

Published in final edited form as:

Curr Opin Neurobiol. 2014 April ; 0: 130–133. doi:10.1016/j.conb.2014.01.003.

Bayesian statistics: Relevant for the brain?

Prof. Konrad Paul Kording

Rm 1483, 345 E Superior St, Chicago, IL 60611, United States, Phone: +1 312 2386966

Konrad Paul Kording: koerding@gmail.com

Abstract

Analyzing data from experiments involves variables that we neuroscientists are uncertain about. Efficiently calculating with such variables usually requires Bayesian statistics. As it is crucial when analyzing complex data, it seems natural that the brain would “use” such statistics to analyze data from the world. And indeed, recent studies in the areas of perception, action, and cognition suggest that Bayesian behavior is widespread, in many modalities and species. Consequently, many models have suggested that the brain is built on simple Bayesian principles. While the brain’s code is probably not actually simple, I believe that Bayesian principles will facilitate the construction of faithful models of the brain.

What is Bayesian statistics?

Bayesian statistics can be seen as a model of the way we understand things. Our sensors are noisy and ambiguous as several possible worlds (or models of how the brain works) could give rise to the same sensor readings. We therefore have uncertainty in our data and cannot be certain which model or hypothesis we should believe in. However, we can considerably reduce uncertainty about the world using previously acquired knowledge and by integrating data across sensors and time. As new data comes in, we update our hypotheses. Bayesian statistics is the rigorous way of calculating the probability of a given hypothesis in the presence of such kinds of uncertainty. Within Bayesian statistics, previously acquired knowledge is called prior, while newly acquired sensory information is called likelihood.

A simple example based on brain-machine interfaces highlights its use. Let’s say we have a monkey opening and closing its hand, while we record from its primary motor cortex [1]. We want to decode how the monkey is moving the hand, maybe to build a prosthetic device. Let us say the monkey wants to open the hand 80% of the time and close it otherwise (prior $p(\text{open})=0.8$). Let us say we record the number of spikes from a neuron related to hand opening which gives 10 ± 3 spikes (mean \pm sd) when the hand is open and 13 ± 3 spikes when the hand is closed. How could we estimate if the hand should be open based on both the spikes and the prior knowledge?

We can then use **Bayes rule**:

$$p(\text{open}|\text{spikes}) = \frac{p(\text{open})p(\text{spikes}|\text{open})}{p(\text{open})p(\text{spikes}|\text{open}) + p(\text{closed})p(\text{spikes}|\text{closed})}$$

© 2014 Elsevier Ltd. All rights reserved.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

If we got 19 spikes and used approximate Gaussian distributions we would have a probability of about 53% that the hand should be closed. This combination of prior and likelihood is a typical application of Bayes rule. All of Bayesian statistics is in some way built upon Bayes rule.

Analyzing neural data with Bayesian statistics gives better results with less data

When it comes to the analysis of the increasingly complex datasets in neuroscience, Bayesian statistics is used ubiquitously [2]. This should be no surprise, after all Bayesian statistics is just the calculus of variables about which we have uncertainty.

I want to go through one illustrative example. In many experiments, the experimentalist is showing visual stimuli while measuring spikes from a neuron. They need to know the receptive field of the neuron, the (to first approximation) linear transfer function from input to spikes. However, the input is high dimensional. For example, a spatial pattern may be described with the brightness values of 10*10 pixels. Estimating the underlying 100 free parameters takes an awful lot of trials, even when choosing the stimuli in an optimal way [3]. However, not all potential receptive fields are equally likely. In fact, from previous experiments we know that receptive fields tend to be small (sparse in space), smooth (sparse spatial derivatives), and localized in frequency space (sparse in frequency). Putting these ideas into a Bayesian prior, it is possible to obtain the same quality of receptive field mapping with a lot less data [4]. In a way, a theory borne by previous experiments can radically simplify subsequent experiments enabled by Bayesian ideas.

This example shows the nature of prior knowledge that is typically used. Sparseness in a few dimensions was used, and combined in a soft way with data. The “theory”, if we should call it that, proposed that localized receptive fields are more likely than non-localized. It did not posit that receptive fields should be maximally localized. Formulating theories in such a soft way allows them to be readily combined with data.

Similar applications can be found in many areas of neuroscience. In the field of brain machine interfaces, we have priors about how people want to interact with machines and need to combine this with data from the brain [5]. When analyzing high dimensional data, such as data from multiple simultaneously recorded neurons [4] or imaging [6], Bayesian approaches usually simplify the data analysis by allowing the use of prior knowledge about the relations of variables. When analyzing interactions between molecules there is the need to use prior knowledge about data and to combine data across experiments [7]. Regardless the subfield of neuroscience, when data analysis gets complicated, Bayesian statistics tend to be useful.

Why Bayesian data-analysis matters

Being able to combine data, using previous experiments to set priors for following experiments is the very essence of the scientific method. Being able to quantify such relations in a world characterized by progressive data integration and increasing use of databases is crucial for current and necessary for much of future neuroscience.

Assuming that the brain solves problems close to the Bayesian optimum often predicts behavior

If progressively more complex datasets force neuroscientists to use Bayesian statistics to make sense of their data – it should come as no surprise that animals would have to go down

a similar route. After all, the world is complex, perception is noisy and ambiguous and data is scarce. In fact, it is known that without prior information learning is entirely impossible, an idea that is known since the work of Hume [8] and that has been recently formalized in the so-called no free lunch theorems [9]. While there is no doubt that the brain needs to combine noisy information with relevant prior knowledge, how close the way it does this to the optimal Bayesian solution is an issue of intense debate.

Before moving to the discussion of concrete applications we need to clarify what we mean by Bayesian behavior. The sole essence of Bayesian models of behavior is that they predict that behavior is close to the best possible solution to a problem encountered by the animal. Nothing about likelihoods or priors needs to be explicitly represented. For example, chemical stability of synapses over multiple timescales[10], may implement the belief that the past is similar to the present and that the world evolves over multiple timescales [11]. As far as models of behavior are concerned, they do not make any statement about the nature of the implementation of the underlying statistical calculations[12].

Bayesian models are frequently used to model information integration. Sensory signals get combined with one another or with priors. They are integrated over time or space. Recent research using simple models for such integration has compared actual behavior with predictions of models of optimal integration. Behavior is often close to optimal for cue combination in motor settings, auditory, visual and sensorimotor behavior [13].

Bayesian models have also recently been used to model many high level cognitive phenomena. Some studies construct Bayesian models of how subjects estimate the values of continuous variables [e.g. 14]. Other studies construct Bayesian models of how subjects estimate the structure of the world [15–18]. A thriving community tries to relate actual behavior to predictions of optimal Bayesian behavior.

However, there is an active debate about the explanatory power of Bayesian models. A central criticism of Bayesian theories is that the degrees of freedom in the models allow a new kind of over-fitting, basically allowing the modeler to achieve any predictions they like [19]. The defense by these Bayesian modelers is that there are far fewer degrees of freedom because knowledge about the actual structure of the world makes the set of permissible Bayesian models rather small[20]. Constraining Bayesian models, e.g. by measuring the statistical properties of the real world, seems crucial to strengthen their predictive power.

Why Bayesian models of behavior matter

The question of how “Bayesian” or optimal the brain is, is important in and by itself. Beyond this, understanding the problem that the brain needs to solve is important. Any credible model of how the brain or a part of it works must be able to solve the problems that are defined by the environment. This is a surprisingly high bar. For example, while many Bayesian or optimal control algorithms are used to control robots [21], we find few neuromorphic implementations of such algorithms. Lastly, building Bayesian models of information integration leads to an understanding of the relevant variables and their interplay. This understanding is inspiring a good number of experiments in experimental neuroscience.

Asking how the nervous system could achieve Bayesian behavior leads to new insights about neural codes

Many theories are competing about the representation of uncertainty and computation with it and they generally derive from simple assumptions. For example, in linear Probabilistic Population Codes [PPC, 22] the dearth of spikes in a Poisson model represents uncertainty.

In sampling models uncertainty is represented in the sequence of activities [23]. In yet other work it's the firing rate of specialized neurons that represents uncertainty [24]. Many different codes of uncertainty can be unified through the use of Bayesian statistics and the idea of population codes [25]. Each of the proposed models of the representation of uncertainty is elegant by focusing on just one easily communicable code of uncertainty.

However, uncertainty is just one of the many pieces of knowledge about the world that the brain is representing. And in most cases, when probed carefully, the brain exhibits complicated codes [e.g. 26]. Existing studies that ask how uncertainty is represented all point into the direction of complicated codes [27–32]. Admittedly, there have been too few studies asking how uncertainty is represented so it is hard to be sure. Still, I am aware of no dataset that suggests that, across brain areas, there is a simple, conserved, representation of uncertainty.

Each of the proposed models predict some measured data. Some neurons lower firing rates when uncertainty is higher [22] as predicted by probabilistic population codes. Other neurons increase their firing rate [24]. The fact that activities during sleep resemble activities during natural scene viewing [23] is predicted by sampling models. Moreover, there is a broad range of fMRI findings about the role of uncertainty, that do not clearly support any one model [but see 27–28]. And in some cases where neural responses were probed during a task where uncertainty can be well controlled have shown a rather complex set of neural responses [30]. It seems that the overall representation of uncertainty is complicated.

If we find some data disagreeing with a given theory of neural representation of uncertainty (which would be very easy) should we declare it falsified and move on? I disagree with this conclusion [33]. Instead we can view theories as priors, allowing us to deal with approximately true theories. Instead of a theory that posits that no information is lost (as is the case for all theories of optimal processing) we should use a theory that states that high information loss is less probable than low information loss. Instead of assuming that the neural code for uncertainty is the same in all brain areas (as is generally implied) we might prefer to assume that it is similar in related brain areas. We should work on finding ways of making approximately correct theories useful, by using them as part of the data analysis process.

Future promise

Reformulating existing theories as priors and combining them with other theories and with data should allow making these theories more useful, even if they do not capture the full amount of variance. In fact, such combinations between Bayesian data analysis and computational neuroscience modeling promises a more constructive process for improving theories and for analyzing and optimizing experiments.

References

1. Ethier C, Oby E, Bauman M, Miller L. Restoration of grasp following paralysis through brain-controlled stimulation of muscles. *Nature*. 2012; 485:368–371. [PubMed: 22522928]
- 2*. Paninski L, Eden U, Brown E, Kass R. Statistical analysis of neurophysiological data. forthcoming. An introductory textbook to the statistical analysis of neural data.
3. Lewi J, Butera R, Paninski L. Sequential optimal design of neurophysiology experiments. *Neural Computation*. 2009; 21:619–687. [PubMed: 18928364]
- 4**. Park M, Pillow JW. Receptive field inference with localized priors. *PLoS computational biology*. 2011; 7:e1002219. Shows how priors over potential receptive fields based on past experiments can dramatically improve the process of receptive field estimation. [PubMed: 22046110]

5. Corbett EA, Perreault EJ, Kording KP. Decoding with limited neural data: a mixture of time-warped trajectory models for directional reaches. *Journal of Neural Engineering*. 2012; 9:036002. [PubMed: 22488128]
6. Woolrich MW. Bayesian inference in fMRI. *NeuroImage*. 2012; 62:801–810. [PubMed: 22063092]
7. Jansen R, Yu H, Greenbaum D, Kluger Y, Krogan NJ, Chung S, Emili A, Snyder M, Greenblatt JF, Gerstein M. A Bayesian networks approach for predicting protein-protein interactions from genomic data. *Science*. 2003; 302:449–453. [PubMed: 14564010]
8. Hume, D. Courier Dover Publications. CreateSpace Independent Publishing Platform; 1739/2012. A treatise of human nature.
9. Wolpert DH. What the no free lunch theorems really mean; how to improve search algorithms. Santa fe institute working paper. 2012:12.
10. Fusi S, Drew PJ, Abbott LF. Cascade models of synaptically stored memories. *Neuron*. 2005; 45:599–611. [PubMed: 15721245]
11. Berniker M, Kording K. Estimating the sources of motor errors for adaptation and generalization. *Nat Neurosci*. 2008; 11:1454–1461. [PubMed: 19011624]
12. Nachman I, Regev A, Friedman N. Inferring quantitative models of regulatory networks from expression data. *Bioinformatics*. 2004; 20 (Suppl 1):i248–256. [PubMed: 15262806]
13. Trommershauser, J.; Kording, K.; Landy, MS. *Sensory Cue Integration*. Oxford: Oxford University Press; 2010.
14. Griffiths TL, Tenenbaum JB. Optimal predictions in everyday cognition. *Psychol Sci*. 2006; 17:767–773. [PubMed: 16984293]
- 15**. van den Berg R, Vogel M, Josic K, Ma WJ. Optimal inference of sameness. *Proceedings of the National Academy of Sciences of the United States of America*. 2012; 109:3178–3183. Correct treatment of decision making where subjects have to decide if two stimuli are identical in the presence of uncertainty. [PubMed: 22315400]
16. Ma WJ, Navalpakkam V, Beck JM, Berg R, Pouget A. Behavior and neural basis of near-optimal visual search. *Nature Neuroscience*. 2011; 14:783–790.
17. Beierholm, U.; Shams, L.; Kording, K.; Ma, WJ. Comparing Bayesian models for multisensory cue combination without mandatory integration. *Neural Information Processing Systems*; Vancouver: 2008.
18. Perfors A, Tenenbaum JB, Griffiths TL, Xu F. A tutorial introduction to Bayesian models of cognitive development. *Cognition*. 2011; 120:302–321. [PubMed: 21269608]
- 19**. Bowers JS, Davis CJ. Bayesian just-so stories in psychology and neuroscience. *Psychological bulletin*. 2012; 138:389. Summarizes a broad range of possible criticism of many Bayesian approaches to the modeling of behavior. [PubMed: 22545686]
20. Griffiths TL, Chater N, Norris D, Pouget A. How the Bayesians got their beliefs (and what those beliefs actually are): comment on Bowers and Davis (2012). *Psychological Bulletin*. 2012; 138:415–422. [PubMed: 22545687]
21. Dvijotham, K.; Todorov, E. Reinforcement Learning and Approximate Dynamic Programming for Feedback Control. 2012. *Linearly Solvable Optimal Control*; p. 119-141.
22. Ma WJ, Beck JM, Latham PE, Pouget A. Bayesian inference with probabilistic population codes. *Nat Neurosci*. 2006; 9:1432–1438. [PubMed: 17057707]
23. Fiser J, Berkes P, Orban G, Lengyel M. Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn Sci*. 2010; 14:119–130. [PubMed: 20153683]
24. Schultz W. Updating dopamine reward signals. *Current opinion in neurobiology*. 2012
25. Pouget A, Beck JM, Ma WJ, Latham PE. Probabilistic brains: knowns and unknowns. *Nature Neuroscience*. 2013; 16:1170–1178.
26. Carandini M, Demb JB, Mante V, Tolhurst DJ, Dan Y, Olshausen BA, Gallant JL, Rust NC. Do we know what the early visual system does? *J Neurosci*. 2005; 25:10577–10597. [PubMed: 16291931]
- 27*. Vilares I, Howard JD, Fernandes HL, Gottfried JA, Kording KP. Differential Representations of Prior and Likelihood Uncertainty in the Human Brain. *Current biology* : CB. 2012 Asks how

prior and likelihood uncertainty are represented in the nervous system in a visual position estimation task.

- 28*. O'Reilly JX, Schuffelgen U, Cuell SF, Behrens TE, Mars RB, Rushworth MF. Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proceedings of the National Academy of Sciences of the United States of America*. 2013 Shows how priors can be combined with other pieces of information that is integrated over time and analyzes the involved brain regions.
29. Schultz W. Dopamine signals for reward value and risk: basic and recent data. *Behavioral and Brain Functions*. 2010; 6:24. [PubMed: 20416052]
30. Gu Y, Angelaki DE, DeAngelis GC. Neural correlates of multisensory cue integration in macaque MSTd. *Nature neuroscience*. 2008; 11:1201–1210.
31. Fetsch CR, DeAngelis GC, Angelaki DE. Bridging the gap between theories of sensory cue integration and the physiology of multisensory neurons. *Nature Reviews Neuroscience*. 2013; 14:429–442.
32. O'Reilly JX, et al. Brain Systems for Probabilistic and Dynamic Prediction: Computational Specificity and Integration. *PLoS Biology*. 2013; 11:e1001662. [PubMed: 24086106]
33. Fernandes HL, Kording KP. In praise of “false” models and rich data. *Journal of motor behavior*. 2010; 42:343–349. [PubMed: 21184351]

Highlights

- Bayesian statistics is a standard tool for modeling and data analysis
- Allows combining past data with new evidence
- In many cases people behave as predicted by Bayesian statistics
- Models should progressively be combined with data

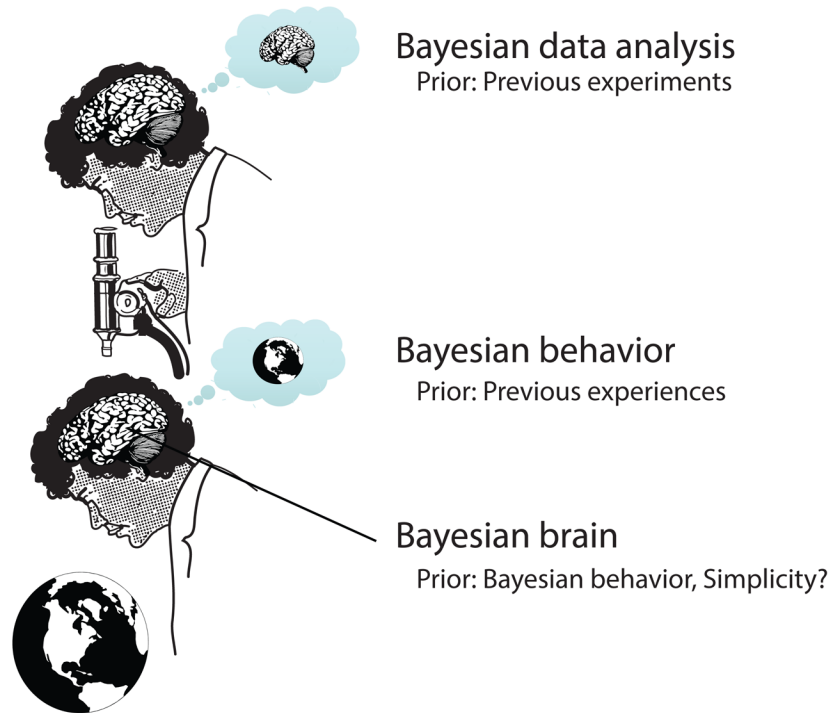


Figure 1. Scientists try to make sense of the brain. Brains try to make sense of the world. Both processes are statistical. But current theories are focused on simplicity of the neural code, the assumption of which may be wrong.