

On What Embodiment Might Have to do with Cognition

Tim Smithers

Euskal Herriko Unibertsitatea
Informatika Fakultatea
649 Postakutxa
20080 Donostia
Euskal Herria, Espainia
ccpsmsmt@si.ehu.es

Abstract

Without widely accepted and acceptable characterisations of what is to be meant by the terms cognition and embodiment, to ask what the role of embodiment in cognition is remains an ill-formed question. The connection between cognition and physical embodiment is however an important issue, and investigating it may well bring some enlightenment to the concepts of cognition and embodiment. In this short paper I offer some observations on what seem to be some important consequences of physical embodiment on the nature of cognition, at least cognition as it might need to be in physically embodied agents. These are all derived from my own work on physically embodied mobile robots and point to the need for a dynamical systems approach to understanding cognition, at least in physically embodied agents.

Introduction

The announcement and description of this symposium tells us that “the role of physical embodiment in cognition has long been the subject of debate.” This may be true, but it has often seemed that this debate is more about matters of faith than about well argued positions or experimental results. Why should embodiment have anything to do with cognition, or, if this is a different question, why should cognition have anything to do with embodiment? Not all physically embodied things are cognitive things, and, perhaps, not all cognitive things are physically embodied things, at least not in any important sense of embodiment. Knowledge-based Systems, or world class Chess Playing machines, for example, are not embodied in any important sense, though at least some people are happy to say they have cognitive capacities, albehem particular and limited.

This question, what has embodiment got to do with cognition (assuming for now that it is one question) is not easy to answer for it is not clear what we want or need to mean by embodiment or cognition: it is not clear what form an answer could or should take. There is no single widely accepted (or acceptable) concept of cognition. Nor is it so clear what embodiment is to mean. In this symposium it is “physical embodiment”

that is at issue. Why the “physical” here? Is it possible to have non-physical embodiment? If so, does this have the same kind of role in the same kind of cognition, or the same kind of role in a different kind of cognition, or a different kind of role in a different kind of cognition?

In this short paper I do not propose to offer an answer to this question, or to the many that follow from it. I really don’t know how to, though I do think there is a need to try. What I do offer are some observations on what seem to be important consequences of physical embodiment on the nature of cognition, or at least the way we should try to understand and explain cognition. In so doing, I hope to offer some pointers to what physical embodiment might have to do with cognition.

Physical Embodiment: Mass and Time

A straightforward observation, though one that seems often to be overlooked by the computational modellers of robots (so called simulation builders), is that with physical embodiment comes physical mass. It cannot be otherwise, according to accepted physical Law. Furthermore, if something has mass then whatever it does or whatever happens to it TAKES time to do or happen! Again, it cannot be otherwise, according to accepted physical Law.

That things take time is thus a direct, fundamental, and unavoidable consequence of physical embodiment. We can then observe that if cognition has anything to do with what an embodied agent does or to what happens to it—as it surely somehow does—then just as time is of fundamental importance in understanding the effect and consequences of embodiment, so too will time be of fundamental importance in understanding cognition. To put this another way, to properly understand physically embodied cognition, the concept of time will have to be in there from the start, not added on at some arbitrary later stage, as it often is in current information processing and computational approaches to cognition.

Taking time: IN-time and OVER-time

It is not quite so simple as this though. We need to make a further distinction if we are to make progress on picking out quite how embodiment might have something to do with cognition.

As we have just observed, all events and happenings in the physical world take time, everything that happens takes some time, nothing that happens takes no time: walking, thinking, computation, whatever. We need, however, to make a distinction between things that happen IN-time and things that happen OVER-time.

For things that happen IN-time the time taken matters. If they happen in different amounts of time they are importantly and distinguishably different happenings. In these cases, the time in which they happen is a fundamental characteristic of the happening. To properly understand and explain them we need to introduce the concept of time. These kinds of things are, we might say, *embedded in time*. For example, in walking, the time taken to move legs matters. If the leg movements involved took different amounts of time we would not have the same kind of walking—we may not even have walking at all! It is the same for talking. If the movements involved in making the vocal sounds of talking took longer or shorter amounts of time than they do, we would not understand them to be talking. They'd just be funny noises. Perhaps it is the same for thinking too? Would thinking be the same if the underlying neuronal processes took different amounts of time than they do?

For things which happen OVER-time the time taken to happen is of no fundamental importance, though it may well be of practical importance. In these cases, if the amount of time taken changes or varies we still have the same thing happening. The time-taking aspect of these kinds of things can be safely disregarded in any proper understanding or explanation of them: they are NOT embedded in time; they are just *contained in time*. For example, computation, what real computers do, is not changed by the amount of time it takes to get done: it is still the same computation. This is of no great surprise since the formal specification of any particular computation does not involve anything that says or implies how much time it should or must take to be done correctly. The result of the computation does not change just because it take one second as opposed to two days to do: the rate of doing computation has no impact on the computation being done. (Though, as mentioned above, this may well be of practical importance.)

This distinction between things that happen IN-time and things that happen OVER-time is important. It helps us identify when physical embodiment is important or not. For animals and robots it is. Their behaviour happens in time: their physical embodiment embeds them IN-time. For computers it is not. Their behaviour happens over time: their physical embodi-

ment does *not* embed them in time.

It is also the basis for deciding between those things that we can sensibly attempt to describe in terms of the mathematical theory of dynamical systems, and those things which are properly described in other (non-dynamical) terms, such as computation, logic, or information processing, for example. We can also see that this distinction is important in any discussion of cognition because the time-neutral concepts of computation, logic, and information processing are popular ones in attempts to understand and characterise the nature of cognition. The more so since, as I have suggested above, the actions of embodied things happen IN-time, NOT just OVER-time. All this should thus lead us to suspect that the way time may be involved in understanding cognition, and the role of embodiment, will be in IN-time terms, not in OVER-time terms. In other words, we should be led to expect that we will need a dynamical systems approach to understanding physically embodied cognition, NOT a computational, logic-based, or information processing one.

A tempting alternative might be to have the processes concerned with physical embodiment as IN-time processes—dynamical systems—and the processes concerned with cognition as OVER-time processes—computational process, for example. This would leave us with the need to have, and to explain, an effective interface between IN-time processes and OVER-time processes, where effective here would mean that no IN-timeness gets across from the IN-time side to the OVER-time side. This is the real-time control engineers dream. And it is only a dream and must forever remain so. If, for any reason whatsoever time counts, so that we have any amount of IN-timeness, we cannot escape. We can of course make things feel like such an escape by making the OVER-time computations go so fast that they never cause problems on interfacing with IN-time processes, but this is only an illusion, albeit a useful one.

Agents as the Connection Between Embodiment and Cognition

As I observed above, not all physically embodied things are cognitive things, and, perhaps, not all cognitive things are embodied things. So what connects these two concepts? The answer is, I think, the concept of an agent: physically embodied agents are, I suspect, necessarily cognitive things.

So, we need to say more about what it is to be an agent, and in a way that will sustain the natural observation that not all embodied things are cognitive things.

An agent can be characterised as something which has the capacity to interact with its natural (typically changing, variable, and unpredictable) environment in such a way as to (at least) sustain and maintain its capacity to interact in this way. (A deliberately circular characterisation in the style of Maturana and Varela.)

Thus agents are things which interact with their environments in a self-sustaining and self-maintaining way. I call this type of interaction effective interaction. (More often agents are said to act, but I prefer the term and notion of interaction since I think it better reflects the actual relationship between agents and their environments.) This is a (working) characterisation of an agent, it is NOT a definition—it does not identify necessary and sufficient properties—and it is a rather broad characterisation at that. It is intended to be a minimal characterisation, one that does not introduce unwanted or (currently) unjustified extras. Particular examples of agents may engage in more than self-sustaining and self-maintaining interaction. Some effects of their interactions may be of consequence to other agents: their interactions may involve agent-agent interaction, but I don't want to get involved in these complications here.

For physically embodied agents all interaction with the environment involves energy transduction and the generation of forces: it involves doing physical work, what I call *interactive work*, which in turn requires a supply of energy. It cannot be otherwise, according to accepted physical Law. Elsewhere (Smithers, 1994) I have suggested that this leads to a need for a distinction to be made between the types of processes that a working physically embodied agent implements: between what I have called infrastructure processes and interaction processes.

Infrastructure processes are processes which are, in combination, responsible for sustaining and maintaining an agent's capacity for effective interaction, the capacity for doing interactive work. Interaction processes, on the other hand, are, in combination, responsible for sustaining and maintaining an agent's effective interaction with its environment—supported by the capacity to interact, as sustained and maintained by the infrastructure processes. They do this by suitably modifying or modulating the dynamics of interaction established and maintained by the infrastructure processes of the agent. Do to this, interaction processes must also use, and thus be supplied with, energy. This is the means by which they modify or modulate the workings of infrastructure processes. There is, however, a significant difference between the amount of energy involved in the functioning and energy transductions of infrastructure and interaction processes. Infrastructure processes, because they have to do significant work—move the agent around—require relatively large amounts of energy: they are significant power consumers. Interaction processes need only trade in relatively small amounts of energy and energy transductions, sufficient for signalling and sensing purposes.

It should be noted that low energy signalling and sensing is not unique to interaction processes. Infrastructure processes will also typically employ signalling and sensing systems in their sustaining and maintaining responsibilities. I would also add that,

contrary to what I have previously suggested—see (Smithers, 1994)—this distinction between infrastructure processes and interaction processes making up an agent *do not* map onto the biological notions of the metabolism and nervous system in an animal. Both the infrastructure and interaction processes must be supplied with energy and so both require metabolic support, though this may be different in detail. In an agent this metabolic support, of both kinds of processes, needs also to be regulated (sustained and maintained). For the moment I lump all the responsibility for this in with the infrastructure processes, but I admit that this does seem a little ragged.

The distinction between infrastructure and interaction processes in an agent is, I think, nonetheless important in understanding what physical embodiment might have to do with cognition because it allows us to see what aspects of embodiment are important in cognition. It sustains another natural observation that if embodiment is important in cognition it does not seem that all of embodiment is, or at least not equally so. This distinction follows from the agent characterisation presented above in the case of physically embodied agents, and thus allows us to hypothesise why not all physically embodied things are cognitive. We might propose that those things which do not have interaction processes are not cognitive things, and are not agents.

We can put this the other way round to. If cognition is to do with the nature and behaviour of the interaction processes in an agent, in physically embodied agents, at least, these interaction processes necessarily need the support of (and to interact with) infrastructure processes, and since both involve energy and its transduction, both need physical embodiment for proper implementation.

Some Final Comments

This short paper is not intended to be anything more than a brief presentation of what seem to me to be some important observation that can be made on the what physical embodiment might have to do with cognition. Very little, if any, of the above can be taken as definitional, and all of it needs further working out to become useful—if it is not to be thrown out first. Two important, and I believe, significant conclusions can be drawn from these observations however. First, that to understand cognition in physically embodied agents we need a dynamical system approach which can deal with the necessarily IN-timeness of cognition. Second, that we could usefully associate cognition with what the interaction processes of an agent do. Indeed we might even define cognition in physically embodied agents as what the their interaction processes do. The problem of cognition then becomes: what are the interaction processes need to sustain and maintain the effectiveness of an agent's interaction with its environment. A thoroughly dynamical physical and problem!

References

Smithers, T. 1994. What the Dynamics of Adaptive Behaviour and Cognition Might Look Like in Agent-Environment Interaction Systems. In Smithers, T. and Moreno, A. (Eds.) Workshop Notes for the 3rd International Workshop on Artificial Life and Artificial Intelligence: On the Role of Dynamics and Representation in Adaptive Behaviour and Cognition, December 1994, San Sebastián, pp 134-152.

Acknowledgments

The development of the ideas presented here have benefited much from many discussions with numerous people. I am particularly grateful to Randy Beer, Amaia Bernaras, Tony Bell, Mark Bickhard, Sunil Cherian, Wayne Christensen, Cliff Hooker, Leslie Kaelbling, Alvaro Moreno, Rolf Pfeifer, John Stewart, Javier Torrealdea, Wade Troxell, Tim van Gelder, and William Warren, for their various suggestions, corrections, support, encouragement, and their willingness to engage in constructive discussion. Nevertheless, the misconceptions and confusions that remain are all very much my own. I am also grateful to Alvaro Moreno who read and usefully commented on earlier versions of this paper. Lastly, but not least, I am happy to acknowledge the financial support of the University of the Basque Country for my current position, and the Faculty of Informatics, in particular, for providing a very supportive academic home. This paper has been written as a part of a project funded by the University of the Basque Country, UPV 003.230-HA203/95.