

Minds, Selves and 21st Century Technology

23rd and 24th June, 2016



Lisbon, Portugal

Predicting Me, Experiencing Us: Predictive Processing, Big Data and the Mind of Society

Paul Smart, Kieron O'Hara and Wendy Hall

Electronics and Computer Science, University of Southampton, Southampton, UK

An increasingly popular view in the sciences of the mind sees the biological brain as a hierarchically-organized prediction machine, one in which higher-order neural regions are continuously attempting to predict the activity of lower-order regions at a variety of (increasingly abstract) spatial and temporal scales (Friston, 2005, 2009). The role of the brain, if this view is correct, is simply to minimize the error that results from the brain's attempt to predict its own patterns of neural activity (many of which, of course, are driven by information originating from the sensory surfaces). Such a view seems to afford a great deal of explanatory leverage when it comes to a broad swath of seemingly disparate psychological phenomena (e.g., learning, memory, perception, action, emotion, attention, planning, reasoning, and imagination) (see Clark, 2016), and this has led to the suggestion that predictive processing models might provide our "best clue yet [as] to the shape of a unified science of mind and action" (Clark, 2013, p. 181). Some have even suggested that predictive processing models can help to shed light on the neural bases of conscious experience (Seth, 2013; Seth et al., 2011). Hohwy et al. (2008), for example, provide a predictive processing account of binocular rivalry, which has often been used as a tool to probe the neural correlates of conscious (visual) experience (see Frith et al., 1999).

The ability to account for a broad range of human psychological phenomena highlights one reason why predictive processing models are the focus of current theoretical and empirical attention within the sciences of the mind. A similar interest, however, can also be found in attempts to advance the current state-of-the-art in machine intelligence and machine learning. Interestingly, the vision of the brain as a hierarchical prediction machine is one that establishes contact with work that goes under the

heading of 'deep learning' (Arel et al., 2010; Bengio, 2009). Deep learning systems thus often attempt to make use of hierarchically-organized predictive processing schemes and (increasingly abstract) generative models as a means of supporting various forms of intelligent response. As is indicated by the recent successes of Google Deepmind (a company devoted to the commercial exploitation of deep learning systems), such architectures have the potential to achieve (and sometimes even surpass) human-level competence in a variety of disparate problem domains (e.g., Mnih et al., 2015).

Based on their ability to capture deep statistical regularities in complex bodies of data, deep learning architectures are often seen as particularly suited to the analysis of big data (see Najafabadi et al., 2015). The sources of big data are many and varied; however, two sources of big data are of particular interest in the present context. The first relates to the data that arises as the result of self-tracking activities and the increasing use of wearable computing devices (see Swan, 2013). Let us call such forms of big data 'personal data'. Another source of big data is the data that arises as a result of the direct or indirect monitoring of large-scale forms of social activity. Let us call this kind of big data 'social data'. Crucially, both kinds of data have become increasingly common as a result of our interactions with the Web and Internet (i.e., the online realm). This is particularly so in the case of social data, where online Web-based systems are sometimes said to function as 'social machines' that provide insight into the nature of social interactions and exchanges (see Smart & Shadbolt, 2014).

A broad array of philosophical issues begin to emerge once we consider the relationship between predictive processing models, the analysis of big data, and the possible emergence of artificial intelligence

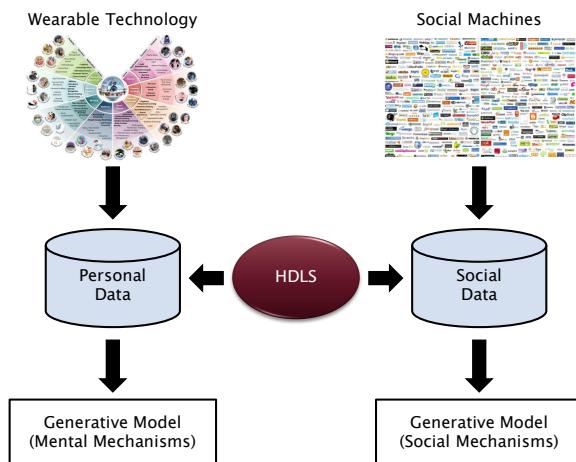


Figure 1: Homomimetic deep learning systems (HDLS) help to support the creation of generative models that are of predictive relevance to both personal and social data. Personal data (derived from, for example, wearable devices) provides information about the behavioral and physiological responses of the individual, as well as information about the kinds of environments in which an individual is embedded. Social data (derived from, for example, social machines) provides information about social processes.

systems, especially those that have as their focus of interest the predictive modeling of either the individual (in the case of personal data) or social groups (in the case of social data). Inasmuch as predictive processing schemes are able to support the sort of capabilities we associate with advanced cognition (including aspects of conscious experience), it seems fair to ask to what extent a hierarchically-organized predictive processing system might help to support analytic efforts that lie at the heart of attempts to understand both ourselves (as individuals) (e.g., Lupton, 2013) and the societies in which we live (e.g., Lupton, 2015). Beyond this, however, we want to raise the possibility that a particular kind of deep learning system—dubbed a ‘homomimetic deep learning system’ (HDLS)—could serve as the material basis for experientially-potent forms of machine cognition (see Figure 1). In the case of personal data, we suggest that predictive processing systems could help to capture elements of the self based on the data generated by our biological bodies. Such claims help to inform debates concerning the possibility of ‘mind uploading’, as well as the prospects for future forms of ‘digital immortality’ and the ‘digital resurrection’ of the self. In the case of social data, we speculate that predictive models could serve as the basis for novel forms of intelligent system that are informed by the dynamics of the social world. Such systems, we suggest, also provide something of a

novel analytical tool for those who are concerned with mechanism-based explanations in the social sciences (Hedström, 2005; Hedström & Bearman, 2009; Hedström & Ylikoski, 2010).

Acknowledgements

This work is supported under SOCIAM: The Theory and Practice of Social Machines. The SOCIAM Project is funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/J017728/1 and comprises the Universities of Southampton, Oxford and Edinburgh.

References

Arel, I., Rose, D. C., & Karnowski, T. P. (2010) Deep machine learning – A new frontier in artificial intelligence research. *IEEE Computational Intelligence Magazine*, 5(4), 13-18.

Bengio, Y. (2009) Learning deep architectures for AI. *Foundations and Trends in Machine Learning*, 2(1), 1-127.

Clark, A. (2013) Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science. *Behavioral and Brain Sciences*, 36(3), 181-253.

Clark, A. (2016) *Surfing Uncertainty: Prediction, Action and the Embodied Mind*. Oxford University Press, New York, New York, USA.

Friston, K. (2005) A theory of cortical responses. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 360(1456), 815-836.

Friston, K. (2009) The free-energy principle: A rough guide to the brain? *Trends In Cognitive Sciences*, 13(7), 293-301.

Frith, C., Perry, R., & Lumer, E. (1999) The neural correlates of conscious experience: An experimental framework. *Trends In Cognitive Sciences*, 3(3), 105-114.

Hedström, P. (2005) *Dissecting the Social: On the Principles of Analytical Sociology*. Cambridge University press, Cambridge, UK.

Hedström, P., & Bearman, P. (Eds.) (2009) *The Oxford Handbook of Analytical Sociology*. Oxford University Press, Oxford, UK.

Hedström, P., & Ylikoski, P. (2010) Causal mechanisms in the social sciences. *Annual Review of Sociology*, 36, 49-67.

Hohwy, J., Roepstorff, A., & Friston, K. (2008) Predictive coding explains binocular rivalry: An epistemological review. *Cognition*, 108(3), 687-701.

Lupton, D. (2013) Understanding the Human Machine. *IEEE Technology and Society Magazine*, 32(4), 25-30.

Lupton, D. (2015) *Digital Sociology*. Routledge, Abingdon, Oxon, UK.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A., Veness, J., Bellemare, M., Graves, A., Riedmiller, M., Fidjeland, A., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015) Human-level control through deep reinforcement

learning. *Nature*, 518, 529-533.

Najafabadi, M. M., Villanustre, F., Khoshgoftaar, T. M., Seliya, N., Wald, R., & Muharemagic, E. (2015) Deep learning applications and challenges in big data analytics. *Journal of Big Data*, 2(1), 1-21.

Seth, A. K. (2013) Interoceptive inference, emotion, and the embodied self. *Trends In Cognitive Sciences*, 17(11), 565-573.

Seth, A. K., Suzuki, K., & Critchley, H. D. (2011) An interoceptive predictive coding model of conscious presence. *Frontiers in Psychology*, 2(395), 1-16.

Smart, P. R., & Shadbolt, N. R. (2014) Social Machines. In M. Khosrow-Pour (Ed.), *Encyclopedia of Information Science and Technology*. IGI Global, Hershey, Pennsylvania, USA.

Swan, M. (2013) The quantified self: Fundamental disruption in big data science and biological discovery. *Big Data*, 1(2), 85-99.