

Research Contributions in Human-Computer Interaction

(submitted feature article)

Jacob O. Wobbrock
The Information School | DUB Group
University of Washington
wobbrock@uw.edu

Julie A. Kientz
Human Centered Design & Engineering | DUB Group
University of Washington
jkientz@uw.edu

Introduction

All scholarly fields strive to contribute new knowledge. In the field of human-computer interaction (HCI), this new knowledge increasingly comes in rich forms, like videos and demos, but the archival research paper still remains the most widely used and accepted capture and delivery mechanism for research knowledge. The knowledge contribution made by a research paper—or more precisely, made by the work that a research paper describes—is any research paper’s central feature. For example, a theoretical physics paper may contribute a new mathematical model for the behavior of light near black holes. A civil engineering paper may contribute a new method for stress-testing bridges. A social anthropology paper may contribute an account of people’s reactions to teen pregnancies in rural religious communities. Whatever the field of inquiry, whatever the phenomena of interest, every research paper strives to make a research contribution by offering new knowledge. In an effort to distinguish this kind of knowledge from everyday know-how, some scholars even capitalize the term: Knowledge.

In the whole of human inquiry, there are, of course, countless specific research contributions to be made. But the *types* of these contributions—the general forms this new knowledge takes—are relatively few. The three examples in the opening paragraph each make a different type of contribution: the first is *theoretical*, the second is *methodological*, and the third is *empirical*. These are three different research contribution types—the knowledge they contribute comes to us from three different “ways of knowing” [4].

As is often observed, the field of HCI is highly inter- and multi-disciplinary. It is also young. It has taken a few decades for the types of knowledge in HCI to emerge, converge, and stabilize, and new ways of knowing still swirl about [2]. But over time, an identifiable pattern of research contribution types has evolved. These types were recently refined and put to use in the submission process to the flagship conference in HCI, the ACM CHI conference, for the year 2016.

In this article, we offer an encapsulation of the research contribution types in HCI, endeavoring to provide definition to our field and the types of knowledge it produces. We also describe how different contribution types have been codified in the CHI submission process. For the upcoming CHI 2016 conference, we show how the submitted and accepted papers were distributed across contribution types, providing insight into the kinds of knowledge the HCI field develops and disseminates.

Seven Research Contribution Types in HCI

In 2012, the first author posted an unpublished whitepaper to his academic website describing seven research contribution types in HCI [6]. This whitepaper was shared widely and ultimately became adopted as the basis for the contribution types used in the submission process to CHI 2016. The

whitepaper not only described seven research contribution types, but also described how they are judged and gave specific examples of each. Here we summarize those contribution types, with examples.

1. EMPIRICAL CONTRIBUTIONS

Empirical research contributions are the backbone of science. They provide new knowledge through findings based on observation and data-gathering. Data may be qualitative or quantitative, aspiringly objective or unapologetically subjective, from the laboratory or from the field. In HCI, empirical contributions arise from a variety of sources, including experiments, user tests, field observations, interviews, surveys, focus groups, diaries, ethnographies, sensing, log files, and many others.

Empirical research contributions are evaluated mainly on the importance of their findings and on the soundness of their methods. If empirical findings are uninteresting or unimportant, or if the methods by which those findings arose are sloppy, imprecise, or confounded, then empirical contributions are judged unfavorably.

Examples:

- *Interview study.* Burke, M., Kraut, R. and Williams, D. (2010). Social use of computer-mediated communication by adults on the autism spectrum. Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW '10). New York: ACM Press, 425-434. <http://dx.doi.org/10.1145/1718918.1718991>
- *Diary study.* Czerwinski, M., Horvitz, E. and Wilhite, S. (2004). A diary study of task switching and interruptions. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '04). New York: ACM Press, 175-182. <http://dx.doi.org/10.1145/985692.985715>
- *Quantitative lab experiment.* Lee, S.C. and Zhai, S. (2009). The performance of touch screen soft buttons. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '09). New York: ACM Press, 309-318. <http://dx.doi.org/10.1145/1518701.1518750>
- *Qualitative field study.* Poltrock, S.E. and Grudin, J. (1994). Organizational obstacles to interface design and development: Two participant-observer studies. ACM Transactions on Computer-Human Interaction 1 (1), 52-80. <http://dx.doi.org/10.1145/174630.174633>

2. ARTIFACT CONTRIBUTIONS

HCI is driven by the creation and realization of interactive artifacts. Whereas empirical contributions arise from descriptive discovery-driven activities (science), artifact contributions arise from generative design-driven activities (invention). Artifacts, often prototypes, include new systems, architectures, tools, toolkits, techniques, sketches, mockups, or envisionments that reveal new possibilities, enable new explorations, facilitate new insights, or compel us to consider new possible futures. New knowledge is embedded and manifested in artifacts and the supporting materials that describe them.

Artifact research contributions are evaluated according to the type of artifact that gave rise to them. They are often accompanied by empirical studies but do not have to be, and sometimes should not be [1]. New systems, architectures, tools, and toolkits are evaluated in a holistic fashion according to what they make possible and how they do so [3]. New input and interaction techniques, by contrast, are evaluated precisely and quantitatively so as to isolate their human performance benefits. New design expressions, including sketches, mockups, and envisionments, are evaluated by how insightful, compelling, and richly-painted is their expression. Of particular importance is how well designs embody tradeoffs and hold competing priorities in balance.

Examples:

- *Input device.* Baudisch, P., Sinclair, M. and Wilson, A. (2006). Soap: A pointing device that works in mid-air. Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '06). New York: ACM Press, 43-46. <http://dx.doi.org/10.1145/1166253.1166261>
- *System.* Dixon, M. and Fogarty, J.A. (2010). Prefab: Implementing advanced behaviors using pixel-based reverse engineering of interface structure. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '10). New York: ACM Press, 1525-1534. <http://dx.doi.org/10.1145/1753326.1753554>
- *Hardware toolkit.* Greenberg, S. and Fitchett, C. (2001). Phidgets: Easy development of physical interfaces through physical widgets. Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '01). New York: ACM Press, 209-218. <http://dx.doi.org/10.1145/502348.502388>
- *Input technique.* Grossman, T. and Balakrishnan, R. (2005). The Bubble Cursor: Enhancing target acquisition by dynamic resizing of the cursor's activation area. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '05). New York: ACM Press, 281-290. <http://dx.doi.org/10.1145/1054972.1055012>
- *Envisionment.* Ishii, H. and Ullmer, B. (1997). Tangible bits: Towards seamless interfaces between people, bits and atoms. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '97). New York: ACM Press, 234-241. <http://dx.doi.org/10.1145/258549.258715>

3. METHODOLOGICAL CONTRIBUTIONS

Methodological research contributions create new knowledge that informs how we carry out our work. Such contributions may improve research or practice. They may influence how we do science or how we do design. They may improve how we discover things, measure things, analyze things, create things, or build things.

Methodological research contributions are evaluated on the utility, reproducibility, reliability, and validity of the new method or method enhancement. Useful methods that can be reproduced and provide reliable, valid results are judged favorably. Knowing that a method has these properties requires repeated validation.

Examples:

- *Method application.* Consolvo, S. and Walker, M. (2003). Using the experience sampling method to evaluate ubicomp applications. IEEE Pervasive Computing 2 (2), 24-31. <http://dx.doi.org/10.1109/MPRV.2003.1203750>
- *Method innovation.* Druin, A. (1999). Cooperative inquiry: Developing new technologies for children with children. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '99). New York: ACM Press, 592-599. <http://dx.doi.org/10.1145/302979.303166>
- *Method adaptation.* Millen, D.R. (2000). Rapid ethnography: Time deepening strategies for HCI field research. Proceedings of the ACM Conference on Designing Interactive Systems (DIS '00). New York: ACM Press, 280-286. <http://dx.doi.org/10.1145/347642.347763>
- *New measures.* Soukoreff, R.W. and MacKenzie, I.S. (2003). Metrics for text entry research: An evaluation of MSD and KSPC, and a new unified error metric. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '03). New York: ACM Press, 113-120. <http://dx.doi.org/10.1145/642611.642632>

4. THEORETICAL CONTRIBUTIONS

Theoretical research contributions consist of new or improved concepts, definitions, models, principles, or frameworks. They are vehicles for thought. Whereas methodological contributions inform how we do things, theoretical contributions inform what we do, why we do it, and what we expect from it. Theories may be qualitative or quantitative. They may have both descriptive and predictive power—by distilling the essential features of a phenomenon, they are also able to suggest

how that phenomenon will behave. Fully developed theories offer explanatory accounts, not simply observing *that* but explaining *why*. Theories should be testable and falsifiable; if they are not, they do not qualify as scientific theories [5].

Theoretical research contributions are evaluated based on their novelty, soundness, and power to describe, predict, and explain. A theory that accounts well for observed data from a specific situation but has no ability to generalize to new situations is of limited use. Conversely, a theory that is so broad it can account for anything probably does not contain any descriptive power. Theory validation is almost always achieved through empirical work.

Examples:

- *Thought framework*. Bellotti, V., Back, M., Edwards, W.K., Grinter, R.E., Henderson, A. and Lopes, C. (2002). Making sense of sensing systems: Five questions for designers and researchers. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '02). New York: ACM Press, 415-422. <http://dx.doi.org/10.1145/503376.503450>
- *Design space*. Card, S.K., Mackinlay, J.D. and Robertson, G.G. (1990). The design space of input devices. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '90). New York: ACM Press, 117-124. <http://dx.doi.org/10.1145/97243.97263>
- *Design criteria*. Nathan, L.P., Friedman, B., Klasnja, P., Kane, S.K. and Miller, J.K. (2008). Envisioning systemic effects on persons and society throughout interactive system design. Proceedings of the ACM Conference on Designing Interactive Systems (DIS '08). New York: ACM Press, 1-10. <http://dx.doi.org/10.1145/1394445.1394446>
- *Quantitative model*. Wobbrock, J.O., Cutrell, E., Harada, S. and MacKenzie, I.S. (2008). An error model for pointing based on Fitts' law. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '08). New York: ACM Press, 1613-1622. <http://dx.doi.org/10.1145/1357054.1357306>

5. DATASET CONTRIBUTIONS

A dataset contribution provides a new and useful corpus, often accompanied by an analysis of its characteristics, for the benefit of the research community. Benchmark tests may accompany datasets to standardize comparisons. Datasets enable evaluations of shared repositories by new algorithms, systems, or methods.

Dataset research contributions are judged favorably the extent to which they supply the research community with a representative corpus against which to test and measure. Often, datasets are published with new tools that enable researchers to work with the new corpus. Dataset contributions often accompany methodological contributions where a new method is meant to exercise the dataset.

Examples:

- *Test corpus*. MacKenzie, I.S. and Soukoreff, R.W. (2003). Phrase sets for evaluating text entry techniques. Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI '03). New York: ACM Press, 754-755. <http://dx.doi.org/10.1145/765891.765971>
- *Benchmark tasks*. Myers, B. et al. (1997). Using benchmarks to teach and evaluate user interface tools. Unpublished. Available at <http://www.cs.cmu.edu/~amulet/papers/benchmarks.pdf>
- *Corpus procedure*. Paek, T. and Hsu, B.-J.P. (2011). Sampling representative phrase sets for text entry experiments: A procedure and public resource. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '11). New York: ACM Press, 2477-2480. <http://dx.doi.org/10.1145/1978942.1979304>

- *Repository*. Plaisant, C., Fekete, J.-D. and Grinstein, G. (2008). Promoting insight-based evaluation of visualizations: From contest to benchmark repository. *IEEE Transactions on Visualization and Computer Graphics* 14 (1), 120-134. <http://dx.doi.org/10.1109/TVCG.2007.70412>

6. SURVEY CONTRIBUTIONS

Survey research contributions and other meta-analyses review and synthesize work done on a research topic with the goal of exposing trends and gaps. Survey contributions are appropriate after a topic has reached a certain level of maturity. It is not uncommon for surveys to have references numbering in the hundreds. The journal *ACM Computing Surveys* is solely devoted to publishing survey contributions. In HCI, the journal *Foundations and Trends in HCI* regularly publishes survey contributions.

Survey research contributions, and meta-analyses in general, are evaluated based on how well they organize what is currently known about a topic and reveal opportunities for further research. To be effective, survey contributions must not be mere laundry lists of prior work. Good surveys exhibit completeness, depth, maturity, and good organization.

Examples:

- *Technique survey*. Balakrishnan, R. (2004). "Beating" Fitts' law: Virtual enhancements for pointing facilitation. *International Journal of Human-Computer Studies* 61 (6), 857-874. <http://dx.doi.org/10.1016/j.ijhcs.2004.09.002>
- *Tool survey*. Johnson, G., Gross, M.D., Hong, J. and Do, E.Y.-L. (2009). Computational support for sketching in design: A review. *Foundations and Trends in Human-Computer Interaction* 2 (1), 1-93. <http://dx.doi.org/10.1561/1100000013>
- *Domain survey*. MacKenzie, I.S. and Soukoreff, R.W. (2002). Text entry for mobile computing: Models and methods, theory and practice. *Human-Computer Interaction* 17 (2-3), 147-198. <http://dx.doi.org/10.1080/07370024.2002.9667313>
- *Technology survey*. Shaer, O. and Hornecker, E. (2009). Tangible user interfaces: Past, present and future directions. *Foundations and Trends in Human-Computer Interaction* 3 (1-2), 1-137. <http://dx.doi.org/10.1561/1100000026>

7. OPINION CONTRIBUTIONS

Opinion research contributions seek to change the minds of readers through persuasion. Although the term "opinion" might suggest a less-than-scientific effort, in fact, opinion contributions draw upon many of the above contribution types to make their case. Opinion contributions are considered a separate research contribution type not because they lack a research basis, but because their goal is to persuade, not just inform. Along with persuasion, the goal of opinion contributions is to compel reflection, discussion, and debate.

Opinion research contributions are evaluated on the strength of their argument. Strong arguments credibly use supporting evidence and fairly consider opposing perspectives. They focus on topics of wide interest, and should be broadly accessible.

Examples:

- *Evaluation*. Bernstein, M.S., Ackerman, M.S., Chi, E.H. and Miller, R.C. (2011). The trouble with social computing systems research. *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI'11)*. New York: ACM Press, 389-398. <http://dx.doi.org/10.1145/1979742.1979618>

- *Prioritization*. Dourish, P. (2006). Implications for design. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '06). Montréal, Québec (April 22-27, 2006). New York: ACM Press, 541-550. <http://dx.doi.org/10.1145/1124772.1124855>
- *Application*. Greenberg, S. and Buxton, B. (2008). Usability evaluation considered harmful (some of the time). Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '08). New York: ACM Press, 111-120. <http://dx.doi.org/10.1145/1357054.1357074>
- *Vision*. Newell, A. and Card, S.K. (1985). The prospects for psychological science in human-computer interaction. Human-Computer Interaction 1 (3), 209-242. http://dx.doi.org/10.1207/s15327051hci0103_1

Contribution Types at the ACM CHI Conference

The ACM CHI conference has increasingly come to consider research contribution types as part of its paper submissions. From 1996 – 2001, authors were asked to indicate a “paper type,” although it was never referred to as a “contribution.” Examples from 1996 included “Experience papers” and “Systems papers.” The year 1998 was the first that the call-for-papers (CfP) included the word “contribution.” By 2000, lengthy contribution descriptors appeared, like “a thought-provoking, well substantiated analysis of an HCI-related issue,” or “experience gained in adapting designs and applying other HCI contributions to real-world conditions, presented in the form of a design briefing or case history.” From 2002 – 2008, authors were asked to submit a 30-word freeform “contributions and benefits statement.” From 2009 – 2013, authors were asked to choose from a list of contribution types, but from 2014 – 2015, contribution types were listed in the CfP but not selectable when submitting. In 2016, Wobbrock’s whitepaper contribution types [6] were adopted, modified, and made optionally selectable as part of the keywords list upon submission. Aside from a few terminology changes, an important change was that the empirical contribution type was split in two, one for studies of systems and one for studies of people. Tables 1, 2, and 3 show how research contribution types were codified over the last decade of CHI.

CHI 2006 – 2008 (“Lengthy descriptors”)
Case studies of the development and use of interactive systems that provide new insights for HCI research [2007, 2008: or practice].
Interaction technologies—A new technique or device [2006: or other component of the user interface].
Interactive systems—Descriptions of the architecture, interface, and evaluation of a new interactive system.
Methods [2006: Methodologies] and tools—New methods, processes, techniques, and tools for use in interactive system design, development, and deployment.
Metrics by which the performance of interactive technologies can be measured and compared more meaningfully.
Reflective Analyses—Thought-provoking well substantiated analyses of HCI issues.
Reports of Fieldwork and Ethnography—Findings, guidelines, and so on, from studies of real world settings, or of technology use in such settings, with clear relevance to the design and deployment of interactive systems.
Laboratory Studies Reports—Tests of theory, explorations of new phenomena using well-designed controlled studies of human-computer interactions. [2006: Results from Laboratory Studies—Findings, techniques, methods, and so on, from controlled studies of systems, techniques, and other phenomena relevant to HCI.]
Theories and Models—Presentation and critical analysis of HCI theories, including but not limited to formal approaches. [2006: Theories and Models—Descriptions and evaluation of HCI theories, models, and other formal approaches.]

Table 1. Contribution types for CHI 2006 – 2008. Authors also submitted a 30-word contributions and benefits statement.

CHI 2009 – 2015 (“Quasi contribution types”)
Understanding users
Development or refinement of interface artifacts or techniques
Systems, tools, architectures and infrastructure
Methodology
Theory
Innovation, creativity and vision
[2009, 2010: Experience]
[2014, 2015: Validation and refutation]
Opinions [2011 – 2015: Argument]

Table 2. Contribution types for CHI 2009 – 2015. Through 2013, authors were asked to choose a type when submitting.

CHI 2016 (“Contribution types”)	Wobbrock (2012)
Empirical study that tells us about how people use a system	Empirical
Empirical study that tells us about people	
Artifact or system	Artifact
Method	Methodological
Theory	Theoretical
Dataset	Dataset
Meta-analysis / Literature survey	Survey
Essay / argument	Opinion

Table 3. Contribution types for CHI 2016.

The earlier CHI contribution types embodied a mix of concerns—not only types of knowledge, but also types of work, substantive concerns, and methodological approaches. Eventually, lengthy descriptors gave way to pithy terms focused on the *types of knowledge created*.

Most recently, the CHI 2016 submission process made it optional for authors to indicate their papers’ research contribution type. Of 2316 submissions, only 130 opted not to indicate any contribution (5.6%); 1481 indicated exactly one contribution (63.9%); 705 indicated more than one (30.4%). On average, a submission had 1.29 contribution types indicated ($SD=0.66$). Figure 1 shows how contribution types were presented to authors during the CHI 2016 submission process.

What is the primary contribution type of this paper?

- | | |
|--|--|
| <input type="checkbox"/> Artifact or System | <input type="checkbox"/> Essay/Argument |
| <input type="checkbox"/> Dataset | <input type="checkbox"/> Meta-Analysis/Literature Survey |
| <input type="checkbox"/> Empirical study that tells us about how people use a system | <input type="checkbox"/> Method |
| <input type="checkbox"/> Empirical study that tells us about people | <input type="checkbox"/> Theory |

Figure 1. Optional checkboxes for the eight CHI 2016 contribution types. Authors could select none, one, or more than one.

So how has the HCI research community engaged with these eight types of knowledge? The distribution of submitted and accepted papers across all eight of CHI 2016’s research contribution types is shown in Figure 2.

CHI 2016 by Contribution Type

(2316 submissions, 546 acceptances, 23.6%)




Figure 2. CHI 2016 submissions and acceptances by contribution type, sorted by descending number of submissions.

Of the 2316 submissions to CHI 2016, nearly half (44.0%) were empirical studies of system use, and almost the same percentage of such papers comprised the final program, with an acceptance rate (23.5%) nearly identical to that of the overall conference (23.6%). Empirical studies of people also dominated submissions (28.4%) and had the best acceptance rate of any contribution type (26.3%). Artifact or system papers were also prevalent in submissions (24.2%) and had an acceptance rate (23.9%) similar to the overall conference.

In contrast, acceptance rates for the less prevalent contribution types—Method, Theory, Essay/Argument, Meta-Analysis/Literature Survey, and Dataset—were lower than the overall conference acceptance rate of 23.6%. Most were around 17%, with Dataset research contributions having the lowest acceptance rate (16.7%) as well as being the least submitted (1.3%). Essay/Argument contributions enjoyed the highest acceptance rate of this bunch (21.6%), yielding 19 accepted papers in the final program (3.5%).

It is clear from Figure 2 that empirical studies and technology artifacts dominated the research activities of the CHI 2016 community, and that other research contribution types were not only less common, but also less likely to be accepted for publication. We stress that the level of research activity should not be confused with the *value* placed on particular contribution types once they are published. For example, Method and Theory contributions may be hard-won publications, but may be highly cited and relied upon for years to come.

Conclusion

As the HCI field matures, it is important to reflect on the knowledge it produces and the forms that knowledge takes. Such reflection can be challenging in a broad and diverse field, and while we should be wary of constraining our imaginations, we should embrace giving definition to the knowledge we produce. Doing so provides a valuable map for navigating the field of HCI, and helps newcomers take their first steps. It is, after all, those newcomers who will uncover new possibilities and push the frontiers of our knowledge even further.

Acknowledgements

We thank James Fogarty for his feedback on early drafts.

References

- [1] Greenberg, S. and Buxton, B. (2008). Usability evaluation considered harmful (some of the time). *Proceedings of CHI 2008*. New York: ACM Press, 111-120.
- [2] Höök, K., Bardzell, J., Bowen, S., Dalsgaard, P., Reeves, S. and Waern, A. (2015). Framing IxD knowledge. *Interactions* 22 (6), 32-36.
- [3] Olsen, D. (2007). Evaluating user interface systems research. *Proceedings of UIST 2007*. New York: ACM Press, 251-258.
- [4] Olson, J.S. and Kellogg, W.A. (2014). *Ways of Knowing in HCI*. New York: Springer.
- [5] Popper, K. (2002). *The Logic of Scientific Discovery (2ed)*. London: Routledge Classics.
- [6] Wobbrock, J.O. (2012). Seven research contributions in HCI. Unpublished. Available at <http://faculty.washington.edu/wobbrock/pubs/Wobbrock-2012.pdf>

About the Authors



Jacob O. Wobbrock is an Associate Professor in the Information School and, by courtesy, in Computer Science & Engineering at the University of Washington. He chairs the Master of Human-Computer Interaction & Design program and is a founding member of the DUB Group.

wobbrock@uw.edu



Julie A. Kientz is an Associate Professor in Human Centered Design & Engineering and, by courtesy, in the Information School and in Computer Science & Engineering at the University of Washington. She directs the Computing for Healthy Living & Learning Lab and is a founding member of the DUB Group. jkientz@uw.edu