Ethical know-how. Action, wisdom, and cognition

Preface

The original invitation to deliver these lectures came with the explicit — and initially surprising — request to venture into the territory of ethical thought. I quickly realized that an opportunity to present a whole spectrum of thoughts and personal explorations I had been pursuing for some time was a temptation too great to resist.

Let the reader be warned. I tackle this subject in the style I believe is most fruitful as a philosophical ecology for today, as I have done in my recent book The Embodied Mind (MIT Press 1991). I draw, on the one hand, from scientific work as a necessary reference point for philosophical reflection; on the other, I extend the philosophical horizon to include non-Western traditions. Even so, ethics is new terrain for me, and what I have to say here must be taken in the spirit of adventure more than anything else. Yet it is terrain I am eager to explore, because I firmly believe that an understanding of ethics in a nonmoralistic framework such as the one I propose is crucial for our confused and confusing modern world.

My thanks go first and foremost to the Fondazione Sigma Tau and Edizione Laterza for making these Italian Lectures possible, and to Lorena Preta and Pino Donghi for doing me the honor of inviting me to be among the very first lecturers. My sincere wish is that whatever there is of value in what I have to say may help to further the visionary cultural work they already have begun.

My thanks also to University di Bologna and to Paolo Fabbri for graciously hosting this event. A final thanks is due to a numerous, faithful, and mostly young

audience in Bologna that made the experience of delivering these lectures enormously rewarding.

The first lecture: Know-How and Know-What

Setting out the question

Ethics is closer to wisdom than to reason, closer to understanding what is good than to correctly adjudicating particular situations. I am not alone in thinking this, for it seems that nowadays the focus has moved away from meta- ethical issues to a much sharper debate between those who demand a detached, critical morality based on prescriptive principles and those who pursue an active and engaged ethics based on a tradition that identifies the good.

This can be seen as a reenactment of the classical opposition between morality and situatedness. On the side of morality, we have such eminent representatives of the Kantian tradition of moral judgment as Jürgen Habermas and John Rawls. On the side of situatedness, we have the descendants of Hegel, whose position is ably represented by philosophers like Charles Taylor, who clearly explains the differences between the two camps in his recent Sources of the Self:

Much contemporary moral philosophy, particularly but not only in the English-speaking world, has given such a narrow focus to morality that some of the crucial connections I want to draw here are incomprehensible in its terms. This moral philosophy has tended to focus on what is right to do rather than on what it is good to be, on defining the content of obligation rather than the nature of good life; and it has no conceptual place left for a notion of the good as the object of our love or allegiance or as the privileged focus of attention or will. [Note 1-1]

Although I draw heavily on recent contributions to this debate in the literatures of phenomenology and pragmatism, I find equally interesting the enormous body of thought about what it is good to be that comes from the three wisdom traditions of the East: Confucianism, Taoism, and Buddhism. In what follows I highlight these non-Western contributions and thus take a comparative view of ethical experience.

As a first approximation, let me say that a wise (or virtuous) person is one who knows what is good and spontaneously does it. It is this immediacy of perception and action which we want to examine critically. This approach stands in stark contrast to the usual way of investigating ethical behavior, which begins by analyzing the intentional content of an act and ends by evaluating the rationality of particular moral judgments.

Philosophers are not only ones who have lost sight of the distinction pointed at by Taylor. For instance, none other than the psychologist Jean Piaget in The Moral Judgement of the Child opens his discussion by saying that "it is the moral judgement that we propose to investigate, not moral behavior," only to conclude that "logic is the morality of thought just as morality is the logic of action… Pure reason is the arbiter both of theoretical reflection and daily practice" (404). [Note 1-2] But we have to ask ourselves: why should one conflate ethical behavior with judgment? Most people answer this question by repeating the received (Western) opinion on this matter, not by describing what they do in everyday life. This is crucial. Consider a normal day in the street. You are walking down the sidewalk thinking about what you need to say in an upcoming meeting and you hear the noise of an accident. You immediately see if you can help. You are in your office. The conversation is lively and a topic comes up that embarrasses your secretary. You immediately perceive that embarrassment and turn the conversation away from the topic with a humorous remark. Actions such as these do not spring from judgment and reasoning, but from an immediate coping with what is confronting us. We can only say we do such things because the situation brought forth the actions from us. And yet these are true ethical actions; in fact, in our daily, normal life they represent the most common kind of ethical behavior.

And yet the present tendency is to contrast this pervasive mode of being ethical with situations in which one experiences a central I performing deliberate, willed action. For example, I read a newspaper article about the devastating civil war in Yugoslavia and call a friend to join forces in a relief campaign for the victims. Or I learn that my child is having trouble in school and after pondering a course of action I solemnly decide to help him with his homework. In situations of this kind, we feel that the action is "ours.- We can explain what we do in terms of the goal we expect to achieve.

It is quite clear that one aspect of our moral and ethical behavior is grounded in such judgments and justifications, but we cannot and should not pass quickly over the first, more pervasive mode of ethical behavior as being merely '`reflexive." Why not start with an investigation of this pervasive mode and see whether it leads us into an understanding of the difference between know-how and know- what, between spontaneous coping and rational judgment? [Note 1-3]

We have now in front of us two interdependent questions that are central to these lectures:

How can one best understand ethical know-how?How does it develop and flourish in human beings?

Immediate coping in cognitive science

In order to pursue the first question we need first to ascertain how this neglect of the study of immediate coping manifests within the very sciences dedicated to the study of mind and knowing: the cognitive sciences, to which we now turn.

New Forms for Old Problems

"Rationalistic," "Cartesian," "objectivist": these are some terms used to characterize the dominant tradition of recent times. Yet when we reexamine our understanding of knowledge and cognition, I find that the best expression to use for our tradition is abstract: nothing characterizes better the units of knowledge that have been deemed most "natural." It is this tendency to find our way toward the rarefied atmosphere of the general and the formal, the logical and the well defined, the represented and the foreseen, which characterizes our Western world.

However, there are strong indications that within the loose federation of sciences dealing with knowledge and cognition — the cognitive sciences — the conviction is slowly growing that this picture is upside down and that a radical paradigm shift is imminent. At the very center of this emerging view is the conviction that the proper units of knowledge are primarily concrete, embodied, incorporated, lived; that knowledge is about situatedness; and that the uniqueness of knowledge, its historicity and context, is not a "noise" concealing an abstract configuration in its true essence. The concrete is not a step toward something else: it is both where we are and how we get to where we will be.

Perhaps nothing illustrates better this tendency than the gradual transformation of ideas in the very pragmatic field of artificial intelligence. In its first three decades (1950–1980) research in artificial intelligence (and cognitive science in general) was based entirely on the computationalist paradigm, according to which knowledge is a manipulation of symbols by logic-like rules, an idea that finds its fullest expression in modern digital computers. Initially, researchers in artificial intelligence concentrated on solving the most general problems, such as natural language translation or devising a "general problem solver." These attempts, which tried to match the intelligence of a highly trained expert, were seen as tackling the interesting, hard issues of cognition. After years of consistent and persistent failure, researchers began to seek out more modest challenges. Soon it became clear that even the most ordinary tasks, even tasks performed by tiny insects, lie beyond the grasp of a computational strategy. Early optimism has given way to the recent and growing conviction that artificial intelligence worthy of the name will not be achieved without first understanding the situated embodiments of simple acts. [Note 1-4]

The fact is that some researchers have always believed that cognition can only be understood in terms of how significance arises out of the autonomous totality that is the organism. Paradoxically, a good example is Jean Piaget, who has transformed this intuition into good research. By studying how children shape their world through sensorimotor actions, he has done nothing less than study how the constitution of a perceptual object is grounded in ontogeny.

Piaget successfully introduced the notion that cognition even at what seems to be its highest level — is grounded in the concrete activity of the whole organism, that is, in sensorimotor coupling. In short: the world is not something that is given to us but something we engage in by moving, touching, breathing, and eating. This is what I call cognition as enaction since enaction connotes this bringing forth by concrete handling.

Microworlds and Microidentities

Picture yourself walking down the street, perhaps going to meet somebody. It is the end of the day and there is nothing very special in your mind. You stop at a kiosk and buy a pack of cigarettes, then continue on your way. You are in a relaxed mood. You put your hand into your pocket and suddenly you discover that your wallet is missing. Your mood is shattered. Your thoughts are muddled. And before you know it, a new world has emerged. You see clearly that you left your wallet at the kiosk. Your mood shifts again to one of concern about losing your documents and your money. The only thing on your mind now is getting back to the store as quickly as possible. You ignore the surrounding trees and passers-by; all your attention is directed at avoiding further delays.

Situations like this are the very stuff of our lives, and they involve the most ordinary situations as well as the more interesting ethical stances. We always operate in some kind of immediacy of a given situation. Our lived world is so ready-at-hand that we have no deliberateness about what it is and how we inhabit it. When we sit at the table to eat with a relative or friend, the entire complex know-how of how to handle our utensils, how to sit, how to converse, is present without deliberation. We could say that our having lunch-self is transparent. [Note 1-5] You finish lunch, return to the office, and enter into a readiness that has its own mode of speaking, moving, and making assessments.

We have a readiness-for-action proper to every specific lived situation. Moreover, we are constantly moving from one readiness-for-action to another. Often these transitions or punctuations are slight and virtually imperceptible. Sometimes they are overwhelming, as when we experience a sudden shock or come face-to-face with unexpected danger.

I call any such readiness-for-action a microidentity and its corresponding lived situation a microworld. Thus, "who we are" at any moment cannot be divorced from what other things and who other people are to us. We could engage at this point in a bit of phenomenology and identify some typical microworlds within which we move during a normal day. The point, though, is not to catalog these microworlds but to notice their recurrence: the ability to take appropriate action is, in some important sense, how we embody a stream of recurrent microworld transitions. I am not saying that recurrence always applies. For example, when we arrive in a foreign country for the first time, we face it virtually "empty-handed." Many simple social interactions have to be done deliberately or learned outright. In other words, microworlds and microidentities are historically constituted. But in general, "who we are" — the pervasive mode of living — consists of already constituted microworlds.

When we leave the realm of human experience and shift to that of animals, the same kind of analysis applies as an external account. The extreme case is illustrative: biologists have known for some time that invertebrates have a rather small repertoire of behavior patterns. For example, a cockroach has only four fundamental modes of locomotion: standing, slow walking, fast walking, and running. Nevertheless, this basic repertoire makes it possible for these animals to navigate appropriately in any possible environment known on the planet, be it natural or artificial. A key question for the biologist is then: how does the animal

decide which motor action to take in a given circumstance? How does it select the appropriate behavior? How does it assess a given situation as requiring, say, running as opposed to slow walking?

In the two extreme cases — human experience during breakdowns, and simple animal behavior at moments of transition — we are confronted, in vastly different manners to be sure, with a common issue. At each such breakdown point, the manner in which the cognitive agent will next be constituted is neither simply determined nor simply planned. Instead, its constitution is a matter of the commonsensical emergence of an appropriate stance from the entire history of the agent's life. Once a behavioral stance is selected or a microworld is brought forth, we can more clearly analyze its mode of operation and its optimal strategy. In fact, the key to autonomy is that a living system finds its way into the next moment by acting appropriately out of its own resources. And it is the breakdowns, the hinges that articulate microworlds, that are the source of the autonomous and creative side of living cognition. Such common sense, then, needs to be examined on a microscale, for it is during breakdowns that the concrete is born.

Knowledge as Enaction

Let me now illuminate these issues further by explaining better what I mean by the word "embodied." Embodiment entails the following: (1) cognition dependent upon the kinds of experience that come from having a body with various sensorimotor capacities; and (2) individual sensorimotor capacities that are themselves embedded in a more encompassing biological and cultural context. These two points have already been touched on, but here I wish to explore further their corporeal specificity, to emphasize once again that sensory and motor processes, perception and action, are fundamentally inseparable in lived cognition, and not merely contingently linked as input/output pairs.

In order to make my ideas more precise, let me now give a preliminary formulation of what I mean by an enactive approach to cognition. [Note 1-6] In a nutshell, the enactive approach underscores the importance of two interrelated points: (f) that perception consists of perceptually guided action; and (2) that cognitive structures emerge from the recurrent sensorimotor patterns that enable action to be perceptually guided. These two points will become clearer as we proceed.

Let us begin with the notion of perceptually guided action. According to the dominant computationalist tradition, the point of departure for understanding perception is typically abstract: the information-processing problem of recovering pre-given properties of the world. According to the enactive approach, however, the point of departure for understanding perception is the study of how the perceiver guides his actions in local situations. Since these local situations constantly change as a result of the perceiver's activity, the reference point for understanding perception is no longer a pre-given, perceiver-independent world, but rather the sensorimotor structure of the cognitive agent, the way in which the nervous system links sensory and motor surfaces. It is this structure — the manner in which the perceiver is embodied — and not some pre-given world, that determines how the perceiver can act and be modulated by environmental events. Thus the overall concern of an enactive approach to perception is not to determine how some perceiver-independent world is to be recovered; it is, rather, to determine the common principles or lawful linkages between sensory and motor systems that explain how action can be perceptually guided in a perceiver-dependent world.

This central concern of the enactive position stands in contradistinction to the received view that perception is fundamentally the truthful reconstruction of a portion of the physical world through a registering of existing environmental information. In the enactive approach reality is not a given: it is perceiver-dependent, not because the perceiver "constructs" it as he or she pleases, but because what counts as a relevant world is inseparable from the structure of the perceiver.

A classical illustration of the perceptual guidance of action is the study of Richard Held and Alan Hein, who raised kittens in the dark and exposed them to light only under controlled conditions.[Note 1-7] A first group of animals was allowed to move around normally, but they were harnessed to a simple carriage and basket that contained the second group of animals. The two groups, therefore, shared the same visual experience, but the second group was entirely passive. When the animals were released after a few weeks of this treatment, the first group of kittens behaved normally, but those who had been carried around behaved as if they were blind: they bumped into objects and fell over edges. This beautiful study illustrates the — enactive — view that objects are seen not by the visual extraction of features, but rather by the visual guidance of action. Similar results have been obtained under various other circumstances and studied even at the level of the single cell.

If the reader feels that this example is fine for cats, but irrelevant for humans, let us consider another case. Paul Bach y Rita designed a video camera for blind persons that can stimulate multiple points on the skin by electrically activated vibration.[Note 1-8] Thus images formed with the camera were translated into patterned tactile sensations — with the following results. Patterns projected onto the skin had no "visual" content if the subject remained motionless. However, if the subject directed the camera by moving his head, hands, or body for a few hours, a remarkable transformation occurred. The tactile sensations became visual perceptions, the patterns of vibration on the skin were not felt but seen as images projected into the space being explored by the bodily directed "gaze" of the video camera. Thus in order to experience "real objects out there," it was enough for the person to actively direct the camera. This experience is an excellent example of the perceiver-dependent nature of what otherwise seems an internal representation of a perceiver-independent world of features.

From Sensorimotor Patterns to Cognitive Capacities

Let us now turn to the idea that the more familiar cognitive structures of human life emerge from the kinds of recurrent sensorimotor patterns that enable action to be perceptually guided. What we need to examine now is how this sensorimotor coupling can be linked with other kinds of typically human cognitive performance. Otherwise, we might be tempted to deny to the "low" level event of sensing and acting the significance that we grant to the "higher" cognitive levels.

In fact, the idea that the cognitive structures of human life emerge from recurrent sensorimotor patterns is at the very core of Piaget's program [Note 1-9] and has been argued for in recent works by cognitive linguists George Lakoff and Mark Johnson.[Note 1-10] We will present the idea of embodied cognitive structures with special reference to their work. Once again we must move out of the abstract and emphasize an experientialist approach to cognition. As Lakoff says, the central claim of their approach is that meaningful conceptual structures arise from two sources: (I) the structured nature of bodily experience; and (2) our capacity to project imaginatively from certain well-structured aspects of bodily and interactional experience to conceptual structures. Rational and abstract thought is the application of very general cognitive processes — focusing, scanning, superimposition, figure-ground reversal, and so on — to such structures.[Note 1-11] The basic idea is that embodied (sensorimotor) structures are the substance of experience, and that experiential structures "motivate" conceptual understanding and rational thought. Since I have emphasized that perception and action are embodied in self-organizing sensorimotor processes, it is natural to postulate that cognitive structures emerge from recurrent patterns of sensorimotor activity. In either case, the point is not that experience strictly determines conceptual structures and modes of thought; it is, rather, that experience both makes possible and constrains conceptual understanding across a multitude of cognitive domains.

Lakoff and Johnson provide numerous examples of cognitive structures that are generated from experiential processes. To review all of these examples here would take us too far afield. Let me discuss briefly just one of the most significant kinds: basic-level categories. Consider most of the middle-sized things with which we continually interact: tables, chairs, dogs, cats, forks, knives, cups, and so on. These things belong to a level of categorization that is intermediate between lower (subordinate) and higher (superordinate) levels. If we take a chair, for example, at the lower level it might belong to the category "rocking chair," whereas at the higher level it belongs to the category "furniture." Eleanor Rosch and her colleagues have shown that this intermediate level of categorization (table, chair, and so on) is psychologically the most fundamental or basic one.[Note 1-12] Among the reasons why these basic-level categories are considered to he psychologically fundamental are the following: (1) the basic level is the most general level where category members have similar overall perceived shapes; (2) it is the most general level where a person uses similar motor actions for interacting with category members; and (3) it is the level where clusters of correlated attributes are most apparent. It would seem, therefore, that what determines whether a category belongs to the basic level depends not on how things are arranged in some pre-given world, but rather on the sensorimotor structure of our bodies and the kinds of perceptually guided interactions this structure makes possible. Basic-level categories are both experiential and embodied. A similar argument can be made for image-schemas emerging from basic forms of sensorimotor activities and interactions.

Know-how and know-what reconsidered

Cognitive science is waking up to the full importance of the realization that perception does not consist in the recovery of a pre-given world, but rather in the perceptual guidance of action in a world that is inseparable from our sensorimotor capacities, and that "higher" cognitive structures also emerge from recurrent patterns of perceptually guided action. Thus cognition consists not of representations but of embodied action. Thus we can say that the world we know is not pre-given; it is, rather, enacted through our history of structural coupling, and the temporal hinges that articulate enaction are rooted in the number of alternative microworlds that are activated in every situation. These alternatives are the source of both common sense and creativity in cognition.

Thus it seems more and more compelling to look at knowledge — to understand understanding — in a manner that can only be called post-Cartesian: that is knowledge appears more and more as being built from small domains composed of microworlds and microidentities. Behavioral repertoires vary throughout the animal kingdom, but what all living cognitive beings seem to have in common is know- how constituted on the basis of the concrete. Thus what we call general and abstract are aggregates of readiness-for-action.

In other words, cognitive science is waking up to the simple fact that just being there, immediate coping, is far from simple or reflexive. Immediate coping is, in fact, the real "hard work," since it took the longest evolutionary time to develop. The ability to make intentional, rational analyses during breakdowns appeared only recently and very rapidly in evolutionary terms. (This point of view is also gaining ground in the related fields of modern robotics and artificial life research.);[Note 1-13]

My interest in immediate coping does not mean that I deny the importance of deliberation and analysis. My point is that it is important to understand the role and relevance of both cognitive modes. It is at the moments of breakdown, that is, when we are not experts of our microworld anymore, that we deliberate and analyze, that we become like beginners seeking to feel at ease with the task at hand.[Note 1-14] In this light one can say that computationalist cognitive science has been mostly concerned with the behavior of beginners and not with that of experts.

This distinction, as Hubert Dreyfus reminds us, was recognized very clearly by John Dewey in Human Nature and Conduct, from whom we borrow in fact the distinction between know-how and know-what:

We may be said to know how by means of our habits. … We walk and read aloud, we get off and on street cars, we dress and undress, and do a thousand useful acts without thinking of them. We know something, namely, how to do them. … [If] we choose to call [this] knowledge … then other things also called knowledge, knowledge of about things, knowledge that things are thus and so, knowledge that involved reflection and conscious appreciation. remains of a different sort. [Note 1-15]

In summary, then, my main point is that most of our mental and active life is of the immediate coping variety, which is transparent, stable, and grounded in our personal histories. Because it is so immediate, not only do we not see it, we do not see that we do not see it, and this is why so few people have paid any

attention to it until phenomenology and pragmatism, on the one hand, and new trends in cognitive science, on the other hand, brought it to the fore. Yet the question remains: how can this distinction between coping behaviors and abstract judgment, between situatedness and morality, be applied to the study of ethics and the notion of ethical expertise?

The second lecture: On Ethical Expertise

The ethical expert

My main point in the First Lecture was that philosophers and scientists who study the mind have grossly neglected skilled behavior, which is immediate, central, and pervasive, in favor of exploring deliberate, intentional analysis. To counter this imbalance we must impress on ourselves how much of our lives is spent in skilled behavior — working, moving, talking, eating — and how little is spent in deliberate, intentional analysis. Yet it is this latter category that we notice. It is this latter category which has been the focus of philosophers and scientists alike.

It is also clear that we can add responding to the needs of others to our list of skilled behaviors without doing violence to our concept of ordinary life. And if that is so, then it should also be clear that the situations in which we exercise ethical expertise far outnumber those in which we must exercise explicit ethical deliberation.

Nevertheless, even the most subtle of modern writers on ethics continue to tell us that the central issue is reasoning. For instance, no less a light than Alaisdair MacIntyre, in After Virtue, concludes from a reading of Aristotle's Nicomachean Ethics that the moral agent is best described as a competent performer deliberately choosing among maxims: "In practical reasoning the possession [of an adequate sense of the tradition to which one belongs] … appears in the kind of capacity for judgement which the agent possesses in knowing how to select among the relevant stack of max- ims and how to apply them in particular situations." [Note 2-1] The fact that this pronouncement translates easily into information-processing jargon betrays the inability of this kind of approach to appreciate immediate coping, which lies beyond the grasp of computationalist description, regardless of the details of its terms of analysis. Only the enactivist view-point, which I introduced in the First Lecture, can account for immediate coping.

Furthermore, we acquire our ethical behavior in much the same way we acquire all other modes of behavior: they become transparent to us as we grow up in society. This is because learning is, as we know, circular: we learn what we are supposed to be in order to be accepted as learners. This socialization process has roots too profound for analysis here.

Still, we can say here that when the matter is viewed in this light, it is clear that an ethical expert is nothing more or less than a full participant in a community: we are all experts because we all belong to a fully textured tradition in which we move at ease. In traditional communities, furthermore, there are models of ethical expertise who can be singled out as even more expert than the common run (the "wise ones"). In our modern society, however, such role models for ethical expertise (unlike, say, role models for athletic expertise) are more difficult to identify. This is, I claim, one important reason why modern ethical thinking has such a nihilistic flavor, a point to which I will return later.

The view from the teaching traditions

This neglect of ethical coping as a central locus for concern is not universal. Some of the great teaching traditions of the East — Taoism, Confucianism, Buddhism — see things otherwise, and that is precisely why I want to turn to them here.

I am aware that this is a bold step, but I believe that it is one we need to take. In my opinion, we need to enlarge our horizon to encompass non-Western traditions of reflection upon experience. If philosophy in the West no longer occupies a privileged, foundational position with respect to other cultural activities, such as science or art, then a full appreciation of philosophy and its importance for human experience requires that we examine the role of philosophy in cultures other than our own. In our culture, cognitive science has caused great excitement among philosophers (and the public at large) because it has enabled them to see their tradition in a new light. Were we to entertain the idea that there is no hard and fast distinction between science and philosophy, then philosophers such as Descartes, Locke, Leibniz, Hume, Kant, and Husserl would take on a new significance: they could be seen as, among other things, proto-cognitive scientists. (Or as Jerry Fodor puts it: "In intellectual history, everything happens twice, first as philosophy and then as cognitive science." [Note 2-2]) Might this not also be the case for philosophical traditions with which we are less familiar?

I focus on the three traditions I have mentioned because they belong to a constellation with much in common, and because it is my contention that the rediscovery of Asian philosophy, particularly of the Buddhist tradition, is a second renaissance in the cultural history of the West. [Note 2-3]

Let me start by concentrating on Meng-tzu or Mencius, an early Confucian from around the fourth century B.C.E. who holds a position of authority in this tradition comparable to that of Thomas Aquinas in the Christian tradition.

Mencius's view of ethics and the development of a virtuous person rests on the assumption that human nature is capable of such flourishing, and that people can strive for such growth. His view is that a person's natural disposition, joined with appropriate developmental conditions, determines a person's emotional responses. In this he follows a relatively simple developmental model: basic capacities exist and when they are nurtured in an unobstructed manner they generate the qualities one seeks. This is important, for it stands in total opposition to our Western Christian tradition of the fall and original sin. It means that when

Mencius declares that human nature is good, he is not referring to a hidden ontological substratum, but to human capacities. As he puts it: "As far as what is genuinely in him is concerned a man is capable of becoming good. This is what I mean by good. As for his becoming bad, that is not the fault of his native environment." [Note 2-4] What is essential to highlight at this point is Mencius's understanding of how people develop their basic capacities, that is, how people actively a-titivate the proper dispositions. There are some key pragmatic issues here that I shall return to later. For now, however, and in order to examine more fully Mencius's ideas on ethics, we need to examine three interrelated concepts central to what virtue is all about. These three key notions are extension (t'ui or ta), attention (ssu), and intelligent awareness (chih). Let us start by considering extension.

According to Mencius, people actualize virtue when they learn to extend knowledge and feelings from situations in which a particular action is considered correct to analogous situations in which the correct action is unclear. This assumes, of course, that we wish to make such an extension. The example he cites to illustrate this point is easy to understand:

Suppose a man were, all of a sudden, to see a young child on the verge of falling into a well. He would certainly be moved to compassion, not because he wanted to get in the good grace of the parents nor because he wished to win the praise of his fellow villagers or friends, nor yet because he disliked the cry of the child. [Note 2-5]

Given that this a normal starting point, the idea is to extend the feelings that arise from this situation in an appropriate way to other situations. Notice that both the kind of example Mencius uses as his basis for ethical training, a very ordinary know-how we all share, and the core of the method of extension resemble the learning we apply to all skills. One starts from a simple situation we can all handle and then extends one's skill in widening circles to situations that are more complex. This process, however, presumes that people can and will attend to what needs to be done using intelligent awareness. To extend feelings is both to see that one situation resembles another and to have these feelings "break through" into the new situation.

Now the specific capacity of the mind that underlies this process is the ability to attend (ssu). Again here one uses a natural capacity to focus on concrete objects. Failures in attention underlie misdirected self-cultivation. As Mencius puts it: "If one attends one gets it; if one does not, one does not." Mencius thus conceives of ethical training as a process that depends on perceiving clearly and identifying correspondences or affinities. He is especially opposed to the notion supported at the time by Mohism (and also evident, as we said, in Western ethical thinking) that ethical reasoning consists mainly in the application of rules or principles. For Mencius, rules become apparent to people either only after careful reflection or only in situations where duress forces them to make an assessment. The most significant difference lies in the importance given to a true description of the situation. To gather a situation under a rule a person must describe the situation in terms of categories we may call cognitive. Instead, if we try and see correspondences and affinities, the situation at hand becomes much more textured. All relevant aspects are included, not just those which fit the reduction of a categorical analysis. In this, then, Mencius coincides closely with my earlier assertion that reflection and analysis are most prominent when our immediate microworld breaks down.

For Mencius only truly virtuous people attend to their nature sufficiently well to understand an event in terms of their experience and thus ensure that the appropriate extension follows easily. For the truly virtuous then, moral judgment that results in immediate and spontaneous moral action is not different from true description. This approach permits Mencius to distinguish truly virtuous actions from those which only seem virtuous. An action is frilly virtuous only if it flows from an activated disposition. An action may be right but cannot be fully virtuous if it is not properly motivated. This is an important point.

Mencius clarifies this distinction by contrasting the man of virtue with the "village honest man," which one translator renders as the "bourgeois righteous simulation of excellence." According to Mencius,

[The village honest man thinks the following]:"Being in this world one must behave in a manner pleasing to this world. So long as one is good it is all right." … If you want to censure him you cannot find anything. He shares with others the practices of the day and is in harmony with the sordid world. He is liked by the multitude and is self-righteous. It is impossible to embark on the way of Yao and Shun [two famous sages] with such a man. Hence the name "enemy of virtue." Confucius said: "I dislike the village honest man for fear he might be confused with the virtuous." [Note 2-6]

In order to distinguish virtues from their semblances, Mencius identifies four kinds of human action, of which only one manifests truly ethical behavior, the other are semblances at best or downright counterfeits. These four, in ascending order of excellence, are (1) actions that arise from a desire for gain, (2) actions that arise from habitual response patterns, (3) actions that arise from following rules, and (4) actions that arise from extension. Those people who act from habitual response patterns rather than with intelligent awareness fail to perceive situations accurately. Those people whose actions are generated by adherence to rules are like beginners learning a motor skill. To use Mencius's language, such rules will always remain external to the agent, for they will always differ at least in some ways from the agent's internal inclination.

For Mencius, only people who act from dispositions they have at the very moment of action as a result of a long process of cultivation merit the name of truly virtuous. Such a person can be said to have "acted through benevolence and rightness. It was not that he put into action benevolence and rightness." Such a person does not act out ethics, but embodies it like any expert embodies his know- how; the wise man is ethical, or more explicitly, his actions arise from inclinations that his disposition produces in response to specific situations.

Thus truly ethical behavior does not arise from mere habit or from obedience to patterns or rules. Truly expert people act from extended inclinations, not from precepts, and thus transcend the limitations inherent in a repertoire of purely habitual responses. This is why truly ethical behavior may sometimes seem

unfathomable to the untrained eye, why it can be what is called in the Vajrayana tradition "crazy wisdom." [Note 2-7] This suppleness reveals the key elements in the person who has cultivated his expertise, for his expertise contains the intelligent awareness that Mencius calls chili. One cannot overemphasize the importance of this learning dimension in the teaching traditions. In particular it corresponds to the cultivation of prajtia in Mahayana Buddhism. Mencius himself highlights it when he distinguishes the excellence of Confucius as being due to his highly trained intelligent awareness. We can understand the character of this sort of excellence more clearly if we first confront the two extremes of how virtue is misunderstood.

On one extreme are those who consider crazy wisdom virtuous but insist that it is spontaneous expression unfettered by reason. And on the other extreme, are those who despise crazy wisdom and insist that people should rely on rational calculations about goals and means. The intelligent awareness — which only occasionally manifests as "crazy wisdom" — that Mencius describes takes a middle way between these two extremes: intelligence should guide our actions, but in harmony with the texture of the situation at hand, not in accordance with a set of rules or procedures.

And because truly ethical behavior takes the middle way between spontaneity and rational calculation, the truly ethical person can, like any other kind of expert, after acting spontaneously, reconstruct the intelligent awareness that justifies the action. And, like any other kind of expert, the truly ethical person can use such a postiori justification as a stepping-stone for continued learning. Indeed, even the beginner can use this sort of deliberate analysis to acquire sufficient intelligent awareness to bypass deliberateness altogether and become an expert.

In summary, then, we see that the interplay of intelligent awareness, attention, and extension is how a virtuous person becomes truly virtuous from even modest beginnings, and how truly ethical behavior differs from that of "the village honest man."

The pragmatic key to ethical expertise

Mencius's view of ethical expertise is at the same time very far from the dominant Western tradition of rational judgment and very close to the teachings of both Taoism and Buddhism. In all three traditions ethical behavior is approached pragmatically and _progressively.

To say that ethical behavior must be approached pragmatically and progressively is to put in positive terms what the Tao Te Ching of Lao-tzu explains with negatives, among them the well-known but untranslatable expression wu-wei, _which is sometimes, but inadequately, rendered as "not-doing":

A man of the highest virtue does not keep to virtue and that is why he has virtue …A man of the lowest virtue never strays from virtue and that is why he is without virtue …Thus the wise man deals with things through wu-wei and teaches through no words …Then the thousand things flourish without interruption …Less and less is done until wu-wei is achieved.When wu wei is done, nothing is left undone.

To us this formulation sounds like a paradox, and that it is, but not a vicious, circular one. To resolve it we must combine both sides of it; we must exercise a metalevel of understanding beyond the reach of logical analysis alone, as many a frustrated scholar has discovered. My point is that wu-wei points to a journey of experience and learning, not to a mere intellectual puzzle that one solves. It points to the process of acquiring a disposition where immediacy precedes deliberation, where nondual action precedes the radical distinction between subject and object.

It is hardly a coincidence that we find the same apparent paradox in all other traditions that point to the need for progressive ethical expertise. This is very clear in all main traditions of Buddhism. The third Ch'an patriarch Seng-ts'an, for instance, says:

When rest and no rest cease to be,then even oneness disappears.From small mind comes rest and unrestbut mind awakened transcends both.

This little poem echoes the discussion in that most famous of Buddhist thinkers, Nagarjuna, who in his Mulamadhyamikakarikas deconstructs all pairs of opposites, such as action and inaction, rest and motion, and finds that both elements of each pair are empty — sunya — that is to say, each exists only in relation to the other. His work is best understood against the larger background of the Mahayana Buddhist teachings and the human ideal of the boddhisattva to which we will return later. But like Mencius's truly virtuous person and like Lao-tzu's man of the highest virtue, a true boddhisattva neither comes nor goes, but rather "deals with things through tvu-wei." From the very first stage of the ten-stage boddhisattvic path (and it is a learning journey!), which is called acala, the immovable, the boddhisattva works without making any effort, just as the moon illuminates everything impartially. Again here the paradox of non-action in action vanishes when the actor becomes the action, that is to say, when the action becomes nondual.

As Martin Buber has put it: "This is the activity of the human being who has become whole: it has been called not doing, for nothing particular, nothing partial is at work in man and thus nothing of him intrudes in the world." [Note 2-8] When one is the action, no residue of self-consciousness remains to observe the action externally. When nondual action is ongoing and well established, it is experienced as grounded in a substrate both at rest and at peace. To forget one's self is to realize one's emptiness, to realize that one's every characteristic is conditioned and conditional. Every expert knows this sensation of emptiness well; in the West, for example, athletes, artists, and craftsmen have always insisted that self-consciousness interferes with optimal performance. This is one important aspect of what the Heart Sutra (a key text in Mahayana Buddhism) extols when it says that one who has realized the emptiness of all actions acts freely because he is "without hindrance in the mind." Needless to say, there are major differences between the expert athlete and the boddhisattva, not the least of which being the range of their expertise. We should not confuse the one with the other, but the comparison is sufficiently appropriate to illustrate that what the teaching traditions are getting at is not mere mystical mumbo-jumbo.

Thus we can distinguish between self-conscious or intentional action and self-less or intentionless action. At first the idea of action without intention seems absurd, but in fact our lives are full of intentionless actions. We dress, we eat, and more important, we exercise consideration for others. We do all these things without intention, but we do not do them randomly or purely spontaneously. We do them without intention because we are experts at them. Through appropriate extension and attention and by training over time we have transformed these actions into embodied behavior.

But just what is the key element that makes such intentionless learning possible? The answer is right in front of us. Our microworlds and microidentities do not come all stuck together in one solid, centralized unitary self, but rather arise and subside in a succession of shifting patterns. In Buddhist terminology this is the doctrine, whose truth can be verified by direct observation, that the self is empty of self-nature, void of any graspable substantiality. Once we are fully able to ride with the enormous openness contained in this sunya of self, the possibilities for further self- understanding become both vast and immediately accessible. This point is crucial. It is the golden thread that unites our self-understanding with an external and scientific account of mental functioning.

On non-unitary cognitive selves

To make this non-unitary self meaningful in terms of our own tradition and from our (Western) perspective, I need only turn to modern cognitive science. Yet we need not confine ourselves to any single trend within the field of cognitive science, for even the more conservative viewpoints in the field, the classical computationalist perspective, for example, deny the existence of a solid, centralized, unitary self.

Computationalism in cognitive science embraces the idea that the self or cognizing subject is fundamentally fragmented or non-unified simply because it postulates mental or cognitive processes of which we are not only unaware, but of which we cannot be aware. In fact, computationalism postulates mental (not just physical and biological) mechanisms and processes that are not accessible to the "personal level" of consciousness, especially self-consciousness. In other words, one cannot discern in conscious awareness or self- conscious introspection any of the cognitive structures and processes that are postulated to account for cognitive behavior. Indeed, if cognition is fundamentally symbolic computation, this discrepancy between "personal" and "sub-personal" immediately follows, since presumably none of us has any awareness of computing in an internal, symbolic medium when we think.

It is possible to overlook the depth of this challenge to our self-understanding, largely because of our post-Freudian belief in the unconscious. There is a difference, however, between what we usually mean by "unconscious" and the sense in which computationalism means that mental processes are unconscious. We usually suppose that what is unconscious can be brought to consciousness — if not through self-conscious reflection, then through a disciplined form of introspective analysis such as Freudian psychotherapy. Computationalism insists on the existence of mental processes that cannot be brought to consciousness at all. Thus it is not that we are simply unaware of the rules that govern the generation of mental images or of the rules that govern visual processing; we cannot in principle ever be aware of these rules.

One computationalist account asserts that these processes cannot be brought to consciousness without ceasing to function, because by its very nature consciousness is slow and deliberate, not fast and automatic as, say, vision must be to function properly. Another account describes these processes as "'nodular," that is, as comprising distinct subsystems that cannot be "penetrated" by conscious mental activity. [Note 2-9] Thus in this sense computationalism challenges our conviction that consciousness and the mind amount to the same thing, or that there is any essential or necessary connection between the two.

Of course, Freud also challenged the idea that the mind and consciousness are the same. Furthermore, he certainly realized that to distinguish between the mind and consciousness entails the disunity of the self or cognizing subject, a point to which we shall turn shortly. It is not clear, however, whether Freud took the further step of calling into question the idea that there is an essential or necessary connection between the mind and consciousness. Freud, in his argument for unconscious beliefs, desires, and motivations, left open the possibility that these unconscious processes belonged to a fragment of ourselves hidden in the depths of the psyche." [Note 2-10] Although it is not clear the extent to which Freud meant such a fragmentation literally, it is clear that when cognitive scientists postulate a collection of fragmentary, nonunifiable processes, they mean exactly what they are saying. As Dennett puts it: "Although the new [cognitivist] theories abound with deliberately fanciful homunculus metaphors — subsystems like little people in the brain sending messages back and forth, asking for help, obeying and volunteering — the actual subsystems are deemed to be unproblematic nonconscious bits of organic machinery, as utterly lacking in point of view or inner life as a kidney or kneecap."[Note 2-11] In other words, the characterization of these "subpersonal" systems in "fanciful homunculus metaphors" is only provisional, for eventually all such metaphors are "discharged" — they are traded in for the storm of activity among such selfless processes as neural networks or AI data structures. [Note 2-12]

Our pretheoretical, everyday conviction, however, is that cognition and consciousness — especially self-consciousness — belong together in the same domain. Cognitivism runs directly counter to this conviction: in determining the domain of cognition it explicitly cuts across the conscious/ unconscious distinction. The domain of cognition consists of those systems which must be seen as having a distinct representational level, not necessarily of those systems which are conscious. Some representational systems are, of course, conscious, but they need not be to have representations or intentional states. Thus, for cognitivists, cognition and intentionality (representation) are the inseparable pair, not cognition and consciousness.

This theoretical division of the domain of cognition was considered by early cognitive scientists to be "an empirical discovery of no small importance"[Note 2-13] and indicates, again, the remarkable transformation wrought by the cognitive sciences altogether. But now a problem arises: we seem to be losing our grip on something that is undeniably close and familiar — our sense of self. If consciousness — to say nothing of self-consciousness — is not essential for cognition, and if, in the case of cognitive systems that are conscious, such as ourselves, consciousness amounts to only one kind of mental process, then just what is the

cognizing subject? Is it the collection of all mental processes, both conscious and unconscious? Or is it simply one kind of mental process, such as consciousness, among all the others? In either case, our sense of self is challenged, for we typically suppose that to be a self is to have a coherent and unified "point of view," a stable and constant vantage point from which to think, perceive, and act. Indeed, this sense that we have (are?) a self seems so incontrovertible that its denial — even by science — strikes us as absurd. And yet, if someone were to turn the tables and ask us to look for the self, we would be hard-pressed to find it. Dennett, as usual, makes this point with flair: "You enter the brain through the eye, march up the optic nerve, round and round the cortex, looking behind every neuron, and then before you know it, you emerge into daylight on the spike of a motor nerve impulse, scratching your head and wondering where the self is."[Note 2-14]

Our problem, however, goes even deeper. It is one thing to be unable to find a coherent and unified self amid the furious storm of "subpersonal" activity. This inability would certainly challenge our sense of self, but the challenge would be limited. We could still suppose that there really is a self, we simply cannot find it in this fashion. Perhaps, as Jean-Paul Sartre held, the self is too close, and so we cannot uncover it by turning back upon ourselves. The computationalist challenge, however, is much more serious. According to computationalism, cognition can proceed without consciousness, for there is no essential or necessary connection between the two. Now whatever else we suppose the self to be, we typically suppose that consciousness is its central feature. It follows, then, that computationalism challenges our conviction that this most central feature of the self is needed for cognition. In other words, the cognitivist challenge does not consist simply in asserting that we cannot find the self; it consists, rather, in the further implication that the self is not even needed for cognition.

At this point, the tension between what science affirms and our own immediate experience seems to insist upon is tangible. If cognition can proceed without the self, then why do we nonetheless have the experience of self? We cannot simply dismiss this experience without explanation. Until recently, many scientists and philosophers of mind nonchalantly shrugged off this problem by arguing that the perplexities surrounding it are just not relevant to the purposes of cognitive science. [Note 2-15]

To make any further headway in our inquiry we must look more closely at the nature of this fragmentation. As I will discuss in the Third Lecture, the nature of this fragmentation is that of emergent (or self-organizing) properties from brain mechanism, giving rise to what I shall term a virtual self, a mode of analysis which is very recent in cognitive science and Western thought altogether.

The Third Lecture: The Embodiment of Emptiness

More on non-unitary selves and cognitive agents

I want to pursue further the new approach to the notion of a cognitive subject by focusing on the cognitive activity that occurs in that very special space I call the hinges of the immediate present. For it is in the immediate present that the cognitive subject actually lives. But before we proceed we need to revise, as we did for the question of skills and their importance, some entrenched assumptions concerning cognitive mechanisms inherited from the computationalist orthodoxy. [Note 3-1]

The main point is this: the sensorimotor coherencies of microworlds and microidentities we discussed in the First Lecture disguise the arising of a large set, or ensemble, of transiently correlated neurons within the brain. These ensembles are both the source and the result of the activity of the sensory and effector surfaces.

Consider, for example, Aplysia, a water mollusk with a "small" nervous system (a few thousand neurons). When Aplysia brings its siphon in touch with a surface (or when the siphon is independently touched), it contracts its gill. This is the so-called gill-withdrawal reaction, one among many of the behavioral patterns normally present in these animals. Traditionally, this sort of behavior is described as being mediated by a "reflex-arc." A recent study has determined that in Aplysia the gill-withdrawal reaction activates a significant proportion of the entire nervous system.[Note 3-2] This ensemble of activated neurons arises in a coordinated and mutually influential manner, and their co-activation abates after a few seconds. Thus the neurons of even this invertebrate ganglion must be conceived as a network of overlapping ensembles which arise in various coherent configurations depending on the animal's context.

The lessons learned from such humble mollusks can be applied to animals with larger and more complex nervous systems, in short, to animals with brains. Indeed, all that changes as we move "up" the chain of being is the number of mediating interneurons and the architecture of the nervous system being studied. In humans, for example, some 1011 interneurons interconnect some loo motor neurons which are linked to 107 sensory neurons distributed in receptor surfaces throughout the body. The rise and decay of neuronal self-organization has an even more pronounced effect in larger brains. In the cat's brain, for instance, 5—100 million neurons are active during a simple visuo-motor task of pressing a lever.[Note 3-3] Such neural assemblies arise in a patchwork of regional areas, evincing the enormous distributed parallelism proper to vertebrate brains.

In fact, it is fair to say that a recently established fact of the brain's constitution is what I like to call the Law of Reciprocity: if a region (say a cortical area, or a specific nucleus) A is connected to another region B, then B is also connected to A, but by a different anatomical route. Consider, for instance, the mammalian visual system. Consider, further, the well-known flow of impulses from the retina to the so-called first "relay" station in the visual system, the dorsal thalamus (call this the region A), and then on from the thalamus to the primary visual cortex (call this B), and then on to other cortical regions. There are, in accordance with the Law of Reciprocity, connections from B back to A, from the cortex back to the thalamus, and they are even more numerous than those

from the thalamus to the cortex.[Note 3-4] This bidirectional thalamo-cortical neuronal traffic is not a mere anatomical nicety: the visual performance of an animal depends on the integrity of this feedback loop.

Thus, the dynamic underlying a perceptuo-motor task is that of a network, a highly cooperative, two-way system, and not that of a linear process in which information is abstracted from sense data in a unidirectional sequence of stages. The dense interconnections among the subnetworks of the brain ensure that every active neuron will operate as part of a large and distributed ensemble. For example, although neurons in the visual cortex do have distinct responses to specific "features" of the visual stimuli (position, direction, contrast, and so on), these responses occur only in an anesthetized animal with a highly simplified (internal and external) environment. When more normal sensory conditions are allowed, and the animal is studied awake and behaving, it has become increasingly clear that the stereotyped neuronal responses to "features" are highly labile and context sensitive to effects such as bodily tilt or auditory stimulation.[Note 3-5] Even a change in posture that is not accompanied by any change in sensorial stimulation will alter the neuronal responses, demonstrating that even the supposedly downstream motorium is in resonance with the sensorium.[Note 3-6]

If I may continue to use vision as an example, I can take the previous discussion up one level of generalization, to note that in recent years research has turned to the study not of centralized "reconstruction" of a visual scene for the benefit of an ulterior homunculus, but of a patchwork of visual modalities, including at least form (shape, size, rigidity), surface properties (color, texture, specular reflectance, transparency), three-dimensional spatial relationships (relative positions, three-dimensional orientation in space, distance), and three-dimensional movement (trajectory, rotation).[Note 3-7] It has become evident that these different aspects of vision are emergent properties of concurrent subnetworks, which have a degree of independence and even anatomical separability, but cross-correlate and work together so that a visual percept is this coherency.

This kind of architecture is strongly reminiscent of a "society" of agents, to use Minsky's metaphor.[Note 3-8] This multidirectional multiplicity is counterintuitive but typical of complex systems. I say counterintuitive because we are used to the traditional causal mode of input—processing—output. Nothing in the foregoing description suggests that the brain "processes" information in such a way; such popular, computer-like descriptions of the workings of the brain are simply incorrect. Instead, the architecture of the brain supports a different kind of operation: signals move "back and forth," gradually becoming more coherent until a microworld has been constituted. The entire exercise takes a certain amount of time, which accounts for why every animal exhibits a natural temporal parsing. In the human brain this flurry of cooperation typically takes about 200—500 msec, the "nowness" of a perceptuo-motor unity.[Note 3-9] Contrary to what seems to be the case from a cursory introspection, cognition does not flow seamlessly from one "state" to another, but rather consists in a punctuated succession of behavioral patterns that arise and subside in measurable time. This insight of recent neuroscience — and of cognitive science in general — is fundamental, for it relieves us from the tyranny of searching for a centralized, hornuncular quality to account for a cognitive agent's normal behavior.

Our present concern at this point is with one of the many consequences of this view of the disunity of the subject, understood as a cognitive agent. The question I have in mind can be formulated thus: Given that there is a myriad of contending subprocesses in every cognitive act, how are we to understand the moment of negotiation and emergence when one of the many potential microworlds takes the lead and constitutes a definite behavior? In more evocative terms: How are we to understand the very moment of being-there when something concrete and specific shows up?

The answer I wish to propose here is that within the gap during a breakdown there is a rich dynamic involving concurrent subidentities and agents. This rapid dialogue, invisible to introspection, has recently been revealed in brain studies.

Some key aspects of this idea were first introduced by Walter Freeman, who over many years of research, managed to insert an array of electrodes into the olfactory bulb of a rabbit so that a small portion of the global activity could be measured while the animal behaved freely? [Note 3-10] He found that there is no clear pattern of global activity in the bulb unless the animal is exposed to one specific odor several times. Furthermore, he found for the first time that such emergent patterns of activity are created out of a background of incoherent or chaotic activity by fast oscillations (i.e., with periods of about 5—10 msec) until the cortex settles into a global electrical pattern, which lasts until the end of the sniffing behavior and then dissolves back into the chaotic background.[Note 3-11] The oscillations then provide a means of selectively binding a set of neurons in a transient aggregate that constitutes the substrate for smell perception at that precise instant. Smell appears in this light, not as some kind of mapping of external features, but as a creative form of enacting significance on the basis of the animal's embodied history. What is most pertinent here is that this enaction happens at the hinge between one behavioral moment and the next, via fast oscillations between neuronal populations that can give rise to coherent patterns.

There is growing evidence for this kind of fast resonance to transiently bind neuronal ensembles during a percept. For example, it has been reported in the visual cortex of cats and monkeys; it has also been found in such radically different neural structures as the avian brain and the ganglia of an invertebrate, Hermissenda. [Note 3-12] This universality is important, for it points to the fundamental nature of resonance binding as a mechanism for the enaction of sensorimotor couplings. Had resonance binding been restricted to, say, mammals, it would have been far less interesting as a working hypothesis. It is important to note here that this fast resonance is not linked to sensorial triggers in any simple way: the oscillations appear and disappear quickly and quite spontaneously in various places of the brain.

It seems that between breakdowns these oscillations are the symptoms of very rapid reciprocal cooperation and competition between distinct agents activated by the current situation, vying with each other for differing modes of interpretation for a coherent cognitive framework and readinessfor-action. This dynamic engages all the subnetworks that give rise to the entire readiness-for-action in the next moment. It involves not just sensory interpretation and motor action but the entire gamut of cognitive expectations and emotional tonality central to the shaping of a microworld. On the basis of this dynamic one neuronal ensemble (one

cognitive subnetwork) finally prevails to become the behavioral mode for the next cognitive moment, a microworld.

When I say "prevails," I do not mean to say that this is a process of optimization; it resembles more a bifurcation in a chaotic dynamic. It follows that such a cradle of autonomous action is forever lost to lived experience, since by definition, we can only inhabit a microidentity when it is already present, and not when it is in gestation. In other words, in the breakdown before the next microworld shows up, there are a myriad of possibilities available until, out of the constraints of the situation and the recurrence of history, a single one is selected. This fast dynamic is the neural correlate of the autonomous constitution of a cognitive agent at a given present moment of its life.

Emergent properties and virtual selves

The nature of the identity of the cognitive self just discussed is one of emergence through a distributed process. The emergent properties of an interneural network are enormously rich and merit further discussion at this point. What I wish to underscore here is the relatively recent (and stunning!) conclusion that lots of simple agents having simple properties may be brought together, even in a haphazard way, to give rise to what appears to an observer as a purposeful and integrated whole, without the need for central supervision. We have already touched on this theme when discussing the constant arising and subsiding of neuronal ensembles underlying behavior. I wish at this point to address this issue more generally. I base my conclusions on contemporary studies of various complex systems inspired by biological examples. [Note 3-13]

One of the most compelling of these examples is the social insect colony. The beehive and the ants' nest have long been considered "superorganisms," but this was little more than a metaphor until recently. It was not until the 1970s that detailed experiments were made whose results could not be explained without taking into account the entire colony.[Note 3-14] In one particularly elegant experiment, the most efficient nurses in a Neoponera apicalis colony were removed to form a subcolony. These nurses radically changed social status, foraging more and nursing less. The contrary happened in the main colony: formerly low-level nurses increased their nursing activity. The whole colony, however, showed evidence of both configurational identity and memory. When the efficient nurses were returned to the main colony, they resumed their previous status.[Note 3-15]

What is particularly striking about the insect colony is that we readily admit that its separate components are individuals and that it has no center or localized "self." Yet the whole does behave as a unit and as if there were a coordinating agent present at its center. This corresponds exactly to what I mean by a selfless (or virtual) self a coherent global pattern that emerges from the activity of simple local components, which seems to be centrally located, but is nowhere to be found, and yet is essential as a level of interaction for the behavior of the whole.

The import of this model of how complex systems exhibit emergent properties through the coordinated activity of simple elements is, in my eyes, quite profound for our understanding of cognitive properties. It introduces an explicit alternative to the dominant computationalist tradition, which postulates that sensory inputs are successively elaborated to reconstitute a centralized and internal representation of the external world.

Applied to the brain, this new model explains why we find networks and subnetworks interacting promiscuously without any real hierarchy of the sort typical of computer algorithms. To put this differently, in the brain there is no principled distinction between software and hardware or, more precisely, between symbols and nonsymbols. I raise this point to help the reader break the hold that computationalism has had on our discourse in the area for so many years and resist the consequent tendency to conceptualize the cognitive self as some computer program or high-level computational description, for it is not that sort of thing at all. _The cognitive self is its own implementation: its history and its action are of one piece.

In fact, all we find in modern artificial neural network machines underlying the regularities we call their behavior or performance are interactions between ensembles. We may see that some of these ensembles recur regularly enough to describe them as being program-like, but this is another matter. Although artificially constructed, such emerging ensembles are not "computations" in the sense that their dynamics are formally specifiable as implementations of high-level algorithms. Neural networks even in their fine detail are not _like a machine language, since there is simply no transition from an elemental operational level with a semantics and a higher, emergent level where behavior occurs. If there were, the classical computer wisdom would immediately apply: we could ignore the hardware since it adds nothing of significance to the actual computation (other than constraints of time and space). In contrast, in distributed, network models these "details" are precisely what makes a global effect possible, and why they mark a sharp break with tradition in AI. [Note 3-16] Naturally, this reinforces the parallel conclusions that apply to natural neural networks in the brain, as we discussed before.

Now this demands that we clarify the second aspect of the self to be addressed: its mode of relation with the environment. Ordinary life is necessarily one of situated agents, continually coming up with what to do faced with ongoing parallel activities in their various perceptuo-motor systems. This continual redefinition of what to do is not at all like a plan selected from a repertoire of potential alternatives; it is enormously dependent on contingency and improvisation, and is more flexible than any plan can be. A situated cognitive entity has — by definition — a perspective. This means that it isn't related to its environment "objectively," independently of the system's location, heading, attitudes, and history. Instead, it relates to it in relation to the perspective established by the constantly emerging properties of the agent itself and in terms of the role such running redefinition plays in the coherence of the entire system.

Here we must sharply differentiate between "environment" and "world," for the cognitive subject is "in" both, but not in the same way. On the one hand, a body interacts with its environment in a straightforward way. These interactions are of the nature of macrophysical encounters — sensory transduction, mechanical

performance, and so on — nothing surprising about them. However, this coupling is possible only if the encounters are embraced from the perspective of the system itself. This embrace requires the elaboration of a surplus signification based on this perspective; it is the origin of the cognitive agent's world. Whatever is encountered in the environment must be valued or not and interacted with or not. This basic assessment of surplus signification cannot be divorced from the way in which the coupling event encounters a functioning perceptuo-motor unit; indeed, such encounters give rise to intentions (I am tempted to say "desires"), and intentions are unique to living cognition.[Note 3-17]

To put this another way, the nature of the environment for a cognitive self acquires a curious status: it is that which lends itself to a surplus of signification. Like a jam session, the environment inspires the neural "music" of the cognitive system. Indeed, the cognitive system cannot live without this constant coupling with and the constantly emerging regularities provided by its environment; without the possibility of coupled activity the system would become a mere solipsistic ghost.

For instance, light and reflectance (among many other macrophysical parameters such as edges and textures, but let us simplify for the argument's sake), lend themselves to a wide variety of color spaces, depending on the nervous system involved in that encounter. During their respective evolutionary paths, fishes, birds, mammals, and insects have brought forth various differently color spaces, not only with quite distinct behavioral significance, but with different dimensionalities. Thus differences in color vision from one animal to the next are not a matter of a greater or lesser ability to resolve colors.[Note 3-18] Color is demonstrably not a property that is to be "recovered" from environmental "inputs" in some unique way. Color is a dimension that shows up only in the phylogenetic dialogue between an environment and the history of an active autonomous self that partly defines what counts as an environment. Light and reflectances provide a mode of coupling, a perturbation that triggers, that provides an occasion for, the enormous informative capacity of neural networks to constitute sensorimotor correlations and hence put into action their capacity for imagining and presenting. It is only after all this has happened, after a mode of coupling becomes regular and repetitive, like colors in our — and others' — worlds, that we observers, for ease of language, say that color corresponds to or represents an aspect of the world.

A dramatic and recent example of this surplus signification and the dazzling performance of the brain as a generator of neural "narratives" is provided by the technology of the so-called virtual realities. A helmet fitted with cameras over the eyes and a glove or suit with electrical transducers for motions are linked, not through the usual coupling with the environment, but through a computer. Thus each movement of the hand or body corresponds to images according to principles entirely under the control of the programmer. For example, each time my hand, which appears as a "virtual" iconic hand in my image, points to a place, the image that follows simulates flying to the place pointed at. Visual perception and motions thus give rise to regularities that are proper to this new manner of perceptuo-motor coupling. What is most significant for me here is how quickly this "virtual" world comes to seem real: we inhabit a body within this new world after about fifteen minutes or so "under the headset." And as far as this world is concerned, the experience of flying through walls or diving into fractal galaxies seems perfectly "real." This "reality shift" occurs despite the poor quality of the image, the low sensitivity of the sensors, and the limited bandwidth of the interface between sensory and image surfaces available through a program running on a personal computer. The nervous system is such a gifted synthesizer of regularities that any basic material suffices as an environment to bring forth a compelling world.

Even the very pragmatically oriented field of artificial intelligence is beginning to study the situatedness of agent endowed with progressively richer internal self-organizing modules. [Note 3-19] When the synthesis of intelligent behavior is approached in such an incremental manner, with strict adherence to the sensorimotor viability of an agent, the notion that the world is a source of information to be represented simply disappears. The autonomy of the cognitive self comes fully in focus. Thus in Rodney Brooks's proposal for a new robotics (or, as he says, for a nouvelle Al) his minimal creatures join together in various activities through a rule of cohabitation between them. This engineering strategy is homologous to an evolutionary pathway through which modular subnetworks intertwine with one another in the brain. This new approach to artificial intelligence should result in the creation of devices that are more truly intelligent, autonomous, and sense-giving than the brittle information processors that depend on a pre-given environment or an optimal plan that have been constructed to date.

It is interesting to note that in this paper Brooks also traces the origin of what he describes as the "deception of AI" to the tendency in AI (and in the rest of cognitive science as well) to abstraction, especially for the purpose of factoring out situated perceptual and motor skills. As I have argued here (and as Brooks argues for his own reasons), such abstraction misses the essence of cognitive intelligence, which resides only in its embodiment. It is as if one could separate cognitive problems into two types: those which can be solved through abstraction and those which cannot. Those of the second type typically involve perceptual and motor skills of agents in unspecified environments. When cognitive intelligence is approached from this self- situated perspective, it quickly becomes obvious that there is no place where perception could deliver a representation of the world in the traditional sense. The world shows up through the enactment of the perceptuo-motor regularities. As Brooks puts it:

Just as there is no central representation there is no central system. Each activity layer connects perception to action directly. It is only the observer of the creature who imputes a central representation or central control. The creature itself has none: it is a collection of competing behaviors. Out of the local chaos of their interactions there emerges, in the eye of the observer, a coherent pattern of behavior. [Note 3-20]

The self as virtual person

Even if we like these ideas about selfless selves at both the basic behavioral level and the more elaborate cognitive level, as Dennett puts it: "We want to exempt ourselves (we want to exempt our selves). The problem is that it seems as if we at least are very different: we are top-down, centered, globally directing." [Note 3-21] This is why we feel compelled to project a centralized center or agent, be it a homuncular soul- like entity, or a vaguer sense of "self

as a process."

I think that the radical novelty of our newly acquired and still fragmentary understanding of emergent properties in distributed network processes lies precisely in that they are strong metaphors, nay, exemplars, for what is a selfless self: a coherent whole that is nowhere to be found and yet can provide an occasion for the coordinated activity of neural ensembles. I underline the strength of these metaphors because without the numerous examples worked out recently this apparent paradox of nonlocalization liable to designation as a totality becomes a contradiction, and unless this apparent paradox is addressed on this constructive metalevel, we quickly slide back into the traditional debates about the existence versus the nonexistence of the self and the person. The seeming paradox resides in a two-way movement between levels: "upward" with the emergence of properties from the constituting elements, and "downward" with the constraints imposed by global coherence on local interactions. The result (and the resolution of the paradox) is a nonsubstantial self that acts as if it were present, like a _virtual interface.

_The more we see the selfless nature of our selves in various "regions" of the organism, the more we become suspicious of our feeling of "I" as a true center. Either we are unique in the living and natural world, or else our very immediate sense of a central, personal self is the same kind of illusion of a center, accountable by more of the same kind of analysis as we have already performed on the basic sensorimotor cognitive selves.

Needless to say, my preference is squarely with the second alternative. What we call "I" can be analyzed as arising out of our recursive linguistic abilities and their unique capacity for self-description and narration. As long-standing evidence from neuropsychology shows, language is another modular capacity cohabiting with everything else we are cognitively. Our sense of a personal "I" can be construed as an ongoing interpretative narrative of some aspects of the parallel activities in our daily life, whence the constant shifts in forms of attention typical of our microidentities. Whence also the relative fragility of its narrative construction.[Note 3-22]

If this narrative "I" is necessarily constituted through language, then it follows that this personal self is linked to life because language cannot but operate as a social phenomenon. In fact, one could go one step further: the selfless "I" is a bridge between the corporeal body which is common to all beings with nervous systems and the social dynamics in which humans live. My "I" is neither private nor public alone, but partakes of both. And so do the kinds of narratives that go with it, such as values, habits, and preferences. In purely functionalist logic, "I" can be said to be for the interactions with others, for creating social life. Out of these articulations come the emergent properties of social life for which the selfless "I" is the basic component. Thus whenever we find regularities such as laws or social roles and conceive of them as externally given, we have succumbed to the fallacy of attributing substantial identity to what is really an emergent property of a complex, distributed process mediated by social interactions. Such emergent social properties can be projected as "exogenous" reference points, as is traditionally done, but they can equally well be deconstructed by the kind of analysis I have followed here.

Interestingly, even if we accept a re-interpretation of the "I" as virtual — as the product of linguistic closure and emergent distributed properties — our natural inclination in daily life is to continue as if nothing had changed. This is the best evidence that the process of self-constitution is so entrenched that seeing through it requires more than a convincing analysis. Exploration of the sunya, the virtual nature, of this deeply entrenched and continuously active drive for identity constitution is a matter of learning and sustained transformation.

The pragmatics of the virtual self

The Main Proposition

We now have what we need to grasp the nature of the emptiness of self and its relevance for ethical know-how. Modern Western science teaches us that the self is virtual and empty, and that it arises continuously to cope with breakdowns in our microworlds. Taoism, Confucianism, and Buddhism teach us that ethical expertise is progressive in nature and grounded in the ongoing realization of this empty self in ordinary life and action.

These two strands support each other, and at this point give substance to the following postulate, the core of my proposition in these lectures:

_Ethical know-how is the progressive, firsthand acquaintance with the virtuality of self.

_We normally avoid this aspect of our fragmented, virtual nature, and yet praxis is what ethical learning is all about. In other words, if we do not practice transformation, we will never attain the highest degree of ethical expertise. Learning to embody the empty self is certainly difficult, but all these wisdom traditions agree that the acquisition of this expertise is not only progressive and open-ended but centrally important.

Lessons from Psychoanalysis

We should avoid dismissing a concern with personal transformation as a bit of philosophical chinoiserie. Consider for a moment a Western tradition that arrives at a convergent conclusion: psychoanalysis. For my purpose here, psychoanalysis is particularly important because it is the only Western tradition centrally concerned with a pragmatics of human transformation. As Lacan says, "The unconscious is ethical at its core." [Note 3-23] But this remark must not be interpreted in the received Western tradition of rational deductive principles for action. What Lacan is claiming is closer to our argument here: the ethical implies putting the status of the knowing subject into question, and thus highlighting the need for therapeutic space provided for the person by the analytic situation. As is well known, an analytic "cure" is not a medical treatment, but rather the creation of a sort of suspension for the craziness of desire, which is tantamount to

saying a suspension of the acquired manner of emergence of the virtual self. This position is in itself ethical: it is the ethics of a know-how concerning the unconscious. [Note 3-24]

This sort of ethical know-how requires that we attempt to realize that no moral principle is in itself realizable, since the analytic process makes it progressively clearer that we are doomed to never be satisfied with any set of hopes and expectations, however rational they may seem to be. There cannot be any social order or moral order that is objectively desirable. The root of this is the fundamental psychoanalytic discovery that in contradistinction to theories that posit a unified central self, Freudian theories of the subject explode the self into pieces, both within the person and between persons. The analytic stance on ethics cannot then be based on compassion in the sense of sympathy or caritas, for this is always and already contaminated with the structure of the wildness of desire. Instead, it proposes that we suspend the temptation to be identified with the other and, instead, undertake a journey of learning to see ourselves and others as inescapably transitory and fragmented. The demand for this ethical link is manifested then in the only true love possible in the psychoanalytic world, that of transference.[Note 3-25] Instead of creating a link between the illusion of an ideal center or moral principle, this transferential love is always attempting to reveal to the subject that all the traits he takes as ideal in himself or in others are subsumed entirely in the unrealizable desire to make whole and substantial that which is forever fragmentary and virtual.

This Western tradition professes, at least in some of its schools, an ideal that is close to the wisdom or higher ethical training of the teaching traditions, but with its own twist. Let me return now to those traditions by describing the Buddhist ideal of the boddhisattva, which we touched on briefly in the Second Lecture.

The lesson from the teaching traditions

In all Buddhist traditions, we have seen, the practice of recognizing the emptiness of the self is the very foundation of training, including ethical training. The practice of recognizing emptiness in every moment is known as the practice of mindfulness/awareness or samatha-vispasnya. Essentially a radical not-doing, it is traditionally understood as a universal practice, but despite having been refined and explored for over 2500 years by over half the world, it was never discovered independently in the West. [Note 3-26] Instead of a space for the human in the analyst's studio required in psychoanalysis, samatha-vipasnya creates the space through non-action, which includes nonresponse to language.

The mindfillness/awareness student first begins by observing, in a precisely prescribed fashion, what the mind is doing, its restless, perpetual grasping, from moment to moment. This beginning enables the student to free himself from some of his habitual patterns of thought, which leads to further mindfulness. Eventually he begins to realize that there is no self in any of his actual experiences. This discovery can be disturbing. It can cause the student to lose heart and tempt him to swing to the other extreme, that of nihilism. This flight into nihilism demonstrates that the reflex to cling to a substantial self is so strong and deep-seated that we reify its absence as a sort of abyss.

As the student-practitioner continues, however, and as his mind relaxes further into awareness, a sense of warmth and inclusiveness begins to appear quite naturally. The street-fighter mentality of watchful self-interest slips away gradually to be replaced by an interest in others. We are already other-directed even at our most negative, and we already feel warmth toward some people, such as family and friends, in the same way that Mencius begins his ethical training with an examination of our spontaneous concern for the child falling into a well. The conscious realization of the sense of relatedness and the development of a more impartial sense of warmth is encouraged in the mindfulness/ awareness tradition by various practices that strikingly resemble the element of extension we discussed in Mencius and the Taoist tradition.

It is said that the full realization of groundlessness cannot occur if there is no warmth. For this reason, in the Mahayana tradition, which we have so far presented as being centrally concerned with groundlessness as sunyata, there is an equally central and complementary concern with groundlessness as compassion. The Sanskrit term translated here as "compassion" is karuna. This translation has some shortcomings, but there is no other satisfactory English term. In fact, most of the traditional Mahayana presentations do not begin with groundlessness, but rather with the cultivation of karuna. Nagarjuna, the equivalent of Mencius for the Buddhist tradition, for example, states in one of his works that the Mahayana teaching has "an essence of emptiness and compassion." [Note 3-27] This statement is sometimes paraphrased by saying that emptiness (sunyata) is full of compassion (karuna). Keiji Nishitani, a contemporary writer from the Zen tradition, echoes this statement when he succinctly states that "the nature of the task of the ought is the otherdirectedness of the is." [Note 3-28]

Thus sunyata, the loss of a fixed reference point or ground in either self, other, or a relationship between them, is said to be as inseparable from compassion as the two sides of a coin or the two wings of a bird. Our natural impulse, in this view, is one of compassion, but it has been obscured by habits of ego-clinging like the sun obscured by a passing cloud.

This is by no means the end of the path, however, for there is a further step to be made in understanding beyond the sunyata of primarily negative terms: no self, ego/essness, no world, nonduality, emptiness, groundlessness. In actual fact, the majority of the world's Buddhists do not speak of their deepest concerns in negative terms; these negatives are preliminaries — necessary to remove habitual patterns of grasping, unsurpassably important but nonetheless preliminaries — that are pointing toward the realization of a positively conceived state.

To be sure, this state can be threatening and paradoxical, as we already evoked in discussing wu-wei in Taoism. It is no ground whatsoever; it cannot be grasped as ground, reference point, or nest for a sense of ego. It does not exist — nor does it not exist. It cannot be an object of mind or of the conceptualizing process; it cannot be seen, heard, or thought of. Thus many traditional images exist for it: the sight of a blind man, a flower blooming in the sky. When the

conceptual mind tries to grasp it, it finds nothing, and so it experiences it as emptiness. It can be known (and can only be known) directly. It is called Buddha nature, no mind, primordial mind, absolute bodhicitta, wisdom mind, All Goodness, Great Perfection, That which cannot be fabricated by mind, Naturalness. It is not a hair's breadth different from the ordinary world; it is that very same ordinary, conditional, impermanent, painful, groundless world experienced (known) as the unconditional, supreme state. And the natural manifestation, the embodiment, of this state is karuna — unconditional, fearless, "ruthless," spontaneous compassion. As a contemporary Tibetan teacher puts it poignantly: "When the reasoning mind no longer clings and grasps, … one awakens into the wisdom with which one was born, and compassionate energy arises without pretense"[Note 3-29] How can we understand further this unconditional compassion? We need to backtrack and consider the development of ethical and virtuous action from the more mundane point of view of the student. The possibility for concern for others present in all humans is usually mixed with the sense of ego and so becomes confused with the need to satisfy one's own cravings for recognition and self- evaluation; these were the categories that included the village honest man of Mencius among others. Here, instead, I am referring to the spontaneous gestures that arise when one is not caught in the habitual patterns — when one is not performing volitional actions out of acquired habitual patterns. In other words, actions that embody and express the realization of the emptiness of self in a nondual manifestation of subject and object. It is the de-centered self ideal also expressed in the ethics of psychoanalysis. When action is done without the business-deal mentality, there can be relaxation and this is called supreme (or transcendental) generosity or prajñaparatnita. [Note 3-30]

When discussing wisdom from the point of view of this highest ethical expertise, the Sanskrit term often used is bodhicitta, which has been variously translated as "enlightened mind," "the heart of the enlightened state of mind," or simply the "awakened heart." Bodhicitta is said to have two aspects, one absolute and one relative. Absolute bodhicitta is the term applied to what is considered fundamental in most Buddhist practices, the experience of the groundlessness of sunyata, or the (positively defined) sudden glimpse of the awakened state itself, which reminds us of the demand for intelligent awareness that Mencius demanded for the truly virtuous man. Relative bodhicitta, on the other hand, is that fundamental warmth toward the world which practitioners report arising from the absolute experience, and which manifests as concern and appropriate action for the welfare of others beyond naive compassion. Conversely (in the order in which we have previously described these experiences), it is said that it is the progressive extension and development of a sense of unproblematic warmth toward the world which leads to the experience of the flash of absolute _bodhicitta.

Practitioners in these teaching traditions obviously do not realize any of these things all at once. Like any learning journey, it takes time and a sustained discipline, with many semblances of progress and retrogress along the way. But practitioners report that they catch glimpses that encourage them to keep striving. One of the most important steps consists in developing compassion toward one's own grasping fixation on the ego-self or maitri. _The idea behind this attitude is that confronting our own grasping tendencies is a friendly act toward ourselves. As this friendliness develops, our awareness and concern for those around us enlarges as well. It is at this point that we can begin to envision a more open-ended and nonegocentric compassion.

It should not be surprising at this point that one of the main characteristics of spontaneous compassion, which is not a characteristic of volitional action based on habitual patterns, is that it follows no rules. It is not derived from an axiomatic ethical system or even from pragmatic moral injunctions. Its highest aspiration is to be responsive to the needs of the particular situation. Nagarjuna conveys this attitude of responsiveness:

Just as the grammarian makes one study grammar,A Buddha teaches according to the tolerance of his students;Some he urges to refrain from sins, others to do good,Some to rely on dualism, others on non-dualism;And to some he teaches the profound,The terrifying, the practice of enlightenment,Whose essence is emptiness that is compassion. [Note 3-31]

Unrealized practitioners, of course, cannot dispense with rules and moral injunctions. At the beginners' level, there are many ethical rules in Buddhism whose aim is to put the body and mind into a state that imitates as nearly as possible how genuine compassion might manifest in that situation (just as the meditative sitting posture is said to imitate enlightenment). By following these rules, beginning Buddhists learn to actualize compassion the same way followers of Mencius are exhorted to actualize virtue, by extending knowledge and feelings from situations where a particular action is considered correct to analogous situations where correct action is unclear. Most interestingly for our discussion, compassionate action is also called skillful means (upaya) in Buddhism. Skillful means are considered to be inseparable from wisdom.

We must not, however, identify skillful means with ordinary skills like learning to drive a car or play the violin. In some ways skillful means in Buddhism are like our more familiar notion of a sensorimotor skill: the student practices ("plants good seeds"), that is, avoids harmful actions, performs beneficial ones, meditates, and extends his behavior to a wider and wider range. However, unlike mastery of an ordinary skill, mastery of the skillful means of ethical expertise results in the elimination of all habits so that the practitioner can realize that wisdom and compassion can arise directly and spontaneously out of wisdom. It is as if one were born already knowing how to play the violin and had to practice with great exertion in order to remove the habits that prevented one from displaying that virtuosity. Thus the true wu-wei of the wise is not manufactured, but uncovered. In Buddhism this is the image of the fully accomplished boddhisattva.

We touch here on an extremely important and philosophically delicate point: Is there a ground underlying the nonsolidity of the self? Or more succinctly, What is left in sunyata? The Tibetan Buddhist tradition talks about the constituents of virtual mind as being transformed by the continued boddhisattva journey into wisdom. This sense of transformation does not mean going away from the world and getting out of mental functioning, since the very constituents on which the inaccurate sense of self and world are based are also the basis of wisdom. The means of transforming mental constituents into wisdom is intelligent awareness, that is, the moment-to-moment realization of the virtual self as it is — empty of any egoistic ground whatsoever, yet filled with wisdom. Here one is positing

that authentic care resides at the very ground of Being, and can be made fully manifest in a sustained, successful ethical training. A thoroughly alien thought for our nihilistic Western mood, indeed, but one worthy of being entertained.

How can such an attitude of all-encompassing, de- centered, responsive, compassionate concern be fostered and embodied in our culture? It obviously cannot be created merely through norms and rationalistic injunctions. It must be developed and embodied through disciplines that facilitate the letting-go of ego-centered habits and enable compassion to become spontaneous and self-sustaining. It is not that there is no need for normative rules in the relative world — clearly such rules are a necessity in any society. It is that unless such rules are informed by the wisdom that enables them to be dissolved in the demands of responsivity to the particularity and immediacy of lived situations, the rules will become sterile, scholastic hindrances to compassionate action rather than conduits for its manifestation.

Perhaps less obvious but even more strongly enjoined by the mindfulness/awareness tradition is that practices undertaken simply as self-improvement schemes will only strengthen the very egotism they are intended to dispel. Because of the strength of egocentric habitual conditioning, there is a constant tendency, as practitioners in all contemplative traditions are aware, to try to grasp onto, possess, and become proud of the slightest insight, glimpse of openness, or understanding. Unless such tendencies become part of the path of letting-go which leads to compassion, then insights can actually do more harm than good. Buddhist teachers have often written that it is far better to remain an ordinary person and believe in ultimate foundations than to cling to some remembered experience of groundlessness without manifesting compassion.

Talk alone will certainly not suffice to engender spontaneous non-egocentric concerns and ethically developed persons. Even more than experiences of insight, words and concepts can be easily grasped at, taken as ground, and woven into a cloak of egohood. Teachers in all contemplative traditions warn against taking fixed views and concepts as reality. We simply cannot overlook the need for some form of sustained, disciplined practice or pratique de transformation de sujet, to use Foucault's apt term. [Note 3-32] This is not something that one can make up for oneself — anymore than one can make up the history of Western science for oneself. Nothing will take its place. Individuals must personally discover and grow into their own sense of virtual self.

To conclude: I have tried to weave together themes from science of mind and from the depth of the teaching traditions to illuminate my central concern about what ethical know-how is and how it is acquired. My presentation is, more than anything, a plea for a re-enchantment of wisdom, understood as non-intentional action. This skillful approach to living is based on a pragmatics of transformation that demands nothing less than a moment-to-moment awareness of the virtual nature of our selves. In its full unfolding it opens up openness as authentic caring. These are radical ideas and strong measures for the troubled times we have at hand, and the even more troubled ones we are likely to have.

Endnotes

1-1

Charles Taylor, Sources of the Self: The Making of the Modern Identity (Cambridge: Harvard University Press, 1989), p. 3.

1-2

Jean Piaget, The Moral Judgement of the Child (Glencoe, Ill.: The Free Press, 1935); cited in H. Dreyfus, "Towards a Phenomenology of Ethical Expertise," in J. Ogilvy, ed., Revisioning Philosophy (Albany: SUNY Press, 1990.)

1-3

The inspiration for this crucial observation comes mostly from my ideas on the central role of sensorimotor coordination in cognition (cf. H. Maturana and F. Varela, The Tree of Knowledge, 2d ed. [Boston: Shainbhala, New Science Library, 1992], and from my own experience in the wisdom traditions described below. However, I am very indebted to Hubert Dreyfus's recent work on the phenomenology of skills and their ethical importance. See H. Dreyfus and S. Dreyfus, ''What Is Morality? A Phenomenological Account of the Development of Ethical Expertise," in D. Rassmussen, ed., Universalism versus Communitarianism (Cambridge: MIT Press, 1990), and H. Dreyfus, "Towards a Phenomenology of Ethical Expertise." In what follows I draw heavily from these papers for some key philosophical quotations chosen by Dreyfus.

1-4

For a recent collection of papers explicitly dealing with this critique, see L. Steels and R. Brooks, eds., The Artificial Life Route to Artcial Intelligence: Building Embodied, Situated Agents (New Haven, Conn.: Lawrence Erlbaum Associates, 1995), and my own contribution therein, "The Re-enchantment of the Concrete," pp. 11-20.

1-5

I am grateful to Fernando Flores for sharing his insights on these topics with me; they are described in documents for internal use in his consulting firm Business Design Associates, Emeryville, California.

1–6

See E Varela, Connaître: Les Sciences Cognitives (Paris: Seuil, 1989), and Varela, Thompson, and Kosch, The Embodied Mind.

1–7

R. Held and A. Hein,"Adaptation of Disarranged Hand-Eye Coordination Contingent upon Re-afferent Stimulation," Perceptual-Motor Skills 8 (1958): 87–90.

1–8

P. Bach y Rita, Brain Mechanisms in Sensory Substitution (New York: Academic Press, 1962).

1–9

See, for instance, Piaget's classic Biologie et Connaissance (Paris: Gallimard, 1969).

1–10

See George Lakoff, Women, Fire, and Dangerous Things (Chicago: University of Chicago Press, 1983), and Mark Johnson, The Body in the Mind (Chicago: University of Chicago Press, 1989).

1–11

George Lakoff, "Cognitive Semantics," in Umberto Eco et al., eds., Meaning and Mental Representations (Bloomington: Indiana University Press, 1988), p. 121. This article provides a concise overview of Lakoff and Johnson's experientialist approach.

1–12

E. Rosch, C. B. Mervis, W. D. Gray, D. M. Johnson, and P. Boyes-Braem, "Basic Objects in Natural Categories," Cognitive Psychology 8 (1976): 382–439.

1–13

See P. Bourgine and F. Varela, eds., Towards a Practice of Autonomous Systems: The First European Conference on Artificial Life (Cambridge: MIT Press, Bradford Books, 1992).

1–14

H. Dreyfus discusses this point in H. Dreyfus and S. Dreyfus, Mind over Machine (New York: Macmillan, 1986).

1–15

J. Dewey, Human Nature and Conduct: An Introduction to Social Psychology (London: G. Allen & Unwin, 1922), p. 177.

2–1

Alaisdair MacIntyre,4fter Virtue (Notre Dame, Ind.: University of Notre Dame Press, 1981), p. 140; cited in Dreyfus and Dreyfus,"What Is Morality?"

2–2

Jerry Fodor, "The Present Status of the Innateness Controversy," in RePresentations: Philosophical Essays on the Foundations of Cognitive Science (Cambridge: MIT Press, Bradford Books, 1981), p. 298.

2–3

For a recent discussion of the ethnocentrism in Western philosophy from an insider's perspective, see Roger Pol-Droit, L'Amnesie Philosophique (Paris: Presses Universitaires de France, 1989). For an extensive study of non-Western thought, see David Loy, Non-Duality (New Haven: Yale University Press, 1989).

2–4

All quotations from Mencius come from Lee Yearly, Mencius and Aquinas: Theories of Virtue and Concepts of Courage (Albany: SUNY Press, 1991), p. 60. I am infinitely indebted to Yearly for his timely book and his perceptive selection of Mencius as a paradigmatic case; as is clear I borrow heavily from this work.

2–5

Ibid., p. 62.

2-6

Ibid., p. 67.

2-7

See, for instance, Chögyam Trungpa, Crazy Wisdom (Boston: Shambhala, 1990).

2-8

M. Buber, I and Thou (Edinburgh: T. Clark, I970), p. 125.

2-9

Jerry Fodor, The Modularity of Mind (Cambridge: MIT Press, Bradford Books, 1983).

2-10

Douglas R. Hofstadter and Daniel Dennett, eds., The Mind's Eye: Fantasies and Reflections on Self and Soul (New York: Basic Books, 1981), p. 12.

2-11

Ibid., p.13.

2-12

See Dennett's essays "Toward a Cognitive Theory of Consciousness" and "Artificial Intelligence as Philosophy and Psychology," in his Brainstorms (Cambridge: MIT Press, Bradford Books, 1978).

2-13

Zenon Pylyshyn, Computation and Cognition (Cambridge: MIT Press, 1989),p. 265.

2-14

Daniel Dennett, Elbow Room: The Varieties of Free Will: Evil, Wanting (Cambridge: MIT Press, Bradford Books, 1984), PP- 74-75.

2-15

See Jerry Fodor, The Language of Thought (Cambridge: Harvard University Press, 1975), p. 52.

3-1

Many of the ideas in this lecture are taken from my "Organism: A Meshwork of Selfless Selves," in A. Tauber, ed., Organism and the Origin of Self (Dordrecht: Kluwer, 1991).

3-2

See T. Carew and C. Sahley,"Invertebrate Learning and Memory: From Behavior to Molecules," Aim. Rev. Neurosci. 9 (1983): 435-87, and for voltage-sensitive dyes, see D. Zecevic, J. Wu, L. Cohen, J. London, H. Hopp, and C. Falk, "Hundreds of Neurons in the Aplysia Abdominal Ganglion Are Active during the Gill-Withdrawal Reflex,"Neurosci. 9 (1989): 3681-89.

3-3

E. John, Y Tang, A. Brill, R. Young, and K. Ono, "Double-Labeled Metabolic Maps of Memory," Science 233 (1986): 1167-75.

3-4

W. Singer,"Extraretinal Influences to the Thalamus," Physiol. Rev. 57 (1977): 386-420; F. Varela and W Singer, "Neuronal Dynamics in the Visual Cortico-Thalamic Pathway Revealed through Binocular Rivalry" Exp. Brain Res. 66 (1987): 10-20.

3-5

G. Horn and R. Hill, "Modifications of Receptive Fields of Cells in the Visual Cortex Occurring Spontaneously and Associated with Bodily Tilt," Nature 221 (1974): 185–87; M. Fishman and C. Michael, "Integration of Auditory Information in the Cat's Visual Cortex," Vision Research 13 (1973): 1415–19; F. Morell,"Visual System's View of Acoustic Space," Nature 238 (1972): 44–46.

3–6

M. Abeles, Local Cortical Circuits (Berlin: Springer Verlag, 1984).

3–7

See, for example, E. DeYoe and D. C. Van Essen, "Concurrent Processing Streams in Monkey Visual Cortex," Trends Neurosci.II (1988): 219–26.

3–8

M. Minsky, The Society of Mind (New York: Simon & Schuster, 1987).

3–9

R. Llinis,"The Intrinsic Electrophysiological Properties of Mammalian Neurons: Insights into Central Nervous System Function," Science 242 (1988): 1654–64; A. Gevins, R. Schaffer, J. Doyle, B. Cutillo, R. Tannehill, and S. Bressler, "Shadows of Thought: Shifting Lateralization of Human Brain Electrical Patterns during a Brief Visuo-MotorTask," Science 220(1983): 97–99; F. Varela, A. Toro, E. John, and E. Schwartz,"Perceptual Framing and Cortical Alpha Rhythms," Neuropsychologia 19 (1981): 675–86.

3–10

W. Freeman, Mass Action in the Nervous System (New York:Academic Press, 1975).

3–11

W Freeman and C. Skarda, "Spatial EEG Patterns, Nonlinear Dynamics, and Perception: The Neo-Sherringtonian View," Brain Research Reviews 10 (1985): 145–75.

3–12

For a recent review, see S. Bressler, "The Gamma Wave: A Cortical Information Carrier," Trends Neurosci. 13 (199o): 161–62; C. Gray and W. Singer's "Stimulus-Specific Neuronal Oscillations in Orientation Columns in Cat Visual Cortex," Proc. Natl. Acad. Sci. 86 (1989): 16981702, has been largely responsible for the wider acceptance of this hypothesis; for Hermissenda, see A. Gelperin and D. Tank, "Odour-Modulated Collective Network Oscillations of Olfactory Interneurons in a Terrestrial Mollusc," Nature 345 (199o): 437–39: and for the results on the bird's brain, see S. Neuenschwander, A. Engel, P. König, W. Singer, and F. Varela, "Synchronization of Neuronal Responses in the Optic Tectum of Awake Pigeons," Visual Neurosci. 13 (1996): 575–84.

3–13

C. Langton, Artificial Life (Reading, Mass.: Addison-Wesley, 1989); B. Goodwin and P. Saunders, eds., Theoretical Biology: Epigenetic and Evolutionary Order from Complex Systems (Edinburgh: Edinburgh University Press, 1989); J. Farmer, A. Lapedes, N. Packard, and B. Wendroff, eds., Evolution, Games and Learning (Amsterdam: North-Holland, 1986); S. Wolfram, Theory and Applications of Cellular Automata (Singapore: World Scientific, 1986); Bourgine and Varela, Towards a Practice of Autonomous Systems.

3–14

E. O. Wilson, The Insect Societies (Cambridge: Harvard University Press, 1971).

3–15

D. Fresnau and I. Lachaud,"La regulation sociale: Donnas preliminaires sur les facteurs individuals controlant l'organisation des taches chez Neoponera apicalis," Actes Coll. Insects Sociaux 2 (1985): 185–93.

3–16

P. Smolensky, "On the Proper Treatment of Connectionism," Beh. Brain Sci. i 1 (1988): 1–74; D. Dennett, "Mother Nature versus the Walking Encyclopedia: A Western Drama," in S. Ramsey, D. Rummelhart, and S. Stich, eds., Philosophy and Connectionist Theory (forthcoming).

3–17

D. Dennett, The Intentional Stance (Cambridge: MIT Press, 1987).

3-18

For more on this see E. Thompson, A. Palacios, and F. Varela, "Ways of Coloring: Comparative Color Vision as a Case Study in Cognitive Science," Beh. Brain Sci. 15 (1992): 1-45.

3-19

Ph. Agree, The Dynamic Structures of Everyday Life, Report No. AI-TR 1085 (Cambridge: MIT Artificial Intelligence Laboratory, 1988); R. A. Brooks, "Achieving Artificial Intelligence through Building Robots," A.I. Memo 899 (Cambridge: MIT Artificial Intelligence Laboratory, May 1986); idem, "Intelligence without Representation," MIT Artificial Intelligence Report, 1987.

3-20

Brooks, "Achieving Artificial Intelligence," p. 11.

3-21

D. Dennett, review of Artificial Life, Biology and Philos., 2 (1991): 23-25.

3-22

For an extensive discussion of the self as the center of narrative, see D. Dennett, Consciousness Explained (New York: Little, Brown, 1991).

3-23

"Le statut de l'inconscient est ethique." J. Lacan, Quatre concepts fundamentaux de la psychoanalyse (Paris: Editions du Seuil, 1979), p. 35.

3-24

John Rachman, Le savoir faire avec l'inconscient: Ethique et psychanalyse (Bordeaux: William Blake, 1986).

3-25

Jonathan Lear, Love and Its Place in Nature (New York: Farrar Strauss, 1990).

3-26

For more on the relation between this tradition of mind training and analysis and cognitive science, see Varela, Thompson, and Rosch, The Embodied Mind. In this section I borrow some of the ideas from chapter II.

3-27

Nagarjuna, Precious Garland and The Song of the Four Mindfulnesses, trans. P. J. Hopkins (London: Allen & Unwin, 1975), p. 76.

3-28

Keiji Nishitani. Religion and Nothingness (Berkeley: University of California Press, 1982), p. 260.

3-29

Chögyam Trungpa, Sadhana of Mahamudra (Boulder: Vajradhatu Press, 1986).

3-30

The classic exposition is by the Indian philosopher Shantideva (c. eighth century c.E.). See his Guide to the Bodhisattva 's Wily of Life, trans. Stephen Batchelor (Dharamsala, India: Library of Tibetan Works and Archives, 1979). For an extensive commentary and discussion of this text by a contemporary Tibetan teacher, see Geshe Kelsang Gyatso, Meaningfid to Behold: View, Meditation, and Action in Mahayana Buddhism (London: Wisdom Publications, 1980).

3-31

This translation is Robert Thurman's. For Hopkins's translation, see Precious Garland, p. 76.

M. Foucault, Histoire de la Sexualite, vol. 3 (Paris: Gallimard, 1986).