

A novel information theoretic measure of surprise

Mohammad Javad Faraji
EPFL, Lausanne, CH
mohammadjavad.faraji@epfl.ch,

Kerstin Preuschoff
Univ. of Geneva, Geneva, CH
kerstin.preuschoff@unige.ch,

Wulfram Gerstner
EPFL, Lausanne, CH
wulfram.gerstner@epfl.ch

Surprise is informative because it drives attention [IB09] and modifies learning [SD00]. Correlates of surprise have been observed at different stages of neural processing, and found to be relevant for learning and memory formation [RR03]. Although surprise is ubiquitous, there is neither a widely accepted theory that quantitatively links surprise to observed behavior, such as the startle response, nor an agreement on how surprise should influence learning speed or other parameters in iterative statistical learning algorithms.

Building on and going beyond earlier surprise measures [BI10, Sha01, Pal12], we propose a novel information theoretic measure for calculating surprise in a Bayesian framework so as to capture uncertainty of the world as well as imperfections of the subjective model of the world, two important aspects of surprise. The principle of future surprise minimization leads to a learning rule that can be interpreted as a surprise modulated belief update suitable for learning within changing environments.

Importantly, we do not need an assumption on how quickly the world changes. We apply our surprise-modulated learning rule to an exploration task in a maze-like environment. Our results are consistent with the behavioral finding that surprising events induce humans and animals to learn faster and to adapt more quickly to changing environments [PH80].

Information content [Sha01] captures the inherent unexpectedness of a piece of data for a given set of models (uncertainty of the world), whereas Bayesian surprise [BI10] measures the change in belief caused by a new data point (observer dependent). These are two complementary approaches for calculating surprise. The former is about data and the latter is about a model. In our approach both aspects are combined with a third aspect: if we are uncertain about what to expect, receiving a low-probability data sample is less surprising than in a situation when we are almost certain about the world. A surprise minimizing learning (SMiLe) rule is derived by solving a constrained optimization problem defined as follows: the objective is to maximally reduce surprise when facing the same data again in the not so far future, under the constraint that the posterior belief (after the update step) is not too different from the prior.

The resulting SMiLe rule balances the influence of newly acquired data with prior knowledge where the balance depends on surprise. In case of a fundamental change in the world signaled by surprising samples, data acquired before the change is downgraded as less informative about the current state of the world. A simultaneous increase of the influence of newly acquired data on learning leads to a fast adaptation of the model to an environmental change. While in a stationary environment our algorithm approaches the known Bayesian update rule, it also allows the model to react to changes in the environment.

In summary, surprising data increases the uncertainty we have about our current model of the world and gives a bigger influence of newly acquired data on belief update. The interaction between surprise and uncertainty is important for modeling the behavior of humans and animals in changing environments. The surprise signal could be broadly transmitted in the brain by a neuromodulator with widespread axonal ramifications (e.g., norepinephrine (NE) released from locus coeruleus (LC) neurons) and influence synaptic plasticity rules.

Acknowledgments: This project is funded by the European Research Council (grant agreement no. 268 689)

References

- [BI10] Pierre Baldi and Laurent Itti. Of bits and wows: a bayesian theory of surprise with applications to attention. *Neural Networks*, 23(5):649–666, 2010.
- [IB09] Laurent Itti and Pierre Baldi. Bayesian surprise attracts human attention. *Vision research*, 49(10):1295–1306, 2009.
- [Pal12] Gunther Palm. *Novelty, information and surprise*. Springer, 2012.
- [PH80] John M Pearce and Geoffrey Hall. A model for pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological review*, 87(6):532, 1980.
- [RR03] Charan Ranganath and Gregor Rainer. Neural mechanisms for detecting and remembering novel events. *Nature Reviews Neuroscience*, 4(3):193–202, 2003.
- [SD00] Wolfram Schultz and Anthony Dickinson. Neuronal coding of prediction errors. *Annual review of neuroscience*, 23(1):473–500, 2000.
- [Sha01] Claude Elwood Shannon. A mathematical theory of communication. *ACM SIGMOBILE Mobile Computing and Communications Review*, 5(1):3–55, 2001.