Pupil dilation indexes statistical learning about the uncertainty of stimulus distributions

Francesco Silvestrin (F.Silvestrin@uea.ac.uk)
School of Psychology, University of East Anglia, Norwich Research Park
Norwich NR4 7TJ, United Kingdom

William D Penny (w.penny@uea.ac.uk)
School of Psychology, University of East Anglia, Norwich Research Park
Norwich NR4 7TJ, United Kingdom

Thomas H B FitzGerald (t.fitzgerald@uea.ac.uk)
School of Psychology, University of East Anglia, Norwich Research Park
Norwich NR4 7TJ, United Kingdom

Abstract:

Learning about the uncertainty of environmental stimuli is a fundamental requirement of adaptive behaviour. In this experiment we probe whether pupil dilation in response to brief auditory stimuli reflects statistical learning about the underlying stimulus distributions. Specifically, we consider whether pupil dilation reflects automatic (task-irrelevant) learning about the precision of Gaussian distributions of tones. By comparing responses to perceptually identical outlier and standard tones in low and high precision blocks, we provide clear evidence that subjects do indeed learn about precision, as reflected by increased responses to surprising (outlier) tones during high precision blocks. This extends previous work looking at electrophysiological effects of precision learning, and provides new evidence that the putatively noradrenergic processes underlying pupil dilation reflect learning about the uncertainty of stimulus distributions. In addition, we use our data to test a new convolution-based approach for analysing pupillometry data, which we believe has considerable promise for this and future studies.

Keywords: Auditory Oddball; Predictive Coding; Surprise; Statistical learning; Pupillometry

Introduction

In recent years, probabilistic accounts of brain function have gained considerable popularity (for example, Knill & Pouget, 2004). Arguably, the most popular of these frameworks is Predictive Coding (see Clark, 2013), according to which the brain represents all information probabilistically, in the form of Gaussian (normal) probability distributions, and uses previous knowledge to disambiguate the incoming signal. Internal models of the world are used to make predictions about the stimuli that the sensorium is going to receive in any given moment. The brain is always trying to anticipate what is coming next, and it does so in a hierarchically organized fashion (see Felleman & Van Essen, 1991; Friston, 2008).

A key concept in Predictive Coding is that of precision (see Kanai et al., 2015). More precise (less uncertain) information is represented as a narrower Gaussian (with a lower variance) compared with less precise information. Following this principle, a precise model of the world should lead to the formation of precise predictions, the violation of which should be perceived as more surprising. This hypothesis has been tested by Garrido and colleagues (2013), who used a modified version of the auditory oddball task (e.g. Boly et al., 2011) in which they presented participants with a series of tones whose frequency was sampled either from a wide or narrow normal distribution in different blocks. They found greater responses to outliers in the context of a more precise stimulus distribution both behaviourally (in the form of reaction times) and with magnetoencephalography (mismatch negativity).

We sought to extend this approach to consider whether the putatively noradrenergic mechanisms underlying pupil dilation responses to surprising stimuli (Preuschoff et al., 2011; Lavín et al., 2014: Alamia et al., 2019) reflect learning about precision. Additionally, we used our data to test a novel, model-based approach to analysing pupillometry data, which we believe has considerable promise for exploring automatic statistical learning in humans.

Materials and Methods

Procedure

16 participants (12 females), aged 18 to 34 (mean = 21.1) were asked to look at a fixation cross in the centre of a computer screen while listening to a series of tones through headphones. Their task consisted only in pressing the space bar when the sound came only from one of the two headphones' noise sources (i.e. when they heard it coming only from one side). Importantly,

the pitch of the tones was entirely task irrelevant, deconfounding outlier and target tones, which is a concern with several previous studies (Liao et al., 2016; Hong et al., 2014), as well as indexing automatic, rather than task-dependent statistical learning processes.

The experiment was divided into 4 sessions, during each of which subjects were presented with 800 pure tones, each lasting 50 ms, with and interstimulus interval of one second. 4 blocks (2 high variance blocks and 2 low variance blocks) were present in each session, with 200 tones each and no breaks between blocks. The order of the blocks was counterbalanced and participants were not aware of the presence of different blocks within each session. For each session the frequency of 688 out of 800 tones was sampled from a Gaussian distribution in log-frequency, with mean μ = 500 Hz and standard deviation σ_l = 0.5 octaves for low variance blocks and σ_h =1.5 octaves for high variance blocks. Out of the 112 remaining tones, 56 were standard probes (500 Hz, corresponding to the mean of the distribution) and 56 were odd probes (2000 Hz, two octaves above the mean), which slightly distorted the probability distribution, adding two point-masses of 7% probability each. The number of unilateral, target tones varied across sessions (82, 80, 75, and 78 respectively). Both probes and targets were pseudo-randomly inserted in the stream, and targets were made to never coincide with a probe, or occur immediately after it. This was done because targets are very likely to elicit a strong, long lasting pupil response, which would confound the effect of surprise.

Pupillometry data was recorded at 500 Hz using an EyeLink 1000 eye-tracking device. Linear interpolation was used to remove artefacts relating to eyeblinks and saccades, and the data were low pass filtered at 20 Hz.

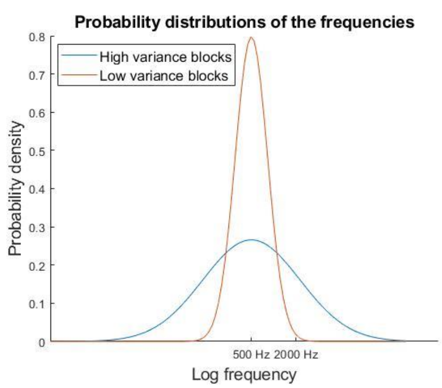

Figure 1. Probability distributions from which frequencies were sampled. The task alternated between high and low variance blocks. In line with previous work (Garrido et al., 2013) probes were added at 500 Hz and 2000 Hz.

Data analysis

Classical probe-tone analysis

We first analysed our data using a classical model-free approach, where we epoched, baseline-corrected and averaged the responses to the probe tones in each variance condition for each subject. Figure 3 shows the average responses for odd and standard probes in both high and low variance blocks.

We then averaged all data-points from 500ms to 1000ms after stimulus onset and performed a two-way repeated measures ANOVA with variance (high vs low) and probe type (odd vs standard) as within-subjects factors.

Convolution modelling

Additionally, we analysed our data using a novel convolution based approach, derived from similar techniques developed for analysing fMRI data (Penny et al., 2003). Briefly, regressors modelling tone onsets and key task parameters are convolved with a Gamma function governed by shape, scale and delay parameters. (For a related approach see Korn & Bach, 2016). These are then entered into a General Linear Model with a first-order autoregressive error process. Thus observation y_t made at time t is modelled in terms of predictors $\mathbf{x}_t$, regression weights $\mathbf{w}$ and and error term e_t that has both autroregressive (ae_{t-1}) and Independent and Identically Distributed (IID) (z_t) components (for a fuller discussion, see Penny et al., 2003):

$$\begin{aligned} y_t &= \mathbf{x}_t\mathbf{w} + e_t, \\ e_t &= ae_{t-1} + z_t. \end{aligned} \tag{1}$$

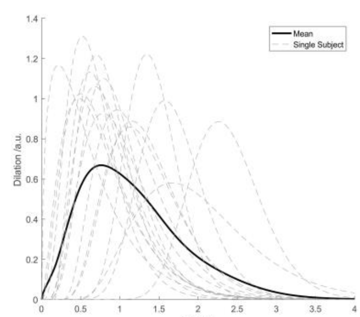

Figure 2: Estimated pupil dilation responses derived from the convolution model. (Single subject responses in grey, and the mean in black). These strongly resemble averaged responses from tasks using slower designs (for example Hong et al., 2014)

To model learning about variance, we used a simple algorithm in which trial-by-trial variance estimates (v_t) are updated based on a 'variance prediction error' based on the difference (in log space) between the pitch of the current tone and the mean of the distribution (500 Hz). The speed of learning is governed by a learning rate parameter α .

$$v_t = v_{t-1} + \alpha\left(s_t - v_{t-1}\right),$$
$$s_t = \left(\ln\left(p_t\right) - \ln\left(500\right)\right)^2. \quad (2)$$

The regression model we used contained the following regressors: *Events* indicating tone onset, *Variance* modelling trial-by-trial variance estimates, *Deviance* modelling the absolute distance in log space between the tone and the mean frequency of the distribution (500 Hz), *Variance*Deviance*, indicating an effect of variance learning on deviant tones, in keeping with our hypothesis, *Pitch* indicating the log frequency of each tone, and *Target* indicating whether a tone was a target or not.

Model parameters were fitted using Variational Laplace (Friston et al., 2007), with uninformative prior distributions for the parameters, which were appropriately transformed where necessary. (Weak shrinkage priors were used for the regression coefficients). Group-level statistics were based on the maximum a posteriori (MAP) estimates of each parameter.

Results

Classical probe-tone analysis

A two-way repeated measures ANOVA revealed a main effect of variance (F = 5.9195, p = 0.0280), with low variance blocks having a bigger pupil response overall, but not of probe type (F = 2.3627, p= =0.1451). Most importantly, the interaction resulted significant (F = 15.7016, p = 0.0013), with odd probes eliciting greater pupil dilation only in low variance.

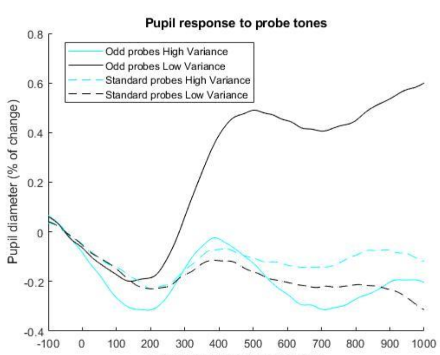

Figure 3: Averaged pupil response to odd and standard probes in low and high variance blocks.

This suggests that participants learned the distributions of the stimuli and correctly assigned a low probability to the occurrence of 2000 Hz tones when they were embedded in a more precise (with less variance) distribution.

Convolution modelling

In keeping with our probe-tone analysis, the results from our convolution modelling provided clear evidence that dynamic inferences about precision modulated pupil responses, as manifested by a strongly significant negative *Variance*Deviance* response ($t(15)$=-5.31, p<0.001, Figure 4). Additionally, the results of our convolution modelling also clearly showed a main effect of deviance itself ($t(15)$=4.74, p<0.001, Figure 4), which it was not possible to clearly demonstrate in out probe-tone results. This is likely to reflect greater sensitivity in our model-based approach, which mitigates the effect of overlapping responses to successive stimuli. As might have been expected, no clear effect of *Variance* was observed ($t(15)$=1.09, p=0.295, Figure 4), though it seems possible that it might have affected baseline responses that we do not consider in this work.

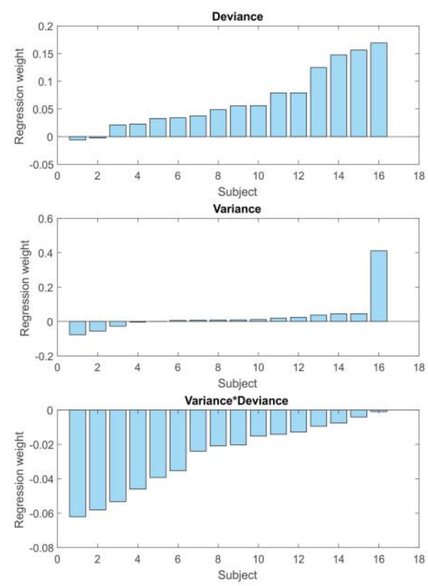

Figure 4: MAP estimates of single subject regression weights for the *Deviance*, *Variance* and *Deviance*Variance* regressors. In keeping with our hypotheses, our results indicate a positive effect of *Deviance (top)* and a negative effect of *Deviance*Variance* (bottom). No clear effect of *Variance* was observed.

Learning rates showed consider variability (mean: 0.13, variance: 0.03, range 0.01 – 0.75), which may reflect individual differences in statistical learning. However, without further investigation, it is hard to be certain about this. We will address this in future analyses.

Discussion

The results obtained in this study provide clear evidence that pupil dilation reflects statistical learning about the precision of stimulus distributions, in keeping with Predictive Coding, and other theories of probabilistic cognition. This extends previous work showing evidence of such learning in reaction times and MEG responses (Garrido et al., 2013), and suggests that pupillometry can be a useful tool for examining statistical learning about higher order properties of stimulus distributions, something we will consider further in future work.

The convolution modelling approach that we present opens up a number of exciting avenues for future analysis. In the first place, as with convolution-based approaches to fMRI analysis (Penny et al., 2003), it supports analysis of studies where stimuli are presented close together in time, relative to the natural time course of the pupil dilation response. Second, adopting a model-based approach allows one to fit and compare different models, and thus, potentially at least, to infer in a more computationally-informed way on the processes underling pupil dilation responses. However, much work needs to be done to address the reliability and accuracy of parameter estimates derived using this approach, something we will consider in future work.

In sum, our work represents a contribution both to understanding the statistical learning processes underlying pupil dilation responses to surprising stimuli, and towards analysing pupillometry data more generally.

Acknowledgments

THBF is supported by a European Research Council (ERC) Starting Grant under the Horizon 2020 program (Grant Agreement 804701).

References

Alamia, A., VanRullen, R., Pasqualotto, E., Mouraux, A., & Zenon, A. (2019). Pupil-linked arousal responds to unconscious surprisal. *Journal of Neuroscience*, 3010-18.

Boly, M., Garrido, M. I., Gosseries, O., Bruno, M. A., Boveroux, P., Schnakers, C., ... & Friston, K. (2011). Preserved feedforward but impaired top-down processes in the vegetative state. *Science*, *332*(6031), 858-862.

Brainard, D. H., & Vision, S. (1997). The psychophysics toolbox. *Spatial vision*, *10*, 433-436.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, *36*(3), 181-204.

Felleman, D. J., & Van, D. E. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral cortex (New York, NY: 1991)*, *1*(1), 1-47.

Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., & Penny, W. (2007). Variational free energy and the Laplace approximation. *Neuroimage*, *34*(1), 220-234.

Friston, K. (2008). Hierarchical models in the brain. *PLoS computational biology*, *4*(11), e1000211.

Garrido, M. I., Sahani, M., & Dolan, R. J. (2013). Outlier responses reflect sensitivity to statistical structure in the human brain. *PLoS Computational Biology*, *9*(3), e1002999.

Hong, L., Walz, J. M., & Sajda, P. (2014). Your eyes give you away: Prestimulus changes in pupil diameter correlate with poststimulus task-related EEG dynamics. *PLoS One*, *9*(3), e91321.

Liao, H. I., Yoneya, M., Kidani, S., Kashino, M., & Furukawa, S. (2016). Human pupillary dilation response to deviant auditory stimuli: Effects of stimulus properties and voluntary attention. *Frontiers in neuroscience*, *10*, 43.

Kanai, R., Komura, Y., Shipp, S., & Friston, K. (2015). Cerebral hierarchies: predictive processing, precision and the pulvinar. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *370*(1668), 20140169.

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *TRENDS in Neurosciences*, *27*(12), 712-719.

Korn, C. W., & Bach, D. R. (2016). A solid frame for the window on cognition: Modeling event-related pupil responses. *Journal of Vision*, *16*(3), 28-28.

Lavín, C., San Martín, R., & Rosales Jubal, E. (2014). Pupil dilation signals uncertainty and surprise in a learning gambling task. *Frontiers in behavioral neuroscience*, *7*, 218.

Penny, W., Kiebel, S., & Friston, K. (2003). Variational Bayesian inference for fMRI time series. *NeuroImage*, *19*(3), 727-741.

Preuschoff, K., t Hart, B. M., & Einhauser, W. (2011). Pupil dilation signals surprise: Evidence for noradrenaline’s role in decision making. *Frontiers in neuroscience*, *5*, 115.