

Academic Research Statement

Scott Carter
scott@madpickle.net
madpickle.net/scott

The goal of my research is to bring new technologies and techniques to bear on problems facing modern knowledge workers. Specifically, I have led the design, development, and evaluation of tools for 1) helping researchers discover user needs and analyze and organize data; 2) making GUI and mobile interfaces more accessible; and 3) archiving, searching, and reusing digital content in the context of distributed systems.

Tools for research

We fail more often because we solve the wrong problem than because we get the wrong solution to the right problem. –Russell L. Ackoff

While there are countless frameworks, toolkits, code editors, and processes designed to help engineers build solutions to problems, finding the right problem to solve often gets short shrift. My PhD work focused on addressing this problem by building tools to help researchers better leverage needfinding methods using modern technology. One such method was the diary study, which is a method of understanding participant behavior and intent in situ that minimizes the effects of observers on participants. I analyzed ongoing diary studies and ran several on my own to understand how they were being used and where new (at the time) phone-based messaging and media recording technologies could improve them. My initial work led to new methods and tools for diary studies, involving media capture, annotation, researcher feedback, and elicitation [1].

I later extended these early technologies to a system that could support a wider variety of needfinding methods, including experience sampling (in which study participants are asked a battery of questions over an extended period of time) and early stage prototyping. Using a confederation of desktop and mobile tools that I built (collectively called Momento), experimenters could respond to participant requests, ask participants to manually capture or record data, or automatically gather context data from mobile devices [2].

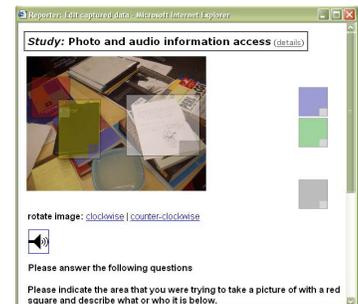
The needfinding methods I designed are still in use in many institutions today, over a decade later. The tools I initially built were based on early (pre-iPhone) mobile technologies, however, recently I built from scratch a new experience-sampling system for modern desktop- and mobile-web systems. Researchers at the University of Cambridge used this new system as recently as this year to conduct needfinding studies of hundreds of participants across the UK (a publication is in preparation).

I have collaborated with other researchers on a variety of other tools and methods to help researchers, including interactive visualizations of topics and relationships in computational documents [3] as well as toolkits and methods for building, designing, and evaluating peripheral and ambient displays [4,5].

Tools for guidance and access

One of the most critical aspects of designing tools with a human-centered approach is understanding that people often have limited attention, training, or abilities because of physical or context-based impairments. I have found that building tools for these scenarios is a fascinating challenge as it requires designing and evaluating tools that incorporate both output and input adaptations.

SeeReader is one such tool, a text-to-speech-based document reader for mobile devices that adapts output (the way a document is displayed) to help people understand rich documents. The tool automatically analyzes a document to link callouts in text to the visual components to which they refer. For example, while reading the text “as shown in Figure 2” aloud to the user the visual display automatically frames Figure 2 in the document. It also includes a skimming



mode that reads section titles as the user moves their finger around a circular widget. With this approach, people with visual impairments, or who need to divide their attention between tasks, can use audio to listen to a document while having access to important and more visually salient content at-a-glance [6].

The Input Adapter Tool (IAT) is another application that automatically analyzed and adapted GUI components to work with available input interactors. For example, IAT can adapt a painting interface that depends on mouse input to work with switch- or speech-based controls for users with motor impairment issues [7].

Relatedly, NudgeCam is a mobile application I built that analyzes media and context data in real time to help people shoot higher quality videos, take better photos, and make more professional audio recordings. The tool combines well known capture heuristics and task-specific guidelines with live feedback to guide users who might have poor visual or audio understanding due to lack of experience, distractions, or other impairments [8].

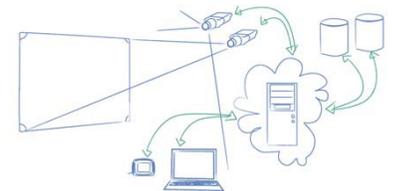


Tools for workers

Global warming, pandemics, and other ecological issues are pushing knowledge workers into a future in which collocated work will be necessarily curtailed, or in some cases, eliminated entirely. Now more than ever, research into tools to support distributed workers is paramount.

Communicating knowledge is a key concern for distributed teams. Without systems to help users capture and access information, methods for conducting work can become siloed in disparate sections of an organization. ShowHow was a set of technologies I built to help users capture and transfer knowledge. Based on extensive formative studies, we created a variety of methods for people to record and access information using mobile devices, head-mounted displays, and an interactive tutorial creation tool. Our findings suggested that head-mounted displays in particular have a key role to play in the transfer of institutional knowledge [9].

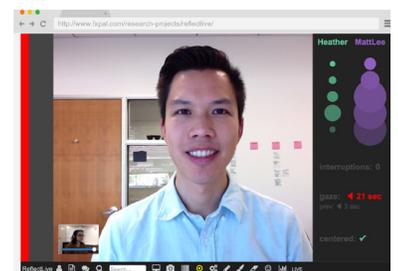
Many conversations in work places take place informally using props such as notepads and whiteboards. Another system I helped build, ReBoard, automatically detected and captured content drawn on whiteboards or other display surfaces. Content was then made available to distributed workers via desktop and mobile interfaces [10].



Public information displays can also help connect distributed work teams. In earlier work, I added an annotation system to a network of public displays to explore how to best augment informal, workaday discussions [11]. Lessons learned from evaluations of this system helped re-design the system to support a wider set of distributed teams and installation locations.

More recently, I led a team of user experience researchers investigating current teleconferencing practices. Our field work showed that current tools do not provide adequate support for document sharing and search [12]. To address the issues we discovered, I designed MixMeet, a WebRTC-enabled web platform takes advantage of live content analysis to support real-time search and content reuse [13]. The platform also included a mobile app and watch-based interface [14]. After deploying and testing this platform, I worked with another colleague to extend it to support another teleconference application, ReflectLive, which provides real-time feedback to doctors to help them improve their communication behaviors with patients [15].

I have collaborated with colleagues on a wide variety of other projects, including an investigations of tagging technologies [16], active reading devices [17], interactive headware [18], etc. I am also excited to work with other researchers and engineers on interactive technologies and techniques for the workplace and beyond.



Next steps

One of the most important barriers to knowledge work is that digital media lacks presence, making it difficult for team members to maintain a shared understanding of ongoing activities and problems. Unlike work that involves physical artifacts, there are few objects readily available to serve as props for the type of informal discussions that can increase team cohesiveness and spur informal conversations that lead to creative solutions. This issue is

exacerbated in remote work. Historically, this problem has inspired a broad array of research programs that I and others have explored, from glanceable and peripheral displays, semi-automated information visualizations, and more recently deep learning-based visual and textual summarizations. Still, imbuing digital artifacts with presence enough to anchor conversations remains a challenge. In more recent work, I have investigated methods to analyze digital content to make key components ready-to-hand for sharing and reuse in remote work scenarios [19,20]. Going forward, I am curious to understand how sharing media can play a role in helping people of all abilities and dispositions reach their potential as knowledge work continues to unfurl across time and space.

Service

I have served as the Administrative Editor for the HCI Journal since 2017. Human-Computer Interaction is one of the top journals for computer science and electronics and (as of August 2020) has the highest impact factor of any HCI-focused journal (<http://www.guide2research.com/journal/human-computer-interaction>).

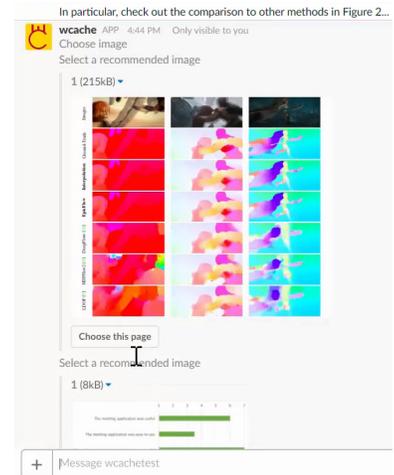
I am the treasurer for UIST 2021 and have also served as a committee member for Mobile HCI, Multimedia, WWW, ICWSM, AAAI, and other conferences. I have reviewed papers for many other conferences, symposiums (including DIS, CHI, UIST, Ubicomp, CSCW), and journals (including IEEE Computer Graphics and Applications, IEEE Pervasive Computing, MultiMedia Tools and Applications).

I enjoy mentoring interns and I have been lucky to have completed several successful research projects with a wide diversity of excellent students, including Christine Dierk [18], Stacy Branham [8,10], Ville Mäkelä [9,14], and Heather Faucett [15].

I plan to expand my role in academic service in the future, focusing in particular on venues that foster work at the intersection of distributed work, user experience, and technical tool building.

References

1. Carter, S. and J. Mankoff. When Participants Do the Capturing: The Role of Media in Diary Studies. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI), Pages: 899-908. 2005.
2. Carter, S., J. Mankoff, and J. Heer. Momento: Support for Situated Ubicomp Experimentation. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI), Pages: 125-134. 2007
3. Wenskovitch, J., J. Zhao, S. Carter, M. Cooper, and C. North. Albireo: An Interactive Tool for Visually Summarizing Computational Notebook Structure. Symposium on Visualization in Data Science (VDS), 2019.
4. Matthews, T., T. Rattenbury, and S. Carter. Defining, Designing, and Evaluating Peripheral Displays: An Analysis Using Activity Theory. Journal of Human-Computer Interaction, 22(1/2). Pages: 221-261. 2007.
5. Matthews, T., A. K. Dey, J. Mankoff, S. Carter, and T. Rattenbury. A Toolkit For Managing User Attention in Peripheral Displays. Proceedings of the ACM Symposium on User Interface Software and Technology (UIST), Pages: 247-256. 2004.
6. Carter, S. and L. Denoue. SeeReader: An (Almost) Eyes-Free Mobile Rich Document Viewer. International Journal of Computer Science Issues (IJCSI). Vol. 1. Pages: 36-41. 2009.
7. Carter, S., A. Hurst, J. Mankoff, and J. Li. Dynamically Adapting GUIs to Diverse Input Devices. Proceedings of The ACM Conference on Computers and Accessibility (ASSETS), Pages: 63-70. 2006.
8. Carter, S., J. Doherty, J. Adcock, and S. Branham. NudgeCam: Toward targeted, higher quality media capture. Proceedings of the ACM International Conference on Multimedia (MM), Pages: 615-618. 2010.



9. Carter, S., M. Cooper, P. Qvarfordt, and V. Mäkelä. Creating expository documents with web-based authoring and heads-up capture. *IEEE Pervasive Computing*, 14(3). Pages: 44-52. 2015.
10. Branham, S., G. Golovchinsky, S. Carter, and J. Biehl. Let's Go from the Whiteboard: Supporting Transitions in Work through Whiteboard Capture and Reuse. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, Pages: 75-84. 2010.
11. Carter, S., E. Churchill, L. Denoue, J. Helfman, and L. Nelson. Digital Graffiti: Public Annotation of Multimedia Content. *Extended abstracts of the ACM Conference on Human Factors in Computing Systems (CHI)*, Pages: 1207-1210. 2004.
12. Marlow, J., S. Carter, N. Good, and J.-W. Chen. Beyond Talking Heads: Multimedia Artifact Creation, Use, and Sharing in Distributed Meetings. *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW)*, Pages: 1703-1715. 2016.
13. Denoue, L., M. Cooper, and S. Carter. Searching Live Meetings: "Show me the Action". *Proceedings of the ACM Symposium on Document Engineering (DocEng)*, Pages: 195-198. 2015.
14. Carter, S., J. Marlow, A. Komori, and V. Mäkelä. Bringing mobile into meetings: Enhancing distributed meeting participation on smartwatches and mobile phones. *Proceedings of the International Conference on Human-Computer Interaction with Mobile Devices and Services (Mobile HCI)*, Pages: 407-417.
15. Faucett, Heather A., Matthew L. Lee, and Scott Carter. I Should Listen More: Real-time Sensing and Feedback of Non-Verbal Communication in Video Telehealth. *Proceedings of ACM Human Computer Interaction*, 1(CSCW), Article 44, 19 pages. 2017.
16. Carter, S., C. Liao, L. Denoue, G. Golovchinsky, Q. Liu. Linking digital media to physical documents: Comparing content-based and marker-based tags. *IEEE Pervasive Computing*, 9(2). Pages: 46-55. 2010.
17. Golovchinsky, G., S. Carter, and A. Dunnigan. ARA: The Active Reading Application. *Proceedings of the ACM International Conference on Multimedia (MM)*, 2011.
18. Dierk, C., S. Carter, P. Chiu, T. Dunnigan, D. Kimber. Use Your Head! Exploring Interaction Modalities for Hat Technologies. *Proceedings of the on Designing Interactive Systems Conference (DIS)*, Pages: 1033-1045. 2019.
19. Denoue, L., S. Carter, J. Marlow, and M. Cooper. DocHandles: Linking Document Fragments in Messaging Apps. *Proceedings of the ACM Symposium on Document Engineering (DocEng)*, Pages: 81-84. 2017.
20. Carter, S., L. Denoue, and D. Avrahami. Documenting Physical Objects with Live Video and Object Detection . *ACM International Conference on Multimedia (MM)*, Pages: 1032-1034 . 2019.