CBMM Memo No. 118 September 1, 2021

From Marr's Vision to the Problem of Human Intelligence

Tomaso Poggio

This work was supported by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF-1231216.

From Marr's Vision to the problem of human intelligence[1]

Tomaso Poggio
McGovern Institute for Brain Research, Center for Brains, Minds and Machines,
Massachusetts Institute of Technology,
Cambridge, MA 02139, U.S.A.

1. The Vision book

The link between computation and neuroscience -- the realization that the brain is a computer -- is old. Turing wrote about it. McCullock, Pitts and Lettvin followed the idea from the perspective of both computation and of neuroscience. Seeds of the idea can be found in centuries-old writings. Though it may not be true that Marr's book *Vision* started the field known as computational neuroscience, it is certainly true that it had a key role in its beginning and rapid growth. A few years ago at Cosyne, the main conference for computational neuroscience, I mentioned David's work in my keynote talk. In the days afterwards, a surprising number of well-known researchers came to me to recount how they entered the field after reading David's book and how their career was due to that book!

2. Forty years later

We still do not understand the brain. Of course, this is not surprising. The problem of intelligence -- of how intelligence is created by the brain and of how to make intelligent machines -- is one of the greatest problems in science, possibly the most fundamental of all.

In the '76 when I was working with David for a three-months period at MIT, we fully realized that a satisfactory understanding of human intelligence was far away because the problem was so deep and so difficult. We hoped however that

[1] This is the preface to a Chinese translation of Vision, curated by JiaJun Wu.

computational ideas could help decrypt puzzles in the neuroscience, in particular neuroscience of the visual system.
At the time, even a proof of concept that 3D reconstruction could be done via 2.5D sketches was not possible. However, the situation has changed drastically over the last decade with machine intelligence making significant progress. We have MarrNet by JiaJun Wu which achieves 3D shape reconstruction via 2.5D sketches. We also have Alexa, AlphaZero, AlphaFold and MobilEye. We have transformers and great progress in NLP. The engineering community feels that machine learning and its recent network-based architectures are a powerful paradigm, potentially leading to the creation of intelligent machines. So…what about human intelligence, the real interest of David and mine?

3*. Many forms of intelligence*

I believe that there are many forms of "intelligence". Are computers that beat humans at chess and Go more inteligent than human players? Were computers in the 50's more intelligent than mathematicians because they could perform numerical integrations much faster? Are Residual Networks more intelligent than us because they may get better classification of certain image database? Intelligence is a pretty vague word, also because we apply it to systems, like future machines, that do not exist. Human intelligence is well defined: it is a natural phenomenon produced by biological brains. Science can study both brains and the behavior they produce, with experimental and theoretical techniques.

4. *Solving biological intelligence*

The previous arguments implies that networks performing well in object recognition are not by themselves the solution to the problem of how visual cortex works, though they may help. A recent trend in neuroscience is to fit activities of neurons in visual cortex with activity of units in RELU networks — such as AlexNet — trained with backpropagation. The reasonable agreement that has been reported in this optimization process is encouraging, but there is still a long way to go before claiming that these networks may lead to a plausible model of cortex. We will need to clarify what are the biophysical correlates of RELU nonlinearities, where are they in visual cortex, where are the weights, how are they modified and how activities of spiking neurons map into the static units of today's deep networks. More importantly, back propagation and batch learning of labeled data are almost certainly biologically implausible. Thus we would need to replace gradient descent with online learning rules based on known biophysics of synapses. All of this and much more is required. It would be a long and difficult but feasible endeavor. It would be, however, very surprising if the final models — at a level of biological plausibility on par with the Hodgkin-Huxley equations — would behave exactly as the engineering models of today. It would be equally surprising if, in the process of understanding how visual cortex really works, we will fail to uncover new interesting algorithms. In fact, Marr wrote, perhaps exagerating somewhat *"... a neural net theory, unless it is closely tied to the known anatomy and physiology of some part of the brain and makes some unexpected and testable predictions, is of no value." (Marr 1975)*

5. Levels of understanding

I just emphasized that progress in understanding human vision and human intelligence requires a very close interaction between experiments and models. The ability of artificial vision systems to fit neural activity data is unlikely by itself to close the gap in our understanding of biological vision. Progress will require models that make non-trivial falsifiable predictions and are tightly linked not only to visual performance but especially to the underlying neural circuitry and biophysics.

If the argument above seems to contradicts the levels of understanding framework described in Marr's *Vision*, the reason is that I do not believe that the levels of understanding framework strictly applies to the brain. The simple observation David and I wrote about, was that a complex system -- like a computer and like the brain -- should be understood at several different levels such as the hardware, the algorithms and the computations. In the *Vision* book, David emphasized that explanations at different levels are largely independent of each other: a software engineer does not need to know the hardware in any great detail. Forty years ago the message was novel and important: the study of the problems to be solved -- and of the associated computations -- is relevant in its own right and is needed for a full understanding of the brain. The section, however, in the *Vision* book about levels of understanding was based on an argument on the visual system of the fly by Werner Reichardt and myself, where we stressed that one ought to study the brain at different levels of organization, from the behavior of a whole animal to the signal flow, ie the algorithms, to circuits and single cells. In particular, we claimed that it is necessary to study nervous systems at all levels simultaneously. From this perspective, the importance of coupling experimental and theoretical work in the neurosciences follows directly: without close interaction with experiments, theory is very likely to be sterile.

6. Why human intelligence?

For many researchers, a reason to study human intelligence and the human brain was the bet that this is the best way to develop artificial intelligence. Over the last 10 years, however, the situation, as I mentioned earlier, changed somewhat. My friend Demis Hassabis said a few months ago that in his estimate, the probability of engineers winning the race for intelligence without help from neuroscience went from 10% to about 50%. I think these are still pretty good odds for a bet on neuroscience as a good approach to solve AI!

But …let us assume that neuroscience will not help AI. Is then still a good idea to "invest" in computational neuroscience? My answer is yes. Here I am telling you why.

The first point is obvious: understanding how our brain creates human intelligence is at least as fascinating and exciting and important as understanding the cosmos with its planets, stars and black holes. Curiosity-driven science simply demands an exploration of the universe of our brain and our mind. After all, they are the very tool we use to understand everything else.

The second point is based on my belief that there exist various forms of intelligence, probably an infinite number of them. It is then quite unlikely that we could build an AI identical to human intelligence unless we know all the constraints under which evolution created the human mind. This is exactly what we see happening with machine learning today. CNNs and transformers are obviously different from us: both leverage large data produced by humans, something evolution could not have done. In summary, such machines may be "intelligent" but in a way that is quite different from human intelligence.

7*. How to understand human intelligence*

Emphasizing the connections between levels is also a recognition that the problem of understanding the brain is very difficult and we need to use every approach and every technique we have. It is important also to recognize, as I mentioned, that the emphasis on coupling the different levels, *de facto* implies an emphasis on a very close interaction between experiments and models as a necessary condition for progress in understanding the brain through a computational lens.

REFERENCES

Marr D, 1969 “A theory of cerebellar cortex.” The Journal of Physiology, 202 437-470.

Marr D, 1975 “Approaches to Biological Information Processing” Science, Vol. 190, No. 4217, pp. 875-876.

Marr D, Poggio T, 1977 “From Understanding Computation to Understanding Neural Circuitry” In: Neuronal Mechanisms in Visual Perception, E. Poppel, R. Held and J.E. Dowling (eds.), Neurosciences Research Program Bulletin, 15 470-488.

Reichardt W, Poggio T, 1976 “Visual Control of Orientation Behavior in the Fly: A Quantitative Analysis” Quarterly Review of Biophysics, 3 311-375.

Wu Jiajun et al., 2017 “MarrNet: 3D Shape Reconstruction via 2.5D Sketches”, NIPS