HCI Research as Problem-Solving

Antti Oulasvirta

Aalto University, Finland

Kasper Hornbæk

University of Copenhagen, Denmark

ABSTRACT

This essay contributes a meta-scientific account of human-computer interaction (HCI) research as problem-solving. We build on the philosophy of Larry Laudan, who develops *problem* and *solution* as the foundational concepts of science. We argue that most HCI research is about three main types of problem: empirical, conceptual, and constructive. We elaborate upon Laudan's concept of *problem-solving capacity* as a universal criterion for determining the progress of solutions (outcomes): Instead of asking whether research is 'valid' or follows the 'right' approach, it urges us to ask how its solutions advance our capacity to solve important problems in human use of computers. This offers a rich, generative, and 'discipline-free' view of HCI and resolves some existing debates about what HCI is or should be. It may also help unify efforts across nominally disparate traditions in empirical research, theory, design, and engineering.

Author Keywords

Human-computer interaction; Problem-solving; Scientific progress; Research problem; Larry Laudan

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous

INTRODUCTION

The spark for writing this essay comes from feelings of confusion, and even embarrassment, arising in describing our field to students and other researchers. What is human-computer interaction (HCI) as a field? As numerous ideas and disciplines contribute to HCI, its unique character is elusive. Although HCI is in intellectual debt to many other fields, few would agree that it reduces to them. It has its own subject of enquiry, which is *not* part of the natural or social sciences. It does not belong to engineering, computer science, or design either. So what is it?

The essay has a grand ambition: to develop a conceptually coherent account of the '95% of HCI research'. We know of no other paper offering an attempt to address the field as a whole. We are motivated first and foremost by the intel-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CHI'16, May 07-12, 2016, San Jose, CA, USA

© 2016 ACM. ISBN 978-1-4503-3362-7/16/05...\$15.00

DOI: <http://dx.doi.org/10.1145/2858036.2858283>

lectual enigma pertaining to what HCI *is*: There is no accepted account that would tell how HCI's numerous approaches contribute to pursuit of shared objectives. In contrast, HCI has been criticised for lack of 'motor themes, mainstream topics, and schools of thought' [25] and for being fragmented 'across topics, theories, methods, and people' [38]. Consequently, some have called for 'a hard science' [36], others for 'strong concepts' [19] or an 'inter-discipline' [3]. These are serious concerns with serious implications for the field.

Why bother with a meta-scientific paper at a technical conference? Because the stakes are high. Philosophies of science are at worst an impotent topic worthy of hallway conversations. But if the critics are right, our field is seriously crippled, from the project level to the larger arenas of research *Realpolitik*. Lacking a coherent view of what HCI is, and what *good* research in HCI is, how can we communicate results to others, assess research, co-ordinate efforts, or compete? In addition, as we show, philosophical views offer thinking tools that can aid in generating ideas and generally enhance the quality of research.

The contribution here lies in describing HCI as *problem-solving*. An overview is given in Figure 1. The view originates from Larry Laudan's philosophy of science [28]. Laudan describes scientific progress in terms of two foundational concepts: *research problem* and *solution*. Laudan's 'problem' is not what we mean by the term in ordinary language. It is defined via inabilities and absences occurring in descriptions; knowledge; or, as often in HCI, constructive solutions. For example, a research problem may involve lack of understanding of how colour schemes on a web page affect the aesthetic experience of its use. More generally, Laudan's research problem subsumes what we traditionally understand in HCI as a 'design problem' but also problems to do with theory and empirical research.

Most of our argumentation builds on a concept put forth by Laudan that links problems with solutions: *problem-solving capacity*. For Laudan, a solution is something special, too. In the above-mentioned case of aesthetic perception of web pages, possible solutions range from descriptions of self-reports to models of aesthetic impressions. These solutions change the status of the inabilities and absences but in different ways. Laudan qualifies this in terms of *improvements to problem-solving capacity*. This is counter to some traditional notions of progress [28, p. 14]:

In appraising the merits of theories, it is more important to ask whether they constitute adequate solutions to significant problems than it is to ask whether they are 'true',

Figure 1. This paper analyses HCI research as problem-solving. Scientific progress in HCI is defined as improvements in our ability to solve important problems related to human use of computing. Firstly, a subject of enquiry is defined and its improvement potential analysed. Then, a research problem is formulated. The outcome of the research (i.e., the solution) is evaluated for its contribution to problem-solving capacity defined in terms of five criteria.

‘corroborated’, ‘well-confirmed’ or otherwise justifiable within the framework of contemporary epistemology.

With this definition, the benefit of problem-solving is that it allows covering a wider scope of research than previous accounts, which have been restricted to certain disciplines, topics, or approaches (e.g., research-through-design [53], interaction criticism [2], usability science [15], or interaction science [21]). However, because Laudan developed his view with natural and social sciences in mind, he missed design and engineering contributions. Extending Laudan’s typology to propose that research problems in HCI include not only empirical and conceptual but also constructive problems, we present the first typology developed to encompass most recognised research problems in HCI. It is now possible to describe research contributions regardless of the background traditions, paradigms, or methods. The seemingly multi- or, rather, *hyper*-disciplinary field is—in the end—about solving three types of problem. This reduces the number of dimensions dramatically when one is talking about HCI.

Having built the conceptual foundation, we return to answer four fundamental questions: 1) What *is* HCI research, 2) what is *good* HCI research, 3) are we doing a good job as a field, and 4) could we do an even better job?

We aim to show through these discussions that Laudan’s problem-solving view is not just ‘solutionism’. It offers a useful, timeless, and actionable non-disciplinary stance to HCI. Instead of asking whether research subscribes to the ‘right’ approach, a system is ‘novel’, or a theory is ‘true’, one asks how it advances our ability to solve important problems relevant to human use of computers. Are we addressing the right problems? Are we solving them well? The view helps us contribute to some longstanding debates about HCI. Moreover, we show that the view is generative. We provide ideas on how to apply it as a thinking tool. Problem-solving capacity can be analysed for individual papers or even whole sub-topics and the field at large. It al-

so works as a springboard for generating ideas to improve research agendas.

We conclude on a positive note by arguing that HCI is neither unscientific nor non-scientific (as some have claimed [40]) or in deep crisis [25]. Such views do not recognise the kinds of contributions being made. Instead, on many counts, HCI has improved problem-solving capacity in human use of computing remarkably and continues to do so. However, as we show, these contributions tend to focus on empirical and constructive problem types. In a contrast to calls for HCI to be more scientific [21], interdisciplinary [3], hard [36], soft [9], or rigorous [40], the systematic weakness of HCI is, in fact, our inability to produce conceptual contributions (theories, methods, concepts, and principles) that link empirical and constructive research.

THREE TYPES OF RESEARCH PROBLEM IN HCI

Our first point is that the key to understanding HCI as problem-solving is the recognition that its research efforts cluster around a few recurring problem types. We effectively ‘collapse’ the (apparent) multiplicity of research efforts under a few problem types. This not only simplifies HCI but also transcends some biasing presumptions arising from methodology, theory, or discipline. One can now see similarities and differences between, say, an observational study of a novel technology and a rigorous laboratory experiment, without being bound by their traditions.

In this section, we 1) introduce Laudan’s notion of research problem briefly, 2) extend his typology to cover engineering and design contributions to HCI, and 3) argue that contributions in HCI can be classified via this typology.

Laudan originally distinguished only two types of research problem—empirical and conceptual. These are defined in terms of absence or inability to understand or achieve some ends. As we argue below, the two types are applicable also to HCI. However, to not let design ‘off the hook’, HCI should cover engineering and design contributions. This aspect is clear in almost all definitions of HCI as a field, in-

Figure 2. The problem-solving view ‘collapses’ research problems in HCI into three main categories, each with three subtypes.

cluding that of the 1992 ACM Curriculum [18]. We therefore propose adding a *constructive problem* type. An overview is given in Figure 2. This typology is orthogonal to the well-known Pasteur’s Quadrant, which constitutes an attempt to bridge the gap between applied and basic research by suggesting ‘use-inspired basic research’ as an acceptable type. In our view, in HCI, all problems are (somehow) use-inspired and the quadrant offers little insight.

Empirical Problems

The landscape is replete with *empirical problems*, across all HCI venues, from studies of how people use mouseover to embarrassing experiences with technology and effective ways of crowdsourcing contributions. Nevertheless, this is perhaps the most straightforward type to define:

Definition: Empirical research is aimed at creating or elaborating descriptions of real-world phenomena related to human use of computing.

Laudan cites three characteristic *subtypes*:

1. unknown phenomena
2. unknown factors
3. unknown effects

Qualitative research, ethnography in particular, is an approach often followed to shed light on novel *phenomena*. An example is the 1996 TOCHI article ‘A Field Study of Exploratory Learning Strategies’ [41], which reported observations of how users explore software. The constituent *factors* of phenomena, however, can be exposed only after the ‘carrier’, the phenomenon, has been identified. Consider, for example, the paper ‘Distance Matters’ [37]: it catalogues phenomena and factors that affect mediated human-to-human communication. Finally, after identifying factors, one can measure and quantify their *effects* on something of interest. A common example is evaluative studies wherein statistical inference is used to quantify the most potent effects. One could cite fisheye menus here—though there is a great deal of knowledge about the technique and how to implement it, a study that evaluated its usability found no benefits of this technique [20].

Conceptual Problems

Conceptual problems are non-empirical; they involve issues in theory development in the most general sense. They are also what Laudan calls second-order problems: their substance does not pertain to the world directly, unlike empiri-

cal problems. Conceptual problems might involve difficulties in explaining empirical phenomena, nagging issues in models of interaction, or seeming conflicts between certain principles of design. Fitts’ law [45] is perhaps the most well-known example. It is a statistical model connecting aimed-movement performance (speed and accuracy) to two properties of a user interface that designers can affect: distance to and width of selection areas such as buttons. The research problem it solves is how performance in aimed movement is connected to task demands imposed by a UI.

We offer the following, more general definition:

Definition: Work on a conceptual research problem is aimed at explaining previously unconnected phenomena occurring in interaction.

Responses to this type of problem include theories, concepts, methods, principles, and models. Furthermore, Laudan distinguishes among three characteristic subtypes:

1. implausibility
2. inconsistency
3. incompatibility

We discuss each subtype with well-known examples from HCI literature. *Implausibility* means that the phenomenon is unreasonable, improbable, or lacking an explanation. Consider the 1985 paper in *HCI Journal* entitled ‘Direct Manipulation Interfaces’ [22], whose authors sought to explain why GUIs felt more direct and command-language interfaces felt more indirect. *Inconsistency* means that a position is inconsistent with data, with itself, or with some other position. For example, empirical research on privacy in HCI led to an account of privacy as a reciprocal process among two or more parties to communication [11]. This observation countered the then-more-common view that privacy is a state or property attributable to a technological system. Finally, *incompatibility* means that two positions have assumptions that cannot be reconciled. The debate [52] about using throughput (TP) as a metric for pointing performance falls into this category. Two scholars proposed two metrics that entailed partially incompatible interpretations of the concept and guidance on how to analyse data.

Constructive Problems

We extend the typology of problems with a third type:

Definition: Constructive research is aimed at producing understanding about the construction of an interactive artefact for some purpose in human use of computing.

We put emphasis on *understanding*: the objective is not the construction itself but the ideas or principles it manifests. This problem type covers some of the sub-areas of HCI showing the most vitality at conferences, including interactive systems, interactive applications, interface and sensor technology, interaction techniques, input devices, UI design, interaction design, and concept design. Importantly, this problem type cuts across design and engineering, both extensive topics. We further distinguish three subtypes:

1. no known solution
2. partial, ineffective, or inefficient solution
3. insufficient knowledge or resources for implementation or deployment

To provide examples of each, we look at recent research on tangible and surface computing. The CHI '97 paper 'Tangible Bits' [23] contributed both a novel technical concept for interaction and the first ideas of technical solutions to a problem for which there were none at the time of writing (subtype 1). Generally, work innovating novel concepts for interaction falls into this category. The second subtype is more typical of engineering papers but also found in design-driven papers aimed at improving existing interactions. For example, the CHI '01 paper 'SenseTable' [39] offered a new electromagnetic tracking method with better performance than computer-vision-based methods for tangible interaction. Generally, papers presenting improved solutions for some aspect of interaction fall into this category. The third subtype involves inability to implement or deploy, which can be caused by lack of knowledge or resources. For example, a CHI'09 paper described the iterative design and deployment of WeSpace [49], a collaborative multi-surface system.

PROBLEM-SOLVING CAPACITY OF HCI RESEARCH

The three-part typology would be sterile if it did not offer some way to assess also the results of research. That is, it should enable us to answer the question 'what is good HCI research?'. This is where we arrive at Laudan's key concept for understanding solutions: *problem-solving capacity*.

To Laudan, the outcomes, findings, and results of research are *solutions* that solve research problems but with capacities. One can informally think of problem-solving capacity as 'solution strength': A weak paper addresses an insignificant problem and solves it inefficiently, while raising concerns about validity. A strong one offers a generalisable and efficient solution to an important, recurring problem. Because this notion can be used to assess research and generate ideas for its improvement, it elevates the problem-solving view from a descriptive status to prescription.

Five Aspects of Problem-Solving Capacity

Laudan discussed four criteria for problem-solving capacity. To account for concerns about validity and reliability in HCI, we propose adding a fifth, to produce this list:

1. significance
2. effectiveness
3. efficiency
4. transfer
5. confidence

Significance means that a solution addresses a problem that is important to the stakeholders of the research—be they researchers or practitioners or end users. Early research on cognitive models is a well-known example. While it successfully addressed the common types of interaction in the 1980s, it was criticised in the 1990s for being mired in insignificant problems as new interactions and contexts of use

emerged [42]. Laudan did not present a metric for significance, but there are many reference points used in HCI. A significant *intrinsic* problem for researchers, for example, might be the problem of what 'interaction' is. It is, however, quite characteristic of HCI that the discipline's significance is gauged by reference to broader issues in society and industry. Our experience is that such significance is argued for by reference to survey data from users and sometimes in terms of *avoiding* some costs (e.g., accidents) or of profitability. It is common, nonetheless, for numerical estimates not to be available, especially when one is considering novel technology. In such cases, arguments for significance are often speculative and assume the risk that the technology or phenomenon in question may not materialise. This feature distinguishes HCI from some neighbouring fields, such as human factors, that put more emphasis on realism.

Effectiveness means that the solution resolves the essential aspects of the stated problem. A weak contribution omits or misconstrues influential aspects of the problem. Proponents of the distributed cognition view, for example, criticised the cognitive science prevailing at the time for inability to explain how individuals and organisations perceive, attend, or remember in real-world environments. The new theory was an attempt to explain how environmental constraints and resources, together with habits and practices, contribute to faculties that were previously attributed to the mind.

Efficiency refers to the costs of applying a solution relative to the gains achieved. Mathematical models, design heuristics, and reports of errors in usability tests are relatively inexpensive to apply, even if not easy to come by. In contrast, while detailed second-by-second interaction analyses may be comprehensive (and thereby effective), they are often cumbersome to obtain and hard to apply.

Laudan also talks about *transfer*, or how well the solution transfers to neighbouring problems or other instances of the problem. Perhaps the most 'transferable' solutions in HCI are the user-centred design method and usability testing (see Lindgaard [30]). Both are almost universally applicable in design projects, although one may question the other aspects of their problem-solving capacity.

We add the criterion of *confidence*, referring to the probability that the proposed solution holds. As do the other four criteria, this one too cuts across the three problem types. In empirical research, confidence is affected foremost by validity and reliability. Perhaps the wrong statistical test was used, or missing data were ignored. Such flaws increase the risk that the result does not hold beyond the study in question. In theoretical work, omissions, such as forgetting to address certain prominent factors, counter-arguments, or assumptions about unlikely conditions, decrease confidence. In constructive research, confidence is affected by arguments as to how well the solution addresses issues that may work against it. For example, engineers would talk about the 'robustness' of a solution and a designer might argue that a given design is suitable for different contexts.

HCI RESEARCH AS A PROBLEM-SOLVING FIELD

To summarise, problem-solving capacity refers to our ability to solve important research problems effectively, efficiently, and with high confidence in the solutions' validity. With this concept, HCI can now be characterised as a problem-solving field with its own unique 1) subject of enquiry, human use of computing; 2) research problems; 3) types of problem-solving capacity pursued; and 4) achievements in improving problem-solving capacity. We offer a very high-level view of the first three below.

Firstly, building on the three problem types outlined above, we can define a research problem in HCI thus:

Definition: A research problem in HCI is a stated lack of understanding about some phenomenon in human use of computing, or stated inability to construct interactive technology to address that phenomenon for desired ends.

This definition recognises at least 1) designers and innovators and their emphasis on the construction; 2) empirical researchers and their emphasis on methods and reliable knowledge; and 3) scientists with their emphasis on theories, concepts, and models. We have already stated that what makes HCI unique among the fields looking at human use of computing (others include human factors and information systems) is that it places a real and strong emphasis on constructive problems. In this respect, it is closer to design and engineering fields.

This view of HCI as a field differs from what might be formed from reading some traditional philosophies of science. As in Thomas Kuhn's notion of *scientific paradigms*, sometimes HCI is described in terms of parallel paradigms [17]. This, however, may justify fortresses and thereby prevent critical evaluation of work: 'You don't understand where I come from; therefore, you are not competent to critique my results.' Moreover, this does not encourage the various camps to work together to solve problems that, especially in HCI, require contributions from multiple angles. Thinking in terms of problem-solving capacity puts all traditions on equal footing by ignoring them and considering 1) the problems and 2) the problem-solving capacity achieved. Similarly, Karl Popper's *critical rationalism* calls for subjecting theories to rational critique and decisive empirical tests. Similar ideas can be seen in recent debates around HCI—for example, in the discussion surrounding reproducibility or research and the call for more theorising and scientific discovery in HCI [21]. In the problem-solving view, such activities are important but should not preclude others, such as construction or identification of problems.

Mixing Problem Types in Papers

The most immediate observation about HCI papers is that they often involve two problem types. Consider a paper presenting a novel construction. In addition, it may describe evaluative studies whose purpose is to understand interac-

tion allowed by the new solution. In fact, HCI literature displays all possible pair-wise combinations of problem types:

- Empirical–constructive (e.g., an empirical study with implications for design)
- Empirical–conceptual (e.g., an empirical study to validate a theory)
- Constructive–conceptual (e.g., prototypes employed to explore principles of design)
- Constructive–empirical (e.g., a technique with a study that not only evaluates but contributes to understanding of relevant phenomena, factors, and effects)
- Conceptual–empirical (e.g., a theory or model validated or refined on the basis of a decisive test)
- Conceptual–constructive (e.g., novel ways to design for theoretically predicted phenomena in interaction)

High Tolerance for Risk

Our third observation is that HCI is willing to accept relatively high uncertainty and risk to achieve large gains. In principle, the improvement potential represented by a given design idea may be unknown and only estimated retrospectively—for example, after market launch. Many visionary papers in HCI have entailed great leaps of faith. Consider 'Put That There' [6], which presented a large-display system driven by gestures and speech already in the late 1970s, decades before the necessary technical capability emerged.

Society and Industry Shaping 'Significance'

HCI also puts strong emphasis on *practical* improvement potential, responding to everything from the issues of stakeholder groups to world problems such as inequality. External stakeholders' problems are often cited with reference to the *significance* of research problems. Some solution types aimed at addressing practical problems present design principles, models and simulations predictive of real-world problems, policy advice, and methodology advice.

More generally, *practical* problem-solving capacity can be thought of in terms of how and how much better the problem can be solved by the relevant stakeholders of our research, whether end users, practitioners, designers, developers, or policy-makers. What can they now achieve that they were not able to before reading the paper? The five criteria can be brought to bear for practical efforts too. One can think about the significance of a solution for such a group and the significance for them of the problem at hand. Inversely, one can assess the cost for the stakeholder if the problem remains unsolved. Consider research on gaze-based input for special user groups: A strong argument is that lack of such input methods may prevent full participation in work or social life. One could also consider quantitative aspects of importance, such as how many users the improvement may affect and how frequently they encounter a problem. For example, in applications of Fitts' law to keyboard optimisation, it was estimated that the QWERTY layout causes unacceptably many excessive 'travel miles' for the fingers, hampering productivity and causing repetitive-strain injuries.

Multiple Forms of Progress

Research on any two HCI topics may differ greatly in their ‘distribution’ of types of problem-solving capacities. This can confound comparison of outcomes. Are we progressing more on topic A or topic B? It may be easy for an outsider to dismiss a topic simply by using a different standard for problem-solving capacity.

As an analytical exercise, let us consider two topics, with different problem-solving capacities: 1) interruptions and disruptions and 2) interaction techniques.

The research challenges in interruptions research are to define and characterise the typical interruptions in HCI, explain the mechanisms that link them to detrimental effects, and suggest how to improve HCI by means of design. Hence, the topic is a cluster of empirical, conceptual, and constructive problems. When it emerged, in the 1990s, most papers reported effects of interface designs and conditions on interruption costs. Regrettably, these results were elusive: they could disappear upon change in context. In other words, the problem-solving capacity was low, as an HCI researcher would not know which solution might hold in a given context. In the 2000s, papers started to emerge that linked interruptability to known capacity limitations of human memory and cognition. These papers referred to mechanisms such as working memory capacity and threads in cognition. These could better explain the *preconditions* for interruptability and the efficacy of design solutions. This work has culminated in a series of papers demonstrating superior interruption-tolerance for theoretically motivated designs in common use contexts, such as driving or desktop computing. Work emerged, such as a CHI '15 Best Paper we discuss below, that offered resolutions to some outstanding questions distinguishing theories. Thus, from the problem-solving angle, one can say that the research in this area has been successful in increasing problem-solving capacity over two decades of work. It has reached the point where it can propose *theoretically grounded design solutions for real, complex contexts such as driving*. However, on the negative side, the area has not fared as well in the exploration of new design opportunities. The design examples addressed in conceptual papers are close to present-day designs in the commercial market.

Compare this to the sub-field of interaction techniques. It has been, in broad terms, interested in developing techniques that change means of input and output during interaction to enhance user performance and satisfaction. Although it has been successful in innovating a *plethora of new techniques*, some of which are actually in use, some researchers have lamented that the research is driven by ‘point designs’ and that it cannot explain and generalise principles going beyond individual techniques. It is clear that in comparison to interruptions research, work in this area has addressed fewer conceptual problems. However, it is safe to conclude that it has explored a much larger design space.

A SNAPSHOT OF TODAY'S HCI: CHI BEST PAPERS

To delve more deeply into HCI as a problem-solving field, we now look at a sample of recent *top* papers. We do this to provide an overview of where the field is heading currently and to emphasise that we are doing a good job as a field, when examined through the lens of problem-solving. We analysed the 21 Best Papers from the proceedings of the 2015 ACM Conference on Human Factors in Computing Systems (CHI). While Best Papers are a curated, special sample, unrepresentative of HCI at large, they give an idea of where the field may be heading. The CHI Best Papers are the top ‘1% of submissions’, nominated on the basis of review scores and chosen by a committee.

The two co-authors categorised each paper by the most prominent problem type as determined from the key parts of the paper. They read central parts of each paper to understand the problem definition and claimed problem-solving capacity. Their coding of four papers differed in the first round; these disagreements were resolved to arrive at a consensus categorisation.

The observations can be summarised under the following three points, which we expand upon below:

1. Best Papers focus mostly on empirical and constructive types. The conceptual type is under-represented.
2. All five criteria for problem-solving capacity are mentioned as motivations for contributions. However, we could not identify any consensus on the criteria, and each paper followed its own strategy.
3. There is a split between addressing practical and theory-oriented research problems. This contributes to incommensurability of HCI research.

Predominance of Two Problem Types

We could immediately observe that the majority of papers (12) tackle *empirical problems*. Most describe unknown phenomena in human interaction with computers. Vashistha et al. [46] studied how community moderation in a voice forum for rural India affected uptake and use; Buehler et al. [8] surveyed and analysed the assistive technology available on a 3D printing platform. In these and other cases (e.g., the work of Semaan et al. [43]), the stated motivation lies in describing a *significant phenomenon* in HCI. In addition, some empirical papers describe unknown *factors* shaping these phenomena. For instance, Block et al. [4] wrote that ‘we know very little about factors that contribute to collaboration and learning around interactive surfaces’ (p. 867).

Eight of the papers are mainly about *constructive problems*. Some constructive papers are tied to a specific technology. *BaseLase* [35] presents a specific mirror design for interactive laser projection on the floor. Other papers focus more on the interaction-design principles or concepts behind construction [29,31]. *Affordance++* [31] demonstrates a way of using electrical muscle stimulation to guide people in how to use an object. Some constructive papers are explicit about the problems. *ColourID* [14], a tool for improving colour

identification by visually impaired users, is an attempt to rectify the fact that existing tools are ‘often slow to use and imprecise’ (p. 3543). However, some constructive papers just present the solution, leaving the problem implicit (e.g., those by Lopes et al. [31] and Weigel et al. [32]).

Only one Best Paper deals with a *conceptual problem*. Borst and colleagues [7] presented a model of interruptions with the explicit aim of reconciling earlier findings through an integrated theory. Their paper notes the anomaly that earlier results on interruptions need to be reconciled and that ‘to improve our understanding of interruptions, these studies should be integrated into a cognitive theory’ (p. 2971) and existing theories explain only some findings (e.g., ‘other interruption effects cannot be easily explained within memory-for-goals theory’).

The Best Papers often mix problem types, so the discussion above considers the main type of research problem addressed. Eslami et al. [13], for instance, investigated the extent to which users are aware of Facebook’s newsfeed algorithm. To do so, they conducted interviews but also constructed a system that showed users the consequences of the algorithm for their newsfeed. Simm et al. [44] used *Tiree Energy Pulse*, a prototype developed as part of the research, to aid in understanding energy forecasting, thereby combining an empirical problem with a constructive one.

The glaring absence of Best Papers focusing on conceptual problems should cause alarm. Do such papers not get selected as among the best, or is this gap a general one in HCI? We also scanned the 98 papers that received an honourable mention and found only a few among those (e.g., by Egelman and Peer [12] and Vatavu and Wobbrock [47]). It seems that we rarely identify and address anomalies in research.

Problem-Solving Capacities: Diverse and Under-defined

Criteria for justifying a claim of increased problem-solving capacity vary. *Significance* is often clearly spelled out and solid arguments for it presented—for instance, in studies of fatherhood in social media [1], gender bias in image search [24], or perceptions of the newsfeed algorithm on Facebook [13]. *Effectiveness* is often easy to show for constructive papers, wherein the demonstration of a particular UI concept or technical realisation often is sufficient to show effectiveness (this is sometimes called an existence proof [16]). For instance, ‘Velocitap’ [48] shows that it is technically possible to perform sentence-level decoding of touchscreen text entry at interactive rates. Other papers tackling construction instead evaluate the technical performance of the solution or its user experience (there are various examples [29,31,35,50]). For empirical problems, showing effectiveness convincingly is difficult because most pertain to unknown phenomena. Only a few papers (such as Semaan et al.’s [43]) go beyond describing those phenomena and showing how the descriptions may be used. *Efficiency* is not a big topic in Best Papers. Many constructive papers argue instead for a proof of concept. ‘Velocitap’

[48], for instance, states that sentence-based decoding requires immense storage and processing power (a 2 GB language model; recognition on an eight-core server), but this is not a key concern.

For constructive problems, *transfer* is often shown through application examples. For *Acoustruments* [27], a set of passive plastic devices that can extend the sensing capabilities of mobile phones, generalisability is illustrated through nine examples of mechanisms. For *Affordance++* [31], the use of electrical muscle stimulation was applied across several types of task. In the paper on *iSkin* [32] too, various applications are shown. This seems much more difficult for empirical problems, for which arguments about transfer are found in only a few studies. *Confidence* for the solution is addressed in multiple ways, including conducting several studies [7,24,48,50], using large samples [4], performing long-term follow-up [13], and reporting on implementation [13,44]. The sample covers this aspect quite well.

Research and Practical Impacts: Equally Common

It was striking to learn that about half of the CHI Best Papers were written to address practical problems—in particular, constructing interactive technology for real-world use. Many provide design guidelines [51], concepts [29], and ideas for how to improve existing systems [8]. Most of these are about empirical problems; it was rare to see a paper tackling a construction problem that offered outputs that practitioners could take up and use (though exceptions exist [14,50]).

About half of the papers are explicit about increasing the problem-solving capacity of researchers or advancing the state of knowledge. For instance, Menking and Erickson [33] studied the work of women on Wikipedia, noting how ‘Wikipedia’s gender gap may relate to prevailing feeling rules or participation strategies; at the same time this work contributes to advancing Hochschild’s theory of emotions work [...]’ (p. 208). Some papers also describe implications for research methodology [4] and modelling [7,51].

MOVING HCI FORWARD

In light of the discussion so far, we now turn to implications for our field. The question we want to address is ‘can we do an even better job in HCI research?’.

More Work on Actionable Theories of Interaction

Even if HCI puts strong emphasis on construction, its conceptual work has not been powerful enough to drive it. We can carry out ever more sophisticated studies, with larger samples and more complex set-up, *ad infinitum*, but without conceptual contributions that link empirical findings and the design of technology, the results will remain unactionable. From a problem-solving perspective, without conceptual ‘glue’ to anchor them, constructive contributions readily remain point designs and empirical studies point studies. Laudan talks about conceptual problems as second-order problems, in the sense that we are dealing not directly with the world but with descriptions of the world.

Criterion	Evaluation Criteria	Heuristics for Refining Ideas
Significance	Number of stakeholders involved; importance of the improvement for stakeholders; costs incurred when the improvement is <i>not</i> achieved	Target a different stakeholder group or a larger number of stakeholders; aim at a greater improvement over the present baseline; report on direct comparisons against baseline solutions
Effectiveness	Capture the essential aspects of the problem; match between evaluation metrics and priorities	Use multiple evaluation criteria and richer evaluation contexts; validate evaluation criteria; address unnoticed real-world difficulties
Efficiency	How much effort or resources it takes to create or deploy the solution; scalability; size	Develop tools for practitioners; share datasets and code; reduce price/cost
Transfer	Number of users, tasks, and contexts for which the solution can be applied; qualitatively new contexts wherein the solution can be applied	Identify and target new user groups, contexts, or tasks; demonstrate broad-based generalisability
Confidence	Empirical validity; reliability; replicability; reproducibility; robustness	Replicate the result in different contexts; report on different metrics for judging validity and reliability; allow reanalysis

Table 1. Some heuristics for assessing and contributing to evolution of problem-solving capacity in a research project.

Our strong recommendation is to put more effort into *integrative* concepts, theories, methods, and models that can *link empirical and constructive solutions*. This, we believe, is required for the ‘motor themes’ to emerge that are called upon to fill the ‘big hole’ in HCI research [25]. Without such ‘glue’, our research continues to have lower problem-solving capacity than desired. Empirical research should be done in such a way that its hypotheses inform design, and designs should embody and be driven by empirically validated hypotheses. However, while it is fruitful to strive for integrative types of knowledge, it is healthy to remember that work on constructive problems can advance also without any hypothesis. And, vice versa, there are numerous examples of theories that lack direct relevance.

Improvement of Writing Culture

Our writing culture does not support the problem-solving view. The impression from the Best Paper sample was that many papers could do a better job in describing the problem they are tackling. This is essential from the problem-solving angle. Some papers make explicit only the solution (e.g., a new technology) or approach (e.g., what they did), neither of which is about the research problem. These papers only rarely explain how the result would improve our problem-solving capacity and instead just use language such as ‘we know little about’, ‘significant gap in knowledge’, and ‘no researchers have developed systems that’.

Systematic Improvement of Problem-Solving Capacities

Problem-solving is not merely a description. It offers a ‘thinking tool’ for refining research ideas and generating better ones. This sets it apart from some previous attempts to state the qualities desired in HCI research, which have often been normative or silent with regard to idea-generation.

Firstly, to improve an individual research effort, the five criteria for problem-solving capacity can be used ‘prescriptively’ to generate ideas for how to improve. In Table 1, we have listed heuristics to assess and nurture problem-solving capacity for the problem being considered. There is a row for each property of problem-solving capacity, and the columns present related assessment criteria and development strategies. These refer to the definitions and criteria given above. The table can be applied by assessing the solution *obtained* (if the research has ended) or *desired* (for planned research) and considering whether it could be improved fur-

ther. The list is meant not to be complete but to show that metrics and constructive ideas can be generated for each of the aspects. The authors of this paper have used these criteria internally to develop and refine research ideas.

Secondly, problem-solving capacity can be applied to whole sub-topics also, to assess them and see opportunities to improve. Let us discuss Fitts’ law as an example. It is one of the few thoroughly studied models in HCI and addresses a pervasive phenomenon in interaction. Fitts’ law is also reasonably transferable: it has been found to apply across a wide variety of devices and contexts (even underwater) [45]. Thus, from the perspectives of significance and transfer, Fitts’ law has increased our problem-solving capacity. However, it can be criticised from the angles of effectiveness, efficiency, and confidence. First, Fitts’ law does not *completely* solve the problem of aimed-movement performance, because it relies on heavy aggregation of data at the task level. It dismisses cognitive factors (e.g., performance objective) and dynamics of motion (e.g., trajectory, variability, and force used). It is not an efficient solution either, because its free parameters must be calibrated for each task and context. Also, these parameters are fragile. One can go so far as to claim that these shortcomings limit Fitts’ law to *interpolation* within a set of empirical data and it fails to be a truly *predictive* model. To advance problem-solving capacity in this line of research, the effectiveness and efficiency of the modelling approach should be improved.

Thirdly, although the problem-solving view does not encompass a notion of pseudo-science, it can steer the researcher to avoid pathological practices. These are defined by Irving Languir as wishful thinking, fraud, exaggeration of effects, and *ad hoc* excuses. The problem-solving view may aid in avoiding these via three means: 1) by asking researchers to explicate their research problems, as opposed to just presenting results; 2) by providing criteria for outcomes that entail going beyond ‘point designs’, ‘novelty’, and ‘existence proofs’; and 3) by driving researchers to present more solid evidence and thereby increase confidence.

Rethinking What Constitutes ‘Good’ Research

HCI has tended to develop and adopt superficial criteria for evaluating its research and for its goal-setting. Some of these may have been outright damaging. While HCI has been called an interdisciplinary or trans-disciplinary field, a par-

adox exists: these views tend to lock us into the belief systems of the informing disciplines. The problem-solving view is inherently trans-disciplinary or, rather, *non-disciplinary*, because it makes no reference to such demarcations and belief systems. The empirical problem type does not differentiate among psychological, sociological, quantitative, qualitative, and other stances that have divided the ranks in the past. The conceptual problem type does not distinguish between knowledge produced by design probes and that generated via cognitive models. The constructive problem type does not differentiate among problems of ‘design’, ‘engineering’, and ‘computer science’ type. What matters is how problem-solving capacity is improved. In its ignorance of disciplines, problem-solving runs counter to many existing conceptualisations of HCI that address its development and its *modus operandi*. We believe problem-solving offers a fresh counter-argument to some superficial criteria held up in HCI.

For example, some analyses of HCI view its progress as *dialectical*, from thesis to antithesis to synthesis. Consider the ‘three waves’ [5] or ‘four epochs’ of HCI [42] or the various ‘turns’ (e.g., to practice [26] or to the wild [10]). According to these analyses, HCI research started with classical cognitivism but, on account of its limitations, evolved through other stages. The problem-solving view does not clash with an appreciation for dramatic changes in technology, user groups, contexts, and activities; such would shape the landscape of important problems. As new problems emerge, old solutions may lose some of their significance and generalisability. However, it is a categorical mistake to suggest that the old problems have *completely* lost their significance or solutions their capacity. There are many recurring phenomena in interaction that will stay, and have stayed, central to interaction—even with radical evolution of technologies and contexts. Various practical problems have persisted since the early days of HCI—for instance, writing (a document such as this). Sudden ignorance of previous results decreases the problem-solving capacity we possess as a field.

Further damaging dogma has involved design implications, which have at times been taken as a token of our multi-disciplinarity. Whilst Dourish [11] has argued that not all papers must present implications, informal observations suggest that some reviewers and research assessment criteria still emphasise them. From the problem-solving angle, design implications might be one useful way of informing practical efforts in an efficient and transferable manner; they might also be a way of spelling out the practical implications of the solutions a paper describes. However, the drawbacks of stating design implications may overshadow the benefits. Design implications are little more than incoherent lists of if-then rules. Their capacity to change design, which consists of multiple, interrelated decisions, is limited. At the same time, they trivialise empirical findings. So, while they should be allowed, there should be no requirement to present them.

We also see no reason to favour *models* over other forms of theorising. One previously prominent view was that HCI should be model-driven and aim at prediction and control. Key proponents of this argument were Newell and Card [36]; early critics included John Carroll [9]. The problem-solving view both agrees and disagrees with it—*without* committing a logical fallacy. On one hand, models are effective representations of hypotheses, they succinctly explain empirical phenomena, and they allow deriving rich implications for design. They are a powerful form of theorising that does not let design ‘off the hook’. Models may increase conceptual, empirical, and constructive problem-solving capacity. On the other hand, later criticism has been correct in pointing out that the ‘hard science’ stance is too narrow. From the Laudanesque standpoint, it limits the problem-solving capacity of HCI as a field. Models simply cannot address *all* types of research problems we *must* address.

Finally, one damaging criterion in assessment of HCI papers has been novelty. This has a strong role in the call for papers and discussions at programme committee meetings. However, novelty is only *correlated* with problem-solving capacity: A paper can improve that capacity tremendously without being novel. A paper may be novel without increasing problem-solving capacity at all! We admit that novelty is relevant for subtype 1 contributions. The prevalence of this type of contribution tends to reflect our evaluation criteria but also the rapidly evolving technological landscape. From the CHI’15 sample, it is clear that sometimes demonstrations of novel concepts for interaction help us engineer new solutions (as with Affordance++); in some cases, empirical studies are the first to reveal a phenomenon that will inspire later research and inform practice (e.g., gender bias in image search). The issue is confusion of a surface feature with progress. Novelty drives researchers to maximise the number of subtype 1 problems but not the capacity to solve them. One ends up with a strong landscape of subtype 1 but leaving the rest unaddressed. A healthier balance must be found, with the intensification of efforts around problems that are found to be important. ‘Grand challenges’, or commonly agreed upon problems, offer one vehicle for such co-ordinated efforts.

LIMITATIONS AND CRITIQUES OF PROBLEM-SOLVING

Problem-solving offers a comprehensive and actionable account of the majority of HCI research. However, the view comes with its own subscriptions and limitations. To conclude, we review and respond to objections collected from presentations and interactions with other HCI researchers.

– *Problem-solving takes HCI down the road of solutionism, as described and critiqued by Morozov [34].*

We have gone to some length to define problem and solution to avoid confusion with narrower, more pragmatic concepts—for instance, *design problem* and *user problem*. We are not calling for research centred on these. Our view acknowledges the need to study interaction purely because of conceptual interests, to develop designs purely to ex-

plore, and to study empirical phenomena purely to learn. These are all subsumed within the notion of ‘research problem’.

– *It does not establish HCI as a discipline that has clear overlaps with and boundaries to other disciplines.*

We reject the notion that HCI can be defined by enumerating which *other* fields it involves. Rather, it should be delineated by its subject of enquiry, purposes, and characteristics.

– *It ignores the role of art in HCI.*

Some artistic endeavours can be described as problems and solutions. Consider the problem of designing an installation that triggers some experience in visitors. Aesthetic objectives in design may be analysed as constructive problems, too. However, artists may resist this description.

– *Many scientific discoveries and innovations arise from curiosity, not problem-solving.*

Problem-solving does not preclude curiosity. It values research—and even ‘blue sky research’—that *identifies* problems, which is often curiosity-driven. Identifying problems is a precondition for solutions’ appearance later. However, we agree that problem-solving does not *encourage* curiosity in any way, simply because it is silent about those activities that lead researchers to formulate their problems.

– *Certain topics, such as user experience, cannot be described as problems and solutions, because they pertain to unmeasurable and subjective qualities.*

Several good examples in HCI illustrate that subjectivity and non-measurability may slow increases in problem-solving capacity at first but need not prevent improvements. One such example is the notion of cognitive workload, which has now been defined and instrumented to such a level that we routinely use this concept in our research (e.g., the TLX questionnaires). Addressing subjective qualities in computer use is a requirement for any serious theory of HCI.

– *HCI problems tend to be messy. How does problem-solving fare with ill-defined or ‘wicked’ problems?*

Although new topics tend to be vague and poorly understood at first, we disagree with the contention that HCI research should *stay* messy. Better-formulated problems and better solutions improve clarity, too.

– *Many research contributions have been visions, not solutions. Consider Memex and ubiquitous computing.*

We have argued that such visions can be described as conceptual and constructive problems of subtype 1 (i.e., implausibility or no known solution). While problem-solving recognises this contribution type as key, it is true that it says nothing about *how* they should be generated.

– *The view is iffy and leads to a lot of ‘on one hand’ and ‘on the other hand’. Does it allow a stronger stance on HCI?*

Perhaps the weakest aspect of Laudan's work is that without consensus on what is important, we cannot define problem-solving capacity, because we cannot assess ‘significance’. More generally, we must accept that some elements of problem-solving capacity are going to be subjective and debatable. What may be significant and efficient for one stakeholder may be very different for somebody else.

– *An important goal of HCI research is not problem-solving but impact on society and industry.*

Problem-solving capacity can be assessed also for stakeholders outside the research community. The challenge is to *translate their problems into research problems*. It is regrettable that Laudan offers little guidance for this translation process. One might ask, though, *why* universities should attempt to solve industry's problems.

CONCLUDING REMARKS

What we think HCI research is and is not greatly affects our conferences, journals, papers, funding applications, supervision, thesis topics, and careers. This paper has advanced the view that HCI research is about solving problems related to human use of computing. We have shown how the extent to which HCI does this can be used to analyse individual papers as well as entire research programmes. The problem-solving view can also generate ideas for research and provides a fresh view of longstanding debates on what HCI research is. We hope it will generate new debate, too.

The problem-solving view should be judged as any other HCI research contribution: by looking at a) the problem it tackles and b) the increase in problem-solving capacity it offers. We have argued, and given initial evidence, that the problem-solving view aids in addressing such problems in HCI; indeed, it helps us begin tackling some of the grand conceptual problems of current HCI, including what HCI research is (problem-solving), what good HCI research is (solutions that increase problem-solving capacity), and how to move our field forward (bridge the empirical and the constructive). We make no pretence that problem-solving applies to all HCI research or solves *all* problems. But we do believe it provides some great first questions for any paper or research programme in HCI: *Which problems does it tackle, and how does it increase our capacity to solve them?*

ACKNOWLEDGEMENTS

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreements 637991 and 648785).

We thank several colleagues for discussions, including Susanne Bødker, Stuart Reeves, Barry Brown, Antti Salovaara, Giulio Jacucci, Vassilis Kostakos, Pierre Dragicevic, Andrew Howes, and Pertti Saariluoma.

REFERENCES

1. Tawfiq Ammari and Sarita Schoenebeck. 2015. Understanding and Supporting Fathers and Fatherhood on Social Media Sites. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 1905–1914.
2. Jeffrey Bardzell. 2011. Interaction criticism: An introduction to the practice. *Interacting with Computers* 23, 6: 604–621. <http://doi.org/10.1016/j.intcom.2011.07.001>
3. Alan F. Blackwell. 2015. HCI As an Inter-Discipline. *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, ACM, 503–516. <http://doi.org/10.1145/2702613.2732505>
4. Florian Block, James Hammerman, Michael Horn, et al. 2015. Fluid Grouping: Quantifying Group Engagement Around Interactive Tabletop Exhibits in the Wild. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 867–876. <http://doi.org/10.1145/2702123.2702231>
5. Susanne Bødker. 2006. When Second Wave HCI Meets Third Wave Challenges. *Proceedings of the 4th Nordic Conference on Human-computer Interaction: Changing Roles*, ACM, 1–8. <http://doi.org/10.1145/1182475.1182476>
6. Richard A Bolt. 1980. “Put-that-there”: Voice and gesture at the graphics interface. ACM.
7. Jelmer P. Borst, Niels A. Taatgen, and Hedderik van Rijn. 2015. What Makes Interruptions Disruptive?: A Process-Model Account of the Effects of the Problem State Bottleneck on Task Interruption and Resumption. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 2971–2980. <http://doi.org/10.1145/2702123.2702156>
8. Erin Buehler, Stacy Branham, Abdullah Ali, et al. 2015. Sharing is Caring: Assistive Technology Designs on Thingiverse. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 525–534. <http://doi.org/10.1145/2702123.2702525>
9. John M Carroll and Robert L Campbell. 1986. Softening up hard science: reply to newell and card. *Human-Computer Interaction* 2, 3: 227–249.
10. Andy Crabtree, Alan Chamberlain, Rebecca E. Grinter, Matt Jones, Tom Rodden, and Yvonne Rogers (eds.). 2013. Introduction to the Special Issue of “The Turn to The Wild.” *ACM Trans. Comput.-Hum. Interact.* 20, 3: 13:1–13:4. <http://doi.org/10.1145/2491500.2491501>
11. Paul Dourish. 2006. Implications for design. *Proceedings of the SIGCHI conference on Human Factors in computing systems*, ACM, 541–550. Retrieved August 21, 2015 from <http://dl.acm.org/citation.cfm?id=1124855>
12. Serge Egelman and Eyal Peer. 2015. Scaling the Security Wall: Developing a Security Behavior Intentions Scale (SeBIS). *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 2873–2882. <http://doi.org/10.1145/2702123.2702249>
13. Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, et al. 2015. “I Always Assumed That I Wasn’t Really That Close to [Her]”: Reasoning About Invisible Algorithms in News Feeds. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 153–162. <http://doi.org/10.1145/2702123.2702556>
14. David R. Flatla, Alan R. Andrade, Ross D. Teviotdale, Dylan L. Knowles, and Craig Stewart. 2015. ColourID: Improving Colour Identification for People with Impaired Colour Vision. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 3543–3552. <http://doi.org/10.1145/2702123.2702578>
15. Douglas J Gillan and Randolph G Bias. 2001. Usability science. I: foundations. *International Journal of Human-Computer Interaction* 13, 4: 351–372.
16. Saul Greenberg and Bill Buxton. 2008. Usability evaluation considered harmful (some of the time). *Proceedings of the SIGCHI conference on Human factors in computing systems*, 111–120. <http://dl.acm.org/citation.cfm?id=1357074>
17. Steve Harrison, Deborah Tatar, and Phoebe Sengers. 2007. The three paradigms of HCI. *Alt. Chi. Session at the SIGCHI Conference on Human Factors in Computing Systems San Jose, California, USA*, 1–18.
18. Thomas T Hewett, Ronald Baecker, Stuart Card, et al. 1992. *ACM SIGCHI curricula for human-computer interaction*. ACM.
19. Kristina Höök and Jonas Löwgren. 2012. Strong concepts: Intermediate-level knowledge in interaction design research. *ACM Transactions on Computer-Human Interaction (TOCHI)* 19, 3: 23.
20. Kasper Hornbæk and Morten Hertzum. 2007. Untangling the usability of fisheye menus. *ACM Transactions on Computer-Human Interaction (TOCHI)* 14, 2: 6.
21. Andrew Howes, Benjamin R Cowan, Christian P Janssen, et al. 2014. Interaction science SIG: overcoming challenges. *Proceedings of the extended abstracts of the 32nd annual ACM conference on Human factors in computing systems*, ACM, 1127–1130.
22. Edwin L. Hutchins, James D. Hollan, and Donald A. Norman. 1985. Direct manipulation interfaces. *Human-Computer Interaction* 1, 4: 311–338.
23. Hiroshi Ishii and Brygg Ullmer. 1997. Tangible bits: towards seamless interfaces between people, bits and atoms. *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*, ACM, 234–241.
24. Matthew Kay, Cynthia Matuszek, and Sean A. Munson. 2015. Unequal Representation and Gender Stereotypes in Image Search Results for Occupations. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 3819–3828. <http://doi.org/10.1145/2702123.2702520>
25. Vassilis Kostakos. 2015. The Big Hole in HCI Research. *interactions* 22, 2: 48–51. <http://doi.org/10.1145/2729103>
26. Kari Kuutti and Liam J. Bannon. 2014. The Turn to Practice in HCI: Towards a Research Agenda. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 3543–3552. <http://doi.org/10.1145/2556288.2557111>
27. Gierad Laput, Eric Brockmeyer, Scott E. Hudson, and Chris Harrison. 2015. Acoustruments: Passive, Acoustically-Driven, Interactive Controls for Handheld Devices. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 2161–2170. <http://doi.org/10.1145/2702123.2702414>

28. Larry Laudan. 1978. *Progress and its problems: Towards a theory of scientific growth*. Univ of California Press.
29. Moon-Hwan Lee, Seijin Cha, and Tek-Jin Nam. 2015. Patina Engraver: Visualizing Activity Logs As Patina in Fashionable Trackers. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 1173–1182. <http://doi.org/10.1145/2702123.2702213>
30. Gitte Lindgaard. 2014. The usefulness of traditional usability evaluation methods. *Interactions* 21, 6: 80–82.
31. Pedro Lopes, Patrik Jonell, and Patrick Baudisch. 2015. Affordance++: Allowing Objects to Communicate Dynamic Use. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 2515–2524. <http://doi.org/10.1145/2702123.2702128>
32. Martin Weigel, Tong Lu, Gilles Bailly, Antti Oulasvirta, Carmel Majidi, and Jürgen Steimle. 2015. iSkin: Flexible, Stretchable and Visually Customizable On-Body Touch Sensors for Mobile Computing. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*.
33. Amanda Menking and Ingrid Erickson. 2015. The Heart Work of Wikipedia: Gendered, Emotional Labor in the World's Largest Online Encyclopedia. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 207–210. <http://doi.org/10.1145/2702123.2702514>
34. Evgeny Morozov. 2014. *To save everything, click here: The folly of technological solutionism*. PublicAffairs.
35. Jörg Müller, Dieter Eberle, and Constantin Schmidt. 2015. BaseLase: An Interactive Focus+Context Laser Floor. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 3869–3878. <http://doi.org/10.1145/2702123.2702246>
36. Allen Newell and Stuart K. Card. 1985. The prospects for psychological science in human-computer interaction. *Human-computer interaction* 1, 3: 209–242.
37. Gary M. Olson and Judith S. Olson. 2000. Distance Matters. *Hum.-Comput. Interact.* 15, 2: 139–178. http://doi.org/10.1207/S15327051HCI1523_4
38. Gary M. Olson and Judith S. Olson. 2003. Human-computer interaction: Psychological aspects of the human use of computing. *Annual review of psychology* 54, 1: 491–516.
39. James Patten, Hiroshi Ishii, Jim Hines, and Gian Pangaro. 2001. Sensetable: a wireless object tracking platform for tangible user interfaces. *Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM, 253–260.
40. Stuart Reeves. 2015. Human-computer interaction as science.
41. John Rieman. 1996. A field study of exploratory learning strategies. *ACM Transactions on Computer-Human Interaction (TOCHI)* 3, 3: 189–218.
42. Yvonne Rogers. 2012. HCI theory: classical, modern, and contemporary. *Synthesis Lectures on Human-Centered Informatics* 5, 2: 1–129.
43. Bryan Semaan, Heather Faucett, Scott P. Robertson, Misa Maruyama, and Sara Douglas. 2015. Designing Political Deliberation Environments to Support Interactions in the Public Sphere. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 3167–3176. <http://doi.org/10.1145/2702123.2702403>
44. Will Simm, Maria Angela Ferrario, Adrian Friday, et al. 2015. Tiree Energy Pulse: Exploring Renewable Energy Forecasts on the Edge of the Grid. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 1965–1974. <http://doi.org/10.1145/2702123.2702285>
45. R William Soukoreff and I Scott MacKenzie. 2004. Towards a standard for pointing device evaluation, perspectives on 27 years of Fitts' law research in HCI. *International journal of human-computer studies* 61, 6: 751–789.
46. Aditya Vashistha, Edward Cutrell, Gaetano Borriello, and William Thies. 2015. Sangeet Swara: A Community-Moderated Voice Forum in Rural India. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 417–426. <http://doi.org/10.1145/2702123.2702191>
47. Radu-Daniel Vatavu and Jacob O. Wobbrock. 2015. Formalizing Agreement Analysis for Elicitation Studies: New Measures, Significance Test, and Toolkit. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 1325–1334. <http://doi.org/10.1145/2702123.2702223>
48. Keith Vertanen, Haythem Memmi, Justin Emge, Shyam Reyall, and Per Ola Kristensson. 2015. VelociTap: Investigating Fast Mobile Text Entry Using Sentence-Based Decoding of Touchscreen Keyboard Input. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 659–668. <http://doi.org/10.1145/2702123.2702135>
49. Daniel Wigdor, Hao Jiang, Clifton Forlines, Michelle Borkin, and Chia Shen. 2009. WeSpace: the design development and deployment of a walk-up and share multi-surface visual collaboration system. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 1237–1246.
50. Wesley Willett, Bernhard Jenny, Tobias Isenberg, and Pierre Dragicevic. 2015. Lightweight Relief Shearing for Enhanced Terrain Perception on Interactive Maps. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 3563–3572. <http://doi.org/10.1145/2702123.2702172>
51. Pamela Wisniewski, Haiyan Jia, Na Wang, et al. 2015. Resilience Mitigates the Negative Effects of Adolescent Internet Addiction and Online Risk Exposure. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 4029–4038. <http://doi.org/10.1145/2702123.2702240>
52. Shumin Zhai. 2004. Characterizing computer input with Fitts' law parameters—the information and non-information aspects of pointing. *International Journal of Human-Computer Studies* 61, 6: 791–809.
53. John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through design as a method for interaction design research in HCI. *Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM, 493–502. <http://dl.acm.org/citation.cfm?id=1240704>