Mind in matter

Version Dec 4, 2005

© Ton Rullmann, Utrecht Contact: mail@tonrullmann.nl

Preface

I have wanted to write this paper for many years. Being trained as a scientist and specializing in simulations of biological systems, I became intrigued by the capacity of these systems to organize themselves. Would it be possible to find rules for how this organization occurred, and could these processes be simulated in the computer? My own work developed in a more application-oriented direction. In the work of people like Prigogine, Kauffman, Ray and many others exciting vistas were explored. Partially through these studies I became more interested in philosophical questions. The core theme probably was: what is life? Somehow it became increasingly important to understand what a living system actually is, and whether it can be mimicked by a machine. These questions became linked with a certain unease about our place in the world. Where in a mechanical picture is room for freedom, beauty, wonder and awe? What can be said about the meaning of life, of *my* life, and of *our* lives? Can we find meaningful ways to talk about our inner life, that are still connected to the scientific ways of studying nature? In other words: can we look upon the biological machines that we are, and simultaneously (rather than alternatively) see and appreciate our inner spiritual richness? Can we be mind *and* matter?
[Note: I don't try to distinguish here between concepts such as mind, spirit, soul, consciousness. Suffice it to say that in using the word "mind" I do not refer to the brain or to intellectual activity as such, but to a rather more ephemeral quality of being, that somehow expresses a unique, personal core. Obviously more needs to be said.]

In the past few years many people have written about quantum mechanics and its implications for our image of the world. I certainly do not want to argue with these viewpoints here. Yet I have felt for a long time that the restoration of unity in our world view and the re-celebration of connectedness (if you wish, the re-enchantment of the world) does not have to be based solely on the strange and wondrous form of mechanics upon which this world is ultimately based (but who knows what else might still be discovered). I sincerely think that many of the problems we are faced with, do not stem from using the wrong kind of mechanics, but from misunderstanding the nature of the mechanical world picture. Also it seems to me that "meaning" and "mind" have a lot to do with the context in which meaning is assigned. Meaning does not exist in a vacuum, nor does mind. They depend on pre-existing notions, ideas and minds that set the stage for a new phase, and on interactions with the real world that shape, destroy or qualify meaning. Our knowing of the world is much less a passive obedience to eternal, mechanical laws, than a process in which we actively *assign* meaning to phenomena. I will argue that this active process is not a uniquely human property, but is shared by all complex systems.

I doubt whether the following ideas are original in any sense, yet the composition is certainly my own. I am greatly indebted to the work of Ilya Prigogine, Stuart Kauffman, Tom Ray, Ken Wilber, Itzhak Bentov, Ervin Laszlo and Gregory Bateson, whose books inspired me and deserve much more attention than I have given them here. Intentionally I left out all references because I needed to stay with my own flow, and feared I would become stuck in scholarly argumentation if I started looking up quotations and re-reading book chapters. I loved the experience of writing freely (but, I trust, not incoherently), following the meanderings of my mind, and often not knowing whether I travelled old or new roads.

In the following some lines of thought are unfinished. These appear in square brackets.

From the ashes

Resurrection. In a closed system information has to be destroyed for something new to appear. It is the eternal cycle of life and death. But think now about a small seed on the floor of a burned forest. Amidst the ashes it is sprouting new life. It grows roots and a few tiny leaves, a small plant is emerging. Then the leaves become bigger, the stem becomes longer, more side-shoots appear, which gradually grow into stronger twigs, each bearing a bunch of leaves. The process continues and a whole, complex structure emerges: a tree, with multiple branches, supporting hundreds of smaller twigs, and tens of thousands of leaves. The tree is covered with bark, channels in the inside carry nutrients upwards from the roots, sunlight is captured by green molecules in the cells of the leaves. Many layers of organization together form a huge, living structure. The biophysicists tell us: life is far from equilibrium. Order can be created out of chaos, if a stream of free energy is flowing, and a little can be tapped from it to keep us away from the death of boring equilibrium, the uninteresting, bland average. OK. So far so good. We can sort of understand this. Isn't it a fact of life that all things have a tendency to wear down, and slowly go to pieces. We all have to spend energy (free or not), to go against the accumulating dust and cobwebs. But how is it then that life seems to be creating so much "newness" and structure and complexity all the time? Think about this seed growing up from the ashes.

It is easy to say: oh, it's the DNA of course. The seed has all the information about how to grow into a tree. First of all, that's an amazing conclusion! The information contained in the little seed has the power to organize tons of carbon from the air and water from the earth into this complex shape of a tree. Truly a case of mind-over-matter; or let's be a bit more modest for now: information shaping matter.

But more is going on. Because the information is not complete: there is not enough DNA to specify every little bit of the tree. The seed does not contain a miniature copy of the tree. The exact shape of every leave and branch, down the last barely visible undulation of the surface, has not been pre-fixed. But what about that slightly bigger bulge? The distinct shape of the leave tip? The "ovalness" and the speckles? It is obvious that some of these features are encoded in the gene material, because we recognize them as being "characteristic" of the tree species, so almost by definition, they must be in the DNA (although I'm not even sure this is true in all cases, at least not in the direct sense. Because: how can shape be encoded in linear DNA). Anyway, the boundary line is fuzzy: to some extent, the properties of the tree are laid down in the original information contained in the seed, and to some extent they are determined by the quirks of growing and living. But is hard to define where the boundary is. All trees of one species differ somewhat, and a few trees may even be quite a lot different. Yet the information they all had at the beginning of their life was the same. And we are able to recognize the similarities between the different individuals enough, so we can make a description of the species, i.e. a summary of what makes an oak – or rather, what makes an oak different from a chestnut (because often it is easier to describe things by comparing them to other things; but I am digressing). We may be able to summarize how an ideal oak looks like, but we cannot describe a real oak: there's no end to the amount of detail we would need, the number of differences between the ideal thing and the real example is literally countless. So not only has information conquered matter, but information has also multiplied itself, and it has done so in a way that can not be copied exactly.

[Objection 1: look at proteins, not shape. This only shifts the problem: what is a tree? Focus on processes.]
[Objection 2: not all details are important. But how do we know which ones? Differences between ideal and real are not very important. But how do you define the ideal? It seems possible that the ideal shifts, e.g. smaller shape in a colder climate zone. Maybe the genome mutates. Do we know it, until multiple changes have occurred and the species have separated? We cannot define species, except in a self-referential way: a species is that which does not breed with other species.]

The point is that life creates order (information), and is "driven by" [need to define that] information (and needs energy and materials). Secondly, it seems impossible to precisely delineate the information contained in a living system, because there are always variations, and we need patient observation to deduce what is going on, and even then we have to invoke self-reference to come up with an answer: this is an oak tree because it looks so much like other oak trees; this is a C. elegans nematode, because its genome is the same as that of 10 other C. elegans critters, except in 1013 places but I know these are just point mutations; this is the European house fly, because it does not mate with the American house fly.

Fit-fitter-fittest
[To be written: evolution is a powerful concept; but is it possible to evaluate the fitness beforehand and make predictions? Isn't fitness highly dependent on the interaction between a species and its environment, which is changing as a result of evolution?]

The devil is in the details
[To be written: complex systems frequently show near chaotic behaviour, yet can be characterized by a number of stable states (attractors), i.e. certain features are recognizable and return from time to time, yet the overall dynamics is nearly impossible to predict. The type and structure of the attractor states may depend on global features of the system (e.g. the topology of the interaction network). Some features may be very important (i.e. can not be changed without breaking the system), other features may be more change tolerant. Complex systems, such as living cells, are a curious mix of rugged features, great plasticity, and near-chaotic phenomena.]

Scope for freedom
What is going on here? It seems that living systems have too many variables for us to ever fully "know" them. Yet, they are living and thriving, so they appear to know what they are doing! Could they be more intelligent than us?

There seem to be two properties by which we can describe complex systems. One is related to the *magnitude* of the complexity: the size of the system, the number of connections relative to the number of entities, etc. The other property measures the *degree of fuzziness*: the amount of arbitrariness, the lack of knowledge, the magnitude of the unknown aspects. [This is dynamic, may change over time, depends on our interaction with the system].

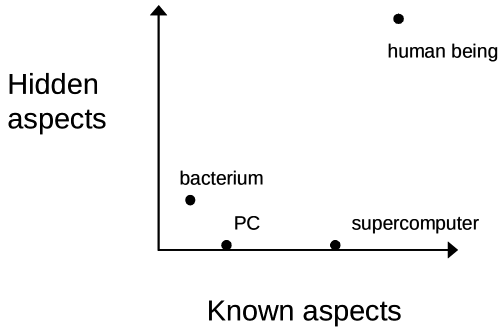

One could say: the first axis represents organization, the second represents self-organization. The distinction is somewhat similar to that between energy and entropy, but it's not the same. Rather, I try to distinguish here between an "external" and an "internal" aspect of information. The horizontal axis represents the known organization of the system, e.g. represented by its entropy. The vertical axis represents the "hidden" complexity, the internal freedom the system has to respond according to its fancies, the degree to which it is outside our control. There is an element of inner richness, of many layers of meaning.

> *A man walks by the seaside. It is sunny, and a light breeze touches his skin. The sun warms his body, and the wind cools it down. Perfect, dynamic harmony. Suddenly the man sees a crab coming out of the sea. "Hey little crab, what are you doing? Don't you know this is a dangerous place? The seagulls may see you, and pick you up for lunch." Only slightly perturbed – because he does not understand the man, who unfortunately is a foreigner speaking a strange language – the crab continues going about his crab-business. The man comes closer, wanting to…. Yes, what exactly does he want to do? Typically human to be so mixed up. He wants to warn the crab, but hey isn't that ridiculous, so he starts to feel a little bit annoyed with himself, yet is attracted to the fascinating, if slightly grim appearance of the little arthropod. Blissfully unaware of the ruminations of this undoubtedly larger consciousness, the crab decides to explain to the man in his best crabbese that he really is on an important business trip and cannot be disturbed. He quickly moves towards the man, and starts waving his claws. To his surprise the man gives a shriek and hurries away.*

The self-aspect of organization is the part that is not so easy to characterize from the outside. It represents the context in which signals are received, interpreted, and answered. It represents the system's freedom to respond. It represents its history and determines its uniqueness. In a statistical sense we may know the response in advance. That is, if we have done "sufficient" experiments, and the system is "sufficiently" predictable, we will be able to say what will happen, at least in some average sense. But we cannot know all the details.

Consider a fir tree, growing cones. Late in the year the cones drop to the forest floor, at least most of them. The force that propels the cones is gravitation, the pull of the earth (actually, the mutual pull between the cone and the earth). The force that holds the cones in place are the attractions between the large molecules in the stalk connecting the cone to the mother tree. As the stalk tissue shrivels and shrinks, this force becomes weaker, and at a certain moment the gravitational force wins, and tears the stalk, letting the cone move towards the earth. So far so good. The tree just knows how to arrange that most of its cones drop down at the right time of the year. We cannot predict, except through experience, when this moment will be. We can also say: evolution has fine-tuned the cellular metabolism in such a way that the cells of the stalk respond to the changing temperature in the autumn and slowly die. However, this statement (which is a gross simplification) does not tell us anything more about the exact moment than simply saying "the tree knows how to do it". Maybe we can develop a more detailed theory. We can set up a calculation of the forces within the stalk, taking into account how the stalk tissue is built from various cell types and fibres. We can simulate the life processes of a cell, and simulate how cells die when it gets colder. We can calculate how that affects the strength of the stalk, and what the breaking point will be, given the weight of the fir cone and the strength of the wind pulling at it. Still, we won't be able to predict precisely when the fir cone will drop! For one, we do not know the weight of the cone. It's not trivial to weigh it without removing it from the tree. But hey, we are scientists, so we will find a clever way. Secondly, we cannot predict the wind force with any great accuracy. And finally, breaking of the stalk is a molecular process; we need to know the precise location and shape of all molecules inside in order to accurately predict the breaking of thís particular stalk.

But does it matter? Do we need to pinpoint the moment to the last second? Probably not. Probably the precise moment does not matter, unless you happen to be below the fir cone when it falls. Maybe the tree cares. (For example it might matter whether the cones fall all in one day, or in the same week, or the same month. Or maybe it is important that some stay on the tree. I just don't know). Maybe the earth cares. But let's not go into that now. We can say that we sort of understand what is going on physically, and that all that is needed to achieve more complete understanding, is reducing the uncertainty by doing more, and more detailed experiments. How complete our understanding can be is a matter of speculation, but in this case not a matter of great practical importance.

But wait. This was only one tiny example. One insignificant (or so it seems) phenomenon in the natural world, that we can only approximately know. One decision taken by a living system, which we can more or less anticipate, although not completely. Life as we know it, is full of such imprecisions. A long chain of decisions is being made all day long by each and every living organism on the planet and its surroundings. Probably the great majority of these "points of change" cannot be predicted with absolute rigor. Moreover, all these changes that organisms are going through, are affecting each other. Instead of a single chain of events, we have a tangled web of events, with myriad strands of cause and effect running in all directions. Can we ever say we really know? Can we ever allow to forget about contingency and individuality?

The recipe that science has developed is called reductionism. Scientific models operate within a specific domain, and at a specific level of detail. For example, in chemistry effects from nuclear quantum physics do not play an important role, and in physiology many biochemical details are ignored. In scientific research a carefully selected part of reality is isolated, investigated and represented by a theoretical model. Taken together, the scientific model domains form an interlocking hierarchy, mimicking the hierarchy observed in nature: quarks form elementary particles, from which atoms are built; atoms form molecules, and many molecules together build supramolecular complexes and cells; layers of cells form tissues; tissues form organs, and organs are what organisms are made of (can you supply a better definition of organ?). At each level we can more or less forget about the details of the underlying level: only its average properties count. The method of science is to aim for predictability. Taken together with the hierarchical structure, this naturally leads to emphasizing at every level that which is constant or which repeats itself, and de-emphasizing individual variations.

Several problems present itself. In the first place: what happens at the boundaries of these scientific domains? Do the domains really nicely match and overlap? Is the cell as described by a molecular biologist the same as the cell studied by a physiologist in his animal models? Does the hierarchy of models really exist as a close-knit unity, or are there gaps and mismatches? Secondly, and this is related to the previous problem, the hierarchical relationship is dual: the units at the lower level together form a new unit at a higher level, but at the same time the higher level also partially determines the lower level. For example, a cell suffering from diabetes does not exist. Diabetes is a disease of the system. Yet, somehow the cells must be affected as well, because the cells and the connections between them together *constitute* the system. So the system behaviour called diabetes is somehow showing up in the behaviour of some of the cells and some of the inter-cellular connections. So we are forced to enlarge the scope of our inquiries. The third problem is the order-out-of-chaos problem. How does something new, beautifully complex arise, and where does the direction come from? How can the emerging system determine its own course, and like baron von Munchhausen pull itself out of the morass? Various answers have been provided: non-equilibrium physical chemistry shows the conditions in which this is possible, in the language of thermodynamics i.e. requirements in terms of energy and matter fluxes. Many types of computer simulations have shown that relatively simple sets of relationships between variables, can give rise to surprisingly complex patterns, often depending on a critical balance between several parameters governing the shape and strength of the relationships. In this case the emerging complexity is not a mystery but built into the equations. Nevertheless the fact that complexity cán emerge may be a riddle in certain cases. How is it possible that the physical constants have precisely the correct proportionality for life to arise in the universe? And how did life evolve so quickly, was there really enough time for it to happen by chance? These and similar questions have led some to propose the existence of an information field, that may actually guide dynamical processes. For now I want to look at the order-chaos problem from a different perspective: how do we define order?

The meaning of random

Consider a string of characters, such as: jfjewfjfkafkv +m efslcdeg. What does it mean? There might be an intelligible message here, hidden by a particular encryption scheme. Or it might just be rubbish. Without the context (who wrote it? for whom was it intended? are there more messages like this, etc.) it is impossible to tell.

Now consider 0000111100001111 and 0110100101010010. The first one looks more ordered, the second one more disordered. From one perspective, the first one is more special, because it is non-random. The second one looks like one of many; another one of its kind is 0110100101001010 (spot the difference). They look like a random sprinkling of 1’s and 0’s, and one could be exchanged for the next, as long as the 1’s and 0’s are more or less equally spaced. They are all shades of grey. From a different perspective, however, the situation is the reverse. The ordered sequence is the more simple one. Look, it can be written much shorter: 4x04x14x04x1. Or even shorter: 2x(4x04x1). Try that for any of the other sequences, and you will not gain much. So, the “random” sequences are more special. It all depends on your focus: the ordered sequence deviates more from the average of all possible series of 0’s and 1’s, and there are only a few of such sequences; but the unordered sequences (of which there are many more with the same degree of similarity to the average), are each very specific in themselves.

What is the reference point here? Is it how much the sequence deviates from the average of eight 0’s and eight 1’s? In that case all listed sequences are pretty alike. Is it about local structure, i.e. how much equal symbols follow in a row? Then the first sequence is indeed the more special one. Or do the symbols at position 12 and 13 have special meaning? In that case the first and third sequence have something important in common! Without knowing the context, that is without knowing how these sequences will be interpreted by the machines that ingest them, we cannot say which one is more special or unique. That is: information can not be defined without context, and the context is supplied by the whole machinery that processes that information. Our visual information processing system can easily identify certain levels of contrast, but has problems with others (e.g. one 0 in a sea of 1's will stick out, whereas detecting the difference between the second and the third sequence above is more difficult). Other processing systems may well work differently!

Similarly, the science of statistical mechanics derives the collective properties of large assemblies of molecules such as a gas or a crystal, by investigating how likely it is that molecules arrange themselves in patterns that deviate from the average. This works fine for systems that are more or less homogeneous. When however the starting situation is highly non-homogeneous, such as is the case for any living cell, and when in fact the processes that go on inside the cell highly depend on the relative positioning of the various parts (think about the structure of the nucleus, the cell-membrane, the ribosome etc.) it becomes much more difficult to set up a statistical theory.

Process

A system progressing to order is defining itself. We attach meaning to it, by interpreting it in a certain context. We define it by comparing it to other systems and noting similarities and dissimilarities, and by noting the parts it appears to consist of. Frequently we describe the relationships between these parts, and observe specific spatial ordering, symmetries etc. When a temporal quality is added the description becomes in effect a report of the ontogeny – perceived or imaginary – of the system. Our knowledge is couched much more in relational than in absolute terms. Objects are constantly compared to objects we encountered previously, and qualities of the previous encounter tend to colour the present experience. [In a way this is precisely what science is about: making a systematic and consistent inventory of phenomena, linked to a precise inventory of the conditions under which they occur; and these conditions in turn are linked to a consistent inventory of other phenomena, etc.]

This is the human brain level. The nerve system may be operating in much the same way: It is geared to detect differences. Constant signals lead to desensitization. And how is this at lower, cellular levels, and for other, non-neurological systems? A humble E.coli, for example, is sampling its environment using its receptors on the outside of its cellular body. When one of these receptors repeatedly sniffs the chemical it is designed to detect, and which represents either food or a toxicant, it influences the motion of the bacterium in such a way that – although it is moving erratically – on average the bacterium swims toward or away from the apparent source of the chemical. Over time, however, the effect erodes away, i.e. the E. coli has a form of chemical memory based on differences that helps it to move where it has to go in order to survive.

Maybe we can say: physical interactions are being processed (integrated, combined) into information. Some of this information is recognizable from the outside as a "change of state" of the system, and some is hidden inside, as random fluctuations or as meaningful variations that may affect the behaviour later on and that we have not yet learned to discriminate. And again, what is meaningful or not, may only become clear later on, depending on the development of the system (is the fluctuation driving the system to a new equilibrium) and the systems this particular system is interacting with or is a subsystem of (the whole which is determining – or should one say, colouring – the parts).

Summarizing: we as human beings are living systems that are constantly acquiring information, comparing it to previous data, placing it in the context of what has gone before, and thereby assigning meaning to the events that shape our lives. It seems to me that the same is true of other biological systems.

> *The dolphin jumps out of the water. Imperceptibly he changes the world. The world will never be the same anymore.*
> *The dolphin jumps, the water flows from his back, and stirs the ocean. A trickle of water flows in a new direction, carrying some bacteria that will infect a cockle closer to shore. Who says this doesn't make a difference? A fisherman comes along, picks up the cockle along with a score of its relatives. He prepares his dinner, and gets sick. Who says ...*
> *The dolphin jumps, and jumps again. His actions attract a female dolphin, who is admiring in the dolphin-way his strength or playfulness or whatever it is that attracts her. They mate, a new dolphin is born. Who says this new creature is not carrying the seed of a future generation of dolphins that are particularly fond of, and capable of jumping out of the water? The point is: we don't know, and never will.*
> *By reading this, you have been changed. Who says it will not affect you five years from now in an unexpected way?*
> *The other point is: the same event is interpreted in many different ways. For the cockle it is a matter of currents, of opening up and closing, of being able to feed from the current, resist the onslaught of micro-organisms etc. The dolphin creates a wave, and the cockle is dealing with the consequences, the risks and opportunities. It deals with it in the best way possible for it in its present state. For the female dolphin it is the spectacle of the male, the sensory images, the reflexes generated in her body, the impulses and desires set up through an ancient and personal system of nerves, neurotransmitters, hormones, receptors and so on. She deals with it according to her present state. For a human it is similar. Yet, since our body chemistry and brain wiring is different, we will not feel the same impulses. But we may feel the desire to swim with the dolphin, to be close to him. Maybe this is even not all that different from what the female dolphin experiences.*

Life is not just about being, but equally – or possibly even more – about becoming. A useful definition of life may actually be: those processes of becoming that never end. In the classical scientific view, processes of becoming can be described by fixed laws. It now appears that such laws only apply to idealized entities. Processes of becoming are highly dependent on the internal states (which we imperfectly know) of the things that are. These processes are generating “newness” in the world, simply because they appear new to us, or to any other bystander interacting with the “things that are and things that become”. They appear new because we have no way to determine with certainty in a finite life time that they are not, and before we die they have changed us.

There seems to be something very fundamental here. The fact that we do not know completely the present state of the world. The fact that – even if the world were deterministic, which thanks to quantum mechanics, it is not – we would not be able to reduce that uncertainty. In fact, constant newness is being created. The fact that we as beings with a finite mind and finite tools, can only map the world to a finite representation. Therefore, unavoidably, we project our fancies and limitations into the world. We assign meaning, and obscure information while doing so. Is this random, or does it have meaning? We only have a finite time to decide, we act upon it, and change the world including ourselves – irreversibly.

Who is “we” in the previous paragraph? Could not the same be said about any biological system?

Conjecture

Two things have been said: information about biological systems is incomplete; this information has a relational aspect. I want to speculate a little bit further.

Incompleteness is related to finiteness. With growing system size, the number of elements of the system does not keep up with the number of possible connections between those elements. Fuzziness sets in. It becomes difficult to differentiate between all possible configurations. We tend to perceive critical systems with attractor states. Otherwise we do not recognize them as a system because we don’t recognize the system properties. This is related to the next point.

The relational aspect means that systems respond to differential signals from the outside. Differences below a threshold are ignored, differences above a threshold lead to a change. This corresponds to digitizing the signal.

So the existence of a finite set of stable “macroscopic” states emerging from a multitude of “microscopic” states is paralleled by the digitization of the signals that connect the system to its environment. Any coding system must be finite, otherwise the business of decoding would never be completed. At the same time this process of mapping a large number of smaller units to a smaller number of bigger units, creates the potential for evolution: constellations may change, and new groupings, corresponding to new stability conditions (in other words: new meanings) may arise. It is tempting to mention here Gödel’s theorem, stating that a formal body of propositions is either incomplete or contradictory. If mathematical-logical systems have this fundamental property, what does this tell us about our knowledge of the world?

Mind

Conclusion: Life creates information. We cannot delineate all of it. We cannot fully interpret it or predict it *a priori*. Meaning is intricately linked to interactions, i.e. to a dialogical process between complex systems.

Let’s return to the diagram representing the two aspects of organization. The horizontal axis represents the externally visible organization, the agreed-upon sets of meanings, the quantifiable properties and interactions, in other words: the realities of physics. The vertical axis represents the hidden organization, reflecting the unknowability and non-ideality of the system, capturing its individuality and history, the scope for growth, in other words: the potentialities of mind. In this view the mind of a system is the bearer of its uniqueness. It is that which interacts, and decides, according to its partially obscured logic. If the reactions become totally predictable, one could say the system is behaving totally mindless, i.e. like it has no mind of its own.

In fact, physics and mind can not really be separated. How else can mind manifest itself than through interacting in the physical world? A mind that does not open up and connect to other minds, is literally a closed mind, in which nothing comes in and from which nothing goes out, a barren mind. In reality a constant interplay between “inward” and “outward” is going on: mind expressing itself in the world, the world imprinting itself on the mind. Integrated interactions become information. Mind provides a context for interpreting the information, and generates change. Change may be an adaptation (mind projecting another aspect of itself), a synthesis (systems combining to form a super-system) or a dissolving (systems breaking apart).

This may seem rather abstract and obscure. What I am trying to argue here is that at a fundamental level the world is more akin to us than we imagined. It is also more wonderful than we imagined. We have grown accustomed to a view of the world as a mechanical place, operating according to laws. At the same time, in particular through the machines that we have learnt to build, we have grown weary of mechanical things – even though we are also intoxicated by the might and pleasure they bring – because they are so mind-bogglingly mindless! Machines are dead things. And the more the world is exposed by science as operating by mechanical laws, the more we feel isolated. And even worse: we ourselves are part of the same, mechanical world. Our bodies, our brains operate according to those same natural laws. So who are we then? Are we just Darwinian machines, and are mind and soul empty concepts?

But I think that this picture is based on false ideas. The picture of a mechanical world is far too simple. To begin with, types of mechanics may exist that allow for a far greater degree of interconnectivity than we normally imagine, such as explored in the A-field theory. What I have discussed here is something else: real systems at all levels have intrinsic meaning and knowledge. Mind is an integral part of being, at every level. This is because the structure of knowledge is deeply connected to the structure of the world through its context-dependence and hierarchical growth. Form builds on form, thought builds on thought, form shapes thought, thought informs form. There is a fundamental duality here, maybe comparable to the particle-wave duality of quantum mechanics: a duality between form and thought, between physis and psyche, or physics and mind. Knowing presupposes context, and context is determined by process i.e. interaction between observer and observed. Interaction leads to change, and change is represented by altered knowing. It is an eternal dance, both upward and downward – systems emerging and systems dissolving – as well as outward and inward – systems expressing themselves through their physical aspect, and systems acquiring an inner complexity and individuality.

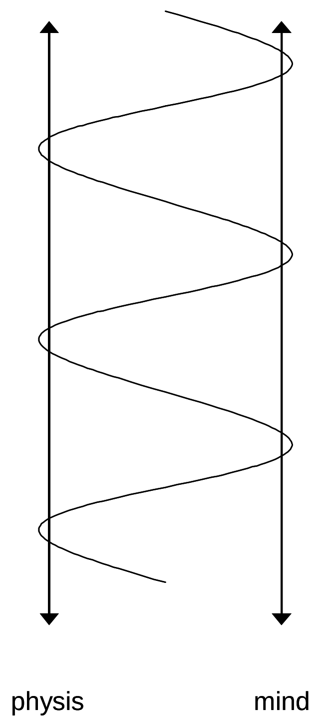

Natural science as we know it is limited to studying physis: the repeatable, the externally visible, the non-unique, the a-historical. We can objectify and generalize what we see, but only at the loss of individuality. We can make models of phenomena, but a model is not the thing itself; the best model is the thing itself.

I have spoken mostly about biological systems. But there is no reason to treat non-biological systems any different with regard to the mind-in-matter issue. Why can't a rock have a mind? Or mother earth? In fact, even electrons and atoms seem to have a history that is undeniably their own. Also man-made artefacts can have, if not a life, than certainly a mind of their own. Don't we all talk to our machines from time to time, in particular when they are behaving badly and we have no clue what is going on? Computers are prime examples. Now you may doubt whether a PC really has an individuality that could be called mind, and whether it's not all a matter of tiresomely predictable transistors, wires and bytes. Yet, is there really anyone on the planet who can say she truly understands the complete interaction of all the hardware and software components of a PC, and not just as an abstraction – the prototype design of the PC in the engineer's manual – but as the real thing, model number YN09245? I admit the margin for individuality in this case is quite small. We have constrained our machines so much that they can only deviate in the direction of error. Your average household appliance is not going to accidentally produce a more beautiful result than intended. This means they have a rather poor mind, if one could call it a mind at all. Maybe machines would be more interesting if they were allowed more freedom. Somebody once proposed to write a self-learning word-processor program: it would interact with the user, and learn what it had to do from the user's feedback. Actually, good feedback would keep it alive. If the program proved not up to its task, it would die. Wouldn't such a program have a mind of its own? Now think about the Internet. What kind of a system is that?

Two other concepts should be mentioned, although this is not the place to explore them fully. Maybe some other time. The first concept is "consciousness". Right now I would say that consciousness is mind that is aware of itself. Most minds are probably not very much aware of themselves, i.e. their information does not contain a concept to describe itself. This needs further discussion; other people may view the relationship between mind and consciousness differently. The second concept is "life". I have given one definition above: processes of becoming that never end. Other definitions equate life with stemming of biological origin (which is akin to systems having a carbon- and oxygen-based metabolic chemistry) and/or being able to reproduce itself (although sexually reproducing animals do not produce exact copies of themselves --- it is so hard to rigorously define these things!). Based on the arguments given above I do not think that mind is restricted to life according to any of these definitions. The old distinction of animate vs. inanimate may come closest to the difference between possessing a mind and being mind-less. This also brings in connotations of spirit. Can we sensibly talk about nature spirits and devas as representing the minds of the various systems around us in nature? Can we speak about a hierarchy of minds or consciousnesses, that parallels the hierarchy of natural forms? I leave this for another time.

Connecting

Can we know other minds? That depends on what is meant by "knowing". Science offers one road to knowledge, which as we have seen, is based upon extracting information in the form of physical, measurable interactions. Mind is viewed in, and to some extent pushed into, its outward aspect. Some hidden context is lost, some subtle meaning disappears. Never mind, one could say, that's OK if we know what we are doing. Are there alternative ways of connecting to minds? I think so. These ways I would collectively call "empathy". They are about placing oneself in the other girl's position: to transpose one's frame of mind, i.e. begin to approach and appreciate the context that is meaningful to the other. At some level it might become identification with the other. Empathy is about feeling my way towards the other. I need my senses, my reasoning, but most of all my intuition, which I would say is the capacity of my own mind to find within its own large store of potentialities some form of common ground with the other, some synchronicity or correspondence. It is never going to be a definitive knowing; rather, it will be changing all the time. It will adapt to circumstances, and most of all it will be uniquely personal. But that does not mean unintelligible. We may have to find out how to be empathic. It could well be we have to start listening to parts of us that we did not listen to before, but actually may know the job better than we do. Our feet know about grounding and running, our hearts know about love, our bodies know about pain.

In some sense language is the ultimate structure of mind: words can embody whole worlds. This may be an anthropocentric, or even cerebral-cortex-centric view of mind. Yet, we can at least try to listen to the secret whisperings of starfish, cats, trees, planets, our body and celestial beings as well as of our fellow human beings. And also there is symmetry in referring to language as the structure of mind: just as the study of physis (i.e. science) requires language to formulate its findings, so does the study of minds (i.e. empathy) require real-world action to amount to anything: form and thought cannot be separated.

What are the practical consequences of all this? Simply this: in areas such as medicine and ecology, can we afford to stay within the physis-only paradigm, and not include mind – i.e. individuality, spontaneity and creativity – within our studies and practices? Can we risk to only utilize the narrow window of physically manifest interactions, and assume that it is not worth minding the array of still hidden potentialities? Can we evolve practices that develop mindful attention into dialogue, and from there into co-creation of new realities that truly honour the diversity in the world and allow all of us to be participant creators?