White Rose Research Online

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk
Universities of Leeds, Sheffield and York

Deposited via The University of Leeds.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/id/eprint/99260/

Version: Accepted Version

Proceedings Paper:
Evans, MH, Fox, CW and Prescott, TJ (2014) Machines Learning - Towards a New Synthetic Autobiographical Memory. In: Duff, A, Lepora, NF, Mura, A, Prescott, TJ and Verschure, PFMJ, (eds.) Biomimetic and Biohybrid Systems. 3rd International Conference on Biomimetic and Biohybrid Systems: Living Machines 2014, 30 Jul - 01 Aug 2014, Milan, Italy. Lecture Notes in Computer Science, 8608. Springer Verlag, pp. 84-96. ISBN: 978-3-319-09434-2. ISSN: 0302-9743.

https://doi.org/10.1007/978-3-319-09435-9_8

Reuse
Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown
If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Machines learning - towards a new synthetic autobiographical memory

Mathew H. Evans, Charles W. Fox and Tony J. Prescott

Sheffield Centre for Robotics (SCentRo), University of Sheffield, Sheffield, S10 2TN, U.K.
{mat.evans, charles.fox, t.j.prescott}@shef.ac.uk

Keywords Synthetic, Autobiographical, Memory, Episodic, Hippocampus, Robotics, Predictive, Coding, Deep, Learning

Abstract. Autobiographical memory is the organisation of episodes and contextual information from an individual's experiences into a coherent narrative, which is key to our sense of self. Formation and recall of autobiographical memories is essential for effective, adaptive behaviour in the world, providing contextual information necessary for planning actions and memory functions such as event reconstruction. A synthetic autobiographical memory system would endow intelligent robotic agents with many essential components of cognition through active compression and storage of historical sensorimotor data in an easily addressable manner. Current approaches neither fulfil these functional requirements, nor build upon recent understanding of predictive coding, deep learning, nor the neurobiology of memory. This position paper highlights desiderata for a modern implementation of synthetic autobiographical memory based on human episodic memory, and proposes that a recently developed model of hippocampal memory could be extended as a generalised model of autobiographical memory. Initial implementation will be targeted at social interaction, where current synthetic autobiographical memory systems have had success.

1 Introduction

We receive a continual and very high band-width stream of sensory data during our waking hours. Our autobiographical systems process this data in a highly specific and adaptive fashion so as to provide quick access (within seconds) to relevant information experienced from hours to decades earlier. Autobiographical memory (AM) is defined as the recollection of events from one's life. Though similar in conception to Tulving's episodic memory [52], AM goes beyond simple declarative facts of an event to recall of rich contextual details of a scene [13]. AM is a prerequisite for developing the narrative self [32], related to and conceptually similar to the temporally extended self [43]. This version of the self is the individual's own life story, developed through experience, remembered, and projected into the future.

For robots to behave in a flexible and adaptable manner, and succeed in complex sensorimotor tasks, it is essential that they store their experience appropriately and use this information during online processing. For example, in human-robot interaction a memory of a familiar person or game would be intuitively advantageous especially in developing trust or attachment. This concept is familiar in probabilistic robotics, where an informative prior is essential for accurate inference. At a purely computational level, it has been shown that learning through 'episodic control' is more efficient than building a forward model or developing habits when experience is limited and tasks are complex, as is often the case in real-world robotics applications [24].

Innovative sensor designs, miniaturised HD cameras and affordable hard drives have endowed modern robots with an impressive capacity for gathering and storing information from the world. This information can be pooled across a range of modalities (e.g. vision, audition, touch, LIDAR,

depth) at high bandwidth. However, robots remain poor at extracting or retrieving task-relevant information when needed, either offline from their vast archives, or online during streaming [54].

In practice, data streams are filtered through feature detectors, processed by machine learning black boxes, or passively compressed into annotated histories where no such tools are available. As a result, despite having the ability to encode virtually everything that happens to them and access to vast stores of additional information online, robots are poor at determining which aspects of their history are important for making decisions and performing actions, or for framing engagement with people.

How does the brain solve this problem? By forming autobiographical memories that evolve over the life-time of an agent that places events in the context of the self and its goals. In this position paper we outline common approaches to autobiographical memory modelling in robotics, and contrast this with more modern understanding of memory function and organisation (see Figure 1). Four main principles are identified: compression, pattern completion, pattern separation and unitary coherent perception. To address each of these principles in a single architecture we outline a modelling framework based on predictive coding: deep learning for hierarchical representation and compression of sensory inputs (modality specific hierarchies in Figure 1, thought to reside in the sensory processing circuitry of the brain), and episodic memory formation through a particle filter and Boltzmann machine hippocampus model (Following Fox and Prescott [16], see Figures 2 and 3). We close with a discussion of our specific implementation goals, centring our episodic memory system within a rich sensorimotor system to aid ongoing processing.

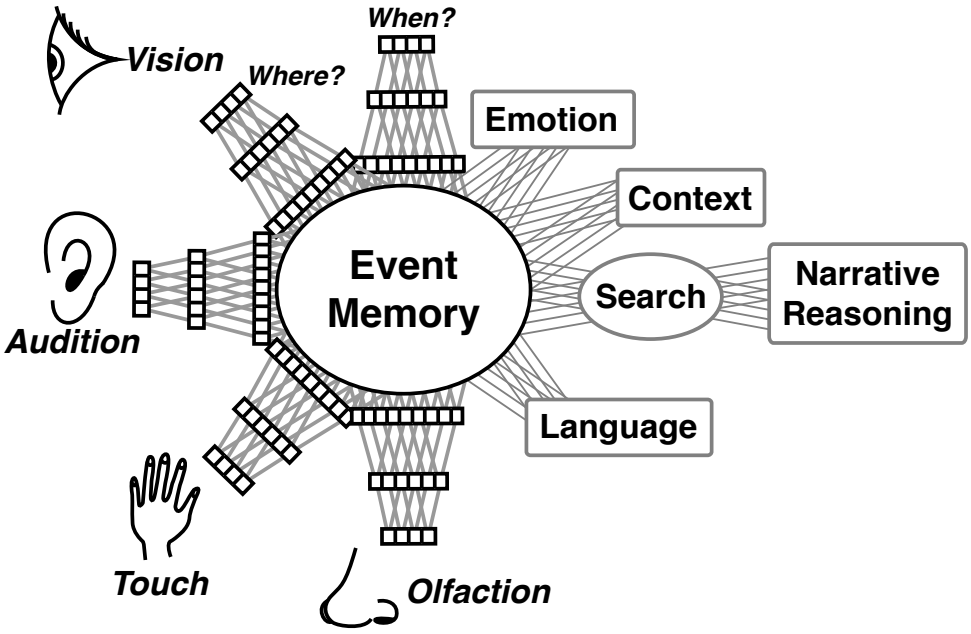

Fig. 1. An overview of our proposed model of biological autobiographical memory. Based on the basic systems model of Rubin [44], event memories are formed from convergent activity across disparate brain systems. Algorithmically we turn to deep networks for sensory processing [3], and predictive coding to frame learning [11]. Storage and recall is modulated by emotional, contextual and linguistic information. Narrative reasoning seeds searches, recognitions, reconstructions and predictions (the core of pattern separation and completion) by re-activating modality-specific schemas [21].

2 Current approaches to synthetic autobiographical memory

What is the state of the art in this field? Synthetic autobiographical memory (SAM) broadly has two flavours; (i) storage of unedited data streams (e.g. for maximum likelihood estimation [25]), or (ii) annotated high-level sequences (e.g. of social interactions [36]). The initial problem with approach (i) - storage capacity - is being alleviated to a degree by improvements in hard-drive capacity. However, storing a complete sensory history becomes un-wieldy quickly and the ability to search and utilise this vast store of information to inform ongoing processing becomes seriously problematic.

Recent success has been had in demonstrating the effectiveness of approach (ii) such as in cooperative tasks [36], or learning through social interaction [37]. These models are largely symbolic, similar in spirit to classic models of cognition like ACT-R [1] and others (reviewed in [54]), with hand-set higher level representations and action scripts coordinated into event memories [38]. As in early cognitive science, these models of memory have appealed to the computer metaphor for the mind, though thorough critiques of this approach have been articulated [10]. Pointedly, short-term and long-term storage of homogenous information is an idea from the 1960s but continues today in many models of memory [29, 5].

There remain fundamental gaps between cognition, memory and learning in these models. The compression of experience into memories is not handled in an adaptive way as it is in nature, and these models neither capture the rich feature-set of biological memory nor exploit advances in machine learning to increase the power and flexibility of stored information. Our position is that memory, and as a result our model of AM, is central to action and formed adaptively through experience with the world.

3 Characteristics of biological autobiographical memory

As noted in Wood et al. [54], the cognitive science and neurobiology of memory has changed drastically over recent years, but synthetic approaches have not kept pace. In particular, models of episodic and autobiographical memory are not compatible with theories of active, distributed memory systems [44] or ideas such as the predictive brain [11]. Forming new memories is not only about data compression, though this is important, but about selectivity and efficiency in the mechanisms of memory coding and retrieval. Cognitive scientists and neurobiologists assert that long-term memory formation is not a passive process of logging data into generic storage [2][44], but a highly active process depending on factors such as the depth to which an event is processed, or the wider context of the current task [26][42].

Biological memory is highly distributed and tightly-woven into ongoing cognition [54]. The neural underpinnings of different kinds of memory and imagery cannot be separated from the circuitry of perception and action. For example, when a person is asked to imagine rotating an object, neural activity is elevated in the visual and somatosensory cortices [12].

In order to match the function of human AM, a SAM system should have the following properties.

– **Compression.** Human memory systems are vast, and the detail to which events can be recalled is extensive. However, we certainly do not remember everything that happens in our lives, and items are not stored with equal detail. Reducing the volume of data to be stored through active data compression involves attenuating redundant information (efficient coding [48]), and prioritising attention and storage resources to information most critical for achieving goals.

- **Pattern Completion**. Reconstructing an event from brief exposure to part of that event, or from an impoverished, noisy or degraded version of the full scene e.g. when experiencing a familiar environment in the dark. This operation follows the idea of a *schema* [21][40]. Schemas can be thought of as generative models for a particular basic system [44]. Such schemas could be seeded with a small piece of information - e.g. you were at a children's party - and a rich scene can then be filled in - i.e. there was a cake, gifts, guests, and games were played.
- **Pattern Separation.** The process of transforming similar representations or memories into highly dissimilar, nonoverlapping representations [31]. This is important as we have many experiences that are similar to each other but nonetheless must be remembered as distinct. For example in discriminating edible from inedible plants. This function also relates to chunking of distinct event sequences from longer ongoing experiences, a core feature of episodic memory (see Figure 2). During later retrieval it is important only to recall relevant information for a particular task, thus separating relevant from irrelevant memories in a given context.
- **Unitary Coherent Perception.** Intuition attests that we experience a single *unitary* version of the world, not a probability distribution or blur over possible world states. Neither do we experience a world where independent percepts in a scene conflict with or contradict one another; our experience is *coherent*. These seemingly obvious features of experience, while having some recent experimental support [18], are in opposition to optimal decision making theories [4] and some forms of the Bayesian brain hypothesis (for a recent review see [39]). At the least, a maximum a posteriori approximation must be made at the 'percept' stage moment-to-moment even if the underlying computations are Bayesian.
  Helpfully, as a heuristic, unitary coherent (UC) perception may avoid the NP-hard computational complexity of full Bayesian inference of having to consider every possible interpretation of a current scene [18]. The existence of immediate perception as UC places strong constraints on models of AM: if AM is to store aspects of immediate percepts, then the recall and storage of memories would also be UC rather than fully Bayesian. This in turn will place constraints on theories of reconceptualisation of past episodes, which may otherwise require wider storage of probabilities.

3.1 The role of the hippocampal network in autobiographical memory

What is known about the neural circuits thought to underly AM in the brain? The extended hippocampal network, and its coordinated interaction with neocortex and subcortical structures, is thought to be the neurobiological substrate of the four AM functions outlined above. At a gross level, damage to the hippocampal-entorhinal cortex network famously causes anterograde amnesia [30]. At a cellular level the microstructure of the CA3 region suggests it likely serves as a 'convergence zone enabling information from different sensory modalities to be associated together (an idea originated by Marr [27]).

Hippocampal neural activity shows attractor dynamics [53] enabling content-addressable reactivation of the entire stored representation through activation of a particular memory by the dentate gyrus (DG) [28]. This role of older DG granule cells in associative pattern completion is complemented by the finding that neurogenesis supports pattern separation by young DG granule cells [31]. Critically, for pattern completion in AM, it has been shown that patients with hippocampal lesions cannot recall nor imagine spatially coherent events [22], and this reconstructive process can be modelled at the level of neurons [9]. Together these results highlight the importance of this structure in the combination and coordination of memory sequences, tying together activity across distributed circuits in the brain. We base our model of episodic memory on the hippocampal system.

These findings from the study of biological memory suggest that to create effective SAM we could adopt the following strategy. Compression and pattern completion can be accounted for by an implementation of the predictive brain hypothesis [11]. This allows both efficient and accurate recognition of events from partial information, as well as affording decision making, planning and reconstruction through inference in the model. High level unitary representations of a scene or event are then coordinated into coherent sequences by iterative pattern completion and separation operations in a model based on the hippocampus.

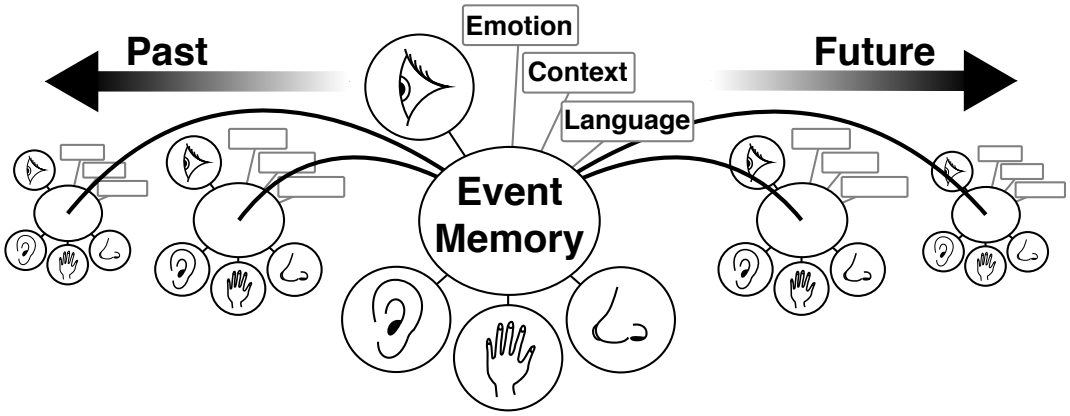

Fig. 2. Episodic memory in our model functions by combining disparate information across modalities into a single coherent percept from moment to moment. These transient percepts are monitored and coordinated over time into event sequences and episodic memories through iterative pattern completion and separation operations.

4 Capturing the characteristics of biological autobiographical memory in a single modelling framework

4.1 Compression and Pattern Completion: The Predictive Brain Hypothesis

In distributed episodic memory the individual sensor or motor systems encode their own memories in a domain-appropriate format [44]. Encoding of sensory input has been shown to be very efficient in biological systems, with neural populations being tuned to the scene statistics of the world [48]. A further efficiency is not to simply encode the world as it is, but to develop a generative model of the world and encode only events that could not be predicted.

So called Predictive Brain hypotheses [11][19]have become popular in recent years, claiming to provide a unified computational account of a range of cognitive capacities and seeking to explain an increasing number of neuroscience phenomena e.g. [20][41][50].

Reconstructing a past event can be understood as one of the functions of such a prediction engine, using past experience to anticipate and make sense of events as they happen. Seeding such a mechanism with appropriate clues will allow retrieval of a past episode. The same system, operating continuously, can also serve to fill-in and enrich the representation of the current situation, preparing the platform for more informed and appropriate action. It is well known that people can fill in sensory scenes with expected information, and the reconstruction of memories

from experienced or suggested fragments can lead to false memories even in individuals with 'highly superior' AM [33].

This powerful and efficient learning strategy would be highly appropriate for SAM and has not previously been applied to that domain.

4.2 Deep Learning: Towards a Synthetic Predictive Brain

A recent development in the field of machine learning is the arrival of practical deep learning systems [3] for abstracting information from data in an unsupervised way. Though neural networks have been investigated for decades [7, 45], recent developments of efficient training methods for deep (i.e. more than three layer) neural networks (DNNs) finally provide existence proofs of algorithms that can reconstruct (predict) complex sensory scenes in the manner required by predictive brain hypotheses.

Current excitement around DNNs can be ascribed to two prominent features. Firstly, they are now able, for the first time, to match or exceed human performance in certain benchmark pattern recognition tasks due to large-scale implementation on GPUs (see [3] for a review). Secondly, the kinds of invariant higher order representations developed by neural networks bare striking similarity to the tuning curves of neurons in higher order sensory cortex [55].

The unique processing architectures that result from training a DNN on data from a particular sensory domain could be described as a schema [21], an important concept in memory as discussed earlier. We propose that it is precisely the interaction of a compact episodic memory, and the sequential re-activation of an appropriate schema that underlies the compression and pattern completion capabilities of autobiographical memory. Therefore, it is to these DNN methods that we turn to compress sensory data streams and provide inputs to the episodic memory system.

4.3 Pattern Separation and Unitary Coherent Perception: Hippocampus as a Unitary Coherent Particle Filter

The vast majority of invasive neuroscience research uses rodent models. As a result, spatial reasoning or navigation tasks are often used to probe the biological foundations of memory. Many models of hippocampal function that are faithful to circuit-level details are therefore based on navigation or spatial reasoning [6]. Our model [16] has been developed as an effective algorithm for spatial inference (an idea recently expanded upon by Penny et al. [35]), while at the same time accounting for other known hippocampus-circuit phenomena such as the existence of sensory-pattern-specific and object-specific ('grandmother') cells. In robotics too, navigation has historically been a widely studied problem leading to widely-cited solutions to the Simultaneous Localisation And Mapping (SLAM). Our model may be viewed as a particular bio-inspired implementation of the SLAM algorithm, with EM algorithm steps mapping to neural activation and Hebbian learning respectively.

Our approach combines and extends research in both robotics and computational neuroscience fields: mapping well studied algorithms that have been shown to perform spatial reasoning and memory in navigation onto a neural circuit known to perform this function. Here we propose an update of this navigation model to autobiographical memory, echoing the well-articulated arguments of Buzsáki and Moser [8], based on the theory that the underlying computations for spatial and semantic memory are fundamentally the same.

The problem of locating oneself in an environment involves associating a given percept with a particular *place* (localisation) and encoding the transitions from one place to another by monitoring sensory differences, odometry and memories of previous traversals of the same environment. Autobiographical memory can be seen as essentially the same computation, but associating a

given percept to an *event*, while transitions from event to event can be handled by recursive pattern separation and completion operations. The analogy between space and AM is seen clearly in the well-known mnemonic technique of the 'memory palace' which explicitly uses spatial memory to organise sequences of object percepts [49].

It had been proposed before that the behaviour of animals during learning could be modelled as a particle filter [15], and elsewhere that sequential learning in the hippocampus could be modelled with a Temporal Restricted Boltzmann Machine (TRBM) [51], a machine learning algorithm. Fox and Prescott [16] took a TRBM model for navigation and extended it to include a mapping to hippocampal-entorhinal cortex circuitry that included: unitary coherence; sensor/odometry inputs from entorhinal cortex to DG; CA3 associative learning of input-place mappings; CA1 decoded localisation posteriors (place cells); and a subiculum - septum 'lostness detection' loop (see Figure 3). In a series of papers the model was extended to include online learning with a biomimetic sub-theta cycle after-depolarisation [17], and scaled up to complete a real-world navigation task by processing visual inputs through SURF feature extraction and K-Nearest Neighbour clustering [46].

We consider decoding localisation posteriors (DG-CA3-CA1 in the model) to be a pattern completion loop, whereas lostness detection and correction (Sub-Sep) to be pattern separation. By compressing an agent's sensory history through deep generative networks instead of SURF features, and by expanding the range of inputs in the model to the full gamut of sensory modalities, we aim to develop a more complete account of synthetic autobiographical memory.

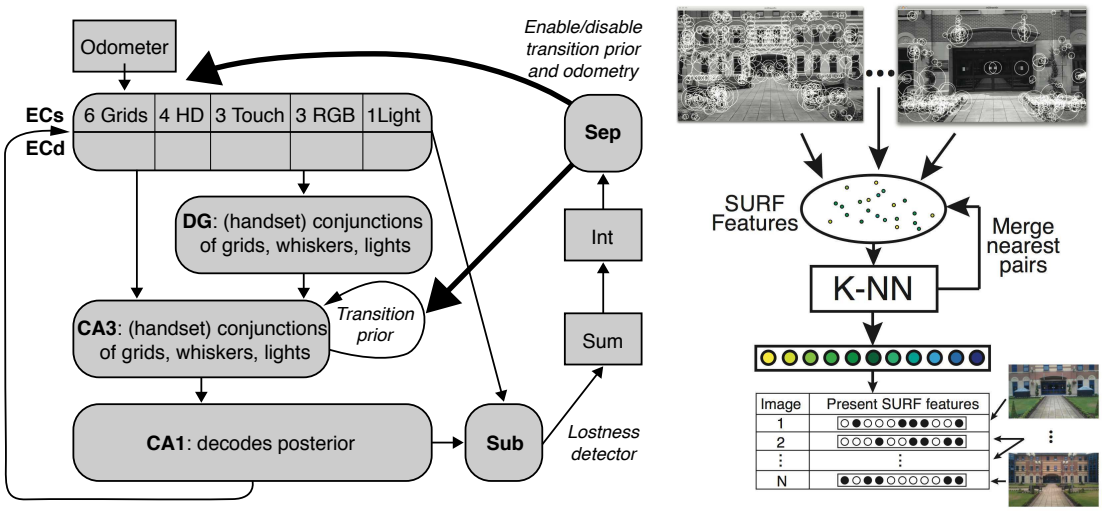

Fig. 3. (a) The Bayes filter hippocampus model, adapted from Fox and Prescott [16]. See main text for a synopsys of components. (b) A schematic of the SURF visual feature extraction step, constructing feature vector inputs to the Bayes filter hippocampus model, adapted from Saul et al. [46]. We aim to replace this feature extraction step with a range of modality-specific deep generative networks.

5 Conclusions and future directions

5.1 Advantages of our approach

The Deep Learning and UC hippocampus architecture presented here, based on basic systems and predictive coding theories of brain function, has a number of advantages over previous approaches. Using DNNs to learn predictive models of sensory scenes allows the agent to capture the full explanatory power of the data. Current approaches constrain how memories are stored at design time, limiting the practical utility of that information for new tasks in the future.

Our approach leverages state of the art machine learning approaches, bringing benchmark performance in speech and image processing to new modalities and new tasks. Code for many of these methods have been optimised for execution on GPUs, allowing the exploitation of cheap high-performance computing for real-time operation on robots. In addition, the UC hippocampus model has already been shown to work on real robot platforms with real-world sensory data, which is encouraging going forward.

5.2 Implementation objectives: social interaction with an iCub robot

An ideal testbed of our approach to SAM would be to embed the model on a mobile or humanoid robot to compress, parse and store rich streams of incoming data to improve ongoing processing performance. Our particular interest is in developing robots for social interaction, both for the inherent technical challenges and strong history of SAM in this task. Building on benchmarks established with existing SAM approaches [38][37], we hope to develop a system that can recognise people, remember interactions, and adapt interactions to specific people. Initial results of this work will be presented at the conference. Extensions to other problems in autonomous robots, and task domains such as life-long learning (as in [34]), will follow in due course.

Potential difficulties of implementing our approach are many and include: securing sufficient computational resources to run the algorithms (especially the DNNs) in real time; storage considerations during development of the data compression algorithms; interfacing the output of the SAM system with motor planning and narrative reasoning, which would ultimately involve encoding the memories themselves in terms of motor-control consequences; interfacing the UC episodic memory hippocampus model with rich predictive DNNs is non trivial; and event separation which would ultimately require a hierarchical implementation - a large problem in itself but one that is an interesting current area of research e.g. [47].

An additional active area of research is to advance machine learning algorithms to better capture the rich statistical structure of the world. Standard Deep Learning approaches are based on Restricted Boltzmann machines and their variations [3]. The latent representational structure in these models are relatively simple. The relationships between representations and latent variables within the learnt hierarchy, relationships that exist in the world and the data, are not captured by these models. Deep Gaussian processes [14], a recently proposed extension of Gaussian Process Latent Variable models [23], potentially provide a method for capturing rich hierarchical statistical relationships between latent variables. Integrating a Deep Gaussian Process with our UC hippocampus could substantially increase the power of this model.

6 Acknowledgements

The authors would like to thank the EU taxpayer for funding this research (EU grant no. 612139 WYSIWYD - "What You Say Is What You Did"), and our colleagues at the Sheffield Centre for Robotics for interesting discussions that stimulated this work.

Bibliography

[1] Anderson, J.R.: Act: A simple theory of complex cognition. American Psychologist 51(4), 355 (1996)

[2] Baddeley, A.: Essentials of Human Memory (Classic Edition). Psychology Press (2013)

[3] Bengio, Y.: Learning deep architectures for ai. Foundations and Trends in Machine Learning 2(1), 1–127 (2009)

[4] Bernardo, J.M., Smith, A.F.: Bayesian theory, vol. 405. John Wiley & Sons (2009)

[5] Berntsen, D., Rubin, D.C.: Understanding autobiographical memory: Theories and approaches. Cambridge University Press (2012)

[6] Bird, C.M., Burgess, N.: The hippocampus and memory: insights from spatial processing. Nature Reviews Neuroscience 9(3), 182–194 (2008)

[7] Bryson, A.E., Denham, W.F., Dreyfus, S.E.: Optimal programming problems with inequality constraints. AIAA journal 1(11), 2544–2550 (1963)

[8] Buzsáki, G., Moser, E.I.: Memory, navigation and theta rhythm in the hippocampal-entorhinal system. Nature neuroscience 16(2), 130–138 (2013)

[9] Byrne, P., Becker, S., Burgess, N.: Remembering the past and imagining the future: a neural model of spatial memory and imagery. Psychological review 114(2), 340 (2007)

[10] Carello, C., Turvey, M.T., Kugler, P.N., Shaw, R.E.: Inadequacies of the computer metaphor. Handbook of cognitive neuroscience pp. 229–248 (1984)

[11] Clark, A.: Whatever next? predictive brains, situated agents, and the future of cognitive science. Behavioral and Brain Sciences 36(03), 181–204 (2013)

[12] Cohen, M.S., Kosslyn, S.M., Breiter, H.C., DiGirolamo, G.J., Thompson, W.L., Anderson, A., Bookheimer, S., Rosen, B.R., Belliveau, J.: Changes in cortical activity during mental rotation a mapping study using functional mri. Brain 119(1), 89–100 (1996)

[13] Conway, M.A.: Autobiographical memory: An introduction. Open University Press (1990)

[14] Damianou, A.C., Lawrence, N.D.: Deep gaussian processes. arXiv preprint arXiv:1211.0358 (2012)

[15] Daw, N., Courville, A.: The pigeon as particle filter. Advances in Neural Information Processing Systems 20, 369–376 (2008)

[16] Fox, C., Prescott, T.: Hippocampus as unitary coherent particle filter. In: Neural Networks (IJCNN), The 2010 International Joint Conference on. pp. 1–8. IEEE (2010)

[17] Fox, C., Prescott, T.: Learning in a unitary coherent hippocampus. In: Artificial Neural Networks–ICANN 2010, pp. 388–394. Springer (2010)

[18] Fox, C., Stafford, T.: Maximum utility unitary coherent perception vs. the bayesian brain. In: Proceedings of the 34th annual conference of the Cognitive Science Society (2012)

[19] Friston, K.: The free-energy principle: a unified brain theory? Nature Reviews Neuroscience 11(2), 127–138 (2010)

[20] Friston, K., Kiebel, S.: Predictive coding under the free-energy principle. Philosophical Transactions of the Royal Society B: Biological Sciences 364(1521), 1211–1221 (2009)

[21] Ghosh, V.E., Gilboa, A.: What is a memory schema? a historical perspective on current neuroscience literature. Neuropsychologia 53, 104–114 (2014)

[22] Hassabis, D., Maguire, E.A.: The construction system of the brain. Philosophical Transactions of the Royal Society B: Biological Sciences 364(1521), 1263–1271 (2009)

[23] Lawrence, N.: Probabilistic non-linear principal component analysis with gaussian process latent variable models. The Journal of Machine Learning Research 6, 1783–1816 (2005)

[24] Lengyel, M., Dayan, P.: Hippocampal contributions to control: The third way. In: Advances in neural information processing systems. pp. 889–896 (2007)

[25] Lepora, N., Fox, C., Evans, M., Diamond, M., Gurney, K., Prescott, T.: Optimal decision-making in mammals: insights from a robot study of rodent texture discrimination. Journal of The Royal Society Interface 9(72), 1517–1528 (2012)
[26] Lisman, J., Grace, A.A., Duzel, E.: A neohebbian framework for episodic memory; role of dopamine-dependent late ltp. Trends in neurosciences 34(10), 536–547 (2011)
[27] Marr, D.: Simple memory: a theory for archicortex. Philosophical Transactions of the Royal Society B 262, 23–81 (1971)
[28] McClelland, J.L., McNaughton, B.L., O'Reilly, R.C.: Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychological review 102(3), 419 (1995)
[29] Miller, G.A.: The cognitive revolution: a historical perspective. Trends in cognitive sciences 7(3), 141–144 (2003)
[30] Milner, B., Corkin, S., Teuber, H.L.: Further analysis of the hippocampal amnesic syndrome: 14-year follow-up study of hm. Neuropsychologia 6(3), 215–234 (1968)
[31] Nakashiba, T., Cushman, J.D., Pelkey, K.A., Renaudineau, S., Buhl, D.L., McHugh, T.J., Barrera, V.R., Chittajallu, R., Iwamoto, K.S., McBain, C.J., et al.: Young dentate granule cells mediate pattern separation, whereas old granule cells facilitate pattern completion. Cell 149(1), 188–201 (2012)
[32] Neisser, U., Fivush, R.: The remembering self: Construction and accuracy in the self-narrative. No. 6, Cambridge University Press (1994)
[33] Patihis, L., Frenda, S.J., LePort, A.K., Petersen, N., Nichols, R.M., Stark, C.E., McGaugh, J.L., Loftus, E.F.: False memories in highly superior autobiographical memory individuals. Proceedings of the National Academy of Sciences 110(52), 20947–20952 (2013)
[34] Paul, R., Rus, D., Newman, P.: How was your day? online visual workspace summaries using incremental clustering in topic space. In: Robotics and Automation (ICRA), 2012 IEEE International Conference on. pp. 4058–4065. IEEE (2012)
[35] Penny, W.D., Zeidman, P., Burgess, N.: Forward and backward inference in spatial cognition. PLoS computational biology 9(12), e1003383 (2013)
[36] Petit, M., Lallée, S., Boucher, J.D., Pointeau, G., Cheminade, P., Ognibene, D., Chinellato, E., Pattacini, U., Gori, I., Martinez-Hernandez, U., et al.: The coordinating role of language in real-time multimodal learning of cooperative tasks. Autonomous Mental Development, IEEE Transactions on 5(1), 3–17 (2013)
[37] Pointeau, G., Petit, M., , Dominey, P.F.: Successive developmental levels of autobiographical memory for learning through social interaction. IEEE Transactions on Autonomous Mental Development (Forthcoming)
[38] Pointeau, G., Petit, M., Dominey, P.F.: Embodied simulation based on autobiographical memory. In: Biomimetic and Biohybrid Systems, pp. 240–250. Springer (2013)
[39] Pouget, A., Beck, J.M., Ma, W.J., Latham, P.E.: Probabilistic brains: knowns and unknowns. Nature neuroscience 16(9), 1170–1178 (2013)
[40] Preston, A.R., Eichenbaum, H.: Interplay of hippocampus and prefrontal cortex in memory. Current Biology 23(17), R764–R773 (2013)
[41] Rao, R.P., Ballard, D.H.: Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nature neuroscience 2(1), 79–87 (1999)
[42] Rennó-Costa, C., Lisman, J.E., Verschure, P.F.: The mechanism of rate remapping in the dentate gyrus. Neuron 68(6), 1051–1058 (2010)
[43] Rochat, P.: Criteria for an ecological self. The self in infancy: Theory and research 112, 17 (1995)
[44] Rubin, D.C.: The basic-systems model of episodic memory. Perspectives on Psychological Science 1(4), 277–311 (2006)

[45] Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by back-propagating errors. MIT Press, Cambridge, MA, USA (1988)
[46] Saul, A., Prescott, T., Fox, C.: Scaling up a boltzmann machine model of hippocampus with visual features for mobile robots. In: Robotics and Biomimetics (ROBIO), 2011 IEEE International Conference on. pp. 835–840. IEEE (2011)
[47] Schapiro, A.C., Rogers, T.T., Cordova, N.I., Turk-Browne, N.B., Botvinick, M.M.: Neural representations of events arise from temporal community structure. Nature neuroscience (2013)
[48] Simoncelli, E.P.: Vision and the statistics of the visual environment. Current opinion in Neurobiology 13(2), 144–149 (2003)
[49] Spence, J.D.: The memory palace of Matteo Ricci. Penguin Books Harmondsworth (1985)
[50] Srinivasan, M.V., Laughlin, S.B., Dubs, A.: Predictive coding: a fresh view of inhibition in the retina. Proceedings of the Royal Society of London. Series B. Biological Sciences 216(1205), 427–459 (1982)
[51] Taylor, G., Hinton, G., Roweis, S.: Modeling human motion using binary latent variables. In: Schölkopf, B., Platt, J., Hoffman, T. (eds.) Advances in Neural Information Processing Systems. vol. 19 (2007)
[52] Tulving, E.: Elements of episodic memory. Oxford Psychology Series (1985)
[53] Wills, T.J., Lever, C., Cacucci, F., Burgess, N., O'Keefe, J.: Attractor dynamics in the hippocampal representation of the local environment. Science 308(5723), 873–876 (2005)
[54] Wood, R., Baxter, P., Belpaeme, T.: A review of long-term memory in natural and synthetic systems. Adaptive Behavior 20(2), 81–103 (2012)
[55] Yamins, D.L., Hong, H., Cadieu, C., DiCarlo, J.J.: Hierarchical modular optimization of convolutional networks achieves representations similar to macaque it and human ventral stream. In: Advances in Neural Information Processing Systems. pp. 3093–3101 (2013)