Linköping Studies in Arts and Science

Dissertation No. 695

Distributed cognition in home environments

The prospective memory and cognitive practices of older adults

By

Mattias Forsblad

LINKÖPING UNIVERSITY

Department of Computer and Information Science
Linköping University
SE-581 83 Linköping, Sweden

Linköping 2016

Chapter 3. Distributed cognition

In this chapter I introduce the theoretical foundations of distributed cognition. This introduction is primarily based on the description of distributed cognition developed by Hutchins (1995a) and colleagues (Hollan, Hutchins, & Kirsh, 2000). However, I will also embrace theoretically overlapping perspectives, such as situated cognition and embodied cognition, under the heading of distributed cognition[3].

When it is suitable I also address studies that have used distributed cognition to specifically understand prospective memory or everyday environments. At the end of the chapter I introduce the methodological foundations for distributed cognition that have guided my field studies.

3.1 Foundational principles

The foundation of distributed cognition can be understood in the light of the critique that Hutchins directs against deep-rooted assumptions in traditional cognitive science. The critique can be generalized to a characteristic of cognition that is often taken for granted: cognition as an internal process of natural and artificial agents. The problem with this assumption is described by Hutchins' (1995a, pp.359-370) re-analysis of the history and the dawn of classical theories in cognitive science.

Influenced by the development of computers Newel and Simon (1976) suggest in their *Physical Symbol System hypothesis* (PSS-hypothesis) that symbol transformations are necessary and sufficient means for intelligence. A predication from this is that humans do symbol transformation because they have intelligence, and computers, since they do symbol transformations, can also have intelligence. However, Hutchins proposes that cognitive science made a fundamental mistake in using the computer as the metaphor for cognition. In the course of using computers to understand human cognition, cognitive science started to cram every part of cognitive accomplishment into the brain, and in turn focus almost all research endeavors on individuals solving problems in resource-stripped experimental settings. Aspects of human cognition like social resources, cultural

[3] I find no theoretical problems with doing so because distributed cognition is a perspective on all cognition. However, I know that that this is not the standard way of categorizing. Instead, I find several authors who treat situated cognition as the genus for distributed cognition. See, for instance, Robbins and Aydede, 2008, and Roth (2013) both of which provide a thorough account of the development and expansion of the idea of situated cognition, where ideas that can be derived from distributed cognition are described as ideas within situated cognition.

resources, and sensory organs were nowhere to be found in early descriptions of cognitive accomplishment.

Through Hutchins' perspective, cognition is still the processing of symbols and information. Specifically, cognition is "...computation realized through the creation, transformation, and propagation of representational states."(1995a, p.49). He also notes that both humans and computers have sufficient architectures to do symbol transformation. However, distributed cognition holds that to understand cognitive accomplishment researchers should not *a priori* assume a specific unit or center of analysis (e.g. the brain or the computer, as early cognitive science did). Instead the main theoretical principle of distributed cognition is to focus the analysis of cognition on functional relationships between parts across disparate architectures that are involved in a given activity. The parts investigated are therefore not chosen according to some *a priori* preference of architecture. This general assumption about cognition is also the main reason that distributed cognition is a perspective on all cognition (Hutchins, 2013).

As an example of disparate architectures, consider the following example by Lave (1988). There are a number of ways to solve a given math problem, for instance, 75x114. One way is to use pen and paper and the placeholder strategy that is taught in school. A second way is to use a calculator, and a third way could be to ask a friend to help remember the results of steps of a calculation that you do in your head. Lave's point is that "the product may be the same – but the process has been given structure – ordered, divided into units and relations, in action – differently in each case." (p.98). Hence the cognitive processes are, from a distributed cognitive perspective, different.

In Hutchins' description cognition is also a type of cultural process. Cognition is a cultural process that provides solutions to frequently encountered problems (Hutchins, 1995a, p.354). As a consequence of viewing cognition as a cultural process, the first unit of analysis should always be the cultural-cognitive ecosystem (Hutchins, 2010a) because it is the cultural processes that have established solutions to problems and situations.

Hutchins (1995a) states that a consequence of embracing a distributed perspective is that what we know about individual cognition must be forgotten and described anew through what he calls cognitive ethnographies. In cognitive ethnographies, practices in real-life settings are the focus (this will be thoroughly described later). Only by doing research in *the wild*, in naturally occurring settings, can we account for what individuals actually do to accomplish goals. Despite this rather negative view of past research it is important to understand that distributed cognition also comprises individuals with abilities.

Soon after the release of *Cognition in the Wild,* Hutchins (1995a) was criticized for (among other things) minimizing the role of individuals, and attributing too much cognitive accomplishment to extra-individual sources (see description by Latour, 1995; and critique by, for instance, Nardi, 1996). Ironically, distributed cognition, and specifically Hutchins, was suddenly criticized for doing the same thing that Hutchins had criticized cognitive science for doing in its first decades, but in the opposite direction. But this was largely a misinterpretation of what distributed cognition was, namely a perspective on all cognition (Garbis, 2002; Hutchins, 2010b). Hutchins also explicitly suggests that an analysis according to distributed cognition should not start on an individual level if there are reasons to believe that some larger unit of analysis is more appropriate. But to have a complete description of a large socio-cultural system, the analysis should always at some point venture to the individual level (see Hutchins, 1995a, p.50), which is something Hutchins does in his analysis of the accomplishment of the navigation activity on a larger vessel at sea.

What past research on cognition, and so also prospective memory research, has focused on is individual abilities used in a specific cultural setting, the laboratory. This knowledge is not without benefits but should be used cautiously before assuming that the same mechanisms apply in other settings, for instance settings often ascribed to real-life. Therefore, in cultural settings for which an individual perspective traditionally has been applied, the research strategy of distributed cognition is instead to start assuming that individual abilities have potentially equal roles for cognitive accomplishment as any other resources in the setting under study. Only by doing this we can at the end understand individual contributions for information flow and transformation.

Even though distributed cognition does not assume a center or boundary for cognition in advance, it is both an empirical question and a relative matter if a specific center or boundary exists for some cognitive system, or if the system lacks a clear center, or has multiple centers. It is an empirical question because it is a function of the density of information flow between parts. It is a relative matter because the threshold for saying that information flow is enough for something to be included in the functional system is a matter of judging the relevance in relation to current research questions. As a consequence, the unit of analysis could therefore be confined to the brain, the brain and some tool, or some other social being, groups of people, or even environments.

Hutchins (2013) formulates two general principles (or, in fact, hypotheses) of cultural-cognitive ecosystems. First, a cognitive system does not have uniform connectivity. All intelligent systems (for instance, brains) have connections between their parts with a variety of higher and lower information densities. Second, there are formal principles that "operate at multiple scales in cognitive systems" (p.12). An example of such a principle is suggested by Clark (2013) where he argues that the brain is a hierarchical prediction machine. Hutchins (2013) suggests another such formal principle, namely that there are structures in the cultural world that can be exploited by individuals across time and generations.

Below I go further into principles for information flow. Most of the information flow principles mentioned below originate from Hutchins' (1995a) seminal book *Cognition in the Wild,* where he presents an extended analysis of a cognitive ethnography conducted onboard a U.S. Navy ship. Instead of limiting the analysis to competent individuals, as a traditional cognitive science study would have done, Hutchins finds the ship with its crew members and cognitive tools, specifically on the navigational bridge, to be a socially distributed cognitive system. The core conclusion is that how information flows between parts of the system, and how representations transform in the interaction between parts, is the essence of how activities on the navigational bridge are accomplished.

3.2 Information flow principles

There are a number of information flow principles (see Blandford & Furniss, 2006 for a condensed list of principles) which determine information flow and which can be used to determine cognitive accomplishment. The most basic principle is information movement or propagation. This is the most basic principle because information movement is what allows parts in a cognitive system to coordinate operations. Another principle is information transformation, which is when representations change. This can, for instance, occur when information moves from one part to another, or within parts of a system. An example of how information can be transformed across social architectures is the way that humans filter information when they re-tell information that they have previously heard or read from some source. How information is filtered can have direct consequences on cognitive performance, both good and bad.

Additional information principles include information buffering, information hubs, and information bandwidth. As mentioned previously, intelligent systems seldom have an even distribution of information content and flow, and therefore *information buffering* and *hubs* functionality can often be found. An information buffer is a medium that holds incoming information for later retrieval. Hutchins (1995a, p.194) describes how, for instance, information buffering can be used to communicate between individuals whose roles do not allow for instantaneous communication. Buffering can therefore be a crucial part of complex environments where information flows and transforms at a high pace.

Information buffering can be necessary in a cognitive system that has *information hubs* (Blandford & Furniss, 2006). Information hubs are parts in a cognitive system where several information channels meet. Buffering can therefore be necessary to manage large amounts of incoming information. Because of their access to, or gathering of information, hubs are potentially where most information of a system's cognitive states can be found. In Hutchins book the navigational bridge can for instance be seen as an informational hub that coordinates information from several sources.

Another principle is *information bandwidth,* which has an analytical focus on the richness of information flowing between parts in a system. Consider, for instance, a face-to-face communication situation in contrast to communication through a telephone line. Through the use of both verbal and non-verbal information, communication face-to-face can potentially be richer, in other words it can have a larger information bandwidth. Hutchins (1995a, pp.227-228) discusses how, for instance, communication about two landmarks between an operator and a recorder through restricted media can create problems. Specifically, if only one of the individuals can see the landmarks, the landmarks need to be translated into symbolic representations and communicated to the other individual.

Note that richer information should not be understood as something that is only positive. Hutchins describes (pp.229-230) how bandwidth should not have analytical precedence over informational content and representational shapes. On the ship most of communication is verbal, and therefore there are established practices for how to communicate verbally that establish robust information flow. Therefore, in many situations there is no need for face-to-face communication. In the case of the landmarks there is no fundamental reason to suppose that the possibility of using gestures to point at a landmark would lead to better cognitive accomplishment. But Hutchins also provides examples in which rich bandwidth and a shared world are important for performance. In one example (p.232), during a mismatch between navigation plots, a discussion between two individuals with a shared chart becomes necessary. In Hutchins' example, face-to-face communication in relation to the shared world was necessary to negotiate meaning, and to establish trust between parts of the system.

The examples above all operate on a relatively restricted temporal dimension. However, the functional relationships between parts in a system can additionally be viewed as operating on different spatial and temporal scales, which in practice means that a distributed cognitive analysis continuously shifts between different scales of description to explain cognitive accomplishment (Hutchins, 2013). Hollan, Hutchins, and Kirsh (2000) suggest three general spatial and temporal scales: across internal and physical resources, across the social environment, and across time. I will consider these three ways with concrete examples below, in turn. Note that the three ways in which cognition can be distributed are analytical tools for describing cognitive accomplishment through informational flow and transformation. In reality, they can be highly dependent on each other. For instance, for some information to be transmitted across time the information needs some physical (e.g. book) or social vehicle (e.g. a teacher).

I start by describing distribution across time, which sets the premises for cognition as a process across different temporal scales.

3.3 Distribution across time

That cognition can be distributed across time means that the products of past events can shape future cognitive events (Hollan et al., 2000). From an individual developmental life-span perspective this is not strange and has been acknowledged for some time (Craik & Bialystok, 2008). But when viewing distribution across time from a distributed perspective it is not just individuals that change. It is also important to acknowledge how the cultural, material and social world which humans learn from and act upon changes. Rajkomar, Blandford, and Mayer, 2013 present a summary of the heterogenous uses of temporality within distributed cognition research. Below I consider a few of these uses.

Distribution across time can, for instance, be viewed as a social and material *individual* learning process. By borrowing the idea of *internalization* from Activity theory (Vygotsky, 1978) distributed cognition suggests that to a large extent cognition is, through development, an inter-individual process where individuals internalize parts of this cognitive process. Parts and processes that previously existed only in the external world become *internalized,* i.e. primarily intra-individual in nature. Similar ideas about the internalization of culturally-situated cognitive processes have also been seen in the work of Halbwachs (1980/1950) on social collective memory and Donald (1991) on cultural knowledge accumulation through the use of exograms. Halbwachs (1980/1950), for instance, discusses expertise among classical musicians, where he argues that their musical ability has not just developed as an ability to read the conventional musical notation system. Musicians have also assimilated the representational notation system, and through the process of assimilation transformed their intracranial cognitive processes. As mentioned in these cases, symbolic structures are first found in the cultural setting and only later (through, for instance, some learning process) located in the brain.

Distribution across time can also be viewed on a *collective or environmental* level, where it is the cognitive ecosystem that internalizes and learns new ways to handle situations. According to Hutchins, distribution across time works to "solve frequently encountered problems" that groups of individuals experience (Hutchins, 1995a). Consider, for instance, Hutchins and Hazlehurst (1991, see also Hazlehurst & Hutchins, 1998) who use a computer simulation to explore how the individual and collective reciprocate through the mediation of cognitive tools. The case for the simulation was how a community could learn the phases of the moon across generations. They conclude that information that is built into tools is the main contribution to collective learning. But for this to happen, the cultural process is dependent on individuals that build knowledge into tools, tools that keep on existing when individuals die (see also Salomon, 1993). Hazlehurst and Hutchins show that the collective could assimilate the phases of the moon better when individuals interacted with smart tools that earlier generations had developed.

The example above explores distribution across generations. But distribution across time can also be understood as changes across various temporal scales. De Léon (2003) makes an explicit point of this when focusing on the development of tools across three different temporal scales: 30 minutes (development of workspaces while cooking), 30 years (development of a spice shelf), and several hundred years (development of pistols). By conducting what he calls *cognitive biographies of* *things,* a great deal can be said about the cognitive process across time. For instance, in one section de Léon focuses extensively on changes to a spice shelf and its cognitive consequences for cooking.

Different dimensions of distribution across time is also something that is highlighted in Hutchins' introductions to distributed cognition. Hutchins (1995a, p.372) describes how any human practice in a given moment can be understood across three developmental dimensions: the conduct of an activity, the development of the practitioners, and the development of the practice. For examples of each, consider Hutchins' seminal case: navigation.

In any moment when humans navigate there are plenty of interactions between humans and tools before reaching a destination. These interactions coordinate representational media and information in such a way that sub-goals, and eventually the goal of reaching a destination, are met. In Hutchins' model this is the conduct of the activity.

In the same moment, we can understand the conduct of the activity as deriving from the cultural specification of how to navigate. This specification has been transferred across generations of navigators through, for instance, the tools that are available to the navigators and the conventions that exist in a community of navigation. This is the development of the practice. The development of the practice is not as informationally reciprocal in the given moment as the conduct of the activity is. As a result, states of the practice do not change as fast as states across the conduct of an activity, but the practice still shapes any navigation activity.

Also in the same moment, we can understand the conduct of activity as deriving from the expertise of the navigators, which is also something that has developed across time. The three dimensions, according to Hutchins, all account for the development of human practice. When a goal or sub-goal has been reached in the conduct of navigation the practitioners and the practice may have also developed. However, the three dimensions do not develop equally fast. Practitioners and practices do not develop at the same rate as activities are accomplished. It takes years to become an expert navigator and it takes an even longer time to change the nature of navigation. A navigation activity, however, does not usually take years.

A recent empirical example of distribution across time that includes observations of prospective memory processes comes from Rajkomar, Blandford, and Mayer (2013), who report on a cognitive ethnography on home hemodialysis. They do not focus on human practice over longer timescales. Instead, they specifically look at the cognitive rationale of using time as a medium to reduce cognitive complexity across the conduct of the activity, and to some extent the development of the practice. They observed and interviewed nineteen patients and their home dialysis systems.

The authors note that time can serve a number of cognitive functions in short-term activities. For instance, repetitive points in time can be used as conceptual (and also to some extent external) cues for prospective memory tasks. Many participants assigned their weekly or monthly disinfection routine to a specific day, weekday, or day of the month. Some participants also distributed the temporal placeholder by marking their calendars. Marking the calendar seemed to be more important when the routines were not weekly, for instance, when they were every second week. The authors also note that temporal distribution of activities can be used to avoid omission of steps. One caregiver, for instance, initially started preparing to take the patient off the machine five minutes before the end of treatment. But after several experiences where steps were omitted he started initiating preparation 20 minutes in advance. This, he argued, lead to fewer omissions because he was less time-constrained and could give the steps more thought. Finally, the authors also note that space was used as a medium for distributing time. For example, people were observed to place time-sensitive medicine in their home environment based on where they expected to be when it was time for each medicine.

Another perspective on cognition that also views the process of cognition across time on a systemic and collective level is Cognitive Systems Engineering (Hollnagel & Woods, 2005, 1983). CSE was developed in parallel to Hutchins' distributed cognition and shares many aspects with distributed cognition. Most notably, they share an expanded unit of analysis, and specifically a dynamic unit of analysis. But CSE, even more than distributed cognition, focuses its empirical efforts on complex high-risk environments and human-machine interaction. CSE developed in response to a desire for new theoretical tools as a consequence of major accidents, most notably Three Mile Island. Compared to distributed cognition this perspective has a more explicit focus on failures and negative incidents, and a motivation to proactively avoid larger failures. From a societal perspective, failures and negative incidents are less acceptable in complex, high-risk environments than in many of the domains that are studied under the heading of situated and distributed cognition.

In CSE the principal concept is "control" (borrowed from Cybernetics, Ashby, 1956) and how the cognitive system maintains control within acceptable boundaries over necessary functionalities. Interestingly, CSE initially (Hollnagel & Woods, 1983) kept to the information-processing paradigm of cognition. But in recent years (Hollnagel & Woods, 2005) the information-processing paradigm is no longer the primary one. The new paradigm uses a notion of joint cognitive systems that does not primarily focus on informational interaction between humans and machines, focus is instead on how they maintain control together, irrespective of whether there is information involved or not.

In CSE, control is analytically embodied through a number of theoretical models. The most common is COCOM (Contextual control model). COCOM views processes as control loops where each control loop represents a function of the cognitive system. Since CSE takes a systemic perspective, the parts of a cognitive system and the interactions between parts that account for a specific function often extend beyond a single individual. Functions are also hierarchical, meaning that that each overarching function has several minor functions that are needed to sustain control of the overarching function. The notion of levels of a cognitive system is not necessarily different from Hutchins' version of distributed functionality. For instance, in Hutchins' (1995b) article *How a cockpit remembers its speeds* the function specifically under scrutiny is part of a series of functions that are necessary for successful landing.

Another common element is that accountability cannot normally be assigned to specific individuals, either through a contemporary CSE or through a distributed cognitive perspective. CSE often holds that what is called the human factor or human error is in fact often not human error. Instead, the norm within CSE is to seek systemic properties to explain exceedances of accepted safety or performance boundaries. This model can also be used to understand prospective memory in everyday life. People's individual abilities to remember are likely to affect outcomes of attempts to remember something. However, when cognitive processes take place in dynamic physical and social environments, the outcomes are also derived from systemic properties larger than the individuals. It is in part an empirical question to determine which properties are relevant.

To summarize, cognition across time can be understood across a continuum of spatial and temporal ranges. Below I consider two types of spatial dimensions: physical and social architectures. Above I have already considered examples of how cognition can be distributed across social and physical architectures. In the example of how a community could learn the phases of the moon, cognition was socially and physically distributed through the residual knowledge in cognitive tools that traveled across generations. Below I go into more detail of each, in turn.

3.4 Distribution across social architectures

The principles above are descriptive parts of how information flows in any cognitive system. The arenas and media through which information flows make up an important part of information flow as a whole. They establish the ecologies that allow information to be distributed and coordinated.

Blandford and Furniss (2006), for instance, acknowledge *informal communication* situations as social arenas that can facilitate information propagation across a cognitive system, even if the arena is not explicitly related to a task at hand. Communication ecologies that extend outside the pre-defined task execution is something that Hutchins also acknowledges. Hutchins (1995a, p.268) describes the notion of an agent being within a *horizon of observation*. Hutchins specifically mentions two ecological factors that widen the horizon of observation. First, the openness of communication lines or arenas, which allows novices to listen in on communication between more experienced crew members. Second, the openness of tools, which allows individuals to perceive current cognitive states without needing to communicate with others and possibly interrupt important ongoing processes. The navigation chart on the bridge is discussed as a general tool that provides a wider horizon of observation than smaller computer screens do. The horizon of observation is a basic characteristic that describes the potential for information to move across a system. If, for instance, an agent cannot hear or see some information, the agent cannot acquire that information.

Another type of ecology, and perhaps the clearest example of social distribution, is a team that works to solve problems. Teams are the subject of Hutchins' seminal example and have continued to be a common type of unit to study from a distributed perspective. Teams are interesting from a distributed perspective for several reasons. One reason has to do with the characteristics of teams. Each member of a team has a specific role in the pursuit of their team's objective (Orasanu & Salas, 1993). Because no single individual in a team can accomplish the goals by themselves, the members need to have some social and communication structure to accomplish the goals. Hence their overall cognitive accomplishments need to be distributed across the individuals. Note that Hutchins (1995a, p.198, and forward) mentions examples of what can be called the use of *behavioral trigger factors*, where members of a team do not always need shared goals, or need to communicate to perform in accordance with the team. Instead they are triggered to act in certain ways by the circumstances. As an analogy, they can be described as acting like an ant colony, where each individual is triggered by features of the environment in ways that are for the benefit of the colony.

An interesting case of distribution across disparate architectures, which includes teams and triggering factors, and which takes place outside of the highly technical and modern domain, appears in Tribble and Sutton's historical studies of cognition in early modern theatre companies (Tribble, 2005, 2011; Tribble & Sutton, 2011). They describe how individuals could perform up to six different plays in a week without the need for regular practice time or complete and constant access to manuscripts. How did individuals cope with such excessive memory demands? To some extent, ecological pressure shaped certain individual skills. She describes how individuals needed to practice what is referred to as *cognitive thrift*, which is "to learn exactly what was needed, and no more." (Tribble, 2011, p.58). But the theatre as a cognitive ecology is also shown to alleviate individuals from cognitive burden.

Characteristics such as designated entrance and exit doors, and an elaborate computational device in the form of a written plot of the ongoing play, which hung in preparatory areas, are described as powerful social and physical tools that "maximize individual contributions" (Tribble, 2005). Tribble (2011, p.49) also describes how what are called internal exits are rarely mentioned in historical sources. An internal exit is when an actor exits a scene without a cue. Rather than this, the most common case is that actors are either cued by the dialogue to exit, or helped off the stage by fellow actors. It is also important to note that the use of metrics and verbal patterns in the manuscripts were necessary parts of the remembering process. The most important thing was not to present everything verbatim from the manuscript, rather, actors needed to follow the patterns of the manuscript and recognize cues for when to speak; in other words, they needed to deliver a smooth performance. In relation to this Tribble describes how Shakespeare's later plays used more irregular rhythms of lines, and more short and shared lines than his earlier plays. The task for the actor in his later plays therefore became more of a pattern-recognition task, which involved responding to the lines of fellow actors and the rhythm of a situation, rather than recalling long passages by oneself.

Tribble's examples show, first, that team characteristics are important for individual performance, but also that those team characteristics were built into the structures of manuscripts and plots of plays. To use an analogy from Hutchins' (1995b) paper *How a cockpit remembers its speeds*, it is the play itself and the physical and social structure of the theatre that remembers what actors should do and say.

The analogy of the *remembering theatre* can also be said to work in contemporary times. But by taking a distributed perspective Tribble also shows that contemporary theatres are a very different cognitive ecology than the theatres in Shakespearean times. In her own words (Tribble, 2005, p.148): "[the] 'general model of acting cognition' is bound very specifically to late-twentieth-century acting practices, which are in turn based on assumptions about character and subtext derived from modern acting theory. Moreover, these practices are the results of institutional conditions such as long rehearsal periods, a relative scarcity of new plays, and, finally, the exigencies of memorizing prose rather than verse." (See Noice et al., 2004 for descriptions of contemporary theatre practices.) Tribble's comparison emphasizes the general importance of understanding cognitive ecologies as socially and historically situated. Information flow and transformations can take significantly different forms as a consequence of institutional and social structures.

Schwartz and Martin also consider the various forms of distributed systems (Schwartz and Martin, 2008). They take learning situations as their primary case. Specifically, they consider the ability to learn by categorizing distributed cognition as something that can take place between stable/adaptable environments and stable/adaptable individuals, which they represent as a grid (see table 1 below).

Table 1: Four Quadrants of Distributed Cognition (adapted from Schwartz & Martin, 2006)

Individual	***Adaptable***	(1) Induction	(4) Mutual Adaption
	Stable	(2) Symbiotic Tuning	(3) Repurposing
		Stable	***Adaptable***

Physical and Social Environments

Using this grid they define types of learning processes. The authors give examples of each type of learning process.

An example of induction (quadrant 1) is the way that a student (adaptable) tries to make sense of a teacher's words (stable). Examples of symbiotic tuning (quadrant 2) come from Hutchins, and are, in fact, the primary type of learning situation presented in most of the literature on distributed cognition. In these examples the individual is a professional (stable) and the environment has been specifically designed to accomplish a specific task (stable). In such cases, learning is characterized as a process of increasing efficiency within the system and the interdependence of its components. Examples of repurposing (quadrant 3) are, for instance, what Kirsh (1996) talks about when someone arranges the environment to serve some function. For example, imagine a professional cook (stable) who plans a dinner in a non-professional kitchen (adaptable) in which she has never cooked before. In this case, she is likely to rearrange features of the environment so that the process becomes closer to symbiotic tuning. An example of mutual adaptation (quadrant 4) might be a home gardener (adaptable) constantly learning new things as the garden changes and some plants survive while others die (adaptable). (Schwartz & Martin, 2008)

The grid adds to traditional distributed cognition research because it gives structure to cognition as more than problem-solving that is done by professionals in intelligently-adjusted environments (Schwartz & Martin, 2008). In other words, it can be used as the basis for a discussion of types of cognitive systems. This is relevant for the purpose of this thesis since I analyze an environment that is not typical for a distributed cognitive analysis.

In distributed cognition, social structures are in themselves forms of cognitive architectures (Hollan et al., 2000), both when considering immediate direct communication and social structures that individuals have experienced in the past. This is because, like any cognitive architecture, they determine information trajectories and transformations. Consider the following example. In the paper *Professional Vision* Goodwin (1994) explores how members of a profession, through the examples of archeological field excavation and legal argumentation, observe the world in specific ways. At the foundation of Goodwin's concept is that the perception of reality is determined by properties of belonging to a social group (or groups). A profession is a kind of social group where social structures shape perception according to what it means to be a member of the profession. To become a member of a profession it is not just about adopting a way of thinking, it is also about adopting a way of perceiving the world.

Goodwin presents three ways professional vision is manifested in a profession: Coding, highlighting and graphic representations. To exemplify coding, Goodwin discusses the Munsell coding chart that is used world-wide by archeologists to color-code dirt. In the example, archeologists-in-training are scrutinized. Goodwin's point is that novices learn how to perceive dirt like a member of the profession. Despite the fact that the chart does not make the coding process trivial, the chart guides perception according to the profession. Highlighting can be exemplified in the same context. In other examples Goodwin observes how archeologists highlight aspects in the dirt by drawing lines to circumscribe relevant areas. Doing so causes the task of perceiving to be explicitly shared between members, and novices learn to see what the professionals see. For the final manifestation Goodwin discusses the practices within a profession that create and articulate graphical representations. In an example Goodwin describes how two archeologists create a map of the excavation grounds. The map includes several layers of information that are distinct to the archeologists' ability to see in their profession. The graphical representation is in itself a product of a specific profession, a specific visualization of the excavation grounds.

The social structures that determine and constitute professional vision are an architecture of cognition because they determine how information flows and transforms in those specific situations when a member of a profession perceives the world. Even when a member of a profession is alone when perceiving, that person perceives the world as a member of the profession, and not only as an individual.

Hutchins (1995a) also has several sections that cover the specifics of social distribution, for instance, the occurrence of hierarchies and the reaching of consensus between members. Hierarchies and authority can be important in organizations where the nature of reality is socially defined, as in the case of courts (p.256, see also Goodwin, 1994). In other situations consensus might be required. Hutchins shows that the type of organizational structure (horizontal or vertical) has direct consequences on cognitive processing. Specifically, Hutchins (p.261) writes: "Where the power to define the reality of situations is widely distributed in a 'horizontal' structure, there is more potential for diversity of interpretation and more potential for indecision. Where the power is collected in the top of a 'vertical' structure, there is less potential for diversity of interpretation, but also more likelihood that some interpretation will find a great deal of confirmation and that disconfirming evidence will be disregarded." These different forms of distribution of power and cognition are something many organizations need to deal with, which can have important consequences for performance.

The social dimensions of distributed cognition also involve the specifics of communicative practices between agents. This is a vast area, but one study that looked specifically at older adults is Harris, Barnier, Sutton, and Keil (2014, see also Harris, Keil, Sutton, Barnier, & Mcilwain 2011). They studied a number of older couples, specifically investigating communication practices between couples in a sit-down situation, where the couples were supposed to collaboratively remember certain facts and past experiences together. Through an analysis of this process they find that the following communication practices inhibit performance: (a) Incompatible reference to expertise regarding a past experience, where, for instance, a partner is seen a priori as the memory expert. (b) Strategy disagreements, where, for instance, each partner interprets a question differently. (c) Corrections, where, for instance, one partner interrupts a story to correct details. They also find that the following three factors facilitate recall: (a) The occurrence of cueing. (b) Production of new information in response to cues. (c) Repetitions. (d) The occurrence of failed cues. For this last factor, Harris et al. (2011) propose that the willingness to cue is itself an important contributor to successful collaborative remembering.

These findings give some indication of the complexity of the process involved in having a successful socially-distributed memory system in non-work environments. But importantly, as the authors also note, the situation in the above case is experimental. Despite the fact that the couples being investigated are naturally-occurring couples and that we can therefore assume that their communication practices are somewhat similar to how they communicate in real life, we do not know how similar these experimental settings are to communication in normal everyday settings. Here it should also be noted that the social dimensions of distributed remembering in real life everyday situations is, to my knowledge, far more frequently researched than physical dimensions (see, for instance, Bietti, Stone, & Hirst, 2014; Hirst & Levine, 1985 for examples of such foci)[4].

Although Hutchins' original work includes many examples of how information flows and transforms that are social in their architecture, the same principles can be applied to understand flow and transformations across physical architectures. Principles such as information hubs and bandwidth can be seen and exploited in physical structures. In other words, when using a

[4] Since this thesis focuses primarily on the physical aspects of distributed cognition there is a large category of literature which I omit from this review.

distributed perspective there is often nothing special about social distribution in terms of principles of information flow, despite the fact that there can be special aspects of social interaction (as, for instance, was noted in the study by Harris et al (2014) described above). In many of the above examples the social distribution is also mediated by technological means. For example, as previously mentioned, the layout of the physical setting can have direct consequences for social distribution and cognitive accomplishment. The horizon of observations is one such principle. The physical ecology can also establish other practical situations. Hutchins (1995a, p.197), for instance, mentions an example where certain operations are impossible to manage for one individual because of the physical allocation of equipment, even though the operations would have been cognitively possible for the individual to manage if the equipment had been arranged differently.

A study on everyday life that combines the use of social and physical architectures is Wu et al. (2008), which uses a distributed cognitive perspective to understand how families cope with activities when a member has diagnosed memory problems. The authors report a number of practices that seem to have positive outcomes for performance: the use of redundant information, frequent and repetitive synchronization of information between family members, and instant and continuous awareness of updates. This list of practices is close to what is sometimes considered important for cognitive outcomes in professional settings (see, for instance, Hutchins, 2000; Lützhöft & Dekker, 2002).

Below I consider versions of physical distribution of cognition in more detail by considering the physical environment in relation to individuals.

3.5 Distribution across physical architectures

The descriptions that appear in the literature of the physical mode of distribution have most often not been described as distributed cognition. Instead, situated cognition has been the principal concept. This is because the empirical center is the individual and her immediate physical environment. There is general agreement between distributed and situated cognition that physical arrangements can alleviate internal cognition, and therefore can become coupling structures. However, there are a number of theoretical concepts within this notion that clarify cognitive mechanisms.

This section will, for instance, focus on the continuum of explicitness in representational functionality. Note that in the examples below the external structures will be material, but the principles could, in theory, also be used to understand distribution across the immediate social environment.

3.5.1 Situated and contextualizing agents

The core principle of situated cognition is situatedness, which has a specific focus on individual brains' directedness toward, and coordination with, external structures (Clark, 1997). This principle emphasizes moment-to-moment coupling with the external environment, which is reflected in the recent paper by Clark (2013) mentioned previously (3.1), where he argues that brains are, above all, prediction machines that constantly match incoming information with previously held expectations in order to support actions and perception (see also Clark, 1997, 2006)[5].

[5] Related concepts to the coupling idea include extended mind (Clark & Chalmers, 1998), embedded cognition (Rupert, 2004), and enactivism (Varela, Thompson, & Rosch, 1991), all which I omit from my review because I have not used them in my empirical analysis. However, some of these ideas are mentioned in my description of embodieness.

Humans are context-framing beings, which means that they are constantly in a dynamic relationship with their settings and the experience of being in those settings. This idea can, for instance, be seen in Clancey's (1997) definition of situated cognition: "To summarize, cognition is *situated*, on the one hand, by the way conceptualizing relates to sensorimotor coordination and, on the other hand, by the way conceptualization, in conscious beings, is about the agent's role, place, and values in society. Thus, situated cognition is both a theory about *mechanism* (intellectual skills are also perceptual-motor skills) and a theory about *content* (human activity is, first and foremost, organized by conceptualizing the self as a participant-actor, and this is always with respect to communities of practice)" (pp. 27–28).

To my knowledge, one of the first researcher to elaborate extensively on this idea of context-based framing from a cognitive science perspective was Jean Lave (see e.g. Lave, 1977, 1982, 1988; Lave, Murtaugh, & de la Rocha, 1984). From her perspective the mechanisms of cognition can never be understood if we refrain from studying them in their everyday habitats, where cognitive processes are seen as coupled with cultural and social orderings. This also means that how individuals interpret situations and settings is a consequence of their social and cultural experiences. This notion of individuals' connections to cultural context is close to schema-theories that, within psychology, were separately originally elaborated on by Piaget and Bartlett. This notion is also similar to Hutchins and Goodwin's cultural perspective on cognition, which states that humans perceive and think according to experiences they acquired through belonging to a certain group and physical setting.

Since humans are mentally connected with their immediate environment they are in the position of coordinating features of the environment with internal resources. The immediate coordination between internal and external processes has been extensively characterized and specified by David Kirsh and Andy Clark, in their respective writings. Clark (2008) pictures human thinking through what he calls the Principle of Ecological Assembly: a human being is a "canny cognizer [that] tends to recruit, on the spot, whatever mix of problem-solving resources will yield an acceptable result with a minimum of effort." (p.13). The principle suggests that the human (the canny cognizer) utilizes internal and external resources that are available in the specific situation without showing preference for the kind of resources used. If a human decides, deliberately or non-deliberately, that it is efficient to use the external world to manage information processing, then they will do so. At this point, we can already note that Clark's description of the situated human is a sketchy account that does not address the issues of coordinating internal and external resources. However, Clark's theoretical ideas emphasize in situ problem-solving and understanding of cognitive situations.

The idea of in situ problem-solving can be traced back to Lucy Suchman's (1985) concept of *situated action* and Jean Lave's (1988) studies on routines and arithmetic abilities in supermarkets. They criticize previous research within cognitive science, and specifically research on artificial intelligence, for emphasizing the importance of internal plans. In one example Suchman describes two humans trying to make sense of a copying machine. By contrasting this process against how problems were supposed to be handled according to the manual, she concludes that the path towards the objective was not met by following a pre-specified plan. Instead, the objective was met by acting on the changing physical and communication situations that took place during the process of understanding the machine. Through her analysis, Suchman re-specifies what plans (internal or external) are by stating that plans are a resource for, and not a controlling structure of, cognitive accomplishments. Suchman also specifies plans as weak resources that are necessarily vague, and that in retrospect filter out what actually happens in specific situations. I interpret this description primarily as a reaction against traditional cognitive science. I do not think that plans are irrevocably weak or vague resources in principle. Rather, the weakness and vagueness is relative to situations, problem definition, and experiences of agents involved, all in accordance to foundational principles of situated cognition defined by, among others, Suchman.

Lave and colleagues (Lave, 1982, 1988; Lave et al., 1984) make extensive cases on the topic of education and arithmetic abilities. Specifically, she compares arithmetic abilities in more formal settings with less formal settings, such as apprentice-based tailor settings in Liberia and American women shopping in supermarkets. She concludes that arithmetic abilities, specifically how we cognitively solve arithmetic problems, cannot be generalized across these settings. When comparing in situ supermarket math tasks with formal math tasks she finds that the cognitive process is shaped, as in the case of in situ problem-solving in the previous section, by the resources available in each setting.

For instance, in formal settings an economic best-buy problem is usually solved by, for each product, dividing quantity by price and seeing which product has the lowest unit price. But in the supermarket it is observed, through think-aloud protocols, that when most participants compare two (or sometimes three) products they instead *estimate* if the price^1/price^2 quotient is larger or smaller than the quantity^1/quantity^2 quotient. Consider an example where product A costs 25 SEK and weighs 200g and product B costs 30 SEK and weighs 275g. In a formal setting the correct way of calculating the best buy is to calculate 200/25 and compare the result with the result of 275/30. Instead, what Lave finds is that people compare the ratio between 200 and 275 with the ratio between 25 and 30. In other words, the heuristic is: does the increase in cost somewhat pay for the increase in weight?

One possible reason for this strategy is that this is how information is presented in supermarkets. Price is compared with price, and quantities are secondary information that shoppers need to search for. This strategy works as long as the number of products is low, which is often the practical case in supermarkets. But what is taught in formal school is a universal method that works independently of the number of quotients quantified. Lave's message is that arithmetic performance, and cognitive performance in general, is significantly determined by the settings in which they operate. In another example she describes how a 12-year-old vendor (M) selling coconuts in Recife solves simple arithmetic problems in a market (Carraher, Carraher, & Schliemann, 1982 as qouted and discussed in Lave, 1988, p.65).

Customer: *How much is one coconut?*
M: *35.*
Customer: *I'd like ten. How much is that?*
M: (Pause.) *Three will be 105; with three more, that will be 210.* (Pause) *I need four more. That is…* (pause) *315… I think it is 350.*

The base of three used by the vendor is not what is being taught in formal schools, but it is a strategy that works in the market setting for this vendor because three is a common number of coconuts sold in a group. In the market it is also shown that five vendors together had a 99% correctness rate on arithmetic problems similar to the one above. In contrast, when they are measured on equivalent arithmetic tasks in formal settings with pen and paper the vendors have only a 74% correctness rate. Lave concludes that it is not just the available resources that shape and frame the cognitive process. The cognitive processes people employ are driven by people's experiences of being in similar situations. People tend to use pen and paper in school settings because that is how math is conducted in school.

In Lave's reasoning there are similarities to how distributed cognition has depicted professional settings. Similar to professional settings, the experience in a specific environment shapes how situations are perceived and which cognitive protocols that are used. The difference is that what we do in supermarkets is often not considered as a profession, because there is no formal education for supermarket shopping. But certainly, the participants Lave describes act and think in accordance with a social collective, which still shapes cognitive processing. She notes, for instance, that in supermarkets the best math solution is not the only determinant for strategy use. Specifically, she notes that math problems in supermarkets are often abandoned, and that strategies people employ can be determined by the socio-economic background individuals have.

An overall conclusion from Lave's research is that when we want to understand cognition we need to study activities in which cognitive situations occur, where the activity and setting-specific socio-interactional protocols of the people we study determine the cognitive process.

Lave makes an extensive point about the expectancies and routines of everyday life. In her definition of everyday life, people's expectations and routines are central. "*It is the routine character of activity, rich expectations generated over time about its shape, and settings designed for those activities and organized by them, that form the class of events which constitutes an object of analysis in theories of practice.*" (Lave, 1988, p.15). Because of the rich experiences people have in their everyday environments outside of work, they are also involved in a rich cognitive coupling with their environments.

In relation to Suchman's notion of plans Lave raises issues of how we understand routines in everyday life. Lave (1988) notes that routines, as we experience them in everyday life, are in some part an illusion, and in fact always to some extent a complex improvisation. Most people have some kind of morning routine that can be more or less specific. But even if you have a very detailed and exact morning routine, the specifics of each day's morning activities are not constant. Lave (1988, p.188) uses the concept of "continuity by fiat" to describe that (cognitive) processes and products are often the same only because they are regarded as the same. For example, making juice from frozen fruit rather than fresh fruit could be regarded by some as routine, while by someone else it might be a violation of the normal activity of making breakfast, which might in turn significantly change the ways things are done cognitively in the morning and during the day. Naturally, Lave's and Suchman's perspectives have methodological consequences. For instance, retrospective descriptions of how an activity unfolds become problematic because they do not necessarily tap into what is going on.

Making use of Lave's notion of people's expectations about situations, it is possible to theorize about when in Clark's ideas of the human prediction machine (as mentioned earlier) the deliberate thought processes step in. According to my interpretation of Lave, it would be either when the routine is not as fixed from the start, or when violations of routines occur. A methodological consequence of this is that it is not just intra-individual variation that should be studied. Intra-activity-contextual variation in relation to the person acting must also be studied.

Both Lave and Suchman bolstered the development of new trajectories in cognitive science by explicitly contrasting their research to traditional cognitive science, and calling for situated theories of cognition[6]. But according to my reading, they did not formulate specific principles for information flow between agents and environments. Kirsh, however, has focused on providing

[6] This is also true for movements of enactivism, extended mind and embedded cognition which, as previously noted, I do not review.

principles for how and why cognition can be distributed within an individual's immediate environment, which I consider in the next section.

3.5.2 Types of resources

In one article Kirsh (2010) describes something that specifies Clark's idea of a canny cognizer. He describes the basics of internal-external coordination in terms of cost-structures. The reasoning goes like this. To solve a specific problem, humans can use internal processing. But they can also couple themselves with external structures to conduct external processing. Internal processing and external processing both come with costs, and a cognitive process will under normal circumstances flow wherever the total cost is lowest. As Clark also notes, for many situations humans will allocate processing to their immediate environment, but here Kirsh is more specific than Clark. Kirsh notes that the coupling process itself also comes with a cost. Kirsh's cost structures can be understood in relation to the foundational principles of distributed cognition where these three types of cost structures can be found in any cognitive system. Kirsh focuses on the intra- and extra-cranial integration processes, but we can also see that the same features exist within the brain. Consider, for example, what in recent years has been coined type 1 and type 2 thinking (Kahneman, 2011) where, for instance, type 1, the fast and automatic thinking, costs less in terms of internal resources, but can have unwanted consequences if used in the wrong cognitive situation. It is also reasonable to imagine that switching between automatic and deliberate processing would come with a cost.

A question that is related to the reasoning behind cost-structures is why individuals should couple themselves with their immediate environments, that is, why is it that the cost can be lower across external structures? Kirsh and others have, in a number of publications, addressed this and formulated a number of principles regarding how people reshape their environments and why the physical (social) environment can be coupled with the brain in a way that is efficient from a cost-perspective.

One mechanism is *jigging* (Kirsh, 1995). A jig is something in the environment that physically or informationally structures the environment in such a way that it reduces degrees of freedom in the process of meeting a goal. An example of informational structuring mentioned by Kirsh is how groceries and goods in supermarkets are structured so that some groceries will be perceived more than others. In this way visitors are provided with cues to act on. To exemplify physical structuring, we can also picture a supermarket where the visitors are constrained to walk a certain path to get to the checkout. In real life these two kinds of jigging often go hand in hand. Kirsh uses the example of the prevention of unhealthy snacking to illustrate this. One way to prevent unhealthy snacking is to constrain the number and kinds of choices present in the fridge. The individual is thus cued to eat healthily, but is also constrained from eating unhealthily. This works because human thinking works according to the very situated and cost-efficient principle "out of sight is also out of mind" (p.38).

This idea of "out of sight out of mind" is also the substance of the previously-mentioned principle *horizon of observation* (Hutchins, 1995a, p.268), which focuses on the fact that information can be physically available or out of perceptual reach. Maintaining a larger horizon of observation was shown by Hutchins to have positive outcomes for team accomplishment. In one example (1995, p.273-274), a communication error between two parties is spotted only because a third party is included in the phone line communication. In the example the third party notes that the second communication member misunderstands an utterance of the first member, whereupon the third party steps into the chain of communication with his own interpretation of the first member's utterance. This time the second member notices his misunderstanding and corrects himself. In a similar way, this principle is applicable to the physical settings of everyday life, where, for instance, putting things in locations where you normally use and perceive them allows them to be within the horizon of observation, thereby increasing the likelihood that an object will act as a cue for an intention.

Intelligent uses of space

As shown above, physical arrangements can, depending on the specifics of the arrangement, tap into both team cognition and individual processes. Kirsh (1995), in the article *Intelligent Use of Space*, focuses specifically on individual cognition when he describes three general and several specific cognitive effects that spatial arrangements can have. They are as follow.

Spatial arrangements can through highlighting (see cueing) and displaying affordances (see constraining) simplify choice by (a) reducing perceived actions, (b) eliminating decision-points, and (c) off-loading heuristic properties. Imagine putting all the vegetables to be cut for dinner on one side of the sink. Since the vegetables to be cut have been grouped in space, the spatial arrangement has reduced the number of perceived vegetables to be cut to those specific vegetables. At the same time, given that the salad has been planned before grocery shopping, there are no decision points in the cooking process for which vegetables should be included in the salad and which in the stew. When the vegetables have been washed, the vegetables that are to be used in the stew are placed on the cutting board on the other side of the sink, while the ones that are to be used for the salad are placed next to the cutting board. By putting categories of vegetables in different locations the arrangement has heuristic properties in the form of temporal orderings of actions (see, for instance, Hydén, 2014 for how the above principles can assist people with a dementia diagnosis).

Spatial arrangements can simplify perception by (a) physically clustering, (b) grouping to categorize according to some relational property between objects, (c) using symbolic marking, and (d) using physical clustering to sharpen perceptual acuity. For instance, imagine solving a large jigsaw puzzle. Common strategies for solving the puzzle include the use of color- and edge-clustering. This is an efficient strategy because human perceptional acuity is otherwise not apt to perceive such a large jigsaw. Going back to the example of the vegetables above, putting the unwashed vegetables on one side of the sink and the washed ones on the other side is one example of using space to simplify perception of categories, because it is a categorization that would have been hard to perceive without the physical arrangement. It is also an arrangement that alleviates internal memory processes since there is no need to remember which vegetables were washed. Moreover, putting the vegetables for the stew on the cutting board simplifies perception since they then share the same relational perceptual property of being on the cutting board, which makes the search process for the vegetables for the stew easier. Symbol marking is characterized by some physical feature that must be interpreted in the form of some kind of understanding of what the symbol means. This could, for instance, be a string on an object that distinguishes it from an otherwise similar object.

Spatial arrangements can simplify internal computation by doing some of the computation in the world. Kirsh's best-known example is how proficient Tetris-players quickly rotate each zoid when it has entered the screen. They do so to offload mental rotation to the physical structure (Kirsh & Maglio, 1992). Kirsh and Maglio show this by measuring rotation time together with perception time and comparing this to what it would take to mentally rotate a zoid. It displays a better cost-structure to physically rotate zoids and perceive their matches, compared to mentally rotating and imagining where they would fit. Importantly, the player may still need to internally compute the best location for the zoid. But a part of the full computational process has been offloaded. Other examples focus on situations where it is shown that humans can think better about novel interpretations when they have a physical representation to work with (Kirsh, 1995). One example of this is how we rearrange letters in Scrabble, where the computation is done both by internal processing and by externally shuffling letters.

A related concept to these forms of intelligent uses of space, especially on the topic of the ability to think in novel ways, is Clark's (2005) notion of *surrogate situations*. Clark describes a surrogate situation as a kind of representational structure that works as a stand-in for a situation that is not physically or temporally present[7]. It could, for instance, be a model, a diagram or a sketch. Clark (2010) discusses two ways that surrogate situations can assist cognitive processes. First, they often highlight key features by suppressing less relevant information. For instance, for certain kinds of decisions, 2D-drawings may be better suited than 3D-drawings (see Clark, 2005 and Kristensson et al., 2009). Second, they relax temporal constraints on reasoning. For instance, when building a house it is easier to figure out the location of the bathroom by using a blueprint before building the house, rather than deciding on the spot at the construction site. Having a surrogate situation allows more cognitive resources to be allocated to taxing tasks, such as, for instance, recalling intentions or producing novel ideas.

Partly overlapping with surrogate situations is Kirsh's (2010a) explanation of what is special about external representations (compared to internal representations) that allows them to empower cognitive processes. (a) They can serve as a more explicit shareable object of thoughts, which individuals can, for instance, refer to when communicating with others. (b) They are more persistent over time. (c) They can more easily be re-represented and changed in ways that make solutions and situations more transparent. (d) Often, an external representation can be a more natural form of representation than its internal counterpart (see naturalness principle above). (e) External representations can support a more explicit encoding of information. (f) They allow for the creation of more complex structures. (g) They help coordinate thought, and (h) they change the cost-structure of the cognitive process (potentially in a positive way for the individual). The reasons described above can, to some extent, be derived from the fact that external representations are not dependent on a brain-like architecture.

Many of his examples are from non-work environments. Kirsh's (1995) paper *Intelligent Use of Space* takes its empirical basis from a mixture of observations of professionals and other people in non-work life situations. The main point is that one essential characteristic of experts is that they shape their physical environment in intelligent ways, and that this is something that needs to be included in the modeling of cognition. The descriptions span a number of non-professional situations in domestic and non-domestic environments. Despite the fact that important theoretical points are made, and that the analysis of intelligently organized spaces in real-life situations is extensive, Kirsh's empirical descriptions are short and do not describe normal people acting in real-life settings. In fact, the actual *use* of space is hard to decipher from the descriptions provided. Since humans tend to have ample experiences of their home environments there are reasons to believe that the phenomena Kirsh reports on will bear out in reality. More extensive descriptions come from Lave, which I have described previously, but also from other studies on domestic environments.

[7] The definition and description of surrogate situations seems to mirror what within the study of the brain is called mental models (c.f. Gentner & Stevens, 1983; Johnson-Laird, 1983). Perhaps surrogate situations can be viewed as the external counterpart to mental models.

Empirical examples

First, empirical examples from everyday life of Kirsh's principles of the intelligent use of space comes from Palen and Aaløkke (2006), who report on a study of older adults and their strategies for managing medication. This study is highly relevant for this thesis because it is one of a few studies that takes the whole home as the unit of analysis to understand a cognitive activity which explicitly involves prospective memory requirements. Ethnographic methods were used, such as shadowing of health care workers and contextual interviews with the elderly participants. They also videotaped two participants preparing their medicine for the next day.

They describe how a number of older individuals proactively shape their physical environment to manage medication adherence. For instance, one participant kept medicine bottles in a linear order in the cabinet to provide a structure for remembering which medicines to take in what order. They also notice how the shapes of a space are congenial only in combination with specific routines. They found that changes in medicine routines reverberated into physical properties of the participants' homes. For instance, mealtime medicines were often located near where the food was eaten, but when the same medicine was meant to be taken at some other time one participant started to use a pillbox. She started to use a pillbox because she needed something that linked spatial and temporal orderings, and her mealtime routine could not do that any longer. In general, participants distributed their medicines across their homes based on the likelihood of being somewhere at a specific time of day.

Palen and Aaløkke (2006) also document what Tribble (see 3.4) calls cognitive thrift. Specifically, some participants only remembered the relevant information about the medicines that was necessary to properly take the medicines at the right point in time, and in combination with the correct circumstances. Other information was seen as superfluous. Palen and Aaløkke (2006) also identify other important aspects of participants' perspectives on their medication management. For instance, keeping medicines in a visible space also means that they openly display their illnesses. Therefore, some opted not to keep their medicines in an open space, even though visibility would decrease their cognitive burden. Instead, some participants preferred having a strong routine, for instance, opening the cabinet as part of their morning routine.

Two other studies, Crabtree and Rodden (2004) and Crabtree (2003), also take the home as their unit of analysis when they describe practices of intelligently shaping spaces. Specifically, they describe information flow in family homes and family routines for managing information and communication. While they do not directly refer to distributed cognition, they apply concepts that are similar to those employed by studies using distributed cognition. For instance, they consider a case of mail handling and how mail enters the home, is collected by a family member, and sorted and placed at some location by the same family member based on aspects such as relevance to other family members. By having various routines for mail handling, information flow between family members and actions among the family members are coordinated. The authors plot the general course for mail within the home ecology and define a number of manifestations of how communication is realized in space. They define three types of communication spaces (paraphrased from Crabtree & Rodden, 2004, p.205):

Ecological habitats are places where communication media live and where residents go in order to locate particular resources. These spaces (for instance the kitchen table) are fixed spaces in plain view, which are also described as spaces that family members do not need to search for. *Activity centers* are places where media are actively produced and consumed and where information is transformed. These are centers that differ to some extent from the previous type. Mail is, for instance, handled on the porch though it is not a communicative habitat for family members. *Coordinate Displays* are places where media are displayed and made available to residents to coordinate their activities. For instance, mail might be placed at the very corner of the kitchen table so that it is not missed by other family members. Again, these displays overlap with previously mentioned spaces and I interpret coordinate displays as a type of ecological habitat that highlights more urgent information.

What is especially compelling with Crabtree and Rodden's (2004) approach is that they view routines and practices in domestic environments as complex and subtle phenomena, which is not very different from what occurs in work-settings.

There are also studies in professional settings that have directly considered the use of space and prospective memory. Grundgeiger, Sanderson, MacDougall, and Venkatesh (2009 see also Grundgeiger, Sanderson, MacDougall, & Venkatesh, 2010) use distributed cognition in their research on nurses and prospective memory. They present an analysis of video and interviews from nurses' practices in intensive care units. They describe three types of support nurses receive from their environment to perform intentions.

Passive representations are described as equipment that the nurses directly associate with a task. As an example, the authors mention how some nurses are reminded of a pre-transfusion procedure by looking at a label on the blood bag.

Active representations are described as features of the environment that have been designed, either by the nurses themselves or by a designer, to serve as a prospective memory reminder. Examples mentioned included post-it notes created by the nurses and visual cues for performing certain sequences at relevant places in the environment created by the organization.

Proactive representations are described as dynamic and adaptive resources that ensure that intentions are conducted. As an example of a human proactive resource they describe the use of an extra monitoring nurse in certain situations, even when only one nurse is necessary for performing in the situation. As an example of a non-human proactive resource they describe how, when setting the threshold for alarms on physiological monitoring equipment, the nurses choose values that are close to normal levels to make sure that the monitors are checked frequently.

Perception, projection, and imagination

In several of the above examples there are instances of what Kirsh calls *projection*. Projection is the ability to project internal representations onto physical structures (Kirsh, 2009). To explain projection, Kirsh describes a continuum for couplings with physical structures that runs from perception to projection to imagination (Kirsh, 2010b). A demonstrative example is three representational ways of playing tic-tac-toe. Imagination is necessary when you have no representational structure in front of you, where, for example, you need to internally remember your previously spoken moves. If, on the other hand, you have the regular tic-tac-toe grid with all previous moves in front you, then you can rely more on perception to help you remember. Projection is necessary when you have only a partial representational structure in front you. In the case of tic-tac-toe, this might be just the grid, but without the previous moves marked on it. The point is, having the grid provides a representational structure that internal representational entities can be projected onto. The grid helps to anchor thought processes.

Another study that demonstrates projection is one on professional dancers, where Kirsh (2010b) describes how the dancers are able to anchor their thought process by using bodily sequences that are simplified versions of the intended, fully-realized sequences; this is called *marking*. By doing this the dancers can devote more energy to other processes, for instance the process of learning and remembering the choreography, before dancing the full routine. Just as when playing tic-tac-toe with a blank grid, the dancers are able to *project* their intracranial thought process onto an extra-cranial architecture, which in this case is the dancer's own body.

Hutchins describes something equivalent to projection when he talks about *projection of internal conceptual structures onto material anchors* (Hutchins, 2005, see also Fauconnier, 1997). An example from everyday life is the technique of remembering the number of days in the months throughout the year by projecting the names of the months onto the hand, using the knuckles and the spaces between. Other examples include how Micronesian navigators navigate by using properties of the stars to anchor mental images (Hutchins, 1983) and how we perceive groups of people standing in certain arrangements as lines and queues (Hutchins, 2013).

Despite the overlap with Kirsh, the analysis from a distributed cognitive perspective goes on to make an additional distinction about how an individual anchors intracranial thought-processes, between anchoring to extracranial structures and anchoring to *properties of the cognitive ecosystem* in which these anchoring processes exist. For instance, seeing a group of people first as *a line* and later as *a queue* is in part an individual accomplishment. Humans have the ability to mentally project a trajectory onto a spatial array. This is at the heart of Kirsh's description of projection. But why humans interpret these spatial arrays as queues is determined by a number of other 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." (Hutchins, 2013, p.6, see also the previously described notion of professional vision proposed by Goodwin, 1994). From Hutchins' description it becomes clear that to understand cognitive processes we need to view cognition on various levels. To describe what is going on in the head we need to account for properties of the larger cognitive ecosystem.

The first message here is again that a reciprocal process is taking place between representational structures in the world and representational structures in the head. Perhaps a more important message is that there are kinds of (representational) structures in the world that, when they are coupled with individuals, create various forms of cognitive processes. Because of this, Kirsh (2010a) criticizes the notion of external memory as something that is about offloading internal memory. The problem is that the notion of offloading memory only tells half of the story of what is going on. It is for instance true that the previous note taking relieves internal memory to some extent, but the note taking also transforms the cognitive process into something else that follow new causes and consequences (see also the previous discussion of the intelligent uses of space).

The reasoning that underlies the concept of projection suggests that there must be some semiotic connection between the physical structure and the solution of a problem. One model that focuses on types of couplings is Charles Sanders Peirce (1932) theoretical model for semiotics, which differentiates between symbols, indexes, and icons. Of the three categories, icons most closely resemble the things they represent. An icon could, for instance, be a picture of something. Indexes do not directly represent something, but they do directly point to some meaning. A dark cloud, for instance, increases the perceived likelihood of upcoming rain. Symbols are things that have become separated from the context to which they refer. Words are examples of symbols that can be used to refer to something which is not directly present. These different types of connections that humans have with their environment are valuable for understanding the situated agent as it makes sense of its surroundings. Similarly to Kirsh, the different types of semiotic connections show that physical structures can play a significant part in the cognitive process whether there is a clear representational structure or not. There can, therefore, still be couplings between individuals and physical structures without symbols that move between (see also Norman's, 2011, description of *social signifiers*).

Although humans have the ability to interpret symbols, indexes, and icons, this does not mean that all types of structures or all versions of each category are equally preferred in all situations. Rather, there are additional principles that suggest which structures are preferable in a given situation. Norman's *perceptual principle* (Norman, 1993) suggests that "perceptual and spatial representations are more natural and therefore to be preferred over nonperceptual, nonspatial representations". Importantly, though, for this to work efficiently there must be a natural mapping between the representation and what it stands for (Norman, 1993, p.72). In accordance with this the *naturalness principle* (Norman, 1993) suggests that "experiential cognition is aided when the properties of the representation match the properties of the thing being represented" (p.72). For example, it might be easier to remember how and when to do something if it is represented as a pictorial schema of a sequence of actions, rather than as symbolic words describing what to do.

Norman's principles suggest that there is a certain potential cognitive empowerment associated with coupling information processes to external structures, but only if it is done right. To exemplify what it means to do it right, consider again the reasoning of cost structures and the costs of coupling when projecting thoughts onto the grid in tic-tac-toe. There are three major factors that influence the cost. First, the complexity of the task: projecting x's and o's onto a 3x3 grid comes with a lower cost than projecting x's and o's onto a 9x9 grid. Second, the expertise/experience of the agent: the more the expert someone is, the more complex computation she can do in her head. This is because if you are an expert you are also able to perceive more possibilities with less physical anchoring and less physical explicitness (see, for instance, Chabris & Hearst, 2003 for an example of chess). Third, the relatedness between structure and problem: in the game of tic-tac-toe the grid matches the structures of the game to a relatively high degree. If, for instance, we were to take away all horizontal lines the match between structure and problem would be lower.

I find projection to be a useful concept because it elucidates much of what it means to be a situated brain-perceptual system. To understand this better note that the examples above are clear problem situations. The dancers need to practice their choreography and tic-tac-toe players aim to win the game. But at the heart of the concept is something more basic than clear problem situations. Kirsh (2009) believes that a project-create-project cycle is at the heart of much sense-making, and that this becomes especially clear when humans are involved in problem-solving. But I also think that the role of projection becomes clear when the problem situation involves physical structures that have a connection to the problem at hand. It is convenient to project dancing onto a body, and it is convenient to project tic-tac-toe onto a tic-tac-toe grid, But there are likely less clear cases out in the real world.

Kirsh's project-create-project cycle can be further compared with Neisser's (1976) perceptual cycle. The perceptual cycle is a theoretical model that emphasizes the ongoing interpretation of input from the surroundings. In this model the agent's perception of the world modifies its mental model of the world, which in turn directs actions, which in turn leads to new information sampling of the world, etcetera. It is a fundamental cognitive mechanism that, in most cases, humans are sensitive to their present circumstances so that they can act on their present circumstances. The projection-cycle, however, suggests something different: that it is the internal mechanics that determine what we sample when perceiving a physical structure, which in turn affects how we understand and act. These two cycles can be seen as complementary. First, humans need to perceive some external structure before projecting thoughts onto it. Second, in the perceptual cycle projection can be viewed as a type of action, a mental one, that is anchored to a physical structure. As the Kirsh's reasoning suggests, in a problem-solving situation we are likely to use a modified version of the perceptual cycle, in which we project thoughts onto the physical structures from which we sample information. When solving problems, there is a directedness of our mental process to make sense of present and future circumstances, because concrete problems often demand that we are directed in such a way.

This combination of cycles is similar to how I previously described the way that CSE views functions in cognitive systems. However, within CSE it is not just individuals that account for the cycles, it can also be a collection of individuals and tools that control specific functionality. In fact, COCOM is influenced by classical cybernetics feedback loops and Neisser's perception-cycle. In COCOM it can be the collection of parts of the system that perceive the world and modify the mental model of the world, which directs actions, and which in turn leads to new information sampling of the world. The same collection of parts is what predicts and projects understanding on top of what the world provides.

3.5.3 Embodiedness and the non-representational

The body is a special kind of cognitive vehicle that is with an organism from birth, which in recent decades has been proven to be intrinsically linked to mental life. Hollan, Hutchins, and Kirsh (2000) include embodiment (together with cognition as a cultural and social process) as a necessary tenet of distributed cognition, because it is what allows individuals to be coupled with their environments.

We have already seen an example of the role of the body in, for instance, how dancers use their bodies, not just for dancing, but also as a vehicle of thought. In this thesis, I will use the notion of bodily cognition to refer to the idea of the body as a necessary means for interaction and an opportunity for interaction with environments. Note, however, that the role of the body in cognition is not a unitary theoretical idea. Embodied cognition is a larger movement within cognitive science with goals that, to some extent, go beyond my use of the concept of body in this thesis. Therefore, I will pave the way with a brief discussion of what is meant by embodied cognition and the versions of this concept that I will use in my analysis.

Shapiro (2011) suggests three kinds of theoretical positions within embodied cognition: (a) conceptualization, (b) replacement, and (c) constitution.

Conceptualization theories can be boiled down to what is known as *the symbol grounding problem*. The symbol grounding problem is a formulation of the problem of how abstract symbols that are used in mental computation get their content and meaning (Harnad, 1990). Conceptualization theories within embodied cognition address this problem by holding that symbols get their content from bodily interactions with the world, which are, among other things, constrained by features of our bodies (see Miller & Johnson-Laird, 1976 for a model of conceptualization). The conclusion is that human thinking is more or less grounded in bodily experiences (see e.g. Barsalou, 1999; Lakoff & Johnson, 1999). An empirical example of conceptualization is Beilock and Goldin-Meadow (2010), who showed that gesturing adds content to mental representations in such a way that subsequent performance can be affected in tasks that use these mental representations. Anderson (2003) argues that solving the symbol grounding problem is what truly delineates embodied cognition theories from situated cognition theories in general. Conceptualization has not been at the heart of my empirical research. Rather, my focus has primarily been on a combination of replacement and constitution aspects. However, understanding how individuals interpret their surroundings is important for any cognitive analysis, symbol grounding or not.

Replacement theories suggest that evolution has equipped humans with energy efficient processes that minimize the need for complex abstract computation. Instead of abstract computation, simple heuristics are used, which involve bodily interactions with the physical environment. A prime example is how humans catch a ball (and how many predators catch prey). Instead of computing the trajectory of the ball by including distance, velocity, angle, air resistance, etcetera. the human fixes her gaze on the ball and adjusts direction and running speed so that the angle of her gaze remains constant (McBeath, Shaffer, & Kaiser, 1995). Replacement theories have mainly focused on explaining situations in which the body is required for solving the problem, such as catching a ball (Shapiro, 2011), but in theory bodily heuristics could also replace non-body cognitive processing.

Constitution theories suggest that the body is part of cognitive processes. This is a less theoretically extreme form of embodied cognition, since it does not refute the notion of complex intracranial computations but instead suggests that the body is a resource that plays the same part in cognitive processes as other external resources. A prime example of this is how humans count on their fingers. Another example is Kirsh's (2010, see also section 3.5.2 above) study on professional dancers, where it is shown that the body can work as a representational vehicle of thought.

As a concept within cognitive science, embodied cognition is the latecomer of the perspectives presented in this chapter, but it originates from earlier philosophical ideas at the first half of the 20th century (see e.g. Heidegger, 1953; Merleau-Ponty, 2004) and Gibson's ecological approach in psychology, which started to take shape in the 1950s (Gibson, 1986). Together, these earlier theories (similar to Suchman, 1985) argued for the notion of actions as key to the understanding of human thinking and perception. Proponents of embodied cognition have also been inspired by connectionism, which started in the 80s and attempted to model cognition as parallel processing between simple units in a network (Churchland & Sejnowski, 1992), a process that to some extent was supposed to resemble the mechanisms of biological brains. Proponents of both connectionism and embodied cognition argue that we do not necessarily need symbols and representations for all cognitive processing. On the surface, this stands in contrast to Hutchins' original idea of distributed cognition, which maintains that cognition is primarily a symbol manipulation process.

However, based on my review of more individualistic perspectives I find it important, from a distributed cognitive perspective, to emphasize the non-symbolic features of cognition. I do not consider this to be at odds with Hutchins' formulation of distributed cognition. Recall that distributed cognition (Hutchins, 1995a) was born through a re-interpretation of the PSS-hypothesis as a sociocultural system. It allowed cognition to flow across more architectures than brains and computers. To some extent this also led to a new cognitive description of architectures. For instance, through this new perspective, social orderings must be seen as architectures for cognition, and as a part of cognition just as the architecture of the brain is. But empirically, research under the heading of distributed cognition has mainly focused on how explicit information *moves* across architectures and changes representational states. This focus on explicit information does not necessarily account for all types of cognitive processes.

Sutton (2006a) also writes that mainstream distributed cognition research treats external representations (or exograms, as coined by Donald, 1991) as "passive, stable, and medium independent" components, even though models of biological brains have, in recent decades, started to see biological representations as active and reconstructive. But as noted in several examples above, there is not always a movement of stable information across external media, nor are exograms stable in Hutchins' original descriptions. Reconstruction and re-interpretation of external information across cognitive tools during the navigational task is shown to be a necessity in order for the team to accomplish the navigational task.

Therefore, the move away from stable representational structures is not antithetical to what was reviewed under types of agent-environment connections, nor is it so to how Hutchins actually uses the symbol paradigm to understand his empirical material. In fact, he explicitly relaxes the concept of computation by specifying that "the actual implementation of many interesting computations is achieved by other than symbolic means. For our purposes, 'computation' will be taken, in a broad sense, to refer to the propagation of representational state across representational media. This definition encompasses what we think of as prototypical computations (such as arithmetic operations), as well as a range of other phenomena which I contend are fundamentally computational but which are not covered by a narrow view of computation" (Hutchins, 1995a, p.118).

Notably, what remains of the traditional definition of cognition after this is representational states across representational media. However, Hutchins also relaxes his definitions for representational media, and in this way actually ends up closer to embodied cognition than one might expect. For instance, in his listing of observable representational media for mapping the relationship of the ship to its environment and the position plotted on the chart, Hutchins (1995a, p.119) lists "the world" with its own heading. Hutchins specifies that the relationship between the ship and every object in the surrounding world is specifiable as a direction and a distance, and hence (my inference) is a representational medium. The way that Hutchins operationalizes representational media is close to how Brooks (1991) argues that for embodied and situated agents, the world is often its own best model for the agent to act upon which does not necessarily need to re-represented in the mind.

Sutton (2006a) describes a broad notion of distributed cognition that is in line with Hutchins' description of distributed cognition as a perspective on all cognition. First, as already noted in this description, cognitive systems can exist on different temporal and spatial scales where, for instance, it is explicit that distributed cognitive systems can exist momentarily before later ceasing to exist. Also, in line with connections described by Kirsh, representational structures in the world (or in the brain) are not always stable entities. Instead they are dynamic, and their functionality is reconstructive in a specific setting with a specific interpreter. A consequence of this perspective is that the symbol-grounding problem is as much a problem for external representations as for internal representations. Why humans come to interpret some representation or structure in the material world as meaningful must also be grounded in something.

Cognition across the body is an important analytical focus for understanding cognitive accomplishment. Constitution theories identify the body as a cognitive resource in general, and therefore constitution theories largely overlap with the theoretical principles mentioned in previous sections. But in combination with replacement theories, the body becomes a vehicle of thought that significantly shifts, or evens out, the center for the previous brain-centered cognitive system. This shift is in line with the foundational principles of distributed cognition. In practice the shift leads to a decrease in the cognitive burdens that high-level cognitive processes in the brain would have without a body.

3.5.4 Epistemic actions and experts

The act of using the body for the creation of intelligent shapes in space can be understood via the notion of *epistemic actions*, which was first discussed by Kirsh and Maglio (1992, 1994) in their studies on expert Tetris players.

Epistemic actions are defined as actions that generate knowledge about solving a specific problem. This stands in contrast to *pragmatic actions*, which are defined as actions that get the individual closer to solving the same problem. Therefore, epistemic actions are often actions that make the world more congenial for solving a problem. In the case of the Tetris players it was an epistemic action to rotate the zoids in the world instead of in the head. In real life these two types of actions are not always easily discerned. A given action can be viewed both as an epistemic and as a pragmatic action at the same time. For example, putting an object at the door to remember to bring when leaving home is both an epistemic and a pragmatic action. It is epistemic because it alleviates internal remembering by replacing it with the simple perceptual task of noticing the object when crossing the boundary of the home. But it could also be argued that it is a pragmatic action because it gets the object in question closer to leaving home.

To understand the notions of epistemic and pragmatic more specifically it is possible to view them in relation to intentions and effects. The expert Tetris players rotated the zoids with the goal of alleviating internal computation, not primarily for the purpose of bringing the zoids closer to the goal. Hence, the action is an epistemic action. This was likely the original intention of Kirsh and Maglios' definition of the concepts (Nils Dahlbäck, personal correspondence). We can use the same reasoning to understand the action of putting an object to bring along in front of the door. It is an epistemic action if the agent puts the object there to alleviate internal computation in some way, for instance to alleviate remembering or perceptual processes. However, I think we can also understand any action from its epistemic consequences. If an individual has habits, without epistemic intentions, to put objects in certain spots that happen to have epistemic effects, those habits are of equal importance in a specific situation for cognitive accomplishment to the intentional epistemic actions. Although they are equally important for accomplishment, it is still important to understand when an action has an epistemic intent. For instance, if a practice exists with an epistemic intent it is likely that the agent has a certain motivation to maintain the positive outcomes of the action across time and circumstances. An epistemic intent is also an indication of expertise with regard to the cognitive environment. On the contrary, if a practice exists without some deliberate epistemic intent, the practice with the positive outcomes also risks giving way to new circumstances.

We have seen that the notion of in situ problem-solving focuses on the mechanisms for reciprocal causation between individuals and their environments to solve cognitive problems. This could consist of opportunistically harnessing resources at hand on a moment-to-moment basis, as emphasized by Clark. But it could also be, as emphasized by Kirsh, about pre-structuring one's environments so that coupling can be established later. Sometimes this pre-structuring is performed by the one who will use the structures, but it could also be performed by someone else, as in the case of the design of grocery stores. Overall, in situ problem-solving informs us that the establishment of couplings between individuals and their environments is sensitive to specific characteristics of the external (representational) structures.

In the characterization of distributed cognition I noted that research most often studies professional settings in which experts perform complex tasks. In this section I observe that certain elements of situated cognition also focus on professionalism. Kirsh states that "cognitive processes flow to wherever it is cheaper to perform" (2010a, p.442). Clark pictures human cognition as a "canny cognizer [that] tends to recruit, on the spot, whatever mix of problem-solving resources will yield an acceptable result with a minimum of effort." (2008, p.13). To make the point that human thinking is tightly coupled with its immediate environment, they describe something that can be interpreted as an ideal agent that cleverly calculates costs. In contrast, using Hutchins' version of distributed cognition, one could say that it is mainly a cultural process that calculates costs, and not primarily agents. Therefore, on an individual level in a given situation it is not necessarily the case that an efficient calculation is made by an individual alone.

To some extent I generalize Kirsh's and Clark's descriptions, but I think it is fair to say that situated cognition has focused to a large extent on showcasing the immediate environment as important, and theoretically part of cognitive processing. However, there is little empirical description of the hassle people have using and interpreting their immediate environments, something that was originally found in the works by Lave and Suchman. In their old theoretical versions of situated cognition nothing goes as planned, if there even is a plan. If there is a plan, or a routine, they are mere shadows of the cognitive mechanisms that explain cognitive accomplishment. Although individuals make use of their immediate environment for cognitive processing there is a constant hassle of sense-making and reframing of situations. Despite this difference, as clarified in the text above, at a glance these disparate strands of situated cognition are consequences of the goals of their respective publications, that is, to contrast against traditional cognitive science, and not necessarily theoretical oppositions. Of course, individuals need to make sense of the physical and social settings they confront.

3.5.5 Cognitive tools

When shapes of the environment such as, for example, a set of representations with specific cognitive functions are instantiated into "a thing", they become something that is often called a cognitive tool or artifact. Norman (1991) defines a cognitive tool (specifically, an artifact) as "those artificial devices that maintain, display, or operate upon information in order to serve a representational function and that affect human cognitive performance" (p.17). Professional settings, where distributed cognition has mostly been employed, are generally full of cognitive tools, and therefore cognitive tools have been given ample theoretical space. Specifically, what cognitive tools are and what they do has been defined in many ways throughout recent decades (see Garbis, 2002; Susi, 2006). I will not go into any theoretical discussion of such issues, but I note that cognitive tools rely on principles of human cognition as described in the previous section and also that they have significant roles in the evolution of human cognition. As was described in section 3.3, cognitive tools travel efficiently across generations and communicate knowledge and practices across and within generations, which determines cognitive processes.

In professional settings cognitive tools are mediators between agents and processes and are necessary tenets for accomplishing intended activities. Everyday life also involves cognitive tools and has, through the evolution of information technology, become more common than in the past. For prospective memory research, as was noted in the previous chapter, calendars and reminder notes are commonly used by most people in their everyday lives. From a distributed perspective these tools come with inherent cognitive properties. Calendars, for instance, come with representations of weeks or months, which delineates chunks of time. From a distributed perspective the properties and content of a calendar, in accordance with how they work in professional settings, are mediators between agents which use the tool in some future, present, or past state of affairs. Tools can therefore be viewed both from the perspective of their inherent properties and also through the ways they are utilized.

An interesting study that looked at cognitive tools in everyday life that accounted for these two perspectives is de Léon (2003). He reports on cooking activities in the home environments of a number of participants. He brings attention to several manifestations of distributed cognition in relation to cooking. For instance, he describes how participants prepare and maintain their workspaces throughout the activity.

He also makes an extensive case regarding a spice shelf as a cognitive tool. He does so by mapping the historical development of one participant's (Robert) spice shelf. The historical development is shown to follow Robert's increasing interest in cooking, but it has also followed incidental factors, such as the shapes of the spice containers. In certain periods the shelf was more of a box where spices could only be viewed from the top without labels. At the time of the study the spices are categorized into cooking genres (for instance, "western"). One category with common spices he also shares with his wife, and therefore these spices are arranged in alphabetical order. Another category is called "on the way out", which are spices that he will not buy again.

The analysis of the different versions of Robert's spice shelf introduces interesting cognitive aspects of different designs. For instance, a deep shelf that can only display a fraction of spices up front taxes the user's cognition differently than one that only has one large, shallow display where all spices are always visible. From the analysis it is apparent that Robert remembers and makes decisions through probing, maintaining, and shaping the shelf. In the past when the shelf was a box Robert needed to act to gain access to information about the spices. Probing, maintaining and shaping was not as cognitively easy because back then Robert could not observe the spices at a glance. I see de Léon's contribution as important because he gives examples of Kirsh's principles of the intelligent shape of space and tool use through extensive empirical descriptions of how these practices can work in real life.

3.6 Distributed cognition as a framework

Several studies have used distributed cognition to understand aspects of prospective memory in professional settings. There are, based on the review above, reasons to believe that the theoretical principles of distributed cognition can also contribute to the understanding of everyday environments, and specifically the home environment as a cognitive system that shapes prospective memory processes. In fact, Hutchins (1995b), before turning to the technical environment of a cockpit, states that "many of the outcomes that concern us on a daily basis are produced by cognitive systems of this sort" (p.266). The general reason for this is that individuals in non-work environments still participate in everyday life. Kirsh (2008) further characterizes people in general as experts, or near experts, of their everyday environments. This is not a far-fetched idea because the word *everyday* signifies experience, and amount of experience in any environment is a factor for expertise (see for instance Ericsson & Charness, 1994). Kirsh's rationale is also that people's expertise has effect on their physical surroundings because people congenially structure their environments to fit their routines and internal cognitive processes in ways I described in previous sections. There are, as seen in the review above, indeed a few studies which look on either prospective memory or everyday life environments, or both, from a distributed perspective. But they are few, and they do not explicitly consider the home as a cognitive system as whole, established to deal with more than one task.

Below I summarize distributed cognition as an analytical framework. This framework is my view of distributed cognition; from which I have picked analytical tools for my analysis. The list is a summary of the review I have made above and refers to aspects of distributed cognition that can be studied.

There are other frameworks in the literature that aim to do something similar. For instance, in the healthcare domain Blandford and colleagues (Blandford & Furniss, 2006) have developed an analytical framework called *Distributed Cognition for Teamwork* (abbreviated DiCot and referred to in section 3.2). DiCot is a structured approach of listing principles of distributed cognition that are relevant to understanding teamwork in small teams (see Berndt et al., 2015; Furniss et al., 2011; Rajkomar, Blandford, & Mayer, 2013; Rybing, Nilsson, Jonson, & Bång, 2015 for examples of where DiCot has been applied).

The name DiCot embodies distributed cognition as research that has mainly been focused on the study of professional settings. It also characterizes distributed cognition as an analytical framework for teamwork activities, and not primarily as a perspective on all cognition. However, first, DiCot performs the important role of substantiating distributed cognition as a theoretical contribution to the study of cognition, and second, few of the principles listed in DiCot are not principles of teamwork, or even groups of people doing things together. In short, they are information-theoretic principles of information flow with an extended focus on how spaces are arranged.

I do not aim to make a formal comparison with, or add to, DiCot. But I note that I have included a few principles that focus on agents' mechanisms for coupling themselves with their environments which are not included in DiCot. Since distributed cognition is a perspective on all cognition, it can be argued that several analytical frameworks can be derived from this perspective, frameworks that are designed for the study of types of cognitive systems.

In the review above I presented a few studies using distributed cognition for the understanding of everyday environments and occasionally also prospective memory. Something that became clear from the study by Palen and Aaløkke (2006) was that participants' processes for managing medication were distributed across their environments, but also that each individual had their own solutions and that individuals could act in different ways on different days.

Everyday life does not necessarily have the same level of structured systems as the professional environments previously studied using distributed cognition. Everyday life can be expected to be less safety-critical, less based on training, less based on rules of conduct, and less based on commonly established procedures. However, this does not mean that everyday life cannot include features of professionalism for a number of reasons, for example, because these features are relevant, are experienced as relevant by an individual, or are regarded as relevant by a community for the management of some part of everyday life. In my analysis I will describe practices of the participants that can indeed be regarded as involving professionalism.

To account for variations across situations and times, I included elements in the framework for capturing the dynamics and stability of practices and resources. As previously stated, the framework below is a summary of my view of aspects that are usually studied with distributed cognition.

Table 2: Distributed cognition as analytical framework

Aspects of agents' practices
a. internal processing, knowledge and abilities (see mainly cognitive psychology but also section 3.5)
b. projection of internal information onto external structures (see, for instance, section 3.5.2)
c. interpretation of external structures and situations (see, for instance, section 3.4 and 3.5.2)
d. the use of the body to move and to find information (see, for instance, section 3.5.1)
e. shaping space and epistemic actions and consequences (see, for instance, section 3.5.2)
f. stability and dynamics of the above practices (see, for instance, section 3.4)
g. the goals for a given activity (see, for instance, section 3.4, 3.5.1 and 3.7.2)
Aspects of the physical and social environment
a. the nature of information flow and transformation (see, for instance, section 3.2)
b. arrangement of resources (see, for instance, section 3.5.2)
c. physical or symbolic constraints of actions (see, for instance, section 3.5.2)
d. stability and dynamics of resources (see, for instance, section 3.4)
Aspects of external resources
a. passive resources/incidental cues (see, for instance, section 3.5.1)
b. intentional (by someone) resources (see, for instance, section 3.5.1 and 3.5.2)
c. cognitive tools and their functionality and use (see, for instance, 3.5.3)
d. mapping between representation and what it stands for (see several sections)
e. complexity of resources and environment (see, for instance, section 3.4 and 3.5.2)
Aspects of time
a. the development of practices (see, for instance, section 3.3, 3.4 and 3.5.3)
b. the development of a single task (a task analysis) (see, for instance, section 3.2 and 3.3)
c. the development of the practitioners (see for instance section 3.3 and 3.5.2)

Based on the framework above I see five possible general media for cognitive processes: brain, body, material world, social world, and time. The body is a special kind of material architecture that was originally viewed as playing some role in distributed cognition (see Hollan et al., 2000). However, it was not explicitly acknowledged as somewhere cognition could flow. Also, the understanding of the brain is not at odds with distributed cognition. However, most research involving distributed cognition does not include an analysis of individuals' cognitive processes. This is possibly because people are educated in the roles and the researcher can thus assume certain practices and basic knowledge. But, as perhaps best represented by Clark (2013 and other publications), the brain is indeed a significant component in extra-cranial distributed processes.

I find it reasonable to recognize that connectivity (type and density) between the brain and the world, as with any cognitive system, can vary across mental states and cultural settings. Therefore, it can also vary to such an extent that the proper unit of analysis for a given research question is the brain alone, all of which is in accordance with the foundational principles of distributed cognition as a perspective on all cognition.

Because distributed cognition has a dynamic and wide notion of where cognitive processing can take place, Hutchins and his colleagues introduce a new methodological approach to address the wideness and dynamism of the perspective. I now turn to describing this approach, which is referred to as cognitive ethnography.

3.7 Cognitive ethnography

Cognitive ethnography is described by Hutchins (1995a) as the most important methodological tool for understanding cognition in real-life situations, and is also the method I have used in my fieldwork. I believe that Hutchins (1995a) regards cognitive ethnography as the most important methodological tool because he has argued that cognition must be re-understood through studying it in real-life situations. However, this does not mean that other methods are unnecessary. For instance, in Hollan et al. (2000) it is clarified that the relationship between ethnography and experiments within distributed cognition is such that results from ethnographies inform the creation of experiments. And in fact, in this thesis I also let results from experiments inform my ethnographic analysis, hence I use a two-way relationship between ethnography and experiments.

Below I describe my view of what a cognitive ethnography is in two steps. First I describe what ethnography is, and then what ethnography on cognition is. Even though cognitive ethnography is frequently used as a concept within distributed cognition, it is seldom specified what the term means. Even more seldom are attempts to describe differences and similarities between cognitive ethnography and contemporary ethnography within anthropology.

3.7.1 Ethnography

Ethnography has its origins within anthropology and can be imprecisely translated from Greek as a description of a culture or a group of people. More specifically, for instance, Sluka and Robben (2012) in a review of the field characterize ethnography as a micro-analytical case-study approach that aims to study people, groups or communities. Irrespectively of which group is studied the overarching object of study in ethnography is always aspects of culture. There are many definitions of culture within anthropology. One comes from Geertz (1973) which, by adopting the philosopher Max Webers description of humans, defines culture as the webs of significance humans collectively spin.

The notion of studying cultures and the establishment of ethnography as a scientific method was, and still is, highly influenced by Malinowski and his publication *Argonauts of the Western Pacific* (Malinowski, 1922). Malinowski declares that the primary goal of ethnography is to document a culture from the perspectives of its inhabitants (also called the emic or inside perspective). The inside perspective is in contrast to the outside perspective (also called the etic perspective). Ethnographic fieldwork is often metaphorically described as a balancing process between the inside and the outside perspectives (Agar, 2008/1996; Emerson, Fretz, & Shaw, 2011). In Malinowski's studies, the outside perspective is the perspective of western societies, in other words Malinowski's perspective.

For the researcher, the goal is not to reach a complete inside perspective, because becoming an insider could risk losing the ability to communicate the inside to the outside community. Nor is the goal to keep entirely to the outside perspective, because that would risk mistaking your own lifeworld for that of the participants. Instead it is an act that Agar, 2008/1996) depicts as one that should be like a *professional stranger*. Someone that systematically translates the inside perspective to outsiders, and hence aims to change the outsiders' view of the inside. This metaphor is at the heart of ethnographic fieldwork. Another way of understanding ethnographic fieldwork and being a professional stranger is to conceptualize it as a process that questions what is familiar to the researcher, and explains what is unfamiliar to the researcher. Although ethnography stresses this, it is not very different from the process of research in general.

To get the inside perspective of a culture Malinowski argues that ethnographers should use participant observations as their primary method. In practice this means that the researcher takes part, as much as possible, in the same activities the people being studied take part in. For Malinowski, and for many anthropologists, this also means that the ethnographers need to learn the language of the culture being studied. The rationale is that only by participating and interacting in the culture's own language can ethnographers properly document how webs of significances are spun (to adhere to Geertz, 1973 terminology), how inhabitants order and understand situations in their lives. Despite Malinowski's method for conducting an ethnography there are other ways of doing ethnography that emphasize other aspects of cultures. For instance, Franz Boas focused on collecting and bringing home stories and artifacts from the societies he studied, spent shorter periods in the field than Malinowski (Sluka & Robben, 2012, pp. 11-12), and also emphasized the interview as an important method (Agar, 2008/1996). The focus was to understand a culture by studying its products. Similar foci also exist within cognitive science. As mentioned, de Léon (2003) for instance scrutinized tools in the pursuit of unraveling underlying cognitive processes.

Another aspect of Malinowski that became a goal within anthropology was the documentation of all the parts of life in whole cultures. This was an ambitious project that, despite being an ideal, was difficult to attain. Ethnographies with similar objectives to Malinowski's are still conducted today, and in contemporary ethnography several aspects of Malinowski's perspective remain as ideals. Participant observations and the attempt to understand the inside perspective remains prominent. Ladner (2014), for instance, describes the goal of understanding the emic perspective as where fieldwork starts. However, there have been a number of developmental tracks within anthropology that clearly complicate other aspects of traditional ethnography.

One track deals with the understanding of culture. When Malinowski ventured to the islands in the Western Pacific he viewed the society as isolated from his and from other societies. Because of this there was something that could be described as an outer boundary of the whole culture. In contrast, for contemporary ethnography, boundaries are not as distinct and the concept of culture is more complex. For instance, ethnography is currently used to study some aspects of culture without attempting to frame a whole culture. The question of which aspects of the culture are interesting is guided by research questions and the purpose of going into the field (Agar, 2008/1996). One practical result of more malleable boundaries and question-oriented ethnographies is that ethnographies today are not necessarily as long as traditional ethnographies were (Sluka & Robben, 2012).

Another related developmental track is that contemporary ethnography no longer only studies cultures that are far away. Instead ethnographers can stay closer to their own homes, and also close to, and part of, their own webs of significances (Cole, 1977; Geertz, 1973). One practical result of this is that the need to learn the tongue of the participants is now often about understanding the vocabulary and the expressions of the respondents. Language is still a route to the meanings and orderings of the participants, but the ethnographer does not need to learn a completely new language.

Another result of studying nearby cultures, or sub-cultures of one's own culture, is that ethnographers nowadays study the mundane and the familiar to a greater extent. To some extent, traditional ethnography benefited from studying far-away cultures, because it focused on studying something (almost) completely new. Everything needed a translation to the outsiders, otherwise the culture would not be understood[8]. In contemporary ethnography, the researcher risks taking the familiar for granted or risks remaining blind to the real inside perspective. Therefore, the role of the ethnographer as translator has partly shifted. Part of the role of the ethnographer is to make the unfamiliar understandable throughout the fieldwork and analysis. But it is also part of the role of the ethnographer in the field to systematically question the apparently familiar in light of new observations (Marcus, 1998).

The role of the ethnographer as a translator has led to ethnographic methods being used to translate between actors in society. Ladner (2014) describes how a result of this is that traditional boundaries between inside and outside perspectives acquire new meanings. There can now, for instance, be several distinct outside perspectives, for example, a company and its customers, where the company wants an ethnographer to grasp customer's perspectives on company matters.

Another track in the development of ethnography is that ethnographies are no longer bound to geography. Traditional fieldwork within anthropology has often meant fieldwork on a single site, where the notion of a culture is a phenomenon that is localized to that site. This made sense in part because traditional research often involved going to a civilization located at a site that had minimal communications with the rest of the world. Today ethnographic fieldwork is instead bound to other dynamic objects of study (Hannerz, 2003; Marcus, 1995). Marcus (1995) explains that this is because ethnography is no longer bound to anthropology. Instead, ethnography is a method that cuts through several interdisciplinary sciences. Marcus (1995) suggests a number of general objects that ethnographers can follow throughout the ethnographies. They are as follows: people, things, metaphors, stories, life, and conflicts. These objects are not bound to specific scientific traditions, and are not necessarily bound to single sites. Instead they are spread across sites.

Hannerz (2003) describes a similar contrast to single-sited research. He notes that multi-sited fieldwork, or in Hannerz's words translocal fieldwork, has become important because the human world has become more connected. Therefore, in contemporary ethnography relationships between sites are as important as relationships within sites. A methodological consequence of multi-sited fieldwork is that researchers need to select sites that grasp the current study objectives. Hannerz (2003) discusses how this selection can be a result of research design, but can also be a result of current circumstances in ongoing fieldwork. Also, site selection is often an impossible task because all sites that are relevant to the current object of study can hardly be scrutinized within a single ethnography. Nevertheless, similarly to Marcus, Hannerz posits that contemporary fieldwork is about following some object of study, to some extent independently of geographically-bounded sites.

There are several descriptions of how to conduct ethnographic fieldwork. As previously mentioned, to be able to move between the outside and inside perspectives, fieldwork in anthropology has traditionally opted for participant observations as its most important methodological approach. Participant observation is a natural method for obtaining these perspectives because participation yields a kind of experience that deepens the inside perspective and contrasts against the outside perspective. Czarniawska (2007) also notes that most ethnographic researchers have used participant observations. She therefore suggests that ethnographers need to contrast participant observations against non-participant observations more explicitly, using the term "shadowing" to refer to the latter type of observation. I note that there are many uses of participant observations within ethnography that also include shadowing

[8] Despite the study of the unfamiliar even Malinowski noted similarities between "primitive" cultures and western societies.

techniques to some extent. In fact, shadowing was likely what Malinowski did in many of the situations he observed. But I do find Czarniawska's categorization helpful for guidance on how to conduct fieldwork in practice. Much of ethnography is about accounting for different perspectives on the same object of study, and to be able to account for different perspectives, different observational techniques are important.

The various observational techniques are means to what Geertz (1973, borrowing the philosopher Gilbert Ryle's concept) argues to be the goal of ethnography: the production of *thick descriptions*. Thick descriptions, and hence ethnography, is defined as what produces "a stratified hierarchy of meaningful structures" (p.7). Consider two people that contracts their right eyelid (Ryle, 1968). At first sight the contractions appear to be very similar, but the contractions can mean several things. The twitch could for instance be a wink, but it could also be an involuntarily twitch. To understand which interpretation is more valid the researcher needs more information. For instance, they need descriptions of what happened before, what happened after the contraction, the possible social relationships between participants, cultural repertoires etcetera. The role of the thick description is to establish accounts for different interpretations. It is only with thick descriptions that ethnographic analysis can be conducted.

There are also several descriptions of how to conduct ethnographic analysis. Agar (2008/1996) describes this process as following a process of dynamic abduction. In other words, the ethnographic analysis iteratively accounts for observations by using the most likely explanation. Occasionally the analysis is followed by more fieldwork, and sometimes the analysis is conducted by consulting more information in the established descriptions. Kvale and Brinkmann (2013, p.268) describe a number of ways qualitative data can be validated, which is also applicable to ethnographic analysis: contrasting an observation against extreme cases, actively seeking contradictions to an observation, and getting feedback from participants in follow-up interviews are all ways to control and check for correctness of observation.

By using thick descriptions, Geertz argues that one goal of ethnography "is to draw large conclusions from small, but very densely textured facts; to support broad assertions about the role of culture in the construction of collective life by engaging them exactly with complex specifics" (p.28). This is an important point with ethnography, but it also an important point for case studies in general. It is a common misunderstanding that case studies cannot be used to draw conclusions about the collective, or be used to draw general conclusions (Flyvbjerg, 2006). Therefore, the ethnographic product is one that is a reasonably coherent story about culture. This also means that ethnographic products that are based on qualitative data can be analytically generalized to other settings (Kvale & Brinkmann, 2013).

In summary, ethnography is a methodological approach that aims to translate a group's perspectives on their situations to the outside world. The product of contemporary ethnography is, as for traditional ethnography, to study the patterns of meaning that exist in a community. But in contrast to traditional ethnography, contemporary ethnography is driven by a wider choice of research questions. This is partly because ethnography is now a common approach in more disciplines than anthropology. Irrespective of discipline, ethnographic analysis always seeks to find the most likely explanation for observed phenomenon.

3.7.2 Ethnography on cognition

Ethnography for the understanding of cognition became relevant for cognitive science in a number of publications. For instance, Lave (1988) made it clear that the perspectives of the participants, and hence ethnography, were highly relevant for the understanding of cognitive processes taking place in supermarkets. Also, when Hutchins (1995a) defines cognition as a cultural process, ethnography become a natural methodological tool to capture cognition. Ethnography for the understanding of cognition is often referred to as cognitive ethnography.

Cognitive ethnography is used as an approach that does micro-analysis of the flow of information in specific activities. Since it is based on ethnography, the participants' perspectives are important. In practice this means that, as with contemporary ethnography, researchers use a mixture of observational and interviewing techniques. But within distributed cognition micro-analytical approaches are also used without necessarily capturing the participants' perspectives through interviews. This is the case when, for instance, observations are used to study non-human cognition. For example, Johnson (2010) discusses how observational techniques can be used to capture cognitive complexity in non-human primates and cetaceans. She describes how a mixture of qualitative and quantitative data can be derived from video documentation of the subjects' behavior. These kinds of data can be used to document detailed patterns of behavior that reveal important mechanisms in a species' cognition.

Despite descriptions in the literature of observational techniques there are a few explicit accounts of what a cognitive ethnography is. I find it relevant to briefly review them because I find that they partly focus on different aspects of cognition, aspects that I have considered in my analysis.

In *Cognition in the Wild,* Hutchins (1995, p.371) briefly describes cognitive ethnography as the descriptive enterprise of documenting cognitive task worlds. As ethnography in general, it is not a predictive enterprise, but also as ethnography, it is certainly explanatory. For cognitive science the point of this enterprise is to create functional specifications of cognitive systems, which can be used, for instance, for design and experiments.

To substantiate this description Hutchins describes how, in any given moment, any human practice can be understood across the three developmental dimensions I described in section 3.3: the conduct of an activity, the development of the practitioners, and the development of the practice.

Although distributed cognition, through cognitive ethnography in its original form, embodies all of the above processes, it is most often only the conduct of activity that is addressed in the contemporary literature. For instance, in their description of the analytical framework Distributed cognition for teamwork (DiCoT), Blandford and Furniss (2006) explicitly de-emphasize changes across time and the role of social architectures. It seems that cognition as a cultural process is often de-emphasized when distributed cognition is used in work settings, which seems to be the very opposite of what Hutchins and his co-workers intended. This could be because distributed cognition, as a theoretical framework in practice, has lost connection with central aspects of ethnography. For instance, one way to understand what a practice is about is to grasp the perspectives and objectives of participants, but such descriptions seldom exist when distributed cognition is applied. Instead, objectives are, for instance, provided in advance by the specification of the operation under scrutiny.

Through the description of Hutchins' three dimensions it becomes clear that cognitive ethnographies do not start with a specification of which cognitive aspect to study. Instead, they start with finding the objectives of the activity under scrutiny. This is in contrast to traditional cognitive science, which usually decides upon which aspect (abilities or components) of mind to pursue, such as memory, perception, decision-making, attention, executive functions, spatial cognition, etcetera. In cognitive ethnographies, descriptions of activities precede the cognitive aspects. Of course, activities can overlap with aspects of mind to various extents. For instance, cognitive tests in laboratory-like settings are a special type of cultural activity that are often shaped to be activities about specific aspects of mind.

Although the descriptions of activities precede the cognitive aspects, several studies have adopted a distributed cognitive perspective on specific traditional aspects. One example is prospective memory in high-risk sociotechnical systems (Dismukes, 2012; Tobias Grundgeiger, Sanderson, & Dismukes, 2014). In such studies safety issues are the primary motivation for understanding a specific cognitive aspect, because forgetting one's intentions in high-risk environments can have severe consequences.

From a cognitive point of view, starting with activities makes sense because activities have goals. Activities exist to reach some concrete or abstract goal that can be evaluated from an information-theoretic perspective. From an ethnographic point of view it is important to understand that the goals for activities can be understood from several perspectives. There are, for instance, several reasons why navigation is meaningful on a US Navy ship. It is a logistical necessity, but its necessity has likely also been taught in the education of the US Navy.

Also note that although cognitive ethnographies start with activities, traditional cognitive aspects are not meaningless. Cognitive ethnographies both document the perspectives people have and practices people use, and how information is (see Neisser, 1967) transformed, reduced, elaborated, stored, recovered, and used. By having a distributed perspective researchers aim to minimize the risk of attributing those properties of cognition to the wrong subsystem (Hutchins, 1995a, p.356). This is also on par with the ethnographic bottom-up approach of going into the field with an open mind about whatever object is being studied.

I also find it important to point out that traditional aspects of mind can themselves be meaningful for people. This is not strange, since to a large extent the faculties of mind are derived from folk psychological ideas. In my ethnographic material I have observed several instances where the ability to remember became its own meaningful activity. I return to this in my analysis.

With the above background I find cognitive ethnography to be a way of doing contemporary ethnography, rather than a type of ethnography. Let me clarify. As we saw in the previous section, contemporary ethnography is not bound to specific physical sites, rather it is bound by some object of study. In principle, this object of study could be information processing in some activity that spans different architectures and sites. Cognitive ethnography is bound to understanding activities or processes because cognitive ethnography is bound to definitions of cognition as some kind of information processing that strives towards some goal. Therefore, I think cognitive ethnography is a way of doing contemporary ethnography because it is an ethnography about cognition. This is also my interpretation of Hutchins' use of the words. Hutchins' use of the words must be understood as a reaction against traditional cognitive science and the laboratory-based studies that had been dominant, but not as a special kind of ethnography. That said, this is not an established interpretation.

Williams (2006) makes the process of doing cognitive ethnography transparent by noting that cognitive ethnography is separate from traditional ethnography in a fundamental way. Traditional ethnography is described as being interested in the meanings that members of a culture create. In contrast, cognitive ethnography is interested in how these meanings and other knowledge "determine important outcomes" (p. 838). He uses the example of kinship. Traditional ethnography investigates the states of kinship. Cognitive ethnography would need to understand the states of kinship before understanding a cognitive aspect, such as group decision-making. Therefore, when making this distinction Williams stresses that traditional ethnography is a necessary starting point for doing cognitive ethnography. I am sympathetic to the notion of making the process of ethnography transparent in this way, but I am not certain that the delineation should be between process and content.

At the beginning of the chapter in which Hutchins introduces cognitive ethnography, he begins by criticizing the distribution of labor in cognitive science (which was proposed by D'Andrade, 1981) for the sub-disciplines of anthropology and psychology. Psychologists were supposed to focus on cognitive processes, and anthropologists on cognitive content. I consider this distribution to be partly equivalent to Williams's separation between ethnography and cognitive ethnography. But there are several cases where anthropologists study processes without using cognition as a theoretical framework (see, for instance, Malinowski, 1922 and his account of the Kula trading ring). And there are several instances where cognitive scientists focus on content (see, for instance, Huth, de Heer, Griffiths, Theunissen, & Gallant, 2016) without focusing on process.

I find that viewing something as the content for some process can be clarified through Hutchins' three developmental processes: practice, practitioner, and activity (which I described in 3.3). If we study an instance of an activity, the nature of the practice can be considered content to the instance. But in this case, the definition of content is only an epistemological perspective. What counts as content is relative to what is considered the process. Recall from previous sections that, from a distributed cognitive perspective, cognition is a cultural process and culture is, to some extent, a cognitive process. Therefore, in retrospect even content can be understood from a process perspective; that is, how the content came to be. To use another example: friendship can both be viewed as a state and as a process. Humans do not only have friendship-based relationships, humans also make, maintain, and change friendship-based relationships.

In ethnography, the meanings that members of a culture create are orthogonal to, or even independent of, practices and activities. There is some unit of analysis that ethnography captures that is not encapsulated within one practice or activity, but that still can be a determinant for practices and activities.

Friendships and decision making can be used an example. The nature of friendships within a group of people is a set of meanings the people in the group have created. These sets of meanings can shape decision-making within the group, but they can still be studied independently of decision-making. I do not think that this suggests a branching of ethnography and cognitive ethnography. Instead, I view this as a natural extension of the original description of cognitive ethnography, i.e. as a description of the task world. Decisions and decision making can be viewed as a cognitive practice, but so can friendships and the creation of friendships. I can imagine a study where establishing friendships is considered a task, and cognitive ethnography is used to describe the task world for this.

In this thesis I use the term cognitive ethnography to describe an ethnography that is about cognition. The concept of "cognitive" in cognitive ethnography refers to distributed cognition as an alternative to traditional cognitive science. Therefore, it is not necessarily a special kind of ethnography. If we look at how ethnography is used today I believe we can see a versatile use of it that spans both the content and cognitive processes of what people do (see, for instance, Crabtree & Rodden, 2004). Sometimes cognitive ethnography is an ethnography that analytically focuses on constructs that are traditionally assigned to the mind, but that, initially in the empirical process, remains open to what architecture those constructs should be ascribed to.

I now turn to my own work of studying cognition through ethnographic methods. I have used a mixture of ethnographic approaches; whereof some have been used to capture participants' perspectives and some have focused more on capturing the details of various practices.