

Cognitive Homeostatic Agents

Blue Sky Ideas Track

Amol Kelkar

i3ai.org

amol@i3ai.org

ABSTRACT

Human brain has been used as an inspiration for building autonomous agents, but it is not obvious what level of computational description of the brain one should use. This has led to overly opinionated symbolic approaches and overly unstructured connectionist approaches. We propose that using homeostasis as the computational description provides a good compromise. Similar to how physiological homeostasis is the regulation of certain homeostatic variables, cognition can be interpreted as the regulation of certain 'cognitive homeostatic variables'. We present an outline of a Cognitive Homeostatic Agent, built as a hierarchy of physiological and cognitive homeostatic subsystems and describe structures and processes to guide future exploration. We expect this to be a fruitful line of investigation towards building sophisticated artificial agents that can act flexibly in complex environments, and produce behaviors indicating planning, thinking and feelings.

KEYWORDS

artificial agent; homeostasis; cognitive homeostasis; hierarchical reinforcement learning; predictive coding; free energy principle; cognitive homeostatic agent; robotic control; lifelong learning; online learning; emotions; rewards; predictive processing; hippocampus; engrams

ACM Reference Format:

Amol Kelkar. 2021. Cognitive Homeostatic Agents: Blue Sky Ideas Track. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021, IFAAMAS*, 5 pages.

1 INTRODUCTION

The principle goal of artificial agents is intelligent behavior[46]. Early research towards building Artificial Intelligent (AI) systems viewed human cognition as analogous to symbolic computation in digital computers[23]. Later, connectionist approaches used artificial neural networks and attempted to implement cognitive features by mapping them to functions, taking advantage of the universal approximation theorem[17]. Recently, deep neural networks, reinforcement learning, meta-learning and related research areas have achieved great success in many areas of AI, such as playing console games[39], robotic object manipulation[42], natural language understanding and generation[6], and so on. Despite these advances, many challenges remain[19] for building robust artificial agents that can act in real world environments, interacting with humans and other agents. First, such agents must adapt continuously to changes in their environment, which requires online and lifelong learning.

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online. © 2021 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

That remains difficult[10] for current deep learning systems for reasons such as catastrophic forgetting[38]. Second, current artificial agents are driven by arbitrarily defined external rewards, or rewards derived from unsatisfactory signals such as novelty. Third, affects are important signals[48] that may help the agents deal with complex interactions with their environment, other agents and with humans. This area remains under-explored. Lastly, mortality and the life regulation to postpone death are some of the fundamental driving forces behind human creativity, industriousness and morality. It is unclear whether artificial agents would achieve those attributes without the meta objective of self-preservation [36].

The computational description of the brain that symbolic and connectionist approaches use are over and under constrained, respectively. To investigate what an appropriate level of computational description would be, we go back to the first principles and explore how intelligence evolved and operates in living organisms.

2 HOMEOSTASIS AND COGNITION

Adaptive behavior is a hallmark of all living organisms. Single-celled organisms such as bacteria can move towards food and away from noxious chemicals[2], and routinely make complex decisions as part of bacterial colonies[45]. Plants share resources[51] with each other through networks of fungi. Social insects lead complex, socially co-dependent lives[43]. Birds can count[20] and dogs can understand many verbal commands[15]. Octopuses can solve puzzles and remember large underwater terrains using landmarks[24]. It is evident that there is a continuum of cognitive capabilities across living organisms. Humans have exceptionally sophisticated versions of some of these capabilities, but the underlying physiological processes seem to be mostly unremarkable compared to other evolutionarily proximate animal species. Building intelligent agents based on our shared evolutionary inheritance is an under-explored approach, which we focus on in this article.

Adaptive behavior of living organisms is grounded on an internal reward system [3, 18] that is tuned to their homeostatic imperatives [41]. Deviation (or the prediction of deviation) of a homeostatic variable has a positive or negative valence[28] based on whether the deviation is towards or away from a homeostatic target state, i.e. whether the deviation is beneficial or harmful for the organism. For example, on a cold winter night, core body temperature starts to fall, moving away from nominal body temperature, which feels bad. Wearing a sweater feels good because it helps move body temperature closer towards the target value. These valences lead to affects and feelings[16], and drive behavior. Can homeostasis then be a good candidate computational description for building artificial agents that can act in complex environments in a grounded manner?

2.1 (Re)defining Homeostasis

Although the term homeostasis is often used to represent a narrow set of reactionary behaviors that move an organism towards a fixed set point, Cannon[7] originally intended the term to represent myriad physiological strategies to sense, maintain and defend regulated variables, as "wisdom of the body". This broader usage includes allostasis[52], the learned anticipatory responses and non-stationary target states, and other related mechanisms such as predictive homeostasis [40]. We use the term in its broadest scope, as explored under the free energy principle[21], where homeostasis includes all processes that keeps an organism within a relatively small set of favourable states that are conducive to its ongoing existence[14]. Under this definition, homeostasis not only regulates near term physiological states of the organism and survival, but also regulates well-being and long term flourishing.

2.2 Hierarchical Homeostatic Systems

Each organism is built as a hierarchy of homeostatic systems[12, 21, 34] such as membranes, sub-cellular organelles, cells, organs and organ systems. Such hierarchies of homeostatic systems remain stable only under specific conditions. At every level in the hierarchy, a controller subsystem must act to ensure their own homeostasis, the homeostasis of all subordinate subsystems and the homeostasis of the combined system. To achieve this, controllers need to have their own homeostatic variables ("hvars"), and hvars to represent the homeostatic state of subordinate subsystems, conveyed by bottom-up, interoceptive signals from those subsystems. Top-down signals allow a controller to influence homeostasis of subordinate subsystems. These top-down signals convey the expectations of future interoceptive signals[44], which imply changes to steady state, i.e. target values, of certain hvars for the subsystems. The changed hvar targets prompt the subsystems to initiate appropriate actions or generate expectations for their own subsystems. This activity cascade[22] results in the engagement of the entire hierarchy towards meeting the top-down expectations. For example, when bacteria or viruses invade the body and cause tissue injury, one of the immune system's responses is to produce pyrogens[35]. These chemicals are carried by the blood to the brain, where they affect the functioning of the hypothalamus by effectively increasing the target value for the core body temperature hvar. In response, the hypothalamus raises the body's temperature above the normal range, thereby causing a fever.

Signalling among subsystems is critical for a hierarchical homeostatic system to operate. Many organisms and bodily subsystems use chemical messengers[53]. It is difficult to convey different types of messages over long distances quickly and at precise targets using chemical signalling. On the other hand, neurons carry signals relatively rapidly and precisely and networks of neurons are capable of performing computations. This allows a nervous system based homeostatic controller to receive homeostatic state updates from a large number of subsystems and send back control signals to many subsystems, resulting in a fine grained mechanism to maintain homeostasis. This provides a clue why, in animals species, the nervous system took up the role as the highest level homeostatic controller.

2.3 Brain as a Homeostatic Controller

Over evolutionary time scales, nervous systems have proven to be highly advantageous. Neural networks can represent homeostatic state information and other useful sensory information, and complex computations can be performed across signals from different modalities to produce highly adaptive control signals. With a nervous system, organisms can optimize actions over longer time horizons, which sometimes requires ignoring or acting against immediate homeostatic drives.

This ability of a nervous system based controller to potentially drive subsystems against their homeostatic imperatives may prove dangerous for the organism. First, subsystems may be driven to states where they are no longer able to maintain their homeostasis, leading to organ failure. Second, the nervous system operates on partial information and may make decisions that are not optimal. Third, the nervous system's control algorithms, either genetically specified or learned, may be insufficient to handle the current situation. For a stable homeostatic system hierarchy, it is imperative that controllers must treat bottom-up signals conveying subsystem homeostatic state, such as a painful sensation, as valenced, i.e. conveying a "good" or "bad" feeling. The controllers must consider these signals as "privileged", compared to other signals that may carry miscellaneous information. Controllers must pay a high metabolic cost if they choose to ignore these privileged signals. For example, ignoring hunger while voluntary fasting requires constant upkeep of the resolve to fast. Note that although privileged signals are present at every level in the hierarchy, they may become "felt" only at certain systems in the brain.

3 COGNITIVE HOMEOSTASIS

Cognitive science typically makes a sharp distinction between mental activity like planning and activity for physiological homeostasis. We propose that what appears as resisting homeostatic drives during planning is in reality following homeostasis of other, higher order variables.

The brain is organized roughly as a hierarchy of subsystems[26]. One question relevant for this discussion is whether the brain represents a single homeostatic system or a hierarchy of homeostatic subsystems. Note that a single homeostatic system can be organized internally as a hierarchy of subsystems, but in that case, the subsystems do not represent homeostatic variables of their own and thus act merely as functional modules. Looking at what hvars the brain represents will help answer this question.

3.1 Cognitive Homeostatic Variables

As a homeostatic regulator, the brain must represent the physiological hvars such as the core body temperature, sodium levels, blood pressure, heart rate and so on. Basic emotions[8] such as fear, anger, joy, sadness, surprise and disgust can be interpreted as a set of hvars because the brain attempts to regulate them. For example, fear is the expectation of future harm and can be interpreted as an hvar whose steady state is the expectation of ongoing safety. Interpreting emotions as hvars is appropriate because emotions carry valence and by definition, valence is available only when a quantity affects homeostasis of the organism. More complex attributes such as the desire for well-being is also a homeostatic variable, because it has a

valence specific to an individual. Long term goals are hvars because achieving those goals has a positive valence. Habits are hvars. Even short term goals such as raising a hand can be cast as a temporary hvar. Successful completion of every thought has a positive valence, while jumbled, incoherent thoughts that do not lead to a resolution have a negative valence. As a result, each thought can also be interpreted as a process that leads to the target state of a temporary hvar. Indeed, all facets of cognition that have valence can be cast as homeostasis via corresponding hvars. We call these presumed hvars Cognitive Homeostatic Variables ("cognitive hvars").

Cognitive hvars may be fleeting like those for thoughts, short term like those for intentional actions, long term like a new year's resolution, or lifelong like career goals or social standing. They may be genetically specified, such as parenting instincts, or learned, such as habits.

3.2 Why Cognitive Homeostasis?

Why re-frame entire cognition as homeostasis? First, doing so gives a normative view of the behavior of all organisms, from single celled organism to humans, and potentially all self-organized systems such as societies and bacterial colonies. Second, it allows for the prioritization across physiological drives and cognitive goals. For example, conflicting goals such as energy homeostasis (hunger) and the desire to fast (not eat) can influence behavior from different levels in the homeostatic hierarchy, and actions result based on the relative influence of the two hvars and their interplay with other hvars in the system. Third, evolution often results in the reuse of components and patterns. In this case, it is conceivable that the neural patterns useful for physiological homeostasis were repeated for cognitive homeostasis as well, which presents an approach to investigate the evolution of nervous systems. Forth, this unification opens up the possibility of new ways of explaining free-will, self-hood and consciousness.

3.3 Modelling Cognitive HVars

Can we find a common computational description of how hvars are implemented in biological systems, so it can be applied in the construction of artificial agents? In nature, the mechanisms used for homeostasis vary substantially. For example, in the mitochondrial membrane, proton gradient across the membrane is maintained and is thus a homeostatic variable. When proton gradient shifts away from desired levels, molecular machines called proton pumps are triggered to restore the gradient. Locusts possess genetically specified visual neurons (the descending contralateral movement detectors, DCMDs) that detect predatory birds' looming motion[47] and trigger escape jumps. In this case, the hvar is implemented as specific afferent connectivity pattern of the DCMDs. Core body temperature hvar of endothermic animals is implemented as a system of sensors, neural and chemical signalling and several populations of neurons.

All these mechanisms can be summarized computationally as a set of desirable trajectories described using partial state[49] of the system. A trajectory is a sequence of changes described using certain partial state of the system over τ time steps, like a short animation clip. An hvar comprises the set of such trajectories that

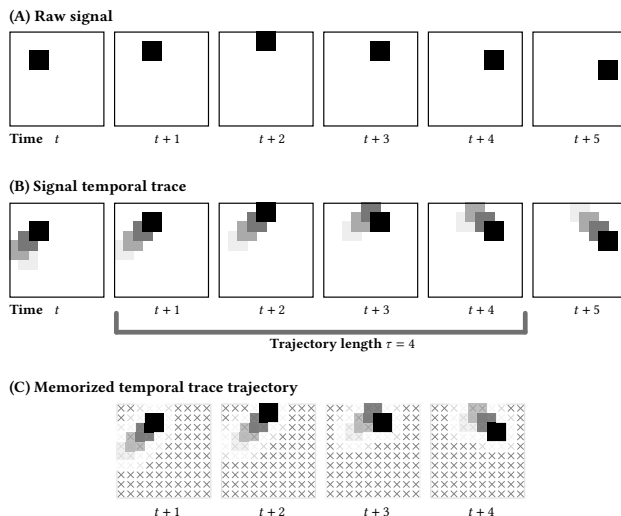


Figure 1: Temporal Trace Trajectory. (A) shows a signal generated from a signal source such as a sensory modality. (B) shows temporal trace, i.e. motion blurred, version of the signal. Temporal Trace Trajectory is a τ time step sequence of the signal's temporal trace. When storing in memory, trajectories can be memorized using a partial state space, indicated by the opaque areas in (C).

use the same set of partial state features. For example, thermoregulation hvar is a set of trajectories where starting state indicates a system that is either too cold or too warm, and ending state is where the system is closer to the optimal temperature. Other state features of the system, say blood glucose level, are not relevant to describe these trajectories, thus "partial state" above.

In other words, an hvar is a human-understandable interpretation of a set of trajectories. In a cognitive homeostatic system, we may not have interpretation for all trajectories. For this reason, when building such a system, we work directly with trajectories and not hvars.

4 COGNITIVE HOMEOSTATIC AGENTS

We propose building Cognitive Homeostatic Agents, artificial agents that are based on the Cognitive Homeostasis principles. The system is organized as a hierarchy-like graph of homeostatic subsystems, with each of those organized as a hierarchy-like graph of either homeostatic or non-homeostatic subsystems. All subsystems are connected to orchestrate a form of predictive processing[33].

Each subsystem contains a memory bank that stores a set of events in the form of temporal trace trajectories (see figure 1). A set of trajectories may represent a homeostatic variable. For instance, the set of trajectories that show core body temperature going from too warm or too cold to nominal temperature represent the core body temperature hvar. Temporal traces are used as a means of encoding time. They are computed as a decaying running sum of signals with a decay factor δ . Higher levels in the hierarchy generally use temporal traces with smaller δ resulting in longer

motion trails and trajectories of longer time duration τ compared to lower levels.

Discussion of specific organization of individual subsystems and modules is beyond of scope of this article.

4.1 System Operation

The system is bootstrapped by adding a set of trajectories to represent known hvars at various levels in the system hierarchy. Over time, other subsystems learn by remembering trajectories that result from the interaction between sensory signals and the bootstrapped trajectories. Once such trajectories are memorized, they may represent new hvars in corresponding subsystems. Thus, the entire system gets wired with higher level hvars that are grounded to the bootstrap hvars.

The system operation depends on predictive processing[33], evidence accumulation[50] and top-down signal generation[1]. Further investigation is needed to identify how a system based on a hierarchy of autoassociative memory modules may produce these characteristics. Here we present a sketch of a plausible, although incomplete, mechanism. In a given homeostatic subsystem, at a given moment in time, select a trajectory matching current conditions. Replay this winning trajectory through top-down signalling over next T time steps. If multiple trajectories are matched with no clear winner, then replay top few matching trajectories over the next few cycles of T time steps each. Repeat till a clear winner emerges by adjusting trajectory strength temporarily through a short term plasticity mechanism.

Confident recall of a higher level trajectory represents an instantiated goal. When such a trajectory is replayed to downstream subsystems, it encourages the whole hierarchy to act according to that goal. Such long term goal trajectories represent cognitive hvars and thus, planned behavior can be interpreted as reflexive behavior corresponding to these higher level cognitive hvars, blurring the boundaries between adaptive, reflexive behavior and intentional, planned behavior.

Cycling through potential future trajectories is seen in the theta band activation of rat hippocampus[32]. This approach also aligns with evidence accumulation[50] and action planning[37] through simulated playback from memory. What modules in the system participate in this mechanism remains an open question.

Unlike reinforcement learning, we propose that using an explicit reward signal is not necessary. The more homeostatically helpful a particular trajectory, the stronger it is imprinted in the memory. High utility trajectories exert a stronger influence on downstream subsystem behavior when replayed. Note that cognitive homeostasis allows hierarchical reinforcement learning[4] to emerge within certain subsystems. If a subsystem learns a cognitive hvar that represents an RL-like reward, trajectories for such hvars would be action policies that maximize corresponding rewards.

4.2 Implementation

Here we do not propose a specific implementation strategy. Instead, we lay out important characteristics that an implementation must have. The system 1) should preferentially remember favorable trajectories, 2) is able to autoassociatively[25] recall and replay matching trajectories based on the current state of the system, 3) is

wired to perform predictive processing[33], 4) has an attention-like mechanisms to identify, instantiate and eventually discard goal trajectories and corresponding cognitive homeostatic variables and 5) has the ability to learn using sequential, non I.I.D., data continuously without catastrophic forgetting, such as using meta-learning approaches[27] to continual learning.

5 CONCLUDING REMARKS

In this article, we described Cognitive Homeostasis, where cognition is interpreted as a homeostatic process, operating over Cognitive Homeostatic Variables. We then described Cognitive Homeostatic Agents that are built as homeostatic hierarchies based on the Cognitive Homeostasis principles.

Much work is needed to explore various aspects of this approach. Is temporal trace trajectory the appropriate unit of memory? What computational framework is suitable for implementing autoassociative trajectory memory? Does the playback of future trajectories use a synchronized clock across various modules? How to integrate bottom-up sensory evidence, the winning trajectories being played back and top-down expectations? How are transient cognitive hvars instantiated? How would planning, curiosity, creativity, emotions and other higher cognitive features emerge? Is this approach biologically plausible, and if so, what are the predictions of this approach that can be investigated through cognitive neuroscience and psychology? Can Cognitive Homeostasis provide normative mechanisms for episodic memory[54] and episodic-like memory[11]? What does this approach imply for free-will, selfhood and consciousness? Also, a mathematical description based on the Free Energy Principle[31] formulation would provide formal grounding for cognitive homeostasis. Finally, implementing cognitive homeostatic agents would provide an empirical validation of this approach.

We are cautiously optimistic that this approach would lead to a fruitful line of investigation towards building sophisticated artificial agents. Some of the reasons for our optimism are - 1) Temporal trace trajectories as the unit of memory aligns with engram as a neural memory unit[29]. 2) Mechanism for sampling alternate futures[9] as a means of planning may enable planning in physical and conceptual[13] spaces. 3) All reward-like signals in the proposed system are internal and grounded on the initial set of homeostatic variables, which obviates the need to define explicit rewards or arbitrary mechanisms to derive rewards. 4) We hypothesize that symbols and symbolic manipulation may emerge through the same mechanisms that instantiate and manage cognitive hvars. 5) our approach produces a sparse factor graph[5] over events that may result in causal reasoning capabilities. 6) Each homeostatic subsystem behaves in "System 1" [30] mode when prediction errors are small. Otherwise it behaves in "System 2" [30] mode, where alternate futures need to be sampled over an extended time period in an attempt to minimize prediction errors. 7) Using homeostasis provides a natural and grounded approach to feelings and emotions.

6 ACKNOWLEDGMENTS

We would like to thank the careful and constructive anonymous reviewers, as well as Lance Hughes for comments and constructive criticism. This work is supported by i3ai.org.

REFERENCES

- [1] Rick A. Adams, Stewart Shipp, and Karl J. Friston. 2013. Predictions not commands: Active inference in the motor system. *Brain Structure and Function* 218, 3 (2013), 611–643. <https://doi.org/10.1007/s00429-012-0475-5>
- [2] Uri Alon, Michael G Surette, Naama Barkai, and Stanislas Leibler. 1999. Robustness in bacterial chemotaxis. *Nature* 397, 6715 (1999), 168–171.
- [3] Andrew G Barto. 2013. Intrinsic motivation and reinforcement learning. In *Intrinsically motivated learning in natural and artificial systems*. Springer, 17–47.
- [4] Andrew G Barto and Sridhar Mahadevan. 2003. Recent advances in hierarchical reinforcement learning. *Discrete event dynamic systems* 13, 1 (2003), 41–77.
- [5] Yoshua Bengio. 2017. The consciousness prior. *arXiv preprint arXiv:1709.08568* (2017).
- [6] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. [arXiv:2005.14165](https://arxiv.org/abs/2005.14165) [cs.CL]
- [7] Walter B. Cannon. 1929. Organization for Physiological Homeostasis. *Physiological Reviews* 9, 3 (1929), 399–431.
- [8] Alessia Celeghin, Matteo Diano, Arianna Bagnis, Marco Viola, and Marco Tamietto. 2017. Basic Emotions in Human Neuroscience: Neuroimaging and Beyond. *Frontiers in Psychology* 8 (2017), 1432. <https://doi.org/10.3389/fpsyg.2017.01432>
- [9] Angus Chadwick, Mark CW van Rossum, and Matthew F Nolan. 2016. Flexible theta sequence compression mediated via phase precessing interneurons. *Elife* 5 (2016), e20349.
- [10] Zhiyuan Chen and Bing Liu. 2018. Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 12, 3 (2018), 1–207.
- [11] Nicola S Clayton and Anthony Dickinson. 1998. Episodic-like memory during cache recovery by scrub jays. *Nature* 395, 6699 (1998), 272–274.
- [12] Patrick Connolly and Vasi van Deventer. 2017. Hierarchical recursive organization and the free energy principle: From biological self-organization to the psychoanalytic mind. *Frontiers in Psychology* 8, SEP (2017). <https://doi.org/10.3389/fpsyg.2017.01695>
- [13] Alexandra O Constantinescu, J. X. O'Reilly, and T. E. J. Behrens. 2016. Organizing conceptual knowledge in humans with a gridlike code. *Science* 352, 6292 (jun 2016), 1464–1468. <https://doi.org/10.1126/science.aaf0941>
- [14] Andrew Corcoran and Jakob Hohwy. 2017. Allostasis, interoception, and the free energy principle: Feeling our way forward. (2017), 1–16. <https://doi.org/10.31234/osf.io/zbqnx>
- [15] Stanley Coren. 1994. *The intelligence of dogs: Canine consciousness and capabilities*. Free Press New York.
- [16] a. D. Craig and a. D. Craig. 2002. How do you feel? Interoception: the sense of the physiological condition of the body. *Nature Reviews. Neuroscience* 3, 8 (2002), 655–666. <https://doi.org/10.1038/nrn894>
- [17] Balázs Csánád Csáji et al. 2001. Approximation with artificial neural networks. *Faculty of Sciences, Eötvös Loránd University, Hungary* 24, 48 (2001), 7.
- [18] Edward L Deci and Richard M Ryan. 2010. Intrinsic motivation. *The corsini encyclopedia of psychology* (2010), 1–2.
- [19] Gabriel Dulac-Arnold, Daniel Mankowitz, and Todd Hester. 2019. Challenges of real-world reinforcement learning. *arXiv preprint arXiv:1904.12901* (2019).
- [20] Nathan J Emery. 2006. Cognitive ornithology: the evolution of avian intelligence. *Philosophical Transactions of the Royal Society B: Biological Sciences* 361, 1465 (2006), 23–43.
- [21] Karl Friston. 2013. Life as we know it. *Journal of the Royal Society Interface* 10, 86 (2013). <https://doi.org/10.1098/rsif.2013.0475>
- [22] Karl J Friston. 2019. Waves of prediction. *PLoS biology* 17, 10 (2019), e3000426.
- [23] James W. Garson. 1991. *What Connectionists Cannot Do: The Threat to Classical AI*. Springer Netherlands, Dordrecht, 113–142. https://doi.org/10.1007/978-94-011-3524-5_6
- [24] Peter Godfrey-Smith. 2016. *Other minds: The octopus, the sea, and the deep origins of consciousness*. Farrar, Straus and Giroux.
- [25] Vladimir I Gritsenko, Dmitri A Rachkovskij, Alexander A Frolov, Ross Gayler, Denis Kleyko, and Evgeny Osipov. 2017. Neural distributed autoassociative memories: A survey. *arXiv preprint arXiv:1709.00848* (2017).
- [26] Claus C. Hilgetag and Alexandros Goulas. 2020. ‘Hierarchy’ in the organization of brain networks. *Philosophical Transactions of the Royal Society B: Biological Sciences* 375, 1796 (2020). <https://doi.org/10.1098/rstb.2019.0319>
- [27] Khurram Javed and Martha White. 2019. Meta-learning representations for continual learning. *arXiv preprint arXiv:1905.12588* (2019).
- [28] Mateus Joffily and Giorgio Coricelli. 2013. Emotional Valence and the Free-Energy Principle. *PLoS Computational Biology* 9, 6 (2013). <https://doi.org/10.1371/journal.pcbi.1003094>
- [29] Sheena A. Josselyn and Susumu Tonegawa. 2020. Memory engrams: Recalling the past and imagining the future. *Science* 367, 6473 (2020). <https://doi.org/10.1126/science.aaw4325> [arXiv:https://science.sciencemag.org/content/367/6473/eaaw4325.full.pdf](https://arxiv.org/abs/https://science.sciencemag.org/content/367/6473/eaaw4325.full.pdf)
- [30] Daniel Kahneman. 2011. *Thinking, fast and slow*. Macmillan.
- [31] Friston Karl. 2012. A Free Energy Principle for Biological Systems. (2012), 2100–2121. <https://doi.org/10.3390/e14112100>
- [32] Kenneth Kay, Jason E. Chung, Marielena Sosa, Jonathan S. Schor, Mattias P. Karlsson, Margaret C. Larkin, Daniel F. Liu, and Loren M. Frank. 2020. Constant Sub-second Cycling between Representations of Possible Futures in the Hippocampus. *Cell* 180, 3 (2020), 552–567.e25. <https://doi.org/10.1016/j.cell.2020.01.014>
- [33] Georg B. Keller and Thomas D. Mrsic-Flogel. 2018. Predictive Processing: A Canonical Cortical Computation. *Neuron* 100, 2 (2018), 424–435. <https://doi.org/10.1016/j.neuron.2018.10.003>
- [34] Michael Kirchhoff, Thomas Parr, Ensor Palacios, Karl Friston, and Julian Kiverstein. 2018. The markov blankets of life: Autonomy, active inference and the free energy principle. *Journal of the Royal Society Interface* 15, 138 (2018). <https://doi.org/10.1098/rsif.2017.0792>
- [35] Matthew J. Kluger. 1991. Fever: Role of Pyrogens and Cryogens. *Physiological Reviews* 71, 1 (1991), 93–127. <https://doi.org/10.1152/physrev.1991.71.1.93> [arXiv:https://doi.org/10.1152/physrev.1991.71.1.93](https://arxiv.org/abs/https://doi.org/10.1152/physrev.1991.71.1.93) PMID: 1986393.
- [36] Kingson Man and Antonio Damasio. 2019. Homeostasis and soft robotics in the design of feeling machines. *Nature Machine Intelligence* 1, 10 (2019), 446–452. <https://doi.org/10.1038/s42256-019-0103-7>
- [37] Torsten Martiny-Huenger, Sarah E. Martiny, Elizabeth J. Parks-Stamm, Elisa Pfeiffer, and Peter M. Gollwitzer. 2017. From conscious thought to automatic action: A simulation account of action planning. *Journal of Experimental Psychology: General* 146, 10 (2017), 1513–1525. <https://doi.org/10.1037/xge0000344>
- [38] Michael McCloskey and Neal J Cohen. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*. Vol. 24. Elsevier, 109–165.
- [39] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [40] M. C. MC Moore-Ede. 1986. Physiology of the circadian timing system: Predictive versus reactive homeostasis. *The American Journal of Physiology* 250, 5 (may 1986), 737–752. <https://doi.org/10.1152/ajpregu.1986.250.5.R737>
- [41] Tobias Morville, Karl Friston, Denis Burdakov, Hartwig R. Siebner, and Oliver J. Hulme. 2018. The Homeostatic Logic of Reward. *bioRxiv* (2018), 242974. <https://doi.org/10.1101/242974>
- [42] OpenAI, Ilge Akkaya, Marcin Andrychowicz, Maciej Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, Jonas Schneider, Nikolas Tezak, Jerry Tworek, Peter Welinder, Lilian Weng, Qiming Yuan, Wojciech Zaremba, and Lei Zhang. 2019. Solving Rubik’s Cube with a Robot Hand. [arXiv:1910.07113](https://arxiv.org/abs/1910.07113) [cs.LG]
- [43] George F Oster and Edward O Wilson. 1978. *Caste and ecology in the social insects*. Princeton University Press.
- [44] Martin P Paulus and Murray B Stein. 2006. An insular view of anxiety. *Biological psychiatry* 60, 4 (2006), 383–387.
- [45] Adin Ross-Gillespie and Rolf Kümmerli. 2014. Collective decision-making in microbes. *Frontiers in microbiology* 5 (2014), 54.
- [46] J Russell Stuart and Peter Norvig. 2009. *Artificial intelligence: a modern approach*. Prentice Hall.
- [47] Roger D. Santer, F. Claire Rind, and Peter J. Simmons. 2012. Predator versus Prey: Locust Looming-Detector Neuron and Behavioural Responses to Stimuli Representing Attacking Bird Predators. *PLoS ONE* 7, 11 (2012). <https://doi.org/10.1371/journal.pone.0050146>
- [48] Matthias Scheutz. 2012. The affect dilemma for artificial agents: should we develop affective artificial agents? *IEEE Transactions on Affective Computing* 3, 4 (2012), 424–433.
- [49] Emmanuel Senft, Séverin Lemaignan, Paul Baxter, and Tony Belpaeme. 2017. Toward supervised reinforcement learning with partial states for social HRI. (2017).
- [50] Michael N N. Shadlen and Daphna Shohamy. 2016. Decision Making and Sequential Sampling from Memory. *Neuron* 90, 5 (2016), 927–939. <https://doi.org/10.1016/j.neuron.2016.04.036>
- [51] Suzanne W Simard, David A Perry, Melanie D Jones, David D Myrold, Daniel M Durrall, and Randy Molina. 1997. Net transfer of carbon between ectomycorrhizal tree species in the field. *Nature* 388, 6642 (1997), 579–582.
- [52] Peter Sterling. 1988. Allostasis: a new paradigm to explain arousal pathology. *Handbook of life stress, cognition and health* (1988).
- [53] Ann Gill Taylor, Lisa E Goehler, Daniel I Galper, Kim E Innes, and Cheryl Bourguignon. 2010. Top-down and bottom-up mechanisms in mind-body medicine: development of an integrative framework for psychophysiological research. *Explore* 6, 1 (2010), 29–41.
- [54] Endel Tulving. 1983. Elements of episodic memory. (1983).