Thoughts with Speech

Termeh!

For me...

RESEARCH-ARTICLE | **OPEN ACCESS** | 

Seeing, Hearing, and Knowing Together: Multimodal Strategies in Deepfake Videos Detection

For Anna...

RESEARCH-ARTICLE | **OPEN ACCESS** | 

High Accuracy and Hidden Disparities: Investigating Foundation Model Performance in Clinical Cognitive Assessment

Anna!

For me...

RESEARCH-ARTICLE | **OPEN ACCESS** |  

The Voice of Endo: Leveraging Speech for an Intelligent System That Can Forecast Illness Flare-ups


For Jamie...

RESEARCH-ARTICLE | **OPEN ACCESS** |  

The Role of ICTs in the Maintenance and Reproduction of Digital Border Assemblages

Types of research contributions in HCI



7 types of contributions:

Empirical

Artefacts

Methodological

Theoretical

Datasets

Survey

Opinion

What are the contributions of your papers?

Joshua:	Framing Responsible Design of AI for Mental Well-Being	Empirical
Coco:	The AI Memory Gap	Empirical
Jenny:	Filtering the Invisible	Empirical
Meryem:	Does my Chatbot have an agenda?	Empirical / Artefact
Clodnah:	“I Don’t Think RAI Applies to My Model”	Empirical
Adam:	Cocoa	Empirical / Artefact
Tianrui:	PaperTok	Empirical / Artefact
Evelyn:	A shared look	Empirical
Jamie:	The Privacy Paradox of LLMs	Empirical
Termeh:	Seeing, Hearing, and Knowing Together	Empirical
Anna:	The Voice of Endo	Empirical

Break time – 5 mins

Can “things” have values?

LANGDON WINNER

Do Artifacts Have Politics?


IN CONTROVERSIES ABOUT TECHNOLOGY AND SOCIETY, there is no idea more provocative than the notion that technical things have political qualities. At issue is the claim that the machines, structures, and systems of modern material culture can be accurately judged not only for their contributions of efficiency and productivity, not merely for their positive and negative environmental side effects, but also for the ways in which they can embody specific forms of power and authority. Since ideas of this kind have a persistent and troubling presence in discussions about the meaning of technology, they deserve explicit attention.¹

Writing in *Technology and Culture* almost two decades ago, Lewis Mumford gave classic statement to one version of the theme, arguing that “from late neolithic times in the Near East, right down to our own day, two technologies have recurrently existed side by side: one authoritarian, the other democratic, the first system-centered, immensely powerful, but inherently unstable, the other man-centered, relatively weak, but resourceful and durable.”² This thesis stands at the heart of Mumford’s studies of the city, architecture, and the history of technics, and mirrors concerns voiced earlier in the works of Peter Kropotkin, William Morris, and other nineteenth century critics of industrialism. More recently, antinuclear and prosolar energy movements in Europe and America have adopted a similar notion as a centerpiece in their arguments. Thus environmentalist Denis Hayes concludes, “The increased deployment of nuclear power facilities must lead society toward authoritarianism. Indeed, safe reliance upon nuclear power as the principal source of energy may be possible only in a totalitarian state.” Echoing the views of many proponents of appropriate technology and the soft energy path, Hayes contends that “dispersed solar sources are more compatible than centralized technologies with social equity, freedom and cultural pluralism.”³


An eagerness to interpret technical artifacts in political language is by no means the exclusive property of critics of large-scale high-technology systems. A long lineage of boosters have insisted that the “biggest and best” that science and industry made available were the best guarantees of democracy, freedom, and social justice. The factory system, automobile, telephone, radio, television, the space program, and of course nuclear power itself have all at one time or another been described as democratizing, liberating forces. David Lilienthal, in *T.V.A.: Democracy on the March*, for example, found this promise in the phos-



Value Sensitive Design (VSD)



“Value Sensitive Design is a theoretically grounded approach to the design of technology that accounts for human values in a principled and comprehensive manner throughout the design process.”



What are values?

“the word “value” refers simply to the economic worth of an object [...] we use a broader meaning of the term wherein a value refers to what a person or group of people consider important in life.”

Friedman et al., 2013. Value Sensitive Design and Information Systems

What examples of values can we think of?

What are values?

“the word “value” refers simply to the economic worth of an object [...] we use a broader meaning of the term wherein a value refers to what a person or group of people consider important in life.”

Friedman et al., 2013. Value Sensitive Design and Information Systems

Human welfare

Universal Usability

Accountability

Ownership and property

Trust

Courtesy

Privacy

Autonomy

Identity

Freedom from bias


Informed consent

Calmness

Environmental sustainability

What values are “in” a technology?

What are the values of ELM?: <https://elm.edina.ac.uk/>



Privacy

Trust

Accountability

Sharing Resources

Environmental
Sustainability

Whose values?

NatWest's Cora Chatbot Agent

Customers


Customers

Friends & Family

Shareholders

Technologists

Designers




Revenue Models

Financial
Regulations

Technology
Standards

Institutional
Responsibilities

Whose values?: tensions and contradictions!



How values get examined in VSD

Conceptual investigations

Who are the direct and indirect stakeholders affected by the design at hand? How are different stakeholders affected? What values are implicated? How should we engage in trade-offs among competing values in the design, implementation, and use of information systems? Should moral have greater weight than, or even trump, non-moral values? *Desk research, synthesis of prior studies, synthesis of ethical / moral frameworks.*

Empirical investigations

How do stakeholders understand individual values in the design context? How do they prioritize competing values? How do they prioritize individual values and instrumental considerations? Are there differences between espoused practice (what people say) compared with actual practice (what people do)? *Surveys, observations, interviews, workshops. Studies can inform design work, or evaluate already existing designs.*

Technical investigations

How does a given technology support certain values? How do they hinder certain values? How can the interactive and underlying architecture of a system be (re)designed to support these? *Technical experiments, iterative testing, “research through design”.*

VSD Envisioning Cards

<https://vsdesign.org/toolkits/envisioningcards/>

ENVISIONING CARDS.

A VALUE SENSITIVE DESIGN TOOLKIT

Use the Envisioning Cards to create ethical technology and improve your design practice. The 2nd Edition consists of 42 Envisioning Cards downloadable in PDF format.



Image

Changing Hands

Title

Envisioning Criterion

Each Envisioning Card is associated with one of five envisioning criteria: *Stakeholders, Time, Values, Pervasiveness, and Multi-lifespan.*

Stakeholders · Time · Values · Pervasiveness · Multi-lifespan

Changing Hands

Title

A single product can change hands once, twice, or more times during its lifecycle. It may be passed among family members (e.g., coming of age gift) or across town (e.g., consignment). How might use of the system change as the technology changes hands?

Design a scenario of your product changing hands. Imagine a specific challenge an individual, a family, or a community might face when wanting to shift ownership. What features might make this process smoother?

Design

Theme

Describes the theme of this Envisioning Card.

Design Activity

Suggested activity for exploring the theme of this Envisioning Card.

What values matter to you...?

Move into your ALP groups – discuss the following:

What five values should drive your work in ALP?


Are any of these values in tension with one another?

OPEN DESIGN MANIFESTO

- #01 DESIGN IS OF, BY & FOR THE PEOPLE
- #02 LEARNING IS EVERYTHING
- #03 OPEN COLLABORATION
- #04 STOP THINKING, START MAKING
- #05 DOCUMENTATION
- #06 MENTORING DEFIES AGEISM
- #07 NO DISCRIMINATION AGAINST PERSONS OR GROUPS
- #08 FREE REDISTRIBUTION
- #09 DERIVED WORKS
- #10 NO DISCRIMINATION AGAINST FIELDS OF ENDEAVOUR
- #11 DISTRIBUTION OF DESIGN
- #12 DESIGN MUST NOT RESTRICT OTHER DESIGNS
- #13 DESIGN MUST BE TECHNOLOGY-NEUTRAL



elepraying



<https://designjustice.org/principles-overview>

Any questions?

If you have any questions about this week, contact me at : john.vines@ed.ac.uk



THE UNIVERSITY *of* EDINBURGH

design
informatics

**Designing
Responsible**
Natural
Language
Processing

Extra Reading

Empirical Research

Studies that give insights about the design and use of technologies based on observation and data gathering.

“In HCI, empirical contributions arise from a variety of sources, including experiments, user tests, field observations, interviews, surveys, focus groups, diaries, ethnographies, sensors, log files, and many others.” (p.40)

Contributions are evaluated based on the level of insights they provide over prior work (e.g., deepened understanding of a user group, studying a novel context), and on the soundness of methods of data gathering and analysis.

“They only care to show us the wheelchair”: Disability Representation in Text-to-Image AI Models

Kelly Avery Mack
kmack3@uw.edu
Google Research and Paul G. Allen
School of Computer Science,
University of Washington
Seattle, WA, USA

Rida Qadri
ridaqadri@google.com
Google Research
San Francisco, CA, USA

Remi Denton
dentone@google.com
Google Research
New York, NY, USA

Shaun K. Kane*
shaunkane@google.com
Google Research
Boulder, CO, USA

Cynthia L. Bennett*
clbennett@google.com
Google Research
New York, NY, USA

ABSTRACT

This paper reports on disability representation in images output from text-to-image (T2I) generative AI systems. Through eight focus groups with 25 people with disabilities, we found that models repeatedly presented reductive archetypes for different disabilities. Often these representations reflected broader societal stereotypes and biases, which our participants were concerned to see reproduced through T2I. Our participants discussed further challenges with using these models including the current reliance on prompt engineering to reach satisfactorily diverse results. Finally, they offered suggestions for how to improve disability representation with solutions like showing multiple, heterogeneous images for a single prompt and including the prompt with images generated. Our discussion reflects on tensions and tradeoffs we found among the diverse perspectives shared to inform future research on representation-oriented generative AI system evaluation metrics and development processes.

1 INTRODUCTION

Generative AI is growing in capability and popularity, promising wide-ranging improved utility. However, literature points out how generative AI replicates existing biases in the world in its outputs [7, 12, 15, 26, 33, 34, 49, 64, 76, 84]. Consequently, there is a growing call to focus on the ethics of these technologies, and especially their impacts on minoritized groups [12, 33, 64], by centering the expertise of impacted communities [63, 64, 80]. While some work has investigated how to quantify these issues in large-language models (LLMs), less has focused on disability and text-to-image (T2I) models. Meanwhile, users have already demonstrated that these models can replicate existing disability stereotypes [22]. People with disabilities are particularly well-situated to identify ableism and provide insights into more respectful AI system development and outputs [33]. As such, this work sought expertise from disabled people with a variety of experiences, and presents recommendations regarding disability representation in AI-generated images. Thus we invited a variety of people who self-reported having a

Empirical Research

Studies that give insights about the design and use of technologies based on observation and data gathering.

“In HCI, empirical contributions arise from a variety of sources, including experiments, user tests, field observations, interviews, surveys, focus groups, diaries, ethnographies, sensors, log files, and many others.” (p.40)

Contributions are evaluated based on the level of insights they provide over prior work (e.g., deepened understanding of a user group, studying a novel context), and on the soundness of methods of data gathering and analysis.

“They only care to show us the wheelchair”: Disability Representation in Text-to-Image AI Models

Kelly Avery Mack
kmack3@uw.edu
Google Research and Paul G. Allen
School of Computer Science,
University of Washington
Seattle, WA, USA

Rida Qadri
ridaqadri@google.com
Google Research
San Francisco, CA, USA

Remi Denton
dentone@google.com
Google Research
New York, NY, USA

Shaun K. Kane*
shaunkane@

How to Drive a London Bus: Measuring Performance in a Mobile and Remote Workplace

Gary W. Pritchard¹, Pam Briggs², John Vines¹, and Patrick Olivier¹
¹Culture Lab,
School of Computing Science,
Newcastle University, UK
{firstname.lastname}@ncl.ac.uk

²PaCT Lab,
Department of Psychology,
Northumbria University, UK
p.briggs@northumbria.ac.uk

ABSTRACT

This paper examines how London bus drivers have responded to performance monitoring via a telematics device called *Drivewell*. This device calculates a score based on various recordable driving-related events like abrupt braking or irregular turning actions. Our qualitative methodology incorporated semi-structured interviews and ethnographic fieldwork in order to explore drivers' attitudes towards the system and its effect on driving behaviour and working conditions. Our findings illustrate how bus operators simultaneously accommodate and resist the demands *Drivewell* places upon them. Our work also demonstrates how this digital intervention acts in conjunction with other driver-related technologies, creating a unique digital ecosystem on the modern London bus. Our research contributes to HCI understandings of digital surveillance and performance monitoring in the modern workplace.

[27], while minimizing cost and making the entire system more environmentally sound.

These developments have led to some drivers acquiring new responsibilities and being subject to new forms of performance monitoring. For instance, it is now relatively straightforward to install in-vehicle data recorders (IVDR) that monitor driving quality indicators like speed, braking and cornering. Individual drivers can benefit directly from this by registering to either a Pay As You Drive (PAYD) insurance system that charges for distance travelled, or one that rewards good driving by offering lower premiums to those who achieve target driving scores [23]. On an organizational scale, professional drivers may also be signed up to 'fleet telematics' systems [18] that are used to monitor driving across an entire fleet. When integrated into driver appraisals these can become highly effective performance management infrastructures that reward good driving and penalize bad.

Artefacts

“Artifacts, often prototypes, include new systems, architectures, tools, toolkits, techniques, sketches, mockups, and envisionments that reveal new possibilities, enable new explorations, facilitate new insights, or compel us to consider new possible futures. New knowledge is embedded in and manifested by artifacts and the supporting materials that describe them.” (p40)

Artefact contributions might be evaluated as part of empirical research. But they do not always need to be – well documented, justified and described artefacts can be a contribution in themselves (see e.g., Gaver and Hook for a discussion!)



Artefacts

“Artifacts, often prototypes, include new systems, architectures, tools, toolkits, techniques, sketches, mockups, and envisionments that reveal new possibilities, enable new explorations, facilitate new insights, or compel us to consider new possible futures. New knowledge is embedded in and manifested by artifacts and the supporting materials that describe them.” (p40)

Artefact contributions might be evaluated as part of empirical research. But they do not always need to be – well documented, justified and described artefacts can be a contribution in themselves (see e.g., Gaver and Hook for a discussion!)

Designing with Data: An Annotated Portfolio

WILLIAM GAVER and **ANDY BOUCHER**, Interaction Research Studio, Northumbria University, Newcastle, England

This article presents an account of the Interaction Research Studio's distinctive approach to incorporating data into design, unpacking a vital but previously unexamined facet of our two-decade-long research programme. We demonstrate how data have served as a pragmatic tool in creating and engaging contextually aware computational products that provide alternatives to dominant task-oriented and consumerist narratives by promoting interpretation, imagination, spirituality and socio-cultural reflection. Leveraging an annotated portfolio approach, we revisit our diverse overarching themes. This article foregoes challenges that have arisen throughout our more theoretical treatments of design and practitioners navigating the landscape.

CCS Concepts: • **Human-centred computing**

Additional Key Words and Phrases: annotated portfolio with data, annotated portfolio



Methodological

“Methodological research contributions create new knowledge that informs how we carry out our work. Such contributions may improve research or practice.”
(p40)

Contributions that are methodological might help us study users of design better systems by, e.g., helping us make use of new forms of data, analyse data with new tools, help us generate better or more diverse ideas, etc.

Usually methodological contributions are evaluated based on its usefulness and novelty to other methods, and how reproducible or how straightforward to adopt it is for others.

Sketching AI Concepts with Capabilities and Examples: AI Innovation in the Intensive Care Unit

Nur Yildirim
Carnegie Mellon University
Pittsburgh, PA, USA
yildirim@cmu.edu

Susanna Zlotnikov
Carnegie Mellon University
Pittsburgh, PA, USA
susanna.zlotnikov@gmail.com

Deniz Sayar
Izmir University of
Economics
Izmir, Turkey
sayardeniz@gmail.com

Jeremy M. Kahn
University of Pittsburgh
Pittsburgh, PA, USA
jeremykahn@pitt.edu

Leigh A. Bukowski
University of Pittsburgh
Pittsburgh, PA, USA
lab108@pitt.edu

Sher Shah Amin
University of Pittsburgh
Pittsburgh, PA, USA
amin.shershah@gmail.com

Kathryn A. Riman
University of Pittsburgh
Pittsburgh, PA, USA
kathrynrیمان@pitt.edu

Billie S. Davis
University of Pittsburgh
Pittsburgh, PA, USA
bid8@pitt.edu

John S. Minturn
University of Pittsburgh
Pittsburgh, PA, USA
jsm120@pitt.edu

Andrew J. King
University of Pittsburgh
Pittsburgh, PA, USA
andrew.king@pitt.edu

Dan Ricketts
University of Pittsburgh
Pittsburgh, PA, USA
d.r@pitt.edu

Lu Tang
University of Pittsburgh
Pittsburgh, PA, USA
lutang@pitt.edu

Venkatesh Sivaraman
Carnegie Mellon University
Pittsburgh, PA, USA
vsivaram@andrew.cmu.edu

Adam Perer
Carnegie Mellon University
Pittsburgh, PA, USA
adamperer@cmu.edu

Sarah M. Preum
Dartmouth College
Hanover, NH, USA
sarah.masud.
preum@dartmouth.edu

James McCann
Carnegie Mellon University
Pittsburgh, PA, USA
jmccann@cs.cmu.edu


John Zimmerman
Carnegie Mellon University
Pittsburgh, PA, USA
johnz@cs.cmu.edu

Methodological

“Methodological research contributions create new knowledge that informs how we carry out our work. Such contributions may improve research or practice.”
(p40)

Contributions that are methodological might help us study users of design better systems by, e.g., helping us make use of new forms of data, analyse data with new tools, help us generate better or more diverse ideas, etc.

Usually methodological contributions are evaluated based on its usefulness and novelty to other methods, and how reproducible or how straightforward to adopt it is for others.



Theoretical

“Theoretical research contributions consist of new or improved concepts, definitions, models, principles, or frameworks. They are vehicles for thought. Whereas methodological contributions inform how we do things, theoretical contributions inform what we do, why we do it, and what we expect from it.”
(p41)

Theories can be descriptive in nature (i.e., they could describe how certain groups of users behave with certain systems) or predictive (i.e., if you introduce a certain design feature this is how people will behave). Good theoretical contributions tell us why things happen, not just what happens.

Expanding Explainability: Towards Social Transparency in AI systems

Upol Ehsan
Georgia Institute of Technology
Atlanta, GA, USA
ehsanu@gatech.edu

Q. Vera Liao
IBM Research AI
Yorktown Heights, NY, USA
vera.liao@ibm.com

Michael Muller
IBM Research AI
Yorktown Heights, NY, USA
michael_muller@us.ibm.com

Mark O. Riedl
Georgia Institute of Technology
Atlanta, GA, USA
riedl@cc.gatech.edu

Justin D. Weisz
IBM Research AI
Yorktown Heights, NY, USA
jweisz@us.ibm.com

ABSTRACT

As AI-powered systems increasingly mediate consequential decision-making, their explainability is critical for end-users to take informed and accountable actions. Explanations in human-human interactions are socially-situated. AI systems are often socio-organizationally embedded. However, Explainable AI (XAI) approaches have been predominantly algorithm-centered. We take a developmental step towards socially-situated XAI by introducing and exploring Social Transparency (ST), a sociotechnically informed perspective that incorporates the socio-organizational context into explaining AI-mediated decision-making. To explore ST conceptually, we conducted interviews with 29 AI users and practitioners grounded in a speculative design scenario. We suggested constitutive design elements of ST and developed a conceptual framework to unpack ST's effect and implications at the technical, decision-making, and organizational level. The framework showcases how ST can potentially calibrate trust in AI, improve decision-making, facilitate organizational collective actions, and cultivate holistic explainability. Our work contributes to the discourse of Human-Centered XAI by expanding the design space of XAI.

ACM Reference Format:

Upol Ehsan, Q. Vera Liao, Michael Muller, Mark O. Riedl, and Justin D. Weisz. 2021. Expanding Explainability: Towards Social Transparency in AI systems. In *CHI Conference on Human Factors in Computing Systems (CHI '21)*, May 8–13, 2021, Yokohama, Japan. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3411764.3445188>

1 INTRODUCTION

Explanations matter. In human-human interactions, they provide necessary delineations of reasoning and justification for one's thoughts and actions, and a primary vehicle to transfer knowledge from one person to another [65]. Explanations play a central role in sense-making, decision-making, coordination, and many other aspects of our personal and social lives [41]. They are becoming increasingly important in human-AI interactions as well. As AI systems are rapidly being employed in high stakes decision-making scenarios in industries such as healthcare [63], finance [76], college admissions [79], hiring [19], and criminal justice [37], the need for explainability becomes paramount. Explainability is not only sought by users and other stakeholders to understand and develop appropriate trust of AI systems, but also to support discovery of new knowledge and make informed decisions [52]. To respond to

Theoretical

“Theoretical research contributions consist of new or improved concepts, definitions, models, principles, or frameworks. They are vehicles for thought. Whereas methodological contributions inform how we do things, theoretical contributions inform what we do, why we do it, and what we expect from it.”
(p41)

Theories can be descriptive in nature (i.e., they could describe how certain groups of users behave with certain systems) or predictive (i.e., if you introduce a certain design feature this is how people will behave). Good theoretical contributions tell us why things happen, not just what happens.

Expanding Explainability: Towards Social Transparency in AI systems

Upol Ehsan
Georgia Institute of Technology
Atlanta, GA, USA
ehsanu@gatech.edu

Q. Vera Liao
IBM Research AI
Yorktown Heights, NY, USA
vera.liao@ibm.com

Michael Muller
IBM Research AI
Yorktown Heights, NY, USA
michael_muller@us.ibm.com

Mark O. Riedl
Georgia Institute of Technology
Atlanta, GA, USA
riedl@cc.gatech.edu

Justin D. Weisz
IBM Research AI
Yorktown Heights, NY, USA
justin_d_weisz@ibm.com

ABSTRACT

Strong Concepts: Intermediate-Level Knowledge in Interaction Design Research

KRISTINA HÖÖK, Royal Institute of Technology
JONAS LÖWGREN, Malmö University

Design-oriented research practices create opportunities for constructing knowledge that is more abstracted than particular instances, without aspiring to be at the scope of generalized theories. We propose an intermediate design knowledge form that we name *strong concepts* that has the following properties: is generative and carries a core design idea, cutting across particular use situations and even application domains; concerned with interactive behavior, not static appearance; is a design element and a part of an artifact and, at the same time, speaks of a use practice and behavior over time; and finally, resides on an abstraction level above particular instances. We present two strong concepts—social navigation and seamfulness—and discuss how they fulfil criteria we might have on knowledge, such as being contestable, defensible, and substantive. Our aim is to foster an academic culture of discursive knowledge construction of intermediate-level knowledge and of how it can be produced and assessed in design-oriented HCI research.

Categories and Subject Descriptors: H.5.2 [Information Interfaces and Presentation]: User Interfaces

General Terms: Design

Additional Key Words and Phrases: Design research, strong concepts, social navigation, seamfulness, bare-skin connection

Datasets

“A dataset contribution provides a new and useful corpus, often accompanied by an analysis of its characteristics, for the benefit of the research community.” (p41)

It's common to see a dataset be contributed alongside a new tool or method of analysis.

Datasets have been historically rare – however, we are increasingly seeing research funders pushing researchers to publish datasets from their research under open access policies, which is increasing their presence.

It's still very rare for a dataset to be a contribution on its own though.

CoAuthor: Designing a Human-AI Collaborative Writing Dataset for Exploring Language Model Capabilities

Mina Lee
minalee@cs.stanford.edu
Stanford University
United States

Percy Liang
pliang@cs.stanford.edu
Stanford University
United States

Qian Yang
qianyang@cornell.edu
Cornell University
United States

ABSTRACT

Large language models (LMs) offer unprecedented language generation capabilities and exciting opportunities for interaction design. However, their highly context-dependent capabilities are difficult to grasp and are often interpreted. In this paper, we argue that by curating HCI community can generative capabilities CoAUTHOR, a dataset in assisting creativity rich interaction GPT-3 across 144 THOR can address collaboration co-“collaborator” nally, we discuss discussion and tion design. sessions are

GPT-J [57], Jurassic-1 [36], Megatron-Turing-NLG [31], and Gopher [46]) can generate a wide variety of prose and dialogues with an unprecedented level of fluency out of the box. Through fine-tuning, these models can further become specialized at particular tasks, such as composing emails [8] or providing health consultation [58]. As a result, the HCI community has become interested in the opportunities surrounding LMs' generative capabilities. Some have started leveraging off-the-shelf LMs for rapid prototyping of novel language interactions [64]; others have started crafting end-to-end systems with fine-tuned LMs directly¹ [69], even as they can become production-ready interaction abilities not doing writing ideas to Does this al stories



Datasets

“A dataset contribution provides a new and useful corpus, often accompanied by an analysis of its characteristics, for the benefit of the research community.” (p41)

It's common to see a dataset be contributed alongside a new tool or method of analysis.

Datasets have been historically rare – however, we are increasingly seeing research funders pushing researchers to publish datasets from their research under open access policies, which is increasing their presence.

It's still very rare for a dataset to be a contribution on its own though.

CoAuthor: Designing a Human-AI Collaborative Writing Dataset for Exploring Language Model Capabilities

Mina Lee
minalee@cs.stanford.edu
Stanford University
United States

Percy Liang
pliang@cs.stanford.edu
Stanford University
United States

Qian Yang
qianyang@cornell.edu
Cornell University
United States

ABSTRACT

Large language models (LMs) offer unprecedented language generation capabilities and exciting opportunities for interaction design. However, their highly context-dependent capabilities are difficult to grasp and are often subjectively interpreted. In this paper, we argue that by *curating and analyzing large interaction datasets*, the HCI community can foster more incisive examinations of LMs' generative capabilities. Exemplifying this approach, we present CoAUTHOR, a dataset designed for revealing GPT-3's capabilities in assisting creative and argumentative writing. CoAUTHOR captures rich interactions between 63 writers and four instances of GPT-3 across 1445 writing sessions. We demonstrate that CoAUTHOR can address questions about GPT-3's language, ideation, and collaboration capabilities, and reveal its contribution as a writing “collaborator” under various definitions of good collaboration. Finally, we discuss how this work may facilitate a more principled discussion around LMs' promises and pitfalls in relation to interaction design. The dataset and an interface for replaying the writing sessions are publicly available at <https://coauthor.stanford.edu>.

GPT-J [57], Jurassic-1 [36], Megatron-Turing-NLG [31], and Gopher [46]) can generate a wide variety of prose and dialogues with an unprecedented level of fluency out of the box. Through fine-tuning, these models can further become specialized at particular tasks, such as composing emails [8] or providing health consultation [58]. As a result, the HCI community has become interested in the opportunities surrounding LMs' generative capabilities. Some have started leveraging off-the-shelf LMs for rapid prototyping of novel natural language interactions [64]; others have started crafting end-user-facing applications with fine-tuned LMs directly¹ [69], even though how soon such applications can become production-ready remain highly disputable [1, 24].

Harnessing LMs' generative capabilities to power interaction designs begins with a *holistic* understanding of these capabilities [5, 68]; this includes understanding what LMs can and cannot do under diverse interaction contexts. For example, when designing the mode of interaction between writers and GPT-3 for writing assistants, designers may ask: Can GPT-3 contribute new ideas to one's writing, or does it merely expand on existing ideas? Does this ideation capability differ in the context of writing fictional stories

Survey

“Survey research contributions and other meta-analyses review and synthesize work done on a research topic with the goal of exposing trends and gaps. Survey contributions are appropriate after a topic has reached a certain level of maturity.” (p42)

Notably, surveys were very rare in HCI conferences and journals until relatively recently – perhaps as they field is so diverse, and relatively new.

Surveys are evaluated on how well they analyse and organise existing work, and how well their analysis of literature reveals insights for new research and design work.

Understanding the LLM-ification of CHI: Unpacking the Impact of LLMs at CHI through a Systematic Literature Review

Rock Yuren Pang
University of Washington
Seattle, Washington, USA
ypang2@cs.washington.edu

Solon Barocas
Microsoft Research
New York, New York, USA
sbarocas@cornell.edu

Hope Schroeder
MIT
Cambridge, Massachusetts, USA
hopes@mit.edu

Ziang Xiao
Johns Hopkins University
Baltimore, Maryland, USA
ziang.xiao@jhu.edu

Danielle Bragg
Microsoft Research
Cambridge, Massachusetts, USA
danielle.bragg@microsoft.com

Kynneddy Simone Smith
Columbia University
New York, New York, USA
simone.kynneddy@columbia.edu

Emily Tseng
Microsoft Research
New York, New York, USA
etseng42@gmail.com

Abstract

Large language models (LLMs) have been positioned to revolutionize HCI, by reshaping not only the interfaces, design patterns, and sociotechnical systems that we study, but also the research practices we use. To-date, however, there has been little understanding of LLMs' uptake in HCI. We address this gap via a systematic literature review of 153 CHI papers from 2020-24 that engage with LLMs. We taxonomize: (1) domains where LLMs are applied; (2) roles of LLMs in HCI projects; (3) contribution types; and (4) acknowledged limitations and risks. We find LLM work in 10 diverse domains, primarily via empirical and artifact contributions. Authors use LLMs in five distinct roles, including as research tools or simulated users. Still, authors often raise validity and reproducibility concerns, and overwhelmingly study closed models. We outline opportunities to improve HCI research with and on LLMs, and provide guiding questions for researchers to consider the validity and appropriateness of LLM-related work.

Systematic Literature Review. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26–May 01, 2025, Yokohama, Japan. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3706598.3713726>

1 Introduction

Large language models (LLMs) are poised to transform the landscape of Human-Computer Interaction (HCI) research. Already, researchers have been using LLMs across the HCI research pipeline, from ideation and system development to data analysis and paper-writing [76]. Past work has shown rapid growth in the raw count of LLM-focused paper preprints, especially in HCI topics [118]. The explosion of LLM-related research has also led to rising discourse in HCI on the opportunities and challenges of LLM usage, including interview and survey studies with researchers to understand their practices [76], and workshops [4, 131] and social media commentary [67] in which scholars debate how the field ought to respond. The surge in LLM-related papers and discussions indicates a growing need to support scholars in understanding the potential and

Survey

“Survey research contributions and other meta-analyses review and synthesize work done on a research topic with the goal of exposing trends and gaps. Survey contributions are appropriate after a topic has reached a certain level of maturity.” (p42)

Notably, surveys were very rare in HCI conferences and journals until relatively recently – perhaps as they field is so diverse, and relatively new.

Surveys are evaluated on how well they analyse and organise existing work, and how well their analysis of literature reveals insights for new research and design work.

Understanding the LLM-ification of CHI: Unpacking the Impact of LLMs at CHI through a Systematic Literature Review

Rock Yuren Pang
University of Washington
Seattle, Washington, USA
ypang2@cs.washington.edu

Solon Barocas
Microsoft Research
New York, New York, USA
sbarocas@cornell.edu

Hope Schroeder
MIT
Cambridge, Massachusetts, USA
hopes@mit.edu

Ziang Xiao
Johns Hopkins University
Baltimore, Maryland, USA
ziang.xiao@jhu.edu

Danielle Bragg
Microsoft Research
Cambridge, Massachusetts, USA
dbragg@microsoft.com

Kynneddy Simone Smith
Columbia University
New York, New York, USA
simone.kynneddy@columbia.edu

Emily Tseng
Microsoft Research
New York, New York, USA
etseng42@gmail.com

Literature Reviews in HCI: A Review of Reviews

Evropi Stefanidi
University of Bremen
Bremen, Germany
evropi@uni-bremen.de

Thomas Kosch
HU Berlin
Bremen, Germany
thomas.kosch@hu-berlin.de

Stefan Schneegass
University of Duisburg-Essen
Essen, Germany
stefan.schneegass@uni-due.de

Marit Bentvelzen
Utrecht University
Utrecht, Netherlands
m.bentvelzen@uu.nl

Mikołaj P. Woźniak
University of Oldenburg
Oldenburg, Germany
mikolaj.wozniak@uni-oldenburg.de

Heiko Müller
University of Oldenburg
Oldenburg, Germany
OFFIS - Institute for Information
Technology
Oldenburg, Germany
heiko.mueller@uni-oldenburg.de

Paweł W. Woźniak
Chalmers University of Technology
Gothenburg, Sweden
pawel.wozniak@chalmers.se

Thomas Mildner
University of Bremen
Bremen, Germany
mildner@uni-bremen.de

Jasmin Niess
University of St. Gallen
St. Gallen, Switzerland
University of Oslo
Oslo, Norway
jasmin.niess@unisg.ch


Review. In CHI Conference on Human Factors in Computing Systems, May 01, 2025, Yokohama, Japan. ACM, 2025. 45/3706598.3713726

Opinion

Opinion contributions might also be referred to as essays or arguments – their purpose tends to be to persuade the reader about the arguments the paper presents.

These types of contributions tend to tackle emergent or complex problems in the field, often drawing on ideas and concepts from outside of HCI and bring them into contact with challenges in the field of HCI. They are very hard papers to write – and just as hard to evaluate and review.

Opinion papers get evaluated on the strength of their argument – strong arguments typically are those that make use of credible evidence, and consider different perspectives on the topic.




Opinion

Opinion contributions might also be referred to as essays or arguments – their purpose tends to be to persuade the reader about the arguments the paper presents.

These types of contributions tend to tackle emergent or complex problems in the field, often drawing on ideas and concepts from outside of HCI and bring them into contact with challenges in the field of HCI. They are very hard papers to write – and just as hard to evaluate and review.

Opinion papers get evaluated on the strength of their argument – strong arguments typically are those that make use of credible evidence, and consider different perspectives on the topic.



What are the contributions of your papers?

Working on your own, choose one of your papers:

Read the introduction, and skim the sections of the rest of the paper:

- How does it motivate the research – what is it responding to? Does it refer to a “lack” or “gap” in knowledge?
- How does it convey its “novelty”?
- How is it structured – what types of sections does it have?
- How does it convey its “contribution”?
- Is there anything unusual about its content based on papers in other fields you may have read?