

**Meditation and Gender Affect the Attentional Blink:
An ERP and Bayesian Learning Analysis**

Jørgen Østmo-Sæter Olsnes



MAPSYK360, masters program in Psychology,

Behaviour and Neuroscience

At

UNIVERSITY OF BERGEN

INSTITUTE FOR BIOLOGICAL AND MEDICAL

PSYCHOLOGY

Fall 2020 - Spring 2021

Word count: 12445

Supervisor: Dr Magdalena-Lucia Dumitru, Institute for Biological and Medical Psychology

Abstract

Recent models view the brain as a Bayesian probabilistic inference machine using free-energy principles to update internal models about its surroundings. This implies that top-down learning functions in the brain are caused by internal predictions about outcomes during tasks and that learning is a product of model updates based on errors in these predictions. This has given rise to Bayesian learning models like the Hierarchical Gaussian Filter that are able to efficiently model learning at an individual level and predict responses better compared to older models like the Rescorla Wagner. We modelled learning rates during an EEG Attentional Blink (AB) task where participants had no prior meditation training (N = 32). We compared learning rates between genders for two meditation types: Open Monitoring Meditation (OMM) and Fixed Attention Meditation (FAM). We assumed that responses to the second AB target within- versus outside-blink responses Target 2 would generate prediction errors, yielding an implicit learning effect. We found that females in the FAM group showed a higher ω levels for volatility estimates (third level), implying that they perceived a higher environmental changeability. Participants in the OMM group scored higher on T2 accuracy, and that this effect was mainly driven by the female participants. EEG data support this conclusion, showing that females in the OMM group exhibited higher ERP amplitudes of the P300 component for outside-AB trials. This indicates that females are instantly affected by mindfulness meditation during selective attention tasks, and that this effect can be explained through volatility estimates.

Keywords: Bayesian learning models, HGF, EEG, Meditation, Gender

Sammendrag

Nyere modeller anser hjernen som en Bayesiansk probabilistisk inferensmaskin som bruker fri-energiprinsipper til å oppdatere internaliserte modeller om omgivelsene. Dette impliserer at top-down læringsfunksjoner i hjernen blir forårsaket av interne prediksjoner om utfallet av en situasjon, og at læring er et produkt av oppdaterte modeller basert på prediksjonsfeil. Dette gav opphav til Bayesianske læringsmodeller som 'Hierarchical Gaussian Filter'. Disse er i stand til å mer effektivt modellere læring på individnivå, og predikere adferd når de sammenlignes med eldre modeller som Rescorla Wagner. Vi modellerte oppnådd læring i en EEG Attentional Blink (AB) oppgave for deltakerne uten tidligere meditasjonstrening (N = 32). Vi sammenlignet grad av læring mellom kjønn for to meditasjonstyper: «Open Monitoring Meditation» (OMM) og «Focused Attention Meditation» (FAM). Vi antok at responser til det andre AB-målet, i innen- versus utenfor-AB svar 'Target 2', ville generere prediksjonsfeil som fører til en implisitt læringseffekt. Vi fant at kvinner i FAM-gruppen viste høyere ω nivåer for 'volatility'-estimer (tredje nivå), noe som antydte at de oppfattet en høyere endringsrate i miljøet. Deltakere i OMM-gruppen scoret høyere på T2-treffsikkerhet, og denne effekten var hovedsakelig forårsaket av kvinnelige deltagere. EEG-data støtter denne konklusjonen, og viser at kvinner i OMM-gruppen viste høyere ERP-amplituder av P300-komponenten for utenfor-AB-gjennomføringer. Dette indikerer at kvinner blir raskt påvirket av 'mindfulness'-meditasjon under selektive oppmerksomhetsoppgaver, og at denne effekten kan forklares gjennom antagelser om hvor stabile omgivelsene er.

Nøkkelord: Bayesianske læringsmodeller, HGF, EEG, Meditasjon, Kjønn

Preface

There are a number of people who need to be commended for their help throughout this thesis. Most prominent among them is my supervisor. I am both grateful and indebted to Dr Magdalena-Lucia Dumitru for her patience and hard work while helping me grasp the foundation of Bayesian models of learning. This paper has been a journey. The impact by hard societal restrictions in the midst of a pandemic, and the abrupt decision by my original supervisor to leave Academia for the private sector was a big challenge. It left me unanchored and working in a field of study that did not seem to have a strong foothold within our institute. I would have struggled had it not been for Magda, who has motivated, pushed, and encouraged me beyond what I would expect was required from her as a supervisor. I would not have learned half of everything that I retain from having written this thesis without her support. A thesis that I expected would be purely theoretical and based on data already collected, received grounding in a new EEG project, with exciting and novel applications of learning models to an Attentional Blink task. Getting to work and gain experience in a lab has been enjoyable, challenging, and rewarding at the same time. I could think of no better supervisor. Thank you.

Further, I would like to thank Dr Andreea Diaconescu (University of Toronto) for her help and suggestions throughout our modelling phase. Her experience with conceptual modelling has been invaluable for the progress of this work, and I am highly grateful for adapting the logRT model to our experimental task, and for her taking the time to answer questions and giving suggestions throughout the data analysis.

I would also like to thank Jørgen A. Haug for the time spent at the EEG lab. Together, we have gathered data for both his and my thesis. It was a pleasure, and I sincerely hope we will be able to work together again in the future. Finally, a thank you goes out to everyone around me that has given me support and encouragement. It is what kept me going!

Table of contents

Abstract.....	3
Sammendrag	4
Preface	5
Table of contents	6
Contemporary Models of Learning	8
Frequentist versus Bayesian statistics	10
Frequentist statistics	10
Bayesian Statistics	14
Attention and learning	15
The attentional blink.....	16
Gender differences in attention	18
Meditation and attention.....	19
Gender differences in meditation	21
Measuring brain activity in our study.....	22
ERP components: the P300	23
Measuring learning rates: the HGF family of models	24
Getting the HGF model to work.....	25
HGF and the AB task	29
Hypotheses of our study	31
Methods	31
Participants	31

Ethics approval	31
Equipment.....	32
Stimuli	32
Procedure	33
Analyses	36
Data modelling	36
EEG data.....	36
Results	37
Accuracy.....	37
HFG Modelling Results.....	39
Comparison of parameters between meditation types.....	39
Model fit comparison	40
Discussion.....	45
Meditation impact on AB accuracy	45
Meditation effects on attention.....	46
Meditation effects on learning.....	47
Using meditation to explore cognitive performance	49
Meditation and modelling.....	50
Limitations of the Study	51
Concluding remarks and future studies	52
Reference list	53

Contemporary Models of Learning

Learning through the prism of contemporary models is an active process based on top-down predictions, not a passive storage device. Indeed, there is increasing evidence that the brain works as a Bayesian predictive learner and hypothesis tester, with information being processed at several hierarchical levels where the salience of a stimulus and the uncertainty of the environment are subject to change. These changes are dependent on each other and form a top-down predictive system, where each level affects the size and effect of perceived prediction errors (Behrens, Woolrich, Walton & Rushworth, 2007; Wacongne, Labyt, van Wassenhove, Bekinschtein, Naccache, & Dehaene, 2011). This goes against older behaviourist theories of learning that viewed learning as a passive process based on associative strengths between stimuli and responses (Chance, 2003, Siegel & Allan, 1996; Skinner, 1985; 1987). However, behaviourism has brought about a strong focus on observable and measurable responses to perceptual stimuli instead of the internal states of the learner. This led to the development of various models, often using mathematical formulas to predict behaviour. In this way, psychology became aligned with the core values of modern science that demands solvable problems, predictive hypotheses, and predicting the outcomes of hypotheses in accord with demarcation and falsification (Kruschke, 2003; Pigliucci & Boudry, 2013 Resnik, 2000).

A criticism to the behaviourist learning paradigm, however, is that it ignores the learner as an agent with intent (Danks, 2003; Mathys, Daunizeau & Stephan, 2011). To overcome this issue, more recent models tend to view the human brain not simply as a response-giver, but as a hypothesis generator that uses Bayesian statistical probability judgements to continuously update its own predictions based on inferences of the environment. It updates information and generates internal models of the world, models that it then uses to predict the future on a situation-basis, and learning is a product of errors in our generated predictions. This gives a

model of the brain that allows for hidden states of the mind such as learning rates, uncertainty, and estimates of the environment (Behrens et al., 2007). Given that the brain is a hypothesis tester, and that learning is the process where hypotheses are tested and corrected against the perceived outcomes, learning can be reduced to updates of an internal model based on prediction errors. (Diaconescu et al, 2014; Schultz & Dickinson, 2000).

A prediction error occurs when the current hypothesis for a specific outcome is disproven. This will generate a response, which will in turn readjust the overlying hypothesis to better fit the environment. If an observed response happens exactly as predicted, there is no need to update the hypothesis, since it can explain the outcome in a particular situation. In contrast, if facts go against the original hypothesis, the learner needs to revise and adjust their expectations. This readjustment of expectations to the observed environment may be viewed as the foundation of learning. (Behrens et al., 2007; Jones & Love, 2011; Mathys et al., 2011; McKiernan, 2017; Seligman, Railton, Baumeister & Sripada, 2013). By mapping estimations onto inferences and by adjusting predictions based on expectations of environment changeability, one can generate precision weighted prediction errors (pwPE). Weighing the prediction errors reduces the net surprise of the learner, which will in turn reduce the cognitive strain required to keep the learner's model updated about their environment. This is in accord with the principle of 'free energy', sometimes also known as 'active inference'. This principle states that learning agents minimize the difference between observations and an internal model of the world. Learners close this gap through updates to the model whenever prediction errors occur by either adapting to the surroundings or by optimizing the surroundings themselves to better suit the preferred outcomes (Friston 2009; 2011).

Iglesias and colleagues (2013) compared learning occurring at a high level with learning occurring at a low level by drawing a distinction between responses to reward directed prediction errors and more subtle and general updates to the predictions made. The latter

happened based on updating and improving expectations about the environment. Prediction errors resulting in reward were associated with basal forebrain activation that impacted inferences about the reward, whereas the midbrain encoded prediction errors independently of the other system on more abstract probabilities concerning the outcomes. This formed the foundation of a meta-Bayesian principle, where a dichotomous system performed different forms of learning, which was further taken as evidence that learning happens at several hierarchical levels, with one system monitoring the other (Iglesias et al., 2013). The levels help maintain an internal overview of both variance and volatility and generate prediction errors about the two systems independently. The more familiar frequentist statistics might not be able to properly model the implied view of the brain as a hypothesis tester that employs environmental statistical cues to model expectations about outcomes based on prediction errors.

Frequentist versus Bayesian statistics

There has been an increasing need for psychology as a field to embrace ecologically valid experimental settings to better capture the organic interactions happening beyond a traditional laboratory. Although statistical competency has increased over the last decades, it is still impossible to apply complex processes without grasping the complex dynamics underlying them (Dempster & McCorry, 2009; Smaldino, 2021).

Frequentist statistics. The most common way of using statistics is through the frequentist method. This method works well for many forms of analyses and requires a relatively low amount of computational effort (Eddy, 2004). The idea behind the frequentist method is to generate a hypothesis, and then calculate the probability of getting a certain result in a specific number of trials. These trials will not be dependent on each other, hence the result of one trial will not influence the next. In other words, in the frequentist method the parameters explored will not be subject to change based on the observed responses or data across trials. An example here could be a psychological experiment where a number of participants are assigned a novel

task with conditions that are based on a prior hypothesis. Each participant completes the task with a certain accuracy. With a sufficiently large sample of participants, the researcher can calculate a mean distribution, a standard deviation of the mean, perhaps also a probability value of whether different groups of participants are similar or not. In other words, the frequentist method gives a probability estimate of outcomes in future trials based on the data from trials already completed. It does not assign a probability to the hypothesis itself (Anderson, 2020; Nickerson, 2000). Earlier models of learning like the Rescorla Wagner Response-Learning model (RW) were able to generate good predictions starting from these premises. Learning happened based on the observed outcomes of stimulus conditions, and the model was altered each time the expected outcome did not match the observed outcome. If no information was observed contradicting the data, net learning would be “old expectations minus new expectations”, which amounts to zero. If someone were to observe a confirming outcome, that would represent zero net gain for the observer. This only makes sense if learning happens as a bottom-up process that reorients the agent towards a specific goal, where learning happens purely at the observational level. (Siegel & Allan, 1996). The RW model was originally applied in animal studies, such as those involving Pavlovian learning, and has had a widespread and strong impact on theories of learning overall (Cools, Clark, Owen, & Robbins, 2002; Siegel & Allan; 1996). An important feature of the Rescorla Wagner model is the assumption of constant learning parameters with incremental changes towards the desired behaviour. This effort is made by the agent in order to achieve their goal, which can be defined through a measurable stimulus. This stimulus needs to be salient, and there needs to be an opportunity for the agent to adjust their behaviour. Formally, this can be rendered through a set of equations:

$$\Delta V_x^{n+1} = \alpha\beta(\lambda - V_{tot}) \quad (1)$$

And then

$$V_x^{n+1} = V_x^n + \Delta V_x^{n+1} \quad (2)$$

In equation (1), the change (Δ) represents how well a stimulus (V_x) keeps a predictive value at a future point ($n+1$) and equals the learning rate (α) towards the stimulus multiplied with the salience of the stimulus (β), dependent on the level of surprise. Surprise is defined as the difference between what is happening (λ) minus what is expected (V_{tot}). This formula is essentially a prediction of how the expectancy of a stimulus or experience will be altered by exposure to it. Equation (2) states that future learning (V_x^{n+1}) is the sum of the present understanding (V_x^n) and the altered expectancy (V_x^{n+1}). For the behaviourist, this gives a solid formula for learning, which can predict future behaviour (Boehme et al., 2015). It also circumvents the learner's internal states apart from a learning value (β) that is based on the salience of the stimulus and the amount of influence this value has on the change in behaviour (α). While these parameters represent the learner's internal states, they are fixed throughout the learning experience. This means that the model moves away from the learner as a hypothesis tester that makes predictions, and towards the learner as a passive input receiver for information through errors in predictions and through updates (Danks, 2003; Siegel & Allan, 1996).

Frequentist statistics is a very robust method. With fixed parameters, there is high certainty that the observed measures are due to independent variables rather than to the environment, and it works well with most types of analyses. A particularly well-known analysis is the null-hypothesis-probability test (NHPT), used to determine the likelihood of the data observed given the hypothesis. Interestingly, while this is the most common way to report significant results within psychology and social studies, over 80% of the researchers, including methods instructors answered wrong on at least one assumption regarding what one can and

cannot know about the hypothesis given a frequentist NHPT (Anderson, 2020; Zimprich, 2012). This was taken as further proof of the brain being a Bayesian learner, since participants intuitively consider the likelihood of the hypothesis given the data instead of the likelihood of the data given the hypothesis. In other words, the intuitive choice is based on Bayesian probability estimates and not on parametric frequentist logic, which further supports the idea of the brain as a Bayesian problem solver rather than as a frequentist hypothesis tester (Haller & Krauss, 2002).

While frequentist statistics does not translate well the intuitions of the human mind, it is less computationally demanding than Bayesian statistics. Thus, while both forms of statistics have been around for a long while, the discrepancy in processing demands has made frequentism more viable because it operates with fewer unknown variables. However, more powerful modern computers are increasingly able to handle free variables and a greater workload, creating less of a bottleneck for researchers who want to employ Bayesian statistics (Eddy, 2004). Within the field of learning, this means that new models without rigidly set parameters can be used to model inferences of hidden states in the environment (Behrens 2007; Mathys et al., 2011; 2014). However, other limitations may arise that must be taken into account. Lab settings, for instance, usually fail to factor in the volatility of natural situations, and this could result in bad data unless the research questions and methods are well fitted to the environment (Stauffer, 2007). As briefly mentioned above, the RW model could not account for learning that occurs following a confirmatory prediction that is, for cases where prediction errors do not occur. However, later versions of the models tried to compensate for it by including additional steps and parameters (Danks, 2003). They introduced the idea of equilibrium, such that learning would not simply be a product of behavioural change from prediction error, but a function of an equilibrium between occurring and non-occurring stimuli. The equilibrium would shift, depending on how many targets are present or absent, resulting in

a change in behaviour based on the sum of observed effects. These principles help RW models account for an environment where there is variance and outcome uncertainty (Behrens 2007; Rushworth & Behrens 2008).

Nowadays, precision, variance, and volatility are terms most commonly used in finance analysis, where they indicate the likelihood and consequences of prediction errors during choice and judgement (Fleming, Kirby & Ostdiek, 2003). Yet these terms can be used in much the same way within the field of learning. Precision works as an inversion of variance (precision = $1 - \text{variance}$), which is the squared standard deviation from the mean in a distribution. An example could be a scenario where there are two outcomes, a green or a red light, where variance would be the measure of how often the light is estimated to flash either red or green, and how much it varies. The precision would be the certainty or assumption of how often it will not vary. Volatility is an estimation of how often changes in variance distribution occur. At this point, a hierarchical trend emerges, where estimations of precision are impacted by estimations of volatility, and beliefs about stimulus outcome are weighted by precision estimates (Wacogne, Labyt, van Wassenhove, Bekinschtein, Naccache, & Dehaene, 2011). With the emergence of this hierarchical trend, the RW model hit a wall, as its algorithmic nature could not cope with the changing parameters that represent internal values. In contrast, Bayesian statistics is particularly well-suited for explaining changes in assumptions, as the inference machine grounded in brain processes uses prediction errors to adjust an internal model of the world (Friston, 2009).

Bayesian statistics. The clearest way in which Bayesian statistics differs from frequentist statistics is in the use of the word ‘probability’. While for the frequentist the word refers to the likelihood of groups being different given the underlying hypothesis, Bayesian statistics is based on the probability of the hypothesis being correct given the observed dataset. The formula capturing the Bayesian theorem is the following:

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)} \quad (3)$$

where A and B are events, $P(A|B)$ is the probability of event A given that event B is true, $P(B|A)$ is the probability of event B given that event A is true, and $P(A)$ and $P(B)$ are independent probabilities of A and B respectively. The formula includes prior knowledge of conditions associated with the event (Austin, Brunner & Hux, 2002; Joyce, 2003). This prior knowledge is simply referred to as ‘priors’, representing the posterior assumptions based on previous observation (Behrens, 2007). So, rather than having several independent occurrences of a simple hypothesis that is repeated a number of times, each observation in the dataset changes the underlying hypothesis to better fit the situation observed. As noted by Haller and Krauss (2002), Bayesian probability is a more intuitive way to observe the world. When observing frequentist statements, humans are prone to generate Bayesian conclusions despite lacking any formal training, which is one of the reasons why the Bayesian brain has gotten so much traction.

Attention and learning

The concept of attention covers several task-dependent categories. These include subcategories like focused attention, where heightened attentional resources are used over a longer time, divided attention, where more than one salient target is present, and selective attention, where distractors need to be ignored in favour of the salient stimulus (Cohen, 2014; Dayan, Kakade, & Montague, 2000; Posner & Peterson, 1990).

Attention stands central to most learning theories. In behaviourist models, attention depends on the salience of the stimulus, which in turn depends on the perceived reward (Miller, Barnet & Grahame, 1995; Chance, 2013). In predictive coding models, an attended stimulus generates the strongest prediction errors. However, some researchers argue that less salient updates in the environment might also induce learning effects, implying that information is also

gained from unattended stimulus (Iglesias et al., 2013). In these cases, precision updates are generated to readjust general information about the world rather than to achieve a fixed goal, as a function of the free energy principle. According to Friston (2009), attention is viewed as a synaptic gain control that optimizes the relative precision of top-down priors and bottom-up sensory observations. This is particularly important in hierarchical inference models because it controls the extent to which prior expectations will influence different levels. In other words, attention is a means through which the brain optimizes the precision of generated inferences. Thus, instead of merely selecting sensory channels, attention becomes a property of predictions themselves, such that higher-precision prediction-errors enjoy greater gain in attentive resources (Desimone, 1996; Friston, 2009; Schroeder, Mehta, & Foxe, 2001).

A recent study looked at difficulties in learning by individuals with autism spectrum disorder (ASD) and reported a connection between expectations from the environment and learning rate (Lawson, Mathys & Rees, 2017). The authors used pupillometry and behavioural measures (RT, accuracy) to measure how well those with ASD handled an uncertain environment compared to controls. They found that ASD was correlated with a higher assumption of environmental volatility, which was argued to cause a more challenging learning environment due to improper weighing of prediction errors when generating pWPEs.

Level of attention is a strong factor in determining whether stimuli are registered on a conscious level, and therefore many different tests have explored the phenomenon. In our study, we investigated a well-known phenomenon where one of the target stimuli is not registered, despite the participant being engaged in focused selective attention.

The attentional blink. We explored a common selective-attention effect known as the ‘attentional blink’ (AB), which is a term coined by Shapiro, Arnell and Raymond as late as 1997. While it is hard to determine exactly when the AB was first discovered, one of the earliest reports of the phenomenon was given by Broadbent in 1987 (Broadbent & Broadbent, 1987).

The effect is observed when two target stimuli, T1 and T2, appear one after the other in a rapid serial visual presentation (RSVP) task. If T2 occurs within 250-500 milliseconds after T1, most people are unable to report it (Dux & Marois, 2009). However, if T2 appears before 200 ms, it will usually be detected, barring cases where T1 and T2 are presented at different locations from each other. This phenomenon is known as “lag1 sparing”.

The AB could happen within as well as across perceptual modalities, usually auditory or visual (Arnell & Jolicoeur, 1999). Several theories have been put forward to explain its nature. Thus, early studies concluded that it must result from depleted attentional resources, as the brain is still engaged in processing T1 by the time T2 is presented. Attention-depletion theories include the inhibition/interference models proposed by Raymond, Shapiro, and Arnell (1992; 1994) and the bottleneck models proposed by Chun and Porter (1995). The inhibition and the interference models were cognitive system theories incorporating working memory, where AB was considered the result of optimization of the first stimulus at the expense of the second. In order to ensure correct encoding of T1, a gating mechanism would keep out further salient stimuli until the process is complete. Further, lag1 sparing would be due to the proximity in time of T1 and T2, which would be insufficient for the gate to close.

Well-known bottleneck models are two-stage process models, where a stimulus quickly activates the conceptual representations of information from an RSVP sequence stored in long term memory. The information is volatile however, and prone to overwriting and decay. To prevent stimuli from being overwritten, they are again processed at stage two, which operates on a limited capacity and serves to encode information into working memory. In brief, at stage one, stimuli are chosen based on feature identification, whereas at stage two AB occurs due to capacity limitations. Here again, the sparing effect is assumed to result from the temporal proximity of T1 and T2 (Raymond, Shapiro & Arnell, 1992; Shapiro, Raymond & Arnell, 1994; Chun & Porter, 1995). While a great deal of further formal and informal theories have been

proposed (see Dux & Marois, 2009 for a full review), most of them revolve around the idea of attention as a resource-limited faculty and often refer to the two-systems model of attention first presented by Broadbent in 1958 (Broadbent, 2013) and later expanded on by Chun and Porter (1995). Two-stage models have gained strong support from neuroimaging studies such as the one by Gross and colleagues (2004). This study reported that activity associated to processing both targets had been observed in visual areas, but that activity associated to attention allocation in parietal-frontal regions occurred selectively in response to T2 depending on whether it happened inside or outside the AB. More specifically, the authors investigated within-AB and outside-AB effects on connections in the fronto-parieto-temporal network. Target visual stimuli resulted in strong activation across the network, whereas non-target stimuli resulted in a weaker, more widespread activation in recognition-relevant locations. Outside-AB activity across connections was significantly stronger, but within-AB there was increased sensitivity to distractors that worked as noise masking the T2 (Gross et al., 2004).

Gender differences in attention. Bayliss and colleagues (2005) ran a study comparing male and female participants both on exogenous and on endogenous attention cues. No differences were found for exogenous attention, but there was evidence suggesting that males respond to non-informative cues differently than females do in a Posner task. Thus, females were more sensitive to information from eye-direction as well as from central arrows. These findings were echoed in another study by Feng and colleagues (2011) using a similar methodology. They found differences between genders, where larger ERP component amplitudes were elicited in female participants compared to male participants. The stimuli were pictures of faces with eyes glancing either to the left or to the right, followed by a target stimulus appearing either on the correctly cued side or on the wrongly cued side. The authors argued that the increase in amplitude was due to more intensive use of resources during the task for directing attention to cued information. A third study using the Posner cuing task found similar results (Merritt,

Hirshman, Wharton, Stangl, Devlin, & Lenz, 2007). Thus, male participants were overall less inclined to use cues during the task, which resulted in higher accuracy for female participants during helpful-cue trials, while male participants did better in nonhelpful-cue trials.

In brief, there are many factors, both external and internal, that might impact the level of attention, from distractors and noise to stimulants and wakefulness (Alhola & Polo-Kantola, 2007; Gilbert, Dibb, Plath, & Hiyane, 2000). Our study aimed to examine how meditation impacts attentive resource allocation in a selective attention task.

Meditation and attention

Meditation is a widely used source of relaxation and stress-reduction. Recent studies have found that practicing meditation extensively or for short periods of time can have a strong impact on attention (Braboszcz, Cahn, Balkrishnan, Maturi, Grandchamp & Delorme, 2013), which is likely to be related to changes in the level of attention resources.

Behavioural findings indicate that meditation can increase both selective attention and sustained attention and also has the potential of improving cognitive flexibility in automatic responding and faster reallocation of attentional resources (Valentine & Sweet, 1999; Carter et al., 2005; Cahn & Polich, 2006; Slagter et al., 2007). Most meditation types require attention to shift away from distractors and/or towards a focus, thus drawing on executive functions. This could include monitoring and conflict resolution among thoughts, feelings and mental plans (Braboszcz et al., 2013). In addition, meditation is thought to help the selection of information from a flow of sensory inputs (Hodgins & Adair, 2010) and can lead to long lasting neuronal and attentional plasticity (Colzato, Sellaro, Samara, Baas & Hommel, 2015). More specifically, meditation impacts the way people perceive and process their physical and social environment, and the way they regulate their attention and emotional response (Slagter et al., 2007). In sum, meditation appears to have a robust effect on attention, with the possibility of affecting top-down processing during a learning process, which often underlies cognitive tasks.

Most meditation practices are classified along a continuum, according to the level of attentional engagement required, from mindfulness to focused attention. Depending on where on this continuum a specific meditation practice taps into, it can have different effects on attention tasks. Lippelt and colleagues (2014) provided a succinct review of the literature dedicated to two main meditation types, focused attention meditation (FAM) and open monitoring meditation (OMM), which exert distinct effects on attentional control by differentially affording local versus global processing (Braboszcz, 2013). The FAM technique requires sustained attention on a selected internal or external object of awareness. Isha Yoga is one form of FAM, which requires constant monitoring of attention to maintain focus on a specific place or event, which could be anything from a candle flame to a breathing rhythm, or the feeling of the air moving through the nose (Manna et al., 2010). In a local/global processing task, intensive training in Isha Yoga resulted in practitioners developing a preference for local processing, which confirms earlier results (Van Leeuwen et al. 2012), thus suggesting that Isha Yoga meditation enhances concentration abilities. In contrast, the OMM technique, which relates to mindfulness meditation, Vipassana, and open awareness, involves adopting an attentive and non-elaborative perspective on anything that occurs in the present moment, without favoring future or past events and without focusing on anything specific. It is a meta-cognitive state where a situation is observed without a specific purpose. The goal is to experience the situation as it is, as a form of directing one's focus to their own consciousness, to what, why, and where one is conscious of. FAM and OMM have been associated with different effects on cognition and attention, especially when target stimuli are unexpected. OMM participants had a significantly better performance during a sustained attention task compared to FAM participants (Valentine & Sweet, 1999), but there were no significant differences in performance between the two groups when faced with expected stimulus. Overall, both OMM and FAM participants performed better than non-meditators.

Yet another study found evidence that Zen meditators trained in both FAM and OMM had faster reaction times to all types of stimuli and a decreased bias towards global attention when compared to controls. Contrasting this result, a group of practitioners who primarily had experience in FAM alone showed a stronger local attention bias, which is the more common bias. Thus, practice of OMM was found to reduce this bias to some degree. These effects can be explained by the FAM practice sustaining a bias towards local features and the OMM practice leading to a reduction of both local and global biases (Van Leeuwen et al., 2012).

The two meditation types also have different effects on cognitive-control styles (Lippelt et al., 2014). FAM promotes processing through a single-channel that strengthens top-down support for information relevant to the task and increases competition between information that are relevant and non-relevant on a local level. OMM, in contrast, creates a stronger parallel processing effect, which reduces efficiency of top-down support as well as local competition. In a study by Slagter et al. (2007), AB magnitude was impacted by intensive OMM sessions. In addition, there was a reduced T1-elicited P3b, which is an ERP component that indexes allocation of attentional resources. The authors argued that improved T2 could be a result of improved deployment of T1 attentional resources and that meditation and other forms of mental training impact resource allocation through practice. Thus, meditation may improve upon learning effects such as those involving mental noise reduction and plasticity/adaptation.

Gender differences in meditation. Considering the differences in attention between genders already mentioned, especially with regard to noise suppression and distractor salience in selective attention tasks, we aimed to take a closer look at the impact of gender on meditation effects. As far as we can tell, there has to date been no study assessing gender differences in meditation effects on attention and learning tasks.

One pioneering study by Rojiani and colleagues (2017) compared female and male participants through self-report on mindfulness and self-compassion after a 12-week OMM

course. They found that female participants showed a high decrease in negative affect. Responses were also correlated with improved skills in mindfulness and self-compassion. Contrasting this, male participants showed a non-significant increase in negative affect. The authors concluded that women tend to respond more favourably to guided meditation training. Katz and Toner (2012) conducted a review of articles exploring the impact of meditation on drug and alcohol dependency. Gender differences were mainly absent, but in one of the studies reviewed, females who completed a mindfulness-based meditation intervention showed a temporary decrease in substance use. This contrasted to control males and females, who showed uninterrupted increase in substance use. (Britton et al. 2010). Further, a study by Simpson et al (2007) found that, when looking at inmates with severe PTSD and substance abuse, both genders benefited from a 10-day OMM course and reported reduced substance use after completion, with females more likely to seek out meditation-related help.

Measuring brain activity in our study

We used Electroencephalography (EEG), and in particular event-related potentials (ERP) to measure brain activity. ERPs measure the evoked electrical potential in the extracellular environment around the pyramidal cells of the active cortex. One can read this activity due to dipolar activity generated by the ion channels present in each neuron while they are active (Light et al., 2010). Electrodes placed at specific locations around the scalp can pick up slight changes in electric current, which are sent to an amplifier. The voltage is measured continuously as a single wave representing changes in amplitude. EEG can be used when screening for neurological conditions that affect activation patterns in the brain such as epilepsy and sleep disorders. Furthermore, states of wakefulness (Purves et al., 2008) or other well-defined theoretical topics have also been researched (da Silva, 2013; Engel Jr, 1984; Schaul, 1998).

While it is possible to gain some knowledge about the neural source of activations, by far the best advantage of EEG is its temporal resolution. By allowing us to measure the environment around the neurons and the changes occurring as a result of neuronal activity, EEG can encode time-based data down to milliseconds after stimulus onset, fast outperforming the temporal specificity of functional Magnetic Resonance Imaging fMRI (Näätänen, 1975). By mapping the recorded activity using timestamps relative to stimulus onset, one can gain precise temporal references of neural activity for all stimuli in a sequence. Various software packages can be used to carve out events and calculate average evoked activation for trials across experimental conditions and allow for the identification of ERP components, among which those associated to cognitive and perceptual tasks: the N400, the Mismatch Negativity, and the P300 (Brown & Hagoort, 1993; Polich & Kok, 1995).

ERP components: the P300. In our study, we focused on the P300 component. As the name suggests, it is associated with a positive increase in amplitude near the 300ms mark post stimulus onset, peaking between 300 and 400 ms or beyond, before returning to baseline at around 600 ms. Several studies have associated this component with working memory, labelling it as a cognitive component. It is often elicited during stimulus discrimination, and is assumed to capture basic attention allocation, information processing, and immediate memory (Polich & Kok, 1995). Given that the P300 and the AB both happen at similar times post stimulus onset, the authors hypothesised that AB could be caused either by inhibitory processes such as the closure of the attentional gate or by prolonged inhibition (McArthur, Budd & Michie, 1999). Evidence for extended reaction times and smaller startle reflexes at peak P300 amplitude would suggest that inhibition of neural networks unrelated to processing significant events might be at stake, rather than increased processing. The authors concluded that the AB and the P300 were likely related, possibly as a by-product of wide-spread inhibition of cortical networks, which in turn would signal optimization of local processing.

Measuring learning rates: the HGF family of models

For our AB task, we used EEG as well as predictive models of humans behavior to map subjects' continuous learning rates. Indeed, predictive models are an important part of current psychological research, as they need to be tested empirically and subsequently be either confirmed or falsified. By doing so, they affording a top-down approach, which some argue would be the foundation of modern research (Stanovich & Stanovich, 2013), instead of the usual bottom-up experimental approach.

We contrasted two perceptual models of learning: the Rescorla Wagner (Siegel & Allan 1996) (RW) and the Hierarchical Gaussian Filter (HGF) (Mathys et al., 2011) to track the participants' updates in beliefs regarding the structure of the task in the process of learning the spatial and temporal distribution of the stimuli (i.e., within- versus outside the AB). HGF operates with a perceptual model and a response model. The perceptual model can be based on a generative model of participants' environmental states based on Bayesian principles or on the frequentist RW model. The function of the perceptual model is to describe the values and states that determine observed responses probabilistically, and model the learner's assumptions about the hidden states of the world such as volatility and variance. (Mathys et al., 2011; 2014). It is a reinforcement learning algorithm that tracks participants' understanding and learning of the task's structure, thus explaining how learning occurs and which parameters shape the learning rates. The response model, also known as the 'decision model', describes how data generated by the perceptual model map onto actual responses given by participants on a trial-wise basis. While the perceptual model generates a representation of the internal states of the learning agent, the response model uses this representation to map predicted responses onto actual responses by making assumptions about task noise and other factors impacting the learning rate. An example of a response model is the linear log-RT model developed by Marshall and colleagues (2016). The model uses response times changes as variables when mapping data, starting from

the assumption that increased learning also increases response speed (Marshall et al., 2016). This is actually what happens when high attentional resources need to be deployed to salient targets, thereby impacting response times (Jonides & Mack, 1984). When analysing our behavioral data, we used the log-RT model further modified by Diaconescu for the AB task to map participants' beliefs onto their responses, as well as a version of the Broyden-Fletcher-Goldfarb-Shanno Algorithm (BFGS), a quasi-newton optimization algorithm to smooth the data and generate parameters measuring learning rates and model fitness.

Getting the HGF model to work. Earlier models of Bayesian inferential statistics involve cumbersome numerical calculations, which explains why they are not the alternative of choice for explaining ongoing learning. Moreover, early models are hard to generalize and often require algorithms to be tailored to specific tasks. These arguments plead against the probabilistic learner and against the Bayesian brain in general (Mathys, Daunizeau and Stephan, 2011). In contrast, HGF is well fitted for the task at hand.

Figure 1 below is a schematic representation of how HGF models work. Data generated by known inputs and observed responses return parameters for selected models capturing the agent's assumptions about hidden states. These assumptions are then weighed by the response model, generating precision weighted prediction errors pwPE through the parameters ζ and β . Of the two, ζ is the weighing parameter and determines the degree to which the agent uses their observations when updating their models of the world, whereas β is a parameter that determines how tightly choices are mapped onto beliefs and is usually affected by decision noise. Using these parameters, an agent can generate new posterior belief estimates from an event, which in turn become the new priors for the next observation. In this way, they help shape the internal assumptions about volatility and predict the next prediction of the learning agent (Marshall et al., 2016). An optimization algorithm determines maximum-a-posteriori (MAP) estimates both for the perceptual parameters and for the decision model parameters. The function of the

optimization algorithm is to calculate unnormalized log-posteriors for these parameters as a function of finding the optimum solution or at least a satisfactory solution (for further review of these steps, see Mathys and colleagues, 2011; 2014).

HGF is usually deployed over three levels (X1, X2, X3). X1 is the perceptual level, where the learner observes stimuli across trials or observations. At this point, there is no learning happening, as the subject merely observes the environment. An example would be a light flashing either green or red, or two buttons that one could press in order to receive a reward or a non-reward. In other words, here prediction is either correct or wrong. The second level X2 is the learner's assumption of the result of X1. Over time, the learner will get a probabilistic overview of the likelihood of the light flashing green or red, or of the selected choice yielding reward or not (e.g., 70% chance of green light and 30% chance of red light). The third level, X3 is an estimate of phasic volatility. A pattern emerges while creating estimates of stimulus likelihood, and X2 will start to generate stronger predictions for the outcome of X1. If the environment is stable, it will foster a strong predictive model. If the environment is volatile, it will foster uncertainty, which is explained by X3. If, for example, the probability of getting a red light suddenly changes from 30% to 90%, the distribution of X2 estimates needs to be updated. X3 is the estimate of how likely it is for the environment to change, which affects the rate of X2 switching (Mathys et al., 2011; Mathys, 2014; Marshall et al., 2016).

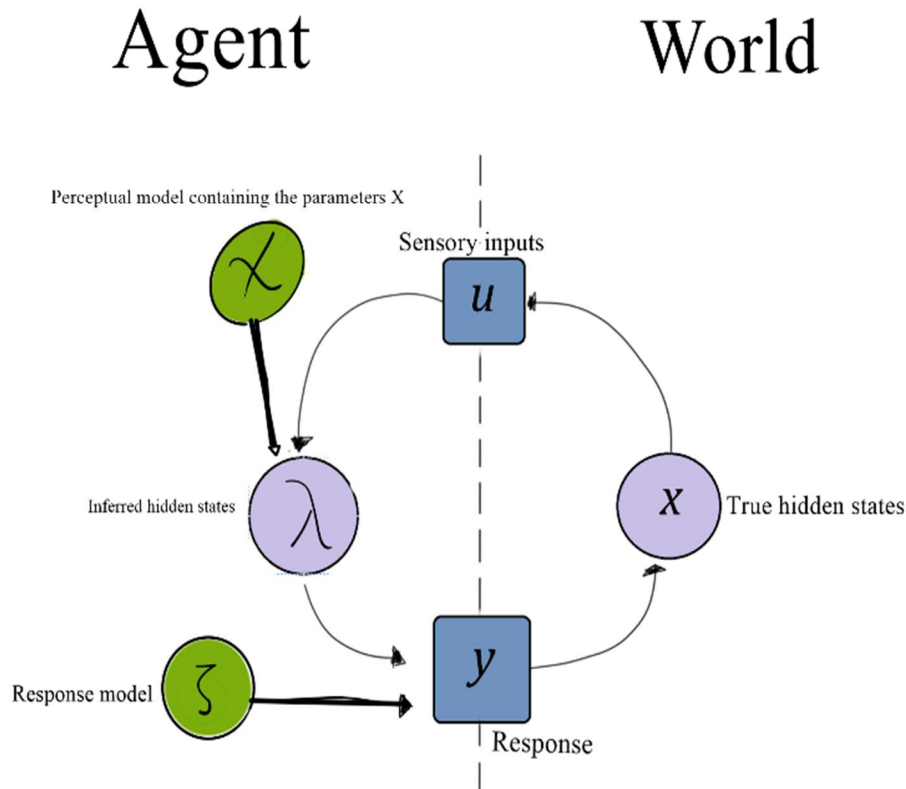


Figure 1. Perceptual model and response model fitted onto data. By knowing inputs (u) and observing responses (y), we can generate parameters for the two models (χ and ζ). We can then use these parameters to make assumptions about the inferred hidden states (λ) via the perceptual model and generate predictions about current/future responses. The true hidden states (x) are not known to the learner and consist of the volatility and variance of the task or environment.

The top-down nature of HGF implies that X_3 impacts X_2 , which in turn generates predictions for X_1 . These predictions are manifested through Gaussian random walks. A random walk is defined as a stochastic process with a succession of steps, usually forming an integer line. Thus, one could start at zero and then either move to +1 or -1 completely at random, and then

take another step, again at random. A Gaussian distribution of these steps determines step sizes. In HGF, random walk step sizes are determined by the parameter ω (for X2), which indicates beliefs about variance independent of volatility, and by the parameter ϑ (for X3), which indicates meta-volatility, or how quickly volatility estimates are updated (Reiter, Diaconescu, Eppinger & Li, 2021). In other words, the generative model is determined at the third level by a random walk that uses a constant step size of ϑ . This affects the X2 level through the coupling parameter k , which together with ω helps determine the step size of the Gaussian random walk. If we were to add further levels to the HGF, the ϑ parameter would be influenced by the step higher up, and so on. (Mathys et al, 2011). Below is an output of the Tapas HGF toolbox describing the three different stages of learning.

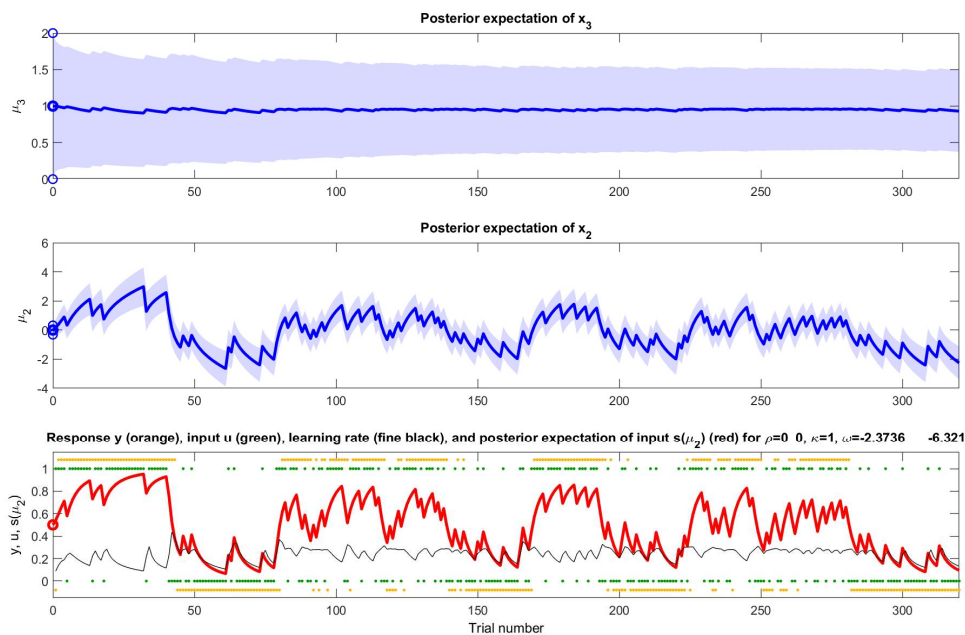


Figure 2. An example of HGF output (demo dataset of the TAPAS toolbox). The lowest level shows inputs received and responses given (y , u , $s(\mu_2)$), overlaid by the distribution of posterior expectations of the agent regarding the outcome. The second level up shows the assumptions of the agent about the variance defined by μ_2 and changing with a step size of ω . The upper level shows how stable these assumptions of variance are and

illustrates the agent's expectancy that sudden changes to the task will happen, defined by μ_3 and θ .

An important feature of the HGF model is its generalizability. Unlike early models of Bayesian learning, it is versatile and can be applied to each learning task due to variable priors for mapping prediction updates. Eliminating the cognitive strain generated by performing a true Bayesian analysis during model update also sits well with the idea of the brain as an energy conserver. The HGF ensures generalizability by using the principle of 'sufficient statistics' (Foley & Marjoram, 2017; Mathys, 2014). Thus, the learner generates close enough approximations of the actual parameters to generate prediction errors, while at the same time keeping them loose enough to be easily computable. An added advantage of being a hierarchical model is that HGF allows for more generalisation and fits many learning paradigms. The model can then also account for individuality between learning agents, hence for different learning rates across participants based on behavioural responses. This makes it a better fit for real-life application, as learning rates may differ greatly across participants (Diaconescu, Litvak, Mathys, Kasper, Friston, & Stephan, 2017; Mathys et al., 2011; 2014).

HGF and the AB task. Our use of the HGF model assumed that variances in our experimental procedure would impact participants as volatility of their environment, and that modelling of the environment would happen for changes in lag during the AB task that is, changes in the interval at which T2 appears after T1, for example one slide apart (lag 1) or three slides apart (lag 3). More specifically, we expected that, based on the principles of free energy and the view of the brain as a probabilistic inference machine (Friston, 2008; 2009; 2010), assumptions would be generated about whether the target would appear early on in a string of stimuli (within-AB), thus having a small chance of being detected, or later on (outside-AB), thus having a higher chance of being detected. These assumptions would, in turn, generate pwPEs that optimize participants' adaptation, allowing them to alter the level of attentional resources they

can dedicate to the task (Friston, 2009). In order to determine how well the two models eHGF and RW performed when predicting behaviour, comparisons were made using the output generated by the optimization algorithm. These values were then run through a Bayesian Model Selection (BMS) module to compare the models against each other. The model selection uses the LME of the candidate models to assesses the relative goodness of fit, as well as to quantify the degree of heterogeneity in the studied sample. (Pitt & Myung, 2002; Rigoux, Stephan, Friston, & Daunizeau, 2014; Stephan, Penny, Daunizeau, Moran, & Friston, 2009; Zeugner, 2010; Vossel et al., 2013).

Given the previously reported different effects of FAM and OMM meditation on top-down processing and attentional control, we further expected them to impact the efficiency of the weighing of precision errors, and thereby the learning outcomes observed through differences in parameter estimates. (Deserno et al., 2020; Hein, de Fockert, & Ruiz, 2021). Further, given that gender differences have been shown to alter attention resources in selective visual attention tasks such as the visual cuing task (Feng et al., 2011), we also expected gender differences in the AB task. As far as we can tell, no study to date has compared genders in an AB RSVP task, despite previous research showing differences in P300 ERP amplitude associated to selective attention (Conroy & Polich, 2007).

Hypotheses of our study

1. Different brief meditation type will create a difference in the ability to efficiently perceive salient stimulus among distractors during the RSVP task.
2. Different brief meditation type will create a difference in learning rate when adapting to a stimulus location placement during the RSVP task through difference in top-down efficiency.
3. We would be able to model this difference in learning rate, and there will be significant differences when comparing parameters for gender and Lag.
4. Based on previous research, the HGF model is likely to have a significant advantage over the RW model when fitting the behavioural data.

Methods

Participants

A total of 38 participants were recruited (22 male) through fliers, social media, and word of mouth. Ages was between 20 and 38 years old ($M = 24.70$, $SD = 2.64$). All participants were right-handed, free of psychological and neurological disabilities, and not taking any medication. Furthermore, participants declared not having practiced for a long period of time any form of meditation previously. Six participants were excluded based on these criteria, for technical reasons, or for providing clear outlier responses. The remaining 32 participants (16 male, $M = 24.39$, $SD = 3.61$) engaged in the meditation session and subsequent AB task.

Ethics approval

This study was approved by the Regional committee for medical and health research ethics (REK-case 60748/2020), and the university of Bergen. Participants signed an informed consent form before they were allowed to start the study, in accordance with the declaration of Helsinki. At the end of the study, participants were given a gift card of 200 NOK to compensate for travel expenses.

Equipment

The equipment used to record the EEG data was a Brainamp DC BrainProducts Amplifier, Powerpack BrainProducts and a 64 Channel Brain products Electrode Input Box. The monitor was a 24 inches ThinkVision with the model number T24-D10. Resolution was 1080p (1920 x 1080). It had a 60 Hz refresh rate with a 32-bit color depth. The hardware used to record the data was a Microsoft Windows XP 2002 stationary computer with an Intel(R) Core (TM)2 CPU 6400 2.13GHz. The software used to create and display the experiment was E-prime 2.0. The software used to record the EEG data was EEG Brainvision Recorder, and the software used to analyse was EEG Brainvision Analyser. EEG comparisons were performed using the EEGlab (v.2021) package for Matlab (Delorme & Makeig, 2004). Data from the statistical models (HGF, RW) were obtained through the open-source resource ‘Translational Algorithms for Psychiatry-Advancing Science’ (TAPAS) toolbox for Matlab (<https://www.tnu.ethz.ch/de/software/tapas>).

Stimuli

For the FAM and OMM meditation sessions, two audio recordings were made by a male native speaker of Norwegian following scripts translated from Dutch to Norwegian (cf. Colzato et al. 2015). The recordings had an average duration of 17 minutes.

For the AB task, we prepared 18 practice trials followed by 288 trials, equally distributed between three conditions: Lag 1 (T2 immediately followed T1), Lag 3 (T2 followed T1 after one distractor), and Lag 8 (T2 followed T1 after 6 distractors). Distractors were letters of the alphabet A to Z, with the exclusion of I, O, Q and S. Targets were single digit numbers from 2 to 9. T1 appeared either as slide 5 or slide 7. The T2 lag components will be referred to as Lag 1, Lag 3 and Lag 8 respectively. T2 appeared among distractors on slide 6, 8, 13 or 15.

Furthermore, at each lag, T2 appeared in the central position, 1.5 degrees to the left, or 1.5 degrees to the right, in equal proportion. Given the 60 cm distance from the screen this

would keep the participants from having to make ocular saccades to observe the stimulus (Bergerbest et al, 2017). The resulting experimental design was 2 (Group: FAM vs. OMM) x 3 (Lag: 1 vs. 3 vs. 8) x 3 (Location: Center vs. Left vs. Right).

Procedure

We conducted this study at the EEG laboratory of the Institute for Biological and Medical Psychology at the University of Bergen. Upon entering the lab, the participants were asked to fill out a consent form detailing the goal of the experiment and their rights as volunteers. They were then directed into the testing area which was a faraday caged room containing the EEG equipment, monitor, and keyboard. A camera was placed in the corner for us to monitor the situation, should anything happen, or should the participant fall asleep.

The preparation of the EEG cap took between 40 and 70 minutes per participant. Our aim was to keep the prep time below 60 minutes. Passive electrodes were applied to a standardized 64-placement cap. A total of 59 electrode were placed on the scalp, in addition to two lateral and two horizontal eye electrodes, one reference on either mastoid, and one ground at the nape of the neck. Scalp electrodes were fastened to the cap, and then treated with Abralyt conductive gel to both increase contact with the scalp, and to scrub away dead cell matter and other layers on the scalp itself. All of this was done to reduce the impedance (how much the electric current is opposed or restricted) as much as possible.

Once all electrodes were at a level below 10 $\mu\Omega$ (micro-ohm), the participant was installed in front of the monitor at a pre-marked spot 60 cm away. Participants were both given on-screen visual instructions and oral instructions from the experiment assistants prior to experiment-onset on how to operate the equipment and how the experiment would proceed. They were also instructed to keep as still as possible and avoid excessive blinking or facial movements during experimental trials. All instructions appeared on the monitor in front of them. Participants were randomly assigned to either the FAM or the OMM group and were asked to

relax and listened to the guided meditation instructions. Following the meditation session, the participants rested for one minute before starting the AB task.

Each trial began with a fixation cross lasting 516 ms, followed by 18 slides each shown for 66 ms, with an inter stimulus interval of 33 ms. The task was to accurately report the observed target after each trial using a numpad in front of them followed by ENTER. If participants were unable to answer, they were instructed to respond by pressing '0' followed by the ENTER key. Participants initiated each trial by pressing the SPACE key, allowing them to rest and prepare before each RSVP sequence. From start to completion the AB task took an average of 30 minutes. For an example of the RSVP task see figure 3. Participants were debriefed after completing the task.




Figure 3 shows a typical trial. Here, the RSVP features T2 at Lag 3.

Analyses

Data modelling. Both RW and the eHGF modelling were performed on Matlab R2020a and the TAPAS toolbox running the enhanced hierarchical Gaussian filter eHGF model. Accuracy data ('1' for hits and '0' for misses) were re-converted into binary values and entered as inputs of the perceptual model. The within- versus outside-AB condition was also converted into binary (0 and 1) values, and was coded as responses in the perceptual model. The eHGF is identical to the HGF in every practical term with the only difference being that it adds an extra parameter $Kappa_1$, which scales the first level with respect to the second. In our study, we kept this parameter fixed to 1, resulting in the classical HGF model. The reason for using the eHGF then is that it can handle more forms of data without running into exceedance errors. The eHGF and RW response models were modified by A. Diaconescu specifically for this experiment. Subsequently, we compared model fit for the two candidate models using the SPM BMS function (Ashburner et al., 2014).

EEG data. The raw data from the EEG were first controlled by eye, and labels were added showing timestamps for both T1 and T2 onset, including lag type and lateralization. The DC offset was then removed, and a Finite Impulse Response (FIR) filter was added for basic cleanup. Then data were again inspected by eye for obvious artefacts and excessive noise. Usually, this would be the point where bad channels were also removed, but there was none apparent in the current study. After initial pruning, the data were decomposed through an Independent Component Analysis (ICA) in order to remap activity onto recurring components. Subsequently, components were evaluated by eye and removed if clearly not brain related (Delorme & Makeig, 2004). The data were then re-referenced to the mastoid electrodes, and epoched following target timestamps. The range was set between -200 and 1000 ms, with a baseline set between -200 and 0. Group comparisons were based on ERPs using both meditation groups and lag (within- and outside-AB).

Results

Accuracy

The analysis yielded a main effect lag, $F(2, 56) = 82.13, p < 0.001, \eta^2 = 0.746$. Bonferroni-adjusted paired comparisons revealed significantly higher scores for lag 8 ($M = 0.904$) compared to both lag 1 ($M = 0.709$) and lag 3 ($M = 0.792$), and also higher scores for lag 3 compared to lag 1, $F(2, 27) = 100.04, p < 0.001, \eta^2 = 0.881$. Response accuracy was higher for participants in the OMM group compared to participants in the FAM group, $F(1, 28) = 4.42, p = 0.045, \eta^2 = 0.136$. However, there was no significant difference in responses between male and female participants, $p = 0.776$.

RT

When comparing reaction time scores for T2 across lags, with Group and Gender as between factors, we obtained no significant results (all p -values > 0.05).

Next, we compared accuracy scores between males and females after averaging across within-AB trials (lags 1 and 3) and contrasting the means with outside-AB trials (lag 8) in a one-way ANOVA for each meditation group. A robust test of equality of means (Welch) revealed higher response accuracy for females than for males in the OMM group only,

$F(1, 8.206) = 6.47, p = 0.034$. No significant differences were found for response time across genders.




Figure 4 shows distribution of behavioural data across meditation, lag and gender for accuracy

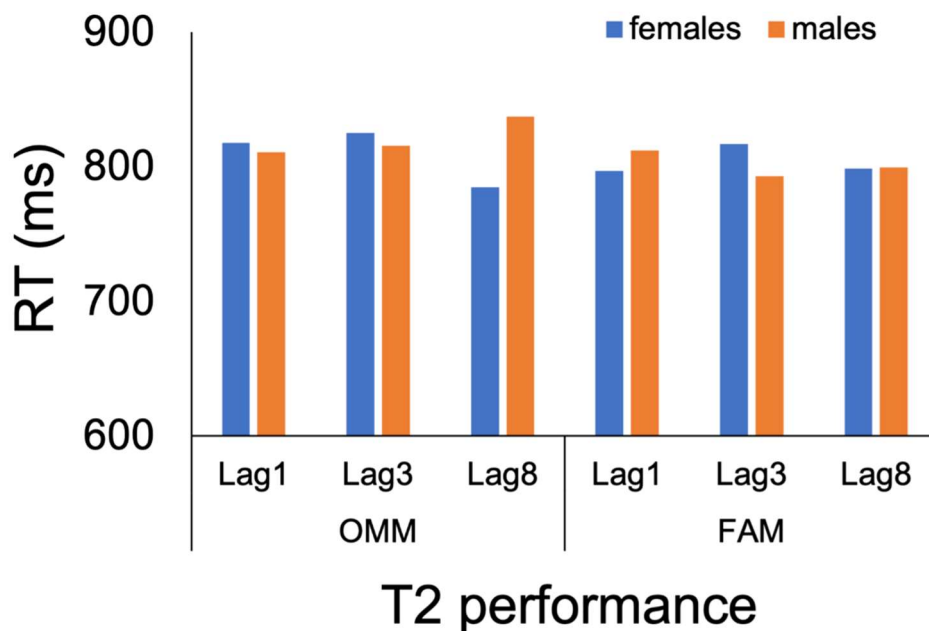


Figure 5 shows distribution of behavioural data across meditation, lag and gender for RT

HFG Modelling Results

Comparison of parameters between meditation types. When Comparing the parameters between the two models we found no significant differences between the Meditation groups when comparing for within- and outside-AB responses on their own ($p > 0.05$) When taking gender into account however there was a significant difference between males and females on the third level ϑ values $t(30) = -2.74, p = 0.01, d = 0.96$. The β_3 parameter was also significantly different, at $t(30) = 2.09, p = 0.045, d = 0.74$. When comparing for the FAM group specifically we found that again there was a significant difference between males and females both for the ϑ level and for the β_3 level. $t(14) = -2.67, p = 0.02, d = 1.34$ and $t(14) = 2.23, p = 0.043, d = 1.12$ respectively.

Figure 6 shows the output from one of the participants. The black line is the overall learning rate, and while this individual has a comparatively high ω_2 level compared to other participants, the projected learning rate is still consistent with the rest in that it shows an effect of learning at first, before evening out as the learning agent becomes increasingly certain about the distribution of responses.




Figure 6. Output generated from the eHGF perceptual model and the LogRT response model. The lowest level (y , u , $s(\mu_2)$) shows inputs received and responses given, overlaid by the distribution of posterior expectations of the agent regarding the outcome. The second level shows the assumptions of the agent about the variance defined by μ_2 and changing with a step size of ω_2 . The upper level shows how stable these assumptions of variance are and illustrates the agent's expectancy of sudden changes in the task, defined by μ_3 and ϑ .

Model fit comparison. The LME for the RW model had a mean of -301.35 (SD : 75.89) across participants, while the eHGF reached a total of -291.45 (SD : 85.5). Using the SPM BMS function, we determined that the HGF with LogRT had an xp advantage of 0.64 to 0.36 over the RW LogRT, indicating that the HGF was a better fit for the data. There was however a Bayesian Omnibus Risk (BOR) factor of 0.81, indicating that there is a considerable risk of confounding factors. The more robust *protected_xp* result yields an xp value of 0.53 versus 0.47, which still indicates an advantage.

Model type	Mean LME (higher is better)	<i>SD:</i>	<i>Xp</i>	<i>Protected_Xp</i>	<i>Bor</i>
eHGF	-291.45	85.5	0.64	0.53	0.81
RW	-301.35	75.89	0.36	0.47	0.81

Table 1: LME and XP values when comparing the eHGF and the RW model.

EEG results

Figure 7 shows significant differences based on permutations using a FDR filter with a threshold level of $p = 0.05$ measured across the entire scalp at all three lags around the P300 timestamp (300-500ms). The spectral map suggests stronger differences at Lag 8 when comparing the FAM male, FAM female, OMM male, and OMM female groups. The effect extends over the frontal, occipital and right-parietal area, suggesting that female participants in the OMM group had the highest neural activity. Figure 7 and 8 further present topographic maps illustrating main effects and marginal means.

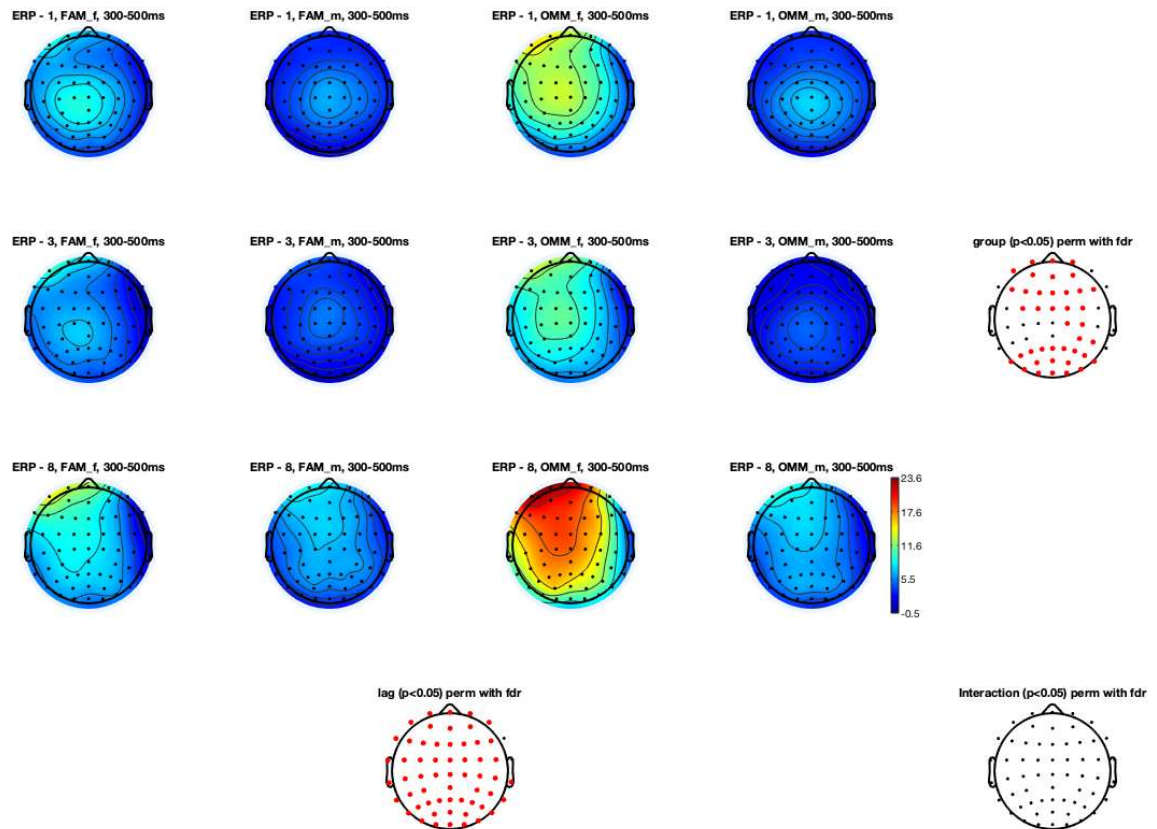


Figure 7. Average topographical map over electrode placements and recorded activity. Red dots mark loci of significant main effects in a 4 x 3 ANOVA with the four groups (FAM males, FAM females, OMM males, and OMM females) crossed with three lags (1, 3, and 8) at $p < 0.05$, FDR corrected. The time range corresponds to the P300 component (300-500 ms after T2 onset).

Figure 8 shows significant differences between genders and lag within the P300 window at Lag 1 and 8. The former implies a left-hemispheric effect, with the spectral map showing increased activity mainly for the female OMM participants. A slightly weaker effect is apparent for female FAM participants. At Lag 8, a pattern of frontoparietal activity emerges that extends to the occipital area and the left temporal lobe. Again, the spectral map confirms that the main effect is driven by the female OMM group, as shown more clearly in the topographic map across

groups: there was significant activation across the entire scalp, with maxima towards the left frontoparietal areas.

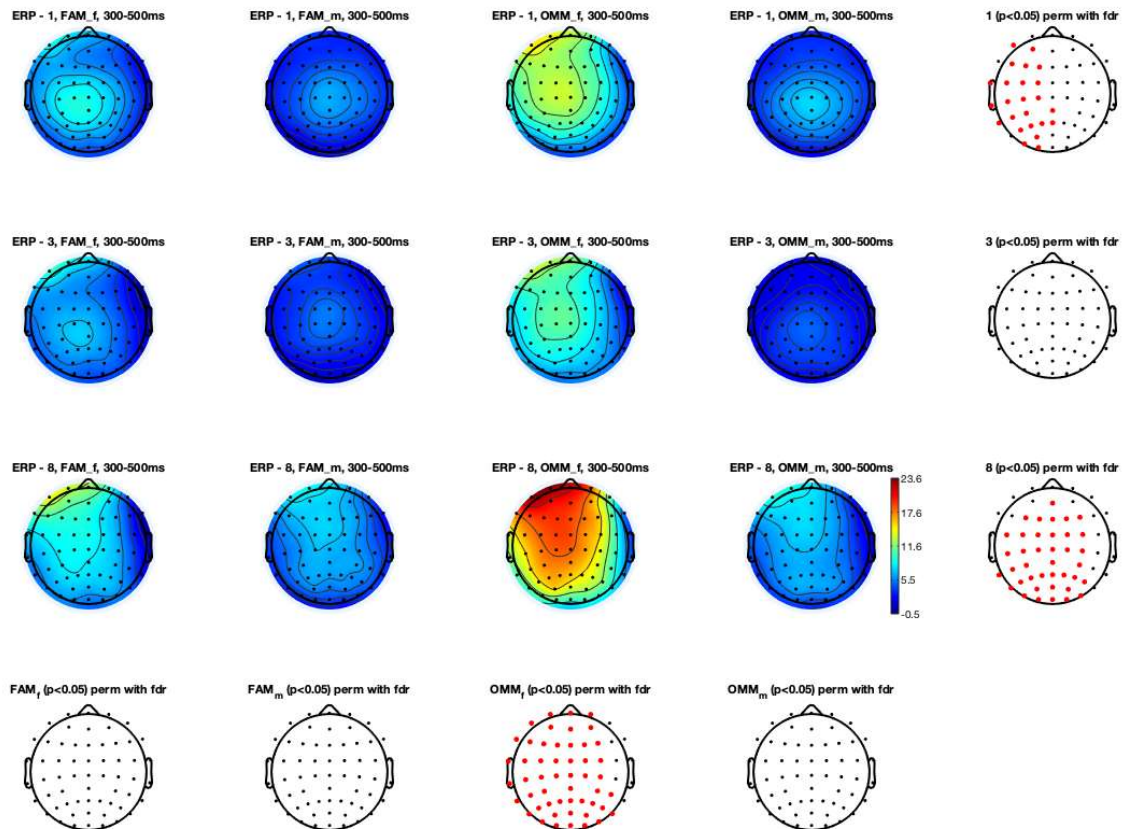


Figure 8. Average topographical map over electrode placements and recorded activity. Red dots mark significant main effects in a 4 x 3 ANOVA with the four groups (FAM males, FAM females, OMM males, and OMM females) crossed with three lags (1, 3, and 8) at $p < 0.05$, FDR corrected. The time range is 300-500 ms after T2 onset.

Figure 9 shows a marked difference in amplitude between lags, highest for the female OMM participants. The effect peaks by the P300 component, with a max amplitude of $12.74\mu\text{V}$ compared to the baseline at 418ms. The average activation of the FAM female group reaches a peak of $8.02\mu\text{V}$ at 440ms. For male participants in the FAM and the OMM groups, peaks have

lower amplitude and occur somewhat later, $5.5\mu\text{V}$ at 475ms . There also appears to be a sustained P3b effect moving up until the 600ms mark.

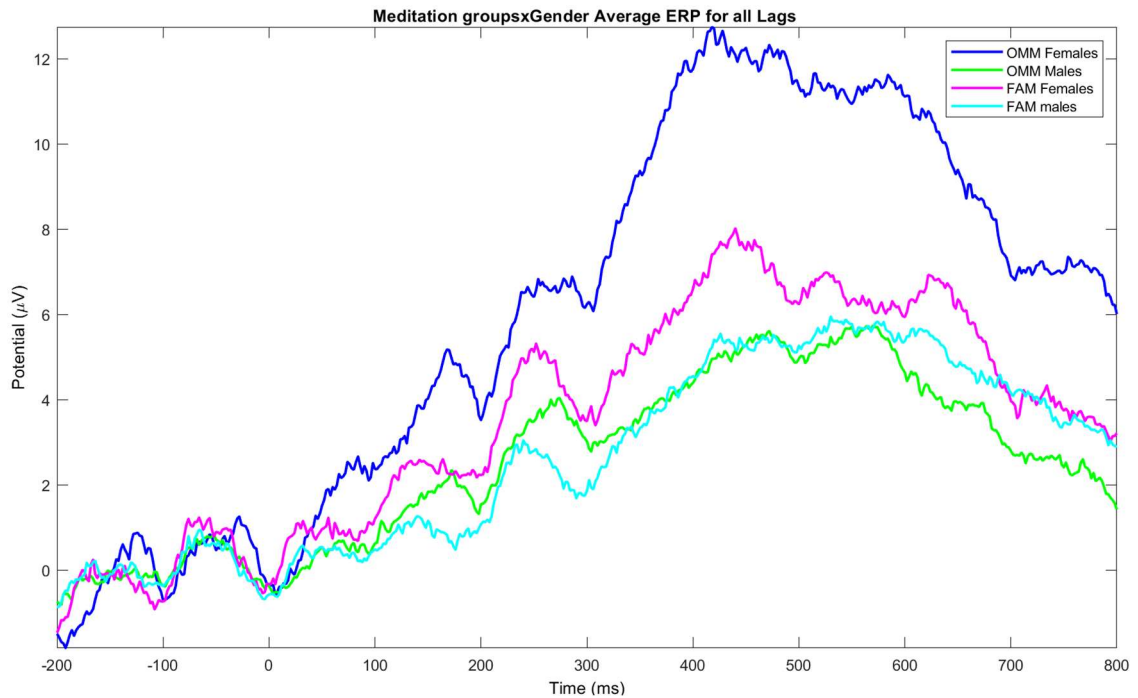


Figure 9. Average T2 amplitude for the four groups (FAM males, FAM females, OMM males, and OMM females) across lags. Amplitude is measured in microvolt, time is measured in milliseconds ranging between -200 ms before T2 onset up til 800 ms after T2 onset.

Discussion

Meditation Impact on AB Accuracy

Previous studies have found various types of meditation to impact attention in a number of ways (Bergerbest et al., 2017; Carter et al., 2005; Colzato et al., 2015; Slagter et al., 2007; Valentine & Sweet, 1999; Van Leeuwen et al., 2012). Importantly, with the exception of Colzato et al. (2015), these studies have explored the behavioral performance of well-versed meditators and have not looked at EEG correlates of performance in brief meditation sessions or attempted to model their effect on the AB. Our study found that participants in the OMM group scored significantly higher than those in the FAM group for T2 accuracy. These results support the findings of a previous study where OMM was shown to increase performance more than FAM in a sustained attention task with unexpected stimuli (Valentine & Sweet, 1999). The results can also be interpreted from the perspective of learning as an update of environmental models of the world using pwPEs, because increased volatility is associated with reduced learning (Behrens, 2007; Friston, 2009). We found that the FAM group also scored higher on estimates about volatility in the experimental environment.

Previous studies comparing genders in a selective attention cuing task also found a significant difference in response times between male and female participants (Merritt et al., 2007). We found no such difference in our AB task, which may be due to non-specific instructions given to participants. Specifically, we did not ask them to respond as fast as possible, which might explain the lack of differences found. However, we allowed participants to take a rest between trials as often as they wished to. This measure, together with the overall fast pace of stimuli presentation, prevented participants from entering responses in a leisurely fashion. Moreover, the fact that RT coupled with accuracy was entered into the logrt model

yielding significant differences between males and females suggests the existence of a trend differentiating participants by gender.

Meditation effects on attention. Previous studies found attentional amplitude difference between genders (Conroy & Polich, 2007; Feng 2011), and so did our study, with female groups showing increased amplitude both within and outside the AB. This effect helped distinguish between meditation types, with brief OMM sessions affording stronger attentional resources than brief FAM sessions. Interestingly, the difference between participants was almost negligible in the male population across all lags, and the effect between meditation groups seems to be driven by females in the OMM group, as shown in average topographical maps. This differentiation between male and female responses was observed by a number of studies with the Posner task (Bayliss et al, 2005; Feng et al., 2011), but in these instances it was argued to be because of responsiveness to explicit cuing, where males were found to be consistently less willing to accept the cued advice, and showed weaker ERPs as a result of less cognitive appraisal during the task. The gender differences in our non-cued task suggests that these differences might move beyond gender specific attention to explicit direction cues and might also extend into more abstract updates about the environment in general through the improved efficiency of model updates.

During our AB task, participants received no explicit information about the structure of the task. Thus, they relied solely on individual learning and internal inferences on the experimental environment. Despite this absence of cues, male participants still had consistently lower ERP amplitudes. Interestingly, female OMM participants scored higher than male OMM and female FAM participants on accuracy during the Lag 8 trial, which was the one least impacted by the AB. An interpretation here could be that when female participants, especially those in the FAM group, were outside of the attentive scope of the AB, were able to deploy higher attentive resource allocation during the RSVP task.

Our results showed a difference in the ability to attend to rapid visual stimuli depending on brief meditation type, with OMM participants scoring better, particularly outside the AB. This in turn would suggest that there is possibility of an endogenous modulating of either the number of attentional resources people are able to distribute, or a way to impact the way we use them.

Meditation type was shown to impact local and global biases (van Leeuwen et al., 2012). These biases can be fitted into a Bayesian brain setting by thinking of them as influences on models for situations or environments. A local bias might encourage the learning agent to put more weight on prediction errors from environmental cues with high assumed salience, even if they might end up being distractors. This would tilt the precision of the pwPEs, and potentially causing an impact on learning rate by obscuring the hidden states of the environment. By the same token, we could explain these results by a more global bias or a reduction of bias in general that might keep the learning agent less exposed to these distractors, and in turn allow them to gain a better model of the actual hidden states. FAM practitioners have been known to be more prone to local bias (van Leeuwen et al., 2012), and females have also been known to put a stronger emphasis on contextual cues (Bayliss et al., 2005; Feng et al., 2011; Meritt et al, 2007), which could result in a higher learning rate for all female participants.

Meditation effects on learning. To our knowledge, there has been no previous experiment that uses the eHGF to model learning rates during the attentional task. The HGF and eHGF has been shown to consistently generate better models for learning than traditional models like the RW (Hein, de Fockert & Ruiz, 2020; Iglesias et al., 2013; Diaconescu et al., 2014), and have been able to generate strong perceptual models based on the theory of the brain working through Bayesian Inference. Commonly these learning paradigms usually have a clear reward focus and instructions that frame the paradigm. The learning agent is made aware of potential volatility in the environment, and they are implicitly incentivised to start generating inferences about the

environment. Iglesias and colleagues (2013) still found that while larger, reward-directed prediction errors was modulated by dopaminergic activity in the basal forebrain, the midbrain was associated with smaller updates based on the environment. These were acetyl-cholinergic and had no distinct reward for the learner beyond the updating of their generative model of the world. On these premises we decided to model a generative model based on the learning factor based on Within- versus Outside-AB observations during the AB task.

The data from female OMM participants supports the assumption that, while there is no reward for the participant during the AB task that would motivate the precision weighing of the prediction errors, there would still be an effect present based on the general updating of environment estimates in order to optimize responses during a performance task. This is further supported by the concept of free energy in the brain, with attention as a predictor of error size. In this scenario, the predictions of within- versus outside- AB would increase the response accuracy depending on the attentional resources allocated, through forming stronger or weaker prediction errors (Friston, 2009). In turn, these prediction errors help generate a more accurate model of the experimental environment, optimizing learning and improving proficiency on the AB-task. The different learning rate estimates for the variation in prediction errors suggest a modest learning rate in our AB task, as indicated by moderate ω values. However, when looking at individual learners, we found a different trend in volatility estimates (ω_3). This was further confirmed when comparing ω_3 values between genders, where a significant effect was found, particularly in the FAM group between male and female participants. Females had higher ω_3 values than males, which suggests a stronger impact from volatility on their estimates. This was further accompanied by significant differences in β levels, which are the parameters representing noise and fitness of a perceptual model as determined by the observational model. In this case, β levels represent how RTs map onto the priors of prediction errors at the next level, with females being significantly more impacted by their estimates about changes in the

environment (Marshall et al., 2016). The effect of a brief OMM session in female participants would thus be to reduce top-down effects on attention (Lippelt et al, 2014). Given the generally high uncertainty and weak guiding cues during an AB task, this could in turn generate less noise, and allow for more accurate predictions regarding the AB-task based on lower-level updates about the environment compared to female FAM participants (Iglesias et al, 2013).

These findings corroborate the results reported for the Posner task, where female participants were more strongly influenced by environmental data (cues). This would imply that the strong effect seen between female and male participants in our task is based on gender-different impact from volatility estimates, where female participants engage in stronger top-down processes of stimulus inferences and trying to generate accurate models with very little information. No significant differences between female FAM participants and female OMM participants were found in ω^3 values.

The RW model revealed no significant differences across groups. The low α levels suggest that the model failed to capture learning during the AB task in our study. Despite low variance estimates, the eHGF outperformed the RW, as indicated by LME comparisons using the SPM BMS function. Better performance by the eHGF model indicates that it could map learning rates better than the RW model, thereby confirming previous findings. (Hein, de Fockert & Ruiz, 2020; Iglesias et al., 2013; Diaconescu et al., 2014).

Using meditation to explore cognitive performance. Meditation studies have mainly focused on well-being benefits such as stress reduction and affective control. However, recent studies have used meditation to investigate the impact it might have on cognition, and in particular on attentive resources. The findings of our study contribute to this line of research. We have attempted a first modelling of meditation effects on the AB phenomenon and we have shown that even brief OMM sessions free up increased attentional resources. The opportunities of exploiting this result are many and varied. For example, educators might consider including

meditation as a welcome addition to teaching environments. Based on our findings, future studies should still look more closely into the relationship between gender and meditation efficiency, as our results could indicate that meditation is not a 'one size fits all' phenomenon. This is further supported by the fact that males seem to show less effect of OMM when it comes to affective regulation and self-compassion (Rojiani et al, 2017).

Meditation and modelling. Generalisability is the degree to which data can be expanded beyond the laboratory and onto more general behaviour. A common way to achieve generalisability is through formulating theoretical models. Thus, abstract predictions can be structured into a system so they can be applied to every situation associated with that particular behaviour or process (Treagust, Chittleborough & Mamiala, 2002). An example of a well-known model is the multiple-component model of working memory first introduced by Baddeley and Hitch in 1974. They conceptualized several parts of working memory into functional components that explained how various senses were combined and incorporated into long term memory from short term memory. The theory has generated a large body of research with basis on these concepts and has been updated several times (Baddeley & Logie 1999). Although models are useful for explaining complex concepts while being easily accessible and providing a framework for what to expect given the underlying hypothesis, they can always be further optimized. Like the working memory model, they are often subject of corrections and revisions (Frigg & Hartmann, 2006). Therefore, it is important when faced with multiple models for the same process or situation, to be able to select the model that best explains a phenomenon, preferably also in the simplest way possible. This implies using a rule of thumb known as Occam's Razor, that achieves parsimony by removing unnecessary assumptions. In our case, this would mean that "A simple model that explains the data should be selected in favour of a complex model that explains the same data" (Lazar, 2010; Balasubramanian, 1997). Nevertheless, more is at stake when selecting models than simply using Occam's Razor. Wears

and Lewis (1999) argued that focusing too much on parsimony could result in reduced accuracy, especially for statistical models. These models are governed by mathematical equations that predict how sample data will be generated based on parameters used to shape and frame the data generated. Within psychology these models can be used to generate generalisable predictions about behaviour.

The behaviourist movement used simple models to explain complex behaviour, and were able to predict outcomes to a certain degree. Parsimonious models like RW should then, from this perspective, always be held in advantage over more complex models like the HGF, for being able to calculate a learning effect using less parameters. The issue is that, while both models explain the same phenomenon, the omission of important parameters and strong assumptions lower model accuracy, thus increasing the likelihood of missing their target phenomenon (Wears & Lewis, 1999). With RW models, one may observe the output and the input, but lacks the necessary tools to explain the internal process generating the result, because these models ignore the agent's input, without the possibility of adjustment, optimization, and adaptation to the environment. In contrast, HGF uses input and response in the same way as the RW does and based on the principle of the brain as a probabilistic inference machine it creates parameters that can be adjusted and optimized. Lacking the added third layer of the HGF, the RW was not able to find clear patterns of learning in the AB, which would support the argument that accuracy and power of explanation should be taken into account when considering the rule of parsimony. We mainly found notable effects in the ω_3 levels that represent internal estimates of the environment, a parameter that is beyond the scope of the RW.

Limitations of the Study

One limitation of the experimental design used in our study is maintaining constant the randomness level of stimuli within and outside the AB. Specifically, in two thirds of the trials, T2 occurred within the AB (lags 1 and 3) and in one third of the trials T2 occurred outside the

AB (lab 8). It would be interesting to observe further differences in learning rates for within versus outside AB targets while modifying the proportion of within and outside AB trials, thereby varying the environmental volatility.

Another limitation of our study is the use of a male voice in the recording of the guided meditation session, which might account for some of the response difference between males and females. The gender of the guide should be taken into account as a potential confound in future meditation studies, which could explore the effects on meditation efficiency by male versus female narrators and how the gender of the narrator affects male and female participants. In our study, we found differences in performance between the FAM and OMM group, despite being recorded by the same speaker, which pleads for there being a well-justified significant effect of meditation type.

Concluding remarks and future studies

We modelled AB-related learning rates using ERP and Bayesian models for the first time. We compared meditation type and gender and found significant differences. The novel results of both gender differences and brief meditation type suggest a great number of ways to move from here. The first venue that should be explored is a variation of this study using a reward-learning task and controlled volatility, allowing us to observe the impact of meditation on higher-level prediction errors and how the altered top-down impact alters the learning rates. Moreover, given the lack of research on gender differences in the AB-task, an experiment controlling for AB between genders like this one while also including a non-meditation setting, checking for significant difference between male and female participants also outside a meditation paradigm. This would further provide a better foundation for the literature on differences in attention between genders in the Posner task.

Reference list

- Alhola, P., & Polo-Kantola, P. (2007). Sleep deprivation: Impact on cognitive performance. *Neuropsychiatric disease and treatment*.
- Anderson, S. F. (2020). Misinterpreting p: The discrepancy between p values and the probability the null hypothesis is true, the influence of multiple testing, and implications for the replication crisis. *Psychological methods*, 25(5), 596.
- Arnell, K. M., & Jolicoeur, P. (1999). The attentional blink across stimulus modalities: Evidence for central processing limitations. *Journal of Experimental Psychology: Human Perception and Performance*, 25(3), 630.
- Ashburner, J., Barnes, G., Chen, C. C., Daunizeau, J., Flandin, G., Friston, K., ... & Penny, W. (2014). SPM12 manual. *Wellcome Trust Centre for Neuroimaging, London, UK, 2464*.
- Austin, P. C., Brunner, L. J., & Hux Md Sm, J. E. (2002). Bayeswatch: an overview of Bayesian statistics. *Journal of evaluation in clinical practice*, 8(2), 277-286.
- Balasubramanian, V. (1997). Statistical inference, Occam's razor, and statistical mechanics on the space of probability distributions. *Neural computation*, 9(2), 349-368.
- Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nature neuroscience*, 10(9), 1214-1221.
- Bergerbest, D., Shilkrot, O., Joseph, M., & Salti, M. (2017). Right visual-field advantage in the attentional blink: Asymmetry in attentional gating across time and space. *Attention, Perception, & Psychophysics*, 79(7), 1979-1992.

- Boehme, R., Deserno, L., Gleich, T., Katthagen, T., Pankow, A., Behr, J., ... & Schlagenhaut, F. (2015). Aberrant salience is related to reduced reinforcement learning signals and elevated dopamine synthesis capacity in healthy adults. *Journal of Neuroscience*, *35*(28), 10103-10111.
- Braboszcz, C., Cahn, B. R., Balakrishnan, B., Maturi, R. K., Grandchamp, R., & Delorme, A. (2013). Plasticity of visual attention in Isha yoga meditation practitioners before and after a 3-month retreat. *Frontiers in psychology*, *4*, 914.
- Broadbent, D. E., & Broadbent, M. H. (1987). From detection to identification: Response to multiple targets in rapid serial visual presentation. *Perception & psychophysics*, *42*(2), 105-113.
- Broadbent, D. E. (2013). *Perception and communication*. Elsevier.
- Brown, C., & Hagoort, P. (1993). The processing nature of the N400: Evidence from masked priming. *Journal of cognitive neuroscience*, *5*(1), 34-44.
- Cahn, B. R., & Polich, J. (2006). Meditation states and traits: EEG, ERP, and neuroimaging studies. *Psychological bulletin*, *132*(2), 180.
- Chance, P. (2013). *Learning and behavior*. Cengage Learning.
- Cohen, R. A. (2014). Models and mechanisms of attention. In *The neuropsychology of attention* (pp. 265-280). Springer, Boston, MA.
- Conroy, M. A., & Polich, J. (2007). Normative variation of P3a and P3b from a large sample: Gender, topography, and response time. *Journal of Psychophysiology*, *21*(1), 22-32.
- Cools, R., Clark, L., Owen, A. M., & Robbins, T. W. (2002). Defining the neural mechanisms of probabilistic reversal learning using event-related functional magnetic resonance imaging. *Journal of Neuroscience*, *22*(11), 4563-4567.

- Colzato, L. S., Sellaro, R., Samara, I., Baas, M., & Hommel, B. (2015). Meditation-induced states predict attentional control over time. *Consciousness and Cognition*, *37*, 57-62.
- Corbetta, M., Miezin, F. M., Shulman, G. L., & Petersen, S. E. (1993). A PET study of visuospatial attention. *Journal of Neuroscience*, *13*(3), 1202-1226.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature reviews neuroscience*, *3*(3), 201-215.
- Danks, D. (2003). Equilibria of the Rescorla–Wagner model. *Journal of Mathematical Psychology*, *47*(2), 109-121.
- da Silva, F. L. (2013). EEG and MEG: relevance to neuroscience. *Neuron*, *80*(5), 1112-1128.
- Dayan, P., Kakade, S., & Montague, P. R. (2000). Learning and selective attention. *Nature neuroscience*, *3*(11), 1218-1223.
- Deserno, L., Boehme, R., Mathys, C., Katthagen, T., Kaminski, J., Stephan, K. E., ... & Schlagenhauf, F. (2020). Volatility estimates increase choice switching and relate to prefrontal activity in schizophrenia. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, *5*(2), 173-183.
- Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of neuroscience methods*, *134*(1), 9-21.
- Dempster, M., & McCorry, N. K. (2009). The role of previous experience and attitudes toward statistics in statistics assessment outcomes among undergraduate psychology students. *Journal of Statistics Education*, *17*(2).
- Desimone, R. (1996). Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Sciences*, *93*(24), 13494-13499.

- Diaconescu, A. O., Litvak, V., Mathys, C., Kasper, L., Friston, K. J., & Stephan, K. E. (2017). A computational hierarchy in human cortex. *arXiv preprint arXiv:1709.02323*.
- Diaconescu, A. O., Mathys, C., Weber, L. A., Daunizeau, J., Kasper, L., Lomakina, E. I., ... & Stephan, K. E. (2014). Inferring on the intentions of others by hierarchical Bayesian learning. *PLoS Comput Biol*, *10*(9), e1003810.
- Engel Jr, J. (1984). A practical guide for routine EEG studies in epilepsy. *Journal of clinical neurophysiology: official publication of the American Electroencephalographic Society*, *1*(2), 109-142.
- Gilbert, D. G., Dibb, W. D., Plath, L. C., & Hiyane, S. G. (2000). Effects of nicotine and caffeine, separately and in combination, on EEG topography, mood, heart rate, cortisol, and vigilance. *Psychophysiology*, *37*(5), 583-595.
- Dux, P. E., & Marois, R. (2009). The attentional blink: A review of data and theory. *Attention, Perception, & Psychophysics*, *71*(8), 1683-1700.
- Eddy, S. R. (2004). What is Bayesian statistics?. *Nature biotechnology*, *22*(9), 1177-1178.
- Foley, B. R., & Marjoram, P. (2017). Sure enough: efficient Bayesian learning and choice. *Animal cognition*, *20*(5), 867-880.
- Feng, Q., Zheng, Y., Zhang, X., Song, Y., Luo, Y. J., Li, Y., & Talhelm, T. (2011). Gender differences in visual reflexive attention shifting: Evidence from an ERP study. *Brain research*, *1401*, 59-65.
- Fleming, J., Kirby, C., & Ostdiek, B. (2003). The economic value of volatility timing using “realized” volatility. *Journal of Financial Economics*, *67*(3), 473-509.
- Frigg, R., & Hartmann, S. (2006). Models in science.

- Friston, K. (2008). Hierarchical models in the brain. *PLoS Comput Biol*, 4(11), e1000211.
- Friston, K. (2009). The free-energy principle: a rough guide to the brain?. *Trends in cognitive sciences*, 13(7), 293-301.
- Friston, K. (2010). The free-energy principle: a unified brain theory?. *Nature reviews neuroscience*, 11(2), 127-138.
- Gross, J., Schmitz, F., Schnitzler, I., Kessler, K., Shapiro, K., Hommel, B., & Schnitzler, A. (2004). Modulation of long-range neural synchrony reflects temporal limitations of visual attention in humans. *Proceedings of the national Academy of Sciences*, 101(35), 13050-13055.
- Haller, H., & Krauss, S. (2002). Misinterpretations of significance: A problem students share with their teachers. *Methods of Psychological Research*, 7(1), 1-20.
- Hein, T. P., de Fockert, J., & Ruiz, M. H. (2021). State anxiety biases estimates of uncertainty and impairs reward learning in volatile environments. *NeuroImage*, 224, 117424.
- Hodgins, H. S., & Adair, K. C. (2010). Attentional processes and meditation. *Consciousness and cognition*, 19(4), 872-878.
- Iglesias, S., Mathys, C., Brodersen, K. H., Kasper, L., Piccirelli, M., den Ouden, H. E., & Stephan, K. E. (2013). Hierarchical prediction errors in midbrain and basal forebrain during sensory learning. *Neuron*, 80(2), 519-530.
- Jones, M., & Love, B. C. (2011). Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and brain sciences*, 34(4), 169.
- Jonides, J., & Mack, R. (1984). On the cost and benefit of cost and benefit. *Psychological Bulletin*, 96(1), 29.

- Joyce, J. (2003). Bayes' theorem.
- Katz, D., & Toner, B. (2013). A systematic review of gender differences in the effectiveness of mindfulness-based treatments for substance use disorders. *Mindfulness*, 4(4), 318-331.
- Kruschke, J. K. (2003). Attention in learning. *Current Directions in Psychological Science*, 12(5), 171-175.
- Lawson, R. P., Mathys, C., & Rees, G. (2017). Adults with autism overestimate the volatility of the sensory environment. *Nature neuroscience*, 20(9), 1293.
- Lazar, N. (2010). Ockham's razor. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(2), 243-246.
- Light, G. A., Williams, L. E., Minow, F., Sprock, J., Rissling, A., Sharp, R., ... & Braff, D. L. (2010). Electroencephalography (EEG) and event-related potentials (ERPs) with human participants. *Current protocols in neuroscience*, 52(1), 6-25.
- Lutz, A., Slagter, H. A., Dunne, J. D., & Davidson, R. J. (2008). Attention regulation and monitoring in meditation. *Trends in cognitive sciences*, 12(4), 163-169.
- Manna, A., Raffone, A., Perrucci, M. G., Nardo, D., Ferretti, A., Tartaro, A., ... & Romani, G. L. (2010). Neural correlates of focused attention and cognitive monitoring in meditation. *Brain research bulletin*, 82(1-2), 46-56.
- Marshall, L., Mathys, C., Ruge, D., De Berker, A. O., Dayan, P., Stephan, K. E., & Bestmann, S. (2016). Pharmacological fingerprints of contextual uncertainty. *PLoS Biology*, 14(11), e1002575.
- Mathys, C., Daunizeau, J., Friston, K. J., & Stephan, K. E. (2011). A Bayesian foundation for individual learning under uncertainty. *Frontiers in human neuroscience*, 5, 39.

- Mathys, C. D., Lomakina, E. I., Daunizeau, J., Iglesias, S., Brodersen, K. H., Friston, K. J., & Stephan, K. E. (2014). Uncertainty in perception and the Hierarchical Gaussian Filter. *Frontiers in human neuroscience*, 8, 825.
- McArthur, G., Budd, T., & Michie, P. (1999). The attentional blink and P300. *Neuroreport*, 10(17), 3691-3695.
- McKiernan, P. (2017). Prospective thinking; scenario planning meets neuroscience. *Technological Forecasting and Social Change*, 124, 66-76.
- Merritt, P., Hirshman, E., Wharton, W., Stangl, B., Devlin, J., & Lenz, A. (2007). Evidence for gender differences in visual selective attention. *Personality and individual differences*, 43(3), 597-609.
- Miller, R. R., Barnet, R. C., & Grahame, N. J. (1995). Assessment of the Rescorla-Wagner model. *Psychological bulletin*, 117(3), 363.
- Näätänen, R. (1975). Selective attention and evoked potentials in humans—A critical review. *Biological Psychology*, 2(4), 237-307.
- Nickerson, R. S. (2000). Null hypothesis significance testing: a review of an old and continuing controversy. *Psychological methods*, 5(2), 241.
- Pigliucci, M., & Boudry, M. (2013). Why the demarcation problem matters.
- Pitt, M. A., & Myung, I. J. (2002). When a good fit can be bad. *Trends in cognitive sciences*, 6(10), 421-425.
- Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual review of neuroscience*, 13(1), 25-42.

- Polich, J., & Kok, A. (1995). Cognitive and biological determinants of P300: an integrative review. *Biological psychology*, *41*(2), 103-146.
- Purves, D., Cabeza, R., Huettel, S. A., LaBar, K. S., Platt, M. L., Woldorff, M. G., & Brannon, E. M. (2008). *Cognitive neuroscience*. Sunderland: Sinauer Associates, Inc.
- Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink?. *Journal of experimental psychology: Human perception and performance*, *18*(3), 849.
- Resnik, D. B. (2000). A pragmatic approach to the demarcation problem. *Studies in History and Philosophy of Science Part A*, *31*(2), 249-267.
- Reiter, A. M., Diaconescu, A. O., Eppinger, B., & Li, S. C. (2021). Human aging alters social inference about others' changing intentions. *Neurobiology of Aging*, *103*, 98-108.
- Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studies—revisited. *Neuroimage*, *84*, 971-985.
- Rojiani, R., Santoyo, J. F., Rahrig, H., Roth, H. D., & Britton, W. B. (2017). Women benefit more than men in response to college-based meditation training. *Frontiers in psychology*, *8*, 551.
- Rushworth, M. F., & Behrens, T. E. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. *Nature neuroscience*, *11*(4), 389-397.
- Schaul, N. (1998). The fundamental neural mechanisms of electroencephalography. *Electroencephalography and clinical Neurophysiology*, *106*(2), 101-107.
- Schultz, W., & Dickinson, A. (2000). Neuronal coding of prediction errors. *Annual review of neuroscience*, *23*(1), 473-500.

- Schroeder, C. E., Mehta, A. D., & Foxe, J. J. (2001). Determinants and mechanisms of attentional modulation of neural processing. *Front Biosci*, 6, D672-D684.
- Seligman, M. E., Railton, P., Baumeister, R. F., & Sripada, C. (2013). Navigating into the future or driven by the past. *Perspectives on psychological science*, 8(2), 119-141.
- Shapiro, K. L., Raymond, J. E., & Arnell, K. M. (1994). Attention to visual pattern information produces the attentional blink in rapid serial visual presentation. *Journal of Experimental psychology: Human perception and performance*, 20(2), 357.
- Siegel, S., & Allan, L. G. (1996). The widespread influence of the Rescorla-Wagner model. *Psychonomic Bulletin & Review*, 3(3), 314-321.
- Simpson, T. L., Kaysen, D., Bowen, S., MacPherson, L. M., Chawla, N., Blume, A., ... & Larimer, M. (2007). PTSD symptoms, substance use, and vipassana meditation among incarcerated individuals. *Journal of Traumatic Stress: Official Publication of The International Society for Traumatic Stress Studies*, 20(3), 239-249.
- Skinner, B. F. (1985). Cognitive science and behaviourism. *British Journal of psychology*, 76(3), 291-301.
- Skinner, B. F. (1987). Whatever happened to psychology as the science of behavior?. *American psychologist*, 42(8), 780.
- Slagter, H. A., Lutz, A., Greischar, L. L., Francis, A. D., Nieuwenhuis, S., Davis, J. M., & Davidson, R. J. (2007). Mental training affects distribution of limited brain resources. *PLoS Biol*, 5(6), e138.
- Stanovich, K. E., & Stanovich, K. E. (2013). *How to think straight about psychology* (p. 256). London: Pearson.

- Stauffer, H. B. (2007). *Contemporary Bayesian and frequentist statistical research methods for natural resource scientists*. John Wiley & Sons.
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *Neuroimage*, *46*(4), 1004-1017.
- Treagust, D. F., Chittleborough, G., & Mamiala, T. L. (2002). Students' understanding of the role of scientific models in learning science. *International journal of science education*, *24*(4), 357-368.
- Valentine, E. R., & Sweet, P. L. (1999). Meditation and attention: A comparison of the effects of concentrative and mindfulness meditation on sustained attention. *Mental Health, Religion & Culture*, *2*(1), 59-70.
- Vossel, S., Mathys, C., Daunizeau, J., Bauer, M., Driver, J., Friston, K. J., & Stephan, K. E. (2014). Spatial attention, precision, and Bayesian inference: a study of saccadic response speed. *Cerebral cortex*, *24*(6), 1436-1450.
- Wacongne, C., Labyt, E., van Wassenhove, V., Bekinschtein, T., Naccache, L., & Dehaene, S. (2011). Evidence for a hierarchy of predictions and prediction errors in human cortex. *Proceedings of the National Academy of Sciences*, *108*(51), 20754-20759.
- Wears, R. L., & Lewis, R. J. (1999). Statistical models and Occam's razor. *Academic Emergency Medicine*, *6*(2), 93-94.
- Zeugner, S. (2011). Bayesian model averaging with BMS. *Tutorial to the R-package BMS 1e30*.
- Zimprich, D. (2012). Attitudes toward statistics among Swiss psychology students. *Swiss Journal of Psychology*, *71*(3), 149.