



HUMAN-COMPUTER

INTERACTION LAB



4th ANNUAL SYMPOSIUM

THURSDAY MAY 25th, 2023

PRESENTATION SUMMARIES
AND OTHER INFORMATION



Human-Computer Interaction Lab (HCIL) University of Maryland, College Park

The Human-Computer Interaction Lab (HCIL), launched in 1983 at the University of Maryland, has a rich history of transforming the experience people have with new technologies. From understanding user needs to developing and evaluating these technologies, the lab's faculty, staff, and students have been leading HCI research for the last 40 years.

The HCIL was established as an interdisciplinary effort within the University of Maryland's Institute for Advanced Computer Studies (UMIACS). Today, HCIL participants include faculty, staff, and students from the following units on campus: Information Studies (iSchool), Computer Science, Psychology, Education, English, Engineering, Journalism, and American Studies. In addition, HCIL faculty are members in a number of centers and institutes on campus, including UMIACS, the Maryland Institute for Technology in the Humanities (MITH), the Applied Research Laboratory for Intelligence and Security (ARLIS), the Maryland Cybersecurity Center (MC2), the Social Data Science Center (SoDa), and the Dingman Center for Entrepreneurship.

This booklet contains Symposium presentation summaries and selected highlights of our faculty and students' news and accomplishments over the past year.

Please visit our website (<https://hcil.umd.edu>) for more information about the research happening in the HCIL.

If you would like more information, please contact:

Human-Computer Interaction Lab
2117 Hornbake Building, South Wing
4130 Campus Drive
College Park, Maryland 20742
hcil-info@umd.edu

HCIL Director's Letter

HCIL's 40th Annual Symposium | May 2023

I want to start this letter by winding back time a bit. For those reading this of a certain age, feel free to make the Wayne's World dream scene sound effect now.

The year was 1983, and technological advancements abounded. NASA was continuing to push forward in space exploration. ARPANET (the predecessor of the modern internet) moved to TCP/IP. Apple released Lisa, the first commercial computer with a GUI, and the wildly popular Apple IIe personal computer. C++ was released. The first cellular phone call was made. Compact disc (CD) players were developed. We saw our first glance of Mario in a new arcade game and our first reference to a "computer virus."

From a technological perspective, 1983 must have been incredibly exciting. But it also raised many questions regarding the design of these many new technologies – how would they be used? Were they user-friendly? Were they accessible? Had developers considered privacy, security, and ethical questions related to the design of these technologies? How might these technologies reshape everything from healthcare to education?

With this in mind, it's somewhat unsurprising that Ben Shneiderman launched the Human-Computer Interaction Lab in 1983. Just a year after the first ACM Conference on Human Factors in Computing Systems (CHI), there was an urgent need for research on the usability, utility, and accessibility of computers and other new technologies. In those early years, Ben, Catherine Plaisant, and their collaborators conducted important research that informed the development and use of hypertext, touch screen displays, and information visualization. If you've never checked out the history of the HCIL, I highly encourage you to visit the history section of our website and our YouTube page, which includes a number of our older video presentations.

Perhaps more importantly, the HCIL's core values were less focused on solving hard problems (although that was still important!) and more on building a supportive training environment for junior researchers and facilitating collaboration among the diverse group of researchers who came through the lab. These values are encapsulated in what Ben has called "The Maryland Way" – seven "sparks" that foster innovation. The Maryland Way provides an initial framework through which to identify a problem, iterate on research ideas, move beyond failures, and work hard to achieve a goal.

Beyond that, I view the core of The Maryland Way as an indicator of the lab's values. Yes, we strive hard to advance HCI research, and we've had significant success across a large number of domains. More importantly, however, The Maryland Way reflects our focus on community and the strength we can derive from it. This community can be found in events like the symposium, but also in more mundane activities throughout the year, ranging from our weekly brown bag lunches to paper clinics, retreats, group outings, and even our shared space on the second floor of Hornbake.

As we celebrate our 40th anniversary, I'm happy to report that the HCIL community remains strong. This year we welcomed three new faculty to the lab: Pablo Paredes (CS), Sheena Erete (iSchool), and Heera Lee (iSchool). I'm excited to share that two of our faculty – Hernisa Kacorri and David Weintrop – received tenure this spring, which is testament to the high quality of their research. We've also seen 12 PhD students successfully defend their dissertations, and 10 students finish

their master's degree. Many of our faculty and students have won awards and grants this past year, including five papers that received a best paper or honorable mention at CHI last month. I'd also like to congratulate Hal Daumé III, who is the lead investigator on a newly awarded \$20 million NSF grant focused on trustworthy AI.

This symposium, we also send a fond farewell to Niklas Elmqvist, the eighth director of the lab and a professor in the iSchool. In the near decade Niklas spent at UMD, he made major contributions to HCI research, particularly in data visualization. He oversaw the lab throughout the first year of the pandemic and worked hard to maintain our community during a time of tremendous uncertainty. I'm grateful for Niklas' leadership and guidance over the years, and I wish him well as he heads across the Atlantic – and much closer to home – to launch the (new) Center for Anytime & Anywhere Analytics (CA3) at Aarhus University.

Today, you'll hear from a wide range of HCIL faculty and students who will be sharing the latest results from research spanning AI, natural language processing, and data visualization to accessibility, privacy & security, education, and more. The program is stuffed full of content, with three plenary talks, two panels, and 38 lightning talks! During the lunch hour, you can join a tutorial on HuggingFace, talk shop at a themed lunch table, or chat with other attendees. Make sure to stay for the reception, where you can munch on delicious hors d'oeuvres while checking out demos and posters.

The work we do in the lab would not be possible without support, including from campus units, funding agencies, and our sponsors, so I want to offer my sincere gratitude to our industry sponsors (VEX Robotics, Adobe, and Rice Murtha Psoras) and our campus sponsors (College of Information Studies, Philip Merrill College of Journalism, Computer Science Department, the Applied Research Laboratory for Intelligence and Security, and the Institute for Advanced Computer Studies). I also want to thank Luis Cortes, Joel Chan, Rachael Bennett, Sarah Grun, and Craig Allen Taylor for their help organizing this event, as well as all the students, faculty, and staff who volunteered this week.

As we spend time together today to share ongoing research from the lab, I hope you are as inspired as I am by the people who have made HCIL one of the preeminent HCI research labs for the last 40 years.

Welcome to the 40th annual HCIL symposium!

A handwritten signature in black ink that reads "Jessica Vitak". The signature is fluid and cursive, with a large, sweeping flourish at the end of the name.

Jessica Vitak
Director, Human-Computer Interaction Lab
Associate Professor, College of Information Studies
University of Maryland, College Park



40th Annual Human-Computer Interaction Lab (HCIL) Symposium

May 25, 2023

Brendan Iribe Center, Room 0324, University of Maryland
All times listed in Eastern Time (ET)

Detailed Program

Time	Details
9:00am	Registration Opens (atrium) <ul style="list-style-type: none">Pick up your program, swag bag, and name tag, and sign up for a lunch option (coffee & tea available).
9:30am	Welcome Address <ul style="list-style-type: none">Jessica Vitak, Director, HCIL
10:00am	Plenary Talks <ul style="list-style-type: none">Sheena Erete: <i>A Method to the Madness: Applying an Intersectional Analysis of Structural Oppression and Power in HCI and Design</i>David Weintrop & Jimmy Lin: <i>Learning Computer Science with Robots: Supporting students across the K-12 Continuum</i>Susannah Paletz: <i>Emotions in Social Media: Computational Social Science Needs Good Social Science</i>
11:00am	Lightning Talks: Session 1 <ul style="list-style-type: none">13 short talks on topics spanning natural language processing, AI, and data visualization.
12:30pm	Lunch <p>We have multiple lunch options for you (signup required for some):</p> <ul style="list-style-type: none">Sign up for our short tutorial on HuggingFace (IRB-4105)Join HCIL faculty at one of our themed lunch tablesTake lunch inside or outside with other attendees
1:30pm	Panel <ul style="list-style-type: none">Topic: Teaching Design to UndergraduatesParticipants: Caro Williams-Pierce (moderator), Jason Aston, Tamara Clegg, Heera Lee, Victoria Van Hying



40th Annual Human-Computer Interaction Lab (HCIL) Symposium

May 25, 2023

Brendan Iribe Center, Room 0324, University of Maryland
All times listed in Eastern Time (ET)

Detailed Program (continued)

Time	Details
2:05pm	Lighting Talks: Session 2
	<ul style="list-style-type: none">• 13 short talks on topics addressing design challenges in accessibility and health.
3:45pm	Coffee Break
	<ul style="list-style-type: none">• Coffee, tea, and light snacks
4:05pm	Panel
	<ul style="list-style-type: none">• Topic: The promises and pitfalls of human-AI coordination: language translation as an emerging case• Panelists: Joel Chan (moderator), Marine Carpuat, Ge Gao, Michel Simard
4:40pm	Lightning Talks: Session 3
	<ul style="list-style-type: none">• 12 short talks on topics addressing the intersection of HCI and older adults, education, usable privacy/security, and more.
6:00pm	Reception (atrium)
	<ul style="list-style-type: none">• Stay for cake, hors d'oeuvres, and conversation!• Check out demos and student posters.

Sponsorships

The research that happens in the lab—and the lab itself—cannot happen without financial support, and we are grateful to our corporate, campus, and research sponsors, especially those who have contributed to HCIL and its members in the last year.

Industry Sponsors:

VEX Robotics
Adobe
Rice Murtha Psoras

Campus Sponsors:

College of Information Studies (iSchool)
Institute for Advanced Computer Studies (UMIACS)
Philip Merrill College of Journalism
Computer Science Department
Applied Research Laboratory for Intelligence and Security (ARLIS)

Research Sponsors:

Adobe
Army Research Lab (ARL)
Army Research Office
Google
Intelligence Advanced Research Projects Activity (IARPA)
Meta
National Institute on Disability, Independent Living, and Rehabilitation Research (NIDILRR)
National Institutes of Health (NIH)
National Science Foundation (NSF)
Office of Naval Research (Minerva Research Initiative)
Protocol Labs
Sigma Computing
Sloan Foundation
University of Maryland



**PLENARY SPEAKERS
&
PANEL DESCRIPTIONS**



Plenary Speaker: Sheena Erete, Associate Professor, College of Information Studies, University of Maryland

Talk Title: A Method to the Madness: Applying an Intersectional Analysis of Structural Oppression and Power in HCI and Design

Abstract: With increased focus on historically excluded populations, there have been recent calls for HCI research methods to more adequately acknowledge and address the historical context of racism, sexism, gendered racism, epistemic violence, classism, and so on. In this talk, I'll share results from a paper recently published in *ACM Transactions on Computer-Human Interaction*. The paper makes the following contributions: (1) identify the saturated site of violence; (2) identify the intersecting systems of power and who holds power (past and present); (3) describe the “conceptual glue” that binds these intersecting systems together and the assumption(s) that those who hold power are employing to guide their interactions; (4) examine the ways in which Black people are subjugated, surveilled, and/or expected to assimilate to “normative” ways of being and behaving; and (5) identify acts of resistance. This work contributes an alternative to traditional HCI and design methods that falsely perpetuate a lens of neutrality and colorblindness that centers on whiteness, innovation, and capitalism and ignores the history of State-sanctioned violence and structural oppression.



Bio: Sheena Erete is an associate professor in the College of Information at the University of Maryland, College Park and the founder and director of the Community Research and Design Collective. Her research explores and designs technologies used by geographically-bound communities to address social issues by considering social, cultural, and economic contexts as well as socio-technical infrastructures. She works to co-design sustainable technologies, practices, and policies with community organizations that aim to counter structural oppression using equity-centered, justice-oriented, assets-based approaches to research and design. Currently, she is focusing on issues such as equity in AI/ML tools, community safety, education, political efficacy, and economic development in communities that have been historically oppressed and resourced-constrained due to unfair policies and State violence.




Plenary Speakers: David Weintrop (Assistant Professor) and Yuhan Lin (PhD Student), Technology, Learning and Leadership Program, College of Education, University of Maryland

Talk: Learning Computer Science with Robots: Supporting Students Across the K-12 Continuum

Abstract: Programming and its associated skills are becoming foundational skills that all students need to succeed. Despite this growing reality, significant open questions remain related to how best to introduce learners to programming at different grades and, critically, how to connect learning experiences as youth move from basic, introductory experiences in early grades to more powerful programming tools and languages as they advance in their studies. In this plenary talk, we will share our experiences designing tools to support computer science learning across the K-12 continuum. This begins with learning environments designed for kindergarten-age learners and continues on through high school. This work is being conducted in collaboration with VEX Robotics and investigates design strategies and tools to scaffold learners at various stages. In particular, this talk will highlight the VEX 123 and Switch Mode.



Bio: David Weintrop is an Assistant Professor in the Department of Teaching & Learning, Policy & Leadership in the College of Education with a joint appointment in the College of Information Studies at the University of Maryland. His research focuses on the design, implementation, and evaluation of effective, engaging, and equitable computational learning experiences. His work lies at the intersection of design, computer science education, and the learning sciences. In support of his work, David received a National Academy of Education/Spencer postdoctoral fellowship and an NSF CAREER award. David has a Ph.D. in the Learning Sciences from Northwestern University and a B.S. in Computer Science from the University of Michigan.



Bio: Yuhan (Jimmy) Lin is a doctoral student in the Technology, Learning and Leadership PhD program at the University of Maryland's College of Education. His recent work involves creating a taxonomy for block-based programming environments. His research focuses on exploring ways to bridge the gap between block-based and text-based programming and using physical computing as a means to support computer science programming education.



**Plenary Speaker: Susannah B. F. Paletz, Associate Professor,
College of Information Studies, University of Maryland**

***Talk Title: Emotions in Social Media: Computational Social Science
Needs Good Social Science***

Abstract: A multitude of studies measure emotions in social media to better understand topics such as post sharing. However, much of this work draws on outdated, inaccurate, incomplete, and/or contested models of emotion, leaving the field with findings that may fail to replicate or capture the full picture of emotions in social media. In response, I and a multidisciplinary team created an annotation scheme based on the latest psychological science in terms of both emotion theory and cross-cultural psychology methods. I will describe this scheme, which covers all the multimedia content of a post and relies on native, trained annotators living and working in the country of study. I will then show comparisons to some popular automated text-based measures of emotion. I strongly encourage computational scientists to partner with social scientists, and vice versa, to bring the best theory and methods to any kind of computational social science.



Bio: Susannah B. F. Paletz is an Associate Professor at the College of Information Studies and affiliate of both the Social Data Science Center and the Applied Research Laboratory of Intelligence and Security (ARLIS). Her research focuses on teams, creativity, and culture, as well as applying psychology methods and theory to problems in the national interest. She earned her Ph.D. in Social/Personality Psychology from the University of California, Berkeley (2003), and has been a Research Psychologist at NASA Ames Research Center, a Postdoctoral Fellow/Research Associate at the University of Pittsburgh, and a Research Scientist at ARLIS. She is on the editorial board of *Small Group Research*, is an elected board member of the Interdisciplinary Network for Group Research (INGRoup), and is the founding Director of the Organizational Teams and Technology Research Society (OTTRS), a cross-UMD/DC-area research group. Her work has been funded by NSF, the Army Research Laboratory, the Office of Naval Research, and the Minerva Research Initiative.



PANEL: Teaching Design to Undergraduates

Abstract: Teaching design to undergraduates is a unique pedagogical puzzle. We have to (re)teach students how to learn a creative process rather than a series of facts, how to decentralize themselves and center the user instead, how to recognize failure as a learning opportunity, and how to value the messiness of post-it notes and hasty sketches. Our classes range in size from 10 to 150, require managing teams of undergraduate and graduate instructional support, occur along a spectrum of fully in-person to fully online asynchronous, and include students with a variety of interests (from ‘I love designing!’ to ‘It’s a required course, and I’d rather be building databases’).

We five panelists have considerable but different experiences with design as well as teaching design to undergraduates, and have brought this panel together to share our expertise with the HCIL Symposium attendees (and each other!). Below, we briefly describe our backgrounds and teaching experiences, a question we love answering, and a question we want answered at the panel.

Jason Aston has been a lecturer in HCI and information design at UMD since Fall 2021. His current topic areas center on visual design, data and information visualization as well as courses involving exploring different modes of interaction using physical computation tools and novel storytelling methods. His experiences in collaborative prototyping, experiential/digital design and the arts has led him to explore the human-technology relationship beyond one of just efficiency and convenience. *He asks (himself and his students):* “What’s so wrong with friction after all?”

Tamara Clegg has done research leveraging participatory design methods for over a decade. Since 2017 she has taught undergraduate design courses as well. Her background is in designing community-based technologies and experiences for youth and young adults from minoritized, resource constrained communities to engage in STEM in their everyday lives. *A question she loves answering:* How can we promote creativity and innovation in design processes? *A question she wants answered:* How do design thinking approaches fit into current industry practices?

Heera Lee has been teaching courses in User-Centered Design, UX research methods, and Capstone projects since Fall 2022. Her research focuses on affective computing, where she employs biosensors to investigate anxiety in English language learners from diverse cultures, including Deaf and intellectually disabled students. She also develops educational technologies that help instructors understand the emotions of these learners. These interests stem from her teaching experience as an English language instructor and her work with the Kennedy Krieger Institute. *A question that she loves answering is:* What steps can we take to ensure that design outcomes have a lasting positive impact on the community, rather than just serving as a one-time school project? *A question she wants answered:* How can we stay up-to-date with the constantly evolving trends in design?

Victoria Van Hying joined UMD in Fall 2020, and spent her first year here teaching from home. She has worked at the Library of Congress and Zooniverse on crowdsourcing projects and platforms for 6 years, which taught her valuable lessons in project management, product ownership, agile teamwork, identifying and meeting the needs of volunteer users, and cultural heritage organizations. This industry experience informed teaching large INST 362 classes— asynchronous for 90 students, and online synchronous version for 60. She worked with a team of UTAs and TAs each semester, and learned about scaling the design cycle learning experience in large classes. *A question she loves answering:* what didn’t work well and what did you learn from failure? *A question she wants answered:* How can we streamline our core course teaching to reduce the need/impulse for new teachers to reinvent courses?

Caro Williams-Pierce learned how to teach design to undergraduates in the midst of a worldwide pandemic, so instead of shifting from in-person teaching to online teaching, she has gone in the opposite direction! Her background is in designing videogames for mathematical play, so her experience with learner-centered design translated fairly easily into teaching user-centered design. She has taught design in a variety of contexts, from small in-person courses to fully online synchronous courses with 130 students and a large instructional team. *A question she loves answering:* How do you design a new course from scratch? *A question she wants answered:* What tips and tricks do you have for lessening student dependency on grades?

We are delighted to be conducting this panel on the complexity of teaching design to undergraduates, and we hope you will bring questions YOU want answered to our event!



PANEL: The Promises and Pitfalls of Human-AI Coordination: Language Translation as an Emerging Case

Abstract: With the advancement of AI technology, machine translation (MT) systems have become increasingly sophisticated and can produce natural sounding outputs in a wide range of languages. The promise of MT to enable communication across language barriers seems within reach. We can envision a near future, where people of diverse backgrounds can seamlessly interact, collaborate and exchange knowledge in any language using MT. However, fulfilling this promise will require addressing some important challenges to design reliable systems that truly support people's needs. The current panel brings together multidisciplinary researchers to discuss the existing challenges in human-MT coordination, as well as possible approaches to address them. With MT as a point of departure, the discussion will also consider how AI-powered language processing tools of the next generation should be (re)designed so that they can serve the social good through effective and responsible ways.

Moderator:

Joel Chan

Panelists:

Marine Carpuat is an Associate Professor of Computer Science at the University of Maryland, College Park. Her research focuses on multilingual natural language processing and machine translation. She is the recipient of an NSF CAREER award, research awards from Google and Amazon, best paper awards at *SEM and TALN, and an Outstanding Teaching Award. Marine also served as a Program Chair for NAACL 2022. Before joining the faculty at Maryland, she was a Research Scientist at the National Research Council Canada. She received a PhD in Computer Science and a MPhil in Electrical Engineering from the Hong Kong University of Science & Technology, and a Diplome d'Ingenieur from the French Grande Ecole Supelec.

Ge Gao is an Assistant Professor at the University of Maryland with joint appointments across the College of Information Studies (INFO) and the Institute for Advanced Computer Studies (UMIACS). Her research pays close attention to scenarios where people speaking different native languages interact with each other as members of the same task force, institution, or society. The overarching goal of this research is to advance the technology design for future workplaces and communities that are better connected, more effective, and smarter. Ge has published her work regularly at top-tier HCI venues, such as CHI, CSCW, and Ubicomp. Prior to Maryland, she obtained her PhD in Communication at Cornell University.

Michel Simard is a Senior Research Officer at the National Research Council of Canada (NRC), where he leads research on Machine Translation. He has expertise on various aspects of MT and other applications of natural language processing techniques to support communication across languages. He has recently served as area chair for NAACL, as program chair for AMTA and has co-organized a series of workshops on post-editing of Machine Translation. Before joining NRC, he worked as a researcher at Xerox Research Centre Europe (now Naver Labs Europe) and the Université de Montréal's RALI lab. He received a PhD in Computer Science from the Université de Montréal, an MSc in Computer Science from McGill University and a BSc in Mathematics & Computer Science from the Université de Montréal.



LIGHTNING TALKS: ABSTRACTS

[talks listed in order of presentation]

Clusters:

Session 1: AI, NLP, & Data Viz

Session 2: Accessibility & Health

Session 3: Potpourri

Bridging Human-AI Communication Gap for Visual-language Navigation

Lingjun Zhao, Khanh Nguyen, Hal Daumé III

Department of Computer Science

Contact: lzhao123@umd.edu

Problem. We study the problem of generating natural-language navigation instructions that can guide humans in a 3D visual environment, as well as communicate model uncertainty to enable those followers to adapt to errors. Improved solutions to this problem can potentially improve a wide range of real-world applications, including interfaces that aid visually impaired people, virtual assistants for travel, and smart houses that help residents manage and locate possessions.

Previous work trains models to mimic human-generated instructions. These models capture the grammatical structures and grounded meanings of human languages. However, because human communication is fundamentally a cooperative act, training agents based on such static datasets does not facilitate an agent’s ability to reason pragmatically, i.e. to anticipate how their utterances would be interpreted by humans. Our work provides a mathematical interpretation that allows for computational evaluation of pragmatic communication capabilities.

An important aspect of pragmatic communication that has largely been overlooked is to ensure that a listener understands when the speaker is uncertain. We aim to endow our agent with the ability to communicate its uncertainty at the word level. Unfortunately, previous work on estimating confidence of generative models requires word-level supervision, which can be unreliable for speaker models, or does not differentiate the uncertainty between meanings and forms. Our approach uses a discriminative follower model to predict the semantic uncertainty of generated instructions.

Methods and Findings. Using the Room-to-Room dataset, we learn speaker models that describe paths in photo-realistic environments. The models take as input a sequence of visual observations and actions along a path, and outputs a sequence of words. The speaker models are learned by fine-tuning a pretrained GPT-2 model, and training transformer- and LSTM- based encoder-decoder models.

To improve the accuracy and clarity of our speaker models’ instructions to better communicate with humans, we characterize the cognitive capabilities that enable humans to pragmatically use natural language to instruct others. We show that probabilistic instruction generation agents possess

similar capabilities, and design an evaluation scheme for probing those capabilities. Applying this scheme to examine the vanilla speaker models, we find that their theory-of-mind capability is severely lacking, as these agents poorly anticipate the listener’s interpretations of their instructions.

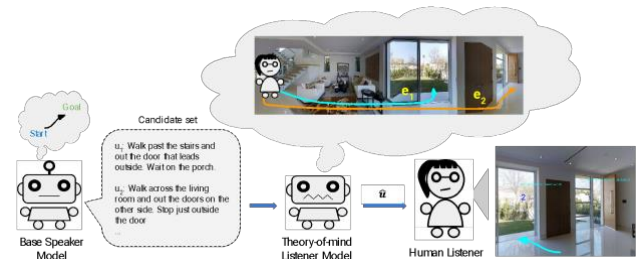


Figure 1 – The *base speaker* generates a small set of relevant candidate instructions to describe a trajectory. Then, it employs the *theory-of-mind model* to simulate how the human listener would follow each instruction in the candidate set. The speaker finally elects the candidate instruction that causes the theory-of-mind listener to generate the trajectory most similar to the intended trajectory. The output instruction is finally sent to the human listener for a real execution in the environment.

We augment the speaker models with better models of the listener and obtain a significant performance boost in guiding real humans, as shown in Figure 1. We find that using ensemble agents as a listener model consistently outperforms a single agent. Yet, there remains a considerable gap between our best pragmatic speaker and humans. We present empirical evidence to argue that closing this gap necessitates constructing better models of human behavior when interacting with AI-based agents. Building such models requires taking different approaches than imitating human-to-human interactions.

To ensure human listeners understand when the speaker is uncertain, we aim to provide humans the word-level confidence information of the speaker generated instructions. We construct a hallucination dataset and adapt a large-scale pretrained visual-language navigation model to estimate the word-level semantic uncertainty of the speaker generated instructions. We are conducting experiments to investigate whether the uncertainty information can help better guide humans for navigation.

Enhancing Interdisciplinary Information Foraging and Synthesis with Multilingual NLP

Yow-Ting Shiue¹, Joel Chan²

¹Department of Computer Science, ²College of Information Studies

Contact: ytshiue@cs.umd.edu

Interdisciplinary scholars explore literature both near and far from their home community of study. When crossing disciplinary boundaries, terminology problems become a major challenge. Researchers often have difficulties coming up with effective query terms for searching papers related to a specific concept they have in mind in an unfamiliar disciplinary or scholarly community, because that concept may be referred to as a totally different term.

In this project, we aim to explore how multilingual NLP techniques—which aim to help people overcome language barriers—may be adapted to enhance interdisciplinary information foraging. Machine Translation can not only be used to translate between two very different languages, but also quite similar ones such as dialects. It may also be applied to “translation” within the same language but diachronically (e.g. ancient to modern Mandarin) or synchronically (e.g. analyzing polarization).

As a first step, we adapt MUSE¹, a word embedding alignment algorithm that can map a source embedding space to a target space, to bridge the semantic gap between different research communities. For each community, we collect a corpus of papers based on a set of seed papers and their incoming/outgoing citations. The corpus is then utilized to train word embeddings, which form the community-specific semantic space. MUSE alignment between two semantic spaces allows us to compute word similarities across the two communities.

Table 1 shows some example results from our current prototype, tuned to a user exploring creativity research across different scholarly communities, such as psychology and management. MUSE returns interesting terms in the Management community which can be considered related to the concept of “examples” in Psychology. However, they are not direct equivalents of the query term. In our pilot testing, we observed that whether mappings are interesting and useful was uniquely defined by every researcher’s interest and perspective. This leads us to now formulate multilingual NLP-assisted interdisciplinary information foraging from an information retrieval perspective instead of a straight one-to-one dictionary translation task from one scholarly community to another.

Table 1 – Example query and similar terms with context in MUSE-aligned semantic space

Query: “examples” (Psychology) → ? (Management)
“substitutes” (Management)
Nor will the incumbent choose the mixed model as its response because when [...], monetization intensities are strategic substitutes .
This literature examines the temporal sequencing of experiential and vicarious learning (Bingham and Davis 2012) and conditions under which they act as substitutes or complements [...].
“pockets” (Management)
Kauffman's complexity theory seems to apply equally well to firms in coevolutionary pockets .
As part of the translation, possible parallels between the application of complexity catastrophe theory to coevolutionary pockets and studies by institutional theorists and social network analysts are discussed.

To more systematically test our ideas, we are planning an evaluation experiment akin to a query expansion scenario. Our MUSE system will expand a user’s query term from a home community by retrieving the top 10 most similar terms from a given target community, based on cosine similarity to the source query term in the MUSE aligned space; we also retrieve 5 context sentences for each expanded term. We will compare this against several other baseline systems, including 1) expanding with 10 terms based on cosine similarity in a universal embedding space trained with papers from all communities, 2) retrieving top 50 most similar context sentences based on contextualized sentence embeddings, and 3) a random baseline that returns a random sample of 50 sentences.

We predict low overlap between our MUSE system and the baselines due to the cross-disciplinary “language” barriers. We also predict that, compared to our MUSE system, sentences from the baselines will have higher *relevance* but lower *interestingness* to users. Based on the analysis results, our next step is to build an interactive system to support scholars to navigate and understand relevant concepts in an unfamiliar community.

¹ <https://github.com/facebookresearch/MUSE>

TutoAI: AI-assisted interactive mixed-media tutorial creation

¹Yuexi Chen, ²Vlad Morariu, ²Anh Truong, ¹Zhicheng Liu

¹Department of Computer Science, University of Maryland ²Adobe Research

Contact: ychen151@umd.edu

People turn to instructional videos to acquire in-demand skills, including cooking, assembling furniture, applying makeup, etc. Though instructional videos contain visual details to complete a task, users still need to exert great mental effort to follow. For example, users tend to skip less interesting parts, but the timeline does not have a clear boundary of steps, and they may miss important information. When users get stuck, they may want to replay a video part, but scrubbing the timeline is tedious and imprecise. Moreover, steps in instructional videos usually have dependencies, and it relies on users to figure them out. Recently, mixed-media tutorials, which unify videos, images, text, and diagrams in an interactive user interface, have shown promise as a more browsable alternative. For example, YouTube introduced Chapters to provide a high-level summary of a video: each chapter corresponds to a video segment with a short text description, a thumbnail, and a starting timestamp. Researchers have also proposed diverse mixed-media tutorials that present extracted informational units in non-linear representations for tasks such as applying makeup and cooking [1]. These mixed-media tutorials significantly reduced navigation effort by providing essential step and object information upfront and received enthusiastic responses in user studies.

However, creating such tutorials remains a challenge. On the one hand, manual creation is tedious, as the creator needs to spend a significant amount of time breaking down the video into steps, writing text descriptions, and extracting object information. On the other hand, AI-generated results are far from perfect and still rely on human annotators; existing methods to automatically generate mixed-media tutorials thus often employ domain-specific heuristics and cannot be generalized across multiple domains. Furthermore, mixed-media tutorials involve components across different modalities, including image, video, audio, and text, and the space of applicable machine learning (ML) methods is vast. In particular, we identify two challenges in choosing and assembling ML models for generating mixed-media tutorials. First, the same type of component can be scattered in multiple modalities and extracted by different ML methods (e.g., ingredients in cooking videos appear in both

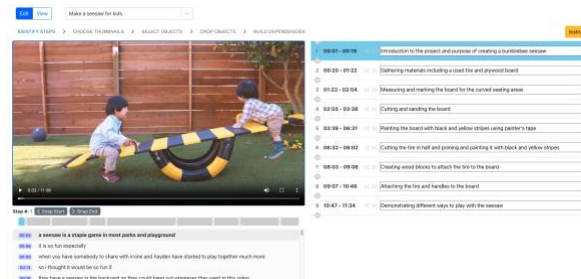


Figure 1 – TutoAI interface. We’re at stage 1: identify steps. This stage aims to break down the video into several steps and provide text descriptions and time boundaries for each step. On the left is a video player and its transcript (“Make a seesaw for kids”); on the right are the AI-generated video steps.

audio narrations and video frames). Second, different types of components extracted by different ML methods need to be integrated (e.g., after identifying the high-level steps and parts in furniture building videos, we need to associate the relevant parts for each step). Currently, there are no principles or guidelines on how to choose and assemble ML models to enhance mixed-media tutorials. To address these challenges, we present TutoAI, an AI-assisted tool for creating mixed-media tutorials in multiple domains. We make the following contributions: we identify essential components of mixed-media tutorials and present an approach to select and integrate ML methods into a pipeline to create mixed-media tutorials; an interactive prototype that integrates AI-generated results from a set of ML models; an empirical evaluation of the usability of the UI and the quality of mixed-media tutorials.

PAPERS

1. Anh Truong, Peggy Chi, David Salesin, Irfan Essa, and Maneesh Agrawala. 2021. Automatic generation of two-level hierarchical tutorials from instructional makeup videos. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–16.

Augmenting Human-AI Co-Writing with Interactive Visualization

Md Naimul Hoque, Niklas Elmqvist

College of Information Studies, University of Maryland, College Park

Contact: nhoque@umd.edu, elm@umd.edu

Writing is a fundamental human activity, but today we have the opportunity to leverage computational methods to help in this endeavor. Modern creative writers use a range of computational tools. The first kind is text editor-like writing software such as Microsoft Word or Google Docs. There are several professional and paid software available to writers. For example, Scrivener allows writers to organize a story in sections, add synopsis and notes to each section, and easily merge or swap between sections. Granthika is a similar sort of paid service where writers can outline a more detailed narrative world, including character descriptions, major events in a timeline, and causal constraints on the events. These tools primarily enhance the organizational capabilities of writers with limited feedback on the actual writing.

More recently, researchers have proposed several emerging writing support tools. A dominant trend is the use of LLMs. These tools can generate text based on a prompt, often helping writers explore alternate narrative worlds and creative angles [1]. Another set of tools uses NLP to extract patterns from the text that are otherwise difficult to notice. These tools typically focus on providing feedback while writers revise their text. For example, Du et al. [2] proposed a tool to help writers in their iterative revising process. Serman et al. [3] proposed an analytic model that allows writers to interact with the literary style of an article.

While existing tools are useful to writers, there exist several challenges for human-AI co-writing. For example, one critical challenge is enabling human autonomy and reducing ownership tensions in the process [4, 5]. Another relevant problem is the lack of trust among writers and AI as LLMs and other NLP models are typically black boxes and are difficult to interpret and explore intuitively. Other challenges include identifying plagiarism and bias in the AI-generated text and applying NLP to more abstract and complex narrative components such as characterization, and narrative structure [4, 5].

In this work, we hypothesize that some of these challenges arise from a lack of communication interface between writers and AI. We further propose *interactive data visualization* to be that interface (Figure 1). There are abundant works on using visualization for text analytics,



Figure 1 – *Visualization for co-writing*. A workflow for supporting human-AI co-writing with visualization.

human-AI collaboration, explainable AI (XAI), and digital humanities. However, it is still unknown how to translate and integrate these techniques to human-AI co-writing.

Our work aims to fill that void. To demonstrate our proposal, we present two case studies where we combine NLP and interactive visualization to support creative writing. The first case study is on mitigating social biases in fiction writing and the second is on the design of dynamic characters and scenes. We conclude by outlining our future work and broader impact.

REFERENCES

1. Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. Wordcraft: Story Writing With Large Language Models. In Proceedings of the ACM Conference on Intelligent User Interfaces. ACM, New York, NY, USA, 841–852. <https://doi.org/10.1145/3490099.3511105>
2. Wanyu Du, Vipul Raheja, Dhruv Kumar, Zae Myung Kim, Melissa Lopez, and Dongyeop Kang. 2022. Understanding Iterative Revision from Human-Written Text. In Proceedings of the Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics, Stroudsburg, PA, USA, 3573–3590. <https://doi.org/10.18653/v1/2022.acl-long.250>
3. Sarah Serman, Evey Huang, Vivian Liu, and Eric Paulos. 2020. Interacting with Literary Style through Computational Tools. In Proceedings of the ACM Conference on Human Factors in Computing Systems. ACM, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376730>
4. Piotr Mirowski, Kory W. Mathewson, Jaylen Pittman, and Richard Evans. 2022. Co-Writing Screenplays and Theatre Scripts with Language Models: An Evaluation by Industry Professionals. CoRR abs/2209.14958 (2022), 102 pages. <https://doi.org/10.48550/arXiv.2209.14958>
5. Ben Shneiderman. 2022. Human-Centered AI. Oxford University Press, Oxford, United Kingdom.

Fluid Transformers and Creative Analogies: Evaluating Large Language Models' Capacity for Augmenting Cross-Domain Analogical Creativity

Zijian “Jason” Ding¹, Arvind Srinivasan¹, Stephen MacNeil², Joel Chan¹

¹University of Maryland, College Park, ²Temple University

Contact: ding@umd.edu

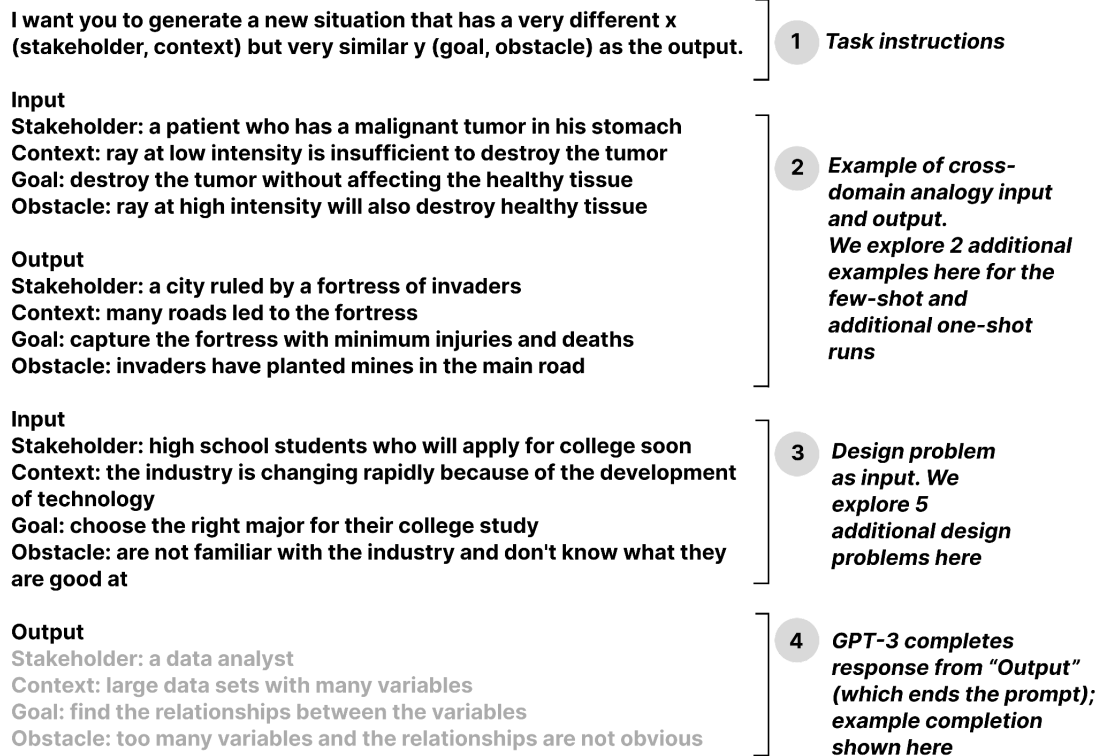


Figure 1 Our prompt-based learning prompts for generating cross-domain analogies..

Cross-domain analogical reasoning is a core creative ability that can be challenging for humans. Recent work has shown some proofs-of-concept of Large language Models' (LLMs) ability to generate cross-domain analogies. However, the reliability and potential usefulness of this capacity for augmenting human creative work has received little systematic exploration. In this paper, we systematically explore LLMs capacity to augment cross-domain analogical reasoning by generating analogies with multiple learning paradigms (zero-shot/one-shot/few-shot learning, as example shown in Figure 1). Across three studies, we found: 1) LLM-generated cross-domain analogies were frequently

judged as helpful in the context of a problem reformulation task (median 4 out of 5 helpfulness rating), and frequently (~80% of cases) led to observable changes in problem formulations, and 2) there was an upper bound of ~25% of outputs being rated as potentially harmful, with a majority due to potentially upsetting content, rather than biased or toxic content. These results demonstrate the potential utility --- and risks --- of LLMs for augmenting cross-domain analogical creativity.

ACKNOWLEDGMENTS

This project was supported by ONR N000142012506.

AnalogiLead: Improving Selection of Analogical Inspirations with Chunking and Recombination

Arvind Srinivasan, Joel Chan
 University of Maryland, College Park
 Contact: arvindrb@umd.edu

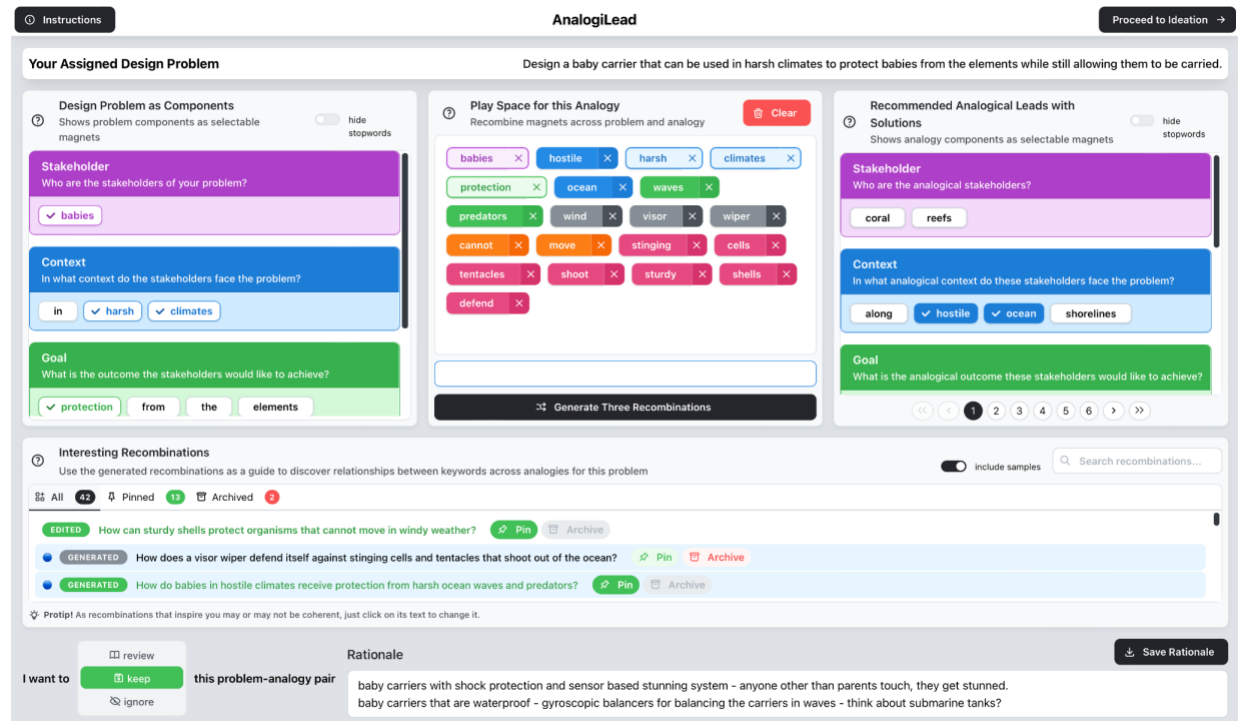


Figure 1 - AnalogiLead. This Proposed Interface facilitates divergent thinking through mixing & adding magnets to generate insightful questions & recognize beneficial analogies for selection.

Innovation is vital in various fields, and analogical thinking is a powerful tool for generating creative solutions to complex problems. However, recognizing analogies can be time-consuming, and successful recognition doesn't guarantee their adoption in innovation. In this thesis, A novel computational support system for analogical innovation is proposed that employs the cognitive mechanisms for chunking and recombination as mediums of interaction. *Chunking* involves identifying and extracting meaningful chunks or segments from a design problem into interactive tiles called *magnets* while *Recombination* involves combining these *magnets* to generate insightful questions that elicit divergent thinking. In this way, the proposed system aims to streamline the process of recognizing and selecting analogical inspirations for innovation while avoiding premature rejection and design fixation. To evaluate the effectiveness of the system, a within-subjects study involving 23

participants was conducted, comparing the proposed interface with a baseline. The study found that using chunking and recombination as interactive mechanisms helped reduce reflexive rejection of useful analogical leads: participants were 4x less likely to ignore analogical leads, and 12x less likely to change their decisions with our system compared to the baseline, at the cost of a small (~0.6x) increase in processing time per analogy. Overall, these results suggest that our proposed intervention may be an effective tool for facilitating the selection of beneficial analogies, fostering analogical innovation through computational support.

ACKNOWLEDGMENTS

This project was supported by ONR N000142012506.

The Impact of AI Explanations on Fairness in Human-AI Decision Making

Navita Goyal^{*1}, Connor Baumler^{*1}, Tin Nguyen¹, Hal Daumé III^{1,2}

¹University of Maryland, ²Microsoft Research

Contact: navita@umd.edu, baumler@umd.edu

AI systems have been known to amplify biases in real world data. Human-AI teams have the potential to control for these biases for fairer decision-making, and there is hope that explanations can help humans understand and combat model biases. Traditionally, explanations focus on the input features that are salient to the model’s predictions. If a model is biased against some protected group, explanations may include features that demonstrate this bias. However, the relationship between a proxy feature and the protected one may be less clear to a human. In this work, we consider whether explanations are sufficient to alleviate model biases due to proxy features in human-AI decision-making teams.

We study the effect of the presence of protected and proxy features on participants' perception of model fairness and their ability to improve demographic parity over an AI-only model. Further, we examine how different interventions—model bias disclosure and proxy correlation disclosure—affect fairness perception and parity.

We consider the task of loan outcome prediction. Participants are shown relevant features such as loan amount or occupation as well as a protected feature (binary gender, which we construct to be correlated with outcome in the “historic” training data) or a proxy feature (university, including women’s colleges). Based on these features, a model prediction, and an input-influence explanation of which features contributed to the model’s prediction, participants are asked to predict whether the applicant will complete their loan on time or be late.

In line with previous work [1], we find that explanations indeed help humans identify fairness concerns when biases are direct; however, when model biases are indirect, explanations alone are not sufficient in flagging fairness concerns (Figure 1, top, Phase 1). Bias disclosure helps improve fairness in human-AI teams to some degree in the case of indirect bias; however, participants are not able to specifically catch the biased predictions in the Proxy condition with bias disclosure alone (Figure 1, top). Correlation disclosure significantly improves parity in the cases most affected by indirect bias (i.e., applicants who went to a women’s college) (Figure 1, bottom). This calls for more efforts towards designing appropriate interventions for

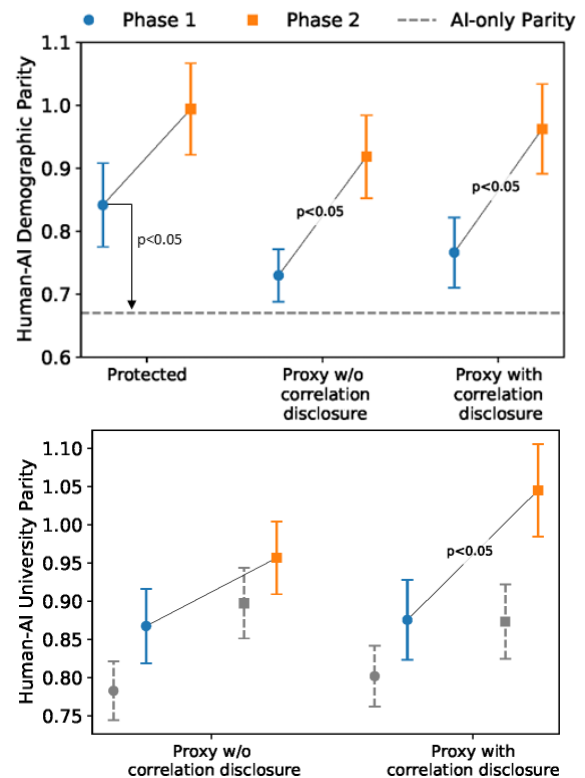


Figure 1 – Fairness measures across conditions without human teammate (AI-only), before disclosures (Phase 1), and after disclosures (Phase 2).

complex cases. Our work highlights the need to expand explanation beyond the influence of specific input features and incorporate underlying correlations that models may rely on.

REFERENCES

- [1] Jonathan Dodge et al. 2019. Explaining Models: An Empirical Study of How Explanations Impact Fairness Judgment. In *Proceedings of the 24th International Conference on Intelligent User Interfaces (IUI '19)*. Association for Computing Machinery, New York, NY, USA, 275–285.

* Denotes equal contribution

DrillVis: Adaptive Visualization Dashboards for Supporting Varying Expertise and Effort

Sungbok Shin*, Inyoup Na*, Brian W. Davis, Steven Mirsky, Niklas Elmquist
University of Maryland, USDA

Contact: sbshin90@umd.edu, elm@umd.edu

Visualization dashboards are becoming increasingly popular for communicating important facts, even complex ones, to the general public as well as supporting decision-making in virtually every industry. A recent study surveyed 83 existing dashboards from “the wild” and qualitatively categorized this corpus based on purpose, audience, visual features, and semantics [1]. However, central to this study as well as the use of visualization in general is that each and every dashboard is designed for a specific audience and purpose, making them essentially immutable in nature even if the data being visualized is dynamic and streaming. What if the user wants to do something different with the dashboard than it was intended for? What if the dashboard is intended for a different audience than the current user?

We propose DrillVis: a technique for adaptive visualization dashboards that can be drilled down to dynamically accommodate different purposes, audiences, and effort for each specific user. Instead of a static and immutable dashboard structure consisting of a grid of coordinated interactive charts, a Drillboard is essentially a hierarchy of charts. At the bottom level of this tree is the baseline dashboard, which consists of all the different charts displaying the data at its highest detail. Each consecutive hierarchy level above the baseline merges two or more charts, yielding a less detailed and more abstracted representation. The root of the Drillboard tree is thus a single chart that represents the entire dashboard and dataset in a compact and easily glanceable representation. Drilling down to any level of the tree allows a user to adapt their view of the dashboard depending on their expertise, purpose, and desired effort. A Drillboard can also provide aggregation levels for specific predefined labels, such as “novice”, “intermediate”, and “expert”.

While a Drillboard aggregation hierarchy can be generated in an ad hoc manner, imposing structure on the process will make the resulting Drillboard easier to navigate and use. To this end, we present a formal vocabulary of dashboard chart representations as well as rules for merging two or more charts into one chart (either using the same visual representation or a different one). A Drillboard can thus be created by applying these rules iteratively starting from the baseline dashboard. Of course, authoring a Drillboard is significantly more complex than authoring a normal visualization dashboard with a static structure. To support

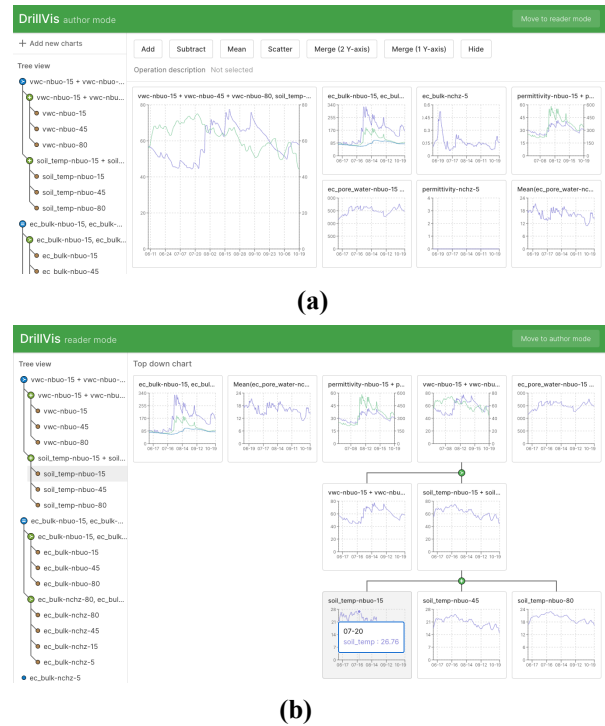


Figure 1 – DrillVis in Action. (a) *Author mode* of DrillVis. The user can formulate the Drillboard aggregation hierarchy. (b) *Reader mode* of DrillVis. The user can drill down to explore the details of the lower levels of the hierarchy.

this task, we present an authoring environment for creating web-based Drillboards for visualizing multidimensional data. The tool lets the user load a multidimensional dataset and then create a normal static visualization dashboard (the baseline). The choice for which rules to apply is left to the user. The user can also specify predefined labels in the hierarchy, allowing the end-user to easily navigate to specific views of the Drillboard.

REFERENCES

1. A. Sarikaya, M. Correll, L. Bartram, M. Tory and D. Fisher. 2019. What Do We Talk About When We Talk About Dashboards? In *IEEE Transactions on Visualization and Computer Graphics*, 25(1), pp.682-692. <https://doi.org/10.1109/TVCG.2018.2864903>

(*=Both authors contributed equally to this work)

Evaluating the Effectiveness of Event Sequence Visual Summarization Techniques

Kazi Tasnim Zinat & Zhicheng Liu

Department of Computer Science

Contact: kzintas@umd.edu

Event sequence datasets are ubiquitous in many application domains, such as application logs and electronic healthcare. These datasets are often complex and heterogeneous, making visualizations based on simple visual encoding and aggregation inadequate. Extensive research has thus focused on techniques to generate visual summaries of event sequences by showing only important events and salient structures that serve as overviews.

Despite advances in novel visual summarization techniques, to-date, there have been no empirical studies comparing these techniques through controlled experiments. The lack of systematic evaluation is problematic: researchers have no established baselines and methods to measure and innovate new techniques; practitioners have no guidance on choosing a suitable technique for their data and analytic needs.

To address this gap, we present the design and results of an insight-based crowdsourcing experiment evaluating three existing visual summarization techniques: CoreFlow, SentenTree, and Sequence Synopsis. In the experiment, the participants evaluate how closely visual summaries generated by the techniques match known ground truths. We compare the visual summaries generated by these techniques across three tasks, on six datasets, at six levels of granularity, and collected Likert Scale Ratings, Completion time of rating each visualization, and justification as text response from 180 participants.

Results from Analysis

Our analysis shows that Sequence Synopsis produces the highest-quality visual summaries for all three tasks, but understanding Sequence Synopsis results also takes the longest time. We also find that the participants evaluate visual summary quality based on two aspects: content and interpretability. Our experiment analyzes the effectiveness of existing visual summarization techniques and can be developed in future to guide researchers and practitioners in selecting suitable techniques for their data and analysis needs.




Fig 1 - Likert scale ratings distribution across Technique

PAPERS

- Zinat, Kazi Tasnim, Yang, Jinhua, Gandhi, Arjun, et al. "A Comparative Evaluation of Visual Summarization Techniques for Event Sequences". Computer Graphics Forum (Proceedings of EuroVis 2023).

During the UMD offense, Maryland made a shot after missing a shot, this pattern happens 7 times

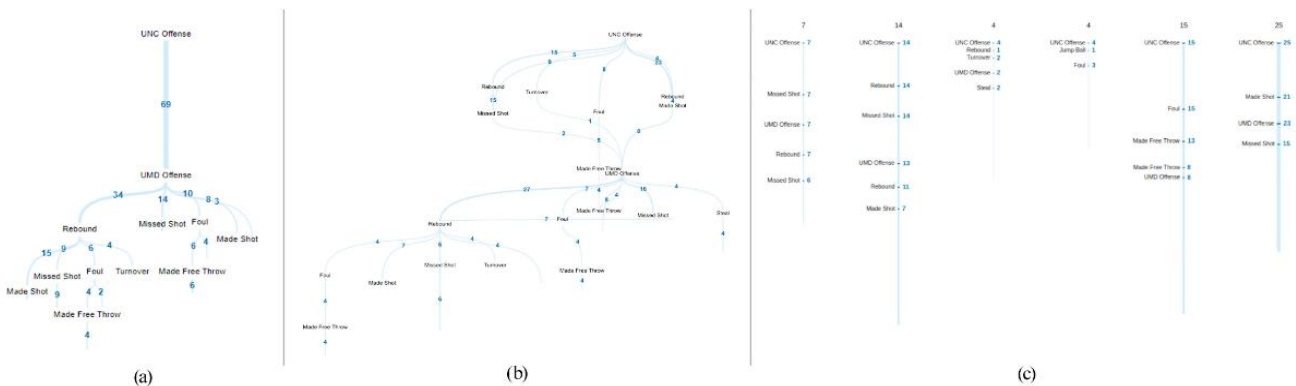


Fig 2 - The Three Techniques we implemented: CoreFlow, SentenTree and Sequence Synopsis

Intents, Techniques, and Components: a Unified Analysis of Interaction Authoring Tasks in Data Visualization

Hyemi Song, Sai Gopinath, Zhicheng Liu

Department of Computer Science

Contact: hsong02@umd.edu

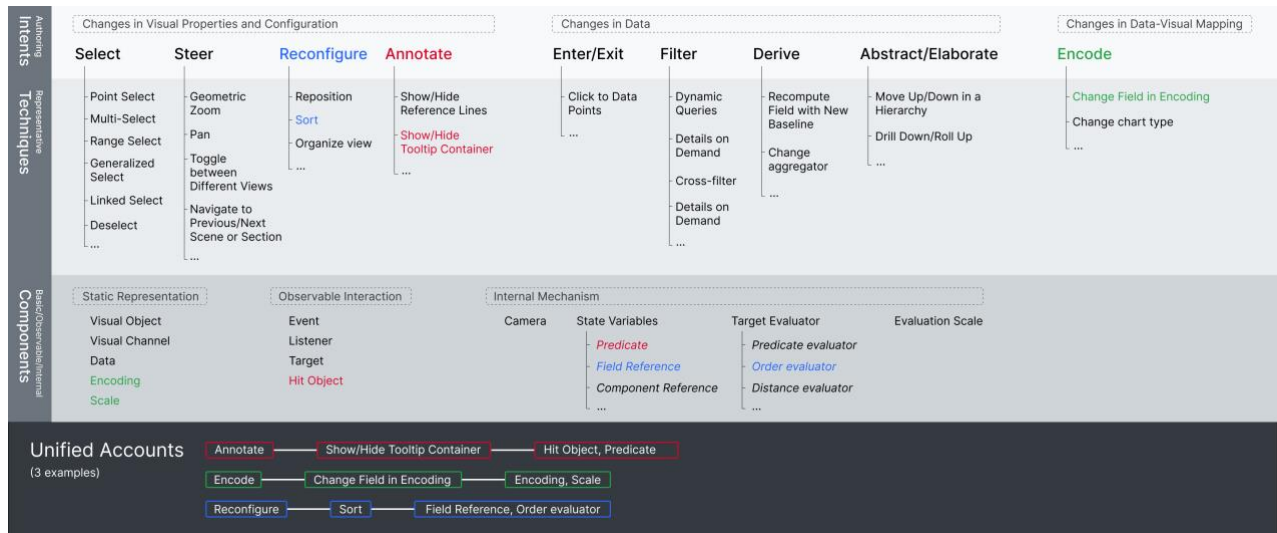


Fig. 1: A summary of the categories in our unified analysis

There is a growing interest in designing tools to support interactivity specification and authoring in data visualization. To understand the expressiveness of extant interaction authoring systems and to inform the design of new languages and tools, it is important to have theories and models that describe the task space of interaction authoring in data visualization. While previous research has provided taxonomies and frameworks to describe tasks in interactive visualization, they primarily focus on how visualizations are used for analytical or communication purposes, rather than how interactivity is composed. In this paper, we address this gap by providing a unified framework that describes interaction authoring tasks at multiple levels of abstraction.

To analyze multifaceted interaction authoring tasks, we first identify three levels of abstraction in interaction authoring: intents, techniques, and components. We then classify 592 interaction units from 47 real-world visualization applications into categories at the intent and component levels. With the resulting intent and component taxonomies, we identify representative techniques for each intent category and required components for each technique in a unified analysis. Through the analysis, we make the following contributions:

- We interpret and consolidate existing taxonomies on user intent in the context of interaction authoring, and define nine types of authoring intent: select, steer, reconfigure, annotate, encode, enter/exit, filter, abstract/elaborate, and derive.
- We identify a set of low-level primitives for composing interactivity, organized into three categories: representation components (e.g., visual object, channel, encoding, scale), observable interaction components (e.g., event, hit object, target), and internal components (e.g., camera, state variable, target evaluator).
- We contribute a unified framework of the task space of interaction authoring across three levels of abstraction: intents, techniques, and components. For each type of authoring intent, we identify representative interaction techniques; for each technique, we describe how the desired behaviors can be achieved through the low-level components.

We discuss the descriptive, evaluative, and generative powers of our analysis. In particular, we discuss the implications of our work on the design and evaluation of interaction authoring tools.


The State of the Art in Creating Visualization Corpora for Automated Chart Analysis

Chen Chen, Zhicheng Liu
 Department of Computer Science
 Contact: cchen24@cs.umd.edu

Recent advances in automated chart analysis and techniques seek to enable more effective retrieval, interpretation, creation, and transformation of data visualizations. Typically, these research efforts require a corpus of charts collected from the wild. Such corpora are essential for developing and evaluating chart analysis methods, and for providing real-world examples that end users can modify and repurpose.

There has been, however, little research on 1) the common practices for creating the corpora, 2) what constitutes a good chart corpus for various tasks and applications, and 3) the potential pitfalls and gaps in existing corpus-based research for automated chart analysis. Based on our preliminary observation, many relevant papers do not use corpora contributed by prior work; instead, they build their own corpora. There are many possible reasons for this: previous corpora are not publicly available, the corpora are not of high quality, the corpora do not have the labels required for specific tasks, or the existing corpora do not contain visualizations representing the desired design space. The current state of corpora creation and usage seems ad hoc making it difficult to compare chart analysis techniques, measure scientific progress, and identify unsolved research problems.

This survey aims to contribute a comprehensive understanding of the state of the art in creating corpora for automated chart analysis research. We collect 56 research papers from areas including AI, HCI, NLP, and Visualization that either contribute a new chart corpus, or a technique or system that takes charts in a corpus as inputs, or a model trained on a corpus. We first identify the automated chart analysis tasks along three dimensions: why (the goal), how (the method), and what (the outputs). We then extract five main properties of chart corpora used in these research works: chart format, corpus scope, collection method, annotations, and diversity. Along these task dimensions and corpus properties, we present results on the current patterns and practices of corpora creation and usage.



Making Sense of Waiting Time Activities

Jian Zheng, Ge Gao
 College of Information Studies
 Contact: jzheng23@umd.edu

People often find themselves in the status of waiting on a daily basis. Previous research in HCI tends to consider *waiting time* as moments that are wasted. Following such a mindset, various technical tools [e.g., 1, 2] have been invented so that people could harness their waiting time for productivity. However, little of this research has grounded the interventions upon in a comprehensive understanding of people’s waiting time activities under a natural setting.

In the current study, we investigated how knowledge workers make spontaneous use of their daily waiting time. Our sample consisted of 19 working adults located in the United States. Participants used an experience sampling app to generate self-reports about their waiting time activities. For a subset of these participants, cellphone and computer usage logging data was also collected to contextualize their self-reports. The data collection with each participant lasted for two weeks.

Our findings so far revealed that, on average, participants allocated approximately 60% of their waiting time to leisure, 20% to productive, and another 20% to maintenance activities (Figure 1). There were several situational factors shaping a person’s choice of the activity at each specific waiting moment (Figure 2).

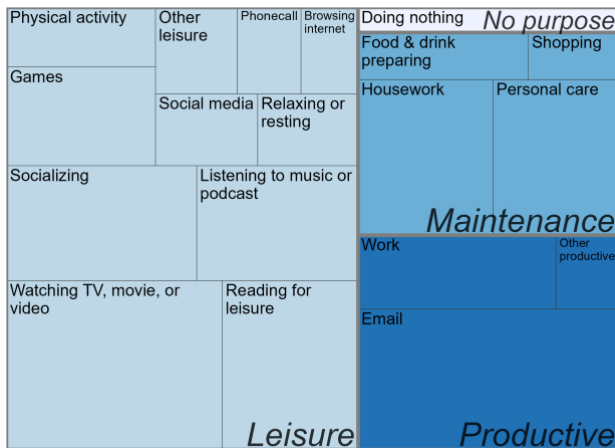


Figure 1. Waiting time activities could be categorized into leisure activities (60%), productive activities (20%), and another maintenance activities (20%). The size of each areas indicates the amount of time spent on the corresponding activity.

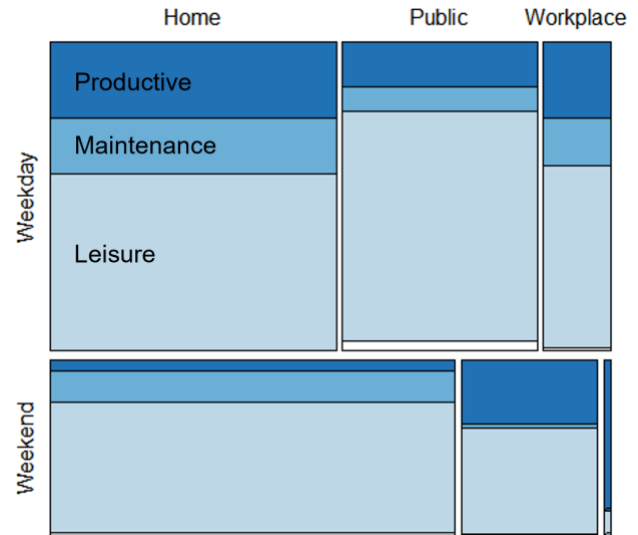


Figure 2. A person’s decision of what activities to perform while waiting depended on situational factors, such as the time (weekdays or weekend) and location (home, workplace, or other public places) of waiting.

Further, we found waiting scenarios vary along the dimension of solid vs plastic [3]. It was more difficult to anticipate the ending time of plastic waiting scenarios, which adds challenges to system design targeting the management of those moments.

Based on these findings, we envision the design of future recommendation systems that enable people to (re)purpose their waiting time activities against the situational features of each given life scenario.

Reference:

1. Cai, C. J., Ren, A., & Miller, R. C. (2017). WaitSuite: Productive Use of Diverse Waiting Moments. *ACM Trans. Comput.-Hum. Interact.*, 24(1), 1–41.
2. Kang, B., Min, C., Kim, W., Hwang, I., Park, C., Lee, S., Lee, S.-J., & Song, J. (2017). Zaturi: We Put Together the 25th Hour for You. *Create a Book for Your Baby*. CSCW, 1850–1863.
3. Rattenbury, T., Nafus, D., & Anderson, K. (2008). Plastic: a metaphor for integrated technologies. *Proceedings of the 10th International Conference on Ubiquitous Computing*, 232–241.

Sensemaking Sans Power: Interactive Data Visualization Using Color-Changing Ink

Biswaksen Patnaik¹, Huaishu Peng^{1,2}, Niklas Elmqvist^{2,1}

¹Department of Computer Science, ²College of Information Studies

Contact: bpatnaik@umd.edu

We present an approach for interactively visualizing data using color-changing inks without the need for electronic displays or computers. Color-changing inks are a family of physical inks that change their color characteristics in response to an external stimulus such as heat, UV light, water, and pressure. Visualizations created using color-changing inks can embed interactivity in printed material without external computational media. We survey current color-changing ink technology and then use these findings to derive a framework for how it can be used to construct interactive data representations. We also enumerate the interaction techniques possible using this technology. We then show some examples of how to use color-changing ink to create interactive visualizations on paper [1]. While obviously limited in scope to situations where no power or computing is present, or as a complement to digital displays, our findings can be employed for paper, data physicalization, and embedded visualizations.

In data visualization, a mark is a geometric primitive that forms the basic graphical elements, and a channel controls the appearance of a mark. In this work, we propose leveraging the potential of color-changing inks to design marks that enable dynamically changing channels in print media visualizations (Fig. 1).

Unfortunately, color-changing inks are not compatible with current consumer printers, and commercial printing using them is still prohibitively expensive (all our examples in this paper were hand-drawn). However, we hope that our work in this paper can help motivate research and development into color-changing ink printers and designing their algorithms for computer-generated prints.

In summary, our work proposes the following contributions:

(i) a novel design framework on the use of color-changing inks for data visualization; (ii) a practical approach for implementing visualizations made using color-changing inks, including solutions for actuation, interaction, and data

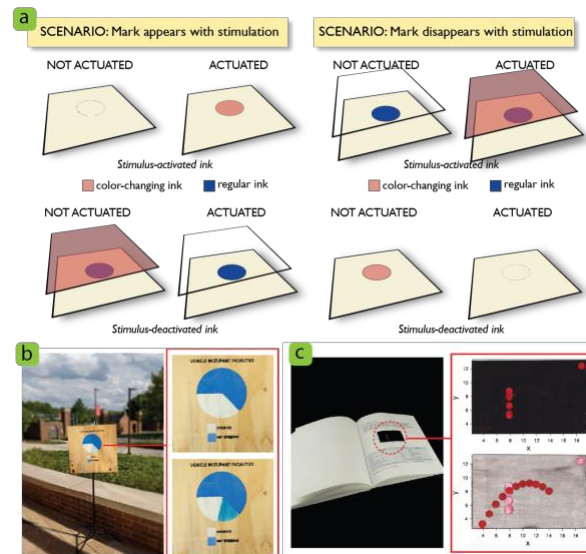


Figure 1 – (a) Controlling mark appearance. (b) Photochromic billboard chart. (*top*: night; *bottom*: day) (c) Thermochromic interactive book. (*top*: no actuation; *bottom*: hand is placed on figure)

masking; and (iii) several example prototypes of visualizing data using color-changing inks

ADDITIONAL INFORMATION (VIDEO)

<https://doi.org/10.1109/TVCG.2022.3209631/mm1>

PAPERS

1. B. Patnaik, H. Peng and N. Elmqvist, "Sensemaking Sans Power: Interactive Data Visualization Using Color-Changing Ink," in IEEE Transactions on Visualization and Computer Graphics, 2022, doi: [10.1109/TVCG.2022.3209631](https://doi.org/10.1109/TVCG.2022.3209631).



Supporting Blind Users with Visual Question Answering

Yang Trista Cao, Kyle Seelman, Kyungjun Lee, Hal Daumé III

Department of Computer Science

Contact: ycao95@umd.edu

Much research has focused on evaluating and pushing the boundary of machine “understanding”—can machines achieve high scores on tasks thought to require human-like comprehension, e.g. image captioning. Recently, with the advancement of deep learning, we saw improvements in machines’ capabilities in accomplishing these tasks, raising the possibility for deployment. However, adapting machine systems in real-life is non-trivial as real-life situations and users can be significantly different from synthetic and crowdsourced dataset examples. We use the visual question answering (VQA) task as an example to call attention to shifting from development on machine “understanding” to building machines that can make positive impacts to society.

VQA is a task that requires a model to answer natural language questions based on images. This idea dates back to at least the 1960s in the form of answering questions about pictorial inputs, and builds on “intelligence” tests like the total Turing test. Recently, the task was re-popularized with new modeling techniques and datasets [1]. However, VQA systems could also be potentially beneficial for visually impaired people in answering their questions about the visual world in real-time. Since most research in VQA focuses on the machine understanding view, it is unclear how well the VQA systems for machine understanding can be adapted for blind users. Thus, we aim to investigate 1) whether VQA model architectures developed and evaluated on machine understanding datasets can be easily adapted to the accessibility setting, and 2) what kinds of model explanations are helpful for blind users to decide whether they should trust the model’s predictions.

Gaps between VQA for Machine Understanding versus for Accessibility

We investigate the gap between the machine understanding VQA and the accessibility VQA by uncovering the challenges of adapting machine understanding VQA model architectures on an accessibility VQA dataset. Here, we focus on English VQA systems and datasets; for machine understanding VQA, we use the VQA-v2 dataset [1], while for accessibility VQA, we use the VizWiz dataset [2]. Through performance assessments of seven machine understanding VQA model architectures that span 2017-2021, we find that model architecture advancements on machine understanding VQA also improve the performance on the accessibility task, but that the gap of the model performance between the two is still significant and is

increasing. This increasing gap in accuracy indicates that adapting model architectures that were developed for machine understanding to assist visually impaired people is challenging, and that model development in this area may indicate architectural overfitting.

We further investigate what types of questions in the accessibility dataset remain hard for the state-of-the-art (SOTA) VQA model architecture. We find some particularly challenging classes within the accessibility dataset for the VQA models as a direction for future work to improve on. Additionally, we observe that many of the questions on which state-of-the-art models perform poorly are not due to the model not learning, but rather due to a need for higher quality annotations and evaluation metrics.

Model Explanations for Blind Users

Beyond model performance, model explanations are shown to be essential in enhancing model usability. For most visual tasks, the standard approach of model explanation is to use color-coded saliency maps over the image to highlight the image pixels which affect the model’s prediction. However, blind users cannot access this visual information as highlighted image pixels. Furthermore, it is unclear what kinds of information is helpful in assisting blind users to make decisions with the model’s prediction. We propose using machine-generated clarifying questions to improve the VQA system with blind users. On one side, through the context of the clarifying question asked, blind users will have a better understanding if the machine comprehended their original question. On the other hand, the clarification information from users may improve the model’s prediction accuracy. Overall, we aim to inspect whether these clarifying questions are helpful in assisting blind users in verifying correct model predictions and catching incorrect ones, as well as improving model accuracy.

ACKNOWLEDGEMENTS

Partial funding for this research was provided by the NSF (Grant No.2131508). We are grateful to Hernisa Kacorri and Jordan Boyd-Graber who provided pointers and suggestions.

REFERENCES

1. Agrawal et al. 2017. Vqa: Visual question answering. *Int. J. Comput. Vision*, 123(1):4–31.
2. Gurari et al. 2018. Vizwiz grand challenge: Answering visual questions from blind people. *2018 IEEE/CVF Conference*, pp. 3608–3617.

TangibleGrid: Tangible Web Layout Design for Blind Users

Jiasheng Li, Zeyu Yan, Ebrima Jarjue, Ashrith Shetty, Huaishu Peng
Department of Computer Science
Contact: jqli@umd.edu

Assistive technologies have greatly changed the lives of blind and visually impaired people. Beyond Internet consumers, blind users are now able to share stories and life events on social media sites such as YouTube and Instagram; some blind users have also created and maintained their own web pages for blogging and knowledge sharing. Indeed, the stories and daily experiences of the blind media influencers have become an important source of support to the blind community. Mastering skills like building web pages has also led to new employment opportunities for blind and visually impaired people.

Unfortunately, creating a web page is still challenging for many blind users despite the strong need for it. For one, web page design often requires blind developers to write code in HTML and CSS, which has a series of accessibility challenges. Responding to this issue, researchers have proposed workshops and online courses that help blind users learn web programming using screen readers, as well as developed assistive programming tools to help blind users in understanding the semantic meaning of code structures.

While these efforts support blind users in writing a program or coding web page content, a second barrier is preventing many blind users from having their own web page. Few accessible tools can help blind users understand and design the graphical layout of a web page, where visual semantics such as the size, shape, and location of the content matter.

TangibleGrid

In responding to this problem, we present TangibleGrid, a working prototype that allows blind users to understand and design the layout of a web page with real-time tangible feedback. With TangibleGrid, a blind user can place multiple visual elements, such as a textbox, a figure, or a video on a web page canvas by directly snapping the corresponding tangible brackets onto a custom baseboard. Each type of bracket has a unique tactile pattern on its top that blind users can understand. The bracket can also be resized while remaining as a rectangle so that a blind user can alter the web page layout by directly resizing or relocating these brackets. Changes are registered to the baseboard immediately so that the brackets' location, dimension, and type can be read to the user in real-time. An HTML web page will also be rendered automatically to the user.

The system design was initiated with insights gained through a semi-structured interview with blind users. It went through




Figure 1 – TangibleGrid overview: a) The complete system of Tangible grid; b) a participant is exploring a web page layout; c) designing a new layout by resizing and placing a bracket to the baseboard.

3 rounds of co-design sessions with a blind developer. The final prototype was evaluated in-person with ten blind participants. All blind participants were able to create a web page layout using TangibleGrid, despite having no previous experience in web page design and editing.

ACCOMPANYING VIDEO

<https://www.youtube.com/watch?v=pHvY-N3Cc8I>

PAPERS

1. Jiasheng Li, Zeyu Yan, Ebrima Haddy Jarjue, Ashrith Shetty, and Huaishu Peng. 2022. TangibleGrid: Tangible Web Layout Design for Blind Users. In Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22). Association for Computing Machinery, New York, NY, USA, Article 47, 1–12. <https://doi.org/10.1145/3526113.3545>

Toucha11y: Making Inaccessible Public Touchscreens Accessible

Jiasheng Li, Zeyu Yan, Arush Shah, Jonathan Lazar, Huaishu Peng

Department of Computer Science

Contact: jsli@umd.edu

From self-service food in restaurants to ID renewal in the Department of Motor Vehicles (DMV), touchscreen kiosks have been increasingly popular in workplaces and public access areas. The emergence of the COVID-19 pandemic further accelerates the adoption of touchscreen devices, as they provide customers with an autonomous experience while promoting social distancing. Unfortunately, the widespread use of touchscreen technology on kiosks can be challenging for blind people. Being primarily visual and nonhaptic, many touchscreen devices—especially those in public—are not accessible to them, preventing blind users from performing tasks independently and potentially introducing feelings of embarrassment.

Much effort has recently been made to improve the accessibility of public touchscreen devices. For example, several recent lawsuits in the United States have compelled large corporations to make their public touchscreen devices accessible. The Americans with Disabilities Act (ADA) has also required that public kiosks (such as ATMs and movie rental kiosks) be accessible to blind users. Indeed, despite the slow pace, we are seeing an increasing number of public touchscreen devices outfitted with accessibility features. Unfortunately, because the guidelines for designing accessible kiosks are not clearly defined, the majority of public touchscreens are still not useful to blind users. Even devices with touchscreens and physical keypads, for example, may not have any audio output; others with audio output may simply speak the information publicly, ignoring the privacy of blind users. Moreover, a significant number of inaccessible touchscreens have already been deployed worldwide. Waiting for device manufacturers to update all of them will not solve today's accessibility problems.

TangibleGrid

We present Toucha11y, a working prototype that aims to enable blind users to access arbitrary public touchscreen devices independently and with little effort. The key to Toucha11y is to bridge the gap between blind users and a public touchscreen device with a set of hardware and software tools, allowing them to explore touchscreen content from their familiar smartphone devices without having to deal with the unfamiliar, inaccessible public touchscreen directly. To make the bridge work, Toucha11y's hardware—




Figure 1. Toucha11y robot

a small mechanical bot—must be placed on top of a public touchscreen device by a blind user (Figure 1). Once placed, the bot's onboard camera will photograph the screen, with its corresponding interface (which can be generated through crowdsourcing and reverse engineering) sent to the user's smartphone. The blind user can freely explore and select content using the smartphone's built-in accessibility features, such as Apple's VoiceOver and Android's TalkBack. These selections will be sent back to the bot, which will physically register the corresponding touch event for the blind user using an extendable reel.

ACCOMPANYING VIDEO

<https://www.youtube.com/watch?v=dqfhE42zB1I>

PAPERS

1. Jiasheng Li, Zeyu Yan, Arush Shah, Jonathan Lazar, and Huaishu Peng. 2023. Toucha11y: Making Inaccessible Public Touchscreens Accessible. In Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 748, 1–13. <https://doi.org/10.1145/3544548.3581254>

Calico: Relocatable On-cloth Wearables with Fast, Reliable, and Precise Locomotion

Anup Sathya, Jiasheng Li, Tauhidur Rahman, Ge Gao, Huaishu Peng

Department of Computer Science

Contact: jsli@umd.edu

Wearable devices such as fitness trackers and smartwatches have been widely adopted in our daily lives for both health monitoring and digital interactivity. We are also witnessing the trend of wearable device minimization, where both the size of these devices and their distance from our bodies are reducing, suggesting a future where miniature wearable devices may seamlessly integrate with us for interaction, actuation, and sensing. However, most wearables to date are resigned to a single location on our body. In reality, the human body offers a multitude of interaction modalities at various places. For example, the collar of one's clothing can be good for voice commands; the forearm offers ample area for spatial feedback and can thus be used for hand input. The areas suitable for sensing also vary. For example, it is best to monitor breathing from the front or back of one's upper body, while the wrist can be ideal for sensing activities such as typing and writing. Thus, if a user wants to interact with physiological information or needs haptic or visual feedback from different areas, they are destined to wear multiple wearable devices in various form factors, on numerous areas of the body.

Calico

We present Calico, a relocatable wearable prototype with a reliable on-cloth track system to address the aforementioned challenges. With Calico, robots can traverse long distances across different areas of the body, such as from one's forearm to the back, or from the thigh to the chest, regardless of the material deformation or the seams between different pieces of cloth. The locomotion is up to 227 mm/s, and an accuracy of up to 4 mm. In doing so, Calico can potentially accommodate many common activities by reliably traversing to various locations on the body. For example, Calico can move along the user's arm to provide feedback about their running progress, then move back and forth on the user's back to provide feedback on their posture when sitting down.

The key to Calico is an on-cloth track system — inspired by the railways — on which the wearable moves. In a railway system, rail tracks are used as dedicated pathways to navigate unfriendly terrain; railway turntables are used as central hubs to allow trains to switch tracks. In a similar vein,




Figure 1. Calico on the track

Calico stitches custom soft tracks directly onto clothing, which then function as expressways to overcome the material's natural deformations. Calico also includes rotational switches, like railway turntables, to allow on-cloth track switching and to ensure that all the critical areas of the human body are reachable. As adding tracks to clothing affects its appearance, weight, and performance, throughout the iterative design process, we carefully consider the system's comfort and ensure that the aesthetics of the system can be modified according to the user's needs and preferences.

ACCOMPANYING VIDEO

<https://youtu.be/RIMcj5uil6Q>

PAPERS

1. Anup Sathya, Jiasheng Li, Tauhidur Rahman, Ge Gao, and Huaishu Peng. 2022. Calico: Relocatable On-cloth Wearables with Fast, Reliable, and Precise Locomotion. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 6, 3, Article 136 (Sept 2022), 32 pages. <https://doi.org/10.1145/3550323>

Creating and Evaluating FormA11y, a tool for remediating PDF forms for accessibility

Sparsh Paliwal¹, Joshua Hoeflich², J Bern Jordan¹, Rajiv Jain³, Vlad I Morariu³, Alexa Siu³, Jonathan Lazar¹

¹iSchool, Trace R and D Center, HCIL, University of Maryland

²Northwestern University, ³Adobe Research,

Contact: sparsh@umd.edu, joshuahoeflich2021@u.northwestern.edu, bjordan@umd.edu, rajivjain@adobe.com, morariu@adobe.com, asiu@adobe.com, jlazar@umd.edu

The Portable Document Format (PDF) enables documents to maintain consistent layout and styling across platforms and thus they are widely used in education, healthcare, government, and financial settings. PDF documents are often not born-accessible, and document creators need to remediate for accessibility after document creation. Unfortunately, document authors often skip this part and as a result, a large number of PDF documents remain inaccessible for users with disabilities.

As an increasing number of research articles document the need for PDF accessibility [1] and there has been an increase in the work done to improve the tooling for PDF remediation [2]. The existing research has mostly focused on the accessibility of scientific PDF documents meant for reading. This research and development project focuses on PDF forms, which are different from other PDF documents because the end-user is expected to interact and enter data.

Existing research and tools do not yet address the following challenges in remediating PDF forms. 1) *the process of remediation is non-intuitive*: remediating a single form field can involve multiple steps which are difficult without prior knowledge, 2) *the process is repetitive*: a form can have hundreds of fields, and thus a person would need to perform the same operations repetitively to make form accessible, 3) *forms have high information density*: many forms are densely packed with text and graphical information, so their remediation is cognitively demanding. Our research addresses these issues through the development of a form remediation tool – FormA11y.

FormA11y

FormA11y leverages a linearized, step-by-step process to simplify PDF Form remediation – fields, groups and tooltips steps. Each step of the tool is designed to be intuitive and user-friendly, rooted in the principles of user experience. The tool assists users in performing basic interactions such as drawing fields and grouping them together to make a PDF form accessible. The division into steps also helps with tackling high information density issues, for example, the user is only presented with

checkboxes and radioboxes in the group step which are groupable together.

We interviewed PDF form remediation experts to understand their needs and used that feedback to iterate over the tool’s design while incorporating engineering and pilot study feedback. We performed a within-subjects study involving 20 participants, evaluating the effectiveness and efficiency of FormA11y compared to the industry standard Adobe Acrobat Pro tool for remediating PDF forms for accessibility. The 20 participants were able to remediate forms for accessibility more than two times faster with fewer tagging errors and greater user satisfaction with FormA11y, compared to Acrobat Pro. Figure 1 shows a screenshot of the FormA11y prototype we tested.

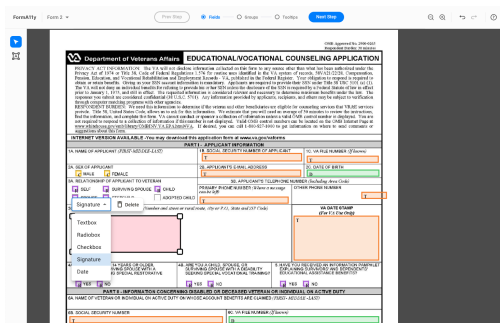


Figure 1 – FormA11y tool used to add fields to a form.

ACKNOWLEDGMENT

Partial funding for this research was provided by Adobe Research and by the National Institute on Disability, Independent Living and Rehabilitation Research (NIDILRR) at the Administration for Community Living, HHS, grant #90REGE0008. The opinions herein are those of the authors and not necessarily those of the funding agencies.

PAPERS

1. Bigham, J. P., Brady, E. L., Gleason, C., Guo, A., & Shamma, D. A. (2016, May). An uninteresting tour through why our research papers aren't accessible. In Proceedings of the 2016 CHI conference extended abstracts on human factors in computing systems (pp. 621-631).
2. Pradhan, Debashish, et al. "Development and Evaluation of a Tool for Assisting Content Creators in Making PDF Files More Accessible." ACM Transactions on Accessible Computing (TACCESS) 15.1 (2022): 1-52.

Contributing to Accessibility Datasets: A Perspective by Blind People

Rie Kamikubo¹, Kyungjun Lee², Hernisa Kacorri¹

College of Information Studies¹, Department of Computer Science²

Contact: rkamikub@umd.edu

To ensure that AI-infused systems work for diverse users including disabled people, we need to bring accessibility datasets sourced from this community in the development lifecycle. However, many ethical and privacy concerns limit greater data inclusion, making such datasets not readily available and the decision to share really tricky.

In this work, we bring potential data contributors from the disability community to the forefront of discussions around increasing the availability of accessibility datasets. Specifically, we focus on the blind community and data that include photos, a challenging use case since one may not be able to visually inspect the photos before deciding to share. In a pair of studies, 13 blind participants engaged in data capturing activities (Figure 1); they were testing in their homes our teachable object recognition app deployed on smartglasses, where the camera could capture surroundings. Then, in a follow-up interview, they reflected with and without probing scenarios on what might influence their decision to share their study data via an AI dataset.

Our findings revealed various motivating and challenging factors that could play into blind participants' decisions to share their study data (i.e., photos and labels of objects). Many perceived potential risks such as re-identification as minimal and supported sharing practices to improve AI-infused technology for greater benefit to both disabled and non-disabled people. Yet, they were hesitant to contribute their data for commercial purposes and companies handling the use of their data, even though almost all frequently used AI-infused applications (e.g. SeeingAI) to read text and identify objects or shared their camera view with sighted helpers (e.g. via Aira). They also expressed concerns for sharing demographic metadata (e.g., age, gender, race) along their study data, due to privacy and safety and the ambiguity of its use for AI.

We see the main contribution of this work being empirical. By investigating data sharing from the perspectives and experiences of blind people, we contribute to the larger call-to-action for the research community, industry, and policy makers in shaping future data practices that are inclusive of disability. We also see how our approach of eliciting participants' perspectives before asking them to decide on whether they want their study data to be shared



(a)

(b)

Figure 1: (a) A blind participant using an object recognition app installed on smartglasses in their home and taking multiple photos of a soda can. (b) Examples of photos of objects taken and labeled by our participants: “Cereal” (top) and “Panera Cup” (bottom).

and how, could be leveraged by future researchers who want to engage participants in decisions around sharing of their study data. We thus share our questionnaire with scenarios and interview questions to facilitate replicability.

PUBLICLY ACCESSIBLE QUESTIONNAIRE

https://go.umd.edu/datasharing_questionnaire

ACKNOWLEDGEMENTS

This work was funded by National Institute on Disability, Independent Living, and Rehabilitation Research (NIDILRR), ACL, HHS (#90REGE0008).

PUBLICATION

Rie Kamikubo, Kyungjun Lee, and Hernisa Kacorri. 2023. Contributing to Accessibility Datasets: Reflections on Sharing Study Data by Blind People. In Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 827, 1–18. <https://doi.org/10.1145/3544548.3581337>

Multimodal Activity Tracking to Support Stroke Rehabilitation

Jong Ho Lee*, Sunghoon Ivan Lee⁺, Eun Kyoung Choe*

*College of Information Studies, University of Maryland College Park

⁺College of Information and Computer Sciences, University of Massachusetts Amherst

Contact: jlee29@umd.edu

Stroke is a leading cause of disability in the United States, with over 795,000 individuals experiencing a stroke annually. A stroke occurs when blood supply is blocked to a part of the brain, which can result in long-term disability such as motor impairments and aphasia. Various rehabilitation techniques, such as stroke-specific physical and occupational therapy, can help stroke survivors regain independence in their daily living. Establishing a goal-setting process that generates personally meaningful goals is critical for maintaining motivation and monitoring progress. However, previous work suggests that engaging stroke survivors in the goal-setting process can be challenging [1], which can lead to diminished motivation during therapy.

To address this issue, we prototyped an Android app (Fig. 1) that enables stroke survivors to journal daily activities using a combination of speech and touch input. Using this prototype, we investigated (1) how journaling daily activities can help stroke survivors develop actionable goals, and (2) how multimodal input can help improve the accessibility of activity tracking apps.

App Prototype

Figure 1. shows two screenshots of the app prototype, *GoalTrack*. *GoalTrack* allows each activity to be logged with four different types of information: *activity details*, *activity difficulties*, *labels*, and *time-duration*. Users have the option to use both voice and touch input to log activity details and difficulties (Fig. 1, left) and time-duration (Fig. 2, right).

Methods and Future Work

We designed an in-person usability study paired with semi-structured interviews to understand stroke survivors' perspectives in using multimodal input to track activities. We employ a mixed-methods approach in analyzing our data, which involves conducting quantitative analysis on generated interaction data on different measures, and reflexive thematic analysis on video and interview transcripts. To date, we have conducted the study with eight stroke survivors, and we plan to continue data collection

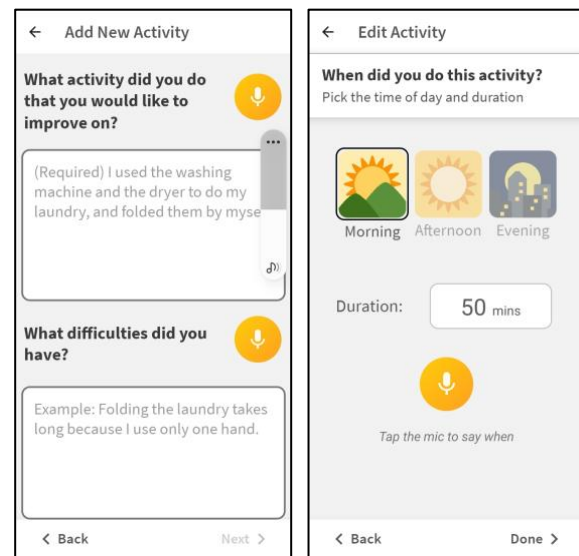


Figure 1 – *GoalTrack*, the prototype app used for the study. In addition to the default keyboard, participants can touch the yellow microphone button to use the speech interface.

with additional participants. Preliminary findings suggest that participants found voice interfaces to be more accessible than touch input, where they often found that it was more efficient than using the default keyboard. Participants also noted that activity journaling helped prompt self-reflection when setting stroke therapy goals.

REFERENCES

1. Saito, Y., Tomori, K., Nagayama, H., Sawadai, T., & Kikuchi, E.. Differences in the occupational therapy goals of clients and therapists affect the outcomes of patients in subacute rehabilitation wards: A case-control study. *Journal of Physical Therapy Science*, 31(7), 521–525. (2019)

Show and Tell: Exploring How Audio Narratives Can Complement Visualizations of Stroke Survivors' Personal Health Data

Aishwarya Shettigar¹, Sunghoon Ivan Lee², Eun Kyoung Choe¹

¹College of Information Studies, University of Maryland

²College of Information & Computer Sciences, University of Massachusetts

Contact: ashettig@umd.edu

Wearable technology in healthcare could give individuals awareness and independence in rehabilitation. In this qualitative work, we investigate how using speech-based, audio narrative contents alongside graphical visualizations affect users' understanding of their personal data. We conducted this work in the context of stroke recovery, where stroke survivors experiencing upper-limb hemiparesis (one-sided loss of function in the shoulders, arms, and hands) can monitor their physical progress using a wearable ring sensor. Using a co-design approach, we engaged with stroke survivors and caregivers to elicit recommendations for the design of multimodal (speech and visual) feedback from the wearable ring data.

Reflexive thematic analysis of co-design sessions (n=10) showed that multimodal feedback can potentially lend therapeutic support and amplify understanding of the personal health data. Audio narratives helped to reinforce the visual feedback. Narrative content framed positively, offering context beyond the visualization's capabilities, could support stroke survivors on their independent recovery journeys. Such content, which encourages self-reflection, motivates progress in rehabilitation, and provides improvement suggestions, enhances the overall experience for stroke survivors. Insights from stroke survivors regarding the complex stages of their recovery revealed potential opportunities for effective multimodal feedback. For instance, during the early chronic stage (post 6 months), audio narratives can complement visualizations without negatively impacting the recovery process because stroke survivors would have already experienced some therapy-related successes and could buy into new therapy approaches. These stroke survivors would also not be "desensitized" to the content, unlike people 10+ years into recovery who may have heard the messaging many times before from clinicians, caregivers, or themselves. Additionally, intrinsic drawbacks of each modality could be compensated for by the other to give a more enhanced experience.

Co-designing with Stroke Survivors and Caregivers

As a human-centered design approach, co-design was able help us facilitate discussion and generate rich ideas in collaboration with stroke survivors (see Figure 1). Including caregivers in the sessions sometimes led to stroke survivors disengaging from the activity, which required us to intervene. However, caregivers helped to ease conversation and provided a valuable perspective on recovery.




Figure 1 – Co-design activity using Microsoft PowerPoint where stroke survivors could interact with embedded visual and audio probes to design their own narratives.

ACKNOWLEDGEMENTS

Funding for this research was provided by the NIH (NIBIB 5R01EB027777 "Achieving Optimal Motor Function in Stroke Survivors via a Human-Centered Approach to Design an mHealth Platform").

REFERENCES

1. Kim, H.-T. Jung, J. Park, Y. Kim, N. Ramasarma, P. Bonato, E. K. Choe, and S. I. Lee. (2019). Towards the design of a ring sensor-based mhealth system to achieve optimal motor function in stroke survivors. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 3(4):1–26.
2. H.-T. Jung, Y. Kim, J. Lee, S. I. Lee, and E. K. Choe. (2022). Envisioning the use of in-situ arm movement data in stroke rehabilitation: Stroke survivors' and occupational therapists' perspectives. *PLoS ONE*, 17(10):e0274142.

How University of Maryland Students are Revolutionizing Voice Technology for Accessibility

Dr. Galina Reitz, Scott Mobarry

College of Information Studies

Contact: ben@cs.umd.edu

University of Maryland students, in a first of its kind experience in the iSchool, are being trained to be the next generation of voice technology developers in INST 208B: Voice Technology Today, taught by Dr. Galina Reitz. Students in this course have been innovating in the rapidly growing voice technology space for accessibility with their semester-long projects.

During the entirety of the semester, students worked in groups to design a conversational virtual assistant using voice interactions that proposed an innovative solution to a problem for individuals who are part of the accessibility community. Students conducted interviews with end users to identify a need with a specific group within the accessibility community that would serve them through an Alexa skill.

The project groups engaged in extensive research of Alexa skills that have already been developed and what gaps existed that needed to be addressed in the accessibility space and voice technology. Once a problem was identified, the teams utilized design thinking concepts to propose a viable solution by building an Alexa skill with the Voiceflow platform.

Throughout the user testing process of their prototypes, project groups learned state-of-the-art techniques for the iterative process of debugging and quality assurance for voice interfaces. The Alexa skill prototypes were created using the Voiceflow platform and show the interaction between the user and the Amazon Echo device.

Project teams developed sophisticated voice user interfaces ranging from supporting children better express themselves who have speech impediments, helping older adults keep track of their medications, users who are hard of hearing interact with voice technology, and improve bilingual user experiences with Alexa and improve Alexa's understanding of different accents.

With this presentation, we will showcase some of the most innovative projects and highlight the talented students who developed the skills.



Figure 1 - Source <https://www.youtube.com/watch?v=4gpFPOfJxI>

A Preventive mHealth App to Support Survivors of Domestic Violence

Julio Poveda*, Mengying Li†, Jing Wang†, Sanchita Chugh†, and Eun Kyoung Choe†

*Department of Computer Science and †College of Information Studies

Contact: jpoveda@umd.edu

Domestic violence (DV) is a widespread societal issue that affects millions worldwide. Many initiatives to tackle DV are reactive (i.e., they help survivors after they have suffered from abuse). We aim to assist DV survivors with a proactive approach, enabling to detect DV at an early stage, in order to minimize the long-term physical and physiological damage the survivors may encounter.

We conducted a literature review on initiatives aimed at empowering survivors and talked with various DV researchers and caseworkers. Based on the insights we obtained, we designed SANA (the term *Heal* in Spanish), an mHealth app that assists people to self-reflect on the potential presence of an abusive relationship.

At first sight, SANA appears to be a fitness app designed to monitor physiological factors such as heart rate. However, it has a hidden mode that allows users to access prompts and visualizations for self-reflection on their feelings and relationship.

User safety is our top priority; thus, we have incorporated best practices when designing technology for survivors, as

well as trauma-informed principles into our app design. To further refine SANA, we plan to interview DV caseworkers and survivors.

PAPERS

1. Glass, N., Clough, A., Case, J., Hanson, G., Barnes-Hoyt, J., Waterbury, A., ... & Perrin, N. (2015). A safety app to respond to dating violence for college women and their friends: the MyPlan study randomized controlled trial protocol. *BMC public health*, 15(1), 1-13.
2. Chen, J. X., McDonald, A., Zou, Y., Tseng, E., Roundy, K. A., Tamersoy, A., ... & Dell, N. (2022, April). Trauma-informed computing: Towards safer technology experiences for all. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (pp. 1-20).
3. Uusitalo, L. (2018). Designing for women experiencing intimate partner violence.



Figure 1 - Views of the SANA mHealth app.

Decorative, Evocative, and Uncanny: Reactions on Ambient-to-Disruptive Health Notifications via Plant-Mimicking Shape-Changing Interfaces

Jarrett G.W. Lee¹, Bongshin Lee², Eun Kyoung Choe¹

¹College of Information Studies, University of Maryland; ²Microsoft Research

Contact: jgwl@umd.edu



Figure 1: One of our artifacts used in the study, with the ability to move its stems up and down at a variety of speeds. Various scripted movements using this mechanism were presented to interviewees to capture their impressions on ambient-to-disruptive notifications using Shape-Changing Interfaces. Please see our supplemental video for animated examples of all movement types used in our study.

Ambient Information Systems (AIS) [1] have shown some success when used as a notification towards users' health-related activities. But in the actual busy lives of users, such notifications might be forgotten or even missed, nullifying the original notification. Could a system use multiple levels of noticeability to ensure its message is received, and how could this concept be effectively portrayed?

To examine these questions, we took a Research through Design approach and created plant-mimicking Shape-Changing Interface (S-CI) [2] artifacts, then conducted interviews with 10 participants who currently used a reminder system for health-related activities. We report findings on acceptable scenarios to disrupting people for health-related activities, and participants' reactions to our design choices, including how using naturalistic aesthetics led to interpretations of the uncanny and morose, and which ways system physicality affected imagined uses. We offer design suggestions in health-related notification systems and physical systems such as S-CIs, and discuss future work in ambruptive (ambient-to-disruptive) technology.

ACCOMPANYING VIDEO

<https://go.umd.edu/3KIO423>

ACKNOWLEDGEMENTS

This research was supported by the National Science Foundation under Award Number #1753452.

PAPERS

Jarrett G.W. Lee, Bongshin Lee, and Eun Kyoung Choe. 2023. Decorative, Evocative, and Uncanny: Reactions on Ambient-to-Disruptive Health Notifications via Plant-Mimicking Shape-Changing Interfaces. In Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23), April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3544548.3581486>

REFERENCES

1. Zachary Pousman and John Stasko. 2006. A taxonomy of ambient information systems: four patterns of design. In Proceedings of the Working Conference on Advanced Visual Interfaces. 67–74.
2. Marcelo Coelho and Jamie Zigelbaum. 2011. Shape-changing interfaces. *Personal and Ubiquitous Computing* 15, 2 (2011), 161–173.



Redefining Behavioral Interventions: A Domain-Centered Value Engineering Approach

Priyanka Mehta, Nhi Tran, Pablo Paredes

Department of Computer Science

Contact: pmeh1@umd.edu, nhitran@terpmail.umd.edu

Recent advances in health technologies, such as mobile health applications and fitness wearables, have made significant strides in improving human well-being. The data collected from these technologies is analyzed to provide personalized suggestions on what habits users can adopt to improve overall well-being [2], also known as interventions. By providing personalized interventions, mobile health applications and technologies are able to empower individuals to take control of their health and make positive changes to their lifestyles.

However, little research has been done to understand the impact of a user's intended priorities and values on their behavior, which can play a significant role in whether or not recommended interventions are successful. The approach in this paper aims to fill this gap by using a mobile application to reconcile this information, which can then be leveraged to give users interventions targeted to achieve the specific goals communicated to researchers.

Following this approach will also help users develop a greater self-awareness of which habits add to or detract from their goals. A priority-guided framework has the advantage of being both intrinsically and externally motivating, making it a valuable tool for promoting healthy behavior change. Ultimately, this approach has the potential to help individuals lead more fulfilling lives by aligning their behavior with their core values. Thus, we arrive at our first research question: **How can an understanding of a user's intended priorities and values assist researchers in developing wellness-oriented interventions?**

This leads us to the second contribution of this work: the application of value domains [3], which are not only integral to understanding individual decision-making but also impact group dynamics, social norms, and institutions. Primarily utilized in a psychological setting, the value domain framework differs from a traditional understanding of human values and behavior by providing interpretable categorizations of human actions into domains. For example, Shalom H. Schwartz's Value Theory posits ten key values in the form of abstract concepts such as "Hedonism", "Security", and "Conformity".

While these values are inextricably linked to human behavior, they can be difficult to conceptualize and apply to

mobile applications. Our value domain approach provides an alternative that is easily understood and reconciled with human actions by users and researchers alike. These domains are intended to be exhaustive and directly associated with the different ways in which people can spend their time. Below is a figure detailing the 9 possible domains, intending to answer our second research question: **Does the application of a domain-centered approach improve our understanding of human priorities?**

Values Domains:

1. Family Relationships
2. Friendships/Social Relationships
3. Couples/Romantic Relationships
4. Work/Career
5. Education-Schooling/Personal Growth and Development
6. Recreation/Leisure/Sport
7. Spirituality/Religion
8. Community/Citizenship
9. Health/Physical Well-Being

The next step is to test our mobile application on a wider audience, prompting users to submit their screen-time and value journaling to help understand their behavior, as well as increase their self-awareness. By taking into account an individual's values and intended priorities to produce targeted interventions, we believe that mobile health applications can be more effective in promoting healthy behavior change and overall well-being.

REFERENCES

1. Y. Vaizman, K. Ellis and G. Lanckriet, "Recognizing Detailed Human Context in the Wild from Smartphones and Smartwatches," in IEEE Pervasive Computing, vol. 16, no. 4, pp. 62-74, October-December 2017. doi: 10.1109/MPRV.2017.3971131.
2. Schwartz, S. H. (2012). An Overview of the Schwartz Theory of Basic Values. Online Readings in Psychology and Culture, <https://doi.org/10.9707/2307-0919.1116>
3. Blackedge, J.T., & Ciarrochi, Joseph. (n.d.) Personal Values Questionnaire.



Designing a Just-in-time Prevention and Intervention Method for Onychophagia

Yoon Kyung Shon, Pablo Paredes

Department of Computer Science

Contact: ykshon@umd.edu

Onychophagia is a commonly referred to a nail-biting behavior and can occur in individuals of various ages. It happens repetitively and compulsively, and several factors play a role in the development of nail biting such as genetic components and underlying psychiatric conditions. Onychophagia is also considered as a type of body-focused repetitive behaviors (BFRBs), and many psychologists have studied it along with other BFRBs symptoms, such as hair-pulling or skin-picking. Despite extensive research, it's exact cause remains unsolved, but it is found that stress is the most significant factor, with boredom and frustration also having a significant impact. Moreover, people with Onychophagia often have difficulty regulating their emotions and may use nail-biting as a means of emotional regulation, whether consciously or unconsciously.

Several methods, including bitter-tasting nail polish, gloves, and fidget toys, have been proposed to understand and prevent Onychophagia, but habit reversal training (HRT) has been proven to be the most effective treatment. However, some researchers argue that HRT for Onychophagia is still incomplete since understanding BFRBs is insufficient and it does not consider emotion regulation. Moreover, HRT is difficult when the behavior occurs unconsciously, and it lacks the function to provide immediate feedback. To overcome the limitations of HRT, a ubiquitous approach is necessary to prevent by providing just-in-time prevention and supporting individuals in regulating their emotions, which would enable them to maintain changed behavior over the long term.

Emotion Regulation

Emotional regulation (ER) is the process of identifying and responding to emotional experiences and influencing the expression of emotions. The ER model proposes that individuals with BFRBs have difficulty managing negative emotions and as a result, they engage in body-focused behaviors to avoid or alleviate negative emotions, such as boredom, tension, and anxiety, which can trigger feelings of shame, guilt, and relief. Deficits in ER may distinguish individuals with BFRBs from those without.

Just-in-time Intervention

Our goal is to provide real-time prevention as well as just-in-time intervention to stop Onychophagia and aid in emotion regulation. This can be particularly helpful for individuals who are not aware that they are engaging in the behavior, or who have difficulty controlling the behavior once it starts, or for those who have difficulty regulating their emotions when they stop the behavior.

To detect Onychophagia, we use camera in laptop and leverage Google's Teachable Machine to distinguish the behavior. We are currently conducting a survey to gather ideas for the further design. We will use Electrical Muscle Stimulation (EMS) to provide prevention methods such as stopping the arm in midair or folding the fingers, but the method of providing them will be decided through the survey. Additionally, since there may be difficulties in emotion regulation when nail-biting is not possible, we plan to intervene to alleviate this to some extent. The intervention method will be determined in detail through the survey, but our current plan is to provide a stimulation similar to nail-biting, such as vibration at the fingertip, to reduce the magnitude of emotional changes. After implementing these two methods, we plan to compare and evaluate their effectiveness.

PAPERS

1. Roberts S, O'Connor K, Aardema F, Bélanger C. The impact of emotions on body-Focused repetitive behaviors: evidence from a non-treatment-seeking sample. *J Behav Ther Exp Psychiatry*. 2015;46:189-197. doi:10.1016/j.jbtep.2014.10.007
2. Halteh P, Scher RK, Lipner SR. Onychophagia: A nail-biting conundrum for physicians. *J Dermatolog Treat*. 2017;28(2):166-172. doi:10.1080/09546634.2016.1200711
3. Ghanizadeh A, Bazrafshan A, Firoozabadi A, Dehbozorgi G. Habit Reversal versus Object Manipulation Training for Treating Nail Biting: A Randomized Controlled Clinical Trial. *Iran J Psychiatry*. 2013;8(2):61-67.



Analytical Concepts of Technology Use in HCI: A Case of Older Adult Users

Pooja Upadhyay, Daniel M Greene, Amanda Lazar

College of Information Studies

Contact: poojau@umd.edu

Over the past 20 years, researchers have consistently reported on different generations of older adults facing negative experiences with technologies. Such experiences are typically attributed to a low technology proficiency. We argue that a major shift is needed in how we think about and study user groups like older adults, where characterizations of a “lagging user” informs much of technological R&D. We begin by reviewing how “technology use” is envisioned through changes in commercial technology development to show how the concept of intuitive out-of-the-box systems may be an illusion.

<p>Technology development for workplace environments; Beginning in the 1980s, computing applications were developed specific to workplaces, e.g. call center systems. Users were trained to become familiar with new systems to use them efficiently for work.</p>	<p>Technology development for everyday life environments: Many technologies initially developed for the workplace (e.g. search systems for research libraries) were commercialized into everyday realms (e.g. online search), and conceived of as intuitive.</p>
<p>“Expected technology use”</p> <ul style="list-style-type: none"> • Here, the <u>user</u> is a stakeholder, e.g. a support technician taking a customer call, or a manager overseeing a call center team. • The design of <u>system functions</u> uses the <u>work context</u> of the user (workplace practices, activities) as design material to develop functions (e.g. to log customer calls). • With systems <u>designed to roughly match the work context</u>, users could predict how system functions can serve work activities. • During implementation of these technologies, initial training <u>familiarizes users with the design of system’s functions, before using the system.</u> <p>Here, a context-design match gives users hints about system functions, and training gives familiarity with system design. Usability tests capture uses of the system in ways expected by the designer (or design).</p>	<p>“Situated technology use” (Nardi, 1996)</p> <ul style="list-style-type: none"> • Now, the <u>user</u> has expanded far beyond narrow work contexts, into <u>varied past and present situated contexts</u> • <u>System functions</u> are often ready to to be used <u>out-of-the-box</u> (e.g. online search), shaped by the workplace versions that came before (e.g. research libraries) • <u>Out-of-the-box</u> systems <u>designed for “everyone” may or may not match the contexts</u> of many. • Without training or instruction manuals, users have <u>no scaffolds to build familiarity</u> with the system’s design, making its functions unclear, or appear as black-boxed. <p>Here, the user must rally their own resources to discover not only the utility of systems, but also interpret the design of its functions. They apply proficiency in unpacking the blackboxed system, tapping into resources like past technical familiarity, asking social networks, social observations of use etc.</p>

Unpacking the notion of “Technological proficiency”

Popular scales, such as Computer Proficiency Questionnaires (CPQ)([Boot, Czaja et. al, 2015](#)) **interpret proficiency as being familiar** with a list of “21st century” expected technology uses. As an example, the CPQ asks if the user uses an “online calendar” function of the computer to “*enter events; check schedules; set up event alerts*”. In our data, participants would check “no” on some or all of these items, not necessarily because they are unfamiliar with the online calendar. Some don’t prefer to use it. Others may be using certain features already but not recognize the vocabulary (e.g., “bookmark websites”), leaving these survey items as an invalid measure of proficiency.

For example, P1 only used her phone’s online calendar when out to log new event information, deleting and moving the information to her paper calendar when back home. P2 despised online calendars, and instead configured a digital system involving a workflow around her email. P3 used webpages bookmarked by her daughter, but referred to it using its position on the browser rather than the vocabulary of “bookmarks” (“netflix up here”).

There is an alternative where we aren’t conceptually limited by familiarity of “expected use” as a proxy for proficiency, like the assumption that scheduling only happens through online calendars. **Users’ proficiency, or their sensemaking of the system, is visible in action as they draw on resources they have.** P1 and P2 display proficiency in bringing in specific “technical familiarities” (e.g. phone, email, computer) to configure efficient uses. P1 uses the phone calendar to add value to her paper calendaring. P2’s email scheduling system works even better than online calendars; 1) takes less time to enter more granular information about events, 2) non intrusively, visually reminds her to check her schedule in her always-open inbox, without persistent notifications.

Such sensemaking is sometimes limited by system design assumptions. Not knowing what Google search can do, P4 actively gathers information from her son by talking to him. She now knows that “search” is for “look[ing] up something or someone”, and it is found within the “safari” browser. She sets out to explore, going into safari to type “google”, following the right steps! Yet, she never realizes that the “magnifying glass icon”, and the list of links she interacted with are very much the search system she wanted to use.



Studying Older Adults' Use of Community-Based Technology

Ruipu Hu, Alisha Pradhan, Elizabeth Bonsignore, Amanda Lazar

College of Information Studies

Contact: rhu12@umd.edu

ABSTRACT

This project studied the ways that older adults sustained the use of a makerspace in their retirement community. Our focus on technology use from a *community* perspective is important, because older adults, like others, often use technology as an integral part of being in a community. For example, community spaces such as libraries and senior centers can support older adults in using and learning to use technology. Past work often focuses on supporting older adults in adopting technologies: the project we describe here focuses on the ways that older adults arranged and co-created sociotechnical resources to sustain community use of a new technology environment.

METHODS

Our ethnographically informed involvement took place over 32 months and involved interviews, observations, and co-creation.

FINDINGS

We identified three ways that community members sustained the use of the makerspace:

- 1) **Allocating human resources**, including paid staff, volunteer residents, and even research team members. The goal of allocating these resources was to ensure residents had a “director on duty” to support use of any makerspace machine at any given time. However, community participants often found that they had to navigate constant shifts in makerspace staff support. *These constant shifts in staffing availability appear to be a key feature of technology environments for older adults in these kinds of community settings.*
- 2) **Structuring makerspace classes** for groups of older residents. Classes often involved making specific things (e.g., greeting cards) over independent use – this was one strategy to accommodate the limitations and instability of staffing. *At times, individuals described how it was “not appealing to continue” working in the makerspace because the focused steps taught in classes were unengaging or unsatisfying.*

- 3) **Providing reference materials** that could serve different purposes. Community members leveraged the research team’s interest in studying the space to co-create manuals for makerspace machines. These materials varied with factors such as characteristics of machines and shared past experiences of residents. *The four roles that these materials served include: reminders, specific task-oriented, troubleshooting, and a multipurpose overview.*

CONTRIBUTION

In conducting this work, we make the following contributions. First, we add to the body of knowledge on the challenges and opportunities that older makerspace users face, providing additional detail and nuance regarding the technical and social resources that older adults may need to sustain long-term use of a community makerspace. This includes leveraging staffing, *and working with constant transitions between staff*, to manage the complexity of machines and different technology backgrounds of participants.

Second, our methods involved the *co-creation* of tools (reference materials) *with* older adults in the community. From doing this, we learned ways to support older adults in learning and using technology within a shared community technology environment such as providing reference materials that can accommodate the constant updates that are characteristic of modern technologies and working towards materials that can be personalized to unique community dynamics.

ACKNOWLEDGEMENTS

This work was supported in part by the National Science Foundation under award IIS-1816145 and a NIDILRR Rehabilitation Engineering Research Center grant, #90REGE0008, from the U.S. Admin. for Community Living, Dept. of Health & Human Services



Domain-based Embeddings for Characterizing Online Sharing Behaviors

Ruizhe Fan, Cody Buntain

College of Information Studies, University of Maryland

Contact: ruizhe@umd.edu, cbuntain@umd.edu

As we leverage online behavior to study political engagement, research shows models of online actors benefit when they integrate behaviors across modalities [1-3]. Network embeddings in particular have emerged as a popular method for characterizing these multimodal interactions [4]. Despite the variety of modalities available in social media and online information sharing, the majority of these approaches focus on a small subset, primarily text and network structure, with a few that integrate temporal dynamics and engagement artifacts, such as “likes” [4]. While these approaches have demonstrable value, they omit one of the most popular forms of online engagement: hyperlink (or just “link”) sharing. Entire platforms have been developed around link sharing (e.g., Slashdot, Digg, and Reddit), and link-sharing accounts for a large portion of behavior on Twitter and Facebook, so excluding link-sharing as a modality omits a large portion of online behavior. Similarly, studies have shown the domains an individual chooses to share in their social media feeds are highly indicative of their political preferences [5,6] and can be used to identify inauthentic accounts/disinformation agents [7,8].

This paper addresses this gap in the literature by 1) introducing a general method for extending network-based embedding methods to include domain-sharing behaviors, 2) demonstrating this method across several common modeling tasks in social media and 3) making available pre-trained embeddings that can be readily employed for a diverse range of tasks. Focusing on Twitter, we show how one can construct dense, network-based embeddings of user interactions via a bipartite graph of Twitter users and the domains they share and apply them to characterize online user behaviors. By exploring this gap and evaluating the potential benefits of expanding modalities to include domains, our approach improves the accuracy and versatility of network-based embeddings for social network analysis.

Specifically, we investigate three tasks in our study: 1) comparing retweet and content-sharing networks to infer politicians’ ideologies, 2) assessing our method’s ability to gauge news sources’ political leaning, and 3) characterizing and identifying disinformation agents from their retweeting

and domain-sharing histories. These tasks effectively demonstrate the robustness of our proposed method, and upon completion of this ongoing work, we are hoping to provide valuable insights into integrating content-sharing information with network embeddings to characterize user behaviors in online social networks.

REFERENCES

1. Barberá, P. Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data. *Political Analysis* 23, 76–91 (2015).
2. Joo, J. & Steinert-Threlkeld, Z. C. Image as Data: Automated Content Analysis for Visual Presentations of Political Actors and Events. *Computational Communication Research* 4, (2022).
3. Williams, N. W., Casas, A. & Wilkerson, J. D. Images as Data for Social Science Research: An Introduction to Convolutional Neural Nets for Image Classification. *Elements in Quantitative and Computational Methods for the Social Sciences* (2020).
4. Pan, S. & Ding, T. Social Media-based User Embedding: A Literature Review. in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence* 6318–6324 (International Joint Conferences on Artificial Intelligence Organization, 2019).
5. Messing, S., van Kessel, P. & Hughes, A. Sharing the News in a Polarized Congress. *Pew Research Center - U.S. Politics & Policy*.
6. Eady, G., Bonneau, R., Tucker, J. A. & Nagler, J. News Sharing on Social Media: Mapping the Ideology of News Media Content, Citizens, and Politicians.
7. Golovchenko, Y., Buntain, C., Eady, G., Brown, M. A. & Tucker, J. A. Cross-Platform State Propaganda: Russian Trolls on Twitter and YouTube during the 2016 U.S. Presidential Election. *The International Journal of Press/Politics* 25, 357–389 (2020).
8. Alizadeh, M., Shapiro, J. N., Buntain, C. & Tucker, J. A. Content-based features predict social media influence operations. *Science Advances* 6, eabb5824 (2020).

Switch Mode: Step up from Block-based Programming

Yuhan Lin, David Weintrop, Jason McKenna

Department of Teaching and Learning, Policy and Leadership

Contact: jimmylin@umd.edu

Block-based programming (BBP) is an effective way to introduce young learners to programming and the field of computer science more broadly [2,4,5]. With BBP environments designed for learners of all ages, youth as young as kindergarten are learning to program [1,2]. BBP reduces syntax errors by using visual cues to denote how commands fit together and not allowing incompatible commands to connect [4]. However, BBP is only a steppingstone to text-based programming for learners who are interested in pursuing future computer science coursework or interested in a career that will rely on programming. As learners progress to high school and beyond, they will have to shift from BBP to conventional text-based programming (TBP) languages such as Python or Java. This transition is not always smooth as learners face additional steps and conceptual challenges when they leave BBP [3].

Although the transition from block-based to text-based programming can be challenging [6], relatively few resources currently exist to help learners navigate the transition. Helping learners move from block-based to text-based programming includes helping them recognize the conceptual consistencies between the two modalities and connect productive practices from BBP that are still useful in TBP, while also developing new skills and practices unique to expressing ideas in text-based programming languages. Along with a shift in modality, practices such as debugging also differ between the two programming approaches, as each modality has different affordances and supports to help resolve issues that arise with programs.

To help learners move from BBP to TBP, we have designed Switch mode, a programming approach that blends block-based and text-based programming features within a virtual robotics platform. This lightning talk introduces Switch mode and discussions key design features and their rationale.




Figure 1: The Switch Mode interface and the drag and drop to add Switch mode blocks to a program:




Figure 2: Converting conventional blocks to Switch mode blocks.

Reference:

- [1] Marina Umaschi Bers. 2018. Coding, playgrounds and literacy in early childhood education: The development of KIBO robotics and ScratchJr. In *2018 IEEE Global Engineering Education Conference*, 2094–2102.
- [2] Louise P. Flannery, Brian Silverman, Elizabeth R. Kazakoff, Marina Umaschi Bers, Paula Bontá, and Mitchel Resnick. 2013. Designing ScratchJr: support for early childhood learning through computer programming. In *Proceedings of the 12th International Conference on Interaction Design and Children (IDC '13)*, Association for Computing Machinery, New York, NY, USA, 1–10.
- [3] Michael Kölling, Neil C. C. Brown, and Amjad Altadmri. 2015. Frame-Based Editing: Easing the Transition from Blocks to Text-Based Programming. In *Proceedings of the Workshop in Primary and Secondary Computing Education on ZZZ - WiPSCE '15*, ACM Press, London, United Kingdom, 29–38.
- [4] John Maloney, Mitchel Resnick, Natalie Rusk, Brian Silverman, and Evelyn Eastmond. 2010. The Scratch Programming Language and Environment. *ACM Trans. Comput. Educ.* 10, 4 (November 2010), 1–15.
- [5] Mitchel Resnick, John Maloney, Andrés Monroy-Hernández, Natalie Rusk, Evelyn Eastmond, Karen Brennan, Amon Millner, Eric Rosenbaum, Jay Silver, Brian Silverman, and Yasmin Kafai. 2009. Scratch: programming for all. *Commun. ACM* 52, 11 (November 2009), 60–67.
- [6] David Weintrop and Uri Wilensky. 2019. Transitioning from introductory block-based and text-based environments to professional programming languages in high school computer science classrooms. *Comput. Educ.* 142, (December 2019), 103646.

“I Feel Like I’m Teaching in a Gladiator Ring”: Barriers and Benefits of Live Coding

Caroline Berger¹, David Weintrop^{1,2}, Niklas Elmqvist¹
 College of Information Studies¹, College of Education²
 Contact: cberger2@umd.edu




Figure 1 – Live coding as a performance. Our interview study found that many instructors feel live coding to be a high-stakes and cognitively taxing activity more akin to a live performance. (Illustration by MidJourney version 5.)

“And then I’m teaching [...] in a classroom that feels like a gladiatorial ring. 200 seats in a wall up in front of me. And I have to lean back to see the top. And really the only constraint in that classroom is that it’s terrifying. It is the most terrifying experience I’ve ever had.” – Participant 08 (Computer Science instructor)

Live coding is defined as “*the process of writing code live on a computer in front of students during class*” [1, p. 164]. Live coding models the process of programming [2] and gives students insight into practical programming skills and practices. Additional work is needed to gather the perspectives of teachers on active learning during live coding activities [1].

To address this gap, we designed a prototype (available at <https://gc9vfa.axshare.com/>) and conducted semi-structured interviews with two instructors, seven teaching assistants, and six students. During each session, we solicited feedback on a live coding prototype tool designed to scaffold the practice by enabling students to follow along on their own devices (Figure 2) as well as give instructors a high-level overview of student programming activity. Our thematic analysis of these sessions enabled us to identify factors that serve as **barriers** against effective live coding, as well as those that **benefit** the practice.

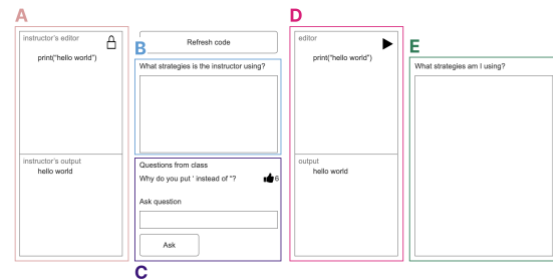


Figure 2 – Student coding along with the instructor view. Student view has the instructor’s editor (A), their own editor (D), and notepads (B, E) where the student articulates the strategies they use and the instructor uses while coding. Refresh code updates the student’s code to match the instructor’s code. The student can ask questions (C).

We found that constraints in the teaching environment, negative feelings, and communication breakdown make live coding difficult. In particular, teachers and students reported intentional and unintentional mistakes, ability to involve the class, and flexibility of scaffolding as benefits of live coding. Based on our findings, we recommend considering how personal computers might detract from live coding lectures and how to direct student attention to important parts of the instructor's code. To support student-teacher collaborative live coding, we recommend tools to support a unified digital space accessed via multiple devices. When a large classroom of students are live coding, errors that students encounter and peeking into student editors may help teachers monitor student progress.

REFERENCES

- [1] Ana Selvaraj, Eda Zhang, Leo Porter, and Adalbert Gerald Soosai Raj. 2021. Live Coding: A Review of the Literature. In *Proceedings of the ACM Conference on Innovation and Technology in Computer Science Education*. ACM, New York, NY, USA, 164–170. <https://doi.org/10.1145/3430665.3456382>
- [2] Sanwar Ali. 2005. Effective Teaching Pedagogies for Undergraduate Computer Science. *Mathematics & Computer Education* 39, 3 (2005), 243–257.

Connecting Interest Spaces to Learning Places with Stickers

Daniel Pauw

College of Information Studies

Contact: dpauw@umd.edu

Educators often seek to guide youth in their informal Science, technology, engineering, and mathematics (STEM) learning explorations. Connecting these interest-driven explorations is a complex process that often seeks to connect STEM learning content to potential interests of the learners. Understanding the context of individual youths while they learn is important to assist youth with connecting learning to their everyday lives. Making connections between interest and learning resources can be difficult though. Communication of their interests can be nonobvious, especially if there is little shared contextual knowledge between interest and learning resources. Using symbolic communication like stickers provide ways for youth to initiate and illustrate conversations about their interests.

Visual Sharing

One way to take youth identity and disposition into account is to look at how youth are interacting and communicating in their everyday lives. Sharing images online goes beyond simple sharing of photographs. Youth today are using significantly more visual means of communicating in general (e.g., emojis and memes) (Figure 1). Therefore, we need to better understand how youth are using these visual means of communicating so that we can help them use these communication practices to support life-relevant learning [1]. We need to understand the ways that visual or pictorial communication can provide valuable information for supporting learning practices [2].

Stickers as Language

I explore the affordances and aesthetics of stickers that youth use to convey their interests. For instance, in co-design sessions, youth demonstrated patterns for selecting and using stickers when communicating. I also examine the interaction experiences stickers promote among their social and family groups (e.g., friends/peers and intergenerational relationships) that provide examples of everyday sticker usage to connect interest with learning resources.

Understanding how digital stickers are utilized in social, online environments will help those facilitating the youth's learning to adapt digital stickers to support interest-driven learning in face-to-face learning environments. Specifically, stickers serve as a kind of in-between zone between the structured, emotional language of emoji and the more

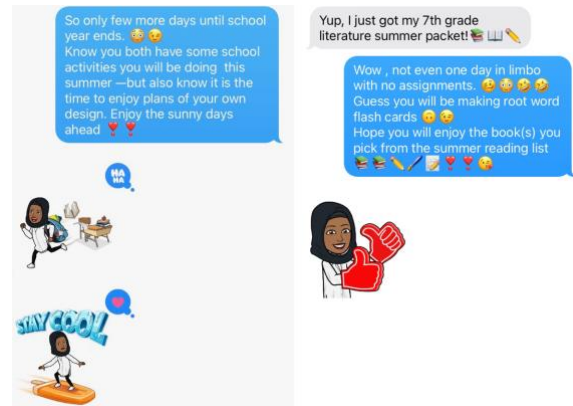


Figure 1. An exchange between a grandmother and her granddaughter using a combination of text, emoji, Tapbacks, and stickers to communicate learning goals. The grandmother's words are in blue, while the granddaughter uses customizable stickers to respond.

freeform expressions of interest of memes. Stickers having structure in the form of packs while being able to potentially be customized or even created by youth themselves – rather than a standard body like emoji – present an opportunity to build a new way of communicating between youths and their older generations.

ACKNOWLEDGEMENTS

Partial funding for this research was provided by NSF Grant No. 1441523.

PAPERS

1. Clegg, T. L., Bonsignore, E., Yip, J. C., Gelderblom, H., Kuhn, A., Valenstein, T. et al. (2012). Technology for promoting scientific practice and personal meaning in life-relevant learning. *Interaction Design and Children*.
2. Pauw, D., Warrick, E., Boston, C., Preece, J., & Clegg, T. L. (2017). Connecting Affinity Spaces to Places and Back: A Look at Pokemon Go. *Computer Supported Cooperative Work and Social Computing*, 275-278.



The Landscape of Introductory Data Science Environments and Curricula

Rotem Israel-Fishelson, Peter F Moon, Rachel Tabak, David Weintrop

Department of Teaching and Learning, Policy and Leadership

Contact: rotemisf@umd.edu

In our increasingly data-driven world, understanding the role of data and the technologies and algorithms built around them, is becoming an essential literacy. This is especially critical for individuals from populations historically excluded from computational and technological fields (e.g., BIPOC individuals, women, neurodiverse individuals, non-native English speakers) as they are disproportionately likely to be negatively impacted by biases and predatory uses of data [1, 4]. It is thus imperative to provide opportunities for all youth to learn foundational data science concepts and practices as part of K-12 education [2].

Given the increasingly data-rich world youth are growing up in, there is tremendous potential for situating data science instruction to draw on learners' interests, experiences, and cultures [3]. Selecting engaging datasets related to learners can foster agency and ownership and is especially important for learners from populations historically excluded from computing. Therefore, it is crucial to consider the provenance and cultural resonance of the datasets used in modern data science curricula, as these are critical for creating inviting, engaging, and equitable data science learning opportunities for all students.

Given the critical role dataset selection plays in engaging learners and signaling whose ideas, values, and cultures are valued by the field, selecting what data learners will engage with should be carried out with great thought and care. Selecting the right datasets to situate data science instruction can be the difference between generative, engaging, and equitable instruction that welcomes learners to the field or instruction that alienates and further perpetuates existing racial, gender, and socioeconomic gaps. Although the academic literature often highlights the importance of reflecting on youth experiences in the datasets they study, little work has been done to understand these datasets more broadly. Additionally, as a core practice of data science is manipulating and interrogating these datasets, it is also important to consider the tools students use to carry out this work, as the tools themselves play an essential role in shaping learners' experiences [5].

In this work, we investigate the datasets in four of the most widely used high school data science curricula:

Bootstrap:Data Science, CodeHS, Introduction to Data Science, and YouCubed Explorations in Data Science. In conducting this analysis, we recognize how these curricula shape the emerging landscape of K-12 data science and seek to highlight the importance of choosing datasets that draw on the cultural knowledge and lived experiences of the youth who will be learning with them. Our analysis identified 296 distinct datasets used across them, which we qualitatively analyzed using a series of analytical lenses that considered various dimensions of the datasets. The result of this analysis is a comprehensive understanding of the data being used to introduce the current generation of learners to data science through four of the most widely used curricula. The findings from applying these analytic lenses provide direction for future revisions and innovations in data science instruction that better situate instruction in the lived experiences of today's students. This work contributes to the larger goal of preparing today's students for the data-rich futures that await them.

ACKNOWLEDGMENTS

This research is supported by the National Science Foundation (Award # 2141655).

PAPERS

- [1] Benjamin, R. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code*. Polity.
- [2] LaMar, T. and Boaler, J. 2021. The importance and emergence of K-12 data science. *Phi Delta Kappan*. 103, 1 (Sep. 2021), 49–53.
- [3] Lee, V.R. and Delaney, V. 2021. Identifying the Content, Lesson Structure, and Data Use Within Pre-collegiate Data Science Curricula. *Journal of Science Education and Technology*.
- [4] O'Neil, C. 2016. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books.
- [5] Pimentel, D.R. et al. 2022. *Tools to Support Data Analysis and Data Science in K-12 Education*. National Academies Workshop on Foundations of Data Science for Students in Grades K-12.

Virtual Space and Social Interactions for Professional Goals

Victoria Chang, Ge Gao
College of Information Studies
Contact: vchang90@umd.edu

Overview

The evolution of virtual interactive technology has created opportunities to explore and implement novel ways of supporting for social interaction in virtual settings. In the current research project, we draw on observations and in-depth interviews to understand how the spatial design of social virtual reality (SVR) platforms can affect people's behaviors at professional events hosted via those platforms.

Participants of our research include 15 individuals who were attendees, hosts, and/or designers of one or more SVR-based professional events. During the data collection phase, our research team shadowed each participant at their self-selected event and took notes about how they interacted with others at the event. This observation was followed with interviews that aimed to understand the participant's subjective experience at the event. In total, 19 events were observed.

Background

Our previous survey-based study sought to understand how people using general SVR platforms may leverage the affordances of the platforms to initiate, maintain, and conclude interactions [1]. In short, our findings suggested that, while face-to-face interactions generally utilized proximity-based cues (e.g., adjusting distance), people skipped certain non-verbal cues in VR and, instead, made verbal interactions and communication more immediately.

Our current work extends the above one by looking into how people perform social interactions with each other at SVR-based professional events. In particular, since VR is designed with spatiality in mind, we want to understand what aspects of virtual spatial design may influence people's interpretations and choices of appropriate ways to behave. This inquiry echoes what other HCI researchers have discussed such that the "appropriate behavioral framing is often defined by the place designed for inside of a space [2]." A person's planning of the next action to perform are tied to their interpretations of the space design, based on their understanding of how they would be expected to behave in physical reality.

Work-in-Progress

Through the ongoing data analysis, we aim to identify the cues leveraged by participants to interact with each other,



Figure 1. These two images show examples from two separate platforms designed for hosting different professional events. The left image is from an instructor-led coding event in AltSpace VR, and the right image is from an academic networking event happening in Venu.

as well as how spatial design of SVR platforms may impact participant's perception of the potential and outcomes of professional goal-oriented interactions.

We have found that the event hosts and platform designers of an SVR-based professional event often formulated their design of the virtual space by referring to their experiences from congruent events in the physical reality. Some hosts and designers focused on creating event spaces that have separated rooms, such as an auditorium and a poster room as seen in **Figure 1**. Others incorporated designs to intentionally signal flexibility in avatar choice, movement, and exploration of the space. These designs served the goal of supporting the attendees' social interaction and networking in the virtual space.

Further, many attendees at those events did not act as how the hosts and designers hoped because they were not clear about potential consequences of behavior. Those event hosts, then, usually ended up taking a crucial role in guiding the attendees toward understanding and predicting what behaviors and code of conduct would be considered as appropriate in each sector of the virtual space.

REFERENCES

1. Jingchao Fang, Victoria Chang, Ge Gao, and Hao-Chuan Wang. 2021. Social Interactions in Virtual Reality: What Cues Do People Use Most and How. *CSCW '21*. ACM, 49-52.
2. Steve Harrison and Paul Dourish. 1996. Re-Place-Ing Space: The Roles of Place and Space in Collaborative Systems. *Proceedings of the ACM Conference on Computer Supported Cooperative Work*. 67-76.



How Library IT Staff Navigate Privacy and Security Challenges and Responsibilities

Noel Warford, Alan F. Luo, Samuel Dooley,
Rachel Greenstadt*, Michelle L. Mazurek, Nora McDonald⁺

University of Maryland, *New York University, ⁺George Mason University

Contact: nwarford@umd.edu

Libraries provide critical IT services to patrons who lack access to computational and internet resources. We conducted 12 semi-structured interviews with library IT staff to learn about their privacy and security protocols and policies, the challenges they face implementing them, and how this relates to their patrons. We frame our findings using Sen's capabilities approach, which centers freedom to achieve wellbeing as a moral imperative and posits that this freedom must be understood in terms of people's capabilities. We find that library IT staff are primarily concerned with protecting their patrons' privacy from threats outside their walls—police, government authorities, and third parties. Despite their dedication to patron privacy, library IT staff frequently must grapple with complex tradeoffs between providing easy, fluid, full-featured access to Internet technologies or third-party resources, protecting library infrastructure, and ensuring patron privacy.

ROLES AND RESPONSIBILITIES

A near-universal responsibility of our participants was to manage public-facing PCs, which patrons use for both recreation and essential tasks. Library IT staff also manage access to digital resources—often through third-party vendors—to give patrons access to ebooks, academic journals, and other services. Occasionally, library IT staff provide basic digital skills education, either formally through scheduled classes or informally through ad-hoc troubleshooting.

GOALS

Participants were dedicated to making the library a safe and trusted space where patrons can be free to seek information without fear of surveillance by government institutions. Participants often framed this desire in relation to the core library values of patron privacy and intellectual freedom. Another goal of library IT staff members was to support and improve patrons' capabilities by ensuring that use of computing resources was easy, and that patrons could accomplish their goals using library resources.

THREATS

Library IT staff view the greatest threats to patrons' capabilities as institutions like governments and third parties. When asked about security and privacy concerns, participants rarely mentioned their patrons posing any malicious or intentional security or privacy risk to the library or other patrons. Some pointed out that when patrons do create risks, it's most often unintentional. Several participants expressed concerns about third-party vendors or law enforcement having access to patron data. Participants generally felt ill-equipped to fully understand or counteract threats from third parties.

TENSIONS

Ensuring patron privacy can sometimes be at odds with taking maximum precautions to protect the security of library infrastructure. For example, libraries could choose to install monitoring software to detect misbehavior or enforce strong login requirements in order to attribute any problems to particular users, but monitoring and logging inherently pose a threat to privacy by tracking user activities. Our participants almost always prioritized protecting patrons' privacy rather than protecting library infrastructure.

Measures taken to protect patron privacy—like fully resetting public PCs between sessions or blocking software downloads—can also conflict with the utility of library resources. One participant describes a patron losing all their documents because they did not realize that anything "saved" to the computer would be erased at the end of the session. Blocking downloads can also limit patrons' access to modern social and work-related activities, like using Zoom to join a video call. These are critical problems for library patrons who may not otherwise have access to computing resources.

PAPER

Alan F. Luo, Noel Warford, Samuel Dooley, Rachel Greenstadt, Michelle L. Mazurek, and Nora McDonald. How Library IT Staff Navigate Privacy and Security Challenges and Responsibilities. In *Proc. USENIX Security*, 2023 (to appear).



Design Insights from Smart Home Power Users on Enhancing Data Transparency, Visibility, and Control

Sunyup Park*, Anna Lenhart*, Michael Zimmer**, Jessica Vitak*

*College of Information Studies, University of Maryland

**Department of Computer Science, Marquette University

Contact: sypark@umd.edu

OVERVIEW

Smart home devices have gained popularity in recent years as they offer a convenient and efficient way to control and monitor various aspects of one's home remotely. While these devices offer significant benefits—ranging from convenience and cost efficiency to added security and accessibility—they have also led to an increased focus on data privacy risks. In this study, we provide insights from smart home power users who make heavy use of existing privacy controls and, in some cases, use third-party platforms to monitor traffic, build custom dashboards, and create custom domain block lists. We argue that this population is uniquely positioned to provide feedback and insights that everyday users of smart home technologies may not consider.

METHOD

We recruited 32 participants from across the U.S., using criterion sampling to identify those who used a wide variety of devices, customized those devices, and expressed some concerns regarding data generated from devices. During virtual focus group sessions, participants shared details on how they used and customized their smart homes to maximize utility while addressing privacy concerns. Participants also completed two brainstorming activities using Google Jamboard, including one focused on the features they wanted to better visualize and manage their smart home data.

FINDINGS AND DISCUSSION

Our findings highlight the drawbacks and limitations our sample of smart home power users identified with their current setups, as well as the design features they wanted to enhance their ability to see and manage their data. Specifically these recommendations centered on three core areas around enhancing data transparency, visibility, and control.

Transparency: Participants discussed two primary ways to make smart home data and device information more transparent: improved product labeling detailing data practices and improved notifications. Multiple participants

referenced Apple's privacy nutrition labels and suggested ways to enhance them (e.g., more exposure, cross-checking information, and adding information about the relationship between device manufacturers and brands to understand how data might flow across other platforms and/or services).

Visibility: Moving beyond transparency, participants further underscored the importance of ongoing visibility into their smart home system and data flows, including device status; types of data being sent through, in, and out of their network; where data is going; and how frequently data is moving or being collected. They emphasized flexibility in sorting data and user-friendly interfaces to understand transaction logs.

Data controls: Building from the need for increased transparency and visibility, participants noted that these features were useful but still limited, especially when many devices and apps follow a "take it or leave it" approach (i.e., agree to full terms of use or don't use device/app). While many of our participants used their advanced networking skills to customize their smart home ecosystems, they identified gaps to be addressed, including the need for controls to create and manage block lists and customizable flags, as well as different settings for different users.

Overall, these findings highlight the challenges to managing data flows from smart devices. If our participants—who were highly motivated to customize their homes to maximize benefits while minimizing external data flows—expressed significant frustration with currently available options, then this signals an urgent need to develop more user-friendly features and controls that are accessible for everyday smart home users.

RELATED WORK BY AUTHORS

Lenhart, A., Park, S., Zimmer, M., & Vitak, J. (2023). "You shouldn't need to share your data": Perceived privacy risks and mitigation strategies among smart home power users. *Proceedings of the 26th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW'23)*.



Who and What Influences Privacy Discourse? Defining Problems, Approaches, and Solutions

Min Cheong Kim, Zothan Mawii, Anna Lenhart, Katie Shilton, Jessica Vitak
College of Information Studies, University of Maryland
Contact: mckim@umd.edu

OVERVIEW

While HCI research largely focuses on questions around the design of tools and how users interact with those tools, technology policy can play a tremendous role in shaping how and for whom technologies are built. Likewise, when building safe and secure devices that minimize privacy risks to end users, we must also consider wider conversations about privacy policy; in the US, this has been an increasing focus on Capitol Hill in recent years.

Privacy policy conversations in the US are shaped by five key actors: (1) technology companies who undertake new data collection practices; (2) academics who research privacy issues from a multitude of perspectives; (3) non-profit organizations and advocacy groups who fight to get citizens' concerns heard; (4) the media who report on harms and risks to citizens; and (5) policymakers and legislators who represent their constituents' interests. That said, the outsized influence of tech giants and a weakened enforcement infrastructure have made privacy policymaking a challenge.

This paper investigates how privacy ideas—particular definitions of privacy, data sharing issues, and/or potential solutions for privacy concerns—become influential in shaping regulations. It also examines how the influences of particular privacy definitions, issues, and solutions are reflected in emerging national privacy policy. We seek to identify how privacy definitions, problems, and solutions become mainstream and impact privacy law and regulations, and to identify factors influencing privacy conversations and regulation in the US. From this, we can examine policy's impact on design conjointly with exploring ways for HCI researchers to have a bigger voice in discussions of the design of new technologies.

METHODS

Through a qualitative study with 23 US-based privacy experts from academia, non-profit organizations, and the

media, we trace the story of privacy ideas, from diffuse public and academic conversations to the narrower content of legislative hearings and bills. Interview questions captured how they defined “influence” in the US privacy landscape, strategies and activities they saw as influential, and current/future privacy issues they felt deserved more attention.

FINDINGS

We identified three major themes covering the importance of different engagement strategies with members of the privacy community and the wider public; privacy issues arising out of new and current technologies; and social factors that contribute to being influential. Importantly, we did not define influence to our participants, but rather had them tell us the factors that make a person, idea, or action influential in privacy discourse. A plurality of experts we interviewed defined influence in the privacy space as being able to influence privacy legislation and conversations in legislative spaces.

Most experts agreed that being able to communicate and translate complex ideas for diverse audiences was critical to influential privacy work. In addition to expected factors like prestige and social capital, experts placed high value on being able to shift the way people frame or conceive of privacy issues. Legal scholarship's robust engagement with privacy research was well-demonstrated in participants' emphasis on legislative impact, as well as in the current national-level interest in advancing privacy legislation. However, the most influential privacy ideas, including contextual integrity and differential privacy, emerged from disciplines outside of law, reminding us of the importance of being inclusive of different disciplines when it comes to influential privacy discourse.

ACKNOWLEDGEMENTS

This project is funded through a research grant from Meta.

Investigating How Survivors of Domestic Violence Seek Digital Safety Advice Online

Julio Poveda, Sridevi Suresh, Wentao Guo, Rushil Dandamudi, Nora McDonald*, and Michelle L. Mazurek

University of Maryland, *George Mason University

Contact: jpoveda@umd.edu

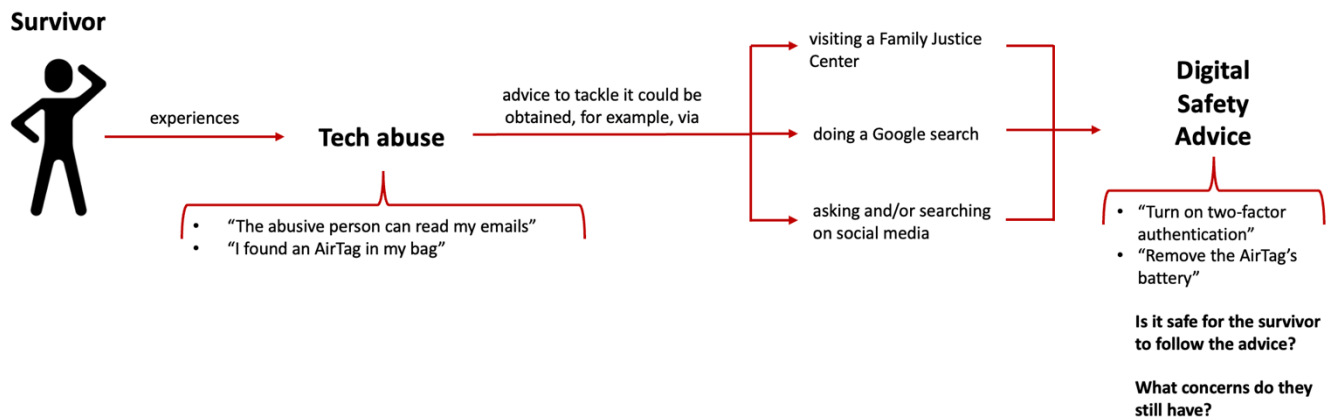


Figure 1 – The tech advice-seeking process DV survivors follow.

Domestic Violence (DV) is a widespread societal issue. Prior work has shown how abusers utilize smart home devices, spyware, and financial apps among others to monitor and harass their current or former partners [1, 2, 3]. However, it has not thoroughly been studied how DV survivors seek digital safety advice online to prevent and or combat technology-facilitated abuse.

In this project, our goal is to study how DV survivors find tech advice and understand how they decide which advice to follow. We are interested in answering questions like what is the quality of the advice they receive via social media (if they receive so)?

We plan to conduct semi-structured interviews with DV primary survivors, as well as secondary survivors (a relative or friend who has helped a primary survivor). We have considered and incorporated various safety measures to avoid participants' retraumatization (e.g., not asking direct questions about their abuse history, sharing support resources) and ensure that they benefit from participating in the project [4].

REFERENCES

- Stephenson, S., Almansoori, M., Emami-Naeini, P., Huang, D.Y., & Chatterjee, R. (2023). Abuse Vectors: A Framework for Conceptualizing IoT-Enabled Interpersonal Abuse In 32nd *USENIX Security Symposium (USENIX Security)*.
- Chatterjee, R., Doerfler, P., Orgad, H., Havron, S., Palmer, J., Freed, D., ... & Ristenpart, T. (2018). The spyware used in intimate partner violence. In *2018 IEEE Symposium on Security and Privacy (SP)* (pp. 441-458).
- Bellini, R. (2023). Paying the Price: When Intimate Partners Use Technology for Financial Harm. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. (pp. 1-17).
- Bhalerao, R., Hamilton, V., McDonald, A., Redmiles, E. M., & Strohmayer, A. (2022). Ethical Practices for Security Research with At-Risk Populations. In *2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)* (pp. 546-553).
- Leitão, R. (2021). Technology-Facilitated Intimate Partner Abuse: a qualitative analysis of data from online domestic abuse forums. *Human-Computer Interaction*, 36(3), 203-242.

HCIL Honors and Awards

On the following pages, we celebrate some of the awards and honors that HCIL faculty and students have received in the last year.

Promotion & Tenure	
Hernisa Kacorri	Received tenure in 2023 (iSchool)
David Weintrop	Received tenure in 2023 (iSchool & Education)

Awards for research papers	
Hal Daumé III, Kyungjun Lee & colleagues	Best paper award, ACL 2022: “What’s Different between Visual Question Answering for Machine “Understanding” Versus for Accessibility?”
Hal Daumé III & colleagues	Test of Time award, ACL: “Midge: Generating Image Descriptions From Computer Vision Detections”
Catherine Plaisant, Jesse Grosjean and Ben Bederson	Test of Time award, IEE VIS: “SpaceTree: Supporting Exploration in Large Node Link Tree, Design Evolution and Empirical Evaluation”
Nitzan Koren; David Weintrop; Mega Subramaniam	ALISE/ProQuest Methodology Paper Award, 2022: “Using Design-Based Implementation Research Method to Create Computational Thinking Assessment Tools for Youth Programs in Public Libraries”
Sheena Erete & colleagues	Best Paper award, ACM CHI 2023: “Deceptive Design Patterns in Safety Technologies: A Case Study of the Citizen App”
Deepthi Raghunandan, Niklas Elmqvist, Leilani Battle	Honorable Mention, ACM CHI 2023: “Code Code Evolution: Understanding How People Change Data Science Notebooks Over Time”
Jiasheng Li, Zeyu Yan, Arush Shah, Jonathan Lazar, Huaishu Peng	Honorable Mention, ACM CHI 2023: “HonorToucha11y: Making Inaccessible Public Touchscreens Accessible”
Hanuma Teja Maddali & Amanda Lazar	Best Paper award, ACM CHI 2023: “Understanding Context to Capture when Reconstructing Meaningful Spaces for Remote Instruction and Connecting in XR”
Amanda Lazar & colleagues	Best Paper award, ACM DIS 2022: “Control Matters in Elder Care Technology: Evidence and Direction for Designing It In”
Yuhan Luo & Eun Kyoung Choe	Honorable Mention, ACM ISS 2022: “NoteWordy: Investigating Touch and Speech Input on Smartphones for Personal Data Capture”

Jonggi Hong, Jaina Gandhi, Ernest Essuah Mensah, Farnaz Zamiri Zeraati, Ebrima Jarjue, Kyungjun Lee, & Hernisa Kacorri	Best Paper Nominee, ACM ASSETS 2022: “Blind Users Accessing Their Training Images in Teachable Object Recognizers”
Zijian Ding, Arvind Srinivasan, Stephen MacNeil, & Joel Chan	Honorable Mention, ACM C&C 2023: “Fluid Transformers and Creative Analogies: Exploring Large Language Models’ Capacity for Augmenting Cross-Domain Analogical Creativity”

Books Published

Gregg Vanderheiden, Jonathan Lazar, Amanda Lazar, Hernisa Kacorri, & Bern Jordan	“Technology and Disability: 50 Years of Trace R&D Center Contributions and Lessons Learned” (Springer, 2022)
Jonathan Lazar	Co-editor, “Proceedings of the Cambridge Workshop on Universal Access and Assistive Technology” (Springer, 2023)

Other Awards & Achievements

Jessica Vitak	Named to “Top 100 Brilliant Women in AI Ethics” list
Jessica Vitak	Named a non-resident fellow at the Center for Democracy and Technology (CDT)
Niklas Elmqvist	Villum Investigator, Villum Foundation
Michelle Mazurek	Carnegie Mellon CyLab Distinguished Alumni Award
Salma Elsayed-Ali	Ann G. Wylie Dissertation Fellowship University of Maryland
Salma Elsayed-Ali	ArtsAMP Interdisciplinary Graduate Student Research Grant – Arts for All, University of Maryland

HCIL Student Graduation

Finally, we celebrate the many HCIL students who have graduated over the last year. Students are an integral part of the lab's success, and these students have worked with faculty on a range of important research projects in recent years. Congratulations and good luck to each of these students!

Student Name	Advisor	Grad Date	Degree
Andrea Batch	Niklas Elmqvist	May 2022	PhD (iSchool)
Yuhan Luo	Eun Kyoung Choe	May 2022	PhD (iSchool)
Lautaro Cabrera	Tammy Clegg & Diane Jass Ketelhut	May 2022	PhD (Education)
Kyungjun Lee	Hernisa Kacorri	August 2022	PhD (Computer Science)
Eric Newburger	Niklas Elmqvist	May 2023	PhD (iSchool)
Xiaoyun Huang	Jessica Vitak	May 2023	PhD (iSchool)
Nihal Katirci	Caro Williams-Pierce	May 2023	PhD (iSchool)
Ekta Shokeen	Caro Williams-Pierce	May 2023	PhD (iSchool)
Nekabari Sigalo	Vanessa Frias-Martinez	May 2023	PhD (iSchool)
Deepthi Raghunandan	Niklas Elmqvist	May 2023	PhD (Computer Science)
Kelsey Fulton	Michelle Mazurek	May 2023	PhD (Computer Science)
Alisha Pradhan	Amanda Lazar	Summer 2023	PhD (iSchool)

Student Name	Advisor	Grad Date	Degree
Jarrett Lee	Eun Kyoung Choe	May 2022	Masters (HCIM)
Sarah DiPasquale	Joel Chan	May 2022	Masters (HCIM)
AJ Rudd Jr	Joel Chan	May 2022	Masters (HCIM)
Katelyn DeValk	Niklas Elmqvist	December 2022	Masters (CS)
Caroline Berger	Niklas Elmqvist	May 2023	Masters (HCIM)
Abhinav Kannan	Niklas Elmqvist	May 2023	Masters (HCIM)
Amelia Short	Amanda Lazar	May 2023	Masters (HCIM)
Aishwarya Shettigar	Eun Kyoung Choe	May 2023	Masters (HCIM)
Sparsh Paliwal	Jonathan Lazar	May 2023	Masters (HCIM)
Arvind Srinivasan	Joel Chan	May 2023	Masters (HCIM)

THANK YOU TO OUR WONDERFUL SPONSORS



UNIVERSITY OF
MARYLAND

hcil.umd.edu