A brief review of Dennett's article
"True Believers: The Intentional Strategy and Why it Works"

Presented by Pedro da Fonseca
to Professor PAUL BERNIER in 14 / 02 / 00

Dennett's article can be divided in three parts: *i*) an introduction where the problem of belief attribution is sketched and two extreme solutions are delineated; *ii*) a description of the intentional stance or strategy (IS); and *iii)* the defence that, if a given system *x* is 'reliably and voluminously' predicted by IS, then there is no further reason (not) to attribute intentionality[1] to *x*. It is specially the last part that which seems to deserve more criticisms so we shall concentrate more on it.

1. "Interpretationism" vs realism.

Dennett distinguishes two ways of looking at belief: as something existing objectively and unambiguously in the brain (or body) of the 'agent'; something whose existence could, in principle, be determined by a purely physical inspection (analysing brain patterns, etc); *or* as the result of an interpretation. The difference is quite big: seeing a Necker cube we can say that it is a fact that it is formed by 12 lines, but if we ask which face of the cube is at the forefront we will obtain not one but two different and equally valid answers. Does this mean that there are two objective facts about the forefront face of the Necker cube, or that there is none objectively, but only subjectively, as an appearance, something made up by our brains? Usually we may be inclined to think that when facts depend on the point of view of the observer they are not really there but only in the observer's mind. This poses an interesting problem for belief attribution because we usually think that all animals can be described as purely physical systems while at least some of them can *also* be described as intentional systems. If this is true then the attribution of beliefs to a person, for instance, seems the result of the uptake of a point of view (either physical or intentional), and therefore a subjective interpretation of a physical fact that, in itself, has no such characteristics (this would consistently lead either to solipsism either to a mysterious epiphenomenalism). Dennett's *tour de force* is to show that *a*) beliefs only appear in higher-level interpretations (like the 3D appearance of the Necker cube), but, *b*) nevertheless, a "belief is a perfectly objective phenomenon" (76).

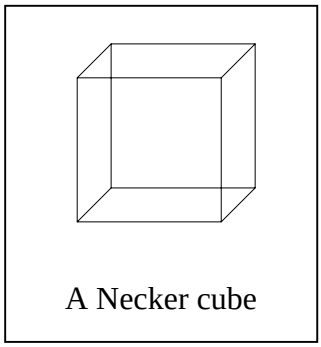
A Necker cube

Difficulties: This first part seemed to us the most obscure of the entire article. Not only did we not understood the allusion to the social scientists made in the beginning but the focus on belief (instead of intentionality) was also not understood. Finally, Dennett presents several circumstances where belief attribution is supposed to be 'a dark, risky, and imponderable business'. In p.75 this depends on plausibility, namely, either on the beliefs themselves ('exotic' beliefs), on the subjects that have them (e.g. non-human animals), or a combination of both (healthy member with awkward belief). But in the next page Dennett seems to connect the subjective character of belief attribution to judgements of *value* since the attribution is 'linked' to questions such as "whether a person is immoral, or has style, or talent, or would make a good wife." Also, it did not seem clear to us why does the realist view on beliefs depends on the beliefs being more 'familiar' (instead of being a metaphysical point of view).

[1] Dennett uses either intentionality, belief (and 'true believers), internal representation, etc. In *Kinds of Minds*, he also uses the expression 'mind-havers', and others. But it seems that the idea directly connected to the IS is just to know when we are justified in ascribing (first person) experiences to a given system.

2. Describing the Intentional Stance

Dennett distinguishes three strategies for predicting an object behaviour: the physical, design, and intentional strategies (or stances). To use the physical strategy we describe the behaviour of the components of a certain physical object and use the mathematical laws of physics to compute their integrated behaviour. Although Dennett links this to the strategy of Laplace it is far from obvious that the components of the system have a maximum size. For instance, to calculate the future behaviour of a galaxy through computer simulations, the galaxy is decomposed in its billions of stars and detailed computations are made to see the result of the interaction between each and every star in successive moments (these calculi are obviously very slow). This seems to be a physical strategy although in this case the smallest component under analysis is much bigger than our entire planet.[2]

The design strategy appears "where one ignores the actual … details of physical constitution of an object, and … predicts that it will behave *as it is designed to behave*" (77). The design stance can have several levels, for instance a clock's behaviour can be predicted by its design but also the behaviour of its internal mechanisms when "their material is not specified" (if it was specified we would be presumably at the physical stance).

Dennett now turns to the IS. He starts by giving a short description, as we will see, there are at least two main problems with this description:

> "Here is how it works: first you decide to treat the object whose behaviour is to be predicted as a rational agent; then you figure out what beliefs that agent ought to have, given its place in the world and its purpose. Then you figure out what desires it ought to have, on the same considerations, and finally you predict that this rational agent will act to further its goals in the light of its beliefs." (77)

The first problem has to do with the concept of 'rational agent'. The second problem applies to all the four steps: it is not clear that the attribution of rationality, agency, beliefs, desires, and purported action always follow the sequence described by Dennett. The idea that the IS is applied in a more holistic fashion is perhaps not implausible, specially when we consider that recognition, of visual scenes for instance, seems also to function in an holistic manner (for instance we don't start by seeing the object, to recognise it we must already now that other parts of visual scene do not belong to it).[3]

Most of what Dennett says next applies to "the principles of believe and desire attribution" (79), which we will see in some detail. The question is now: why do we have to

[2] The predictability of the upper levels of reality given a complete description of the atomic level is considered by Dennett "a dogma of the physical sciences." (77) A dogma implicitly accepted in the argument of the martians discussed below. However, the 'minor' complications of quantum physics completely destroy this possibility. There is NOT a single interpretation of quantum mechanics that allows this possibility of predicting with absolute certainty the macro from the micro. This can easily be seen: imagine we would have a complete description of the Universe (say that the martians would have that), this description would be made in terms of a function call the psi wave, which determines the evolution of the system. Now, this function describes the future as a collection of incompossible events (for instance particle *a* being at point 1 and 2 at the same time), but when we observe the present we never detect these incompossible events, they do not make part of the observation. Now, the function does not tell us which will be the events observed, at most it can be *interpreted* as a measure of the probability of a future event. Therefore, if our best physical theories show that predicting a system through a quantum description with absolute precision is impossible, how can it be a 'dogma' of science? Perhaps of Newtonian science, or science as seen by philosophers. But no more than that.

[3] There is here also another somewhat marginal point: Dennett speaks of "what the agent ought to do; that is what you predict the agent *will* do."(77) The word 'ought' here seems a little out of context since it may suggest moral implications. Is Dennett subtly suggesting that what an agent ought to do (morally) can be determined according to its set of beliefs and desires?

assume, to use the IS, a rational agent? We should note at this point that Dennett says that the IS can be correctly used on a variety of 'intentional systems' like persons, chess playing computers, thermostats or even lightning! Therefore, in each of these cases we must assume that a rational agent is having beliefs and desires and striving to achieve its goals. But why not simply an agent? Perhaps our puzzlement derives from a lack of acquaintance with relevant literature, in any case it should be noted that 'rational agent' here seems (although perhaps we misunderstood) to mean only that the agent will try to achieve the goals regarding his set of beliefs and desires.[4] It seems it does not imply that the agent have abilities like the ability to speak or even think, etc.

Lets turn now to the principles of belief and desire attribution. As we have seen we start with an object and then we consider 'its place in the world and its purpose'. From this we derive both its set of beliefs and desires. Dennett says that the "strategy becomes clearer with a little elaboration" (77), but in fact he only gives one rule both for beliefs and desires, a list of proper desires, and a few exceptions:

- Rule: attribute those beliefs / desires the system ought to have.
- List of Desires: "survival, absence of pain, food, comfort, procreation, entertainment." (78) Other desires might be added if we think the intentional object has certain beliefs. For instance, belief that p will bring the (proper) Desire x, will introduce p into the set of desires.
- List of Beliefs: *All* beliefs except the ones irrelevant for the system or the ones which demand different experiences from the ones the system has been exposed to.
- Exceptions: False desires and false beliefs can also enter the list of attributable beliefs (also by exposure) but they demand a special story.
- Language: Allows for more precise beliefs and desires but sometimes will also constrain desires and beliefs since language is sometimes 'more particular' than the desire or belief we are trying to express.

There are obviously several interesting problems here, both regarding the explanation of false propositions and the implicit critic on Fodor's LOT proposal. But the thing that strikes us most is that, in fact, there is no exposition of the IS. The Rule is simply ambiguous, the list of desires seems at best arbitrary and not at all applicable to plants or chess playing games, the exposure principle ("*Exposure* to *x* … is the *normally sufficient* condition to know (or having true [and relevant] beliefs) about x." (77)) seems incomprehensible since different organisms, for instance different persons, being exposed to the same state of affairs get quite different knowledge; and nonetheless no one is able to get "all the truths relevant to the system interests" (78). But even if all these rules were true it would be difficult to see their application to particular examples. For instance when playing chess we do not suppose that our opponent knows everything relevant to chess winning. We will attribute some set of beliefs to him about chess playing (rules of the game, openings and endings, etc) but only during the game will we be able to make more precise attributions. The clamp: "once a clamp believes there is danger about, it will not relax its grip on its closed shell until it is convinced that the danger has passed" (79). But do we really apply to the clamp a list of desires and beliefs to before we use the IS? It seems to us much more likely that we only attribute few things two the clamp, perhaps even only two: the belief that the shell is being forced, and the

[4] In fact we found out that Dennett does make precisely this point in *Kinds of Minds*, p.45. Then he uses the term 'smart agent'.

desire to keep it close; we can also attribute the clamp a desire to stay alive but it does not seem essential to use the IS. In the first case there is already an attribution of intentionality.

3.1 True Believers as Intentional Systems

We will now enter into the last (and more controversial) part of Dennett's article. There are two controversial claims that appear in this last section: one is to consider the IS as the *only* way to attribute beliefs and desires to an object: we cannot say that beliefs or desires are in the object (independently of our observation of them), we just attribute them if we choose as predictive model the IS. The other is that Dennett seems to apply the IS to objects that may easily be predicted by the design stance or even by the physical stance. We will analyse this claims separately and then point out their connection. Lets start by the first claim:

> "It is not that we attribute (or should attribute) beliefs and desires only to things in which we find internal representations, but rather that when we discover some object for which the intentional strategy works, we endeavour to interpret some of its internal states or processes as internal representations." (84, original in italics)

The main point that this passage suggests is that we do not 'find' or see an object's internal representations. Instead, we try to force them on the object and see if it works out. By itself this does not seem very striking. We don't see that the earth moves, and yet, she does move. If we think about it, almost nothing we see or know comes about without interpretation. Interpretation, or better, reasoning and hypothesis making, are excellent ways of probing the objects, of allowing us to see deeper into them. And yet the general tone of Dennett's presentation seems to suggest that something is lacking in the observation of internal representations that is not lacking in observations in more common (physical?) objects. We will try to clarify this assertion.

Let us return no the example of the Necker cube. Compare the two following cubes. One is an example of a Necker cube while the other might be called a shaded cube. The relevant difference between these two cubes is that, when you try to interpret these figures as 3D objects, only the first allows for two different interpretations. In the second cube the forefront face blocks the others behind it so that there is no doubt on the face which is at front. On the contrary, regarding the Necker cube, we might want to say that there is not 'a matter of fact' about which is the forefront face in the Necker cube, but there would apparently be such a fact in relation to the shaded cube.

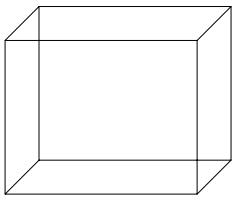

A Necker cube

A shaded cube

On the other hand we might say that neither of the images here have forefront faces. They just 'represent' 3D objects that have forefront faces. If we were looking at the 3D analogues of these cubes we would be obliged to disagree from the 'interpretation dependent' attribution. Imagine a 3D analogue of a Necker cube, for instance a wire-frame cube (all faces transparent) standing at a distance. We might then have the same difficulties in finding out which were the forefront structures of the cube, but, in this case, we know that only one of the structures can be at front of the others. Even if we can interpret it in several ways (by lack of

knowledge) we will know that only one of the ways is the correct one, even if we never know which is it.

The point we are making is that if we want to say that some properties belong to the interpretation or to the objects in themselves the ambiguous character of the interpretation is of no help. If we want we can rotate a 2D representation of a transparent cube, stretch it to the infinity, etc. In those ways we will no longer be able to have an ambiguous interpretation of which 3D object it represents, but this will not give it 3D, it will not turn it into a cube, it will be just ink on paper.

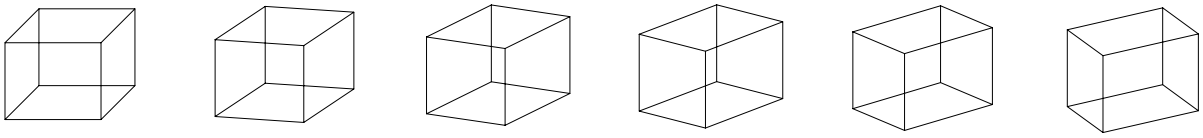

A rotation of a Necker cube (either to the right or left according to interpretation)

Now, in a way, every 2D representation of a cube is just a simulation (a 2D drawing of a 3D object). When we stretch the cube, rotate it, paint its sides, etc, we are in fact *simulating* the way a real 3D object would behave.

Now, what Dennett is defending is that facts about intentional systems are facts about interpretations; interpretations of physical systems that, in themselves, cannot be said to have those characteristics. For instance, consider a bright spot in the ski. We might be unable to see if it is a star, a planet, a galaxy, a binary system, a satellite, a nebula, etc. But it must be something of the sort. If it is a star it is *really* a star, and we know it would be a star even if we had no way of seeing it as a star (for instance if the Earth atmosphere would never allow for astronomical observations, or if no one had the concept of a star). But in the case of an intentional system, and according to Dennett, this does not happen. Consider, for instance, that we observe a person whose behaviour cannot be exactly predicted by the IS. In a realistic interpretation we would have to say that such a person either has or does not have beliefs and desires although we might never be able to know it exactly. Dennett explicitly denies this kind of reasoning; according to Dennett, in this case "there is no fact of the matter of exactly which beliefs and desires a person has" (82-3). As he says:

> "The perverse claim remains: *all there is* to being a true believer is being a system whose behaviour is reliably and predictable via the intentional strategy" (83).

There is a simple way of seeing what Dennett is saying. He proposes a method for making predictions (the IS) that, he says, we use abundantly, either to predict the behaviour of plants, non-human animals or persons. But it seems obvious that, looking at a thermostat and seeing it as having beliefs (it's too hot, it's too cold) and desires (to turn on / off the boiler) does not mean that the thermostat – as a real, independent object – has beliefs and desires. It means that it is a suitable object for we to see it *as if* he/it (?) had beliefs and desires. The same with the representations of the cube, we can see it as more or less ambiguous representations of a cube, but, in fact, they are just pieces of ink on paper. They are able to receive the interpretations we want to make of them, but the properties result from the interpretation, they are not in the objects themselves. At least with persons it seems different, it seems that we are not merely making metaphors but, by the contrary, discovering real (underived) intentionality in them. It would seem, for instance, that a person would continue to act as a person, to have beliefs and desires, even if we saw her as a kind of 'it', a robot for instance.

Dennett says, by the contrary, that it is impossible to distinguish between these systems on the basis of possession of *real* (underived) intentionality. Dennett says that to distinguish

> "those intentional systems that *really* have beliefs and desires from those we may find it handy to treat *as if* they had beliefs and desires … would be a Sisyphean labor, or else be terminated by *fiat*." (80)

In this view intentionality is always derived, not just the words in the books but the beliefs in people's minds. But is it really difficult (or even impossible) to distinguish between intentional and non-intentional (or metaphorically intentional) systems? (It is to this question we will now turn.)

3.2 The scope of the IS

What defines the set of intentional systems? Are there perhaps two sets of intentional systems or just one? Dennett introduces the design stance in cases it would be 'more effective' then the physical stance (77) and the IS is introduced because sometimes 'even the design stance is practically inaccessible" (77). When analysing the IS he even says:

> "What should disqualify the lectern [as an intentional system]? For one thing, the strategy does not recommend itself in this case, for we get no predictive power from it that we did not antecedently have. … <u>In the case of people, or animals, or computers</u>, however, the situation is different. In these cases often the <u>only strategy</u> that is at all practical is the intentional strategy; it gives us predictive power we can get <u>by no other way</u>." (80) (our underlines)

So it would seem that only the systems whose behaviour was not practically predictable by other strategies (this is what would happen in relation to animals and persons) could be considered intentional systems. Computers were one of the examples of design systems (Cf. p.77), but in certain cases (like in chess playing programs, but not in calculators) we might (also) use the intentional stance.[5] In this article Dennett is not completely clear about the scope of the set of intentional systems. In these passages it would seem that he is (very) clearly advancing that it is composed *only* of those systems that are predicted only by the IS. But the only system that he explicitly rejects is the lectern. The thermostat, by the contrary, is constantly used, and in *Kinds of Minds* (more recent) Dennett explicitly says that:

> "*Intentional systems* are, by definition, all and only those entities whose behaviour is predictable/explicable from the intentional stance.[6] Self-replicating macromolecules, thermostats, amoebas, plants, rats, bats, people, and chess playing computers are all intentional systems". (*KM*, p.45)

This should not be surprising or seen just as an accessory part of the theory. In fact, it is precisely because there is no acknowledged important difference between thermostats and people that Dennett is able to extend conclusions easily applicable to reasonings on thermostats (only a *'façon de parler'*) to animals and persons. This, it seems to us, is the most central aspect of Dennett's argument and perhaps the only basis[7] for asserting the thesis

[5] Even in the case of chess playing programs it is not very easy to see why the IS is necessary. Is the IS also necessary when we play tic-tac-toe with a computer? It would seem that in both cases the design stance would do just fine: the computer is designed to play x. It is not very clear that the attribution of beliefs and desires (instead of simply design) are in fact *necessary* to predict (better) the behaviour of a chess playing program.

[6] Notice that Dennett does NOT say 'predictable/explicable **only** from the intentional stance.'

[7] Notice that we have dismissed the possibility of different interpretations as a support for the non-realism of the specific properties of intentional systems. Dennett is not very clear on the role that the possibility of incompatible interpretations plays in his theory. Nevertheless he explicitly accepts that there is a relativism of

presented in our previous chapter (that there is nothing more to a true believer then being predicted by the IS). But are we able to defend this «extended» set? Note that if we make a more restricted use of the IS (using it *only* in situations where no other predicting models are applicable) we would have difficulty in using it in beings other than animals and humans. For instance, a program like Eliza might be able to grasp our attention as an intentional system, but once we discovered the simple set of rules that it uses to generate responses, another and more reliable strategy (a design strategy) would become readily available and we would stop using the IS, and also stop considering Eliza as an instance of an intentional system. This limited use, on the other hand, would seem (at least apparently) the sensible thing to do, since in other areas we choose always the simplest model to predict facts (if we can predict the behaviour of the solar system just with the help of gravity laws, why should we include invisible dogs that kept the planets in their orbits (or any other innocuous complication of the theory)?).

3.3 Differences in complexity

Dennett establishes an important difference between simple systems like thermostats and complex systems like human beings. It is not a 'magical' difference, there

"is no magic moment in the transition from a single thermostat to a system that *really* has an internal representation of the world around it. … The principles, and problems, of interpretation that we discover when we attribute beliefs to people are the *same* principles and problems that we discover when we look at … a thermostat. The differences are of degree" (84)

According to Dennett this difference can be expressed as a difference in complexity. Dennett starts by analysing the kinds of beliefs and desires we can attribute to a thermostat. The main point is that the internal mechanisms of the thermostat depend only on very few factors (in fact only of temperature). It can be connected to a boiler or to anything else, but the responses and internal states will be the same. As we add further sophistication to an intentional system its internal states will become more complex and connected to more aspects of the environment. For instance a system controlled by a computer can check your agenda and predict when you'll get home, if you add it a simple camera and a thermostat and a few things it is easy to make a program that only connects the heating at the proper time (when it is cold and you are arriving or at home). This would make the internal states of the system dependent on much more aspects of the environment. For instance, the internal states of a person are so much fine tuned on the environment that any perceptual difference in the environment would yield a difference in the internal state of the person.

This detailed dependency allows for more precise beliefs to be attributed. If a thermostat reacts the same way either it is connected to a boiler or to a train it would be difficult to attribute 'it' the belief that in fact it is a boiler that 'he' is commanding. Its connection to the boiler "is too impoverished a link" (83) to grant a specific belief. Of course the connection between internal and external differences is not sufficient to create a link (for instance temperature and metal dilatation). But to have the link this connection between internal and external changes must occur.

According to Dennett an increase in connections allow for a more precise, or rich, 'semantics'. The word 'semantics' is not discussed for but it is obviously introduced through

---

interpretations. For instance one culture can "attribute to a member might be quite different from the beliefs and desires another culture would attribute to that very same person." (83) Dennett also accepts that "it may be that [such] two rival schemes are about equally good" (83).

the IS. When we apply a belief to a thermostat we are attributing it some 'semantics'. So, if we have a more complex thermostat we can apply more precise semantics to it. If we have a really good machine we can programme it so that it will have so much information that we can attribute it even more precise semantics. In limit we can have a machine that, as Dennett says '*mirrors* the environment' since it would have "become so perceptually richer and behaviourally more versatile" that to any change in the environment "it will notice, in effect, and make a change in its internal state in response". (84)

Now, if this was an account of the way in which people use the IS it would be simply wrong. But, although Dennett departs from a common sense perspective, this last part of his paper does not hang on to common sense,[8] so to simply note that this is not the way people use the IS is by itself insufficient. But what would be sufficient to show Dennett is wrong (if common sense can't)? We think it could clarify things a bit if we look at how people use the IS and why they use it that way. Now, it is difficult to have (examples with) highly developed thermostats, but there is all sorts of complicated machinery in aeroplanes, spacecraft and so on. One example is the B1 bomber (so-called 'stealth bomber'). One of the most difficult aspects of making this plane had to do with driving a single-wing plane. In fact, the plane is completely unstable and any little breeze is able to send it way from its right course. So, in fact, the plane is driven by a computer who is constantly making flight corrections. The pilot normally controls the plane as he would control a normal plane, but instead of his instructions going directly to the flaps, ailerons, and so on, they go instead to the computer who then decides what should be done. Now, to work correctly, this computer has to have a lot of information about the altitude, atmospheric pressure, winds, how the plane behaves, etc, etc. In fact, because the plane is so sensitive to minor atmospheric changes, he must be sensitive to any little variation and respond accordingly. In fact, we could say that he 'mirrors' his (atmospheric) environment in the relevant sense. But, following common (good?) sense, who would attribute literally any desires or 'beliefs' to such computers? On the other hand we do attribute some degree of intentionality to old people who sometimes are almost incapable of noticing much around them: we keep on saying that 'there is someone there'.

A B1 Stealth Bomber

Now, when studying AI (artificial intelligence) we are introduced to concepts such as 'information' and 'knowledge'. Knowledge, in AI terms, is obtained not only by a gigantic database but by a well organised, all connected, system of concepts and situations that enables a computer to understand a situation and respond in an intelligent way. We can use scripts and things like that to obtain that knowledge. But where is the connection to the IS? Once again let's return to the Necker cube example. If you've got a good algorithm you can make a very good simulation of a rotating / stretching / shading cube. That does not

[8] Cf. p.76 "and my arguments [with respect to this last section of the article] have so far garnered few converts".

turn the computer into a cube. It can simulate a crystal cube for instance, but you don't have to have crystal to do it, just to know how crystal normally behaves. Sometimes it is expressly said that a machine that could simulate a human being in perfection (the Turing test) should be considered as having the same 'phenomenal' (in lack of a better word) experiences. But a computer can as well simulate two persons, or even a crowd, or ten billion people. There is no reason, in principle, why that should not be possible. In fact, a Turing Machine could, in principle, even simulate the entire planet, perhaps even the entire cosmos. Now, what would that mean, that it had all the feelings of all the people it simulated? That it felt all the colours, sensations, etc? And why not, by the way, to say that it had all the temperatures, all the pressures, all the locations, etc. In a sense it is true. It would *simulate* all these properties, just in the same way it simulated the crystal reflections of a crystal cube. Nothing more, nothing less. If we do not attribute to the computer the same physical properties that the simulated system has (for instance a galaxy) what reasons do we have for applying a different reasoning with respect to mental properties?

So, when we look at a computer driving a B1 bomber, it might be mirroring his environment, it might have much more knowledge and 'intelligence' then we will ever have, but why should we apply to *it* the IS? Conversely we must say that the IS is applied to people very dumb, even to animals, or to people with severe mental problems. Does this express some misunderstanding on common sense on what makes a system to have 'beliefs' and 'desires'?[9] We think another example might give us an important insight on the basis of the common sense use of the IS. One of the most incredible, fast, accurate, 'semantically' rich (in Dennett's sense), sensitive, etc, computers is the brain, more specifically the cerebellum. For instance, it allows us to ride a bike, or to walk, or to stand without falling, or to coordinate our muscles, etc, which are exceedingly difficult tasks. To be able to do that the cerebellum needs all sorts of informations on the world and on the body. However, intentionality is not attributed to automatic movements, no matter how intelligent they are. For instance if you see me on a bike skipping a car on the last microsecond you will not think: "gosh, he though about moving the hand, and the arm, and the legs, and he wanted to contorted the whole body because he believed that the centre of gravity of himself and the bike would change which would allow… etc, etc". You just think: "late again for classes…" In the same way we don't think that the chameleon changes his colour because he desires to escape its predators and beliefs that, because the stone under him has the same colour and the visual acuity of its preys, etc, etc." No, people just think that these are *automatic* procedures. That these processes do not involve consciousness, beliefs, desires, 'thinking', whatever, is what they can feel in their own skin. In general, we know from our experience, that it is not the degree of intelligence of a process that requires an 'intentional system'. If it is a process that can be done automatically then we don't need an intentional system to control it, if it cannot be made automatically then an intentional system will be required.

We can indeed love computers (like I do), we can even predict or hope that one they can replace us, or at least help us to achieve some kind of immortal life. This is perhaps not impossible, but if we want to do *that*, we have to do more than to create an universal simulator. We have to do more then to create intelligent behaviour (intelligent behaviour already appears in ants and stuff). We have to create machines (or whatever) that are aware. To do that we cannot put our head in the sand, we cannot repeat ad nauseam that we already created intentional systems just because we have created intelligent ones, as far as we know, the two things mix, but they are not the same.

[9] We should note once more that when Dennett speaks in beliefs, desires, internal representations, etc, he is really speaking about *consciousness.* This is very clear in *Kinds of Minds,* when the connection between the use of the IS and pain is strongly made (Cf. pp.5-7).