4. The cognitive roles of action and world

> Thinking is one of the most notoriously intractable parts of psychology since the thought process is not easily observed. ... The designer, however, has never resembled Rodin's "Thinker" who sits in solitary meditation, but has in contrast always externalized his thoughts, not only as an end-product in the form of a design, but as an integral part of the process itself in the form of drawings and sketches. (Lawson 1980, p. 96)

4.1 Introduction

The topic of this book can be formulated as a question: Why do designers work the way they do, when the traditional theories of cognition and design say that designers should be doing something quite different? This chapter and the next bring the question down to the level of actual action, looking at what happens moment by moment, when the designer is sitting there working on her design, with pencil, clay, balsa wood, or whatever, in hand. On this level, the question becomes: Why do designers work out their designs physically, in the world, when the cognitive theories we have say that design should be done in the head? The starting point here is that conventional wisdom in cognitive science holds mental simulation, planning, etc. to be vastly superior to physically working on a problem, because it allows you to make predictions, test alternatives, and so forth. So why do designers not do what cognitive scientists say they should?

The answer I present in this chapter says that cognition is not an activity going on inside the mind, but an *interactive* process between mind and world. I present a theory that I tentatively call *interactive cognition* which is an effort to explain why designers do what they do, on this "action" level. My main point is to show that an interactive cognition has important advantages over following intramental principles, being both simpler and more effective at the same time.

Sketching

A prototypical design activity to explain in this way is *sketching*: It takes place in a simple setting and with only very simple tools; with the designer sitting at her desk with paper and pencil. As it turns out,

there is little need for more sophisticated technology, because there is little emphasis on the resulting sketch as such—sketching should not be mistaken for the process that produces the final drawings, the detailed and carefully produced drawings that will be the final result of the designer's work. The aesthetic impression of sketches is therefore often far from the sophistication and elegance of the drawings we usually associate with design; in this respect the one on the title page (p. ii) is unrepresentatively beautiful (also see figures 6.1–6.3, & Black 1990). Robbins 1994 contains drawings of the latter kind.

Instead, sketching is the process by which the designer works on her problem, and as such it serves several purposes. She sketches to understand her design problem and what it requires of her, to explore its particular circumstances and problems that must be tackled, to experiment with different approaches to a solution, and to eventually work out her final design, among other purposes. Therefore, the emphasis is on the process of sketching in itself rather than on the product, and for the process paper and pencil are highly suitable, and in their own respect quite sophisticated tools.

Beside sketching, there is a whole ecology of design techniques; this topic is covered in chapter 6. Still, sketching with pen and paper is widely regarded as a characteristic design activity, often even as the very essence of what design is about (e.g. Schön 1983, 1987). This is partly because it is so ubiquitous and typical of design in its various forms; architecture, industrial and graphic design, and so on, and partly since it is representative also of the other techniques: Architects use models of building sites and buildings, as for example Petra had done in Schön's protocol which I will discuss again below. Prototypes of the developing product are ubiquitous among industrial designers. The various techniques exist because their differences make them variously suitable for different types of task: While paper and pen are useful for designing floor plans, other media express tactile and three-dimensional qualities better, for example.

Quist's demonstration of sketching

To illustrate what designers' sketching is like, I will return to the episode with Quist and Petra from chapter 3 (Schön 1983), where I analyzed the inquiring structure of Quist's actions. Petra had set her problem as trying to fit the school building into the slope of the hill, but hadn't been able to solve it. Quist instead reframed the problem as imposing a geometry of his own, to bring "discipline" to the site:

Q: You should begin with a discipline, even if it is arbitrary, because the site is so screwy—you can always break it open later.

After having done this, Quist starts to develop a solution to the problem he has set. He does so by starting to sketch over Petra's drawings to work out the consequences of his framing, at the same time describing to Petra what he is doing:

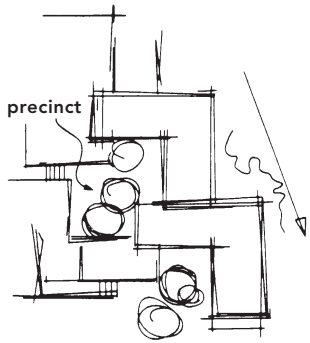

Q: Now in this direction, that being the gully and that the hill, that could then be the bridge, which might generate an upper level which could drop down two ways.

We get a total differential potential here from one end of classroom to far end of the other. There is 15 feet max, right?—so we could have as much as 5-foot intervals, which for a kid is maximum height, right? The section through here could be one of nooks in here and the differentiation between the unit and this would be at two levels.

Now you would give preference to that as a precinct which opens out into here and into here and then, of course, we'd have a wall—on the inside there could be a wall or steps to relate in downward. Well, that either happens here or here, and you'll have to investigate which way it should or can go. If it happens this way, the gallery is northwards—but I think the gallery might be a kind of garden—a sort of soft back area to these.

The kindergarten might go over here—which might indicate that the administration over here—just sort of like what you have here—then this works slightly with the contours—then you might carry the gallery level through and look down into here—which is nice.

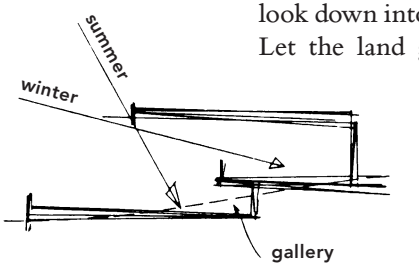

Let the land generate some sub-ideas here, which could be very nice. Maybe the cafeteria needn't be such a formal function—maybe it could come into here to get summer sun here and winter here.

(Also cf. the figures in chapter 3; unfortunately only redrawn versions of the figures are provided with Schön's texts.) When Quist

here arranges the L-units on a line down along the slope, he determines the locations of these units in a manner that is typical of sketching. According to an intramental account, he would conceive their placement first, and then afterward make a sketch to document his decision. But the protocol does not bear this out. Instead of following such a distinct and simple pattern—first thinking and then drawing—Quist's actions make up a more intricate structure.

Schön (*ibid.*) called this pattern *moving–seeing–moving*. From Quist's descriptions of what he draws, we can see that his solution develops as he is drawing. When he starts to draw, he makes a first "move", imposing the geometry of the three "L":s on the slope. In this case, the move originates in his framing. This move then allows him to "see" or visually appreciate the consequences of his move on paper, and hence of his reframing. Then this appreciation informs the next move, and so on. In this way he continues to work out his ideas by sketching, step by step, where each step suggests what the next step might be. The cycles of seeing and moving repeat, they are incremental by nature—hence the term moving–seeing–moving.

The way in which Quist verbally connects his descriptions of what he is doing also shows that his own process of drawing makes him think and have new ideas: Expressions like *that could then be...*, *which might generate...*, and *which might indicate that...* attest to his stepwise reasoning-by-drawing. He also says *There is 15 feet max, right?—so we could...*, and so on. These connecting phrases indicate that he is sketching to work out a solution, not merely to record his progress.

What he draws is then clearly not just the "output" of something he has already conceived in his mind; his words are not an after-the-fact report of something he has already thought out. The increments instead indicate that his reasoning takes place as he is drawing. He is using the drawing and the seeing as the basis for his next move, using visual feedback instead of trying to visualize each step in his mind's eye, and he uses physical drawing of concrete solutions instead of abstract reasoning about requirements and constraints.

The protocol thereby gives no grounds for making a clear distinction between thinking and drawing. The intricate pattern of Quist's activity gives little justification for treating them as separate activities, but rather as two aspects of one single activity. Thinking and sketching go on in parallel and mutually enable one another to move forward. Drawing enables thinking to proceed and vice versa. It is the entire physical activity of sketching, not merely thinking but the

whole *inquiry*, which serves the function that we usually call reasoning. This is the means by which Quist develops a solution, and a *function* that we generally regard as cognitive. This function does not become less cognitive by involving other elements than pure thought.

When he comments that his solution "works slightly with the contours", he is referring back to his original problem-setting. In Schön's words, he has been conducting a drawing experiment, and this comment shows that he is at least partly satisfied with the outcome of his experiment.

Quist also instructs Petra to use drawing experiments in the same manner. He tells her to "investigate which way it should or can go". He says that she "might carry the gallery level through", implying that she will have to try it to find out.

The aspects to make sense of

In summary, there are a number of points about sketching that we need to make sense of. One main question is why designers work in this way; my answer is that it provides definitive advantages over the intramental style. This translates into some more specific questions:

- What does sketching do that you can't do intramentally?
- What are the advantages; how can there be advantages?
- How do activities and materials contribute to cognition?
- How can they contribute; how can they have a cognitive function? What does it mean that they contribute?

A second problem concerns the structure of sketching activity, and why it looks like it does:

- How can sketching have the highly integrated role that it seems to have in the *cognitive* process of developing a solution?
- How can you explain that there seems to be one single process, sketching, going on; not two, sketching and thinking?
- What are the functions of the stepwise dialectical and moving–seeing–moving structures?
- What are the functions of drawing experiments and "investigations"?

As it happens, these features of the working process contain the seeds for answering the above questions.

4.2 Conversation vs. writing

Before going into the theory in full detail, I want to present a preliminary idea of what interactive cognition is, and how it can have important advantages over intramental cognition. I will do so by making an analogy with two ways in which people use language: spoken conversation, and written communication via an intermediating text.

In this analogy, interactive cognition corresponds to ordinary, everyday conversation, where both speaker and recipient are co-present when the communication takes place—this corresponds to an interactive cognition (speaker) having access to the world ("recipient"). The intramental model corresponds to writing. There, the writer (intramental cognition) typically produces her message without having the eventual reader (world) co-present with her at the time of writing. (And symmetrically, the person writing is not available to the reader when she eventually gets the message.) This reflects that in the intramental model, cognition operates as if it were completely isolated and remote from the world and thus its object of concern, even in those cases when it is readily available and close at hand, physically speaking.

There are two reasons why I use this analogy here, and will keep returning to it in what follows: The first is that ordinary language use is a domain that everyone can relate to, regardless of whether they are familiar with linguistics or not. The second is that conversation is practically the only domain where interaction has been studied, so that there is some scientific work that I can draw upon. However, as this is merely an analogy, I have taken care to avoid the aspects of conversation that have no correspondence in interactive cognition, e.g. the other party having a cognitive capacity independent of the speaker, and her ability to take an active role in the interaction.

As it happens, the traditional model of speech production is based on the model rational action from chapter 1, which is intramental in itself. Thus, speech in this view consists in output, controlled by a plan, which is derived from an intention. Hence, no feedback or influence from the recipient goes into speech production. This shows the connection between writing and intramental cognition, where

a model already exists; it then falls on myself to provide a corresponding model of interaction/conversation.

In an argument similar to the present one, Clark & Wilkes-Gibbs (1986) pointed out that the writing-based model doesn't apply to ordinary spoken language, since it works as if the recipient weren't even there. Thus, it assumes that also co-present conversation works as if it were writing, where no reader is available. The authors therefore called this the "literary model" of language production.

The same thing can be said about intramental theories of cognition: They make the same assumptions about cognition: it works as though the world weren't there, even when it is:

> For the crucial activities, at least of human problem solving of any complexity, ... take place centrally. *This is true even when the desired object and the required activity are physical*.
>
> (Newell & Simon 1972, p. 72, my italics)

Here I will be making an argument similar to that of Clark & Wilkes-Gibbs: Cognition does not work according to the "literary"model if it doesn't have to, just as conversation does not.

For the sake of comparison, consider what a person has to do to communicate: Her problem is to make sure that her message will make sense to the addressee. The problem for a *writer* is that her eventual addressee will be at a remote time and place when she reads the message. Because the writer is *separated* from her addressee in this manner, she can neither find out, nor ensure whether the text will indeed be understood when it is read.

Often, the writer does not even know *who* the eventual reader will be. Therefore she must try to make a prediction of who her audience will be, and adapt her message to that prediction instead. Moreover, even if she does know who her eventual reader is, she still cannot predict very accurately what will make sense to her and what will not. So the writer can neither find out nor know for sure that her text will have its desired effect.

For a *speaker* the situation is entirely different, because the listener is present with her right then and there, as she is producing her message. Thereby she can communicate the same message in a completely different manner. First of all, the speaker can get feedback directly from her addressee. Thereby she *can* find out whether what she says makes sense or not. Thus, she can also know for sure that her mes-

sage works as it should, and so she doesn't need to make any predictions as the writer does.

On the one hand, the conditions for written communication here are analogous to the basic structure of intramental cognition: Instead of operating directly on the world, it uses a representation as a stand-in. The representation is a substitute like the writer's estimate of her reader, and for the same reason it adds a layer of uncertainty. On the other hand, just like a speaker, an interactive cognition instead uses feedback from the source by interacting directly with the world. Thereby it attains higher certainty with less work.

Intramental cognition can of course use such direct feedback, too; that is, without having to make predictions and estimates. But to do so requires the same capabilities as interactive cognition uses, in addition to the representational system—which then just adds extra effort without any gain. It becomes a superfluous intermediate step, standing between individual and world. Hence it is a literary model of cognition: It treats the environment as if it were not directly available even when it is, just as the literary model of speech works as if the listener were not there even when she is.

I would like to stress that the relevant distinction is not between written and spoken language *per se*, but whether or not the speaker/writer has the recipient directly available to her. There are situations where speakers are remote from their addressees, and others where writers do have their audience with them. Still, the prototypical circumstances for speaking and writing capture the relevant distinction in an intuitive fashion.

In addition to making predictions, a writer must deal with her uncertainty by making *compensations* in the message. Typically, this compensation amounts to making background information and context explicit in the written text. This will make her text longer, typically a great deal longer than it strictly would have to be. This may be understandable if we consider that the writer cannot after the fact make up for misunderstandings or problems that occur during reading. Such problems will have to be averted in advance, in the text itself, but still without guarantees of success.

In the literary model of cognition, these compensations correspond to the elaborate plans and action sequences that an intramental cognition must prepare, since it has to provide for all eventualities that may occur—and then some more, to be on the safe side.

Because a speaker has her listener available right in front of her,

she doesn't stop at passively listening to whatever feedback the listener will give her; there is also a second purpose with her speech. Beside the obvious function of saying what she wants to say, she can also use her speech to *actively inquire into* her listener's understanding. She can actively make the listener give her the feedback she wants. This way she can determine much more precisely and directly whether her message works the way she wants it to, getting precisely the information she needs, instead of being content with whatever response the listener gives her. Asking questions is the most obvious way of doing this, but as we will see there are other ways that are much more sophisticated.

In written communication, the different activities are clearly grouped into separate phases, happening at different points in time: first writing, then reading. Conversation however does not consist of such phases, where first the speaker's questions explore the addressee's background knowledge, then followed first by the listener's answers, and then the speaker giving the necessary background information that the addressee lacks, before finally delivering the message itself. Conversation is not divided into such large separate chunks, with different roles for the participants in each of them, for instance with regard to who speaks and who listens. The power of everyday, informal conversation lies precisely in that it *isn't* restricted by the separation that makes writing problematic. The ability to alternate back and forth is what makes the compensatory measures of writing literally redundant in conversation. And the more often the speaker gets feedback, the less must she be concerned with getting her speech right the first time, since she will know right away if there is a problem. And if the listener can interrupt at any time, then less time will be wasted on talk that doesn't make sense to her.

Therefore, in conversation the participants' roles are much more democratically assigned than in written communication, with frequent chances for each party to take the floor throughout the conversation, and with shifts between speakers occurring frequently (Sacks, Schegloff & Jefferson 1974). In conversation, parties tend to speak in small, brief contributions, and the floor alternates between speakers frequently in an intricate yet smoothly coordinated web of contributions from several participants (*ibid.*).

Finally, there is no need for a speaker to anticipate and try to prevent any trouble that an addressee might have, since she can give immediate feedback if it happens, and the speaker can then make the

necessary adjustments accordingly. Instead of starting by providing a great amount of background, a speaker can therefore be radically concise and to the point, backing up and becoming elaborate only if and when she has to, and even then no more than necessary. For example, in the following passage from a casual dinner conversation, a very terse question only requires a small elaboration to be understood properly (Tannen 1984):

A: Do you *read?*
(1.0)
B: Do I *read?*
(0.5)
A: Do you read *books?*

The result of the interactive relation between participants in conversation is that communication becomes much more efficient. There is no need for supplanting each phrase with an introductory lecture, since it is straightforward and simple to determine what needs to be said. A speaker can find out whether she is making herself understood, and also make sure and be certain that she does, neither of which a writer can. And all of this with fewer words than a writer would need. For these reasons, conversation can achieve a higher quality than can writing, all the while doing less work, using fewer words, and in shorter time.

The comparison I have presented here implies that having cognition connect directly with the world, without intermediary representations of it, leads to greater certainty. In addition to that, it also suggests that action can have a second purpose beside the usual one of producing the desired result, just like talk in conversation: Manipulating the environment will give cognition richer and more relevant feedback. This makes a successful outcome not only more likely, but also easier to reach.

Further, the separation of intramental cognition from input/perception and output/action would in interactive cognition be replaced by fine-grained, intertwined interaction. In this way, the activity can evolve in close coordination with the world, obviating redundancy and making it lean and adapted to the specific circumstances. Separated, "literary" cognition, not being able to get any feedback, must instead produce the elaborate, redundant, better-safe-than-sorry action schemes that are necessary to provide for all eventualities. 'Interactive' actions will not have to be carefully crafted from the beginning, since they can swiftly adapt to feedback.

So in summary, the efficiency of conversation translates into promising prospects for a similarly organized, interactive cognition. It would enable an individual to act with more precision and greater success, while at the same time using considerably less effort, and doing so in fewer words. Without having seen that this is possible, and does happen in the case of conversation vs. writing, it would seem impossible for interactive cognition to improve quality and decrease effort at the same time compared to the traditional, intramental, mode of operation.

The advantages of interaction in the properties of conversation translate into the corresponding cognitive organization that I will call interactive cognition. In order to work out the details of this scheme, we first need to look at the concept of interaction as such.

4.3 Preliminaries: The meaning of 'interactive'

According to a dictionary, interaction means "mutual or reciprocal action or influence". *Interactive cognition* is meant to indicate that mind, action, and world *mutually* determine an individual's doings, in interaction. Cognition, of course, effects changes upon the world. In conventional cognitive theory, it is the mind alone that determines what an individual does, in a simple causal relation. In interactive cognition, the cognizing individual on the one hand, and the world on the other, *reciprocally* influence each other. In other words, mind and world interactively determine each other, and in particular they interactively determine cognitive performance.

The essence of interaction is that both (or all) participants give and take; speak and listen; act and perceive. This mutual influence is what breaks up simple causal schemes. A second point is that the traffic back and forth comes in a frequent exchange of small and concentrated, effective pieces, rather than as a few monolithic chunks. It corresponds in part to the point about the use of less code and less background information in conversation. When feedback allows speakers to reduce redundancy, it follows naturally that more, smaller exchanges lead to less 'dead weight'. To put it differently, the frequent exchanges make the interacting parts more closely adapted to each other.

Another important point is that both parties make *crucial* contributions to the conversation, and in some sense on an equal level, so

that neither party is above or in control of the other. For example, this would be the case in conversation if there were a clear and rigid question–answer structure, where the speaker controls the direction of discourse, and the listener's contributions were subordinate to and regulated by the speaker's actions.

But there is also a derived, second sense of interaction, the one used for example when experimental psychologists speak of "interaction effects". With this, they mean that they have not been able to isolate a single, simple cause for a certain effect, but multiple ones that also influence or *interact* with each other, so that it is impossible to establish a simple causal relation.

In this case, interaction means approximately "a complex relation". This is the extended sense of interactive cognition, saying that the relation between cognition, action, and the world is not as simple as traditional theories have it. This relation cannot be reduced into a simple causal one, from cognition to action to the world. A linear relation is definitely too simplistic, but also a circular model is too restrictive. In interaction, transitions may come from anywhere and go anywhere, at any point. What really happens is determined by the contents and circumstances of each case, not by some general organizational scheme.

This complex relation also means that cognition and action cannot be reduced to two distinct phenomena. In my use, action retains more of its everyday sense, like "activity", or doing something, rather than mechanical motor behavior.

So at the core of interactive cognition is a process where cognition and action, or knowing and doing, are closely tied together, so as to realize the tight interaction between mind and world. The resulting process is so tightly integrated that it cannot be broken down into well-defined components with simple relations between them. Terms like "doing" and "knowing" can only emphasize and contrast particular aspects of this integrated whole, they do not correspond to distinct sub-elements.

Including action and world in cognitive explanations

Traditional cognitive theory is simple and therefore can explain simple things like experimental tasks and well-defined problems, but when more complicated cognitive domains are considered, typically more realistic tasks such as design, then a more complex relation between cognition and action (and so forth) is required for explaining

what is going on. In realistic cases, cognition, action, and the world interact with each other in intricate ways that cannot be made to fit into conventional explanations, because they sacrifice too much detail in achieving their simplicity, leaving out too many aspects of what they are supposed to explain.

As a consequence, intramental cognition is a simple theory that can explain simple things, but to explain realistic cognitive phenomena we need a theory where the world itself is included in the explanation. One reason why intramental theory can be simpler is that it needs to make no references to the world. Thereby, for example, both action and perception can be left out of its explanations of cognitive phenomena.

In-the-world explanations, on the other hand, are not restricted to entities in the mind only. The world is not replaced by a problem space including only carefully chosen aspects of the world, and where these selected aspects have also been carefully coded so that a simple search algorithm will reach the right solution.

Instead, the interactive explanation includes such aspects of the world as artifacts and their properties. Thereby it can explain more complex phenomena, but it also requires the theory to explain how the head and the world can work together. It is a theory of how the mind *interacts* with the world and the things therein, and how individual and world jointly determine cognitive activity.

Interaction is the best way of using the resources of both mind and world to their fullest, just as conversation can draw upon the resources of both speaker *and* listener, unlike writing. The difference is that the writer *cannot* involve the reader, while the 'literary' model of cognition has *chosen* to leave out the world, as it has been held as an advantage to keep the world separate from cognition:

> Perception and motor behavior are assumed to take place in additional processing systems *off stage*. Input arrives in working [memory], which thus acts as a *buffer* between the *unpredictable* stream of *environmental events* and the cognitive system. (Newell, Rosenbloom & Laird 1989, p. 117, my italics)