

Review of: "Strong Machine Learning: a Way Towards Human-Level Intelligence"

Luís Ferreira Pires¹

¹ University of Twente

Potential competing interests: No potential competing interests to declare.

General remarks

This paper discusses the so-called 'strong machine learning' approach, meant to make the learning process of machine learning algorithms more efficient and consequently less demanding on resources. This paper addresses a very timely topic and gives an overview of the available alternatives, with pointers for future developments. The style of writing is quite informal, with many claims not being substantiated, but the contribution is valuable as a position paper.

Detailed remarks (per section)

Introduction:

- I don't think it makes sense to talk about 'the maximal acceptable levels' because you will probably never know what they are. Any attempt to make ML models more efficient is valuable.
- A statement like 'there is a possibility of a much smaller model...' is speculative and should be avoided. This is actually the hypothesis of further research in this area.
- The comparison with how the human mind performs multiplication is flawed. People learn the whole multiplication table (for digits) by heart, and a trick to apply it iteratively. I don't see a resemblance here.
- It is unclear what the author means with 'more human-like technologies'. Is this about the difference between 'designed models' (models defined by humans) against 'learned models' (models learned by machines)? This should be clarified.
- Textual: asstrong -> as strong

Weak vs. strong machine learning:

- Textual: in order to developed -> in order to develop

Inductive biases: The key to strong machine learning:

- The whole paragraph starting with 'The current best models...' is speculative and should be rewritten and/or get more references to substantiate the claims. The same holds at the end of the section, with possible references to developments that learn inductive biases. Alternatively, the author can refer to the section in which such developments are discussed.

A brief overview of several existing strong machine learning techniques:

[Section title is unnecessarily too long]

- In part 3 (Zero-shot learning), when the knowledge to recognise zebras as 'striped horses' is introduced into the system, I

guess this needs to be done in a controlled way to preserve semantics, which would require, for instance, the use of ontologies. The author should clarify whether he is considering ontologies to be relevant or even essential here.

The tradeoff between generality and specificity

- A sentence like 'the higher the specificity, the lower the generality' sounds trivial to me. This is not a 'tradeoff', but how these concepts are defined and related to each other.