An Evolutionary Approach to the Training of Human-Like Computing systems

David Frost

London
david@kulperfrost.com
www.linkedin.com/in/drdavidfrost

Abstract. The functional requirements of Human-Like Computing (HLC) are similar to the features of behavioural modernity which emerged c. 80k years ago. Compositional symbolic language allowed observational learning to be enhanced by explanation and verbal instruction. Humans become human-like by growing up in the nurturing environment of a human family. HLC systems will similarly need to boot-up within a virtual socio-cultural environment equivalent to that of humans. The degree of abstraction used to create this environment needs to preserve sufficient level of detail to enable relevant behaviour to emerge. A key objective function in selecting benchmarks is the transmission of cultural objects, including language.

Keywords: Symbolic Language, Language Transmission, Human-Like Computing, Virtual Socio-Cultural Environment, Artificial Life.

Introduction

The functional requirements of Human-Like Computing include the ability to learn from small numbers of examples, to construct their own models of the world, to communicate in natural language, and to provide explanations [1]. These requirements are similar to the features of behavioural modernity which emerged c. 80k years ago. Compositional language and symbolic thought were key to this new functionality and allowed social (observational) learning to be enhanced by explanation and verbal instruction. They also facilitated social cooperation and the creation of cultural objects [2].

Pre-symbolic human evolution was an adaptation to the existing physical and social worlds. In its later stages, proto-languages included indexical reference. In the transition to symbolic compositional language, the cognitive processes adapted to physical and social environments were co-opted for symbolic representation [3]. An understanding of this tipping point, the symbolic threshold, is key to the development of HLC systems.

Intergenerational (vertical) transmission of language occurs in the early sensori-motor and pre-operation stages of child development. Language is acquired in parallel with spatio-temporal skills, as if it were just another feature of the natural world (even though it creates a unique inter-subjective space). Humans become human-like by growing up in the nurturing environment of a human family. HLC systems will similarly need to boot-up via “nurturing” interactions within a human-like family.

This paper discusses how an HLC system would be trained. What would the training dataset look like? Where would we get it from? How would we define benchmarks?

Training

Humans are "trained" through the process of child development, a sequence of sensori-motor, linguistic, cognitive and emotional changes which take place through immersion in a human psycho-socio-cultural environment. In the initial stages, the representational world is kept simple through a combination of parental protection and sensory immaturity. Over time, children are exposed to increasingly complex social situations. The process is not that dissimilar to simulated annealing.

Human learning is not passive but requires interaction with both a physical and social world that provide feedback (reinforcement). In particular, it is the ability to understand the complex social roles and behaviours of others, rather than just the physical world, that distinguishes humans from other species.

At an abstract level, child development is the synchronisation of two identical chaotic systems, culture being the driving system and the child being the response system. The systems are coupled though carefully phased changes in the linking of their dynamic variables [4]. Autism and schizophrenia are examples of where that synchronisation is impaired.

Although Haeckel's phrase "ontogeny recapitulates phylogeny" has been widely challenged, there are indirect parallels between cognitive development in children and the evolutionary stages of man [5]. A similar equivalence can be argued between language acquisition and language evolution, except that they also require a complex boot strapping process. The language to be learnt is itself created and sustained by the iterative behaviour of individuals within the cultural group.

The transition between indexical reference and symbolic language involves the creation of persistent conceptual representations and associated compositional structures. The acquisition of symbols is contingent on theory of mind through which the meta-perceptions of others coalesce into a shared inter-subjectivity [6]. These conceptual representations are grounded since they are based on knowledge generated within the enclosing culture [7].

The training of HLC systems therefore involves the gradual exposure to increasingly complex social situations, the acquisition and transmission of indexical reference, and the emergence of symbolic language along with other cultural objects.

Virtual socio-cultural environment

HLC systems need to "grow up" in a socio-cultural environment equivalent to that of humans. The degree of abstraction used to create that environment has to preserve sufficient level of detail to enable relevant behaviour to emerge. The environment can be virtual, as in the Artificial Life paradigm in which multiple autonomous virtual agents interact and evolve. Low level sensori-motor functions, such as vision and movement can be finessed.

The virtual socio-cultural environment needs to be able to model different stages of human evolution, for example: 1. The physical world (e.g. sensori-motor processes, basic animal behaviour), 2. The social world (e.g. communication, theory of mind, proto-language), and 3. The cultural world (symbolic language, behavioural modernity).

The transitions between these stages all occurred in hunter-gather communities. In the first transition, between sensori-motor and social worlds, there was a co-evolution of neural structures and proto-languages [8]. In the subsequent transition, symbolic language appeared too suddenly for evolutionary changes in the brain.

Each stage needs to be defined at various levels of abstraction, for example: 1. A macro-level model of group behaviour (informed by ethology and anthropology), 2. A meso-level model of personal and inter-personal behaviour (informed by psychology), and 3. A micro-level model of neural processes (informed by machine learning and cognitive science).

The objective functions of each level are set by the behaviour of the level above. The hyper-parameters shape the overall evolutionary path from a macro-level model of the sensori-motor world to a micro-level model of symbolic culture.

Benchmarking

Benchmarks are set at various points along the evolutionary path. A key objective function is language transmission. In exploring the emergence of language, the question is not when (or how) when this occurred, but rather what is the nature of a language-ready system that can sustain the acquisition and transmission of meaningful utterances between generations.

One approach is to seed a subset of the virtual population with an algorithmic-based natural language in which each utterance has associated rules about when and why it is used. The objective function measures the degree to which the utterances are learnt by the rest of the population and transmitted to subsequent generations. (This is similar to a paradigm used in experiments with the great apes, except that the algorithmic-based individuals were human.) At subsequence stages, the utterances can be linked to more complex social concepts like trust or deception.

The key point is that the benchmarks must be defined from within a relevant socio-cultural environment. Behavioural modernity emerged within hunter-gathers, small groups of individuals trying to survive within a very challenging world [9]. The archaeological record provides evidence of art, religion and complex social structures. Rationality plays some part in the rhetoric of explanation but has not been the defining feature of being human. In our choice of benchmarks for HCL systems, there is a danger of cherry-picking objective functions which create a super-rational but emotionally limited version of ourselves.

References

1. A Strategy Roadmap for Human-Like Computing, https://epsrc.ukri.org/newsevents/pubs/human-like-computing-strategy-roadmap, last accessed 2019/06/16.
2. Mithen, S.J.: The prehistory of the mind: a search for the origins of art, religion, and science. Thames and Hudson (1996).
3. Deacon, T.: The Symbolic Species: the co-evolution of language and the brain. Allen Lane (1997).
4. Abarbanel, H.D.I.: Analysis of Observer Chaotic Data. Springer (1996).
5. Borchert, C.M., Zihlman, A.: The ontogeny and phylogeny of symbolizing. In: The Life of Symbols ed. M.L. Foster et al. Westview Press (1990).
6. Hudson, R.: Review of Terrence Deacon: The Symbolic Species. Journal of Pragmatics, 12/08 (1999).
7. Wilden, A.: System and Structure: Essays in Communication and Exchange, 2nd ed., Tavistock Publications (1980).
8. Eccles, J.C.: Evolution of the Brain: Creation of the Self. Routledge (1989).
9. Hamilton, M.J. et al.: The complex structure of hunter-gatherer social networks. Proc. R. Soc. B 274, 2195-2202 (2007).