The ultimate trick? Comment on : The Markov blanket trick: On the scope of the free energy principle and active inference by Raja et al

Karl Friston

The Wellcome Centre for Human Neuroimaging, UCL Queen Square Institute of Neurology, London, UK WC1N 3AR. Email: k.friston@ucl.ac.uk

---

Keywords: *free energy principle; random dynamical systems; sparse coupling; Markov blankets*

It is a pleasure to comment upon Raja et al. [1]. Responding to philosophical deconstructions of the FEP has now become a familiar part of my monthly routine. I am starting to realise that philosophy is quintessentially adversarial, much like good practice in litigation. Philosophical deconstructions of the FEP are unique in this regard; for example, the FEP has been applied constructively in mathematics [2], quantum physics [3], ethology [4], psychiatry [5], robotics [6], governance [7], et cetera. However, in (argumentative) philosophical discourse, I have to remind myself that I am a "proponent of the FEP" and a "defender" of its claims and commitments.

My slight problem, in this instance, is that I tend to agree with many of the points made in [1]. This is exemplified by an excerpt from a recent introduction to the FEP [8]:

> "Before starting, it might help to clarify what the free energy principle is—and why it is useful. Many theories in the biological sciences are answers to the question: "what must things do, in order to exist?" The FEP turns this question on its head and asks: "if things exist, what must they do?" More formally, if we can define what it means to be something, can we identify the physics or dynamics that a thing must possess? To answer this question, the FEP calls on some mathematical truisms that follow from each other. Much like Hamilton's principle of least action[1], it is not a falsifiable theory about the way 'things' behave—it is a description of 'things' that are defined in a particular way."

Indeed, the core claim of Raja et al. [1] is that the FEP is a principle or method, which (heuristically) allows us to say that, if a system has a Markov blanket, then we can model its dynamics as a path of least surprisal or, equivalently, as minimizing a free energy functional of Bayesian beliefs. This seems uncontroversial to me.

Having said this, in the spirit of adversarial exchange—and to entertain readers—I will do my best to expose some issues with the arguments in Raja et al; in the style of responses to a curmudgeonly reviewer.

Response to Raja et al

"Direct experimental support for predictive processing in general, and for FEP and active inference in particular is, at best, scarce." P51

I did not know this. I thought the empirical evidence for predictive processing and active inference was overwhelming, ranging from hierarchical connectomes in the brain [9-11] through to the synaptopathies that underwrite psychopathology; e.g., [12, 13]. Figure 1 shows the number of publications (excluding reviews and conference proceedings) referring to active inference or predictive processing, with currently over 400 scientific articles per year, attracting over 13,000 citations per year. The majority are in the empirical sciences, with 5.6% in philosophy. Clearly, these metrics only provide evidence of evidence but speak to the fecundity of active inference, in terms of generating "experimental support".

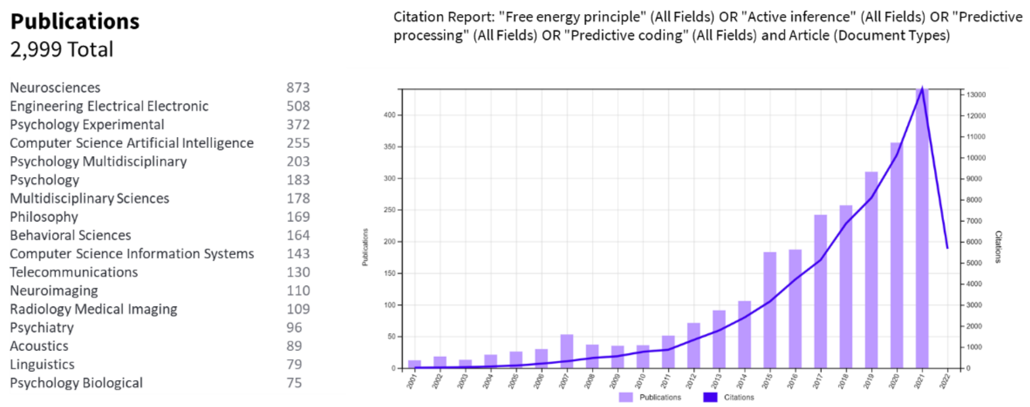

Figure 1: the left panel lists the number of scientific articles referring to the free energy principle, active inference or predictive processing, sorted according to Web of Science categories. The right panel shows the number of publications and citations from 2001 through to 2021. Source: Document search - Web of Science Core Collection.

"To date, the best existing empirical evidence for FEP amounts to no more than a set of simulations that bear, perhaps, metaphorical resemblance to actual biological and cognitive processes." P51

I found this statement upsetting and unsettling. I was upset because I was impressed by the fact that one could simulate things like insight [14], curiosity [15] and narrative exchange [16] from first principles. I now realise that not everyone is quite so impressed.

I was unsettled because demeaning simulations in this way takes the pressure off people to provide proof of principle that this is how the world (or brain) works. I have in mind here Feynman's notion: "What I cannot create, I do not understand". Dismissing simulations of cognitive processes licences a return to narrative formulations of 20th-century cognitive science. Informal narrative accounts represent the opposite direction of travel in cognitive neuroscience, as illustrated by the advent of computational phenotyping [17, 18]; namely, fitting simulations or 'digital twins' to an individual's behavioural or neurophysiological responses; e.g., [19].

"What follows can be read as a series of suggestions for issues that proponents of FEP and active inference will need to pay further attention to, especially if they wish to maintain the more ambitious claims they have made." P51

Thank you. I undertake to pay attention to these issues.

"Some things or systems cannot be described in terms of a Markov blanket. This is acknowledged in the FEP literature. The canonical example of a thing that cannot possess a Markov blanket is a candle flame [28,29]. … Therefore, it is impossible to apply the Markov blanket formalism and to say that candle flames are things although they are clearly things." P58

This made me smile: "although they [flames] are clearly things" would not really count as a killer argument in my world. We ask slightly more fundamental questions, such as "is a state a thing?" The answers usually rest upon the apparatus of the renormalisation group [20], in the following sense: things are defined in terms of particular partitions (i.e., groups) of states. And states are attributes of things at a finer scale. Here, particular partitions are Markov blanket partitions, obtained by the renormalisation group operator. In short, things are constituted by states and states inherit from things. Crucially, as we go from one scale to the next, things get bigger and slow down.

On this reading of Markov blankets, questions about candle flames could be about why we (or at least philosophers) perceive a flame as a thing—i.e. Markov blanket—when the constituent molecules—i.e., constituent Markov blankets—are transient and effervescent (e.g., the collapse of the Markov blanket of an oxygen molecule as it becomes part of a water molecule). The answer—currently entertained by theorists in the field—rests upon the scale-free nature of the FEP [3, 21, 22]. In brief, it is perfectly admissible for a Markov blanket at one scale to persist and thereby be perceived, inferred or measured, while its constituent states derive from Markov blankets at a finer (e.g., molecular) scale that come and go very quickly.

So, does a flame have a Markov blanket? I think the answer depends upon the scale at which the question is posed. From the perspective of the flame—or the philosopher observing flame—the blanket states could be the states of incandescent plasmas that influence (and are influenced by) states beyond the flame, including the philosopher's sensory states. Conversely, from the perspective of the molecular constituents of the plasma there is no flame, because the coupling of concern is at the fast scale of molecular fluctuations. In short, the Markov blanket is just a way of formalising the coupling among the states of things at a particular scale. The implicit separation of temporal scales—and how Markov blankets at one scale couple to blankets at another—is probably what we should be paying attention to.

"[T]here seems to be no reason to choose the beam and not the atoms in the surface of the pendulums as the blanket states of the system or vice versa, other than that the beam might be a better selection to model the coupling event. Namely, the selection of the Markov blanket partition in the case of coupled pendulums does not prescribe the way they count as things." P59

If one wants to understand or simulate how two things are coupled, then drawing a Markov blanket between two clocks is the way to go: e.g., [23]. If you wanted to understand and realise the inner workings of a clock, you can apply your Markov blanket to the things that constitute its inner machinations. The claim of the FEP is that, for every thing, at every scale, there is a Markov blanket. This is just a definition of what it is to be a thing. If there were alternative formal (i.e., mathematical) definitions of things, then one could argue about who has the best claim to 'thingness'.

"The application of the Markov blanket formalism under FEP entails an (at least) implicit commitment to the idea that the relevant aspects of the systems of interest can be modeled in terms of these networks—i.e., in terms of directed and acyclic interactions." P61

This is a foundational error in Raja et al. The application of the Markov blanket formalism entails a commitment to the opposite; namely, *cyclical* interactions. Indeed, it is precisely the circular causality inherent in the two-way traffic across the Markov blanket that lends the FEP its explanatory power. In other words, it takes us from the world of closed systems and equilibrium physics into the world of open systems and nonequilibria, necessary to describe the exchange of something with everything else.

It may help to understand that the FEP (and physics in general) is not concerned with the directed acyclic graphs (DAG) that underwrite causal inference [24]. The FEP concerns itself with the emergence of conditional independencies from sparsely coupled dynamics. Causality is baked into dynamics, in the sense that states *cause* the motion of other states. In other words, causality is inherent in universes that admit time in their metric geometry. The FEP simply shows that inference is an emergent property of causal structure. If one wanted to apply the FEP to make inferences about causality, one would not use a DAG, rather, some form of dynamic causal modelling (DCM) [25].

"The problem pointed out by [26] … cast doubt on the ability of Markov blankets to capture all possible conditional independences and, more importantly, likely reduces the general applicability of the formalism." P39

But the formalism is applied generally (see Figure 1). So, is there a problem? The reason these critiques—e.g., [26-28]—have found little purchase beyond philosophy (and Twitter) is that they deal with the wrong kind of systems (usually, linear edge cases) [26] or ask the wrong questions [27]. A question that could have been asked is: "what is the probability of Markov blankets emerging, as a system grows in size or scale?" An answer could be that any system, of sufficient size, will feature Markov blankets, because the probability of there not being a blanket tends to zero with the number of states [22].

Raja et al. grant that the FEP explains all things as defined FEP-theoretically; they just argue for things are not things in this sense. Conversely, I would submit that to be a thing is to have a Markov blanket. But nuance this claim by noting that not all things are created equal: for example, only a certain kind of thing can show sentient behaviour (e.g., things that possess active states).

"The property of 'heading' seems to be difficult to describe in terms of just internal states of the agent or just external states of the environment. …Thus, it is not clear that the partition entailed by the Markov blanket formalism is the best way to characterize this kind of situation." P61

The best way to characterise this kind of situation is to apply the method of the FEP to relational aspects of the sensed world. Perhaps the most obvious example is navigation and the emergence of 'place' and 'direction [path] cells' in realisations of active (planning as) inference [29]. I found this part of the argumentation in Raja et al. naïve: the FEP explains why "the internal states of the agent" can be read as describing "the external states of the environment".

“However, the example of the “climbability” of a step clearly shows that affordances themselves do not have to do with selecting one action or another but with the very possibility of action given a relational property in the organism-environment system. The re-definition of affordances in terms of action-selection preferences is a good demonstration of FEP’s inability to capture relational properties.” P62

The authors seem to be adopting a position of physical realism, which would please radical enactivists and, possibly, Gibson. However, remember the candle is in the philosopher’s head just in the same way as the consequences of her action. These consequences underwrite the epistemic and pragmatic affordances—celebrated in decompositions of expected free energy—which “capture relational properties” that matter for action [30]. Put simply, in active (planning as) inference, affordances have to be recognised or inferred before plans can be realised through action.

“The question, however, is whether these principled ways to derive formalisms of physics are something more than a mathematical exercise. It is well known that MaxEnt suffers from a lack of results and empirical predictions … Wolfram’s proposal likely shares this problem.” P64

I was curiously ambivalent about this assertion. On the one hand, it is pleasing to see that people appreciate that the FEP is an instance of Jaynes’ maximum entropy principle (specifically, the maximum entropy principle under constraints afforded by a Markov blanket) [2]. One might hope that Wolfram’s proposal can be brought into the same fold, possibly via variational message passing on normal style factor graphs [31].

On the other hand, I was mortified to realise that the fundaments of the natural sciences, such as the maximum entropy principle, are nothing more than a mathematical exercise. I presume that, in philosophy, gauge theoretic treatments and variational principles of stationary action are equally vacuous, and have no currency in terms of how we make sense of our world. Perhaps the authors meant that maximum entropy principle—from which the FEP inherits—is a method and therefore cannot be used to make claims or predictions? [which is difficult to reconcile with its use in climate and ecology modelling: e.g., [32]].

“If FEP holds, so the story goes, the Bayesian mechanics instantiated by biological or cognitive systems can be described in terms of active inference, which can be understood as a specific model within the reinforcement learning framework.” P66

I had always understood reinforcement learning (RL) as a specific model within Bayesian mechanics (sometimes referred to as Bayesian RL). Philosophically, active inference tries to distance itself from RL, which is regarded as a distant cousin of ill repute. In brief, active inference admits RL as a special case, when there is no reducible uncertainty in play. The problem with RL is that it does not offer a mechanics of (Bayesian) belief updating. This is because it optimises *functions of states* as opposed to *functionals of beliefs about states*. The requisite functionals are the variational and expected free energy, nicely summarised in [1].

I note that Raja et al. made no reference to reward in their derivations. Perhaps they have in mind a reward-free kind of reinforcement learning?

“The more recent literature tells a story that goes like this: if you have random dynamical systems with a NESS and a Markov blanket, then you have FEP, and then you have a principled justification for Bayesian inference.” P65

I am not sure the FEP ‘justifies’ Bayesian inference. All the FEP says is that you can realise or reproduce sentient behaviour using an enactive form of Bayesian inference. Alternatively, you can describe self-organisation in terms of maximising Bayesian model evidence; i.e., self-evidencing [33]—should you want to.

“However, the historical development of FEP suggests that the actual logical flow is: if you are able to model any system as if it were an autoencoder/a Helmholtz machine, you can describe any system as engaging in Bayesian inference; Markov blankets permit you to model almost anything as if it were an autoencoder/a Helmholtz machine; thus we can model anything as engaging in Bayesian inference; therefore FEP holds.” P65

I think this is correct. However, autoencoders and Helmholtz machines have nothing to say about action or active inference. In consequence, if you were observing something that never moved, you would never know it was “engaging in Bayesian inference”. This speaks to another issue that the “proponents of the FEP” could pay attention to, namely, the particular kinds of things permitted under the Markov blanket formalism and the different kinds of behaviour and sentience they exhibit.

I appreciate that Raja et al. are trying to stake out an adversarial position for themselves, by telling a story about the historical development of the FEP. However, I think there is a greater (non-adversarial) utility in considering the legacies on which the FEP builds. Foregrounding variational autoencoders is one story, but this story does not speak to enactive aspect of the FEP and ensuing (normative) accounts of sentient behaviour. A better story would appeal to the observation of Kalman—that optimal control has the same functional form as Bayesian filtering [34]—an observation that portended control and planning as inference [35-37].

However, planning as inference is not the end of the story. The (current) dénouement rests upon expected free energy [38] that inherits from a third foundational theme; namely, the principles that underwrite optional Bayesian design [39], that re-emerged in the form of active learning [40] and, subsequently, active inference [41]. This story rests upon the decomposition of expected free energy into *expected information gain* and *expected cost*—see Equation 20 in [1]—thereby reconciling optional Bayesian design and Bayesian decision theory, under a first principle account.

There are other foundations that could usefully be exposed; for example, the relationship between compression, algorithmic complexity and universal computation [42-44] on the one hand, and formulations of salience and curiosity on the other [45-48]. There are also legacy stories that could be told about the good regulator theorem [49], perceptual control theory [50], reinforcement learning [51], predictive coding [52, 53], autopoiesis [54] and so on. All these stories lead to the FEP.

“Additionally, although formally plausible and regularly used in data analysis, the practicality of empirical Bayes for guiding inference in active agents remains to be established.” P68

Empirical Bayes is thoroughly established in computational neuroscience, machine learning and neurorobotics, in the form of predictive processing; e.g., [37, 55, 56]. I am reading predictive processing here as an enactive version of perceptual inference. Perhaps it would help to note that deep inference and learning are exactly hierarchical inference and learning, which just is empirical Bayes (where the likelihood at one level of a hierarchical model plays the role of a prior on a lower level).

"If otherwise the NESS is assumed to be present but not known, how can we even start knowing anything about the generative model that is not fully arbitrary?" P69

This kind of question suggests that Raja et al. have a slightly magical conception of the FEP. The FEP provides a more deflationary account of things. It is just a method; heuristically, it allows us to say that should something have characteristic states of being, it is more likely to be found in those states. This is the ultimate deflationary trick of the FEP. If you accept this, then you can interpret the accompanying self-organisation and dynamics in terms of self-evidencing or active inference.

On this view, formulating the Jaynesian constraints above—that define the characteristics states of something—in terms of a generative model has the following advantage: you can build or realise things under specified constraints. Furthermore, by optimising the constraints one can explain the behaviour of other things—or people—in terms of their generative models, in the spirit of computational phenotyping [5, 19]. In short, the FEP supplies a methodology for realising self-organisation or self-evidencing under an "arbitrary" generative model.

"There are both theoretical and empirical questions that must be addressed before we can consider FEP and active inference as serious contenders in the field." P70

This leaves me feeling suitably admonished. It also makes me realise that I was rather hoping that the FEP would attain "serious contender" status, at some point.

Acknowledgements: KF is supported by funding for the Wellcome Centre for Human Neuroimaging (Ref: 205103/Z/16/Z) and a Canada-UK Artificial Intelligence Initiative (Ref: ES/T01279X/1). I would like to thank Maxwell Ramstead for help in formulating this response.

References

[1] V. Raja, D. Valluri, E. Baggs, A. Chemero, M.L. Anderson, The Markov blanket trick: On the scope of the free energy principle and active inference, Phys Life Rev, 39 (2021) 49-72.
[2] D.A.R. Sakthivadivel, A Constraint Geometry for Inference and Integration, 2022, pp. arXiv:2203.08119.
[3] C. Fields, K. Friston, J.F. Glazebrook, M. Levin, A free energy principle for generic quantum systems, 2021, pp. arXiv:2112.15242.
[4] M.J.D. Ramstead, P.B. Badcock, K.J. Friston, Answering Schrodinger's question: A free-energy formulation, Phys Life Rev, DOI 10.1016/j.plrev.2017.09.001(2017).
[5] R. Smith, S.S. Khalsa, M.P. Paulus, An Active Inference Approach to Dissecting Reasons for Nonadherence to Antidepressants, Biological Psychiatry-Cognitive Neuroscience and Neuroimaging, 6 (2021) 919-934.
[6] P. Lanillos, C. Meo, C. Pezzato, A.A. Meera, M. Baioumy, W. Ohata, A. Tschantz, B. Millidge, M. Wisse, C.L. Buckley, J. Tani, Active Inference in Robotics and Artificial Agents: Survey and Challenges, 2021, pp. arXiv:2112.01871.

[7] B. Khazri, Governing continuous transformation, 1 ed., Springer Cham 2022.
[8] K. Friston, L. Da Costa, N. Sajid, C. Heins, K. Ueltzhöffer, G.A. Pavliotis, T. Parr, The free energy principle made simpler but not too simple, 2022, pp. arXiv:2201.06387.
[9] K. Friston, Hierarchical models in the brain, PLoS Comput Biol., 4 (2008) e1000211.
[10] K. Friston, T. FitzGerald, F. Rigoli, P. Schwartenbeck, G. Pezzulo, Active Inference: A Process Theory, Neural Comput, 29 (2017) 1-49.
[11] K. Friston, T. Parr, B. de Vries, The graphical brain: Belief propagation and active inference, Network neuroscience (Cambridge, Mass.), 1 (2017) 381-414.
[12] S. Shipp, Neural Elements for Predictive Coding, Front Psychol, 7 (2016) 1792.
[13] A.R. Powers, C. Mathys, P.R. Corlett, Pavlovian conditioning–induced hallucinations result from overweighting of perceptual priors, Science, 357 (2017) 596-600.
[14] K.J. Friston, M. Lin, C.D. Frith, G. Pezzulo, J.A. Hobson, S. Ondobaka, Active Inference, Curiosity and Insight, Neural Comput, 29 (2017) 2633-2683.
[15] P. Schwartenbeck, J. Passecker, T.U. Hauser, T.H. FitzGerald, M. Kronbichler, K.J. Friston, Computational mechanisms of curiosity and goal-directed exploration, eLife, 8 (2019) e41703.
[16] K.J. Friston, T. Parr, Y. Yufik, N. Sajid, C.J. Price, E. Holmes, Generative models, linguistic communication and active inference, Neuroscience & Biobehavioral Reviews, 118 (2020) 42-64.
[17] P. Schwartenbeck, K. Friston, Computational Phenotyping in Psychiatry: A Worked Example, eNeuro, 3 (2016).
[18] T. Parr, G. Rees, K.J. Friston, Computational Neuropsychology and Bayesian Inference, Front Hum Neurosci, 12 (2018) 61.
[19] R. Smith, N. Kirlic, J.L. Stewart, J. Touthang, R. Kuplicki, T.J. McDermott, S. Taylor, S.S. Khalsa, M.P. Paulus, R.L. Aupperle, Long-term stability of computational parameters during approach-avoidance conflict in a transdiagnostic psychiatric patient sample, Scientific reports, 11 (2021).
[20] K. Friston, A free energy principle for a particular physics, eprint arXiv:1906.10184, 2019.
[21] C. Fields, J.F. Glazebrook, M. Levin, Minimal physicalism as a scale-free substrate for cognition and consciousness, Neuroscience of consciousness, 2021 (2021) niab013.
[22] D.A.R. Sakthivadivel, Weak Markov Blankets in High-Dimensional, Sparsely-Coupled Random Dynamical Systems, 2022, pp. arXiv:2207.07620.
[23] K. Friston, C. Frith, A Duet for one, Conscious Cogn, 36 (2015) 390-405.
[24] J. Pearl, Probabilistic Reasoning In Intelligent Systems: Networks of Plausible Inference, Morgan Kaufmann, San Fransisco, CA, USA, 1988.
[25] P.A. Valdes-Sosa, A. Roebroeck, J. Daunizeau, K. Friston, Effective connectivity: influence, causality and biophysical modeling, Neuroimage, 58 (2011) 339-361.
[26] M. Biehl, F.A. Pollock, R. Kanai, A technical critique of the free energy principle as presented in 'Life as we know it', arXiv e-prints, DOI (2020) arXiv:2001.06408.
[27] M. Aguilera, B. Millidge, A. Tschantz, C.L. Buckley, How particular is the physics of the free energy principle?, Physics of Life Reviews, 40 (2022) 24-50.
[28] J. Bruineberg, K. Dolega, J. Dewhurst, M. Baltieri, The Emperor's New Markov Blankets, Behavioral and Brain Sciences, DOI 10.1017/S0140525X21002351(2021) 1-63.
[29] R. Kaplan, K.J. Friston, Planning and navigation as active inference, Biological cybernetics, 112 (2018) 323-343.
[30] T. Parr, K.J. Friston, Working memory, attention, and salience in active inference, Scientific reports, 7 (2017) 14678.
[31] J. Dauwels, On Variational Message Passing on Factor Graphs, 2007 IEEE International Symposium on Information Theory, 2007, pp. 2546-2550.
[32] S.J. Phillips, R.P. Anderson, R.E. Schapire, Maximum entropy modeling of species geographic distributions, Ecological Modelling, 190 (2006) 231-259.
[33] J. Hohwy, The Self-Evidencing Brain, Noûs, 50 (2016) 259-285.
[34] E. Todorov, General duality between optimal control and estimation, IEEE Conference on Decision and Control, 2008.

[35] H. Attias, Planning by Probabilistic Inference, Proc. of the 9th Int. Workshop on Artificial Intelligence and Statistics, 2003.
[36] M. Botvinick, M. Toussaint, Planning as inference, Trends Cogn Sci., 16 (2012) 485-488.
[37] B. Millidge, Deep Active Inference as Variational Policy Gradients, arXiv e-prints, DOI (2019) arXiv:1907.03876.
[38] T. Parr, K.J. Friston, Generalised free energy and active inference, Biological cybernetics, 113 (2019) 495-513.
[39] D.V. Lindley, On a Measure of the Information Provided by an Experiment, Ann. Math. Statist., 27 (1956) 986-1005.
[40] D.J.C. MacKay, Information-Based Objective Functions for Active Data Selection, Neural Computation, 4 (1992) 590-604.
[41] K.J. Friston, J. Daunizeau, S.J. Kiebel, Active inference or reinforcement learning?, PLoS One., 4 (2009) e6421.
[42] C.S. Wallace, D.L. Dowe, Minimum Message Length and Kolmogorov Complexity, The Computer Journal, 42 (1999) 270-283.
[43] D.J. MacKay, Free-energy minimisation algorithm for decoding and cryptoanalysis, Electronics Letters, 31 (1995) 445-447.
[44] M. Hutter, Universal Artificial Intellegence : Sequential Decisions Based on Algorithmic Probability, Springer-Verlag Berlin and Heidelberg & Co. KG, Dordrecht, 2006.
[45] T. Parr, K.J. Friston, Attention or salience?, Current Opinion in Psychology, 29 (2019) 1-5.
[46] A. Barto, M. Mirolli, G. Baldassarre, Novelty or surprise?, Front Psychol, 4 (2013) 907.
[47] J. Schmidhuber, Curious model-building control systems, In Proc. International Joint Conference on Neural Networks, Singapore. IEEE, 2 (1991) 1458–1463.
[48] Y. Sun, F. Gomez, J. Schmidhuber, Planning to be surprised: optimal Bayesian exploration in dynamic environments, Proceedings of the 4th international conference on Artificial general intelligence, Springer-Verlag, Mountain View, CA, 2011, pp. 41-51.
[49] R.C. Conant, W.R. Ashby, Every Good Regulator of a system must be a model of that system, Int. J. Systems Sci., 1 (1970) 89-97.
[50] W. Mansell, Control of perception should be operationalized as a fundamental property of the nervous system, Top Cogn Sci, 3 (2011) 257-261.
[51] R.S. Sutton, A.G. Barto, Toward a modern theory of adaptive networks: expectation and prediction, Psychol Rev., 88 (1981) 135-170.
[52] P. Elias, Predictive coding–I, IRE Transactions on Information Theory, 1 (1955) 16–24.
[53] R.P. Rao, D.H. Ballard, Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects, Nat Neurosci., 2 (1999) 79-87.
[54] H.R. Maturana, F. Varela, Autopoiesis: the organization of the living, in: V.F. Maturana HR (Ed.) Autopoiesis and Cognition, Reidel, Dordrecht, Netherlands, 1980.
[55] T. Matsumoto, J. Tani, Goal-Directed Planning for Habituated Agents by Active Inference Using a Variational Recurrent Neural Network, Entropy, 22 (2020) 564.
[56] J. Kiverstein, E. Rietveld, H.A. Slagter, D. Denys, Obsessive Compulsive Disorder: A Pathology of Self-Confidence?, Trends in Cognitive Sciences, 23 (2019) 369-372.