

Crowdsourcing Accessibility: Human-Powered Access Technologies

Erin Brady

Department of Human-Centered Computing
School of Informatics and Computing
Indiana University – Purdue University Indianapolis
brady@iupui.edu

Jeffrey P. Bigham

Human-Computer Interaction Institute
Carnegie Mellon University
jbigham@cmu.edu

now

the essence of knowledge

Boston — Delft

Foundations and Trends[®] in Human-Computer Interaction

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
United States
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is

E. Brady and J. P. Bigham. *Crowdsourcing Accessibility: Human-Powered Access Technologies*. Foundations and Trends[®] Human-Computer Interaction, vol. 8, no. 4, pp. 273–372, 2014.

This Foundations and Trends[®] issue was typeset in L^AT_EX using a class file designed by Neal Parikh. Printed on acid-free paper.

ISBN: 978-1-68083-035-4

© 2015 E. Brady and J. P. Bigham

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

**Foundations and Trends[®] in
Human-Computer Interaction**
Volume 8, Issue 4, 2014
Editorial Board

Editor-in-Chief

Desney S. Tan
Microsoft Research
United States

Editors

Ben Bederson
University of Maryland

Sheelagh Carpendale
University of Calgary

Andy Cockburn
University of Canterbury

Jon Froehlich
University of Maryland

Juan Pablo Hourcade
University of Iowa

Karrie Karahalios
*University of Illinois
at Urbana-Champaign*

Youn-Kyung Lim
*Korea Advanced Institute of Science and
Technology*

Nuria Oliver
Telefonica

Orit Shaer
Wellesley College

Kentaro Toyama
University of Michigan

Editorial Scope

Topics

Foundations and Trends[®] in Human-Computer Interaction publishes surveys and tutorials on the foundations of human-computer interaction. The scope is broad. The list of topics below is meant to illustrate some of the coverage, and is not intended to be an exhaustive list.

- History of the research community
- Design and evaluation
- Theory
- Technology
- Computer supported cooperative work
- Interdisciplinary influence
- Advanced topics and trends

Information for Librarians

Foundations and Trends[®] in Human-Computer Interaction, 2014, Volume 8, 4 issues. ISSN paper version 1551-3955. ISSN online version 1551-3963. Also available as a combined paper and online subscription.

Foundations and Trends[®] in Human-Computer
Interaction

Vol. 8, No. 4 (2014) 273–372

© 2015 E. Brady and J. P. Bigham

DOI: 10.1561/1100000050



Crowdsourcing Accessibility: Human-Powered Access Technologies

Erin Brady

Department of Human-Centered Computing
School of Informatics and Computing
Indiana University – Purdue University Indianapolis
brady@iupui.edu

Jeffrey P. Bigham

Human-Computer Interaction Institute
Carnegie Mellon University
jbigham@cmu.edu

Contents

1	Introduction	2
1.1	Human-Powered Access Technology	4
2	Disability, Accessibility and Access Technologies	7
2.1	Definitions and Models of Disability	7
2.2	A History of Accessibility	9
2.3	Assistive and Access Technologies	10
3	Crowdsourcing and Remote Workers	13
3.1	Composition of the Crowd	14
3.2	Designing Crowd Work	20
3.3	Incentives for Crowd Work	23
4	Examples of Human-Powered Access Technologies	26
5	Design Dimensions for Human-Powered Access Technology	34
5.1	Dimensions of Users	34
5.2	Dimensions of Tasks	38
5.3	Dimensions of Workers	42
5.4	Application of Design Dimensions	46

6 Case Study 1: Making Web Pages Accessible with Social Accessibility	50
6.1 Social Accessibility	51
6.2 Design Dimensions of Social Accessibility	52
6.3 Evaluation of Social Accessibility	54
6.4 Conclusions	55
7 Case Study 2: Answering Visual Questions Quickly with VizWiz	57
7.1 VizWiz	57
7.2 Design Dimensions of VizWiz	59
7.3 Interface and Implementation	60
7.4 Evaluation of VizWiz	63
7.5 Conclusions	72
8 Case Study 3: Combining Inputs for Real-Time Captioning with Scribe	74
8.1 Scribe	75
8.2 Design Dimensions of Scribe	77
8.3 Evaluation of Scribe	78
8.4 Conclusions	79
9 Challenges And Opportunities Going Forward	81
9.1 Challenges	81
9.2 Opportunities	85
10 Conclusions	87
Acknowledgements	89
References	90

Abstract

People with disabilities have always engaged the people around them in order to circumvent inaccessible situations, allowing them to live more independently and get things done in their everyday lives. Increasing connectivity is allowing this approach to be extended to wherever and whenever it is needed. Technology can leverage this human work force to accomplish tasks beyond the capabilities of computers, increasing how accessible the world is for people with disabilities. This article outlines the growth of online human support, outlines a number of projects in this space, and presents a set of challenges and opportunities for this work going forward.

E. Brady and J. P. Bigham. *Crowdsourcing Accessibility: Human-Powered Access Technologies*. Foundations and Trends[®] Human-Computer Interaction, vol. 8, no. 4, pp. 273–372, 2014.

DOI: 10.1561/1100000050.

1

Introduction

People with disabilities have always leveraged the assistance of people around them to help them get things done. Supporters are drawn from family members, friends, employees, or nearby strangers. This can work well when someone is available, but it breaks down if a supporter isn't there when needed. Constant access to support has always been possible for those who could afford it, but the cost may quickly grow prohibitive because the supporter needs to be available at all times even if they are only occasionally needed.

Technology has changed this tradeoff because in the past few decades the people who provide assistance no longer need to be physically proximate for many critical tasks or services. This has opened a rich new era in technology design —creating tools that can incorporate human intelligence when and how it is needed to flexibly support people with disabilities, rather than relying on machine intelligence alone, which may be incomplete or limited to specific situations.

These crowd-powered systems, as they have come to be called [Bernstein, 2012, Lasecki et al., 2014b], address difficult problems, many of which are of particular and practical interest to people with disabilities, by combining human intelligence and computation in new ways. People



Figure 1.1: A diagram of the Faxcess system from the early 1990s for helping blind people interpret visual information. In this example, a blind person has taken a picture of the box for a frozen dinner with a digital scanner, and sent it to a human supporter who is reading the information to him over the phone. It is interesting to note that the supporter herself has a physical disability. While this was not called human computation or crowdsourcing at the time, it presaged many of these ideas.

with disabilities are some of the first true users of this technology, so understanding how they use it and manage their expectations around it can be more broadly informative. For instance, nearly 10,000 blind users have asked VizWiz a question about an image that they took, giving us insights not only into what visual information blind people would like to know about but also how they approach asking workers online for answers about what is around them.

In order to understand how to design these types of technologies, we draw from several successful systems utilizing the paradigm of *human-powered access technologies*. In this article, we first trace how people with disabilities went from being early adopters of technology for re-

mote support from other people to becoming some of the first users of technology powered by the crowd. We discuss methods used not only to obtain high-quality, reliable work with these systems, but also how people with disabilities have actively managed the work and workers to ensure that their expectations were met. We discuss who the people in the crowd are, how they are recruited, what incentives they receive, the ethics of crowdsourcing, and what connections (if any) they have to the people that they support. We explore these ideas through a set of examples of technology in this space, and then describe three case studies of human-powered access technologies, which we analyze through the lens of our design dimensions. The first, Social Accessibility, allows people with disabilities to request accessibility meta-information about webpages, which can then be reused by others visiting the website after them. The second case study describes a technology that we have developed called VizWiz that answers visual questions for blind people in nearly-realtime by sending them to sighted crowd workers. The third case study covers Scribe, a real-time captioning tool that combines novice transcriptions from multiple workers into a coherent stream of captioning. From these prior examples, we extract design dimensions that we believe can help characterize technology in this space, provide insights to designers of new technology so that their work can build on what has come before, and start a framework to allow new technologies to be compared and contrasted to prior work. We then conclude with a list of remaining challenges for work in this area.

1.1 Human-Powered Access Technology

The term human-powered access technology is intended to capture the broad range of technologies supporting access for people with disabilities by leveraging human intelligence, effort, and perception [Bigham et al., 2011]. This term is intended to cover a variety of different kinds of systems powered by humans that support people with disabilities by making something accessible that was not previously. Many of these systems include substantial computational aspects, but we purposefully exclude those technologies that rely only on computation.

Access technologies have often included human intervention to make content accessible to people. Initially, this was due to a lack of existing technologies to solve access problems. For instance, human readers were used by blind people to access text before optical character recognition (OCR) tools and audio books became common. While automated tools now exist to aid users with these access problems, many are in early stages and do not function perfectly. OCR works well for clear text, but does not perform well on handwriting; automatic speech recognition (ASR) similarly works well in ideal conditions if it's been trained for a particular speaker, but is not sufficient for recognizing noisy fragments of speech. Because of these limitations, humans are often used to train, supplement, or replace automated solutions, as many tasks that are hard for a computer to do, *e.g.*, visual tasks, recognizing and understanding human speech, are comparatively easy for people.

Human computation and crowdsourcing are two related terms used to discuss work that is performed by people and mediated by technology [Quinn and Bederson, 2011]. While the two concepts have some overlap, they take advantage of human labor in distinct ways. In instances of human computation, people are used to perform tasks too difficult for a computer to do, often in conjunction with the computer. One well-known example of this is CAPTCHAs, the “Completely Automated Public Turing Test to Tell Computers and Humans Apart” [Von Ahn et al., 2004], which use AI-hard problems that are easy for humans to solve as a way to distinguish between real users and automated bots. Two facets typically define human computation tasks: tasks are structured in a way that, if technology “caught up” and could solve them easily, the human labor could be replaced by computers; and a computational structure exists that distributes the tasks and aggregates the results intelligently [Quinn and Bederson, 2011].

Crowdsourcing is a paradigm of distributing jobs, which traditionally may have been completed by a fixed employee, to remote people online. While similar in nature to human computation, crowdsourcing is not limited to tasks that are hard for computers to do automatically, and in fact can serve as a replacement for technological solutions if they are expensive or unwieldy [Quinn and Bederson, 2011].

Though some tasks may have a single component and need only one worker, many workflows have been developed to allow crowdsourcing to be performed with multiple workers simultaneously or in sequence, allowing them to supplement or exceed a single worker's output [Little et al., 2010]. Crowdsourcing is also related to collective intelligence. Generally, for human-backed access technologies we are interested in “directed crowdsourcing”, in which someone, *e.g.*, the person with a disability, commissions the work to be done.

Human computation lets users access information that may not yet be accessible by technology, while crowdsourcing enables fast access to humans who can provide assistance. This article describes the advent of technologies that leverage the power of humans recruited to improve the support available for people with disabilities. Our focus is on crowdsourcing systems, those that rely on open calls to pools of workers, because this seems to be one of the better ways of recruiting workers improving accessibility support, but we include discussion of systems across the range of human-powered access technology to help fill out the space.

In the next sections, we discuss more about *people with disabilities* —the history of disability, accessibility, and access technologies —and *the crowd* —the people within it, their motivations, and how crowds composed of different people can be used for different purposes.

References

- Judd Antin and Aaron Shaw. Social desirability bias and self-reports of motivation: a study of amazon mechanical turk in the us and india. In *Proceedings of the 30th Annual ACM Conference on Human Factors in Computing Systems*, CHI '12, pages 2925–2934. ACM, 2012.
- Shiri Azenkot, Sanjana Prasain, Alan Borning, Emily Fortuna, Richard E Ladner, and Jacob O Wobbrock. Enhancing independence and safety for blind and deaf-blind public transit riders. In *Proceedings of the 29th Annual ACM Conference on Human Factors in Computing Systems*, CHI '11, pages 3247–3256. ACM, 2011.
- Colin Barnes and Geof Mercer. *Independent futures: Creating user-led disability services in a disabling society*. Policy Press, 2006.
- Michael S. Bernstein. *Crowd-Powered Systems*. PhD thesis, MIT, 2012.
- Michael S. Bernstein, Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackerman, David R. Karger, David Crowell, and Katrina Panovich. Soylent: a word processor with a crowd inside. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology*, UIST '10, pages 313–322. ACM, 2010.
- Michael S. Bernstein, Joel Brandt, Robert C. Miller, and David R. Karger. Crowds in two seconds: enabling realtime crowd-powered interfaces. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, UIST '11, pages 33–42, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0716-1. . URL <http://doi.acm.org/10.1145/2047196.2047201>.

- Jeffrey P. Bigham. Increasing web accessibility by automatically judging alternative text quality. In *Proceedings of the 12th International Conference on Intelligent User Interfaces*, IUI '07, pages 349–352. ACM, 2007.
- Jeffrey P. Bigham. Making the web easier to see with opportunistic accessibility improvement. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, UIST '14, pages 117–122, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-3069-5. URL <http://doi.acm.org/10.1145/2642918.2647357>.
- Jeffrey P. Bigham and Anna C. Cavender. Evaluating existing audio captchas and an interface optimized for non-visual use. In *Proceedings of the 27th Annual ACM Conference on Human Factors in Computing Systems*, CHI '09, pages 1829–1838. ACM, 2009.
- Jeffrey P. Bigham and Richard E. Ladner. Accessmonkey: a collaborative scripting framework for web users and developers. In *Proceedings of the 2007 International Cross-Disciplinary Conference on Web Accessibility*, W4A '07, pages 25–34. ACM, 2007.
- Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C. Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, and Tom Yeh. Vizwiz: nearly real-time answers to visual questions. In *Proceedings of the 23th Annual ACM Symposium on User Interface Software and Technology*, UIST '10, pages 333–342. ACM, 2010.
- Jeffrey P. Bigham, Richard E Ladner, and Yevgen Borodin. The design of human-powered access technology. In *Proceedings of the 13th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '11, pages 3–10. ACM, 2011.
- Jeffrey P. Bigham, Michael S. Bernstein, and Eytan Adar. Human computer interaction and collective intelligence. 2014.
- John Brabyn, William Crandall, and William Gerrey. Remote reading systems for the blind: A potential application of virtual presence. In *Proceedings of the 14th Annual IEEE International Conference of Engineering in Medicine and Biology Society*, volume 4, pages 1538–1539. IEEE, 1992.
- Erin Brady, Meredith Ringel Morris, Yu Zhong, Samuel White, and Jeffrey P. Bigham. Visual challenges in the everyday lives of blind people. In *Proceedings of the 31rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '13, pages 2117–2126. ACM, 2013a.

- Erin Brady, Meredith Ringel Morris, and Jeffrey P. Bigham. Gauging receptiveness to social microvolunteering. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15*, New York, NY, USA, 2015. ACM. URL http://research.microsoft.com/pubs/238319/socialmicrovolunteering_chi2015.pdf.
- Erin L. Brady, Yu Zhong, Meredith Ringel Morris, and Jeffrey P. Bigham. Investigating the appropriateness of social network question asking as a resource for blind users. In *Proceedings of the 16th ACM Conference on Computer Supported Cooperative Work and Social Computing, CSCW '13*, pages 1225–1236, New York, NY, USA, 2013b. ACM. ISBN 978-1-4503-1331-5. . URL <http://doi.acm.org/10.1145/2441776.2441915>.
- Michele A Burton, Erin Brady, Robin Brewer, Callie Neylan, Jeffrey P. Bigham, and Amy Hurst. Crowdsourcing subjective fashion advice using vizwiz: challenges and opportunities. In *Proceedings of the 14th International ACM SIGACCESS Conference on Computers and Accessibility, ASSETS '12*, pages 135–142. ACM, 2012.
- Stefan Carmien, Rogerio DePaula, Andrew Gorman, and Anja Kintsch. Increasing workplace independence for people with cognitive disabilities by leveraging distributed cognition among caregivers and clients. In *Proceedings of the 2003 International ACM SIGGROUP Conference on Supporting Group Work, GROUP '03*, pages 95–104. ACM, 2003.
- Anna Cavender, Shari Trewin, and Vicki Hanson. General writing guidelines for technology and people with disabilities. *ACM SIGACCESS Accessibility and Computing Newsletter*, (92):17–22, 2008.
- Anna C. Cavender, Jeffrey P. Bigham, and Richard E. Ladner. Classinfo-cus: enabling improved visual attention strategies for deaf and hard of hearing students. In *Proceedings of the 11th International ACM SIGACCESS Conference on Computers and Accessibility, ASSETS '09*, pages 67–74, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-558-1. . URL <http://doi.acm.org/10.1145/1639642.1639656>.
- Anna C. Cavender, Daniel S. Otero, Jeffrey P. Bigham, and Richard E Ladner. Asl-stem forum: Enabling sign language to grow through online collaboration. In *Proceedings of the 28th Annual ACM Conference on Human Factors in Computing Systems, CHI '10*, pages 2075–2078. ACM, 2010.
- Dana Chandler and Adam Kapelner. Breaking monotony with meaning: Motivation in crowdsourcing markets. *Journal of Economic Behavior & Organization*, 90:123–133, 2013.
- James I. Charlton. *Nothing about us without us: Disability oppression and empowerment*. University of California Press, 1998.

- Charles L. Chen and T.V. Raman. Axsjax: A talking translation bot using google im: Bringing web-2.0 applications to life. In *Proceedings of the 2008 International Cross-Disciplinary Conference on Web Accessibility*, W4A '08, pages 54–56. ACM, 2008.
- Lydia B. Chilton, John J. Horton, Robert C. Miller, and Shiri Azenkot. Task search in a human computation market. In *Proceedings of the ACM SIGKDD Workshop on Human Computation*, HCOMP '10, pages 1–9, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0222-7. . URL <http://doi.acm.org/10.1145/1837885.1837889>.
- Steven Dow, Anand Kulkarni, Scott Klemmer, and Björn Hartmann. Shepherding the crowd yields better work. In *Proceedings of the 15th ACM Conference on Computer Supported Cooperative Work*, CSCW '12, pages 1013–1022. ACM, 2012.
- Julie S Downs, Mandy B Holbrook, Steve Sheng, and Lorrie Faith Cranor. Are your participants gaming the system?: screening mechanical turk workers. In *Proceedings of the 28th Annual ACM Conference on Human Factors in Computing Systems*, CHI '10, pages 2399–2402. ACM, 2010.
- Jan C Galvin and Marcia J Scherer. *Evaluating, Selecting, and Using Appropriate Assistive Technology*. ERIC, 1996.
- Stephanie Hackett, Bambang Parmanto, and Xiaoming Zeng. Accessibility of internet websites through time. In *Proceedings of the 6th International ACM SIGACCESS Conference on Computers and Accessibility*, number 77-78, pages 32–39. ACM, 2004.
- Awni Y. Hannun, Carl Case, Jared Casper, Bryan C. Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, and Andrew Y. Ng. Deep speech: Scaling up end-to-end speech recognition. *CoRR*, abs/1412.5567, 2014. URL <http://arxiv.org/abs/1412.5567>.
- Jonathan Holman, Jonathan Lazar, Jinjuan Heidi Feng, and John D'Arcy. Developing usable captchas for blind users. In *Proceedings of the 9th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '07, pages 245–246. ACM, 2007.
- Hwajung Hong, Jennifer G. Kim, Gregory D. Abowd, and Rosa I. Arriaga. Designing a social network to support the independence of young adults with autism. In *Proceedings of the 15th ACM Conference on Computer Supported Cooperative Work*, CSCW '12, pages 627–636. ACM, 2012.
- John J. Horton. *Online labor market*. Springer, 2010.

- World Health Organization ICF. Towards a common language for functioning, disability and health, 2002.
- Panagiotis G. Ipeirotis. Mechanical turk: The demographics, 2008. URL <http://www.behind-the-enemy-lines.com/2008/03/mechanical-turk-demographics.html>.
- Panagiotis G. Ipeirotis. The new demographics of mechanical turk, 2010. URL <http://www.behind-the-enemy-lines.com/2010/03/new-demographics-of-mechanical-turk.html>.
- Panagiotis G Ipeirotis, Foster Provost, and Jing Wang. Quality management on amazon mechanical turk. In *Proceedings of the ACM SIGKDD Workshop on Human Computation*, HCOMP '10, pages 64–67. ACM, 2010.
- Lilly C Irani and M Silberman. Turkopticon: Interrupting worker invisibility in amazon mechanical turk. In *Proceedings of the 31st Annual ACM Conference on Human Factors in Computing Systems*, CHI '13, pages 611–620. ACM, 2013.
- Dhruv Jain. Path-guided indoor navigation for the visually impaired using minimal building retrofitting. In *Proceedings of the 16th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '14. ACM, 2014.
- Chandrika Jayant, Matt Renzelmann, Dana Wen, Satria Krisnandi, Richard Ladner, and Dan Comden. Automated tactile graphics translation: in the field. In *Proceedings of the 9th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '05, pages 75–82. ACM, 2007.
- Chandrika Jayant, Hanjie Ji, Samuel White, and Jeffrey P. Bigham. Supporting blind photography. In *Proceedings of the 13th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '11, pages 203–210. ACM, 2011.
- Shaun K. Kane, Chandrika Jayant, Jacob O. Wobbrock, and Richard E. Ladner. Freedom to roam: a study of mobile device adoption and accessibility for people with visual and motor disabilities. In *Proceedings of the 11th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '09, pages 115–122. ACM, 2009.
- Shashank Khanna, Aishwarya Ratan, James Davis, and William Thies. Evaluating and improving the usability of mechanical turk for low-income workers in india. In *Proceedings of the 1st ACM Symposium on Computing for Development*, DEV '10, page 12. ACM, 2010.

- Aniket Kittur, Jeffrey V. Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton. The future of crowd work. In *Proceedings of the 16th ACM Conference on Computer Supported Cooperative Work and Social Computing, CSCW '13*, pages 1301–1318, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1331-5. . URL <http://doi.acm.org/10.1145/2441776.2441923>.
- Richard L. Kline and Ephraim P. Glinert. Improving gui accessibility for people with low vision. In *Proceedings of the 13th Annual ACM Conference on Human Factors in Computing Systems, CHI '95*, pages 114–121, New York, NY, USA, 1995. ACM Press/Addison-Wesley Publishing Co. ISBN 0-201-84705-1. . URL <http://dx.doi.org/10.1145/223904.223919>.
- Raymond Kurzweil. The Ray Kurzweil Reader, 2003. URL <http://www.kurzweilai.net/pdf/RayKurzweilReader.pdf>.
- Richard E. Ladner, Melody Y. Ivory, Rajesh Rao, Sheryl Burgstahler, Dan Comden, Sangyun Hahn, Matthew Renzelmann, Satria Krisnandi, Mahalakshmi Ramasamy, Beverly Slabosky, Andrew Martin, Amelia Lacenski, Stuart Olsen, and Dmitri Groce. Automating tactile graphics translation. In *Proceedings of the 7th International ACM SIGACCESS Conference on Computers and Accessibility, ASSETS '05*, pages 150–157. ACM, 2005.
- Walter S. Lasecki, Kyle I. Murray, Samuel White, Robert C Miller, and Jeffrey P. Bigham. Real-time crowd control of existing interfaces. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology, UIST '11*, pages 23–32. ACM, 2011.
- Walter S. Lasecki, Christopher D. Miller, Adam Sadilek, Andrew Abumoussa, Donato Borrello, Raja Kushalnagar, and Jeffrey P. Bigham. Real-time captioning by groups of non-experts. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology, UIST '12*, pages 23–34. ACM, 2012.
- Walter S. Lasecki, Christopher D. Miller, and Jeffrey P. Bigham. Warping time for more effective real-time crowdsourcing. In *Proceedings of the 31st Annual ACM Conference on Human Factors in Computing Systems, CHI 2013*, pages 2033–2036. ACM, 2013.
- Walter S. Lasecki, Mitchell Gordon, Danai Koutra, Malte F. Jung, Steven P. Dow, and Jeffrey P. Bigham. Glance: Rapidly coding behavioral video with the crowd. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology, UIST '14*, pages 551–562, New York, NY, USA, 2014a. ACM. ISBN 978-1-4503-3069-5. . URL <http://doi.acm.org/10.1145/2642918.2647367>.

- Walter S. Lasecki, Chris Homan, and Jeffrey P. Bigham. Architecting real-time crowd-powered systems. *Human Computation Journal*, September 2014b.
- Walter S. Lasecki, Raja Kushalnagar, and Jeffrey P. Bigham. Helping students keep up with real-time captions by pausing and highlighting. In *Proceedings of the 11th Web for All Conference*, W4A '14, pages 39:1–39:8, New York, NY, USA, 2014c. ACM. ISBN 978-1-4503-2651-3. . URL <http://doi.acm.org/10.1145/2596695.2596701>.
- Walter S. Lasecki, Juho Kim, Nicholas Rafter, Onkur Sen, Jeffrey P. Bigham, and Michael S. Bernstein. Apparition: crowdsourced user interfaces that come to life as you sketch them. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, New York, NY, USA, 2015. ACM. URL <http://www.cs.cmu.edu/~jbigam/pubs/pdfs/2015/apparition.pdf>.
- Jonathan Lazar, Alfreda Dudley-Sponaugle, and Kisha-Dawn Greenidge. Improving web accessibility: a study of webmaster perceptions. *Computers in Human Behavior*, 20(2):269–288, 2004.
- Jonathan Lazar, Aaron Allen, Jason Kleinman, and Chris Malarkey. What frustrates screen reader users on the web: A study of 100 blind users. *International Journal of Human-Computer Interaction*, 22(3):247–269, 2007.
- Kimberly Ling, Gerard Beenen, Pamela Ludford, Xiaoqing Wang, Klarissa Chang, Xin Li, Dan Cosley, Dan Frankowski, Loren Terveen, Al Mamunur Rashid, et al. Using social psychology to motivate contributions to online communities. *Journal of Computer-Mediated Communication*, 10(4):00–00, 2005.
- Greg Little, Lydia B Chilton, Max Goldman, and Robert C Miller. Turkit: human computation algorithms on mechanical turk. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology*, UIST '10, pages 57–66. ACM, 2010.
- Jack M Loomis, Reginald G Golledge, and Roberta L Klatzky. Navigation system for the blind: Auditory display modes and guidance. *Presence: Teleoperators and Virtual Environments*, 7(2):193–203, 1998.
- Jennifer Mankoff, Gillian R. Hayes, and Devva Kasnitz. Disability studies as a source of critical inquiry for the field of assistive technology. In *Proceedings of the 12th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '10, pages 3–10. ACM, 2010.

- Marc Marschark, Greg Leigh, Patricia Sapere, Denis Burnham, Carol Convertino, Michael Stinson, Harry Knoors, Mathijs P.J. Vervloed, and William Noble. Benefits of sign language interpreting and text alternatives for deaf students' classroom learning. *Journal of Deaf Studies and Deaf Education*, 11(4):421–437, 2006.
- David Martin, Benjamin V Hanrahan, Jacki O'Neill, and Neha Gupta. Being a turker. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work and Social Computing*, CSCW '14, pages 224–235. ACM, 2014.
- João Martins, José Carilho, Oliver Schnell, Carlos Duarte, Francisco M Couto, Luís Carriço, and Tiago Guerreiro. Friendsourcing the unmet needs of people with dementia. In *Proceedings of the 11th Web for All Conference*, W4A '14, page 35. ACM, 2014.
- Tara Matthews, Scott Carter, Carol Pai, Janette Fong, and Jennifer Mankoff. Scribe4me: Evaluating a mobile sound transcription tool for the deaf. In *Proceedings of the 8th International Conference of Ubiquitous Computing*, UbiComp '06, pages 159–176. Springer, 2006.
- Shigeki Miyoshi, Hayato Kuroki, Sumihiro Kawano, Mayumi Shirasawa, Yasushi Ishihara, and Masayuki Kobayashi. *Support technique for real-time captionist to use speech recognition software*. Springer, 2008.
- Iftexhar Naim, Daniel Gildea, Walter S. Lasecki, and Jeffrey P. Bigham. Text alignment for real-time crowd captioning. In *Proceedings of the North American Chapter of the Association for Computational Linguistics Conference*, NAACL '13, pages 201–210, 2013.
- Alan F. Newell, Peter Gregor, Maggie Morgan, Graham Pullin, and Catriona Macaulay. User-sensitive inclusive design. *Universal Access in the Information Society*, 10(3):235–243, 2011.
- Oded Nov, Ofer Arazy, and David Anderson. Technology-mediated citizen science participation: A motivational model. In *The 5th International AAAI Conference on Web and Social Media*, ICWSM '11, 2011.
- Abiodun Olalere and Jonathan Lazar. Accessibility of us federal government home pages: Section 508 compliance and site accessibility statements. *Government Information Quarterly*, 28(3):303–309, 2011.
- Ellie Pavlick, Matt Post, Ann Irvine, Dmitry Kachaev, and Chris Callison-Burch. The language demographics of amazon mechanical turk. *Transactions of the Association for Computational Linguistics*, 2:79–92, 2014.
- Mark Pilgrim. *Greasemonkey Hacks: Tips & Tools for Remixing the Web with Firefox*. " O'Reilly Media, Inc. ", 2005.

- Alexander J. Quinn and Benjamin B. Bederson. Human computation: a survey and taxonomy of a growing field. In *Proceedings of the 29th Annual ACM Conference on Human Factors in Computing Systems*, CHI '11, pages 1403–1412. ACM, 2011.
- M Jordan Raddick, Georgia Bracey, Pamela L Gay, Chris J Lintott, Phil Murray, Kevin Schawinski, Alexander S Szalay, and Jan Vandenberg. Galaxy zoo: Exploring the motivations of citizen science volunteers. *Astronomy Education Review*, 9(1):010103, 2010.
- Jakob Rogstadius, Vassilis Kostakos, Aniket Kittur, Boris Smus, Jim Laredo, and Maja Vukovic. An assessment of intrinsic and extrinsic motivation on task performance in crowdsourcing markets. In *The 5th International AAAI Conference on Web and Social Media*, ICWSM '11, 2011.
- Joel Ross, Lilly Irani, M Silberman, Andrew Zaldivar, and Bill Tomlinson. Who are the crowdworkers?: shifting demographics in mechanical turk. In *Proceedings of the 28th Annual ACM Conference on Human Factors in Computing Systems*, CHI '10, pages 2863–2872. ACM, 2010.
- Niloufar Salehi, Lilly C. Irani, Michael S. Bernstein, Ali Alkhatib, Eva Ogbe, Kristy Milland, and Clickhappier. We are dynamo: Overcoming stalling and friction in collective action for crowd workers. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 1621–1630, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3145-6. . URL <http://doi.acm.org/10.1145/2702123.2702508>.
- Daisuke Sato, Hironobu Takagi, Masatomo Kobayashi, Shinya Kawanaka, and Chieko Asakawa. Exploratory analysis of collaborative web accessibility improvement. *ACM Transactions on Accessible Computing (TACCESS)*, 3(2):5, 2010.
- Marcia Joslyn Scherer. *Living in the state of stuck: How technology impacts the lives of persons with disabilities*. Brookline Books, 1996.
- Tom Shakespeare and Nicholas Watson. The social model of disability: an outdated ideology? *Exploring theories and expanding methodologies: Where we are and where we need to go*, (2):9–28, 2001.
- Ben Shneiderman. The limits of speech recognition. *Communications of the ACM*, 43(9):63–65, 2000.
- Solona. What's happened to solona?, a. URL <https://www.freelists.org/post/access-uk/whats-happened-to-solona>.

- Solona. Solona blog, b. URL <http://web.archive.org/web/20091220221455/http://www.solona.net/wordpress/2009/08/important-update-for-solona-users-please-read>. Accessed via the Internet Archive WaybackMachine.
- Michael S. Stinson, Sandy Eisenberg, Christy Horn, Judy Larson, Harry Levitt, and Ross Stuckless. *Real-time speech-to-text services*. Rochester Institute of Technology, National Technical Institute for the Deaf, North-east Technical Assistance Center, 1999.
- Molly Follette Story. Maximizing usability: the principles of universal design. *Assistive Technology*, 10(1):4–12, 1998.
- Hironobu Takagi, Shinya Kawanaka, Masatomo Kobayashi, Takashi Itoh, and Chieko Asakawa. Social accessibility: achieving accessibility through collaborative metadata authoring. In *Proceedings of the 10th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '08, pages 193–200. ACM, 2008.
- Hironobu Takagi, Shinya Kawanaka, Masatomo Kobayashi, Daisuke Sato, and Chieko Asakawa. Collaborative web accessibility improvement: challenges and possibilities. In *Proceedings of the 11th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '09, pages 195–202. ACM, 2009.
- Gregg Vanderheiden. Fundamental principles and priority setting for universal usability. In *Proceedings on the 2000 Conference on Universal Usability*, CUU '00, pages 32–37. ACM, 2000.
- Marynel Vázquez and Aaron Steinfeld. Helping visually impaired users properly aim a camera. In *Proceedings of the 14th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '12, pages 95–102. ACM, 2012.
- Fernanda B. Viegas, Martin Wattenberg, Jesse Kriss, and Frank Van Ham. Talk before you type: coordination in wikipedia. In *Proceedings of the 40th Annual Hawaii International Conference on System Sciences*, HICSS '07, pages 78–78. IEEE, 2007.
- Luis Von Ahn and Laura Dabbish. Labeling images with a computer game. In *Proceedings of the 22nd Annual ACM Conference on Human Factors in Computing Systems*, CHI '04, pages 319–326. ACM, 2004.
- Luis Von Ahn, Manuel Blum, and John Langford. Telling humans and computers apart automatically. *Communications of the ACM*, 47(2):56–60, 2004.

- Luis Von Ahn, Shiry Ginosar, Mihir Kedia, Ruoran Liu, and Manuel Blum. Improving accessibility of the web with a computer game. In *Proceedings of the 24th Annual ACM Conference on Human Factors in Computing Systems*, CHI '06, pages 79–82. ACM, 2006.
- William Yang Wang, Dan Bohus, Ece Kamar, and Eric Horvitz. Crowdsourcing the acquisition of natural language corpora: Methods and observations. In *Proceedings of the Spoken Language Technology Workshop*, SLT '12, pages 73–78, 2012.
- Michele A Williams, Amy Hurst, and Shaun K Kane. Pray before you step out: describing personal and situational blind navigation behaviors. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '13, page 28. ACM, 2013a.
- Michele A Williams, Callie Neylan, and Amy Hurst. Preliminary investigation of the limitations fashion presents to those with vision impairments. *Fashion Practice: The Journal of Design, Creative Process & the Fashion Industry*, 5(1):81–106, 2013b.
- Michele A Williams, Kathryn Ringland, and Amy Hurst. Designing an accessible clothing tag system for people with vision impairments. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '13, page 46. ACM, 2013c.
- Jacob O. Wobbrock, Shaun K. Kane, Krzysztof Z. Gajos, Susumu Harada, and Jon Froehlich. Ability-based design: Concept, principles and examples. *ACM Transactions on Accessible Computing (TACCESS)*, 3(3):9, 2011.
- Haofeng Zhou, Denys Baskov, and Matthew Lease. Crowdsourcing transcription beyond mechanical turk. In *The 1st AAAI Conference on Human Computation and Crowdsourcing*, HCOMP '13, 2013.
- Kathryn Zyskowski, Meredith Ringel Morris, Jeffrey P. Bigham, Mary L Gray, and Shaun Kane. Accessible crowdwork? understanding the value in and challenge of microtask employment for people with disabilities. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work and Social Computing*, CSCW '15, 2015.