The Evils of Lego™*

Ian Grant†

18 May 2010

There are two kinds of boys in the Western world: those that play with Lego™and those that play with Meccano™. The ones that have lego become accountants and computer programmers and the ones that have meccano become chemists, doctors and engineers.[1]

The difference between lego and meccano is this: A lego brick is essentially more abstract than a piece of meccano. This is because lego bricks are rigid[2] and meccano pieces are flexible and malleable. A lego brick can be characterized exactly by its colour and shape which is one of a finite set for which there have ever been moulds in the lego factory.[3] If there is a space of a certain number of studs, then only a brick that size or a smaller brick will fit it. Given a set of lego bricks there are only a finite number of constructions one can make without irreversibly altering the individual pieces by melting or cutting them. The constraints on the construction are rigid and partly determine the possible constructions.

A meccano piece on the other hand cannot be so easily described. There are of course only a finite number of meccano pieces in existence and they each fall into one of a finite number of shapes and colours. Each piece has a certain number of holes through which bolts etc. can be threaded, but the relation between these holes is not fixed as it is in lego. I had lego and I remember seeing a friend's meccano car he and his father had built. I was horrified. They'd bent one of the pieces to make an engine cowling! Once you bend the metal or plastic pieces the holes no longer bear the fixed, discrete spatial relationship and the number of possible constructions with a given set of pieces is not constrained to anything like the degree of a lego construction. It is essentially infinite. It is also dynamic: for example, in a meccano construction a bent metal piece will act as a spring under compression and the metal will expand and contract with heat to a much greater degree than the plastic that lego is made of. I am sure one could construct in basic meccano a device which responded to temperature of its surroundings. To do this in lego you would need to use a purpose-made temperature sensor brick.

*North Americans should read this as *The Evil of Legos*.

†e-mail: `iang@pobox.com`

[1] Of course I do not mean this seriously, nevertheless I would not be at all surprised if there was in fact a positive correlation like this.

[2] I am referring here to the old-fashioned 'classical' lego bricks, not the numerous modern varieties.

[3] My friend Brian Jones told me once that the tolerances to which lego bricks are made are within one hundredth of a millimeter in order to maintain their grip with age (they need to deform, but only to a minimal amount under which the the material is essentially perfectly elastic and recovers its original shape exactly when released). The Lego engineers probably all played with Meccano when they were kids!

When I grew out of lego, I made electrical circuits, but the lego mentality obtained and the way I made my electrical circuits was by putting together basic building blocks. I couldn't develop any intuitive understanding of transistors, capacitance or inductance because I had no dynamical sense of the systems. This is what is absent in lego constructions. The complexity of the electronic circuits I could make was therefore limited by the number of parts I had and this was never enough. Not surprisingly I quickly got bored with my limited ability to do electronics. This was around the time that the very first personal computers appeared. Computers are the ultimate electronic lego blocks. Though finite, the computer could be re-programmed again and again. Maybe, I thought, just maybe, one could program a computer to program itself and then the problem of running out of parts would be solved and I would be able to build something of limitless complexity: something alive. This imagined unlimited potential was very beguiling, but I had to go to school.

I can still recall the smell of the school laboratory where we put together a sodium vapour lamp, a collimator and a soot-blackened glass slide on which we'd scored two thin lines. Then using a little microscope we measured the width of the interference fringes to find the wavelength of the sodium emission lines. Interference effects could be explained perfectly if light were a wave in a medium, but it's not. The velocity of light is the same whatever the state of motion of the observer measuring it and so there seems to be no medium for it to be a wave in. Furthermore it seems as if the energy carried by light is in discrete quanta so that a very dim light is not a uniform low level of illumination. Instead it is a series of randomly occurring intermittent flashes of a characteristic brightness according, not the intensity, but to the *colour* of the illumination! If light were composed of individually extant particles then it seems as if these particles would have to conspire with each other in the most intricate and elaborate ways in order to produce the interference phenomena we could see with our own eyes. A conspiracy on such a scale could hardly be called an explanation.

Yet light is almost the only way we know about anything. The whole Universe is revealed to us almost exclusively through the medium of light in one form or another; light in the form of electromagnetism is what holds the substance of all sensible things together including our own bodies, light pervades everything and *yet we don't have the faintest idea how it can possibly exist!*

The fact that this was so openly accepted as a mystery was hugely appealing to me. It was like being shown the entrance of a cave and being told that no-one knows what is inside here, or how far it goes. But this wasn't just a neighbourhood myth, it was a mystery on a huge scale. *No-one in the whole world knows the answer to this question.* It was obvious to me that there was not much point in thinking about anything else until this was solved.

However I was seduced by the seemingly limitless potential of computers. Sitting in front of a computer is like being given an infinite supply of white paper and a box of crayons with an infinite number of colours. The possibilities are limitless and exceed anyone's imagination. This goes beyond mere analogy: I think it is a fact. The universal nature of abstract computation means that a computational process can represent anything that it is possible to represent. Understanding computation and information, I thought, might be a way to solve this problem of the nature of light. Now twenty-five years later I think I was right, but for reasons which were deeply misguided and totally wrong.

Light and optical phenomena in general are described very accurately by quantum mechanics[4]. The formal mathematical theory is supremely useful and superbly accurate: it works and in this sense it is true. The only problem is that it makes no sense; we don't know how to interpret it. Richard Feynman, who proved to be rather good at it, once said something to the effect that if anyone tells you that they understand quantum mechanics then all you really know is that you have met a liar.

My friend Adrian Kent, whose job it is to understand quantum mechanics[5], tells me that he and most other researchers in his field do not discuss the topic via e-mail. His experience is that it invariably leads to a severe reduction in the quality of his life. Now e-mail is just a mechanism for transmitting information, mostly words, from one person to another and the history of science and of mathematics shows that personal communication has always been a very important part of scientific consensus. The traditional peer-review process that constrains the published body of scientific writing is relatively slow and by its nature conservative. This is not a failing, it is completely necessary. Science would not constitute a body of knowledge if it was not coherent and coherence is necessarily conservative when there is a large body of extant knowledge. Personal communication on the other hand is much more rapid and the private context means that it is possible to tentatively explore ideas that are not fully formulated and would not pass even the most liberal reviewing process. Why then is it that the subject of the interpretation of quantum mechanics cannot bear direct discussion?

As any teacher will confirm, it is only through discussion and debate that we actually know that we know anything at all.[6] The process of examination of students is really an attempt to emulate this discussion on a large scale: the examiner asks a series of questions based on her understanding of the subject and the student answers them according to his. The person marking the exam then interprets the student's answer according to his own understanding and the coherence of the whole exchange is the basis on which the student is judged. At no point during the teaching or the examination do either the student or the examiner or the marker have direct access to the other's understanding.

What then is knowledge? It is not clear to me that any of us even has direct access to their own understanding of a subject: we merely have evidence for or against it by our experience in dialogue with others. Often I have felt that I understand something, only to find that in discussion my ideas are quite incoherent.[7] This experience is probably more common than a lot of people would be comfortable in admitting to, however I don't think it is anything to be particularly ashamed of because knowledge is in fact nothing more than the coherence of the different representations of it.

The absence of coherence in the different representations of the meaning of quantum mechanics is undoubtedly the source of the unease which we feel when we try to discuss it. It is bound to be difficult for honest people to discuss something if they each feel that they don't know what the other is talking about. But I think that this problem is itself a huge clue to the problem of interpreting quantum mechanics. This

[4]Strictly speaking, quantum electro-dynamics or *QED*. See Richard Feynman's beautiful book *QED: The Strange Theory of Light and Matter* for a short, accessible and quite spectacular description.

[5]But to his credit has never claimed to have succeeded!

[6]It is often said that the best way to understand a subject is to teach it to someone else and this is essentially the same observation.

[7]Yeah, yeah, yeah ...

is because any interpretation of a formalism must be in terms of something else we understand. That's what the word *interpretation* means: it is a re-statement of something in a language we already know. In Tarski's terminology [2], we describe the semantics (i.e. the meaning) of a formal language in a *metalanguage*. This metalanguage is one we take as implicitly understood. The metalanguage cannot be another formal language unless that language in turn has an interpretation in some other metalanguage of which we have an intuitive understanding. Thus we have established this fundamental fact:

Any interpretation is an act of judgement by some experiencing subject.

Now acts of judgement ought to be conscious acts, so interpretation presupposes consciousness. This is in itself nothing remarkable, unless one expects a formalism to be able to explain consciousness, because then there is a problem. If we expect a formal language like quantum mechanics to explain something, then it can only be in terms of something we already understand and since understanding presupposes consciousness, consciousness is not one of those things we will be able to explain with quantum mechanics.

By what right do we expect to be able to explain all phenomena in terms of something somehow more fundamental? In other words, why should we expect the Universe to be made of lego and not meccano? If the Universe is in fact made of meccano then we will find that the properties of things (the phenomena) do not reduce to the properties of the parts because the way things are assembled changes their properties: a flat piece of meccano has no particularly good aerodynamic properties but it could possibly be bent into a kind of aerofoil and function as a turbine blade. In this sort of construction there need not be a clear functional decomposition between the different parts: the bolt that holds the blade to the axle may also be holding the curve in the flexed plate to form the aerofoil. In contrast, the function of a particular lego brick is *logically* determined by the studs and the holes that connect it to the other pieces. In a lego construction all the engineering goes on in the factory. In biological organisms, why should we expect the abstract parts we identify to have clear functional decompositions like the machines we make out of lego? Typically they don't. Organisms are not made in a factory, they assemble and maintain themselves, so we should expect every part of an organism to have some ontogenic and metabolic aspect as well as its obvious functions.

Now I am not suggesting that the Universe is in any sense a construction, but what I am saying is that to assume that it can be explained by a process of deconstruction is to make some very strong claims about the properties of the composite entities within it: things like people and trees and rocks. The claim that is implicit in reductionism is that all phenomena are reducible to the properties of the parts of which they are composed, but this is manifestly not the case. Take a fluffy white cloud, for example. The cloud has the property of reflecting a great deal of the sunlight incident upon it. Thus its internal temperature is much lower than would be the temperature of the same amount of water lying in a muddy puddle. The reflectivity of the cloud is an irreducibly macroscopic property of the whole collection of water droplets of which it is composed[8] and the existence of those droplets is conditional

[8] One water droplet alone will not reflect the same proportion of incident light that falls upon it as does a cloud.

on some quite specific atmospheric temperatures and pressures. But those temperatures and pressures are themselves properties of the macroscopic arrangement of the water droplets which are reflecting the sunlight. So when we come to 'explain' the cloud we find we need to explain the presence of all the water droplets in that particular place and time. The explanation is entirely *contingent*, which is to say that it is specific to that particular cloud in that particular place at that particular time, and to explain the cloud we quickly find we need to explain the whole Universe. These are not the kinds of explanations that we can reasonably expect from physics which is a collection of universal laws.

If we admit the idea that consciousness is one such irreducible phenomenon then we immediately have to accept the irreducibility of all the phenomena of conscious experience. This is the philosophical position of *phenomenalism* which was first expounded by the physicist and philosopher Ernst Mach in the late nineteenth century. Mach held that we abstract the phenomena we experience from the continuum of our conscious sensation and that the task of science was not to explain, but just to *describe* these phenomena. Because any description of the world must be realized within it, said Mach, we have every reason to proceed as economically as possible.

According to Mach we should not attribute the universal laws to Nature itself, but to our descriptions of it. This is because everything in Nature is specific and happens for reasons that are essentially unique. The types of phenomena which are the subject of the physical laws on the other hand are abstract representations and it is only through our conscious judgement of the actual unique events as being of the type to which the laws apply that we can understand the laws as explaining the phenomena.

In the phenomenalist's view, because of the irreducibility of the phenomena, the objects of experience and the conscious phenomena are one and the same. There is no need, nor indeed any way, to explain the mysterious connection between objects in the world and our conscious experience of them. That is to say that there is no *mind-body problem*[9].

Phenomenalism has always been an element of some interpretations of quantum mechanics. Niels Bohr [1] was quite explicit about it, as quoted in [3]:

> ...the purpose [of our description of nature] is not to disclose the real essence of the phenomena but only to track down, so far as it is possible, relations between the manifold aspects of our experience.

and in the last few decades or so there have appeared — to my optimistic eye at least — some signs of an emerging consensus ([3], [4] and references therein). This consensus is to the effect that the ultimate physical reality described by quantum mechanics is not a world populated by interacting objects with some properties that explain their collective behaviour; rather it is a world in which we know things only by their relations to other things.

The most explicit along these lines is the relational interpretation of Carlo Rovelli [4]. According to this, the physical reality which quantum mechanics describes is nothing more than the statistical correlations between the states of interacting quantum mechanical systems. In the relational interpretation it makes no sense to speak

[9]Note this does not imply idealism — the view that everything exists only in some sort of mind — because mind itself is a phenomena arising in conscious experience.

of the actual state of a quantum system S except in relation to that of some other system, O, say, and then this statement must be made by yet another system P which has undergone some physical interaction with the combined $S + O$ system. This is very abstract: statistical correlations are not phenomena we directly experience, rather they are features of our abstract representations of ensembles of events. This explains why the strange Schrödinger-cat nature of some quantum states is not something that we directly experience: it is not real.

In the real world, every event is unique and occurs for its own peculiar reasons. In our individual direct experience we never have access to ensembles of identical events which could constitute a test of the statistical correlations. On what basis then do we claim that quantum mechanics is correct?

Quantum mechanics was not an idea that popped up fully-formed one day, and which turned out to be correct when tested against experiments. Rather it was cooked up piecemeal by dozens of people over a period of several decades in order to explain a particular range of observed phenomena. This is the reason for its essentially phenomenological character. It is also the reason why many of its predictions are correct: the phenomena it was put forward to explain were already being observed at the time. Nevertheless it has subsequently proved to successfully describe many more essentially different phenomena which were unheard of at the time it was first developed. Therefore there must be some characteristic features of the original set of phenomena which adequately represent all those which it has subsequently been tested against.

The reason we say quantum mechanics is accurate is that we represent the phenomena we directly experience as statistical distributions of their constituent quantum events, and these distributions are exactly the ones that quantum mechanics describes. For example, we see (with our eyes!) Young's fringes as a certain variation in the intensity of light on a screen according to the distance from the pair of slits in the glass slide. We represent this phenomenon as a numerical function of position on the screen. We can then normalise this function to get a probability distribution which we interpret as the probability of the constituent photons arriving at certain positions and, lo and behold, this is exactly the distribution that QED predicts.

Thus it is not the phenomena of experience themselves that quantum mechanics describes, but our representations of them as probability distributions. Now any formal symbolic representation is an interpretation and interpretation is a conscious act so all the measurement interactions described by quantum mechanics involve conscious experience in an essential way. Rovelli would not agree with this: he is quite explicit in stating that the systems that take the role of observers in the formalism are not necessarily 'conscious or animate in any special way'. However consciousness enters inevitably whenever we consider the *meaning* of the measurement interactions. This is because measurement interactions are always described as the establishment of correlations between systems, and the description is always with respect to some particular observer, so if it has a meaning it can only be described in a metalanguage intuitively understood by that observer.

If we think of consciousness as a property only of brains then we seem to have been driven into a kind of solipsism. This is what reductionism[10] demands. However if we don't accept the premiss of reductionism — that all phenomena can be reduced to properties of their constituent parts — then consciousness is not a phenomenon that

[10] Combined with the principle that interpretation is always a conscious act.

has any location in space-time, it is a fundamental condition of the whole Universe.

This is not as bizarre an idea as it may at first seem. Consciousness is a very abstract notion; in a sense it is the most abstract of all because consciousness is the idea we arrive at when we abstract from all the particular phenomena of experience. Consciousness is not something we experience: it is experience itself. The fact that consciousness is a condition of the Universe is just the fact that events are meaningful, independent of human culture and our knowledge of them.

Neither is this a new idea. *Pratītyasamutpāda*, the Buddhist notion of causality is described like this:

> When this exists, that comes to be;
> with the arising of this, that arises.
> When this does not exist, that does not come to be;
> with the cessation of this, that ceases.

This is one of the central principles taught by Gautama Buddha from around 500 BCE, but some element of it may be present in earlier Hindu philosophy. This is the Buddhist description of the whole of *Samsara*, the sensible world: it is solely the result of the ceaseless interaction of sentient beings.

The idea that we abstract meaning from information is upside-down. The meaning is there in the world, and information is an abstract idea we arrive at when we consider all the different possible representations of it. This is the connection between computation and physical reality. A computational process is a re-encoding of representations. No new information arises within any computation: computation is deterministic so the most information we can get out is what we put in. Any deterministic dynamical explanation of a quantum system[11] is a computation and cannot yield any new information. It is only the conscious act of judgement that yields new information. On this view, probability and information are the same thing. A probability distribution that is the result of a quantum mechanical description of some aspect of Nature is nothing more than an expression of the knowledge we are assuming as the fixed background of the experiment. Thus the proper concept of probability is that of Bayesian inference which takes the role of probabilities as strictly epistemological. There is no need for the mythical 'objective interpretation of probability'.

Quantum mechanics is spectacularly effective in describing the physical world, with one significant exception, which is gravity. From the phenomenological perspective this is not so surprising. Like the quantum state, space-time is not something that can be meaningfully said to have an independent existence. Space-time is nothing more than the partial order of all the events [12] and events are always representations arising in experience. Therefore I would not expect a quantum mechanical description of gravity as space-time curvature. There is no way, nor any need, to unite the description of gravity given in general relativity, a description which is in terms of direct experience of observers, with the quantum mechanical description of fundamental reality as statistical correlations of quantum state. Space-time, like consciousness, is a condition of the Universe, not a phenomenon that can be explained in terms of some more fundamental phenomena. If consciousness is a condition of the Universe then it

[11] Unitary evolution in von Neumann's terminology.

[12] Which quantum mechanics describes only by probability distributions which are conditional on the fixed background assumptions and knowledge.

has no location in space-time and conscious experience may be non-local provided it does not contradict causality in the form of the partial order of events that constitute space-time.

From the point of view of humanity as a whole the world seems to be in a bit of a mess. Reductionism leads us to believe we can apply technological solutions to what we perceive as specific problems but the problems we perceive are abstract representations and when we solve them we inevitably do much more than just that. Lego-style thinking is what got us into this mess, and lego-style thinking is not going to get us out of it. We can't 'move into cyberspace' because cyberspace is just information: abstract representation. So we don't need any more computer programmers, accountants or quantum mechanics. What we are going to need in spades though are clever doctors and engineers. So mothers, please throw away your children's lego. Don't get them meccano, let them make real things like go-carts and rafts and tree houses. These things might be a bit dangerous, but this is nothing compared to the sort of dangers they may have to face later in life, and they'll be better prepared to meet them.

Ian Grant
La Paz, Bolivia

References

[1] Niels Bohr. *Atomic Theory and the Description of Nature*. Cambridge, Cambridge, England, 1934.

[2] Wilfrid Hodges. Tarski's truth definitions. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Stanford, fall 2008 edition, 2008.

[3] David Mermin. What Is Quantum Mechanics Trying to Tell Us? *arXiv eprint quant-ph/9801057*, 1998.

[4] Carlo Rovelli. Relational Quantum Mechanics. *arXiv eprint quant-ph/9609002*, 1995.