Can we not Copy the Human Brain in the Computer? [1]

by Luís Moniz Pereira[2]
Universidade Nova de Lisboa

Thank you very much for inviting me to comment on Prof. Idan Segev's presentation.

The more I try to comment though, the more thoughts spurt forth, and the more comments surge. This being a never-ending subject, commentary leading to further commentary, perforce and unwillingly, I am bound to give this topic short shrift.

Of course, I much enjoyed Segev's excellent presentation, though I must play here a critical role, and help you make your own evaluation, in particular about the just presented Human Brain Project (HBP)[3], aiming at the computer simulation of the human brain, and about its requested FET EU funding[4], to the total tune of one billion euro in ten years, which in that respect concerns us all in Europe. An issue being: Much as it may be an excellent project, is one billion too much? Will it dry out other projects? Will brain science as a consequence become too "monotheistic", too one-track minded so to speak?

Here is some inspiration from Lao Tzu. He said, in the *Tao Te Ching*: "Nature can never be completely described, for such a description of Nature would have to duplicate Nature." That is, one must duplicate nature to understand it completely, to describe it completely. And Jorge Luis Borges, in a story in *Collected Fictions*, wrote that "In that Empire, the Art of Cartography attained such perfection that the Map of the Empire coincided point for point with it." Actually, Borges was inspired by a prior Lewis Carroll's story, in his very last book *Sylvie and Bruno*, wherein "The map went from six feet to the mile, then six yards to the mile, next a hundred yards to the mile, and finally a mile to the mile." So again one would have a map the size of reality, but "farmers said if the map were spread out it would block the sun and crops would fail, so the project was abandoned."

Likewise, one might say: If a simulation is as complex as reality, can we ever understand the simulation?

---

1 Invited follow-up commentary on Idan Segev's presentation "*Can we Copy the Human Brain in the Computer?*" (http://www.gulbenkian.pt/section54artId3704langId2.html), on the occasion of the Calouste Gulbenkian Foundation's *Brain.org* forum (http://www.gulbenkian.pt/section54artId3704langId2.html), 9-10 October 2012, Lisbon, Portugal. My presentation slides are downloadable from here:
http://centria.di.fct.unl.pt/%7Elmp/publications/slides/brain-org/Copy Human Brain.pdf

2 Home page: http://centria.di.fct.unl.pt/~lmp/

3 http://www.humanbrainproject.eu/files/HBP_flagship.pdf

4 http://www.humanbrainproject.eu/flagship_call.html

But first let me jump to my last slide. In it, since we are also bridging science and art in this forum, there figures a drawing by M. C. Escher.

M. C. Escher's model of the brain modeling the brain

One can see that the drawing shows a city being represented in a picture, and that there is a person looking at the picture inside a museum. In this picture a port is seen in the foreground, and the port leads to the city. In the city itself, a street is shown with an entrance to the very museum where the person is looking at the picture. The drawing's title is "Picture Gallery", and is but one of many such drawings by Escher.

One wonders, by analogy: Can we model the brain whilst in the process of modeling the brain? And of course, the incumbent recursion reminds us of paradoxes such as Gödel's Incompleteness Theorem, and other paradoxes arising from self-reference. Indeed, can one correctly simulate a wrongful simulation of oneself? Or will it be perfect from the start?

Let me now jump back to the very beginning again. I will continue by saying that, in my opinion, the HBP project is not a question of reverse engineering, as hinted just now by the previous speaker. In the space shuttle disaster case mentioned, scientists had first of all already designed the shuttle and then, for fault-finding, reverse-engineered a part of it into mathematical equations in order to perform a simulation of the disaster. But we did not design the brain. Can we reverse-engineer it? It appears we do not have a construction specification as we did for the space shuttle...

Well, let’s take something simpler then, a computer. I am an alien, just arrived on a spaceship, when I find this computer made by humans. I try to make a simulation copy of it using my own alien technology, in order to see, to understand how it works. And I make a perfect copy because I know physics well. But when I want to run it, it doesn’t launch. There is no software in it. How does one reverse-engineer the possible software?

My point is that the brain, in its biological evolution, evolved so that it could execute any kind of mind software: personhood, art, whatever. That the brain has bootstrapped itself into generality.

From computer science we know that, as soon as one starts the computer working (by pressing the turn-on button), then what the hardware does – simplifying a bit – is that it goes to the first instruction in memory, which is a software instruction (the first instruction of the program), and loads and executes it in the hardware. What does this mean? It means that that instruction configures the hardware, i.e. the CPU (the processor), and obliges the hardware to execute that instruction. It means that the CPU obtains instruction-specified data from memory, combines the data according to the software instruction, and puts the result back in a memory location specified by the instruction. Then the hardware looks up the next software instruction, and so on and so forth. That is, the software becomes the master of the hardware. And of course, in the brain, that happens a lot too. It is actually called "free will" by some people.
The process of life bootstraps itself into just that: life on the hardware. Likewise, the software process bootstraps itself into just that: mind on the hardware.

So, I grant we can copy the brain to some extent. But I think the extent it is useful to so copy it, at this stage at least, is at the hardware level; and, of course, then to look into diagnoses of diseases, much as we can diagnose a bad computer board: We can then substitute this VLSI circuit here, substitute that graphics chip there, whatever. And that’s all very well and fine. But the question of the software remains lurking. It is not enough to do some so-called reverse engineering. That's too much hype, too much over-spilling. It is too "monotheistic" to say we are going to simulate the brain hardware and that this will amount to the be all of our ability to understand whatever the brain might be doing. There is a lot being left out.
Without a brain, human beings would know nothing, feel nothing, and experience nothing. Knowledge of the brain may not be enough to understand what this experience means, though being quite relevant to it. But what does it mean to simulate the human brain? Can the human brain be conceived independently of its context and external environment? Unlike in Escher's drawing?

I was glad to see that in the HBP there appears to be now somewhat more emphasis on the need for a Theoretical Neurology Institute, and I really think that's the crux. We need more top-down thinking about the brain as opposed to just bottom-up thinking. Reconstituting the brain bottom-up – putting all the little pieces on the map – means that suddenly we will have a very big map, and we will be in the map. The map will have become as big as reality, and what are we looking for in it?
Top-down that tells us what to look for. And I submit that Artificial Intelligence, beyond just the mathematical models mentioned by HBP, will tell you what kind of things: how you might do planning, how you can hypothesize about the future, how you can do this and can do that.

Nowadays we have a pretty good idea of at least some alternative algorithms, or combination of algorithms, that we might look for inside the brain, inside this hardware, because the brain, as much as our robots, exist in the same world and both have to solve pretty much a lot of the same kind of problems. So it's natural that some algorithms are candidates for trying to find them in the brain. Otherwise, we would be rather lost about whatever is going on in the brain. We'll see these nice brain-imaging colors spreading all over the place. But then what is the algorithm, what is being computed? Is the problem solving within the neuron? Is the problem solving going on in the connections? Or, first of all, is the problem solving in the computation, in the algorithm being processed? And this seems to be under-stressed in the project.

By now, I have said almost everything to cover a lot of the gist of what I wanted to say.
So let me go into some detail. For instance, and this I took directly from the 108-page report submitted to the European Union: "By contributing to better understanding of biological mechanisms underlying decision-making, knowledge, feelings, and values, the Human Brain Project (HBP) will inevitably influence conceptions of what it means to be a conscious, sentient being with a subjective view of the world." This sentence by itself appears rather innocuous, if high-minded wishful thinking. A follow-up sentence takes it further: "Whatever the results of the HBP project, they will profoundly influence current beliefs about the human mind, identity, personhood, and our capacity for control." That seems to me a clear over-spilling, be it with a good intention to do so (besides getting funds), but perhaps in 500 years or so we'll get there, to a unifying view that duly takes all influences into account to reach some overall understanding.

Again – everything taken from the report – it says: "The HBP will also make broader contributions to high performance computing. [...] Studies of the neural code will contribute to brain-inspired communications protocols with the potential to dramatically improve efficiency of internal communications in high performance computing systems. [...] Results from HBP will facilitate development of novel approaches to information storage and retrieval."
It goes on: "Information and Computing Technology (ICT) has made great progress addressing intractable problems, but there are still huge areas of the economy and daily life where ICT has limited impact. For ICT to meet these needs, it will require capabilities lacking in current

systems: massive parallelism; very low-power consumption; ability to learn new skills without explicit programming. Many computer scientists believe that the best way to create these capabilities is to imitate the brain."

As a computer scientist I don't believe this, and I don't think other computer scientists will likely believe that these brain studies will so contribute and lead to imitation, except perhaps in some specialized operations, like pattern recognition or vision or hearing – low level sensorial processing, and even then not at the low speed with which the brain functions. They will never, I think, be good enough for improving the high-end technology required of computers and their internal communications. I wish they could and would, but I don't think it will happen ever.

Returning to their statement above that "Many computer scientists believe that the best way to create these capabilities is to imitate the brain", I don't think so. Indeed, I have never heard colleagues of mine subscribing or even hinting at that. Maybe we haven't talked enough with the neurologists and I think there is a problem here. I wish there was more of a communication. I don't know of research centers where we can find computer scientists, artificial intelligence researchers, neurologists, evolutionary biologists, philosophers, what not, working together. And I think one of the reasons is that – and this is a general point about this very meeting, actually – many people are afraid of embarking into software. They want to be grounded on the hardware: hardware is where we stand, they have to feel securely well-founded, they want to have their feet on the ground. And the software is something that bootstraps, emerges, and exists at another level up. (Now it exists in "clouds" even...) So, they don't want to fly with recursions or wings too high. That is the resistance I perceive. Maybe I'm wrong.

And again, here's some more high-flying technological impact spiel: "Simple devices could be integrated into compact, low-power systems with the ability to control complex physical systems (e.g. vehicles, industrial machinery) with many degrees of freedom. Like the brain, such systems would have the ability to create implicit models of their environment, including their own actions, plus representations of those of other agents, and predict likely consequences of their decisions, so as to choose the actions most likely for their goals."
I work in this area of artificial intelligence. The latter are very high-level cognitive abilities, that I really don't see how we will be able, using computer-brain hardware inspired models, to improve on our ability to symbolically represent and reason about these issues – that's how we do it now. In fact, speaking of symbolic processing, computer scientists that work with so-called Neural Nets – having nothing to do with the actual neurons by the way – and with Bayesian Networks and other low level mechanisms, have long found out that low level is not enough. Nowadays, they are looking into hybrid systems, combining symbolic computation with sub-symbolic computation and mechanisms. Again, this is a completely open problem for the Human Brain Project: It says nothing about studying how symbolic computation goes on in the brain. And the fact is we invented symbolic computation: we the brain-holders.

Next, the spiel turns into excessive, overarching, science unifying policy promises: "Neuromorphic computing is a new paradigm for ICT, opening the prospect of compact, low-power systems, with an unlimited range of applications, inaccessible to current technology. HBP will integrate this research on brain modeling and simulation, giving European laboratories and companies a critical competitive edge. HBP tools will open new vistas for research into the biological mechanisms of human consciousness."

I guess I should be wrapping up. I went to the internet and found out what critics of the project are saying[5], and it's good there is an open discussion on the matter. I selected, of course according to my own inclination, what I thought more congenial to my own point of view. And here is one of the things I think is quite valid: "We need more variance in neuroscience"[6]. We need variability, you cannot bet everything on one horse or bet it on one God. Right now we need several gods. And so we need as many different people, expressing different opinions. These diversities are threatened if scarce neuroscience research money is diverted mostly to a single and ever more absorbing project: the HBP.

Many neuroscientists, moreover, think HBP is ill-conceived: they see it as an idiosyncratic approach to brain simulation, which strikes them as cumbersome, and over-detailed. They see HBP as overhyped, thanks to breathless media reports about what will be accomplished. And they are unsure HBP is a project truly open to other ideas.

At the core of HBP, others say, is the conviction that a unifying model has to assimilate data from the bottom up. Models start at the most basic level, and specify its workings before moving to the next one. Gaps are filled with data as new experiments are done. The alternative, approximating and abstracting the biological detail, does not ensure that the model's behavior relates to how the brain works. A combination with top-down paths is needed.
Some computational neuroscientists already use models of individual neurons to explore high-level functions, such as pattern recognition. HBP's bottom-up approach risks missing the wood for the trees: such a detailed model may be no easier to understand than the real brain. That is, if HBP can build the model at all. Judging by its leading up project – the Blue Brain Project[7], which we heard about – and by its results in the past six years, that seems unlikely, they go on.

Critics question: In the Blue Brain Project what have they achieved? They don't think much has been achieved, even in the mouse cortex. The tiny simulated rat cortex has no sensory inputs, nor

---

5 http://www.scientificamerican.com/article.cfm?id=massive-brain-simulators-seung-conntectome
http://www.scientificamerican.com/article.cfm?id=brain-project-neuron-to-whole-brain

6 Kevan Martin, co-director, Institute of Neuroinformatics, Zurich, Switzerland.

7 Blue Brain Project - Wikipedia, the free encyclopedia

outputs to other brain parts, and produces almost no interesting behavior. They don't think that Blue Brain has simulated the cortical column, one that works with the whole animal. And an animal lives in the environment, in a context, is embedded in some situation – is situated.

Critics cannot see how Blue Brain's level of detail, very incomplete after considerable labors, will ever be obtained for more than a few regions of the rat brain, in the next decade, even with the start of HBP. Let alone the brains of the *drosophila, bee, zebrafish, octopus, songbird, mouse, or monkey.*

And so they say: Why don't we start with the drosophila, with the bee? – there is actually an ongoing so-called Green Brain project[8], studying the brain of the bee and other animals, including octopuses if ethics permits. And so on and so forth: Why not start with much simpler brains? And when we can show we can actually simulate these simple things, then we can make a better case for moving on to the bigger things.

So this is the problem: It's a Great project. And all would be just a storm in professors' teacups if HBP had not raised the stakes enormously. The million dollar question – actually the billion euro question – is this: Does Europe want to spend a billion on HBP?

It is easy to imagine other areas of neuroscience research being starved for resources by HBP. Especially in Switzerland, as the host country must provide a substantial fraction of the co-funding. Should Europe spend €1 billion on this passionate quest? Visionaries are necessary to drive progress, but what if they're passionately wrong?

Even now the Swiss government committed, or is in the process of committing, 81 million dollars to the continuation of the Blue Brain Project. So critics say: Is this not enough to keep this line of research alive?

In fact, the leader of the project, Henry Markram, has stated[9]: "Well, if HBP is not funded that's not a real problem; we'll just continue with Blue Brain. Though it will take more time to reach a full brain simulation, it is inevitable to do it."

I think so too: Science should perform ever more sophisticated simulations of the brain, mathematically and in the computer. It is an excellent scientific initiative, with an important impact on brain health concerns. But there exist more ways than one to go about it and, moreover, a wider interdisciplinarity should be brought to bear on it.

Thank you for your attention.

---

[8] https://www.shef.ac.uk/news/nr/green-brain-honey-bee-model-sheffield-university-1.212235

[9] http://www.nature.com/news/computer-modelling-brain-in-a-box-1.10066