CIFMA 2022

Author: Maria Raffa

Ph.D. Student *Visual and Media Studies* – IULM University, Milano (Italy)

maria.raffa@studenti.iulm.it

Presentation for position paper: Markov blankets for sustainability.

Keywords: Markov blankets – free energy principle – active inference – AI – sustainability

Introduction

This contribution's aim is twofold: on one hand, to provide an overview of the state of the art of some kind of Bayesian networks, i.e. Markov blankets, focusing on their relationship with the mathematical frameworks of the free energy principle and active inference. On the other hand, to sketch how these concepts are practically applied to AI, with special regard to their possible use in the field of sustainable development. Indeed, this issue has been addressed from several points of view, but the literature still lacks an overall discourse including both the theoretical description of Markov blankets inside FEP and active inference, and their several different employments for AI, in order to tackle the question of sustainable development.

In order for these purposes, the contribution is structured as follows: the first part sketches briefly what a Markov blanket (MB) is, and how it is linked to the concepts of free energy principle (FEP) and active inference. The second part provides an overview of AI applications of MBs, joint with FEP and active inference, and lastly, in the third part, it is wondered if all this can be useful to address the issue of sustainability.

1. MB, FEP and active inference

The concept of MB was first introduced by Judea Pearl (1988) in the following terms: the Markov blanket of a variable X is the set consisting of the parents of X, the children of X, and the variables sharing a child with X (Koski, Noble, 2009, 50). Specifically, a MB is a kind of Bayesian network, i.e. a statistical multivariate model where its graphical structure allows us to represent and reason about an uncertain domain (Korb, Nicholson, 2004): each node is associated with one variable of the domain, and the direct links between the nodes represent informational or cause-effect relationships.

These dependencies are quantified by conditional probability distributions, meaning the extent to which one node is likely to be affected by the others. The nodes forming the MB of X simplify a lot the computation of the value of a variable: «so, for instance, if X is embedded in a graph, with a hundred variables but its MB consists only of five variables, one can safely ignore ninety-five variables in the computation» (Facchin, 2021, 8).

Therefore, the concept of MB is born in the context of statistics, and later it has been adopted in different fields. It has been used especially in the philosophy of mind and the study of cognition, tied to Karl Friston and colleagues' mathematical frameworks of free energy principle (FEP) and active inference for the study of life (Friston, 2010). FEP, shortly, claims that agents that exist do so because they can persist, maintaining their equilibrium through free energy minimization: an organism's existence defines a probability distribution over the space of all its possible states, and such a probability distribution has low information-theoretic entropy. Entropy is the long-term average of surprisal (i.e. the negative log evidence), so, minimizing surprisal over time will ensure an organism's prolonged existence (Friston, 2010). Free energy, though, can be minimized through active inference, which means self-generated changes in status and the suppression of prediction errors through the information derived from the history of previous interactions with the environment (Friston *et al.*, 2011; Kirchoff *et al.*, 2018). In this framework, MBs are introduced as a tool to describe a specific form of conditional independence between a dynamical system and its environment and are addressed as *real* boundaries of living systems (Friston, 2013; Kirchhoff *et al.*, 2018). Specifically, this means that although living systems need to interact with the environment, as they are open systems, they also need to distinguish themselves from the external environment, so their states must form a set of states that is distinct from the set of environmental states. In this context, conditional independence occurs: when the state of the organism/environment boundary, i.e. the MB, is fixed, what happens on one side of the boundary does not influence what happens on the other side. In this circumstance, all the necessary information for explaining the behavior of the internal system is given by the state of the blanket (Palacios *et al.*, 2020; Howhy, 2019).

All given that, it looks like in Friston's works MB has been affected by a sort of epistemic leap, going from being a statistical model to a real object. This has been criticized, among others, by Jelle Bruineberg and colleagues, as it forms «an example of reification fallacy: treating something abstract as something concrete» (Bruineberg *et al.*, 2021, 19). Bruineberg stresses the difference between the so-called Pearl blankets, i. e. the statistical model, and the Friston blankets, asserting that literature often is filled with confusion about these two souls of MB, that should be kept separate and employed properly depending on the circumstances.

2. MB, FEP, active inference and AI

Taking into account MBs' problematic ontological status, this second part of the contribution will try to provide an overview of MBs' applications to AI, and how they are linked to FEP and active inference.

Broadly, FEP and active inference have been used for several applications, to explain complex processes in different fields: in psychology, they have been used to ground a computational account of neuropsychological syndromes; in economics, FEP has been employed to reformulate some specific optimization processes in terms of the agents' belief; moreover, active inference has been also used to model smooth and saccadic eye movements, and to frame memory and attention. In biology, too, FEP is used to explain how all the organisms and processes that subserve perception and action emerge and are constantly adjusted through a natural model selection process (Mazzaglia *et al.*, 2022).

Moreover, concerning machine learning, much of the current applications are still based on feedforward architectures trained by backpropagation (LeCun *et al.*, 2015), but recently the so-called generative models are becoming more and more popular. The idea behind generative models is the same as predictive perception based on Bayesian inference, i.e minimization through errors. Indeed, these algorithms can generate sensory signals corresponding to predicted causes, as they learn from small quantities of data and generalize to new situations (Seth in Mendoça *et al.* eds., 2020). Generative models are used for several tasks, such as image generation, text prediction and video modeling, and they are especially useful to predict the dynamics of a system, such as an environment, so they have been studied for control, exploration and anomaly detection (Mazzaglia *et al.*, 2022). Specifically, generative models joint with active inference have implications in robotics. Indeed, recent studies have shown how from active inference can emerge active vision (Mirza *et al.*, 2016; Daucè, 2018), and this is useful to address the problem of finding and reaching a certain object in a robotic space (Lanillos *et al.*, 202): recently, has been implemented a generative model using deep neural networks that is able to fuse multiple views into an abstract representation, and is trained from data by minimizing variational free energy. Empirically, it has been proven that a robotic agent, which starts moving without any knowledge about the workspace, at each step chooses the next view pose by evaluating the expected free energy (Van de Maele *et al.*, 2021).

Lastly, strictly concerning MBs, it's worth mentioning their employment to implement models for music generation. A special case of MB, i.e. Markov chain, is a model describing a sequence of possible events, where the probability of the next state depends on a previous state and not on *all* previous states (Kirchhoff *et al.*, 2018). This kind of network has the right features for the task of music generation, as the idea is to take into account, for instance in a corpus of chords, which is the probability for a chord to follow another particular chord (Cruz, 2019).

3. Sustainable employment?

On these grounds, the third part of this contribution tries to consider whether MB, FEP and active inference can be useful to face the urgent issue of sustainability, as environmental exploitation and the lack of energy resources are becoming more and more dramatic nowadays society. This hypothesis comes from the trivial fascination that a sustainable organism – or a society – should be, as a matter of fact, independent, and manage to maintain self-equilibrium, maybe through the minimization of surprise in the prediction of its next states, such as FEP and active inference claim.

As previously stressed, Bayesian networks in a certain way allow knowing the effects given the causes and the causes given the effects, meaning that they are useful to compute the probability of a particular hypothesis, and, consequently, to support decision-making. Furthermore, Bayesian networks are useful for incorporating data from different sources and domains, and since nodes in Bayesian networks are modeled by means of probability distributions, uncertainty can be estimated more accurately than in models where commonly only mean values are taken into account. For these features, Bayesian networks have already been used to address environmental issues (Bromley, 2005) and water issues (Phan *et al.*, 2016). On this path, it is worth mentioning the work of David Requejo-Castro (2021), who has proposed a data-driven Bayesian network approach, taking into reference the Sustainable Development Goal 6 of the 2030 Agenda, i.e. ensuring access to water and sanitation for all. Requejo-Castro's work combines expert opinion and quantitative data to support informed decision-making. Specifically, he uses Bayesian networks-based structure learning algorithms to replicate composite indicators–based conceptual frameworks, which represent expert knowledge, and identifies interlinkages associated with a complex context, coupled with bootstrapping, to reduce results uncertainty, and with a comprehensive result robustness analysis. His results, validated on SDG 6, shows that this combined approach improves model inference capacity, identifies the interlinkages among the variables considered, and can be useful for analysis of the complexities also in different contexts (Requejo-Castro, 2021).

All the sources mentioned above show the reliability of Bayesian networks to support decision-making processes, and their efficacy if used in the context of the concrete issue of nowadays lack of environmental resources. Following up, the final proposal of this contribution is that understanding exactly to which extent could Bayesian networks, especially in the reduced form of

MBs, be framed in the context of FEP and active inference, could be useful to implement concrete tools to support decision-making processes for research addressed to sustainable development.

Conclusion

So far, this contribution has sketched the notion of MBs inside the contexts of FEP and active inference, and how these concepts have been used to implement AI models. Moreover, it has been shown how Bayesian networks have already been used for tasks related to sustainable development. All that considered, this contribution aims to be the very starting point for research that deep into the possible employment of MBs, accurately framed into the contexts of FEP and active inference, to assess the crucial problems of sustainability. Indeed, on one hand, it seems very useful to follow up on the blossoming and promising stream of research that aims to implement AI tools to support decision-making processes to accomplish the Agenda 2030 Sustainable Development Goals; on the other, the theoretical frameworks of FEP and active inference seems still to be clarified and explored, in order to understand which may be their concrete applications in the field of AI.

References

2015, *Transforming our World: The Sustainable Development Agenda to 2030*, United Nations General Assembly.

AA. VV., 2020, Mendoça *et al.* (eds.), *The philosophy and science of predictive processing*, Bloomsbury.

Bruineberg J. *et al.*, 2021, *The Emperor's New Markov Blankets*, «Behavioral and Brain Sciences», 1-63 [preprint], doi: 10.1017/S0140525X21002351.

Bromley J., 2005, *Guidelines for the use of Bayesian networks as a participatory tool for water resource*, Walliford, United Kingdom.

Cruz J., 2019, *Deep Learning vs Markov Model in Music Generation*, Honors College Theses [graduate thesis].

Daucé E., 2018, *Active fovea-based vision through computationally-effective model-based prediction*, «Frontiers in Neurorobotics», 12:76, doi: 10.3389/fnbot.2018.00076.

Facchin M., 2021, *Extended predictive minds: do Markov Blankets matter?*, «Review of Philosophy and Psychology», doi: 10.1007/S13164-021-00607-9.

Friston K., 2010, *The free-energy principle: a unified brain theory?*, «Nature Reviews. Neuroscience» 11 (2), pp. 127-38.

2013, *Life as we know it*, «Journal of the Royal Society Interface», 10, doi: 10.1098/rsif.2013.0475.

et al., 2011, *Action understanding and active inference*, «Biological Cybernetics», 104, 137-160.

Hohwy J., 2019, *Quick'n'Lean or Slow and Rich? Andy Clark on predictive processing and embodied cognition*, in M. Colombo, E. Irvine, M. Stapleton (Eds.) *Andy Clark and His Critics* (pp. 191-205), New York, Oxford University Press.

Kim J. *et al.*, 2018, *Sustainable technology analysis of artificial intelligence using Bayesian and social network models*, «Sustainability», 10 (1).

Kirchhoff M. *et al.*, 2018. *The Markov Blankets of Life: Autonomy, Active Inference and the Free Energy Principle*, «Journal of the Royal Society Interface», 15 (138).

Korb, K., Nicholson, A., 2004. *Bayesian Artificial Intelligence*, Chapman and Hall/CRC. Florida, United States.

Koski, T., Noble J. M., 2009, *Bayesian Networks: an Introduction*, Chichester, Wiley and Sons.

Lanillos P. *et al.*, 2021, *Active Inference in Robotics and Artificial Agents: Survey and Challenges*, arXiv 2021, arXiv:2112.01871.

LeCun Y. *et al.*, 2015, *Deep Learning*, «Nature», 52, 7553, pp. 436-444.

Mazzaglia P. *et al.*, 2022, *The Free Energy Principle for Perception and Action: A Deep Learning Perspective*, «Entropy», 24 (2).

Mirza M. B., 2016, *Scene construction, visual foraging, and active inference*, «Frontiers in Computer Neuroscience», 10:56, doi: 10.3389/fncom.2016.00056.

Palacios E. E., *et al.*, 2020, *On Markov Blanket and hierarchical self-organization*, «Journal of Theoretical Biology», 486:110089.

Parr T., Pezzulo G., Friston K., 2022, *Active Inference: The free energy principle in Mind, Brain and Behaviour*, MIT Press.

Pearl J.,1988, *Probabilistic reasoning in intelligent systems: Networks of plausible inference*, Morgan Kauffman ed., 77-110.

Phan Th.D. *et al.*, 2016. *Applications of Bayesian belief networks in water resource management: A systematic review*, «Environmental Modelling and Software», 85, 98-111. https://doi.org/10.1016/j.envsoft.2016.08.006.

Requejo Castro D., 2021, *Data Driven Bayesian Networks modelling to support decision-making: Application to the context of Sustainable Development Goal 6 on water and sanitation* [PhD thesis].

Van De Maele T. *et al.*, 2021, *Active Vision for Robot Manipulators Using the Free Energy Principle*, «Frontiers in Neurorobotics», 15, doi: 10.3389/fnbot.2021.642780.