Research in the Wild

Synthesis Lectures on Human-Centered Informatics

Editor
John M. Carroll, *Penn State University*

Human-Centered Informatics (HCI) is the intersection of the cultural, the social, the cognitive, and the aesthetic with computing and information technology. It encompasses a huge range of issues, theories, technologies, designs, tools, environments, and human experiences in knowledge work, recreation and leisure activity, teaching and learning, and the potpourri of everyday life. The series publishes state-of-the-art syntheses, case studies, and tutorials in key areas. It shares the focus of leading international conferences in HCI.

Research in the Wild
Yvonne Rogers and Paul Marshall
April 2017

Designing for Gesture and Tangible Interaction
Mary Lou Maher and Lina Lee
March 2017

From Tool to Partner: The Evolution of Human-Computer Interaction
Jonathan Grudin
December 2016

Qualitative HCI Research: Going behind the Scenes
Ann Blandford, Dominic Furniss, and Stephann Makri
April 2016

Learner-Centered Design of Computing Education: Research on Computing for Everyone
Mark Guzdial
November 2015

The Envisionment and Discovery Collaboratory (EDC): Explorations in Human-Centered Informatics with Tabletop Computing Environments
Ernesto G. Arias, Hal Eden, and Gerhard Fischer
October 2015

Humanistic HCI
Jeffrey Bardzell and Shaowen Bardzell
September 2015

The Paradigm Shift to Multimodality in Contemporary Computer Interfaces
Sharon Oviatt and Philip R. Cohen
April 2015

Multitasking in the Digital Age
Gloria Mark
April 2015

The Design of Implicit Interactions
Wendy Ju
March 2015

Core-Task Design: A Practice-Theory Approach to Human Factors
Leena Norros, Paula Savioja, and Hanna Koskinen
March 2015

An Anthropology of Services: Toward a Practice Approach to Designing Services
Jeanette Blomberg and Chuck Darrah
February 2015

Proxemic Interactions: From Theory to Practice
Nicolai Marquardt and Saul Greenberg
February 2015

Contextual Design: Evolved
Karen Holtzblatt and Hugh Beyer
October 2014

Constructing Knowledge Art: An Experiential Perspective on Crafting Participatory Representations
Al Selvin and Simon Buckingham Shum
October 2014

Spaces of Interaction, Places for Experience
David Benyon
September 2014

Geographical Design: Spatial Cognition and Geographical Information Science
Stephen C. Hirtle
March 2011

User-Centered Agile Methods
Hugh Beyer
2010

Experience-Centered Design: Designers, Users, and Communities in Dialogue
Peter Wright and John McCarthy
2010

Experience Design: Technology for All the Right Reasons
Marc Hassenzahl
2010

Designing and Evaluating Usable Technology in Industrial Research: Three Case Studies
Clare-Marie Karat and John Karat
2010

Interacting with Information
Ann Blandford and Simon Attfield
2010

Designing for User Engagement: Aesthetic and Attractive User Interfaces
Alistair Sutcliffe
2009

Context-Aware Mobile Computing: Affordances of Space, Social Awareness, and Social Influence
Geri Gay
2009

Studies of Work and the Workplace in HCI: Concepts and Techniques
Graham Button and Wes Sharrock
2009

Semiotic Engineering Methods for Scientific Research in HCI
Clarisse Sieckenius de Souza and Carla Faria Leitão
2009

Common Ground in Electronically Mediated Conversation
Andrew Monk
2008

Research in the Wild
Yvonne Rogers and Paul Marshall
www.morganclaypool.com

ISBN: 9781627056922 print
ISBN: 9781627058780 ebook
DOI 10.2200/S00764ED1V01Y201703HCI037

A Publication in the Morgan & Claypool Publishers series
SYNTHESIS LECTURES ON HUMAN-CENTERED INFORMATICS, #37

Series Editors: John M. Carroll, Penn State University

Series ISSN: 1946-7680 Print 1946-7699 Electronic

Research in the Wild

Yvonne Rogers and Paul Marshall
University College London

SYNTHESIS LECTURES ON HUMAN-CENTERED INFORMATICS #37

MORGAN & CLAYPOOL PUBLISHERS

ABSTRACT

The phrase "in-the-wild" is becoming popular again in the field of human-computer interaction (HCI), describing approaches to HCI research and accounts of user experience phenomena that differ from those derived from other lab-based methods. The phrase first came to the forefront 20–25 years ago when anthropologists Jean Lave (1988), Lucy Suchman (1987), and Ed Hutchins (1995) began writing about cognition being in-the-wild. Today, it is used more broadly to refer to research that seeks to understand new technology interventions in everyday living.

A reason for its resurgence in contemporary HCI is an acknowledgment that so much technology is now embedded and used in our everyday lives. Researchers have begun following suit—decamping from their usability and living labs and moving into the wild; carrying out *in-situ* development and engagement, sampling experiences, and probing people in their homes and on the streets.

The aim of this book is to examine what this new direction entails and what it means for HCI theory, practice, and design. The focus is on the insights, demands and concerns. But how does research in the wild differ from the other applied approaches in interaction design, such as contextual design, action research, or ethnography? What is added by labeling user research as being in-the-wild? One main difference is where the research starts and ends: unlike user-centered, and more specifically, ethnographic approaches which typically begin by observing existing practices and then suggesting general design implications or system requirements, in-the-wild approaches create and evaluate new technologies and experiences *in situ* (Rogers, 2012). Moreover, novel technologies are often developed to augment people, places, and settings, without necessarily designing them for specific user needs. There has also been a shift in design thinking. Instead of developing solutions that fit in with existing practices, researchers are experimenting with new technological possibilities that can change and even disrupt behavior. Opportunities are created, interventions installed, and different ways of behaving are encouraged. A key concern is how people react, change and integrate these in their everyday lives. This book outlines the emergence and development of research in the wild. It is structured around a framework for conceptualizing and bringing together the different strands. It covers approaches, methods, case studies, and outcomes. Finally, it notes that there is more in the wild research in HCI than usability and other kinds of user studies in HCI and what the implications of this are for the field.

KEYWORDS

human-computer interaction, HCI, *in situ* studies, research in the wild

Contents

Acknowledgments

We would like to thank Eva Hornecker, Jesper Kjeldskov, and Erik Stolterman for their insightful comments on an earlier draft. We would also like to thank all our colleagues at Sussex University, Open University, and University College London who collaborated with us in our forays into the wild.

CHAPTER 1

Introduction

1.1 RESEARCH GONE WILD

It is now quite common to see the phrase "in the wild" inserted into the title of a human-computer interaction (HCI) paper. Examples include "Doing innovation in the wild" (Crabtree et al., 2013a), "Being in the thick of in the wild" (Johnson et al., 2012), and "A robot in the wild" (Williams et al., 2014), as well as abbreviated versions such as "Leaving the wild" (Taylor et al., 2013) and "Calls from the wild" (Cappadonna et al., 2016). Besides attracting eyeballs ("the wild" sounds more intriguing than the more prosaic "An *in situ* study of…" or "An Investigation into…") this trend reflects a shift in how research is being carried out in HCI. Increasingly, researchers are going into people's homes, the outdoors, and public places, to study their reactions to, use, and appropriation of a diversity of technologies that researchers have provided them with or placed in that location. Examples include exploring the co-creation of a street graph depicting changes in electricity consumption for a community (Bird and Rogers, 2010), the use of mobile devices for tracking people's health (Consolvo et al., 2008), and exploring how robots can assist the well-being of visitors in hospital wards (e.g., Dahl and Bolous, 2014). In addition, researchers are working and participating more with communities, designing and deploying technologies *in situ* that address the latter's concerns or needs. Theory has also been rethought in terms of how it can inform, extend, or develop accounts of behavior that is *situated* in naturalistic settings and in the context of *socio-technical* practices.

Research in the wild (RITW) is generally considered as an umbrella term to refer to how, what, and where research is conducted in naturalistic settings (Crabtree et al., 2013b). *Its overarching goal is to understanding how technology is and can be used in the everyday/real world, in order to gain new insights about: how to engage people/communities in various activities, how people's lives are impacted by a specific technology, and what people do when encountering a new technology in a given setting.* The output can be used to inform the development of new understandings, theories, or concepts about human behavior in the real world. This includes rethinking cognitive theories, in terms of ecological concepts (e.g., situated memory) and socio-cultural accounts (e.g., the effects of digitalization on society). More specifically, RITW can be concerned with investigating an assumption, such as whether or not a technology intervention can encourage people to change a behavior (e.g., exercising more). It can be operationalized in terms of a research question to be evaluated in the wild, such as: will providing free activity trackers to employees to encourage them to develop new social

practices at work (e.g., buddying up, competing with each other) that will help them to become fitter and healthier? The perspective taken for this kind of RITW is to observe how people react, change and integrate the technology in question into their everyday lives over a period of time.

RITW is broad in its scope. Some have questioned the need for yet another term for what many HCI researchers would claim they have been doing for years. Indeed, applied research has been an integral part of HCI, addressing real-world problems, by conducting field studies, user studies and ethnographies. The outputs of which are intended to inform system design, often through community engagement. So, what is the value of coining another label? We would argue that, first, it is now widely used not just in HCI, but also in a number of other disciplines, including biology and psychology, reflecting a growing trend towards pursuing more research in naturalistic settings. Second, the term is more encompassing, covering a wider range of research compared with other kinds of named methodological approaches, such as Action Research, Participatory Design, or Research Through Design. Initial ethnographic research, followed by designing a new user experience, together with the application and/or development of theory, technology innovation, and an *in situ* evaluation study are often conducted all in one RITW project.

Hence, while the various components involved in RITW are not new, a single project often addresses several of them. Rather than focus on one aspect, e.g., developing a new technology, advancing a new method, testing the effects of a variable or reporting on the findings of a technology intervention—research in the wild typically combines a number of interlinked strands. Technology innovation can initially inspire the design of a new learning activity that in parallel is framed in terms of a particular theory of learning. Together, they inform the design of an *in situ* study and the research questions it will address.

RITW is agnostic about the methods, technologies, or theories it uses. Accordingly, it does not necessarily follow one kind of methodology, where one design phase follows another, but combines different ones to address a problem/concern or opportunity, as deemed fit. Sometimes, theory might be considered central and other times only marginal; sometimes, "off-the-shelf" technology is deployed and evaluated in an *in situ* study. Other times, the design and deployment of a novel device is the focus. In other settings, the focus of a project is how best to work alongside a community so that a democratic design process is followed.

The multiple decisions that have to be made when operationalizing a problem are often the main drivers, shaping how the proposed research will address identified questions, what methods/technologies to use and what can be learned. In summary, RITW is broadly conceived, accommodating a diversity of methodologies, epistemologies and ways of doing research. What is common to all RITW projects is the importance placed on the setting and context, conducting research in the everyday and in naturalistic environments.

1.2 HOW DOES RESEARCH IN THE WILD DIFFER FROM LAB EXPERIMENTS?

A long-standing debate in HCI is concerned with what is lost and gained when moving research out of a controlled lab setting into the wild (Preece et al., 2015). An obvious benefit is more ecological validity—an *in situ* study is likely to reveal more the kinds of problems and behaviors people will have and adopt if they were to use a novel device at home, at work, or elsewhere. A lab study is less likely to show these aspects as participants try to work out what to do in order to complete the tasks set for them, by following instructions given. They may find themselves having to deal with various "demand characteristics"—the cues that make them aware of what the experimenter expects to find, wants to happen or how they are expected to behave. As such, ecological validity of lab studies can be less reliable, as participants perform to conform to the experimenter's expectations.

A downside of evaluating technology in situ, however, is the researcher losing control over how it will be used or interacted with. Tasks can be set in a lab and predictions made to investigate systematically how participants manage to do them, when using a novel device, system, or app. When in the wild, however, participants are typically given a device to use without any set tasks provided. They may be told what it can do and given instructions on how to use it but the purpose of evaluating it in a naturalistic setting is to explore what happens when they try to use it in this context—where there may be other demands and factors at play. However, this can often mean that only a fraction of the full range of functionality, that has been designed as part of the technology, is used or explored, making it difficult for the researchers to see whether what has been designed is useful, usable, or capable of supporting the intended interactions.

To examine how much is lost and gained, Kjeldskov et al. (2004) conducted a comparative study of a mobile system designed for nurses in the lab vs. in the wild. They found that both settings revealed similar kinds of usability problems but that more were discovered in the lab than in the wild study. However, the cost of running a study in the wild was considerably greater than in the lab, leading them to question "Was it worth the hassle?" They suggest that in the wild studies might be better suited for obtaining initial insights for how to design a new system that can then feed into the requirements gathering process, while early usability testing of a prototype system can be done in the confines of the lab. This pragmatic approach to usability testing and requirements gathering makes good sense when considering how best to develop and progress a new system design. In a follow-up survey of research on mobile HCI using lab and in the wild studies, Kjeldskov and Skov (2014) concluded that it is not a matter of one being better than the other but when best to conduct a lab study vs. an in the wild study. Furthermore, they conclude that when researchers go into the wild they should "go all the way" and not settle for some "half-tame" setting. Only by carrying out truly wild studies can researchers experience and understand real-world use.

Findings from other RITW user studies have shown how they can reveal a lot more than identifying usability problems (Hornecker and Nicol, 2012). In particular, they enable researchers to explore how a range of factors can influence user behavior *in situ*—in terms of how people notice, approach, and decide what to do with a technology intervention—either one they are given to try or one they come across—that goes beyond the scope of what is typically able to be observed in a lab-based study. Rogers et al. (2007) found marked differences in usability and usefulness when comparing a mobile device in the wild and in the lab; the mobile device was developed to enable groups of students to carry out environmental science, as part of a long-term project investigating ecological restoration of urban regions. The device provided interactive software that allowed a user to record and look up relevant data, information visualizations, and statistics. The device was intended to replace the existing practice of using a paper-based method of recording measurements of tree growth when in the field. Placing the new mobile device in the palms of students on a cold spring day revealed a whole host of unexpected, context-based usability and user experience problems. Placing the device in the palms of students on a hot summer day revealed a quite different set of unexpected, context-based usability and user experience problems. The device was used quite differently for the different times of year, where foliage and other environmental cues vary and affect the extent to which a tree can be found and identified.

Other studies have also found how people will often approach and use prototypes differently in the wild compared with in a lab setting (e.g., Brown et al., 2011; Peltonen et al., 2008; van der Linden et al., 2011). People are often inventive and creative in what they do when coming across a prototype or system, but also can get frustrated or confused, in ways that are difficult to predict or expect from lab-based studies (Marshall et al., 2011). Van der Linden et al. (2011) also observed different behaviors—not evident from their lab-based studies—when investigating how haptic technology could improve children's learning to play the violin at school. An *in situ* study of their Music-Jacket system showed how real-time vibrotactile feedback was most effective when matched to tasks selected by their teachers to be at the right level of difficulty—rather than what the researchers thought would be right for them. Similarly, Gallacher et al. (2015) discovered quite different findings when they ran the same in the wild study in different places. Based on the differing outcomes from lab studies and in in the wild approaches, Rogers et al. (2013) questioned whether findings from controlled settings can transfer to real-world settings.

In summary, *in situ* studies can provide new ways of thinking about how to scope and conduct research. Compared with running experiments and usability studies, where researchers try to predict in advance performance and the likelihood or kind of usability errors, running *in situ* studies nearly always provide unexpected findings about what humans might or might not do when confronted with a new technology intervention. Even when experiments are run in the wild, non-significant findings can be most informative. Part of the appeal of RITW is uncovering the unexpected rather than confirming what is hoped for or already known.

1.3 A FRAMEWORK FOR HCI RESEARCH IN THE WILD

RITW is eclectic in what it does and what it seeks to understand. Such an unstructured approach to research might seem unwieldy, lacking the rigor and commitment usually associated with a given epistemology. However, this broad church stance does not mean sloppiness or lowering of standards; rather, it can open up new possibilities for conducting far-reaching, impactful, and innovative research. To help frame RITW we have developed a generic framework. Figure 1.1 depicts RITW in terms of four core bases that connect to each other. These are regarded as starting places from which to scope and operationalize the research, in terms of:

1. technology,
2. design,
3. *in situ* studies, and
4. theory.

Each can inform the others to situate, shape, and progress the research. For example, *designing* a new activity (e.g., collaborative learning) can be done by working alongside others (e.g., participatory design), leading to the development of a new *technology*. The findings from an *in situ* study (e.g., how people search for information on the fly using their smartphones) can inform new *theory* (e.g., augmented memory). An existing theory (e.g., attention) can inform the design of a new app intended to be used to measure how people multitask in their everyday lives when using smartphones, tablets, and laptops. The design of a new *technology* (e.g., augmented reality) can be used to enhance a social activity in the wild (e.g., how families learn about the ecology of woodlands together). It should be stressed, however, that the RITW framework is not meant to be prescriptive, in terms of which base to start from, or what methods and analytic lens to use, when conducting research. The selection of these depends on the motivation for the research, its scoping, the available funding and resources, and expected outcomes.

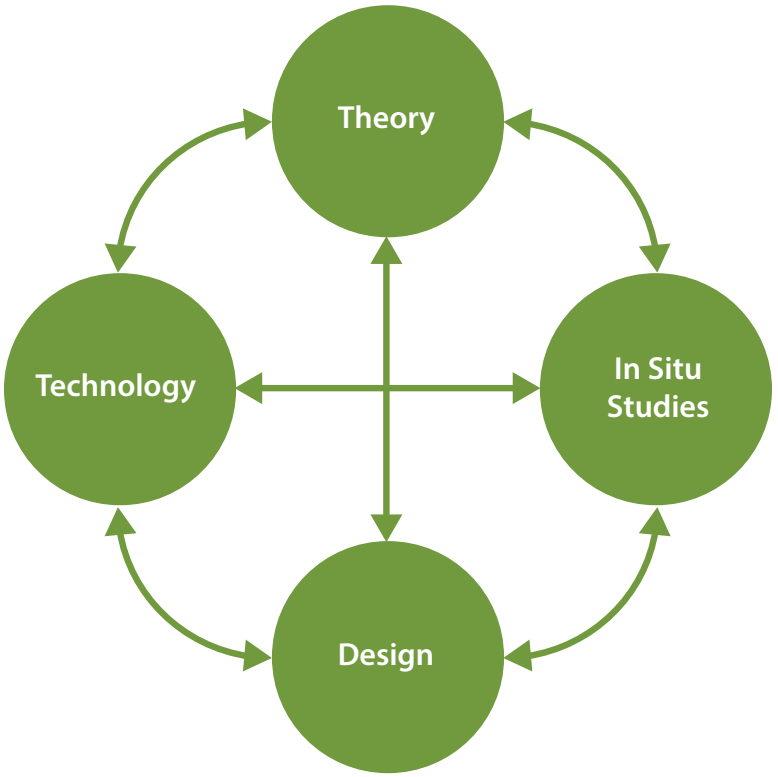

Technology: Concerned with appropriating existing infrastructures/devices (e.g., Internet of Things toolkit, mobile app) *in situ* or developing new ones for a given setting (e.g., a novel public display).

Design: Covers the design space of an experience (e.g., iteratively creating a collaborative travel planning tool for families to use or an augmented reality game for playing outdoors).

***In situ* study:** Concerned with evaluating *in situ* an existing device/tool/service or novel research-based prototype when placed in various settings or given to someone to use over a period of time.

Theory: Investigating a theory, idea, concept or observation about a behavior, setting or other phenomenon using existing ones or developing a new one or extending an existing one.

Figure 1.1: Research in the wild (RITW) framework.

1.4 SCOPING RESEARCH IN THE WILD

There are many ways of conducting research in the wild. An initial challenge is to scope the research to determine what can be realistically discovered or demonstrated, which methods to use to achieve this and what to expect when using them. Sometimes, it might involve deploying hundreds of prototypes in people's homes (e.g., Gaver et al., 2016) to observe the varied adoptions and appropriations of many people rather than those of a few. Other times, it entails months of community-building and stakeholder engagement in order to build up trust and commitment before studying the outcome of an intervention they propose or a disruption on behavior (e.g., changing habits to enable communities to reduce their energy or increase their exercise). In other contexts, it can involve running a longitudinal study across geographical boundaries to determine how new tools encourage participation in different cultures, such as citizen science projects. The scoping will depend a lot on practical concerns, such as how much funding is available, the time of year, logistics and gaining the trust of and acceptance in a community in order to get people on board to see the potential value of a proposed technology.

A number of methods are typically used in RITW, including observation, surveys, remote logging of people's use of technology (e.g., monitoring their activity), and engagement with community members in a variety of contexts through the use of focus groups, co-design sessions, and town hall meetings—in order to hear their opinions and let them voice their concerns. Data that is collected using different methods is typically aggregated to provide a combination of quantitative and qualitative results. However, collecting multiple streams of data over several months can quickly multiply the outputs, making it difficult to tease out what might be causing particular effects or why people behave (or not) in certain ways. Much skill is involved in making sense of the different kinds of data without jumping to conclusions. There may be many factors and interdependencies at play that might be causing the observed effects or observed phenomena.

Despite this increase in uncertainty and lack of control, what is discovered and interpreted from RITW can be most revealing about what happens in the real world (Rogers et al., 2007; Marshall et al., 2011; Hornecker and Nicol, 2012). A benefit of RITW is greater ecological validity compared with extrapolating results from lab studies. Most significantly, RITW studies can show how people understand and appropriate technologies in their own terms and for their own situated purposes. Accordingly, RITW is increasingly being used to show 'impact' in terms of how new interventions have made a difference to a community (e.g., Balestrini et al., 2017), or how in the wild findings can provide empirical evidence for changing behavior or policy in society.

Thought Box: Beyond the Interface

Even though many of us still struggle to get the proverbial photocopier to copy (indeed our computer science department was offering tutorials to all staff, from professors to Ph.D. students, earlier this year with the arrival of a new machine), the pressing problems HCI researchers are increasingly concerned with are how people interact with an ecology of interfaces. A *core challenge is to enable people to be able to switch between multiple interfaces and multiple devices*. This framing requires understanding the context for why and how someone moves between them. Rather than being concerned with how best to support X (where X might be learning, working, socializing) using an individual device (e.g., a laptop, tablet or smartphone) it is necessary to work out how to design across platforms so that people can fluidly use multiple tools and devices, as they go about their everyday lives—picking up one, putting another down, or using several together in unison, by themselves or when interacting with others (Coughlan et al., 2012). What might seem obvious to do in a lab setting may not be obvious and may even be counter-intuitive in a real-world setting. A question this raises is how to frame, and which methods to use, when researching such multi-device settings across time and place in the wild?

1.5 AIM OF THE BOOK

The aim of this book is to provide an overview of HCI research in the wild, illustrating how it can traverse theory, design, technology, and *in situ* studies. It covers the motivations, concerns, methods and outcomes. As part of this endeavor, it addresses the challenges of conducting RITW, including the questions asked, the expectations, the trade-offs, the uncertainties, the form of analyses adopted, the role of the researcher, and their conduct when in the wild settings.

The book is targeted at both students and researchers who are new to the field of HCI and more generally, research methods, or for someone who simply wants to learn more about research in the wild. It covers RITW by charting and critiquing the what, when, where, why and how questions. In subsequent parts of the book, it examines the tools, methods, and platforms that have been imported, adapted and developed to study user-interactions in the wild, and how researchers have grounded concerns, problems, and new opportunities through their framing. It also outlines the benefits, limitations, impacts, and advances that have resulted from research in the wild.

1.6 SUMMARY

One of the motivations for conducting research in the wild is to demonstrate how a technology intervention can engage a community in a participatory manner. Underlying motivations include enabling people to collaborate, connect with each other or join forces in order to raise awareness, and act upon an issue. Another rationale for conducting RITW is to deploy novel technologies in a setting in order to provoke a response (e.g., getting people to comment on a new display in a street), a new kind of interaction (exploring how one looks in an augmented public mirror) or social engagement (e.g., encouraging strangers to talk with one another in a public place). A further reason is to develop new understandings and theorizing about how people use technology in their everyday lives—based on the body of empirical work that demonstrates how behavior differs or is the same as when using "older" and other kinds of technologies. In summary, RITW is becoming more widely accepted as a de facto way of conducting research for HCI, complimenting but also questioning the validity of traditional lab-based research approaches.

CHAPTER 2

Moving Into The Wild: From Situated Cognition to Embodied Interaction

2.1 INTRODUCTION

The phrase "in the wild" first came to the forefront, in the late 1980s and early 1990s when anthropologists Lucy Suchman (1987), Jean Lave (1988), and Ed Hutchins (1995) began writing about cognition in the wild. Collectively, they critiqued the fledgling field of cognitive science, which was concerned with how the mind worked. The accepted theorizing at the time focused on information processing in the head, and the construction of rational models of behavior as the execution of plans. In sharp contrast to this classical view, they explain cognition—as observed in everyday practice—being distributed and situated in the moment. Moreover, in their respective books (see Figure 2.1), they cogently argue that cognition can only be studied in the wild. Their approach was to present a social anthropology of cognition and cognition in practice, respectively.

Figure 2.1: Suchman's, Lave's, and Hutchins' classic "In the Wild" books.

2.2 PLANS AND SITUATED ACTION

The first in the wild classic was Lucy Suchman's (1987) *Plans and Situated Actions* book. It took the fields of HCI and computer science by storm—and its insights were quickly adopted by a new generation of researchers and students. For many, it resonated with their discontent and worries about the limitations of traditional cognitive models. For others, it opened their eyes to new ways of thinking about human-machine interactions. The accepted view at the time was that scientific models were needed to explain how the mind works and that these should form the basis of user models used in machine-human dialogues. Folk theories or common sense explanations were dismissed as inadequate. Suchman, however, argued the opposite: *common sense* notions of planning should *not* be viewed as inadequate versions of scientific models of action, but taken as resources people use in their practical deliberations. To support her contrarian view of how to conceptualize and understand human behavior, she described how people use these resources along with various constraints in the environment in their everyday planning and action. Instead of developing so called scientific models to develop human-computer interfaces, developers should draw from accounts of how people act and react in their everyday lives.

Much of Suchman's early research was to provide detailed situated accounts of the relations among people and between people and technology. One of her most cited examples is of a study she conducted of pairs of users trying to fathom out how to use a Xerox photocopier. While not an in the wild study (since it was conducted in a Xerox Research lab), she noted how its complicated help system did not match the way the pairs understood how it worked or what to do when it did not work in the way they thought it should. The outcome of her detailed analyses of the mismatched photocopier-user interactions led many programmers and developers to rethink how they should structure and what to include in their human-computer models, replacing the simplified process models, that followed sets of rules, such as "if x then y" with alternative kinds of situated models of action (Dourish, 2001).

An analogy that she used in her book to illustrate what she meant by situated action is a description of what it is like to ride the rapids in a canoe. She notes how a great deal of deliberation and reconstruction goes into a canoeist's plan both before they begin and in their account of what happened after the event, but from then how they actually navigate the rapids, depends on embodied skills in responding to whatever comes their way. This powerful image resonated with many as to why models of plans as a control structure that specify behavior were inadequate when designing user interfaces. Despite its impact on a generation of researchers, however, this example, itself, has been somewhat parodied and often misunderstood. Many took it to mean plans are irrelevant to how we act. Suchman never claimed this (and goes to great length to explain what she meant in her later revised version of the book), arguing that what happens in practice is the interaction of both

the contingencies and the projected course of action. A legacy from her pioneering work is the commonly accepted view that users don't follow instructions and plans as simply as had been assumed.

2.3 COGNITION IN PRACTICE

Jean Lave's (1988) book *Cognition in Practice*, published a year later, was primarily concerned with debunking the academic snobbery associated with "common sense explanations and real-world contexts." Similar to Suchman's critique of cognitive science models of everyday planning, she went to great lengths to explain how experimental lab research wasn't superior to everyday people's accounts of what they do in their lives. Moreover, her program of research showed how it was more valuable and legitimate to study people's cognitive behavior in everyday contexts, which she described as "cognition in the wild." To demonstrate how her approach could provide new understandings, she studied adults practicing math in a variety of real-world contexts. Some of the examples she described in her book, which are most illuminating, are of people working out the best price for groceries when shopping in the supermarket and for how dieters measure unusual quantities of ingredients when making a dish at home while following a recipe. Similar to Suchman's book, she compellingly demonstrates, through her detailed case studies, how people often use opportunistic structures in the real world in their everyday cognition.

The legacy of Lave's work was to show how it was possible and necessary to move one particular form of cognitive activity—arithmetic problem-solving—out of the laboratory back into the realm of everyday life. In so doing, she showed how mathematics in the real world is the same for all kinds of thinking, shaped by the reflexive encounter between human minds and the context people find themselves in. A salient example that has been much cited—as illustrative of doing math in practice—is the "cottage cheese" problem; a male dieter, preparing a meal, was faced with having to measure out 3/4 of 2/3 of a cup of cottage cheese that was stipulated in the recipe he was using. How did he work it out? Not by multiplying 3 x 2 and dividing that by 4 x 3, resulting in ½, as would be expected if using algebra in school, but instead by using the available structures in the environment in a situated way. He first measured out 2/3 of a cup, and then spread it on a chopping board in the shape of a circle. Next, he divided the circle into 4 quarters, removed one of them and returned it to the container, leaving on the board the desired 3/4 of 2/3 a cup.

Similar to Suchman's canoe example, the dieter example of using external resources to solve a math problem paved the way for rethinking cognition in practice rather than in abstraction, and the insight that mathematics is *for* something; the mathematical abstractions taught in schools don't necessarily transfer well to use outside the classroom. A number of other examples in Lave's book are used to emphasize how people use the resources from the context they are in to solve problems. Together, the examples convincingly demonstrate how activities in settings are complex improvisations that have much variability. Doing math when out there takes a different form in

different situations. One of the outcomes of this early form of in the wild research was to make studying everyday and common sense reasoning acceptable, by giving it credibility and respectability (cf. Rogers, 1993).

2.4 COGNITION IN THE WILD

Ed Hutchins published *Cognition in the Wild* a few years later in 1995. His seminal book was also very much a reaction against the status quo; but more broadly than either Lucy Suchman's or Jean Lave's efforts. His beef was very much a rally against "cognition in captivity" and "disembodied cognition." He argued that much of mainstream thinking about cognitive science for the past 30 years had resulted in systematic distortions of our understanding of the nature of cognition. Instead, he proposed that cognition should be studied in its natural habitat and that, in doing so, it would change our ideas about its nature. He argued that what was problematic with the classical cognitive science approach, was not its conceptual framework *per se*, but its exclusive focus on modeling the cognitive processes that occurred within one individual. As an alternative, he argued that what was needed was for the same conceptual framework to be applied to a range of cognitive systems, including socio-technical systems at large (i.e., groups of individual agents interacting with each other in particular environments). To do this, he proposed studying cognition beyond the skin of the individual, encompassing the distributed nature of cognitive phenomena across individuals, artifacts, and internal and external representations.

Hutchins also argued that in order to reveal the properties and processes of such a cognitive system required conducting an ethnographic study of a setting. Paying close attention to the activities of people and their interactions with material media was considered fundamental to understanding how such a cognitive system works. Hutchins' intricate analyses of what happens inside a cognitive system at both the micro and macro levels were at the time groundbreaking. One of his most well known examples is an account of how the cockpit plus air traffic control tower system work together as a cognitive system. He illustrates this eloquently by describing the joint activity and accomplishment of a situation when a pilot and co-pilot fly their plane to a higher altitude in conjunction with listening to and talking with air traffic controllers. It demonstrates just how much coordinated activity depends on the orchestration of mechanisms through which co-located and distributed people make small signals to each other in order to progress a sequenced activity, and the levels of inter-subjectivity involved for different states of the system. A missed cue can easily result in a misunderstanding, especially if someone is not expecting it. These can happen even for the most routine of activities which then requires the cognitive system to engage in various forms of repair work and sometimes the adoption of workarounds to get the activity back on track.

The legacy of the distributed cognition approach is its demonstration of how insightful it can be to analyze the complex interdependencies between people and artifacts in their work activities—