

## PETER'S ABSTRACT

All knowledge is statistical, only more or less certain. Often there are conditions under which a face might be hard to identify, or an object might be considered to be two different colours at once, more or less. Meanings in natural language have similar properties. In English, for example, we might say that an object is a book, or like a book. The set of discriminating properties underlying the application of the word 'book' can each be applied more or less, so there can be borderline cases where something is a bit like a book, but also like something else. This has consequences for the use of the word 'logic'.

If we are honest about the way we assess discriminating properties of objects and events, logic is actually more about the extent to which fuzzy discriminants have overlapping references. For people, discriminations and creations of references are the same intuitive (or sometimes subconscious) act. For example, the extent to which we think an object is red depends on how many of our 'red-detecting' neurons are firing. We can then think of red as being a predicate composed of some cumulative function (in reality, the red-detecting neurons), and the redness of an object being determined by an interval on the domain of that function. When we look at an object our redness cumulative function might register 60% +/- 3%, or at night 30% +/- 20%. So we have to pay heed to the statistical ideas of a mean and confidence intervals. Thus a reference becomes akin to a confidence interval on the cumulative function of some predicate (whether the predicate is simple or compound).

A compound predicate will be composed of confidence intervals along the domains of many cumulative distribution functions at once (a multidimensional vector of confidence intervals). Some of these contributing functions may themselves be the output from some compound of functions. So, for example, a compound predicate like 'book' might be a set of functions that somehow maps to a single output - the degree of 'bookness' measured on the cumulative function of 'book'. It is unclear how our minds make these mappings, or learn to make them, and we may have to invent rough Ansätze (plural) for such mapping and learning processes within knowledge representation schemes.

Based on the above, the knowledge that a predicate represents might be efficiently represented in a knowledge representation scheme by capturing mathematical representations of the appropriate cumulative functions and appropriate parameterization of confidence intervals.

Both IS-A and A-KIND-OF relations might have improved representations under this scheme. If we describe a sentence as a chain of such predicates, and logical entailments as overlaps of predicate codomains between sentences, then knowledge federation becomes more about overlap of

codomains (discrete or continuous). One possible advantage of this approach is that its fidelity should improve and become more robust as the available data increases. Another avenue is exploring mappings between the intensity and configuration of outputs of neural networks and generation of statistical profiles. In any case, it is argued that an urgent project for the future of knowledge federation is the investigation of computationally efficient manipulations of parameterizations over sets of cumulative distribution functions.

It's just another way of saying that our thoughts ride (supervene) on neurons, but tackling the idea that alternative (processes of construction of) representations might provide for two things:

- 1) An ability to have mergeable knowledge representation schemes 'on top of' /any/ underlying perceptron/neuron models - plug and play.
- 2) The possibility that knowledge representation schemes might be decoupled from the pressure to model entire brains. So that e.g. AI could have transducers to sense the world at the outer layer, but then skip loads of the neural preprocessing that we do, and go direct to symbolic/computational manipulation as a key level of cognition that might prove faster on silicon for some reason yet to be discovered.

Immanuel Kant argued well that our (normed notion of) logic is a function of the way we are (in general) made. And if our perceptual and mental faculties have variation distributed throughout the population, then maybe there is a (multi-dimensional) Gaussian distribution for all the ways humans might think 'logically'. E.g. creativity might revolve around an ability to uninhibit some small degree of variance in the way some part of the brain handles overlapping of codomains.

I guess another way of putting this is that we shouldn't be too surprised to find that silicon based intelligence (if and when it arrives) genuinely thinks differently from ourselves - often reaching different conclusions. This has, as I see it, profound implications for the role of knowledge representations within our social mechanisms. Hence it is important to begin exploring these issues.

Knowledge federation and knowledge gardening raise, in the sense that they are proto-foundations, the possibility that we might eventually require an ISO standard that enforces mathematical constraints on KF subject mergers in order to bound the 'rationality' of agents that employ the knowledge representations.

I agree that this kind of speculation is not yet science, but I'm willing to place a (small) bet that this is the way things will head very soon.

I see this kind of paper as sitting at the fringe (lunatic?) of discourse, which is precisely where the new stocks & flows (power of pull) people would have us migrate to. The edge.

Yep, loony fringe - but perhaps the back edge of the fringe. If I can see this sort of thing coming that means it is probably nearer than folks in general will consider.