

A minimum relative entropy principle for AGI

Antoine van de Ven* and Ben A.M. Schouten

Fontys University of Applied Sciences

Postbus 347, 5600 AH Eindhoven, The Netherlands

*Antoine.vandeVen@fontys.nl

Abstract

In this paper the principle of minimum relative entropy (PMRE) is proposed as a fundamental principle and idea that can be used in the field of AGI. It is shown to have a very strong mathematical foundation, that it is even more fundamental than Bayes rule or MaxEnt alone and that it can be related to neuroscience. Hierarchical structures, hierarchies in timescales and learning and generating sequences of sequences are some of the aspects that Friston (Fri09) described by using his free-energy principle. These are aspects of cognitive architectures that are in agreement with the foundations of hierarchical memory prediction frameworks (GH09). The PMRE is very similar and often equivalent to Friston's free-energy principle (Fri09), however for actions and the definitions of surprise there is a difference. It is proposed to use relative entropy as the standard definition of surprise. Experiments have shown that this is currently the best indicator of human surprise (IB09). The learning rate or interestingness can be defined as the rate of decrease of relative entropy, so curiosity can then be implemented as looking for situations with the highest learning rate.

Introduction

Just like physics wants to find the underlying laws of nature, it would be nice to find underlying principles for intelligence, inference, surprise and so on. A lot of progress has been made and many principles have been proposed. Depending on what principles or foundations are used, it is possible to come to different theories or implementations of intelligent agents. A good example is AIXI (Hut04) which combines Decision Theory with Solomonoff's universal induction (which combines principles from Ockham, Epicurus, Bayes and Turing). It uses compression and Kolmogorov complexity, but unfortunately this makes it uncomputable in this form. The ability to compress data well has been linked to intelligence and compression progress has been proposed as a simple algorithmic principle for discovery, curiosity and more. While this has a very strong and solid mathematical foundation, the problem is that it is often very hard or even impossible to compute. Often it is also assumed that the agent stores all data of all

sensory observations forever. It seems unlikely that the human brain works like that.

In (vdV09) the principle of minimum relative entropy (PMRE) was proposed to be used in developmental robotics. In this paper we want to propose it as a fundamental principle and idea for use in the field of AGI. We compare it with other principles, relate it to cognitive architectures and show that it can be used to model curiosity. It can be shown that it has a very solid and strong mathematical foundation because it can be derived from three simple axioms (Gif08). The most important assumption and axiom is the principle of minimal updating: beliefs should be updated only to the extent required by the new information. This is incorporated by a locality axiom. The other two axioms are only used to require coordinate invariance and consistency for independent subsystems. By eliminative induction this singles out the logarithmic relative entropy as the formula to minimize. This way the Kullback-Leibler divergence (KLD) (KL51) has been derived as the only correct and unique divergence to minimize. Other forms of divergences and relative entropies in the literature are excluded. It can be shown (Gif08) to be able to do everything orthodox Bayesian inference (which allows arbitrary priors) and MaxEnt (which allows arbitrary constraints) can do, but it can also process both forms simultaneously, which Bayes and MaxEnt cannot do alone. This has only been shown recently and is not well known yet. The current version of the most used textbook on Artificial Intelligence (RN02) doesn't even include the words relative entropy or Kullback-Leibler divergence yet.

Free-energy

While in our approach the PMRE with the KLD is the most fundamental, in other approaches exact Bayesian inference is often taken as most fundamental, and the KLD is then used to do approximate inference. The variational Bayes method is an example of this. It tries to find an approximate distribution of a true posterior distribution by minimizing the KLD between the approximate distribution and the true posterior distribution. Sometimes a free-energy formulation is used which yields the same solution when minimized, but which can

make the calculations easier. In fact the free-energy formulation is the same as the KLD with an extra term (Shannon surprise) that doesn't depend on the approximate distribution, so it doesn't influence the search for the best approximate distribution. In the field of neuroscience, Friston (Fri09) has proposed the minimum free-energy principle as a fundamental principle that could explain a lot about how the brain functions. For perception it is equal to minimizing the KLD, so it is equivalent to the PMRE in that respect. Friston showed that many properties and functions of the brain can be explained by using the free-energy principle, such as the hierarchical structure of the brain, a hierarchy of timescales in the brain and how it could learn and generate sequences of sequences. This is in agreement with the memory prediction framework (GH09). Note that this not only relates these principles to the brain, but that it can also guide the design and choice of cognitive architectures for artificial general intelligence.

Currently the brain is the only working proof that general intelligence is possible, so these principles and results could help and guide biologically inspired AGI. These results seem to confirm the foundations of biologically inspired frameworks which use hierarchical structures, spatio-temporal pattern recognition and the learning and generating of sequences of sequences.

Biologically plausible

The fact that the PMRE only does minimal updating of the beliefs makes it more biologically plausible than some other theories. For example AIXI (Hut04) isn't based on minimal updating, because it uses global compression including all historical data. Brains don't seem to work that way. When observing and learning there are physical changes in the brain to incorporate and encode the new information and new beliefs. Such physical changes are costly for an organism and should be avoided as much as possible, because of limited energy and limited resources. The PMRE avoids this by doing only minimal updating of the beliefs. It is related to compression because in this way it stores new information and beliefs in an efficient way.

A new definition of surprise

Besides the theoretical arguments we can also refer to experiments. Itti and Baldi (IB09) proposed a definition of Bayesian Surprise that is equal to the KLD between the prior and posterior beliefs of the observer. This again is the same formula as used by the PMRE. In experiments they showed that by calculating this they could predict with high precision where humans would look. This formula and definition was found to be more accurate than all other models they compared it with, like Shannon entropy, saliency and other measures. In their derivation Itti and Baldi picked the KLD as the best way to define Bayesian Surprise by referring to the work of Kullback. While we agree on this definition, it would also have been possible to pick another diver-

gence as a measure, because the KLD is just one out of a broader class of divergences called f-divergences. The benefit of the derivation of the PMRE is that it uniquely selects the KLD as the only consistent measure that can be used. So in this way the PMRE helps to select and confirm this definition of surprise.

Relative entropy and curiosity

Relative entropy can also be used to implement curiosity and exploration. In (SHS95) it was used for reinforcement driven information acquisition, but it can also be implemented in different ways. The rate in which the relative entropy decreases can be seen as the learning rate. Curiosity can then be implemented as looking for and exploring the situations with highest learning rate (interestingness). This can be compared with implementations of curiosity which use decrease of prediction errors or compression progress (Sch09).

References

- Karl Friston. The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*, 13(7):293–301, 2009.
- Dileep George and Jeff Hawkins. Towards a mathematical theory of cortical micro-circuits. *PLoS Comput Biol*, 5(10):e1000532, oct 2009.
- Adom Giffin. *Maximum Entropy: The Universal Method for Inference*. PhD thesis, Massey U., Albany, 2008.
- Marcus Hutter. *Universal Artificial Intelligence: Sequential Decisions Based On Algorithmic Probability*. Springer, 1 edition, November 2004.
- Laurent Itti and Pierre Baldi. Bayesian surprise attracts human attention. *Vision Research*, 49(10):1295–1306, June 2009.
- S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, March 1951.
- Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 2nd edition, December 2002.
- Jürgen Schmidhuber. Driven by compression progress: A simple principle explains essential aspects of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes. In *Anticipatory Behavior in Adaptive Learning Systems: From Psychological Theories to Artificial Cognitive Systems*, pages 48–76. Springer-Verlag, 2009.
- Jan Storck, Sepp Hochreiter, and Jürgen Schmidhuber. Reinforcement driven information acquisition in Non-Deterministic environments. In *Proc. International Conference on Artificial Neural Networks (ICANN95)*, 1995.
- Antoine van de Ven. A minimum relative entropy principle for the brain. In *Proceedings of the Ninth International Conference on Epigenetic Robotics*. Lund University Cognitive Studies, 145, 2009.