


O'REILLY®



Understanding Context

ENVIRONMENT, LANGUAGE, AND INFORMATION
ARCHITECTURE

Andrew Hinton
Foreword by Peter Morville

Praise for *Understanding Context*

“Engaging, never shying away from tackling and unraveling the complexity that lies behind the simplest turns of language, this book explains plainly and clearly why designers should pay attention to much-misused concepts such as sense and place, and simultaneously provides sound and elegant foundations for a new and embodied approach to the architectures of information spaces. A necessary read for both those who want to understand the interplay of language, place, information architecture, and design practice and for those who create products and services spanning the digital and the physical, where context is everything.”

**—ANDREA RESMINI, PHD—AUTHOR, PERVASIVE
INFORMATION ARCHITECTURE; SENIOR
LECTURER IN INFORMATICS, JÖNKÖPING
INTERNATIONAL BUSINESS SCHOOL**

“For all those times when someone says something like, “this is not my beautiful life” when they’re using your design, you’ve probably ignored their context. Andrew masterfully makes the case for meeting users where they are, and putting context in its proper place: at the center of making meaningful design.”

—DANA CHISNELL

“Andrew Hinton has created a rigorous and wide-ranging framework for thinking about how we perceive and interact with our linguistic and digital environments. This framework is based on cutting edge cognitive science, and the result is an invaluable guide and common language for all the different people who create and want to understand these relatively new human environments places.”

**—ANDREW D. WILSON, PHD—SENIOR LECTURER
IN PSYCHOLOGY, LEEDS BECKETT UNIVERSITY,
LEEDS UK**

“Understanding Context is a door to knowledge that will allow an entire generation of digital designers to more properly consider the context in which their work is used. This book is easy to read, but also full of important academic concepts more designers should be talking and thinking about in this messy cross-channel world. Hinton’s writing is like attending a master’s program in Gibsonian psychology while having the smartest kid in class sitting next to you explaining what the heck is going on in everyday language. I applaud O’Reilly in bringing this work to the world and Mr. Hinton for pouring his heart into writing it.”

**—ABBY COVERT—AUTHOR, HOW TO MAKE SENSE
OF ANY MESS**

“Context is hard for humans. It’s even harder for the machines that we oddly hope will guide us through the growing chaos and complexity of modern life. Andrew is a better guide; his book helps us perceive, make sense of, and engage with the contexts that surround us—from cities to kitchens to Minecraft.”

**—LOUIS ROSENFELD—AUTHOR, INFORMATION
ARCHITECTURE FOR THE WORLD WIDE WEB;
FOUNDER OF ROSENFELD MEDIA**

“Andrew Hinton’s well-organized, useful, conversational approach makes this vast ontology of context not only accessible, but indeed like a long walk with J.J.Gibson himself. Come along for a day, come away better situated in a world remade (but not replaced) by technology, and get ready to give it better architecture.”

**—MALCOLM MCCULLOUGH—PROFESSOR OF
ARCHITECTURE, UNIVERSITY OF MICHIGAN;
AUTHOR, AMBIENT COMMONS AND DIGITAL
GROUND**

“As computers become smaller they are becoming more ubiquitous. Computers and computing are not only in our mobile phones, but they are rapidly being embedded in everything. Cars, kitchens, street corners and shopping malls are becoming smart and connected. In this connected world, understanding context is more important than ever before. Andrew Hinton has written a thoughtful, well-researched and insightful book, full of key ideas to help you navigate the connected future.”

**—DAVE GRAY—AUTHOR, THE CONNECTED
COMPANY AND GAMESTORMING; FOUND, XPLANE**

Foreword

IT'S IMPOSSIBLE TO UNDERSTAND context. There's always something we're missing. Not long ago, during a rim-to-rim hike of the Grand Canyon, I was thinking deep thoughts about two-billion-year-old rocks. They made me feel small. Although we're more stable than a tornado or a sandbar, we belong in the same category. We are delicate, imperfect patterns that come and go in the blink of an eye. Yet, we're also more ancient than rocks. We are made of stardust, indestructible matter as old as the universe.

That's when I heard the rattle.

Lost in thought, I nearly stepped on a snake. In unfamiliar territory, it's impossible to understand context, but it's still vital that we pay attention.

In the 1990s, I helped to grow a company called Argus. Over the course of seven years, we pioneered the practice of *information architecture* and bootstrapped our way from two to forty souls. Unfortunately, when the tech bubble burst at the end of the decade, we sunk the ship. We didn't see it coming. Later, while packing books into boxes, I suddenly realized what I'd lost. It wasn't just a company; Argus was a part of me. We'd built an organization of people, systems, and information that embodied and extended ourselves. That's the thing about context. It's impossible to see until it's gone.

A year after we closed Argus, I met Andrew Hinton. A group of us were gathered on the beautiful conference grounds of

Asilomar to discuss how we might advance the practice of information architecture. At the time, our work was tied to websites, but Andrew told us to embrace “the structural design of information environments.” So, we wrote those words into the bylaws of the *Information Architecture Institute* and into the new edition of the “polar bear book,” and that became the definition of information architecture that’s celebrated by thousands of people in dozens of countries each year on *World IA Day*. Our words and actions have unforeseeable consequences beyond our current context.

There’s a new ship in town by the name of TUG. It’s reframing information architecture. The Understanding Group was founded by Dan Klyn and Bob Royce, and I’m a strategic advisor. It’s the perfect place for Andrew to be an information architect. He gets to tackle massive projects while surrounded by amazing people. And he’s able to continue what he began in Asilomar: building out the “architecture school” of information architecture.

In articles and talks and in this book, Andrew is helping us all to realize that we’re not designing software or websites. Because “language is infrastructure” and “the map is the territory,” the things we build and inhabit are “places made of information.” From the perspective of experience, these digital ecosystems are as real as the Grand Canyon. This unfamiliar territory can engage, inspire, or overwhelm. It’s easy to become lost, and there are snakes. That’s why this book is important. It’s a map for mapmakers. It won’t explain everything from here to there. That’s impossible. Still, if you’re brave enough to hike its crags and canyons, you will be better at making places and understanding contexts. This book is not a straightforward journey, nor is it short, but it’s one I highly recommend.

—PETER MORVILLE

Preface

A PREFACE CAN BE MANY THINGS; this preface is *two* things. First, it gives a practical introduction to what this book is and what you can learn from it. Second, it tells a personal story about my reasons for writing the book and what I hope it will mean to you.

The Practical Bit

This book is about how information shapes and changes the way people experience context in the products and services we design and build. It's not only about how we design *for* a given context, but how design participates in *making* context. It begins with how people understand context in any environment. Then, it explores how language takes part in that understanding, and how information architecture helps to shape context, and to make it better. It's also an exploration, where "understanding" is more verb than noun; it's less about defining the right answers than discovering the right questions.

Context is an abstract idea, but it brings concrete challenges. What defines the "place" a customer is in, if he's shopping "online" and "in a store" at the same time? What determines the boundaries of a user's identity if her social network has multiple layers of privacy controls? How does a user know if something is a button and not just a label? What does it mean if you put something "in the cloud," but it's also "on your

phone” and “in your laptop”? When we say we are “here” what does that actually mean now that we can be interacting and talking in many places at once? From accidentally hitting Reply All to an email, to discovering that Facebook shared embarrassing photos with your in-laws, we’re facing the challenges of being immersed in contextually confusing environments.

Design needs handles we can grasp and manipulate to make context do what we intend: form understandable environments where users can meet their needs. To get there, we need to do some digging to understand what those handles are and what we’re changing in the world when we use them. So, although the book does cover a lot of *theory*, there’s nothing more *practical* than understanding the nature of our materials. If context is a material or medium that we can affect in user experiences, we should know what it’s made of and how it works.

Who Should Read This Book?

If you design or make products or services that connect one part of our lives to another part, you will get value out of this book. In coming years, context is going to be an increasingly critical part of any design project. And yet, we don’t have a mutual understanding of what context is, or what happens to context when we design something one way versus another. We talk a lot about “information” and “experiences” and “environments” and even “context,” but they’re amorphous, fuzzy concepts. We’ve reached a point in design practice at which we can’t rely on such foggy notions anymore. This book provides a way to understand and work with context, using information as a medium for making.

That said, I did write this book with an assumption that you have some experience as a designer and are interested in exploring these strange, complex questions that underlie the surface of our work. If you're a beginner, you can get a lot out of it, but you won't get "the basics"—even though the concepts here are, in some ways, more foundational than what we usually think of as "basic" theory and practice.

So This Book Teaches Methods for Designing Context?

Not really. This is definitely an "understanding how it works" book more than a "how to make stuff" book. It touches on methods and materials and has a lot of concrete examples. But really it's meant to inform the methods we already use and to suggest some new ways of looking at them or doing them.

I also hope the book helps reframe some important aspects of what and how we design. For myself, the process of researching and writing this book has fundamentally rewired the way I see the world, especially when I plan and design environments for my clients. I've found this new perspective to be immensely valuable in my own work.

Why Information Architecture?

Although "information architecture" is in the title, this book isn't only for self-described "information architects." The structures we make and depend on for all sorts of design work involve information architecture. Architecture is the starting place for figuring out foundations, boundaries, and

connections. Still, the book discusses areas of focus for other disciplines in some detail throughout, especially in the first half. If you're an interaction designer, content strategist, usability engineer, researcher or other such profession, you will find plenty here that relates directly to your main work. And in the second half, hopefully it becomes clearer how intermingled and interdependent these practices are with one another and with information architecture.

Additionally, the reason why I became interested in context to begin with was information architecture, which is also my "home" community of practice. As a community (and as a relatively young, forming discipline) we have a complicated history of figuring out what information architecture is, what it actually does in the world, and what all that means. Although context is a much bigger topic than information architecture, and information architecture is about more than context, I decided the concerns of information architecture were where I'd spend most of my time in this exploration.

What Will You Learn from This Book?

Here is a sampling of the sorts of things you should take away from reading this book:

- How people experience and comprehend context
- Principles for designing environments in which context is more understandable and trustworthy
- How *affordance* works, and how it informs everything users perceive and do

- What *placemaking* and *sensemaking* mean, and how digital information can both enhance and disrupt how they work
- How language, in all its forms, works as an important “raw material” for context design
- Also how language has *semantic function* similar to the way that physical things have affordance, and what that means for digital interfaces and other environments
- Models for understanding the personal context users bring to the environments we design—their situations, motivations, and narratives

These are just some highlights. Overall, the main take-away from this book will be a fresh perspective on what it means to design in a time when digital technology is saturating everything around us.

A Tour Through the Book's Six Parts

The book is made up of 6 parts, each a sort of small book in itself. The parts build on each other through 22 continuously numbered chapters.

Part I The Context Problem: What It Is and How To Think About It

This first part explains what the book means by “context” and introduces some core ideas about context that will be explored throughout. Using an everyday travel scenario and a bit of historical background, it explains what the challenges are and how they came to be. It also sets up the three-part model we work with through the rest of the

book.

Part II Physical Information: The Roots of Context

This part provides a foundation for how users comprehend environments. It explains the theory of affordance, the essential dynamics of perception and embodied cognition, and a framework for describing the structural elements of environments, including the Principle of Nesting, and how Surfaces, Objects, and Events combine to make Places. Although this part is mainly about nontechnological topics, it includes examples to show how these concepts are relevant to designed products and services.

Part III Semantic Information: Language as Environment

In this part, we investigate how language (speech, gestures, text, and pictures) works as an additional environmental layer of “semantic function,” and how that affects context. This part touches on essential concepts about signs and symbols and how we use forms of semantic information for simulating physical environments in user interfaces. It also discusses placemaking and sensemaking and how language structures our contextual experience.

Part IV Digital Information: The Pervasive Influence of Code

Part IV shows how digital information is, at its core, meaningful to computers more than humans. It also shows how that dynamic influences how we experience physical and semantic information. This part demonstrates how digital information makes it possible to

create environments and places that don't behave like the physical world, and how that can be both good and bad.

Part V *The Maps We Live In: Information Architectures for Places and People*

This part brings the ideas from the previous four parts together to show how information works as systems of meaning. The chapters within it explore examples from different kinds of placemaking, using “maps” as a framing device for how we change our territories, or make new ones. It explains how these environments affect our social relationships, conversations, and identities. This part focuses more on the architectural concerns of how *place* works and less on the object-level concerns of a particular interaction.

Part VI *Composing Context: Making Room for Making Meaning*

The last part introduces *composition* as a useful way of looking at how we use information to shape context, including how the composed environment requires a stable ground to build and act upon. **Part VI** describes how we use materials of semantic function in the form of Labels, Relationships, and Rules—aligned with ontology, taxonomy, and “choreography.” It explores how people construct their experience through narrative and story, and how they participate in creating those narratives. And finally, **Part VI** looks at some principles for accommodating meaning-making, and some techniques for understanding context and modeling new environments.

The Personal Bit

When I began this writing project, my aim was to make a slender, deft volume of focused essays. A “thoughts about context, for design” kind of book. Three years later, I’m putting the final touches on a tome collecting, in essence, six little books, trying to solve the mysteries of the universe.

Funny how things like this work out.

* * *

Here’s what got me started:

For a long time—since before I worked with technology for a living—I’ve wondered about how it is that language makes experience. How is it that a novel can captivate us and make us feel as if we’ve lived through those events? How does a table-top game construct a shared place in which events occur that we might remember and talk about years later, even though nothing “real” happened? Why is it that a poet can break a line—in just such a way that it breaks the reader’s heart?

When I eventually found myself in a software-related profession, my obsession only grew. When logging in to a system to move files around, what is it I’m affecting with commands such as “get” or “put.” It’s just bits, being rearranged on the same disc, but somehow those words make it relevant to my body. When exploring early social architectures such as UseNet, Internet Relay Chat, or later, LiveJournal, I noticed how deeply I sometimes felt about my conversations there, and how these places were meaningful *as places*. How is that possible? They’re just virtual marks on virtual surfaces, which aren’t even as real as the printed type in a paper book.

Soon, I found a community where many people were wondering similar things. They were talking to one another under a loosely shared label: *information architecture*. Something about that phrase clicked for me—yes, I thought, that’s what’s going on: information that somehow feels as if we live in it; structures, rooms, passageways. Not virtual reality, exactly. It’s information that shares some of the qualities of space, whose places become as meaningful for us as any other places in our lives.

As someone who now identifies as an “information architect,” I kept working at these questions, until it occurred to me that so much of what I was doing for a living was repairing problems with *context*. Software is doing something to the world that is detaching and rupturing context from whatever helped it make sense before we had computers, networks, and hyperlinks. Eventually, with encouragement, and after many conference talks and articles, I decided to write a book about that.

* * *

When you start writing a book about something, suddenly ideas that you thought you had all figured out seem flimsy and unformed. After finishing what I thought was the first third of the book, I grew worried—what if none of what I think about this is true? So, I did some further research to validate my hunches.

I discovered my hunches were sometimes on track, and often really wrong, and that this thing we call context is actually not well understood. Among academics, there is rigorous work being done, but it isn’t exactly settled science. Yet, even the best of that work wasn’t making it into the general conversation of design practitioners and in the popular “UX” literature.

I realized that if I were to take this book seriously, I couldn't just think aloud on the page about what the answers *might* be. So, six months into my work, I had to set aside the chapters I'd already drafted and take the time to really learn the subject as best I could, while writing about what I was discovering. A couple years and quite a few pages later, here I am revising the preface for publication.

Here's the thing: I'm not finished. The more I learned, the more I saw there would be to discover. Hence, as I mentioned at the beginning, "understanding" is something I hope we can do together, both in this book and beyond it. To that end, I invite you to visit this book's home site (www.contextbook.com), where additional content and links will accumulate, including a bibliography.

My wish for these ideas isn't that they be absolutely right (though some of that is nice), but that they help move along the work we do together toward making better places, good and human places, for the people who dwell in them.

So, now, let's dig in.

Acknowledgments

I WANT MOST OF ALL TO THANK MY WIFE, Erin, who has been a patient supporter through the long, often angst-ridden process of writing this book. She not only went without a husband for many weekends, evenings, and so-called vacation days; she then weathered my verbose ruminations and bouts of self-doubt when I was in her presence. She doesn't believe it, but it's true that I could not have done this without her.

Thanks also to Madeline, my daughter, who has also endured my authorial tribulations, and who has been such an inspiration to me as I've watched her grow up, tackle huge challenges, and already become so much more than I could've imagined.

Thanks to Peter Morville, who told me a long time ago, regarding a completely different subject, "you should write a book," and who then helped me find a path to get it done, with wise counsel along the way. Also thanks to Lou Rosenfeld, who has also given me such generous encouragement and advice over the years. And thanks to Christina Wodtke, especially for that email invitation circa 2001, and all the invaluable conversations since.

A special thanks to Dan Klyn and Bob Royce, who both have enriched and influenced my perspective on this book's subject, and who invited me along on their joint mission to bring "making things be good" to the world—something they called The Understanding Group. I can't imagine any other vocational home giving me the room to work out these ideas

and be the self I needed to be while writing this book.

There are many other wonderful people who have contributed their energy and care in ways large and small toward making this book happen; some have been sources of conversations and knowledge that have become part of the book's fabric, and some have even taken time out of their lives to review drafts and help make the book better. I can't possibly list them all, but I should especially mention, Jorge Arango, Andrea Resmini, Abby Covert, Marsha Haverty, Andrew Wilson, Sabrina Golonka, Karl Fast, Dave Gray, Christian Crumlish, Richard Dalton, Lis Hubert, Malcolm McCullough, and Don Norman.

Thanks also to a mentor from what seems like a previous life, poet and teacher Jeffrey Skinner, for telling me I'm a writer and showing me what it means to be one. When I catch myself just talking about something I could make or do, I hear his rightfully impatient voice from decades ago, saying, "Stop talking about it and just do it!" He also taught me it's better to be done than perfect, advice without which I would still be researching, writing, and revising this tome.

Thanks to my publisher, meaning everyone there who played a part in making this book a reality. When I started working with O'Reilly, I wasn't sure how it would turn out. I knew I wasn't writing something that fit a standard "technology book" mold, but I also knew that O'Reilly was one of the first publishers that really mattered to me—starting with the copy of *The Whole Internet User's Guide and Catalog* I found in 1992—and that O'Reilly's deeply humanist vision would make a good home for my wandering, philosophical ideas about information technology. Sure enough, I found everyone involved to be thoughtful and welcoming, patient with my journey but responsible in prodding me along. So, in

particular to Mary Treseler, Simon St. Laurent, Amy Jollymore, and Meghan Blanchette, I want to thank you all for your belief in this project, your encouragement when I struggled, and your kind guidance through the entire process.

Finally, to my mother, Mary; my dad, George; and (in memoriam) my stepfather, Paul: thank you for getting me to adulthood, showing me what was possible if I took the time to do something right, and teaching me that—with hard work—I could accomplish what I set out to do.

Part I. The Context Problem

What It Is and How to Think About It

PEOPLE HAVE BEEN TRYING TO FIGURE OUT HOW WE UNDERSTAND THE WORLD FOR A VERY LONG TIME; and from all evidence, there are no certain answers for many of the most important questions. So, it would seem that writing a book about “context” would be a fool’s errand. Why try to tackle it to begin with? What do we mean by it when we say it? And after going to all this trouble, what real-world problems will it help us to solve?

Part I helps to answer these questions. It introduces the basic challenges of a growing “context problem” to establish reasons why we should bother with understanding context at all. It also presents common scenarios to help illustrate those challenges. Finally, it introduces a working definition of context, and some models we will use to explore how information works to create and shape context in the chapters to come.

^[1] *Utriusque cosmi maioris scilicet et minoris [...] historia, tomus II* (1619), tractatus I, sectio I, liber X, De triplici animae in corpore vision (from Wikimedia Commons)

Chapter 1. Everything, Yet Something

All and everything is naturally related and interconnected.

—ADA LOVELACE

Birds in Trees, Words in Books

IT WOULD BE GREAT to say there's a simple secret that lets us figure out a subject such as context quickly and easily. But let's not kid ourselves; context is a big, hairy, weird topic. By nature, it's about everything. And if something is about everything, how can we even begin to understand it? When we think we've caught it and we try to crack it open to see what's inside, all we seem to get is more context. It's the stuff of late-night dorm conversations among philosophy students, or tedious debates among philosophers and physicists. Context shifts and dances, it slips and slides. It insists on its mystery, yet it demands we come to terms with it every single day.

Still, in some ways, context actually isn't so mysterious, because from one moment to the next we seem to know it when we see it. Whenever we're trying to figure out *what one thing means in relation to something else*, we say we're trying to understand its context.

It can be a bird in a tree, or a stone in a stream. It can be a single word in a sentence. It can also be that sentence in a

paragraph, that paragraph in a book, and that book's context in a library or on a bedside table—all of which can influence how we interpret the single word where we started.

It can be the context of a text message I receive on my phone, where a “:-)” emoticon can make the difference between an insult and a friendly jest. It can be a *channel* as in “cross-channel” design, like mobile, desktop, telephone, or broadcast. Or, it can be all those channels, but from the separate perspectives of a child and a senior citizen, or a user on a bus versus a customer in a store. It can be also the context of one's own identity, and how one behaves around family versus the workplace. The more closely we look, the more we see how many parts of our lives depend upon context for meaning anything at all.

Whatever context is in any given situation, though, one aspect remains consistent: we need context to *be clear* and to *make sense*. We know that when context is not clear or it doesn't make sense, it means we don't understand something. And misunderstandings are almost always *not good*.

How do we get beyond this general sense of “good” or “clear” or “understandable”? What makes anything more clear or not clear? These are philosophical questions with practical implications. So let's begin with a practical scenario involving an airport, a calendar, and getting from one place to another.

Scenario: Andrew Goes to the Airport

On the morning before a cross-country business flight, I was already getting packed. I knew a seven-day trip would mean extra effort compressing everything into my small carry-on bag, and an early departure meant no time to pack on the day of the trip. While rolling my shirts into tight bars of cloth to

prevent wrinkles—wrinkles I knew would manifest anyway regardless of my efforts—my iPhone started trilling and buzzing. Two apps on my phone—one from Delta Airlines and the other from the TripIt travel service—were each reminding me I could check in and get my digital boarding pass.

I opened the Delta app and tapped my way through the check-in process. After a few steps, it offered me an “Economy Comfort” seat for just a few dollars more. Because I don’t especially like folding myself into my laptop like human origami, I agreed to pay for the perk. I half wondered if now being “Economy Comfort” would give me any other privileges, but I was busy, so I decided to find out when I got to the airport.

To economize a bit and get a better fare, I’d decided to depart from an airport a bit further from home. On the day of the flight, when I arrived inside the terminal, I realized I wasn’t used to this airport and found myself having to learn a new layout. I scanned the environment in front of me, thinking through the options. I didn’t need to print a boarding pass, and I didn’t need to check my luggage bin-sized bag, so I looked for the right queue for going through the TSA security check.

Similar to my usual airport, this one had queue lanes organized by various categories. I was so used to going to a particular lane at my home airport that I hadn’t thought about how they were labeled in a long time.

One was for fast-track-security TSA Pre-Check passengers, which I was pretty sure I wasn’t (though I also wasn’t sure how I would know). One was for “SkyPriority” flyers, and that one made me wonder: I was a “Medallion” member of Delta’s SkyMiles program, but I only had a “Silver” Medallion. I was

also a holder of Delta's branded American Express Gold card, which provides some of the benefits of Medallion status, but not others. (Even though the card is "Gold," it has nothing to do with the "Gold Medallion" level in Delta's loyalty hierarchy.) In my hurried state, I couldn't remember if Silver Medallion qualified me as "SkyPriority" or if my Gold Card did, or if Economy Comfort got me in, or if none of these conditions applied.

I approached a Delta service representative and asked where I should go. She said, "Let me see your boarding pass; it will say if you're SkyPriority." I got out my iPhone and opened the Delta app to show her. But no dice—I couldn't get it to open the boarding pass. The app was having to recheck the network to update the boarding information, and couldn't reach the cloud for some reason, even though the phone's WiFi indicator had full bars. What was happening? Then, the phone popped up a screen about paying for airport WiFi access. Ah, so my phone was naively assuming it had Internet access because it was connecting to WiFi, but the Internet was not available without payment. How annoying.

I had no time to fuss with that in the moment. So, I fumbled until I turned WiFi off, allowing me to use my carrier's data network instead. But that didn't help either—I couldn't get a signal inside the building. I had to apologize and walk back outside for a cell connection to get my digital boarding pass to update.

I then saw a button on the boarding pass inviting me to "Add to Passbook." I had heard of Apple's Passbook app but just hadn't used it yet. Thinking maybe this would help me avoid the forced-data-refresh issue, I added the Delta pass to the Passbook app. I also noticed there was no mention of SkyPriority or any other special security line status, so I

reentered the terminal and used the slower queue, with all the other standard travelers.

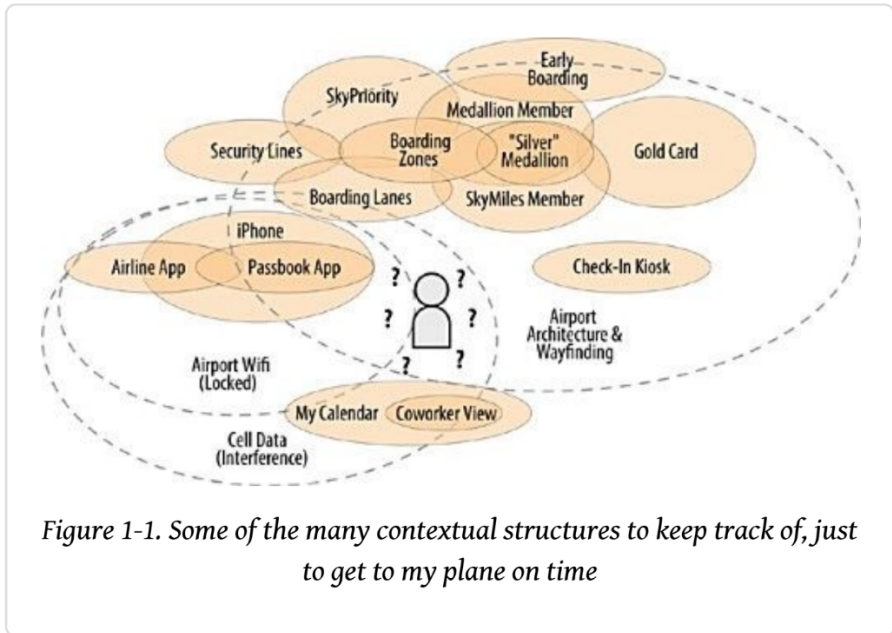


Figure 1-1. Some of the many contextual structures to keep track of, just to get to my plane on time

While waiting in line, I noticed that every time I tapped my phone's Home button to bring up the unlock screen, there was a notification about my flight. I hadn't seen this sort of alert before. I tried swiping it to either dismiss it or open the app that was generating it, but that only brought up a strange-looking version of my boarding pass. It was odd behavior because normally I can't interact with apps when the phone is in locked mode, other than a few things such as audio controls. I wondered if the notification engine on the phone was on the fritz, so I rebooted the phone—no change. By then it was time to go through security anyway, so I just showed my boarding pass and ID and then finished the tedium of the security screening.

After finishing the security ritual and getting my belt and shoes back onto my body, I received a text message from my employer's operations manager, asking about my travel schedule. I answered back that my itinerary should be on my calendar. I could see it on my iOS Calendar app; why couldn't he? I finally just told him the details and said I'd try to figure out the calendar problem later.

TripIt had just alerted me via text message which gate I should use for my flight, so I kept glancing at the message to remind myself where I should go in the airport. Eventually, I made it to the correct gate, following signs and pointers along the airport passageways.

At the gate, I heard the airline associate announce who could board and when, over the public address system. I waited until I heard the category I saw listed on my boarding pass—Zone 1—and boarded when it was called. Then, I had one more label to follow, my seat number. On this flight, I was known as the passenger in 24 C. I finally plopped into it, pondering all the confusing things I'd just encountered.

Breaking It Down

In scenarios like this, we find ourselves in a tangle of digital, semantic, and physical structures. We are *agents* trying to take *action* in our *environments*, and we have to *understand* those environments well enough to take the appropriate actions. For myself, even though I'm a frequent traveler, I was still running into issues that caused me to have to stop, think, and figure out my environment. When I'd try to act out of habit, it would either work or it wouldn't, depending on how the environment accommodated my action. Each element of the environment made a demand on my ability to understand

what I was doing. These moving parts were all in play at a wide variety of environmental levels, from the broad level of a mobile phone network, to the level of a little virtual switch deep in the settings of my phone. Some required almost no effort, but many required a lot, to the point that I was reduced to talking to myself to figure them out.

For example, we don't normally have to wonder about *what* we are. But there were many labels for me in these overlapping systems that made me have to think hard about *what I was to the environment I was in* and what that meant about the actions I could take.

The categories for security and boarding and even my assigned seat presented many overlapping facets that defined what I was to the airport. In each instance, I had to determine the rules represented by the labels involved. Was I a Gold Card traveler? A Silver Medallion traveler? An Economy Comfort traveler? What Zone, what seat? Some labels still made no sense to me.

For example, what did any of these have to do with "SkyPriority"? After I was seated on the plane, waiting for others to board, I checked Delta's site on my phone. I discovered this answer among their Frequently Asked Questions: "SkyTeam Elite Plus members and customers with First and Business Class tickets are eligible for SkyPriority." This sounded nonsensical to me, because I had no idea what SkyTeam Elite Plus meant, or how a passenger who is already "Elite" might possibly benefit from an appendage of "Plus."

Then there was my phone. I tried dismissing that pesky flight notification from the phone's lock screen, but it wouldn't leave. It was only after I'd had a moment to breathe that I realized it was a feature, not a bug. Passbook was taking an action without my prior approval or awareness, trying to do

me a favor by making it possible for me to open my boarding pass without having to unlock my phone. (I've since found this to be a useful feature!) But it was hard for me to learn as such, because the app took agency that broke the structural expectations that I'd learned to date by using iOS, and the alert gave no indication it was coming from Passbook.

Regarding my calendar, it occurred to me that I could see the TripIt itinerary information in my own calendar view because I had subscribed to it, but that subscription must be visible only to me, not those with whom I shared the calendar. I was thinking of "calendar" as the whole thing I was seeing, but in fact it contained a number of calendars, some of which were subscribed to rather than part of the specific calendar I had opened to coworkers. The meaning of "calendar" was disjointed, and the relationships between the various meanings resulted in confusion.

I also mused at the level to which airline travel is now depending on people to use networked devices, even though most airports have complicated network access. Think of the thousands of people who have the same problem with refreshing their boarding passes, not to mention getting gate updates and other information they've come to depend upon. (To Delta's credit, the app avoids this problem in its most recent updates.)

All of these complications were largely problems with context. I struggled to accurately perceive the *meaningful relationships between elements* in my environment. These problems caused extra work and stress for an activity that used to be much more straightforward. Yet, they all happened in an environment that people designed and built. Each encounter was part of a human-made place, composed of physical surfaces, language structures, and digital bits,

woven together into a complex system—a system that should be much less confusing and ambiguous than it turned out to be.

Now, imagine this same journey only a few years from now, when sensors in the terminal will be able to pair with my smartphone, pumping updates to its screen, chatting with me about where I should go through a wearable device, buzzing about which vendors close by serve my favorite foods, and who in my social network might be coming through the same terminal. How will that ecosystem know how to be truly relevant? Will it be overburdened with the noise of advertisements or “pay for more” services? Will it even have the rules figured out any better than I do about what queue I should use or if I can check my luggage for free? Whom will it be alerting to actions that I assume are private, or at least confined to the walls of the airport? How many more “calendars” will be overlapping and intermingling with mine by then?

Humans are much better at sorting out the vagaries of cultural meaning than machines. Yet, if we trip over these contextual conundrums so frequently, how are digital systems going to understand them any better? Moreover, how are we to keep absorbing so much contextual ambiguity and complexity from the multiplying layers of information we’re expected to comprehend just to finish basic tasks?

Chapter 2. A Growing Challenge

Where is it, this present? It has melted in our grasp, fled ere we could touch it, gone in the instant of becoming.

—WILLIAM JAMES

Early Disruptions

THE AIRPORT SCENARIO PRESENTED IN [Chapter 1](#) IS JUST A MUNDANE EXAMPLE OF A BIGGER ISSUE: context is a lot more complicated than it used to be, and it is only going to become more disrupted and detached from the physical clues we've relied on until now. For most of human existence, the context of people, places, and objects has been pretty straightforward. If you're in a field surrounded by trees, that's where you are. The field won't magically transform into a bustling village square in the blink of an eye. If you pick up a tool, a hammer perhaps, it does what its form suggests it will do. Physical laws dictate that the concrete world behaves in certain ways that we evolved to comprehend, usually with little or no explicit thought. So bodies and brains developed to prefer environments in which we don't have to think so hard about what we're doing.

But contemporary life is more complex. Now, we're surrounded by stuff that requires conscious thought for us to get what we need from it. For a long time, technology has

been chipping away at the immediate clarity of context. The invention of writing meant that something said (written) in one place could be read—and therefore, “said” all over again—anywhere the document might go, separate from its original utterance. Writing thus set the stage for later technologies to disrupt the connection even further. The nineteenth and twentieth centuries transformed the way we communicate by “radically separating the contexts of message transmission and reception.”^[2] With the telegraph, tapping an apparatus in one place could send messages across continents and oceans. And, we still hear of celebrities and politicians being caught saying embarrassing things when they don’t realize their clip-on microphones are still live.

In a sense, this is just the extension of all mechanical technology, which introduced contraptions that separate a specific action from its effect. Turning a crank can lift heavy objects via pulleys. Pulling a lever can cause a railroad track to change the course of massive locomotives. Eventually these interactions became so numerous and complicated that they required new fields of study and specialized training. Hidden complexities enter our environment, and the further the specific physical act is from its effect, the more the *context* of cause-and-effect requires explaining and learning.

Most of these changes happened slowly enough that we could keep up with them. Complex machines were mainly run by specialists, and major new communications devices emerged at a slow rate of only a few every generation. It helped that much of it was tied to geography, with all of the physical and cultural context that implies. However, digital networks quickly detached us even from physical location. As William J. Mitchell explains in *City of Bits: Space, Place, and the Infobahn* (MIT Press):

Unlike telephone calls or fax transmissions, which [used to] link specific machines at identifiable locations (the telephone on your desk and the telephone on my desk, say), an exchange of electronic mail (e-mail) links people at indeterminate locations.... You will not know whether I transmitted it from my office or typed it in at home while sipping a glass of wine or entered it into my laptop on a trans-Pacific flight and then sent it from a public telephone at Narita airport.^[3]

The .edu domain of a university might imply that the addressee has some connection to that school's campus, but even early dial-up access broke that connection because it allowed access from anywhere there was a phone line. The "where" of an email address is just the text in the address itself, a string of characters with only the barest literal connection to the sort of physical "place" in which we evolved. Of course, nowadays, we appropriate domain extensions such as .tv—created for the country of Tuvalu—as meaning "television," instead. Geography and "place" have drifted into a strange, turbulent relationship. Back when we were online only part of the time (when dialing through a modem or sitting at a desk), that relationship seemed manageable. But now that we're personally connected to a global network 24 hours per day, we find ourselves both walking on the ground and living in a cloud, all the time, all at once.

The Role of the Web

A big reason why digital networks became so ubiquitous was the advent of the World Wide Web. The Web became the petri dish in which the culture of "being digital" explosively grew. The Web meant that we didn't have to worry about what server we were on or to which directories we had access. It meant that we could just make links and think about

structure later.

The principle driving the original development of the Web was to add a protocol (HTTP) to the Internet that facilitated open sharing. In the phrasing of its creators—in the Web’s founding document—its purpose was “to link and access information of various kinds as a web of nodes *in which the user can browse at will.*”^[4] When you give people the capability to create environments with more ease and flexibility than before, they will use it, even beyond its intended boundaries.

The Web has now become something that has far outstripped what we see in dedicated “web browsers” alone. The characteristics of hyperlinks that once were only about linking one metaphorical “page” to another are now fueling all manner of APIs for easy, fluid syndication and mashing-up of information from many different sources. The spirit of the hyperlink means everything can be connected out of context to everything else. We can link enterprise resource management platforms with loading docks, map software with automobiles, and radio frequency ID (RFID) chips injected into pet dogs that include the dog’s records in licensing databases. Even our sneakers can broadcast on a global network how far we run, for anyone to see. The Web is now more a property of human civilization than a platform. It is infrastructure that we treat as if it were nature, like “shipping” or “irrigation.” HTTP could be retired as a network layer tomorrow, but from now on, people will always demand the ability to link to anything they please.

Additionally, these technologies have allowed us to create a sort of space that’s made of bits, not atoms. This space is full of places that aren’t just supplementary or analog versions of physical environments; they are a new species of place that we visit through the glowing screens of our devices. Writing

about one of those places—YouTube—cultural anthropologist Michael Wesch describes how users sitting in front of a webcam struggle to fully comprehend the context of what they’re doing when communicating on “the most public space in the world, entered from the privacy of our own homes”:

The problem is not lack of context. It is context collapse: an infinite number of contexts collapsing upon one another into that single moment of recording. The images, actions, and words captured by the lens at any moment can be transported to anywhere on the planet and preserved (the performer must assume) for all time. The little glass lens becomes the gateway to a black hole sucking all of time and space—virtually all possible contexts—in upon itself.^[5]

The disorienting lack of pre-Web context one faces on YouTube is not confined to videos. We’re spending more and more of our lives inhabiting these places, whether it’s Facebook or a corporate intranet. If we measure reality by where meaningful human activity takes place, these places are not merely “virtual” anymore. They are now part of our public infrastructure.

The contextual untethering the Web brought to computer networks is now leaking out into our physical surroundings. Structures we assume have stable meanings from day to day are shot through with invisible connections and actions that change those meanings in ways we often don’t understand. We live among active digital objects that adjust our room temperature, run our economies, decide on our financial fitness, route our trains and car traffic, and advise us where we should eat and sleep.

As Rob Kitchin and Martin Dodge explain in *Code/Space: Software and Everyday Life* (MIT Press), “Software is being embedded in material objects, imbuing them with an awareness of their environment, and the calculative

capacities to conduct their own work in the world with only intermittent human oversight.”^[6] These digital agents introduce rules of cause-and-effect into our environment that happen beyond our immediate perception, like a lever that switches far-away railroad tracks. Or, even more puzzling, we might pull a lever that does something different each time, based on some algorithm; or we watch as the algorithm pulls the lever itself, based on its own mysterious motivations.

At the center of all this disruption is how we understand basic elements of our environment: What place am I in? What objects does it contain, and how do they work? Who am I, and who can see me, and what I am doing? What used to be clear is now less so.

Case Study: Facebook Beacon

Some of the infrastructure we take for granted now was almost unimaginable only a decade ago. And perhaps no digital “place” is more ubiquitous in more people’s lives than Facebook. With billions of registered users, it’s become the “telephone network” of social interaction online.

Back in 2007, Facebook launched a service it called *Beacon*, which tracked what users purchased on participating non-Facebook sites, and published that information to the recently introduced *News Feeds* seen by their Facebook “friends.” It took many people by surprise, and sparked a major controversy regarding online privacy.

Facebook is an especially powerful example of context disruption, partly because of how it has shape-shifted the sort of place it is since it began as a closed network for Harvard students alone.

In fact, much of Facebook’s architectural foundation was structured based on the assumption that a user’s network would be limited to people she had already met or could easily meet on her campus. The intrinsic cultural structures of one’s college provided natural boundaries that Facebook re-created in code form.

Over time, Facebook grew rapidly to include other schools, then businesses, and then finally it was opened to the entire Web in 2006. Yet, it wasn’t until much later that it introduced any way of structuring one’s contacts into groups beyond the single bucket of “Friends,” as if everyone you could connect to was the equivalent of someone you met during freshman orientation.

So, for users who had started their Facebook memberships back when their Friends included only their classmates, the sudden shift in context was often disorienting. With pictures of college parties still hanging in their galleries—meant for a social context of peers that would understand them—they were suddenly getting friend invitations from coworkers and family members. Facebook had obliterated the cultural boundaries that had normally kept these facets of one’s personality and personal life comfortably separate.

Before Beacon, the introduction of the News Feed had already caused a lot of concern when users realized it was tracking what they did within Facebook itself and publishing an ongoing status report of those activities to their friends. Actions and changes that had once been quiet adjustments to their profile had been turned into News, published to everyone they knew.

Take, for instance, changes in relationship status. Breaking up with a partner is an intimate, personal event that one might prefer to treat with some subtlety and care. Facebook’s

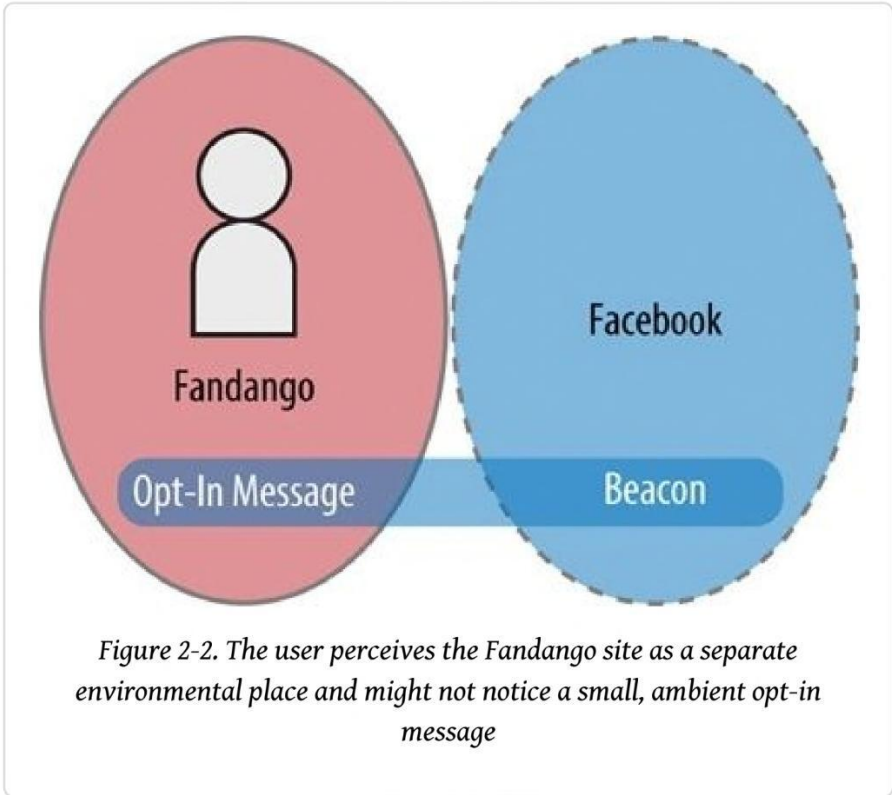
structure made it seem users were changing relationship status within a particular place, separate from other places. Consequently, it was horrifying to discover that changing the setting in a drop-down list in one's personal profile was simultaneously announcing it to everyone he knew. Facebook broke the expectations of cause-and-effect that people bring to their environment.

Just as users were getting used to how the News Feed worked, Beacon launched, publishing information about actions users were taking *outside of Facebook*. Suddenly, Facebook was indiscriminately notifying people of purchases (books about personal matters, medicines for private maladies, or surprise gifts for significant others) and other actions (playing on a video game site during the workday; signing up for a dating site), with confusing contextual clues about what was going on. For example, **Figure 2-1** shows a small opt-out pop-up window that the system used, which was easy to overlook. In addition, it quickly defaulted to “Yes” and disappeared if you didn't acknowledge it in time.^[7]



Figure 2-1. The small Beacon opt-in message that would appear in the lower corner of the screen (from MoveOn.org)^[8]

Unlike one's Facebook profile, this was not information that was already available to your friends; this was information that, in the physical dimension, has always been assumed to be at least implicitly contained within a "store" or "site."



The result? User revolt, widespread controversy, and the eventual dismantling of the Beacon program. And to top it off, it prompted a \$9.5 million class-action lawsuit that was finally settled in February 2013.^[9]

Facebook has notoriously and publically struggled with these issues of place confusion since its founding. But what is true of Facebook is just as true of nearly every networked

environment. Although Beacon was the metaphorical equivalent of having networked cameras and data feeds for your every action available for public consumption, that breakdown of context is no longer merely metaphorical. As our every action and purchase is increasingly picked up by sensors, cameras, brand-loyalty databases, and cloud-connected smartphones, Beacon's misstep seems almost primitive in comparison.

[2] Mitchell, William J. *Placing Words: Symbols, Space, and the City*. Cambridge, MA: MIT Press, 2005: 13, Kindle edition.

[3] ———. *City of Bits: Space, Place, and the Infobahn (On Architecture)*. Cambridge, MA: MIT Press, 1995: 68–9, Kindle locations.

[4] <http://www.w3.org/Proposal.html> (emphasis mine)

[5] <http://mediatedcultures.net/projects/youtube/context-collapse/>

[6] Kitchin, Rob, and Martin Dodge. *Code/Space: Software and Everyday Life*. Cambridge, MA: MIT Press, 2011: 47.

[7] Martin, Kirsten E. "Facebook (A): Beacon and Privacy," *Business Roundtable: Institute for Corporate Ethics*. <http://bit.ly/1nAWYkK>

[8] http://civ.moveon.org/facebookprivacy/beacon_demo.html

[9] Cohen, David. "Ninth U.S. Circuit Court Of Appeals Upholds Facebook Beacon Class-Action Settlement," February 27, 2013. <http://bit.ly/1nAX6Rh>

Chapter 3. Environments, Elements, and Information

Context is worth 80 IQ points.

—ALAN KAY

A Wall and a Field

SOCIAL NETWORKS AND AIRPORTS don't exist in a vacuum. They're part of a wider world. And humans didn't evolve with mobile phones in hand. Our bodies and brains grew in predigital environments that shaped the way we understand our surroundings. If context in digital environments is so hard to get our heads around, maybe we need to begin by establishing up front what an environment actually is. Take a look [Figure 3-1](#), an idyllic landscape in Derbyshire, England.



Figure 3-1. A bucolic Derbyshire landscape^[10]

We might look at this landscape and assume that there's not much information here, but there actually is. This is as much an "information environment" as any website or city intersection. That is, for us to just get around in a place like this, there has to be information about the structures in the environment that our bodies somehow understand well enough to take action. Where can we walk? What can we eat? What can we hold in our hands?

Most of this environment happened all on its own, growing naturally. But also note the stone wall in the field. Some of the earliest structures humans ever added to their environment were of this sort: stones stacked to create barriers and boundaries. Such structures have a physical effect of stopping or slowing terrestrial motion, but they also carry a cultural meaning, such as "this land belongs to someone" or "keep your sheep on the other side, away from

my sheep.” The wall changes the context of the field, transforming an open, undifferentiated vista into a specific human *place* with additional layers of significance.

Keep this field in mind as a starting point, because from here on out, we will be looking at everything we make—digital devices or websites, phone apps or cross-channel services—as structure we add to the environment around us, not unlike this wall, which is deceptively simple at first glance, but full of meaningful layers upon closer reflection.

A Conventional Definition of Context

One challenge we have when grappling with contextual issues comes from our conventional assumptions about what context is to begin with. Yes, we would all agree it has to do with relationships between things. But, we tend to focus on the things, not the context. I mean, we tend to want to understand what a thing is in relation to what contains it, as if that container is “the context.” When you ask whether a wall is an urban wall or a country wall, the subject (wall) is supposedly informed by the setting—a wall “in a field in the country” is a country wall.

This convention is baked into our official dictionary definitions. For example, the *Oxford English Dictionary* defines context as “the circumstances that form the setting for an event, statement, or idea, and in terms of which it can be fully understood.” This definition contains essentially the same elements of most definitions of context:

Circumstances

The setting or situational factors that surround the subject.

Agent

The (implied) entity that is trying to understand the subject.^[11] Note that in this book, you will see the term “agent” used at times, but also words such as “person,” “perceiver,” or “user” somewhat interchangeably, depending on the point being made. They all reference the same concept, except that agents are not always people, or what we normally think of as a “user.”

Subject

The event, statement, or idea that is in the circumstances, and that is the focus of the agent’s attempted understanding. From any particular point of view, there’s always something that’s nested within something else. But, as we’ll see, that relationship is situationally dependent.

Understanding

An apprehension (or effort toward apprehension) of the meaning of the subject and its relationship to its circumstances.

So, the typical way of understanding context looks something like that presented in **Figure 3-2**.

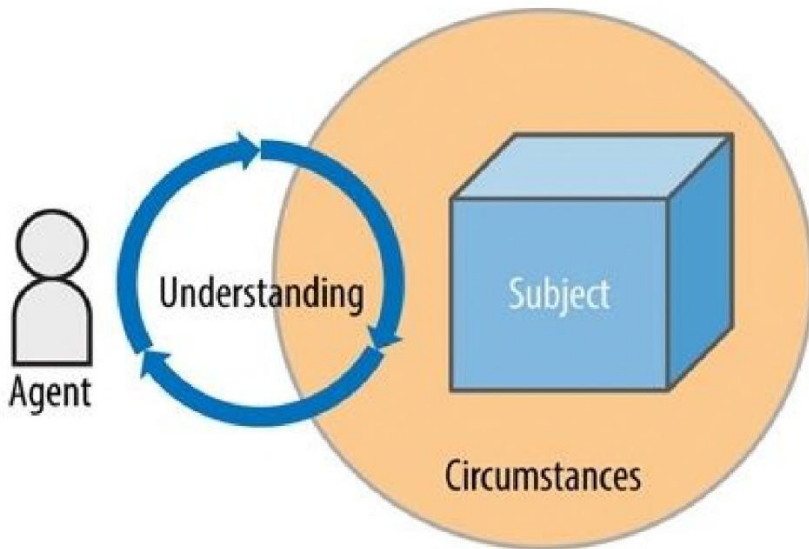


Figure 3-2. The conventional way we think of context

That's not a bad starting point. But the truth is, context is much more complicated. For one thing, as soon as something is "in a context," it changes the context. The same elements are present, but the number of the elements and their relationships are part of a multifaceted reality.

Perhaps the field has only a stone wall in it. But, the wall actually changes the nature of the field, which is now not just "a field" but "a field with a stone wall." Then, it's only a matter of time before more walls are built, roads are added, hotels and pubs spring up, and eventually it's a town with a road called "Field Avenue" as the only reminder that there was ever a field there at all.

As Malcolm McCullough explains in *Digital Ground*, context is bound up in our interaction with our environment.

“Context” is not the setting itself, but the engagement with it as well as the bias that setting gives to the interactions that occur within it.

“Environment” is the sum of all present contexts.^[12]

Context isn't just the surrounding circumstances, because it includes and interacts with the subject that is surrounded, and the agent that tries to comprehend it all.

Paul Dourish, in his seminal paper on context and human-computer interaction, “What We Talk About When We Talk About Context,” argues for a model in which, “Context isn't something that describes a setting; it's something that people do....It is an emergent feature of the interaction, determined in the moment and in the doing. In other words, context and...activity...cannot be separated. Context...arises from and is sustained by the activity itself.”^[13]

This is all sort of brain-bendy stuff, but unless we grapple with it, we run the risk of designing environments that assume too much about how agents understand them.

For example, my Google calendar via website has one set of information in it, including the TripIt calendar I've subscribed to, so that my travel information shows up among all my other scheduled events. However, that's only from my perspective. Others on my team can see only my main calendar, not my calendar subscriptions, a state of affairs that isn't clearly apparent to me. The only indication of this difference is a setting in Google Calendar that takes some effort to discover: it says “Anyone can: See Nothing.”

Even though it's just one setting, the rule has implications that are harder to perceive—it means the circumstances are different depending on which agent is logged in and viewing the interface. The sort of place my calendar is to me is not the sort of place my calendar is to someone else.

As mentioned earlier, the contextual problems have to do not just with the calendar’s settings or even its graphical interface, but the meaning of “calendar” itself. The label we use for the digital object sets an expectation that it is a singular calendar such as a paper calendar hanging on a wall, not a multidimensional, virtual object with many “calendars” in aggregate. The language we use for the environment is part of the environment as well.

So, we at least have to allow that the circumstances are not separate from the agent that’s perceiving and trying to understand the subject; instead, the circumstances contain all these elements, as depicted in **Figure 3-3**.

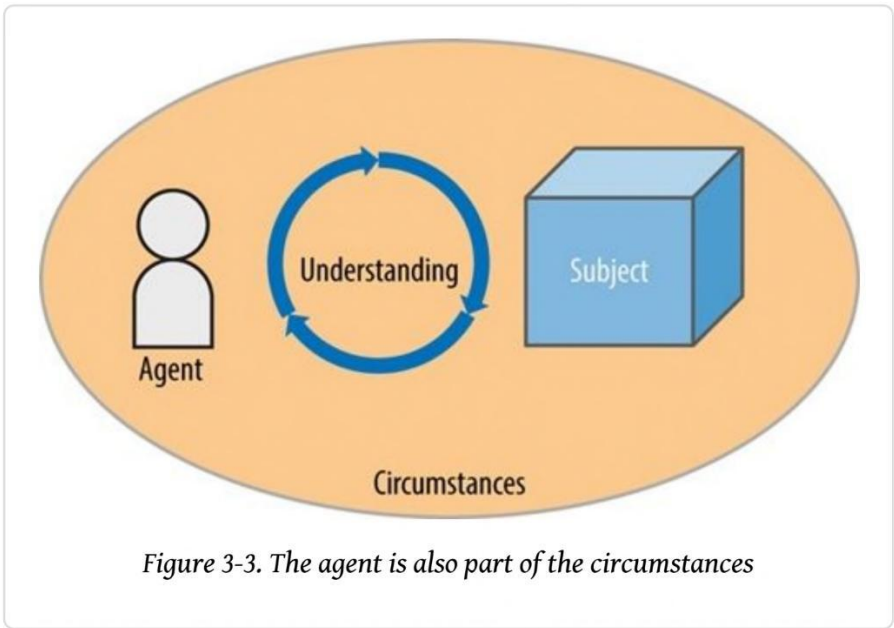


Figure 3-3. The agent is also part of the circumstances

But, what about circumstances in which the agent is also the subject (**Figure 3-4**)? For example, that’s what I was trying to understand when I was at the airport, working out which

label applied to me in the overall system of labeling. And it's what any of us are doing when we're navigating within a space, trying to get somewhere—whether on-screen or off.

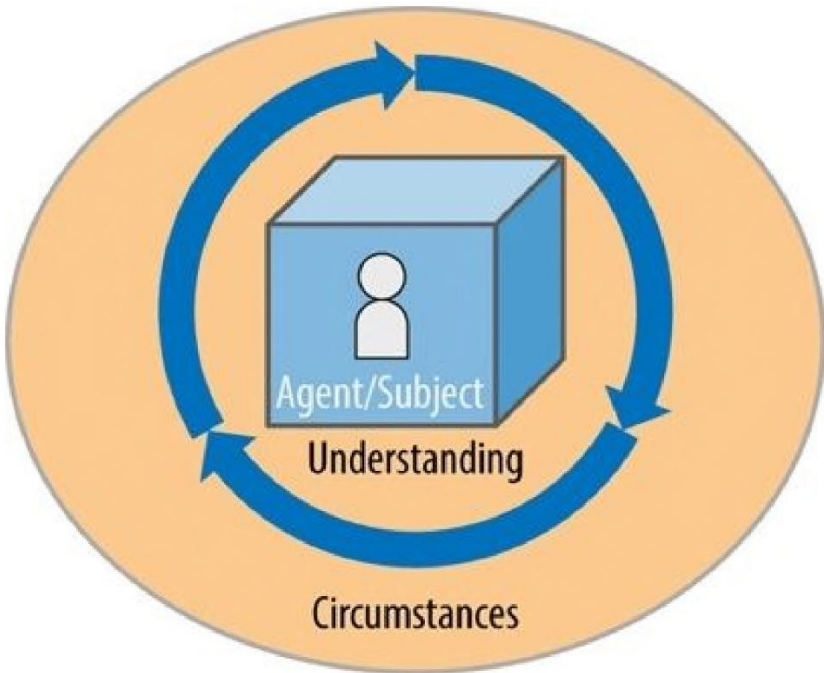


Figure 3-4. The agent is often also the subject

Add to this the fact that we're pretty much always trying to understand who and where we are in relation to our environment while simultaneously needing to understand the context of a lot of other things happening around us.

One moment, something is the *subject* but at another it's just background, part of the *circumstances*. In addition, we're almost never trying to take in only one subject at a time; instead, we're absorbing a shifting, roiling torrent of them, as

illustrated in **Figure 3-5**. It all starts to get pretty overwhelming.

What a mess, right? Humans are pretty smart, but we're still finite creatures. We can really think hard about only one or two things at a time. So, a factor in all this complexity is the amount of attention we can summon and the cognitive effort we have to bring to bear on our environment and everything it's made of, including ourselves.

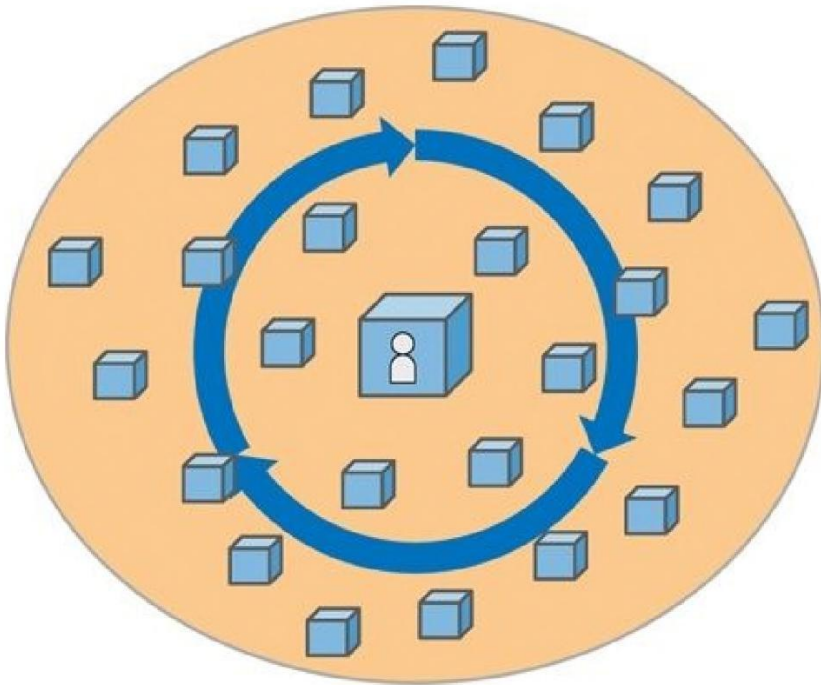


Figure 3-5. An agent in the midst of multiple subjects, and complicated circumstances

A New, Working Definition of Context

So, moving forward, we will stray from the conventional definition and use a new, working definition. It's just a bit more technical than the one we started with in **Chapter 1**:

Context is an agent's understanding of the relationships between the elements of the agent's environment.

In this case, the parts of the definition are as follows:

Agent

A person or other object that can act in the environment. Not all agents are persons, and not all of them perceive or act the same way.

Understanding

An agent's cognitive engagement with, and making sense of, its surroundings. Although there are other non-agent-bound ways of thinking about context, our purpose here has to do with the subjective, first-person experience of the agent, even when it's not a person.

Relationships between the elements

Everything is made up of parts; context is all about how those parts relate to one another.

The agent's environment

Context is always about the entire environment, because the environment is what informs the meaning of anything the agent is trying to understand. Note we aren't saying "environment" in general—a setting observed from some omniscient, god's-eye view. It is "the agent's environment," because context is a function of the agent's

own first-person perception. Perception is what undergirds cognition, experience, and understanding.

This definition makes context a function of how an agent perceives and understands the environment, not a property that exists outside of that understanding. It also doesn't specify that one element is the subject, special from everything else. Any element can be the subject at any moment and then drift or quickly switch to mere "circumstances."

So, does that mean we can't pin context down? Well, in a way, yes: if by pinning it down we mean we try to fix it in one state. That is, we can still map the elements out as agent, subjects, and circumstances; we just have to be careful that we aren't assuming that only one snapshot represents the whole contextual experience.

In my airport scenario in [Chapter 1](#), we could do a pretty decent job of mapping the most important contextual relationships by identifying the elements:

- Each element that can reasonably be called an *agent*: myself, my coworker, each application on my phone that's making some kind of decision and taking action, and even the customer service representative.
- Each important element that at any moment could be a *subject* in the scenario: everything that's an agent, plus my boarding pass, the security queues, my calendar, and so on.
- Finally, the *circumstances*, which are basically each influential element of the environment, including (of course) the subjects and agents. Any element could be a circumstance or could be a subject, and vice versa. It's all

a matter of what the agent needs to perceive and understand in order to act.

We can get a lot of value out of just listing out these different relationships, and how they crisscross and overlap. There's no way to map every single factor in even a simple real-world environment, but it's possible to take snapshots from different perspectives, at various key moments, and bring them together into something more like a *collage* of snapshots that come nearer to telling the entire story. The important thing is to include more than a single, static perspective. What does the Passbook app understand about my situation that causes it to show a notification in the lock screen? What does my coworker understand about my calendar and the absence of any trip information there? And so on.

I offer this way of thinking about modeling context as a starting point. But, before anything else, we still have to get a handle on what we mean by things such as “understanding,” “environment,” and “agent.” It turns out that understanding the elements of the environment depends on the way *information* works.

Modes of Information

Even though it's not listed in our working definition of context, *information* is required before any perceiving or understanding can occur. Before around 1960, people didn't use the word nearly as much as they do today.^[14] The rise of computing technologies influenced how we talk about the way we communicate. Now, we say “information” a great deal. We use it as a catchall for just about anything that is intended to communicate a message or meaning or

knowledge. It's not as general as a word like "stuff," but it's close. Is it expected to inform? Then it's information. Does it do a good job of informing? If it does, it's "good information." If not, it's "bad information."

And it would seem that bad information is everywhere, because the expectations we bring to information are often met with disappointment. Most people would categorize the legalese in a software license or the instructions for a tax form as information—but they would be hard-pressed to agree they fully understand those texts. They just say, "this information is awful" or "I don't get this information at all."

It does seem odd that we'd even use the word "information" for something that doesn't effectively do the work of informing. In *Information Anxiety*, Richard Saul Wurman explains how the term gained an additional conventional meaning around the time of World War II, with the invention of information theories behind electronic transmission. The word "became part of the vocabulary of the science of messages. And, suddenly, the appellation could be applied to something that didn't necessarily have to inform. The definition was extrapolated to general usage as something told or communicated, whether or not it made sense to the receiver."^[15] This slippage influenced our techno-fueled mid-twentieth-century culture to the point that we now talk of living in the "Information Age." Now, *information* can mean the coding of DNA, the credits at the end of a movie, or verbal instructions from a customer service desk. Like so many words that can mean so many things, information is now a muddle.

For regular people in everyday situations, rigorous definitions for the word might not be necessary. But for people whose job it is to make environments with

information, a specialized understanding is called for. Psychologists need to understand what elements make up sadness so that they can help people who are troubled by it. Artists and art historians are responsible for creating and curating art, so they have to think about the materials, techniques, and cultural meanings behind what makes something “art.” And game designers make a living from being deeply interested in what constraints, goals, and mechanisms make something “fun.” Similarly, it’s the responsibility of those of us who make “information things”—from software and digital gadgets to diagrams and newsletters—to have a practical, working model for how information works, and how it can meet the everyday expectations of its users.

You can find many sophisticated models and explanations of information in academic theories and scientific literature.^[16] They all bring valid and valuable perspectives to bear. For our purposes here, though, I’ve developed a simple model that describes three different modes for information that provide lenses for understanding how context works in the environments we design. We’ll begin with a summary, here, but then delve more deeply into each mode in the next three parts of the book.

Figure 3-6 shows the three modes: Digital, Semantic, and Physical.

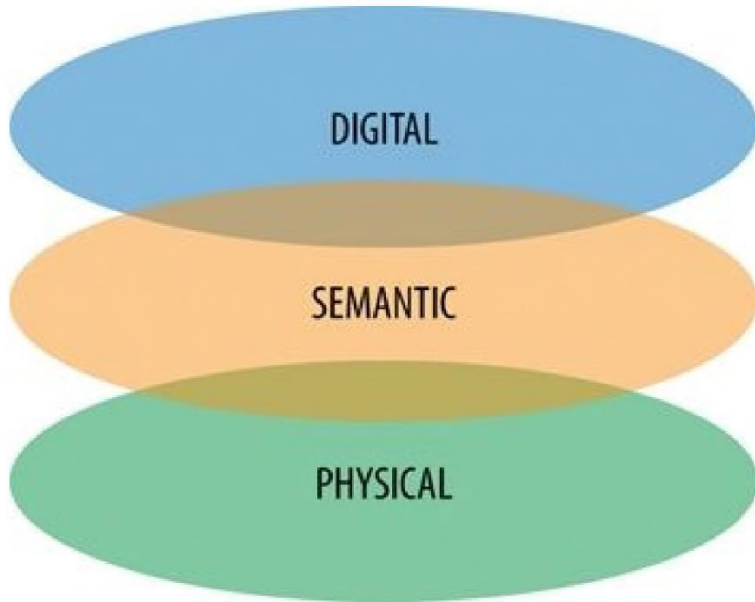


Figure 3-6. The three modes of information

Let's begin from the bottom and work our way up.

Physical

This is the information animals use to perceive their environment for purposes of taking physical action; of course, humans are a part of the animal realm, so we do this, too. It's information that functions "ecologically"—that is, in the relationship between a creature and its environment. It's about surfaces, edges, substances, objects, and the ways in which those things contain and relate to one another to support animals' behaviors and survival.

Semantic

This is information people create for the purpose of communicating meaning to other people. I'll often refer to this as "language." For our discussion, this mode includes all sorts of communication such as gestures, signs, graphics, and of course speech and writing. It's more fluid than physical information and harder to pin down, but it still creates environmental structure for us. It overlaps the Physical mode because much of the human environment depends on complementary qualities of both of these modes, such as the signage and maps positioned in physical locations and written on physical surfaces in an airport.

Digital

This is the "information technology" sort of information by which computers operate, and communicate with other computers. Even though humans created it (or created the computers that also create it), it's not natively readable by people. That's because it works by stripping out people-centric context so that machines can talk among one another with low error rates, as quickly and efficiently as possible. It overlaps the Semantic mode, because it's abstract and made of encoded semantic information. But even though it isn't literally physical, it does exist in physical infrastructure, and it does affect our physical environment more and more every day.

I should mention: like many other models I'll share, this one isn't meant to be taken as mathematically or logically exact. Simple models can sometimes work best when they are clear enough to point us in the right direction but skip the complexities of precision. So, for example, the overlapping parts of the modes are there to evoke how they are seldom

mutually exclusive, and actually influence one another.

Starting from the Bottom

I began with the Physical mode for a reason. Context is about whole environments; otherwise, we are considering parts of an environment out of context. And when we take an environmental view, we have to begin from first principles about how people understand the environment, whether there are digital networks or gadgets in it or not. To illustrate this, **Figure 3-7** presents another informal model that illustrates the layers involved in the discussion moving forward.

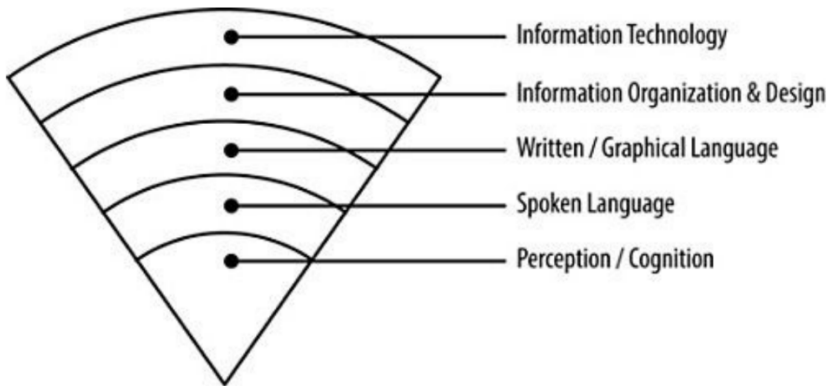


Figure 3-7. Pace layers of information

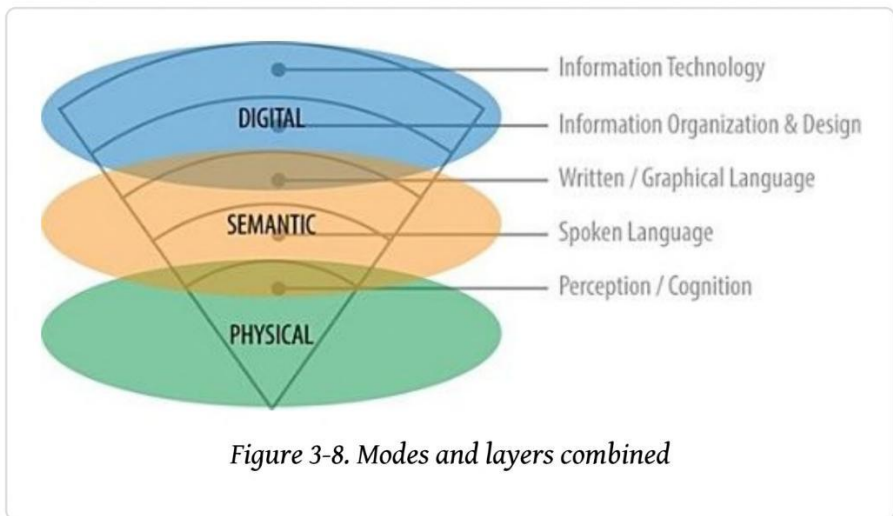
These are based on a concept known as *pace layers*—where the lower-level layer changes more slowly over time than the next, and so on. I've adapted the approach so that it also implies how one layer builds on another.^[17]

- *Perception and cognition* change very slowly for all organisms, including humans, and these abilities had to evolve long ago in order to have a functioning animal to begin with. Perception here means the core faculties of how a body responds to the environment. This is the sort of perception lizards use to climb surfaces and eat bugs or that humans use to walk around or duck a stray football.
- *Spoken language* is next; as we will see, it has been around for a very long time for our species, long enough to at times be hard to separate from the older perception and cognition capabilities of our bodies and brains (as mentioned earlier, I'm lumping gestures in with speech, for simplicity's sake). Even though particular languages can change a lot over centuries, the essential characteristics of spoken language change much more slowly.
- *Written/graphical language* is the way we use physical objects—up until very recently, the surfaces of those objects—to encode verbal language for communicating beyond the present moment. Although spoken language is more of an emergent property of our species, writing is actually more of an ancient *technology*. Writing is also a way of encoding information, which is a precursor to digital code.
- *Information organization and design* arose as identifiable areas of expertise and effort because we had so much stuff written down, and because writing and drawing enabled greater complexity than is possible in mere speech. The ability to freeze speech on a surface and relate it to other frozen speech on the same surface opened up the ability to have maps, diagrams, broadsides, folios, all of which

required organization and layout. Our methods for organizing and designing written information have also been precursors to how we've designed and organized digital information for computing.

- Last, there's *information technology*, which is quite recent, and (as I'm defining it here) depended on the invention of digital software. We've seen this mode change rapidly in our own lifetimes, and it's the layer that has most disrupted our experience of the other two modes, in the shortest time. It didn't happen on its own, however; the ideas behind it originated in writing, linguistic theory, and other influences from further down the model.

If we place the three modes of information on top of these layers as demonstrated in **Figure 3-8**, it gives a rough idea of how these models relate to each other.



In reality, the boundaries are actually much more diffuse and intermingled, but the main idea is hopefully clear: the ways in

which we use and perceive information have evolved over time; some aspects are foundational and more stable, whereas other aspects are more variable and quick to change.

In my experience, most technological design work begins with information technology first and then later figures out the information organization and design and other communicative challenges lower down. Yet, starting with technology takes a lot for granted. It assumes X means X, and Y means Y; or that here is here, and there is there. What happens when we can no longer trust those assumptions? The best way to untangle the many knotted strands that create and shape context is to understand how the world makes sense to us in the first place—with bodies, surfaces, and objects—and build the rest of our understanding from that foundation.

[10] Wikimedia Commons: <http://bit.ly/1uDL7m6>

[11] Using the term *agent* gives us the ability to include nonpersons, such as software or other systems that try to determine context. It's also the term used most often in the scholarly literature for this element.

[12] McCullough, Malcolm. *Digital Ground: Architecture, Pervasive Computing, and Environmental Knowing*. Cambridge, MA: MIT Press, 2004: 48, Kindle edition.

[13] Dourish, Paul. "What We Talk About When We Talk About Context." *Personal and Ubiquitous Computing*. London: Springer-Verlag, February 2004; 8(1):19–30.

[14] Based on a search for "information" in books from 1800 to 2000, using Google's Ngram Viewer (<https://books.google.com/ngrams/>).

[15] Wurman, Richard Saul. *Information Anxiety*. New York: Doubleday, 1989: 38.

[16] I especially recommend Bates, M. “Fundamental Forms of Information.” *Journal of the American Society for Information Science and Technology* 2006; 57(8):1033–45, and ongoing work on a taxonomy of information by Sabrina Golonka, (<http://bit.ly/1ySrrik> and <http://bit.ly/1CM2ti6>).

[17] Borrowed and adapted from the work of Stewart Brand, particularly in *How Buildings Learn*, who adapted his approach from a concept called *shearing layers* created by architect Frank Duffy.

Part II. Physical Information

The Roots of Context

THE PRODUCTS AND SERVICES WE DESIGN ARE PART OF A GREATER ENVIRONMENT, but they have the capacity to change that environment as well as the behaviors of people who use them. Smartphones influence user behavior in a different way than older cell phones, which in turn changed behavior from when only phone booths and land lines were available. Obviously, right? But, did you know that the separation between the environment, the object, and the user is mostly artificial and that they're all part of one dynamic system? Have you also noticed that when we use software, our perception seems to expect that environment to behave according to the same laws we rely on in the physical world?

The mechanics of physical life shape the way we understand abstractions such as language, social systems, and software. That's why we're going to spend some significant quality time together looking at what I'm calling *Physical Information*, the mode at the bottom of the diagram shown in Figure II-1.

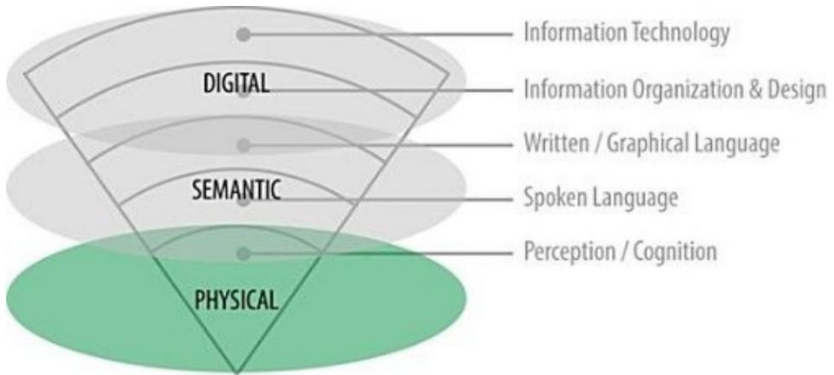


Figure II.1. Physical Information

Part II introduces some essential theories about perception and action. It explores what *affordance* really is, with special attention to how it was originally conceived by its creator, James J. Gibson. It also covers how the environment influences behavior, how memory and learning work, and offers models for breaking down the elements of any environment. Finally, this part shows how physical information principles translate into more-complex parts of our world such as social culture and organizations.

Chapter 4. Perception, Cognition, and Affordance

In the Universe, there are things that are known, and things that are unknown, and in between there are doors.

—WILLIAM BLAKE

Information of a Different Sort

IF WE ARE TO KNOW HOW USERS UNDERSTAND THE CONTEXT OF OBJECTS, people, and places, we need to stipulate what we mean by *understand* in the first place. The way people understand things is through *cognition*, which is the process by which we acquire knowledge and understanding through thought, experience, and our senses. Cognition isn't an abstraction. It's bound up in the very structures of our bodies and physical surroundings.

When a spider quickly and gracefully traverses the intricacies of a web, or a bird like the green bee-eater on this book's cover catches an insect in flight, these creatures are relying on their bodies to form a kind of coupling with their environments—a natural, intuitive dance wherein environment and creature work together as a system. These wonderfully evolved, coupled systems result in complex, advanced behavior, yet with no large brains in sight.

It turns out that we humans, who evolved on the same planet among the same essential structures as spiders and birds, also

rely on this kind of body-to-environment coupling. Our most basic actions—the sort we hardly notice we do—work because our bodies are able to perceive and act among the structures of our environment with little or no thought required.

When I see users tapping and clicking pages or screens to learn how a product works, ignoring and dismissing pop-ups with important alerts because they want to get at the information underneath, or keeping their smartphones with them from room to room in their homes, I wonder why these behaviors occur. Often they don't seem very logical, or at least they show a tendency to act first and think about the logic of the action later. Even though these interfaces and gadgets aren't natural objects and surfaces, users try using them as if they were.

This theory about the body-environment relationship originates in a field called *ecological psychology*, which posits that creatures directly perceive and act in the world by their bodies' ability to *detect information* about the structures in the environment. This information is what I will be calling *physical information*—a mode of information that is at work when bodies and environments do this coupled, dynamic dance of action and perception.

Ecological psychology is sometimes referred to as *Gibsonian* psychology because the theory started with a scientist named James J. Gibson, whose theory of information uses neither the colloquial meaning of information nor the definition we get from information science.^[18] Gibson explains his usage in a key passage of his landmark work, *The Ecological Approach to Visual Perception*:

Information, as the term is used in this book (but not in other books), refers to specification of the observer's environment, not to specification of the observer's receptors or sense organs....[For discussing perception, the term] information cannot have its familiar dictionary meaning of knowledge communicated to a receiver. This is unfortunate, and I would use another term if I could. The only recourse is to ask the reader to remember that picking up information is not to be thought of as a case of communicating. The world does not speak to the observer. Animals and humans communicate with cries, gestures, speech, pictures, writing and television, but we cannot hope to understand perception in terms of these channels; it is quite the other way around. Words and pictures convey information, carry it, or transmit it, but the information in the sea of energy around each of us, luminous or mechanical or chemical energy, is not conveyed. It is simply there. The assumption that information can be transmitted and the assumption that it can be stored are appropriate for the theory of communication, not for the theory of perception.^[19]

Gibson often found himself having to appropriate or invent terms in order to have language he could use to express ideas that the contemporaneous language didn't accommodate.^[20] He's having to ask readers to set aside their existing meaning of *information* and to look at it in a different way, when trying to understand how perception works. For him, "To perceive is to be aware of the surfaces of the environment and of oneself in it."^[21] In other words, perception is about the agent, figuring out the elements of its surroundings and understanding how the agent itself is one of those elements. And information is what organisms perceive in the environment that informs the possibilities for action.

Even this usage of "perception" is more specific than we might be used to: it's about core perceptual faculties, not statements such as "my perception of the painting is that it is pretty" or "the audience perceives her to be very talented." Those are cultural, social layers that we might refer to as

perception, but not the sort of perception we will mainly be discussing in Part 1.

Even though we humans might now be using advanced technology with voice recognition and backlit touch-screen displays, we still depend on the same bodies and brains that our ancestors used thousands of years ago to allow us to act in our environment, no matter how digitally enhanced. Just as with the field and the stone wall presented in **Chapter 3**, even without language or digital technology, the world is full of structures that *inform* bodies about what actions those structures afford.

I'll be drawing from Gibson's work substantially, especially in this part of the book, because I find that it provides an invaluable starting point for rethinking (and more deeply grasping) how users perceive and understand their environments. Gibson's ideas have also found a more recent home as a significant influence in a theoretical perspective called *embodied cognition*.

JAMES. J. GIBSON

James J. (“JJ”) Gibson (1904-1979) was an American experimental psychologist, author, and theorist. He and his wife, Eleanor J. Gibson (1910-2002),—a major scientific figure in her own right—developed an extensive theoretical body of work on what they called ecological perception and learning.



Figure 4-1. James J. Gibson and Eleanor Gibson^[22]

James Gibson developed his theories partly during research funded by the United States Air Force around the time of World War II while studying how pilots orient themselves during flight. ^[23] Gibson realized that his insight would mean overturning more than a century of established scientific research to get to the bottom of the problem, and insisted that a “fresh start” was required. ^[24] What resulted was decades of work dedicated to changing the way science understood perception.

Gibson particularly subscribed to the perspective of American Pragmatism, and the radical empiricism developed by William James. ^[25] As a radical empiricist himself, Gibson insisted on understanding perception based on the facts of the natural world, versus cultural assumptions or artificially contrived experiments.

Eleanor Gibson made major contributions to the science of childhood cognitive development as well as how people in general learn new knowledge. Her work has been foundational to later social science and psychology work on education and communities of practice. She was awarded the National Medal of Science in 1992. A famous experiment she created was the *Visual*

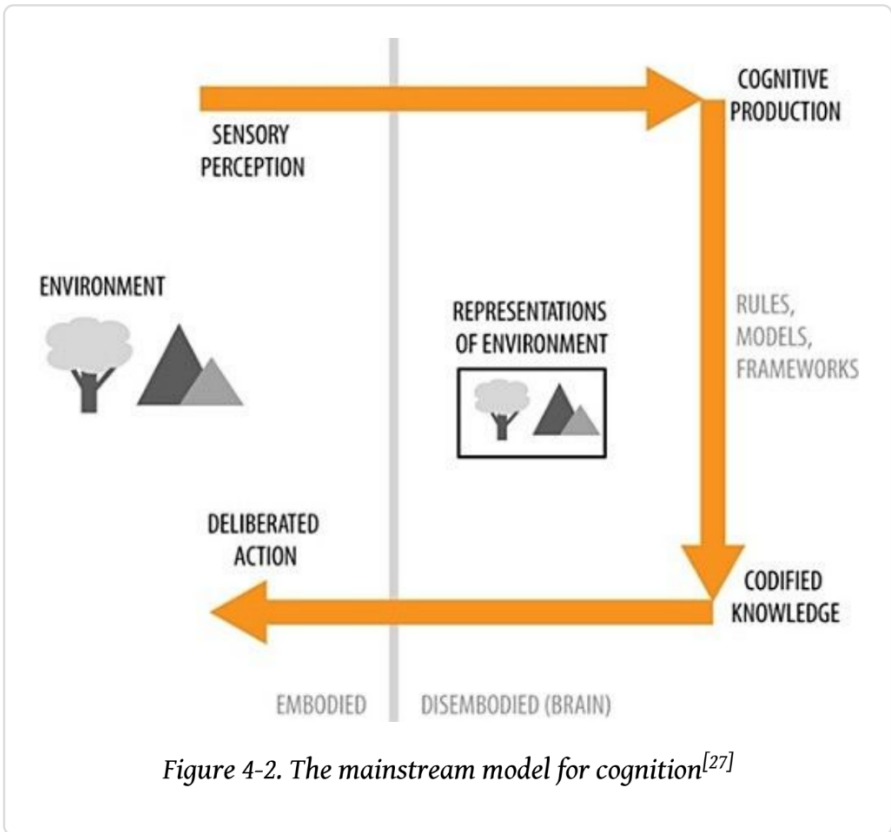
Cliff, in which infants were placed on a wooden table whose surface was extended by a long portion of plate glass. She discovered that infants reacted to the perceived drop-off with caution or anxiety, but she also observed how they would adapt to their perceptions by patting the glass and learning about their environment through action.^[26]

Most of the references to “Gibsonian psychology” in this book are specifically to James Gibson’s work; but it’s important to remember that this amazing couple jointly established some of the most important insights in psychological science in the twentieth century.

A Mainstream View of Cognition

Since roughly the mid-twentieth century, conventional cognitive science holds that cognition is primarily (or exclusively) a brain function, and that the body is mainly an input-output mechanism; the body does not constitute a significant part of cognitive work. I’ll be referring to this perspective as *mainstream* or *disembodied cognition*, though it is called by other names in the literature, such as *representationalism* or *cognitivism*.

According to this view, cognition works something similar to the diagram in **Figure 4-2**.



The process happens through inputs and outputs, with the brain as the “central processing unit”:

1. The brain gathers data about the world through the body’s senses.
2. It then works with *abstract representations* of what is sensed.
3. The brain processes the representational data by using disembodied rules, models, and logical operations.
4. The brain then takes this disembodied, abstract

information and translates it into instructions to the body.

In other words, according to the mainstream view, perception is indirect, and requires representations that are processed the way a computer would process math and logic. The model holds that this is how cognition works for even the most basic bodily action.

This computer-like way of understanding cognition emerged for a reason: modern cognitive science was coming of age just as information theory and computer science were emerging as well; in fact, the “cognitive revolution” that moved psychology past its earlier behaviorist orthodoxy was largely due to the influence of the new field of information science.

[28]

So, of course, cognitive science absorbed a lot of perspectives and metaphors from computing. The computer became not just a metaphor for understanding the brain, but a literal explanation for its function. It framed the human mind as made of information processed by a brain that works like an advanced machine.^[29] This theoretical foundation is still influential today, in many branches of psychology, neuroscience, economics, and even human-computer interaction (HCI).

To be fair, this is a simplified summary, and the disembodied-cognition perspective has evolved over time. Some versions have even adopted aspects of competing theories. Still, the core assumptions are based on brain-first cognition, arguing that at the center is a “model human processor” that computes our cognition using logical rules and representations, much like the earliest cognitive scientists and HCI theorists described.^[30] And let’s face it, this is how

most of us learned how the brain and body function; the brain-is-like-a-computer meme has fully saturated our culture to the point at which it's hard to imagine any other way of understanding cognition.

The mainstream view has been challenged for quite a while by alternative theories, which include examples such as *activity theory*, *situated action theory*, and *distributed cognition theory*.^[31] These and others are all worth learning about, and they all bring some needed rigor to design practice. They also illustrate how there isn't necessarily a single accepted way to understand cognition, users, or products. For our purposes, we will be exploring context through the perspective of *embodied cognition theory*.

Embodied Cognition: An Alternative View

In my own experience, and in the process of investigating this book's subject, I've found the theory of embodied cognition to be a convincing approach that explains many of the mysteries I've encountered over my years of observing users and designing for them.

Embodied cognition has been gaining traction in the last decade or so, sparking a paradigm shift in cognitive science, but it still isn't mainstream. That's partly because the argument implies mainstream cognitive science has been largely wrong for a generation or more. Yet, embodied cognition is an increasingly influential perspective in the user-experience design fields, and stands to fundamentally change the way we think about and design human-computer interaction.^[32]

Generally, the embodiment argument claims that our brains

are not the only thing we use for thought and action; instead, our bodies are an important part of how cognition works. There are multiple versions of embodiment theory, some of which still insist the brain is where cognition starts, with the body just helping out. However, the perspective we will be following argues that cognition is truly environment-first, emerging from an active relationship between environment, body, and brain.^[33] As explained by Andrew Wilson and Sabrina Golonka in their article “Embodied Cognition Is Not What You Think It Is”:

The most common definitions [of embodied cognition] involve the straightforward claim that “states of the body modify states of the mind.” However, the implications of embodiment are actually much more radical than this. If cognition can span the brain, body, and the environment, the “states of mind” of disembodied cognitive science won’t exist to be modified. Cognition will instead be an extended system assembled from a broad array of resources. Taking embodiment seriously therefore requires both new methods and theory.^[34]

The embodied approach insists on understanding perception from the first-person perspective of the perceiving organism, not abstract principles from a third-person observer. A spider doesn’t “know” about webs, or that it’s moving vertically up a surface; it just takes action according to its nature. A bird doesn’t “know” it is flying in air; it just moves its body to get from one place to another through its native medium. For we humans, this can be confusing, because by the time we are just past infancy, we develop a dependence on language and abstraction for talking and thinking about how we perceive the world—a lens that adds a lot of conceptual information to our experience. But the perception and cognition underlying that higher-order comprehension is just about bodies and structures, not concepts. Conscious reflection on our

experience happens after the perception, not before.

How can anything behave intelligently without a brain orchestrating every action? To illustrate, let's look at how a Venus flytrap "behaves" with no brain at all. Even though ecological psychology and embodiment are not about plants, but terrestrial animals with brains and bodies that move, the flytrap is a helpful example because it illustrates how something that seems like intelligent behavior can occur through a coupled action between environment and organism.

The Venus flytrap (**Figure 4-3**) excretes a chemical that attracts insects. The movement of insects drawn to the plant then triggers tiny hairs on its surface. These hairs *structurally* cause the plant to close on the prey and trap it.



Figure 4-3. A Venus flytrap—complex behavior without a brain^[35]

This behavior already has some complexity going on, but there's more: the trap closes only if more than one hair has been triggered within about 20 seconds. This bit of conditional logic embodied by the plant's structure prevents it from trapping things with no nutritional value. Natural selection evidently filtered out all the flytraps that made too many mistakes when catching dinner. This is complex

behavior with apparent intelligence underpinning it. Yet it's all driven by the physical coupling of the organism's "body" and a particular environmental niche.

Now, imagine an organism that evolved to have the equivalent of millions of Venus flytraps, simple mechanisms that engage the structures of the environment in a similar manner, each adding a unique and complementary piece to a huge cognition puzzle. Though fanciful, it is one way of thinking about how complex organisms evolved, some of them eventually developing brains.

In animals with brains, the brain enhances and augments the body. The brain isn't the center of the behavioral universe; rather, it's the other way around. It's this "other way around" perspective that Gibson continually emphasizes in his work.

Figure 4-4 illustrates a new model for the brain-body-environment relationship.

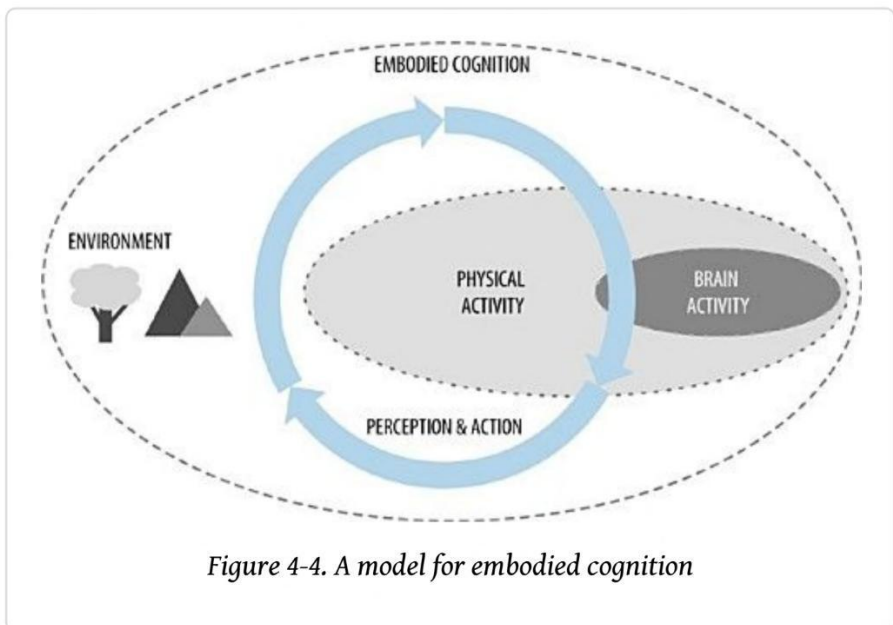


Figure 4-4. A model for embodied cognition

In this model, there's a continuous loop of action and perception in which the entire environment is involved in how a perceiver deciphers that environment, all of it working as a dynamical, perceptual system. Of course, the brain plays an important role, but it isn't the *originating source* of cognition. Perception gives rise to cognition in a reciprocal relationship—a resonant coupling—between the body and its surroundings.

This perception-action loop is the dynamo at the center of our cognition. In fact, perception makes up most of what we think of as cognition to begin with. As Louise Barrett puts it in *Beyond the Brain: How Body and Environment Shape Animal and Human Minds* (Princeton University Press), “Once we begin exploring the actual mechanisms that animals use to negotiate their worlds, it becomes hard to decide where ‘perception’ ends and ‘cognition’ starts.”^[36] Just perceiving the environment's information already does a lot of the work that we often attribute to brain-based cognition.

Embodiment challenges us to understand the experience of the agent not from general abstract categories, but through the lived experience of the perceiver. One reason I prefer this approach is that it aligns nicely with what user experience (UX) design is all about: including the experiential reality of the user as a primary input to design rather than relying only on the goals of a business or the needs of a technology. Embodied cognition is a way of understanding more deeply how users have experiences, and how even subtle changes in the environment can have profound impacts on those experiences.

USING THE ENVIRONMENT FOR THINKING

Designing for digital products and services requires working with a lot of abstractions, so it's helpful to bring those abstractions out of the cloudy dimension of pure thought and into the dimension of physical activity. This is why we so often find ourselves using our bodily environment for working out design problems and why it's emerged as a recognized best practice.

As an example let's consider how we use an office stalwart: the sticky note. By using sticky notes, we can move language around on a physical surface. As we'll see, language makes it possible for us to use bits of semantic information (labels, phrases, icons) as stand-ins for what they represent—anything from simple objects to large, complex ideas. By using the physical surface and the uncannily just-sticky-enough adhesive on the notes ([Figure 4-5](#)), we not only make use of the spatial relationships between notes to discover affinities and create structures, but also engage our bodies in thinking through the problem.

Sketching is another way we can externalize thought into bodily engagement with our environment—whether we're working through diagrammatic models to discover and rehearse abstract structural relationships, or we're informally playing around with representations of actual objects and interfaces. Sketching isn't only about what's being put on paper or drawn on-screen; sketching also engages our bodies in working through the contours of structure and potential action. Sketching can come in many forms, from chalk on a blackboard to CAD drawings or “wireframes” to making quick-and-cheap physical prototypes.



Figure 4-5. Using sticky notes to work through abstractions

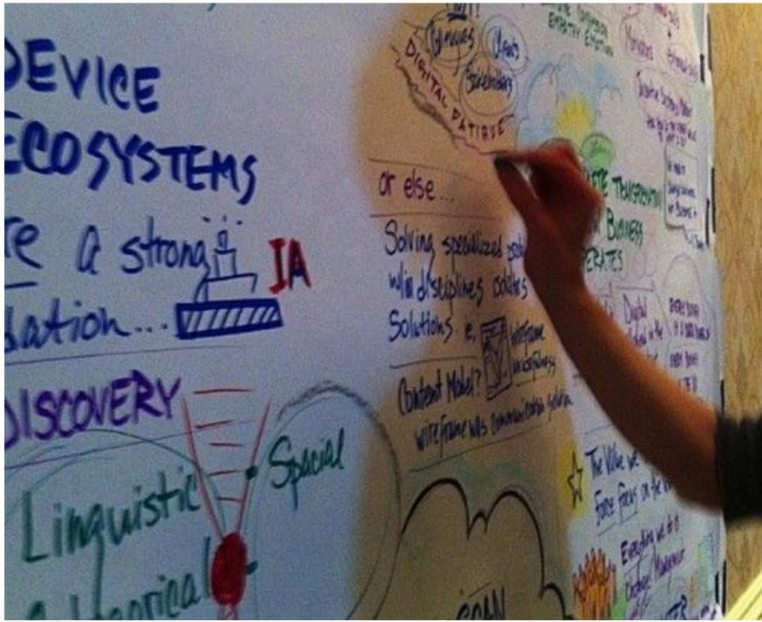


Figure 4-6. Kate Rutter, live-sketchnoting Karen McGrane's closing plenary at IA Summit 2013^[37]

Action and the Perceptual System

As shown in the perception-action loop of **Figure 4-4**, we understand our environment by *taking action* in it. Gibson stresses that perception is not a set of discrete inputs and outputs, but happens as a *perceptual system* that uses all the parts of the system at once, where the distinction between input and output is effaced so that they “are considered together so as to make a continuous loop.”^[38] The body plays an active part in the dynamical feedback system of perception. Context, then, is also a result of action by a

perceiving agent, not a separate set of facts somehow insulated from that active perception. Even when observing faraway stars, astronomers' actions have effects on how the light that has reached the telescope is interpreted and understood.

It's important to stress how deeply physical action and perception are connected, even when we are perceiving "virtual" or screen-based artifacts. In the documentary *Visions of Light: The Art of Cinematography*, legendary cameraman William Fraker tells a story about being the cinematographer on the movie *Rosemary's Baby*. At one point, he was filming a scene in which Ruth Gordon's sly-yet-sinister character, Minnie, is talking on the phone in a bedroom. Fraker explains how director Roman Polanski asked him to move the camera so that Minnie's face would be mostly hidden by the edge of a doorway, as shown in **Figure 4-7**. Fraker was puzzled by the choice, but he went along with it.

Fraker then recounts seeing the movie's theatrical premiere, during which he noticed the audience actually lean to the right in their seats in an attempt to peek around the bedroom door frame. It turned out that Polanski asked for the odd, occluding angle for a good reason: to engage the audience physically and heighten dramatic tension by obscuring the available visual information.

Even though anyone in the theater would have consciously admitted there was no way to see more by shifting position, the unconscious impulse is to shift to the side to get a better look. It's an intriguing illustration of how our bodies are active participants in understanding our environment, even in defiance of everyday logic.



Figure 4-7. Minnie's semi-hidden phone conversation in *Rosemary's Baby* (courtesy Paramount Pictures)^[39]

Gibson uses the phrase *perceptual system* rather than just “the eye” because we don’t perceive anything with just one isolated sense organ.^[40] Perception is a function of the whole bodily context. The eye is made of parts, and the eye itself is a part of a larger system of parts, which is itself part of some other larger system. Thus, what we see is influenced by how we move and what we touch, smell, and hear, and vice versa.

In the specific case of watching a movie, viewers trying to see more of Minnie’s conversation were responding to a virtual experience as if it were a three-dimensional physical environment. They responded this way not because those dimensions were actually there, but because that sort of information was being mimicked on-screen, and taking action—in this case leaning to adjust the angle of viewable surfaces—is what a body does when it wants to detect richer information about the elements in view. As we will see in

later chapters, this distinction between directly perceived information and interpreted, simulated-physical information is important to the way we design interfaces between people and digital elements of the environment.

This systemic point of view is important in a broader sense of how we look at context and the products we design and build. Breaking things down into their component parts is a necessary approach to understanding complex systems and getting collaborative work done. But we have a tendency (or even a cognitive bias) toward forgetting the big picture of where the components came from. A specific part of a system might work fine and come through testing with flying colors, but it might fail once placed into the full context of the environment.

Information Pickup

Gibson coins the phrase *information pickup* to express how perception picks up, or detects, the information in the environment that our bodies use for taking action. It's in this dynamic of information pickup that the environment *specifies* what our bodies can or cannot do. The information picked up is about the mutual structural compatibility for action between bodies and environments.

In the same way a weather vane's "body" adjusts its behavior directly based on the direction of wind, an organism's biological structures respond to the environment directly. A weather vane (**Figure 4-8**) moves the way it does because its structure responds to the movement of air surrounding it. Similarly, the movements of the elbow joint also shown in **Figure 4-8** are largely responses to the structure of the environment. When we reach for a fork at dinner or prop

ourselves at the table, the specifics of our motion don't need to be computed because their physical structure evolved for reaching, propping, and other similar actions. The evolutionary pressures on a species result in bodily structures and systems that fit within and resonate with the structures and systems of the environment.^[41]

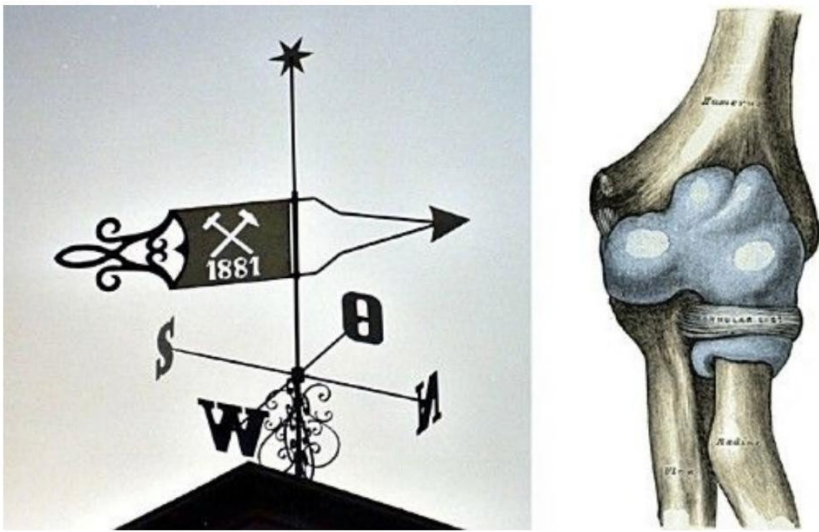


Figure 4-8. Weather vanes^[42] and the human elbow joint^[43] respond to their environments.

Information pickup is the process whereby the body can “orient, explore, investigate, adjust, optimize, resonate, extract, and come to an equilibrium.”^[44] Most of our action is calibrated on the fly by our bodies, tuning themselves in real time to the tasks at hand.^[45] When standing on a wobbly surface, we squat a bit to lower our center of gravity, and our arms shoot out to help our balance. This unthinking reaction

to maintain equilibrium is behind more of our behavior than we realize, as when moviegoers leaned to the side to see more of Minnie's phone conversation. There's not a lot of abstract calculation driving those responses. They're baked into the whole-body root system that makes vision possible, responding to arrays of energy interacting with the surfaces of the environment. Gibson argued that all our senses work in a similar manner.

Affordance

If you've done much design work, you've probably encountered talk of *affordances*. The concept was invented by J.J. Gibson, and codeveloped by his wife, Eleanor. Over time, affordance became an important principle for the Gibsons' theoretical system, tying together many elements of their theories.

Gibson explains his coining of the term in his final, major work, *The Ecological Approach to Visual Perception*:

The affordances of the environment are what it offers the animal, what it provides or furnishes, either for good or ill. The verb to afford is found in the dictionary, but the noun affordance is not. I have made it up. I mean by it something that refers to both the environment and the animal in a way that no existing term does. It implies the complementarity of the animal and the environment.^[46]

More succinctly: Affordances are properties of environmental structures that provide opportunities for action to complementary organisms.^[47]

When ambient energy—light, sound, and so on—interacts with structures, it produces perceivable information about the intrinsic properties of those structures. The

“complementary” part is key: it refers to the way organisms’ physical abilities are complementary to the particular affordances in the environment. Hands can grasp branches, spider legs can traverse webs, or an animal’s eyes (and the rest of its visual system) can couple with the structure of light reflected from surfaces and objects. Affordances exist on their own in the environment, but they are partly defined by their relationship with a particular species’ physical capabilities. These environmental and bodily structures fit together because the contours of the latter evolved within the shaping mold of the former.

And, of course, action is required for this information pickup. In nature, a tree branch has structural properties that we detect through the way light and other energy interact with the substance and texture of the branch. We are always in motion and perceiving the branch from multiple angles. Even if we’re standing still, our bodies are breathing, our eyes shifting, and able to detect the way shadows and light shift in our view. This is an unconscious, brain-body-environment dynamic that results in our bodies’ detection of affording properties: whether a branch is a good size to fit the hand, and whether it’s top-heavy or just right for use as a cane or club.

Perception evolved in a natural world full of affordances, but our built environment also has these environmental properties. Stairs, such as those shown in [Figure 4-9](#), are made of surfaces and materials that have affordances when arranged into steps. Those affordances, when interacting with energy such as light and vibration, uniquely structure that energy to create information that our bodies detect.

Our bodies require no explanation of how stairs work, because the information our bodies need is intrinsic in the

structure of the stairs. Just as a hand pressed in clay shapes the clay into structure identical to the hand, affordance shapes energy in ways that accurately convey information about the affordance. Organisms rely on this uniquely structured energy to “pick up” information relevant to what the organism can do.



Figure 4-9. Stairs at City Lights Bookstore in San Francisco (photo by author)

For a person who never encountered stairs before, there might be some question as to why climbing up their incline would be desirable, but the perceiver’s body would pick up the fact that it could use them to go upward either way. Humans evolved among surfaces that varied in height, so our bodies have properties that are complementary with the affordances of such surface arrangements.

For most of us, this is a counterintuitive idea, because we’re so used to thinking in terms of brain-first cognition. In

addition, some of the specifics of exactly what affordance is and how Gibson meant it are still being worked out among academics. The principle of affordance was always a sort of work in progress for Gibson, and actually emerged later in his work than some of his other concepts.^[48] He developed it as an answer to Gestalt psychology theories about how we seem to perceive the *meaning* of something as readily as we perceive its physical properties; rather than splitting these two kinds of meaning apart, he wanted to unify the dualism into one thing: affordance.^[49]

We won't be exploring all the various shades of theory involved with affordance scholarship. But, because there is so much talk of affordances in design circles, I think it's valuable to establish some basic assumptions we will be working from, based on Gibson's own work.

Affordance is a revolutionary idea

Gibson claimed in 1979 that affordance is “a radical departure from existing theories of value and meaning.”^[50] And in some ways, it is still just as radical. As an empiricist, Gibson was interested in routing around the cultural orthodoxies and interpretive ideas we layer onto our surroundings; he wanted to start with raw facts about how nature works. If we take Gibson's theory of affordances seriously on its own terms, we have to take seriously the *whole* of his ecological system, not pick and choose parts of it to bolt onto mainstream, brain-first theory.^[51]

Affordances are value-neutral

In design circles, there is sometimes mention of *anti-*

affordances. But in Gibson's framework there's no need for that idea. Affordances offer the animal opportunities for action, in his words, "for good or ill." He writes of affordances having "positive" or "negative" effects, explaining that some affordances are inconvenient or even dangerous for a given agent, but they are still affordances.^[52] An affordance isn't always good from the perspective of a particular organism. Water affords drowning for a terrestrial mammal, but it affords movement and respiration for a fish.^[53] Fire affords comforting heat, light, and the ability to cook food, but it also affords injury by burning, or destruction of property. Affordance is about the structural and chemical properties that involve relationships between elements in the environment, some of which happen to be human beings. Separating what a structure affords from the effect on the perceiver's self-interests helps us to remember that not all situations are the same from one perceiver to the next.

Perception of affordance information comes first; our ideas about it come later

Gibson argued that when we perceive something, we are not constructing the perception in our brains based on preexisting abstract ideas. He states, "You do not have to classify and label things in order to perceive what they afford." When I pick up a fork to eat food, my brain isn't first considering the fork's form and matching it to a category of eating utensils and then telling my arm it's OK to use the fork. The fork affords the stabbing of bites of food and bringing them to my mouth; my body extends its abilities by using the fork as a multipointed extension of my arm. That is, my body appropriates the fork based on its structure, not its category. The facts that it is a *dinner* fork and part of a set of *flatware*

are based on categories that emerge later, from personal experience and social convention.

Affordances exist in the environment whether they are perceived or not

One contemporary theoretical stance argues that we do not perceive real things in the world, but only our brain's ideas and representations of them. Gibson strongly disagreed, insisting that we couldn't perceive anything unless there were a real, physical, and measurable relationship between the things in the world and our bodies. He allowed room for a sort of "life of the mind" that might, in a sense, slosh about atop these real foundations, but it exists only because it was able to emerge from physical coupling between creature and world.

The properties that give something affordance exist whether they are perceived in the moment or not; they are latent possibilities for action. Affordance is required for perception; but an affordance doesn't have to be perceived to exist. I might not be able to see the stairs around the corner in a building, but that doesn't mean the stairs' ability to support climbing doesn't exist. This idea complicates the commonly taught concept that a mental model drives behavior.

Affordance means the information we need for action does not have to be "mental" and is actually in the structures of the environment. I perceive and use the stairs not because I have a mental model of them; no model is needed because all the information necessary is intrinsic to the shape and substance of the stairs. Prior learned experience might influence my usage in some way, but that's in addition to perception, not perception itself.

Affordances are there, whether they are perceived *accurately* or not

A Venus flytrap exists because it can get nutrition. And it gets nutrition because it “tricks” its prey into thinking it is a source of food for the prey, rather than the other way around. For a fly, the affordance of the flytrap is being caught, dissolved, and absorbed by a plant. What the fly perceives, until it is often too late, is “food.” Likewise, we might perceive ground where there is actually quicksand, or a tree branch that is actually a snake. Perception is of the information created by the affordance, not the affordance itself. This is an important distinction that has often been misunderstood in design practice, leading to convoluted discussions of “perceived” versus “actual” affordances. The affordance is a property of the object, not the perception of it.

Affording information is always in a context of other information

No single affordance exists by itself; it’s always nested within a broader context of other affording structures. For example, even if we claim to “add an affordance” by attaching a handle to a hammer head, the hammer is useful only insofar as it can bang on things that need to be banged upon. Stairs afford climbing, but they’re always part of some surrounding environment that affords other actions, such as floors and landings, walls, handrails, and whatever is in the rooms the stairs connect. In digital devices, the physical buttons and switches mean nothing on their own, physically, other than “pushable” or “flippable”—what they actually affect when invoked is perceived only contextually.

Affordances are learned

Human infants are not born understanding how to use stairs. Even if we allow that perception couples with the information of the stairs' surfaces to detect they are solid, flat, and go upward, we still have to learn how and why to use them with any degree of facility. Infants and toddlers not only inspect the stairs themselves, but also were likely carried up them by caregivers and saw others walking or running up and down them long before trying out the stairs for themselves. Learning how to use the environment happens in a densely textured context of social and physical experience. This is true of everything we take for granted in our environment, down to the simplest shapes and surfaces. We learned how to use it all, whether we remember learning it or not. In digital design, there is talk of *natural interfaces* and *intuitive designs*. What those phrases are really getting at is whether an interface or environment has information for action that has already been learned by its users. When designing objects and places for humans, we generally should assume that no affordance is natural. We should instead ask: is this structure's affordance more or less conventional or learnable—keeping in mind that “learnable” is often dependent upon how the affordance builds on established convention.

Directly Perceived versus Indirectly Meaningful

The Gibsons continued to expand their theories into how affordances function underneath complex cultural structures, such as language, cinema, and whole social systems. Other scholars have since continued to apply affordance theory to understanding all sorts of information. Likewise, for

designers of digital interfaces, affordance has become a tool for asking questions about what an interface offers the user for taking action. Is an on-screen item a button or link? Is it movable? Or is it just decoration or background? Is a touch target too small to engage? Can a user discover a feature, or is it hidden? The way affordance is discussed in these questions tends to be inexact and muddled. That's partly because, even among design theorists and practitioners, affordance has a long and muddled history. There are good reasons for the confusion, and they have to do with the differences between how we perceive physical things versus how we interpret the meaning of language or simulated objects.

The scene from *Rosemary's Baby* serves as an apt example of this simulated-object issue. In the scene depicted in [Figure 4-7](#), moviegoers can't see Minnie's mouth moving behind the door frame, even if they shift to the right in their seats, because the information on the screen only simulates what it portrays. There is no real door frame or bedroom that a viewer can perceive more richly via bodily movement. If there were, the door and rooms would have affordances that create directly specifying information for bodies to pick up, informing the body that moving further right will continue revealing more of Minnie's actions. Of course, some audience members tried this, but calibrated by stopping as soon as their perception picked up that nothing was changing. Then, there were undoubtedly nervous titters in the crowd—how silly that we tried to see more! As we will see later, we tend to begin to interact with information this way, based on what our bodies assume it will give us, even if that information is tricking us or simulating something else, as in digital interfaces.

Perception might be momentarily fooled by the movie, but the only affordance actually at work is what is produced by a

projector, film, and the reflective surface of the screen. The projector, film, and screen are quite real and afford the viewing of the projected light, but that's where affordance stops. The way the audience interprets the meaning of the shapes, colors, and shadows simulated by that reflected light is a different sort of experience than being in an actual room and looking through an actual door.

In his ecological framework, Gibson refers to any surface on which we show communicative information as a “display.” This includes paintings, sketches, photographs, scrolls, clay tablets, projected images, and even sculptures. To Gibson, a display is “a surface that has been shaped or processed so as to exhibit information for more than just the surface itself.”^[54] Like a smartphone screen, a surface with writing on it has no intrinsic meaning outside of its surface's physical information; but we aren't interested in the surface so much as what we interpret from the writing. Gibson refers to the knowledge one can gain from these information artifacts as *mediated* or *indirect*—that is, compared to direct physical information pickup, these provide information *via a medium*.

Depending on where you read about affordances, you might see **Affordance** used to explain this sort of mediated, indirectly meaningful information. However, for the sake of clarity, I will be specifying **Affordance** as that which creates information about itself, and I will not be using the term for information that is about something beyond the affordance. Images, words, digital interfaces—these things all provide information, but the ultimately relevant meaning we take away from them is not intrinsic. It is interpreted, based on convention or abstraction. This is a complex point to grasp, but don't worry if it isn't clear just yet. We will be contemplating it together even more in many chapters to

come.

This approach is roughly similar to that found in the more recent work of Don Norman, who is most responsible for introducing the theory of affordance to the design profession. Norman updates his take on affordance in the revised, updated edition of his landmark book, *The Design of Everyday Things* (Basic Books). Generally, Norman cautions that we should distinguish between affordances, such as the form of a door handle that we recognize as fitting our hand and suited for pulling or pushing, and *signifiers*, such as the “Push” or “Pull” signs that often adorn such doors.^[55] We will look at signifiers and how they intersect with affordance in **Part III**.

This distinction is also recommended in recent work by ecological psychologist Sabrina Golonka, whose research focuses on the way language works to create information we find meaningful, “without straining or redefining original notions of affordances or direct-perception.” For affordance to be a useful concept, we need to tighten down what it means and put a solid boundary around it.^[56]

That doesn’t mean we are done with affordance after this part of the book. Affordance is a critical factor in how we understand other sorts of information. Just as a complex brain wouldn’t exist without a body, mediated information wouldn’t exist without direct perception to build upon. No matter how lofty and abstract our thoughts are or how complex our systems might be, all of it is rooted, finally, to the human body’s mutual relationship with the physical environment.

As designers of digitally infused parts of our environment, we have to continually work to keep this bodily foundation in mind. That’s because the dynamic by which we understand

the context of a scene in a movie—or a link on a web page—borrows from the dynamic that makes it possible for us to use the stairs in a building or pick a blackberry in a briar patch. Our perception is, in a sense, hungry for affordance and tries to find it wherever it can, even from indirectly meaningful information. That is, what matters to the first-person perspective of a user is the blended spectrum of information the user perceives, whether it is direct or indirect. It's in the teasing apart of these sorts of information where the challenge of context for design truly lies.

Soft Assembly

Affordance gives us one kind of information: what I'm calling physical information. But, what we experience and use for perception is the information, not the affordance that created it. Cognition grabs information and acts on it, without being especially picky about technical distinctions of where the information originates.

Cognition recruits all sorts of mechanisms in the name of figuring out the world, from many disparate bodily and sensory functions. The way this works is called *soft assembly*. It's a process wherein many various factors of body-environment interaction aggregate on the fly, adding up to behaviors effective in the moment for the body.^[57] Out of all that activity of mutual interplay between environment and perceiver, there emerges the singular behavior. Now that we can even embed sensors and reactive mechanisms into our own skin, this way of thinking about how those small parts assemble into a whole may be more relevant than ever.^[58]

We're used to thinking of ourselves as separate from our

environment, yet an ecological or embodied view offers that the boundary between the self and the environment is not absolute; it's porous and in flux. For example, when we pick up a fallen tree branch and use it as a tool—perhaps to knock fruit from the higher reaches of the tree—the tool becomes an extension of our bodies, perceived and wielded as we would wield a longer arm. Even when we drive a car, with practice, the car blends into our sense of how our bodies fit into the environment.^[59]

This isn't so radical a notion if we don't think of the outer layer of the human body as an absolute boundary but as more of an inflection point. Thus, it's not a big leap to go from "counting with my fingers" to "counting with sticks." As author Louise Barrett explains, "When we take a step back and consider how a cognitive process operates as a whole, we often find that the barrier between what's inside the skin and what's outside is often purely arbitrary, and, once we realize this, it dissolves."^[60] The way we understand our context is deeply influenced by the environment around us, partly because cognition includes the environment itself.

In *Supersizing the Mind: Embodiment, Action, and Cognitive Extension* (Oxford University Press), Andy Clark argues that our bodies and brains move with great fluidity between various sorts of cognition. From moment to moment, our cognition uses various combinations of cognitive loops—subactive cognition, active-body cognition, and extended cognition, using the scaffolding of the environment around us. Clark explains that our minds "are promiscuously body-and-world exploiting. They are forever testing and exploring the possibilities for incorporating new resources and structures deep into their embodied acting and problem-solving regimes."^[61]

Clark also explains how an assembly principle is behind how cognition works as efficiently as it can, using “whatever mix of problem-solving resources will yield an acceptable result with minimal effort.”^[62] We might say that we use a combination of “loops of least effort.”^[63] That’s why audience members leaned to the right in a theater, even though there was no logical reason to do so. We act to perceive, based on the least effortful interpretation of the information provided, even though it sometimes leads us astray. That is, we can easily misinterpret our context, and act before our error is clear to us. Even though we might logically categorize the variety of resources the perceiver recruits for cognition, to the perceiver it is all a big mash-up of *information about the environment*. For designers, that means the burden is on the work of design to carefully parse how each element of an environment might influence user action, because the user will probably just act, without perceiving a difference.

“Satisficing”

Perhaps this loops-of-least-effort idea helps explain a behavior pattern first described by scientist-economist Herbert Simon, who called it *satisficing*. Satisficing is a concept that explains how we conserve energy by doing whatever is just enough to meet a threshold of acceptability. It’s a portmanteau combining “satisfy” and “suffice.” Its use has been expanded to explain other phenomena, from how people decide what to buy to the way a species changes in response to evolutionary pressures.^[64] Satisficing is a valuable idea for design practice, because it reminds us that users *use* what we design. They don’t typically ponder it, analyze it, or come to know all its marvelous secrets. They *act* in the world

based on the most obvious information available and with as little concentration as possible. That's because cognition starts with, and depends upon, continual action and interaction with the environment. Users aren't motivated by first understanding the environment. They're too busy just getting things done, and in fact they tend to improvise as they go, often using the environment in different ways than intended by designers.^[65]

Even the most careful users eventually “poke” the environment to see how it responds or where it will take them by clicking or tapping things, hovering with a mouse, waving a controller or phone around in the air, or entering words into a search field. Just the act of looking is a physical action that probes the environment for structural affordance information, picking up the minimum that seems to be needed to move and then appropriating the environment to their needs. We see this when we observe people using software: they'll often try things out just to see what happens, or they find workarounds that we never imagined they would use. It does no good to call them “bad users.” These are people who behave the way people behave. This is one reason why lab-based testing can be a problem; test subjects can be primed to assume too much about the tested artifact, and they can overthink their interactions because they know they're being observed. Out in the world they are generally less conscious of their behaviors and improvised actions.

The embodied view flips the traditional role of the designer. We're used to thinking of design as creating an intricately engineered setting for the user, for which every act has been accounted. But the contextual meaning of the environment is never permanently established, because context is a function of the *active engagement* of the user. This means the primary

aim of the designer is not to design ways for the artifact to be used but instead to *design the artifact to be clearly understood*,^[66] so the user can recruit it into her full environmental experience in whatever way she needs.

Umwelts

We have a Boston terrier named Sigmund (**Figure 4-10**). He's a brownish-red color, unlike typical Bostons. When walking Sigmund, I notice that no matter how well he's staying by my side, on occasion he can't help going off-task. Sometimes, he stops in his tracks, as if the ground has reached up and grabbed him. And, in a sense, that's what is happening. Sigmund is perceiving something in his environment that is making him stop. It's not premeditated or calculated; it's a response to the environment not unlike walking into a glass wall. For him, it's something he perceives as "in the way" or even dangerous, like an angle of shadow along the ground that could be a hole or something closing in on him. It happens less as he gets older and learns more about the world around him. Stairs used to completely freak him out, but now he's a pro.



Figure 4-10. Sigmund

But often what affects Sigmund's behavior is invisible to me. Like many dogs, he's much more perceptive of sound and smell than hominids. In those dimensions, Sigmund's world is much richer than mine. If I bring him outside and we encounter even a mild breeze, he'll stop with his nose pointed upward and just smell the air the way I might watch a movie at IMAX. Certainly we make friends differently. At the dog park, I'm mostly paying attention to the visual aspects of faces around me, whereas Sigmund gets to know his kind from sniffing the other end.^[67]

For Sigmund and me, much of our worlds overlap; we're both warm-blooded terrestrial mammals, after all. We're just responding to different sorts of information in addition to what we share. Sigmund might stop because of a scent I cannot smell, but I might stop in my tracks because I see a caution light or stop sign. Sigmund and I are walking in somewhat different worlds—each of us in our own *umwelt*.

Umwelt is an idea introduced by biologist Jakob Johann von Uexküll (1874–1944), who defined it as the world as perceived and acted upon by a given organism. Uexküll studied the sense organs and behaviors of various creatures such as insects, amoebae, and jellyfish, and developed theories on how these they experience their environments.^[68]

As a result of this work, Uexküll argued that biological existence couldn't be understood only as molecular pieces and parts, but as organisms sensing the world as part of a system of signs. In other words, Uexküll pioneered the connection of biology with *semiotics*, creating a field now called *biosemiotics* (we will look more closely at signs, signification, and semiotics in **Part III**). His work also strongly influenced seminal ideas in phenomenology from the likes of Martin Heidegger and Maurice Merleau-Ponty.

If we stretch Uexküll's concept just a little, we can think of different people as being in their own *umwelts*, even though they are the same species. Our needs and experiences shape how we interpret the information about the structures around us.

For skateboarders, an empty swimming pool has special meaning for activities that don't register for nonriders; for jugglers, objects of a certain size and weight, such as oranges, can mean "great for juggling," whereas for the rest of us, they just look delicious.

Simulated Objects

paintings representing, **Semantic Function of Simulated Objects**

range of information, **Semantic Function of Simulated Objects**

skeuomorphic vs. flat design, **Semantic Function of Simulated Objects**

using in helping physical bodies comprehend environment, **About Definitions**

semantic information, **Modes of Information**

(see also Semantic mode in information model)

about, **Language as Environment**

behaving as organizational infrastructure, **The Tales Organizations Tell**

computers processing inputs of, **Digital Learning and Agency**

language as

about, **Language as Environment, How Language Works**

adding to confusion, **Physical and Semantic Confusion**

architecture and, **Semantic Architecture**

as infrastructure, **Semantic Function**

as Semantic mode in information model, **Modes of Information**

contextual, **Signification Conflation**

creating meaningful information, **Directly Perceived versus Indirectly Meaningful**

disconnecting from physical environment, **Physical and Semantic Confusion**

environment and (see language: as environment)

grammar (see grammar)

labeling function of, **Tools for Understanding, What Writing Does**

levels of linguistic structure, **How Language Works**

oral, **The Origins of Writing, What Writing Does**