Proceedings of the Master's Programme Cognitive Neuroscience of the Radboud University

Table of Contents

From the Editors in Chief

Dear Reader,

We are glad to present the first issue of the 16th volume of the CNS Journal.

With everything that has been happening in this crazy 2020, it was not easy to adapt the journal to the new lifestyle the year has brought us. Many changes were made in the process of creating this new issue. However, the product of the work done shows that even in these difficult times, gems can still be made. In fact, if we had to summarise with a few words the work that was done throughout these last months, they would surely be: *per aspera ad astra.*

We want to highlight the incredible work and effort of each CNS Journal member. Despite all the circumstances, the CNS Journal team continued their strong work. We, as Editors in Chief, are very proud of this! Therefore, we want to thank each single team: editors, subeditors, layout, webmasters, and public relations. To all of you: thank you! Without you all, this would not have been possible.

We hope you truly enjoy this new piece of the CNS Journal adventure.

Furthermore, we want to spend a moment remembering a member of this journal that sadly will not be able to see and enjoy this new edition. Vaibhav Arya was a brilliant CNS student and a fundamental member of the CNS Journal. First, joining the subediting team (junior and senior) and then becoming an Editor in Chief last year, he embodied all the values that this student-led project incorporates: in-depth appraisal and assessment, creativity, innovation, mentoring, and teamwork. He was one of us. He is and will be deeply missed. Thank you Vab, you were one of the best people we will ever meet. Rest in peace.

Nijmegen, January 2021

Gregorio Borghi and Annika Mordelt

Editors in Chief

From Dr. Saskia Haegens

It is my pleasure to welcome you to the first issue of the 16th volume of the Nijmegen CNS journal. This is, in fact, my second editorial for this journal — back when I was a student myself, I was the editor-in-chief for one of the first editions — so on a personal level this is somewhat of a coming full circle. It is very rewarding to see the journal still going strong after all this time, and to see how the CNS program has grown over the years. Writing this editorial feels even more special now that my own students are publishing in this very journal!

The research published in this issue was conducted in a highly unusual academic year — one during which our lives were turned upside down, seemingly overnight, by a global pandemic. Labs, offices and the entire university closed down. Ongoing experiments had to be abandoned, and virtually all labs had to pivot to working from home. To say this was disruptive to the thesis projects presented here would be an understatement.

It really speaks to the resilience of these students who, despite such extreme circumstances, still managed to produce the excellent work published here. Between the uncertainty, fear, stress, anxiety, the social isolation and worry about the future — especially during such a critical period of their (academic) careers — they managed to pull through. They showed up to zoom lab meetings, ran remote matlab sessions, did online experiments, and made the best of it.

Having life as we knew it come to a sudden halt also presented us with an opportunity for reflection. While fairly disruptive to our work, I want to acknowledge that most of us in academia were pretty lucky in that we can actually relatively easily work or study from home, at least for a good while, without serious consequences. Most of us had the flexibility to switch to data analysis, to writing papers, all from the safety of our homes.
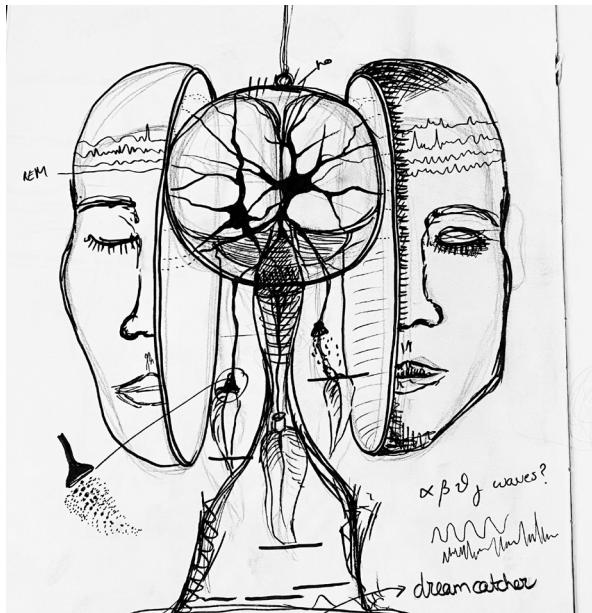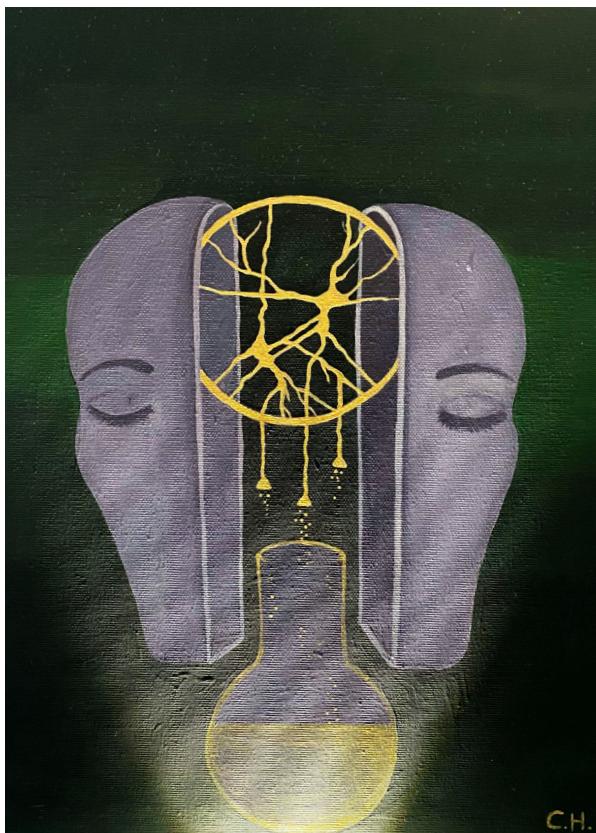
The students I had the honour of working with this past year, all had very different reactions to these highly stressful times, and different coping mechanisms. For some, focusing on their research was a positive distraction, while for others, slowing down and exploring new hobbies and routines was a better way to deal with the situation. Interestingly — and despite my telling them our research is not that important, in the grand scheme of things, and not to worry too much about it — they all finished their projects and produced impressive work. And, perhaps the part I valued most, they managed to support each other in the process. Indeed, resilience takes many forms, and this current edition of the Nijmegen CNS journal is testament to that.

Saskia Haegens

About the Cover

Over the past few decades, the preeminent importance of sleep for people's physiological well-being has become increasingly recognized among scientists. Studies have repeatedly provided strong evidence that dreaming helps people improve their waking life. However, there are dreams that are experienced very similarly as in waking life; lucid dreams. A study in 1985 by Stephen LaBerge at Standford University revealed that, unlike in most dreams, time perception in lucid dreaming is about the same as in waking life. Regardless of how they work and whether they are truly 'dreams', people who experience lucid dreams are able to observe their dreams, recall the waking world, and sometimes control the direction of the dream.

The cover is inspired by the power of lucid dreaming and aims to represent a dreamer who is aware of dreaming while asleep and is sorting the caught dreams using the dreamcatcher. I decided to title the cover "The neuronal dreamcatcher", or simply "The dreamer". Adverse to a brain made out of neurons or a dreamcatcher which is a hoop on which is woven a net or a web, you can observe a person who combines both these qualities in a metaphysical art style: dreamcatcher made out of neurons.

The 'neuronal dreamcatcher' is catching dreams during the night and separating the good dreams from nightmares. Traditionally dreamcatchers were used to hang over a cradle as protection: to encourage good dreams. Despite the dreamcatcher depicting both good and bad dreams, the feathers almost always direct good dreams to the sleeping owner of the dreamcatcher. The feathers in this case are symbolized by synaptic endings which are sending information out. As in dreams mostly everything is imaginable, for example here the synaptic endings are releasing stars instead of neurotransmitters. Stars and constellations are covering the background of this cover. Good memories and dreams are represented by the glowing stars, while nightmares are represented by the black holes and the black background.

The woven net or web, made of neurons, around the hoop is inspired by the beautiful neuron images Camillo Golgi obtained around 1870-1880 with his staining method and by the images the neuroscientist Greg Dunn made.

Carmen Heuvelmans
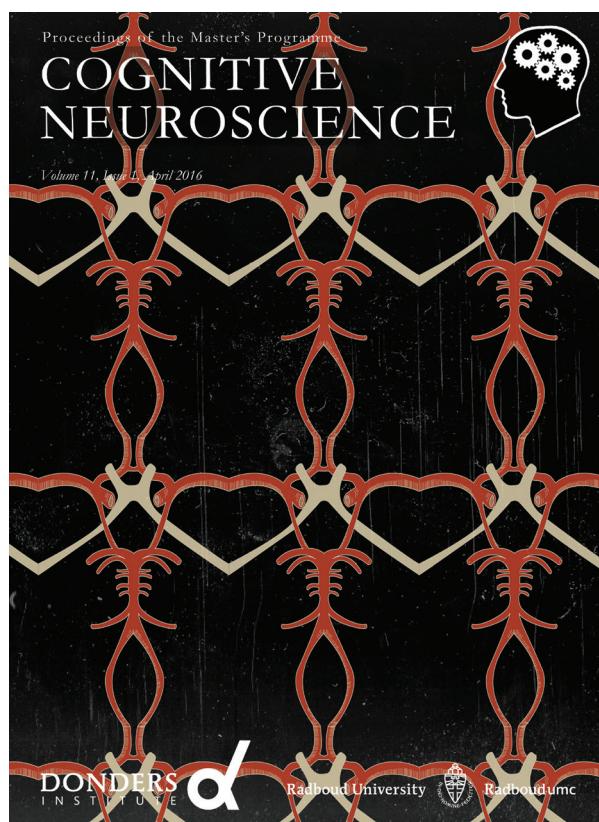
Meet the Journal Team: Layout

What do you take with you from your experience as a part of the CNS Journal team?

The power of collaboration. Everyone in our team has their particular interests and strengths. We all work on a specific part of the journal, and this works together perfectly. Releasing a student journal twice a year only works if you do it together and allow people to work on where they are drawn to.
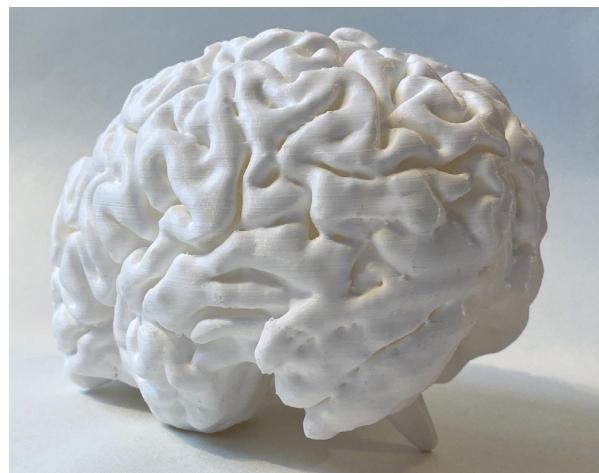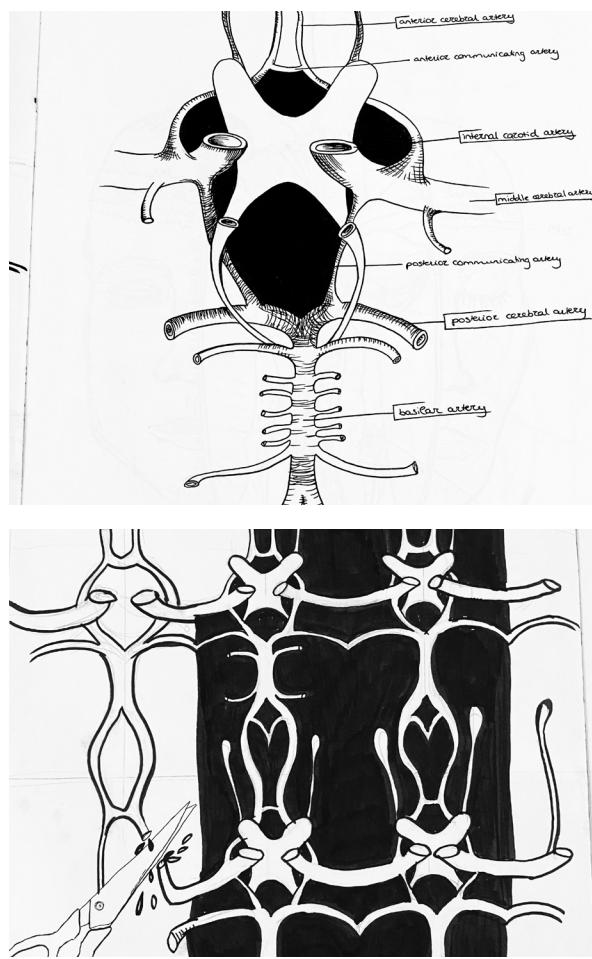
What would be your advice for students that want to join the layout team for later issues?

Dare. Don't be afraid to come up with your own ideas, paintings or sketches. When I saw that in previous versions students used pictures from the web, I was insecure if it was actually a good idea to bring my own drawings. I tried it the first two times, but none of my covers were selected. I was hesitant to try a third time, however this was my lucky time. The journal team voted both of my covers with a slight majority for The dreamer instead of the intriguing pattern of the circle of Willis.

Why in particular did you choose to work on the layout?

I believed it was one of the few possibilities to bridge my creative interests with my passion for neuroscience. Science and art are often considered separate fields. I have been able this year to combine some 3D printing knowledge with real-scale anatomical brain scans of my colleagues but nothing more. I remember seeing a cover of the CNS Journal with a sketch of Leonardo da Vinci and immediately I thought like I can combine my passion for art and science by designing intriguing covers with patterns of neurons or 3D printed brain parts.

Carmen Heuvelmans

Measuring Real-World Head Orientation Priors With Naturalistic Motion Tracking in a Bayesian Multisensory Integration Framework

Lucas Billen[1]
Supervisor: Pieter Medendorp[1]

[1]Donders Centre for Cognition, Nijmegen, The Netherlands

Previous research has shown that the brain integrates multisensory information via Bayesian inference to achieve spatial orientation. The key feature of this approach is that, in addition to vestibular, visual, and somatosensory information, prior knowledge is incorporated in the sensory integration. The effect of such a prior is twofold: Near upright, it improves precision, but with increasing head-tilt, systematic errors in the final head-in-space estimate are induced. The prior is assumed to be based on lifelong experiences, represented as a Gaussian distribution centered on upright. Whether this accurately represents the underlying head-in-space prior was unknown. Here, we used motion tracking to kinematically measure the head orientation distributions of six participants performing naturalistic activities. We investigated whether 1) the resulting head orientation distributions can accurately represent the underlying head-in-space prior and 2) whether performances on tasks of perceived visual verticality (SVV tasks) can be simulated by incorporating the obtained real-world prior into a previously developed multisensory integration model. In line with previous research, we expected the naturalistic head orientation distributions to be best described by Gaussian distributions, accurately simulating SVV task performance. Results showed that head orientation distributions were, in fact, best fitted by t Location-Scale distributions, characterized by fatter tails compared to Gaussian distributions. Simulation of SVV task performance was not in line with previous research regarding both magnitude and direction of the biases. Thus, using a novel motion tracking approach, we provide evidence that the underlying head-in-space prior deviates considerably from normality. Future research should focus on successfully incorporating such a prior in the Bayesian multisensory integration model.

Keywords: Bayesian inference, multisensory integration, prior, motion tracking, spatial orientation, subjective visual vertical

Disclaimer: Due to regulations concerning the Covid-19 pandemic, it was not possible to collect novel data for this study, and the study had to be re-steered. This involved a novel analysis of a previously collected pilot data set of six subjects in the sensorimotor lab. Because of the limited sample size, statistical analysis lacked sufficient power.

Corresponding author: Lucas Billen; E-mail: lucas.billen@gmx.de

Spatial orientation, which is our sense of body orientation and self-motion relative to the environment, is fundamental to numerous basic motor actions, such as balance, locomotion, and the interaction with objects in the environment (MacNeilage et al., 2008). Not being able to tell how we are oriented in space could be fatal in many situations. For example, if a diver or a pilot loses their sense of spatial orientation, this could have severe consequences. However, in most situations, the brain seems to be able to achieve spatial orientation effortlessly. Given that the information from the various sensory organs is inherently noisy and partly ambiguous, how is the brain able to accurately reconstruct the state of the world and the state of the body, such that errors remain minimal?

Research has shown that the statistically optimal way of dealing with the inherent noisiness of the sensory information is to use several sources of information. The multisensory information is then integrated while their respective noisiness is taken into account in a Bayesian inference fashion (Clemens et al., 2011; De Vrijer et al., 2008; Körding & Wolpert, 2004; MacNeilage et al., 2007). According to the *Bayesian multisensory integration model* developed by Clemens and colleagues (2011) and later adapted by Alberts and colleagues (2016), the brain uses three sources of information to achieve an estimate of how the head is oriented in space (Figure 1A). Firstly, as part of the vestibular system, the otoliths directly measure acceleration of the head, and therefore provide a direct signal of how the head is oriented in space. Secondly, information from body sensors providing an estimate of the orientation of the body in space can be combined with the information provided by the neck sensors, which measure the angle between head and body. This yields a second, albeit indirect, measure of how the head is oriented in space. Thirdly, this model allows for the brain to useprior knowledge to estimate the current orientation of the head in space. This so-called *prior* is assumed to be a representation of life-long experiences of how the head is typically oriented in space (Clemens et al., 2011).

The different noisy signals are usually represented as *Gaussian probability distributions*. It is assumed that the sensory signals (i.e. information from the otoliths, and the transformed signals from the body sensors and the neck proprioceptive information) are calibrated unbiasedly, meaning that they are centered on the true head-tilt angle. However, because these signals are inherently noisy, the resulting uncertainty about the sensory information is represented in the width of the Gaussian distribution. The noisier

the signal, the higher the uncertainty and therefore the broader the distribution. In terms of Bayesian inference, this means that during the integration of the sensory information, highly noisy signals are weighted less, while less noisy and therefore more reliable signals are weighted more heavily. The model assumes the noise levels of the body sensors and the neck sensors to be constant, while the noise of the sensory information coming from the otoliths increases rectilinearly with increasing head-tilt (Clemens et al., 2011; De Vrijer et al., 2008). The benefit of this Bayesian inference approach is that all of the available information is used optimally, such that the resulting head-in-space estimate has a lower degree of perceptual uncertainty than can be derived from the individual sources.

In contrast to the sensory signals, the prior is assumed to be centered around a head-tilt of zero degrees (i.e. upright), because the most likely head orientation during everyday life is assumed to be upright, too. The effect of such a prior is twofold: At small head-tilt angles, it improves precision, because it further reduces the uncertainty of the final estimate. However, the prior also induces a bias in the final estimate of head-in-space orientation, which becomes increasingly more pronounced at larger head-tilt angles (see Figure 1B). Thus, at large head-tilt angles, the prior biases the final estimate towards zero and away from the true head-tilt angle, resulting in an underestimation of one's actual head-tilt (so-called Aubert effect; Aubert, 1861; Mittelstaedt, 1983; Van Beuzekom & Van Gisbergen, 2000) (for a complete description of the model, see Methods).

These large systematic errors in one's perception of head orientation can behaviorally be measured with the so-called Subjective Visual Vertical task (SVV), hence providing a methodological approach to indirectly study the underlying multisensory integration processes taking place (Aubert, 1861; Barra et al., 2010; Ceyte et al., 2009; De Vrijer et al., 2009; Eggert, 1998; for a review on the perception of verticality, see Dieterich & Brandt, 2019). The SVV task is conducted in the dark to minimize visual influences. During the task, participants' bodies are roll-tilted while they are sitting in a vestibular chair. They are then presented with luminous bars with varying angles relative to true vertical. The participants' task is to judge whether the bar is rotated clockwise or counterclockwise compared to their perceived gravitational vertical. As was mentioned above, participants are quite accurate at this task at small head-tilt angles, but as the head-tilt increases (in some experiments up to 120°), the systematic error increases to up to
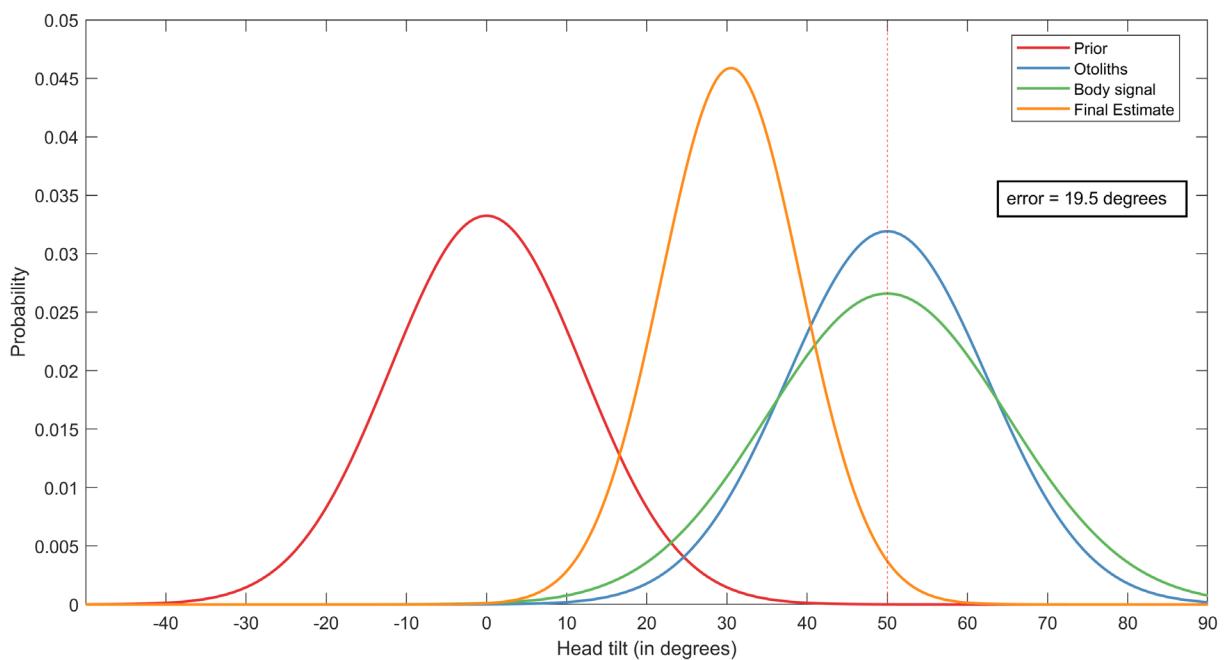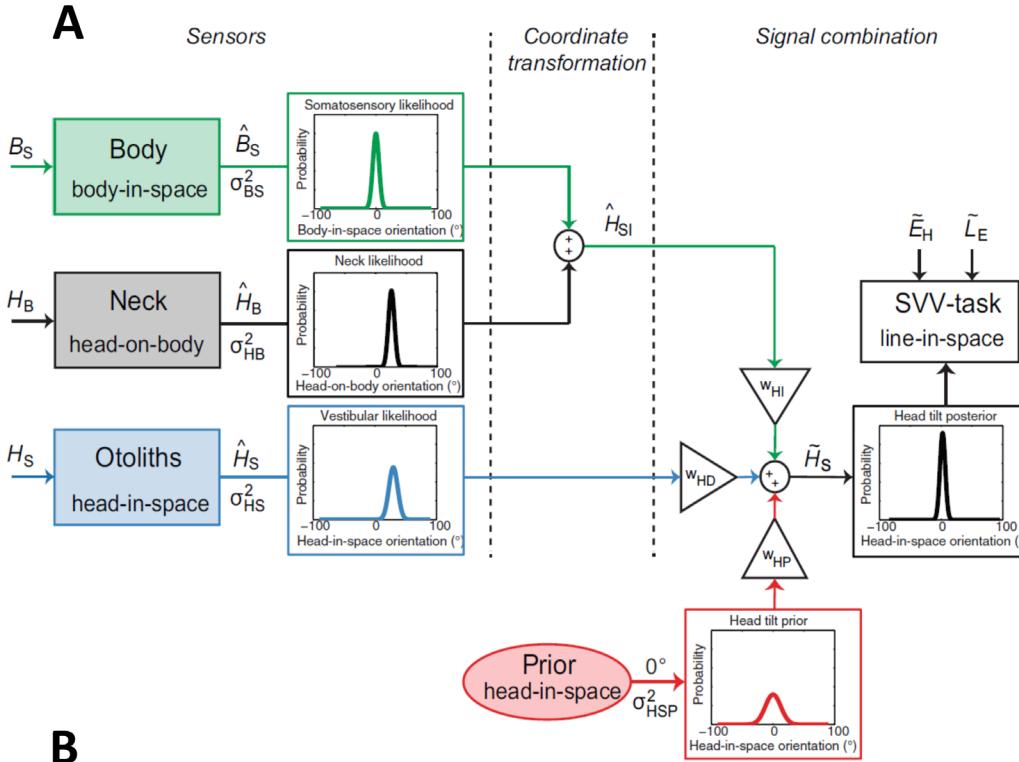
Figure 1. Bayesian multisensory integration. **A.** Bayesian multisensory integration model: The brain can use information from body sensors, neck sensors and the otoliths, represented as Gaussian probability distributions. The otoliths provide a direct measure of head-in-space orientation, while the information from the body and the neck sensors can be combined in the coordinate transformation stage to provide an additional indirect measure. Furthermore, a head-in-space prior that is centred on 0° head-tilt is assumed to be part of the signal combination stage, resulting in a final head-in-space estimate (adapted from Clemens et al. (2011) and Alberts et al. (2016)). **B.** Example of multisensory integration: The probability of the sensory estimate of head orientation in space (otolith and body sensors) can be represented as Gaussians centred on the true tilt angle (i.e. 50°) and corrupted by noise, represented by the width of the Gaussian. The Gaussian prior is centred on upright (i.e. 0°). Thus, the final estimate which is given by the optimal integration of the sensory information and the prior, will be biased toward the prior, but with smaller uncertainty with respect to the individual sources.

60° (which means that the luminous bar has to be tilted 60° to be perceived as vertical), indicating a strong underestimation of one's own head-tilt. The aforementioned multisensory integration can explain the observed behavior on this task adequately, even though performances on the SVV task can differ substantially between individuals (Clemens et al., 2011). The large systematic error can be accurately explained by the prior introducing a bias towards upright.

Even though the model provides a good fit to the data and is intuitively appealing, there is still uncertainty about the true underlying nature of the distributions of the various signals. For example, it is assumed that the prior is based on lifelong experiences of how the head is typically oriented in space. Given the observation that the head's vertical axis is usually aligned with gravity, it is justified to assume that the prior is centered on a head-tilt of 0°. However, in previous research, the prior distribution was always assumed to be Gaussian for reasons of computational convenience and/or simplicity. On the one hand, this makes the modelling easier and more intuitive. However, having a fixed distribution type makes it impossible to determine whether the model reflects the true underlying nature of the prior distribution (Stocker & Simoncelli, 2006). In other words, whether a Gaussian distribution reflects the true underlying nature of the prior, and whether potential differences in the underlying prior can explain the individual differences on SVV task performance, were not addressed in previous research regarding head orientation in space.

The present study aimed to tackle these questions. Because it is assumed that the prior is based on lifelong experiences, *naturalistic motion tracking* might be a viable approach to measure the underlying prior. Thus, we used motion tracking to measure kinematic head movements during typical naturalistic activities. Subsequently, we investigated what type of distribution fits the measured data best in an attempt to test the basic assumption that the prior is of Gaussian nature. We then used the best fitting distributions as representations of the underlying prior, by implementing them in the Bayesian sensory integration model of Clemens et al. (2011). We were able to simulate what the SVV task performance would look like in the same participants, providing us with a novel approach to gain insights into the true underlying distribution of the head-in-space prior. Thus, the current study combined a naturalistic motion tracking approach with a controlled lab-based task of perceived verticality to get a more realistic and complete view

of what the underlying head-in-space prior might look like, and how it can bias perception on tasks of perceived verticality.

Not many studies have used motion tracking to investigate head movements during everyday life activities. Carriot and colleagues were the first to study the natural vestibular inputs that the brain needs to process during naturalistic activities (Carriot et al., 2014). Using a micro-electromechanical systems module, which combines three linear accelerometers and three gyroscopes, they measured participant's head movements during several active and passive movements, such as walking, running, jumping, and riding on a bus. Interestingly, they showed that the probability distributions of the angular velocities that the vestibular system experiences in everyday life also deviate significantly from normality. Instead, the probability distributions were characterized by large excess kurtoses (i.e. fatter tails). Using a similar methodological approach, while focusing more on the head-in-space orientation might therefore be a good approach to quantify naturalistic head orientations. The general study design of Carriot et al.'s study therefore serves as a good basis for the present study. Thus, similar to the study by Carriot and colleagues, participants in this study performed five naturalistic activities: walking, running, going up and down the stairs, standing and sitting. These tasks cover a wide range of activities that predominantly occur in everyday life, therefore providing a relatively realistic representation of activities that the underlying head-in-space might be based on.

In line with the Bayesian sensory integration model developed by Clemens and colleagues (2011), we expected to show that the resulting naturalistic head-in-space distributions measured by the motion trackers will be best captured by Gaussian distributions centered on 0° degrees head-tilt. Furthermore, we expected that upon implementation of those naturalistic priors into the model, a simulation of SVV task performances would closely follow the actual performances observed in past studies.

Materials and Methods

Participants

Six healthy subjects participated in the study (three male and three female). Ages ranged from 23 to 28 yrs ($M = 25.5$; $SD = 1.64$ yrs). They were free of any known neurological or movement disordershad normal or corrected-to-normal vision

and were personally recruited by the experimenter. All participants took part voluntarily and gave written consent after they were informed about the experimental procedure.

Naturalistic Motion Tracking

Experimental setup. To acquire naturalistic motion kinematics, the MVN motion capture suit from Xsens was used (Xsens, 2017). This system consists of 17 sensors for full body motion tracking. For the purpose of this project, only 11 sensors were used to measure upper body and head kinematics. Sensors were placed on the pelvis, shoulders, sternum, upper arms, forearms, hands, and head. The sensors on the torso were attached with a tight-fitting vest, the sensor on the head was attached with a headband and the sensors on the arms were attached with Velcro-straps. This equipment was provided by Xsens and therefore was designed to securely contain the sensors with integrated Velcro pockets. Each sensor is a compact Inertial-Magnetic Measurement Unit (IMMU) (47mm ×30 mm × 13mm, weight: 16 g), containing inertial sensor components, including a 3D rate gyroscope measuring angular velocities and a 3D accelerometer measuring acceleration. Additionally, it comprises a 3D magnetometer, a barometer, and a thermometer. Combined with the internal Xsens signal processing algorithms, 3D drift-free orientation data is provided. The sensors are wirelessly connected to the Awinda Station, which serves as the interface between the laptop running the Xsens-based software (MVN Analyze/Animate) and the IMMUs.

The integrated MVN Fusion Engine calculates the position, orientation, and numerous other kinematic measures of each body segment with respect to an earth-fixed reference coordinate system. By default, the earth-fixed reference is defined as a right-handed Cartesian coordinate system with X being positive when pointing to the local magnetic North, Y being oriented according to the right-handed coordinates (pointing West), and Z being positive when pointing up (Figure 2A).

For each body segment, all kinematic quantities are expressed in a common, local coordinate frame, L, which is also a right-handed coordinate system with X being positive when pointing forward, lying in the horizontal plane, Y being oriented according to the right-handed coordinate system with respect
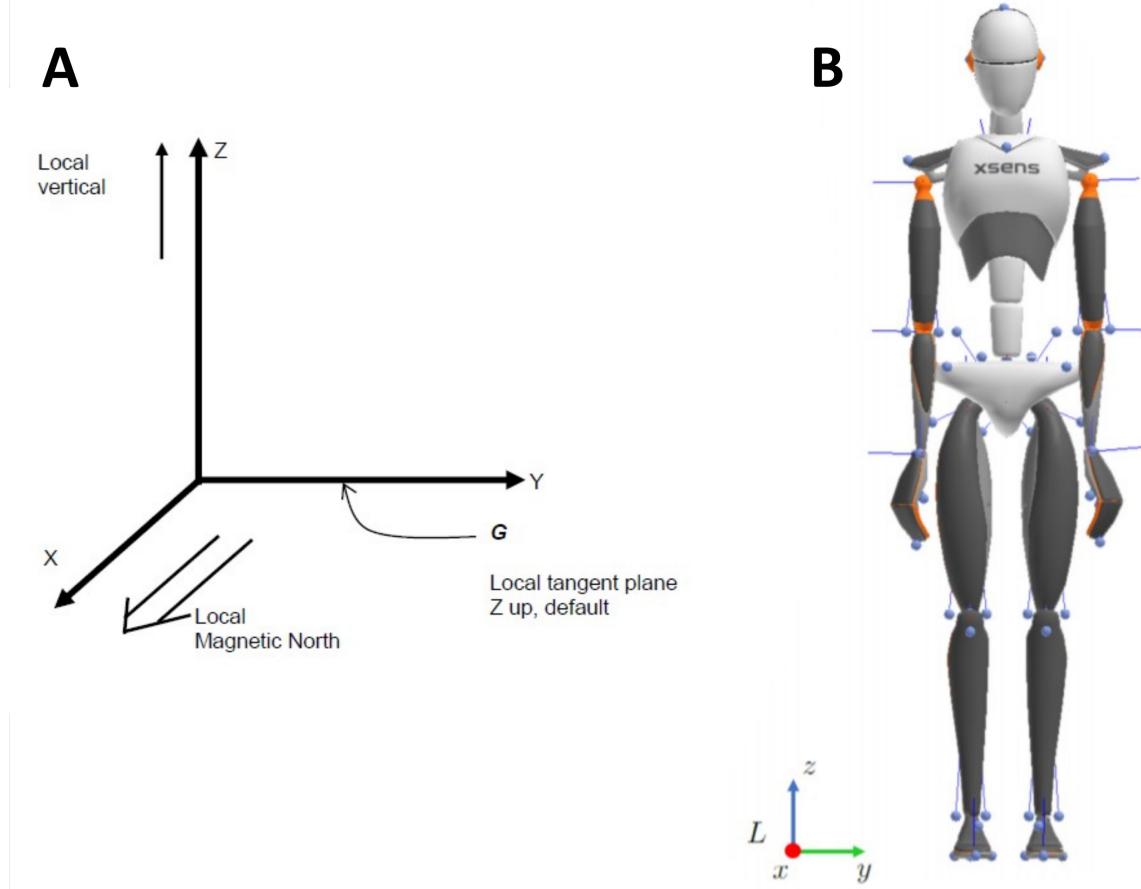

Figure 2. Global and local reference frames. **A.** Representation of the earth-fixed reference coordinate system. **B.** MVN avatar in N-pose.

to X and Z, and Z being positive when pointing upwards along the vertical, gravity referenced, axis. The system was calibrated while the participant was standing in a neutral position ('N-pose'), as shown in Figure 2B. In this pose, the participant is standing in a relaxed, upright position, with the feet being parallel to each other and the arms flat against the body, while looking straight ahead with a natural head position.

Procedure. During the experiment, participants were asked to perform five different tasks, namely walking, running, going up and down the stairs, sitting and standing. The walking and running tasks were performed outside on a standard sidewalk on campus. Going up and down the stairs was done in the university building and the sitting/standing task was performed on a normal office chair in the lab. The first three tasks were adapted from the experiment conducted by Carriot et al. (2014). Each task was repeated three times and each repetition lasted about two minutes, resulting in roughly six minutes of recorded data for each task. The participants were instructed to perform each task at a comfortable speed and while moving and looking around as naturally as possible.

Data Analysis

Pre-Processing. The MVN software saves the orientation data of the motion-trackers in quaternion form. For the purpose of this project, the quaternion data was converted to Euler angles, because we needed to represent the head-in-space prior as a distribution comprised of angles in degrees in the roll-tilt dimension. Therefore, after importing the raw data into MATLAB (version 2019a), the orientation data of the MVN sensor that was attached to the head was converted to Euler angles represented in radians, such that

$$\varphi = atan2\left(2(q_0q_1 + q_2q_3), 1 - 2(q_1^2 + q_2^2)\right) \qquad (1)$$

$$\theta = asin\left(2(q_0q_2 - q_3q_1)\right) \qquad (2)$$

$$\psi = atan2\left(2(q_0q_3 + q_1q_2), 1 - 2(q_2^2 + q_3^2)\right) \qquad (3)$$

in which φ equals the roll-tilt angle, θ equals the pitch angle and equals the yaw angle. q_0, q_1, q_2 and q_3 stand for the four elements that a quaternion is comprised of (Hemingway & O'Reilly, 2018). After the conversion, the data was cleaned by deleting outliers that deviated more than four standard

deviations from the mean. Lastly, the data was converted from radians to degrees.

Distribution fitting. We characterized the data based on their four statistical moments, namely the mean, the variance, the skewness, and the kurtosis. Shortly, in probability and statistics, the mean or expected value is a measure of the central tendency of a probability distribution, i.e. the location of the distribution. The second moment, the variance, provides information about the spread of the distribution. The third moment, the skewness, is a measure of the asymmetry of a probability distribution. A normal distribution (or any other symmetrical distribution) has a skewness of zero. A negative skew indicates that the left tail of the distribution is longer, and a positive skew indicates that the right tail is longer. Lastly, the kurtosis is a measure regarding the tails of a distribution. A normal distribution has a kurtosis of three. A kurtosis greater than three (i.e. excess kurtosis), indicates that the probability distribution has fatter tails, which means that it produces more outliers. Additionally, the peak of the distribution is oftentimes higher and sharper (Brown, 2016). Just like the skewness, the kurtosis is a measure of the shape of the distribution.

To test which distribution type best represents the measured head orientation data and therefore the underlying head-in-space prior, multiple distributions were fitted to the converted roll-tilt data of the MVN sensor. The fitting procedure was performed via the opensource function 'fitmethis' (De Castro, 2020). This function finds the distribution that best fits the data among all distributions available in MATLAB's built-in Maximum Likelihood Estimation function (for a complete overview of the fitted distributions and their respective parameters, see Appendix). Because some distributions can only be fitted to non-negative data (such as the Weibull distribution), we added a constant of 100 degrees to the head orientation data, so that the distributions are roughly centered on 100 degrees head-tilt (instead of 0). This did not affect the actual fitting procedure. The distributions are then ranked according to their Log-Likelihood. This procedure provided us with the necessary distribution-specific parameters which we subsequently used to represent the underlying head-in-space prior. It should be noted that we would formally have to fit circular distributions, because we are dealing with rotation data (De Winkel et al., 2018; Murray & Morgenstern, 2010). However, because the standard deviations of the head orientation data were rather small, differences

between the distributions that were fitted here and circular distributions such as the Von Mises distribution would be negligible (De Winkel et al., 2018). For reasons of computational simplicity, we therefore chose to fit non-circular distributions.

From the resulting fits, it became clear that the normal distribution, the (log)logistic distribution and the t Location-Scale distribution provide the best fits. These three distributions are therefore briefly introduced here.

Normal Distribution. A normal (or Gaussian) distribution is a continuous probability distribution for a real-valued random variable. The parameter μ is the mean of the distribution and σ is its standard deviation, with variance σ^2. The general form of its probability density function is

$$y = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}} \tag{4}$$

During the MLE fitting procedure, the maximum likelihood estimators of μ and σ, respectively, are

$$\bar{x} = \sum_{i=1}^{n} \frac{x_i}{n} \tag{5}$$

$$s_{MLE}^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^2 \tag{6}$$

where (5) is the sample mean, an unbiased estimator of the parameter μ, and (6) is a biased estimator of the parameter σ^2 (*MathWorks - Normal Distribution*, 2020). As was mentioned above, a normal distribution has, by definition, a kurtosis of 3 and is non-skewed.

Logistic Distribution. The logistic distribution is typically used for growth models and in logistic regression. It resembles the normal distribution, but it has longer tails and therefore a higher kurtosis. Its probability density function is defined as

$$f(x; \mu, \sigma) = \frac{e^{-\frac{x-\mu}{\sigma}}}{\sigma\left(1 + e^{-\frac{x-\mu}{2\sigma}}\right)^2} \tag{7}$$

where μ is the mean of the distribution and σ is the scale parameter.

t Location-Scale Distribution. The t Location-Scale distribution is a generalized form of the Student's t distribution. It typically has heavier tails than the both the normal distribution and the logistic

distribution. Its probability density function is given by

$$p(x \mid v, \mu, \sigma) = \frac{\Gamma\left(\frac{v+1}{2}\right)}{\Gamma\left(\frac{v}{2}\right)\sqrt{\pi v}\sigma}\left(1 + \frac{1}{v}\left(\frac{x-\mu}{\sigma}\right)^2\right)^{-\frac{v+1}{2}} \tag{8}$$

where $\Gamma(\bullet)$ is the gamma function, μ is the location parameter, σ is the scale parameter and v is the shape parameter. Compared to the standard Student's t distribution, which only has one parameter, v, the t Location-Scale distribution is more flexible, because here, the scale parameter σ is independent of the shape parameter v, which is not the case in the traditional Student's t distribution. As v increases towards infinity, the distribution approaches the normal distribution.

Bayesian sensory integration model. Figure 1A represents the Bayesian sensory integration model that was used to implement the measured head-in-space orientation to predict the performance on the SVV task. This framework was originally developed by Clemens et al. (2011), although this version of the model is mostly based on the work by Alberts et al. (2016). The model contains three stages of information processing: an input stage, a coordinate transformation stage, and a sensory integration stage.

Sensory input. In the sensory input stage, physical information about the world is transformed to sensory signals, denoted with a hat symbol ($\wedge$). It is assumed that all sensory signals are unbiased but corrupted by Gaussian noise with variance σ^2. Firstly, the otoliths provide the brain with direct information about the orientation of the head in space ($\hat{H}_S$). At small head-tilt angles this information is very precise. However, due to the physiological properties of the otoliths, it can be assumed that the noise level of the sensory information increases rectilinearly with increasing head-tilt (De Vrijer et al., 2008; Tarnutzer et al., 2009, 2010):

$$\sigma_{HS} = \alpha_{HS} * |H_S| + \beta_{HS} \tag{9}$$

Here, β_{HS} reflects the noise level of the otoliths at $0°$ head-tilt and α_{HS} reflects the proportional noise increase with increasing head-tilt. Secondly, neck sensors provide proprioceptive head-on-body information ($\hat{H}_B$) and thirdly, body somato-sensors respond to the orientation of the body in space ($\hat{B}_S$).

Coordinate transformation. In addition to the direct head-in-space information from the otoliths, the brain can use the information from the body somatosensors and the neck sensors to get an *indirect* measure of head-in-space orientation ($\hat{H}_{SI}$). In order to do so, the information from these two sources needs to be combined. This involves a coordinate transformation, such that

$$\hat{H}_{SI} = \hat{B}_S + \hat{H}_B \qquad (10)$$

This means that the Gaussian distributions of the single sources that are centered on B_S and H_B are now combined to one Gaussian distribution centered on $B_S + H_B$.

Sensory integration. At this stage, all available information is statistically optimally combined to a single final head-in-space estimate. As was mentioned before, it is also assumed that, in addition to the sensory information, the brain uses prior information about head orientation in space. In previous research, this prior was represented as a Gaussian distribution that was centered on 0, while the variance of the prior was one of the free parameters. As was described before, we based our head-on-space prior on the motion tracking data, meaning that it is not fitted as a free parameter. This prior is denoted as H_{prior}. When integrating the sensory signals and the prior, the peak of the resulting distribution representing the head-in-space orientation estimation (the posterior) follows from Bayes' rule, and is given by:

$$\tilde{H}_S = w_{HD} * \hat{H}_S + w_{HI} * \hat{H}_{SI} + w_{HP} * H_{Prior} \qquad (11)$$

with

$$w_{HD} = \frac{1/\sigma_{HS}^2}{1/\sigma_{HS}^2 + 1/(\sigma_{HB}^2 + \sigma_{BS}^2) + 1/\sigma_{Prior}^2} \qquad (12)$$

$$w_{HI} = \frac{1/(\sigma_{HB}^2 + \sigma_{BS}^2)}{1/\sigma_{HS}^2 + 1/(\sigma_{HB}^2 + \sigma_{BS}^2) + 1/\sigma_{Prior}^2} \qquad (13)$$

$$w_{HP} = \frac{1/\sigma_{Prior}^2}{1/\sigma_{HS}^2 + 1/(\sigma_{HB}^2 + \sigma_{BS}^2) + 1/\sigma_{Prior}^2} \qquad (14)$$

Here, w_{HD}, w_{HI} and w_{HP}, are the noise-dependent weights of the direct, indirect, and prior information pathway.

Finally, the brain needs to compute the orientation of the luminous line in space. This is achieved by combining the head-in-space information ($\tilde{H}_S$) with eye-in-head information $\tilde{E}_H$ and line-relative-to-eye information ($\tilde{L}_E$). The line-relative-to-eye information itself is assumed to be unbiased, and therefore does not contribute to the resulting error. Regarding the eye-in-head information, it is assumed that the eyes automatically counter-rotate to compensate for small head-tilt angles (i.e. if the head is slightly tilted CW, the eyes rotate CCW).

However, evidence suggests that the brain does not seem to compensate for this counter-rotation, resulting in small errors in the direction opposite of the actual head-tilt at small head-tilt angles (E-effect; Palla et al., 2006). This uncompensated ocular counterroll can be represented as:

$$\tilde{E}_H = -A_{OCR} * sin(\hat{H}_S) \qquad (15)$$

The final systematic error that occurs at different degrees of head-tilt angles (E-effect at small angles, A-effect at large angles) can, thus, be described as:

$$\mu(Error) = (H_S - \tilde{H}_S) + \tilde{E}_H \qquad (16)$$

After having established the type of distribution that best fits the head orientation data, we forward simulated what the SVV task data might look like in these subjects, based on the multisensory integration model. We used both average values for the various parameters from previous research (Clemens et al., 2011) and, for the signal of the prior, the parameters that resulted from the distribution fitting. More specifically, we used the values from Clemens et al. (2011) for the parameters of the otolith signal, the body-in-space signal, the head-on-body signal and the ocular counterroll (α_{HS}, β_{HS}, σ_{BS}^2, σ_{HB}^2, A_{OCR}) (for an overview of the specific parameters, including the parameters for the prior signal, see Table 4 and 5). Upon simulating the SVV task data with the multisensory integration model using the best fitting distributions as the prior, we expected to find that the predicted SVV task performance between head-tilts of $\pm 120°$ would closely match previously observed SVV task behavior.

Results

In the following, the shape of the different head orientation distributions and the best fitting distribution types will be discussed. This will be done by investigating the four statistical moments of a distribution. Figure 3 shows a histogram of the roll-tilt data of the head of one example subject (S4) during all five activities (44 bins). The bottom
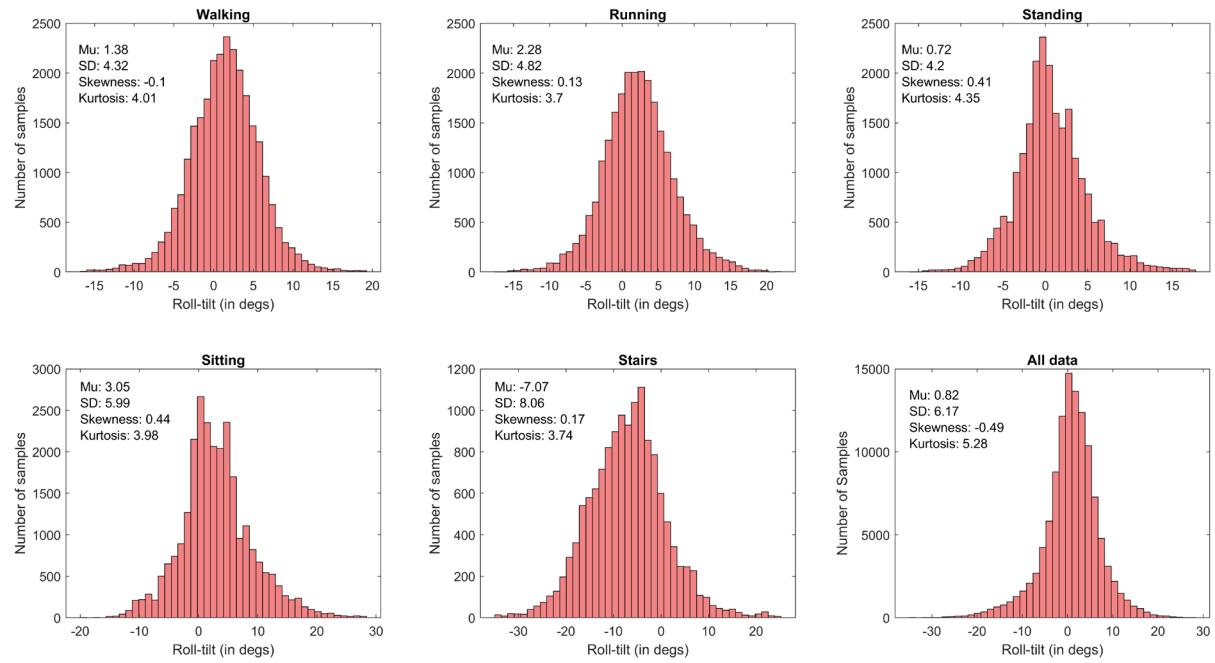
Figure 3. Example data of one subject. The number of samples of the MVN sensors is plotted as a function of roll-tilt of the head. All activities are displayed. The last subplot contains the pooled activities.

right subplot demonstrates a histogram of the data pooled across all activities. Most distributions are roughly centered on 0° head-tilt, while being relatively non-skewed. This indicates that this particular subject tends to hold their head relatively stable in the roll-dimension (i.e. with low variance) without having a bias to either the right or the left roll-tilt side. The standard deviation of the roll-tilt is lowest in the "standing" condition, which is unsurprising. Going up and down the stairs resulted in the highest variance in the roll-tilt dimension, which can potentially be explained by both the frequent gaze shifts that are necessary and by the constant shifts of the body weight when going up and down the stairs. The kurtosis of the last plot (the data of the pooled activities) is equal to 5.28 (see Table 1), indicating that this particular subject has more extreme values in the roll-tilt dimension and a higher peak than would be assumed if the distribution was normally distributed. This pattern is consistent across most subjects (Table 1).

Figure 4 shows the pooled data across all activities for each subject separately. Table 1 shows the four statistical moments of that data.

1. Mean: In all six subjects, the mean roll-tilt angle of the head across all activities was roughly centered on zero. Subject 1 showed the largest bias (M = -2.62 degrees). One-sample Wilcoxon signed-rank tests indicated that the median of the head orientation distributions of all subjects deviated significantly from 0° head-tilt (p < .001). However, effect sizes

were rather small (effect size formula based on Rosenthal (1994); r = .25; r = .17; r = .14; r = .19; r = .20; r = .02, for S1 to S6 respectively), suggesting that the significant results are caused by the large sample sizes (≈120.000 samples per subject). Thus, on average, the participants kept their head upright, without demonstrating considerable biases to either the left or the right roll-tilt side.

2. Standard deviation: Overall, the standard deviations can be considered relatively low. It was highest for Subject 1 (SD = 9.60 degrees) and lowest for Subject 4 (SD = 6.17 degrees), meaning that the majority of the head tilts in the roll dimension were small. Overall, Subject 6 showed the most extreme head-tilts with values up to 59.6 degrees and Subject 3 had the least extreme head-tilts (33.9 degrees).

3. Skewness: The distributions from five out of six subjects were practically non-skewed, with only slight deviations from zero. Subject 2 showed

Table 1. Four statistical moments of pooled activities across subjects.

Subject	Mean	Std	Skew	Kurtosis
S1	-2.62	9.60	-0.85	5.86
S2	-1.98	7.59	-1.04	4.44
S3	0.48	6.44	-0.60	5.29
S4	0.82	6.17	-0.49	5.28
S5	1.16	6.98	0.24	6.32
S6	-0.08	6.26	0.12	5.87
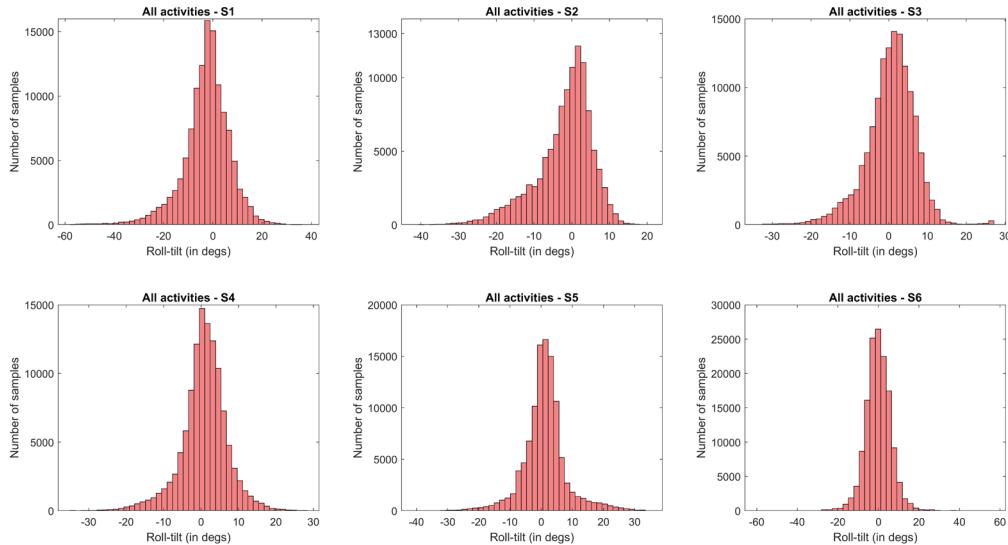
Figure 4. All pooled activities across subjects. *Note.* S = Subject.

Table 2. Best-fitting distributions and their respective log-likelihoods. *Note.* Note that not the absolute values should be interpreted, but the value relative to the other fits. The higher the value, the better the fit (relative to the others).

Subject	Best fits	Log-Likelihood
S1	t Location-Scale distribution	-442350.66
	Logistic distribution	-443412.15
	Loglogistic distribution	-447007.89
	Weibull distribution	-447266.82
S2	Extreme Value distribution	-367709.04
	Weibull distribution	-367739.56
	t Location-Scale distribution	-373448.99
	Logistic distribution	-373756.24
S3	t Location-Scale distribution	-432454.20
	Logistic distribution	-432981.81
	Loglogistic distribution	-434876.83
	Normal distribution	-437678.17
S4	t Location-Scale distribution	-369039.59
	Logistic distribution	-369984.30
	Loglogistic distribution	-371200.89
	Normal distribution	-374895.08
S5	t Location-Scale distribution	-349880.51
	Logistic distribution	-354003.98
	Loglogistic distribution	-354033.10
	Nakagami distribution	-361908.37
S6	t Location-Scale distribution	-454105.05
	Logistic distribution	-454405.78
	Loglogistic distribution	-454651.04
	Normal distribution	-458572.73
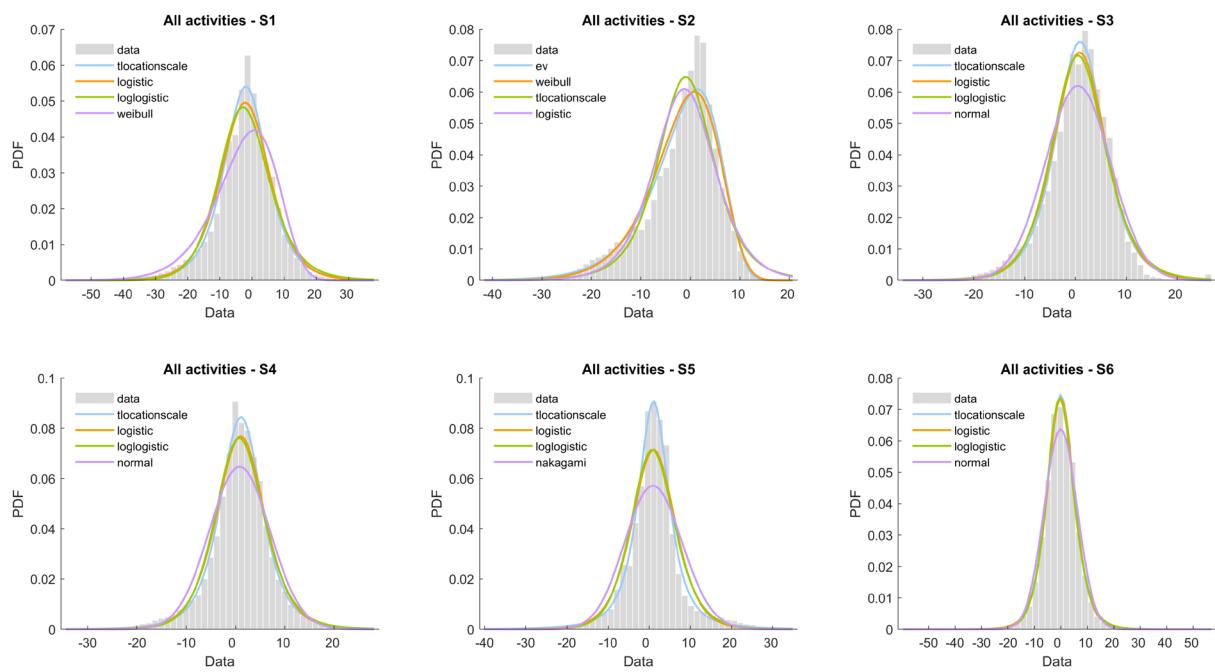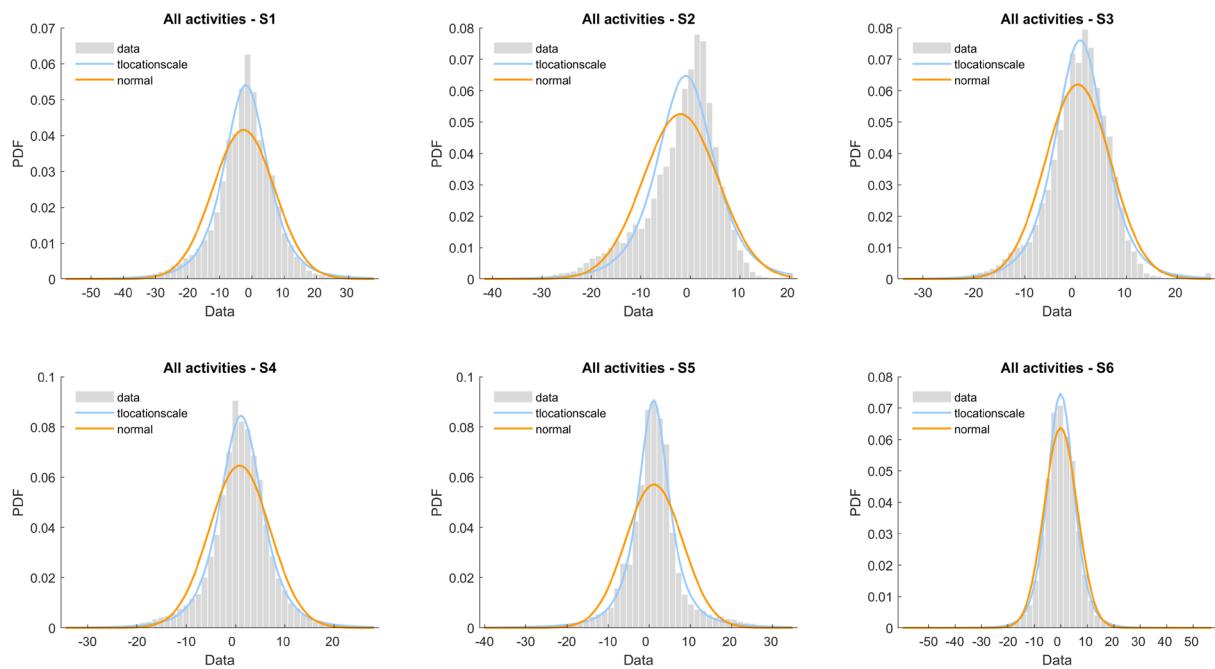
Figure 5. Distribution fitting. S = Subject; PDF = probability density function. **A.** The four best-fitting distribution types for each individual are shown, plotted on top of the data. Distributions can, thus, differ between individuals. Distributions are sorted by best fit (blue = best fit; orange = second-best fit etc.). **B.** Fitted normal distributions vs. t Location-Scale distributions plotted on top of the data for each individual.

the largest asymmetry (-1.04), which is also visible in Figure 4. As a general rule of thumb, skewness values that exceed ±1 can be considered highly skewed (*Normality Testing - Skewness and Kurtosis*, n.d.).

4. Kurtosis: The distributions of all subjects showed excess kurtoses (>3), with values ranging from 4.44 (S2) to 6.32 (S5), which indicates that there were more outliers (i.e. fatter tails) and higher peaks than would be expected if the data were normally distributed.

Summarizing, for all subjects the head orientation distributions were all roughly centered on 0° roll-tilt, with relatively small standard deviations, supporting the assumption that the prior is centered on upright. Furthermore, the distributions of five out of six subjects were practically symmetrical. Only subject 2 demonstrated a considerable skewness of the head orientation data to the left. Notably, all subjects had head orientation distributions with excess kurtosis, which deviate considerably from what would be expected under the assumption that the data are normally distributed.

Distribution fitting

Figure 5A shows the four best fitting distributions, superimposed on the combined roll-tilt head orientation data of all subjects. Table 2 shows the corresponding log-likelihoods for those fits. It should be noted that, in theory, the log-likelihood can lie between and and the values in itself are not meaningful. The values can only be compared to other log-likelihoods. The results show that in five out of six cases, the t Location-Scale distribution provided the best fit to the data, as indicated by the highest log-likelihoods. For subject 2, the extreme value distribution provided the best fit. This is presumably caused by the fact that this head orientation distribution is the most asymmetric one, therefore resulting in worse fits of distributions that are by definition symmetric, such as the normal distribution or the t Location-Scale distribution.

Importantly, the Gaussian distribution does not fit the data well. Figure 5B shows a comparison between the fits of the t Location-Scale distribution and the normal distribution. In contrast to the fit of the normal distribution, the t Location-Scale distribution follows the data much more closely, providing a better representation of the data and, consequently, a more realistic depiction of what the underlying head-in-space prior might look like. Most strikingly, the normal distributions are not able to follow the fat tails of the data, which in turn results in lower-than-optimal peaks. This

means that they underrepresent the amount of the head-tilt data that lies closely around 0°. Hence it follows that the spread of the normal distributions around the inflection points is too large, resulting in an overestimated spread of the data. Furthermore, the excess kurtoses of the data, and therefore its 'tailedness', cannot be captured by the normal distributions, because, by definition, they have a kurtosis of three. It seems logical that the tLocation-Scale distribution provides a better fit compared to the Gaussian distribution, because it has one parameter more that can be flexibly fitted to the data. Therefore, to avoid overfitting, we additionally compared the distribution fits based on the Akaike Information Criterion (AIC).The AIC deals with the risk of over fitting by punishing an increasing number of parameters, therefore providing a more objective measure when comparing models with different numbers of parameters. Generally, a difference in AIC scores of more than 10 means that there is essentially no empirical support for the
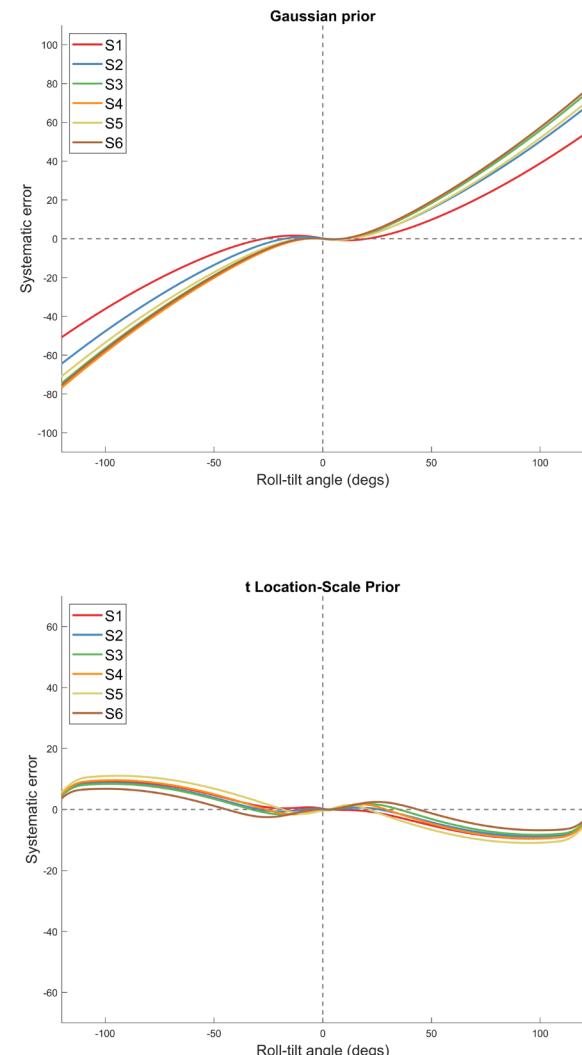

Figure 6. Model simulation of the systematic errors for all subjects. *Note.* S = Subject.
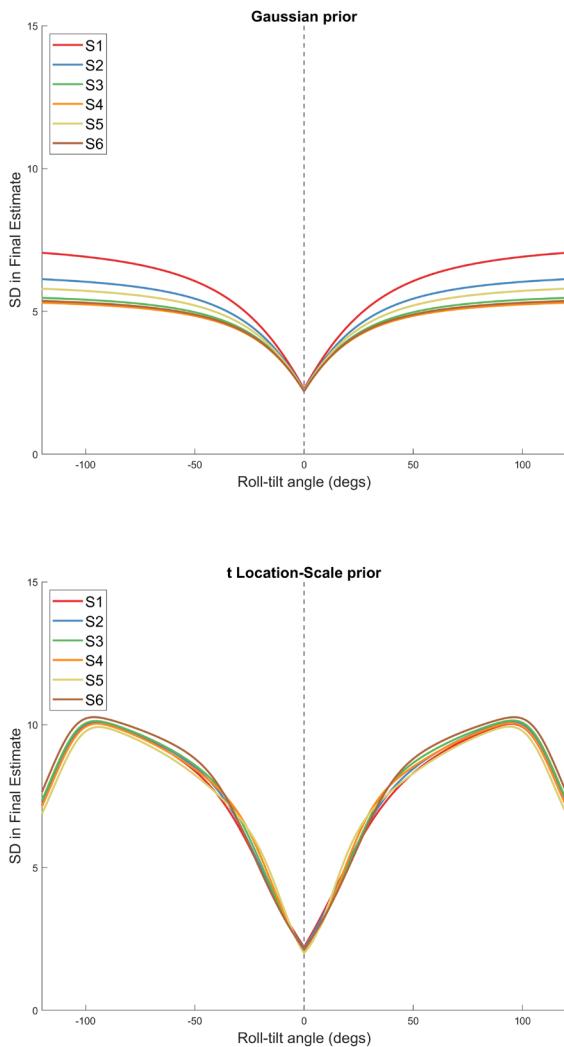
Figure 7. Model simulation of the standard deviations of the final head-in-space estimate. *Note.* S = Subject.

model with the higher AIC (Burnham & Anderson, 2002; Cavanaugh & Neath, 2019). Thus, the smaller the AIC score, the better the fit of the distribution (relative to the other distribution fits). As can be inferred from Table 3, the AIC scores for the t Location-Scale distribution fits are substantially smaller than the AIC scores for the Gaussian distribution, providing further evidence that the t Location-Scale distribution provide a significantly better fit compared to the Gaussian distributions.

To summarize, the normal distributions do not provide a good fit for the naturalistic head orientation data. The data is more peaked and has fatter tails than what a normal distribution is able to capture. Instead, the t Location-Scale distribution provides both the best overall fit in five out of the six subjects and in all six cases a better fit than the normal distribution. Due to its three parameters (location, scale and shape parameter), it is more flexible and therefore better suited to capture this

particular dataset and, consequently, mightbe a better representation of the underlying head-in-space prior.

Model Simulation

We subsequently simulated the Bayesian sensory integration model with both the traditional Gaussian distribution prior and with the best-fit t Location-Scale distribution prior. For the other sensory parameters (see Table 5), we used the best-fit parameters obtained in previous studies (Alberts et al., 2016; Clemens et al., 2011). Thus, only the parameters relating to the prior differed between individuals, while the other parameters were kept constant across participants. Figure 6 shows the expected systematic error in the head-in-space estimate as a function of head roll-tilt, ranging from -120 degrees (CCW) to +120 degrees (CW).

Gaussian Prior

With the Gaussian prior, the systematic errors simulated by the model closely correspond to previous findings of systematic errors on SVV tasks (e.g. Alberts et al., 2016; Clemens et al., 2011; De Vrijer et al., 2008). At the maximum head-tilt of 120 degrees the systematic errors ranged between roughly 50° (S1) and 76° (S4). Thus, when being roll-tilted 120°, a presented line would need to be rotated between 50° to 76° in the head-tilt direction to be perceived as completely vertical, because the participants severely underestimate their own head-tilt. These large individual differences are caused by the different variances of the priors. As is shown in Table 4, S1 and S4 have the largest and smallest variances of the fitted Gaussian prior, respectively (S1: $SD = 9.59$; S4: $SD = 6.17$). Because the prior is

Table 3. AIC scores of t location-scale fits and Gaussian fits. *Note.* The smaller the score, the better the fit. A difference in scores of >10 is considered significant.

| Subject	t Location-Scale	Gaussian
S1	884707.32	898018.93
S2	746903.99	753097.96
S3	864914.41	875360.34
S4	738085.19	749794.17
S5	699767.01	723877.08
S6	908216.10	917149.47
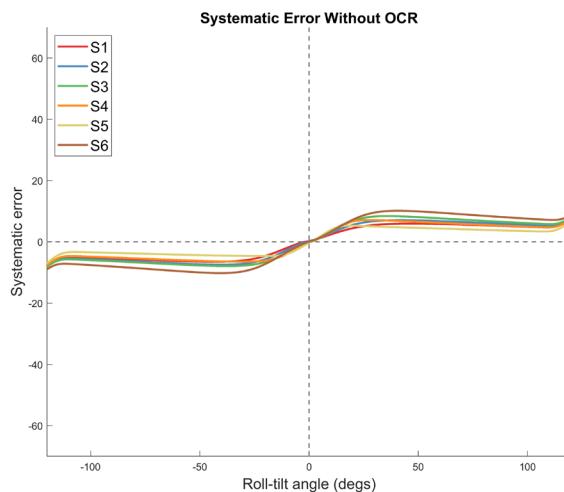
Figure 8. Simulated systematic errors without uncompensated ocular counterroll with t location-scale prior. *Note.* S = Subject; OCR = Ocular counterroll.

weighted according to its variance in the integration stage of the model, the prior is weighted less in S1 and weighted more in S4, resulting in the simulated differences in systematic errors. At small roll-tilt angles, however, the E-effect becomes visible, where small systematic errors in the *opposite* direction of the head-tilt appear, as if the brain overestimates one's own head-tilt. This is caused by the aforementioned uncompensated ocular counterroll (, in the model). Figure 7 shows the variance of the resulting head-in-space estimate as a function of head-tilt. Again, S1 shows the largest variance in the head-in-space estimate at large tilt-angles, which is explained by the high-variance prior. Regarding the magnitudes of the variances, they also correspond to the

Table 4. Parameters of the Gaussian prior and the t location-scale prior. *Note.* Note that the μ and the σ do not correspond to the same thing in the Gaussian and t Location-Scale case. In the latter case, μ = location parameter and σ = scale parameter, not Mean and SD.

Subject	Gaussian prior		t Location-Scale prior		
	μ	σ	μ	σ	ν
S1	-2.62	9.59	-1.92	6.92	3.87
S2	-1.98	7.59	-0.94	5.82	4.42
S3	0.48	6.44	0.91	4.98	4.79
S4	0.82	6.17	1.10	4.41	3.71
S5	1.16	6.98	1.05	3.97	2.37
S6	-0.08	6.26	-0.10	5.13	6.19

equivalent values reported in previous studies. Near upright, the smaller variances of the final estimate indicate lower uncertainty. As roll-tilt increases, so does the uncertainty in the head-in-space estimate. Thus, simulating the Bayesian optimal integration model by using the Gaussian distributions that were obtained from the naturalistic head orientation distributions resulted in systematic errors and head-in-space estimate variances that closely correspond to previously reported results on tasks of perceived verticality.

t Location-Scale prior

For the t Location-Scale distribution prior, this was not the case. Although providing a much better fit to the naturalistic data, the simulated systematic errors do not correspond to both the Gaussian prior version of the simulation and previously measured systematic errors on SVV tasks. In contrast to what would be expected, the systematic errors are 1) much smaller in magnitude (ranging between 6.83° (S6) and 11.06° (S5)) and 2) in the opposite direction of what previous studies have reported. Furthermore, differences between individuals are rather small and do not correspond to the individual differences from the simulation with the Gaussian prior (i.e. S6 and S5 with minimum and maximum systematic errors, compared to S1 and S4 in the Gaussian prior simulation). In contrast to the systematic error, the variance of the final head-in-space estimate corresponds more closely to the variance profile of the Gaussian prior simulation, with higher variances at large head-tilts, compared to upright. However, at head-tilt angles beyond 100°, the variance suddenly decreases significantly, as if uncertainty about one's own head-tilt also decreases. Thus, even though the t Location-Scale distribution fits the naturalistic head-orientation data considerably better than a Gaussian distribution, it does not seem to be able to capture the previously observed systematic errors and variances in the final head-in-space estimate.

Table 5. Best-fit parameter values adapted from previous studies. *Note.* α_{HS} = proportional variance increase of otolith signal; β_{HS} = base signal; σ^2_{BS} = variance body-in-space signal; σ^2_{HR} = variance head-on-body signal; A_{OCR} = uncompensated ocular counterroll.

Parameters	α_{HS}	β_{HS}	σ^2_{BS}	σ^2_{HR}	A_{OCR}
Values	0.16	2.4	10.8	4.9	14.6

Discussion

This study investigated whether motion tracking of naturalistic activities can serve as a viable methodological approach to represent and model the underlying head-in-space prior as part of a Bayesian multisensory integration framework. More specifically, we investigated what type of distribution best fits the real-world head orientation data and whether it can be successfully integrated in the Bayesian optimal integration model to explain performances on tasks of perceived verticality. Based on the assumptions from previous research (e.g. Alberts et al., 2016; Clemens et al., 2011; De Vrijer et al., 2008; MacNeilage et al., 2007)2007, we assumed the naturalistic head orientation distributions to be best described by Gaussian distributions and that we would be able to adequately simulate what the SVV task performance would look like in these subjects. The hypotheses were not supported. We showed that a Gaussian distribution is not able to capture the peaks and tails of the naturalistic head orientation distributions. In fact, the best fit was provided by the tLocation-Scale distributions. However, upon simulating the Bayesian optimal integration model with both the best-fitting Gaussian distribution and the t Location-Scale distribution as the prior, only the Gaussian prior version of the model simulated the biases that were observed in previous studies of verticality perception in a realistic way.

Why the simulated SVV task performances do not correspond to previous findings

Two reasons contribute to the finding that the systematic errors that resulted from the model simulation incorporating the t Location-Scale prior did not correspond to previous findings on tasks of verticality perception. Figure 8 illustrates, what the systematic error would look like if the uncompensated ocular counterroll was *not* added to the final estimate (Equation 16). It becomes clear that in that case, the errors go in the direction of the head-tilt, which is in line with previous research, but they only reach values of about 10°, therefore being far smaller than what was expected. Prior research has shown that systematic errors can reach magnitudes up to 60° at head-tilts of 120° (Clemens et al., 2011). The small magnitude of the systematic error observed here is partly caused by the fat tails of the t Location-Scale prior. The effect this has on the model simulation is remarkable. Even though the

t Location-Scale distributions appear to be smaller in width than the normal distributions, they have far fatter tails. This causes the prior to be weighted less in the signal integration stage of the model. Thus, it only slightly pulls the final head-in-space estimate towards zero, resulting in a much smaller systematic error.

If then, on top of the already small systematic errors, the OCR is added to the error, this results in the observed systematic errors of the model simulation. As can be inferred from Equation 15, the OCR is essentially represented as a relatively large sinusoid that is added on top of the head-in-space error, causing the systematic error to go in the direction opposite of the actual head-tilt. Thus, both the fat tails of the t Location-Scale prior and the OCR that is added on top of the error cause the model simulation to inaccurately predict the SVV task performance.

Can the model work with t Location-Scale priors?

How could the simulation of the SVV task performance be improved in the case of the t Location-Scale prior? One of the main limitations of the present study is that the variances of the various sensory signals were fixed in the model simulation. This is problematic, because the variances were adapted from previous research under the assumption that the prior is of Gaussian nature. Replacing the Gaussian prior for a t Location-Scale prior, while keeping the variances of the sensory signals fixed, thus causes the presented results. To compensate for the different properties of the t Location-Scale distribution (e.g. the fatter tails), the variances of the sensory signals would need to be adjusted, too. This could be achieved by either increasing the variances of the otolith signal and the indirect body signal, therefore increasing the relative weight of the t Location-Scale prior, and/or by reducing the large effect that the uncompensated ocular counterroll has on the results. Ideally, this would be achieved by letting the same participants that took part in the motion tracking experiment perform an SVV task. One could then fit the Bayesian multisensory integration model to the SVV data. In that case, the prior would be fixed, because it is based on the previously measured head orientation distributions, while the sensory signals are fitted to the data as free parameters. This should provide a more accurate fit of the model and, in turn, should result in more realistic parameters for the various signals that are

part of the multisensory integration model. Due to the current situation regarding the COVID-19 crisis, such an experiment was unfortunately not possible, but should be considered in future research.

Implications for future research

In the current study, we were able to extensively quantify head orientations during everyday life activities. We demonstrated that wireless motion tracking can be flexibly used to accurately quantify everyday life activities and that it can serve as a way to combine naturalistic tasks with controlled lab-based measures to investigate spatial orientation. Using a similar naturalistic approach, Carriot and colleagues (2014) have shown that the vestibular system has to deal with angular velocities that deviate considerably from normality. Here we showed that the same applies to probability distributions of head-orientations in the roll-tilt dimension. Participants' head orientation distributions were all roughly centered on upright and, apart from one subject, largely non-skewed. However, the probability distributions were characterized by fatter tails, as indicated by higher kurtoses. Thus, in everyday life, participants experience more extreme head-tilts in the roll dimension than what a normal distribution is able to capture. Assuming that the underlying head-in-space prior is based on lifelong experiences of how the head is typically oriented in space, and further assuming that the activities used in the current study are a realistic representation of everyday life activities, it can be concluded that the underlying prior deviates considerably from normality. Thus, frequently observed individual differences on SVV tasks could be explained by different underlying priors, and, therefore, different naturalistic head orientation distributions. Future research should explore this possibility more extensively.

In the current study, we have challenged the assumption that the head-in-space prior is of Gaussian nature. Apart from the head-in-space prior, the original model (Clemens et al., 2011) also assumed the sensory information from the various sources (otoliths, neck, body somatosensors) to be corrupted by Gaussian noise. Therefore, the question can be asked whether the sensory information from those sources is also non-Gaussian, contrary to what is assumed in the multisensory integration model. Future research could look into the specific statistical characteristics of the different sensory signals. It should be noted though that we would not expect the multisensory integration model to predict significantly different results, even if the

sensory signals are found to be of non-Gaussian nature, too. This is because the sensory signals are assumed to be unbiased and, thus, to be centered on the true head/body tilt, while the systematic errors that the model predicts are solely elicited by the prior. The specific shape of the distributions of the sensory signals would therefore not have a large impact on the multisensory integration process. An alternative multisensory integration model that employs unbiased, albeit non-Gaussian sensory signals, would likely predict very similar systematic errors in the head-in-space estimate compared to the original model.

We assumed the line-relative-to-eye information (i.e. the sensory information of how the line falls onto the retina) to be unbiased. However, Girshick, Landy and Simoncelli (2011) have shown that observers are biased towards perceiving cardinal (i.e. horizontal, vertical) relative to oblique orientations. They argue that these biases might occur due to a prior centered on the two cardinal orientations (i.e. 0° and 90°). Thus, in addition to the head-in-space prior that is solely centered on upright, future work might include an additional "line-relative-to-eye prior" that is centered on 0° and 90° and could therefore account for the biases found by Girshick et al. (2011). In the present study, we chose to use an unbiased line-relative-to-eye representation so that we employ a multisensory integration model that only differs from previous studies in its head-in-space prior, while keeping all other components of the model the same. This made comparisons with previous findings (e.g. Clemens et al. (2011)) substantially easier.

Future projects should also consider the possibility of applying the current approach to patient groups with vestibular disorders, potentially providing novel evidence about the underlying priors in such populations. Previous research has shown that patients with bilateral vestibular function loss tend to show a larger bias in the SVV task at large head-tilts (90°) compared to control subjects (Alberts et al., 2015). In line with the Bayesian optimal integration model (Clemens et al., 2011), this can theoretically be explained by the fact that the brain in those patient groups is not able to use vestibular information. Instead, contributions from other sensory signals regarding head-in-space orientation and the contribution from the prior will be weighted more heavily, resulting in larger systematic errors. Investigating how these patients orient their heads during naturalistic activities and what that indicates with regards to the underlying prior might, therefore, provide valuable evidence

regarding the underlying mechanisms taking place to achieve spatial orientation in patients with vestibular function loss.

Limitations

One of the limitations of the current study is the fact that the chosen naturalistic activities constrain the head movements of the participants, in that it is advantageous to keep the head as stable as possible to maintain spatial orientation, for example during the running task. This leads to relatively small inter-subject differences in both the head orientation distributions and, after integration of those distributions in the model, SVV task performance simulations. In future research it could be considered to introduce activities that force participants to move their heads around more extensively. Even if those activities would not perfectly represent the activities that the underlying prior is based on, they would allow for larger inter-subject differences, which, in turn, might be correlated to inter-subject differences on SVV task performance, providing evidence that everyday life head movements might influence lab-based tasks of verticality perception.

Conclusion

Here we were able to, for the first time, measure and quantify head orientations during naturalistic activities, providing novel evidence of what the underlying head-in-space prior as part of a Bayesian multisensory integration model might look like. In contrast to how the prior has been modelled in previous studies, we showed that head orientation distributions deviate from normality and are characterized by fatter tails, indicating that the underlying head-in-space prior might, too, be of non-Gaussian nature. T Location-Scale distributions provided the best fit to the data in the vast majority of the subjects. However, after integrating those distributions as representations of the prior in the model, simulations of SVV task performance did not correspond to previous behavior on tasks of verticality perception regarding both the direction and magnitude of the systematic errors. This was caused by the properties of the t Location-Scale distribution (i.e. the fat tails) and the fact that the variances of the sensory signals were not adjusted to account for those specific properties of the prior. Future research should investigate this more extensively, for example by adjusting the motion tracked real-world activities to elicit more inter-subject differences in head orientations. Those same

subjects would subsequently participate in a task on verticality perception to investigate correlations between individual differences of naturalistic head orientation distributions and lab-based tasks of perceived verticality.

References

Alberts, B. B. G. T., Selen, L. P. J., Bertolini, G., Straumann, D., Medendorp, W. P., &Tarnutzer, A. A. (2016). Dissociating vestibular and somatosensory contributions to spatial orientation. *Journal of Neurophysiology*, *116*(1), 30–40. https://doi.org/10.1152/jn.00056.2016

Alberts, B. B. G. T., Selen, L. P. J., Verhagen, W. I. M., &Medendorp, W. P. (2015). Sensory substitution in bilateral vestibular a-reflexic patients. *Physiological Reports*, *3*(5), e12385. https://doi.org/10.14814/phy2.12385

Aubert, H. (1861). Eine scheinbare bedeutende Drehung von Objecten bei Neigung des Kopfes nach rechts oder links. *Archiv für pathologische Anatomie und Physiologie und für klinische Medicin*, *20*(3), 381–393. https://doi.org/10.1007/BF02355256

Barra, J., Marquer, A., Joassin, R., Reymond, C., Metge, L., Chauvineau, V., &Pérennou, D. (2010). Humans use internal models to construct and update a sense of verticality. *Brain*, *133*(12), 3552–3563. https://doi.org/10.1093/brain/awq311

Breckenridge, W. G. (1979). *Quaternions proposed standard conventions* [Technical Report]. NASA Jet Propulsion Laboratory.

Brown, S. (2016). *Measures of Shape: Skewness and Kurtosis*. https://brownmath.com/stat/shape.htm#Kurtosis

Burnham, K. P., & Anderson, D. R. (2002). *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach* (2nd ed.). Springer-Verlag. https://doi.org/10.1007/b97636

Carriot, J., Jamali, M., Chacron, M. J., & Cullen, K. E. (2014). Statistics of the vestibular input experienced during natural self-motion: Implications for neural processing. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *34*(24), 8347–8357. https://doi.org/10.1523/JNEUROSCI.0692-14.2014

Cavanaugh, J. E., & Neath, A. A. (2019). The Akaike information criterion: Background, derivation, properties, application, interpretation, and refinements. *WIREs Computational Statistics*, *11*(3), e1460. https://doi.org/10.1002/wics.1460

Ceyte, H., Cian, C., Trousselard, M., & Barraud, P.-A. (2009). Influence of perceived egocentric coordinates on the subjective visual vertical. *Neuroscience Letters*, *462*(1), 85–88. https://doi.org/10.1016/j.neulet.2009.06.048

Clemens, I. A. H., De Vrijer, M., Selen, L. P. J., Van Gisbergen, J. A. M., & Medendorp, W. P. (2011). Multisensory processing in spatial orientation: An inverse probabilistic approach. *The Journal*

of Neuroscience: The Official Journal of the Society for Neuroscience, 31(14), 5365–5377. https://doi.org/10.1523/JNEUROSCI.6472-10.2011

De Castro, F. (2020). Fitmethis. MATLAB Central File Exchange. https://de.mathworks.com/matlabcentral/fileexchange/40167-fitmethis

De Vrijer, M., Medendorp, W. P., &Gisbergen, J. A. M. V. (2009). Accuracy-precision trade-off in visual orientation constancy. Journal of Vision, 9(2), 9–9. https://doi.org/10.1167/9.2.9

De Vrijer, M., Medendorp, W. P., & Van Gisbergen, J. a. M. (2008). Shared computational mechanism for tilt compensation accounts for biased verticality percepts in motion and pattern vision. Journal ofNeurophysiology, 99(2), 915–930. https://doi.org/10.1152/jn.00921.2007

De Winkel, K. N., Katliar, M., Diers, D., &Bülthoff, H. H. (2018). Causal Inference in the Perception of Verticality. Scientific Reports, 8(1), 5483. https://doi.org/10.1038/s41598-018-23838-w

Dieterich, M., & Brandt, T. (2019). Perception of Verticality and Vestibular Disorders of Balance and Falls. Frontiers in Neurology, 10. https://doi.org/10.3389/fneur.2019.00172

Eggert, T. (1998). Der Einfluss orientierter Texturen auf die subjektive visuelle Vertikale und seine systemtheoretische Analyse [PhD thesis]. Technical University of Munich.

Girshick, A. R., Landy, M. S., &Simoncelli, E. P. (2011). Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. Nature Neuroscience, 14(7), 926–932. https://doi.org/10.1038/nn.2831

Hemingway, E. G., & O'Reilly, O. M. (2018). Perspectives on Euler angle singularities, gimbal lock, and the orthogonality of applied forces and applied moments. Multibody System Dynamics, 44(1), 31–56. https://doi.org/10.1007/s11044-018-9620-0

Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. Nature, 427(6971), 244. https://doi.org/10.1038/nature02169

MacNeilage, P. R., Banks, M. S., Berger, D. R., &Bülthoff, H. H. (2007). A Bayesian model of the disambiguation of gravitoinertial force by visual cues. Experimental Brain Research, 179(2), 263–290. https://doi.org/10.1007/s00221-006-0792-0

MacNeilage, P. R., Ganesan, N., &Angelaki, D. E. (2008). Computational approaches to spatial orientation: From transfer functions to dynamic bayesian inference. Journal of Neurophysiology, 100(6), 2981–2996. https://doi.org/10.1152/jn.90677.2008

MathWorks—Normal Distribution. (2020). Normal Distribution. https://de.mathworks.com/help/stats/normal-distribution.html

Mittelstaedt, H. (1983). A new solution to the problem of the subjective vertical. Naturwissenschaften, 70(6), 272–281. https://doi.org/10.1007/BF00404833

Murray, R. F., & Morgenstern, Y. (2010). Cue combination on the circle and the sphere. Journal of Vision, 10(11), 15–15. https://doi.org/10.1167/10.11.15

Normality Testing—Skewness and Kurtosis. (n.d.). GoodData Documentation. Retrieved 4 July 2020, from https://help.gooddata.com/doc/en/reporting-and-dashboards/maql-analytical-query-language/maql-expression-reference/aggregation-functions/statistical-functions/predictive-statistical-use-cases/normality-testing-skewness-and-kurtosis#:~:text=As%20a%20general%20rule%20of,the%20distribution%20is%20approximately%20symmetric.

Palla, A., Bockisch, C. J., Bergamin, O., & Straumann, D. (2006). Dissociated hysteresis of static ocular counterroll in humans. Journal of Neurophysiology, 95(4), 2222–2232. https://doi.org/10.1152/jn.01014.2005

Rosenthal, R. (1994). Parametric measures of effect size. In H. Cooper & L. V. Hedges (Eds.), The handbook of research synthesis (pp. 231–244). Russell Sage Foundation.

Stocker, A. A., &Simoncelli, E. P. (2006). Noise characteristics and prior expectations in human visual speed perception. Nature Neuroscience, 9(4), 578–585. https://doi.org/10.1038/nn1669

Tarnutzer, A. A., Bockisch, C. J., & Straumann, D. (2010). Roll-dependent modulation of the subjective visual vertical: Contributions of head- and trunk-based signals. Journal of Neurophysiology, 103(2), 934–941. https://doi.org/10.1152/jn.00407.2009

Tarnutzer, A. A., Bockisch, C., Straumann, D., &Olasagasti, I. (2009). Gravity dependence of subjective visual vertical variability. Journal of Neurophysiology, 102(3), 1657–1671. https://doi.org/10.1152/jn.00007.2008

Van Beuzekom, A. D., & Van Gisbergen, J. A. (2000). Properties of the internal representation of gravity inferred from spatial-direction and body-tilt estimates. Journal of Neurophysiology, 84(1), 11–27. https://doi.org/10.1152/jn.2000.84.1.11/F

Xsens. (2017). MVN Analyze. https://www.xsens.com/products/mvn-analyze

Appendix

Table 6. All fitted distributions and their respective parameters.

Distribution type	Parameter 1	Parameter 2	Parameter 3
Normal Distribution	μ: mean	σ: standard deviation	-
Exponential Distribution	μ: mean	-	-
Gamma Distribution	a: shape parameter	b: scale parameter	-
Logistic Distribution	μ: mean	σ: scale parameter	-
t Location-Scale Distribution	μ: location parameter	σ: scale parameter	ν: shape parameter
Uniform Distribution	a: lower endpoint (minimum)	b: upper endpoint (maximum)	-
Extreme Value Distribution	μ: location parameter	σ: scale parameter	-
Rayleigh Distribution	b: scale parameter	-	-
Generalized Extreme Value Distribution	k: shape parameter	σ: scale parameter	μ: location parameter
Beta Distribution	a: first shape parameter	b: second shape parameter	-
Nakagami Distribution	μ: shape parameter	ω: scale parameter	-
Rician Distribution	s: noncentrality parameter	σ: scale parameter	-
Inverse Gaussian Distribution	μ: scale parameter	λ: shape parameter	-
Birnbaum-Saunders Distribution	β: scale parameter	γ: shape parameter	-
Generalized Pareto Distribution	k: tail index parameter	σ: scale parameter	θ: threshold parameter
Loglogistic Distribution	μ: mean of logarithmic values	σ: scale parameter of logarithmic values	-
Lognormal Distribution	μ: mean of logarithmic values	σ: scale parameter of logarithmic values	-
Weibull Distribution	a: scale parameter	b: shape parameter	-

Hyperrealistic Neural Decoding:
Reconstruction of Face Stimuli From FMRI Measurements via the GAN Latent Space

Thirza Dado[1], Yagmur Güçlütürk[1], Luca Ambrogioni[1],
Gabriëlle Ras[1], Sander E. Bosch[1], Marcel van Gerven[1] & Umut Güçlü[1]

[1]Radboud University Nijmegen, Donders Institute for Brain, Cognition and Behaviour,
The Netherlands

We introduce a new framework for hyperrealistic reconstruction of perceived naturalistic stimuli from brain recordings. To this end, we embrace the use of generative adversarial networks (GANs) at the earliest step of our neural decoding pipeline by acquiring functional magnetic resonance imaging data as subjects perceived face images created by the generator network of a GAN. Subsequently, we used a decoding approach to predict the latent state of the GAN from brain data. Hence, latent representations for stimulus (re-)generation were obtained, leading to state-of-the-art image reconstructions.

Keywords: deep learning, face generation, functional magnetic resonance imaging, generative adversarial network, neural decoding

Corresponding author: Thirza Dado; E-mail: thirza.dado.1@donders.ru.nl

In recent years, the field of neural decoding has been gaining more and more traction as advanced computational methods became increasingly available for application on neural data. This is a very welcome development in both neuroscience and neurotechnology since reading neural information will not only help understand and explain human brain function but will also find applications in brain computer interfaces and neuroprosthetics to help people with disabilities.

Neural decoding can be conceptualized as the inverse problem where brain responses are mapped back to sensory stimuli via a latent space (Van Gerven, Seeliger, Güçlü, & Güçlütürk, 2019). Such a mapping can be idealized as a composite function of linear and nonlinear transformations. The linear transformation models the mapping from brain responses to the latent space. The latent space should effectively capture the defining properties of the underlying neural representations. The nonlinear transformation models the mapping from the latent space to sensory stimuli.

The systematic correspondences between latent representations of discriminative convolutional networks (convnets) and neural representations of sensory cortices are well established (Yamins et al., 2014; Seyed-Mahdi, Khaligh-Razavi & Kriegeskorte, 2014; Cadieu et al., 2014; Güçlü & Van Gerven, 2015; Güçlü & Van Gerven, 2017; Güçlü, Thielen, Hanke, & Van Gerven, 2016). As such, exploiting these systematic correspondences in neural decoding of visual experience has pushed the state-of-the-art forward (Van Gerven et al, 2019) . This includes linear reconstruction of perceived handwritten characters (Schoenmakers, Barth, Heskes, & Van Gerven, 2013), neural decoding of perceived and imagined object categories (Horikawa & Kamitani, 2017), and

reconstruction of natural images (Seeliger, Güçlü, Ambrogioni, Güçlütürk & Van Gerven, 2018; Shen, Horikawa, Majima, & Kamitani, 2019) and faces (Güçlütürk et al., 2017; VanRullen & Reddy, 2019). Yet, there is still room for improvement since these state-of-the-art results still fall short of providing photorealistic reconstructions.

At the same time, generative adversarial networks (GANs) have emerged as perhaps the most powerful generative models to date that can potentially bring neural decoding to the next level (Brock, Donahue, & Simonyan, 2018; Goodfellow et al., 2014; Karras, Aila, Laine, & Lehtinen, 2017; Karras, Laine, & Aila, 2019). A GAN is a is a deep learning architecture for generative modelling, consisting of two competing neural networks, as described by Goodfellow et al. (2014). In short, a generator network is pitted against a discriminator network that learns to distinguish reconstructed "fake" data samples from real data samples. In turn, the generator's goal is to fool the discriminator by generating new and unique, real-looking data samples from randomly sampled low-dimensional latent features. Competition is the drive between both neural networks to improve their methods in tandem until the generated samples are indistinguishable from the real ones. However, since the *true* latent representations of GANs are not readily available for pre-existing neural data (unlike those of the aforementioned discriminative convnets), the adoption of GANs in neural decoding has been relatively slow (see (Seeliger et al., 2018) for an earlier attempt with GANs and (VanRullen & Reddy, 2019) for a related attempt with variational autoencoders-GAN [VAE-GANs]).

In this study, we introduce a very powerful yet simple framework for HYperrealistic reconstruction of PERception (HYPER), which elegantly
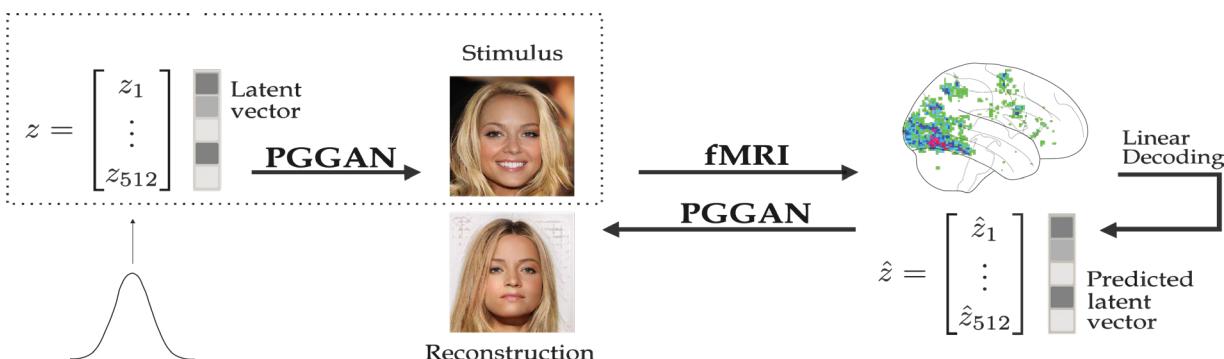
Figure 1. Schematic illustration of the HYPER framework. Face images are generated from randomly sampled latent features $z \in Z$ by a face-generating GAN, as denoted by the dotted box. These faces are then presented as visual stimuli during brain scanning. Next, a linear decoding model learns the mapping from brain responses to the original latent representation, after which it predicts latent features for unseen brain responses. Ultimately, these predicted latent features are fed to the GAN for image reconstruction.

integrates GANs in neural decoding by combining the following components (Fig. 1):

i   GAN. We used a pretrained GAN, which allows for the generation of meaningful data samples from randomly sampled latent vectors. This model is used both for generating the stimulus set and for the ultimate reconstruction of perceived stimuli. In the current study, we used the progressive growing of GANs (PGGAN) model (Karras et al., 2017), which generates photorealistic faces that resemble celebrities.

ii   Functional magnetic resonance imaging (fMRI). We made use of neural data with a known latent representation, obtained by presenting the stimulus set produced using the above-mentioned generative model, and recording the brain responses of participants to these stimuli. In the current study, we collected fMRI recordings in response to the images produced using the PGGAN. We created a dataset consisting of a separate training and test set.

iii   Decoding model. We used a decoding model, mapping the neural data to the latent space of the generative model. Using this model, we then obtained latent vectors for the neural responses corresponding to the stimulus images in the test set. Feeding these latent vectors back into the generative model resulted in the hyperrealistic reconstructions of perception.

Method

Training on synthetic images with known latent features

State-of-the-art face reconstruction techniques use deep neural networks to encode vectors of latent features for the images presented during the fMRI experiment (Güçlütürk et al., 2017; VanRullen & Reddy, 2019). These feature vectors have been shown to have a linear relation with measured brain responses. However, this approach entails information loss since the target images need to be reconstructed from the linear prediction using an approximate inversion network such as a variational decoder, leading to a severe bottleneck to the maximum possible reconstruction quality.

In this paper, we avoid this sub-optimality by presenting photorealistic synthetic images generated using PGGAN to the participants. This allows us to store the ground-truth latents corresponding to the generated images which can be perfectly reconstructed using the generative model after predicting them from brain data.

Neural Decoding

Progressive GAN. To achieve the generation of high-resolution images, a training procedure was developed that grows the generator and discriminator network in a progressive fashion (Karras et al., 2017). More specifically, training on face images from the CelebA-HQ dataset started at a low resolution of 4×4 pixels and layers were added incrementally. To avoid shocks to the well-trained lower-resolution layers, these additional layers were "faded in" smoothly by linear interpolation of the weights from 0 to 1. In the end, a mapping was established from 512-dimensional latent features to hyper-realistic face images with a final resolution of 1024×1024 pixels. At this point, both the generator and discriminator network consisted of nine phases and 23.1M trainable parameters.

Predicting latent vectors from brain data. We adapted the deep generative network of PGGAN by adding a dense layer at the beginning to transform brain data into latent vectors. This layer was trained by minimizing the Euclidean distance between true and predicted latent representations (*batchsize* = 30, *lr* = 0.00001, Adam optimization) with weight decay (*alpha* = 0.01) to reduce complexity and multicollinearity of the model. The remainder of the generative network was kept fixed.

Datasets

Visual Stimuli. High-resolution face images (1024×1024 pixels) were generated by the generator network of a Progressive GAN (PGGAN) model (Karras et al., 2017) from randomly sampled latent vectors. Each generated face image was cropped and resized to 224×224 pixels. In total, 1050 unique faces were presented once for the training set, and 36 faces were repeated 14 times for the test set of which the average brain response was taken. This ensured that the training set covered a large stimulus space to fit a general face model, whereas the voxel responses from the test set contained less noise and higher statistical power.

Brain responses. fMRI data was collected, consisting of blood oxygen level dependent (BOLD) responses that corresponded to the perceived face stimuli. The BOLD responses (TR = 1.5 s, voxel size = 2×2×2 mm³, whole brain coverage) of two healthy subjects were measured (S1: 30-year old male; S2: 32-year old male) while they were fixating

on a target (0.6×0.6 degrees) (Thaler, Schütz, Goodale, & Gegenfurtner, 2013) superimposed on the stimuli (15×15 degrees) to minimize involuntary eye movements.

During preprocessing, the obtained brain volumes were realigned to the first functional scan and the mean functional scan, respectively, after which the volumes were normalized to MNI space. A general linear model was fit to deconvolve task-related neural activation with the canonical hemodynamic response function (HRF). Next, for each voxel, we computed its t-statistic and converted these t-scores to z-statistics to obtain a brain map in terms of z per perceived stimulus. Ultimately, most-active 4096 voxels were selected from the training set to define a voxel mask (Fig. 2). Most of these mask voxels are located in the downstream brain regions. Voxel responses from the test set are not used to create the voxel mask to avoid double-dipping.

The experiment was approved by the local ethics committee (CMO Regio Arnhem-Nijmegen). Subjects provided written informed consent in accordance with the Declaration of Helsinki. The fMRI dataset for both subjects and used models are openly accessible.

Evaluation

Model performance is assessed in terms of three metrics: latent similarity, feature similarity, and structural similarity. First, latent similarity is the Euclidean similarity between predicted and true latent vectors. Second, feature similarity is the Euclidean similarity between feature extraction layer outputs (n=2048) of the ResNet50 model, pretrained for face recognition, which we feed stimuli and reconstructions. Lastly, structural similarity is used to measure the spatial interdependence between pixels of stimuli and reconstructions (Wang, Bovik, Sheikh, & Simoncelli, 2004).

Next, based on the assumption that there exists a hyperplane in latent space for binary semantic attributes (e.g. male vs. female), Shen, Gu, Tang and Zhou (2019) have identified the decision boundaries for five semantic face attributes in PGGAN's latent space: gender, age, the presence of eyeglasses, smile, and pose, by training five independent linear support vector machines (SVMs). We used these decision boundaries to compute feature scores per image, by taking the dot product between latent representation and decision boundary, resulting in a scalar. In this way, model performance with regard to specific visual features could be captured along a continuous spectrum and could be compared across images.
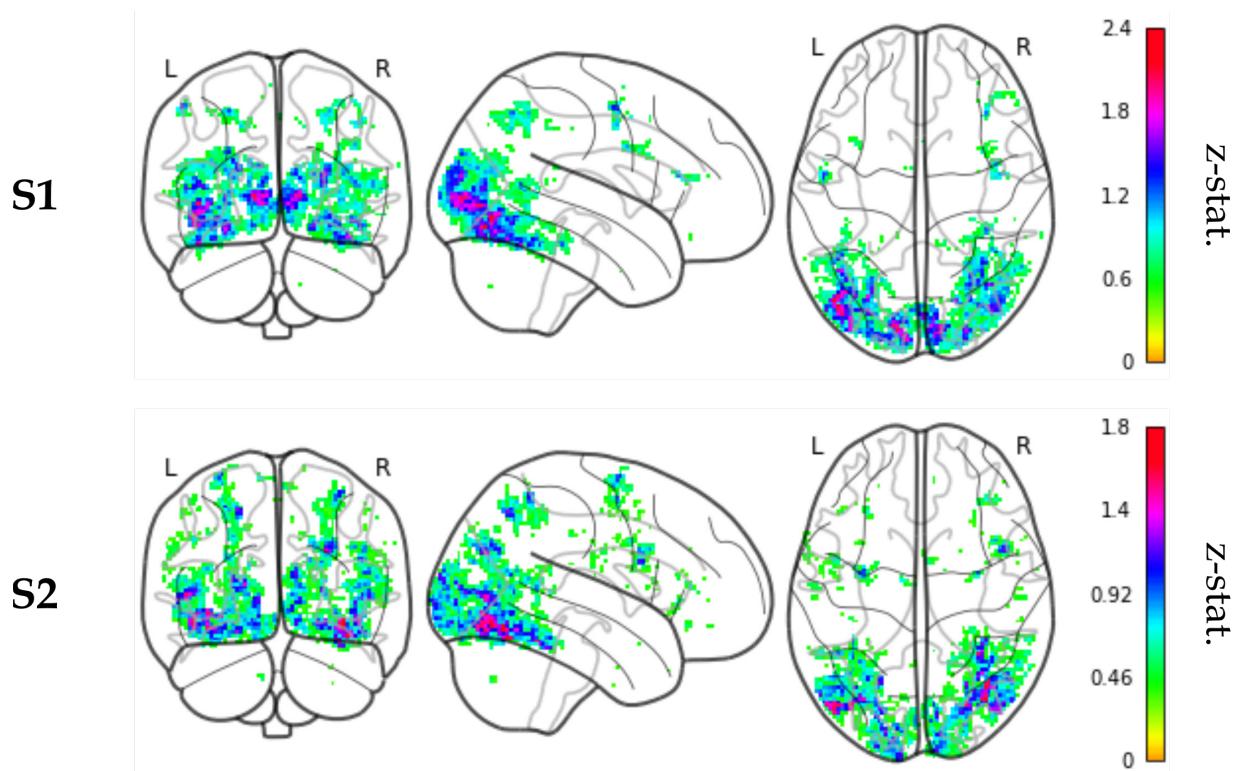

Figure 2. Voxel mask: 4096 most active voxels based on highest z-statistics within the averaged z-map from the training set responses, resulting in a distributed network of activity.

Implementation Details

fMRI preprocessing is implemented in SPM12 after which first-order analysis is carried out in Python's Nipy environment. NVIDIA's PGGAN TensorFlow source code is used in combination with CUDA V10.0.130, CuDNN, and Anaconda3 (Python 3.6). Keras' pretrained implementation of VGGFace (ResNet50 model) is used to evaluate similarities between feature maps of the perceived and reconstructed images. Linear decoding is implemented using ScikitLearn.

Results

Linear decoding of fMRI recordings using PGGAN's latent space has led to unprecedented stimuli reconstructions. Figure 3 presents all the image reconstructions together with the originally perceived stimuli.

To keep the presentation concise, the first half of the images (1-18) are reconstructed from brain activations from Subject 1 and the second half (19-36) from Subject 2. The interpolations visualize the distance between predicted and true latent representations that underlie the (re)generated faces. It demonstrates which features are being retained or change. The bar graphs next to the perceived and reconstructed images show the scores of each image in terms of five semantic face attributes in PGGAN's latent space: gender, age, the presence of eyeglasses, smile, and pose. Looking at the similarities and differences in the graphs for perceived and reconstructed images is a way to evaluate how well each semantic attribute is captured by our model. For most reconstructions, the two graphs match in terms of directionality. A few cases, however, demonstrate that there is still room for improvement (e.g. number 31, 34, and 35). Correlating the feature scores for stimuli and reconstructions resulted in significant ($p < 0.05$; Student's t-test) results for gender, age, eyeglasses, and pose, but not for smile (Fig. 4). We would like to point out that using feature scores quantifies model performance as continuous rather than binary, explaining the significant correlation for eyeglasses despite lack of reconstruction in number 1 and 8.

Next, we compared the performance of the HYPER framework to the state-of-the-art VAE-GAN approach (VanRullen & Reddy, 2019) and the traditional eigenface approach (Cowen, Chun, & Kuhl, 2014) which maps the brain recordings onto different latent spaces. For a fair comparison, we used the same voxel mask to evaluate all the methods presented in this study without any optimization to a particular decoding approach. The VAE-GAN approach predicts 1024-dimensional latent representations which are fed to the VAE's decoder network for stimulus reconstruction (128×128 pixels). The eigenface approach predicts the first 512 principal components (or ' eigenfaces') after which stimulus reconstruction (64×64 pixels) is achieved by applying an inverse principal component analysis (PCA) transform. All quantitative and qualitative comparisons showed that the HYPER framework outperformed the baselines and had significantly above-chance latent and feature reconstruction performance ($p < 0.001$, permutation test), indicating the probability that a random latent vector or image would be more similar to the original stimulus (Table 1).

We also present arbitrarily chosen but representative reconstruction examples from the VAE-GAN and eigenface approach, again demonstrating that the HYPER framework resulted in markedly better reconstructions (Fig. 5).

Figure 3. Results of model 0 that is trained on only the latent vectors. Here, we display the testing set samples 1-18 for Subject 1 and 19-36 for Subject 2. Image reconstructions **(left)** versus perceived images **(right)**. Interpolations visualize similarity regarding the underlying latent representations. Next to each reconstruction and perceived stimulus, a rotated bar graph displays the corresponding feature scores for gender (g), age (a), eyeglasses (e), pose (p), and smile (s).
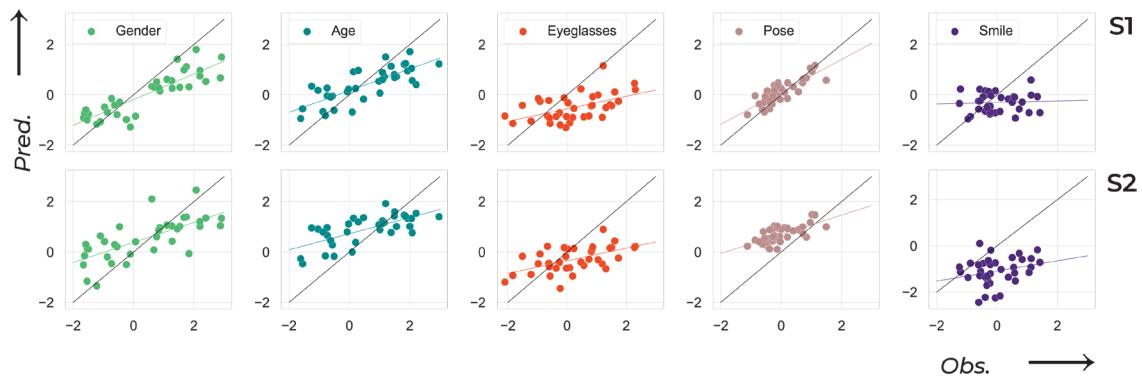
Figure 4. Reconstruction performance on five features. The x-axis denotes the true scores with respect to the perceived stimuli whereas the y-axis represents the predicted scores with respect to the reconstructions. Additionally, the Pearson correlation coefficient (r) and corresponding p-value (p) are displayed.

Discussion

We have decoded brain recordings during perception of face photographs using the presented HYPER method, leading to state-of-the-art stimulus reconstructions. Consequently, this work serves as a proof-of-concept of using generative modelling to approximate neural manifolds of real-world data, possibly bringing our understanding of human brain function forward in the process. The success of this approach is due to the astonishing performance of PGGAN. At the same time, PGGAN puts (potential) bottlenecks on what can be reconstructed: the generator network had to regenerate face images that it had already generated before, guaranteeing its competence. The next step is verifying whether a linear decoding model trained on brain responses with regard to generated face images generalizes

to brain responses to real faces. The true latent representations of real images are not accessible, but would no longer be required if the decoding model has learned to accurately predict them from the artificial data samples. This would result in a great leap forward within the field of neural coding.

Next, the HYPER framework resulted in considerably better reconstructions than the two benchmark approaches. It is important to note that the reconstructions by the VAE-GAN approach appear to be of lower quality than those presented in the original study. A likely explanation for this result could be that the number of training images in our dataset was not sufficient to effectively train their model (8000 vs 1050) and the different voxel selection procedure.

Importantly, image reconstructions by HYPER appear to contain biases. That is, the model predicts

Table 1. Model performance of the HYPER framework compared to the state-of-the-art VAE-GAN (VanRullen & Reddy, 2019) and the eigenface approach (Cowen et al., 2014) is assessed in terms of the feature similarity (column 2) and structural similarity (column 3) between stimuli and reconstructions (mean ± std error). The first column displays latent similarity which is only applicable to the HYPER method because the true and predicted latent vectors are known. Because of resolution differences, all images were resized to 224 × 224 pixels and smoothed with a Gaussian filter (kernel size = 3) for a fair comparison. Also, the backgrounds of the images were removed. In addition, statistical significance of the HYPER method was evaluated against randomly generated latent vectors and their reconstructions.

|		Lat. Sim.	Feat. Sim	Struct. Sim
S1	HYPER	0.4521 ± 0.0026	0.1745 ± 0.0038	0.6663 ± 0.0115
		(p < 0.001; perm. test)	(p < 0.001; perm. test)	(p < 0.001; perm. test)
	VAE-GAN	-	0.1416 ± 0.0025	0.5598 ± 0.0151
	Eigenface	-	0.1319 ± 0.0016	0.5877 ± 0.0115
S2	HYPER	0.4447 ± 0.0020	0.1715 ± 0.0049	0.6035 ± 0.0128
		(p < 0.001; perm. test)	(p < 0.001; perm. test)	(p < 0.001; perm. test)
	VAE-GAN	-	0.1461 ± 0.0022	0.5832 ± 0.0141
	Eigenface	-	0.1261 ± 0.0019	0.5616 ± 0.0097
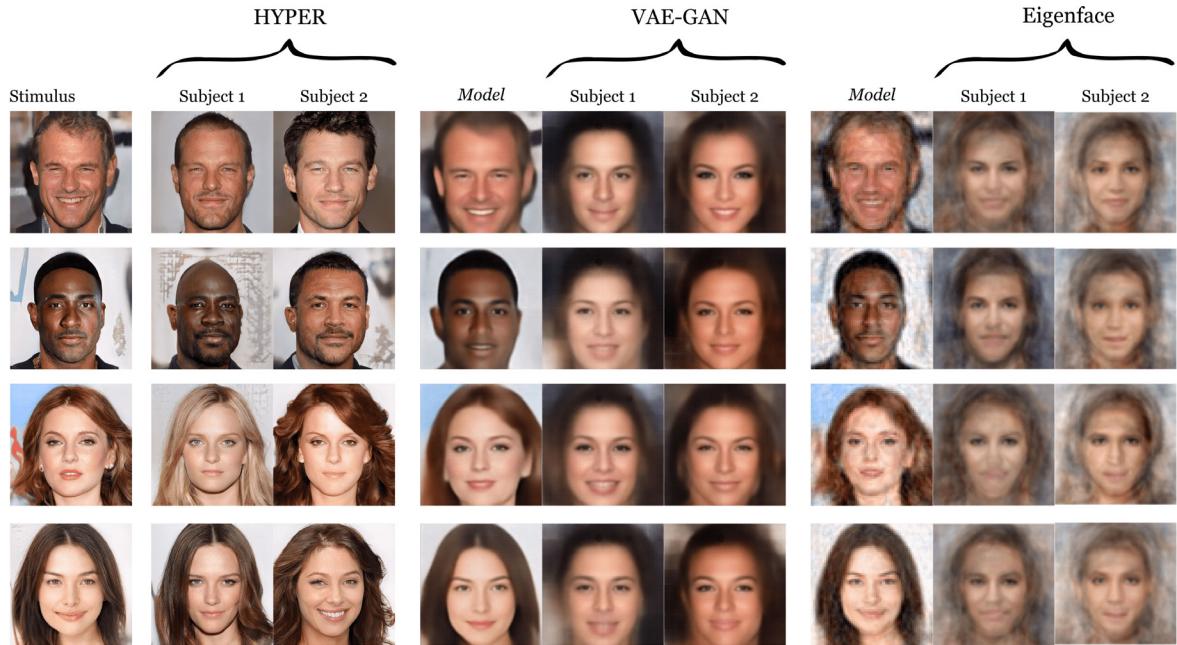
Figure 5. Qualitative results of our approach compared to VanRullen and Reddy (2019) (the VAE-GAN approach) and the eigenface approach in reconstructing image 26, 28, and 36 (arbitrarily chosen). The model columns display the best possible results. For VanRullen and Reddy (2019), this displays reconstructions directly decoded from the 1024-dimensional latent representation of this method. For the eigenfaces approach, this shows reconstructions directly obtained from the 512 principal components.

primarily latent representations corresponding to young, western-looking faces without eyeglasses because predictions tend to follow the image statistics of the (celebrity) training set. PGGAN's generator network is also known to suffer from this problem – referred to as "feature entanglement" – where manipulating one particular feature in latent space affects other features as well (Shen et al., 2018). For example, editing a latent vector to make the generated face wear eyeglasses simultaneously makes the face look older because of such biases in the training data. Feature entanglement obstructs the generator to map unfamiliar latent elements to their respective visual features. It is easy to foresee the complications for reconstructing images of real faces.

A modified version of PGGAN, called StyleGAN (Karras et al., 2019; Karras et al., 2020), is designed to overcome the feature entanglement problem. StyleGAN maps the entangled latent vector to an additional intermediate latent space, thereby reducing feature entanglement, which is then integrated into the generator network using adaptive instance normalization. This results in superior control over the semantic attributes in the reconstructed images and possibly the generator's competence to reconstruct unfamiliar features. Compared to PGGAN, the generated face photographs by StyleGAN have improved

considerably in quality and variation, of which the latter is likely to alleviate current biases. Replacing the PGGAN with StyleGAN would therefore be a logical next step for studies concerned with the neural decoding of faces.

Furthermore, neural decoding can reveal what information is (not) present in the observed brain activations. That is, even though participants are presented with identical stimuli, sensory information is likely to be integrated with subjective expectations and beliefs, causing subjective variations in reconstructions. This may include enhanced, diminished, missing, imagined, or transformed information. Eventually, the HYPER framework might allow us to bridge the gap between objective and subjective experience. However, care must be taken as "mind reading" technologies also involve serious ethical concerns regarding mental privacy. Although current approaches to neural decoding, such as the one presented in this manuscript, would not allow for involuntary access to thoughts of a person, future developments may allow for extraction of information from the brain more easily, as the field is rapidly developing. As with all scientific and technological developments, ethical principles and guidelines as well as data protection regulations should be followed strictly to ensure the safety of (the data of) potential users of these technologies.

Finally, besides the large scientific potential, this research could also have societal impacts when enabling various applications in the field of neurotechnology (e.g. brain computer interfacing and neuroprosthetics) to help people with disabilities. While the current work focuses on decoding of sensory perception, extensions of our framework to imagery could make it a preferred means for communication for locked-in patients.

Conclusion

We have presented a framework for HYperrealistic reconstruction of PERception (HYPER) by neural decoding of brain responses via the GAN latent space, leading to unparalleled state-of-the-art stimulus reconstructions. Considering the speed of progress in the field of generative modelling, we believe that the HYPER framework that we have introduced in this study will likely result in even more impressive reconstructions of perception and possibly even imagery in the near future, ultimately allowing for better understanding the mechanisms of human brain function.

References

Brock, A., Donahue, J., & Simonyan, K. (2018). Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*.

Cadieu, C. F., Hong, H., Yamins, D. L., Pinto, N., Ardila, D., Solomon, E. A., ... DiCarlo, J. J. (2014). Deep neural networks rival the representation of primate it cortex for core visual object recognition. *PLoS Comput Biol*, *10*(12), e1003963.

Cowen, A. S., Chun, M. M., & Kuhl, B. A. (2014). Neural portraits of perception: reconstructing face images from evoked brain activity. *Neuroimage*, *94*, 12–22.

Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., ... others (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, *536*(7615), 171–178.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... Bengio, Y. (2014). Generative adversarial nets. In *Advances in neural information processing systems* (pp. 2672–2680).

Güçlü, U., Thielen, J., Hanke, M., & Van Gerven, M. (2016). Brains on beats. In *Advances in neural information processing systems* (pp. 2101–2109).

Güçlü, U., & van Gerven, M. A. (2015). Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *Journal of Neuroscience*, *35*(27), 10005–10014.

Güçlü, U., & van Gerven, M. A. (2017). Increasingly complex representations of natural movies across the dorsal stream are shared between subjects. *NeuroImage*, *145*, 329–336.

Güçlütürk, Y., Güçlü, U., Seeliger, K., Bosch, S., van Lier, R., & van Gerven, M. A. (2017). Reconstructing perceived faces from brain activations with deep adversarial neural decoding. *Advances in Neural Information Processing Systems*, 4246-4257.

Horikawa, T., & Kamitani, Y. (2017). Generic decoding of seen and imagined objects using hierarchical visual features. *Nature Communications*, *8*(1), 1–15.

Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2017). Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.

Karras, T., Laine, S., & Aila, T. (2019). A style-based generator architecture for generative adversarial networks. In *Proceedings of the ieee conference on computer vision and pattern recognition* (pp. 4401–4410).

Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., & Aila, T. (2020). Analyzing and improving the image quality of stylegan. In *Proceedings of the ieee/cvf conference on computer vision and pattern recognition* (pp. 8110–8119).

Khaligh-Razavi, S.-M., & Kriegeskorte, N. (2014). Deep supervised, but not unsupervised, models may explain it cortical representation. *PLoS Comput Biol*, *10*(11).

Schoenmakers, S., Barth, M., Heskes, T., & Van Gerven, M. (2013). Linear reconstruction of perceived images from human brain activity. *NeuroImage*, *83*, 951–961.

Seeliger, K., Güçlü, U., Ambrogioni, L., Güçlütürk, Y., & van Gerven, M. A. (2018). Generative adversarial networks for reconstructing natural images from brain activity. *NeuroImage*, *181*, 775–785.

Shen, G., Horikawa, T., Majima, K., & Kamitani, Y. (2019). Deep image reconstruction from human brain activity. *PLoS Comput Biol*, *15*(1), e1006633.

Shen, Y., Gu, J., Tang, X., & Zhou, B. (2019). Interpreting the latent space of gans for semantic face editing. *arXiv preprint arXiv:1907.10786*.

Thaler, L., Schütz, A. C., Goodale, M. A., & Gegenfurtner, K. R. (2013). What is the best fixation target? the effect of target shape on stability of fixational eye movements. *Vision Research*, *76*, 31–42.

van Gerven, M. A., Seeliger, K., Güçlü, U., & Güçlütürk, Y. (2019). Current advances in neural decoding. In *Explainable ai: Interpreting, explaining and visualizing deep learning* (pp. 379–394). Springer.

VanRullen, R., & Reddy, L. (2019). Reconstructing faces from fmri patterns using deep generative neural networks. *Communications biology*, *2*(1), 193.

Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, *13*(4), 600–612.

Yamins, D. L., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., & DiCarlo, J. J. (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences*, *111*(23), 8619–8624.

Active Sensing Through Oscillatory Synchronisation: A Possible Mechanism for Filtering and Amplifying Input in Both Humans and Artificial Cognitive Agents

Djamari Oetringer[1]

Supervisors: Dr. Saskia Haegens[1,2], Dr. Hesham ElShafei[1], Dr. Johan Kwisthout[1]

[1]*Radboud University Nijmegen, Donders Institute for Brain, Cognition and Behavior, The Netherlands*
[2]*Columbia University, New York, USA*

Brain oscillations are known to reflect fluctuations of low and high excitability states in neuronal populations. These oscillations can adjust to the surrounding environment such that high excitability states co-occur with relevant sensory information. Such adjustment is a promising mechanism for filtering sensory input and could occur through neural entrainment. Driven by an external rhythmic input, intrinsic oscillations might phase-align with (i.e., entrain to) this input, resulting in the optimal processing of stimuli that are in phase with the rhythm. Oscillatory adjustment could also occur through covert active sensing which entails that the motor cortex drives the signals in the sensory cortex. Thus, covert active sensing and entrainment could explain a novel behavioural effect found in prior work, namely that subjects respond faster in a discrimination task when the external rhythm is faster. 13 subjects performed a visual discrimination task while brain signals were recorded using MEG. Targets were cued by a rhythmic stream of visual stimuli at different frequencies and appeared after one, two, or three cycles, or not at all. In summary, we found support for the aforementioned behavioural effect (i.e., subjects responding faster when cued by faster external rhythms) and covert active sensing, but not for entrainment. We further discuss how the findings of the current study could inspire the development of artificial cognitive agents to tackle the problem of determining which information from the environment is relevant. Importantly, this includes a proposal for how the fields of neuroscience and AI can actively interact with each other, such that both fields benefit.

Corresponding author: Djamari Oetringer; E-mail: d.oetringer@gmail.com

Neural oscillations

Neural oscillations are rhythmic fluctuations in neural activity that can be measured by, among other methods, non-invasive scalp recordings such as electroencephalography (EEG) and magnetoencephalography (MEG).

Oscillations as a filter mechanism. Bishop (1932) suggested that brain oscillations reflect fluctuations of low and high excitability states in neuronal populations. Intracellular peaks over time correspond to high excitability, meaning a state in which a neuron needs less stimulation to generate an action potential. In contrast, intracellular troughs correspond to low excitability states. This proposal by Bishop has been confirmed by multiple studies (see Schroeder and Lakatos (2009) for a recent discussion).

This mechanism as proposed by Bishop could play a key role in the sampling of relevant information and controlling the flow of information through the brain. If phases of high excitability coincide with relevant input, this input receives optimal processing. If instead the input coincides with a state of low excitability, the amplitude of the generated neural activity will be smaller (see Figure 1).

Entrainment. One key question is whether the brain can use rhythmic fluctuation as an active mechanism for sensory sampling (i.e., if the timing of relevant input is predictable, rhythmic fluctuations can adjust to coincide with this input). There is evidence that intrinsic brain rhythms can be synchronised to external rhythmic stimuli in our environment via neural entrainment (see Figure 1; Thut, Schyns, and Gross (2011); Zoefel, ten Oever, and Sack (2018) for reviews). Here, neural entrainment entails that the intrinsic oscillators synchronise with an external rhythm. The advantage of such entrainment is that by predicting the timing of a relevant stimulus, the peaks of excitability can be adjusted such that they coincide with the predicted timing. Then, the stimulus undergoes optimal processing if it indeed occurs at the predicted timing. Without such a prediction, the timing of the stimulus relative to peaks of excitability is random and thus would not necessarily be processed optimally.

However, the behavioural correlates and the underlying neural mechanisms of neural entrainment remain unclear (Haegens & Golumbic, 2018; Obleser & Kayser, 2019). One of the issues is that it is difficult to establish whether neural findings are part of a proactive mechanism, or whether they simply echo the rhythmic environment (Nobre & van Ede, 2018). During a rhythmic input, oscillations could reflect a build-up of rhythmic evoked responses rather than actual entrainment. A
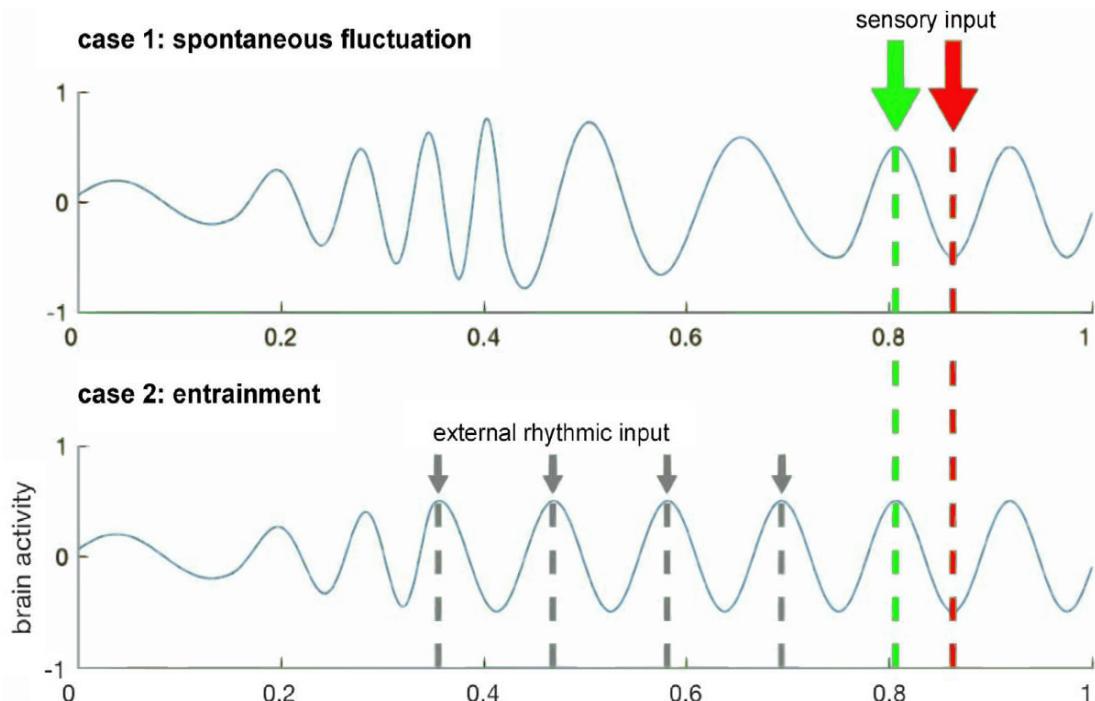
**Figure 1.** Schematic overview of rhythmic facilitation. The green input is facilitated due to the timing relative to the oscillation, as opposed to the red input which is impaired. Here, facilitation entails that the input is processed preferentially or more efficiently. **Top:** intrinsic, spontaneous oscillation. **Bottom:** oscillation entrained to an external rhythm. Figure adapted from Haegens and Golumbic (2018).

stronger case for neural entrainment could be made when one can demonstrate the effects of rhythmic input in the brain signals after the rhythm itself is no longer present (Haegens & Golumbic, 2018). Zoefel et al. (2018) argued that the oscillations could be more than a repetition of evoked responses. They argued that research should disentangle endogenous oscillatory activity, evoked responses, and predictive processes. In the current study, we included time windows in our analyses that did not contain evoked responses for this reason.

Importantly, we can only speak of true entrainment when the following three requirements are met (Haegens, 2020; Haegens & Golumbic, 2018):

1. An endogenous neural oscillator is present apart from rhythmic stimulation.

2. The neural oscillator phase-aligns with the external rhythm. This implies that the frequency of the oscillator exactly matches that of the external rhythm. Crucially, this should only work for a certain range of frequencies rather than for any frequency. Since intrinsic oscillations have a limited range, these oscillations cannot entrain to a rhythm that is outside of this range.

3. This phase-alignment continues for some number of cycles beyond the presentation of the external rhythm.

Active sensing. Active sensing is the process of actively gathering more information about what is being sensed, rather than passively waiting for input to arrive. Examples given by Schroeder, Wilson, Radman, Scharfman and Lakatos (2010) are somatosensory exploration, natural viewing, and sniffing. When identifying an object using tactile information, we actively use our fingers to feel different parts of the object until we have identified the object. When viewing a scene, we use saccades to move our fovea to multiple parts of the scene to create a full picture, rather than passively viewing the centre and processing the relatively scarce information. Finally, when detecting an odour, we, and other animals, start sniffing in order to gather more information about the odour to identify it. Interestingly, high frequency neural oscillations in the olfactory system are heavily involved in the sampling of olfactory input (Pont, 1987). Sniffing could play a major role in actively coordinating the timing between the inflow of input and the phase of the neural oscillations.

The examples given above all involve actual movement. These movements can have an impact on the ongoing neural oscillations, for example by resetting the phase after a saccade (Leszczynski & Schroeder, 2019). This makes the oscillations time-locked to the movement. If then the input is time-locked to this same movement as well, as is the case for saccadic movements, the input arrives consistently at a certain phase of the oscillations. Possibly the brain employs the same neural mechanisms for perceptual selection without overt movements taking place, which we investigated in the current study. This covert active sensing entails that the motor system coordinates the oscillations in the sensory cortex by the use of synchronisation. This would be supported by finding an increase in synchronisation between the motor and sensory cortex, without any actual movement taking place.

In fact, multiple studies have found support for this sensory-motor coupling in the auditory domain. For example, Alho et al. (2014) found a correlation between performance in a phonetic categorisation task and neural synchronisation between the auditory and premotor cortex. Another study used causal connectivity analysis and found that, during continuous speech perception, the motor cortex modulates oscillations in the auditory cortex in the low-frequency range (Park, Ince, Schyns, Thut, & Gross, 2015). Assaneo, Florencia and Poeppel (2018) even showed that, while listening to speech, there is only an auditory-motor coupling when the rate of the speech falls within a certain range. Moreover, this coupling is enhanced at the frequency that corresponds to the mean syllable rate in natural speech (about 4.5 Hz). By the use of neural modelling, they also suggested that the possible underlying neural architecture, namely an intrinsic oscillator, could give rise to such coupling.

In this project we investigated sensory-motor coupling in the visual domain. To our knowledge, this mechanism has not yet been thoroughly explored in this sensory domain.

As explained above, this potential mechanism could explain how we sample the input from our environment. This mechanism is interesting from the perspective of AI, as determining which input is relevant (and thus should be processed after filtering) is one of the open computational problems. The computational problems relevant to the current study are described next.

Computational problems of relevance in AI

The field of AI is already able to achieve impressive results in various domains. For example, the facial recognition system DeepFace is able to

recognise faces at roughly the same performance level as humans (Taigman, Yang, Ranzato, & Wolf, 2014). In addition, the computer system AlphaGo defeated the human European champion in Go (Silver et al., 2016). Moreover, an increasing amount of research is devoted to improving the detection of cancer by the use of artificial neural networks (e.g., Chon, Balachandar, & Lu, 2017). Although these and other examples of successful AI applications can be useful, they require many resources (i.e., time and energy), particularly when training the models. This is unlike natural intelligence, which can learn game rules and possible strategies after few presentations and without the need for much energy. Moreover, the scope of the solvable problems using the methods of conventional AI is limited, especially as compared to the achievements of natural intelligence. Some of the problems that are outside of this scope of solvable problems, are actually solved by natural intelligence. For example, humans are able to adjust their movements after an injury. If an AI system with a robotic arm would play a game of Go against a human opponent, but both players slightly injure their arm due to an accident shortly before the game, the AI system would have trouble adjusting to the new situation (such as a motor in the arm having less power than before). Instead, the human player would use their other arm or adjust the movements of their injured arm such that it is not painful, while still reaching the goal of placing a piece on the board at the right location.

Importantly, one of the main differences between conventional AI and natural intelligence is the use of the time dimension. Conventional AI such as deep neural networks abstract away from natural neural networks, to the extent that the dimension of time almost completely disappears. As discussed in the section on neural oscillations, the timing of input in combination with oscillations could be of great importance in determining which input is relevant. This specific problem of determining what is relevant is actually an unsolved computational problem in AI, as explained below.

The frame problem. In Cognitive Science and AI, the problem of determining what is relevant is also known as the *frame problem*. In the original interpretation, this problem entails that one has to make a computational system determine what information in the world does **not** change after a certain event or action (McCarthy & Hayes, 1981). This is also termed the *inertia problem* (McDermott, 1987). For example, if someone takes a cookie out of a jar, they do not only know that the number of

cookies in the jar decreased by one, but also that this action did not change the stain that is on their trousers, the city that they are currently in, and many other details that were present in the world before taking the cookie out of the jar.

In a broader sense, the frame problem is about "how the relevant pieces of knowledge are found and how they influence one's understanding of the situation" (Haselager, 1997, p.83). Here, knowledge can consist of both current perception and already existing knowledge. Thus, the frame problem can be simplified to *determining what is relevant*, given the vast amount of information in the world. This is exactly what natural intelligence is exposed to as well, given the continuous stream of input from various sensory modalities as described earlier.

The problem of abstraction. Another way of looking at the problem of what is relevant, is looking at how one can abstract away from the world and all of its information to the problem at hand. Computational cognitive scientists and computer scientists often describe a computational problem while assuming that an abstraction from the real world to the computational problem has already taken place. These computational problems then already exclude any input that is irrelevant to the problem. Kwisthout (2012) argued that this abstraction cannot just be assumed to take place correctly and without any computational overload. Additionally, Kwisthout presented a computational framework that pertains to abstracting away from all of the available information in the world to a formal representation of the current problem to be solved. Abstracting away means taking only those pieces of information that are relevant and leaving those that are irrelevant. Kwisthout then showed that finding a subset of relevant pieces of information given all possible subsets is in fact intractable, meaning that it is very unlikely that either natural or artificial intelligence is or will be able to solve this abstraction problem. Nonetheless, humans and other animals seem to somehow solve this problem with ease, as they perform everyday tasks in a world full of incoming stimuli and knowledge.

In order to create an artificial cognitive agent that can dynamically and appropriately react to its environment, the problem of abstraction should be solved, possibly by the use of a heuristic (i.e., strategies that simplify the problem by creating short-cuts). By using a heuristic, the mechanism may not always result in the correct answer, but in practice the results are sufficient. Since the proof by Kwisthout that showed intractability only applies

to an exact solution, there is reason to believe that the use of a heuristic, rather than always finding exact solutions, could still be tractable. Interestingly, Dennett (2006) noted that human beings are not perfect in determining what is relevant. They make mistakes, but in practice it works well enough. Thus, this is a clear indication that these problems are solved by natural intelligence by the use of heuristics and that AI could do that as well. A possible approach that could lead to a solution in AI would be taking inspiration from how the natural brain solves this problem, possibly sufficiently rather than exactly.

The problem of perceptual relevance. In this project we focused on a small aspect of the aforementioned problems: the problem of determining which stimuli in the current environment are relevant and amplifying these stimuli when processing the information. Here we term this problem the *problem of perceptual relevance*. This thus includes the information that is perceived while performing a certain task, but not any already existing knowledge. Note that pre-existing knowledge may still be involved in determining which perceived stimulus is relevant, but the question in the problem of perceptual relevance only pertains to perceived stimuli.

Linking findings in natural intelligence and AI to each other. The fact that humans do not seem to have any trouble with establishing what is relevant indicates that there must be a way to address the problem. Perhaps the strategy or mechanism as used by humans could inspire possible mechanisms to be used by an artificial cognitive agent. That is one of the reasons why the field of neuroscience is important to AI.

We can also learn about natural intelligence by researching AI. If we find a way in which AI could solve the frame problem, then possibly this holds for natural intelligence as well. However, as Dennett (2006) noted, many proposals to solve the frame problem easily become biologically implausible. This is an issue when one tries to understand natural cognition by researching AI. Here we tried to avoid this pitfall by studying literature in the field of neuroscience, as well as conducting an experiment with human subjects.

The notions of active sensing and entrainment, as described in the section on neural oscillations, are potential solutions that can be studied in natural intelligence. These mechanisms cannot solve the problem of perceptual relevance fully, in either natural or artificial intelligence, but it is possibly

a solution to at least part of the problem. We believe that natural intelligence solves the problem of perceptual relevance by means of multiple mechanisms that are integrated together, of which sensory-motor coupling could be one.

Research questions and hypotheses

In this project we investigated entrainment and the role of active sensing in the human brain. As part of investigating active sensing and entrainment, we aimed to answer the following research questions:
- Do brain oscillations in the motor and visual system adapt to different visual external rhythmic streams?
- Does the motor system coordinate active sensing through oscillatory inter-regional phase coupling with the visual system?
- If so, does this coupling correlate with performance?

The motivation for this study stems from a series of psychophysics experiments showing that subjects were faster when task frequency increased (manuscript in preparation). Here, the task frequency was the frequency of the external rhythm that was used to cue participants. In these series of experiments, participants performed an auditory task where they were rhythmically cued regarding the timing of a stimulus probe. Reaction time (RT) decreased when task frequency increased, which was a consistent finding across all experiments. Active sensing could explain this finding. Namely, if the frequency of the external rhythm increases, then active sensing increases the functional connectivity between the motor cortex and visual cortex. This speeds up interregional communication and increases the sampling rate to adjust to higher frequencies, given that the oscillations are entrained. A higher sampling rate implies that input could be sampled earlier than for a lower sampling rate. If stimuli are sampled earlier in time, a reaction can also take place earlier, giving a lower RT for higher task frequency. These unpublished results are in line with the literature about auditory-motor coupling discussed in the section on neural oscillations.

Based on these findings, we expected to find such a relationship between RT and frequency of the rhythmic cue: RT decreases with increasing task frequency. As active sensing is our underlying hypothesis that could explain this finding, we also expected an increase in synchronisation in both the motor cortex and the visual cortex when performing a task that includes a rhythmic cue. More importantly, we expected this synchronisation to somehow differ

between different task frequencies. This difference in synchronisation could manifest itself in two ways: frequency-specific or frequency non-specific (see Figure 2). In the former case, synchronisation peaks shift to higher neural frequencies given a higher task frequency. In the latter case, synchronisation increases at a certain neural frequency given a higher task frequency. Alternatively, synchronisation could increase when being rhythmically cued, but without the task frequency having an effect on the magnitude or location of the peak in the frequency domain. Here, we quantified neural synchronisation as inter-trial coherence (ITC) or power. We expected to at least see a shift in peak given a lower or higher task frequency during the presentation of a rhythmic cue. This would be a consequence of having regular evoked responses. Our main window of interest was that following the external rhythm. Within this window, we expected either of the two possibilities or possibly a combination, meaning that the peak both shifts and increases in amplitude. Furthermore, we expected our findings to be within the delta-to-theta range (1-7 Hz).

Other than the neural effects within the motor and visual cortex, we also investigated the functional connectivity between the two regions to answer our second research question. This is more indicative of active sensing taking place, as we expected a high functional connectivity in anticipation of a task-related stimulus. Here we hypothesised that beta oscillations establish the connection between two sources (Spitzer & Haegens, 2017), and thus we expected to see the effect in the beta frequency range (14-40 Hz).

Both a higher functional connectivity and an increase in synchrony within a brain region would be beneficial for performing the task. We therefore further expected to see a positive correlation between the functional connectivity and performance, and between synchrony measures within regions and performance.

Effects of the corona regulations

From Monday, March 16th onward, all data collection with participants was set on hold at the Donders Centre for Cognitive Neuroimaging. At that point in time we had collected MEG and MRI data of only eight participants. Before we had started this data collection, we also collected pilot data of another five participants that were not supposed to be part of the final analysis. We had initially decided to exclude these because of a slight change in paradigm. Namely, the used baseline period (see the section on experimental design) was lengthened from 1 to 2.5 s. Given the unusual circumstances, we have decided to include those five pilot participants in our analysis in order to have a more substantial dataset for the purpose of this thesis. Furthermore, we reported inconclusive results and additionally mentioned some of the trends that seem to be present based on visual inspection, rather than statistical significance.

Methods

Participants

Thirteen participants were recruited through the SONA subject database of Radboud University (age mean = 26; age SD = 4.40; 9 female, 4 male). Five of those participants performed the task with
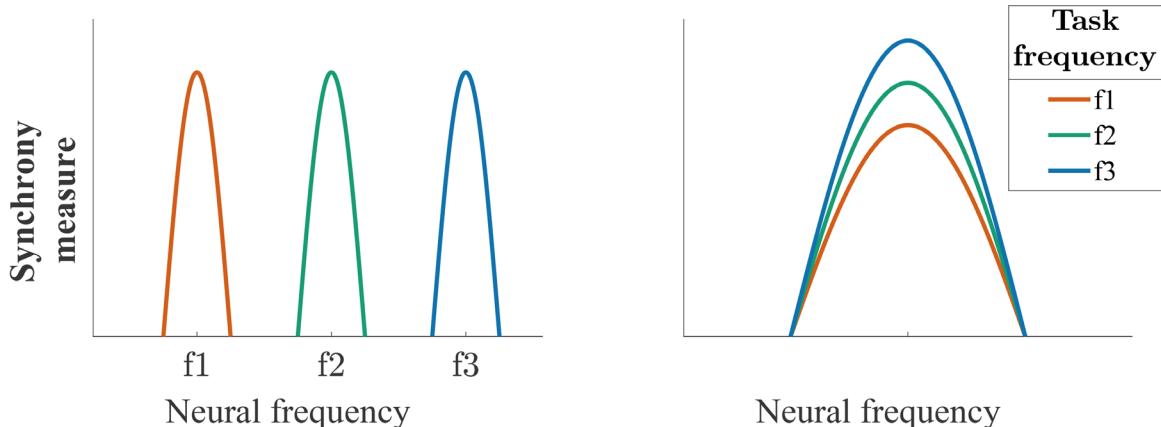
Figure 2. Schematic visualisation of how an increase in synchronisation due to increasing task frequency could manifest itself in the brain signals. Here, f1 is the lowest task frequency and f3 the highest. **Left:** frequency-specific; the peak height of the synchrony measure stays the same, but the peak shifts to a higher neural frequency. **Right:** frequency non-specific; the neural frequency at which synchrony is high does not shift, but the amount of synchronicity increases for higher task frequencies.

a baseline period of 1 s, while the remaining eight had a baseline period of 2.5 s. All participants were either right-handed or ambidextrous, and reported normal or corrected-to-normal vision and no neurological or health problems. At the start of the session, participants signed a consent form. They were rewarded monetarily. Nine participants already had an MRI scan available from a previous MEG experiment.

The study was approved by the local ethics committee and conducted according to the corresponding ethical guidelines (CMO Arnhem-Nijmegen).

Experimental design

Task. Participants performed a visual discrimination task. They were presented with a stimulus and instructed to indicate whether the stimulus was a number or a letter. They responded using a button press with their right index finger. The button mapping was counterbalanced across participants.

Stimuli. The stimuli were adapted from the study by Gwilliams and King (2017) (see Figure 3A) and consisted of "digital-clock" style letters and numbers, each consisting of at most seven line segments. There were four possible letters and four possible numbers. We had chosen pairs of stimuli such that they would differ in only one line segment from each other. These pairs were: 1 and J, 4 and H, 6 and E, and 8 and A. The intensity value of the differentiating line segment could be changed to adjust the difficulty level of the task. An example

of such an adjustment is shown in Figure 3B. The task was made more challenging by making the intensity values of the ambiguous line segments of both stimuli closer to each other. We further had chosen the stimuli such that in two out of four pairs, the letter contained the extra line segment, while the number contained the extra line segment in the remaining two pairs. This was needed to avoid that either of the two categories (letter or number) was harder than the other. It further prevented subjects from basing their decision on mere luminance, which would have been an option if for instance only the letters had one line segment more. Difficulty level was adjusted to individual subjects in order to have the right performance level for each subject. It was determined during the training phase.

The stimuli were presented on a semitranslucent screen (1920 x 1080 pixel resolution, 120 Hz refresh rate) back-projected by a PROpixx projector (VPixx Technologies). They were presented on a grey background and 10.6 by 5.7 cm in size. The non-ambiguous line segments of the stimuli were black, while the ambiguous line segment had a value between the grey background and black, where the exact value depended on the difficulty.

To create a visual rhythm, we presented a cue stream consisting of five zeros at the task frequency before the onset of the stimulus. Participants were asked not to respond to these zeros, but still attend to them since they helped predict the timing of the stimulus.

Conditions and target timings. The frequency of the rhythm at which the cue stream was presented varied across conditions. Each participant was

A   B
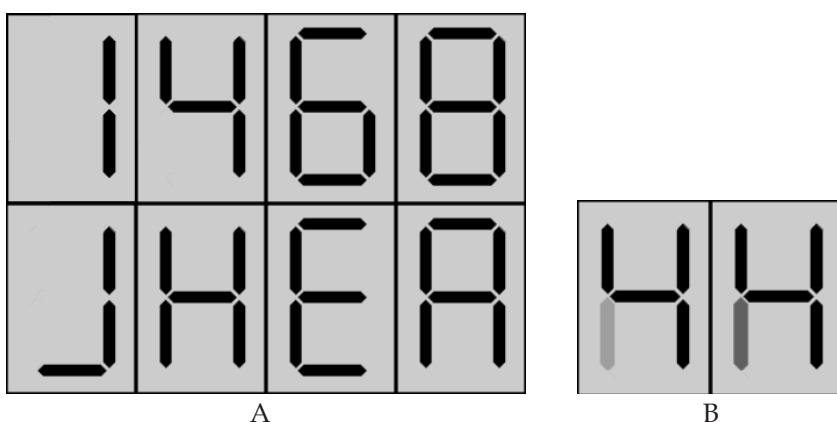
Figure 3. Stimuli. **A.** The eight possible stimuli. **Top:** 1, 4, 6, and 8. **Bottom:** J, H, E, and A. Any letter in the bottom row forms a pair with the number that is shown above it. These images exclude the ambiguity that is added to make the task more difficult. **B.** Example of a pair of stimuli where ambiguity is added by changing the intensity value of the one line segment that differs between the two stimuli. Here, left is closer to 4 and should be identified as a number, while right is closer to H and thus should be identified as a letter.

exposed to three task frequencies in a block-wise matter: 1.3 Hz, 2.1 Hz and 3.1 Hz. In what follows, these task frequencies are denoted by f1, f2, and f3, respectively. We had chosen the task frequencies such that they were within the delta frequency range: low enough to not be irritating to the eye. They were also chosen such that the harmonics of the lower frequencies did not interfere with a higher frequency while taking frequency resolution into account.

The target could appear at four target timings (described in Figure 4). The stimulus occurred either one, two or three cycles after the cue stream, where a cycle equals the inverse of the task frequency. In 40% of the trials, the stimulus did not appear at all, in which case the participant was not supposed to press any button. These trials are termed catch trials. We introduced the various target timings to create big enough windows of interest to perform meaningful frequency analysis with a high enough frequency resolution, for which three cycles need to fit in the window. Our window of interest was defined as the period between the offset of the cue stream and the onset of the stimulus, or the end of the trial in case of a catch trial. The purpose of the various target timings was to avoid biasing the expectation of the participant regarding the timing of the stimulus toward the end of the window of interest. As such, these three target timings were equally likely to occur (20% each). We introduced the catch trials (remaining 40%) to have a stimulus-free window of interest. Additionally, we expected no or very few mistakes during catch trials. Thus, fewer trials had to be omitted in this target timing as compared to any non-catch target timing. Moreover, we could increase the ratio of the catch trials without affecting participants' expectations regarding the timing of the target. In total, 60% of the trials were

used for MEG analysis.

In what follows, the trials with 1, 2 or 3 cycles are referred to as cycle1, cycle2, and cycle3 trials, respectively.

Protocol. Each trial started with a fixation cross, shown for 200 ms. The fixation cross was followed by an empty screen, which remained for 1 s in case of the first five subjects and 2.5 s in case of the other eight participants. In non-catch trials, the response window was 2 s. In case of catch trials, the screen remained empty after the cue stream for 3 cycles plus a jitter of 300-400 ms. We added the jitter to avoid having roughly the same phase in the brain signals at the start of the next trial. All trials were followed by another 500 ms of an empty screen to avoid contaminating the baseline period at the start of the next trial with a motor evoked response.

Participants performed 12 blocks of 40 trials each (480 trials in total). The task frequency varied across blocks, giving 4 blocks per task frequency. Multiple times throughout the experiment, participants were asked two questions. Before and after each block, they were presented with the statement 'Please rate your sleepiness'. They responded using a button press on a four-point scale ranging from 'Very Alert' to 'Very Sleepy'. After each block, participants were also asked about their perceived rhythmicity during the preceding block, again on a four-point scale and using a button press to respond. Here the scale ranged from 'Very Irregular' to 'Very Rhythmic'. They were then shown their performance (accuracy) during the preceding block, and were invited to take a break as long as they needed. At the end of this break, the head position was adjusted to get back to the initial head position, as measured at the start of the experiment, as much as possible.
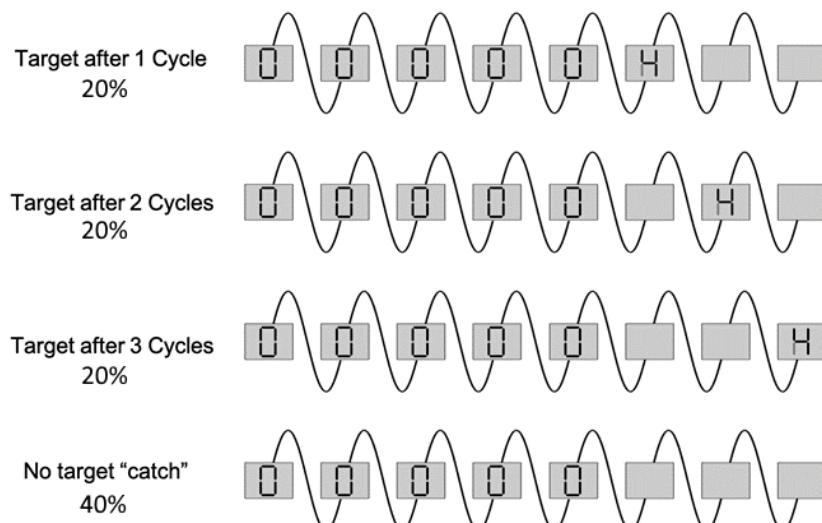

Figure 4. Overview of target timings and ratios.

Before the main experiment started, participants performed the task in the training phase. This training phase was meant to determine the subject-specific level of difficulty. Here, one block consisted of 36 trials. Within such a block, every 12 trials were of a different task frequency. The order of the possible frequencies was randomised. Catch trials were still included, although they did not contribute to learning the task, to avoid surprising the participant at the start of the main experiment.

Participants received feedback after every trial: a green fixation cross if they were correct, and a red one if they were incorrect. This differed from the main part of the experiment, where they only received feedback after a full block by means of an accuracy percentage. After each block, the level of difficulty was adjusted if needed. There were 11 possible difficulties (i.e., pairs of intensity values of the ambiguous line segments). Each participant performed at least 3 training blocks. The goal was to have an accuracy between 70 and 85 percent before starting the main part of the experiment, excluding the accuracy on catch trials. The resulting difficulty was used throughout all 12 blocks of the main experiment.

The session lasted in total about 2 hours. This includes preparation, training, breaks, and measuring the head shape as described in the section on measurements. The experiment was programmed using Psychtoolbox (Brainard, 1997) in MATLAB (The Mathworks, Inc).

Measurements

Subjects were seated in a CTF-275 MEG system with axial gradiometers at a distance of 80 cm from the projection screen. We monitored their head position using three head coils: one in each ear using earplugs and one taped to the nasion. During the main part of the experiment, eye movements were measured using an Eye Link 1000 Eye tracker (SR Research). RT and accuracy were logged for each trial. After the experiment, the shape of the participant's head was acquired with the Polhemus. Information about the head shape was used to increase the quality of source reconstruction.

Finally, those participants that did not yet have a T1-weighted anatomical MRI available, underwent an MRI scan while wearing ear plugs that contained vitamin E, which is useful in source reconstruction analysis described in the section on neurophysiological analysis.

Behavioural analysis

Behaviour was analysed for all trial types except catch trials. RTs were normalized, after which outliers were removed by calculating the Tukey fences. This last step also removed any non-catch trial where the subject did not respond. For all subjects, at most 10 percent of the trials were flagged as outliers. Furthermore, any incorrect trial was removed when analysing RT, while they were still included for computing accuracy.

Repeated-measures ANOVAs were performed on accuracy and median RT with 2 factors: task frequency (3 levels: f1, f2, f3) and target timing (3 levels: cycle1, cycle2, and cycle3).

Neurophysiological data analysis

Pre-processing. MEG data was first down-sampled to 300 Hz. We used three band-stop filters to remove line noise (50 Hz) and its harmonics (100 Hz, 150 Hz). We defined the frequencies of the filters as the frequency to be removed ± 1 Hz. Trials were cut into epochs of 13 s (-1 to 12 s, relative to the onset of the fixation cross).

Trials contaminated with high variance, muscle artefacts, or SQUID jumps were removed through visual inspection on a trial-by-trial basis. In case of muscle artefacts, data was first high-pass filtered (60 Hz) before visual inspection. We performed an independent component analysis (ICA) on remaining trials in order to identify components representing heartbeat, blinks, or saccades. On average, 4 components were removed per subject. Finally, all incorrect trials were removed before analysing the brain signals at both sensor- and source-level.

Source-level data. We computed individual volume conduction models using the single-shell method on the MRI image, supplemented with MEG Polhemus head shape information to further refine co-registration. Individual source models were computed by warping the MNI coordinates of a 5 mm-grid to the individual MRI images using non-linear normalization. The volume conduction model, source model and cleaned MEG data were then used to compute the leadfield.

In order to localise visual and motor regions, we used the evoked responses to the first zero of the cue stream and to the button press respectively. A spatial filter was computed from the leadfield, volume conduction model and covariance matrix (from start of baseline until end of activity window) using

linearly constrained minimum variance (LCMV) beamforming (van Veen, van Drongelen, Yuchtman, & Suzuki (1997); lambda 5%). The spatial filter was then used to compute the average time courses within the activity windows (visual: +80 to +180 ms, motor: -100 to +400 ms) for each voxel, and the corresponding baselines (visual baseline period: -100 to 0 ms, motor baseline period: -600 to -100 ms). We had carefully selected the visual window of interest to include the peak evoked responses for all individuals. Voxels with the highest increase in signal within the activity windows were selected: two for the motor source and two for the visual source. In case of the visual source, one voxel was selected from each hemisphere.

A covariance matrix was computed for the full epoch. Then a spatial filter was computed using LCMV beamforming with the same parameters as before. Finally, the virtual sensors were computed using this spatial filter, and averaging across the two voxels per source.

Window selection

Each trial contained three time periods: baseline, cue stream, and target periods. Here, the target period was the time between the offset of the cue stream and the onset of the target (or the end of the trial in case of catch trials). We defined the trial-level baseline period as the period between the offset of the fixation cross and the onset of the cue stream. In order to avoid contamination with evoked responses due to visual stimulation, we removed the first 200 or 300 ms of the baseline period (starting after fixation cross) and the first 200 ms of the target window (starting at offset of the cue stream).

The maximum length of the resulting baseline window was either 0.8 or 2.2 s. In combination with the variance of task frequency, the full cue stream was sometimes longer and sometimes shorter than the maximum baseline window. For spectral analysis, window lengths were matched. In case of catch trials, the end of the target window was defined as 3 cycles plus 200 ms. Again, window lengths were matched.

Power and inter-trial coherence

For the target window and corresponding baseline, only cycle3 and catch trials were included, because these allow for better spectral resolution. We zero-padded all windows up to 10 s. A Fourier transformation was then performed with a Hanning taper on each window. We defined the frequencies of

interest as 0.1 Hz to 10 Hz (in steps of 0.1 Hz) and 10 to 40 Hz (in steps of 1 Hz). This was done for each trial separately, resulting in a complex Fourier spectrum per trial per time window, from which both the power and ITC could be computed. We further removed any frequency that was below the true frequency resolution. Here, the true frequency resolution equals 1 divided by the length of the window in seconds. In case of ITC, we specifically computed the inter-trial phase coherence rather than the inter-trial linear coherence.

To compute the change in power during the cue stream window or target window as compared to the corresponding baseline, we averaged the power of the baseline windows of all trials with a certain task frequency. This means that all target timings (or only cycle3 and catch in case of investigating the target window) were taken together to compute an overall baseline. Then the percentage change in power was computed per trial, using this average baseline rather than the single-trial baseline. This is to avoid any extreme values in percentage change in power due to an extremely low value in the baseline window that can be present at the trial-level.

ITC was first computed per window and per condition. We then averaged the ITC across the two included timing types. To compute an increase in ITC as compared to the baseline, we subtracted the baseline ITC from the ITC during the window of interest.

Functional connectivity analysis

In order to compute the functional connectivity between the visual and motor source, we used the virtual channel data as computed in the section on neurophysiological data analysis. A Fourier transformation with a Hanning taper was again performed, but now the frequencies of interest were defined as 1 to 40 Hz, with a frequency resolution of 1 Hz. Again, the data per time window were first padded to 10 s, and any results in the range of frequencies below the true frequency resolution were removed.

Connectivity analysis was performed on the remaining data. More specifically, we computed the imaginary part of the coherence between the two virtual channels. As a result, we had a measure of functional connectivity per window. The effect of volume conduction is a common concern in functional connectivity analysis. It implies that perfect source separation is not possible, because of which some signals in our computed motor source may actually originate from the visual source or vice

versa. This would then give a spurious increase in functional connectivity. However, by looking at the imaginary part of coherence, we circumvent the effect of volume conduction (Nobre & van Ede, 2018).

Statistical analysis of neurophysiological data

To determine the statistical significance of ITC, power, and functional connectivity, cluster-based permutation tests (Maris & Oostenveld, 2007) were performed within the frequency range of 1 to 7 Hz or in the range of 1 to 40 Hz in case of functional connectivity. The cluster-level statistic equalled the sum of the sample-specific statistics (t- or F-values) that belong to the cluster. The test statistic that was evaluated by the use of permutation, equalled the maximum of the cluster-level statistic. The alpha value of the clusters was set to 0.05. All tests consisted of 1000 permutations with Monte-Carlo estimates of the significance probabilities.

We first performed the statistical analysis on task-induced signals, meaning that the signal of interest during either the cue window or the target window was contrasted with the corresponding baseline. Here, dependent samples t-statistics were computed for each sample in each permutation. In case of ITC and functional connectivity, the test was one-tailed, while it was two-tailed in case for power. This was done because ITC and functional connectivity were expected to only increase as compared to baseline, while power could decrease as well due to an increase in attention. This analysis was done for all task frequencies together in order to investigate whether the task itself already affected these measures, and then for all task frequencies separately. In case of two-tailed tests, alpha was set to 0.025, rather than 0.05.

These cluster-based permutation tests were performed on both sensor-level data and source-level data. In case of sensor-level data, clusters were two-dimensional (space and neural frequency) and the minimum number of channels per cluster was set to 3. For source-level data, clusters were only computed in one dimension (neural frequency), but then for each source separately.

When taking all task frequencies together, the measurement (i.e., ITC or power) was first computed per task frequency, and then averaged across task frequencies. To test whether the frequency of the external rhythm affected the increase in synchrony, the task frequencies were contrasted by comparing the change in either ITC or power, rather than the raw values. Here, dependent samples F-statistics were computed for each sample in each permutation.

Reaction time and inter-trial coherence

As part of exploratory analyses to investigate the correlation between synchrony and behaviour, the relationship between RT and ITC was studied. Subjects were binned based on their RTs. Specifically, we subtracted the mean RT during f3 trials from the mean RT during f1 trials. We specifically looked at the difference in RT within each subject rather than at the raw RTs per subject, because the variability of the average RT between subjects could mitigate any potential effect, especially given the low number of subjects. Subjects were divided into two bins based on this RT effect, and the average ITC increase during the cue stream was computed.

Questionnaires

Subjects were asked about their sleepiness and rhythmicity per block. To investigate the effect of task frequency on sleepiness, we matched the responses with the task frequency of the preceding block. Thus, the first response to the sleepiness question, taking place before the start of the first block, was omitted. We performed repeated-measures ANOVAs on both the sleepiness and rhythmicity responses with the factor task frequencies and the levels f1, f2 and f3.

Results

Behaviour

When the task frequency increased, subjects responded faster (Figure 5A, $df = 2$, $F = 5.62$, $p = .01$). Post-hoc tests revealed that responses were significantly faster in case of f3 as compared to f1 ($t = 2.81$, $p = .047$). The differences between f1 and f2 ($t = 1.65$, $p = .126$) and between f2 and f3 ($t = 2.207$, $p = .095$) were not significant. Furthermore, the timing of the target affected RT (Figure 5B, $df = 2$, $F = 8.20$, $p = .002$). Here, post-hoc test revealed that subjects responded significantly slower in case of cycle1, relative to cycle2 ($t = 2.82$, $p = .031$), as well as to cycle3 ($t = 4.93$, $p = .001$). There was no significant difference in RT between cycle2 and cycle3 ($t = 0.43$, $p = .676$). Likewise, there was no interaction between target timing and task frequency (Figure 5C, $df = 4$, $F = 0.30$, $p = .88$).

As opposed to RT, accuracy was not affected
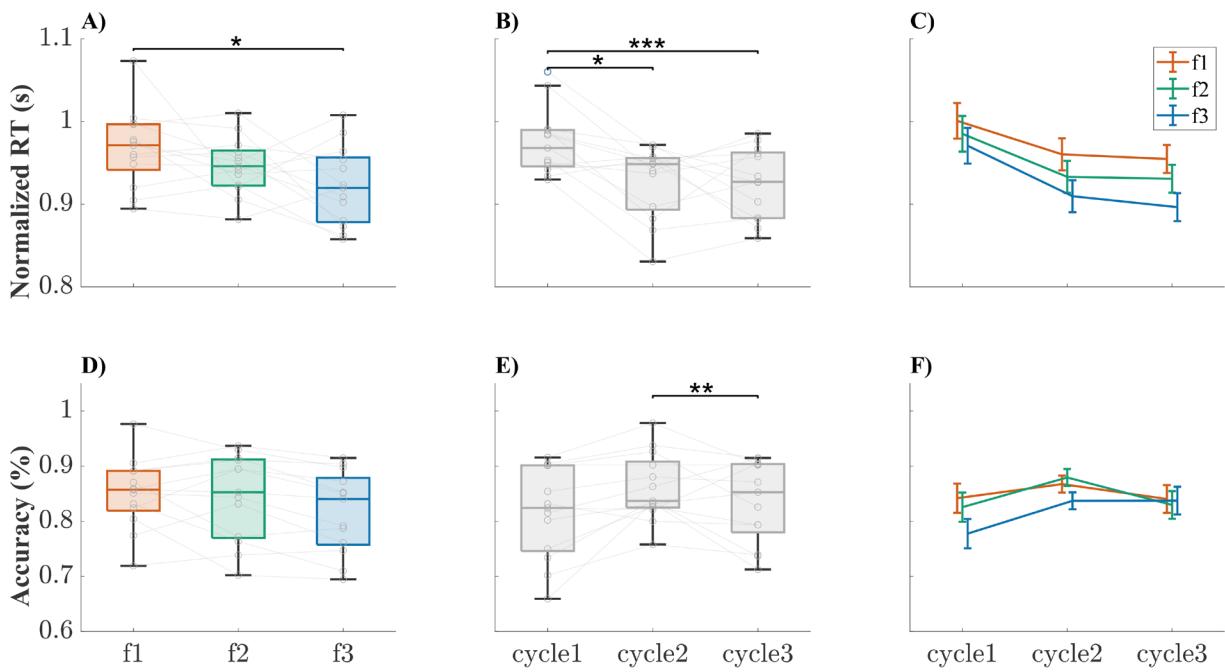
Figure 5. Behavioural results per task frequency **(A,D)**, per target timing **(B,E)**, and interaction between task frequency and target timing **(C,F)**. **Top panel (A to C):** normalized RT. **Bottom panel (D to F):** accuracy. The grey lines between boxplots and the grey circles on the boxplots represent the results of individual subjects. Post-hoc tests followed a repeated-measures ANOVA when main effects were found. Significant differences following post-hoc dependent t-tests are indicated by stars (* p < .05, ** p < .01, *** p < .001). Error bars in C and F represent the standard error of the mean.

by task frequency, although it did trend toward significance (Figure 5D, $df = 1.36$, $F = 3.488$, $p = .069$, Greenhouse-Geisser corrected). Target timing did have a significant effect on accuracy (Figure 5E, $df = 2$, $F = 4.77$, $p = .018$). Post-hoc tests revealed that subjects responded correctly more often during cycle2 trials as compared to cycle3 trials ($t = 3.67$, $p = .010$), while the difference in accuracy between cycle1 and cycle2 ($t = 0.32$, $p = .753$) and between cycle1 and cycle3 ($t = 2.17$, $p = .101$) were not significant. There was again no interaction between task frequency and target timing (Figure 5F, $df = 4$, $F = 1.48$, $p = .223$).

Inter-trial coherence

During the cue window, ITC increased in both the visual and motor source (Figure 6A, E; $p < .001$ in the visual source, and $p = .015$ and $p = .039$ in the motor source) for a broad range of neural frequencies (visual: 1 to 7 Hz; motor: 1.2 to 4.8 Hz and 5 to 7 Hz). When separating the three task frequencies, clear peaks were visible in the visual source (Figure 6B-D). Moreover, the peaks matched the corresponding task frequency and its harmonics. The increase in ITC was significant for a broad range of neural frequencies (all $p < .001$, f1: 1 to 7 Hz, f2: 1.3 to 7 Hz, f3: 1 to 7 Hz). In the motor source,

peaks were visible for the separate task frequencies as well (Figure 6F-H). Significant clusters of increase in ITC were found in the motor source, although for a smaller range of frequencies as compared to the visual source (f1: 1.2 to 2.6 Hz, $p = .030$, 5.8 to 6.9 Hz, $p = .044$; f2: 1.3 to 4.6 Hz, $p = .004$; f3: 1.1 to 4.4 Hz, $p = .016$, 4.6 Hz to 6.9 Hz, $p = .019$).

During the target window, there was no significant increase in ITC when taking all task frequencies together (Figure 7A, E). Furthermore, separating the task frequencies did not reveal any clear peaks (Figure 7B-D, F-H). Based on visual inspection, there did seem to be a higher offset in the low frequency range for higher task frequencies. We therefore contrasted the increases in ITC per task frequency (Figure 8). Indeed, it looked like a higher task frequency gave a higher increase in ITC around 2 Hz, especially in the visual source, but none of these differences were significant. Interestingly, the difference was significant at the sensor-level data (Figure 9), giving one cluster at the right side (1.4 Hz to 1.9 Hz, $p = .008$) and one at occipital sensors (1.4 Hz to 1.9 Hz, $p = .040$). Here, the planar MEG gradients had been computed and combined with the axial data before analysis. This implies that the signals were underneath the sensors that picked them up. Post-hoc tests revealed that in both clusters, the increase in ITC was significantly higher for f3 as
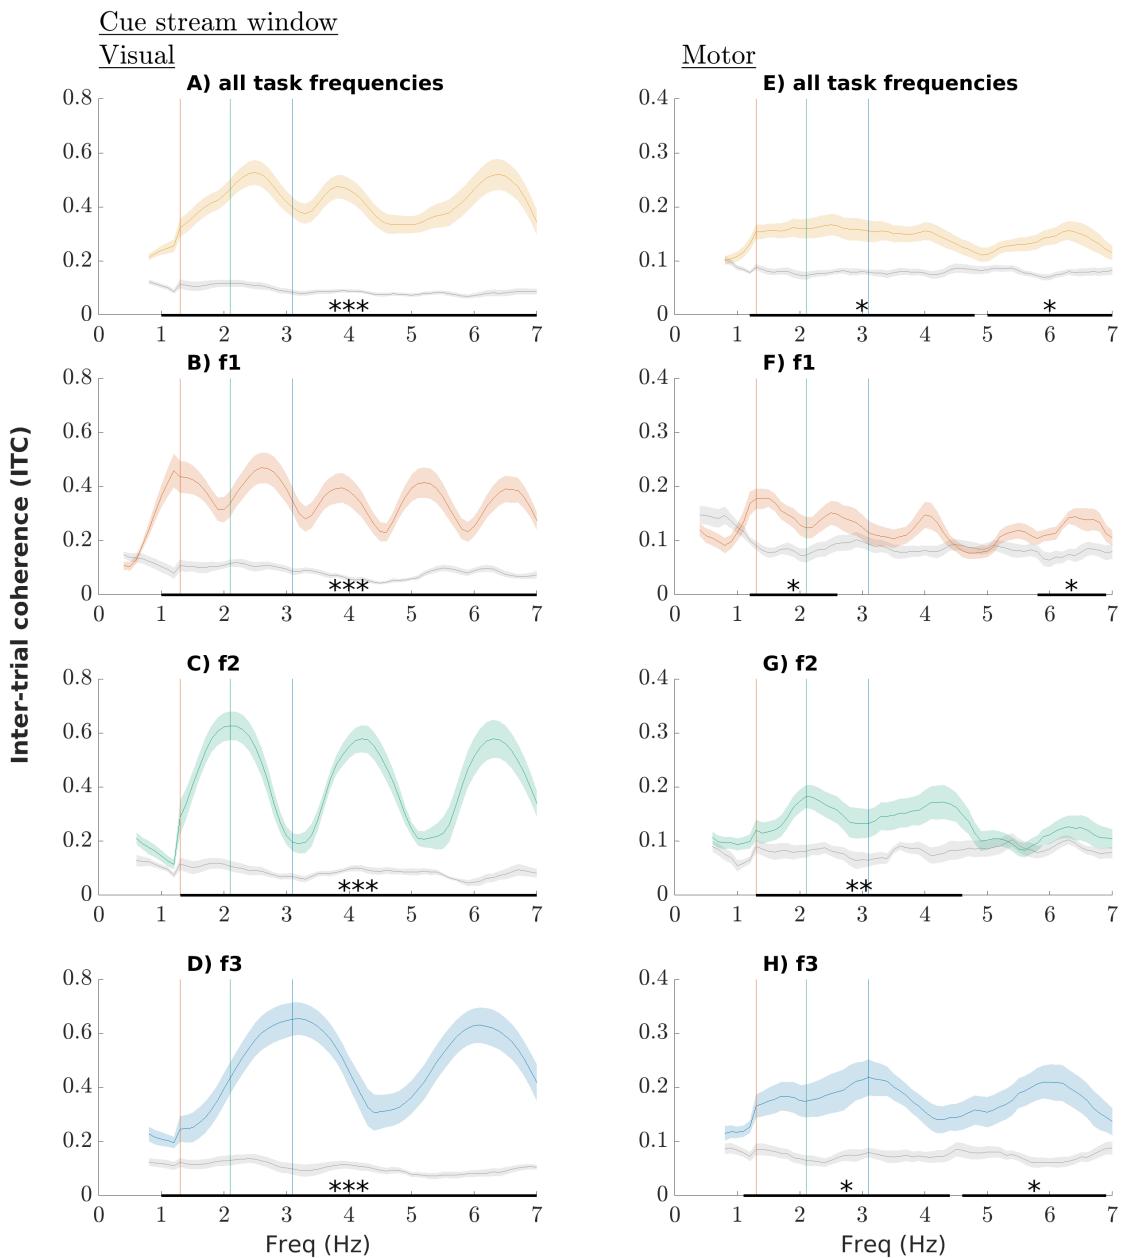
Figure 6. Raw ITC during the cue stream window. The coloured lines represent the ITC during the cue window, while the grey lines represent the ITC during the corresponding baseline window. Shaded areas represent the standard error of the mean. The vertical coloured lines correspond to the task frequencies (red: f1, 1.3 Hz; green: f2, 2.1 Hz; blue: f3, 3.1 Hz). Significant differences following cluster-based permutation tests are indicated by stars (* p < .05, ** p < .01, *** p < .001). **A to D:** visual source ITC for respectively all task frequencies together, f1, f2, and f3. **E to H:** motor source ITC for respectively all task frequencies together, f1, f2, and f3.

compared to f2 (right: *p* < .001, occipital: *p* = .002) and as compared to f1 (right: *p* < .001, occipital: *p* < 0.001). The differences between f1 and f2 were not significant (right: *p* = .542, occipital: *p* = .060)

Power

As compared to baseline, power decreased in the visual source between 2.3 and 4.2 Hz during the cue window (Figure 10A, *p* = .029). There was no

significant difference in the motor source (Figure 10B). When separating the task frequencies, clear peaks in power were present in the visual source that corresponded to the task frequency and its harmonics (Figure 10B-D). These peaks were absent in the motor source (Figure 10F-H). Only the first peak of f3 in the visual source was significant (2.2 to 4 Hz, *p* = .014).

During the target window, power in both the visual source and the motor source generally
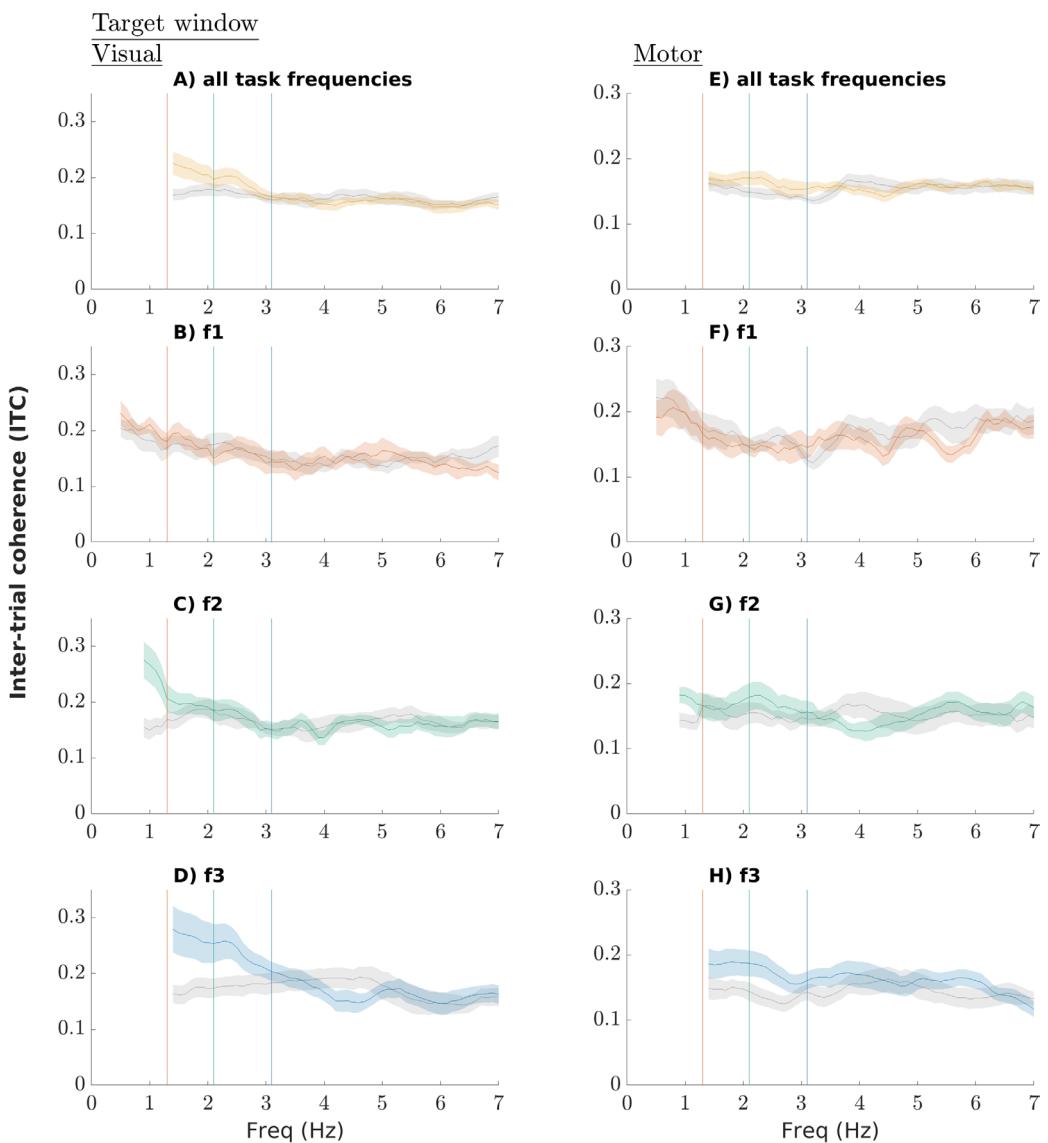
Figure 7. Raw ITC during the target window. The coloured lines represent the ITC during the target window, while the grey lines represent the ITC during the corresponding baseline window. Shaded areas represent the standard error of the mean. The vertical coloured lines correspond to the task frequencies (red: f1, 1.3 Hz; green: f2, 2.1 Hz; blue: f3, 3.1 Hz). No significant differences followed from cluster-based permutation tests. **A to D:** visual source ITC for respectively all task frequencies together, f1, f2, and f3. **E to H:** motor source ITC for respectively all task frequencies together, f1, f2, and f3.
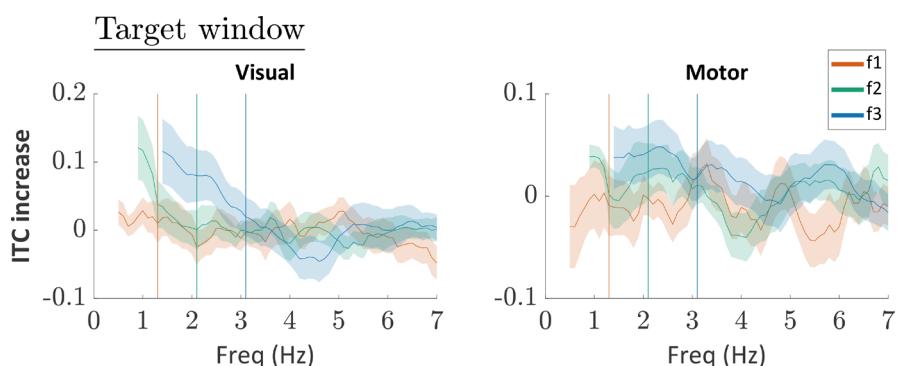
Figure 8. Increase in ITC during the target window at source-level, separately for each task frequency. Shaded areas represent the standard error of the mean. The vertical coloured lines correspond to the task frequencies (red: f1, 1.3 Hz; green: f2, 2.1 Hz; blue: f3, 3.1 Hz). Cluster-based permutation tests did not reveal significant differences. **Left:** increase in ITC in the visual source. **Right:** increase in ITC in the motor source.
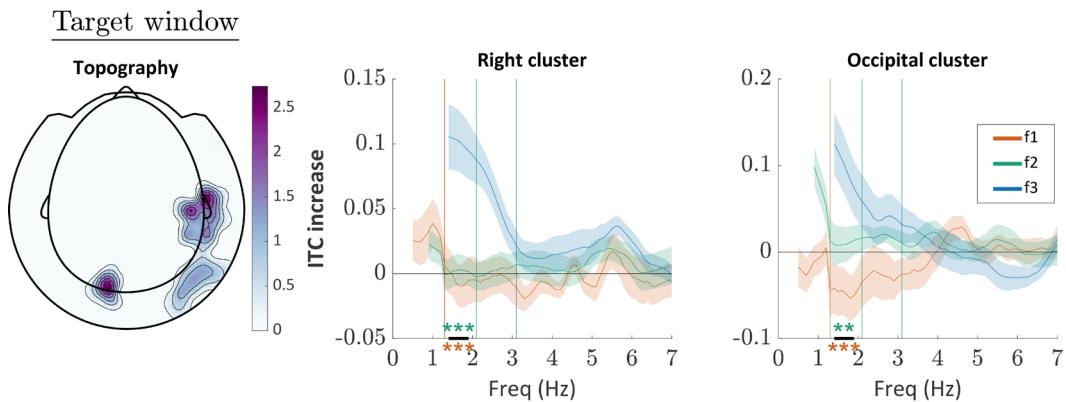
Figure 9. Increase in ITC during the target window at sensor-level, separately for each task frequency. Shaded areas represent the standard error of the mean. The vertical coloured lines correspond to the task frequencies (red: f1, 1.3 Hz; green: f2, 2.1 Hz; blue: f3, 3.1 Hz). **Left:** topography of the masked statistics of the two significant clusters (F-values, cluster-based permutation). **Middle and right:** ITC increase per cluster, averaged across the sensors that are part of the cluster. Significant differences following post-hoc dependent t-tests are indicated by stars (* $p < .05$, ** $p < .01$, *** $p < .001$). The green stars represent the p-values of the f3-f2 contrast, while the red stars represent those of the f3-f1 contrast.
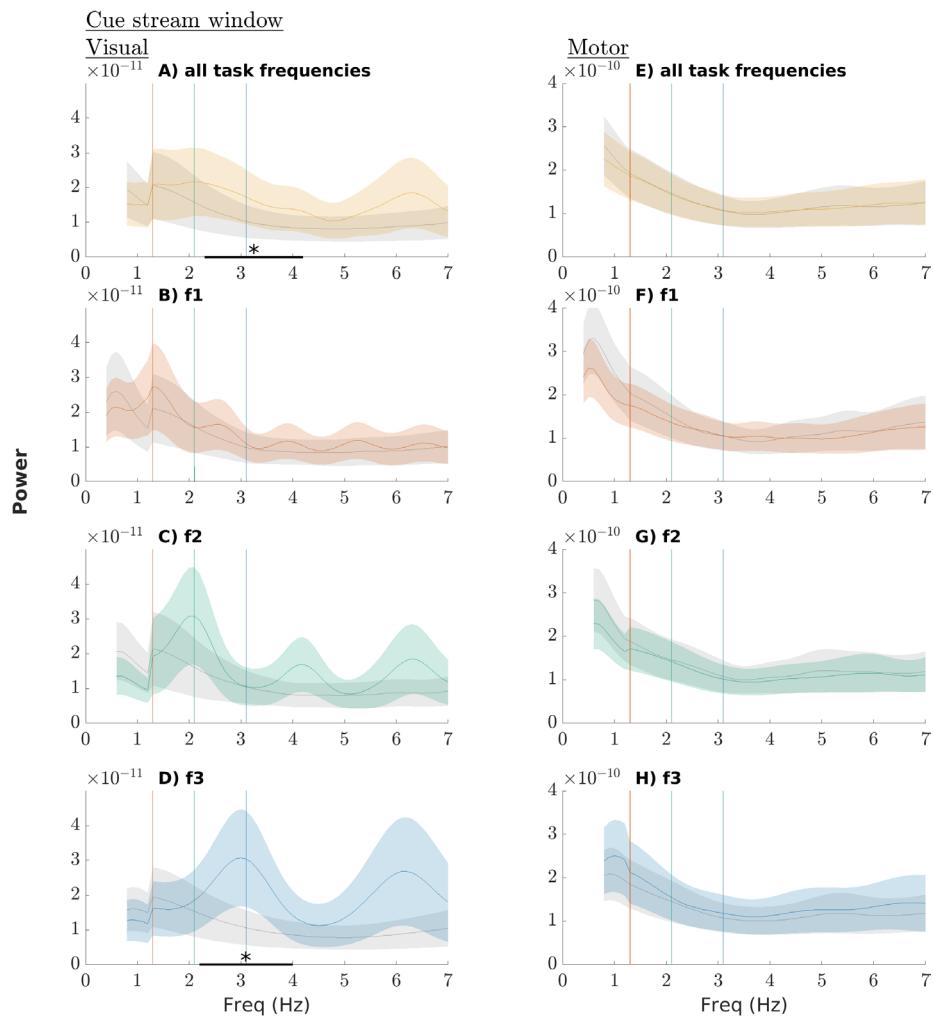
Figure 10. Raw power during the cue stream window. The coloured lines represent the power during the cue window, while the grey lines represent the power during the corresponding baseline window. Shaded areas represent the standard error of the mean. The vertical coloured lines correspond to the task frequencies (red: f1, 1.3 Hz; green: f2, 2.1 Hz; blue: f3, 3.1 Hz). Significant differences following cluster-based permutation tests are indicated by stars (* $p < .05$, ** $p < .01$, *** $p < .001$). **A to D:** visual source power for respectively all task frequencies together, f1, f2, and f3. **E to H:** motor source power for respectively all task frequencies together, f1, f2, and f3.

decreased (Figure 11A, E), with both sources having a significant cluster (visual: 2.8 to 5.3 Hz, p = .024; motor: 1 to 4.3 Hz, p = .007). When separating the task frequencies, this general decrease in power seemed to be present for each task frequency in both the visual source (Figure 11B-D) and the motor source (Figure 11B-D) based on visual inspection, although only the differences in the motor source during f1 trials (1 to 3.9 Hz, p = .005) and in the visual source during f3 trials (1.4 Hz to 7 Hz, p < .001) were significant.

Functional connectivity

Functional connectivity between the visual source and the motor source increased significantly during the cue stream as compared to baseline in the range of 6 to 8 Hz (Figure 12A, p = .025). When separating the task frequencies, such an increase was present at roughly the same frequency range for each task frequency (Figure 12B-D, f1: 6 to 7 Hz, f2: 7 Hz to 8 Hz, f3: 6 to 8 Hz), although only the cluster during f3 trials was significant (p = .024), while those during f1 and f2 trials only showed a trend toward significance (f1: p = .076, f2: p = .094).

There was no significant increase in functional connectivity during the target window when taking all task frequencies together (Figure 12E). However, when separating the task frequencies (Figure 12F-H), functional connectivity increased significantly during f3 trials between 7 and 10 Hz (p = .003). There was no trend toward significance for f1 and f2, unlike during the cue window.
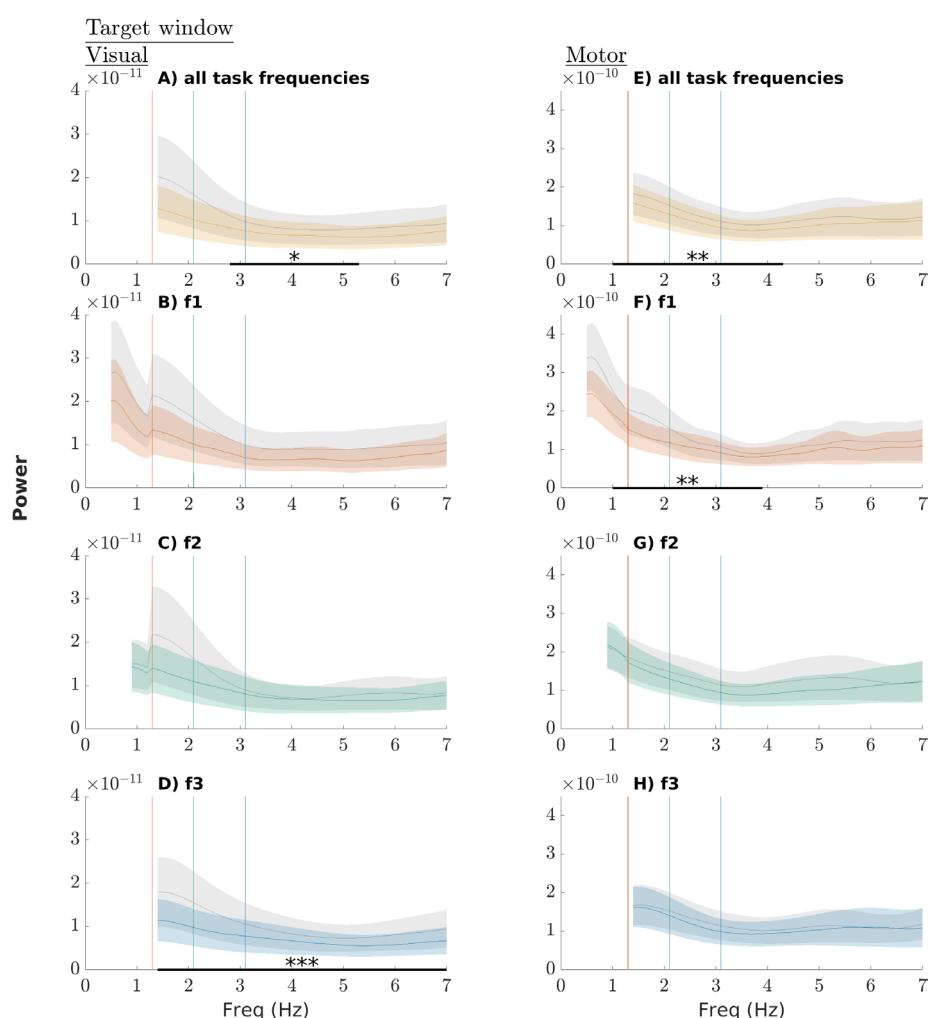

Figure 11. Raw power during the target window. The coloured lines represent the power during the target window, while the grey lines represent the power during the corresponding baseline window. Shaded areas represent the standard error of the mean. The vertical coloured lines correspond to the task frequencies (red: f1, 1.3 Hz; green: f2, 2.1 Hz; blue: f3, 3.1 Hz). Significant differences following cluster-based permutation tests are indicated by stars (* p < .05, ** p < .01, *** p < .001). **A to D:** visual source power for respectively all task frequencies together, f1, f2, and f3. **E to H:** motor source power for respectively all task frequencies together, f1, f2, and f3.

Reaction time and inter-trial coherence

Based on visual inspection, the increase in ITC was consistently higher for subjects with a smaller RT effect than those with a larger RT effect, except in the motor source during f1 trials (Figure 13). Furthermore, this difference seemed to be more extreme for higher task frequencies. No statistical test was performed on this data.

Sleepiness and rhythmicity

Task frequency affected neither sleepiness ($df = 2$, $F = 0.388$, $p = .683$), nor rhythmicity ($df = 2$, $F = 0.889$, $p = .427$).

Discussion

In this study we investigated the notions of entrainment and active sensing, which constitute potential mechanisms for sampling the relevant input from the environment. To study these potential mechanisms, subjects were presented with a visual discrimination task, in which the target was preceded by a visual rhythmic cue stream. This rhythmic cue stream could have three different task frequencies, all falling within the delta range, which is thought to play a role in sensory sampling (Schroeder & Lakatos, 2009). Based on a series of unpublished experiments (manuscript in preparation), we expected to find a decrease in RT with increasing
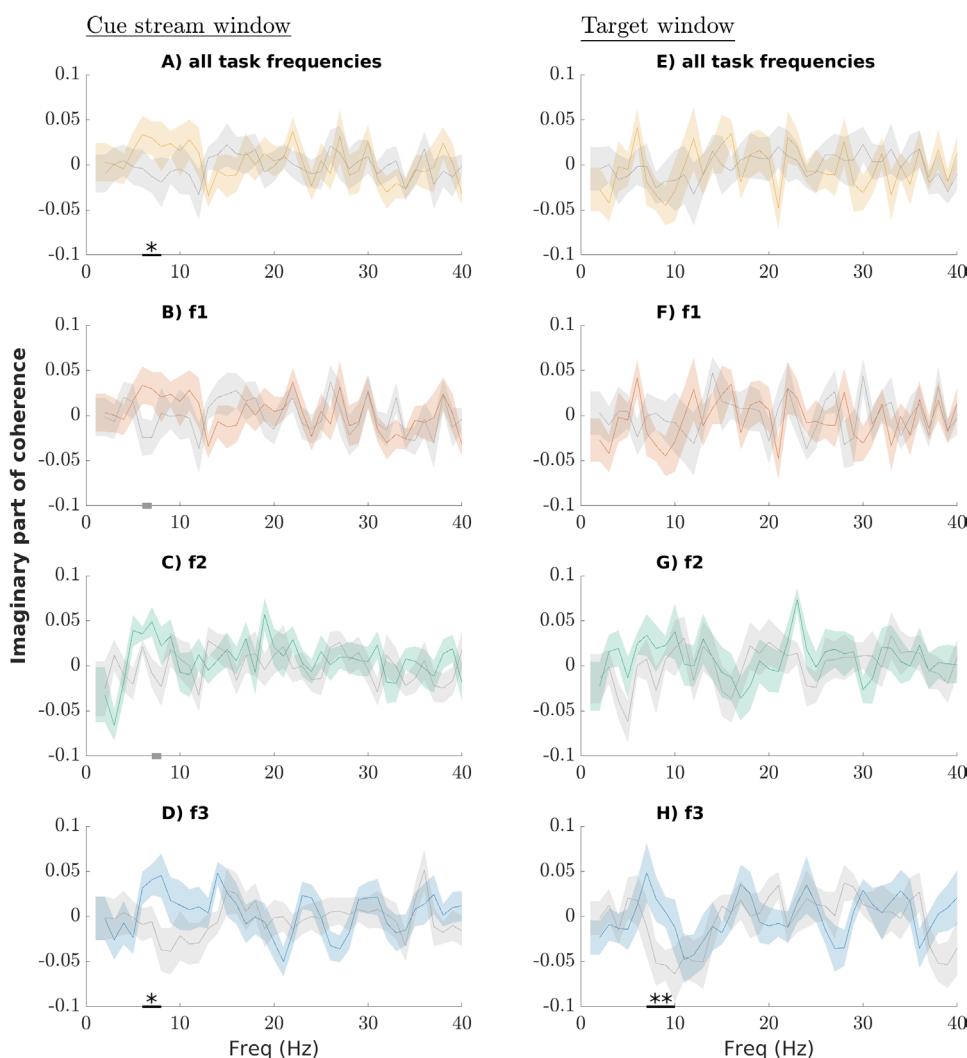
Figure 12. Imaginary part of the coherence (i.e., functional connectivity). The coloured lines represent the functional connectivity during the window of interest, while the grey lines represent the functional connectivity during the corresponding baseline window. Shaded areas represent the standard error of the mean. Significant differences following cluster-based permutation tests are indicated by stars (* p < .05, ** p < .01, *** p < .001). Grey bars on the x-axis are clusters that trend toward significance (p-value between .05 and .1). **A to D:** functional connectivity during the cue stream window for respectively all task frequencies together, f1, f2, and f3. **E to H:** functional connectivity during the target window for respectively all task frequencies together, f1, f2, and f3.
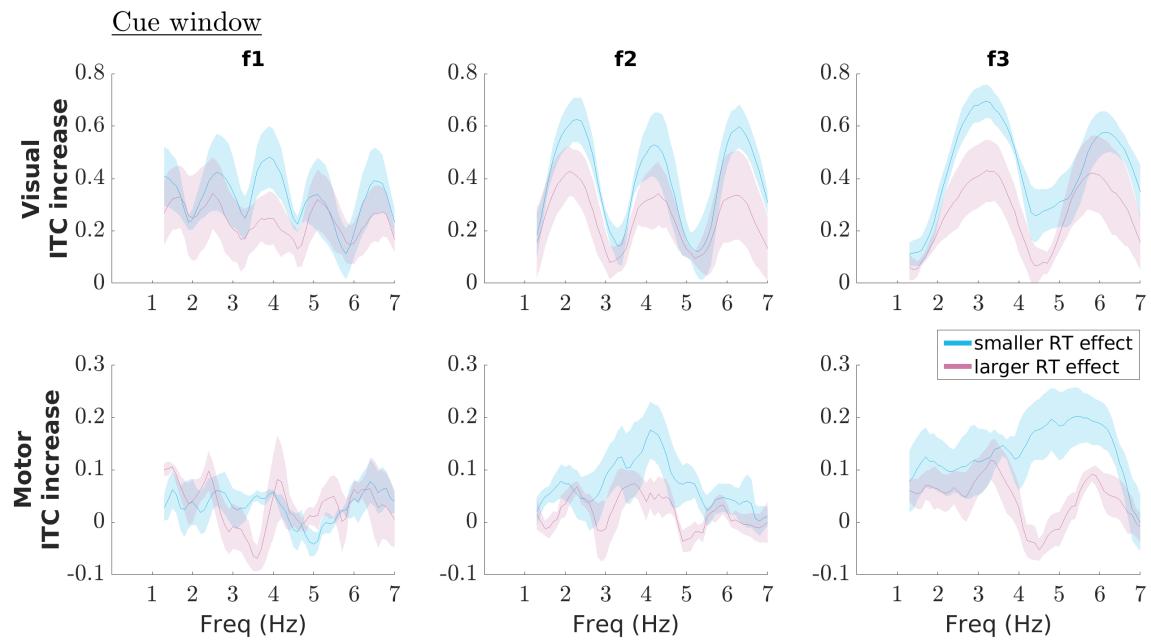
Figure 13. ITC increase during the cue window as compared to the baseline window, separately for subjects with a small RT effect and subjects with a large RT effect. Here, the RT effect is defined as subtracting the mean RT across f3 trials from the mean RT across f1 trials. Shaded areas represent the standard error of the mean. No statistical analysis was performed on this data.

task frequency. As this behavioural effect would be supported by an increase in synchrony, we further expected an increase in phase-alignment in both the visual source and the motor source, and between these two sources. We indeed replicated the finding of decreasing RT with increasing task frequency. We further measured peaks in synchrony to the external rhythm. The neural frequency of these peaks corresponded to the task frequency. However, shortly after the rhythm, these peaks in phase-alignment diminished. We further found an increase in functional connectivity between the visual source and the motor source, but in a different neural frequency range than expected. Unfortunately, we were only able to collect data from 13 subjects. Therefore, more data should be gathered before drawing final conclusions. Importantly, 5 out of 13 subjects had a different baseline period length. This difference could slightly affect the results by giving a lower frequency resolution, but importantly this further implies that these 5 subjects had the same number of trials in less time, potentially affecting their attention with respect to the 8 subjects with a longer baseline period.

Faster responses to faster rhythms

We found that subjects responded faster in a visual discrimination task when the task frequency increased, which is a replication of a series of previous experiments with an auditory discrimination

task (manuscript in preparation). Importantly, here we demonstrated the effect in the visual domain rather than the auditory domain, making this a novel behavioural finding. One possible explanation is that there was a trade-off between RT and accuracy, with RT being more prioritised for higher task frequencies. However, we did not find an effect of task frequency on accuracy. We therefore conclude that the effect of task frequency on RT is not the result of a trade-off between RT and accuracy.

One could further argue that higher task frequencies make subjects more alert, possibly decreasing their RT. However, we did not find an effect of task frequency on the subjects' sleepiness rating and therefore conclude that alertness is not different for different task frequencies.

Another concern is that some task frequencies might be perceived as being more rhythmic than others. Again, we did not find an effect of task frequency on the subjects' responses to the perceived rhythmicity question.

Furthermore, we found an increase in RT when the timing of the target was early, as compared to later time points. Possibly this is an effect of the Hazard rate (Näätänen, 1971), entailing that there is less time-uncertainty for the target at later time points, as the probability of the target occurring increased when more time had passed. More uncertainty in time appears to slow down the response (Niemi & Näätänen, 1981). The accuracy however did not show such an effect. Here we would expect to find

higher accuracy for later time points, but instead we found that accuracy increased when the target was presented after two cycles as compared to when it was presented after three cycles. Possibly subjects were inclined to prioritise RT over accuracy when the target appeared later in time, potentially because of growing impatience for three cycles. Another possible explanation for RT being lower for later time points, is that only one cycle is not yet enough time for the motor system to be fully prepared, while two cycles is enough, since RTs become significantly faster from two cycles onward.

One remaining question is whether an increase of general pace of the task could have this effect on RT, rather than the task frequency itself. Namely, if the task frequency increases, the duration of a full trial decreases. As a result, the target is presented more often in the same amount of time, as compared to lower task frequencies, increasing the general pace of the experiment. It would therefore be interesting in a future study to examine the effect of the task frequency on RT, but instead with a constant time between the start of one trial and the start of the next trial.

Entrainment

During the cue window, we found clear peaks in both ITC and power as compared to baseline. These peaks appeared at the task frequency and its harmonics. This finding is in line with the frequency specific hypothesis as described in Figure 2, in which the peak in power or ITC shifts toward higher neural frequencies for higher task frequencies. The same ITC effect was also found by Will and Berg (2007) in the auditory domain. However, we did not find this tendency for either ITC or power in the period after the rhythm had been presented. At the sensor-level, we did find a higher ITC increase in the delta range for the highest task frequency as compared to the lower ones. We investigated whether this increase could be a result of a difference in evoked responses between task frequencies, but no such difference was found. We also found that delta and theta power decreased after the presentation of the rhythm and before target onset in both the visual source and motor source. This could be an effect of attention (Fries, Reynolds, Rorie, & Desimone, 2001), indicating that subjects attended more when expecting a target stimulus.

We can only speak of true entrainment when three requirements are met (Haegens, 2020; Haegens & Golumbic, 2018; discussed in the section on neural oscillations). First, an endogenous neural oscillator

must be present apart from rhythmic stimulation. In the current study, this would be during the baseline period, but investigating the presence of such oscillations is outside the scope of this study. Second, the neural oscillator phase-aligns with the external rhythm, and this only happens for a limited range of frequencies. In this study, the ITC and power showed peaks at exactly the task frequency and its harmonics during the presentation of the external rhythm, which is in line with this requirement. Because of the limited number of task frequencies used in this study, it is unclear whether this only happens for a limited range of task frequencies. Importantly, when investigating ITC and power during the presentation of a rhythm, it is impossible to distinguish between true oscillations and a series of evoked responses. Third, the phase-alignment must continue for some number of cycles after the presentation of the external rhythm. This however we did not find, violating the last requirement. We therefore conclude that entrainment, as defined above, is not happening under the circumstances of this experiment. This conclusion contradicts the general idea that neural oscillations entrain to external rhythms within a wide range of stimuli, tasks, and neural frequencies (as reviewed by Lakatos, Gross, & Thut, 2019), and is instead more in line with critical reviews such as Helfrich, Breska, and Knight (2019), Obleser and Kayser (2019), Haegens and Golumbic (2018) and Zoefel et al. (2018). Instead, the peaks in ITC and power as found during the presentation of a rhythmic cue, could be a result of evoked responses. Even though the last requirement of entrainment was violated, it would still be interesting to see whether the first two requirements hold or not by studying the oscillations before the onset of the rhythm and increasing the number of task frequencies. Importantly, the lack of support for entrainment only applies to the context of the experiment of this study. Instead, possibly entrainment still plays a role in other contexts. For instance, one can change the sensory domain (e.g., auditory rather than visual), task frequencies (e.g., up to 10 Hz rather than up to 3.1 Hz) and task (e.g., detection rather than discrimination). Another approach is to use near-threshold stimuli, since the effect of the timing relative to the phase of neural oscillations are higher for such stimuli (Schroeder & Lakatos, 2009).

As part of exploratory analysis, subjects were binned into two groups based on the difference between their mean RT for the highest task frequency and that of the lowest task frequency. The 'large RT effect' bin consisted of those subjects that

showed a bigger decrease in RT with increasing task frequency. The ITC increase seemed to be lower for subjects with a large RT effect than those with a small RT effect. An increase in ITC implies that there is less jitter across trials. Possibly a jitter across trials affects to which extent RT changes when the task frequency increases. Namely, if there is zero jitter in the visual system, motor system and in the communication between these two systems, stimuli that are in-phase with a rhythm, as was always the case in the experiment of the current study, will always be optimally processed, which then gives the same RT regardless of the frequency of the rhythm. This is under the assumption that the intrinsic oscillations correctly adjust to the external rhythm by either entrainment or another mechanism, such that high excitability states coincide with stimuli that are presented an integer number of cycles after the rhythm. Instead, if there is jitter, the chances of neural input arriving at a state of high excitability is higher for higher task frequencies due to a higher sampling rate, as argued in the section on research questions and hypotheses. Thus, when there is more jitter, ITC decreases and the RT effect is larger. This possible interaction between RT and ITC should be further investigated.

Important to note here, is that the excitability of the neurons only depend on the phase of the oscillation, and thus in theory the high excitability states are longer for lower task frequencies, but occur more often for higher task frequencies. This would also imply that a small offset of the input with respect to the maximum excitability state (e.g., 50 ms) would decrease the chance of the input arriving during the same excitability state for higher task frequencies. However, for big offsets (e.g., 200 ms) the chances of the input arriving at the next excitability state is higher for higher task frequencies, because the next high excitability state occurs earlier in time, giving the higher sampling rate as mentioned above. An important aspect that should be further investigated before drawing conclusions about the effect of jitter, is the range of this potential jitter (e.g., more than one cycle or not) and the distribution (i.e., flat or biased).

Covert active sensing

While we expected to find an increase in functional connectivity in the beta range, as it could establish the connection between the motor cortex and the sensory cortex (Spitzer & Haegens, 2017), we instead found this increase in the delta-theta range. This increase as compared to baseline was present during both the presentation of the external rhythm and shortly after it, in particular for the highest task frequency. The theta-delta range has been shown to play a major role in sensory sampling (Schroeder & Lakatos, 2009). Delta phases are thought to have a modulating effect on beta bursts (for a recent review, see Morillon, Arnal, Schroeder, & Keitel (2019)), while theta phase may modulate gamma power (Canolty et al., 2006). In this study, we specifically focused on the phase alignment, rather than power, as a measure of functional connectivity by looking at the imaginary part of coherence, which may be why the possible effect on beta power was undetected.

Covert active sensing entails that the motor system coordinates the brain signals in the sensory cortex by the use of synchronisation. We therefore expected to find an increase in functional connectivity during the task, which we found both during the presentation of an external rhythm and shortly after it. To further investigate whether the motor cortex drives the signals in the sensory cortex, one could look at the Wiener-Granger causality (Bressler & Seth, 2011) between the two sources. A next step would be to use neuromodulation methodologies that make it possible to directly control the signals in the motor cortex, such as trans-cranial magnetic stimulation (TMS) in human participants or optogenetics in animals.

Importance for AI

To discuss how our findings and those of other studies could help us to make a step toward tackling the problem of perceptual relevance in AI, we focused on the following questions:
- How could the findings of the current study be used as a source of inspiration when tackling the problem of perceptual relevance in the development of artificial cognitive agents?
- How could oscillatory mechanisms be realised in artificial systems?
- What are some good practices when combining empirical research within the fields of AI and neuroscience, such that both fields benefit?

Time dimension. Before discussing the potential uses of oscillations in AI, the possible use of a time dimension should be discussed first. Conventional AI, such as machine learning, artificial neural networks, and Bayesian networks, have been able to achieve some impressive goals but mostly ignore the dimension of time. The scope of problems that can be solved by AI may increase by including this dimension, as timing plays a major role

in how natural intelligence processes information, as discussed in the current study. Time can still roughly be used by these conventional methods, as done by for example Kucherenko, Hasegawa, Henter, Kaneko and Kjellström (2019). Their artificial neural network took auditory time series (speech) as input and predicted which gestures over time would fit those time series. It contained gated recurrent units (GRUs), which uses information from previous input, such as earlier time points, and combines them with the current input. Thus, even though both input and output contain a time dimension, the underlying computations only look at the order in which the input is presented, rather than actual timing.

As opposed to these conventional AI techniques, spiking neural networks (SNNs; Maass, 1997) resemble natural neural networks much more by outputting 'spikes' over time depending on the input, just like natural neurons fire action potentials over time. This resemblance makes it possible to incorporate the time dimension within the network. This way, the same input to an artificial neuron can elicit a different response in the network when presented at a different moment in time.

Simulating SNNs on conventional computers is computationally expensive. However, because of more recent developments in the field of neuromorphic hardware, efficiency is greatly increasing. In neuromorphic hardware, information is represented by the use of spikes over time, rather than a series of zeros and ones as in conventional computers. For instance, Intel recently developed a neuromorphic chip called Loihi. This chip was used by Imam and Cleland (2020) to implement rapid online learning. Their algorithm, implemented on Loihi, was able to learn new odours and reliably recognise them after only one presentation. Such achievements are not possible for conventional methods, as they require a substantial amount of data to train the parameters. These developments in the field of neuromorphic computing therefore make the exploration of timing and specifically oscillations in AI an interesting next step in research.

Oscillations. Neural oscillations are rhythmic fluctuations in neural activity. They include the intra- or extracellular electrical currents, which are affected by action potentials.

The concept of oscillations could still be used in conventional AI without having a real time dimension. For instance, Neil, Pfeiffer and Liu (2016) proposed that instead of (conventional) long short-term memory networks (LSTMs), phased

LSTMs should be used. Just like GRUs, LSTMs combine the current input with information about previous input. Their new phased LSTM model contained extra gates that could be either open or closed. Whether a gate was open or closed, was determined by an oscillation, in the sense that it was open for some time points, then closed for other time points, and this pattern was repeated over time. Information through those gates could only pass when the gate was open. This implementation is thus in line with Bishop's proposal of oscillations reflecting fluctuations in excitability. The new model is able to integrate input from sensors of different sampling rates and has a general improvement in performance. However, the implementation is vastly simplified, as is evidenced by having binary gates (i.e., having only 'open' and 'closed' states), while Bishop's excitability states are not binary.

Such novel implementations can increase the scope of problems that can be solved by AI or increase the performance on problems that are already addressed to some extent. In the following discussion, we will focus on more biologically plausible implementations to be able to make use of the exact timing at which input arrives.

The next question is whether and how the notion of oscillations can be used in SNNs and neuromorphic computing. Pfeiffer and Pfeil (2018) already discussed that SNNs should be able to model the timing of reference spikes relative to network oscillations. We are not aware of any implementation that includes such network oscillations. However, since these networks and hardware closely mimic the electrical activity of natural neurons, oscillations may already be present as an emergent property. Before implementing specifically active sensing or entrainment, these potential properties of SNNs and neuromorphic hardware should be studied. For the remaining part of this discussion, we will assume that oscillations can be present in SNNs, either intentionally implemented or as an emergent property. We further assume that these oscillations are either within the artificial spiking neurons or outside of them while being able to influence them.

Perceptual relevance. Starting with Bishop's proposal of oscillations representing fluctuations in excitability states, artificial spiking neurons should combine the information of this oscillation with the input it receives. The same input to the neuron should be more likely to elicit a spike for some timings relative to the phase of the oscillations, as compared to other timings. This does not have to hold for all input, but at least for input that is close

to threshold. If artificial oscillations are present within the artificial neuron, the artificial electrical current of that neuron fluctuates. We then expect the property of fluctuations in excitability to already be present as an emerging property, as Bishop's proposal was based on precisely this property of intracellular fluctuations in electrical current.

In order to use this proposal to tackle the problem of perceptual relevance, which is defined as determining which sensory input is relevant in the current context, the oscillations should be adjusted accordingly. In this study, we hypothesised that entrainment could underlie this adjustment by first exactly matching the intrinsic oscillation to the external rhythm, continuing this adjusted oscillation for some time. However, our results only showed clear peaks in power and ITC during the rhythmic input. These peaks were not present anymore shortly after the presentation of this external rhythm. We therefore believe that other potential mechanism could be more promising in tackling the problem of perceptual relevance. Instead of entrainment, the notion of rhythmic tracking could be further investigated. Rhythmic tracking is defined as a rhythmic neural response to rhythmic input (Haegens, 2020). This term does not imply exact phase-alignment of an intrinsic oscillation to an external rhythm and neither that the response should continue for some amount of time after the presentation of the external rhythm. Perhaps neural responses at frequencies higher than the delta and theta range could have a function here, which could be further investigated using the current dataset.

The second potential mechanism we studied was covert active sensing. This mechanism is more promising, as we did find an increase in functional connectivity between the motor source and the visual source both during the presentation of an external visual rhythm and briefly after it. A potential reason is that perhaps the presence of a rhythm increases the communication rate between the motor cortex and the visual cortex. Moreover, a faster external rhythm elicits higher communication rates. This difference in communication rate could be caused by oscillations adjusting differently to the external rhythm. Therefore, making the artificial oscillations in an SNN adjust to an external rhythm is an interesting first step, followed by an increase in communication rate. Here, this adjustment does not have to exactly match the frequency of the external rhythm, as is implied by entrainment. Instead, the oscillations could acquire an increase in functional connectivity within a certain frequency range (e.g., 6-10 Hz, as in our findings) but with a higher increase

in synchrony for faster external rhythms.

This proposal of implementing active sensing in SNNs is thus not so much about filtering perceptual input before it is processed in the network, but rather about changing the communication between multiple artificial neuronal populations within the network. This could still contribute to tackling the problem of perceptual relevance, as the filtering of input does not have to happen at the first level of the network. Filtering may take place while activity goes from one neuronal population to the other.

The frame and abstraction problem. Above weonly discussed possible implementations in AI that could tackle the problem of relevance. Therefore, this only includes information that is currently being perceived by an agent, and not any already existing knowledge. However, there are types of knowledge that have to be filtered other than the information that is currently being perceived by an agent. These types of knowledge are factual knowledge that was gained in the past, information about the current context, and predictions about events that are not directly perceived. By extending the scope of information to include these aspects as well, the problem of abstraction and the frame problem, explained in the section on computational problems of relevance in AI, are included.

In case of purely perceptual relevance, the proposal of oscillations reflecting fluctuations in excitability gives a straightforward possible mechanism for filtering input: if the sensory input arrives at the wrong timing, it is filtered out. However, when including factual knowledge about the past, the current context and expectations, there is no direct sensory input that can be filtered depending on the timing. Instead, all of the information comes from within the system. The information could be filtered by the use of oscillations when the information flows from one area to another, but then a new important question arises: when does the information start to flow and how is this flow elicited? As this idea of using oscillations as a filter mechanism gives rise to more questions than potential answers, further research is needed. How to effectively perform such research in which AI and neuroscience are combined, is discussed in the next section.

Combining AI and neuroscience research

AI and neuroscience are two closely related fields. They can learn from each other but without an active interaction, progress is slow. In this section,

we discuss how the two fields can contribute to each other and how to effectively combine research in the fields of AI and neuroscience.

The use of AI research in neuroscience. Interestingly, computational problems such as the original frame problem (i.e., how to efficiently determine which aspects in the world do *not* change) are not obvious when looking at natural intelligence but instead are brought to light by the novel methodology that is used by AI (Dennett, 2006). Namely, in natural intelligence we start with already existing intelligence when trying to understand it. Instead, in AI one has to build intelligence from the ground up. This way, philosophical and computational problems are encountered that are not too clear when looking at the finished picture (i.e., natural intelligence). This is one of the ways in which AI research can contribute to the field of neuroscience: defining problems and questions that are undetected when studying natural intelligence.

AI research can further contribute to the field of neuroscience by testing and developing hypotheses from a different perspective. By implementing the hypotheses, one can investigate the behaviour of the system given the hypothesis. Furthermore, parameters can be controlled to learn more about the effects of these parameters. Not only does this give more controlled studies, which is often not ethically, or practically, possible on subjects, but it also provides a more efficient way of testing and adjusting hypotheses before testing on subjects, which requires more resources. AI can go further than merely virtual simulations by having physical platforms, such as robots, embedded in a rich environment. This change in environment is important, as virtual environments are simplified, possibly to the extent that pitfalls or obstacles remain undetected. Perhaps in the future this can be done by having a physical robot running on neuromorphic hardware and placed in the real world.

Another contribution of AI in the field of neuroscience is the development of data analysis methods, such as machine learning, to further investigate brain signal recordings. This kind of contribution is however not further discussed in this study. The fields of AI and neuroscience can benefit most from each other by having an active interaction between the two fields. This active interaction is described next.

Active interaction between AI and neuroscience. Figure 14 presents a schematic overview of effective interaction between the fields of neuroscience and AI. Each arrow is explained below.

1. If we have results found in the field of neuroscience, one can attempt to implement these findings in AI. For the findings of the current study, this step is discussed in the section on the importance for AI. See Hassabis, Kumaran, Summerfield and Botvinick (2017) for a review of neuroscience-inspired AI studies. The arrow is bidirectional, because active discussion is needed during the development of the implementation in order to stay close to neuroscience.

2. When implementing previous findings, one may encounter an aspect that has to be implemented to complete the implementation, while the current results and discussion in neuroscience do not suffice, creating a gap between neuroscience and AI. This is how the frame problem in the original sense came to light.

3. If such a gap is encountered, there are two ways in which it can be closed.

a. Go back to the field of neuroscience to discuss the gap and hypotheses to close it. Possibly new experiments can be conducted that are designed to answer a question in AI. This discussion can also lead to novel questions and insights within the field of neuroscience. After this discussion and possible new experiments, arrow 1 leads back to improving the implementation.

b. Instead of directly consulting neuroscience, one can also develop the implementation further by investigating the possibilities in the field of AI. This step should still be done with inspiration from neuroscience in order to have a biologically plausible implementation. By having a biologically plausible implementation, the results can contribute to neuroscience in later steps.

4. When the field of AI has found a working implementation that goes beyond the knowledge in neuroscience, the implementation could be studied to investigate why it works and question whether natural intelligence could be using such an implementation as well.

5. The resulting implementation can be used in a simulation to gather results about the behaviour of the system.

6. Then comparisons can be made between the implementation in AI and the findings in neuroscience. This comparison not only pertains to the results of the simulation, but also to the actual implementation itself.

7. The comparison in turn can result in novel insights in the field of neuroscience, creating new research questions and hypotheses that can further
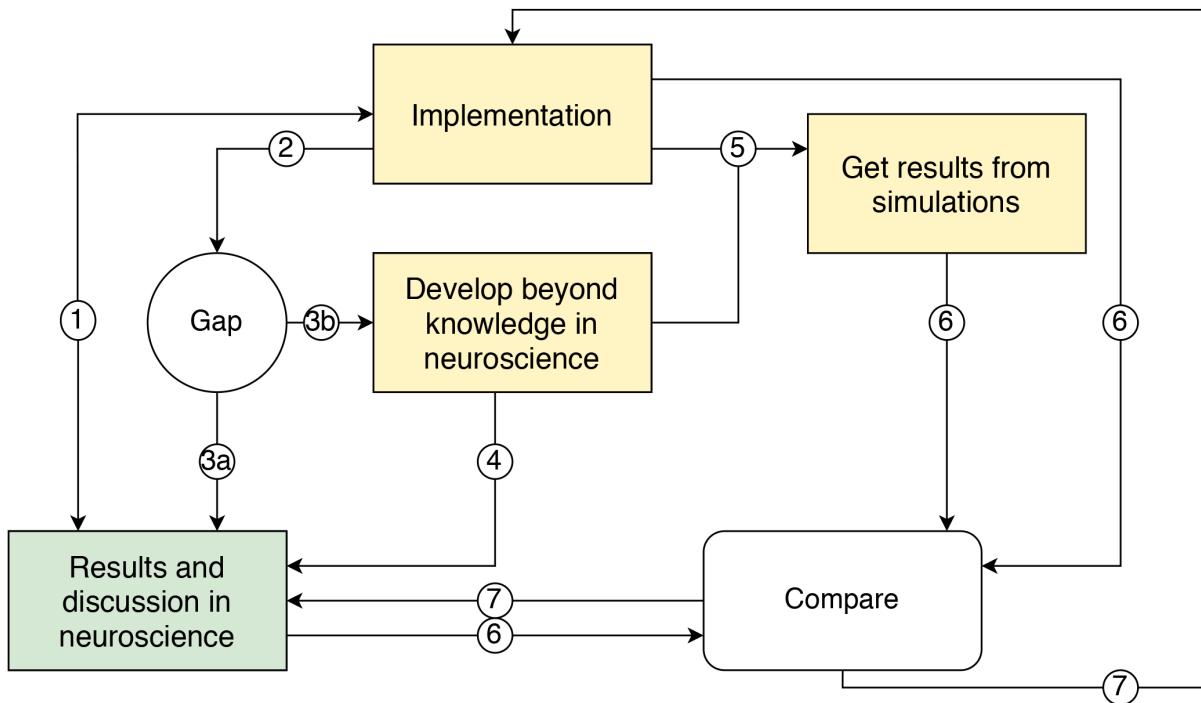
Figure 14. Proposed general flow of research when combining the field of AI and neuroscience to further learn about both fields.

be investigated. Additionally, the same comparison can give new insights on how to improve the current implementation in AI.

This active interaction is preferred over AI researchers merely looking at results and discussions in neuroscience because of three reasons. First, there is no guarantee that neuroscience has useful results to answer the question at hand. By having this interaction, research in the field of neuroscience can be conducted that is designed to answer questions in both AI and neuroscience. Without this interaction, a researcher in AI has to wait for useful findings to occur, and may encounter a hiatus while developing a biologically plausible system. Second, it is of great importance to understand the underlying mechanisms that result in the measurements. For example, in case of MEG, one has to be aware that radial sources (i.e., dipoles in line with the centre of the head) are not detected, and that perfect source separation is not possible. These aspects should be considered when interpreting the results and possibly using them in AI. Third, the discussion in neuroscience that influences the artificial implementation should include various studies and not focus on one or very few studies. The more studies are included, the more likely that the artificial implementation actually resembles findings in neuroscience. In general, researchers in neuroscience will have a better overview of the field than researchers in AI.

Conclusions

In this project we investigated the role of neural oscillations in sensory sampling and how this role could inspire the development of AI. At the behavioural level, we found that subjects responded faster when task frequency increased. Although we conclude that entrainment is not happening, it is possible that active sensing could still contribute to the RT effect in combination with other neural mechanisms. Rather than the oscillations adjusting to the task frequency by entraining to the external rhythm, the oscillations possibly still adjust but, for instance, not with an exact re-alignment to the task frequency. Then, active sensing could still increase the communication between the motor system and the visual system, giving lower RTs when task frequency increases.

Since the decrease in RT with increasing task frequency is a novel finding, particularly in the visual domain, we propose to further investigate this effect. More task frequencies could be added to the task to find where the limit of this effect is. Additionally, various sensory domains could also be studied and compared.

We further discussed how the notion of time and specifically oscillations can be used in the development of AI. By including these concepts in an artificial system, we may be able to tackle the problem of perceptual relevance in general AI.

This is needed to develop artificial cognitive agents that are able to perform the same scope of tasks as humans can in a complex environment, rather than having narrow AI. Promising developments are those in the fields of neuromorphic computing and spiking neural networks, which are more biologically plausible than conventional hardware and AI.

When drawing inspiration from neuroscience to develop AI, an active interaction between the two research fields is necessary to optimally contribute to both fields. We described how such a process could take place and how both fields can benefit from such a collaboration.

In the current project, we reported a novel behavioural effect in the visual domain and found support for the concept of covert active sensing. Moreover, by relating the study to AI, this project creates a step toward further developing general AI.

References

Alho, J., Lin, F. H., Sato, M., Tiitinen, H., Sams, M., & Jääskeläinen, I. P. (2014). Enhanced neural synchrony between left auditory and premotor cortex is associated with successful phonetic categorization. *Frontiers in Psychology, 5*, 394.

Assaneo, Florencia, M., & Poeppel, D. (2018). The coupling between auditory and motor cortices is rate-restricted: Evidence for an intrinsic speech-motor rhythm. *Science Advances, 4*(2), 3842.

Bishop, G. H. (1932). Cyclic changes in excitability of the optic pathway of the rabbit. *American Journal of Physiology-Legacy Content, 103*(1), 213–224.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*(4), 433–436.

Bressler, S. L. &Seth, A. K. (2011). Wiener–Granger causality: a well-established methodology. *Neuroimage, 58*(2), 323–329.

Canolty, R. T., Edwards, E., Dalal, S. S., Soltani, M., Nagarajan, S. S., Kirsch, H. E., Berger, M. S., Barbaro, N. M., & Knight, R. T. (2006). High gamma power is phase-locked to theta oscillations in human neocortex. S*cience, 313*(5793), 1626–1628.

Chon, A., Balachandar, N., & Lu, P. (2017). Deep convolutional neural networks for lung cancer detection. *Standford University*.

Dennett, Daniel (2006). The frame problem of AI. *Philosophy of Psychology: Contemporary Readings* 433, pp. 67–83.

Fries, P., Reynolds, J. H., Rorie, A. E., & Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, *291*(5508), 1560-1563.

Gwilliams, L., & King, J. R. (2017). Performance-optimized hierarchical models only partially predict neural responses during perceptual decision making. B*ioRxiv*, 221630.

Haegens, S. (2020). Entrainment revisited: a commentary on Meyer, Sun, and Martin (2020). *Language, Cognition and Neuroscience*, 1–5.

Haegens, S., & Golumbic, E. Z. (2018). Rhythmic facilitation of sensory processing: a critical review. *Neuroscience & Biobehavioral Reviews, 86*, 150–165.

Haselager, W. F. G. (1997). *Cognitive science and folk psychology: The right frame of mind*.

Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. *Neuron, 95*(2), 245–258.

Helfrich, R. F., Breska, A., & Knight, R. T. (2019). Neural entrainment and network resonance in support of top-down guided attention. *Current Opinion in Psychology, 29*, 82–89.

Imam, N., & Cleland, T. A. (2020). Rapid online learning and robust recall in a neuromorphic olfactory circuit. *Nature Machine Intelligence, 2*(3), 181–191.

Kucherenko, T., Hasegawa, D., Henter, G. E., Kaneko, N. & Kjellström, H. (2019). Analyzing input and output representations for speech-driven gesture generation. *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*, 97–104.

Kwisthout, J. (2012). Relevancy in problem solving: a computational framework. *The Journal of Problem Solving, 5*(1), 4.

Lakatos, P., Gross, J. & Thut, G. (2019). A new unifying account of the roles of neuronal entrainment. *Current Biology , 29*(18), R890–R905.

Leszczynski, M., & Schroeder, C. E. (2019). The Role of Neuronal Oscillations in Visual Active Sensing. *Frontiers in integrative neuroscience,* 13.

Maass, W. (1997). Networks of spiking neurons: the third generation of neural network models. *Neural networks, 10*(9), 1659–1671.

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG-and MEG data. *Journal of Neuroscience Methods, 164*(1), 177–190.

McCarthy, J., & Hayes, P. J. (1981). Some philosophical problems from the standpoint of artificial intelligence. *Readings in Artificial Intelligence*. Elsevier, 431–450.

McDermott, D. (1987). AI, logic, and the frame problem. *The Frame Problem in Artificial Intelligence*. Elsevier, 105–118.

Morillon, B., Arnal, L. H., Schroeder, C. E., & Keitel, A. (2019). Prominence of delta oscillatory rhythms in the motor cortex and their relevance for auditory and speech perception. *Neuroscience & Biobehavioral Reviews, 107*, 136–142.

Näätänen, R. (1971). Non-aging fore-periods and simple reaction time. *Acta Psychologica, 35*(4), 316–327.

Neil, D., Pfeiffer, M., & Liu, S. (2016). Phased lstm: Accelerating recurrent network training for long or event-based sequences. *Advances in Neural Information Processing Systems*, 3882–3890.

Niemi, P., & Näätänen, R. (1981). Foreperiod and simple reaction time. *Psychological Bulletin, 89*(1), 133.

Nobre, A. C., & van Ede, F. (2018). Anticipated moments: temporal structure in attention. *Nature Reviews Neuroscience, 19*(1), 34.

Obleser, J., & Kayser, C. (2019). Neural entrainment and attentional selection in the listening brain. *Trends in Cognitive Sciences, 23*(11), 913–926.

Park, H., Ince, R. A., Schyns, P. G., Thut, G., &Gross, J. (2015). Frontal top-down signals increase coupling of auditory low-frequency oscillations to continuous speech in human listeners. *Current Biology, 25*(12), 1649–1653.

Pfeiffer, M., & Pfeil, T. (2018). Deep learning with spiking neurons: opportunities and challenges. *Frontiers in Neuroscience, 12*, 774.

Pont, J. S. (1987). Firing patterns of bulbar respiratory neurones during sniffing in the conscious, non-paralyzed rabbit. *Brain Research, 414*(1), 163–168.

Schroeder, C. E., & Lakatos, P. (2009). Low-frequency neuronal oscillations as instruments of sensory selection. *Trends in Neurosciences, 32*(1), 9–18.

Schroeder, C. E., Wilson, D. A., Radman, T., Scharfman, H., & Lakatos, P. (2010). Dynamics of active sensing and perceptual selection. *Current Opinion in Neurobiology, 20*(2), 172–176.

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Driessche, G., Schrittwieser, J., Antonoglou, J., Panneershelvam, V., Lanctot, M. et al. (2016). Mastering the game of Go with deep neural networks and tree search. N*ature, 529*(7587), 484–489.

Spitzer, B., & Haegens, S. (2017). Beyond the status quo: a role for beta oscillations in endogenous content (re) activation. *Eneuro, 4*(4).

Taigman, Y., Yang, M., Ranzato, M., & Wolf, L. (2014). Deepface: Closing the gap to human-level performance in face verification. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1701–1708.

Thut, G., Schyns, P., & Gross, J. (2011). Entrainment of perceptually relevant brain oscillations by non-invasive rhythmic stimulation of the human brain. *Frontiers in Psychology, 2*, 170.

van Veen, B. D., van Drongelen, W., Yuchtman, M., & Suzuki, A. (1997). Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. *IEEE Transactions on biomedical engineering, 44*(9), 867–880.

Will, U., & Berg, E. (2007). Brain wave synchronization and entrainment to periodic acoustic stimuli. *Neuroscience Letters, 424*(1), 55–60.

Zoefel, B., ten Oever, S., & Sack, A. T. (2018). The involvement of endogenous neural oscillations in the processing of rhythmic input: More than a regular repetition of evoked neural responses. *Frontiers in Neuroscience, 12*, 95.

Effect of Eye Closure on Anticipatory Somatosensory Alpha Oscillations

Corinne Orlemann[1]
Supervisors: Hesham ElShafei[1], Saskia Haegens[1,2]

[1]*Radboud University Nijmegen, Donders Institute for Brain, Cognition and Behaviour, The Netherlands*
[2]*Department of Psychiatry, Columbia University, New York, USA*

Alpha oscillations (8–13 Hz) have traditionally been associated with a state of cortical idling and show an increase over posterior areas when awake participants close their eyes. In recent years, alpha oscillations have been proposed to reflect a mechanism of active functional inhibition, by suppressing cortical activity over task-irrelevant brain regions and facilitating excitability in task-relevant ones. Many studies investigated this top-down role of alpha, though it remains unclear how eye closure might impact modulation of alpha activity during cognitive tasks, and how this affects subsequent behavioural performance. In this study, we tested 33 participants performing a somatosensory spatial discrimination task in an eyes-open and closed condition, while recording brain activity using magnetoencephalography (MEG). We report an increase of alpha oscillations with eye closure and found further evidence of a functional inhibition role of alpha oscillations during somatosensory attention, reflected by alpha activity decrease over the contralateral somatosensory cortex as well as alpha increase over visual regions. Furthermore, we report the visual alpha modulation to be significantly higher for eyes-open than eyes-closed trials, implicating a necessity of increased alpha inhibition when irrelevant visual input is available. We demonstrated that this posterior anticipatory alpha activity predicted task performance, by associating higher activity with positive behavioural outcome, independent of eye condition. We therefore showed evidence that eye closure alters the general alpha activity profile and furthermore influences the anticipatory posterior alpha modulation during a somatosensory attention task. Eye closure however did not have an effect on the impact of alpha modulation on behaviour.

Keywords: alpha rhythm, functional inhibition, eye closure, somatosensory attention, MEG

Corresponding author: Corinne Orlemann; E-mail: corinneorlemann@gmail.com

Our brains continuously receive a high amount of information and are thus challenged with the task of filtering relevant input and suppressing distracting input. In recent years, cortical alpha oscillations (8-13 Hz) have been proposed to play a role in this filtering process (Jensen & Mazaheri, 2010). Furthermore, eye closure is associated with a general increase of posterior alpha (Adrian & Matthews, 1934) and is theoretically linked to facilitate sensory attentional processes by shutting out distracting visual input (Glenberg, Schroeder, & Robertson, 1998). Here, we aimed to investigate whether and how increased alpha activity due to eye closure impacts anticipatory alpha modulations during a somatosensory discrimination task.

Inhibitory alpha

For a long time the alpha rhythm was associated with a general state of cortical idling. However, more recent research proposes that alpha oscillatory activity reflects a mechanism for functional inhibition (Klimesch, Sauseng, & Hanslmayr, & 2007; Jensen & Mazaheri, 2010; Foxe & Snyder, 2011; Haegens, Nácher, Luna, Romo, & Jensen, 2011). This is supported by studies showing an increase of alpha band power over task irrelevant brain regions and networks, and a decrease of alpha power over task relevant brain regions (e.g., Pfurtscheller & Klimesch, 1991; Worden, Foxe, Wang, & Simpson, 2000; Sauseng et al., 2005). In this way, alpha oscillations are thought to gate the information flow in the brain to increase perceptual performance, which is demonstrated by an increase of alpha oscillations in anticipation of distractors (Bonnefond & Jensen, 2012).

In the visual domain this targeted inhibition of task-irrelevant information occurs in a lateralized pattern (e.g., Thut, Nietzel, Brandt, & Pascual-Leone, 2006; Ikkai, Dandekar, Curtis, 2016). In other words, when a stimulus was presented in either the right or left visual field, alpha activity decreased contralateral to the attended region and increased contralateral to the to-be-ignored location (Worden et al., 2000; Sauseng et al., 2005; Rihs, Michel, & Thut, 2007; Kelly, Gomez-Ramirez, & Foxe, 2009). This alpha activity increase over occipital regions, responsible for processing the not-to-be attended visual information, is assumed to reflect an inhibition of distracting input. The pattern of lateralized alpha in- and decrease can be described by a lateralization index, defined by positive values for right-ward attention and negative values for left-ward attention (Thut et al, 2006). Importantly, the

alpha lateralization index was deterministic of target perception, with negative values preceding a more rapid and accurate detection of left-ward targets and positive-values preceding the detection of right-ward targets (Thut et al., 2006; Kelly et al., 2009; Händel, Haarmeier, & Jensen, 2011).

Somatosensory anticipatory alpha lateralization

The lateralization pattern is not specific to the visual domain but also exists in the somatosensory domain (Haegens, Osipova, Oostenveld, & Jensen, 2010). In this case, a lateralization occurs over the right or left hemispheric somatosensory cortex, with alpha decrease contralateral to the stimulus application, indicating a similar functional mechanism for tactile tasks as the posterior alpha during visual tasks.

This lateralization pattern occurred in an anticipatory fashion prior to the presentation of an expected stimulus, when a cue guided attention towards the left or the right side. The power of this lateralization pattern furthermore decreased when this attentional cue was unreliable (Haegens, Händel, & Jensen, 2011).

Furthermore, increased somatosensory alpha lateralization modulated the subsequent task performance. For example, better performance (i.e., more accurate and faster responses) followed a higher increase of lateralized alpha over the somatosensory cortex (Haegens et al., 2011; Haegens, Luther, & Jensen, 2012). These findings show that alpha induces an anticipatory brain state of sensory regions to optimize processing in a predictive manner. Studies in the visuospatial literature indicated that lateralized alpha modulation is mainly driven by a contralateral decrease (e.g., Sauseng et al., 2005; Thut et al., 2006), while others also reported a role of ipsilateral increase (Rihs et al., 2007; Rihs, Michel, & Thut, 2009). Here, ipsilateral increase seemed especially dependent on distractor strength. Also for the somatosensory domain, ipsilateral alpha increased with the presentation of distractors, though with further increase of distractor strength a global decrease of alpha was shown both contra- as well as ipsilateral (Haegens et al. 2012). Similar to the visual information processing, somatosensory alpha therefore increased over the task-irrelevant sensory cortex, showing an inhibitory mechanism in the context of expected irrelevant information in the form of distractors.

Importantly, during somatosensory discrimination processing the lateralized alpha pattern

is accompanied by an increase of anticipatory posterior alpha power (Haegens et al., 2012). As somatosensory alpha increase is associated with an inhibition of processing of the task-irrelevant hand, the posterior alpha is assumed to reflect a general inhibition of visual input. Furthermore, posterior alpha increase also reflected a modulation on task performance, with more accurate responses emerging after a higher increase of posterior alpha (Haegens et al., 2012).

Eye closure effect

Since the discovery of cortical alpha oscillatory activity by Hans Berger (1929) almost a century ago, it has been known that a general increase of alpha power occurs when awake participants close their eyes. This modulation is especially localized in parieto-occipital regions (Adrian & Matthews, 1934), but has also been observed in a more widespread distribution (Geller et al., 2014). However, little is known about how this general alpha effect might interact with the inhibitory role of alpha during cognitive tasks and whether these alpha activities might originate from the same underlying sources.

Anecdotally, eye closure enhances the concentration on other sensory modalities to improve the cognitive performance by suppressing visual input. This line of thought stems from the assumption of a dominant role of visual processing. In this sense, a functional imaging study reported a shift of processing networks, highlighting the exclusion of the dominant visual mode as a consequence of eye closure (Brodoehl, Klingner, & Witte, 2015). On the somatosensory level, eye closure lead to a more sensitive somatosensory threshold during a simple perception paradigm (Brodoehl, Klingner, Stieglitz, & Witte, 2015). In the context of memory recall, eye closure also facilitated recollection (e.g., Vredeveldt, Baddeley, & Hitch, 2012; Vredeveldt et al., 2015).

So far, a possible eye closure effect on the inhibitory modulation of alpha has only been investigated in the context of auditory attention (Wöstmann, Schmitt, & Obleser, 2019). The prior reported effects of lateralized and posterior alpha modulation during somatosensory attention tasks exist also for the auditory domain, showing alpha modulations for auditory spatial tasks (Banerjee, Snyder, Molholm, & Foxe, 2011) as well as for listening tasks including speech perception (Strauß, Wöstmann, & Obleser, 2014; Dimitrijevic, Smith, Kadis, & Moore, 2017). In the same manner as for the other two modalities, alpha increase over task-

irrelevant and decrease over task-relevant temporal regions improved performance during an auditory task (Dimitrijevic et al., 2017). In their study, Wöstmann and colleagues (2019) reported that eye closure not only increased the general power of alpha oscillations, but also the modulation of alpha during the attentional task. The study presented two auditory streams of numbers, which either had to be attended or ignored. In this regard, alpha activity over parietal and occipital regions increased with the presentation of the attended numbers compared to the ignored numbers, which was shown for both the eyes open as well as the eyes closed condition. Yet, this modulation of alpha significantly increased with eye closure. However, eye closure did not facilitate task performance, indicating a general neural effect of eye closure with no impact on behaviour. So far the eye closure alpha effect has not been investigated in the context of somatosensory alpha lateralization.

Current Study

Here we aimed to investigate whether eye closure affects alpha modulation during a somatosensory discrimination task. We adapted the paradigm used by Haegens et al. (2011) to include eyes-open and eyes-closed conditions.

First, we examined differences in alpha activation for open and closed eyes, expecting higher alpha power for the eyes closed condition. Next, we analysed the pre-stimulus time window regarding the alpha activation pattern of paying attention to the left or the right hand. This pattern was further examined on a possible difference in activity based on eye closure. Lastly, we investigated whether performance was modulated by alpha lateralization and/or posterior alpha power, hypothesizing a possible effect of eye closure on this modulation.

Method

Participants

34 participants (Age: $M = 25$, $SD = 3.86$, range = 20-33 years; 18 female, 16 male; 30 right handed, 2 left handed, 2 ambidextrous) took part in the experiment. One subject was excluded from analysis due to poor data quality. All participants were recruited with the online system SONA and were healthy, free from any neurological or psychiatric disorders, and reported normal hearing and normal or corrected-to-normal vision. Prior to the experiment, participants received a detailed overview of the MEG system and the study and

signed an informed consent form. The study falls under the general ethics approval (CMO 2014/288 "Imaging Human Cognition") in accordance with the Declaration of Helsinki.

Materials

For the presentation of the sensory stimuli, two electrodes were attached to the participant's right and left thumb. These electrodes administered a short electrical stimulus (0.2 ms) in the form of a pulse train. Electrical stimuli were generated using two constant-current high voltage stimulators (Digitimer Ltd, Model DS7A). Through the use of two Digitimer stimulators we were able to present different intensities for each hand, in order to account for the variance of sensory thresholds between hands.

Intensity of the electrical stimuli were determined by acquiring the sensory threshold for each thumb and setting the final stimulus at 150% of this threshold level (M_{right} = 6.4 mA, range = 3.9 − 9.5 mA; M_{left} = 5.5 mA, range = 3.2 − 9.9 mA).

During the experiment the pulse trains were presented in a low and high frequency (frequency represented by the number of pulses). Low (either 25 or 33.3 Hz) and high frequencies (41.7, 50, or 66.7 Hz) were determined for each participant individually to ensure a successful execution of the task, avoiding chance level as well as ceiling performance. Auditory cues and feedback (200 ms length each) were computer generated and presented binaurally through air-conducting tubes.

Experimental paradigm

Participants performed a somatosensory discrimination task while their brain activity was recorded using MEG. Participants received an electrical stimulus of a low or high frequency to either the right or left thumb. Prior to the stimulus presentation, a 100% valid auditory cue guided participants' attention to either the right or left side. Participants were instructed to determine as fast and accurately as possible whether the perceived stimulus was of the low or the high frequency. Answers were
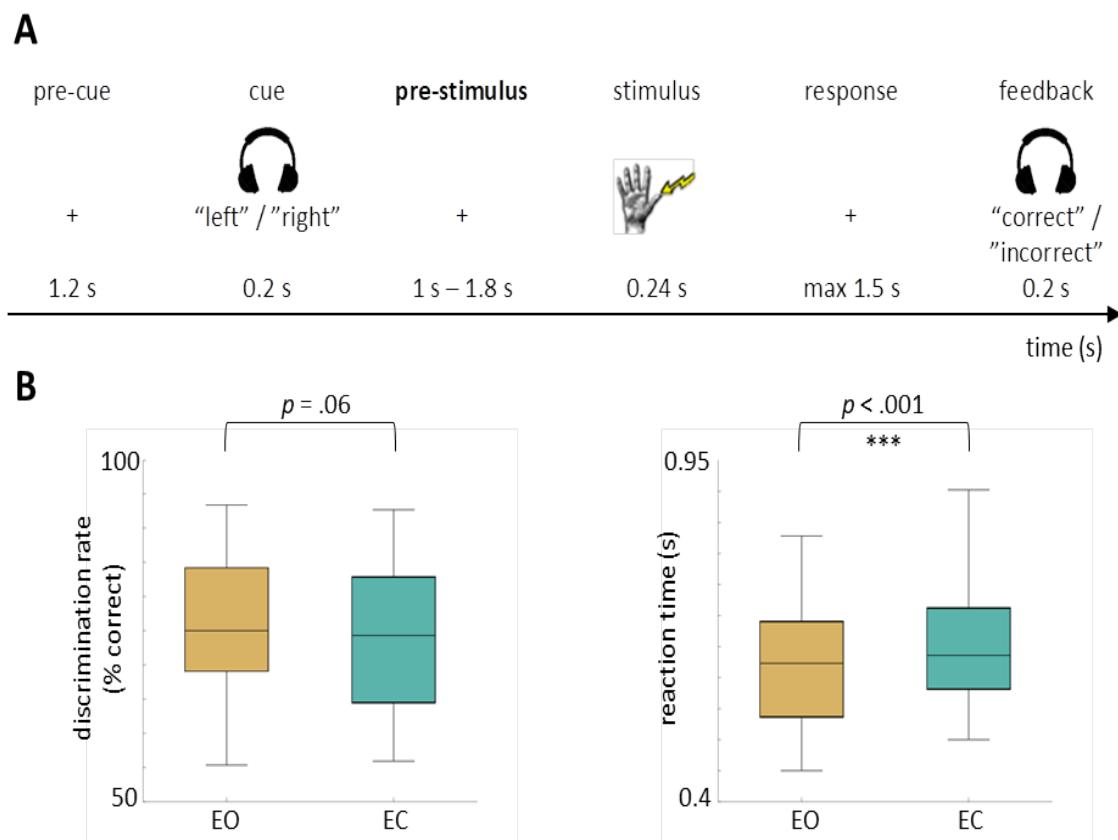
Figure 1. Experimental design and behavioural results. **A.** Typical trial procedure. Participants performed the task in an eyes-open and an eyes-closed condition. The pre-stimulus interval served as the primary window of interest. **B.** Behavioural performance of the discrimination rate (left) and the RTs (right) divided into the two eye conditions. RTs were calculated only on correct trials and showed a significant difference between eyes open and eyes closed.

given via button press with the right index finger (i.e., left button press indicated the low frequency; right button press indicated the high frequency).

A typical trial started with a pre-cue interval of 1.2 s followed by the auditory cue, a jittered 1–1.8 s pre-stimulus interval, the stimulus (240 ms pulse train), a maximum response time of 1.5 s, and finally auditory feedback indicating whether the answer was correct or incorrect (Fig. 1A).

Participants performed this task under two conditions, an eyes-open condition (EO) and an eyes-closed condition (EC). Conditions were presented in a counter-balanced block-design of four blocks per condition with 76 trials each, resulting in a total of 304 trials per condition. During the EO condition participants were instructed to fixate on a fixation cross in the middle of the screen. For the EC condition participants kept their eyes closed for the duration of the block. After each block, participants were presented with a short questionnaire of four answer possibilities to rate their sleepiness (very sleepy, sleepy, awake, very awake).

Prior to the experiment, participants performed four training blocks (two per condition, 12 trials per block), during which they got familiar with the task.

Participants were seated upright in the MEG helmet and instructed to keep their head position as stable as possible for the duration of the experiment. After each block participants were able to take a short break for which they stayed seated in the MEG chair.

Data acquisition

Whole-head brain activity was recorded with a 275-channel CTF MEG system with axial gradiometers at a sampling rate of 1200 Hz (CTF MEG Systems, VSM MedTech Ltd.). The MEG system was positioned in a magnetically shielded room. For a real time representation of the participant's head position, three head localization coils were placed at the right and left ear canals as well as at the nasion. The real time representation of the head position allowed for monitoring of head movements and adjustments to the original position. The three points furthermore served as offline anatomical landmarks. During the experiment eye movement of the left eye was tracked using an Eyelink 1000 eyetracker (SR Research Ltd.). Experimental stimuli were programmed and presented with the software Presentation (Version 18.0, Neurobehavioral Systems, Inc., Berkeley, CA, www.neurobs.com).

In a separate session an anatomical MRI of the

participant's brain was acquired, unless a recent anatomical MRI could be obtained from the database of the institute. MR images were acquired with any of the three available scanners at the Donders Centre of Neuroimaging; The 3T Siemens Magnetom Prisma MR scanner, the 3T Siemens Magnetom PrismaFit MR scanner, or the 3T Siemens Magnetom Skyra MR scanner (Erlangen, Germany). For the MRI scan participants were wearing ear plugs with drops of Vitamin E for improved co-registration of the MRIs and MEG data. The anatomical images served the purpose of source reconstruction of the MEG data; however, note that source reconstruction analysis is not included for this report.

Behavioural analysis

For the behavioural data, we analysed discrimination rate (percentage of correct responses) and reaction times (RTs) of correct trials. A repeated-measures two-way ANOVA was computed for each of these behavioural measures with the factors eye condition (EO and EC) and cue condition (left and right).

Pre-processing

A pre-processing pipeline for the MEG data was developed using the Matlab toolbox FieldTrip (Oostenveld, Fries, Maris, & Schoffelen, 2011). Raw data was downsampled offline to 300 Hz and epoched (-4 − 3 s, relative to the somatosensory stimulus onset). Trials were first visually inspected and rejected based on a trial summary representation of variance. Furthermore, the resulting trials were manually inspected on a trial-by-trial basis and rejected on the basis of muscle artefacts and SQUID sensor jumps. We used independent component analysis to correct for cardiac and eye movement components ($Mean_{removed} = 8$). A last visual inspection of the corrected trials was performed based on variance, range, z-scores and maximum absolute values. This cleaning process was done blind to experimental condition. On average 11% of all trials were rejected during this procedure. Time window descriptions in the following analyses are always in regard to the stimulus onset as time t = 0 s, unless indicated otherwise.

Spectral analysis

Since data was collected with axial gradiometers, a transformation towards a planar representation of the MEG field distribution was calculated using

the nearest-neighbour method. Planar gradient representations allow for a clearer presentation of the data as activity is typically maximal above a source. We computed spectral representations for three time windows: the pre-stimulus window (-1 − 0 s), the post-stimulus window (0 − 1 s), and the pre-cue window (-1 − 0 s of cue-onset). The pre-cue window served as a baseline and was determined by re-defining the data to centre of the cue onset. The post-stimulus time window was specifically used for the determination of the visual regions of interest (as explained below). Fast Fourier transformations were computed for all sensors by segmenting trials into the respective time windows and multiplying them with a Hanning taper. The computation of power distributions ranged from 1 to 30 Hz.

For a time-resolved-representation of the spectral power distribution, we computed an additional time-frequency analysis (TFR) on the pre-stimulus interval (-1.5 − 0.5 s). For this, we chose a constant sliding time window of 0.5 s for each frequency (1 − 30 Hz).

Impact of eye closure on oscillatory activity

To statistically test for power differences between the EO and the EC condition in the alpha and beta range, we computed a cluster-based permutation test on the whole head in the pre-stimulus time-window -1 s to 0 s (Maris & Oostenveld, 2007). For this test the contrast EO − EC was chosen, which leads to positive values for a stronger activation in the EO condition and negative values for stronger activation in the EC condition. This statistical analysis allows dealing with the multiple comparison problem by computing a dependent-samples t-value for each sample-pair (i.e., sensor-frequency-pairs). Samples that exceeded a p-value threshold of .05 were clustered based on their spatial adjacency. Individual cluster statistics were calculated through the summation of t-values within each cluster. Subsequently, the maximum cluster was selected as the test-statistic. Data was randomized across the two conditions and the test-statistic was repeated 1000 times. This allows to obtain a reference distribution of maximum cluster t-values to evaluate the statistic of the actual data.

Alpha peak selection

For the sensor-level analysis, we computed the individual alpha peaks for each participant. These peaks were selected based only on the activity in the individual visual ROIs. The FFT of the pre-stimulus interval (-1 − 0 s) was divided into EO and EC trials. We determined participants' peak frequencies by calculating the maximal absolute activity within a broad alpha range (7 − 14 Hz). As intra-individual alpha peaks did not significantly vary between conditions ($t(32) = 0.53$, $p = .6$) they were averaged for each participant ($M = 10$ Hz, range $= 8.5 − 11.5$ Hz) across EO and EC conditions.

Calculating the individual alpha peak frequencies allows taking into account the inter-individual variability of alpha frequencies between subjects. In this way, we obtained a more accurate estimation of alpha activity by avoiding a possible bias against some participants and contamination of neighbouring frequency bands, as can be the case by choosing a fixed frequency band (Haegens et al., 2014). All further analysis was computed using these individual alpha peaks unless indicated otherwise.

Regions of interest selection

Three regions of interest (ROIs) where selected for further sensor-level analysis: A visual ROI, a left somatosensory ROI and a right somatosensory ROI. We computed these ROIs individually for each participant.

For the selection of the somatosensory ROIs, individual post-stimulus event-related fields (ERFs) were calculated regardless of eye-condition. These ERFs were divided in right-hand-stimulus and left-hand-stimulus trials. A baseline correction was applied based on the baseline window -1.5 to -1 s. The sensors with the maximum activity of the stimulus ERF determined the left and right somatosensory regions of interest.

As the experiment did not include the presentation of strong visual stimuli, the visual sensors were selected with the previously computed post-stimulus FFT (0 s − 1 s). The post-stimulus FFT was divided into EO and EC trials, averaged over trials within the two conditions, and finally contrasted in the form of EO − EC. The sensors with the maximal difference within this contrast were used for the visual ROI.

Therefore, for each participant the individual 20 maximum channels for left somatosensory, right somatosensory and visual cortex were available for further analysis.

Alpha lateralization index

The alpha lateralization index describes the ratio

of alpha distribution over both the ipsilateral and contralateral hemispheres of the somatosensory cortex. This index is calculated based on the activity of individual alpha peaks within individual somatosensory ROIs in the following way: Alpha lateralization index = (alpha-ipsilateral − alpha-contralateral) / (alpha-ipsilateral + alpha-contralateral). The index gives positive values if alpha power activity is higher over the ipsilateral hemisphere and/or lower over the contralateral hemisphere. Negative values arise if alpha power activity is lower over the ipsilateral hemisphere and/or higher over the contralateral hemisphere. The calculation of the alpha lateralization index follows the same rationale as the index used by Thut et al. (2006).

Effect of alpha modulation on performance

To investigate the impact of pre-stimulus alpha modulation on behavioural performance, we conducted an analysis on the alpha lateralization index as well as the visual alpha power and their interaction with the discrimination rate (percentage correct) and reaction times.

For each participant the pre-stimulus FFT (-1 − 0 s) was divided into correct and incorrect trials (omitting no-response trials) for each of the two eye conditions separately. The alpha lateralization index was computed for correct and incorrect trials for each condition following the strategy as explained above. For statistical analysis, we calculated a repeated-measures two-way ANOVA with the factors eye condition (EO or EC) and discrimination rate (correct or incorrect).

We followed the same procedure for the computation of alpha lateralization indices for fast and slow trials. For each participant, a median split of the reaction times of only the correct trials determined the categorisation of fast and slow trials. Similar to the analysis of the discrimination rate, a repeated-measures two-way ANOVA with factors eye condition (EO or EC) and reaction times (fast or slow) was computed for statistical analysis.

We repeated this process for investigating the relationship between visual alpha power and behavioural performance. The pre-stimulus FFT (-1 − 0 s) was divided into the two separate eye conditions and baseline corrected based on the baseline FFT (-1 − 0 s cue-onset, separate for each condition). The visual alpha power was computed based on individual visual ROIs and individual alpha

peaks for correct and incorrect trials, as well as for fast and slow trials. We calculated two repeated-measures two-way ANOVAs, one with the factors eye condition and discrimination rate, and one with the factors eyes condition and reaction times.

Results

Performance over all 33 participants for both eye conditions combined was an average discrimination rate of 73.4% ($SD = 10.2\%$) and an average reaction time (correct trials only) of 673.1 ms ($SD = 92.9$ ms). Figure 1B shows the discrimination rates and RTs for the EO and the EC condition.

Slower performance with eye closure

To test for statistical differences between the conditions, we calculated a repeated-measures two-way ANOVA with the factors eye (EO or EC) and attention condition (left or right) on the discrimination rate. Only a trend was observed for the main effect of eye condition ($F (1, 32) = 3.62$, $p = .06$), indicating higher accuracy for the EO condition. Neither the difference of attention sides ($F (1, 32) = 0.04, p = .85$), nor the interaction of eye and attention condition ($F (1, 32) = 0.22, p = .64$) were significant.

In the same fashion, a repeated-measures two-way ANOVA was computed on the RTs. The main effect of eye condition was significant ($F (1, 32) = 31.58, p < .001$) with faster performance in the EO condition. A main effect of attention side was found as well ($F (1, 32) = 5.89, p = .02$), reflecting faster RTs for the right side. This result can be attributed to the fact that participants were giving their answers through button press with the right index finger. Therefore, participants performed the task faster when the stimulus was applied to the same hand they were giving their answer with. The interaction between eye and attention condition was not significant for the RTs ($F (1, 32) = 0.19, p = .67$).

Eye closure increases oscillatory alpha activity

Generally, the EC condition showed higher alpha and beta power in the pre-stimulus window than the EO condition, with a focus on occipital regions (see Fig. 2A). A cluster-based permutation test was computed to test the difference of oscillatory activity between EO and EC for significance. Contrasting the EO with the EC condition revealed two negative
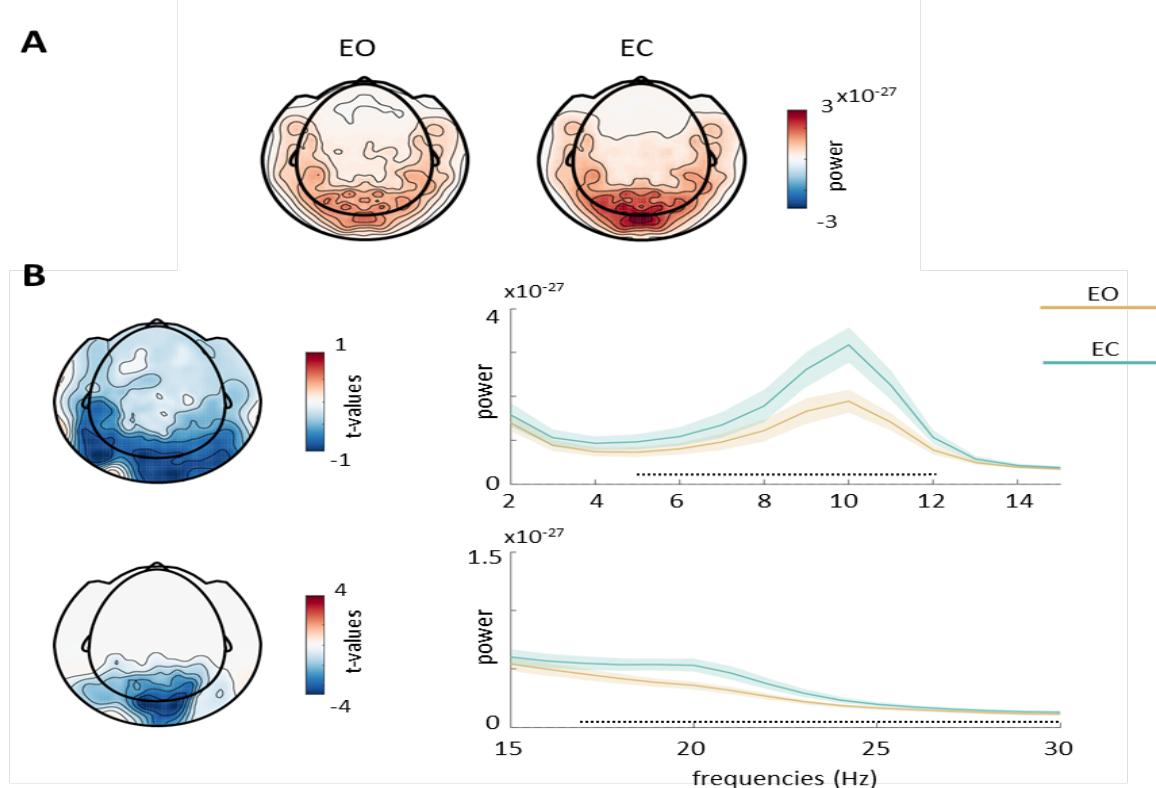
Figure 2. Spectral condition differences. **A.** Non-baseline corrected FFT (1 − 30 Hz) of the pre-stimulus window (t = -1 − 0 s) for eyes open and eyes closed. **B.** Cluster statistics of the EO − EC contrast. **Left:** Topographic representations of cluster distributions. **Right:** Frequency distributions of clusters, dotted lines represent frequencies with significant differences between conditions. **Top:** Widespread alpha cluster with significant frequency differences 5 − 12 Hz. **Bottom:** Occipital centred beta cluster with significant frequency differences 17 − 30 Hz.

clusters ($p < .001$; Fig. 2B), one in the alpha range (7 − 14 Hz) and one in the beta range (15 − 30 Hz). The alpha cluster was distributed over widespread regions with a peak at 10 Hz, while the beta cluster was concentrated towards occipital sensors, showing the highest difference between conditions around 20 Hz. Therefore, power of both, alpha and beta oscillations, significantly increased with eyes closure.

Anticipatory posterior alpha modulation stronger for open eyes

To investigate the alpha modulation during the pre-stimulus interval, we computed spectral representations of the attention conditions, by contrasting the pre-stimulus alpha activation against a baseline (i.e., left-ward attention vs baseline, right-ward attention vs baseline; see Fig. 3A & B). Only for the EO condition an increase of occipital alpha power can be observed, while both conditions show a decrease of alpha power over contralateral central sensors in anticipation of the stimulus. A paired-samples t-test was calculated on the pre-stimulus visual alpha power increase against baseline

between the two conditions to test for differences in posterior alpha modulation due to eye closure. Visual alpha power was based on individual alpha peaks and individual visual ROI sensors. The EO condition had a significantly higher posterior alpha modulation than the EC condition ($t (32) = 6.28$, $p < .001$; see Fig. 3C). This result reflects an increase of posterior alpha power during the pre-stimulus interval vs baseline in the eyes open condition, while the EC condition does not show such modulation. Hence, despite an overall increase of alpha power with eyes closure, the anticipatory posterior alpha modulation during the pre-stimulus interval was higher for open eyes.

No difference of alpha lateralization between eye conditions

To further investigate the alpha power decrease during the pre-stimulus interval over contralateral somatosensory regions, we computed a normalization which takes into account leftward and rightward attention and lead to a visualisation of the previously reported lateralization pattern
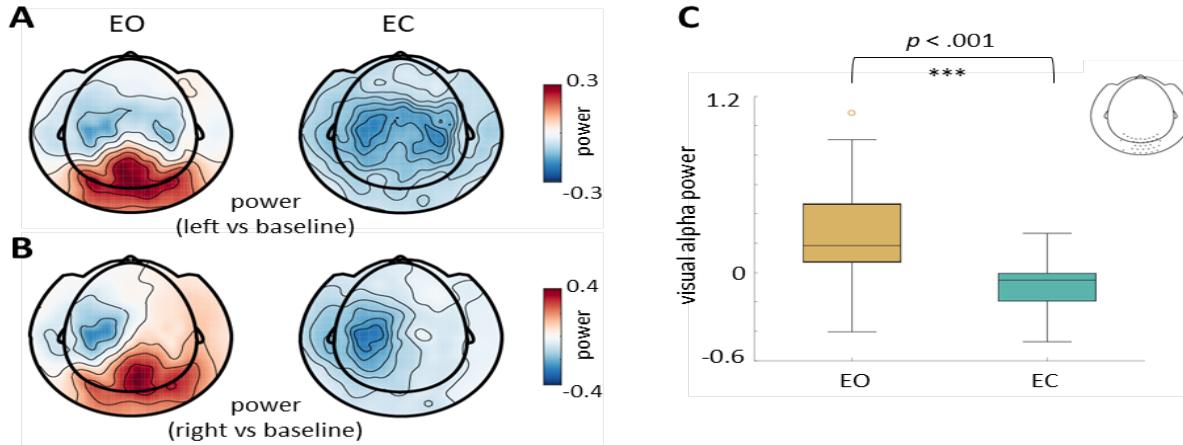
Figure 3. Attention modulation. Topographic plots of pre-stimulus alpha power (individual alpha peaks) modulation following the left cue **(A)** and the right cue **(B)** against baseline activity. **Left:** eyes open condition. **Right:** eyes closed condition. Both conditions show central alpha modulation based on cue direction. Only eyes open condition shows an additional modulation of visual alpha increase. **C.** Difference of visual alpha modulation between conditions (p < .001). The empty helmet layout marks the visual ROIs used for the power estimation (summary individual visual ROIs).

(see Fig. 4A; Haegens et al., 2011; Haegens et al., 2012). We further examined the time course of this modulation of the somatosensory anticipatory alpha through the computation of time-frequency representations (TFRs). These TFRs were computed separately for the EO and EC condition in regard to the activation of lower band frequencies (5 – 30 Hz) for the left vs right normalization (see Fig. 4C; right

hemispheric sensors were mirrored to combine with left hemispheric ones). Through this visualisation we observed a sustained modulation for alpha (around 10 Hz) as well as for beta (around 20 Hz) during the pre-stimulus interval.

To investigate differences between EO and EC alpha lateralization, the alpha lateralization index over somatosensory sensors was computed for both
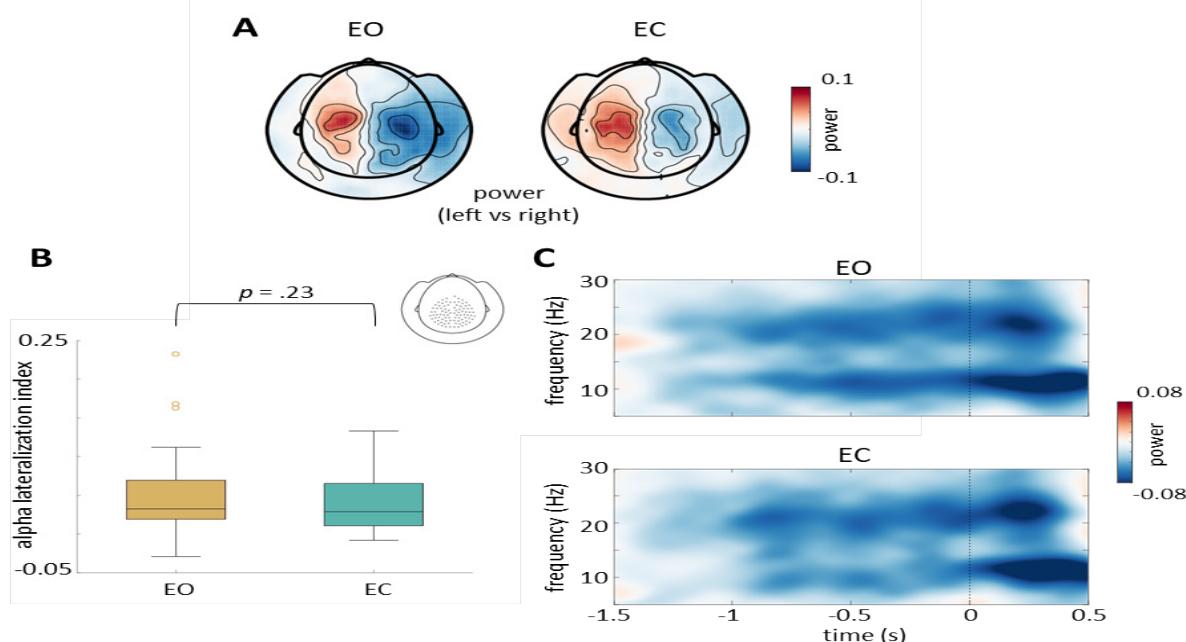
Figure 4. Left vs right attention modulation. **A.** Topographic representation of the attention left vs attention right alpha power modulation (individual alpha peaks) for the eyes open (left) and the eyes closed (right) condition. **B.** Difference of alpha lateralization index between conditions regarding the individual somatosensory ROIs represented in the empty helmet. **C.** Time-frequency representation of the low frequency (5 – 30 Hz) modulation during the pre-stimulus interval (individual somatosensory ROIs). **Top:** Eyes open condition. **Bottom:** Eyes closed condition.

conditions. There was no significant difference of alpha lateralization between the conditions (t (32) = 1.21, p = .23; Fig. 4B). This indicates that eye closure did not have an impact on the anticipatory lateralized alpha modulation.

No performance modulation by alpha lateralization index

To investigate whether performance was modulated by the pre-stimulus lateralized alpha activation, we computed repeated-measure two-way ANOVAs on the lateralization index; one with the factors eye condition (EO or EC) and discrimination rate (correct or incorrect), and one with the factors eye condition (EO or EC) and reaction times (slow or fast). Figure 5 shows a representation of the results. For discrimination rate, neither the main effect of performance (F (1, 32) = 0.004, p = .95), nor the main effect of eye condition (F (1, 32) = 2.77, p = .11), nor the interaction (F (1, 32) =

0.7, p = .41) showed significant results. Therefore, no increased alpha lateralization was found for improved discrimination for either condition.

For reaction times, no significant effects for the main effect of eye condition (F (1, 32) = 0.33, p = .57) and the interaction effect of eye condition and RTs (F (1, 32) = 0.01, p = .94) were found. However, a trend for the main effect of RTs (F (1, 32) = 3.95, p = .06) can be observed, indicating a link between faster RTs and increased alpha lateralization index.

Improved performance with increased posterior alpha power

Similar to the analysis of the lateralization index, we wanted to test for a possible influence of posterior alpha modulation on performance outcome. Therefore, repeated-measures two-way ANOVAs were calculated on the pre-stimulus posterior alpha modulation (i.e., power vs baseline) with the factors eye condition and discrimination
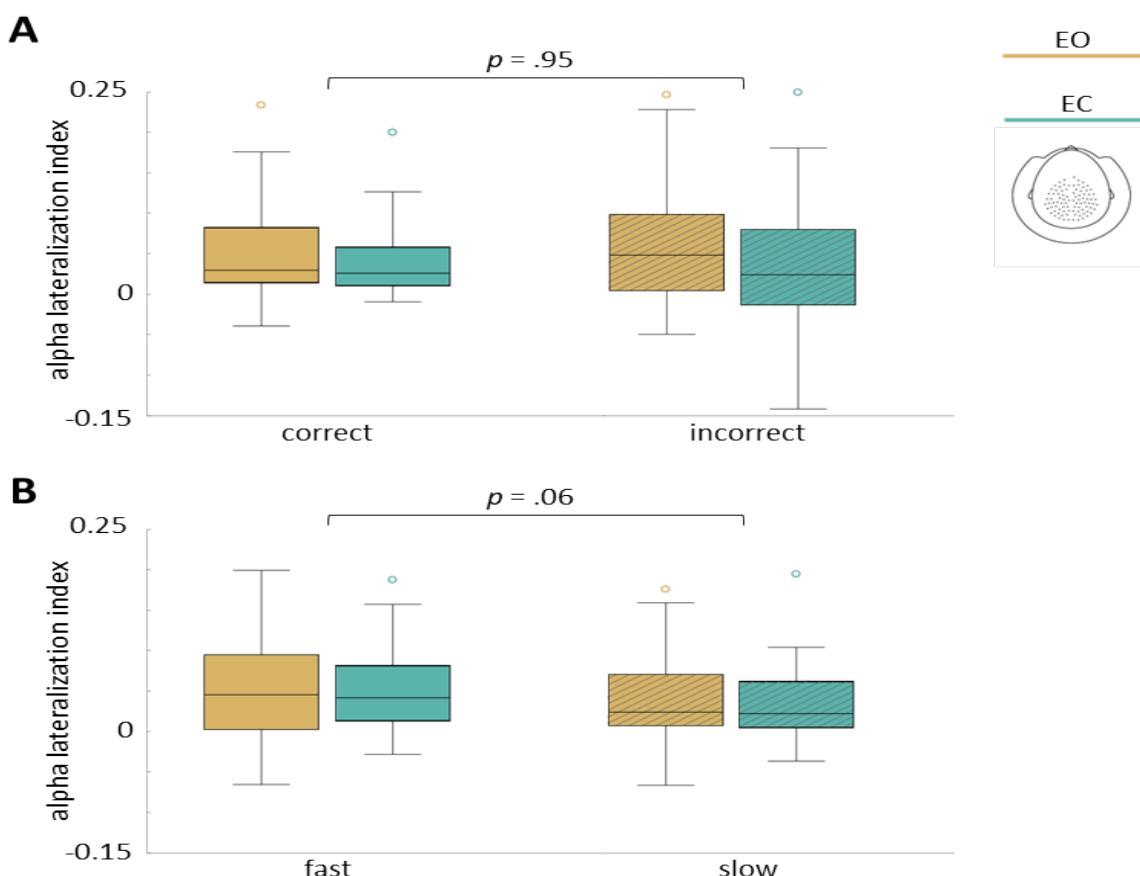
Figure 5. Behavioural performance modulation by alpha lateralization index. The empty helmet represents the somatosensory ROIs used for this analysis (summary of individual ROIs). P-values represent the main effect of performance. **A.** Modulation by alpha lateralization index on discrimination rate (% correct responses). No significant differences between correct and incorrect trials were found for either the eyes open or the eyes closed condition. **B.** Modulation by alpha lateralization index on RTs of correct trials. A significance tendency of reaction times was found, indicating a possible positive impact of increased alpha lateralization index on performance speed independent of eye condition.

rate, and the factors eye condition and RTs (Fig. 6).

For discrimination rate, a main effect of performance (F (1, 32) = 7.28, p < .01) was found, indicating increased accuracy with higher posterior alpha modulation. A main effect of eye condition (F (1, 32) = 42.24, p > .001) was found as well, with higher posterior alpha modulation for the EO condition. The interaction of eye condition and performance did not prove to be significant (F (1, 32) = 0.28, p = .6).

For the RTs, a main effect of performance (F (1, 32) = 11.1, p < .01), reflecting increased posterior alpha modulation for faster trials, and eye condition (F (1, 32) = 37.34, p < .001), with increased posterior alpha modulation for the EO condition, was found. The interaction of eye condition and performance was not significant (F (1, 32) = 0.002, p = .96).

In summary, these results indicate an impact of posterior alpha modulation on behaviour, with better performance following a higher increase of posterior alpha power during the pre-stimulus interval. This effect was furthermore independent of eye condition.

Discussion

In this study we aimed to investigate the effect of eye closure on attentional modulations of alpha oscillations and the subsequent behavioural performance during a somatosensory discrimination task. For this we adapted the paradigm by Haegens et al. (2011) to include an eyes-open and an eyes-closed condition. We replicated the lateralization pattern of attentional anticipatory alpha modulation, however found no significant impact on performance outcome. However, we do report a trend of faster reaction times with increased lateralization. Furthermore, though eye closure lead to a general increase of oscillatory activity, it also resulted in a reduced modulation of the posterior alpha power during the pre-stimulus interval in comparison with the eyes-open condition. The posterior alpha
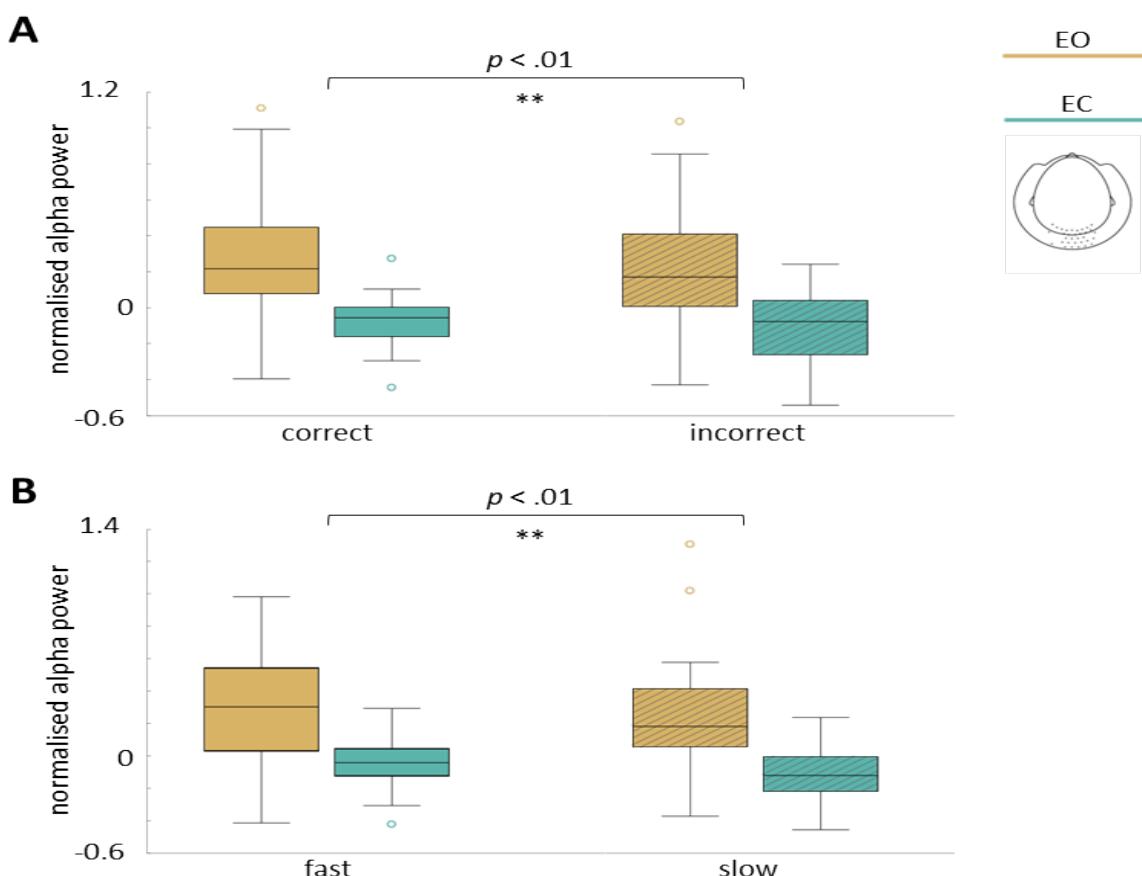
Figure 6. Behavioural performance modulation by posterior alpha power. The empty helmet represents the visual ROIs used for this analysis (summary of individual ROIs). P-values represent the main effect of performance. **A.** Modulation by posterior alpha power on discrimination rate (% correct responses). A significant difference of correct and incorrect trials was found, indicating an improved discrimination rate with higher posterior alpha modulation independent of eye condition. **B.** Modulation by posterior alpha power on RTs of correct trials. A significant difference of fast and slow trials was found, indicating a faster performance speed with higher posterior alpha modulation independent of eye condition.

modulation positively influenced task performance, independent of eye condition.

Eye closure leads to general increase of reaction times

Even though we did not formulate prior hypotheses on the effect of eye closure on general performance, we report here an increase of reaction times for the eyes-closed condition in comparison with eyes open, and a tendency of reduced accuracy. This speaks against the anecdotal assumption of improved performance with eye closure, which is thought to lead to better focus towards other sensory modalities (e.g., Glenberg et al., 1998). It furthermore contradicts previous findings of increased somatosensory perception detection associated with eye closure (Brodoehl et al., 2015a). However, these behavioural results are in accordance with a study conducted by Götz et al. (2017), who reported an impairment of somatosensory discrimination with eye closure accompanied by a reduction of somatosensory evoked fields. The authors proposed a possible trade-off between perception that requires spatial discrimination and perception that does not, implicating a negative effect of eye closure on spatial perception. This line of thinking follows from the assumption that eye closure does not only represent a control of visual input, but also influences processing mechanisms of somatosensory information.

Eye closure increases alpha power but decreases posterior alpha modulation

In agreement with previous studies (e.g., Adrian & Matthews, 1934; Geller et al., 2014; Wöstmann et al., 2019), we found an increase of alpha power with eye closure. We furthermore found an additional increase of beta power with eye closure, which was centred towards occipital sensors, while the alpha power increase showed a more widespread pattern. This widespread pattern of increased alpha power with eye closure is in accordance with previous findings (Geller et al., 2014) and indicates that the eye closure effect is not bound to occipital sources. The global increase of alpha further supports the assumption of an eye-closure effect on neural processing, unlike simply reflecting the disengagement of visual areas.

Though eye closure lead to a general increase of alpha power, we furthermore report a reduction of anticipatory posterior alpha modulation in comparison to the eyes-open condition. In this case, modulation of posterior alpha increase against baseline activity was more prominent, when participants had their eyes open than when their eyes were closed.

Anticipatory alpha power increase has been proposed to reflect a functional mechanism of inhibition to gate information (Jensen & Mazaheri, 2010; Foxe & Snyder, 2011). In the context of somatosensory attention, visual input is not relevant for the completion of the task and therefore increase of alpha activity over occipital regions was proposed to regulate an inhibition of the processing of visual input (Haegens et al., 2012). Our current findings further support this idea, as a much higher increase of posterior alpha synchronization is observed for the condition in which visual input is apparent.

Lateralized alpha modulation for both eyes open and eyes closed

In the context of spatial tactile attention, the inhibitory functional mechanism of alpha oscillations is further reflected by an increase of alpha activity over the ipsilateral and a decrease over the contralateral attentional somatosensory cortex (Haegens et al., 2011; Haegens et al., 2012). We replicated this pattern of contralateral alpha power decrease for both the eyes-open and the eyes-closed condition. Even though we do not observe an ipsilateral increase in our contrasts, this is possibly due to the fact that our study did not include the presentation of distractors. It has been shown, for both the visual (Rihs et al., 2007; Rihs, Michel, & Thut, 2009) as well as the somatosensory domain (Haegens et al., 2012), that ipsilateral increase of alpha is dependent on the presence and the strength of distractors engaging the task-irrelevant side.

The observed lateralization pattern reflected a general decrease of alpha over right somatosensory sensors for left-side attention, and left somatosensory sensors for right-side attention. This further supports the theory that alpha gates the information flow during somatosensory attention tasks, by facilitating excitability in the task relevant somatosensory cortex. However, this interpretation stems purely from the visualisation of spectral sensor activity. Further source analysis is necessary to contrast the activation over somatosensory cortices for a statistical test of significant contralateral alpha activity decrease. Since no source reconstruction was included for this report, no statistical claim and no certainty of the sources of this attentional modulation can be made at this point, as selection of somatosensory ROIs are imprecise on sensor level (see future directions section for more details).

We also investigated the evolution of pre-stimulus alpha modulation over time, which showed a sustained modulation for alpha (around 10 Hz) as well as for beta (around 20 Hz) oscillations for both eye conditions. This indicates that the anticipatory modulation of alpha started after the presentation of the cue and lasted until the presentation of the stimulus, in line with previous studies (e.g., Jensen & Mazaheri, 2010; Haegens et al., 2011; Haegens et al., 2012; Ikkai et al., 2016).

No impact of lateralization index on task performance

The lateralization index reflects the ratio of ipsilateral alpha activity increase and contralateral alpha activity decrease. In previous studies this alpha lateralization index modulated the task performance, leading to better accuracy and faster reaction times with increased lateralization index (Haegens et al., 2011; Haegens et al., 2012). However, we were not able to replicate accuracy effects here, though we did find a trend for RTs, indicating a possible impact of lateralization index on the speed of performance (independent of eye condition). It is worthy to note, that this analysis as well has to be fine-tuned for further interpretational gains (see future directions). Furthermore, for our analysis of RTs we calculated a separation of fast and slow trials based on a median split. This definition of RTs could be improved by binning RTs into multiple ranges from slow to fast (e.g., five instead of two), or by performing a single trial correlation. We plan to address these possibilities in future analysis.

For now, we were not able to replicate an influence on performance through the modulation of alpha lateralization and no influence of eye closure was reported. If these findings translate to later, more precise analyses this would implicate that alpha lateralization is not predictive of task performance.

Posterior alpha modulation increases task performance independent of eye condition

We replicated previous findings of an impact of posterior alpha modulation on task performance both for discrimination rate and RTs (Haegens et al., 2012). In this sense, higher posterior alpha was associated with more accurate and faster responses. This modulation was independent of eye condition. We therefore found evidence that the inhibition of visual processing leads to increased task performance. This is in line with the idea that through alpha increase over visual regions, resources are gated towards task-relevant brain regions, leading to more successful behavioural outcome. This gating process does not seem to be impacted by eye closure.

In summary, we did observe a steady increase of alpha activity with eye closure, yet the anticipatory alpha modulation showed a higher increase for the eye-open condition. Furthermore, even though posterior alpha modulation showed an impact on the behavioural outcome, this effect was independent of eye condition. This implicates, that the eye closure induced alpha increase and the inhibitory alpha modulation during the EO condition, follow the same mechanism to suppress visual regions and improve the somatosensory task performance.

Similar mechanism for alpha and beta?

Interestingly, we observed similar effects in the beta range (15 − 30 Hz) as for the alpha range. First of all, we observed an eye closure effect not only for alpha but also for beta oscillations. This is not surprising as previous research has indicated an effect of eye closure on all frequency bands to some extent (Geller et al., 2014). We here show — in contrast to previous findings (Geller et al., 2014) — that the beta effect of eye closure was limited to occipital sensors and does not show a widespread pattern like alpha. This could implicate, that with eye closure, unlike alpha, beta does not show a further influence on processing mechanisms outside the occipital cortex, but only reflects an impact of the disengagement of the visual domain.

Regarding anticipatory oscillatory modulation, we observed a sustained contralateral decrease of beta activity for the pre-stimulus interval, similar to the alpha activity. Previous studies proposed that beta follows a similar mechanism as alpha during somatosensory attentional tasks, reflected by a modulation of beta activity in the form of contralateral decrease in anticipation of expected stimuli and an impact of this modulation on performance speed (van Ede, Jensen, & Maris, 2010; van Ede, de Lange, Jensen, & Maris, 2011). However, a study by Haegens et al. (2012) also showed a possible dissociation of alpha and beta in the somatosensory domain, with ipsilateral increased alpha being accompanied by decreased beta activity (see Spitzer & Haegens, 2017, for further discussion on a potential alpha-like role of beta oscillations).

For future investigation of the interplay of these two oscillatory bands, we are planning to expand our analysis to also take into account a possible effect of beta modulation during somatosensory attention.

If the current trend of a similar beta activity to the alpha activity holds true, it would speak for further evidence of similar functions of the two oscillatory bands in the context of somatosensory tasks. Furthermore, it will be interesting to see whether the eye closure effect on beta oscillations, which showed to be specific for occipital sensors, shows a different modulation than for alpha oscillations.

Future directions

It has to be noted that the results reported here were all conducted on sensor level, which hold the constraint of a mixture of signals from different sources. This complicates possible assumptions of the origins of a signal. Therefore, the selection of our ROIs also followed a less spatially resolved categorisation, with possible contamination of auditory and visual sources. For future directions we plan to move to a source level representation of the data. For this the individual anatomical MRIs of the participants were collected to enable a source reconstruction using beamformer techniques, which allow to counteract the problem of spatial mixing and allow for more precise interpretation of effects.

Conclusion

We showed a general increase of alpha oscillations with eye closure as long established (Adrian & Matthews, 1934; Geller et al., 2014), with a widespread pattern of alpha, which implicates a possible influence of eye closure on processes outside the visual domain. Furthermore, the findings of this study support the previously proposed mechanism of alpha oscillations to gate information flow in the brain through functional inhibition (Jensen & Mazaheri, 2010; Foxe & Snyder, 2011). In this sense, we showed that alpha power over posterior sensors increased more when available visual input had to be actively inhibited as attention was guided towards the somatosensory domain. This posterior alpha activity increase furthermore has shown to be behaviourally relevant for both eyes conditions, as performance increased with higher alpha activity. This implies that with an increased inhibition of the visual domain, attentional resources are gated towards the somatosensory task, independent of eye condition. We replicated a pattern of contralateral alpha decrease over task-relevant somatosensory sensors, indicating a regulatory role of alpha for facilitation of excitability. However, we did not replicate previous findings of a somatosensory alpha lateralization impact on performance

outcome, which leads to the assumptions that the somatosensory lateralized alpha modulation might only reflect a neural process but does not further influence behaviour. Our behavioural results furthermore contradict an anecdotal assumption of performance improvement with eye closure, as we report an increase of reaction times for the eyes-closed condition. In summary, we provide novel insights into alpha activity during eye closure, showing evidence that while eye closure increased the global alpha activity it did not further impact the inhibitory alpha modulation on performance.

References

Adrian E.D., Matthews B.H.C. (1934). The berger rhythm: potential changes from the occipital lobes in man. *Brain, 57*(4), 355–85.

Banerjee, S., Snyder, A. C., Molholm, S., & Foxe, J. J. (2011). Oscillatory alpha-band mechanisms and the deployment of spatial attention to anticipated auditory and visual target locations: supramodal or sensory-specific control mechanisms? *J Neurosci, 31*(27), 9923-9932.

Berger, H. (1929). Über das Elektrenkephalogramm des Menschen. *Arch Psychiatr Nervenkr, 87*(1), 527-570.

Bonnefond, M., & Jensen, O. (2012). Alpha Oscillations Serve to Protect Working Memory Maintenance against Anticipated Distracters. *Current Biology, 22*, 1969-1974.

Brodoehl, S., Klingner, C., Stieglitz, K., & Witte, O. W. (2015a). The impact of eye closure on somatosensory perception in the elderly. *Behav Brain Res, 293*, 89-95.

Brodoehl, S., Klingner, C. M., & Witte, O. W. (2015b). Eye closure enhances dark night perceptions. *Sci Rep, 5*, 10515.

Dimitrijevic, A., Smith, M. L., Kadis, D. S., & Moore, D. R. (2017). Cortical Alpha Oscillations Predict Speech Intelligibility. *Front Hum Neurosci, 11*, 88.

Foxe, J. J., & Snyder, A. C. (2011). The Role of Alpha-Band Brain Oscillations as a Sensory Suppression Mechanism during Selective Attention. *Front Psychol, 2*, 154.

Geller, A. S., Burke, J. F., Sperling, M. R., Sharan, A. D., Litt, B., Baltuch, G. H., . . ., & Kahana, M. J. (2014). Eye closure causes widespread low-frequency power increase and focal gamma attenuation in the human electrocorticogram. *Clin Neurophysiol, 125*(9), 1764-1773.

Glenberg, A. M., Schroeder, J. L., & Robertson, D. A. (1998). Averting the gaze disengages the environment and facilitates remembering. *Memory & cognition, 26*(4), 651–658.

Götz, T., Hanke, D., Huonker, R., Weiss, T., Klingner, C., Brodoehl, S., Baumbach, P., & Witte, O. W. (2017). The Influence of Eye Closure on Somatosensory Discrimination: A Trade-off Between Simple Perception and Discrimination. *Cerebral cortex (New*

York, N.Y. : 1991), 27(6), 3231–3239.

Haegens, S., Cousijn, H., Wallis, G., Harrison, P. J., & Nobre, A. C. (2014). Inter- and intra-individual variability in alpha peak frequency. *NeuroImage*, *92*(100), 46–55.

Haegens, S., Händel, B. F., & Jensen, O. (2011). Top–down controlled alpha band activity in somatosensory areas determines behavioral performance in a discrimination task. *Journal of Neuroscience, 31*, 5197–5204.

Haegens, S., Luther, L., & Jensen, O. (2012). Somatosensory Anticipatory Alpha Activity Increases to Suppress Distracting Input. *Journal of Cognitive Neuroscience, 24*, 677-685.

Haegens, S., Nácher, V., Luna, R., Romo, R., & Jensen, O. (2011). α-Oscillations in the monkey sensorimotor network influence discrimination performance by rhythmical inhibition of neuronal spiking. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(48), 19377–19382.

Haegens, S., Osipova, D., Oostenveld, R., & Jensen, O. (2010). Somatosensory working memory performance in humans depends on both engagement and disengagement of regions in a distributed network. *Human Brain Mapping, 31*, 26–35.

Ikkai, A., Dandekar, S., & Curtis, C.E. (2016). Lateralization in Alpha-Band Oscillations Predicts the Locus and Spatial Distribution of Attention. *PLoS ONE*, 11.

Jensen, O., & Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: gating by inhibition. *Front Hum Neurosci, 4*, 186.

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods, 164,* 177–190.

Kelly, S. P., Gomez-Ramirez, M., & Foxe, J. J. (2009). The strength of anticipatory spatial biasing predicts target discrimination at attended locations: a high-density EEG study. *The European journal of neuroscience*, *30*(11), 2224–2234.

Klimesch, W., Sauseng, P., & Hanslmayr, S. (2007). EEG alpha oscillations: the inhibition-timing hypothesis. *Brain research reviews*, *53*(1), 63–88.

Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J.-M. (2011). FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational Intelligence and Neuroscience, 2011(156869)*, 1–9.

Pfurtscheller, G., & Klimesch, W. (1991). Event-related desynchronization during motor behavior and visual information processing. *Electroencephalography and clinical neurophysiology. Supplement*, *42*, 58–65.

Rihs, T. A., Michel, C. M., & Thut, G. (2007). Mechanisms of selective inhibition in visual spatial attention are indexed by α-band EEG synchronization. *European Journal of Neuroscience, 25*, 603–610.

Rihs, T. A., Michel, C. M., & Thut, G. (2009). A bias for posterior alpha-band power suppression versus enhancement during shifting versus maintenance of spatial attention. *Neuroimage, 44*, 190–199.

Sauseng, P., Klimesch, W., Stadler, W., Schabus, M.,

Doppelmayr, M., Hanslmayr, S., et al. (2005). A shift of visual spatial attention is selectively associated with human EEG alpha activity. European Journal of Neuroscience, 22, 2917–2926.

Spitzer, B., & Haegens, S. (2017). Beyond the Status Quo: A Role for Beta Oscillations in Endogenous Content (Re)Activation. *eNeuro, 4*(4), ENEURO.0170-17.2017.

Strauß, A., Wöstmann, M., & Obleser, J. (2014). Cortical alpha oscillations as a tool for auditory selective inhibition. *Front Hum Neurosci,* 8:350.

Thut, G., Nietzel, A., Brandt, S. A., & Pascual-Leone, A. (2006). Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. Journal of Neuroscience, 26, 9494–9502.

van Albada, S. J., & Robinson, P. A. (2013). Relationships between Electroencephalographic Spectral Peaks Across Frequency Bands. *Frontiers in human neuroscience*, *7*, 56.

van Ede, F., de Lange, F., Jensen, O., & Maris, E. (2011). Orienting attention to an upcoming tactile event involves a spatially and temporally specific modulation of sensorimotor alpha- and beta-band oscillations. Journal of Neuroscience, 31, 2016–2024.

van Ede, F., Jensen, O., & Maris, E. (2010). Tactile expectation modulates prestimulus beta-band oscillations in human sensorimotor cortex. Neuroimage, 51, 867–876.

Vredeveldt, A., Baddeley, A. D., & Hitch, G. J. (2012). The Effects of Eye-Closure and "Ear-Closure" on Recall of Visual and Auditory Aspects of a Criminal Event. *Europe's Journal of Psychology, 8*(2), 284-299.

Vredeveldt, A., Tredoux, C. G., Nortje, A., Kempen, K., Puljevic, C., & Labuschagne, G. N. (2015). A field evaluation of the Eye-Closure Interview with witnesses of serious crimes. *Law Hum Behav, 39*(2), 189-197.

Worden, M. S., Foxe, J. J., Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. *Journal of Neuroscience, 20,* RC63.

Wöstmann, M., Schmitt, L., & Obleser, J. (2019). Does Closing the Eyes Enhance Auditory Attention? Eye Closure Increases Attentional Alpha-Power Modulation but Not Listening Performance. *Journal of Cognitive Neuroscience, 32*, 212-225.

Supplementary Material

Methods of visualisation

To visualize attention effects within the pre-stimulus interval, we evaluated the pre-stimulus FFT based on whether attention was being guided to the left or the right hand (dependent on the auditory cue). For this, we divided trials into four conditions: EO attention-right, EO attention-left, EC attention-right, and EC attention-left. A baseline correction was applied based on the FFT of the baseline window -1 s to 0 s of cue-onset. Baseline corrections were done separately for EO and EC trials. For each attention side of each condition the whole-head power distributions were computed.

For a visualisation of the lateralization pattern the pre-stimulus FFT was once again divided into the four conditions, based on eye-condition and attention-side, but not baseline corrected. Instead, a normalization in the form of (attention-left − attention-right) / (attention-left + attention-right) was computed for each eye-condition separately. This normalization leads to positive values if there is a stronger decrease in the attention-right condition and negative values for a stronger decrease in the attention-left condition.

To obtain a representation of the modulation of alpha activity over time, the TFR of the pre-stimulus interval (-1 s − 0 s) was divided into EO and EC trials and baseline corrected (with baseline -1.5 s − -1 s) for each condition separately. As before, the left vs right normalization of attention sides was applied for the two conditions. Only the individual somatosensory ROIs were used for this visualisation. Left hemispheric sensors were mirrored to combine them with the right hemispheric sensors, leading to a general representation of the decrease of activity. Power spectra of low frequencies in the range of 5 to 30 Hz were used for this visualisation.

A Study of Laminar fMRI and Directed Connectivity During Sentence Processing: A Pilot Study

Lenno R. P. T. Ruijters[1,2]
Supervisors: Daniel Sharoh[1,2], Kirsten Weber[1,2]

[1]*Radboud University Nijmegen, Donders Institute for Brain, Cognition and Behaviour, The Netherlands*
[2]*Max Planck Institute for Psycholinguistics, The Netherlands*

It is currently unclear how the language network in the brain evokes a concept that is the result of semantic composition. We conducted a pilot study laying the groundworks for a laminar functional magnetic resonance imaging (fMRI) study that aims to investigate this matter further. A paradigm was developed where target words were either primed through compositional semantics, identity primes, neither, or both. A behavioural experiment showed that the identity priming manipulation successfully reduced reaction times in a lexical decision task, which was the predicted behavioural effect. We predicted and found a modulated blood oxygen level-dependent (BOLD) response in the left inferior frontal gyrus and middle temporal gyrus to target words that were compositionally primed. This work will inform a subsequent laminar fMRI study, where the directed connectivity during sentence processing between these regions will be examined.

Keywords: sentence processing, compositional meaning, priming, laminar fMRI, pilot, IFG, MTG

Corresponding author: Lenno Ruijters; E-mail: l.ruijters@spinozacentre.nl

When using a statement like "It's getting cold" as a request to close a window, one can see that an utterance can have meaning beyond the word-level. The example shows that it is possible to produce utterances that, as a whole, evoke a certain image unattributable to the individual words making up the sentence. A sentence such as, "The man followed the woman and lifted the wallet from her bag," would likely evoke the image of a pickpocket. The distinction between the word-meaning and the sentence-meaning becomes clearer when the word order is switched around. In "The woman followed the man and lifted the wallet from her bag," the idea of a pickpocket is largely irrelevant, even though the individual words in both sentences are identical. We will refer to this sentence-level meaning as the compositional meaning of the sentence, and the processes by which it is derived as composition.

Even though speakers constantly derive the compositional meaning of utterances they encounter, this topic has not received a great deal of attention in the neuroimaging and priming literature. Common psychological and neurobiological models of sentence processing acknowledge the presence of some form of semantic composition by including a step where listeners combine words into a coherent whole, but neglect to speculate how this compositional meaning is eventually retrieved (Ferreira & Lowder, 2016; Frazier & Clifton, 1996; 1997; MacDonald et al., 1994; Vosse & Kempen, 2000; Humphries et al., 2006; 2007; Friederici, 2011; Hagoort, 2005; 2013; 2016). If the syntactic and semantic components of an utterance are the recipe for the compositional meaning, then the current models would be akin to a cooking class that ended once eggs, flour and milk are mixed into a bowl. A critical subsequent step is missing to get to a fully baked conclusion. It is currently understudied how the networks responsible for the composition of language evoke a concept or image as a result of complex input.

The common language network

Computational and neurobiological models of sentence processing have both set out to explain the processes that give rise to sentence processing. Computational models set out to study what computations (e.g., statistical inferences or network behaviour) may lay at the root of sentence processing. Neurobiological models in the meantime try to understand which sections of the brain are involved, what the role of each section might be in sentence processing, and how they interplay. Research on the

neurobiology of language has converged on a model that includes a large, mostly left-lateralised network. This network includes a large array of regions: from the visual word-form area on the occipito-temporal border, involved in reading (Dehaene & Cohen, 2011) and bilateral temporal cortex, which has been linked to speech sound processing (McGettigan & Scott, 2012), to the inferior frontal gyrus (IFG; Thompson-Shill et al., 1997) and angular gyrus (AG; Binder et al., 2009). We will focus on the properties of the left middle temporal gyrus (MTG), IFG, and AG, as these appear to play a more prominent role in semantic processing of both words and sentences (Friederici, 2011; Hagoort, 2005; 2013; 2016). Understanding this network is then crucial to understanding how compositional meaning is derived from utterances, and vice versa.

When it comes to developing a map of the language network, modern neurobiological models of language overlap in key aspects, for example the importance of the temporal lobe. Involvement of the MTG seems to be critical in semantic processing, lexical memory retrieval, and conceptualisation (Whitney et al., 2010; Snijders et al., 2010; Visser et al., 2012; Hagoort, 2013; 2016; Davey et al. 2016). Structurally, its spatial proximity to the hippocampus, as well as its pronounced structural connections to the occipital, frontal, and parietal lobes make it a strong candidate region for these functions (Turken & Dronkers, 2011). The connections between MTG and frontal cortex are further supported by Xiang et al. (2010), where a functional connectivity map of IFG based on resting state fMRI data showed links with MTG in the perisylvian language network. Here, functional connectivity shows how regions functionally communicate, beyond the structural pathways that may be in place. Due to its wide connections, the MTG is believed to be involved in conceptual representation processes (Snijders et al., 2010; Hagoort, 2013; 2016). Further frontotemporal connections are proposed by Friederici (2011), who lays out the structural pathways through which the temporal cortex communicates with the frontal lobe. The posterior superior temporal gyrus (STG) and BA44 are connected through a dorsal pathway, and the anterior STG is connected with BA45 and the frontal operculum through two ventral pathways (Friederici, 2011). These connections illustrate its potential for complex language processing.

Hagoort (2013; 2016) and Humphries et al. (2006) seem to agree that the MTG is involved in word-level semantic processing. Their claim is supported by, among others, Snijders et al. (2009), which tested the computational Unification Space model proposed

by Vosse and Kempen (2000), and found that MTG activity is related to lexical retrieval, including word-level semantics. However, Humphries et al. (2006) also agree with Friederici (2011) in suggesting that the MTG (as well as the STG and the anterior temporal lobe regions) are additionally involved in semantic relation calculations, such as ambiguity resolution. The MTG's connectivity profile nominates it as a candidate for this early composition function, as it both centrally connects brain regions implicated in language processing and since it is sensitive to lexical retrieval demands. Despite the different interpretations of MTG's function in language processing, however, it is generally accepted that the MTG is involved in language processing in some capacity.

As the demand for integration of new input increases, so does IFG activity. It is considered critical to composition processing, be it for relation calculations (Friederici, 2011) or for composition directly (Hagoort, 2013; 2016). Hagoort (2016) provides a description of the connections between the MTG and the IFG, where the two regions are in continuous back-and-forth with one another during ongoing integration of a stream of input. Through feedback loops, context of the overall input is repeatedly adjusted (for a more detailed study of the connectivity profile of the IFG, see Xiang et al., 2010). Thompson-Shill et al. (1997) found that response selection on the basis of semantic relations involved the IFG, but simple retrieval processes did not. This suggests that the IFG is indeed involved when semantic properties of disparate items need to be evaluated (see also Roskies et al., 2006). Zhu et al. (2009) found that small semantic violations evoked larger IFG responses (corresponding with longer rejection times) than large violations (which had shorter rejection times). The differential IFG activity in semantic manipulations suggests that the IFG is involved in semantic processing (as opposed to exclusively syntactic processes), where a higher degree of successful integration leads to more activity. However, Zhu et al.'s (2009) results could also be interpreted as resulting from re-evaluation difficulties, where large violations simply are given up on, and small violations require more effort. Regardless, IFG activity appears to be related to the extraction of meaning from input that requires composition to interpret correctly. Studies on the time course of MTG and IFG activity provide converging evidence for a functional hierarchy between these two regions. Frontal activity during sentence processing follows temporal activity (100-300ms in Friederici, 2011; ~150ms in Hagoort, 2013),

suggesting that frontal processes during sentence processing rely on input from the temporal cortex. Again, different models propose different functions of the IFG. Humphries et al. (2006) found IFG activity after both coherent and incoherent complex input, suggesting that the IFG is always active for complex input, and Friederici (2011) suggests that IFG activity in language processing is largely attributable to syntactic processing. Here we will assume that the IFG is the locus for compositional processes, following the Memory, Unification, and Control (MUC) model (Hagoort, 2013; 2016).

The model that we take as a starting-point in this study is Hagoort's (2016) MUC model, due to the fact that it makes clear predictions of feedback between its titular unification and memory components. Following this model, we expect that we can measure both the effect of input in the MTG through retrieval processes, and the effect of composition as a result of a neural feedback system in the MTG (Snijders et al., 2010; Baggio & Hagoort, 2011). The purpose of the present study is to determine whether the IFG communicates with the MTG in feedback fashion during compositional meaning processes, to evoke a unified representation. However, we deviate from the MUC model in one aspect. We would expect that compositional meaning elicitation could also involve other regions, like the AG, as some literature suggests (e.g., Humphries et al., 2006; Binder et al., 2009), which the MUC model proposes is involved in retrieval processes (Hagoort, 2016). Evidence for how or why the AG is involved in this process is currently still lacking, however. Overall, we will maintain an exploratory approach, in order to further develop a fully comprehensive model of sentence (or language) processing in the brain.

Laminar fMRI

The experiment discussed in this thesis will serve as a pilot study for subsequent work investigating the behaviour of laminar structures in relation to processing compositional meaning. It was as such designed with laminar resolution functional magnetic resonance imaging (fMRI) in mind.

Laminar fMRI refers to the acquisition of fMRI data at a high-enough resolution for the blood oxygen level-dependent (BOLD) -response in the grey matter to be separated into several bins over depth (conventionally voxel sizes <1mm³). The responses within these bins are then looked at independently. Based on the assumption that layer-specific activity is attributable to top-down or bottom-up processes (e.g., Kok et al., 2016; Sharoh

et al., 2019), it becomes possible to study how brain regions interact with one another at different levels. Here we define top-down processes as information being relayed from a hierarchically higher region to a hierarchically lower region (i.e., feedback). In more concrete terms, activity in the middle layer of isocortex is generally evoked by signals from brain regions that are involved at a lower level in a given process than a reference region (i.e., bottom-up processes); meanwhile activity in the superficial and deeper layers of isocortex are commonly attributed to input from higher order regions (i.e., top-down processes; Harris & Mrsic-Flogel, 2013; Lefort, et al., 2009; Brown & Hestrin, 2009).

Studying these top-down and bottom-up connections in vivo in this way was inaccessible until the advent of laminar fMRI (Koopmans et al., 2010; Olman et al., 2012; Kok et al., 2016; Sharoh et al., 2019). It is important to understand these connections, since a lot of the brain's functioning relies on attentional and world-knowledge restrictions that are imposed by higher-order areas onto, for example, sensory brain regions (e.g., Hagoort et al., 2004; Kok et al., 2016). One reason why feedback during sentence processing is currently understudied, is because there are notable difficulties in capturing this activity using non-invasive neuroimaging techniques. One way to study feedback processes in the brain would be to look at the functional time-course of different brain regions, and see how activity in higher-order regions temporally correlates to activity in lower-order regions. In fMRI, this is difficult because the low temporal resolution of the BOLD response, on which fMRI relies, makes it difficult to disentangle activity patterns with high temporal overlap. Additionally, the variations in the hemodynamic response make it difficult to compare the temporal relations between activity patterns (Aguirre et al., 1998; Kastrup et al., 1999). These properties of the signals that are measured in fMRI cause problems when inferring which underlying activity preceded which. Electroencephalography (EEG) and magnetoencephalography (MEG) in the meantime suffer from their own restrictions. Even with the high temporal resolution of M/EEG, which allows these methods to capture the fine temporal aspects of communication between the IFG and MTG, and the relatively high spatial resolution of MEG, it is potentially problematic to decide the causal relationship of temporally related activity patterns. This is especially the case in on-going processes such as sentence processing, where continuous activity causes temporal overlap (but see Lyu et al., 2019).

Laminar fMRI, in combination with specialised analysis methods, offers an alternative to the strictly temporally-based method. Laminar fMRI enables the study of directed connections on the basis of structural properties of the in vivo brain (Koopmans et al., 2010; Olman et al., 2012; Kok et al., 2016; Sharoh et al., 2019). By looking at layer-specific activity in laminar fMRI, the issue of low temporal resolution is circumvented when looking at interregional interactions, by looking at which layers of different regions respond to a given stimulus (e.g., Kok et al., 2016; Sharoh et al., 2019). Assuming different layers receive input from different hierarchical directions, by finding an interaction between different layers of two regions, it is plausible to infer how these two regions communicate, especially if there is evidence for a directed correlation between regions. By designing stimuli that differentially evoke bottom-up and top-down activity in a region(s) of interest (ROI), it becomes possible to disentangle which regions communicate with this ROI, in what order, and what proportion of the measured activity within a region relates to bottom-up and top-down processes.

Our manipulation

Behavioural priming is the psychological effect where, once a person is presented with a stimulus, their reaction to a similar stimulus is affected compared to if an unrelated stimulus were presented. An example of these priming effects in language is identity priming, where recognition or production of a word is enhanced if it has been presented before (e.g., Buckner et al., 2000; Raposo et al., 2006). There is also semantic priming, where processing of an item is once again enhanced if it is semantically related to a preceding stimulus (e.g., Hutchison, 2003; Rissman et al., 2003; Rossell et al., 2003; Wible et al., 2006), or syntactic priming, where sentence structures are comprehended or produced more easily if they have been presented earlier (Pickering & Branigan, 1999). Moreover, these effects have been found in neuroimaging experiments in the form of modulated BOLD or event-related potential (ERP) responses (for a review, see Segaert et al., 2013). Our assumption is that identity priming is a mix of lower order priming and semantic priming, and that we can evoke a top-down semantic priming effect through compositional primes. We expect that by manipulating the degree of different forms of priming, we can disentangle the top-down and bottom-up processes that lead to a compositional interpretation of linguistic input. We will further

discuss identity and semantic priming, as these forms of priming were used for our manipulation.

At the neuronal level, these priming effects manifest as a modulated response of the neuron due to repeated activation. While this often goes paired with a behavioural effect as well (for a review on semantic priming, see Hutchison, 2003), we will focus on the neural response in light of this study. The most straightforward and strongest effect is that of identity priming. If an individual is stimulated in an identical fashion twice in a row, the neural response to the second stimulation is modulated by the first. Most commonly, this effect is found in the form of repetition suppression (Segaert et al., 2013). This effect has been found in visual priming (Summerfield et al., 2008), auditory priming (Bergerbest et al., 2006), motor priming (de C. Hamilton & Grafton, 2008), and linguistic priming (Buckner et al., 2000; Raposo et al., 2006). We expect that we can modulate the lower order representations of semantic concepts by repeating the lower level stimulation. It should be noted, however, that repetition priming has also been linked to repetition enhancement effects. Segaert et al. (2013) concluded in their review that this enhancement is likely due to feedback effects as well, where the direction of the effect is task dependant. We nonetheless expect part of this offset to be driven by bottom-up activity modulating the neural response to the presentation of an identity-primed target word.

Semantic priming is one step removed from this. Semantic priming is the pre-activation of semantic representations based on a non-identical, semantically related prime. Behavioural experiments have previously shown that reaction times are faster in a number of tasks when a target is semantically primed against a non-primed condition (for a review, see Hutchison, 2003). At the word-level, the semantic priming effect has been found in several fMRI studies, where the STG and the MTG showed stronger activity in the unrelated condition than in the related condition (Rissman et al., 2003; Rossell et al., 2003; Wible et al., 2006). Notably, the priming effect appears to be stronger for identity priming than for semantic priming (e.g., Rugg, 1985). Moreover, the priming effects of linguistic context and repetition appear to be additive (Den Heyer et al., 1985). While Den Heyer et al. (1985) only studied these priming effects in a behavioural study, it suggests that the effect of repetition priming is not at ceiling and can be enhanced with context priming. As such, the combined effect can be distinguished from the individual effects. In addition, if we can find this additive effect in behaviour, it stands to

reason that these individual effects have their unique (albeit potentially overlapping) neural fingerprints. Whether the semantic priming effect is the result of feedback from higher regions, communication of neurons on the same hierarchical level, or both is currently uncertain. It is worth noting that expectancy, a top-down effect, can influence the size of the semantic priming effect (e.g., Keefe & Neely, 1990), but it is unclear how such effects inform the expected effects of a compositional prime.

As with identity priming, semantic priming may induce enhancement effects, depending on task or ROI (e.g., Rossell et al., 2003; for a review, see Segaert et al., 2013), but again we expect that semantically priming a target evokes a modulation nonetheless. If we assume that a (deliberately constructed) sentence can elicit a unique semantic representation, then this representation should semantically prime related targets (similar to the context effect in Den Heyer et al., 1985, but at the sentence level). We expect a modulation as a result of the IFG communicating the output of composition to the MTG in the case of these compositional primes, without lower order priming effects. In response to the compositional stimuli, we expect IFG activity to be higher compared to a non-compositional condition, and we expect this activity to be related to MTG activity. By manipulating the compositional semantic representations, we predict we are able to study the top-down effects of the IFG onto the linguistic memory network through semantic priming effects (be they in the form of suppression or enhancement). It should be pointed out that identity priming may include a semantic priming effect, since prime and target are semantically related, but it includes a lower order priming effect as well. We do not expect a correlation with IFG activity and the identity manipulation.

Our main interest is to see how the interpretation of utterances would neurally prime a target word. To do this, we want to look at the differences between priming of a target word by an identity token and by a complex, compositionally related token. We assume that priming effects from a compositional token would result from higher-order composition processes. These higher-order processes proceed to activate a combined conceptual representation that would be relayed back to the memory compartment of the language system. This feedback-induced priming effect would be identifiable by its laminar profile. In order to allow the study of the laminar activity profile of top-down processes in the future, we manipulate the bottom-up effect as well. This way, changes in the BOLD response between

layers and conditions may be disentangled. We assumed that the priming effect of the identity prime includes a bottom-up effect in the MTG. Combining these factors, we end up with four conditions: a compositional prime (sentence) that includes an identity prime (i.e., a copy of the target) (+Comp/+ID); a compositional prime (sentence) where the identity prime is replaced with another word, but is otherwise identical to the +Comp/+ID prime (+Comp/-ID); a non-compositional prime, which is a scrambled list of words, one of which is the identity prime for the target (-Comp/+ID); a non-compositional prime (scrambled list of words) that does not contain the identity prime (-Comp/-ID).

Since this is a pilot study, our predictions will remain general. We predict that we find a priming effect in the left MTG for both the identity and compositional priming conditions. We expect that the IFG is more strongly involved in our compositional condition. We expect that MTG activity shows a repetition effect for the identity conditions, and a semantically primed pattern for the compositional conditions. We also might expect some AG activity in response to the compositional primes.

Methods

Design

We designed an experiment to identify potentially distinct language processing mechanisms through different priming methods. Our paradigm consisted of a 2*2 design with compositional priming (on vs. off) by identity priming (on vs. off) manipulations. Primes consisted of strings of words. Targets consisted of single words. The prime strings were either intact Dutch sentences (Sentence), or random strings of Dutch words (Scramble). The sentences were always designed to be compositionally equivalent to the target word (compositional primes), whereas the scrambles were designed not to be (non-compositional primes). Our goal was to semantically prime the target with the compositional primes. In addition, we manipulated whether the target word was present in the prime string (identity prime) or replaced by a word that did not prime that target by itself (non-identity prime). See Table 1 for a schematic overview of the study design. Target words never exceeded 13 characters in length to fit within the reading span limit imposed by Rayner (2009) and consisted only of nouns and (inflected) verbs. Scrambles were created so that there was no intrinsic semantic priming at both the compositional and word levels (see subsection scrambling for the scrambling procedure).

Stimulus creation

Stimuli were constructed in several steps outlined below. Two native Dutch speakers constructed 236 sentences for the compositional prime conditions. 126 of these sentences were adapted from Lai et al. (2015). Each sentence was constructed with a target in mind that semantically summarised the prime (e.g., PRIME: "The man followed the woman and lifted the wallet from her bag"; TARGET: "Pickpocket"). Additionally, the target had to be able

Table 1. The primes for the target "verdrinking" (drowning) in each condition. Original Dutch presented in black, followed by the grey English translation. The translation is intended to reflect the intended message of the sentence and may not accurately reflect the word-count or the semantic relations between individual items and the target of the Dutch item.

Prime-type	Compositional	Non-compositional
Identity	Ze probeerde haar hoofd boven het oppervlak te houden maar kon de **verdrinking** niet meer bevechten.	Door keek naar hadden had linker torso waren de opzocht achter rieten **verdrinking** de op niet.
	She tried to keep her head above the surface, but could not fight the **drowning** any longer.	By looked to had had left torso were de searched behind reed **drowning** the on not.
Non-identity	Ze probeerde haar hoofd boven het oppervlak te houden maar kon de **stroming** niet meer bevechten.	Door keek naar hadden had linker torso waren de opzocht achter rieten **stroming** de op niet.
	She tried to keep her head above the surface, but could not fight the **current** any longer.	By looked to had had left torso were de searched behind reed **current** the on not.

to replace an item in the sentence without altering the interpretation or interpretability of the sentence (e.g., PRIME: "The pickpocket followed the woman and lifted the wallet from her bag"; TARGET: "Pickpocket"). The stimuli based on Lai et al. (2015) contained a valence confound. We subsequently constructed our items so that half of the final set would contain negative valence composition, and the other half neutral valence composition. Valence was considered an orthogonal manipulation but was not included in the pilot analyses.

Nearest neighbour test. Following stimulus construction, we confirmed that the individual lexical items in the primes were not likely to semantically prime the target by themselves. The vector space model for semantic distance snaut was used for this purpose (Mandera et al., 2017) and rank nearest neighbour was taken as our distance metric (Lapesa and Evert, 2013; 2014). Lemma-level rather than word-level relations were used to avoid the influence of syntactic information. Nearest neighbour ranks of lemmas were based on the cosine distance between vector representations (Kenett et al., 2017; Ettinger et al., 2016; Lenci, 2018). Lexical items that fell outside the first 40 nearest semantic neighbours were considered non-priming. We verified both forwards and backwards distance values.

Frequency in the training data of snaut seemed to affect the reliability of the semantic distance values. Vector values that were based on insufficient data were taken to be unstable. We determined an item needed at least 500 exemplars in SoNaR-500 (Oosterdijk et al., 2013) . Word-level distance was considered if snaut reported no instances for the lemma-level distance, and if the frequency threshold was met in SoNaR-500. Where spotted and possible, if the item used in the stimuli had a low frequency (<500), a synonym with a higher frequency was used in snaut (e.g., kots instead of braaksel; both "vomit"). In case the target had a SoNaR-500 frequency of <500 and no reasonable alternative form was available, the target was replaced entirely. Two items were excluded after this step (resulting N = 234).

Independent prime-target associations. Following the nearest-neighbour assessment, one RA and two interns were asked to provide at least three associations for each prime independently. If at least two of these judges agreed on a target, the target was accepted for later steps (be it our target or a new target). If one agreed with our target and/or the rest gave associations close to our target, the target

was checked and either changed accordingly, or left unchanged for later steps. If targets were changed, they were checked against snaut again. If there was no consensus on a target, or if implementing a new target would violate the other criteria for our stimuli, the entire item was removed from the stimulus list. 223 items remained in the stimulus list after this process.

Assessing stimulus frequency. A frequency analysis of the lemmas of our items was performed next using the subtitle corpus SubtLex-NL (Keuleers et al., 2010; see also New et al., 2007; Dimitropoulou et al., 2010). This was done to reduce the influence of item frequency on our measurements. Non-content words were deleted from the stimulus list, and remaining items were replaced with their uninflected forms. The raw lemma frequency for each content word was extracted and the 10log for each item was calculated (Howes & Solomon, 1951; but see also Murray & Forster, 2004). Next, the mean and standard deviation for targets (M = 2.74, SD = 0.712, n = 234) and prime content items (M = 3.60, SD = 1.10, n = 1606) were calculated separately. The outcome indicated that targets were on average 10 times less frequent than prime content items. This was to be expected since target items were often more specific (i.e., less underspecified) than the items they replace in the primes, and thus less frequent. We ensured that no items were more or less frequent than three times the standard deviation from the mean. This was done separately for target items and prime content items. Six items fell outside the 3SD range (one of 234 targets and five of 1606 sentence items). These were replaced and the semantic relations of these new items to the targets/primes were confirmed using snaut. A subsequent identical check was done at 2SD, to ensure that, while all items were within 3SD, the distribution was not weighted towards the tails. Here 52 items stood out (six of 234 targets, 46 of 1606 sentence items). This was considered to be within acceptable limits.

Valence check. As the stimuli contained a valence manipulation, we verified that independent participants replicated our valence labelling. A valence judgement survey was constructed for LimeSurvey (Limesurvey GmbH). Items had to be judged on a 1-to-9 scale (1 being negatively valent, 5 being neutral, 9 being positively valent), using the Self-Assessment mannequin (Bradley & Lang, 1994). Four separate surveys were developed totalling all individual lexical items in our stimuli (N = 1107, 223 Target items), and another survey for

the compositional primes as a whole (N = 2342). Subjects were recruited using the Max Planck Institute for Psycholinguistics participant database until each list had exactly ten complete responses (total N = 57, male = 8). The study was conducted according to the institutional guidelines of the local ethics committees. The survey was conducted online and in an environment unknown to the experimenters. Subjects were encouraged to complete the survey on a laptop or desktop computer (mobile N = 4). Subjects could pause and recommence the survey at any time. Informed consent was collected through a confirmation screen at the start of the questionnaire, and subjects were reimbursed for 30 minutes of their time. Individual responses were excluded on the basis of the item-specific SD and mean, where responses > 2SD from the item mean were excluded. The mean and SD for each item was then recalculated. All three categories showed evidence for high linearity (sentences, targets, and sentence items). Rank correlation for prime valence was ϱ = 0.991 (Fig. 1A), for target valence ϱ = 0.985 (Fig. 1B), and for prime content item valence ϱ = 0.962. When comparing the pre-labelled valence groups (negative vs. neutral) on valence scores using a t-test assuming inequality between groups (negative = 115; neutral = 108), scores showed a significant difference between the negative-labelled items and the neutral-labelled items (P = 2.0e-53, ci = -2.64; -2.20, t-stat = -21.6 and DF = 193.15). Testing for a correlation between prime scores and target scores showed evidence for a correlation, after removing prime-target pairs with a difference in valence scores > 3 SD (= 3 * 0.76) from the mean difference (=

1.0) (Nremoved = 2). Due to an unidentified error, one item was missing from the resulting list. The resulting list included 220 items with a correlation coefficient of ϱ = 0.81 between prime valence and target valance (Fig. 2) at P = 3.7e-49. Prime content items had a mean valence value of 5.2, with an SD of 1.60 after removing outlier responses. Due to the high linearity in the valence scores, we decided to add the prime and target valence scores as regressors for each item for the analysis in the main study, rather than rely on the binary negative/neutral labels.

Semantic relatedness test. The remaining items were then assessed in terms of their semantic relatedness. A semantic relatedness judgment task was created using LimeSurvey (Limesurvey GmbH). Participants who participated in the valence survey were excluded from participating in this survey. Target items were judged on their semantic relatedness to a preceding sentence on a 1-to-5 scale. Sentences could not be revisited once the target was presented. Our stimuli were split into two counterbalanced lists. In each list, half the target items remained paired with their intended prime, and half the targets were swapped randomly between primes. These new pairs were then checked to make sure they were not judged as semantically related at the sentence level by one of the researchers. Next, two versions of both lists were created, where each version contained ~50% +Comp/-ID primes and ~50% +Comp/+ID primes, to ensure that the identity manipulation did not have a detrimental effect on the semantic relatedness judgment. This resulted in a total of four conditions (related vs. unrelated target * +Comp/-
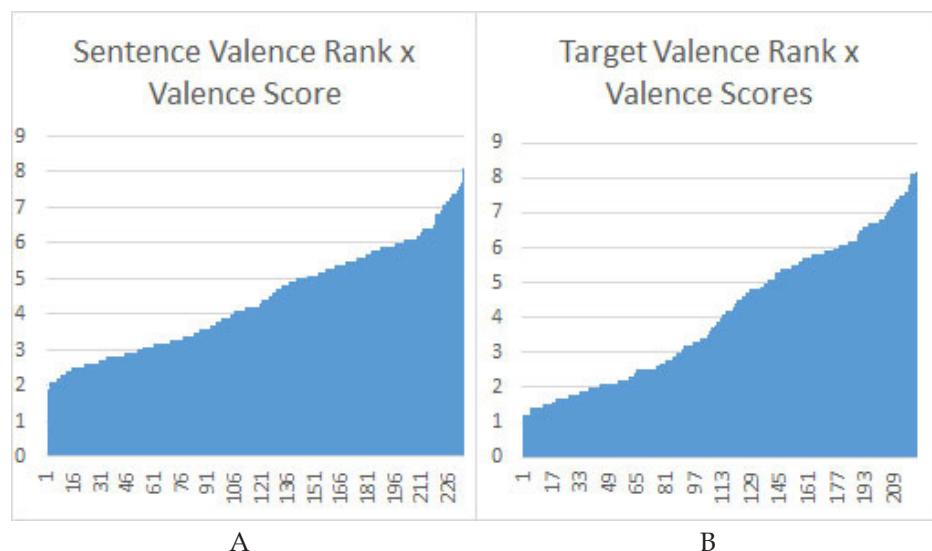

Figure 1. A. Valence scores ordered by rank for each prime. **B.** Valence scores ordered by rank for targets. Both primes and targets show high linearity between score and rank (ρ = 0.991 and ρ = 0.985 respectively).
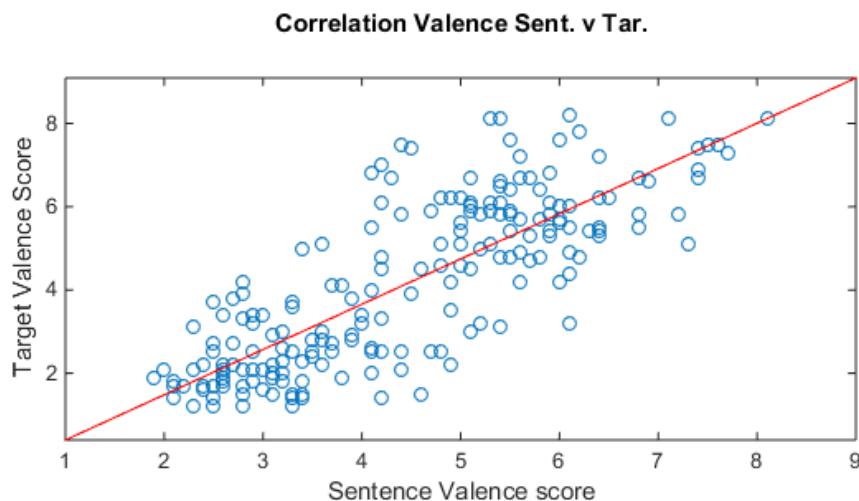
Figure 2. The correlation between the mean valence score for each sentence and their attributed target words, $\rho = 0.81$.

ID vs +Comp/+ID prime) for each stimulus across lists (stimuli per list = 221). Subjects were recruited using the Max Planck Institute for Psycholinguistics participant database until each list had a total of five complete responses (total N = 22). The study was conducted according to the institutional guidelines of the local ethics committees. The survey was conducted online and in an environment unknown to the experimenters. Subjects were encouraged to complete the survey on a laptop or desktop computer (tablet N = 1, mobile N = 1). Subjects could pause and recommence the survey at any time. Informed consent was collected through a confirmation screen at the start of the questionnaire, and subjects were reimbursed for 30 minutes of their time. All complete responses were included in the analysis. We chose an average relatedness score of >3 in the related conditions as a cut-off point to exclude items. This point coincided with 2SD from the mean of the +Comp/-ID primed targets. Thirteen items were removed from the stimulus list as a result (all on the basis of the +Comp/-ID prime relation), leaving a total of 208. Once these items were removed, no further statistics were performed on the results. Visualisation of the mean scores after exclusion indicated a clear separation between the related and unrelated conditions (Fig. 3).

Scrambling. After the two sentence lists (+Comp/+ID and +Comp/-ID primes), scramble primes for each target were generated (-Comp/+ID and –Comp/-ID). Two non-compositional prime lists were created, where one contained the identity prime and one did not. To create these lists, the +Comp/+ID prime lists were randomly divided into two groups. Each group was converted into a

string array, where each prime was represented as a row vector. Individual strings in each array were first jumbled within columns. Gaps formed within rows. To repair these gaps, all strings were moved to the top of the array, and empty values were inserted between strings so that each row was of the same length as it was previously, matching the length of its associated compositional prime. Then all prime content items were jumbled within rows, and all strings were aligned to the left to remove any gaps. Identity tokens were swapped back to their original row and column, so that they matched the location in the original primes. Each resulting group was copied, and for each copy the identity token was replaced with the prime-matched token in the non-identity prime condition (e.g., verdrinking with stroming in Table 1). Some primes contained split verbs in the -ID condition but not in the +ID condition. Missing words in these non-identity scrambles were added to the prime at a location that matched the number of words between the prime-matched token and the target in the non-identity sentence. Each non-compositional prime was checked with snaut again (Mandera et al., 2017). In case of a violation (rank nearest neighbour <40), items were manually swapped between primes. The result was four non-compositional prime groups, each matching half of the compositional primes. These consisted of two -Comp/+ID prime groups, with two matching -Comp/-ID prime groups.

Next, all four prime conditions for each target were divided into presentation lists. Each list was created by combining non-compositional primes for one half of the targets with compositional primes for the other half. For both prime-types, half of the stimuli included an identity prime, and half did not,
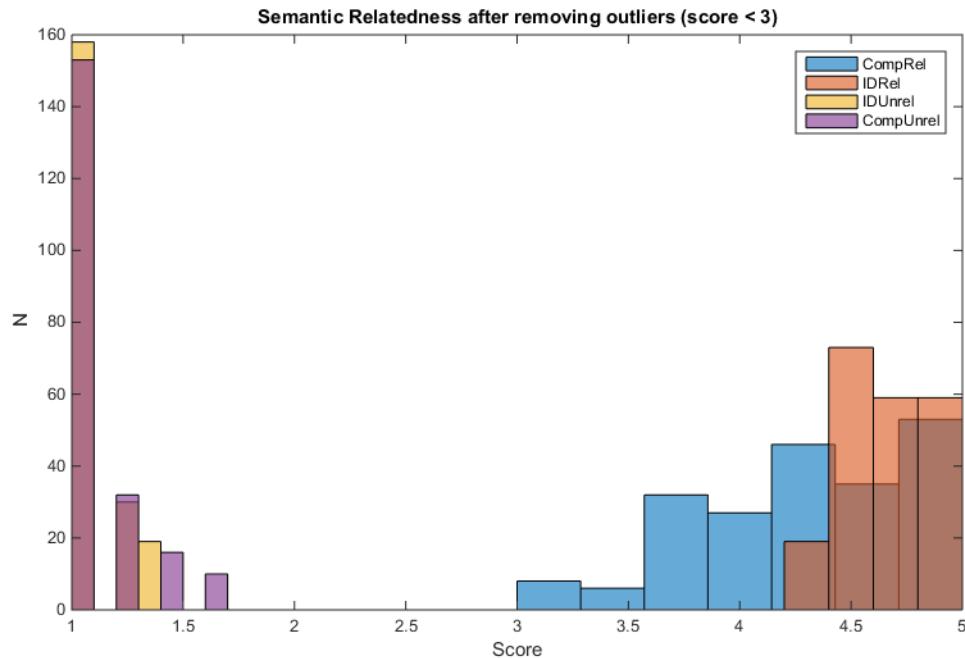
Figure 3. Histogram of the spread of mean semantic relatedness scores between sentences and targets. Purple and yellow indicate the condition in which an unrelated target was presented. Blue and orange show the Compositional/Non-identity and Compositional/+Identity relatedness scores respectively. Items in the relatedness condition with a score of < 3 were removed.

so that each subject saw all four conditions. This resulted in four unique, counter-balanced lists (see Supplementary information for the final lists used in the fMRI experiment). At this point, duplicate targets were still present in the stimulus list. Duplicate targets were removed to avoid between-trial identity priming. A final run-through of the stimuli revealed 15 items that still needed to be removed, leaving the final count at 193 prime-target pairs. The final number of Lai et al. (2015)-based stimuli was 102.

Behavioural experiment

Design. The behavioural experiment was still part of the pre-piloting, designed to see if a behavioural priming effect could be detected before attempting to find it in an fMRI experiment. The experiment consisted of a lexical decision task in order to provide participants with a task that did not explicitly draw attention to the different prime types. For this purpose, the main study design was extended to a 2*2*2 design, where a word-pseudoword factor was added. For each list, half the targets for all four conditions were replaced with a matched pseudoword. Pseudowords were constructed using Wuggy (Keuleers & Brysbaert, 2010). Verbal and written consonant-vowel structures as well as assumed stress patterns were maintained based on native speaker intuition. In addition, transitional

probabilities were maintained where possible . This resulted in a total of eight unique, counter-balanced lists. Items were initially ordered randomly, before being reordered manually to avoid identity priming of targets by earlier unrelated primes. That is to say, specific primes would only appear once a particular target had already been presented. Some of these cases survived for the behavioural experiment but were filtered out in subsequent experiments.

The experiment was performed using Presentation® software (Version 20.2, Neurobehavioral Systems, Inc., Berkeley, CA, www.neurobs.com). Primes were presented word-by-word in white text on a black background, with an on-time of 300ms and an off-time of 150ms. After each prime, a fixation cross was presented for 1000ms, before a 6000ms presentation of the target. Finally, a fixation cross was presented again for 1500ms before the next trial began. The task was divided into 13 blocks, between which short breaks were planned.

Data acquisition. Participants were recruited using the Max Planck Institute for Psycholinguistics participant database, and subjects received a reimbursement for 45 minutes of their time. The study was conducted according to the institutional guidelines of the local ethics committees. Written consent was collected before the start of the experiment. Participants were instructed to indicate

upon target presentation whether the target was an existing Dutch word or not with a button press, and participants were informed that there was a time limit involved. A total of 33 healthy, right-handed native Dutch-speakers, aged 18-40 without diagnosed reading or language disorders, were recruited for this experiment. One subject was excluded due to a pre-screening error. One subject showed an error-rate well below chance (85-100% per condition), and their responses were taken to be inverted. This subject was still included in the analysis after reversing their responses. Consequently, we collected four complete responses for each stimulus list.

Analysis. For the analysis, responses were filtered to only include correct responses to real-word trials, leaving only the original two factors (+Comp prime vs. -Comp prime * +ID prime vs. -ID prime). Data were analysed in R (R core team, 2018) using a linear mixed effects (LME) model from the LME4 package (Bates et al., 2015). Our model included the fixed interaction effect between the two factors, as well as random effects of subject and stimulus/trial number. Contrasts for both factors were set to -0.5 and 0.5. This was the most complete model we could use for the analysis. Adding random slopes to the model prevented convergence and were therefore not included.

FMRI experiment

Design. For the fMRI experiment, we utilised the original 2*2 design (+Comp prime vs -Comp prime * +ID prime vs -ID prime). The pseudo-word condition was dropped, which meant that the amount of trials for the remaining conditions doubled, and thus the statistical power of the experiment was increased. The experiment was performed using Presentation® software (Version 20.2, Neurobehavioral Systems, Inc., Berkeley, CA, www.neurobs.com). Primes were presented word-by-word in white text on a black background, with an on-time of 300ms and an off-time of 150ms. After each prime, a fixation cross was presented for an amount of time that extended the prime duration to the next multiple of 650ms plus 650 or 1300ms. Afterwards, the target was presented for 1300ms. Finally, a fixation cross was presented again for 1000-2000ms (in steps of 50ms) before the next trial began. The task was divided into seven runs, between which short breaks were planned. During each break, a light-blue screen was presented to reduce participant fatigue. Catch trials were added every 6-8 trials. During catch trials, a question mark

was presented, and participants had to indicate by means of button press whether the preceding prime was a sentence or a scramble. The motivation behind this task was to ensure participants read our primes as intended. Catch trial duration was intended to be 2600ms, but due to a programming error, catch trial duration was 3000ms for eleven of twelve subjects. This difference was accounted for in data analysis. Stimuli were projected with an EIKI LC-XL100 beamer at a resolution of 1024x768 (4:3 aspect ratio) and a refresh rate of 60Hz on a 451x338mm screen. Participants viewed stimuli at the back of the scanner through a mirror mounted on the head coil. Before the start of the experiment, subjects had an opportunity to adjust this mirror for full view of the projection. All four lists that resulted from the stimulus creation described above were included. Items were reordered again after target order concerns remained after the behavioural experiment.

Data acquisition. The experiment was conducted using a Siemens Prisma 3T MRI scanner at the Donders Institute for Brain, Cognition, and Behaviour. For each participant, we collected a 0.8mm isometric anatomical map using MP2RAGE (Marques et al., 2010) with a FoV of 256mm and a flip angle of 8°. A TR of 2400ms, TE of 2.22ms, and TI of 1000ms, as well as a GRAPPA acceleration factor of 2 meant a full anatomical image was acquired in just under seven minutes.

Functional data was not acquired with a laminar analysis in mind during the pilot. As such, only 1.3mm isometric fMRI data was collected. Our protocol entailed a 3D-EPI sequence based on Stirnberg et al. (2017) with a FoV of 210mm and a flip angle of 20°. The TR was 2600ms, TE1 32ms, multi-echo dTE 60ms, and TI 900ms. Acquisition was accelerated with CAIPIRINHA at a factor 6. SPAIR fat-saturation was enabled. FoV positioning was determined by means of a head-scout. The full scan-time was approximately 45 minutes, excluding breaks.

Heart rate and respiratory data were collected with a BrainProducts BrainAmp system and were added as regressors in the fMRI analysis. Heart rate was collected using a pulse sensor made MRI-tolerable, and respiratory data by means of a respiration belt. Physiological data were converted to workable files before they were checked and cleaned using HERA (Hermans, 2018), and finally converted to regressors using RETROIcorplus (Glover et al., 2000). The first through fifth order sine and cosine heart rate and respiratory regressors, as well as the first order sine and cosine interactions between heart rate and

respiration were modelled. Eye-tracking was on during scanning to monitor participant drowsiness, but eye-movements were not recorded.

Thirteen healthy, right-handed native Dutch-speakers, aged 18-30 without diagnosed reading or language disorders, were recruited for this experiment. One subject was excluded due to a pre-screening error. The study was conducted according to the institutional guidelines of the local ethics committees, and subjects were reimbursed for 90 minutes of their time. Written informed consent was collected before the start of each experiment. For each stimulus list, three responses were taken into the analysis.

Pre-processing. Pre-processing was performed using AFNI tools (Cox, 1996; all AFNI programmes discussed are prefixed with 3d). Pre-processing consisted of realignment of the functional volumes using 3dvolreg to a base of the fifth volume of the first run. Motion correction parameters were also collected from this realignment. The anatomical images were skullstriped using 3dSkullStrip. A mask was applied to the base volume of each subject using 3dautomask, and the anatomical image was aligned to this mask using the 3dallineate tool with a quintic interpolation.

For one participant, FoV was shifted for the first run compared to the other runs. In this participant, the base volume for motion correction was taken from the second run. For five participants, the FoV was shifted for the last two runs compared to the rest. An extra pre-processing step was added for these runs. Transformation matrices were calculated by catenating the transformation matrices of an initial large realignment shift to bring the volumes into the same space as the base (AFNI 3dallineate) and a subsequent small motion correction shift (AFNI 3dvolreg). This matrix was then applied to the original NIfTIs using 3dallineate to reduce the number of interpolation steps by one. For all participants except one, this resulted in good realignment. Runs six and seven were excluded from analysis for the remaining participant. Heart rate data for run three in one participant were unusable. The design matrix for this participant was adjusted by only including the respiratory regressors for run 3, and adding the regressor arrays of runs 1-2, 3, and 4-7 separately.

For group-level comparisons, our data were aligned to MNI152 space. The anatomical scan of each subject was normalised using 3dUnifize in AFNI, and brain extraction was performed with the mri_watershed program in freesurfer (Ségonne et al.,

2004; Fischl, 2012). For each subject, eleven brain-extracted images were generated where the height parameter of mri_watershed varied between ten and 20. For each subject, the best brain-extracted image was selected and manual improvements were applied as necessary. A non-linear transformation was then applied to these brain-extracted anatomicals to align them in MNI152 space using 3dQwarp. The align_epi_anat.py script in AFNI (Saad et al., 2009) calculated the affine transformation matrices for the automasked functional data to the brain-extracted anatomical, which were then catenated with the non-linear transformation parameters to bring the functional data to MNI152 space. The resulting images were then smoothed using an 8mm smoothing kernel in 3dmerge (Cox, 1996).

Analysis. Design matrices for the group-level analysis were modelled so that the first four columns matched conditions between subjects (column one: +Comp/+ID, column two: +Comp/-ID, etc.). These models were fitted within the GLM framework, and beta-maps and residual-maps were calculated for each subject. Fitting was performed with in-house MATLAB code and an open source fMRI analysis toolbox (https://github.com/TimVanMourik/OpenFmriAnalysis). 3dANOVA3 (Cox, 1996) was used to compare the resulting beta maps in a repeated measures ANOVA. To compare the +Comp/+ID condition to the -Comp/-ID activity patterns (the full priming effect vs. no priming) and the +Comp/-ID vs -Comp/+ID activity patterns (composition priming only vs. identity priming only), separate pairwise comparisons were performed using 3dttest++. To estimate the smoothness of the data, the spatial autocorrelations of the residual time series were calculated with 3dFWHMx and averaged over all runs and subjects. The resulting correlation parameter averages were entered into 3dClustSim to estimate the minimum cluster size required to exclude noise clusters with a likelihood of $\alpha < 0.05$. The uncorrected p-value threshold was set to $P < 0.001$.

Partial η^2 maps were calculated from the F-maps. These were calculated voxel-wise on the basis of voxel-specific F-values and the degrees of freedom (Cohen, 1973). Effect size-maps for pair-wise effects were also calculated using a version of Hedge's g (Hedge, 1981) that accounted for low sample sizes. Average beta-maps and standard deviation-maps for each condition across subjects were calculated using 3dMean (Cox, 1996). A voxel-wise calculation was performed over these maps based on voxel-wise betas and standard deviations per condition, and

sample size. All effect size maps were then masked with a mask of the MNI template dilated by one level in 3dmask_tool to generate the final effect size maps.

Results

Behavioural experiment

The results from the LME analysis indicate a robust difference in lexical decision reaction time as a function of identity prime type compared to non-primed targets, and a potential difference as a function of compositional prime types. The LME showed a robust effect for subjects recognising identity primed words (+ID) 37.21 milliseconds faster than non-identity primed words (-ID) (SE = 7.71, t = 4.83). Additionally, the LME showed a potential effect where compositionally primed words (+Comp) were recognised 20.10 milliseconds faster than non-compositionally primed words (-Comp) (SE = 7.71, t = 2.61). There was no evidence for an interaction of ID-effect * Composition-effect (β = 0.80, SE = 15.42, t = 0.05). T-statistics were not subjected to significance testing owing to the challenges in assessing the degrees of freedom in LME analysis (Luke, 2017). T-to-z conversion is a common method to derive p-values for LMEs but is anti-conservative (as are most alternative methods; see ibid.). Not reporting p-values prevents marginal effects from being interpreted as significant, but the

t-values still offer some insight into the robustness of the effects. See Figure 4 for a visualisation of the reaction time results (Allen et al., 2018; Wickham, 2009). See Table 2 for an overview of the main effects.

Subsequent pairwise comparisons were performed using the emmean package in R (Fig. 4; Lenth, 2019) in order to gain further insight into the potential additive effects of the compositional priming effect. The comparisons revealed a significantly faster reaction time of 37.6ms for identity-only primed (-Comp/+ID) targets compared to non-primed (-Comp/-ID) targets (SE = 10.9, P < 0.005), and a 36.8ms decrease in reaction times for identity and compositional primed (+Comp/+ID) targets versus compositional-only primed (+Comp/-ID) targets (SE = 10.9, P < 0.005). A significant advantage for identity-and-compositional primed (+Comp/+ID) words compared to non-primed (-Comp/-ID) words of 57.3ms was also found (SE = 10.9, P < 0.001). The direct comparison of these conditions did not inform us as to whether the additive effect of composition priming was significant, however, so this effect was not further explored behaviourally. Other pairwise comparisons did not show significant differences between conditions (P > 0.2). All p-values were adjusted using the Tukey method for comparing a family of four estimates (Tukey, 1949). For an overview of the pairwise comparison results, see Table 3.
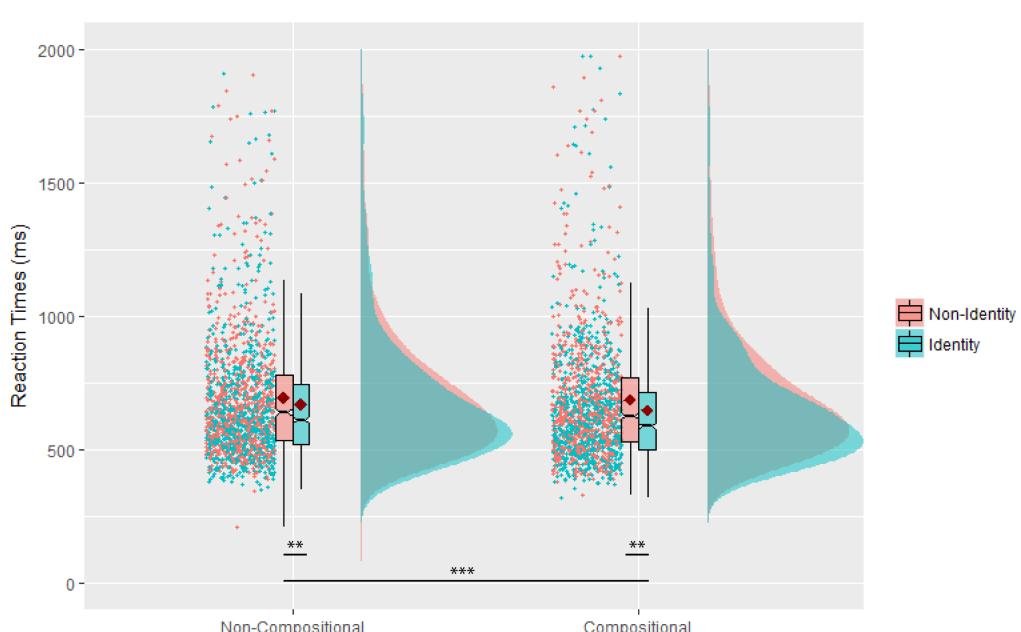

Figure 4. Reaction times are plotted against ID and Comp conditions. Significance bars indicate pairwise comparison effects (** p < 0.005; *** p < 0.001). Twenty data points were cut-off from the top for visibility purposes. Red dots indicate mean reaction times per condition.

fMRI experiment

A three-way ANOVA was used to assess the main effects of the priming conditions in this study. Composition priming and identity priming were modelled as fixed effects, and subject was modelled as a random effect. The ANOVA results were corrected for multiple comparisons with Puncorr = 0.001; α = 0.05. A cluster size > 703 voxels was determined to be significant. Functional data were mapped onto MNI space, and corresponding region labels are based on the Talairach-Tournoux Atlas (Talairach & Tournoux, 1988). Separate F-statistics were calculated to assess the main effects of Compositional and Identity priming. These were thresholded for multiple comparison correction to a value of F = 19.69 (1,11), Puncorr < 0.001. The left MTG, left superior and inferior frontal gyri (S/IFG), and left basal ganglia were sensitive to compositional priming manipulations (Fig. 5). Significant clusters sensitive to identity priming manipulations were found at the right middle frontal gyrus (MFG) and the right lentiform nucleus (see Fig. 6). There was no evidence for an interaction effect after correction. See Table 4 for an overview of the clusters in the main effects.

Group level t-statistics were calculated to assess significance of several additional contrasts. Due to the small sample size and associated Type II error rate inflation, it is challenging to observe small (ηp2 > 0.02) to medium-sized effects (ηp2 > 0.13; values based on Draper, 2011). When comparing the effect size maps to the ANOVA results, it becomes apparent that some large estimated effects (ηp2 > 0.26) were not found either under the current sample size (Fig. 7). Focus remained on the results that survived the initial correction at Puncorr < 0.001, but we endeavoured to improve our qualitative understanding of the results and potentially inform future work related to this pilot study by

also exploring the results at less stringent P-value thresholds (Puncorr < 0.005 and Puncorr < 0.01). Differences in conditions found in this exploration might become statistically significant with a better-powered sample. The minimum cluster size was kept at 703 voxels at these lower P-thresholds. For an overview of all clusters that survived any of the corrections, see Table 5.

Identity-and-compositional versus non-primed. Significant differences in activation were found in response to identity-and-compositional primed versus non-primed targets (+Comp/+ID vs -Comp/-ID) at Puncorr < 0.001 (Fig. 8, orange). Clusters over left S/MTG and bilateral basal ganglia were significantly more active in the identity-and-compositional primed condition over the non-primed condition. At a correction of Puncorr < 0.005 positive clusters over left supramarginal gyrus and IFG also survived correction (Fig. 9, orange).

Compositional-only versus non-primed. Significant differences in activation were found in response to compositional-only primed versus non-primed targets (+Comp/-ID vs -Comp/-ID) at Puncorr < 0.001 (Fig. 8, green). A cluster over the left S/MTG showed significantly more activity in the compositional-only primed condition over the non-primed condition. At a correction of Puncorr < 0.005 a negative cluster over left SFG also survived correction, and the positive cluster over the left MTG extended over to the left IFG (Fig. 9, green). At a correction of Puncorr < 0.01 a positive cluster over the left basal ganglia survived (Fig. 10, green).

Compositional-only versus identity-only. Significant differences in activation were found in response to compositional-only primed versus identity-only primed targets (+Comp/-ID vs -Comp/+ID) at Puncorr < 0.001 (Fig. 11).

Table 2. Overview of the main effects of the mixed model analysis.

Main effect	Beta	SE	t-value
Identity (+ID *vs.* -ID)	-37.21	7.71	-4.83
Composition (+Comp *vs.* –Comp)	-20.10	7.71	-2.61

Table 3. Overview of the pairwise comparisons between the behavioural conditions. Ordered by significance. Italics indicate a significant difference between conditions.

| Comparison	Beta	SE	z-ratio	p-value
+Comp/+ID vs. *-Comp/-ID*	*-57.3*	*10.9*	*-5.260*	*<0.0001*
-Comp/+ID vs. *-Comp/-ID*	*-37.6*	*10.9*	*-3.444*	*0.0032*
+Comp/+ID vs. *+Comp/-ID*	*-36.8*	*10.9*	*-3.382*	*0.0040*
+Comp/-ID *vs.* -Comp/-ID	-20.5	11.0	-1.871	0.2405
+Comp/+ID *vs.* -Comp/+ID	-19.7	10.9	-1.816	0.2658
+Comp/-ID *vs.* -Comp/+ID	17.1	10.9	1.568	0.3969
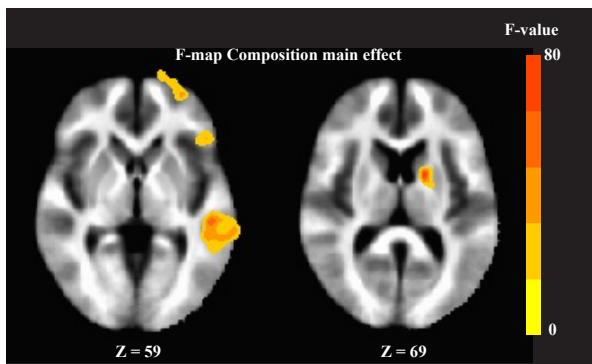
Figure 5. F-map of the Composition priming main effect. Clusters were corrected for multiple comparisons and were significant at Puncorr < 0.001, α = 0.05. Significant clusters after correction appeared over the left MTG, the left IFG, the left SFG (left), and the left basal ganglia (right).
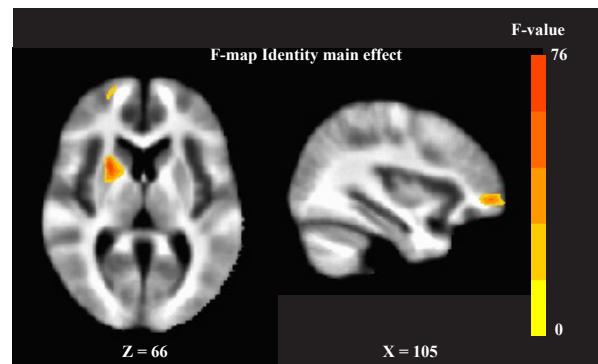

Figure 6. F-map of the Identity priming main effect. Clusters were corrected for multiple comparisons and were significant at Puncorr < 0.001, α = 0.05. Significant clusters after correction appeared over the right lentiform nucleus (left) and the right MFG (right).

Table 4. Overview of significant clusters of the main effects from the rmANOVA, and the locus region of each cluster based on the Talairach-Tournoux Atlas (Talairach & Tournoux, 1988). Hemisphere is given in the third column (L=left, R=right). Peak F-values are given in the fourth column. Peak xyz-coordinates are given in MNI space. No significant clusters were found for the interaction.

| Main effect	Locus	L/R	F-value	x	y	z
Composition	Left MTG	L	70.33	-50.3	-21.7	-10.9
	Left SFG	L	39.14	-17.9	65.9	4.7
	Left IFG	L	37.65	-46.7	28.7	-3.7
	Left basal ganglia	L	79.23	-16.7	3.5	10.7
Idenity	Right MFG	R	48.02	-34.9	-53.9	-2.5
	Right lentiform gyrus	R	76.44	-22.9	-8.3	7.1
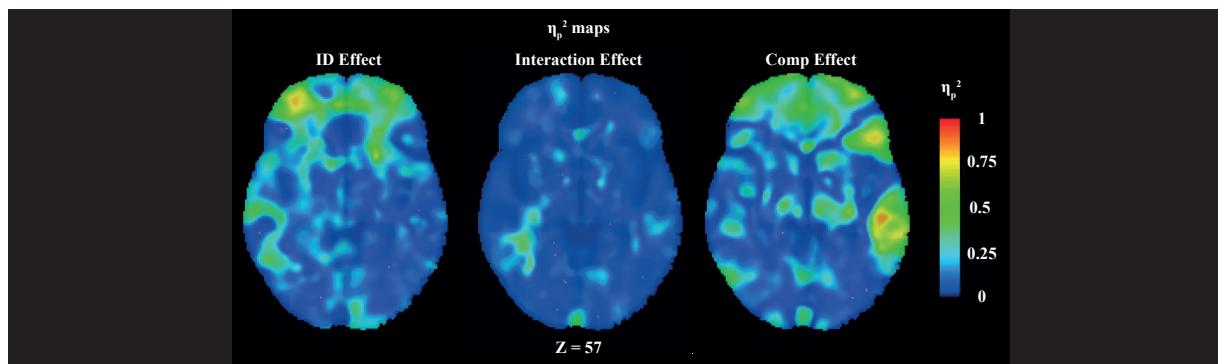
Figure 7. Partial η2 maps for the three main effects. Focussing on the left AG, we see that for the Identity priming main effect (left) and interaction effect (middle), we have a small effect size (~0.25). For the Composition priming effect, we see around a medium effect size (~0.5). N=12 is not enough to reliably detect a potential effect in this region.

Positive clusters over the left MTG and the IFG indicated significantly higher activity in response to compositional-only primed targets compared to identity-only primed targets. Negative clusters over the right AG and the bilateral S/MFG suggest higher activity in these regions in response to identity-only primed targets compared to compositional-only primed targets. At a correction of Puncorr < 0.005 negative clusters over the bilateral MFG dorsal to the significant clusters also survived correction (fig. 12, left). At a correction of Puncorr < 0.01 a negative cluster over the left inferior temporal gyrus (ITG) survived (fig. 12, right).

Identity-and-compositional versus identity-only. No clusters survived the initial correction when comparing activity in response to identity-and-compositional primed targets and identity-only primed targets (+Comp/+ID vs. -Comp/+ID). At Puncorr < 0.005, a positive cluster over the left

Table 5. Overview of surviving clusters in pairwise comparisons. Locus region of the cluster based on the Talairach-Tournoux Atlas (Talairach & Tournoux, 1988) is given in the second column and hemisphere in the third column (L=left, R=right). Peak t-values are given in the fourth column. Peak xyz-coordinates are given in MNI space. The uncorrected P-value thresholds to find the cluster are reported in the last column.

| Comparison	Locus	L/R	t-value	x	y	z	P_{uncorr}
+Comp/+ID vs. -Comp/-ID	S/MTG	L	6.49	-51.5	-26.5	-2.5	0.001
	basal ganglia	L/R	7.23				0.001
	SMG	L	5.48	-59.9	-50.5	21.5	0.005
	IFG	L	6.11	-35.9	4.7	29.9	0.005
+Comp/-ID vs. -Comp/-ID	MTG	L	9.58	-53.9	-28.9	-2.5	0.001
	SFG	L	-4.67	-28.7	68.3	3.5	0.005
	IFG	L	-				0.005
	basal ganglia	L	5.84	-15.5	3.5	10.7	0.01
+Comp/-ID vs. -Comp/+ID	MTG	L	8.84	-52.7	-30.1	-3.7	0.001
	IFG	L	6.62	-46.7	28.7	-1.3	0.001
	AG	R	-6.35	45.7	-68.5	34.7	0.001
	S/MFG (ventral)	L/R	-5.95				0.001
	MFG (dorsal)	L/R	-6.08				0.005
	ITG	R	-4.46	64.9	-20.5	-19.3	0.01
+Comp/+ID vs. -Comp/+ID	MTG	L	5.36	-65.9	-36.1	-0.1	0.005
	SFG	L	-5.15	-16.7	65.9	3.5	0.01
+Comp/+ID vs. +Comp/-ID	MFG	R	4.07	37.3	51.5	-4.9	0.01
	PCC		4.78	-3.5	-34.9	23.9	0.01
-Comp/+ID vs. -Comp/-ID	No effects						
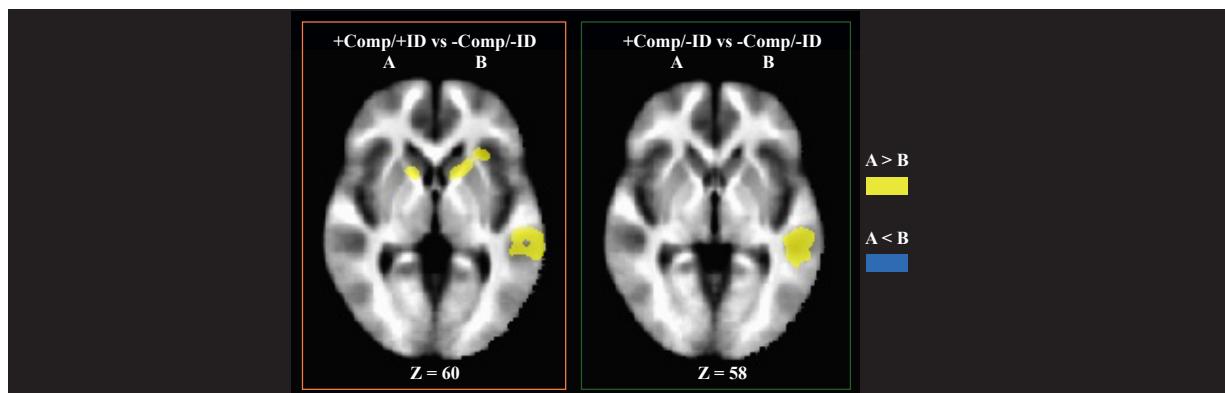
Figure 8. T-maps of the +Comp/+ID vs -Comp/-ID (left; orange) and +Comp/-ID vs -Comp/-ID (right; green) contrasts. Clusters were corrected for multiple comparisons and were significant at Puncorr < 0.001, α = 0.05. Significant clusters appear over the left MTG for both contrasts. In addition, there is a significant cluster over the bilateral basal ganglia in the +Comp/-ID vs -Comp/-ID contrast.
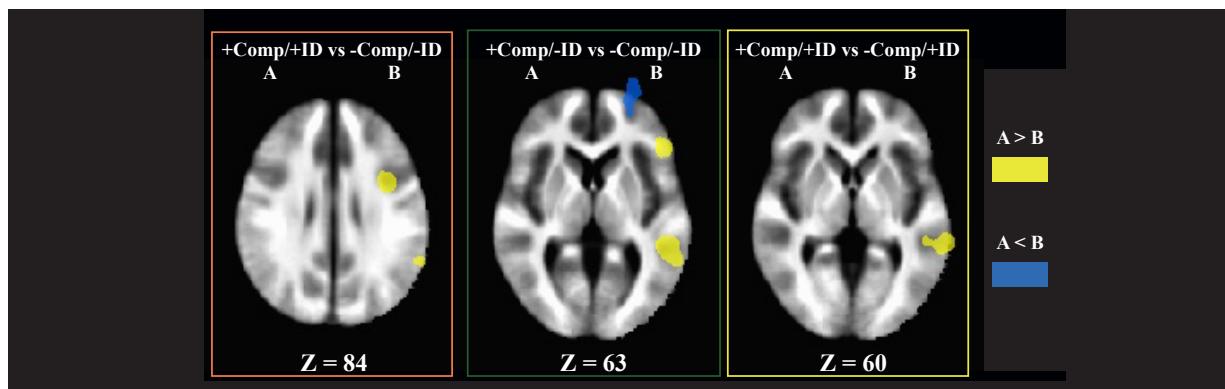

Figure 9. T-maps of the +Comp/+ID vs -Comp/-ID showing the clusters surviving correction at Puncorr < 0.005, α = 0.05 over the left IFG and the SMG (left; orange), +Comp/-ID vs -Comp/-ID showing clusters over the left SFG and the IFG (middle; green), and +Comp/+ID vs -Comp/+ID showing clusters over the left MTG (right; yellow) contrasts.
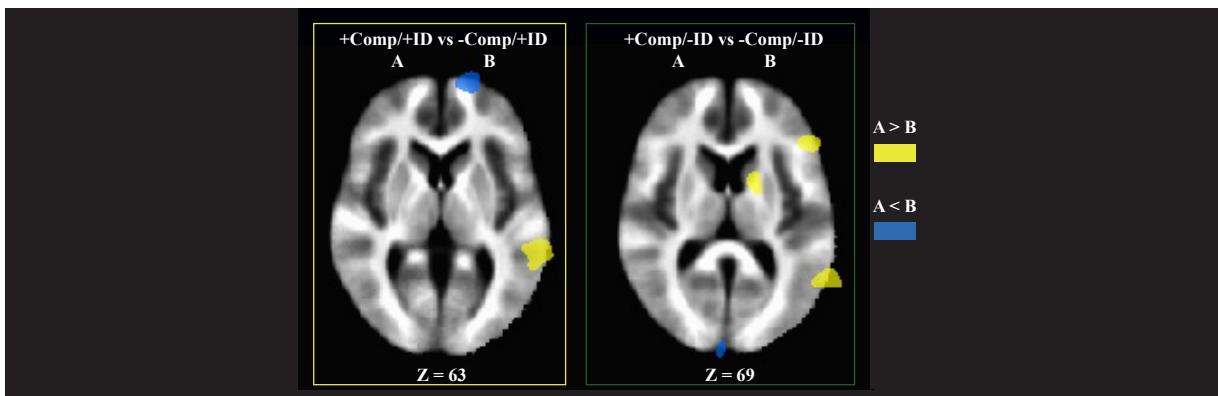
Figure 10. T-maps for +Comp/+ID vs -Comp/+ID (left; yellow) and Com/-ID vs -Comp/-ID contrasts (right; green) corrected at Puncorr < 0.01, α = 0.05. At the left, we see a surviving cluster over left SFG. At the +Comp/-ID vs -Comp/-ID contrast, we see a new cluster over the left basal ganglia.
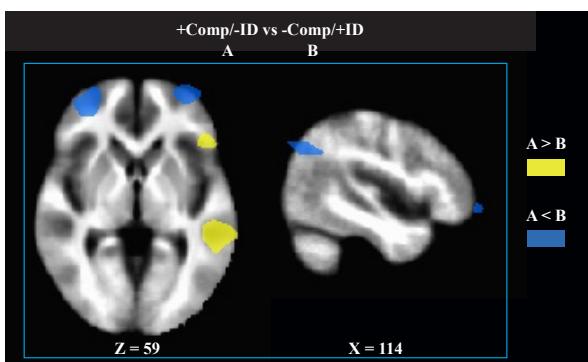

Figure 11. Transversal and sagittal view of the t-map for the +Comp/-ID vs -Comp/+ID contrast (blue). Clusters were corrected for multiple comparisons and were significant at Puncorr < 0.001, α = 0.05. Significant clusters appear over the left MTG, the left IFG, the bilateral STG (left), and the right AG (right).

Figure 12. T-maps of +Comp/-ID vs -Comp/+ID contrast corrected at Puncorr < 0.005, α = 0.05 (left) and Puncorr < 0.01, α = 0.05 (right). At Puncorr < 0.005 we see bilateral S/MFG activity that was not present at Puncorr < 0.001 at a more dorsal Z-slice. At Puncorr < 0.01, we see a cluster over right ITG. In both images, we can also see the right AG cluster from Puncorr < 0.001.

MTG survived correction (Fig. 9, yellow), and at Puncorr < 0.01, a negative cluster over the left SFG survived (Fig. 10, yellow).

Identity-and-compositional versus compositional-only. No clusters survived the initial correction when comparing activity in response to identity-and-compositional primed targets and compositional-only primed targets (+Comp/+ID vs. +Comp/-ID). Only at Puncorr < 0.01 two positive clusters survived over the posterior cingulate cortex and the right MFG (Fig. 13).

Discussion

Behavioural experiment

In the lexical decision task we observed faster reaction times to items that were identity primed compared to items that were not identity primed, both as a main effect (Table 2) and as simple effects (Table 3). These results indicated that the paradigm was successful in eliciting a priming response consistent with previous literature (for an overview, see Segaert et al., 2013). A separate effect was found when comparing items primed by sentences to those primed by scrambled sentences, though this effect was less pronounced (see Table 2). The effect could indicate an effect of compositional priming, where the compositional meaning of the sentence semantically primed the target. As such, it shows promise of the expected BOLD response. However, this reaction time effect could also arise from processing difficulties in interpreting scrambles, rather than from the compositional priming effect. Reaction times might be shorter after any sentence than after any scramble. Scrambles may be harder to process, and any subsequent processes
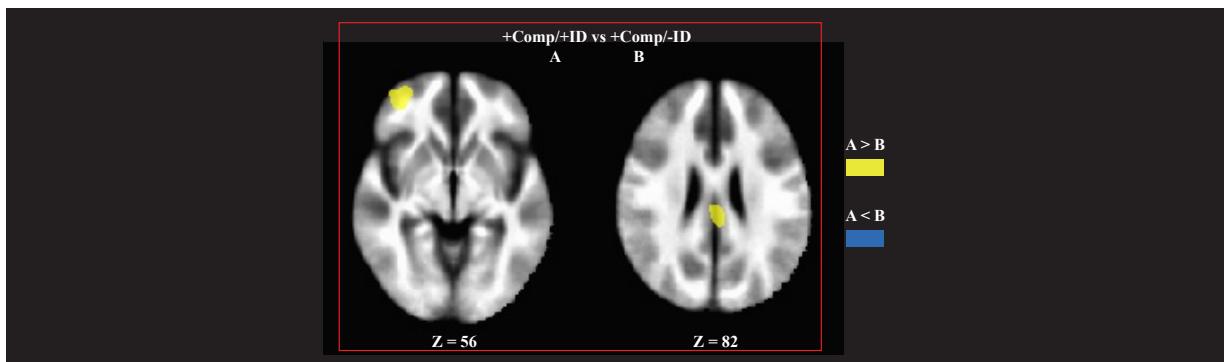
Figure 13. T-maps for +Comp/+ID vs +Comp/-ID corrected at Puncorr < 0.01, α = 0.05. Surviving clusters appear over right MFG and PCC.

could require more time. Further testing is required to comprehend the effects arising from processing hurdles of random word lists to subsequent word processing and compositional priming fully. We believe that, here, we have been able to lay the groundworks for a paradigm that can be utilised to compare the various priming effects that may occur between sentence primes and single-word targets, and could be expanded to study differences between sentence and word-list processing in general.

fMRI experiment

First, it should be noted that any conclusions drawn from this pilot experiment are preliminary and require support from the full experiment. Additionally, the results section discussed findings for corrections at both Puncorr < 0.001 and less stringent P-thresholds. For the discussion, we will primarily focus on the regions that showed significant differences in the pairwise comparisons with a correction at α = 0.05 and Puncorr < 0.001.

When looking at the compositional priming effect, we see that the left MTG strongly responds to the presentation of a target after a compositional prime compared to a non-compositional prime (Fig. 5, 7right, 8 & 11). This is accordance with our prediction that the MTG would be sensitive to compositional priming given its prominent role in the language network (as outlined in the introduction). Responses in the MTG were not sensitive to identity priming (see Table 5), suggesting that this result is more likely attributable to the compositional prime. What is interesting to note is the presence of an enhancement effect in response to compositionally primed targets, rather than a suppression effect. This effect may be similar to the repetition enhancement found in masked prime studies, as explained by Segaert et al. (2013). Schnyer et al. (2002) proposed that a masked prime may lead to enhanced fMRI

signal if "activation from the masked prime continues to spread until the target is identified" (Segaert et al., 2013). Segaert et al. (2013) then propose that the response to the prime and the response to the target are additive, resulting in enhancement. A similar principle may be a at play here, where the spread of activation from the priming sentence increases the fMRI signal related to the target, in addition to the target activation itself, indicating successful compositional priming.

Somewhat surprisingly, we found little evidence for the identity priming effect in the left temporal cortex. There was no main effect in left MTG of +ID vs -ID (Table 4), or a pairwise effect (row 5 & 6 in Table 5). Since repetition priming is quite a well-established effect (e.g., Bergerbest et al., 2006; Schnyer et al., 2002; Elger et al., 1997), we expected that this response would be stronger in the identity comparisons than in the composition comparisons. Especially since the effect is absent from the -Comp/+ID with -Comp/-ID comparison, where no other priming effects should attenuate the identity priming effect, it seems that the current setup is unable to catch both the perceptual and semantic aspects of the identity priming effect. It could be that the task redirected the attention of the subjects away from the identity tokens, and as such, there were no differential responses to the identity tokens (for the effects of attention on visual priming, see Vuilleumier et al., 2006; Thoma & Henson, 2011). For example, Schnyer et al. (2002) found an identity priming effect at the word-level using a lexical decision task, which required lexical access to perform, whereas our task did not necessarily. If subjects had been asked during catch trials whether the target was present in the preceding string, perhaps a stronger identity priming effect would have been found. While a different task for the fMRI experiment should be considered for the main study, it should be pointed out that a task drawing attention

to target words might consequently reduce attention to the unity within the primes, and as such potentially eliminate the compositional priming effect. A task like the lexical decision task would also reduce the power of the experiment by reducing the number of trials of interest. The challenge lies in finding a task that requires attention to both elements of each trials without disqualifying trials from subsequent analysis.

The left IFG also appeared to respond more strongly to the compositionally primed targets than the non-compositionally primed targets (Fig. 5; Table 4). If theleft IFG activity is related to the composition of linguistic input, and if the output of the IFG is subsequently relayed back to the left MTG, then it stands to reason that the IFG activates when a semantically related target to the complex prime is presented. No effect was expected for the identity vs non-identity comparisons in the IFG.

Somewhat surprising was the reduced response in the left S/MFG to compositionally primed targets vs non-compositionally primed targets (Table 5, rows 2-4). The S/MFG have been associated with working memory demands and cognitive control (Du Boisgueheneuc et al., 2006; Li et al., 2013; Emch et al., 2019). This could indicate that a reduced S/MFG activity in response to a compositionally primed target may be due to reduced demands to retrieve an item from working memory, as seen in all but one +Comp vs -Comp contrast (Table 5). Given that our task did not require overt attention to the presence of identity primes, it seems that this is an automatic response. In other words, if the prime is incoherent, the brain automatically allocates more resources to remembering its contents and retrieving it, whether task-relevant or not. Here, this effect is found to be bilateral and mostly seen in the compositional vs. non-compositional comparison. However, we also see this effect in the right MFG when comparing +Comp/+ID to +Comp/-ID. Here it appears that the right MFG responds stronger when an item may be present in working memory than when it is not, and may as such reflect a recollection process related to verbal working memory (Emch et al., 2019). The right MFG has also been associated with redirection of attention (Japee et al., 2015), which might indicate that the presentation of the target causes subjects to reprioritise the focus of the prime.

An unexpected result from this pilot was that when comparing the response to an exclusively compositionally primed target with the response to an exclusively identity primed target (+Comp/-ID vs. -Comp/+ID), we found evidence for the right AG and the ITG response (Fig. 11). While AG activity

was expected in this condition, we expected it in the left hemisphere, as the left hemisphere is usually more strongly associated with linguistic processing. What is also surprising is that activity appears to be stronger for the identity primed condition than the compositionally primed condition. Since compositional priming has led to enhancement of the BOLD response in the other comparisons, it might be expected here. It is perfectly possible that different brain regions respond to the same prime differently, or it might be that this is an effect driven by the identity prime. Since this response does not show up for any of the other comparisons, it is difficult to tell to which condition the right AG is sensitive. It may also be, given the 'pure' priming conditions in our experiment, that the right AG responds to a condition only if the other priming stream is absent, and that once multiple paths of relating input become available, its involvement in the process is affected.

Right ITG activity may be a result of the predictability of the target word. Bonhage et al. (2015) found a correlation between the right ITG activity and word predictions in semantically rich contexts. Contrary to our results however, Bonhage et al. (2015) found that the right ITG was more active for a sentence condition compared to a jabberwocky condition where only word category could be predicted. It is unclear how these results may be reconciled with our findings. Right ITG activity may also be related to a type of repetition. Schweinberger et al. (2002) found that right ITG responses related to familiar face repetitions in an ERP study. Similarly, Li et al. (1993) found an inferior temporal response to familiarity of stimuli in rhesus monkeys. Perhaps this effect could be extended to familiar word repetitions, but this effect does not seem to be reported in linguistic contexts. It would be interesting to see if in a follow-up this effect might pop up in different conditions.

Basal ganglia activity was found bilaterally in the +Comp/+ID vs. -Comp/-ID and left lateralised in +Comp/-ID vs. -Comp/-ID. Basal ganglia involvement was expected, but no strong predictions were made because basal ganglia have been found to be involved in a number of linguistic (Copland et al., 2000; Bacon Moore et al., 2013) and non-linguistic tasks (Rao et al., 1997; Haaland et al., 1997). The effect that was found may be a verbal working memory process (Bacon Moore et al., 2013) involved in complex input processing. Interestingly, the main effects show that left and right basal ganglia are responsive to only compositional and identity priming respectively (Fig. 5 & 6). These

main effects are not reflected in the pairwise comparisons, however. This makes it difficult to draw any conclusions about basal ganglia activity, as it is unclear which condition drives the basal ganglia effects in our analysis.

An alternative interpretation for our results arises not from a compositional priming effect, but rather simply due to participants processing sentence level information while reading the primes. We see that the language network responds in a fashion similar to on-line sentence processing (e.g., as modelled by Hagoort, 2016 or Friederici, 2011). The question is whether the response to the compositionally primed targets should be any different from the activity in on-line sentence processing. We presented participants with a one-word compositional summary of a sentence they previously read. Readers may relate this word back to the previous context, similar to how new words in an input are inserted into the previous context in Hagoort (2016) and Friederici (2011), involving comparable neural processes. For future analyses, it may be worthwhile to model prime presentation as a nuisance regressor to eliminate potential effects from the prime processing. We intend include this regressor in the study based on this pilot, but we do not expect this to have a detrimental effect on the results.

Through this pilot study, we were able to establish brain regions that likely play a role in semantic composition. The regions responsive to the manipulations in this study provide a starting point for a more thorough study on semantic composition processes. The left MTG and the IFG were expected candidates, but we also found bilateral S/MFG responses and potentially right AG and ITG activity related to processing primed targets. In the follow-up laminar study, these regions will hopefully show depth-dependant interactions with one another that would further increase our understanding of sentence processing. It remains to be discovered at what cortical depth the haemodynamic responses originate that drive the IFG and the MTG signal in this study, but the prediction that left IFG activity shows a correlation to deep-bin MTG activity still holds. If this correlation is found in the follow-up laminar fMRI study, it would be indicative of the predicted semantic composition process and top-down priming effect (as proposed by Hagoort, 2016), and would suggest that this stream of information is important for the conceptualisation of composed semantics. Any effect of right AG and bilateral S/MFG correlated to the compositional and especially identity priming effects would support a model for a bilateral language network that involves the right hemisphere at the word-level.

Conclusion

We have conducted a pilot study to see if it is possible to examine the processes involved in compositional meaning extraction of sentences. We used a priming paradigm in order to do so, where the target was either compositionally primed, identity primed, both, or neither. We found a strong effect of identity priming in a behavioural lexical decision task, as well as a smaller potential effect of compositional priming. We generated effect-size maps for each factor and pairwise contrast in our fMRI analysis, which can be utilised to estimate the sample needed for ROI selection in a follow-up laminar fMRI experiment. We found evidence for the existence of a neural compositional priming effect in left MTG, IFG and S/MFG. We also found an effect that could suggest a right-lateralised identity priming effect in S/MFG, AG and ITG. These findings will be taken as a starting point into a follow-up laminar fMRI experiment focussed on studying directed connectivity between these regions.

References

Aguirre, G. K., Zarahn, E., & D'Esposito, M. (1998). The Variability of Human, BOLD Hemodynamic Responses. *NeuroImage, 8*(4), 360–369.

Allen, M., Poggiali, D., Whitaker, K., Marshall, T. R., & Kievit, R. (2018). Raincloud plots: a multi-platform tool for robust data visualization. *PeerJ Preprints, 6*(e27137v1).

Bacon Moore, A., Li, Z., Tyner, C. E., Hu, X., & Crosson, B. (2013). Bilateral basal ganglia activity in verbal working memory. *Brain and Language, 125*(3), 316–323.

Baggio, G., & Hagoort, P. (2011). The balance between memory and unification in semantics: A dynamic account of the N400. *Language and Cognitive Processes, 26*(9), 1338–1367.

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software, 67*(1), 1–48.

Bergerbest, D., Ghahremani, D. G., & Gabrieli, J. D. E. (2006). Neural Correlates of Auditory Repetition Priming: Reduced fMRI Activation in the Auditory Cortex. *Journal of Cognitive Neuroscience, 16*(6), 966–977.

Binder, J. R., Desai, R. H., Graves, W. W., & Conant, L. L. (2009). Where Is the Semantic System? A Critical Review and Meta-Analysis of 120 Functional Neuroimaging Studies. *Cerebral Cortex, 19*(12), 2767–2796.

Bonahge, C. E., Mueller, J. L., Friederici, A. D., & Fiebach, C. J. (2015). Combined eye tracking and fMRI reveals

neural basis of linguistic predictions during sentence comprehension. *Cortex, 68*, 33–47.

Bradley, M. M., & Lang, P. J. (1994). Measuring emotion: The self-assessment manikin and the semantic differential. *Journal of Behavior Therapy and Experimental Psychiatry, 25*(1), 49–59.

Brown, S. P., & Hestrin, S. (2009). Intracortical circuits of pyramidal neurons reflect their long-range axonal targets. *Nature, 457*, 1133–1136.

Buckner, R. L., Koutstaal, W., Schacter, D. L., & Rosen, B. R. (2000). Functional MRI evidence for a role of frontal and inferior temporal cortex in amodal components of priming. *Brain, 123*(3), 620–640.

Cohen, J. (1973). Eta-Squared and Partial Eta-Squared in Fixed Factor ANOVA Designs. *Educational and Psychological Measurement, 33*, 107–112.

Copland, D. A., Chenery, H. J., & Murdoch, B. E. (2000). Processing lexical ambiguities in word triplets: evidence of lexical-semantic deficits following dominant nonthalamic subcortical lesions. *Neuropsychology, 14*(3), 379–390.

Cox, R. W. (1996). AFNI: Software for Analysis and Visualization of Functional Magnetic Resonance Neuroimages. *Computers and Biomedical Research, 29*(3), 162–173.

Dale, A. M. (1999). Optimal experimental design for event related fMRI. *Human Brain Mapping, 8*(2–3), 109–114.

Dale, A. M., Greve, D. N., & Burock, M. A. (1999). Optimal Stimulus Sequences for Event-Related fMRI. Presented at the 5th International Conference on Functional Mapping of the Human Brain., Duesseldorf, Germany.

Davey, J., Thompson, H. E., Hallam, G., Karapanagiotidis, T., Murphy, C., De Caso, I., … Jefferies, E. (2016). Exploring the role of the posterior middle temporal gyrus in semantic cognition: Integration of anterior temporal lobe with executive processes. *NeuroImage, 137*, 165–177.

de C. Hamilton, A. F., & Grafton, S. T. (2008). Repetition suppression for performed hand gestures revealed by fMRI. *Human Brain Mapping, 30*(9), 2898–2906.

Dehaene, S., & Cohen, L. (2011). The unique role of the visual word form area in reading. *Trends in Cognitive Sciences, 15*(6), 254–262.

den Heyer, K., Goring, A., & Dannenbring, G. L. (1985). Semantic priming and word repetition: The two effects are additive. *Journal of Memory and Language, 24*(6), 699–716.

Dimitropoulou, M., Duñabeitia, J. A., Avilés, A., Corral, J., & Carreiras, M. (2010). Subtitle-based word frequencies as the best estimate of reading behavior: the case of greek. *Frontiers in Psychology, 1*, 218.

Draper, S. (2011). Effect size [WWW document]. Retrieved from http://www.psy.gla.ac.uk/~steve/best/effect.html

du Boisgueheneuc, F., Levy, R., Volle, E., Seassau, M., Duffau, H., Kinkingnehun, S., … Dubois, B. (2006). Functions of the left superior frontal gyrus in humans: a lesion study. *Brain, 129*(12), 3315–3328.

Elger, C. E., Grunwald, T., Lehnertz, K., Kutas, M., Helmstaedter, C., Brockhaus, A., … Heinze, H. J. (1997). Human temporal lobe potentials in verbal learning and memory processes. *Neuropsychologia, 35*(5), 657–667.

Emch, M., von Bastian, C. C., & Koch, K. (2019). Neural Correlates of Verbal Working Memory: An fMRI Meta-Analysis. *Frontiers in Human Neuroscience, 13*, 180.

Ettinger, A., Feldman, N., Resnik, P., & Phillips, C. (2016). Modeling N400 amplitude using vector space models of word representation. *CogSci.*

Ferreira, F., & Lowder, M. W. (2016). Prediction, Information Structure, and Good-Enough Language Processing. *Psychology of Learning and Motivation, 65*, 217–247.

Fischl, B. (2012). FreeSurfer. *NeuroImage, 62*(2), 774–781.

Frazier, L., & Clifton, C. J. (1996). *Construal.* Cambridge, MA, USA: The MIT Press.

Frazier, L., & Clifton, C. J. (1997). Construal: Overview, Motivation, and Some New Evidence. *Journal of Psycholinguistic Research, 26*(3), 277–295.

Friederici, A. D. (2011). The Brain Basis of Language Processing: from Structure to Function. *Physiological Review, 91*, 1357–1392.

Geuter, S., Qi, G., Welsh, R. C., Wager, T. D., & Lindquist, M. A. (2018). Effect Size and Power in fMRI Group Analysis. *BioRxiv.*

Glover, G. H., Li, T.-Q., & Ress, D. (2000). Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magnetic Resonance in Medicine, 44*(1), 162–167.

Haaland, K. Y., Harrington, D. L., O'Brien, S., & Hermanowicz, N. (1997). Cognitive–Motor Learning in Parkinson's Disease. *Neuropsychology, 11*(2), 180–186.

Hagoort, P. (2005). On Broca, brain, and binding: a new framework. *Trends in Cognitive Sciences, 9*(9), 416–423.

Hagoort, P. (2013). MUC (Memory, Unification, Control) and beyond. *Frontiers in Psychology, 4*(416).

Hagoort, P. (2016). MUC (Memory, Unification, Control): A Model on the Neurobiology of Language Beyond Single Word Processing. In G. Hickok & S. L. Small (Eds.), *Neurobiology of Language* (pp. 339–347). Academic Press.

Hagoort, P., Hald, L., Bastiaansen, M., & Petersson, K. M. (2004). Integration of Word Meaning and World Knowledge in Language Comprehension. *Science, 304*(5669), 438–441.

Harris, K. D., & Mrsic-Flogel, T. D. (2013). Cortical connectivity and sensory coding. *Nature, 503*, 51–58.

Hedge, L. V. (1981). Distribution Theory for Glass's Estimator of Effect Size and Related Estimators. *Statistics, 6*(2), 107–128.

Hermans, E., & van de Laar, T. (2018). HERA: Heart Rate Analysis (Version 0.06 beta).

Humphries, C., Binder, J. R., Medler, D. A., & Liebenthal, E. (2006). Syntactic and Semantic Modulation of Neural Activity during Auditory Sentence Comprehension. *Journal of Cognitive Neuroscience, 18*(4), 665–679.

Humphries, C., Binder, J. R., Medler, D. A., & Liebenthal, E. (2007). Time course of semantic processes during sentence comprehension: An fMRI study. *NeuroImage, 36*(3), 924–932.

Hutchison, K. S. (2003). Is semantic priming due to association strength or feature overlap? A microanalytic review. *Psychonomic Bulletin & Review, 10*(4), 785–813.

Japee, S., Holiday, K., Satyshur, M. D., Mukai, I., & Ungerleider, L. G. (2015). A role of right middle frontal gyrus in reorienting of attention: a case study. *Frontiers in Systems Neuroscience, 9*, 23.

Kastrup, A., Krüger, G., Glover, G. H., Neumann-Haefelin, T., & Moseley, M. E. (1999). Regional Variability of Cerebral Blood Oxygenation Response to Hypercapnia. *NeuroImage, 10*(6), 675–681.

Keefe, D. E., & Neely, J. H. (1990). Semantic priming in the pronunciation task: The role of prospective prime-generated expectancies. *Memory & Cognition, 18*(3), 289–298.

Kenett, Y. N., Levi, E., Anaki, D., & Faust, M. (2017). The semantic distance task: Quantifying semantic distance with semantic network path length. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 43*(9), 1470–1489.

Keuleers, E., & Brysbaert, M. (2010). Wuggy: A multilingual pseudoword generator. *Behavior Research Methods, 42*(3), 627–633.

Keuleers, E., Brysbaert, M., & New, B. (2010). SUBTLEX-NL: A new measure for Dutch word frequency based on film subtitles. *Behavior Research Methods, 42*(3), 643–650.

Kok, P., Bains, L. J., van Mourik, T., Norris, D. G., & de Lange, F. P. (2016). Selective Activation of the Deep Layers of the Human Primary Visual Cortex by Top-Down Feedback. *Current Biology, 26*(3), 371–376.

Koopmans, P. J., Barth, M., & Norris, D. G. (2010). Layer specific BOLD activation in human V1. *Human Brain Mapping, 31*(9), 1297–1304.

Lai, V. T., Willems, R. M., & Hagoort, P. (2015). Feel between the Lines: Implied Emotion in Sentence Comprehension. *Journal of Cognitive Neuroscience, 27*(8), 1528–1541.

Lapesa, G., & Evert, S. (2013). Evaluating Neighbor Rank and Distance Measures as Predictors of Semantic Priming. *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*, 66–74.

Lapesa, G., & Evert, S. (2014). A Large Scale Evaluation of Distributional Semantic Models: Parameters, Interactions and Model Selection. *Transactions of the Association for Computational Linguistics, 2*, 531–545.

Lefort, S., Tomm, C., Floyd Sarria, J.-C., & Petersen, C. C. H. (2009). The Excitatory Neuronal Network of the C2 Barrel Column in Mouse Primary Somatosensory Cortex. *Neuron, 61*(2), 301–316.

Lenci, A. (2018). Distributional Models of Word Meaning. *Annual Review of Linguistics, 4*, 151–171.

Lenth, R. (2019). emmeans: Estimated Marginal Means, aka Least-Squares Means. Retrieved from https://CRAN.R-project.org/package=emmeans

Li, L., Miller, E. K., & Desimone, R. (1993). The representation of stimulus familiarity in anterior inferior temporal cortex. *Journal of Neurophysiology, 69*(6), 1918–1929.

Li, W., Qin, W., Liu, H., Fan, L., Wang, J., Jiang, T., & Yu, C. (2013). Subregions of the human superior frontal gyrus and their connections. *NeuroImage, 78*, 46–58.

Limesurvey GmbH. (n.d.). LimeSurvey: An Open Source survey tool (Version 3.19.2+191017). Hamburg, Germany: Limesurvey GmbH. Retrieved from http://www.limesurvey.org

Luke, S. G. (2017). Evaluating significance in linear mixed-effects models in R. *Behavior Research Methods, 49*(4), 1494–1502.

Lyu, B., Choi, H. S., Marslen-Wilson, W. D., Clarke, A., Randall, B., & Tyler, L. K. (2019). Neural dynamics of semantic composition. *PNAS, 116*(42), 21318–21327.

MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review, 101*(4), 676–703.

Mandera, P., Keulers, E., & Brysbaert, M. (2017). Explaining human performance in psycholinguistic tasks with models of semantic similarity based on prediction and counting: A review and empirical validation. *Journal of Memory and Language, 92*, 57–78.

Marques, J. P., Kober, T., Krueger, G., van der Zwaag, W., van de Moortele, P. F., & Greutter, R. (2010). MP2RAGE, a self bias-field corrected sequence for improved segmentation and T1-mapping at high field. *NeuroImage, 49*(2), 1271–1281.

McGettigan, C., & Scott, S. K. (2012). Cortical asymmetries in speech perception: what's wrong, what's right and what's left? *Trends in Cognitive Sciences, 16*(5), 269–276.

New, B., Brysbaert, M., Veronis, J., & Pallier, C. (2007). The use of film subtitles to estimate word frequencies. *Applied Psycholinguistics, 28*(4), 661–677.

Olman, C. A., Harel, N., Feinberg, D. A., He, S., Zhang, P., Ugurbil, K., & Yacoub, E. (2012). Layer-Specific fMRI Reflects Different Neuronal Computations at Different Depths in Human V1. *PLOS One*.

Oosterdijk, N., Reynaert, M., Hoste, V., & Schuurman, I. (2012). The Construction of a 500-Million-Word Reference Corpus of Contemporary Written Dutch. In P. Spyns & J. Odijk (Eds.), *Essential Speech and Language Technology for Dutch. Theory and Applications of Natural Language Processing*. Berlin, Heidelberg: Springer.

Pickering, M. J., & Branigan, H. P. (1999). Syntactic priming in language production. *Trends in Cognitive Sciences, 3*(4), 136–141.

R Core Team. (2018). R: A language and environment for statistical computing (Version R version 3.5.2 (2018-12-20), Eggshell Igloo) [X86_64, mingw32]. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from https://www.R-project.org/

Rao, S. M., Harrington, D. L., Haaland, K. Y., Bobholz,

J. A., Cox, R. W., & Binder, J. R. (1997). Distributed Neural Systems Underlying the Timing of Movements. *The Journal of Neuroscience, 17*(14), 5528–5535.

Raposo, A., Moss, H. E., Stamatakis, E. A., & Tyler, L. K. (2006). Repetition suppression and semantic enhancement: An investigation of the neural correlates of priming. *Neuropsychologia, 44*(12), 2284–2295.

Rayner, K. (2009). Eye movements and attention in reading, scene perception, and visual search. *The Quarterly Journal of Experimental Psychology, 62*(8), 1457–1506.

Rissman, J., Eliassen, J. C., & Blumstein, S. E. (2003). An Event-Related fMRI Investigation of Implicit Semantic Priming. *Journal of Cognitive Neuroscience, 15*(8), 1160–1175.

Roskies, A. L., Fiez, J. A., Balota, D. A., Raichle, M. E., & Peterson, S. E. (2001). Task-Dependent Modulation of Regions in the Left Inferior Frontal Cortex during Semantic Processing. *Journal of Cognitive Neuroscience, 13*(6), 829–843.

Rossell, S. L., Price, C. J., & Nobre, A. C. (2003). The anatomy and time course of semantic priming investigated by fMRI and ERPs. *Neuropsychologia, 41*(5), 550–564.

Rugg, M. D. (1985). The Effects of Semantic Priming and Word Repetition on Event-Related Potentials. *Psychophysiology, 22*(6), 642–647.

Saad, Z. S., Glen, D. R., Chen, G., Beauchamp, M. S., Desai, R., & Cox, R. W. (2009). A new method for improving functional-to-structural alignment using local Pearson correlation. *NeuroImage, 44*, 838–848.

Schnyer, D. M., Ryan, L., Trouard, T., & Forster, K. (2002). Masked word repetition results in increased fMRI signal: a framework for understanding signal changes in priming. *Neuroreport, 13*(3), 281–284.

Schweinberger, S. R., Pickering, E. C., Jentzsch, I., Burton, A. M., & Kaufmann, J. M. (2002). Event-related brain potential evidence for a response of inferior temporal cortex to familiar face repetitions. *Cognitive Brain Research, 14*(3), 398–409.

Segaert, K., Weber, K., de Lange, F. P., Petersson, K. M., & Hagoort, P. (2013). The suppression of repetition enhancement: A review of fMRI studies. *Neuropsychologia, 51*(1), 59–66.

Ségonne, F., Dale, A. M., Busa, E., Glessner, M., Salat, D., Hahn, H. K., & Fischl, B. (2004). A hybrid approach to the skull stripping problem in MRI. *NeuroImage, 22*(3), 1060–1075.

Sharoh, D., van Mourik, T., Bains, L. J., Segeart, K., Weber, K., Hagoort, P., & Norris, D. G. (2019). Laminar Specific fMRI Reveals Directed Interactions in Distributed Networks During Language Processing. *PNAS, 116*(42), 21185–21190.

Snijders, T. M., Petersson, K. M., & Hagoort, P. (2010). Effective connectivity of cortical and subcortical regions during unification of sentence structure. *NeuroImage, 52*(4), 1633–1644.

Snijders, T. M., Vosse, T., Kempen, G., van Berkum, J. J. A., Petersson, K. M., & Hagoort, P. (2009). Retrieval and Unification of Syntactic Structure in Sentence Comprehension: an fMRI Study Using Word-Category Ambiguity. *Cerebral Cortex, 19*(7), 1493–1503.

Stirnberg, R., Huijbers, W., Brenner, D., Poser, B. A., Breteler, M., & Stöcker, T. (2017). Rapid whole-brain resting-state fMRI at 3 T: Efficiency-optimized three-dimensional EPI versus repetition time-matched simultaneous-multi-slice EPI. *NeuroImage, 163*, 81–92.

Summerfield, C., Trittschuh, E. H., Monty, J. M., Mesulam, M.-M., & Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nature Neuroscience, 11*, 1004–1006.

Talaraich, J., & Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain*. New York: Thieme.

Thoma, V., & Henson, R. N. (2011). Object representations in ventral and dorsal visual streams: fMRI repetition effects depend on attention and part–whole configuration. *NeuroImage, 57*(2), 513–525.

Thompson-Schill, S. L., D'Esposito, M., Aguirre, G. K., & Farah, M. J. (1997). Role of left inferior prefrontal cortex in retrieval of semantic knowledge: A reevaluation. *PNAS, 94*(26), 14792–14797.

Tukey, J. W. (1949). Comparing Individual Means in the Analysis of Variance. *Biometrics, 5*(2), 99–114. Retrieved from http://www.jstor.org/stable/3001913

Turken, A. U., & Dronkers, N. F. (2011). The Neural Architecture of the Language Comprehension Network: Converging Evidence from Lesion and Connectivity Analyses. *Frontiers in Systems Neuroscience, 5*(1).

Visser, M., Jefferies, E., Embleton, K. V., & Lambon Ralph, M. A. (2012). Both the Middle Temporal Gyrus and the Ventral Anterior Temporal Area Are Crucial for Multimodal Semantic Processing: Distortion-corrected fMRI Evidence for a Double Gradient of Information Convergence in the Temporal Lobes. *Journal of Cognitive Neuroscience, 24*(8), 1766–1778.

Vosse, T., & Kempen, G. (2000). Syntactic structure assembly in human parsing: a computational model based on competitive inhibition and a lexicalist grammar. *Cognition, 75*(2), 105–143.

Vuilleumier, P., Schwartz, S., Duhoux, S., Dolan, R. J., & Driver, J. (2005). Selective Attention Modulates Neural Substrates of Repetition Priming and "Implicit" Visual Memory: Suppressions and Enhancements Revealed by fMRI. *Journal of Cognitive Neuroscience, 17*(8), 1245–1260.

Whitney, C., Kirk, M., O'Sullivan, J., Lambon Ralph, M. A., & Jefferies, E. (2010). The Neural Organization of Semantic Control: TMS Evidence for a Distributed Network in Left Inferior Frontal and Posterior Middle Temporal Gyrus. *Cerebral Cortex, 21*(5), 1066–1075.

Wible, C. G., Han, S. D., Spencer, M. H., Kubicki, M., Niznikiewicz, M. H., Jolesz, F. A., … Nestor, P. (2006). Connectivity among semantic associates: An fMRI study of semantic priming. *Brain and Language, 97*(3), 294–305.

Wickham, H. (2009). *ggplot2: Elegant Graphics for Data Analysis*. New York: Springer-Verlag. Retrieved from

http://ggplot2.org

Xiang, H.-D., Fonteijn, H. M., Norris, D. G., & Hagoort, P. (2010). Topographical Functional Connectivity Pattern in the Perisylvian Language Networks. *Cerebral Cortex, 20*(3), 549–560.

Zhu, Z., Zhang, J. X., Wang, S., Xiao, Z., Huang, J., & Chen, H.-C. (2009). Involvement of left inferior frontal gyrus in sentence-level semantic integration. *NeuroImage, 47*(2), 756–763.

Supplementary information

Lists of the used stimuli can be found in the publication on the CNS Journal website: https://www.ru.nl/master/cns/journal.

The Combined Effects of Alzheimer's Disease, Stroke and Gender on Aβ Burden and Synaptic Density

Minou Verhaeg[1], Klara Lohkamp[1], Maximilian Wiesmann[1], Amanda Kiliaan[1]

[1]*Department of Neuroanatomy, Radboud University Medical Centre Nijmegen, Donders Institute for Brain, Cognition and Behaviour, The Netherlands*

Alzheimer's disease (AD) and stroke are both diseases that have a very high impact on the current society. Prevalence of both conditions is expected to rise due to the increasing life expectancy of the population. There is a close relation between stroke and AD (e.g., prevalence of stroke is increased in AD patients and stroke increases the risk of AD). In the current longitudinal study, the interaction of stroke and AD was investigated in the APPswe/PS1dE9 mouse model with a 30 minute-induced middle cerebral artery occlusion in the right hemisphere. This study focused on the Aβ burden and synaptic density, via synaptophysin expression in the APP model, the stroke model and the combined APP/stroke model. Both male and female mice were used because there are differences in risk factors and pathology in both AD and stroke between the two sexes. An increased Aβ burden, especially visible in the cortex, hippocampus, thalamus and basal ganglia, combined with higher Aβ concentrations in the anterior part of the brain were observed in female compared to male APP mice. Eight months after stroke induction, male APP mice showed a lower Aβ burden in the cortex of the affected hemisphere compared to the unaffected hemisphere. Furthermore, between the hemispheres, asymmetrical expression of synaptophysin was observed in wildtype sham animals. Male mice had a higher synaptophysin expression in the affected hemisphere, while female mice had a decreased synaptophysin expression in the affected hemisphere. Lastly, female APP mice showed lower synaptophysin expression in the affected hemisphere than males. Overall, stroke affected the Aβ burden differently in male and female mice. These gender differences are most likely also the cause of the observed differences in synaptic density. More research into other AD markers, especially vascular pathology, is required to further determine the long-term effects of stroke on AD.

Keywords: Alzheimer's disease, ischemic stroke, neuropathology, animal models, amyloid beta, synaptic density

Corresponding author: Minou Verhaeg; E-mail: minou_verhaeg@hotmail.com

List of abbreviations

AD	Alzheimer's disease
APP	Amyloid beta precursor protein
Aβ	Amyloid beta
B2M	Beta-2 microglobin
CAA	Cerebral amyloid angiopathy
CBF	Cerebral blood flow
CCA	Common carotid artery
CSF	Cerebral spinal fluid
ELISA	Enzyme-linked immune sorbent assay
EpoE	Epolipoprotein E
GLUT-1	Glucose transporter 1
IHC	Immunohistochemical staining
MCA	Middle cerebral artery
MCAO	Middle cerebral artery occlusion
MCI	Mild cognitive impairment
PBS	Phosphate-buffered saline
PSD-95	Post synaptic density marker 95
qPCR	Quantitative polymerase chain reaction
SPB	Systolic blood pressure
SYP	Synaptophysin
WT	Wildtype

Due to increasing life expectancy, dementia has developed into one of the major public health problems worldwide. In America alone it is estimated that 5.8 million people were living with Alzheimer's disease (AD) in 2019 (Association, 2019). AD is the most common form of dementia and starts with occasional problems with episodic memory, referred to as mild cognitive impairment (MCI). Ultimately MCI can develop into dementia with profound memory impairments, disorientation, and deficits in motor function (D. J. Selkoe & Schenk, 2003). From a neuropathological perspective AD is characterized by the accumulation of amyloid beta (Aβ) plaques, neurofibrillary tau tangles, and the loss of neurons and their synapses (Dennis J Selkoe, 2002). The cognitive decline observed in AD is most closely associated with the synaptic loss that is observed in the brain (Dennis J Selkoe, 2002). Hippocampus, amygdala and the frontal, temporal, and parietal lobes undergo significant synaptic loss (DeKosky S.T., 1996). Previous studies have found that a reduction in synaptic density (indicated by synaptophysin (SYP) expression), and overall synaptic degeneration play a critical role in the progression of AD (Robinson et al., 2014; Wuwongse et al., 2013; Zeng et al., 2015). This synaptic dysfunction, along with neuroinflammation, apoptosis, and dysregulation of neurons, is partially mediated by Aβ and tau proteins

(Goedert & Spillantini, 2006; Klyubin, Cullen, Hu, & Rowan, 2012; Shankar & Walsh, 2009).

The accumulated Aβ proteins (plaques) are another characteristic of AD. These plaques are specific to AD and are not present in other types of dementia, or directly linked to dementia symptoms (Jarrett, Berger, & Lansbury Jr, 1993). The Aβ protein is commonly present in two lengths: Aβ40 and Aβ42. Aβ42 is more prone to aggregation, leading to the formation of the Aβ plaques (Jarrett et al., 1993). These protein aggregations are primarily present in the hippocampus, amygdala and cortices of the frontal, temporal and parietal lobes (Gouras et al., 2000; D. J. Selkoe & Schenk, 2003). AD can be divided into two types: familial and sporadic AD. Most AD patients have no genetic predisposition and are classified as sporadic. The smaller group, familial AD, is autosomal dominant and strikes most often before the age of 65. This familial variant of AD is driven by mutations in Aβ precursor protein (APP), apolipoprotein E (EpoE) or presenilin 1 or 2, which cause elevated levels of Aβ40, Aβ42 and the ratio between them (Cacace, Sleegers, & Van Broeckhoven, 2016; Holtzman et al., 2000). One hypothesis, the Aβ hypothesis, states that Aβ deposits are also the causal factor in the sporadic form of AD. However, the Aβ hypothesis is under heavy debate, since Aβ does not directly cause dementia symptoms (de la Torre, 2004). Increasingly more evidence points to a direct involvement of the cerebral vasculature in the development of AD (de la Torre, 2004). Therefore, the vascular hypothesis was introduced, wherein vascular factors contribute significantly to the development of dementia and AD.

Vascular hypothesis

Multiple vascular factors have been proven to contribute to the risk of developing AD (R. Kalaria, 2002; R. N. Kalaria, Akinyemi, & Ihara, 2012). In total, 60 to 90% of AD patients show cerebrovascular pathology at autopsy (R. N. Kalaria et al., 2012). In the first documented case of AD in 1906 by Alois Alzheimer arteriosclerotic changes in cerebral blood vessels already were mentioned (Drouin & Drouin, 2017). Changes in cerebral vasculature are closely linked to AD and Aβ accumulation, since efflux transporters located in the cerebral vessels and drainage through perivascular pathways into the cervical lymph nodes are mainly responsible for the clearance of Aβ out of the brain (Ueno, Chiba, Matsumoto, Nakagawa, & Miyanaka, 2014). Aβ deposits have also been observed

inside the blood vessels in the brain, known as a condition called cerebral amyloid angiopathy (CAA) (Lewis et al., 2006; McGowan et al., 2005). There is a significant association between certain vascular risk factors, such as hypertension in midlife, and development of AD or other forms of cognitive decline later in life (Kivipelto et al., 2001; Skoog et al., 1996). Hypertension leads to atherosclerotic changes in the arteries, lowering the cerebral blood flow (CBF) (Skoog et al., 1996). A lower CBF can cause ischemia, increases the production of Aβ and reduces Aβ clearance, leading to an overall higher concentration of Aβ in the brain (Nelson, Sweeney, Sagare, & Zlokovic, 2016; Ramanathan, Nelson, Sagare, & Zlokovic, 2015; Ueno et al., 2014). The onset of hypertension can already appear decades before the onset of AD (Skoog et al., 1996).

Ischemic stroke

Hypertension is also a major risk factor for further vascular complications, including stroke. Stroke can be divided into two types: ischemic stroke, in which an artery is occluded, and hemorrhagic stroke, in which an artery bursts (Amarenco, Bogousslavsky, Caplan, Donnan, & Hennerici, 2009). During a stroke, the CBF to the afferent brain areas of the affected vessels is severely lowered. Ischemic stroke is the most common variant, which most often occurs in a specific artery. Over 90% of ischemic stroke cases involve the middle cerebral artery (Roger et al., 2011). The affected area can be divided into an area of irreversible injury, the ischemic core, and an area of damaged but salvageable tissue, the penumbra. The penumbra is hypo-perfused, but cells still maintain basic cerebral metabolic rate and can be salvaged when reperfusion to the brain area occurs (Marchal et al., 1996). Due to the loss of CBF and the drop in clearance, accumulation of Aβ is seen in the regions directly affected by the stroke.

There is a clear link between AD and stroke. Stroke has been proven to double the prevalence of AD in elderly (Honig et al., 2003; Sun et al., 2006). Furthermore, AD patients with stroke seem to have more severe dementia compared to those who do not suffer from stroke (Leys et al., 1999; White et al., 2002), due to the tissue and vascular damage caused by the stroke. The interaction between vascularity, especially stroke, and AD pathology works both ways. A stroke does not only strengthen AD symptoms, but people with AD also have a high risk of suffering from a stroke (Chi, Chien, Ku, Hu, & Chiou, 2013). Stroke and AD also share a lot of the same risk factors such as aging (Sahathevan,

Brodtmann, & Donnan, 2012), hypertension, diabetes, hypercholesterolemia, smoking and obesity (de Bruijn & Ikram, 2014; Nelson et al., 2016). More and more studies start to question the clear diagnostic separation between vascular dementia and AD (Agüero-Torres, Kivipelto, & von Strauss, 2006; de la Torre, 2004). Due to the significant contribution of these vascular factors a model on the development of dementia was developed by Sweeney et al. (Fig. 1) (Sweeney, Sagare, & Zlokovic, 2015). This model takes both vascular damage and the contribution of Aβ into account. It also includes the interaction of the two factors, possibly through the interaction between stroke and APP, as this link has been proven in animal models (Jendroska, Hoffmann, & Patt, 1997; Jendroska et al., 1995; Nakamura, Takeda, Niigawa, Hariguchi, & Nishimura, 1992). The precise mechanism behind this interaction however, is still not well understood (de la Torre, 2004).

Gender differences

One highly significant risk factor which needs to be studied in more detail is the effect of gender. Both in AD and stroke sex differences play a profound role in prevalence and progression of the
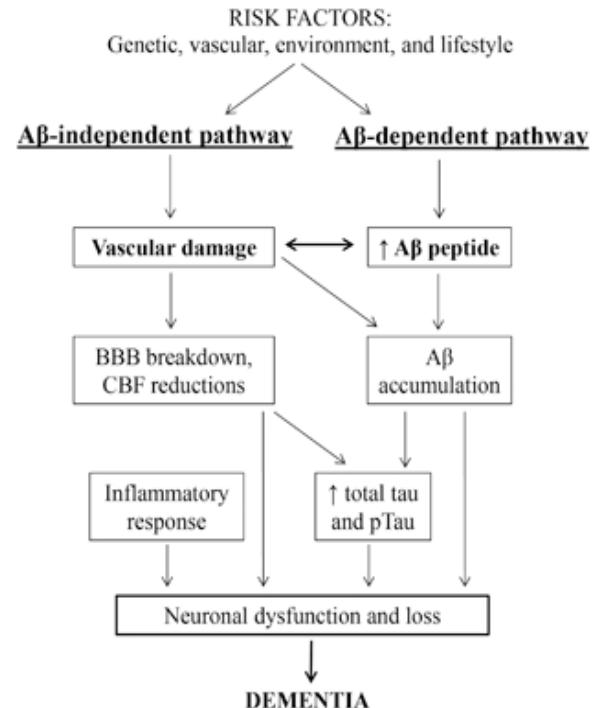

Figure 1. Multifactorial model of developing dementia or AD. Illustrating two pathways, the Aβ dependent and independent pathway, that are triggered by a variety of risk factors. These pathways interact and together cause neuronal dysfunction and neuronal loss leading to AD [37].

disease. In the case of AD, females seem to have a higher susceptibility to the disease compared to age-matched males (Azad, Al Bugami, & Loy-English, 2007; Launer et al., 1999). The higher incidence of AD in females might be explained (partly) by the higher susceptibility to AD risk factors. Females over the age of 75 have a higher prevalence of hypertension, hyperlipidemia and diabetes compared to age-matched men. These are all factors that are associated with the development of AD (Azad et al., 2007). In APP mice, females have a higher level of Aβ40 and Aβ42 in the hippocampus at four months of age. Female mice also have a higher plaque number and heavier Aβ burden at both 12 and 17 months of age (Wang, Tanila, Puoliväli, Kadish, & van Groen, 2003).

Compared to age-matched females males have a higher incidence rate for stroke, but women have an overall higher prevalence rate and more severe strokes (Kim, Lee, Roh, Ahn, & Hwang, 2010). Age is an important risk factor for stroke (Sahathevan et al., 2012). Women live longer and are therefore more likely to suffer from stroke (Appelros, Stegmayr, & Terént, 2009). Furthermore, a division between pre and post menopause should be made. Studies have shown that premenopausal women experience fewer strokes than men, but postmenopausal women experience more strokes than age-matched men. Estrogen has proven to protect against AD (Haast, Gustafson, & Kiliaan, 2012). After menopause, estrogen levels will drop and the chance of developing AD increases. This postmenopausal phenomenon, and the fact that the women on average live longer, lead to the fact that the onset of the stroke and therefore the severity and mortality of the stroke is higher in women (Haast et al., 2012). Furthermore, men seem to be more prone to ischemic stroke then women, while women are more prone to subarachnoid hemorrhage, regardless of age (Appelros et al., 2009). Unfortunately, only few studies on AD, vascular factors, and the interaction use both male and female models. A lot is still unclear about the consequences of sex on either AD, stroke, or the combination of both.

Although a lot of research has been done on the development of AD (de la Torre, 2004), there is much still unclear about the sex differences and the interaction of AD and stroke. Animal research plays an important role in better understanding the interaction between AD, stroke, and gender. One animal model that is widely used to study AD is the APPswe/PS1dE9 (APP) mouse model (Jankowsky et al., 2001). These mice contain human transgenes for both APP and PSEN1 (Goodwin et al., 2019), which are not naturally expressed in mice. Both Aβ42 and Aβ40 levels increase with age, but Aβ42 levels are especially elevated in these mice (Jankowsky et al., 2003). Amyloid deposits appear first in the cortex at 6 weeks of age, in the hippocampus at 3 months of age and finally in the striatum, thalamus and brain stem at around 4 or 5 months of age. The specific Aβ burden differs per sex. Female mice reach the maximum level of deposits around 9 months of age, while male mice reach this maximum at around 12 months of age (Ordóñez-Gutiérrez, Antón, & Wandosell, 2015). Dendritic spine loss appears around the plaques approximately 4 weeks after plaque formation and eventually at 7 months of age, mice start showing cognitive impairments. A gradual reduction in GLUT-1 between 8 and 18 months of age has been reported in this model too (Hooijmans et al., 2007).

Current study

This study focused on the long-term relationship between AD and vascular risk factors, both in male and female mice. The current study will investigate the effects of stroke, AD and gender on synaptic density and Aβ plaques in the mouse brain, both markers for AD pathology. The precise relation between stroke and AD is still unclear. Studies on this topic are limited and often not of a longitudinal nature (Kemppainen, Hämäläinen, Miettinen, Koistinaho, & Tanila, 2014).

To do this, a mouse model of AD, expressing Aβ, was subjected to a transient middle cerebral artery occlusion in the right hemisphere of the brain to mimic a stroke. Since the left hemisphere did not suffer from any reduced blood flow, a within-subject control (left compared to right hemisphere) can be used next to the overall control group. Animals were sacrificed 8 months after stroke induction to investigate the long-term effects of AD and stroke.

Via quantitative polymerase chain reaction (qPCR), synaptophysin was analyzed as an indicator for synaptic density, and Aβ was analyzed via enzyme-linked immune sorbent assay (ELISA) and immunohistochemical stainings (IHC). Multiple areas will be analyzed via Aβ staining.

It is important to investigate the effect of stroke and AD in both males and females, since sex differences play an important role in both AD and stroke. It is expected that the stroke will have an aggravating effect on the AD mouse model, shown by increased Aβ deposits.

Synaptic density could be decreased in APP animals as a result of AD pathology. It is expected

that the synaptic density is further decreased by stoke induction, which will present as a further decrease in synaptic density in the combined AD-stroke mice. When comparing male and female mice, it is expected that the female mice will have a stronger reaction to the stroke, represented by decreased synaptic density, since females seem to recover less than age matched controls over time (Kim et al., 2010). Furthermore, female APP mice are expected to have a heavier Aβ burden compared to male APP mice, as shown before in other studies (Wang et al., 2003).

Methods

Animals

This study involved the use of two strains of mice: the APPswe/Ps1dE9 (APP) mice as a model for AD and the C57B1/6 wildtype (WT) littermates. For both strains, male and female mice were included. The APPswe/PS1dE9 founder mice were obtained from John Hopkins University, Baltimore, MD, USA (Jankowsky et al., 2003). A colony was bred at the Central Animal facility at the Radboud University medical center in the Netherlands. Originally, the line was maintained on a hybrid background by backcrossing to C3HeJxC57BL/6J F1 mice. For this study, the desired mice were created by backcrossing the breeder mice to C57BL/6J for fifteen generations. At 3 months of age, the APP mice start to express amyloid beta plaques in their brain (Radde et al., 2006).

Mice were housed at 21°C, at an artificial 12:12h light-dark cycle (lights on at 7 a.m.) and were housed with a maximum of six animals per cage. After surgery, mice were housed individually in digital ventilated cages (DVC) for 24/7 activity monitoring

(Tecniplast, Buggiate, Italy).

Water and food (Sniff rm/h V1534, Bio Service, Uden, The Netherlands) were available *ad libitum*. Experiments were performed according to Dutch federal regulations for animal protection and were approved by the Veterinary Authority of Radboud University medical center in Nijmegen, the Netherlands, and the Animal Experiment Committee of the Radboud University in Nijmegen, the Netherlands (dierenexperimentencommissie, DEC, 2012-248 & 2015-0079).

Study design

Animals were divided in eight groups, depending on genotype, surgery, and sex (Table 1). A total of 144 mice were included in the study. Ninety-two mice completed all experiments. An estimation of 14 animals per group was made to reach sufficient power. All animals were randomly divided between stroke and sham, considering an equal division of sex.

During postmortem procedures and all data analyses, researchers remained double blinded for experimental groups.

Due to the large number of animals, and the longitudinal aspect of this study, the experiment was divided in multiple cohorts, each containing the same experimental structure and design. Animals from each experimental group were distributed evenly over the different cohorts. The cohorts were distributed over 2 years (cohort 1&2: May 2016-May 2017, cohort 3&4: February 2017-April 2018, cohort 5&6: September 2017-August 2018).

Different tests were performed before and after stroke induction, to assess health and behavior of the mice (Fig. 2). At three and a half months of age general health parameters body weight and systolic
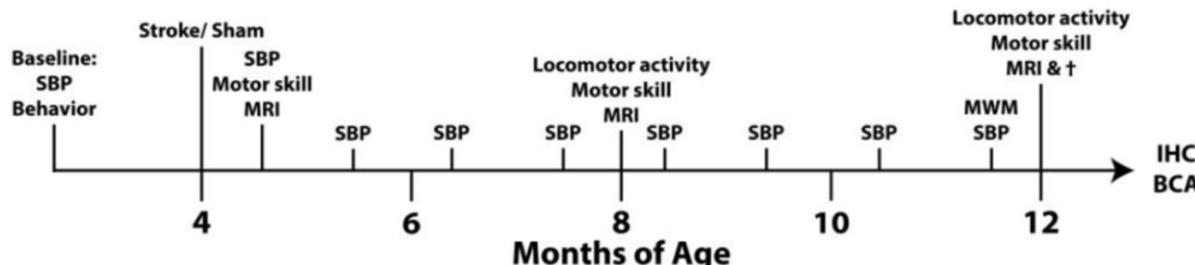

Figure 2. Overview of the current study design. A timeline of the study. Baseline SBP and behavior were assessed at 2.5 months. Stroke was induced at four months of age. At 4.5, 8 and 12 months, behavior, cognition, motor skill and locomotor activity were assessed by performing the pole test, grip test and open field. At the same timepoints, MRI was performed. At 12 months of age, rotarod and MWM were performed. After the last tests at 12 months of age animals were sacrificed, and their brains were obtained for IHC and BCA. SPB: systolic blood pressure. MWM: Morris water maze. IHC: immunohistochemistry. BCA: biochemical analysis.

Table 1. Animal groups. An overview of the animal groups with all differential characteristics, including group size (n). APP = APPswe/PS1dE9 mice model for AD. WT = wildtype.

Strain	APP				WT			
Surgery	Stroke		Sham		Stroke		Sham	
Sex	♂	♀	♂	♀	♂	♀	♂	♀
n	8	9	9	9	12	14	14	15
Group name	1: APP – Stroke – Male	2: APP – Stroke – Female	3: APP – Sham – Male	4: APP – Sham – Female	5: WT – Stroke – Male	6: WT – Stroke – Female	7: WT – Sham – Male	8: WT – Sham – Female

blood pressure (SPB) were measured. Baseline behavior and motor skills were assessed for all mice via the pole test, grip test and open field test.

At four months of age either the stroke or sham operation was performed. Two weeks after the operation all behavioral tests were repeated and imaging was conducted (resting state functional magnetic resonance imaging (rsfMRI), diffusion tensor imaging (DTI) and flow-sensitive alternating inversion recovery-arterial spin labeling (FAIR-ASL), to assess the effects of the surgery. SBP measurements were repeated every month. Imaging, motor skill, and behavioral assessments were repeated at eight and 12 months of age. At 12 months of age, the rotarod test and the Morris water maze were performed in addition to the other tests. All behavioral procedures were performed during daytime. Afterwards, all animals were sacrificed via transcardinal perfusion and tissues were collected for postmortem analyses.

Postmortem analysis was done after tissues of all cohorts were collected. The order of all postmortem analyses was randomized for cohort numbers.

Transient occlusion of the middle cerebral artery

At four months of age, ischemic stroke was induced via a 30 minute transient occlusion of the middle cerebral artery (MCAO) in the right hemisphere (Fig. 3, (Bertrand, Dygert, & Toborek, 2017)). 5 mg/kg Rimadylwas injected before surgery, to prevent inflammation and pain during recovery. Animals were anesthetized with 2-3% isoflurane (Abbott Animal Health, Abbott park USA) and maintained on 1,5% isoflurane during the surgery. The monofilament (Doccol corporation, Sharon USA, 7-0 monofilimanet, 190-200 μm diameter, 2-3 mm coating) was inserted via the common carotid artery (CCA), through the internal carotid artery (ICA) up to the middle cerebral artery (MCA). Cerebral blood flow was monitored with a Laser doppler flow probe (moorVMS-LDF2, Moor

Instruments). Occlusion was maintained for 30 minutes and the surgery was deemed successful if an 80% drop in regional cerebral blood flow was reached. For the control group the filament was introduced shortly, but was retracted immediately after reaching the MCA. During the first week after surgery mice were checked daily for signs of discomfort or stress.

Biochemistry

At 12 months of age, the mice were sacrificed via transcardinal perfusion using 0.1M phosphate-buffered saline (PBS) followed by a solution of 4% paraformaldehyde in 0.1M PBS. Tissues from the ears, eyes, brain, heart, aorta, and CCA were collected, as well as some cerebral spinal fluid (CSF). The brains were cut in a transversal plane, such that part of the stroke was present in both the frontal and ventral part. The frontal brain half was divided in the left and right hemisphere and used to isolate deoxyribonucleic acid (DNA), ribonucleic acid (RNA), and proteins.

RNA was collected and prepared for qPCR (Supplementary protocol 1-4). Next to the RNA concentrations Nanodrop measures also provides measurements which can be used as quality control and purity check. There were some samples that were outside of the advised range. These samples were checked for abnormalities regarding the qPCR results, both for primary outcomes and melting curves. For the cDNA synthesis, 400 ng of total RNA was used. QPCR was performed using SYBRGreen as a reagent, for glucose transporter 1 (GLUT-1), synaptophysin (SYP) and housekeeping gene beta-2 microglobin (B2M) (Supplementary protocol 5). Supplementary protocol 2 to 4 were redone for samples which had abnormal values for the housekeeping gene B2M. Thresholds for the Ct values were set at 0.08.

DNA was collected either via an ear clip or via postmortem isolation of the brain tissue (Supplementary protocol 6). Genotyping was
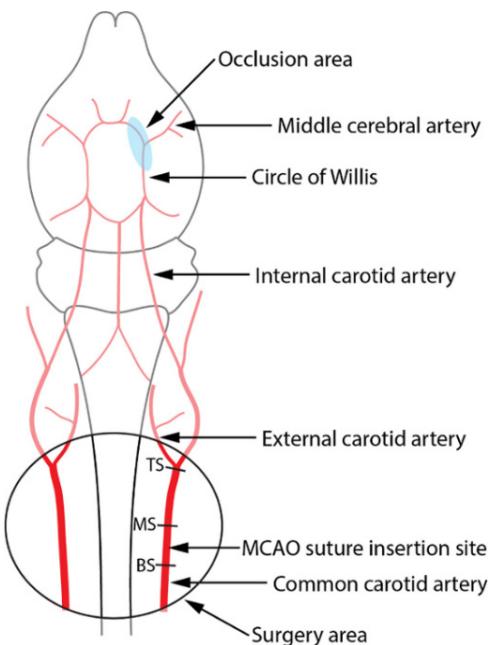
Figure 3. Induction of ischemic stroke. The incision was made in the neck of the animal (surgery area). The filament was inserted in the right common carotid artery and pushed up to the middle cerebral artery (the occlusion area). The filament was held in place for 30 minutes.

performed on an agarose gel (Supplementary protocol 7).

The protocol for protein isolation was derived from Simões et al. (2013) and adjusted to fit this study (Supplementary protocol 7). Proteins were only isolated from APP mice, since the WT mice do not express the proteins of interest, amyloid beta. Protein concentrations were adjusted to each hold 25 ng in 50 μL. ELISA was performed for the amyloid beta 40 and 42 proteins, using the protocols Human Aβ42 ELISA kit (KHB3441, invitrogen) and Human Aβ40 ELISA Kit (KHB3481, invitrogen). Concentrations were determined via calibration samples and calculated to the unit of pg/ng.

Immunohistochemistry

The posterior brain halves were postfixed in 4% paraformaldehyde overnight at 4°C. The next day, brains were separately transferred to a solution containing 0.1M PBS and 0.01% sodium azide. Brains were cryoprotected in 30% sucrose in phosphate buffer for 24h before cutting. This part of the brains was used for immunostaining of GLUT-1 and Aβ. Brains were cut in eight series of 30-μm-thick coronal sections using a freezing sliding microtome (Microm HM 450, Walldorf, Germany). Brain sections were stored in 0.1M PBS with 0,01% sodium azide at 4°C before immunohistochemical

staining. In total two immunohistochemical stainings were performed using standard free-floating labelling procedure at room temperature. For each staining, one subseries of brain sections per animal was used. After every incubation step, brain sections were rinsed 3 times for 15 minutes in 0.1M PBS for the GLUT-1 staining. For the Aβ staining, brain sections were rinsed in tris Buffered Saline + 0,5% Triton (TBS-T).

All free-floating brain sections were rinsed with 0.1M PBS for 15 minutes and treated with 0.3% H_2O_2 in 0.1M PBS for 30 minutes. Next, brain sections were pre-incubated in 2 ml 0.1M PBS-BT for 30 minutes. For Aβ staining, the brain sections were pre-incubated in 2 ml 0.05M Tri-sodium citrate solution at 85°C for 30 minutes. After this pre-incubation step, primary anti-bodies (GLUT-1: rabbit anti-GLUT-1 [1:40.000; Millipore, Billerica, MA, USA] or WO-2: mouse anti-human Aβ4-10 [1:10.000; Centre for Molecular Biology, University of Heidelberg, Germany]) were added overnight. After rinsing, secondary anti-bodies (GLUT-1: donkey anti-rabbit [1:1500; Jackson ImmunoResearch, West Grove, PA, USA], WO-2: donkey anti-mouse biotin [1:1500; Jackson ImmunoResearch, West Grove, PA, USA]) were added for 1.5 h followed by incubation with ABC-Elite (1:800; GLUT-1in PBS-BT, Aβ in TBS-T) for 1.5h. Afterwards, all brain sections were pre-incubated in DAB-Ni for 10 minutes followed by incubation in DAB-Ni with 30% H_2O_2 for 10 minutes. Brain sections were mounted on gelatine-coated slides (0.5% gelatine and 0.05% chrome-alum) and dried at 37°C overnight. The slides were dehydrated with alcohol, cleared with xylol and coverslipped with Entellan.

Quantification

The stained sections were viewed with a 5x objective Axio Imager A2 (Zeiss Germany), and ZEN software was used for image acquisition. Brain sections were preselected according to the mouse brain atlas of Franklin and Paxinos [66]. The regions of interest (ROIs) for Aβ contain the cortex, hippocampus and thalamus (bregma: -1.94), as reference areas which do not directly get blood supply from the MCA, and the cortex, corpus callosum and basal ganglia (bregma: 0.62), areas which are partially or directly supplied by the MCA. The analysis of all images was performed using ImageJ (National Institute of Health, Bethesda, MD, USA). For the Aβ staining the average plaque size, relative positive Aβ area, and number of plaques per area (mm²) were calculated.

Data analysis

All statistical analyses were preformed using IBM SPSS 25 software (IBM Corporation, New York, NY, USA). Data were split on genotype, surgery, and sex and checked for outliers. Two animals were excluded for the analyses of the ELISA for Aβ40 and Aβ42, because of extremely deviant values.

Differences between sex (male/female), surgery (sham/stroke), and hemispheres (left/right) were determined. Repeated measures ANOVA with Bonferroni corrections were used to determine significant differences between the left and right hemispheres. Multivariate ANOVA with Bonferroni corrections to determine significant differences between surgery and/or sex. Data were checked for interaction between the variables, and further split if necessary. GraphPad Prism (6.01, GraphPad Software, La Jolla California, USA) was used to create all presented figures. All data are presented as mean ± SEM. Data were considered significant if $p \leq 0.05$.

$0.08 \geq p \geq 0.05$, * $p \leq 0.05$, ** $p \leq 0.01$, *** $p \leq 0.001$.

Results

An overview of the various results can be found in Supplementary table 1.

Genotype

All animals were genotyped at the start and the end of the experiment. The PrP gene is present in all mice, while the hAPP gene is only present in the APP mice. WT mice only show one band at 750 bp, while the APP mice show two bands, at 750 bp and 400 bp. No abnormalities were seen regarding genotypeandall expression patterns were complementary to the 'assumed' genotype (APP vs WT) of the individual animals (Fig. 4).

Amyloid beta burden in the frontal brain

To determine the level of Aβ, proteins were isolated from the frontal part of the brain, including parts of the ischemic core and the penumbra. Proteins were isolated from the brains of all APP mice. ELISA for Aβ40 and Aβ42 were performed and the ratio between the two Aβ proteins was determined. No significant differences between the left and right hemisphere were found regarding Aβ40, Aβ42 or the Aβ42/Aβ40 ratio (Fig. 5).

Female APP mice showed a higher concentration of Aβ40 in both the left ($p<0.006$) and right ($p<0.024$) hemisphere compared to male APP mice (Figure 6A). No significant differences between females and males were found regarding Aβ42 or the Aβ42/Aβ40 ratio (Figure 6B-C). No significant difference between stroke APP and sham APP mice were found regarding Aβ40, Aβ42 or the ratio Aβ42/Aβ40 (Figure 6A-C).

Amyloid beta burden in ischemic and unaffected regions

To further determine the level of Aβ brain sections of APP mice were immunohistochemically stained for Aβ at bregma 0.62, close to the ischemic core and at bregma -1.94 further away from the ischemic region. At bregma -1.94 the cortex (Fig. 7A), hippocampus (Fig. 7B), and thalamus (Fig. 7C) were analyzed for differences in the average plaque size, relative Aβ positive area, and the number of Aβ plaques per area (mm²). A trend was visible in the cortex in male sham ($p<.067$) and female stroke ($p<.075$) mice, in which the right hemispheres seems to have a higher average plaque size compared to the left hemisphere. No significant differences were found between the left and right hemisphere in the cortex regarding relative Aβ area or number of Aβ plaques per area (mm²). No significant differences were found between the left and right hemisphere in the hippocampus regarding average plaque size. Female stroke mice showed a significant increase ($p<.032$) in relative Aβ area in the right hemisphere compared to the left hemisphere in the hippocampus. A trend ($p<.07$) was visible in the female stroke mice, in which the right hemisphere tended to have an increased number of Aβ plaques per area (mm²) compared to the left hemisphere in the hippocampus. Female sham mice showed a significant increase in average plaque size in the right hemisphere compared to the left hemisphere ($p<.003$). No significant differences were found between the left and right hemisphere in the thalamus in relative Aβ area or number of Aβ plaques per area (mm²).

Brain sections of APP mice were also immunohistochemically stained for Aβ at bregma 0.62, close to the ischemic core. At bregma 0.62, the cortex (Fig. 8A), corpus callosum (Fig. 8B) and basal ganglia (Fig. 8C) were analyzed for differences in the average plaque size, relative Aβ positive area and the number of Aβ plaques per area (mm²). Male stroke mice showed a significant decrease in average plaque size ($p<.03$), relative Aβ area ($p<.004$) and number
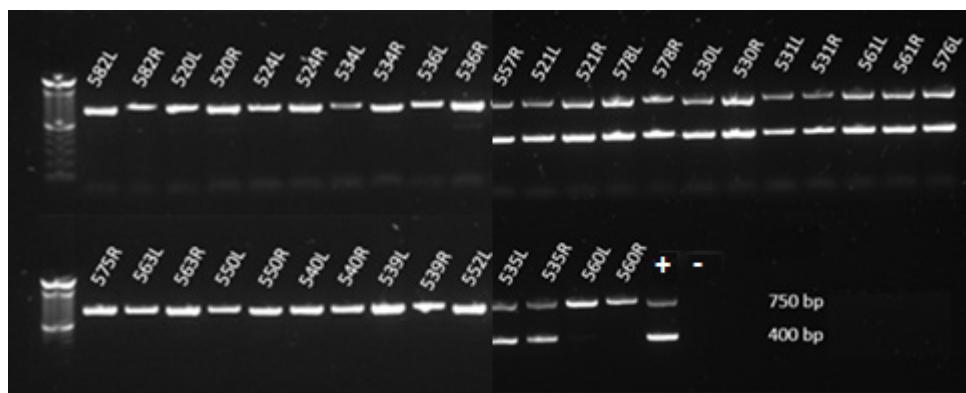
Figure 4. Genotyping by agrose gel. Representative photo of samples of the genotyping performed on an agrose gel. Other gels were performed in a similar way. WT animals only show one band at 750 bp (top band, PrP gene), the APP mice have two bands, at 750 bp (top band, PrP gene) and 400 bp (lower band, APP gene). Numbers refer to the animal number. L: left hemisphere, R: right hemisphere, +: positive control, -: negative control.

of Aβ plaques per area (mm²) (p<.038) in the right hemisphere compared to the left hemisphere in the cortex. No significant differences were found between the left and right hemispheres in the corpus callosum regarding average plaque size or relative Aβ area. Male sham animals showed a significant increase in number of Aβ plaques per area (mm²) in the right hemisphere compared to the left hemisphere (p<.032) in the corpus callosum. No significant differences were found between the left and right

hemispheres in the basal ganglia regarding average plaque size. Female stroke mice showed a significant decrease (p<.042) in relative Aβ area in the right hemisphere compared to the left hemisphere in the basal ganglia. No significant differences were found between the left and right hemispheres in the basal ganglia in number of Aβ plaques per area (mm²).

Data from bregma -1.94, the cortex (Fig. 9A), hippocampus (Fig. 9B), and thalamus (Fig. 9C) were further compared for sex (male vs female) and
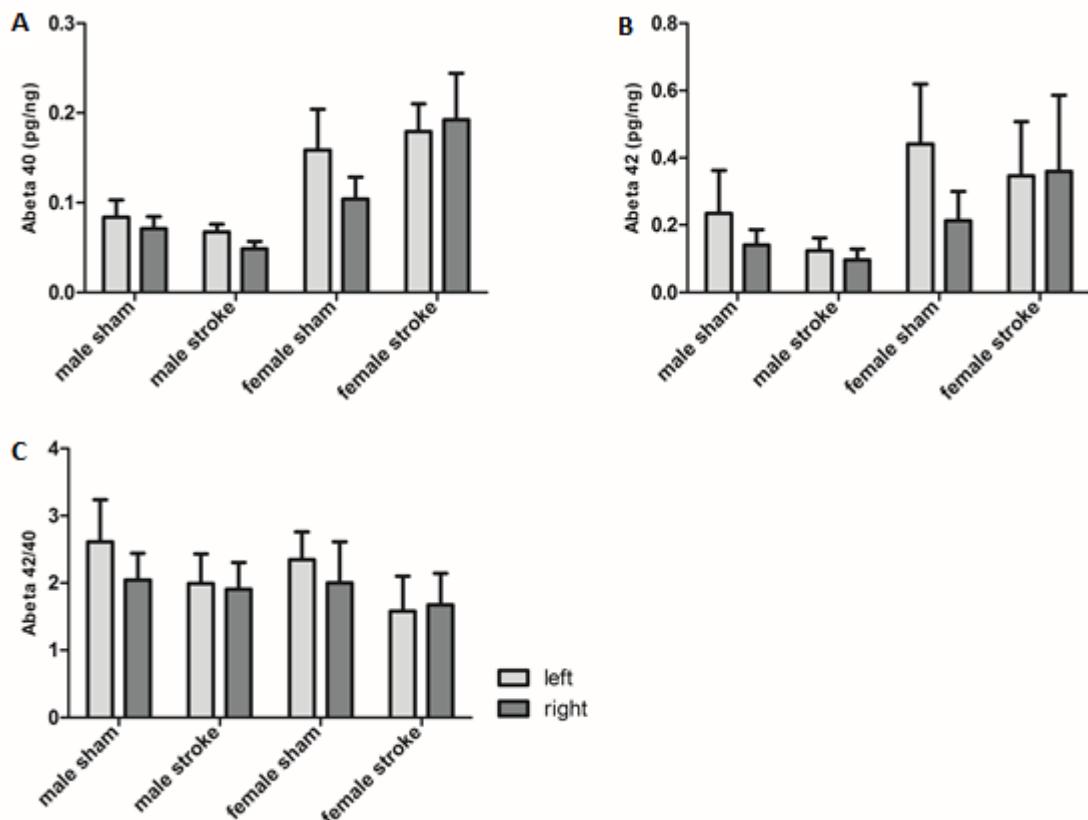

Figure 5. Concentration of Amyloid beta in the frontal part of the brain, L vs R effects. All data are presented as mean ± SEM. No significant differences between the left and right hemisphere were found regarding the concentration of Aβ40 **(A)**, Aβ42 **(B)** or the ratio of Aβ42/Aβ40 **(C)**.
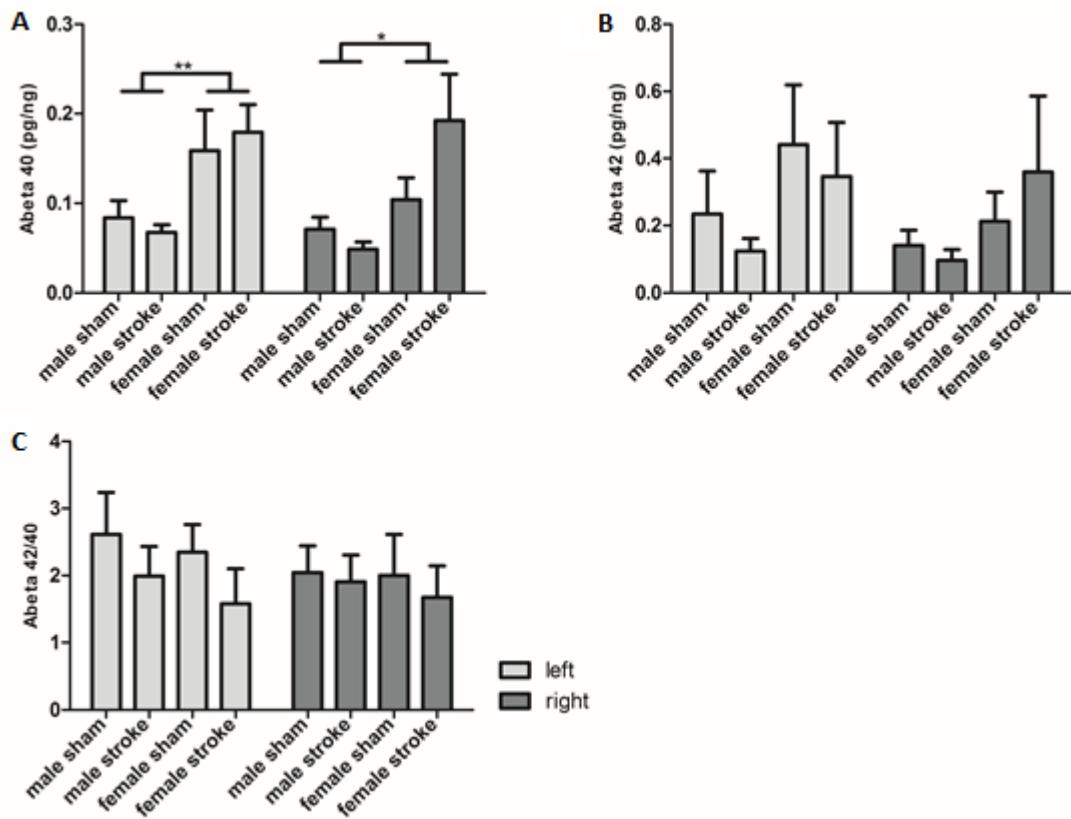
Figure 6. Concentration of Amyloid beta in the frontal part of the brain, sex, and surgery effects. All data are presented as mean ± SEM. The concentration of Aβ40 was increased in female mice compared to male mice in both the left (p<.006) and the right (p<.024) hemisphere **(A)**. No significant sex or surgery differences were found regarding the concentration of Aβ42 **(B)** or the ratio of Aβ42/Aβ40 **(C)**.

surgery (sham vs stroke) differences. No significant differences were found in the cortex between males and females regarding average plaque size, relative Aβ area, or number of Aβ plaques per area (mm²). No significant differences were found in the hippocampus between males and females regarding average plaque size. Female mice showed a significant increase in relative Aβ area in both left (p<.044) and right (p<.012) hemisphere compared to males in the hippocampus. Female mice also displayed a significant increase in number of Aβ plaques per area (mm2) in both left (p<.025) and right (p<.019) hemisphere compared to males in the hippocampus. A significant increase in average plaque size (p<.042) in the right hemisphere and in relative Aβ area in both left (p<.012) and right (p<.003) hemisphere was found in female mice compared to male mice in the thalamus. And furthermore, a significant increase in number of Aβ plaques per area (mm²) (p<.030) in the left hemisphere of the thalamus was seen in female mice compared to male mice. No significant differences between stroke and sham APP mice were found regarding relative Aβ area or number of Aβ plaques per area (mm2). No interactions were found between sex and surgery.

Differences in sex and surgery were also analysed in the cortex (Fig. 10A), hippocampus (Fig. 10B) and thalamus (Fig. 10C) on bregma 0.62. No significant differences were found in the cortex between males and females regarding average plaque size. In the cortex, the relative Aβ area in the right hemisphere (p<.003) and in number of Aβ plaques per area (mm²) in both the left (p<.031) and right (p<.001) hemisphere was significantly increased in female mice compared to male mice. No significant differences were found in the corpus callosum between males and females regarding average plaque size, relative Aβ area, or number of Aβ plaques per area (mm²). Female mice presented with a significant increase in the basal ganglia in average plaque size (p<.033) in the right hemisphere. A trend was visible in the basal ganglia in the right hemisphere, in which stroke mice seem to have a lower average plaque size compared to sham mice (p<.073). A significant increase in relative Aβ area in the left hemisphere (p<.004) was found in female mice compared to male mice. A trend (p<.069) was visible in the right hemisphere where female mice seem to have an increased number of Aβ plaques per area (mm²) compared to the left hemisphere. Female mice showed a significant increase in number of Aβ plaques per area (mm²) in both the left (p<.001) and
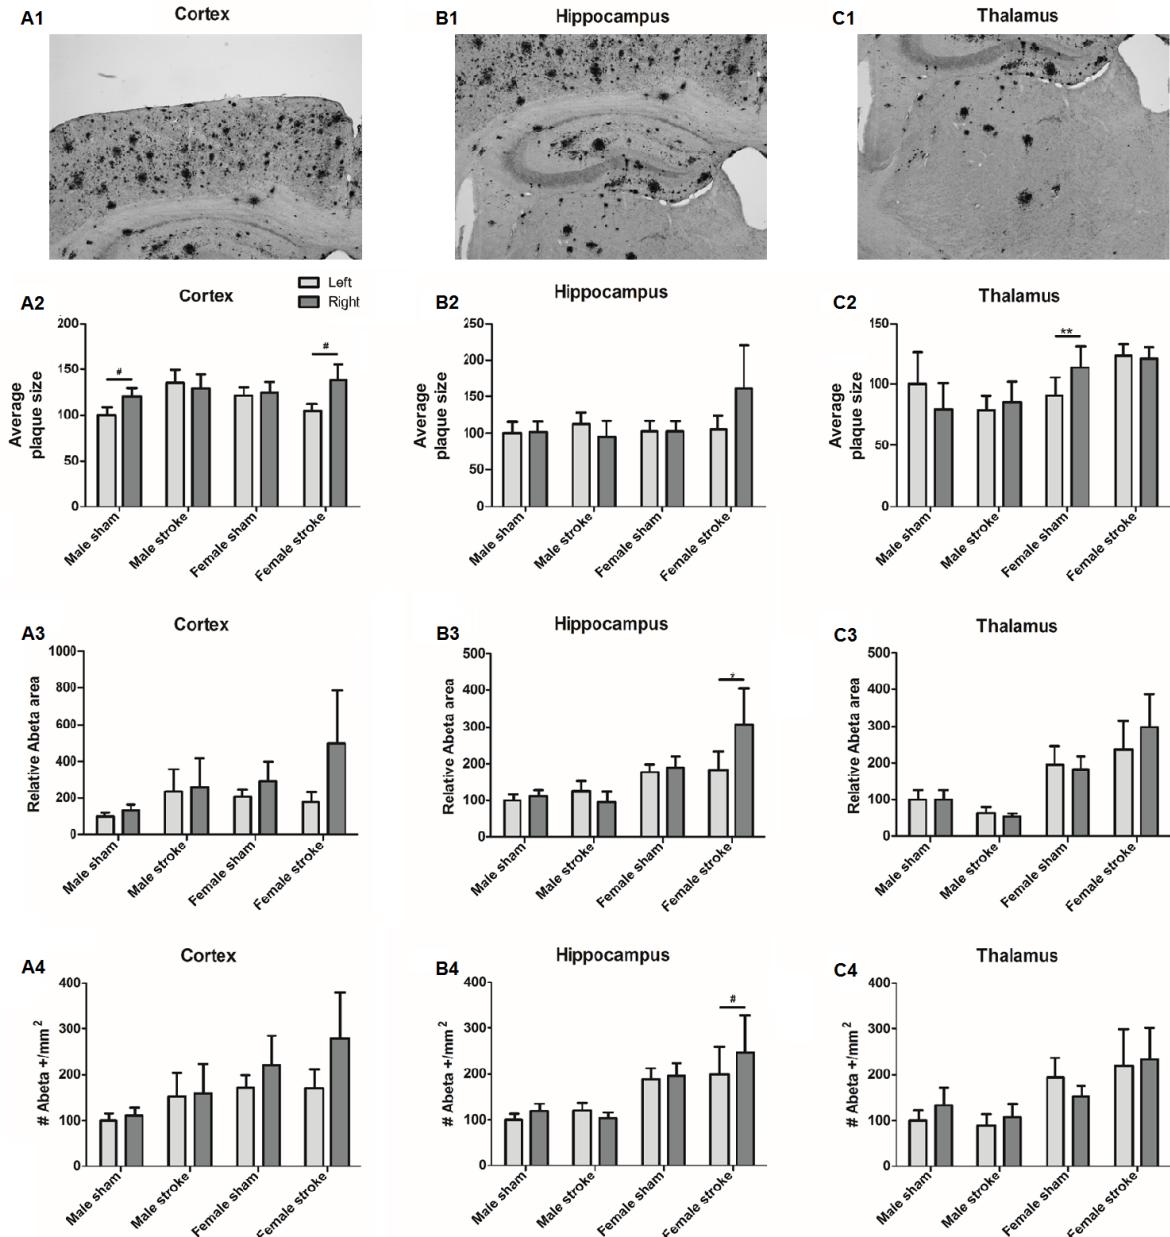
Figure 7. Average plaque size, relative Aβ positive area and number of Aβ plaques per area (mm2) in the cortex, hippocampus, and thalamus (bregma -1.94) of APP mice, L vs R effects. All data were normalized against the left male sham group and presented as mean ± SEM. A representative photo of the Aβ staining of the right hemisphere of the cortex is shown **(A1)**. A trend was visible in the average plaque size in the cortex in which the right hemisphere has a higher average plaque size than the left hemisphere in both male sham (p<.067) and female stroke animals (p<.075) **(A2)**. No significant differences between the left and right hemispheres were found in the cortex regarding relative Aβ positive area **(A3)** or number of Aβ plaques per area (mm2) **(A4)**. A representative photo of the Aβ staining of the right hemisphere of the hippocampus is shown **(B1)**. No significant differences between the left and right hemispheres were found in the hippocampus regarding average plaque size **(B2)**. Relative Aβ positive area was increased in the hippocampus in the right hemisphere compared to the left hemisphere (p<.032) in female stroke mice **(B3)**. A trend was visible in number of Aβ plaques per area (mm2) (p<.070) in the hippocampus, in which the right hemisphere has a higher amount of Aβ plaques than the left hemisphere **(B4)**. A representative photo of the Aβ staining of the right hemisphere of the thalamus is shown **(C1)**. Average plaque size was increased in the thalamus in the right hemisphere compared to the left hemisphere (p<.003). In female sham mice **(C2)**. No significant differences between the left and right hemisphere were found in the thalamus regarding relative Aβ positive area **(C3)** or number of Aβ plaques per area (mm2) **(C4)**.
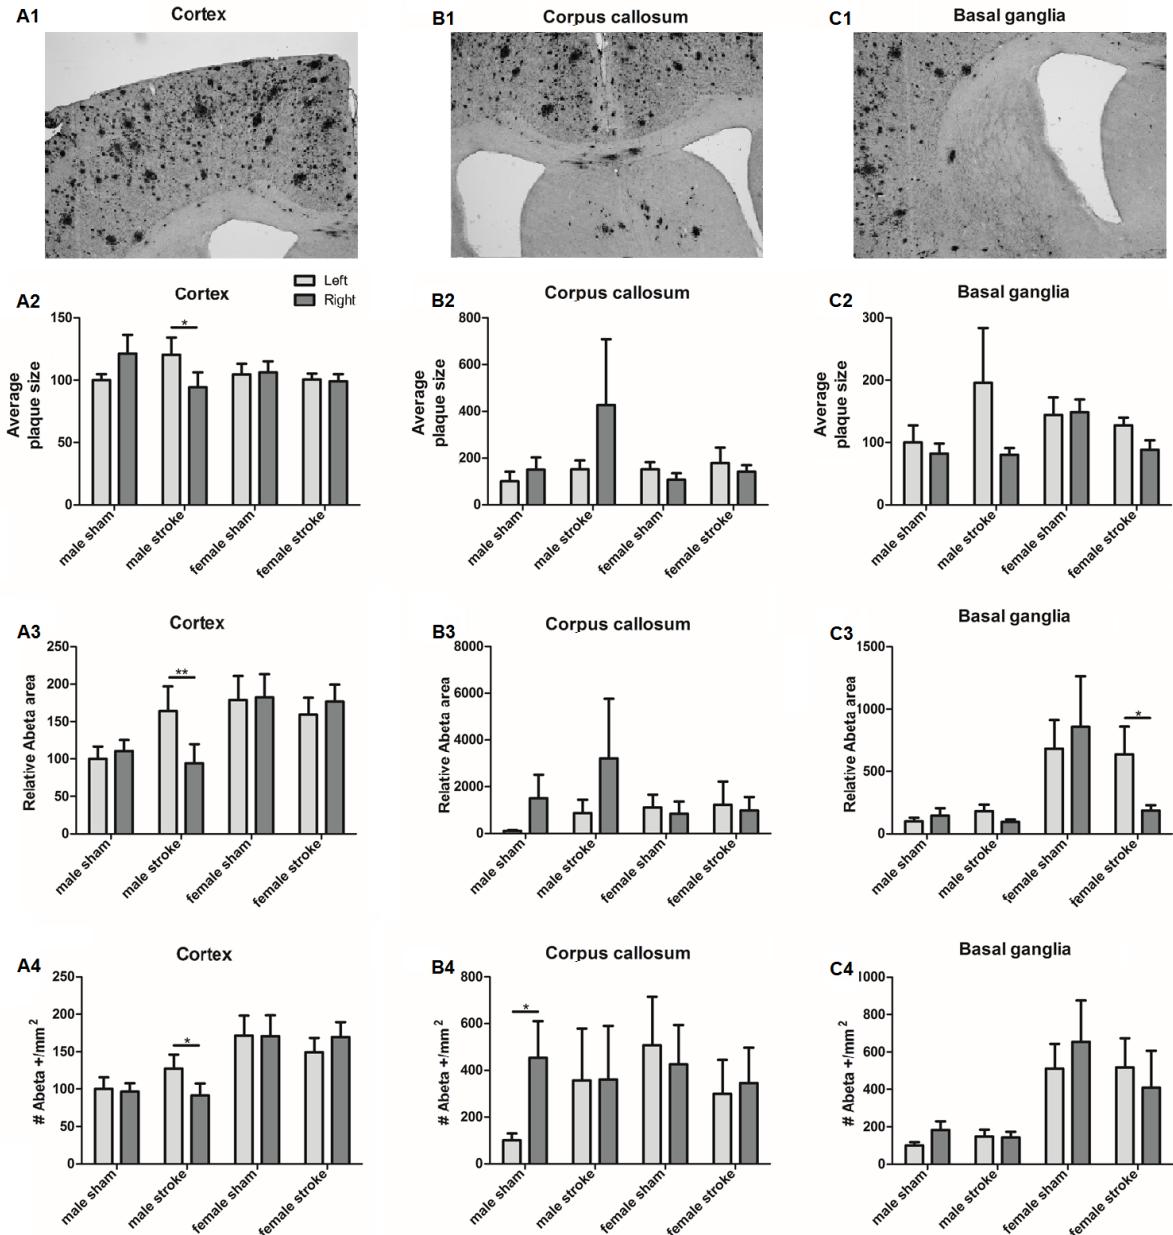
Figure 8. Average plaque size, relative Aβ positive area and number of Aβ plaques per area (mm2) in the cortex, corpus callosum, and basal ganglia (bregma 0.62) of APP mice, L vs R effects. All data were normalized against the left male sham group and presented as mean ± SEM. A representative photo of the Aβ staining of the right hemisphere of the cortex is shown **(A1)**. Average plaque size was decreased in the cortex in the right hemisphere compared to the left hemisphere (p<.030) in male stroke mice **(A2)**. Relative Aβ positive area was decreased in the right hemisphere compared to the left hemisphere (p<.004) in male stroke mice **(A3)**. Number of Aβ plaques per area (mm2) was decreased in the cortex in the right hemisphere compared to the left hemisphere (p<.038) in male stroke mice **(A4)**. A representative photo of the Aβ staining of the right hemisphere of the corpus callosum is shown **(B1)**. No significant differences between the left and right hemisphere were found in the corpus callosum regarding the average plaque size **(B2)** or relative Aβ positive area **(B3)**. Number of Aβ plaques per area (mm2) was increased in the corpus callosum in the right hemisphere compared to the left hemisphere (p<.032) in male sham mice **(B4)**. A representative photo of the Aβ staining of the right hemisphere of the basal ganglia is shown **(C1)**. No significant differences between the left and right hemispheres were found in the basal ganglia regarding average plaque size **(C2)**. Relative Aβ positive area was decreased in the basal ganglia in the right hemisphere compared to the left hemisphere (p<.042) in female stroke mice **(C3)**. No significant differences between the left and right hemisphere were found in the basal ganglia regarding the number of Aβ plaques per area (mm2) **(C4)**.
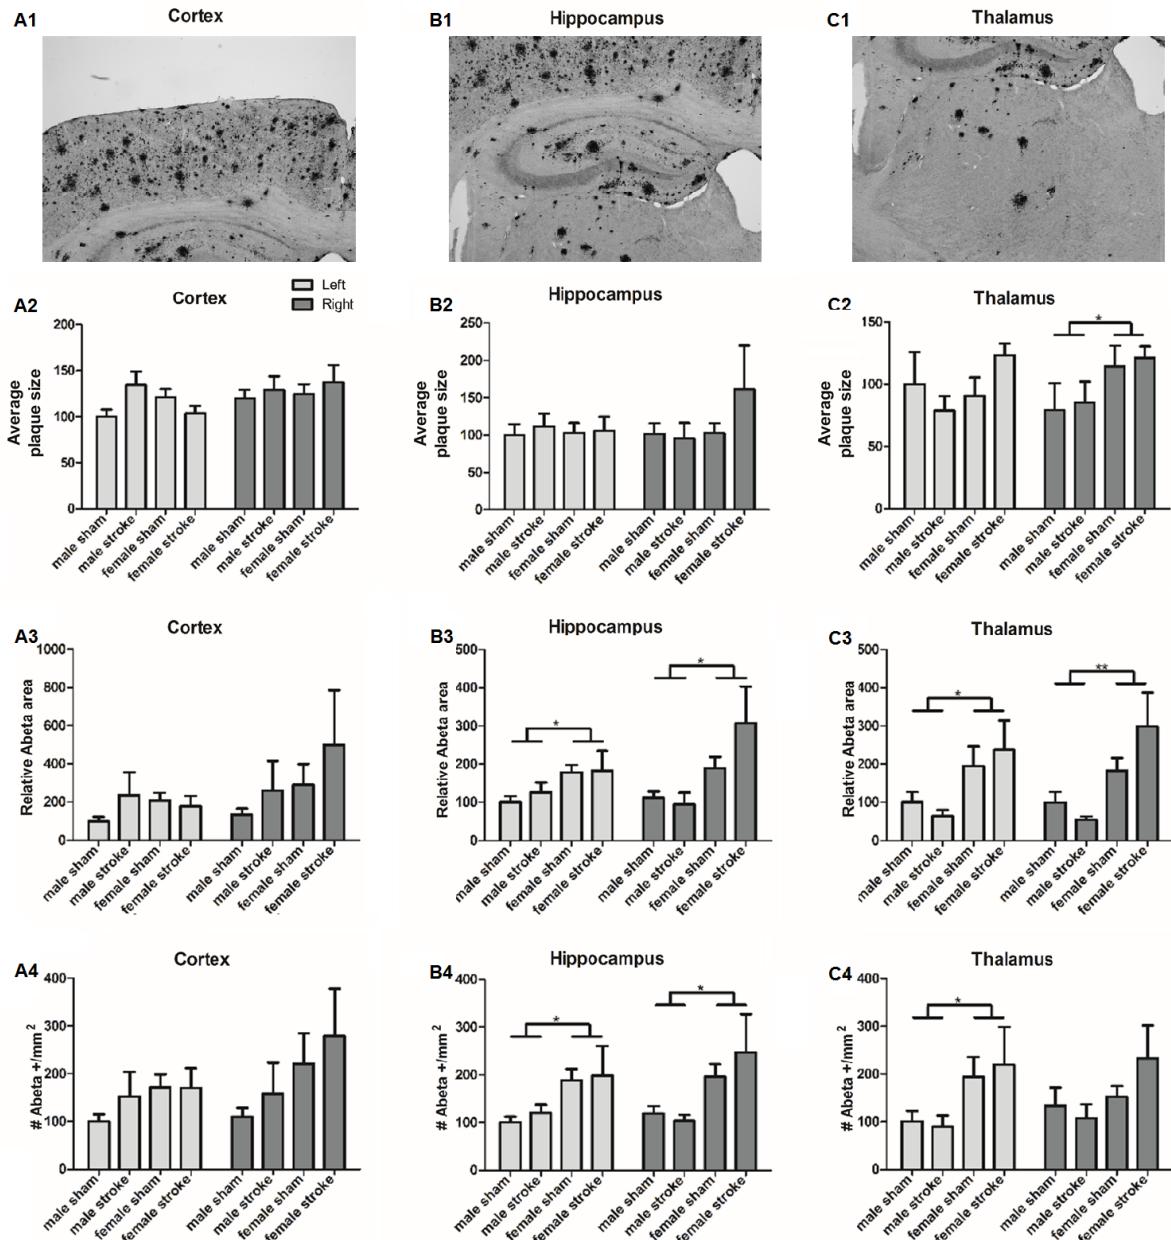
Figure 9. Average plaque size, relative Aβ positive area and number of Aβ plaques per area (mm2) in the cortex, hippocampus, and thalamus (bregma -1.94) of APP mice, sex and surgery effects. All data were normalized against the left male sham group and presented as mean ± SEM. Representative photo of the Aβ staining of the right hemisphere of the cortex is shown **(A1)**. No significant differences between sex or surgery were found in the cortex regarding average plaque size **(A2)**, relative Aβ positive area **(A3)** or number of Aβ plaques per area (mm2) **(A4)**. A representative photo of the Aβ staining of the right hemisphere of the hippocampus is shown **(B1)**. No significant differences between sex or surgery were found in the hippocampus regarding average plaque size **(B2)**. Relative Aβ positive area was increased in the hippocampus in female mice compared to male mice in both the left (p<.044) and right (p<.012) hemisphere **(B3)**. Number of Aβ plaques per area (mm2) was increased in the hippocampus in female mice compared to male mice in both left (p<.025) and right (p<.019) hemisphere **(B4)**. A representative photo of the Aβ staining of the right hemisphere of the thalamus is shown **(C1)**. Average plaque size was increased in the thalamus in female mice compared to male mice (p<.042) in the right hemisphere **(C2)**. Relative Aβ positive area was increased in the thalamus in female mice compared to male mice in both the left (p<.012) and right (p<.003) hemisphere **(C3)**. Number of Aβ plaques per area (mm2) was increased in the thalamus in female mice compared to male mice in the left hemisphere (p<.03) **(C4)**. No significant differences were found between surgery in any of the regions nor parameters.
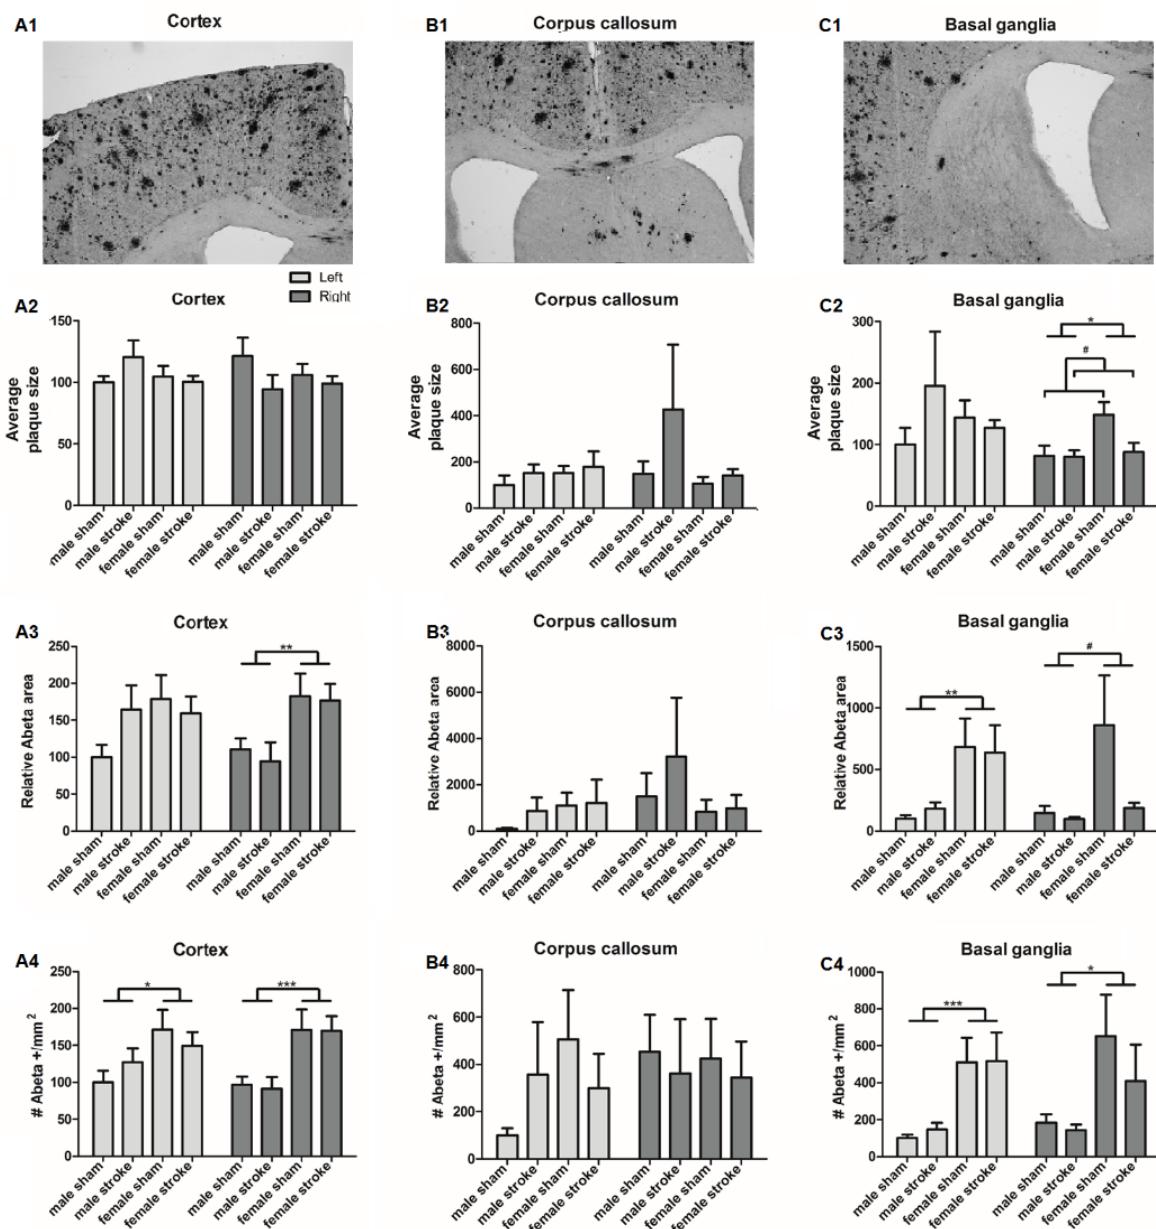
Figure 10. Average plaque size relative Aβ positive area and number of Aβ plaques per area (mm2) in the cortex, corpus callosum and basal ganglia (bregma 0.62) of APP mice, sex and surgery effects. All data were normalized against the left male sham group and presented as mean ± SEM. A representative photo of the Aβ staining of the right hemisphere of the cortex is shown **(A1)**. No significant differences between sex or surgery were found in the cortex regarding average plaque size **(A2)**. Relative Aβ positive area was increased in the cortex in female mice compared to male mice in the right (p<.003) hemisphere **(A3)**. Number of Aβ plaques per area (mm2) was increased in the cortex in female mice compared to male mice in both left (p<.031) and right (p<.001) hemisphere **(A4)**. A representative photo of the Aβ staining of the right hemisphere of the corpus callosum is shown **(B1)**. No significant differences between sex or surgery were found in the corpus callosum regarding average plaque size **(B2)**, relative Aβ positive area **(B3)** or number of Aβ plaques per area (mm2) **(B4)**. A representative photo of the Aβ staining of the right hemisphere of the basal ganglia is shown **(C1)**. Average plaque size was increased in the basal ganglia in female mice compared to male mice (p<.033) in the right hemisphere **(C2)**. Also, a trend was visible in the average plaque size in the basal ganglia, in which stroke animals have lower average plaque size in the right hemisphere than sham animals (p<.073). Relative Aβ positive area was increased in the basal ganglia in female mice compared to male mice in the left hemisphere (p<.004) **(C3)**. A trend was visible in the relative Aβ positive area (p<.069), in which the female mice showed a higher relative Aβ positive area in the right hemisphere than male mice. Number of Aβ plaques per area (mm2) was increased in the basal ganglia in female mice compared to male mice in both the left (p<.001) and right (p<.024) hemisphere **(C4)**. No other significant differences were found between surgery in any of the regions nor parameters.

right (p<.024) hemisphere compared to male mice. No other significant differences between surgeries were found regarding average plaque size, relative Aβ area or number of Aβ plaques per area (mm2). No interactions were found between sex and surgery.

Synaptic density in affected ischemic area

Male WT sham mice showed an increased in relative synaptophysin mRNA expression in the right hemisphere (Fig. 11), compared to the left hemisphere (p<.024). Relative synaptophysin mRNA expression was decreased in the right hemisphere compared to the left hemisphere in the female WT sham group (p<.026).

Data of synaptophysin mRNA expression were split according to genotypeor sex (Fig. 12). A significant decrease (p<.004) in relative synaptophysin expression was shown in the right hemisphere in female WT sham mice compared to male WT sham mice (Fig. 12A). Furthermore, relative synaptophysin expression was decreased (p<.024) in female APP mice compared to male APP mice in the left hemisphere (Fig. 12B). No significant differences were found regarding genotype or surgery in either male (Fig. 12C) or female mice (Fig. 12D).

Discussion

The overall aim of this longitudinal study was to elucidate the long-term effect of vascular damage, caused by stroke, on AD. This report focused on the Aβ burden and changes in synaptic density resulting from stroke and/or AD. Furthermore, the study used both male and female animals to further investigate the sex differences present in both AD and stroke. Many aspects about the effect of stroke and AD on each other remain unclear, especially regarding the long-term effects after the stroke. Therefore, it is very important to further investigate the effect of stroke and AD markers, such as the Aβ burden.

Sex differences Aβ burden

Female APP mice had a higher Aβ burden than male APP mice that is present in many brain areas, including the thalamus and basal ganglia. Furthermore, the protein levels of Aβ40 were also elevated in the anterior part of the brain, which became apparent by the ELISA analysis. These results were largely as expected, as an increased Aβ burden in the hippocampus in female APP/PS1 mice has been reported in AD mice studies before (Callahan et al., 2001; Wang et al., 2003). Wang et al. (2003) reported an increase of Aβ protein levels and Aβ burden in the hippocampus of 12-month-old female APP/PS1 mice. Callahan et al. (2001) demonstrated a similar effect in both enzyme levels and senile plaques in a different APP aging mouse model in both the hippocampus and neocortex. Wang et al. (2003) demonstrated elevated protein levels of Aβ40 and Aβ42, but not the Aβ42/Aβ40 ratio, in the hippocampus. The current results
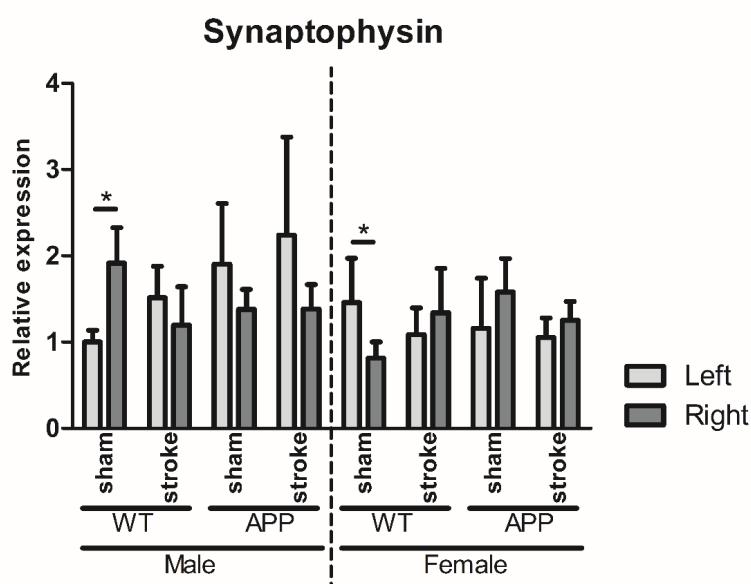

Figure 11. Relative expression of synaptophysin, L vs R effects. All data were normalized against the left male WT sham group and presented as mean ± SEM. Relative expression of synaptophysin was increased in the right hemisphere compared to the left hemisphere in the male WT sham group (p<.024) and decreased right compared to left in the female WT group (p<.026).

showed that the gender effect in Aβ is not limited to the hippocampal area, and possibly the cortex, but widely spread throughout different brain areas. No significant differences in sex regarding Aβ42 were found in the current study. Whether this is due to the mixture of tissues that are affected and unaffected by the stroke in this analysis or due to an insufficient power (alpha = 0.05, power = 0.56) caused by high variation in the experimental group, remains unclear. But overall, such increased Aβ concentration and Aβ burden in female APP mice was found in many brain areas and is not restricted to the hippocampus and cortex.

The effects of stroke on Aβ burden

The Aβ concentration and Aβ burden did not significantly differ between the stroke and sham groups, suggesting that the stroke operation had

no effect on the Aβ burden in APP mice. Notably, there was a trend found in the right hemisphere of the basal ganglia, which indicated a decrease in the averageplaque size in stroke mice compared to sham mice. This effect is most likely due to the high amount of atrophy in the area. The basal ganglia are mostly effected by the occlusion of the MCA, since they are directly supplied of oxygen and glucose by this artery (De Reuck, 1971). Supply of oxygen and glucose is severely decreased for a substantial amount of time, resulting in a high rate of cell death and atrophy in this area. The trend found in this region could therefore be associated with the atrophy and cellular apoptosis in this area.

With exception of this trend in the basal ganglia, the Aβ concentration and Aβ burden did not seem to be affected by the stroke and sham operations implying that stroke had no direct effect on the Aβ burden in the surrounding brain areas at bregma
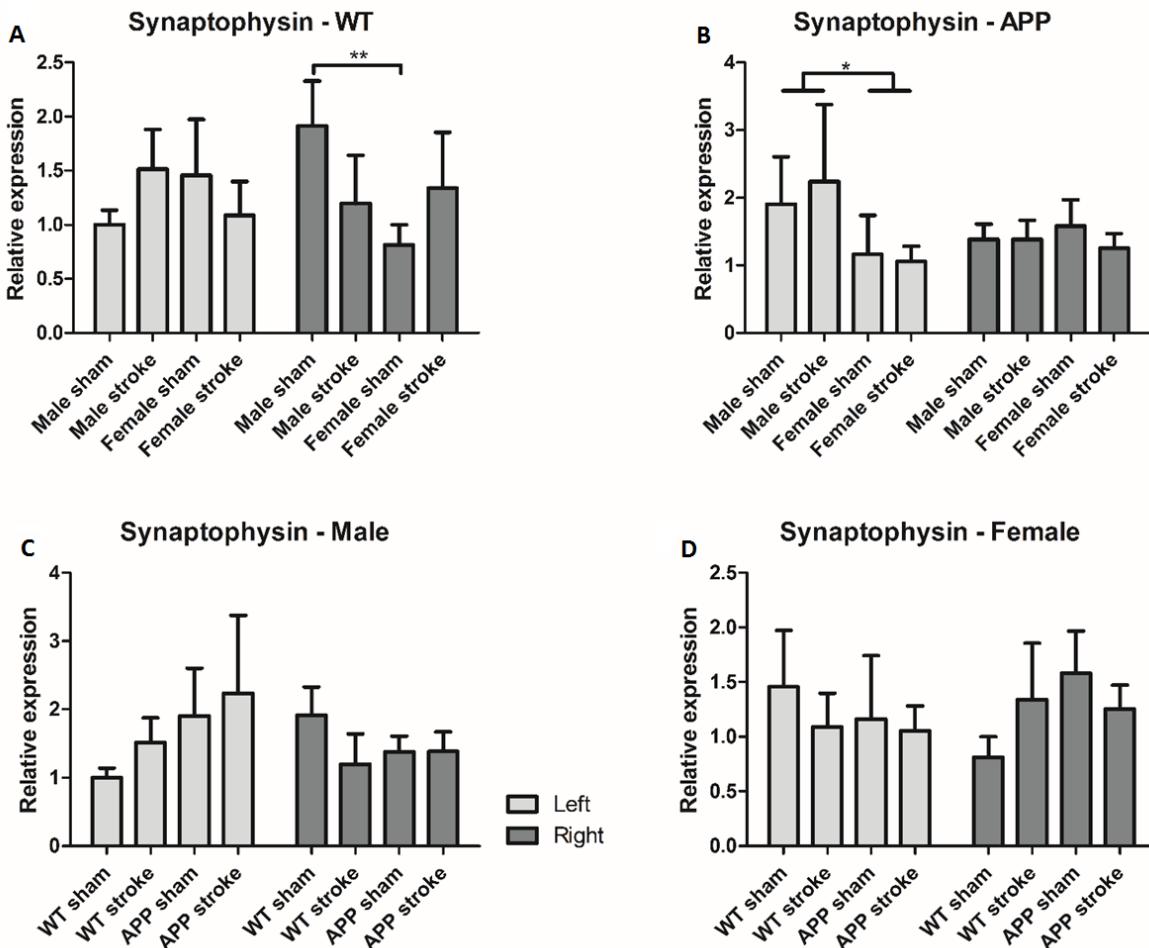

Figure 12. Relative expression of synaptophysin, sex and genotype effects. All data were normalized against the left male WT sham group and presented as mean ± SEM. Data were split on either genotype (A-B) or gender (C-D). Relative expression of synaptophysin was decreased (p<.004) in the right hemisphere in female mice compared to male mice regarding the WT sham group **(A)**. Relative expression of synaptophysin was also decreased (p<.024) in the left hemisphere in female APP mice compared to male APP mice **(B)**. No significant differences were found between genotype or surgery in either male **(C)** or female mice **(D)**.

of 0.62. But this conclusion cannot be drawn only from this comparison, since there were 2 controls present in this study. As mentioned before, next to the sham group serving as control, there was also a within-subject control, namely the left hemisphere of the mice. In each group only the right MCA was targeted, therefore the left hemisphere was not directly affected by the stroke operation. When the left and right hemispheres were compared with each other, stroke effects became visible. Male stroke mice showed a significant decrease in the Aβ burden, average size and number of plaques in the cortex of the affected hemisphere compared to the unaffected left hemisphere. This result was not expected, since studies on short-term effects reported an opposite effect of the stroke on the Aβ burden (Garcia-Alloza et al., 2011; Thiel, Cechetto, Heiss, Hachinski, & Whitehead, 2014) due to disruptions in amyloid clearance pathways the Aβ burden increased (Garcia-Alloza et al., 2011). However, mice in these studies were sacrificed days or weeks after stroke induction (Garcia-Alloza et al., 2011; Thiel et al., 2014). In the current study, mice had 8 months of recovery time. It could be that the disruption in the amyloid clearance pathway was not only minimized, but even increased due to ongoing inflammatory processes. These inflammatory processes are still present 8 months after stroke, indicated by the increased activity of reactive microglia in the male stroke animals compared to the male sham animals, that was not represented in the female animals. The expression patterns of the reactive microglia are very similar to the decrease in Aβ plaque patterns, indicating a correlation between the two. An interaction between microglial activity and Aβ plaques has been established before. However, the precise effect of microglial activity on Aβ clearance is very complex and can easily shift between a stimulating or inhibiting effect (Merlo, Spampinato, Caruso, & Sortino, 2020). Both positive and negative influences of microglia on Aβ clearance have been found, depending on the balance between pro- and anti-inflammatory activity of the microglia and the overall phenotype of the microglial cells (Merlo et al., 2020). In our case, the microglia possibly have increased the Aβ clearance in the affected hemisphere. This would explain why the effect was visible only in male stroke mice, since only the male mice showed an increase in microglial activity in the cortex of the affected hemisphere after stroke. To be sure about the effects of microglial activity on Aβ clearance a more in-depth analysis of the inflammatory response is necessary.

Areas not directly supplied by the MCA and therefore not directly affected by the stroke, showed different patterns of Aβ burden. Female stroke mice showed an increase in Aβ burden in the hippocampus of the affected hemisphere, coupled with a trend in the number of plaques in the hippocampus and a trend in the average plaque size in the cortex (bregma -1.94).

Taken together, the Aβ burden was decreased in regions directly supplied by the MCA, while being increased in other brain areas. It remains unclear how stroke affects these brain regions in different manner and what role microglial activity might play in the Aβ clearance in this case.

Effects sham operation on Aβ burden

Next to the observed stroke effects, the sham groups also showed some significant differences. The female sham group showed a significant increase in average plaque size in the thalamus and there was a trend towards a higher average plaque size in the affected hemisphere of the male sham mice. A possible explanation for the observed results might be that the sham operation actually had a significant effect on the animal and the affected hemisphere. This could be possible since the filament is shortly introduced during the sham operation, blocking blood flow for a couple of seconds. Furthermore, in both sham and stroke animals after the 'occlusion', the right common carotid artery is permanently tied off and it is not clear if this effects the brain. The effect and resulting stress of the sham operation might be large enough to cause differences between the two hemispheres of the sham animals. The animals and the vascular system did undergo significant levels of stress during the surgery. Stress has been reported to contribute to the formation of Aβ plaques before (Han et al., 2017; Huang et al., 2015). The stress caused by the short blockage of blood flow and the overall stress of the anesthesia and operation, could possibly contribute to the formation of the Aβ plaques (Eckenhoff et al., 2004; Marques & Lapa, 2018) and therefore the differences between the hemispheres in the sham animals.

Along the same line, the effect of the sham operation might also explain the lack of significance between the sham and stroke animals. If the sham operation truly affects the brain in a significant manner, similar to the effect of stroke, then the lack of significance between the sham and stroke groups would be logical. It is very important to consider all factors in these sham and within-subject control models before drawing any conclusions regarding the effects or lack of effects of stroke on the Aβ burden. Ideally, an extra control group

without operation would have been added to further determine the effects of stress and stroke on the brain. Unfortunately, due to the already large nature of the study, this was not possible.

Furthermore, neither a sham/stroke nor a within-subject control effect was found in the results of the ELISA analyses of the anterior part of the brain. This could be caused by the chosen methodology. In the ELISA analyses, no division was made between affected and unaffected tissue because the amount of tissue was very small per animal and it was not possible to make a clear division. The effect of the stroke operation might be masked because both ischemic regions and unaffected regions are taken together for the ELISA analyses. This would weaken the differences between the different hemisphere explaining the lack of significance in this analysis compared to the IHC.

Overall, the multiple control groups and methods for analyzing Aβ give multiple perspectives on the research question. However, the opposing results require critical evaluation of different controls and methods used.

Asymmetric synaptic density

Next to the Aβ burden, this study also investigated synaptic density via determining synaptophysin mRNA expression. In this qPCR analysis an asymmetry between the two hemispheres was found in the WT sham animals. Asymmetrical expression of genes has been reported before in many animal species and in humans (Vallortigara, Chiandetti, & Sovrano, 2011). Asymmetry of synaptophysin expression specifically, has also been reported in some species, e.g. developing chicken (in both via pre and post-synaptic markers) (Roy, Nag, Upadhyay, Mathur, & Jain, 2014) and aging chimpanzees (Sherwood et al., 2010). Although asymmetry of synaptophysin has not been found before in mice, they do show some brain asymmetry and for example, paw preference (Waters & Denenberg, 1994), indicating that the mouse brain has asymmetrical properties and hemispheres are not identical. Therefore, asymmetrical expression of synaptophysin could be a possibility. However, the observed asymmetry in synaptophysin has an opposite direction in female compared to male mice: female WT sham mice showed lower synaptophysin expression in their right hemisphere compared to their left hemisphere, while male WT sham mice showed higher synaptophysin expression in their right hemisphere compared to the left one. This opposing effect in different genders has not been

reported before. Furthermore, this asymmetry is not present in any of the other experimental groups, in which the mice suffered from stroke and/or AD pathology. The sham/stroke operation and the AD pathology could influence synaptic density to such an extent that the asymmetry would be hidden. However, this would mean that the effect of stroke on synaptic density would have an opposing effect on male compared to female mice, otherwise the effect could not be concealed in both genders of the other experimental groups. To determine the true nature of this possible asymmetry, it would be beneficial to include mRNA analyses of post synaptic density marker PSD-95. A pattern of PSD-95 expression similar to the synaptophysin expression would strengthen the observed asymmetry, while a lack of similarity could suggest a false positive in the synaptophysin analyses.

Effect of AD on synaptic density

In APP mice, the effect of gender was seen in the left unaffected hemisphere. Female APP mice showed decreased expression of synaptophysin in the left hemisphere compared to male APP mice, regardless of stroke or sham operation. This effect is most likely related to the AD pathology, since AD is also associated with synaptic loss, and the observed differences here were not present in the WT animals.

As mentioned before, female APP mice show a higher Aβ burden than their male counterparts, as shown both in the ELISA analyses and the immunohistochemical staining. This indicates a heavier AD pathology in female mice. Important to note, the Aβ burden is just a marker for AD, it is not a direct comparison for AD pathology. The exact contribution of the Aβ plaques to AD pathology remains unclear, but it is known that there are correlations between the amount of Aβ plaques and other factors associated with AD (Hooijmans et al., 2007). Hooijmans et al. demonstrated that Aβ plaques are related to decreased GLUT-1 and hippocampal atrophy in 18-month-old mice (Hooijmans et al., 2007) and both Aβ plaques and AD directly affect synaptic density negatively (Goulay, Romo, Hol, & Dijkhuizen, 2019; Dennis J Selkoe, 2002). Taken all together, a cautious link can be made to the high of Aβ burden and the overall weight of the AD pathology. Therefore, a more severe AD pathology indicated by a higher Aβ burden could explain the reduction in synaptic density in female mice, since AD negatively effects the synapses.

It would be beneficial to look at other AD markers such as vascularity and neuroinflammation

that may also differ between sexes to confirm the correlation between the synaptic density and heavier AD pathology. This could shed light on the question whetherthe observed effect is solely driven by Aβ burden or by the overall AD pathology.

Future perspective

Longitudinal studies on the combination of AD and stroke have never been done before. This study is the first step to elucidate the long-termeffects of stroke on AD pathology. Both the AD and stroke model have been widely used in research and have also been combined to determine short-term effects of stroke and AD on each other. There are some considerations to be made for the stroke model and its controls, both within and between subjects. Especially, differences of the two control models should be closely considered for their implications. Both controls can be influenced by the surgeries done, and it is vital to keep this in mind when comparing experimental groups.

To investigate the effect of sex, both male and female mice were used. In humans, females are protected via estrogen against the development, as well as the severe adverse effects of stroke (Reeves et al., 2008). Therefore, postmenopausal women are more prone to suffer a stroke. The current mouse model does not take this postmenopausal group into account. Future research could benefit from including groups that represent both pre- and postmenopausal women.

The mice used in this study also underwent multiple other tests, including behavioral studies and MRI analyses. It was not possible to include all data in this report. However, correlation analyses between behavioral studies, MRI data and postmortem analysis are crucial to elucidate the overall effect of stroke on AD. Markers for neuronal andsynaptic density could be correlated to previously performed cognitive test results. Also, neuroinflammatory factors should be corelated to Aβ burden. Originally, neurogenesis was planned to be further investigated by means of qPCR analysis of BDNF mRNA expression. However, due to the age of the mice and the low level of DCX-positive cells found in the hippocampus (other data from current study, not published yet), the concentrations were too low to be properly analyzed.

A crucial factor in AD pathology is the vascularity of the brain. GLUT-1 transporters in the blood vessels are vital for Aβ clearance in the brain (Ueno et al., 2014) and most AD patients show vascular pathologies postmortem (R. N. Kalaria et al., 2012).

A previous study already found a correlation between GLUT-1 and Aβ plaques in the mouse model that was used in the current study (Hooijmans et al., 2007). Postmortem analyses of GLUT-1 and overall architecture of the vascularity should be done and correlated to Aβ burden, MRI data and behavioral test.

Lastly, cell death should be considered as an influencing factor. By looking at the immunohistochemical staining, cell death and atrophy were seen in the basal ganglia, however, the full extent of this atrophy should be further analyzed by MRI. The degree of atrophy could also be correlated to the behavioral tests, especially focusing on memory tests and e.g. hippocampal atrophy. Overall, many factors are still to be considered in future research.

Conclusions

Due to the growing life expectancy, AD and stroke prevalence are expected to rise even more in the coming years (Association, 2019). Studies on the direct interaction but also the long-term effects of stroke and AD are important to improve treatment of the conditions and eventually reduce the prevalence of both. Sex is an important factor in many diseases and disorders, including AD and stroke. Therefore, including both sexes in research is vital for understanding the different presentation of the disorders.

The current longitudinal study replicated the sex differences found in Aβ plaques in the hippocampus and showed that this increased Aβ burden in females also extends to other brain areas such as the thalamus and basal ganglia. Eight months after stroke induction male mice showed a decrease in Aβ plaques in the areas directly affected by the stroke compared to the unaffected hemisphere, possibly due to increased microglial activity in the affected hemisphere. This effect of microglial activity between the two hemispheres was not present in the female mice. Differences between male and female mice were also present in synaptic density, specifically in asymmetrical expression between hemispheres. Furthermore, female APP mice showed an overall lower synaptic density than their male counterpart. This is most likely caused by the increased Aβ pathology in the female mice. Overall, the effects of stroke on AD are profound, but varying between the sexes. More research on AD pathological markers other than Aβ burden and synaptic density should shed more light on the long-term effects of stroke on AD.

References

Agüero-Torres, H., Kivipelto, M., & von Strauss, E. (2006). Rethinking the dementia diagnoses in a population-based study: what is Alzheimer's disease and what is vascular dementia? *Dementia and Geriatric Cognitive Disorders, 22*(3), 244-249.

Amarenco, P., Bogousslavsky, J., Caplan, L., Donnan, G., & Hennerici, M. (2009). Classification of stroke subtypes. *Cerebrovascular diseases, 27*(5), 493-501.

Appelros, P., Stegmayr, B., & Terént, A. (2009). Sex differences in stroke epidemiology: a systematic review. *Stroke, 40*(4), 1082-1090.

Association, A. s. (2019). 2019 Alzheimer's disease facts and figures. *Alzheimer's & Dementia, 15*(3), 321-387.

Azad, N. A., Al Bugami, M., & Loy-English, I. (2007). Gender differences in dementia risk factors. *Gender medicine, 4*(2), 120-129.

Bertrand, L., Dygert, L., & Toborek, M. (2017). Induction of ischemic stroke and ischemia-reperfusion in mice using the middle artery occlusion technique and visualization of infarct area. *JoVE (Journal of Visualized Experiments)*(120), e54805.

Cacace, R., Sleegers, K., & Van Broeckhoven, C. (2016). Molecular genetics of early-onset Alzheimer's disease revisited. *Alzheimer's & Dementia, 12*(6), 733-748.

Callahan, M. J., Lipinski, W. J., Bian, F., Durham, R. A., Pack, A., & Walker, L. C. (2001). Augmented senile plaque load in aged female β-amyloid precursor protein-transgenic mice. *The American journal of pathology, 158*(3), 1173-1177.

Chi, N.-F., Chien, L.-N., Ku, H.-L., Hu, C.-J., & Chiou, H.-Y. (2013). Alzheimer disease and risk of stroke: a population-based cohort study. *Neurology, 80*(8), 705-711.

de Bruijn, R. F., & Ikram, M. A. (2014). Cardiovascular risk factors and future risk of Alzheimer's disease. *BMC medicine, 12*(1), 130.

de la Torre, J. C. (2004). Is Alzheimer's disease a neurodegenerative or a vascular disorder? Data, dogma, and dialectics. *The Lancet Neurology, 3*(3), 184-190.

De Reuck, J. (1971). The human periventricular arterial blood supply and the anatomy of cerebral infarctions. *European neurology, 5*(6), 321-334.

DeKosky S.T., S. S. W., and Styren S.D. (1996). Structural correlates of cognition in dementia: quantification and assessment of synapse change. *Neurodegener J Neurodegener Disord Neuropprotection Neuroregeneration, 5*, 417-421.

Drouin, E., & Drouin, G. (2017). The first report of Alzheimer's disease. *The Lancet Neurology, 16*(9), 687.

Eckenhoff, R. G., Johansson, J. S., Wei, H., Carnini, A., Kang, B., Wei, W., . . . Eckenhoff, M. F. (2004). Inhaled anesthetic enhancement of amyloid-β oligomerization and cytotoxicity. *Anesthesiology: The Journal of the American Society of Anesthesiologists, 101*(3), 703-709.

Garcia-Alloza, M., Gregory, J., Kuchibhotla, K. V., Fine, S., Wei, Y., Ayata, C., . . . Bacskai, B. J. (2011). Cerebrovascular lesions induce transient β-amyloid deposition. *Brain, 134*(12), 3697-3707.

Goedert, M., & Spillantini, M. G. (2006). A century of Alzheimer's disease. *Science, 314*(5800), 777-781.

Goodwin, L. O., Splinter, E., Davis, T. L., Urban, R., He, H., Braun, R. E., . . . Ndukum, J. (2019). Large-scale discovery of mouse transgenic integration sites reveals frequent structural variation and insertional mutagenesis. *Genome research, 29*(3), 494-505.

Goulay, R., Romo, L. M., Hol, E. M., & Dijkhuizen, R. M. (2019). From Stroke to Dementia: a Comprehensive Review Exposing Tight Interactions Between Stroke and Amyloid-β Formation. *Translational stroke research*, 1-14.

Gouras, G. K., Tsai, J., Naslund, J., Vincent, B., Edgar, M., Checler, F., . . . Xu, H. (2000). Intraneuronal Aβ42 accumulation in human brain. *The American journal of pathology, 156*(1), 15-20.

Haast, R. A., Gustafson, D. R., & Kiliaan, A. J. (2012). Sex differences in stroke. *Journal of Cerebral Blood Flow & Metabolism, 32*(12), 2100-2107.

Han, B., Wang, J.-H., Geng, Y., Shen, L., Wang, H.-L., Wang, Y.-Y., & Wang, M.-W. (2017). Chronic stress contributes to cognitive dysfunction and hippocampal metabolic abnormalities in APP/PS1 mice. *Cellular Physiology and Biochemistry, 41*(5), 1766-1776.

Holtzman, D. M., Bales, K. R., Tenkova, T., Fagan, A. M., Parsadanian, M., Sartorius, L. J., . . . Wozniak, D. (2000). Apolipoprotein E isoform-dependent amyloid deposition and neuritic degeneration in a mouse model of Alzheimer's disease. *Proceedings of the National Academy of Sciences, 97*(6), 2892-2897.

Honig, L. S., Tang, M.-X., Albert, S., Costa, R., Luchsinger, J., Manly, J., . . . Mayeux, R. (2003). Stroke and the risk of Alzheimer disease. *Archives of neurology, 60*(12), 1707-1712.

Hooijmans, C. R., Graven, C., Dederen, P. J., Tanila, H., van Groen, T., & Kiliaan, A. J. (2007). Amyloid beta deposition is related to decreased glucose transporter-1 levels and hippocampal atrophy in brains of aged APP/PS1 mice. *Brain research, 1181*, 93-103.

Huang, H., Wang, L., Cao, M., Marshall, C., Gao, J., Xiao, N., . . . Xiao, M. (2015). Isolation housing exacerbates Alzheimer's disease-like pathophysiology in aged APP/PS1 mice. *International Journal of Neuropsychopharmacology, 18*(7), pyu116.

Jankowsky, J. L., Fadale, D. J., Anderson, J., Xu, G. M., Gonzales, V., Jenkins, N. A., . . . Wagner, S. L. (2003). Mutant presenilins specifically elevate the levels of the 42 residue β-amyloid peptide in vivo: evidence for augmentation of a 42-specific γ secretase. *Human molecular genetics, 13*(2), 159-170.

Jankowsky, J. L., Slunt, H. H., Ratovitski, T., Jenkins, N. A., Copeland, N. G., & Borchelt, D. R. (2001). Co-expression of multiple transgenes in mouse CNS: a comparison of strategies. *Biomolecular engineering, 17*(6), 157-165.

Jarrett, J. T., Berger, E. P., & Lansbury Jr, P. T. (1993). The carboxy terminus of the. beta. amyloid protein

is critical for the seeding of amyloid formation: Implications for the pathogenesis of Alzheimer's disease. *Biochemistry, 32*(18), 4693-4697.

Jendroska, K., Hoffmann, O., & Patt, S. (1997). Amyloid β peptide and precursor protein (APP) in mild and severe brain ischemia. *Annals of the New York Academy of Sciences, 826*(1), 401-405.

Jendroska, K., Poewe, W., Pluess, J., Iwerssen-Schmidt, H., Paulsen, J., Barthel, S., . . . DeArmond, S. (1995). Ischemic stress induces deposition of amyloid β immunoreactivity in human brain. *Acta neuropathologica, 90*(5), 461-466.

Kalaria, R. (2002). Similarities between Alzheimer's disease and vascular dementia. *Journal of the neurological sciences, 203*, 29-34.

Kalaria, R. N., Akinyemi, R., & Ihara, M. (2012). Does vascular pathology contribute to Alzheimer changes? *Journal of the neurological sciences, 322*(1-2), 141-147.

Kemppainen, S., Hämäläinen, E., Miettinen, P., Koistinaho, J., & Tanila, H. (2014). Behavioral and neuropathological consequences of transient global ischemia in APP/PS1 Alzheimer model mice. *Behavioural brain research, 275*, 15-26.

Kim, J.-S., Lee, K.-B., Roh, H., Ahn, M.-Y., & Hwang, H.-W. (2010). Gender differences in the functional recovery after acute stroke. *Journal of Clinical Neurology, 6*(4), 183-188.

Kivipelto, M., Helkala, E.-L., Laakso, M. P., Hänninen, T., Hallikainen, M., Alhainen, K., . . . Nissinen, A. (2001). Midlife vascular risk factors and Alzheimer's disease in later life: longitudinal, population based study. *Bmj, 322*(7300), 1447-1451.

Klyubin, I., Cullen, W. K., Hu, N.-W., & Rowan, M. J. (2012). Alzheimer's disease Aβ assemblies mediating rapid disruption of synaptic plasticity and memory. *Molecular brain, 5*(1), 25.

Launer, L., Andersen, K., Dewey, M., Letenneur, L., Ott, A., Amaducci, L., . . . Kragh-Sorensen, P. (1999). Rates and risk factors for dementia and Alzheimer's disease: results from EURODEM pooled analyses. *Neurology, 52*(1), 78-78.

Lewis, H., Beher, D., Cookson, N., Oakley, A., Piggott, M., Morris, C., . . . Kenny, R. (2006). Quantification of Alzheimer pathology in ageing and dementia: age-related accumulation of amyloid-β (42) peptide in vascular dementia. *Neuropathology and applied neurobiology, 32*(2), 103-118.

Leys, D., Erkinjuntti, T., Desmond, D. W., Schmidt, R., Englund, E., Pasquier, F., . . . Chabriat, H. (1999). Vascular dementia: the role of cerebral infarcts. *Alzheimer disease and associated disorders, 13*, S38-48.

Marchal, G., Beaudouin, V., Rioux, P., de la Sayette, V., Le Doze, F. o., Viader, F., . . . Baron, J. C. (1996). Prolonged persistence of substantial volumes of potentially viable brain tissue after stroke: a correlative PET-CT study with voxel-based data analysis. *Stroke, 27*(4), 599-606.

Marques, A. F. V. d. S., & Lapa, T. A. S. C. (2018). Anesthesia and Alzheimer disease–Current perceptions. *Revista brasileira de anestesiologia, 68*(2), 174-182.

McGowan, E., Pickford, F., Kim, J., Onstead, L., Eriksen, J., Yu, C., . . . Das, P. (2005). Aβ42 is essential for parenchymal and vascular amyloid deposition in mice. *Neuron, 47*(2), 191-199.

Merlo, S., Spampinato, S., Caruso, G., & Sortino, M. (2020). The ambiguous role of microglia in Aβ toxicity: Chances for therapeutic intervention. *Current neuropharmacology.*

Nakamura, Y., Takeda, M., Niigawa, H., Hariguchi, S., & Nishimura, T. (1992). Amyloid β-protein precursor deposition in rat hippocampus lesioned by ibotenic acid injection. *Neuroscience letters, 136*(1), 95-98.

Nelson, A. R., Sweeney, M. D., Sagare, A. P., & Zlokovic, B. V. (2016). Neurovascular dysfunction and neurodegeneration in dementia and Alzheimer's disease. *Biochimica et Biophysica Acta (BBA)-Molecular Basis of Disease, 1862*(5), 887-900.

Ordónez-Gutiérrez, L., Antón, M., & Wandosell, F. (2015). Peripheral amyloid levels present gender differences associated with aging in AβPP/PS1 mice. *Journal of Alzheimer's Disease, 44*(4), 1063-1068.

Radde, R., Bolmont, T., Kaeser, S. A., Coomaraswamy, J., Lindau, D., Stoltze, L., . . . Gengler, S. (2006). Aβ42-driven cerebral amyloidosis in transgenic mice reveals early and robust pathology. *EMBO reports, 7*(9), 940-946.

Ramanathan, A., Nelson, A. R., Sagare, A. P., & Zlokovic, B. V. (2015). Impaired vascular-mediated clearance of brain amyloid beta in Alzheimer's disease: the role, regulation and restoration of LRP1. *Frontiers in aging neuroscience, 7*, 136.

Reeves, M. J., Bushnell, C. D., Howard, G., Gargano, J. W., Duncan, P. W., Lynch, G., . . . Lisabeth, L. (2008). Sex differences in stroke: epidemiology, clinical presentation, medical care, and outcomes. *The Lancet Neurology, 7*(10), 915-926.

Robinson, J. L., Molina-Porcel, L., Corrada, M. M., Raible, K., Lee, E. B., Lee, V. M.-Y., . . . Trojanowski, J. Q. (2014). Perforant path synaptic loss correlates with cognitive impairment and Alzheimer's disease in the oldest-old. *Brain, 137*(9), 2578-2587.

Roger, V. L., Go, A. S., Lloyd-Jones, D. M., Adams, R. J., Berry, J. D., Brown, T. M., . . . Ford, E. S. (2011). Heart disease and stroke statistics—2011 update: a report from the American Heart Association. *Circulation, 123*(4), e18-e209.

Roy, S., Nag, T. C., Upadhyay, A. D., Mathur, R., & Jain, S. (2014). Prenatal music stimulation facilitates the postnatal functional development of the auditory as well as visual system in chicks (Gallus domesticus). *Journal of biosciences, 39*(1), 107-117.

Sahathevan, R., Brodtmann, A., & Donnan, G. A. (2012). Dementia, stroke, and vascular risk factors; a review. *International Journal of Stroke, 7*(1), 61-73.

Selkoe, D. J. (2002). Alzheimer's disease is a synaptic failure. *Science, 298*(5594), 789-791.

Selkoe, D. J., & Schenk, D. (2003). Alzheimer's disease: molecular understanding predicts amyloid-based

therapeutics. *Annu Rev Pharmacol Toxicol, 43*, 545-584. doi:10.1146/annurev.pharmtox.43.100901.140248

Shankar, G. M., & Walsh, D. M. (2009). Alzheimer's disease: synaptic dysfunction and Aβ. *Molecular neurodegeneration, 4*(1), 48.

Sherwood, C. C., Duka, T., Stimpson, C. D., Schenker, N. M., Garrison, A. R., Schapiro, S. J., . . . Hof, P. R. (2010). Neocortical synaptophysin asymmetry and behavioral lateralization in chimpanzees (Pan troglodytes). *European Journal of Neuroscience, 31*(8), 1456-1464.

Simões, A. E., Pereira, D. M., Amaral, J. D., Nunes, A. F., Gomes, S. E., Rodrigues, P. M., . . . Thibodeau, S. N. (2013). Efficient recovery of proteins from multiple source samples after trizol® or trizol® LS RNA extraction and long-term storage. *BMC genomics, 14*(1), 181.

Skoog, I., Nilsson, L., Persson, G., Lernfelt, B., Landahl, S., Palmertz, B., . . . Svanborg, A. (1996). 15-year longitudinal study of blood pressure and dementia. *The Lancet, 347*(9009), 1141-1145.

Sun, X., He, G., Qing, H., Zhou, W., Dobie, F., Cai, F., . . . Song, W. (2006). Hypoxia facilitates Alzheimer's disease pathogenesis by up-regulating BACE1 gene expression. *Proceedings of the National Academy of Sciences, 103*(49), 18727-18732.

Sweeney, M. D., Sagare, A. P., & Zlokovic, B. V. (2015). Cerebrospinal fluid biomarkers of neurovascular dysfunction in mild dementia and Alzheimer's disease. *Journal of Cerebral Blood Flow & Metabolism, 35*(7), 1055-1068.

Thiel, A., Cechetto, D. F., Heiss, W.-D., Hachinski, V., & Whitehead, S. N. (2014). Amyloid burden, neuroinflammation, and links to cognitive decline after ischemic stroke. *Stroke, 45*(9), 2825-2829.

Ueno, M., Chiba, Y., Matsumoto, K., Nakagawa, T., & Miyanaka, H. (2014). Clearance of beta-amyloid in the brain. *Current medicinal chemistry, 21*(35), 4085-4090.

Vallortigara, G., Chiandetti, C., & Sovrano, V. A. (2011). Brain asymmetry (animal). *Wiley Interdisciplinary Reviews: Cognitive Science, 2*(2), 146-157.

Wang, J., Tanila, H., Puoliväli, J., Kadish, I., & van Groen, T. (2003). Gender differences in the amount and deposition of amyloidβ in APPswe and PS1 double transgenic mice. *Neurobiology of disease, 14*(3), 318-327.

Waters, N. S., & Denenberg, V. H. (1994). Analysis of two measures of paw preference in a large population of inbred mice. *Behavioural Brain Research, 63*(2), 195-204.

White, L., Petrovitch, H., Hardman, J., Nelson, J., Davis, D. G., Ross, G. W., . . . Markesbery, W. R. (2002). Cerebrovascular pathology and dementia in autopsied Honolulu-Asia Aging Study participants. *Annals of the New York Academy of Sciences, 977*(1), 9-23.

Wuwongse, S., Cheng, S. S.-Y., Wong, G. T.-H., Hung, C. H.-L., Zhang, N. Q., Ho, Y.-S., . . . Chang, R. C.-C. (2013). Effects of corticosterone and amyloid-beta on proteins essential for synaptic function: implications for depression and Alzheimer's disease. *Biochimica et Biophysica Acta (BBA)-Molecular Basis of Disease, 1832*(12), 2245-2256.

Zeng, Y., Zhang, J., Zhu, Y., Zhang, J., Shen, H., Lu, J., . . . Zhou, M. (2015). Tripchlorolide improves cognitive deficits by reducing amyloid β and upregulating synapse-related proteins in a transgenic model of Alzheimer's Disease. *Journal of neurochemistry, 133*(1), 38-52.

Supplementary information

Supplementary protocol 1: RNA isolation

The protocol is done in the fume hood. Always wear gloves and clean them regularly with 70% ethanol. Make sure that there are enough sterile pipette tips and autoclaved 1,5 ml Eppendorf tubes, and enough DEPC-treated milliQ, TRIzol, chloroform, and 2-propanol and 75% ethanol. All pipette tips and tubes that have come into contact with TRIzol have to be collected in a glass container and disposed of in the chemical waste container. The protocol is used to isolate RNA from frozen (-80°) tissue, which is placed in a 2 ml Eppendorf tube. The tubes with MQ and ethanol used for cleaning the homogenizer and the 2-propanol and ethanol supernatants from steps 24 and 28 need to be collected in a 50 ml tube and disposed of in the chemical waste container.

1. Switch on the centrifuge, set it to 4°. Get a bucket of ice. Place the tissue homogeniser in the fume hood. Switch on the thermomixer and set it to 60°.

For each 4 samples you have, fill one marked test tube half full with 70% ethanol and one marked test tube half full with MQ (DEPC-treated).

2. Take the tissue samples from the freezer and put them in the ice bucket
3. In the fume hood, add 1 ml of TRIzol to each sample
4. Dip the tip of the homogenizer into the test tube with 70% ethanol
5. Clean the tip of the homogenizer with a tissue paper
6. Put the homogenizer into the tube with TRIzol and your sample, push the tissue down, switch on the homogenizer for 10 seconds, move the tube up and down during the homogenizing
7. Close the tube and put it back in the ice
8. Put the homogenizer into the test tube with MQ and switch on the machine for about 3 seconds

Clean the tip of the homogenizer with a tissue paper

Dip the tip of the homogenizer into the test tube with 70% EtOH

Clean the tip of the homogenizer with a tissue paper

9. Homogenize the next sample (steps 6-8)
10. Switch to new test tubes with MQ and EtOH, and tissue paper, after each 4 samples
11. Store the samples at RT for 5 minutes

12. Add 200 μl of chloroform to each sample
13. Shake the tubes vigorously for 15 seconds
14. Store the samples at RT for 5 minutes
15. Centrifuge at 12.000 g at 4° for 10 minutes
16. Make a new labeled, sterile 1,5 mlEppendorf tube for each sample
17. Carefully put the centrifuged tubesback into the ice
18. In the fume hood, pipette the top, clear layer (500 μl) into the new tube. Be careful not to touch the intermediate layer with your pipette tip.

It's more important that what you pipette off is clean than that you get the whole 500 μl

19. Add 500 μl of 2-propanol to eachtube
20. Shake the tubes vigorously for 15 seconds
21. Store the samples at RT for 15 minutes
22. Centrifuge at 12.000 g at 4° for 10 minutes
23. Carefully put the tubes back intothe ice. Most of the times, a white pellet is visible at the bottom of the tube
24. Set the pipette to 600 μl, and carefully pipette off the supernatant in two steps. Do not touch the pellet or the place where the pellet should be. Discard the supernatant into the 50 ml tube
25. Add 1 ml of 75% EtOH to each tube. Flick the tubes
26. Centrifuge at 7.500 g at 4° for 5 minutes
27. Carefully put the tubes back into the ice
28. Set the pipette to 600 μl, and carefully pipette off the supernatant in two steps. Do not touch the pellet or the place where the pellet should be. Discard the supernatant into the 50 ml tube. Set another pipette to 50 μl and pipette out the last remaining ethanol
29. Take the tubes out of the ice, open the tubes and let the RNA dry for 10 minutes
30. Add 25 μl of MQ (DEPC-treated) to each sample. Pipette up and down a few times
31. Put the tubes in the thermomixer (no shaking) for 10 minutes
32. Flick the tubes with your fingers and spin them down
33. Store the tubes at -80°C or continue with the protocol for measuring the RNA concentration using the NanoDrop

Supplementary protocol 2: Measuring RNA/DNA with the Nanodrop 2000

Using this protocol, you can measure the concentration and purity of your RNA or DNA samples, using the NanoDrop spectrophotometer.

Many proteins absorb at 280 (aromatic rings

absorb this wavelength) and DNA/RNA absorbs at 260. 260/280 ratio should be 1.80-2.00, 260/230 1.80-2.20. When working with RNA isolated with TRIzol, the 260/230 values are going to be lower (0.3-0.8) because of the high salt concentrations.

1. Switch on the computer and screen. The NanoDrop machine is always on. When working with RNA, get a bucket of ice, and wear clean gloves throughout the protocol

2. Thaw your samples (if needed)

3. Flick your samples using your fingers, spin them down

4. When working with RNA, put your samples on ice

5. Run the NanoDrop 2000 program on the computer

6. Select 'Nucleic acids' on the menu screen

7. Make sure the arm of the NanoDrop is down, click OK

8. In the top right corner of the screen, select if you want to measure DNA or RNA

9. Pipette 1 µl of the solution your DNA or RNA is dissolved in (usually 1x TE or DEPC-MQ) onto the pedestal, put the arm down, and click the 'Blank' button in the program

10. Put the arm up when the blank measurement is done and, using a tissue paper, gently wipe off the liquid from the pedestal and the arm

11. Pipette 1 µl of your DNA/RNA sample onto the pedestal, put the arm down, fill in your sample number in the field in the top right, and click the 'Measure' button in the program

12. Put the arm up when the measurement is done and, using a tissue paper, gently wipe off the liquid from the pedestal and the arm

13. When working with RNA, measure your sample again (duplo) with a new 1 µl

14. Repeat steps 11-13 for all your samples

15. After the first sample, you get a pop-up asking you where you want to save your data

16. After the last sample is measured, click the 'Reports' button in the lower left corner of the screen. Click 'Export' to export your data to an .xml file you can use later on

17. Close the NanoDrop program and switch off the computer and screen

18. Store the tubes at -80°C or continue with the next protocol

Supplementary protocol 3: DNA treatment of RNA

Get a bucket of ice. Thaw the RNA samples (if needed). Flick them and spin them down and put them on ice. Get the DNase, DNase buffer, and DNase stop solution from the -20° freezer. Set the Thermomixer to 37°C (no shaking). All steps are done on ice, unless stated otherwise.

1. Take an autoclaved Eppendorf vial (1,5 ml) for each RNA sample and label it.

2. Add up to 1 µg of RNA to the vial, add MQ to 8 µl.

3. Add 1 µl of DNase buffer to each vial.

4. Add 1 µl of DNase enzyme to each vial.

5. Flick the tubes, spin them down.

6. Put the vials in the 37° Thermomixer for 30 minutes.

7. Take the tubes out of the Thermomixer, put them back on ice. Set the Thermomixer to 65°.

8. Add 1 µl of Stop solution to each vial.

9. Flick the tubes, spin them down.

10. Put the vials in the 65° Thermomixer for 10 minutes.

11. Take the tubes out of the Thermomixer, put them back on ice.

12. Measure the RNA concentration with the NanoDrop and store the RNA in the -80° freezer.

Supplementary protocol 4: cDNA synthesis

This protocol can be used to generate cDNA from RNA. The main idea is that one absolute amount of RNA is added to the cDNA synthesis reaction and that this amount is the same for all the samples. The preparation that needs to be done for this protocol is that for all your samples you have to calculate how many µl of RNA you have to add to the reaction mix to get the absolute amount. The equation for calculating how many µl you have to add is

(absolute amount you want) / (RNA concentration of the sample)

The reaction mix per sample consists of:

- x µl RNA
- (7.5 - x) µl of RNase-free water
- 2 µl of 5x reverse transcription reaction mix
- 0.5 µl of reverse transcriptase enzyme

So, the maximum volume of RNA that can be added is 7.5 µl. Calculate the absolute amount of RNA you want to add based on the sample in your collection with the lowest RNA

concentration.

The maximum amount of RNA that can be added to this 10 µl reaction is 500 ng.

1. Calculate how many µl of RNA you need to pipette to get the absolute amount of RNA (400 ng), and calculate how many µl of RNase-freewater you need to add to that to get to 7.5 µl end volume.

2. Get a bucket of ice

3. Thaw your RNA samples and the 5x reverse transcription reaction mix, flick them with your fingers and put them in the ice

4. Put the reverse transcriptase enzyme directly on ice, it is not frozen

5. Prepare 8-tubestrips

6. Add the RNase-free water to the strips, then the RNA. The total volume in each tube should be 7.5 µl

7. Add 2 µl of 5x reverse transcription reaction mix to each tube, then add 0.5 µl reverse transcriptase enzyme

8. Close the tubestrips, flick them with your finger and spin them down

9. Put the strips in the PCRmachine and run the ISCRIPT program

10. Spin the strips down and dilute the cDNA 1:10 using MQ. Tap the strips to mix

11. Spin the strips down, store them at 4º (short term) or -20º (long term)

Supplementary protocol 5: qPCR

Use this protocol to prepare a QPCR plate, configure the ABI Prism 7900HT or StepOne Plus machine, and run the QPCR protocol. Read up on the background & principles of QPCR using the attached qRT-PCR_Basics PDF file. We use the SYBR Green method here. To make the QPCR master mix (see below), use the Bio-Rad iTaq SYBR Green mix. Use the MicroAmp Optical 96-Well plates. To make the QPCR master mix (see below), use theBio-Rad iTaq SYBR Green mix. Use the MicroAmp Fast 96-Well plates.

First, design the layout of the QPCR plate. All samples must be run in triplicates, so select three adjacent wells for each sample. Three wells are used for the no template control (NTC), a reaction in which MQ will be added instead of a sample. This NTC is used to check whether no cDNA contamination is present in the QPCR master mix.

If you have multiple primer sets in one plate, make a master mix for each pair. Calculate how much master mix you need to make ((number of samples + 1 NTC) x 3). Because of pipette errors and dead volumes of pipettes, make some extra master mix. When you have 24 samples, add 3 samples.

For each well you need these components:
• 5 µl SYBR Green mix (Bio-Rad iTaq)
• 0.4 µl forward primer
• 0.4 µl reverse primer
• 2.2 µl MQ
This is 8 µl in total.

1. Get a bucket of ice.

2. Made the mix in reverse order, so add the SYBR Green mix last. When pipetting this, pipette a few times up and down to mix everything. Master mixes cannot be vortexed. Put the master mixes on ice. Defrost the cDNA (if needed), flick the strips and spin them down. They can be kept at room temperature.

3. Get a plate, put it in a black plate holder (base plate) and put it on ice.

4. Use a manual pipette set to 8.0 µl or the electric repeater pipette set to 8.0 µl aliquots to pipette the master mix into the plate.

5. Use the manual P2 pipette to add 2 µl of cDNA to the correct wells. Add 2 µl of MQ to the NTC wells.

6. Take the plate out of the ice.

7. Affix a cover sticker to the plate, using the gray applicator. Press down the sticker on all sides and corners. Use your fingers to press down the sticker on each well. Spin down the plate using the plate spinner.

8. Take your plate to the StepOne Plus. Switch on the machine and the computer

9. Designing the plate is done on your own zero client. Then you save this and take it to the QPCR machine to run it

10. Start StepOne Software v2.3

11. Click OK

12. Click File -> New Experiment -> Advanced Setup

13. Fill in Experiment Name and Username

14. Click SYBR Green Reagents button & Standard (~ 2 hours to complete a run) buttons

15. Click the Plate Setup button in the top left of the screen

16. Fill in the target name (the name of your gene of interest). If there are more than 1 on your plate, click Add New Target

17. Fill in the names of the samples. Click Add New Sample button for each sample. Also add the NTC as a sample

18. Click the Assign Targets and Samples button

19. Design the plate layout according to your

scheme, selecting the appropriate Target and Sample for each well. When selecting the NTC, click the N button under Task. All other samples are set to U.

20. Click the Run Method button in the top left of the screen

21. Change the first step to 30 seconds 95° instead of 10 minutes. First change the seconds to 30, then change the minutes to00

22. Change Reaction Volume Per Well to 10

23. In the top left of the screen, click the save icon and save your file on your USB stick

24. Take the plate and the USB stick to the StepOnePlus machine at Internal Medicine

25. Switch on the computer if necessary

26. Put your USB stick in the computer, copy your .eds file to the Anatomy folder in the Users folder (located on the desktop)

27. Start the StepOne Software

28. Open the tray of the StepOnePlus machine, put your plate in and close the tray

29. In the top left of the screen, click the Open button and find your file in the Users -> Anatomy folder. Open it

30. Click the START RUN button in the top right corner of the screen

31. After about 2 hours the run is done

32. Click the Save button

33. Open the tray of the machine and remove the plate. Close the tray

34. Copy the .eds file to your USB stick

35. Close the StepOne software

36. Open the .eds file on your zero client

37. Uncheck the Auto box next to Threshold in the bottom om the screen

38. Fill in 0.2 and press Enter

39. Click the Export button in the top of the screen

40. Click the Customize Export tab

41. Uncheck everything except Well, Sample Name, Target Name, Task, Ct, Ct Threshold, Tm1, Tm2, Tm3

42. Click the Export Properties tab

43. Change the Export File Location to your folder

44. Check the Open file(s) when export is complete box

45. Click the Start Export button

46. Do data analyses in Excel

Supplementary protocol 6: gDNA isolation with Trizol

Step 1 to 18 from the RNA isolation protocol

(Labguru, Bram Geenen),

1. Pipette off the RNA phase

2. Add 500 µl 100% Ethanol to the phenol and interphase of the sample

3. Carefully shake until you see the DNA appear

4. Centrifuge at 7.500g for 5 minutes at RT

5. Discard supernatant (with pipette)

6. Add 700 µl Sodium citrate in 10% Ethanol, pH 8,5

7. Wash pellet by carefully moving the tube upside down a few times

8. Centrifuge at 7.500g for 5 minutes at RT

9. Discard supernatant (with pipette)

10. Add 700 µl 70% Ethanol

11. Wash pellet by carefully moving the tube upside down a few times

12. Centrifuge at 7.500g for 5 minutes at RT

13. Discard supernatant (with pipette)

14. Let the pellet dry for ca. 10 minutes

15. Add 500 µl MQ

16. Vortex and spin tubes down

17. Measure gDNA with Nanodrop with the 'Measuring RNA/DNA with the NanoDrop 2000' protocol (Labguru, Bram Geenen)

18. Samples are stored at -20°C until further analysis will be performed

Note: concentration of the DNA has to be below 100 ng/µl for successful genotyping. Dilute if needed.

Sodium citrate:

1. Dissolve 1.455 g sodium citrate (2 H20) in 5 ml 100% ethanol

2. Add MQ to a total of 40 ml

3. Adjust pH to 8,5

4. Add MQ to a total of 50 ml

Supplementary protocol 7: running a agarose gel

1. Set up the gelelectrophoresis system

2. Combine 100 ml of 1x TBE and 1.0 gram of agarose powder in a medium flask.

3. Mix gently.

4. Microwave on HIGH until the mixture starts to boil. Gently mix the contents. Beware of delayed boiling! Repeat another 2 times until agarose is completely dissolved.

5. Allow the contents in the flask to cool until it stops steaming.

6. Avoid introducing any additional air bubbles.

7. Add 10 μl of SYBR Safe DNA stain to the flask using a pipette.

8. Mix gently and carefully pour gel.

9. Remove any air bubbles in gel lanes by gently pushing the bubbles off to the sides with pipette tip.

10. Allow gel to harden at RT for 20 minutes

11. Put gel and TBE buffer in agarose system.

12. Add 5 μl of loading buffer to each sample.

13. Load 5 μl of TrackIt 100 bp ladder into the first slot.

14. Carefully inject 5 μl of each sample into the slots in the gel.

15. Run gel at 100V. for 30 minutes

16. Take a picture of the gel using the GelDoc

Solutions:

• Invitrogen SYBR Safe DNA Gel Stain (400 μl) Life Technologies #S33102

• Invitrogen TrackIt 100 bp ladder (500 μl) Life Technologies #10488-058

• Invitrogen UltraPure™ TBE Buffer, 10X (1 L) Life Technologies #15581-044

• Invitrogen UltraPure™ Agarose (100 gram) Life Technologies #16500-100

Supplementary protocol 8: Protein isolation

Day 1 RNA isolation

1. Step 1 to 18 from the RNA isolation protocol (Labguru, Bram Geenen)

Day 2 DNA en Protein isolation o/n 4ºC

1. Pipette of the RNA phase

2. Add 500ul (1:1) 100% Ethanol to precipitate DNA and mix

3. Incubate for 5 minutes by 4°C

4. Centrifuge 7,500xg at 4°C for 5 min

5. Supernatent 1ml (phenol/ethanol fase) to a 2ml eppendorf tube for protein isolation

6. (Pellet can be used for gDNA isolation)

7. Add 1 mL (1:1) of isopropanol to the phenol-ethanol

8. Mix and Incubate for 10 minutes at RT

9. Centrifuge for 10 minutes at $12,000 \times g$ at 4°C to pellet the proteins

10. Discard the supernatant

11. Wash the pellet in 1.5 mL of wash solution (0.3 M guanidine hydrochloride in 95% ethanol)

12. Incubate for 20 minutes ad RT. (with shaking)

13. Centrifuge for 5 minutes at $7500 \times g$ at 4°C

14. Discard the supernatant. By poring

15. Repeat step 6 step 9 twice (wash three times in total)

16. Air dry the protein pellet for 5–10 minutes

17. Add a tablet Complete protease inhibitor to 10 ml of 5M Guanidine buffer

18. Resuspend the pellet in 500 μL of 5M guanidine hydrochloride in Tris-HCl, pH8) by pipetting up and down

19. Vortex to complete suspension

20. Centrifuge for 10 minutes at $10,000 \times g$ at 4°C to remove insoluble materials

21. Transfer the supernatant to a new 1.5 ml tube (Guanidine Soluble Fraction)

22. Samples are stored at -80ºC until the ELISA will be performed

0.3M Guanidine buffer:

Dissolve 7.164 gr guanidine HCl in 250ml 95ml Ethanol.
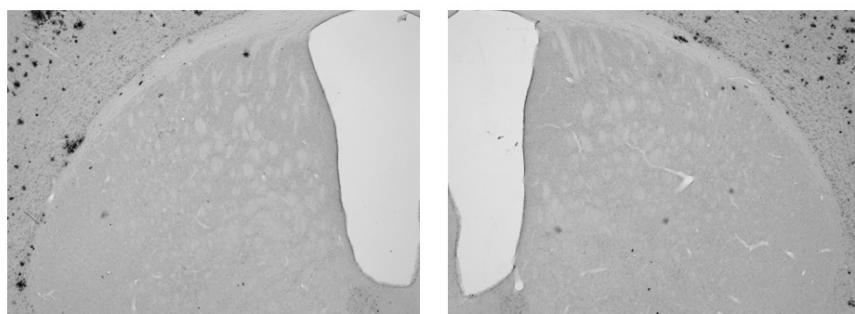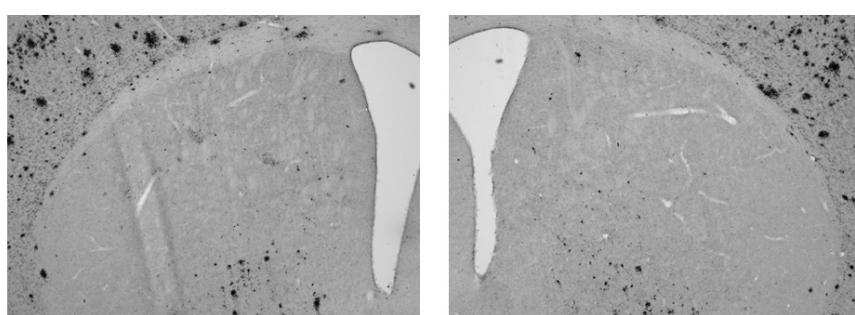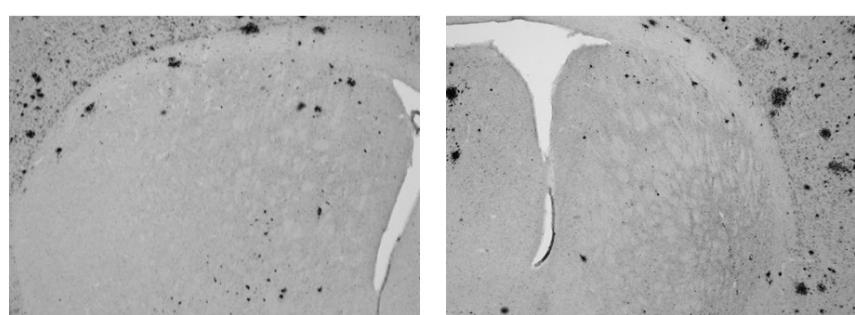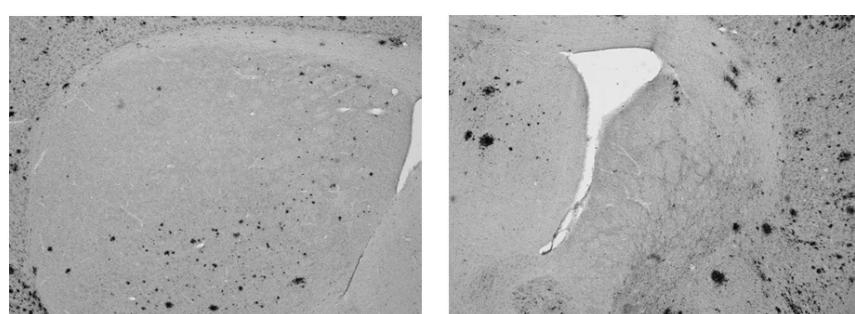
5M Guanidine buffer:

Dissolve 119.4 gr guanidine HCl and 1.5 gr Trizma base in 200 ml MQ

Adjust pH to 8.0

Add MQ to a total of 250 ml (expiration date 2 years)
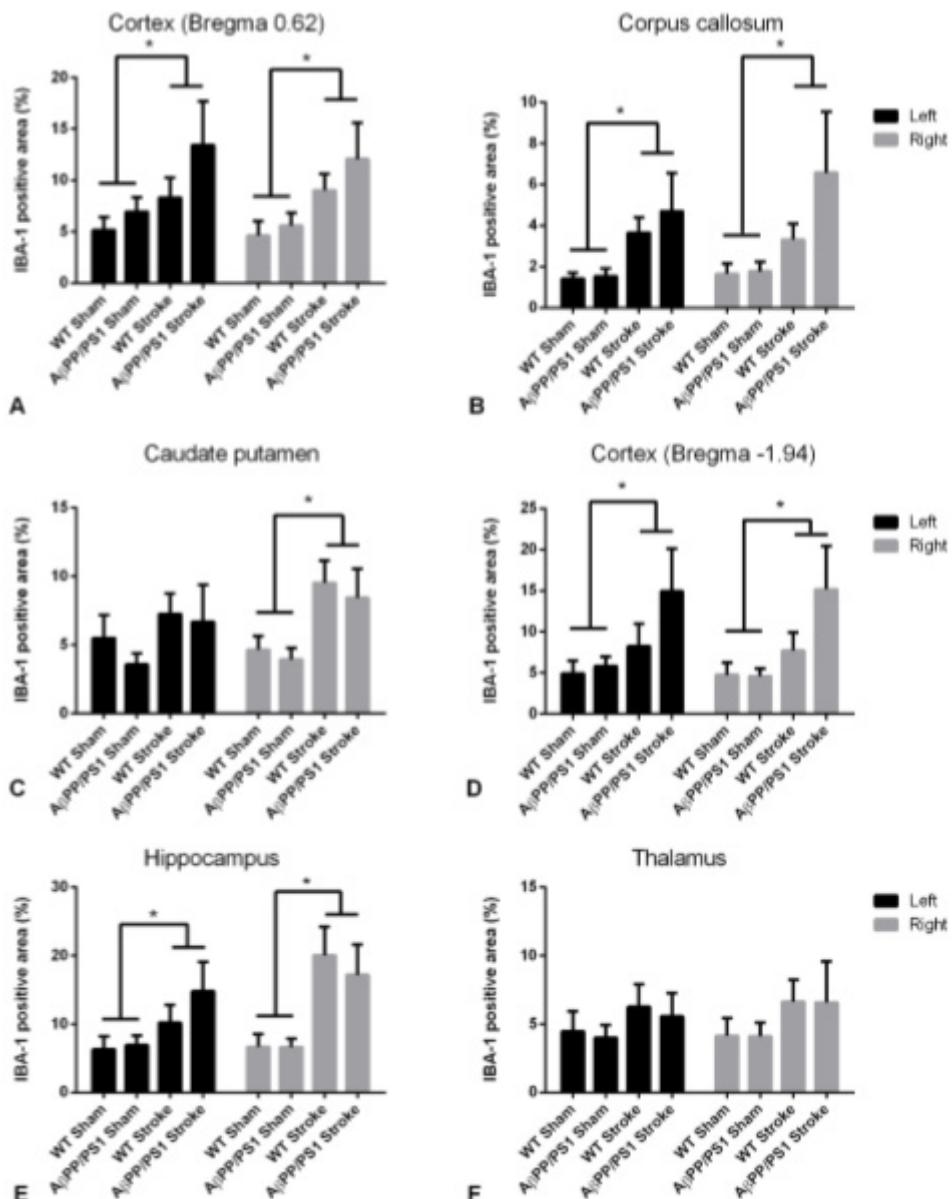
Supplementary table 1. Overview of all results. Summary of all significant results found in Aβ and synaptophysin expression. 0: no significant results, ↑: significant increase, ↓: significant decrease, ↑#: trend which shows an increase, ↓#: trend which shows a decrease.

| Target	Region	Parameter	Effect	Results
ELISA - Aβ	Anterior brain	Aβ40	Right vs Left	0
			Stroke vs Sham	0
			Female vs Male	↑ in right and left hemisphere in female mice
		Aβ42	Right vs Left	0
			Stroke vs Sham	0
			Female vs Male	0
		Aβ42/ Aβ40	Right vs Left	0
			Stroke vs Sham	0
			Female vs Male	0
IHC - Aβ40	Cortex (-1.94)	Plaque size	Right vs Left	↑# in male sham and female stroke mice
			Stroke vs Sham	0
			Female vs Male	0
		% +area	Right vs Left	0
			Stroke vs Sham	0
			Female vs Male	0
		#plaques/mm²	Right vs Left	0
			Stroke vs Sham	0
			Female vs Male	0
	Hippocampus (-1.94)	Plaque size	Right vs Left	0
			Stroke vs Sham	0
			Female vs Male	0
		% +area	Right vs Left	↑ in female stroke mice
			Stroke vs Sham	0
			Female vs Male	↑ in left and right hemisphere in female mice
		#plaques/mm²	Right vs Left	↑# in female stroke mice
			Stroke vs Sham	0
			Female vs Male	↑ in left and right hemisphere in female mice
	Thalamus (-1.94)	Plaque size	Right vs Left	↑ in female shame mice
			Stroke vs Sham	0
			Female vs Male	↑ in right hemisphere in female mice
		% +area	Right vs Left	0
			Stroke vs Sham	0
			Female vs Male	↑ in left and right hemisphere in female mice
		#plaques/mm²	Right vs Left	0
			Stroke vs Sham	0
			Female vs Male	↑ in left hemisphere in female mice
	Cortex (0.62)	Plaque size	Right vs Left	↓ in male stroke mice
			Stroke vs Sham	0
			Female vs Male	0
		% +area	Right vs Left	↓ in male stroke mice
			Stroke vs Sham	0
			Female vs Male	↑ in right hemisphere in female mice
		#plaques/mm²	Right vs Left	↓ in male stroke mice
			Stroke vs Sham	0
			Female vs Male	↑ in left and right hemisphere in female mice
	Corpus Callosum (0.62)	Plaque size	Right vs Left	0
			Stroke vs Sham	0
			Female vs Male	0
		% +area	Right vs Left	0
			Stroke vs Sham	0
			Female vs Male	0
		#plaques/mm²	Right vs Left	↑ in male sham mice
			Stroke vs Sham	0
			Female vs Male	0
	Basal ganglia (0.62)	Plaque size	Right vs Left	0
			Stroke vs Sham	↓# in right hemisphere
			Female vs Male	↑ in right hemisphere in female mice
		% +area	Right vs Left	↓ in female stroke mice
			Stroke vs Sham	0
			Female vs Male	↑ in left and right (trend) hemisphere in female mice
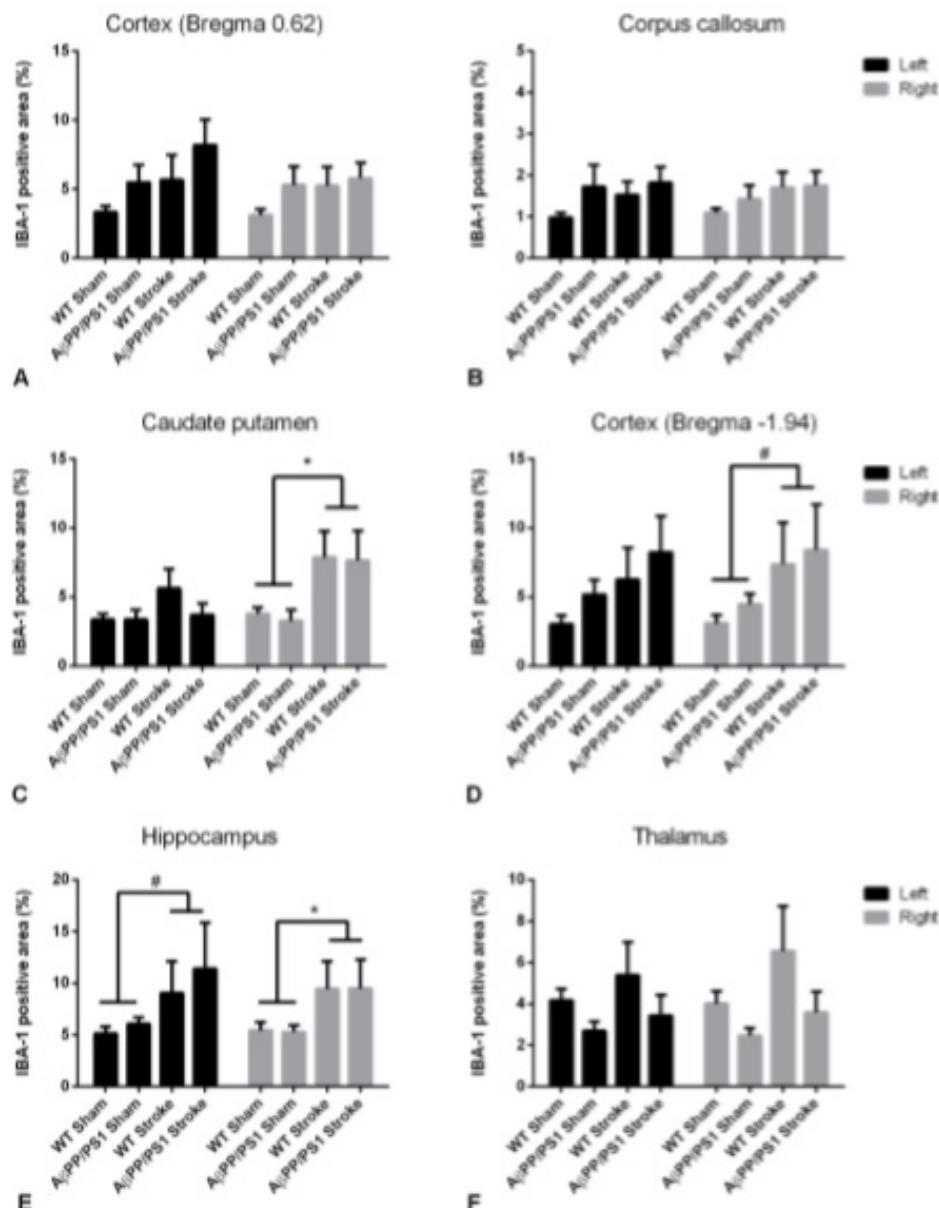
|			Right vs Left	0
		#plaques/mm²	Stroke vs Sham	0
			Female vs Male	↑ in left and right hemisphere in female mice
qPCR - SYP	Anterior brain	Synaptophysin	Right vs Left	↑ in male WT sham mice ↓ in female WT sham mice
			Stroke vs Sham	0
			Female vs Male	↓ in right hemisphere in WT sham mice ↓ in left hemisphere in APP mice
			APP vs WT	0

**Male sham**

Female sham

Male stroke

Female stroke

Supplementary figure 1. Atrophy in the basal ganglia. Representative photos of the left and right thalamus in male sham, female sham, male stroke and female stroke mice.

Supplementary figure 2. Neuroinflammation in male mice in the cortices, corpus callosum, caudate putamen, hippocampus and thalamus. All data are presented as mean ± SEM. At bregma 0.62 the percentage of IBA-1 positive area was significantly increased in stroke mice compared to sham mice in the cortex in both the left ($p<0.036$) and right ($p<0.008$) hemisphere **(A)**, in the corpus callosum in both the left ($p<0.004$) and right ($p<0.014$) hemisphere **(B)** and in the caudate putamen in the right hemisphere ($p<0.003$) **(C)**. The percentage of IBA-1 positive area was also significantly increased at bregma -1.94 in the cortex in both the left ($p<0.027$) and right ($p<0.010$) hemisphere **(D)** and the hippocampus in both the left ($p<0.032$) and right ($p<0.001$) hemispheres **(E)**. No significant differences were found in the thalamus **(F)**.

Supplementary figure 3. Neuroinflammation in female mice in the cortices, corpus callosum, caudate putamen, hippocampus and thalamus. All data are presented as mean ± SEM. No significant differences in the percentage of IBA-1 positive area were found in the cortex **(A)** and corpus callosum **(B)** at bregma 0.62. The percentage of IBA-1 positive area was significantly increased in stroke mice compared to sham mice in the caudate putamen **(C)** in the right hemisphere (p<0.006). A trend was visible in the cortex at bregma -1.94 **(D)** in which the stroke animals seemed to have an increased percentage of IBA-1 positive area compared to sham animals in the right hemisphere (p<0.077). In the hippocampus **(E)** a trend was visible in the left hemisphere in which the stroke animals seemed to have in increased in the percentage of IBA-1 positive area compared to sham animals (p<0.078). In the right hemisphere of the hippocampus, a significant increased in stroke mice compared to sham mice was found (p<0.047). No significant differences were found in the thalamus **(F)**.

Abstracts

Proceedings of the Master's Programme Cognitive Neuroscience is a platform for CNS students to publish their Master theses. Given the number of submissions, we select the articles that received the best reviews, under recommendation of our editors, for the printed edition of the journal. The abstracts of the other articles are provided below, and for interested readers a full version is available on our website: www.ru.nl/master/cns/journal.

Changes in Resting-State Functional Connectivity in Patients With Neuralgic Amyotrophy

Melissa Bakkenes

Neuralgic amoytrophy is a peripheral nervous system disorder, where the brachial plexus is affected. Neuralgic amyotrophy starts with an acute phase, which is likely caused by an acute auto-immune reaction against the brachial plexus. After the acute phase, patients with neuralgic amyotrophy often have motor problems, such as persisting abnormal posture and movement patterns of the scapula in the chronic phase. There are some indications that there could be a central nervous system problem, underlying the persisting long-term consequences of NA. Therefore, the aim of this project is to determine whether there is a difference in resting-state networks between neuralgic amyotrophy patients and healthy controls by comparing the brain activity of several resting-state brain networks between neuralgic amytrophy patients and healthy controls. To test this, we have performed a dual regression analysis with 41 neuralgic amyotrophy patients (mean age 43 ± 11 years, 16 females) with lateralized symptoms in the right upper extremity and 24 age and sex matched healthy controls (mean age 43 ± 8 years, 10 females). In addition, we assessed the specificity of a possible sensorimotor network resting-state fMRI difference in neuralgic amyotrophy patients compared to healthy controls, by considering the functional relevance of this difference on the motor impairments of NA patients. Furthermore, we wanted to evaluate if there are structural changes in the brain of patients with NA compared to healthy controls. In addition, we also wanted to evaluate whether possible resting-state fMRI changes in neuralgic amyotrophy patients are (partially) explained by structural changes in the brain. For these questions, a voxel based morphometry analysis was conducted. We have found that there is a difference in connectivity in the sensorimotor network and left frontoparietal network for neuralgic amyotrophy patients compared to healthy controls. There is no significant correlation between the sensorimotor network and capability of the upper-limb in neuralgic amyotrophy patients. There are also no significant group differences in grey matter volume. This study concludes that a short inflammation of the brachial plexus can lead to task-independent functional reorganization of the central nervous system in patients with neuralgic amyotrophy. This finding could aid in the treatment for the persisting motor deficits in patients with neuralgic amyotrophy.

Laminar fMRI at 3T: A Replication Attempt of Top-Down and Bottom-Up Laminar Activity in the Primary Visual Cortex

J. Karolis Degutis

The emerging field of laminar functional MRI (fMRI) has the capability of non-invasively measuring depth-dependent activity within regions of the neocortex. Most previous studies using laminar fMRI have been done at ultra-high field (7T and above) as it allowed for higher spatial resolution; yet, recent innovations in MR sequences have enabled submillimeter functional resolution at 3T. This study aimed to replicate two previous 7T laminar fMRI findings while using a 3T MRI scanner and thus tested whether laminar fMRI could be used for cognitively-relevant research questions when scanning at a lower field strength. A data quality comparison between the current findings and a previous 7T dataset (Lawrence et al. 2018) found lower tSNR and lower functional t-contrast activation at 3T. Additionally, the study failed to replicate the two previous results: the bottom-up stimulus-contrast yielded non-specific agranular compared to granular layer activation, while the top-down visual working memory main effect could not be replicated due to lower functional t-contrasts. This study provides evidence against the feasibility of laminar fMRI at 3T.

Cocaine Self-Administration and Social Behaviour in Extremes of the Sensory Processing Sensitivity Trait in Rats

Sophie Fennema

Sensory processing sensitivity (SPS) is a trait defined by sensory information processing, emotional reactions, and susceptibility to overstimulation. Individuals scoring high on this trait are differentially susceptible to positive and negative environments. In this study, 22 rats were selected on extremes on the SPS trait. High and low SPS-like rats either underwent cocaine self-administration trials or remained drug naïve. After a training period, where rats had access to self-administration boxes for 1 hour per day, long access exposure trials began, where they were allowed to self-administer cocaine for 6 hours per day. Social interaction and memory was scored before drug exposure, after training, and after long access exposure to the drug. High SPS animals show a greater escalation in cocaine intake compared to low SPS animals, and show more social and less non-social behaviour on the social interaction test. After long access exposure to cocaine, rats show more non-social behaviour compared to naïve rats. This study provides a deeper insight into the nature of addiction in susceptible individuals.

Searching Near and Far: Investigating Depth-Dependent Adaptation of Search Template Size in Naturalistic Visual Search

Maëlle Lerebourg

When searching for an object, current theories of visual search posit that we form a visual representation or template of the target by pre-activating neurons tuned to target features. Within naturalistic scenes, visual features of the target may however change drastically depending on its location in the scene, e.g. its retinal size depends on distance. Across two experiments, we investigated a potential mechanism for template-based search in the real world and its neural basis. We developed a search paradigm requiring participants to take into account depth-dependent size changes and investigated whether the visual system may account for these by creating differently-sized templates based on depth.

In a first experiment, we used breaking continuous flash suppression (b-CFS) to probe the template and test whether size-matching probes were detected faster. Suppression times to probes were however generally not by their match with target features.

Using fMRI, we found overlapping activation patterns for seeing objects of varying retinal and preparing to search for these objects near or far specific to LOC, suggesting depth-dependent templates, an effect which was however not specific to the search instructions per se. Distance-information from scene-selective areas PPA or OPA did not contribute to template-size, but we found a potential contribution of early visual cortex and low-level visual features. While further research is needed to understand whether our findings indeed reflect rescaled templates or a more general process as e.g. distance processing to rescale the object representation itself, these likely still contribute to our ability to account for changes in visual features during search.

Stress and Motivation:
A fMRI Study to the Effects of Stress on Effort-Based Decision-Making

Tim Vriens
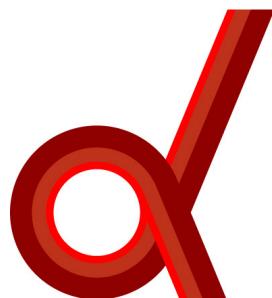
Reaching desired goals generally requires effort. In daily life, we often face important decisions on whether to engage in effortful yet very rewarding actions (e.g., studying hard to obtain a university degree, accepting a demanding job with a great salary). These decisions are often taken under considerable stress (such as when having financial uncertainty). Previous studies have shown that effort-based decision-making implicates a cortico-subcortical network involving the prefrontal cortex and midbrain catecholamine (noradrenaline and dopamine). Stress is known to affect both prefrontal activity and catecholamine levels. Despite these commonalities, the effect of stress on effort-based decision-making, and its neural correlates, are surprisingly understudied. In this study, the effects of stress on effort-based decision-making are investigated, focusing on behavior, neural activity, and noradrenergic release measured by pupil dilation. Using a within-subjects design with acute stress induction, we hypothesized that stress would decrease the willingness to engage in mental effort, and that this would be linked with alterations in prefrontal activity and pupil dilation. Participants performed an effort-based decision-making task in the fMRI scanner during control or during experimentally-induced acute stress. The results showed that stress reduced the willingness to engage in effort. Pupil dilation tracked task difficulty during calculation, but did not change under stress. Unexpectedly, no difference was observed between stress and control at the whole-brain level. An a-priori region-of-interest analysis based revealed no significant effect of stress on effort or reward coding in the anterior cingulate cortex (ACC). We did observe a significant decrease in effort coding under stress in the dorsolateral prefrontal cortex (DLPFC). However, this was specific to the high reward condition and for participants who experienced the control session first. Overall, these results suggest that acute stress reduces the willingness to engage in mental effort, with preliminary evidence of alterations in DLPFC.

A, B, or *C*ontrasting: The Influence of a Learning Task on Neurophysiological Correlates of Feedback Processing

Leonie Weindorf

Contrasting similar items has been proposed to enhance learning by increasing the specificity of mental representations. Since feedback is an important component of such tasks, the learning effect of contrasting might be partially related to the neurocognitive mechanisms of feedback processing. Previous electroencephalography (EEG) studies have demonstrated that the P300 and the feedback related negativity (FRN) can indicate whether an instance of feedback leads to successful learning. To investigate whether contrasting influences the effectiveness of feedback processing, we manipulated the orthographic similarity between the answer options on a three-choice vocabulary learning task. EEG was recorded, while participants learned 50 pseudo-translations to Italian words over six blocks. The learning outcome was determined via an immediate and a one-week delayed posttest. Results show better performance during the task for words that were presented with dissimilar distractors (shuffled condition), whereas posttest performance was higher for words that had been studied with similar answer options (sorted condition). The parietal P3b was larger for the sorted compared to the shuffled condition. A larger P3a and smaller FRN to negative feedback were associated with error correction. A larger P3a and smaller FRN to positive feedback were correlated with accuracy on the delayed recognition test. Most of these learning effects were only found for the sorted condition. This indicates that the similarity training might elicit improved memory encoding and attention reflected by the P300 amplitude, as well as enhanced utilization of valence feedback, reflected by the FRN amplitude. Taken together, the results suggest that the memory advantage of contrasting might partially be due to enhanced feedback processing. However, due to the limited sample size (as a result of the corona crisis), no definite conclusions can be drawn, and additional research is needed to corroborate these findings.

Institutes Associated With the Master's Programme Cognitive Neuroscience

Donders Institute for Brain, Cognition and Behaviour:
Centre for Cognitive Neuroimaging
Kapittelweg 29
6525 EN Nijmegen

P.O. Box 9101
6500 HB Nijmegen
http://www.ru.nl/donders

MAX PLANCK INSTITUTE
FOR PSYCHOLINGUISTICS

Max Planck Institute for Psycholinguistics
Wundtlaan 1
6525 XD Nijmegen

P.O. Box 310
6500 AH Nijmegen
http://www.mpi.nl

Radboudumc

Radboudumc
Geert Grooteplein-Zuid 10
6525 GA Nijmegen

P.O. Box 9101
6500 HB Nijmegen
http://www.umcn.nl

Baby & Child
Research Center

Baby & Child Research Center
Montessorilaan 3
6525 HR Nijmegen

P.O. Box 9101
6500 HB Nijmegen
https://www.babyandchild.nl