Biological Psychology 94 (2013) 490–497

Contents lists available at ScienceDirect

Biological Psychology

journal homepage: www.elsevier.com/locate/biopsycho

An improved algorithm for model-based analysis of evoked skin conductance responses☆

Dominik R. Bach[a,b,c,*], Karl J. Friston[a], Raymond J. Dolan[a,b]

[a] Wellcome Trust Centre for Neuroimaging, University College London, United Kingdom
[b] Berlin School of Mind and Brain, Humboldt University Berlin, Germany
[c] Zurich University Hospital for Psychiatry, Switzerland

ARTICLE INFO

Article history:
Received 4 February 2013
Accepted 15 September 2013
Available online 21 September 2013

Keywords:
Skin conductance responses (SCR)
Galvanic skin response (GSR)
Electrodermal activity (EDA)
General linear convolution model (GLM)
Generative model
Model inversion

ABSTRACT

Model-based analysis of psychophysiological signals is more robust to noise – compared to standard approaches – and may furnish better predictors of psychological state, given a physiological signal. We have previously established the improved predictive validity of model-based analysis of evoked skin conductance responses to brief stimuli, relative to standard approaches. Here, we consider some technical aspects of the underlying generative model and demonstrate further improvements. Most importantly, harvesting between-subject variability in response shape can improve predictive validity, but only under constraints on plausible response forms. A further improvement is achieved by conditioning the physiological signal with high pass filtering. A general conclusion is that precise modelling of physiological time series does not markedly increase predictive validity; instead, it appears that a more constrained model and optimised data features provide better results, probably through a suppression of physiological fluctuation that is not caused by the experiment.

1. Introduction

Recent interest in model-based analysis of skin conductance responses (SCR) (Bach & Friston, 2013) is – in part – motivated by the need to improve the temporal resolution of inference in rapid event-related paradigms (Barry, Feldmann, Gordon, Cocker, & Rennie, 1993). In model-based approaches, generative (forward) models specify how underlying physiological or psychological states generate observed data. Model inversion refers to estimating these (hidden) states from data. It turns out that inversion of probabilistic forward models has fundamental advantages, one of them being a propensity to suppress the effect of measurement noise (Bach, Flandin, Friston, & Dolan, 2009; Bach, Daunizeau, Friston, & Dolan, 2010; Bach, Friston, & Dolan, 2010; Bach, Daunizeau, Kuelzow, Friston, & Dolan, 2011). Statistical inference on the hidden states is generally more powerful than statistical comparisons of observed data because the models are more informed or constrained, leaving greater degrees of freedom in the data for efficient inference. Furthermore, the parameters of generative models provide a quantitative and explicit description of assumptions implicit in operational approaches (Bach & Friston, 2013), thus allowing for rigorous testing of those assumptions. Finally, model-based approaches afford quantitative rather than qualitative description of hidden, psychological processes.

Evoked skin conductance responses (eSCR) that follow a short (less than second) stimulus can be analysed with general linear convolution models – similar to the convolution models widely used in the analysis of functional magnetic resonance images (Friston, Jezzard, & Turner, 1994). In order to estimate the amplitude of evoked sympathetic nerve activity (SNA) from eSCR, we proposed such a convolution model (Bach et al., 2009). This model comprises two parts: a peripheral model incorporating a (standard linear time invariant) convolution operator, thought to be implemented by the sudomotor nerve terminal and sweat gland; and a linear neural model assuming infinitely short neural bursts immediately after each stimulus. We have shown that time invariance assumptions for the peripheral system are largely met (Bach, Flandin, Friston, & Dolan, 2010), while non-linearities in the peripheral system may occur but can easily be modelled within this framework – see (Bach & Friston, 2013) for a discussion. The model is highly regularised by placing informative constraints on the form or shape of the convolution kernel which models the peripheral response function (RF). This enables one to estimate the mean evoked SNA amplitude for each condition of an experimental design – even

* Corresponding author at: Zurich University Hospital for Psychiatry, Lenggstrasse 31, 8032 Zürich, Switzerland. Tel.: +41 44 384 21 11.
E-mail address: dominik.bach@uzh.ch (D.R. Bach).

0301-0511/$ – see front matter © 2013 The Authors. Published by Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.biopsycho.2013.09.010

when observed eSCR overlap in time. This model was designed to optimise the predictive power of the estimates, rather than to precisely reconstruct the observed time series. Indeed, when subjects observe negative-arousing or neutral pictures, picture category can be better predicted from SNA estimates than from SCR peaks, an observation that speaks to its predictive validity (Bach et al., 2009).

Model-based eSCR analysis, based on probabilistic inversion of a general linear convolution models, is thus a potentially powerful method. As with any method, however, the practical implementation makes certain technical assumptions that go beyond the known biophysical properties of the system. Three points deserve particular attention:

Firstly, the peripheral response model uses a canonical skin conductance response function (SCRF) for all experiments and individuals. Such a stereotypical response function is a strong biophysical assumption and unsupported by observation. Indeed in our own validation experiments, we observed large inter-individual variability, accounting for up to 20% of overall response variability (Bach, Flandin, et al., 2010). Therefore, we added orthogonalised Taylor expansions to the SCRF to account for differences between individuals and conditions, thus improving model fit (Bach et al., 2009). Effectively, this enables the model to fit a subject specific RF in terms of a linear mixture of basis functions of peristimulus time, where the basis set is generated by the Taylor expansion. However, because the additional basis functions are orthogonalised to the SCRF, they do not affect the estimation of the parameter controlling the amplitude of the SCRF (Calhoun, Stevens, Pearlson, & Kiehl, 2004; Hopfinger, Buchel, Holmes, & Friston, 2000). Yet, this parameter is taken to estimate the SNA. This means that additional basis functions improve data fit at the within subject level but not comparisons of SNA at the between subject level. Hence, one might ask whether modelling an individual response function (IRF) – rather than a canonical skin conductance response function (SCRF) – improves predictive validity.

There are several ways to model subject specific IRF. First, the SCRF together with the remaining basis functions can be used to estimate a subject and condition specific IRF. That is, we can reconstruct the estimated eSCR, measure the peak amplitude (over peristimulus time) and use this as an estimate of SNA amplitude, instead of the canonical parameter estimate. Other regularised basis sets also provide models of IRFs. An uninformed finite impulse response (FIR) model was proposed in Bach et al. (2009) due to its popularity in fMRI research. A cosine set also used in fMRI analysis serves the same purpose. These basis sets typically have a larger number of basis functions than basis sets built upon truncated Taylor expansions. This means that although they are more flexible, they require greater numbers of parameters to be estimated. In all these approaches, separate IRFs are estimated for each condition within one individual. A more informed approach is to assume the form of the subject specific response function is the same for all experimental conditions. We have implemented this constraint by extracting data from all conditions and fitting a response function to the first principal component of the data. We will refer to this response function as the subject-specific response function (SRF).

A second issue is that skin conductance time series comprise both phasic responses and a slowly drifting tonic component. This drift is why many analysis schemes high pass filter the signal (Boucsein, 2012), including contemporary model-based approaches (Benedek & Kaernbach, 2010a, 2010b). This renders the phasic responses finite and removes slow signal drifts which are difficult to model. In our implementation, we used a bidirectional first-order Butterworth filter with time constant of 10 s (corresponding to a cut off frequency of 0.0159 Hz) (Boucsein, 2012). A bidirectional filter was chosen as it retains peak latencies. Because this filter can slightly distort the shape of the signal, the regressors of the general linear convolution model are subjected to the same filter. The choice of the filter frequency is based on prior experience but not on theoretical considerations. Therefore, one may ask whether there is an optimal filter that provides the best data features for modelling. Generally speaking – when modelling biological time series – data features that cannot be produced by a plausible forward model are probably measurement noise or the product of hidden processes not included explicitly in the model. This usually means they can be discarded with impunity, thereby increasing the signal-to-noise ratio (SNR) of the pre-processed data. Data conditioning is then, effectively, a part of model inversion. The question here is whether there is an optimal high pass filtering of skin conductance timeseries that increases signal-to-noise. In case of a signal with precisely known RF, the matched filter theorem provides a way of theoretically deriving a filter that maximises the SNR. In our case, the true RF is not precisely known, and also varies between individuals, such that we sought to empirically determine the filter characteristics that maximise predictive validity of SN estimates.

Finally, a linear neuronal model makes the strong assumption that SNA evoked by a short stimulus occurs with constant latency. We have previously shown that under this assumption there is no evidence for time-varying responses in the peripheral system (Bach, Flandin, et al., 2010). Here, we revisit this assumption and investigate whether modelling variations in neuronal latency improves predictive validity, under the assumption of an invariant peripheral response. Hence, we compare linear and non-linear models. Two particular non-linear models are considered. First, we used our previous approach that uses Dynamic Causal Modelling (DCM) (Bach, Daunizeau, et al., 2010) to obtain estimates of SNA amplitude per trial by letting response amplitude and onset vary on a trial-by-trial basis. Note that the neural model here is still informed insofar as it specifies a certain response window. Some authors propose uninformed neural models; in other words, they assume that SNA can occur any time, but to only use SNA during post-stimulus time windows for analysis (Benedek & Kaernbach, 2010a, 2010b). We sought to emulate this approach using DCM for spontaneous fluctuations (Bach et al., 2011). Both approaches yield a trial-by-trial estimate of SNA amplitude, which was averaged across experimental conditions for comparison with other approaches.

In some circumstances – e.g. to use neural response estimates as explanatory variables for analysis of independent experimental data, such as fMRI, trial-by-trial estimates of SN amplitudes may be required. Here, we sought to establish whether linear models are sufficient for this purpose or whether the iterative procedures required for non-linear models inherent in DCM are justified.

We have previously discussed how to benchmark methods that estimate hidden variables from observed data (Bach & Friston, 2013). One way is by making certain assumptions about what causes the hidden variable to change. A consensus assumption in the psychophysiology literature is that emotionally arousing events increase sympathetic arousal, as engendered by negative and positive arousing images. This has been demonstrated using operational approaches (Amrhein, Muhlberger, Pauli, & Wiedemann, 2004; Greenwald & Lang, 1989; Johnsen, Thayer, & Hugdahl, 1995; Winton, Putnam, & Krauss, 1984). Here, we assume that negative and positive arousing images would elicit greater sympathetic arousal than neutral non-arousing images, and evaluated different models in terms of their ability to distinguish between image types, using just the observed SCR.

In summary, we evaluated our method empirically, by comparing the predictive validity of different generative models. In a first step, we compared a canonical response function against various forms of an individualised response functions (IRF). Taking the best model from this step, we then compared various filter settings, non-linear methods, and the efficiency of trial-by-trial