

This article was downloaded by: [72.192.152.173]

On: 21 October 2013, At: 15:32

Publisher: Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Philosophical Psychology

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/cphp20>

The cultural ecosystem of human cognition

Edwin Hutchins

Published online: 19 Aug 2013.

To cite this article: Edwin Hutchins , Philosophical Psychology (2013): The cultural ecosystem of human cognition, Philosophical Psychology, DOI: 10.1080/09515089.2013.830548

To link to this article: <http://dx.doi.org/10.1080/09515089.2013.830548>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden. Terms & Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

The cultural ecosystem of human cognition

Edwin Hutchins

Everybody knows that humans are cultural animals. Although this fact is universally acknowledged, many opportunities to exploit it are overlooked. In this article, I propose shifting our attention from local examples of extended mind to the cultural-cognitive ecosystems within which human cognition is embedded. I conclude by offering a set of conjectures about the features of cultural-cognitive ecosystems.

Keywords: Cultural Practices; Cultural-Cognitive Ecosystems; Distributed Cognition; Extended Cognition

1. Introduction

Man's nervous system does not merely enable him to acquire culture; it positively demands that he do so if it is going to function at all. Rather than culture acting only to supplement, develop, and extend organically based capacities logically and genetically prior to it, it would seem to be ingredient to those capacities themselves. A cultureless human being would probably turn out to be not an intrinsically talented, though unfulfilled ape, but a wholly mindless and consequently unworkable monstrosity. (Geertz, 1973, p. 68)

In his 2001 book, *Mindware*, Andy Clark, noting the way that the cultural world can orchestrate thinking, called for "a new kind of cognitive scientific collaboration involving neuroscience, physiology, and social, cultural, and technological studies in about equal measure" (2001, p. 154). In the dozen years since that call, considerable progress has been made on this collaboration, but much work remains to be done. My purpose in this article is to facilitate the collaboration that Clark called for and to nourish it by suggesting a modification of perspective.

Edwin Hutchins is Professor of Cognitive Science at the University of California, San Diego.
Correspondence to: Edwin Hutchins, Department of Cognitive Science, UC San Diego, La Jolla, CA 92093-0515, USA. Email: ehutchins@ucsd.edu

1.1. Plan of the Paper

I will begin by reviewing some relations between distributed cognition and extended mind. I will suggest that distributed cognition and extended mind imply related, but different (hopefully complementary), perspectives on cognition. The extended mind framework is a way to approach distributed cognitive systems in a specific range of spatial and temporal scales. In earlier writings, I have referred to these as “functional systems,” and Clark has called them “ecological assemblies.” Distributed cognitive systems also exist at other spatial and temporal scales. In this article, I propose to shift the focus from ecological assemblies surrounding an individual person to cultural ecosystems operating at larger spatial and temporal scales.

Because our culture is largely invisible to us, making the cognitive ecosystem visible is a challenge. I present a variety of cultural practices that all share a common underlying formal process: the projection of a trajectory onto a spatial array. I intend the discussion of the examples as a sort of intuition pump to help readers shift their attention from particular ecological assemblies to the properties of the cognitive ecosystem in which the assemblies arise. On the basis of these observations, I offer a set of conjectures about the cognitive ecosystem and the implications of the features of the ecosystem for the study of cognitive science.

2. Extended Mind

Clark and Chalmers launched the first wave of extended mind with their seminal 1998 paper, “The extended mind.” Describing some examples of extended mind, Clark and Chalmers said, “in these cases, the human organism is linked with an external entity in a two-way interaction, creating a coupled system that can be seen as a cognitive system in its own right” (1998, p. 8).

The first wave of extended mind theorizing made narrow commitments to the nature of the relationships between internal and external resources, arguing for the parity of internal and external resources. A second wave of extended mind studies now entertains a wide variety of possible relations between internal and external resources. For example, “human cognitive processing, EXTENDED claims, may at times loop into the environment surrounding the organism” (Clark, 2008, p. 111). There is now talk of “hybrid systems,” in which cognitive processes “spread beyond boundaries of skin and skull” (Michaelian & Sutton, 2013, p. 3). Typically, the causal flow in cases of extended cognition is conceived to run from the inside out as the brain “recruits” resources in the environment. However, the causal flow can be conceived to run the outside in as well. Some efforts in extended mind follow Vygotsky (1978) in noting how internal processes are shaped by their coordination with external resources. An immersion metaphor is sometimes used, as in, “capacities for thinking and remembering, on such views, soak in from the socially-available models to which our biological mechanisms are especially attuned” (Michaelian & Sutton, 2013, p. 5).

Extended mind is thus sensitive to the role of the cultural environment in orchestrating cognitive processes. As Michaelian and Sutton say,

Clark saw distributed cognition as a natural extension of the point made in much connectionist literature that order and systematicity in human cognition and action can derive in part from the stability of our environments, rather than as a direct product or reflection of exhaustively-specified internal recipes. Cognition might thus be multiply distributed, both within neural networks and across bodies, artifacts, and social groups. (2013, p. 6)

There is a considerable overlap between the practice of distributed cognition and extended mind. Michaelian and Sutton (2013) point out that both approaches examine interactions with the social and material world, both emphasize the range of processes that implement those interactions, and both aspire to shift explanatory practices in cognitive science. Before turning to a discussion of distributed cognition, however, I wish to note two features that distinguish the extended mind approach from distributed cognition. First, extended mind picks out a *kind* of cognition. In the extended mind view, mind may sometimes extend beyond the brain, and sometimes it does not. Extended cognition refers only to that subset of cognitive events that involve interaction of internal and external resources. The individuation of cases of extended mind depends on empirical claims that are grounded in a container metaphor for mind. The extension of mind is manifest in links and relations that cross the usual boundary of the mind container. Second, extended mind assumes a *center* in the cognitive system: the organism (or the organism's brain), which is the normal mind container with respect to which cognition can be said to extend.

3. Distributed Cognition

Distributed cognition is not a kind of cognition; it is a perspective on all of cognition (see especially Hutchins, 2001, 2006). Distributed cognition begins with the assumption that all instances of cognition *can be seen* as emerging from distributed processes. For any process there is always a way to see it as distributed. In practice this implies that wherever we find cognition, it will be possible to investigate how a process we call cognitive *emerges* from the *interactions* among elements in some system. Of course, the notions of centralized and distributed are always relative to some scale of investigation. Thus, to take the distributed perspective is not to make any claim about the nature of the world. Rather, it is to choose a way of looking at the world, one that selects scales of investigation such that wholes are seen as emergent from interactions among their parts. The boundaries of the unit of analysis for distributed cognition are not fixed in advance; they depend on the scale of the system under investigation, which can vary as described below. Within a particular scale, the boundary of the cognitive unit of analysis may shift dynamically during the course of activity.

Within distributed cognition, the interesting question then is not “is cognition distributed or is it not?” or even “is cognition sometimes distributed and sometimes not distributed?” Rather, the interesting questions concern the elements of the cognitive system, the relations among the elements, and how cognitive processes arise from interactions among those elements. It is possible to pose many empirical hypotheses within this perspective. The hypothesis of extended cognition is an important hypothesis within the perspective of distributed cognition. Such hypotheses

are subject to confirmation or rejection via empirical investigations. However, there is no series of experiments or set of observations that could prove or disprove the distributed cognition perspective, because the perspective itself makes no empirical claims.

The distributed cognition perspective recognizes the emergence of cognitive phenomena in distributed systems at many spatial scales (Hutchins, 1995a, 2001, 2006). For example, a brain is known to be a huge distributed cognitive system. Cognitive processes are believed to emerge from complex interactions among very large numbers of neurons. On a slightly larger spatial scale, recent work on the multimodality of perception highlights the ways that cognitive functions are distributed across areas of the brain and organs of the body. While I do not believe in the extreme modularity of mind proposed by some theorists (Fodor, 1983; Minsky, 1988; Mithen, 1999), the modular mind, as a conceptual construct, is a clear example of taking the distributed cognition perspective.

In this review of spatial scales for distributed cognition systems, the extended mind approach occupies a mid-level scale. It picks out a particular class of distributed cognitive systems that operate on a spatial scale somewhat larger than an individual person. Such systems also have a temporal scale, typically completing operational cycles on the order of seconds or minutes. Cognition in such extended mind systems emerges from the interactions among resources inside an individual agent and resources that are external to the agent. When the focus is on systems that involve the interaction of persons with their immediate material and social environment, the intersection of distributed cognition with extended mind is substantial. Unlike extended mind, however, distributed cognition does not *assume* a center for any cognitive system. Nor does it grant a priori importance to the boundaries of skin or skull. For distributed cognition, the existence of boundaries and centers are empirical questions. Centers and boundaries are features that are determined by the relative density of information flow across a system. Some systems have a clear center while other systems have multiple centers or no center at all. Given the preferred unit of analysis for extended mind, an organism's brain may rightfully be considered the center of that cognitive system and the nature of the transformations that occur at the skin of the organism make the skin a consequential boundary.

The developing field of collective intelligence examines the distribution of cognition on an even larger spatial scale. For collective intelligence, the individual units of the cognitive system are fully constituted autonomous agents. Important topics in the study of collective intelligence include insect societies, economic markets, social media, crowd computation, group memory, public policy design, and the emergence and evolution of intelligence.

Distributed cognition searches for principles that apply at multiple scales and across vastly different kinds of cognitive systems. For example, the emergence of a shared language is clearly a cognitive accomplishment, but it is not one that is accomplished by any individual. The emergence of a language is a cognitive process that takes place in an evolving cognitive ecosystem that includes a shared world of objects and events as well as adaptive resources internal to each member of the community. In the 1990s,

I explored the emergence of language-like behavior using computational models of communities of agents (Hazlehurst & Hutchins, 1998; Hutchins & Hazlehurst, 1991; Hutchins & Hazlehurst, 1995; Hutchins & Hazlehurst, 2002). A few years earlier, Hinton and Becker (unpublished manuscript) tackled the problem of how separate visual modules in the brain could learn to communicate the information needed to recover depth from stereovision without having a teacher. It turns out that the problem of how a community can learn a lexicon without a teacher to specify the elements of a lexicon is very similar to the problem of how various brain areas might learn to communicate without a teacher to specify the form of information to exchange. While the implementations in the two studies were quite different, the formal description of the solution in the two cases was identical. The solutions are methods that incrementally maximize a measure called mutual information.

A second example of the search for principles that operate at different scales comes in the form of the hypothesis that cognition has a fractal nature. That is, regardless of the level of integration of phenomena considered, intelligent systems will be characterized by local regions of high interconnectivity separated by regions of lower interconnectivity. This is closely related to the information theoretic definition of complexity. This hypothesis could be restated as follows: cognitive systems will be found to be complex at all scales.

More recently, I have been using information theoretic measures to explore the hypothesis that cultural practices tend to reduce entropy (increase predictability) at all scales in a cultural cognitive ecosystem (Hutchins, 2012). This is important because a brain that is a prediction machine, as suggested by Clark in his latest work (2013), will require predictable experience. Referring to prediction driven learning in cultural context, Clark says,

At multiple time-scales, and using a wide variety of means (including words, equations, graphs, other agents, pictures, and all the tools of modern consumer electronics) we thus stack the dice so that we can more easily minimize costly prediction errors in an endlessly empowering cascade of contexts from shopping and socializing, to astronomy, philosophy, and logic. (Clark, 2013, p. 195)

I expect that many other principles operate across a wide range of scales in cognitive systems. This is a field that is ripe for exploration.

4. A Family of Trajector-Based Cultural Practices

In this section, I describe a family of practices that are all based on the projection of a trajector onto a real or imagined spatial array. I will refer to this as the family of “imagined-trajector-based” practices. I discussed some members of this family in an earlier paper (Hutchins, 2005). In that paper, my focus was on the consequences for an individual cognizer of anchoring concepts to material patterns or structures. In this article, I hope to shift the focus to the properties of the cognitive ecosystem in which these practices exist (Hutchins, 2010a, 2010b).

4.1. The Cultural Practice of Queuing For Service

In some cultural contexts, people seeking service arrange themselves in a queue as a way to control the sequence of access to services. People in a queue arrange their bodies in a linear array.¹ The practice of queuing for service consists of three interlocking component practices. First, there is a cooperative social practice of forming linear arrangements of bodies. Second, there are spatial material (and perhaps architectural) practices that designate some location as the source of service. Third, there is a socially shared individual mental practice of seeing the linear arrangement of bodies with respect to the service location as a queue. These practices are mutually supportive and depend on one another for their meaning and their very *raison d'être*.

Let us focus for a moment on just the most obviously cognitive element of the queuing practice. Seeing a line as a queue is an example of the mapping of a conceptual structure, what in cognitive grammar is called a *trajector* (Langacker, 1987, p. 231), onto a physical array. This mapping of imagined structure onto perceived structure produces a conceptual blend (Fauconnier & Turner, 2008; Hutchins, 2005) which gives rise to a particular emergent property: a sequential ordering of the bodies of the individuals in the queue. The sequence of access to service is not present in either the physical line of people or in the *trajector*. It emerges only when some particular viewer blends the conceptual *trajector* with the perception of an appropriately ordered and situated physical array. Seeing the line as a queue is a *cognitive* practice because it makes possible a set of inferences. Who is next in line? Who arrived before whom? How far am I in space from (and how long must I wait before) getting service?

The composite queuing practice is also cognitive, but for different reasons than the individual practice, and it implements different cognitive functions than the individual practice. The practice of queuing for service is, above all else, a cooperative public means to record and remember the order of arrival of clients. The queue also manages a forgetting function when people leave the line either before or after receiving service.

This everyday practice often takes place in complex social and institutional settings. Examining such settings can reveal the extent of the network of elements that are related in the cognitive ecosystem. In airports, for example, elaborate material arrangements serve to induce the formation of queues and shape the queues as they form. There may be patterns of lines painted on the floor, guide ropes or tapes, and signs such as “enter here” or “wait here for first available agent.” The practice of forming a queue for service exists in a cultural ecosystem that includes services to be rendered (a set of facts about economic systems); the roles of service provider, who renders services, and client, who accesses services (facts about social organization); and locations in space at which service is rendered (facts about architecture).

Seeing an array of people as a queue integrates these elements in a particular way. It is an example of enacting a meaning by seeing the world in a particular way. A physical pattern that is open to many interpretations is “seen as” a particular, culturally meaningful, phenomenal object. The phenomenon of enacting meanings by “seeing” the world in particular ways (Stewart, Gapenne, & di Paulo, 2010) is absolutely

ubiquitous in human experience and is accomplished via cultural practices. When a line is being *seen as* a queue, other elements of the setting will be *seen as* instances of other roles in the queuing for service practice. This sort of fit suggests that cultural practices are composed of coherent constellations of mutually supportive component practices. In such a system, increasing the likelihood of any component increases the likelihood of the other components.

Forming a line from a group is a physical form of dimensionality reduction (Hutchins, 2012). A group of people occupying two dimensions of a surface approximates a one-dimensional array when they form a queue. This dimensionality reduction does not take place in any person's mind. It takes place in the space shared by the participants to the practice. Once the dimensionality reduction in physical space has emerged, however, it supports or affords the cognitive practices of making inferences on the line "seen as" a queue. The experience of a one-dimensional line is more predictable than the experience of a two-dimensional crowd. The experience of a queue has lower entropy than the experience of a crowd. This increase in predictability and structure is a property of the distributed system, not of any individual mind.

4.2. *More Members of the Extended Family of Trajectory-Based Practices*

The superposition of a trajectory on a spatial array of physical objects is the common constituent in a large family of cultural practices. In each instance, the projection of a trajectory combines with different material and/or social practices to create composite practices. Each of these composite practices consists of a system of interlocking component practices. Each practice increases the predictability of experience by reducing the dimensionality of experience. Each one is located in a network of other practices and has mutually supportive relations to the other practices in the network. In the paragraphs below, I will mention practices and explore the local network of supporting elements for each.

I begin with three practices that are not good examples of extended cognition, but are important members of this family of practices.

Superposition of a trajectory onto visible or imagined objects is often applied to the natural world. Take, for example, the practice of seeing an array of stars as a constellation. One can see points of light with one's visual system. It takes a cultural practice to see a constellation (Hutchins, 2008). Decades ago, when I was trained in celestial navigation, I learned a number of strategies for attending to patterns in the night sky in order to identify specific useful stars. Many people know how to follow the so-called "pointer" stars in Ursa Major (also known as the big dipper) to find a useful star called Polaris. At the other end of the dipper, the stars in the handle suggest an arc. One can "follow the arc to Arcturus and drive a spike into Spica." Calling the two stars on the lip of the big dipper the "pointer" stars and the mnemonic phrase oriented to the arc of stars in the handle of the dipper are ways of speaking that help to organize the cultural practices of seeing. These discursive practices (Goodwin, 1994) exploit and activate the practice of imagining particular trajectories on particular visible arrays of points of light. They are part of the local network in the cognitive ecosystem

surrounding the practice of seeing constellations and using constellations to find particular stars.

In a similar way, the lexicon of sequence relations in general fits the practice of imagining a trajectory on a spatial array. In the case of the queue, the ways of speaking about “first in line,” “next,” “back of the line,” and so on are discursive practices that enter into relations of mutual reinforcement with the conception of the linear spatial array as a queue. Of course, these words and phrases are constrained by their use in other contexts, and these relations increase the tightness of the weave of the fabric of the cognitive ecosystem.

The memory practice known as the *method of loci* proceeds by associating an idea with each of an array of spatial landmarks, and then imposing a trajectory on the array of landmarks. This produces a sequential ordering on the set of ideas. It is also possible to imagine an array of objects and then project a trajectory onto the imagined array. The phenomenon, known as fictive motion in language (Fauconnier, 1997; Langacker, 1987; Talmy, 1996), is produced by the projection of a trajectory onto a real or imagined scene, as when one says, “the road runs down to the beach.” The road is static, but this way of speaking adds a dynamic component to the experience of the static object. Fictive motion is a linguistic phenomenon, but employing it is not strictly a matter of knowledge of language. It is also a matter of knowing how to project a trajectory.

We turn now to some practices that are good examples of extended cognition. I present them here not for their value as examples of extended mind, but to illustrate the range of practices that are members of the trajectory-based family.

The practices of numeracy and literacy also exploit the superposition of a trajectory on real or imagined objects. Consider writing and reading. In the cognitive ecosystem, each of these practices provides the other’s reason for being. How many such pairs are there in our modern cognitive ecosystem? The cultural practices of writing and reading assemble a complex configuration of resources including writing implements, a physical text, physically inscribed in some medium, located in space, ways of moving the body (hands and eyes) with respect to the text, visual perception of words, extensive knowledge of language, etc. Writing and reading follow a conventional imposed trajectory (left to right and top to bottom for English). Imposing a trajectory on an array of written items creates a sequential list.

Learning to read requires the domestication of visual attention (Goody, 1977). Similar processes of domestication of visual attention are at work in other sorts of reading including reading natural phenomena such as the night sky, reading static cultural notations such as those found in mathematics or music, and reading dynamic cultural representations as in flying an airplane on instruments. In all of these activities, the domestication of visual attention produces culturally conventional trajectories of attention across spatial arrays of objects.

Dimensionality reduction is a key component of systems that must coordinate with visual attention. While an infinite number of scan patterns are possible, linear patterns are predictable and computationally inexpensive to describe. This is probably why linear trajectories are so common in the ecosystem. The practices of reading and writing in coordination with the practice of printing lined pages produces yet another form of

dimensionality reduction. Most writing systems include conventions for spatial layout under which the total surface area of a page can be seen as a single sequence of locations.

The location of entries in a bound book may become part of a notation system. The apprehension of enduring relations among locations on different pages is possible only because the pages of the book have been bound together. Combining a numbered page sequence with results of the practices of reading and writing reduces the total two-dimensional surface area of all the pages in a bound book into a single line of locations. If written entries are made in the book following the conventions of the writing system, then this single thread of locations throughout the book can be interpreted as a temporal sequence, which supports inferences about when the annotations were entered into the book. Because both books and cityscapes inherit numbering schemes from the same source, a location in a book with numbered pages is very much like an address of a building along a single street. Some of the practices for navigating a city are shared with the practices for navigating a book. For example, I'm looking at page 48 (or I'm on 48th Street) and I want to see something that I know is on page 53 (or 53rd Street). Which way do I turn the pages (or walk)? This suggests that relationships among practices in the cognitive ecosystem can create possibilities for generalization of skill across activity systems.

The number line is a key component of many systems of mathematics. Projecting a trajectory onto a linear array of number tokens creates a number sequence. Projecting a trajectory on a linear array of numbers in order of magnitude and arranged with a constant interval produces a number line. This practice enables a variety of emergent inferences about relations among numbers. All sorts of scales for the measurement and expression of quantities are members of this family of cultural practices. For example, when the numbers are read as times, a number line becomes a timeline. Reading an analog clock involves seeing a circle as a timeline. In this case, the trajectory is a curve rather than a straight line.

5. Instrumented Digital Cognitive Ethnography

In this section, I describe two additional members of the family of practices. These practices are tuned specifically to the activities on the flight deck of a commercial airliner. As preparation for this discussion, let's take a look at the range of objects that pilots achieve coordination with by projecting trajectories onto spatial arrays. Figure 1 is an image of the first officer's instrument panel captured in an actual airliner cruising at 33,000 feet. In this image, there are 18 distinct scales. The scales come in the form of dials, reticules, and strips. Each scale is read by projecting a trajectory onto the scale, yet each does a different kind of work and is a component of a different functional system. The fact that the scales are related provides an economy of processing. Reading an airspeed indicator (Hutchins, 1995b) shares some practice components with reading an altimeter, which is similar in some ways to reading a compass, a vertical speed indicator, an oil pressure indicator, or a clock. This image captures a slice of an ecosystem in which the projection of trajectories is a dominant species. In the global commercial aviation



Figure 1. Flight instrument panel of a McDonnell Douglas MD-80 airliner.

system, each operating flight deck is a cognitive center, a zone of high density of cognitively consequential cultural practices. The activity system of the airline flight deck is a tightly woven patch in the fabric of the cognitive ecosystem of the modern world.

Over the past few decades, new measurement technologies have made possible fine-grained observation and recording of behavior, first in the laboratory, and now increasingly in real-world settings. The members of my lab have been developing a suite of digital tools to support cognitive ethnography (Fouse, Weibel, Hutchins, & Hollan, 2011; Weibel et al., 2012a; Weibel, Fouse, Emmenegger, Kimmich, & Hutchins, 2012b; Weibel, Fouse, Hutchins, & Hollan, 2011). My team recently collected data on a qualified crew flying a scenario in a high-fidelity simulator of Boeing's new 787 Dreamliner (Hutchins, Weibel, Emmenegger, Fouse, & Holder, forthcoming). We instrumented the simulator and the pilots to collect a rich set of time series data. Our data streams included multiple high-definition video recordings, digital audio recordings, digital pen data (recording notes made by the pilots as well as notes made by observers), wearable eye tracking on both pilots, and digital data from the simulator itself.

Because space is limited, I will touch on only two trajectory-based practices in this setting: reading a navigation chart and reading a waypoint list. These practices appear in just a few seconds of flight deck activity.

While still more than 30 minutes from their destination, the crew programmed the flight management computer for the arrival procedure. This requires the crew to retrieve information from a navigation chart and enter the information into the airplane's flight management computer system. The first officer held the arrival chart in his hands and said, "okay, so after Helens, then we're gonna go to Battle Ground." This is a spoken representation of one leg of the arrival from a waypoint called HELNS to a waypoint called Battle Ground. While producing this utterance, the first officer's eyes first jumped around the chart making several brief fixations before fixating on the depiction of the HELNS waypoint. He then fixated just below the depiction of the HELNS waypoint (along the path of flight) for about half a second, then further down the route of flight for another half second. This was followed by a saccade to the information box for the Battle Ground waypoint where he fixated for almost a full second.

This short vignette allows us to see how the first officer goes about seeing the static marks on the chart as a representation of a dynamic flight route. The pilot's eye gaze enacts the dynamic aspects of the flight route on the chart by projecting a trajectory onto the spatial arrangement of the chart. This simultaneously gives the chart meaning as a representation of the route of flight, places the waypoints in sequence, and enacts the anticipated route of flight in motor activity.

Collecting eye tracking data on both pilots simultaneously allows us to measure the allocation of visual attention by the flight crew system. While the first officer was reading the chart as described above, the captain looked at the blank space on the computer interface where the name of the next waypoint after HELNS should be entered. The first officer's spoken utterance thus coordinated the allocation of visual attention by the two pilots, each using a different cultural practice of reading, to two different representations of the flight route, one on the chart, and the other in the waypoint list on the computer interface. This coordinating effect is only possible because the language of sequential relations bears a synergistic relation to both kinds of reading practices.

While the first officer's visual attention enacted a trajectory on the source of navigation information (the chart), the captain's visual attention was enacting a trajectory on the destination for that information (the computer interface). The captain's eye gaze anticipated the first officer's next action in the activity, which was entering the identifier for the waypoint designated BTG into the list of waypoints that define the route. This complementary allocation of visual attention is evidence of the pilots' joint participation in, and joint construction of, a shared problem-solving activity.

As the first officer then repositioned his body and hands to use the keypad to enter the waypoint identifier, he shifted his eye gaze to the blank space on the waypoint list where the identifier for the next waypoint should be entered. At this point, the two pilots had produced not only congruent eye-gaze point of regard, but congruent use of the practice of projecting a trajectory on the display to read it as a sequential list.

At the moment the first officer read the chart while the captain read the waypoint list, the practices of the pilots were coordinated and complementary. When both looked at the same blank field on the waypoint list, their practices were congruent. This complementarity and congruence of practices determine important performance characteristics of the system,

including such things as the probability of the formation and retention of memories or of noticing an alert or detecting an error.

6. Conjectures Concerning the Properties of the Cultural-Cognitive Ecosystem

The practice of superimposing a trajector on a real or imagined array of objects is very productive and there are many other practices in this family. A wide variety of other kinds of practices enter into symbiotic relations with members of the family: arrangements of bodies; architectural features; ways of speaking and verbal mnemonics; moral principles; arrangements of marks on surfaces; domesticated patterns of visual attention; the physical form of commercially produced writing surfaces; and more. Similarly, the members of the family are embedded in activity systems and participate in the accomplishment of a range of emergent distributed cognitive processes including memory, representation of sequence, sequential computation, moral judgment, planning, time reckoning, arithmetic reasoning, navigation, search, and so on.

The variety of the examples listed above reminds us that the cultural-cognitive ecosystem is heterogeneous and complex. As an object of study, this cognitive ecosystem falls into the cracks among the academic disciplines as they are currently organized. Because no field or discipline has yet taken ownership of cultural-cognitive ecosystems, little is known about their function. Truly understanding how such systems work will require a large and sustained scientific endeavor. However, it is already possible to generate a number of conjectures about the functioning of these systems.

6.1. Conjectures Concerning Cognitive Systems in General

- Connectivity within cognitive systems has a fractal structure. Intelligent systems at all scales have nodes of dense interconnectivity separated by sparser connections. I do not expect intelligent behavior to arise from networks having uniform connectivity.
- Some formal principles will be found to operate at multiple scales in cognitive systems. Perhaps information-theoretic measures will be useful in discovering and expressing these principles. For example, I suspect that cultural practices reduce entropy at multiple scales of a cognitive ecosystem. In a recent paper attempting to sketch a grand vision for the future of cognitive science, Clark (2013) proposed hierarchical prediction machines (HPMs) as a model of human cognitive processing. He proposed that the HPMs act to reduce the internal entropy of experience by learning a generative model of the world. I have tried to show here how other processes act in the cognitive ecosystem to reduce the entropy of experience by putting cultural structure in the world to be experienced by a person.

6.2. Conjectures Concerning Global Features of the Cultural Cognitive Ecosystem

Like any ecosystem, the cultural-cognitive ecosystem can be seen as a constraint satisfaction system that settles into a subset of possible configurations of elements. It is

a dynamical system in which certain configurations of elements (what we know as stable practices) emerge (self-assemble) preferentially. In this perspective, constraints exist in many places and interact with one another through a variety of mechanisms of constraint satisfaction. Some of these are neural mechanisms; others are implemented in material tools; and still others are emergent in social processes of collective intelligence, the development of conventions, for example.

- Cultural practices are emergent structural configurations in a rich network of relationships.
- The development of new practices is constrained by the existing networks in the ecosystem.
- Culture is learnable because the ecosystem of practices has structure.

6.3. *Conjectures Concerning Local Features of the Cultural Cognitive Ecosystem*

We have observed that cultural practices have internal coherence. They consist of mutually supporting component practices and artifacts. If it is assumed that the assembly of functional systems proceeds via constraint satisfaction, then these conjectures follow.

- The alignment of components and the internal coherence of practices is dynamic and adaptive.
- Which practices assemble at any moment depends on the local structure of the ecosystem (which elements are available in local time and space).
- Experience, training, and the design of environments can all be seen as ways to bias the probability of the dynamic formation of particular practices (bias the assembly process).
- The stability, resilience, or persistence of a practice depends on the network of relations to other practices within which it is embedded. This includes membership in a family of practices as described above.
- Interlocking relations among practices may produce conditions of multiple determination (over-determination?) of particular features.
- The dynamics of practice formation and maintenance may include positive feedback loops such that the more prevalent a practice becomes, the more probable its formation.
- Learning in the ecosystem includes changes that are outside of individual persons, in artifacts, for example.

6.4. *Cognitive Consequences for Individuals*

- Cultural practices decrease entropy and increase the predictability of experience.
- The richness of the ecosystem creates conditions of multiple determination that promotes reliable induction of internal elements in individuals.
- Possibilities for individual learning depend on the structure of the ecosystem, both because the local ecosystem determines the inventory of available things to be learned, and because family resemblances among practices reduce the complexity of learning processes.
- The generalization of action across contexts is facilitated by family resemblances among practices.

6.5 Cognitive Consequences for Groups

- Interpersonal coordination of practices, including communicative practices, is facilitated by the fact that families of practices exist.
- The distribution of cognitive skills in a community will be determined by the distribution of practices engaged in by the members of the community.

7. Conclusion

Referring to prediction-driven learning in cultural context, Clark says:

Such learning routines make human minds permeable, at multiple spatial and temporal scales, to the statistical structure of the world as reflected in the training signals. But those training signals are now delivered as part of a complex developmental web that gradually comes to include all the complex regularities embodied in the web of statistical relations among the symbols and other forms of socio-cultural scaffolding in which we are immersed. (Clark, 2013, p. 195)

In this article, I have tried to illuminate this web of cultural regularities in which we are all immersed. I have also tried to motivate a set of conjectures about its operation. I am optimistic that some of these conjectures will soon be rephrased as proper testable hypotheses. Then the science of cultural cognitive ecosystems can begin.

Acknowledgements

I am grateful to Charles Goodwin, Tyler Marghetis, Will Newsome, the members of the DCOG-HCI laboratory at UCSD, and two anonymous reviewers for critical comments on an earlier draft of this paper. The research on airline flight decks was supported by a contract from the Boeing Commercial Airplane Group.

Note

- [1] Of course, a linear array of people is not always a queue. For example, when posing for a group portrait, people may position their bodies in a linear arrangement without any sense that one end of the array precedes the other.

References

- Clark, A. (2001). *Mindware: An introduction to the philosophy of cognitive science*. Oxford: Oxford University Press.
- Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford: Oxford University Press.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36, 181–204.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58, 7–19.
- Fauconnier, G. (1997). *Mappings in thought and language*. Cambridge: Cambridge University Press.

- Fauconnier, G., & Turner, M. (2008). *The way we think: Conceptual blending and the mind's hidden complexities*. New York: Basic Books.
- Fodor, J. (1983). *The modularity of mind: An essay on faculty psychology*. Cambridge, MA: MIT Press.
- Fouse, A., Weibel, N., Hutchins, E., & Hollan, J. (2011). ChronoViz: A system for supporting navigation of time-coded data. In D. Tan (Ed.), *Proceedings of the 2011 SIGCHI conference on human factors in computing systems* (pp. 299–304). New York: ACM.
- Geertz, C. (1973). *The interpretation of cultures*. New York: Basic Books.
- Goodwin, C. (1994). Professional vision. *American Anthropologist*, 96, 606–633.
- Goody, J. (1977). *The domestication of the savage mind*. Cambridge: Cambridge University Press.
- Hazlehurst, B., & Hutchins, E. (1998). The emergence of propositions from the co-ordination of talk and action in a shared world. (K. Plunkett, Ed.). *Language and Cognitive Processes*, 13, 373–424.
- Hinton, G., & Becker, S. (unpublished manuscript). *An unsupervised learning procedure that discovers surfaces in random-dot stereograms*. University of Toronto.
- Hutchins, E. (1995a). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Hutchins, E. (1995b). How a cockpit remembers its speeds. *Cognitive Science*, 19, 265–288.
- Hutchins, E. (2001). Cognition, distributed. In N. Smelser & P. Baltes (Eds.), *The international encyclopedia of the social and behavioral sciences* (pp. 2068–2072). Oxford: Elsevier.
- Hutchins, E. (2005). Material anchors for conceptual blends. *Journal of Pragmatics*, 37, 1555–1577.
- Hutchins, E. (2006). The distributed cognition perspective on human interaction. In N. J. Enfield & S. C. Levinson (Eds.), *Roots of human sociality: Culture, cognition and interaction* (pp. 375–398). Oxford: Berg.
- Hutchins, E. (2008). The role of cultural practices in the emergence of modern human intelligence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363, 2011–2019.
- Hutchins, E. (2010a). Cognitive ecology. *Topics in Cognitive Science*, 2, 705–715.
- Hutchins, E. (2010b). Enculturating the supersized mind. *Philosophical Studies*, 152, 437–446.
- Hutchins, E. (2012). Concepts in practice as sources of order. *Mind, Culture and Activity*, 19, 314–323.
- Hutchins, E., & Hazlehurst, B. (1991). Learning in the cultural process. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial life II: SFI studies in the sciences of complexity* (Vol. 10, pp. 689–706). New York: Addison Wesley.
- Hutchins, E., & Hazlehurst, B. (1995). How to invent a lexicon: The development of shared symbols in interaction. In N. Gilbert & R. Conte (Eds.), *Artificial societies: The computer simulation of social life* (pp. 157–189). London: UCL Press.
- Hutchins, E., & Hazlehurst, B. (2002). Auto-Organization and emergence of shared language structure. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language* (pp. 279–305). London: Springer-Verlag.
- Hutchins, E., Weibel, N., Emmenegger, C., Fouse, A., & Holder, B. (forthcoming). An integrative approach to understanding flight crew activity. *Journal of Cognitive Engineering and Decision Making*.
- Langacker, R. W. (1987). *Foundations of cognitive grammar: Theoretical prerequisites*. Stanford, CA: Stanford University Press.
- Michaelian, K., & Sutton, J. (2013). Distributed cognition and memory research: History and future directions. *Review of Philosophy and Psychology*, 4, 1–24.
- Minsky, M. (1988). *Society of mind*. New York: Simon and Schuster.
- Mithen, S. (1999). *The prehistory of the mind: The cognitive origins of art, religion and science*. London: Thames and Hudson.
- Stewart, J., Gapenne, O., & di Paulo, E. (Eds.). (2010). *Enaction: Toward a new paradigm for cognitive science*. Cambridge, MA: MIT Press.
- Talmy, L. (1996). Fictive motion in language and “ception.” In P. Bloom, M. A. Peterson, L. Nadel, & M. F. Garrett (Eds.), *Language and space* (pp. 211–276). Cambridge, MA: MIT Press.

- Vygotsky, L. (1978). *Mind in society: The development of higher psychological processes*. Cambridge, MA: Harvard University Press.
- Weibel, N., Fouse, A., Emmenegger, C., Friedman, W., Hutchins, E., & Hollan, J. (2012a). Digital pen and paper practices in observational research. In J. A. Konstan (Ed.), *Proceedings of the 2012 SIGCHI conference on human factors in computing systems* (pp. 1331–1340). New York: ACM.
- Weibel, N., Fouse, A., Emmenegger, C., Kimmich, S., & Hutchins, E. (2012b, March). Let's look at the cockpit. In A. Hyrskykari (Chair), *Visual attention: Studies, tools, and methods*. Symposium conducted at the 2012 *Eye Tracking Research and Applications Symposium*. Santa Barbara, CA.
- Weibel, N., Fouse, A., Hutchins, E., & Hollan, J. (2011). Supporting an integrated paper-digital workflow for observational research. In P. Pu (Ed.), *Proceedings of the 16th international conference on intelligent user interfaces* (pp. 257–266). New York: ACM.