ELL457/HSL622: Minor Test I

February 7, 2023

Maximum Marks: 15

Instructions:

- **Please clearly indicate the question number at the start of each response.**
- **Please read all questions carefully.**
- **Please ensure that your responses are to-the-point and that you write only what is asked for on the answer script you submit.**
- **While the exam is open-notes, all your answers must be written entirely in your own words, without any copying from anywhere.**

1. Recall the 3 levels of Marr's hypothesis: *computational/functional*, *algorithmic*, and *implementational*. Thinking of the mind as a kind of computer for now, and restricting ourselves to the domain of language processing, how much similarity do you think there is between the mind and modern AI systems at each of these levels? For each of the three levels, try to give a percentage to roughly quantify your judgement about the amount of similarity, along with a brief explanation of your reasoning behind this. **[6]**

   Hint: It may be helpful to think about specific AI systems you have come across which do some sort of language processing, like Google Translate or ChatGPT or any others, though of course detailed knowledge of any system is not required.

   Key points:

   - *Computational level* – Probably the greatest similarity at this level; in certain domains AI systems are able to produce largely human-like responses. At the same time, human language use clearly interfaces a lot with more abstract reasoning about causal relations etc., and AI systems often seem to falter on tasks or responses which require such reasoning. So AI systems may still largely be doing something more similar to relatively shallow or surface-level human linguistic processing.
   - *Algorithmic level* – Harder to assess, as we don't really understand fully the algorithms the mind uses, especially for more 'subconscious' tasks like much of language processing. There is probably some degree of overlap in terms of simple things like learning co-occurrence probabilities of words etc.; and plausibly networks of biological neurons carry out computation in ways that have some commonalities with the deep learning models prevalent in AI today. But there are also likely to be substantive differences, and especially for language processing there is little evidence of direct correspondence.
   - *Implementational level* – The overlap here would seem to be the least, perhaps none. In general modern digital computers seem to operate very differently in terms of their physical operations/mechanisms than the brain. There are some attempts to build neuromorphic computing systems that emulate some aspects of neuronal operation more directly, but these are still quite limited and have little prevalence in current AI systems.

2. Explain why exactly the traditional idea of 'free will' seems incompatible with physicalism. How does this relate to the principle of the *Causal Closure of the Physical (CCP)*? **[3]**

   Key points:

   - Free will posits that there is a component of our decision-making process which is *sui generis*; entirely autonomous and independent of external influences or conditions.

- Physicalism implies that all cognition is nothing more than physical processes; and every physical event appears to be a consequence of physical causes and circumstances (CCP). So seems to rule out an uncaused 'free will'.
- As per CCP, if there were something like free will which is not determined by physical factors/forces, then it would also be unable to have physical effects or consequences.

3. Explain concisely the relationship between the epistemological nature-nurture debate and the two major paradigms of computing in AI: classical and connectionist. Please *do not* give generic descriptions of any of these things, but only a *specific* answer about how the two sides of the nature-nurture debate have been reflected in the use of these two computing paradigms for developing AI systems. **[3]**

   Key points:

   - Nature-nurture debate is about the sources of human knowledge and behavioural tendencies: whether innate from birth or acquired via experience.
   - Classical or symbolic AI systems largely relied on hardcoded knowledge (rulebases etc.); this paradigm mostly sought to endow AI with the right 'nature' to enable it to accomplish certain kinds of tasks. Scope for the AI learning from experience has been limited in such systems.
   - Connectionist AI proved so successful precisely because it shifted the burden of knowledge largely from 'nature' to 'nurture': These systems could much more effectively leverage data/experience to adapt and learn over time. However it seems difficult to directly encode human-like innate knowledge or reasoning mechanisms into these systems, which may account for some of their limitations.

4. Friedenberg and Silverman say that

   When you bite into a candy bar, you have a subjective conscious experience of tasting it. The candy bar of course has no such experience. There is nothing that "it is like" for the candy bar being bitten.

   In what way is this observation problematic for the functionalist view on the mind-body problem? Explain concisely. **[3]**

   Key points:

   - The functionalist view treats mind as a *functional kind*: as just a name for all the cognitive functions carried out by the brain/body.
   - Seemingly, observing or describing these functions does not require reference to subjective conscious experience or qualia. One can imagine a functionally-identical mind without qualia; hence functionalism fails to account for the source or relevance of qualia.
   - We seem to have a strong intuition that entities with a 'mind' have some form of subjective experience, and that this is a key difference between them and entities without a mind (like candy bars). This intuitive aspect of what we call a 'mind' appears to be left out of a purely functionalist account.