

# Electrophysiology Analysis, Bayesian

Jakob H. Macke\*

Max Planck Institute for Biological Cybernetics and Bernstein Center for Computational Neuroscience, Tübingen, Germany

## Synonyms

[Bayesian modelling of neural recordings](#); [Bayesian neural data analysis](#)

## Definition

*Bayesian analysis of electrophysiological data* refers to the statistical processing of data obtained in electrophysiological experiments (i.e., recordings of action potentials or voltage measurements with electrodes or imaging devices) which utilize methods from Bayesian statistics. Bayesian statistics is a framework for describing and modelling empirical data using the mathematical language of probability to model uncertainty. Bayesian statistics provides a principled and flexible framework for combining empirical observations with prior knowledge and for quantifying uncertainty. These features are especially useful for analysis questions in which the dataset sizes are small in comparison to the complexity of the model, which is often the case in neurophysiological data analysis.

## Detailed Description

### Overview

The Bayesian approach to statistics has become an established framework for analysis of empirical data (Gelman et al. 2013; Spiegelhalter and Rice 2009). While originating as a subdiscipline of statistics, Bayesian techniques also have become associated with the field of *machine learning* (Bishop 2006; Barber 2012). Bayesian statistics is well suited for the analysis of neurophysiological data (Brown et al. 2004; Kass et al. 2005; Chen 2013): It provides a principled framework for incorporating a priori knowledge about the system by using *prior distributions*, as well as for quantifying the residual (or *posterior*) uncertainty about the parameters after observing the data. For many analysis questions in neurophysiology, one needs to make inferences based on datasets which are small in comparison to the dimensionality or complexity of the model of interest. First, this makes it important to *regularize* the parameter estimates such that they favor explanations of the data which are consistent with prior knowledge. The use of priors also makes it possible to automatically control the complexity of the model inferred from data. Second, the fact that the data sizes are small also implies the need to quantify and visualize to what extent the parameters of the model are well constrained by the data. Third, in Bayesian statistics the parameters of the model are themselves treated as stochastic variables. This provides means of defining richer models by using simple models as building blocks of hierarchically defined models. Fourth, Bayesian statistics provides powerful machinery for dealing with the presence of unobserved processes in the model (so-called

---

\*Email: [jakob@tuebingen.mpg.de](mailto:jakob@tuebingen.mpg.de)

latent variables), which are ubiquitous in neurophysiological applications, e.g., arising from internal states or inputs that cannot be measured directly.

In Bayesian statistics, one starts by writing down a probabilistic model  $P(Y|\theta)$  of how data  $Y$  collected in an experiment are related to an underlying parameter  $\theta$ . Regarded as a function of  $\theta$ ,  $P(Y|\theta)$  is sometimes referred to as the *likelihood*. Prior knowledge about the possible values of  $\theta$  is encoded by a *prior distribution*  $P(\theta)$ . Taken together, the prior and the model  $P(Y|\theta)$  define a *generative model* of the data – one models the process of data generation as first picking a set of parameters from  $P(\theta)$ , and then data as being generated from the likelihood model  $P(Y|\theta)$ . In *Bayesian inference*, one then tries to invert this process: Given empirical data  $Y$ , which values of  $\theta$  are consistent *both* with  $Y$  and with the prior assumptions encoded in  $P(\theta)$ ? The trade-off between prior and likelihood will be determined by the amount of available data: For small dataset sizes, the prior will have a strong influence, but for large datasets, the likelihood term which depends on the observed data will dominate. Thus, the use of prior distributions can be seen as a form of *regularization* which protects the model against *overfitting* to the observed data.

The *posterior distribution*  $P(\theta|Y)$  is calculated via Bayes rule

$$P(\theta|Y) = \frac{P(Y|\theta)P(\theta)}{P(Y)}. \quad (1)$$

The posterior distribution  $P(\theta|Y)$  can then be used to make statements about the parameter values  $\theta$ . For example, the *posterior mean*  $E(\theta|Y) = \int \theta P(\theta|Y) d\theta$  is often reported and visualized in analyses of neurophysiological data as a *point estimate* of the parameters. In addition, the posterior distribution also gives insight into which properties of  $\theta$  are well or less well constrained by the data. If, for example, the posterior variance  $\text{Var}(\theta|Y)$  is small, this implies that the posterior distribution is concentrated around the posterior mean and thus that  $\theta$  is well constrained by the data. In general, Bayesian estimators are derived from the posterior distribution, and the focus of Bayesian approaches is always to characterize the distribution parameters  $\theta$  *given a particular dataset*. This is in contrast to classical (or frequentist) statistical approaches, which generally focus on making statements about what will happen – or what is unlikely to happen – if one repeatedly sampled dataset *given a particular parameter setting*.

The denominator  $P(Y)$  in Eq. 1 has to be such that the posterior distribution is normalized, i.e.,  $P(Y) = \int P(Y|\theta')P(\theta')d\theta'$ . As  $P(Y)$  is the likelihood of the data after marginalizing out (i.e., integrating) the parameters  $\theta$ ,  $P(Y)$  is sometimes referred to as the *marginal likelihood* or *evidence*. The evidence provides an estimate of how likely the observed data are for a given model and prior. It is a useful quantity for setting so-called hyperparameters as well as for calculating *Bayes factors*. A Bayes factor is the ratio of the marginal likelihoods of two models and can be used for hypothesis testing and model selection, i.e., for deciding which of two possible models provides a better explanation of some observed data  $Y$  (Gelman et al. 2013; Spiegelhalter and Rice 2009). While the use of Bayes factors is gaining popularity in the field of neuroscience, publishing conventions imply that the majority of statistical reporting of results in neurophysiological studies is based on classical, frequentist tests and definitions of p-values.

### Example: Receptive Field Estimation

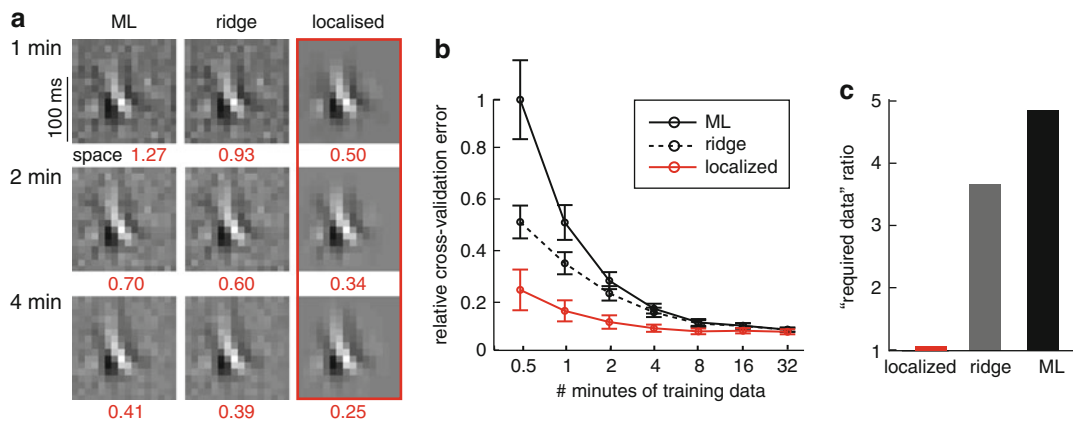
We illustrate the utility of Bayesian approaches for neural data analysis using the example of receptive field estimation for stochastic stimuli using linear models. In a linear encoding model (Paninski et al. 2007), it is assumed that mean firing rate  $\mu(s)$  of a neuron in response to a given  $D$ -dimensional stimulus can be modelled as being a linear function of the stimulus  $\mu(s) = \sum_{i=1}^D \theta_i s_i$ .

In the simplest case, the variability around the mean response is then assumed to be given by a Gaussian distribution  $Y = y|s \sim \mathcal{N}(\mu(s), \sigma^2)$  with variance  $\sigma^2$ . The parameter vector  $\theta$  is called the receptive field, and one tries to estimate  $\theta$  from the responses  $y_1 \dots y_n$  of the neuron to multiple stimuli  $s_1 \dots s_n$ . In classical approaches, one would not place any prior distribution on the values of the parameters  $\theta$ , and this approach would yield receptive field estimates which *overfit* and therefore give noisy estimates especially for small datasets (see Fig. 1a, left column).

In Bayesian approaches one places a prior distribution  $P(\theta)$  on  $\theta$ . A popular choice for  $P(\theta)$  is to chose multivariate normal distribution  $P(\theta) \sim \exp(-\frac{1}{2} \sum_{i,j} \theta_i \theta_j Q_{ij})$  on  $\theta$ , where  $Q$  is the inverse-covariance matrix of the distribution, and different choices of  $Q$  correspond to different priors.  $Q$  is sometimes chosen to be proportional to an identity matrix. In this case, the Bayesian estimate of  $\theta$  penalizes solutions for which the square deviations  $\sum_j \theta_j^2$  are big. However, as this simple prior does not well capture the structure of receptive fields, it only yields slightly improved estimates (Fig. 1a, middle column). It has generally been assumed that receptive fields are smooth and localized, and covariance matrices which reflect these properties have been developed (Sahani and Linden 2003; Park and Pillow 2011). Figure 1b and c shows that using the Bayesian approach developed by Park and Pillow (which favors solutions that are localized and smooth) yields receptive field estimates which have superior quality to those obtained using maximum likelihood and which are identifiable on smaller dataset sizes. It is worth noting this prior (and any appropriately constructed Bayesian prior) only *favors* but does not *enforce* receptive fields which are consistent with its assumptions and therefore would still leave open the possibility of being “overruled” if the data provide strong evidence for a solution which violates the assumptions.

## Algorithmic Challenges

One of the key challenges and practical drawbacks of Bayesian statistics is the fact that computation of the posterior distribution  $P(\theta|Y)$  is often hard. Exact solutions are only available in a small number



**Fig. 1** Illustration of a Bayesian approach for estimating receptive fields (RFs) (Modified, with permission, from Park and Pillow, Plos Computational Biology (2011) (Park and Pillow 2011)). **(a)** Spatiotemporal RFs of neurons in primary visual cortex. A light pixel indicates that the neuron is excited by a dark stimulus at a given spatiotemporal position, a dark pixel that its firing is suppressed, and *gray* that its firing rate is not modulated. **ML:** RFs estimated using maximum likelihood (i.e., with a “non-Bayesian” approach) using 1, 2, or 4 min of data. **Ridge:** RF estimated with a simple prior that favors solutions with small weights. **Localized:** RFs estimated with Bayesian method developed by Park and Pillow which incorporates the prior knowledge that receptive fields are localized and smooth. Localized estimator achieves better receptive field estimates (as indicated by a cross-validation error metric, *red* numbers). **(b)** Advantage of localized estimator persists across different dataset sizes. **(c)** On average, the non-Bayesian method (ML) requires five times more data than localized estimator to achieve a similar cross-validation error

of cases (e.g., when the likelihood of the model is in the exponential family and the prior distribution is *conjugate* to the likelihood (Gelman et al. 2013) ), but not for most models of interest in neurophysiological data analysis. Therefore, in general, approximate methods have to be used to characterize the posterior distribution and its properties (Chen 2013).

Approximate methods can be broadly characterized as being either *deterministic* or *stochastic*. In deterministic approximations, the posterior distribution is approximated by a distribution which has a simpler functional form, and various approaches exist for finding a “good” approximation (such as the Laplace approximation, Expectation Propagation, Variational Inference; see Bishop (2006) for details). In stochastic (or *Monte Carlo*) methods, sampling algorithms are used to generate samples from the posterior distribution  $P(\theta|Y)$ , and these samples can then be used to perform analyses such as calculating the mean and other moments of the distribution or calculating its marginals. While Monte Carlo methods are typically more flexible than deterministic approximations, sampling algorithms such as *Markov Chain Monte Carlo* methods can be computationally intensive (Kass et al. 1998; Gelman et al. 2013; Cronin et al. 2010).

## Example Applications

Bayesian statistical methods have been used extensively on a wide range of analysis questions within neurophysiology, including the following examples:

**Neural Characterization:** To describe how neural spiking activity depends on external stimuli, on its own spiking history as well as on the activity of other neurons, Bayesian methods can be used to estimate receptive fields (Sahani and Linden 2003; Gerwinn et al. 2010; Park and Pillow 2011), tuning curves (Cronin et al. 2010), and spike-history filters (Paninski et al. 2007).

**Spike Sorting and Detection:** Inference in hierarchical Bayesian models has been used to extract putative spikes of single neurons from extracellular recordings (Wood et al. 2004) or calcium measurements (Vogelstein et al. 2009 Jul 22).

**Stimulus Reconstruction and Decoding:** To reconstruct external stimuli and behavior from population activity or to decode intended movements for brain-machine interface applications, Bayesian time series models have been developed (Wu et al. 2006; Gerwinn et al. 2009).

**Estimation of Information-Theoretic Quantities:** Priors over histograms have been proposed in order to reduce the bias in estimating information-theoretic quantities such as entropy or mutual information (Nemenman et al. 2004; Archer et al. 2012).

**Functional Connectivity across Brain Areas:** Functional connections across brain areas have been estimated with a range of different Bayesian approaches. In particular, *Dynamical Causal Models* have enjoyed popularity especially for modelling fMRI and EEG data (Marreiros et al. 2010).

## Cross-References

► [Imaging Analysis, Bayesian](#)

## References

Archer E, Park IM, Pillow J (2012) Bayesian estimation of discrete entropy with mixtures of stick-breaking priors. *Adv Neural Inf Process Syst* 25:2024–2032

- Barber D (2012) Bayesian reasoning and machine learning. Cambridge University Press, Cambridge
- Bishop CM (2006) Pattern recognition and machine learning. Springer, New York
- Brown EN, Kass RE, Mitra PP (2004) Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nat Neurosci* 7(5):456–461
- Chen Z (2013) An overview of Bayesian methods for neural spike train analysis. *Comput Intell Neurosci* 2013(251905), p 17. doi:10.1155/2013/251905
- Cronin B, Stevenson IH, Sur M, Körding KP (2010) Hierarchical Bayesian modeling and Markov chain Monte Carlo sampling for tuning-curve analysis. *J Neurophysiol* 103(1):591–602. doi:10.1152/jn.00379.2009
- Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB (2013) Bayesian data analysis, 3rd edn. Chapman and Hall/CRC
- Gerwinn S, Macke J, Bethge M (2009) Bayesian population decoding of spiking neurons. *Front Comput Neurosci* 3:21
- Gerwinn S, Macke JH, Bethge M (2010) Bayesian inference for generalized linear models for spiking neurons. *Front Comput Neurosci* 4:12. doi:10.3389/fn-com.2010.00012, ISSN 1662–5188 (Electronic); 1662–5188 (Linking)
- Kass RE, Carlin BP, Gelman A, Neal RM (1998) Markov chain Monte Carlo in practice: a roundtable discussion. *Am Stat* 52(2):93–100
- Kass RE, Ventura V, Brown EN (2005) Statistical issues in the analysis of neuronal data. *J Neurophysiol* 94(1):8–25, ISSN 0022-3077 (Print)
- Marreiros AC, Stephan KE, Friston KJ (2010) Dynamic causal modeling. *Scholarpedia* 5(7):9568
- Nemenman I, Bialek W, van Steveninck R d R (2004) Entropy and information in neural spike trains: progress on the sampling problem. *Phys Rev E Stat Nonlin Soft Matter Phys* 69(5 Pt 2):056111, ISSN 1539-3755 (Print)
- Paninski L, Pillow J, Lewi J (2007) Statistical models for neural encoding, decoding, and optimal stimulus design. *Prog Brain Res* 165:493–507. doi:10.1016/S0079-6123(06)65031-0, ISSN 0079-6123 (Print)
- Park M, Pillow JW (2011) Receptive field inference with localized priors. *PLoS Comput Biol* 7(10): e1002219. doi:10.1371/journal.pcbi.1002219
- Sahani M, Linden JF (2003) How linear are auditory cortical responses?. In: Becker S, Thrun S, Obermayer K (eds) *Advances in neural information processing systems*. The MIT Press, Cambridge, Massachusetts, vol 15, p 317
- Spiegelhalter D, Rice K (2009) Bayesian statistics. *Scholarpedia* 4(8):5230
- Vogelstein JT, Watson BO, Packer AM, Yuste R, Jedynak B, Paninski L (2009) Spike inference from calcium imaging using sequential Monte Carlo methods. *Biophys J* 97(2):636–655. doi:10.1016/j.bpj.2008.08.005, ISSN 1542–0086 (Electronic)
- Wood F, Fellows M, Donoghue JP, Black MJ (2004) Automatic spike sorting for neural decoding. In: *Proceedings of the 27th IEEE conference on engineering in medicine and biological systems*, pp 4126–4129
- Wu W, Gao Y, Bienenstock E, Donoghue JP, Black MJ (2006) Bayesian population decoding of motor cortical activity using a kalman filter. *Neural Comput* 18(1):80–118