7. Intermission

7.1 Concerning the general validity of the theory

I have spent the preceding chapters making the case for design as being a cognitive process of a certain kind, but all this would be of limited interest if it applied only to design. Therefore, in this concluding chapter I will ask: How does my argument generalize? Does my exposition hold beyond the cases I have presented here?

It should first be noted, however, that the previous chapters have not been concerned only with design. On the contrary, it has been my aim to build a theory that is formulated on a general cognitive level, while using phenomena from design as the *basis* for this. Chapters 1 and 2 served to show the connection between design and the model of rational action, and thereby that the issues at stake in this book bear directly on the general views of cognition, in folk psychology as well as in scientific theories. Already in Dewey's original formulation, the theory of inquiry concerned cognition in general, and in chapter 3 the inquiring function of action and the "no pure analysis" conclusion were both formulated as general theories of cognition, as was the model of interactive cognition and the related concepts in chapters 4 and 5. The general relevance of chapter 6 will be discussed as part of what follows.

Thus, the preceding pages have not only been of relevance to design, even though one might at first think so. I will now also return to the bigger issues that I mentioned briefly in the Introduction, of whether cognition is intellectual or practical by nature. So how widely *does* my case for the practical view hold?

Let me briefly recapitulate the two contending views: In the view of cognition as basically *intellectual*, our most advanced cognitive capacities are the intramental ones; they comprise the highest achievement of evolution, and they are therefore what sets us apart from other species, and by which we do what no other animal can do. These are the capacities that have always been the main concern of cognitive science: mental representation, mental models, mental simulation; logical inference, case-based reasoning, planning, scripts, and so on.

In the other view, cognition is basically *practical*. There, also intellectual abilities are practical skills of a special kind, where extramental working materials and activities play key roles, and where these in turn are the products of cultural development and refinement. Hence, these abilities are not attributed to certain intramental equipment, as in the intellectual view. Surely the human brain goes beyond those of other animals, but to attribute our leg up only to mental modules that are uniquely human is too simple as an explanation, and too implausible as an evolutionary argument.

Defense #1: Intramental substitutes are used when interaction doesn't work

These are the two antagonists. The intramental model of cognition has always been advocated as superior, but this I have contested. Today it is old news that it doesn't work as advertised, but in addition to this I have sought to explain *why* it is so. And then I have presented a contrasting account of design as interactive and inquiring, but I have also made the case that because of its interactive nature, it is performed more effectively than it could be, even at best, if it followed the principles of traditional theory. These are two good reasons for cognition to be interactive rather than intramental. Thereby I have meant to, as it were, give cognition reason for being interactive.

But from this only, proponents of intramental theory will not accept that cognition is always interactive. The model I have presented gets its superior performance from drawing on working materials, action, and so forth, in the cognitive process. But this is not always possible; in particular, when the object of concern is not physically present and available to cognition. Let us call the favorable conditions *presence*, and the opposite conditions *remoteness*.

Being "representation-hungry"

A defense of the intramental view along these lines has been advanced by Clark & Toribo (1994) who cannot be considered to be conservative in this matter (pp. 418–419):

> The basic trouble ... is that the kinds of problem-domain invoked are just not sufficiently "representation-hungry". Instead they are, without exception, domains in which suitable ambient environmental stimuli exist and can be pressed into service in place of internal representations. ... it is unfair to use these cases to illustrate any more general anti-representationalist claim.

By a representation-hungry problem domain we mean a domain in which one or both of the following conditions apply:

1. The problem involves reasoning about absent, non-existent, or counterfactual states of affairs.
2. The problem requires the agent to be selectively sensitive to parameters whose ambient physical manifestations are complex and unruly (for example, open-endedly disjunctive).

Point one is a straight-forward definition of remoteness; Regarding the second point, the authors clarify that it mainly refers to cognition that deals with abstract properties that have no physically "manifest stimuli" (e.g. visual).

So, the argument goes, whereas cognition may well work interactively under presence, under remoteness it still relies on the full intramental capacities; indeed *must* rely on them, to compensate for these difficulties when remoteness prevents an interactive mode of operation. This is the most natural line of defense, and it is after all how intramentality and representation has always been motivated: It is superior because it can go beyond the here and now; beyond what is immediately present, beyond even what yet exists; to consider alternative, hypothetical situations, and so forth. This argument has been advanced several times, in modern day e.g. by Popper (1935/1959), who originally proposed that our hypotheses can die in our stead (also cf. chapter 5), Craik (1943), and Johnson–Laird (1983).

This *surrogate* capacity (cf. Clark & Grush in press) has a straightforward translation into the evolutionary claim: it is often held to be uniquely human or likewise; other species can purportedly also do the interactive things, but humans can do more, because they can use surrogates when the real thing is in limited supply.

Note, however, that Clark & Toribo have (on purpose) backed down from the fundamentally representational mind (also cf. Clark 1997). Previously, the advantages of mental surrogates have been taken as reasons for their being used by cognition *universally*, also when the original goods are readily available. However, with defense #1, a conservative position might even maintain that cognition still is intramental at the core, while conceding that it may take advantage of interaction when it can. The problem remains of explaining how cognition is sometimes interactive and sometimes intramental, but this is doable and merely a practical matter, even though as always with hybrid models, the result may not be aesthetically pleasing: Perhaps there are two systems that alternate, or maybe even the intramental processes are still running in the background during interactive work, maintaining an internal surrogate model of the environment even though it is not used just then.

By introducing the notion of "representation-hungry" domains, Clark & Toribo have established a clear line of defense, which is good since it identifies a meaningful direction in which to proceed with the argument. These hungry domains then ought to do the intramental capacities more justice, and there their utility should be evident, the authors claim: Whenever conditions degrade far enough, *then* cognition falls back on the representational, modeling, simulating system; when no real food for thought is available, then surrogates are provided for a mind that is starved.

Design is absolutely representation-starved

So arguments against intramentality should not concern interaction-friendly domains, i.e. presence, but representation-hungry ones —remoteness. Let us therefore be fair and do just that, because it can be done quite easily. For as I argued briefly in chapter 6, design *is* a representation-hungry domain; a closer look shows that the conditions for design are *just those* where representations are purportedly needed the most, where intramental cognition should work best, and therefore be of greatest value: Everything of that which is in the designer's chief concern is *twice remote*, in the future situation of use: It is *both* absent *spatially* and non-existent *temporally*. And perhaps it counts also as counterfactual, in that the whole point of design is to bring about something that today is not so. Clark & Toribo state that one of these conditions is enough to make design a hungry domain.

Design then qualifies as a perfect example of when mental surrogates should be of greatest utility, also in the classical view. The designer is concerned with a state of affairs that is distant in both time and space—and which Popper might well have called "hypothetical". And in many accounts, design is concerned with highly abstract properties, as in the second point above. Hence, it would be hard to think of a better match for their criteria on hunger. Design has also many times been explicitly stated as an important domain for intramental theory to explain in general; without reference to this particular issue (e.g. Akin 1986b, Newell & Simon 1972, p. 7, Simon 1973, 1981). It therefore cannot by any measure be regarded as an unfairly chosen domain.

Situating strategies: designers go out of their way to avoid intramental thinking

Hence, design is a perfect match of the conditions where the intramental capacities should come to their very best advantage, but in chapter 6 I showed at length what actually happens: Rather than falling back on intramental capacities, designers use *situating strategies*; they strive to restore or "re-create" presence, so that they can work interactively nevertheless. This goes to show that when hunger sets in, the switch from interactive to intramental cognition doesn't happen. And, after all, isn't it obvious that this is how it has to be? When you are starved, only *thinking* of food just makes you even more hungry.

There are two circumstances that make the force of this argument particularly strong and generally valid. Firstly, this is not a choice between equals, since the interactive alternative is not as readily available as the intramental procedure: To at all make it into an option, designers first have to use the situating strategies to *create* the working materials to interact with. Since they in this manner go *well out of their way* to enable the interactive mode, it means that the advantages of interactivity over intramentality are very strong (it might for example mean that the penalties of working intramentally are very high), which also increases the strength of this fact. So instead of falling back on intramental capacities, cognition spends *extra physical effort* with the only purpose of *avoiding* having to think intramentally.

The second "aggravating" circumstance makes a much stronger case for the *generality* of the interactive model: If cognition is not intramental even here, when the conditions for it are the best imaginable, then when would it be, since all other situations are *less* suitable than this one? Design is the purest possible manifestation of the exact problem that intramental theory was meant to solve; even in the words of the proponents themselves.

And in evolutionary terms, when these intramental abilities are in reason not used to perform the functions that are considered hard and uniquely human—the very functions for which they were once advanced—then when *are* they used? And conversely, when we rely on the older, more primitive interactive capabilities to perform the most difficult, hungry tasks, then what is the reason for having the advanced, intramental ones?

Hence, the "representation-hungriness" of design, and the use of the strategies to re-create presence, are strong cases in point for the *general* plausibility of the interactive model.

Defense #2: When also the strategies fail, then cognition becomes intramental

But this does not exclude quite all possibilities for intramental cognition. An ardent defender would still not yield, but might instead present a second counterargument: Maybe designers work interactively when they can, and when they can't do that, then they use the situating strategies; but when *they* break down, *then* the mental movie starts to roll. *This*, someone might claim, is when the intramental capacities come to their best advantage.

First, cognition would hardly be *fundamentally* intramental if this is the last way out; one might however still claim that intramentality, though not fundamental, provides us with capabilities for handling these extreme situations, which we otherwise wouldn't be able to deal with. But secondly, one then also comes to wonder what these circumstances would be: more extreme than those for which these intramental theories were created? In design for example, more extreme than design-by-drawing, so that the designer is forced to rely on abstraction and thinking only, without working on solutions, without involving users, and so on.

What comes to mind is a designer trying to comply with one of the systematic design methods from chapter 1. In terms of concrete working materials, and favorable conditions for interactive cognition, no approach to design could be more deprived than these methods: Design is inherently remote to begin with, and to a rather substantial extent, too. But on top of this, the approach that design methodology prescribes makes the inherently poor conditions even worse.

The objective of design methods is to focus on the abstract, logical structure of design problems, so from the view I have presented, their effect is essentially to deprive the design task to the fullest possible extent. From the ideal point of view, however, this approach is natural, with its emphasis on abstraction: The procedures that e.g. Alexander, Jones, and Simon advocated are concerned only with logical relations and abstract criteria like requirements and constraints—no drawings, nothing tangible, nothing concrete (Alexander 1963, Jones 1970, Simon 1981).

But from the interactive viewpoint, the goal that design methods aim for comprises the worst possible conditions for cognition, and the

results have also been extremely poor; as seen in chapter 2, it is not only very hard to find successful examples, but even any examples at all, since the methods have proven so thoroughly impossible to use. Thereby the view I have advocated here would explain why design methods have yielded such disappointing results.

But to address this second defense on a general cognitive level requires a more elaborate discussion: it is a tricky matter, because the cognitive science tradition has made a diversive maneuver around the problem, rather than confronting it directly. And the diversion is large enough to require a section of its own.

Explaining (away) poor intramental performance

The evidence from design is quite conclusive that purely intramental performance is very poor, but the evidence is not restricted to design. On the contrary, this is probably the most well-documented fact in all of cognitive science: Innumerable experiments have documented so-called "cognitive limitations", "cognitive strain", etc. Whether the task studied has been concept formation (Bruner, Goodnow & Austin 1956), planning (Hayes-Roth & Hayes-Roth 1979), the comprehension of complex sentences, syllogistic reasoning (Johnson-Laird 1983), attention span, memorizing, mental models and mental simulation (cf. chapter 6 and Norman 1983)—the list could go on forever—whenever purely intramental performance has been studied, the result has always been that people do not perform according to the principles of intramentality.

Limitations on working memory: the tragical number seven, plus or minus two

But what is truly striking is the role this supposed limitation has come to take. It is not regarded as a measure for explaining (away) a theoretical mal-prediction, but as a celebrated scientific finding about the nature of human cognition—that is, the mind as having important limitations in its information processing capacity. This theme is so central that it has even been elevated into a general scientific principle, with its own element in the standard models of the human mind, known as *short-term memory* or *working memory*. Because of its role as the heart of all mental processing, every cognitive function needs to involve working memory in its operation. It is thereby the spider in the web of these models, which can easily be made into the fly in the ointment: If it is attributed with a flaw, then a problem in *any* cognitive function can be ascribed to this single defect.

This theme has been even further reified, by saying that working memory can hold only the legendary 7±2 items at a time. Equally legendary is the paper to which the origin of this fact is attributed, the full title of which is *The magical number seven, plus or minus two: Some limits on our capacity for processing information* (Miller 1956).

This is arguably the most well known, most popularized and most widely disseminated fact that cognitive science has ever produced, as can be seen in any popular scientific writing that touches on the subject of cognitive performance (Csikszentmihalyi 1990 is a prototypical case). Also, innumerable are the design features that have been claimed to "alleviate limitations in working memory"; citing this as the reason, for example, it has been stated that computer menus should hold no more than seven items for the user to choose between; it even seems that phone numbers were made sevenish digits long for this reason, so as to minimize the number of calls to directory assistance (Ellis & Beattie 1986).

Add to this the concept of a "chunk", which means that one of these seven slots can be said to contain anything, however large or small, complex or simple, and the size of which can be chosen so as to suit your purposes. Thereby you can maintain, for any given body of material, that it either fills up, overflows, or barely fits working memory—whichever you prefer (cf. Miller 1956, Simon 1974, 1976, 1979).

In this way the generally held principle, "limitations on human cognitive capacity", has been translated into its own architectural feature in these models, as a "limitation in working memory capacity", and has then been attributed to the specific size of this memory. It could arguably be maintained that the primary purpose of postulating this memory system has been to explain this "limitation".

Competence vs. performance

Closely related to this idea is another explanatory meta-theme which is also very wide-spread. This is the *competence vs. performance* distinction originating in Chomskian linguistics:

> A theory of the former would be a theory of linguistic knowledge and grammar, of what an *idealized* mature speaker-hearer of a language *could* say and understand; a theory of the latter would be a theory of behavior, of what *real* speaker-hearers *actually* do say and how utterances of others are understood. (Reber 1985, p. 137, emphasis added)

A better illustration of the contrast between the ideal and actual perspectives cannot be found. Chomskian theory portrays the human language faculty as based on exact rules that determine with great precision what language should look like. This theory is based on notions like formal grammars, automata theory, and other foundational theoretical principles of computer science. The problem is, of course, that this is not what you find if you look at the language that people use. To accommodate this circumstance, it is said that the theory concerns people's *competence*. This they are held to indeed have deep down below, but it isn't adequately revealed in people's behavior. That instead reflects their *performance*, which is much less sophisticated.

Now the crucial issue in this is to explain the *discrepancy* between competence and performance, and in particular how it can exist. This is done precisely by referring to people's inherent limitations in their "information processing capacity", and stating that the competence is degraded by the cognitive system that cannot fully handle it. This is often done with a direct reference to working memory capacity and the magic number.

Also the competence–performance theme has been used far beyond the original domain of linguistics, since it addresses a widespread and general scientific need. As a result, everything related to people's actual "performance" has attained a distinct derogatory ring to it. And since the study of performance is therefore not a study of cognition *an sich*, it has been regarded as "applied" science and of lesser value, since it tells us little about how cognition *really* is (and which we thus rarely see outside the laboratory).

Psychology as the study of human mental imperfection

Like so many other times, Newell & Simon have seen the weight and scope of also this matter clearer than many others, and they have addressed it on a general, domain-independent level, thereby making the issue very clear. But in so doing, they have also made the seriousness of the problem stand out more clearly, showing that it is not a minor issue with a restricted range of impact. In effect, their conclusion is that psychology is the study of human limitations and shortcomings:

It is precisely when we begin to ask *why* the properly motivated subject does not behave in the manner predicted by the rational model that we recross the boundary again ... to a psychological theory of human rationality. The explanation must lie inside the subject: in limits of his ability to determine what the optimal behavior is, or to execute it if he can determine it. In simple concept attainment experiments, for example, the most important mechanism that prevents the subject from adopting an efficient strategy is usually the limit on the number of symbols he can retain and manipulate in short-term memory. To the extent that this is true, such experiments are experiments to reveal the structure of human short-term memory ...

1. To the extent that the behavior is precisely what is called for by the situation, it will give us information about the task environment. ...
2. To the extent that the behavior departs from perfect rationality we gain information about the psychology of the subject, about the nature of the internal mechanisms that are limiting his performance. (Newell & Simon 1972, pp. 55–56, referring to Bruner *et al.* 1956)

Here, they virtually *define* psychology as the study of how people fail to perform rationally; alternatively, of how they fail to behave as intramental theory says. Adequate performance is not a topic of psychology, it merely reflects the "task environment", but *poor* performance is: "when the subject does not behave as predicted by the rational model, we recross the boundary to a psychological theory". With such an assumption as a prominent part of the underpinnings, the scientific results are surely predisposed to point in a certain direction. There also seems to be a division of labor between different disciplines: one sets up the theories, the other documents how people fail to follow them. No room is then left for the theories to be influenced by how people actually behave; the ideal and the actual should not be mixed.

The same theme is discussed in Simon (1981), and it is also related to Simon's earlier work on "bounded rationality" and the concept of "satisficing", e.g. in Simon (1947). His distinction is probably based on the difference between e.g. economic models which often have assumed that the behavior of the involved agents is rational, and psychology as the study of how it isn't.

With respect to defense #2, the cognate themes discussed here all point in the same direction—"limited information-processing capacity"—from empirical results, to magic numbers, to competence vs. performance, all the way up to meta-scientific principles. The conclusion must then be that it is an *extremely* well-established fact that mental-only performance is quite circumscribed, not only in design, but in cognition in general: the mind alone does not reach any great heights.

Taken together, these "limitations" themes are how the ideal perspective accounts for the fact that cognition doesn't work as the theory dictates. Hence, this is a generally assumed position which appears even more ardent than the second line of defense I proposed above: It states that when the mind is forced to work intramentally, then it performs very poorly—still, it *is* in essence intramental. In this position, there is indeed a discrepancy between competence and performance.

7.2 A matter of choosing the proper yardstick

But now take a step back and look at what the "limited capacity" argument really says. First, cognition is held to work in a certain way; as a computer, intramentally, etc. Then, people are found not to behave as predicted, and quite thoroughly so at that. But the conclusion drawn is that they are information processors nevertheless—only *very bad ones*. The result is the following syllogism, which is the conclusion of defense #2:

The mind is a computer
It does not perform like a computer
∴ The mind is a *malfunctioning* computer

The logic behind this reasoning does comes across as slightly twisted: Is it the mind or the theory that should be sent in for repair? (Other variants are obtained by replacing *is a computer* with *is intramental* or *is rational*.)

Thereby, intramental theory faces two explanatory problems: In addition to the poor forced intramental performance, it should also account for how cognition can come to flourish in the presence of extra-mental factors. The sentiment is after all that the surrogate capacity should make the presence of external materials and so forth an insignificant matter:

> For the crucial activities ... take place centrally. This is true even when the desired object and the required activity are physical. (Newell & Simon 1972, p. 72, cf. chapter 4)

(But you already know the explanation: using the world alleviates the limitations on working memory.) Consider in comparison the explanation I have proposed. Concerning the latter question about the use of external material, I have already presented the interactive account at length. But what about the problematic performance on strictly intramental tasks? Quite simply, natural cognitive performance counts in three contributing parties: mind, world, and action. So what happens if you disable two of these? Of course, the performance breaks down. *But*, more importantly, this constitutes a highly contrived, unnatural task and thus very atypical performance: the tasks that are studied in the laboratory are very poor benchmarks to measure human cognitive capacity by. Therefore they are of limited interest. Psychological experiments do not eliminate the influence of situation and context, thus providing "generalizable" results, as experimental method holds; instead, also an experiment is a situation and context of its own, but of a very peculiar kind, and with very unrepresentative characteristics, yielding equally unrepresentative performance (cf. Whiteside & Wixon 1987, Wixon & Holtzblatt 1990). This can also explain why the experimental results have proven not to generalize to performance elsewhere.

As I see it, the main problem with intramental theory is its choice of *yardstick*: When the standard of measure is badly chosen, then the resulting measurements will necessarily be skewed. In the Introduction I gave a number of examples where human behavior had originally been framed as limited, substandard, and "irrational", but where others have later showed, by taking a different point of view, that human performance in these cases *makes sense*, even being clever and sophisticated. Compare with the examples given in the Introduction, e.g. the quotes regarding Micronesian navigation that were given there. My point is that the same is true here: When intramentality is the badly chosen norm, the resulting accounts of human performance will be biased by predestination; the resulting explanations will be in terms of deviations from the norm, that is, as limitations. Consider the following, admittedly somewhat drastic allusion:

> Our object of concern is human movement. The fundamental form of human movement is by flying. With the advent of modern aviation, we are able to give a detailed model of this capacity: we consider humans to be airplanes. Or more correctly, we consider humans, birds, and airplanes as three instances of flying systems.
>
> However, experimental studies have consistently shown that the human capacity for flying is limited—compared to for instance a Boeing 747–400. Much of this can probably be attributed to the human arms not being at all efficient as wings.
>
> The poor flying capacity is very well documented. For example, experiments off the Tower of Pisa, in the spirit of Galileo, have consistently shown that people can fly for only 7±2 seconds. Also, their landings are really messy. Accordingly, the study of human flying in effect becomes the study of human crashes.
>
> This also explains why we outside experimental settings mostly see people walking. Because of the limited abilities for flight, the legs that are really meant to serve during take-off and landing have become the major means of compensating for these limitations.
>
> This discipline has also, via knowledge of the limitations of human flying capacity, given us important guidelines that are most helpful for design. Some examples are: Build houses on or near the ground; use floors in rooms, particularly if the rooms are above ground level. Also place furniture on the floor, and door openings on the lower parts of walls.

Even though the points here may seem unjust, they have direct parallels in the cognitive literature: almost all of the statements in the first four paragraphs have direct counterparts in Newell & Simon (1972); some of them also e.g. in Miller (1956) and Card, Moran & Newell (1983). As the model for the last paragraph stands the advice found in *Applying Cognitive Psychology to User-Interface Design* (Gardiner & Christie 1987). The following are some examples of guidelines related to working memory:

- Working memory load increases the greater the amount of material that must be remembered temporarily, or "held in mind". (p. 159)
- If the number of referents, or the number of properties and relations ascribed to the referents used in a dialogue exceed the capacity available in working memory, then the probability of

considering all the relevant referents, or relevant properties and relations, will be reduced. (p. 117)

- If the sentences used in a dialogue are open to more than one interpretation, then they are read slower and less accurately than sentences that are open to just one interpretation. (p. 116)
- Do not present information that is *irrelevant* to the task that users are trying to perform. (p. 249)

In the truisms that are stated here, it seems that cognitive psychology is either completely absent as in the last two points, or else its technical terms have served to cloak the underlying self-evident truths with belabored sentences. Interestingly enough, the second point seems to say, "Complicated sentences are more difficult to understand." My personal favorites are however the following guidelines:

- hand-held devices, such as a mouse, light-pen or digitizing tablet-and-stylus combinations lend themselves ideally to drawing, pointing, selecting and moving tasks — in other words, spatial and visual tasks. (p. 271)
- keyboard-based commands are particularly appropriate for word-processing applications ... Similarly, numerical data entry is best served by a keypad (*ibid.*)

But the original purpose of the flying analogy was to show how the chosen perspective can bias your vision, to the point where your view of things becomes outlandish. The whole theme of cognitive limitations was caused by the fatal choice of yardstick, which also lies behind the ill-conceived syllogism above: the yardstick is contained in the first line of the syllogism, which states that the mind is an information processor, and as such this chosen perspective functions as the axiom on which the conclusion will rest.

Also the *change* of yardsticks makes the difference between the two explanations: Given a more appropriate standard of measure, poor intramental performance is no anomaly, so the explanatory problem doesn't even occur. In the interactive view, intramentality is forced and unnatural, and with such restrictions the poor performance is to be expected. Hence, for one perspective the performance is coherent with the theory, but for the other it is an anomaly, and the difference results from the respective yardsticks. What makes the anomaly particularly serious is that it concerns not a minor or peripheral issue, but the main concern of the theory: intramental performance is anomalous to the theory of intramentality.

What the argument by Newell & Simon makes so clear is that the chosen yardstick can put you in a truly absurd situation, where theory forces certain decisions on you; here, having to define psychology as the study of how people fail to match the explanation you have chosen. As I see it, the main failure is that what is really a difficulty for your theory is made into a fault in the object of study. Consider physicists claiming that the deviations from their theoretical predictions were due to the universe committing "physical error", or limitations that give the world a restricted capacity to follow the laws of physics.

Probably the fatal standard of measure was adopted without much deliberation, or even seeing that alternatives existed. Soon, however, this created a situation where theory had to accommodate to the resulting measurements: concepts such as bounded rationality, limitations in cognitive capacity, the competence vs. performance distinction, working memory and its limited size, and so on. However, the premise in the above syllogism was never questioned.

But when we look back today, we can see that the premise is the cause of the problem. Because implicit in the view of the human as a rational, intramental computer was also the yardstick. And it was this yardstick that created the problem to explain: In reality, the question thought to be how people perform so poorly, was in fact the problem of why they perform so poorly *according to this yardstick*, which had been implicit in the view of cognition as rational and intramental.

Hence, it is worse not to realize that you are using a certain yardstick, than it is to deliberately have made a choice that turns out to have been bad. This is much like the matter of problem setting (from chapter 3): Instead of believing that your work is to produce the right answers to the only existing question, the most important insight is to know that what question you are asking will greatly influence the answers that result. The yardstick you are using, like a question and a problem definition, *sets* the types of result your work will produce. When you know that, the greatest part of choosing the right question or yardstick is already done.

7.3 From intramental functions to interactive technologies

I have here advocated a shift in perspective, from an intellectual/intramental view to an practical/interactive view of cognitive phenomena. If we are to make this shift, then what are the changes that will be necessary? How should the theories be changed? What will the results look like?

There already seems to be an emerging trend in which abilities, which were previously assumed to be intramental, are being reinterpreted as interactive and inquiring techniques. These developments are anything but concluded, and I can merely try to convey an idea of what may lie ahead of us.

Planning can serve as a good illustration. Ever since the birth of cognitive science, planning has been thought of as a fundamental cognitive function that is hardwired in the human information processor (Miller *et al.* 1960). According to Camhis (1979), the use of planning as a scientific concept began after World War II, in several different domains, urban planning as well as the philosophy of science, etc., compare with Polya (1945) and chapter 1. But since the emerging critique of cognitive planning theory, by in particular Suchman (1987), planning has begun to lose this status as a privileged and fundamental cognitive capacity. Instead, as for example in Agre & Chapman (1990), plans have come to take on the status of one cognitive technique among many, not more fundamental than, say, writing or riding a bike, but neither any less.

Spatial navigation and the use of maps are another example. For long now, animals as well as humans have been attributed with "cognitive maps" for navigational purposes. Gallistel (1990) represents this view, attributing them to animals as primitive as bees. Now there are indications that this attribution has been done too hastily. This was just what the Åkerlund quote in the Introduction concerned: since the Micronesian seafarers didn't navigate or use maps as we do, it was first concluded that they *couldn't* navigate at all. So just because bees find their way, they needn't use maps, because this is not the only way. For example, Hutchins (1983, 1995) showed that such things as viewer-independent perspective and certain representational techniques, which are required for creating maps, were invented in Renaissance times, are acquired by schooling, and do not even make sense to people in some cultures. Still, as was his point, their navigational feats are remarkable.

Hardware support for cognitive processing

Thus, there is an emerging pattern by which these sophisticated cognitive capacities, once thought to be naturally intramental and "hardware-supported" in the brain's computer, have begun to be reinterpreted as advanced cognitive *technologies*, highly developed working methods that are supported by likewise artifacts; *genuine* hardware support, that is. The procedures and artifacts alike have developed and matured over time, they in turn having been made possible by advances in mathematics and science. Hence, they didn't spring out of nothing, nor have they always been there, *in* there. Also, using them requires acquired, non-trivial skills.

This is what it means that also an intellectual ability is much a practical skill, but of a special variety. Plans are ways for organizing activities ahead of time, and they make use of technologies like writing and linguistic representations of activities, alphabets and so on, and inventions like paper and writing tools—all of which we take for granted. Taken together, these allow us to record, modify and reorganize resources for structuring activity, and later to use the resources to do just that.

The use of maps is even more obviously technological by nature. It involves the same writing tools, but also intellectual inventions like the birds-eye perspective, graphical abstraction and representations of the physical environment, 2D-projections of space, and also such recent advances as geometry, trigonometry, the Mercator projection and its likes, the whole domain of map-making methods and technology, the measurement and numerical representation of distances, and so on and so forth (see Hutchins 1995 for a comprehensive account). The use of maps comes natural to many people, particularly those with academic training, and even more so to those with a background in mathematics, geometry, and related disciplines like computer science—but not to everyone even in Western culture.

Design restored

What I have been attempting on these pages is, in a sense, a similar de-construction and re-construction of design, aiming to show that design is not a purely intramental process closely tied to the fundamental mechanisms of intentionality and planning, but a similarly sophisticated cognitive technology; developed over ages, and relying on subtle but sophisticated co-evolved artifacts and working techniques. Such is the combination of soft lead pencil, drawing paper, and techniques such as thumbnails, which together enable a highly fluid and expressive way of working that computers are far from matching.

A similar reinterpretation can be applied to for example mental simulation vs. hand simulation, as in chapter 6; we can do simulations reasonably well, given the proper supporting tools and techniques, but not so well otherwise. Some people have asked me, "Are you saying that we cannot plan?" No, I am saying, *But look at how we do it*; for example, look at how "mental" simulation is done. In comparison, am I saying that we cannot design, just because I argue that design is not a stage of pure intramental analysis, separated from the other activities of inquiry?

Hence, I don't claim that we cannot plan, but accounts of planning and so forth must be revised, like "mental" simulation. In one sense, design *is* the "restored" view of planning. "Plan" comes from the Middle French *plant*, which means ground plan or map (also influenced by Fr. *plan* as in flat surface, cf. English *plain*). In other words, from drawing a floor plan—i.e. architectural design. The term planning is thus an abstract rendition of design, derived from a process whereby you make plans literally by drawing. It appears that the practice of sketching as a means of design was developed at the same time as when design became a function separate from building (Gombrich 1960, Herbert 1993). Somewhere along the way, as "making plans" became intramentalized, both the working method and the materials and tools were dropped; they were in effect all made into epiphenomena—perhaps the tale of the singular creative idea was born here. Graphic designers also sometimes refer to sketching as "planning" the poster, folder, or other whatever they are producing (e.g. Black 1990). My account has thus merely reinstated the extramental components of the activity (inquiry), the doing and working techniques (sketching), and the materials and tools (paper, pencil), into the cognitive function of planning—such as it was in the original meaning of the word. One purpose of chapter 1 was to show how very closely related design is to general cognition, linked via the model of rationality and rational action.

In the same way, many other important activities have always been interactive and inquiring, but the intramental yardstick has caused this to be ignored. As compensation, mental mechanisms have been invented to handle what is actually done by interaction: representation and mental surrogates of the "outside" world are the paradigmatic examples (cf. Hutchins 1995, ch. 9).

It is such reconsideration that may lie ahead of us, for example regarding cognitive maps: the potential rediscovery of material and procedural factors that have been forgotten, and which would show that navigation, like planning and so forth, may not be as intramental as has been assumed. And if Micronesians navigate quite differently from Western seamen, then why do bees necessarily navigate with Western techniques?

But since these intramental mechanisms are part and parcel of cognitive science, the revisions will have to reach deeply into the groundwork. It will not be sufficient to for example talk about "external representations" and thereby keeping the theory of representation in place, with a new "external" specimen added to it. Hence, human abilities need to be explained differently, but also cognitive theory and explanation need to be reinvented. Instead of proposing new intramental innovations, we need to look for the answers elsewhere, in a very literal sense.

The final question

The topic in this last chapter has been whether cognition is intramental or interactive *in essence*; whether it is basically intellectual or basically practical. I have presented possible arguments and defenses from both camps, but there is one point which they agree on: As I have shown with the "cognitive limitations" theme, both sides agree that mental-only performance is underwhelming, and on this the evidence is decisive—*this* is not a matter of debate. So if you still wish to maintain that cognition is intramental, then you also need to adopt the "cognitive limitations" explanation. In interactive theory in contrast, the mind working on its own is only a circumscribed portion of the full cognitive system, and the unimpressive performance that has been documented so thoroughly is entirely to be expected.

Thus, between the alternative explanations of human performance on cerebral tasks, it doesn't come down to right or wrong, but to a matter of judgment. For example, the above syllogism seems to follow a rationality of its own: Would you say that this strange logic is due to limitations in the information processing capacity of its originators? I personally wouldn't. I'd grant them that they, too, would rather be playing frisbee.