5. Interactive cognition

The theory to be presented in this chapter consists of four steps. They concern the advantages of interactive cognition over intramental theory. Each step can be seen as a layer which is made possible by the layers before it, capitalizing on them to successively add further advantages of the interactive mode:

- Firstly, the advantages of dealing directly with the world instead of a surrogate for it, as the conventional theories do.
- Secondly, the advantages added by action and interaction with the world.
- Thirdly, a fine-grained structure of interaction that maximizes the benefits of involving world and action.
- Fourthly, a set of "shortcuts" made possible by drawing on the specific conditions of a situation rather than the general information a surrogate can only provide.

Step 1: The rediscovery of the world

The first step is to give back to cognition the access to the external world that cognitive science revoked very early. Because it has since then been claimed that blocking out the world was done with good reason, I will also have to motivate why such access is useful. I will therefore argue why and how the world itself can be more useful to cognition than a copy of it. I apologize in advance that the points I will be making here will appear self-evident to many readers. Nonetheless, they have to be stated since they go counter to the conventional positions of cognitive science. The programmatic way in which traditional theory has not "forgotten", or "neglected", but explicitly *kept* the world out of cognitive theories, is reflected in the quotation in chapter 4 about the need to place a buffer around cognition to protect it from the "unpredictable environment".

The existence of mental representations is often motivated by their ability to work as substitutes or stand-ins, as mental models and in mental simulations. Mental models are attributed with capacities similar to computer simulations, enabling them to imitate a physical system or process by embodying laws, equations and principles that describe the workings of that system (Gentner & Stevens 1983, Johnson-Laird 1983):

> If the organism carries a "small-scale model" of external reality and of its own possible actions within its head, it is able to try out various alternatives, conclude which is the best of them, react to future situations before they arise, utilize the knowledge of past events in dealing with the present and future, and in every way to react in a much fuller, safer, and more competent manner to the emergencies which face it. (Craik 1943)

This has been the original official motivation for having cognition operate on a mental representation of the world, instead of directly on the world itself. In particular, this is held to enable cognition to replicate and simulate events in the environment that go beyond the here and now of its immediate surroundings.

Still, the usefulness, and use, of mental models has not been restricted to events that are not directly available to the individual. The classical cognitive science approach to visual perception, to mention but one example, has always had the objective of producing a viewer-independent, complete three-dimensional model of the visible environment (Johnson-Laird 1989, Marr 1982). The reason for this is straightforward: Intramental theories *require* mental representations to work; their using them is not a matter of preference. For example, problem solving requires that the environment be encoded in a problem space, to be able to deal with it at all.

So we have two approaches to keeping cognition informed about and in synch with the world around it: by running an intramental simulation that shadows the events in the environment, as the traditional view does for one reason or another, and by checking the world in itself on a regular basis, not using any intermediate at all, which is what the interactive view proposes. What is the difference between the two? If I claim that checking with the original is simpler and better, intuition is inclined to agree. But since the opposite has been claimed so vigorously, the difference must be examined more closely.

Consider, for the sake of comparison, two different methods in ship navigation for determining the position of one's own ship, where this is done by simulation and measurement, respectively. One is known as *dead reckoning*. It is based on inferring the position of the ship by taking a known position and adding to it the ship's movement from that point. In principle this is an easy and computationally elegant way of determining position, based on simple mathematics. The movement may be computed from the speed and course of the ship and the elapsed time.

In reality, however, a ship is affected by external factors that make dead reckoning very difficult: wind, sea current, and so forth. These could be measured and taken into account, of course, but in practice it cannot be done with sufficient accuracy. And there are other sources of error, e.g. in determining the ship's own course and speed.

The other method, *position fixing*, uses references to elements outside the ship to determine its position. This method has been documented at length by Hutchins (cf. Hutchins 1995). In the specific case he describes, the bearings toward three landmarks are determined and plotted as lines on a map, and the ship's position is where the three lines cross and make a (hopefully) small triangle. This method is repeated at an interval of usually three minutes.

The specific ships that Hutchins describes are helicopter carriers of the so-called "amphibious fleet" (1995, pp. 7–9, 21). Ships of this kind are navigated by position fixing, not dead reckoning, at least when the ship is near hazards (within 8 kilometers) and therefore must be closely coordinated with its surroundings. This in spite of all the technical instruments and sophisticated navigation technology available, and even though all the scientific knowledge accumulated to this day could be used for constructing a model for dead reckoning the ship's movements and position.

Why is dead reckoning not accurate enough, and clearly inferior to position fixing? The purpose of determining the ship's position is to locate it relative to its surroundings, in particular to potential hazards and to where it ought to be or go. The reason why dead reckoning is worse than position fixing is that it is done by simulation based on a model or representation of the actual situation. The problem with simulations and models is that the world doesn't allow itself to be replicated accurately enough. It is simply too complex to be modeled with any precision. When textbooks use the laws of physics to model physical events, this is always done for idealized situations, where many greatly simplifying assumptions are made. Real measurements don't give the calculated results. Even in "real" physics such simplifying assumptions are made to an extent that may surprise the uninitiated.

Hence, computations yield errors, and these deviations from the real thing accumulate with each successive step in the simulation.

Thus the error aggravates rapidly, since each new estimate is based on the previous one, which already was wrong, and so on. So even a good model will begin to drift after a few steps; this is why weather forecasts only work a few days ahead, and become less certain for each day forward. In position fixing this does not happen. First of all, the data come from the real thing, not a model that is approximate at best. This alone makes the fix much more dependable. Moreover, the deviation doesn't accumulate since each fix is determined separately and is not based on the previous ones.

So ship navigation has a bridgeful of sophisticated technology and the accumulated results of science since Euclid to back it up. Still dead reckoning is not accurate enough. The same thing holds for cognition. Intramental cognition dead-reckons the environment, while interactive cognition goes straight to the source, without buffers or models in between. And like position fixes using landmarks, and like a speaker who has her listener available in front of her, her information about the world is not based on an estimate of what the world ought to be like.

The reason why I bring this up is that navigation by dead reckoning faces the same problems as the classical view of how actions are selected, by planning that is. A classical mental plan consists of a sequence of steps, each associated with an action. Each time a step is taken, the world changes from the physical action associated with it, as well as for other reasons. Because actions are not performed as they are selected, they must be selected on the basis of a mental simulation which dead-reckons the state of the world at that point.

This method was used in the robot Shakey in the 60's, and with exactly these negative consequences, where the constant issue of trying to keep an internal representation in synch with the environment became a major problem, as discussed more recently in e.g. Suchman (1987) and Dennett (1991).

This is the consequence of determining actions on the basis of a mental simulation, and this is how classical intramental planning has to be done. Because when the specification of action is separated from the execution of that action, then the consequences and the context of the future action must be simulated.

For example, if a communicated message is prepared in advance, as usually happens with written text, this could be done by "planning", by simulating the addressee's thinking after each sentence, to see if and how she will understand it. It doesn't have to be text, it could also be prepared speech. The risk of drifting further and further off the dead-reckoned course would still remain, however.

Philosophically, dead reckoning goes back to logical deduction from premises, whereas checking with the world is what one might call the pragmatic technique. In fact, some claim that "dead" is derived from "deduced" reckoning. Philosophy has always been concerned with how to reason properly. From syllogisms to formal logic, the aim has been to establish rules for making valid arguments and conclusions. This is the essence of the concept of rationality: the ideal, perfect way of reasoning; ideal thinking if you wish. Philosophers have always wanted to establish how you know what is right; how you know when you are right. This is part of what Dewey called "the quest for certainty" (1929).

So how do you know what is right? For practical purposes, you can simply check with the world to find out if you are right, if that is what you want. But philosophers are almost by definition not interested in practical matters, but are instead usually concerned with matters that cannot be settled by looking at the world. Metaphysics, Plato's world of ideas, ontology (What really exists?), epistemology (How can we know? What is knowledge? What is the relation between knowledge and the world?), What is truth? All of these are topics where the world can give us no answers. Other means are required, such as the principle of *reductio ad absurdum*: Everything that is contradictory and "logically impossible" must be false; something must be true if the opposite leads to a contradiction. In typical philosophical matters, principles like these are the *only* way of finding out, and the mother of all such principles is logical deduction: If something is certain, then other certainties can be deduced from it, step by step, each being perfectly logical, literally.

So logic makes perfect sense in the immaterial domains that philosophy is concerned with. The problem came when other sciences applied the same means to worldly ends. In cognitive science, it could be made to work for idealized domains like games and puzzles. These are also cases where you cannot check with the world to find out what is correct; you have to know the immaterial rules (cf. Zhang 1992, 1994) that are specified by a human and very similar to the rules of logical deduction.

But when applied to real cases, the problem of deduction or dead reckoning becomes a problem of constructing a model—a simulation—of the situation. This is why no dead reckoning model is exact or even good enough, and thus why logical deduction in immaterial domains does not transfer well to material ones—deduction works well in theory but not under authentic conditions. And if even the best model isn't good enough, then imagine how a model of a design problem that is very much under development would perform.

Instead real action must be grounded in feedback. This is what the three-minute fix cycle does. It is in this way analogous to moving–seeing–moving (sailing–fixing–sailing). By relying on feedback, all that remains background becomes harmless; every action is evaluated from its actual effects, not estimates. This is why designers, like Quist above, draw so that they can use their seeing to judge their own ideas, instead of trying to imagine what the consequence of a move is. They frequently see unintended consequences of their moves, and often these are desirable (Schön 1983). By appreciating the consequences of each step, an action sequence develops bit by bit as each step is performed, not in advance, and continual feedback from the world is used to stay on course. This is the pragmatic manner, which is synonymous with inquiry.

The pragmatic manner is a very simple way to find out. Logical deduction, dead reckoning, and so on are techniques that enable you to figure things out when the basic, simple way of finding out by checking is not available, as in metaphysics and so on, just as a writer has to compensate for not having her addressee in front of her. The rational ideal has made the mistake of regarding deduction from premises as the fundamental procedure for finding things out, not a compensatory technique for circumstances beyond the ordinary. It is thus like a literary model, in having been applied also to situations where the special, limiting conditions that motivate it do not apply.

Step 2: Manipulating the world—doing for the sake of knowing

The problems with a pure analysis phase have already been discussed, but the conversation analogy can give an additional angle on this issue: When the reader isn't there at the time of writing, the writer definitely can't ask a question and expect any answer. So analyzing the problem before writing the text seems rational; under these circumstances even the division into separate phases seems to make sense: first think, then write; first analyze the recipient, then design the outline, then write the text.

But when a speaker has her listener available right then and there

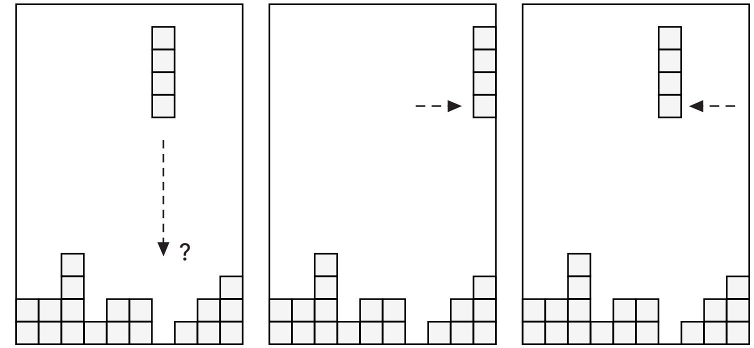

Figure 5.1 Is the Tetris brick lined up correctly? Finding out, and then ensuring that it is, by first moving it three steps to the wall and then three steps back.

as she speaks, why should she be content with passively using her perception for input, and use speech only to output her message? Then she would have to do with whatever feedback the listener was kind enough to give her. The second step in the interactive model concerns employing action for a cognitive purpose, analogous to the use of speech in conversation. The speaker producing the message corresponds to the first or productive purpose of action, from chapter 3—she can sign on speech to also induce feedback and to direct what kind of feedback she gets and when—this corresponds to the second, cognitive, or inquiring purpose.

It is a cognitive purpose since it contributes to performing the cognitive task; its effect achieves what a mental simulation would. If she gives speech this second purpose, then the feedback she gets will be richer, more to the point, and much more useful to her in designing her message.

When Kirsh & Maglio studied subjects playing the video game Tetris, they found that their subjects made moves—physical actions—that could only be explained as serving this second purpose (1992, 1994). In Tetris, bricks in different shapes fall down onto the playing field, and the objective of the game is to build the growing pile on the ground so that it does not fill the playing field all the way up, which will end the game (cf. figure 5.1). The way to do this is to fill the assembling horizontal rows completely; such a full row will disappear, and this is how the player is to keep the pile low. Bricks fall down one at a time, and the player can move the falling brick left and right and also rotate it, so as to fill the bottom pile evenly without leaving holes. Kirsh & Maglio were able to demonstrate that players make certain moves that actually bring the pieces *away* from their goal position. That is, these moves definitely do not serve the productive purpose of action, which is to reach the goal of placing the brick; they are even counterproductive in this sense. Instead these moves have a demonstrably cognitive purpose:

> ... certain cognitive and perceptual problems are more quickly, easily, and reliably solved by performing actions in the world rather than by performing computational actions in the head alone. We have found that some translations [i.e. left and right] and rotations are best understood as *using the world to improve cognition*. (Kirsh & Maglio 1992, my italics)

Their first example of such actions is when the player was to fill a gap say three steps from the wall. Instead of relying on a mental visualization to determine whether the falling brick is lined up correctly, players moved the brick to the wall—that is, away from where it will go—and then moved it back by pressing the proper key one, two, three times. In this way the player could ensure that the brick was over the right position, three steps from the wall (also figure 5.1).

In a second example, players rotated pieces that were not yet completely visible, to determine what kind of brick was coming. This enables them to decide earlier where it should go (figure 5.2). This action cannot have a productive purpose since it is made before the player has decided where to put it. Players also rotated pieces more often the more ambiguous they were (*ibid*.). Kirsh & Maglio also showed that rotating a piece on the screen is much faster than rotating it mentally, and that it is likely that pieces indeed are rotated so as to see where they would fit.

These are two examples of actions that have cognitive purposes, that is, where the individual interacts with the world to perform cognitive functions which traditionally have been attributed to purely intramental processes—in particular mental simulation and prediction. They follow the pattern of a speaker who evokes responses from her listener, and thereby uses speech as a kind of action that serves a cognitive purpose, and that makes intramental cognition both superfluous and inferior.

Dewey elegantly described this as “doing for the sake of know-

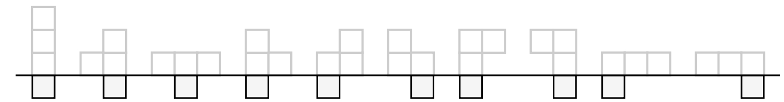

Figure 5.2 By rotating a falling brick that is not yet completely visible, you uncover hidden parts to faster determine its full shape.

ing”. As he also noted, this is a phenomenon that we can find everywhere in ordinary everyday activities, if we only look for them:

> The rudimentary prototype of experimental doing for the sake of knowing is found in ordinary procedures. When we are trying to make out the nature of a confused and unfamiliar object, we perform various acts with a view to establishing a new relationship to it, such as will bring to light qualities which will aid in understanding it. We turn it over, bring it into a better light, rattle and shake it, thump, push and press it, and so on. The object as it is experienced prior to the introduction of these changes baffles us; the intent of these acts is to make changes which will elicit some previously unperceived qualities, and by varying conditions of perception shake loose some property which as it stands blinds or misleads us. (Dewey 1929, p. 87)

It is also through action that our pragmatic knowing is used and put to a test, and it is in this way that we can see whether it serves its purpose or not.

Exploration

Inquiring action can be divided into two kinds: *exploration* and *experimentation*. I will begin with exploration. The quotation from Dewey captures its everyday meaning: “to make changes which will elicit some previously unperceived qualities, and by varying conditions of perception shake loose some property which as it stands blinds or misleads us”. This is why active manipulation betters pure analysis. If you are trying to understand an object, aspects that are not immediately apparent come out if you manipulate it; instead of just passively watching the object, you act upon it to see what happens, such as by rotating a falling Tetris brick to see its obscured parts. Exploration is a fundamental and very common aspect of inquiry, regardless of domain:

This is much of what an infant does when he explores the world around him, what an artist does when he juxtaposes colors to see what effect they make, and what a newcomer does when he wanders around a strange neighborhood. It is also what a scientist does when he first encounters and probes a strange substance to see how it will respond. (Schön 1983, p. 145)

When design specifications are in themselves very incomplete and leave much unsaid, so that just a pure analysis of them would be insufficient, that would correspond to a speaker who merely takes and uses whatever information the listener will give her. But to elaborate the given requirements, and go beyond what is obvious or explicit in them, designers do just like speakers—they use exploration to evoke what the specifications do not mention, and to make out what consequences follow from them. Guindon *et al.* have provided a detailed account of exploratory practice in software design: "[Subject] P8 explicitly acknowledges the need for exploring the problem environment to achieve a good understanding of the requirements before seeking a solution." (Guindon, Krasner & Curtis 1987, p. 69)

Adelson & Soloway (1985) noted that an experienced designer working in an otherwise familiar domain used exploration when he came to an unfamiliar part of the problem. Also Guindon showed that exploration is associated with understanding unfamiliar material. A good understanding allows for systematic work since it gives the designer a map to follow, while little experience requires her to explore without a sense of direction:

[Subject P6] clearly has better design schemas [i.e. understanding] for the communication sub-problem than for the scheduling problem. He successively refines his solution for the communication sub-problem while he performs much more exploratory design for the scheduling problem. By exploratory design, we mean design with many mental simulations of the problem environment and mental simulations of tentative solutions unguided by a plan. (Guindon *et al.* 1987, p. 68)

The authors state that the designer's main method of exploration is through *simulation* of the eventual context that the artifact will enter into. Simulation is prototypical as a physical design activity where the actions involved have an essentially cognitive, inquiring purpose. Guindon has described the use of simulations in great detail,

in her case in the design of an elevator control system (e.g. 1990b). All of her subjects simulated scenarios of how the elevators would be used. Firstly, these served to help them understand the requirements and the problem domain, and thereby to infer requirements and to generate solutions. Secondly, they used simulations to explore the solutions they were developing, to find inconsistencies, incompleteness, or bugs. In other words, to understand the solutions they had developed themselves. Still, both kinds concerned the same setting, only with or without having the design in place. Both kinds of simulation occurred throughout the design process (*ibid.*).

I'm going to imagine one elevator and a few scenarios. Say there's a request from floor 2 to 4. If there is a lift going to 2 on its way up, then stop the lift at 2, open the doors, ... If there is a lift going down from 5 to 1, the lift does not stop at 2 ... What if you press up at the floor, but once in the lift, you press a down button. ... So there's definitively the need for a queue of lift requests for each lift, separate from the floor requests. ... Maybe the floor requests could be handled by a completely separate system from the lift requests. (p. 287)

In this protocol excerpt, the designer imagines a couple of scenarios. By simulating the consequences of certain actions she works out what happens step by step, and this eventually leads her to discover a situation that was not immediately apparent:

By simulating a Lift scenario, the designer realizes that a user may press a floor button to go in one direction, but once inside the lift, may press a lift button to go in another direction. This test case was not mentioned in the problem statement, yet it is critical for the design of a good control algorithm. (p. 288)

Thus this single scenario developed her understanding of the design problem by singly generating a test case, a requirement, and a solution, at the same time. In another example the designer explicitly links her simulations to the need for understanding a certain aspect of the problem, namely scheduling (p. 286):

I'm not sure I understand about scheduling. I'll draw two elevators with a few floors. ... For each lift, I have, say, four buttons that are illuminated or not. And for each lift I also have to know the floor and the direction. Say Lift 1 is at floor 4 and there are requests to go down to floors 3 and 2. ... The floors don't move, the lifts move. It strikes me that I haven't considered enough this idea of having lifts between floors. I'm going to handle that.

Here, to understand the issue better, the designer creates a scenario that involves scheduling, and simulates what would happen. She elaborates one particular case to see what the system should do. In this way she discovers a test case for evaluating solutions, plus an additional requirement.

These simulations enable the designers to explore their problems and thus go beyond the limited information that is directly available to them. In this way they develop their understanding by tinkering with what she has, not merely analyzing the given specifications. Exploration is then a typical case of how action can have an expressly cognitive function.

Experimentation

Experimentation is more powerful than exploration. In fact, exploration can be seen as a limited version of experimentation that lacks certain elements of the "full" process. By making experiments you physically *test your ideas* in the world, instead of trying to figure out in your head what will happen.

If exploratory speech is when a speaker probes the other party by asking her questions, then experimentation in conversation is when a speaker is not just asking questions but actively tries to make her point, and uses the feedback she gets to find out whether her speech is working or not. She has some notion of what she wants to say and how to say it, and by actually saying it she conducts an experiment and can evaluate the outcome, that is, the listener's response. She can use the response she gets (the "consequences" of her talk) to see whether her idea works.

In design, experimentation is the main method for testing and working out ideas. Quist's sketching episode can be regarded as an experiment that he makes to work out his proposal and to test the consequences that come with it. After stating his framing he starts to draw to see where it leads and whether it will work. A string of moves draw out the possibilities contained in his idea, and he speaks and draws simultaneously: "...that could then be the bridge, which might generate an upper level ... we could have as much as 5-foot intervals ... The section through here could be one of nooks". Through this drawing experiment he tests his idea and is able to make something good out of it; it generates possibilities rather than problems. The experiment thereby "confirms" his framing in some sense.

As Quist then goes on to advise Petra on how to proceed, he instructs her to experiment to find her way, in the same way that he just did (Schön 1983):

Q: Well, that either happens here or here, and you'll have to investigate which way it should or can go.

That is, she should make experiments to test her ideas. A little later, he is also explicit about what designers' experimentation consists in—working out your ideas using paper and pencil, *trying them out* by drawing:

Q: Now the calibration of this becomes important. You just have to draw and draw and try out different grids.

Here Quist stresses the cognitive, i.e. non-productive, purpose of drawing. He emphasizes the process of drawing, not the product. She should "draw and draw", and "try out". This is drawing as an investigation; as an inquiry.

Another episode, even though Schön uses it for a different purpose, gives an even clearer illustration of how central experimentation is to design, and of how closely related it is to sketching. Here, Northover is the coach and Judith is the student presenting her problem (1987, pp. 127–132):

Judith: I haven't decided yet whether it's going to be sited right here or right here—I have the feeling it's going to be here and I'm going to make it level.

Northover: Do you have this to a larger scale somewhere?

When Judith describes her design, Northover asks for a scale drawing. He needs to *see* her placements to be able to judge them. Making such scale drawings is seen as an essential drawing experiment (*ibid.*, p. 127). The reason is that the relation between building and site is very important, as previously discussed. By drawing the building to scale, in its location and on a site plan, the designer can examine this relation in detail and work out a proper placement. Above, for example, Quist expresses this importance when he asked Petra if the drawing she has made is to scale, and she answered yes. But Judith answers:

J: Not right now, no. But it works as far as southern orientation —being far enough from here so I don't get drainage problems, being near enough to this flat area so I can set up playgrounds. ...

N: So you don't have it on a site plan at all!

J: No, that didn't seem necessary...

When Judith answers that she doesn't have a scale drawing, Northover expresses his astonishment over not only that she doesn't *have* one, but also since it implies that she hasn't even *made* one. She "feels" that this was not necessary, but to him it is crucial, since it means that she hasn't made the experiment of drawing the site plan to scale, which he sees as essential for evaluating her idea:

> Northover seems to be saying, "You are not designing at all. You are simply having 'ideas' and putting them down on paper. The moves you make have consequences that are testable, but you must draw to scale and in section in order to test them. The whole process of designing is lost to you because you will not do these things." (*ibid.*, p. 130)

Judith continues by describing "a ramp which spirals up". Northover then asks for a floor plan. Again, she says that she thought it was not necessary. She proposes to put "art and cafeteria" on the main level, and she asks him what he thinks. He answers, "That is possible, I guess", and asks about level changes and circulation. Judith expresses her wish: "Most people will use the ramp".

Again, the same clash over drawing to work out ideas: Northover neither approves of nor rejects her proposal—because without having the idea and its implications worked out on a drawing, he simply cannot evaluate it. When he says, "That is possible, I guess", he is really saying, "It might well be a good idea, again it might not; it is just that I cannot tell whether it is without having the idea worked out for me. With anything less I can only guess." He then goes on to give her some constructive advice:

N: I think you have got to really discipline yourself to draw it up to scale and draw a section through it—let's just assume that these ramps do work, that access—if so, this ramp will cut off the views to and from the library.

Here Northover tells her that the drawings that she thought unnecessary are crucial—she has to "discipline" herself to do them, because they are drawing experiments that she must make, to test whether her ramp will work or not. Drawing is the way designers test their conceptions, to see in detail what they lead to, and to develop them from mere ideas to reliable concepts that have been tried and confirmed. Judith does not see this function in drawings. She "decides where it's going to be sited", she "has the feeling it's going to be here" and she "is going to make it level".

I bring up this particular dialogue because it is as if Judith holds the conventional cognitive view of drawing, since she doesn't draw to *work out* the consequences of her ideas and decisions; she draws only to *document her choices* for others to see them. The use of drawing as experimentation is what the conventional view has overlooked. It has regarded drawings as a medium for recording the end product of design, or at the most as an extension of long-term memory (Akin 1986, Goel 1995, Newell & Simon 1972). In this view drawing is not regarded as experimentation but as output or storage, an epiphenomenon of pure thought. Drawing adds nothing to a problem-solving process that is purely intramental.

> Judith simply has no idea what Northover means by drawing, conceived as a process of trying out design moves and discovering their consequences and implications. ... it is clear that she sees drawing not as thought-experimenting but as a way of presenting ideas (Schön 1987, p. 130)

Again the conventional cognitive theories can be said to correspond to novices' beliefs and working styles, and the poorer results that this leads to, whereas the techniques of seasoned designers reflect the interactive approach.

Why are experiments (and simulation) in the physical world superior to models and simulations in the head? The reason is that you want to find out both what you *can* figure out and what you *can't* figure out, i.e. what you cannot simulate mentally. That is, you want to know also about the effects of your actions that you cannot predict or foresee.

Dewey's original purpose with his theory of inquiry was to explain the role of experimentation in science. It was the method of experimentation that had made physicists capable of their monumental advances ever since the Enlightenment. It was also the ancient Greeks' resorting to mere speculation that had made their progress in the natural sciences so marginal, especially in relation to their contributions in immaterial domains—philosophy and so forth—by the use of the same method. Indeed, thought experiments were the single method used in Aristotelian physics. (Or is it the other way round—that philosophy hasn't improved much since those days because of its reliance on such procedures?) In some sense, the non-pragmatist theories of knowledge and scientific reasoning thus remain at the stage of Aristotelian physics.

This shows the limitations of simulation in the head—for that is what a thought experiment is—compared to interactive experimentation in the world. If an idea is tested on a mental model, then the test reveals only those consequences of the idea that are accounted for *by the model*, i.e. the aspects that you have incorporated into it. The remainder is left out. This is why thought experiments never disconfirmed the Aristotelian idea that heavier objects fall faster than lighter ones. An experiment in the world can reveal to you both unanticipated consequences of your idea, and also limitations to your *model*, since surprising consequences indicate shortcomings in the model. Neither of these could come out of testing ideas in a purely intramental way.

This is just as position fixing does not depend on the navigator's understanding of winds, currents, and so on, and thereby is the superior navigational method. Dead reckoning, on the other hand, is only as good as the best available model of these physical phenomena, which obviously isn't good enough yet, and quite probably never will be. And however much it is improved, all the fuss involved can only aspire to eventually become *as good as* the method of simply looking out to see where you are.

Here, I am not trying to say that we are bad at mental models and mental simulation—but that these limitations give cognition good reason for not being intramental; for not using mental models and simulation in the first place. This holds for my argument in general: my purpose is to give cognition reason for not being intramental at all, not to show that it is intramental albeit badly so.

Step 3: Fine-grained interactive structure; economy and efficiency

Having concluded that checking with the world has advantages over intramental cognition, and that manipulating the world brings even further advantages, it is a logical extension that these means should be used in the best possible way. That is to say, if checking with the world is good, then check often and make use of the information you get as *much* and as *often* as possible. And if manipulating the world improves your knowing even further, then you ought to use this technique to its full potential, too. The result would be fine-grained pieces of activity, a continuous attention to feedback that replaces complex pre-planned actions, and simpler and smaller actions that both generate feedback and attend to and adjust to it.

Why is this better? Remember dead reckoning: it starts from a known position, but the shortcomings of prediction yield an accumulating error that makes the computed position deviate more and more from the actual position. Position fixing instead reestablishes accuracy each time a fix is made. So the more often you make fixes, the more often can you make the proper adjustments to your course, and the less will your measured position drift away from the actual one.

This leads to an organization of cognition that is radically different from when cognition is strictly separated from everything else, as happens in traditional theories. The difference in strategy is clear in the following brief example. Here, one person is trying to help another locate a building on a map (Brennan 1990, cited in Clark 1996):

> *no to your right, no over by the quad, right there yah right there*

There are a number of interesting points in this brief and deceptively simple example. First of all, it shows that in normal face-to-face conversation you *can get feedback* on whether what you say works as it should, and that this leads to higher certainty, allowing the speaker to say "yah right there".

This example also shows the predicted *fine-grained* pieces of activity. Instead of a precise description in a complex sentence with embedded clauses and elaborate structure, which would be characteristic of written language, there are several short and very simple segments, which is typical of ordinary conversation. The segments also display a *high sensitivity toward feedback*. Because the speaker continuously attends to visual feedback, she can formulate each segment on the basis of what the addressee is doing. Thus, the speech co-evolves in parallel with the addressee's activity. Thereby it is possible to make out what happened just from what the speaker said.

The speaker can achieve this high sensitivity because fine-grained action has *local control*. When action is specified in a large chunk before it is performed in an equally large chunk, there is no way to make use of feedback. Fine-grained segments, on the other hand, enable local control since they are specified not in advance but as they are performed. Thereby there is a way for feedback to enter into the process.

Conversation is known to have local control (Sacks *et al.* 1974). Among different kinds of spoken language, conversation is on the most-flexible end of the scale, it is completely locally managed, and comes in segments much like those of the present example. Between each segment there is a slot, or what Sacks *et al.* call a "transition-relevance place", where others may give feedback or take the floor, or the current speaker may continue and thereby make a multiple-segment turn, as the speaker did in this example. The size, content, and speaker allocation of each turn are thereby managed locally.

By this single and sole control mechanism, turns at talk are assigned to participants locally, i.e. one turn at a time. More organized scenarios such as debates may instead for instance assign equal shares and ordered turns to speakers in advance. Local management is done in interaction between participants. A speaker may or may not stop voluntarily at a transition point, and others may attempt to initiate a turn at such a point, or they may remain silent, thus encouraging the speaker to continue. For these reasons a locally controlled system is maximally flexible and adaptive to circumstances (*ibid.*).

With fine-grainedness and local control, action is specified—or *designed,* as Clark & Wilkes-Gibbs (1986) call it—as it is being carried out, so it can be altered on the basis of what actually happens. In the example, when the speaker "designs" her instructions she can take advantage of what the addressee is doing right then—where she first puts her finger, just how much too far in some direction she moves it, the exact moment when it passes over the right spot, and so on. Thereby there is no need for the redundancy that would be necessary for a description that is designed in a separate phase, prior to being performed. If the speaker couldn't get feedback, she would have to use an expression that does not depend on how the addressee will use it to locate the spot on the map. Typically, such a phase would

first have to refer by name to something that is easily found, then specify each next step from there, with name and direction, plus for example the number of streets to go past, and so on. The need for all this redundancy would lead to a complex expression that is both much longer and much harder to design—all of this made necessary by having separated the specification of the action from its performance.

The same advantage of conversation over writing has been demonstrated for definite references (Clark 1992, Clark & Wilkes-Gibbs 1986, Krauss & Weinheimer 1964, 1966). Over repeated trials, speakers use feedback to make references shorter, as in the following example (Clark & Wilkes-Gibbs 1986):

1. All right, the next one looks like a person who's ice skating, except they're sticking two arms out in front.
2. Um, the next one's the person ice skating that has two arms?
3. The fourth one is the person ice skating, with two arms.
4. The next one's the ice skater.
5. The fourth one's the ice skater.
6. The ice skater.

If the "grains" of talk are sentences as in this example, then clearly speech becomes fine-grained compared to the writing-like sentence 1: The long reference with several parts in different kinds of relations to each other—part–whole, (i.e. person–arms), activity/function (skating), location (in front of) and so on—is replaced by short and simple ones, both in structure and in syntax.

Also note how the speaker uses rising intonation in line 1 as a question-like prompt for feedback, at the point where she presumably would begin to try to make the initial reference more economical. Krauss & Weinheimer (1966) showed that speakers made their references shorter only when they received feedback: "By monitoring the listener's responses to his encoding, the speaker was able to decrease the number of words needed to code a given figure without running a great risk of being overly cryptic and confusing to the listener." (p. 344) That is, to be more concise without running a risk of not being explicit and elaborate enough.

The authors also examined feedback of two different kinds. When feedback was given concurrently ("mmm", "aha" and so on) there was a much greater difference than when the addressee gave feedback only after the reference had been completed. This supports the argument that when you can get feedback, then the more fine-grained interaction is, the better it works.

Incremental, approximating sequence

In Brennan's (1990) direction-giving example, the segments of talk are formulated on the basis of what is going on at the time when they are spoken. What is said in one segment has effects in the world, and following segments make use of these effects. For example, if the previous segments had not made the addressee move in the right direction, the segments "no over by the quad" and "right there" would not have worked. In other words, these segments make use of their particular context. So the specification of each segment is based on what happens in the world at the same time—speech and action co-evolve in parallel—and the inquiring effects of the spoken segments make the world a particularly valuable resource when the speaker formulates the segments that follow.

As the example shows, changing to a finer grain means that not only action, but also the *specification* of action is broken up into a sequence of smaller pieces. Instead of being planned separately and in advance, specification is done concurrently and together with performance, in a sequence of steps distributed over the whole course of action. The specification of each step builds on the outcome of previous steps, in an incremental fashion.

This fundamental change in procedure means that the method of specification becomes experimental; it becomes an inquiry. In the example from Brennan, the segments of speech serve as *experiments*: the speaker tries an instruction that might work, it has certain effects on the addressee (she moves the pointer). The speaker can then evaluate the outcome of this "experimental" instruction and adapt the following segments accordingly. This experimental procedure gives a role in inquiry to action, as doing for the sake of knowing. The fine-grained, incremental form of her speech is necessary for giving her spoken actions this second, inquiring role.

The traditional theories of action only consider its *productive* effects. The consequence is that according to this theory, actions are only specified so as to produce the desired result; to have a complete productive effect so that they "get it right". Their productive effects are their *only* effects. If specification is incremental, as I am proposing, its objective instead becomes to conduct an inquiry that produces a good outcome *in the end,* not at once. This gives action a different role which is not only productive. The purpose of action is no longer just to give the right result, especially not right away. Instead action is specified to also serve its inquiring purpose, as doing for the sake of knowing. It should manipulate the world to evoke feedback, and serve to test the knowing behind it. So cognitive effort should not mainly be spent on figuring out one large and complex action that produces the right result on the first attempt, but on performing an incremental sequence of simpler actions that concludes in a good result; the result will in this way be firmly grounded through the experimental nature of this procedure.

As a consequence, individual actions become *approximations* instead of perfect, once-and-for-all actions, because they are not specified so as to produce the desired result each on its own. Not approximation in the mathematical sense, but in the sense of being rough and unfinished while moving inquiry forward, because their inquiring effects enable upcoming actions to work better. Thereby the individual actions shouldn't be judged as being correct or wrong, but as parts of a larger sequence that leads to a successful result. That is, individual actions should be evaluated for how they work as parts of an inquiry, not from whether they produce the desired result at once. For example, "no to the right", shouldn't be classified as insufficiently specified, or as a "correction" of an earlier imperfect instruction, but as a concise and thereby efficient part of an incremental sequence. This sequence succeeds rapidly and with little energy, by using an interactive strategy that depends on entities in the external world for its success. In the example that I have been using, the "external entities" are the addressee, her actions, and the objects that she manipulates by these actions.

This approximating model makes one particularly important prediction: that the action that makes up the first attempt at something should be very different from how monolithic pre-specification would have it. There, the first action should be the only one necessary. It should be precise, well-conceived, and have only a productive purpose, that is, only serve to bring about the desired result. Here, in contrast, the productive effect of the first action can be minimal. Instead, the initial objective should be to get an inquiry off the ground. It should be a "starter", exactly what it does is not all-important, not too much cognitive effort should be spent on it. Getting the outcome right, the productive purpose, would become more important further on when the inquiring purpose is fulfilled. It has then worked out how to produce the desired result.

The previous example of direction-giving seems to bear this out. No segment appears to be the result of deep thought—their common purpose rather seems to be to make the other party do the work of finding the right spot! That is, to first get her to start searching, and then merely give her a push in the right direction when needed. The first segment, "no to your right", at first sight does not even look like a starting segment, but on closer consideration it makes perfect sense as such if only the addressee has her finger on the map, or otherwise appears to focus on a particular part of it. So if we only accept it as a first segment, then no doubt the thought behind it was minimal. Instead it fits perfectly into the description of a starter; its obvious motive is to make the addressee begin moving her finger. It is hardly an imperfectly constructed exhaustive description that is repaired afterward.

Clark & Wilkes-Gibbs (1986) also found what I call starters, as well as other inquiring techniques, to be characteristic of collaborative talk: Of the eight kinds of techniques for making references that they discuss, six have an inquiring function: "in three examples [speakers] deliberately drew the addresses into the process; and in three they began by knowingly issuing a questionable or inadequate noun phrase" (p. 113). The inquiring function is clearly present in these cases, sometimes even dominating—the authors themselves conclude that they "do no more than initiate the process" (p. 122). For example, speakers used *try markers* (rising intonation) to bring in the other party, compare with sentence number 2 listed above. These serve as attempts to start an interactive approximating process. They also spoke noun phrases in multiple segments, each inviting affirmation and whose continuation depends on the response (*ibid.*).

Step 4: Pragmatism enables specificity and shortcuts

Shortcuts in inquiring function

In exploration, the inquiring function is explicit—exploration is doing *only* for the sake of knowing. In conversation, this corresponds to a plain question. In experimentation the inquiring function is instead usually *implicit*, as part of an action that also has a productive function. In other words, such an action has both a productive and an inquiring function at the same time. This productive function is typical of experiments in practice, whereas it is usually absent in science (Schön 1983, pp. 147f). A designer, unlike a scientist, understands a situation and changes it at the same time, she is conducting a *productive* inquiry.

The direction-giving example above is a demonstration of the advantage of inquiry and experimentation over intramental specification. All the segments have a productive function: the speaker has an initial idea about how to direct the other, and each of these short instructions serves to produce that result. But each segment is also an experiment (cf. Schön 1983): The speaker's idea of what to do is her "hypothesis". Each segment of speech puts an aspect of her idea into effect by making the addressee move her finger (or hold it still in the final case), and so it tests her hypothesis as a scientific experiment would. It thereby spells out the consequences of the "hypothesis", and you can judge whether the hypothesis works or not. Moreover, it also shows *how* it works and how it does *not* work; each segment—each one a small experiment—lets the speaker develop and adapt her idea further, as the experiment advances her understanding of her situation.

Here we can see how specification and performance proceed together and in parallel: At no point does specification advance far ahead of production. At each step, specification builds on the outcome of the previous step, in an incremental fashion. The first segment adjusts for direction, in the second one the speaker decides to point out the quad, probably from the addressee's finger movements, and then the test that is implicit in "no by the quad" shows that she has succeeded; "yah right there".

This also shows the advantage of experiment over exploration. It is hard to imagine how the speaker could use exploration—"Where do you think it is?" or "What places do you know?"—it seems contrived and hardly very efficient. Questions (and exploration) are too vague, they do not test any specific idea about a solution. Neither do they have any productive effect.

Because actions are also implicit tests of the idea behind them, a designer doesn't have to be explicitly concerned with experimentation; with making tests or evaluating their outcome. Instead, she can rely on the *breakdown* mechanism (this concept originated with Heidegger 1927/1962). For this reason, as long as things go well, as long as there is no trouble, she can simply keep on doing what she does and focus on that without worrying about what might happen, and so forth. Breakdown only occurs when something goes wrong, only then does the designer become aware of the test.

This principle is also at work in conversation, where a speaker ordinarily does not explicitly test whether the audience is understanding what she says. Instead she simply speaks, giving *opportunities* for feedback, and its absence works to confirm that the audience is following her. In our direction-giving example for instance, the spaces between segments serve as such opportunities for indicating trouble that are not taken. Were this not the case, the addressee having trouble following the directions would signal this.

In this way, what a speaker says implicitly tests whether it is understandable, and the absence of feedback works as an implicit confirmation that the addressee is understanding what is said. She does not even notice any implicit tests, constantly thinking, "Since they're not protesting, what I'm saying must make sense to them. Good, then I can continue". Instead, of course, she is busy enough speaking.

For this reason, experimentation can be largely transparent even though it very effectively tests every action and reveals any problematic consequences. It is completely transparent as long as no troubles arise. This enables experimentation to be quite effortless: testing an idea can simply consist of attempting to carry it out. If this succeeds, then you have also conducted an implicit experiment that proves the idea to work, but you have at the same time produced the very result you wanted; experimental verification comes at no extra charge. And you experience your own activity as only being concerned with producing the result.

Perhaps the inquiring function of action has been overlooked precisely because tests are implicit in this way. We do not experience or intuitively see our own actions as tests, even though we experience breakdowns and adjust to them, which shows that actions do work as implicit tests. We do not recognize their inquiring function, only their effect on the world. This would then also be the reason why traditional theory of action similarly only considers its *productive* effects. The consequence is that according to such theory, actions are only specified (planned) so as to produce the desired result; to have a complete productive effect so that they "get it right".

Shortcuts in productive function

Because each attempt also works as an implicit test, speakers can go beyond the first level of laziness, which is to avoid working hard on specifying attempts, and reach a second level. This is what I will call *optimism* for want of a less intentionalistic term. Again consider these minimalist lines from a casual dinner conversation (Tannen 1984):

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

Even with this highly condensed speech, we may safely conclude that A is trying to ask B whether he reads fiction books in his spare time. By starting with an exceptionally brief "Do you *read?*", A is acting like an optimistic high jumper who chooses to enter into a competition at a very high level. If she can make the jump, she will have saved a great deal of effort by skipping all the jumps on lower levels that less optimistic competitors must spend effort on clearing. But if the optimistic high jumper doesn't clear her entry level, she will end up last on the scoreboard, registered as not having cleared any level at all. For an optimistic speaker, on the other hand, the situation looks much more promising: In Tannen's example, A's first optimistic attempt fails, but in the second one she can back down to a lower level and be more elaborate. And as we see, this time it works, even though the second attempt still must be regarded as being very optimistic, as it remains very terse.

But also B must be considered very optimistic. "Do I *read?*" does not reveal much about how B fails to understand what A said, and B doesn't try to make A's repair easier, by for example offering a interpretation (such as "Do you mean…?"). Still, A manages to repair this problem at once.

The advantage of optimism becomes clear if you consider the problem of adapting a message to the addressee's background knowledge, and of being explicit enough while not being redundant. As discussed earlier, a writer who is separated from her audience must provide ample context so that they will understand her, by adding redundancy to her core message. By being optimistic, a speaker skips a maximum amount of redundancy. But her optimistic attempt also works as an implicit test, and as an experiment with the addressee's understanding: If there is a problem, the addressee has several kinds of response available to indicate the problem source in her feedback (Schegloff, Jefferson & Sacks 1977). In this way an optimistic attempt is a powerful way to induce precise feedback which will become helpful if a second attempt is necessary. In the light of Garfinkel's (1967) demonstration that exhaustive descriptions of finite length cannot be made, optimism seems to be a particularly valuable way to determine a proper amount of context.

So, someone might object, it appears that optimism supports the concept of starters and incremental approximation. But couldn't it be that these techniques are used only when they are warranted, and that standard, exhaustive pre-specification still is the standard procedure, although speakers avoid using it when they can? No, because when a speaker appears to sense that her specification will be problematic, she doesn't spend more effort on specifying the first attempt. Instead she does the opposite, and makes the inquiring and incremental nature of speech explicit by marking her attempt as tentative, and by also encouraging the other party to collaborate (taken from Sacks & Schegloff 1979):

A: ...well I was the only one other then then the uhm tch *Fords?*
uh Mrs. Holmes Ford?
You know uh [the the cellist?
B: [Oh yes. She's she's the cellist

A: Yes. Well she and her husband were there.

(The bracket denotes both speakers talking simultaneously.) This example demonstrates almost every phenomenon mentioned so far in steps 3 and 4: A's hesitation ("then then the um tch") displays her trouble and marks her attempt as problematic. It is followed by her starter, *Fords?* This is clearly an optimistic attempt in the present sense, and it also has a rising intonation as a try marker requesting explicit confirmation. (This reverses the role of absent feedback. It comes to mean continuing *trouble*.) In the absence of feedback, A makes two more approximating attempts. Both use the same pattern as the first attempt: hesitation-marker–attempt–try-marker: "uh Mrs. Holmes Ford?" and "You know uh the the cellist?" Here, "Mrs. Holmes Ford" is a second attempt and a more explicit version of "Fords", which shows that this is an optimistic sequence.

So in this case, A seems to be aware of her problem in advance, but she doesn't spend more effort on conceiving her first attempt. On the contrary, a second look at this example even suggests that in the event of trouble, speakers in effect spend *less* effort on getting the first attempt right. This would then mean that they instead emphasize the starter's inquiring function, to invite the addressee to contribute to the interactive process. In other words, it seems as if speakers try to exploit the advantages of the interactive procedure to their fullest when they need it most, and that they therefore *emphasize* the inquiring and interactive aspects of conversation when they sense trouble. This tendency is even more pronounced in the following excerpt (Jefferson 1973, p. 59):

A: I heard you were at the beach yesterday. What's her name,
oh you know, the tall redhead that lives across the street
from Larry? The one who drove him to work the day his
car [was–
B: [Oh *Gina!*

A: Yeah Gina. She said she saw you at the beach yesterday.

Here, both "What's her name" and "oh you know" are minimally informative, so in a sense A appears to be maximally optimistic. It is as if she believes that B can tell who she is thinking of from the context alone. With "oh you know", she even repeats her appeal for help before trying herself. In whichever way a listener interprets what A says, her words serve to start an interactive procedure if not much else. So it seems that she prefers this interactive procedure to pre-specification.

So actions are approximate: they are not specified to be complete and perfect. They are also optimistic: they are specified with little cognitive effort and make shortcuts. Instead, the emphasis in action is on the span of the whole procedure, rather than its start. What about the end of an approximating sequence? When does it reach its conclusion? Also the end of an incremental sequence comes without much commotion. Generally speaking, it ends when it has reached its practical purpose. However, the end comes quietly because there is no explicit evaluation of success, just as tests are not explicit, neither their evaluations. Instead, following the same principle, success is the absence of breakdown. If you try to perform an action and you succeed, then you are finished—automatically and implicitly, without concluding so, or testing whether you are. The end of one action or action sequence is particularly implicit and invisible when it is immediately followed by another action, and this is the normal case under realistic circumstances.

As a result, actions are specified *viably,* not correctly or perfectly. Through attempts, more or less optimistic ones, and trouble that spurs further attempts, specification proceeds until it *works*. That is, until there is no more obstruction, so that nothing remains to be done. It stops because there is nothing more to do, not because an explicit evaluation function has been satisfied (or satisficed).

There are several points were we can find the viability principle at work. There is the phenomenon from conversation that you do not point out or repair a speaker's mistake if it doesn't present any problem; if you can figure out what she meant, or if it is not very important to the purpose at hand, then you simply do not object (Schegloff *et al.* 1977, p. 380). If it ain't broke, don't fix it.

Schober and Clark (1989) also found that speakers and addressees in cooperation settle on referential expressions that are not "correct" or objective, or even *intelligible* to a third person. Instead, they "exploit adventitious commonalities" and settle for the first perspective that makes sense to both. That is, they select the first viable candidate expression. If a third person cannot make sense of it, then she is, as they put it, out of luck. The resulting references are specific and local to them, not "objective" universal descriptions. If that had been the case, an outsider would have no trouble understanding it.

All shortcuts are made possible by viability principle

I have adopted the term *viability* from von Glasersfeld (1982). It is the same explanatory principle as in evolutionary theory and the principle of natural selection. Natural selection is not the survival of the fittest, it is the non-survival of the non-fit. In the same way, only the non-fit actions are improved on. It is not the correct and most fit actions that are prepared and then performed. With this I want to stress that the individual actions in incremental sequences do not approximate a correct action increasingly well, in the same way as evolution does not proceed toward the ideal creature.

Instead, as Bateson (1967) has stated, this kind of explanation is *negative*, because it is not the production of the effective, but the elimination of the ineffective. The principle of implicit tests and breakdown, which I introduced above, is of the same, negative kind.

Breakdown only occurs because further progress is impossible; actions that are good enough, even though still sub-optimal, pass without notice. So actions are viably specified even though there is no evaluation criterion; or more correctly *because* there is no evaluation. This is just as there is no selection mechanism in evolution either—evolution is only the non-survival of the non-fit.

Making the world "cooperate"

When a process is interactive, then it is almost by definition determined by all interacting forces. In Brennan's direction-giving example, even though the addressee doesn't say anything, if we want to explain why the speaker says what she does, it is clear that this depends to a great extent on factors apart from herself. When she tells the addressee, "no over by the quad", or "right there ya right there", these statements are appropriate because of external events happening at the same time. Because of the incremental nature of approximating sequences such as this one, the meaning of an utterance depends on external events, which in turn have been brought about by previous utterances. In this case, "right there" relates to a particular finger movement, which itself was the effect of "no to the left" and "no over by the quad", etc. Similarly, on the first "right there" the finger presumably stops, which makes the repetition "ya right there" appropriate.

Hence, the unfolding actions of each party are highly dependent on the actions of the other party, which in turn depend on your own previous actions. Each individual's actions, here her utterances, cannot have been determined only by her own mental processes, as traditional cognitive explanations would have it. That is to say, we cannot explain why the speaker said exactly this without including elements outside herself in the explanation. An exhaustive description produced in the manner of the literary model, on the other hand, would not have involved external factors in this way. So this interactive manner of specifying the location to the addressee gives these elements a role in determining what is actually said.

As a result, the "passive" external world is promoted to an important role in this process. It must thereby be recognized as having a fundamental and systematic influence on the direction of the process, an influence of the same order as that of the actor herself. Thereby the process is determined interactively by individual and world together. However, saying that the world is a part of cognitive processes appears very alien, both to the cognitive scientist and the nonprofessional.

Still, there has been a similar situation in the study of conversation. Both intuitively and scientifically, speech production has traditionally been thought of as one party, the speaker, producing output and the recipient merely receiving speech—the "literary" model of speech production, based upon how written language is produced. However, studies of talk have shown that speaking in conversation is a fundamentally interactive and collaborative process, i.e. where what is said is determined jointly by speaker and addressee in close cooperation (for example Clark & Wilkes-Gibbs 1986, Goodwin 1979, Sacks *et al.* 1974, Schegloff *et al.* 1977). And this is of course what I have been taking advantage of in comparing interaction to conversation, and intramental cognition to writing. By analogy, traditional cognitive theory is based on a "literary model of action", where actions are the output of the final step; the literary model of speech production is a direct application to language production of the intention-to-action schema, and the rational model of action as a whole (compare with chapter 1).

Pulling the world into the cognitive process is what makes the interactive and incremental strategy so effective. The brief directions ("by the quad", etc.) accomplish their purpose single-handedly, almost lazily, by offloading effort onto outside factors, doing with minimal effort and deliberation what would have required significant energy for traditional models to achieve. The small, simple segments of speech have an effect in the world, and it is these effects that allow the speaker to use much less effort in achieving her end. This is partly due to the transfer of work from one party to another, but also partly because, to an outside observer, this strategy decreases the *total amount* of work done by all involved parties, by eliminating the redundancy that is necessary to compensate for separation, as detailed above. Hence, it works more effectively not only for the speaker, but also as a whole.

The brief, deictic "to the right", "by the quad", and "right there" that are spoken would not work in the literary model. They are so brief because they can just "point" at circumstances in the situation that are available. In effect, these brief expressions specify the location jointly with the situation; they determine it in interaction with each other. The simultaneous presence of both—the absence of separation—allows the spoken specification to become *interdependent*

with the world it is referring to. Separation instead would require a written description to be independent of circumstances. It would have to be self-contained, because it couldn't involve elements that might not be present.

Steps 3 and 4 in relation to sketching

In steps 3 and 4 I have mainly based the presentation on examples from conversation. It can in one way be regarded as a closer cognitive analysis of Schön's (1983) concept of "design as conversation" with the working materials, which is reflected in the dialectical structure of sketching. And it is this dialectical and highly interactive structure of sketching that the analyses in steps 3 and 4 concern.

Step 3 demonstrates the value of the moving–seeing–moving structure of sketching, whose pattern adheres very closely to the fine-grained, interactive structure that is described there. Sketching is made up of very small and simple incremental steps, which yield local control and high sensitivity to feedback. This, in turn, makes sketching into a highly fluid and efficient process, which supports the open-ended and conceptual nature of the design work which sketching is typically used for.

The concepts from step 4, in particular optimism, improve on this by enabling the designer to move forward and test ideas very rapidly. By merely starting to work on a solution by sketching, she can make substantial shortcuts, and there is no need to think first, and then draw the solution. She can just start to work out her ideas, and simply back up and be more careful or elaborate if it turns out she has been too optimistic.

Moreover, the incremental approximating sequence elaborates on the developing dimension from chapter 3. This is the structure which results when you cannot separate the specification of an action from its performance. Instead both of these processes run in parallel throughout the process. The concept of starter serves to point out that in such a process, the quality of the first solution attempts should mainly be regarded as starting the process, which will eventually lead to a satisfactory solution.

In this chapter, I have begun to answer some of the questions that were posed in chapter 4, with a model where activity and working materials are true parts of the cognitive process in itself, for example. As seen in the following quotation, even though the importance of sketching has been recognized, a dichotomy is still made between sketching and thinking:

> 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)