Table of Contents

From the Editors-in-Chief

Dear reader,

We are pleased to present to you the first issue of the tenth volume of *Proceedings of the Master's Programme Cognitive Neuroscience* or in popular terms, the CNS journal. What you hold in your hand is an anniversary edition to celebrate ten years of publishing high quality research conducted by research master students. This issue, themed "Back to the Future," endeavours to highlight past, present, and future aspects of both the CNS journal and the field of Cognitive Neuroscience.

Looking back – Ten years ago students of the Cognitive Neuroscience master programme, in collaboration with prof. Peter Desain, founded the CNS journal, an ambitious project that has now continued solidly for ten academic years. The journal is an exceptional achievement. It is completely run by students, only publishes research conducted by master students, and yet aims to reach the quality of highly-ranked professional scientific journals. We would like to thank prof. Desain for his part in inspiring the start of the CNS journal and we recommend you to read the additional editorial he wrote to commemorate this milestone anniversary.

Looking forward – The six articles published in this issue have topics ranging from conscious decision making to speech perception as, for the anniversary issue, we aimed to publish not only high quality research that spans the entirety of the broad field of Cognitive Neuroscience, but also research that aims to push the limits of our field. For example, within these pages, readers will find the controversial Libet theory challenged and cutting edge research on the cortical mapping of movement through space in the human brain. Additionally, we would like to highlight the use of state-of-the-art techniques like advanced head movement tracking and the use of a 7-Tesla MRI scanner. We believe that, for the progress of the field, challenging long-held hypotheses and publication of results using ever-evolving techniques is essential. Therefore we would like to express our hopes that the Donder's Institute, and with it the CNS journal, will continue to conduct and publish frontline and innovative research for many years to come, a tradition that we, as this year's editors-in-chief, have been honoured to promote.

We would like to thank our dedicated editorial board, the authors, editorial writers, and reviewers for all their efforts that contributed to the publication of this 10th anniversary issue of the CNS journal.

Wishing you a pleasant read,

Nijmegen, January 2015

Dorien Maas & Kelly Woudsma

Editors-in-Chief

From the Director of the CNS Master

Nijmeg You have in your hands the first Issue of the Tenth Volume of the *Proceedings of the Master's Programme Cognitive Neuroscience.*

Ten years – Is this young or old? Perhaps the answer depends on whether you are a student or a professor. Sometimes students are surprised to learn that cognitive neuroscience as a field has already existed much longer than a few decades.

In 1874, two books appeared that helped to lay the foundations of our science. One of the books was the dissertation work of a 26-year old graduate student, Carl Wernicke; the other book was written by a 42-year old professor, Wilhelm Wundt. Wernicke's book describes the functional neuroanatomy of a higher cognitive function (i.e., language) and impairments caused by brain lesions. Wundt's book gives an overview of everything that was known at that time about the anatomy and physiology of the human brain (several of his figures are still amazingly accurate), and about the human mind, based on evidence from psychophysics and reaction time research. With their books, Wernicke and Wundt sought to establish an alliance between the sciences of the brain and those of the mind – in short: cognitive neuroscience.

Wundt also started a new journal, used by his students as the local publication outlet for their work, a bit like the Proceedings of the Master's Programme Cognitive Neuroscience. Wundt's many students included James McKeen Cattell, who did important seminal empirical work in our field but spent most of his career editing journals, including Science, a failing journal that he revived. Cattell was the owner and editor-in-chief of Science for 50 years, until his death in 1944. The history of the early years of our science illustrates that good science begins with good students, and publication outlets for their work that are edited well. I congratulate the student editors and authors with the Tenth Volume of their Proceedings, and hope the journal will flourish for yet another ten years.

Nijmegen, January 2015,

Prof. Ardi Roelofs

Director of the Research Master Cognitive Neuroscience
Donders Center for Cognition

From the origin of the CNS journal

Shortly after the Research Master Cognitive Neuroscience had started in Nijmegen I was on sabbatical in Stanford. There I got to know some students involved in a scientific journal that was completely student-run. I was impressed by the professionalism and quality it radiated and the students told me how much fun it was and how much it taught them hands-on, about the issues involved in scientific publishing. Collecting fair reviews, finding the needed funds, making balanced editorial decisions - even about the work of close friends, all lots of work but really worth it. Slowly I started to suspect that a journal might also be a great addition to the CNS programme in Nijmegen. Before returning home and coining the idea there the Stanford students gave more advice, e.g. on how important it is to grab new students almost directly upon arrival at campus - continuity is quite difficult for any activity being run by students alone.

How well the idea landed in Nijmegen. With the present 10th volume CNS students not only have proven that continuity is possible, but even attained an ever increasing level of quality. Subsequent editorial boards did a tremendous job in building and perfecting the workflow, the website, the layout and all other aspects of the journal.

The enthusiasm for the journal has proven contagious, as students of the Artificial Intelligence programme are in the process of starting their own, in turn being advised by the CNS editorial board. And who knows what will happen next.

In this volume you'll find reports on the thesis work of students, ranging from neurotransmitters and cortex parcellation to consciousness and social behavior. The broad array illustrates how steps are taken in neuro-scientific research at the Donders Institute, connecting layers of scientific understanding that span from molecule to man and beyond. We are very proud of these students and what they already achieved early in their scientific career. And I congratulate all members of the present and past editorial boards with a fantastic achievement.

Nijmegen, January 2015,

Prof. Peter Desain

Artificial Intelligence Department
Donders Center for Cognition

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

Contact Information:
Journal CNS
Radboud University
Postbus 9104
6500 HE Nijmegen
The Netherlands
nijmegencns@gmail.com

Social Bonding: A Matter of Serotonin Transporter Gene Variation and Oxytocin Neurotransmission?

Alexandra Chovsepian[1]
Supervisors: Judith Homberg[1], Tamas Kozicz[1]

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

One genetic factor influencing social coping style is the short ("s") variant of the human serotonin-transporter-linked polymorphic region (5-HTTLPR) that has been associated with higher sensitivity to social stress in a Gene-Environment interaction fashion (GxE). Besides the GxE concept, 5-HTTLPR might be related to the Gene-Environment correlation (rGE) concept, according to which our social environment is heritable since it can be shaped by heritable behaviours. In our approach, the social behaviour of male Wistar 5-HTT+/+ and 5-HTT-/- rats (with the latter modeling the s-carriers of the 5-HTTLPR) is assessed during social interaction between 2 animals of similar (homogenotype interaction) or different (heterogenotype interaction) 5-HTT genotype. Juvenile, adolescent and adult groups are separately tested to assess differences in social interactions during development. In line with previous studies in which 5-HTT-/- mice negatively altered their behaviour when switched from homogenotype to heterogenotype interaction, we hypothesized that, for 5-HTT-/- rats, the formation or lack of social bonds is influenced by the type of social interaction (homo- vs heterogenotype) reflecting a GxE effect and this effect is altered across ages. Furthermore, we examined the hypothesis that the 5-HTT-/- genotype has a significant rGE effect, which can be reflected as increased influence of the 5-HTT-/- rats' behaviour on their partners' social behaviour and thus, their social environment. Regarding the neural underpinnings of this behavioural profile, we hypothesize that in 5-HTT-/- rats the formation of social bonds is characterized by increased oxytocin (OXT) neurotransmission and lower corticotropin-releasing factor (CRF) levels. As expected, we observed a significant 5-HTT-/- GxE effect, with the heterogenotype interaction reducing social bonding in adolescence and adulthood. Interestingly, in juveniles, the 5-HTT-/- GxE effect is expressed in an opposite direction, observed as increased social play during heterogenotype compared to homogenotype interaction. However, we were not able to understand the role of OXT-related activation in social bonding due to anti-OXTR antibody inefficiency, and we found no significant correlation between (lack of) bonding, anxiety and the CRF levels during social interaction.

Keywords: serotonin, transporter, polymorphism, social, rat, behaviour, oxytocin, stress

Corresponding author: Alexandra Chovsepian; **E-mail:** frabalena@gmail.com

1. Introduction

1.1 Serotonin transporter gene variation: gene-environment interaction (GxE) and correlation (rGE)

Serotonin (5-HT) is a highly conserved neurotransmitter that plays a key role in brain function varying from feeding, sleep and nociception to emotion, motivation and cognition (Serretti et al., 2006). An important component of the serotonergic network is the serotonin transporter (5-HTT) which reuptakes the 5-HT from the synaptic cleft into the presynaptic neuron for recycling or metabolic degradation and thus regulates the extracellular serotonin levels by clearing 5-HT from the synaptic cleft (Murphy et al., 2001). Since its identification in the 1990s, the Serotonin-transporter-linked polymorphic region (5-HTTLPR) has been extensively studied in relation to neuropsychiatric disorders (Wendland, Martin, Kruse, Lesch, & Murphy, 2006). Although other subdivisions have been stated, 5-HTTLPR is commonly reported with two variations in humans: A short ("s") and a long ("l"). It has been reported that the s-allele is associated with a relative reduction in serotonin transporter mRNA (Heils et al., 1996). There is a high variation in the allelic frequency of this polymorphism across populations, for example the allelic distribution in Europe is 59% for the l- and 42% for the s-allele (Reneman et al., 2006) while in Japan the s-frequency reaches 81% (Murakami et al., 1999). Therefore, studying the s-allele implication in neuropsychiatric disorders is a matter of high social significance taking into account that the percentage of the population it affects is anything but negligible.

It has been shown that the 5-HTT polymorphism affects depressive responses to life stress in a gene-environment interaction (GxE) manner (Caspi et al., 2003). According to the GxE "for better and for worse" concept as described by Belsky et al. (2009), one allelic variant of a polymorphism (the s-allele) is more sensitive to the environment than the other variant (the l-allele) and influences behaviour negatively under aversive environments, but positively under favorable environments (Belsky et al., 2009). Thus, carriers of the allele that induces higher sensitivity to environmental stimuli might be detrimentally affected by an adverse environment, developing stress and depression, but on the other hand, they might strongly benefit from a positive environment that provides social support. Reversely, an individual's social environment is not constant but is shaped by their own behaviour. Taking into consideration that many components of our behaviour are heritable, our social environment can be heritable as well, and this effect is known as a Gene-Environment correlation (rGE). The rGE reflects genetic influences on environmental exposure (Jaffee & Price, 2007).

1.2 Serotonergic, behavioural and social interaction profile of 5-HTT-/- animal model

The first rodent model for the 5-HTT short allele variant was the 5-HTT-/- mouse, developed in 1998, followed by the development of the 5-HTT-/- rat in 2006 (Homberg et al., 2007a; Smits et al., 2006). 5-HTT-/-mice show a complete lack of 5-HT uptake in brain synaptosomes (Bengel et al., 1998), while in the 5-HTT-/- rat synaptosomes 5-HT uptake is reduced, but not absent (Homberg et al., 2007a). Since a complete absence of the 5-HTT does not exist in humans, except for some rare mutations, the lack of the 5-HTT in 5-HTT-/- rodents may represent the human condition in an exaggerated pattern. However, 5-HTT-/- rodents are regarded as models of the s-allele of 5-HTTLPR, based on evidence that these animals show increased negative emotionality and neurodevelopmental changes that are resembling the effects of the s-allele (Altamura et al., 2007; Hariri & Holmes, 2006; Holmes & Hariri, 2003; Homberg, Van den Bos, Den Heijer, Suer, & Cuppen, 2008; Wellman et al., 2007). However, although the 5-HTT+/- rodents could be considered as a more appropriate model of the human s-allele condition from a molecular and biochemical point of view (Bengel et al., 1998; Hariri & Holmes, 2006; Homberg et al., 2007a), the heterozygous animals do not display such clearly altered endophenotypes compared to 5-HTT+/+ controls (Kalueff, Ren-Patterson, & Murphy, 2007c).

Both 5-HTT-/-rodent models show an altered behavioural phenotype, expressed as higher anxiety, decreased social behaviour, increased negative emotionality and behavioural inhibition. One behavioural component that is influenced by the 5-HTT inactivation is exploratory behaviour. Although 5-HTT-/- models do not have any severe motor impairment, their exploring strategy (running and stopping) is indicative of an anxiety-like endophenotype (Nonkes & Homberg, unpublished data). Another important rodent behaviour to consider is self-grooming, which is highly relevant to anxiety and obsessive compulsive disorder (OCD) (Kalueff & Tuohimaa, 2004; Kalueff, Aldridge,

LaPorte, Murphy, & Touhimaa, 2007a). Although grooming activity was unaffected in 5-HTT-/- mice (Kalueff et al., 2007a), 5-HTT-/-rats show increased self-grooming behaviour (Nonkes & Homberg, unpublished data). In sum, 5-HTT-/- rats are characterised by elevated anxiety levels, expressed not only as higher grooming but also as aberrant exploration strategy.

The involvement of abnormal 5-HT neurotransmission in the induction of negative emotionality, such as fear and anxiety, has been further investigated in the 5-HTT-/-rodent. 5-HTT-/- mice exhibit a robust anxiety-like behaviour assessed by various tests including the open field and Elevated Plus Maze paradigms, as well as the novelty-induced suppressed feeding (Lira et al., 2003) and fear conditioning paradigms (Wellman et al., 2007). 5-HTT-/- rats show higher avoidance of the center of the open field and spend less time on the open arms compared to the closed arms of the Elevated Plus Maze, indicating increased stress (Olivier et al., 2008). These findings collectively confirm that impairment of 5-HTT induces anxiety-like behaviour in 5-HTT-/- Wistar rats (Olivier et al., 2008).

Several studies have observed altered social behaviour in 5-HTT-/- mice and rats. For instance, isolated male 5-HTT-/- mice (Holmes, Murphy, & Crawley, 2002) and 5-HTT-/- rats (Homberg et al., 2007b) showed lower aggression in the resident-intruder test, compared with 5-HTT+/+ animals. Moreover, 5-HTT-/- adult rats showed less initiations of attack than 5-HTT+/+ peers, while the levels of social exploration were unaffected (Homberg et al., 2007b). Importantly, social interaction is reduced in 5-HTT-/- mice (Kalueff, Fox, Gallagher, & Murphy, 2007b) suggesting that 5-HTT-/- mice may represent an animal model of autism spectrum disorders.

These changes in social interaction style might have their roots in earlier developmental stages, since social behaviour is built up during the juvenile period and adolescence. A very interesting type of social behaviour that needs to be considered is social play. Social play behaviour, which is expressed during the peri-adolescence period (post-natal day (PND) 28–35), represents behaviours that occur in adult sexual, affiliative and aggressive encounters (Baenninger, 1967), but displayed out of context or/and in an exaggerated fashion (Poole & Morgan, 1975). According to Homberg et al. (2007c), rat play behaviour like pouncing, pinning and chasing was reduced in juvenile (PND 28-35) 5-HTT-/- rats compared to 5-HTT+/+ rats during social interaction tests within couples of the same genotype (homogenotype condition). However, the 5-HTT-/- rat non-playful social encounters (allo-grooming, social exploration) remained unaltered, suggesting that 5-HTT-/-rats are interested in their play-peer, spend a considerable amount of contact-time with him, but display a markedly reduced initiation to play (Homberg et al., 2007c). This finding suggests that 5-HTT genotype affects the occurrence of social interactions and, therefore, reflects the aforementioned rGE effect.

Furthermore, in a study of adult rat dominant behaviour, it was reported that, in homogenotype groups, 5-HTT-/- mice were equally able to establish and maintain a dominant social position by offensive aggressive behaviour. However, during heterogenotype interaction between 5-HTT+/+ and 5-HTT-/- mice, 5-HTT-/- were significantly inferior to 5-HTT+/+ mice (Lewejohann et al., 2010). Such a 5-HTT-/- genotype-specific change in social behaviour under different social environment (hetero-genotype versus homo-genotype interaction) reflects a GxE effect. Thus, changes in social behaviour between the 5-HTT-/-homogenotype and heterogenotype interactions can be a useful measure of GxE.

1.3 5-HTT implication in CRF and OXT neurotransmission

5-HTT impairment affects the same areas in 5-HTT-/- mice and human s-allele carriers, with the former showing structural changes in the prefrontal cortex (PFC) and amygdala (Wellman et al., 2007) and the latter showing volume changes and increased connectivity between these two brain regions (Hariri & Holmes, 2006; Hariri et al., 2002; Heinz et al., 2005; Pezawas et al., 2005). Specifically, the s-allele is associated with a smaller, but hyper-reactive (Hariri et al., 2002) and hyper-aroused (Canli et al., 2005) amygdala, resulting in higher emotional vigilance. That might explain the higher sensitivity to emotional stimuli that characterizes both s-allele carriers and the 5-HTT-/- animals.

Jiang, Wang, Luo, & Li (2009) demonstrated that the HPA axis and its feedback regulation are altered 5-HTT-/- mice, which could account for their increased sensitivity to stress. Under basal conditions, corticotrophin-releasing factor (CRF) mRNA levels in the hypothalamic paraventricular nucleus (PVN) were lower in 5-HTT-/- mice compared to 5HTT+/+ mice, but under stressful conditions (elevated plus maze test), the HPA axis in 5-HTT-/- mice is super-activated, displaying elevated CRF levels and

significant increase of adrenocorticotropic hormone (ACTH) release (Jiang et al., 2009; Li et al., 2004; Li et al., 1999). The increased sensitivity to stress may be due to a reduction in the glucocorticoid receptor-mediated feedback regulation, as GR expression was significantly reduced in the hypothalamus, pituitary and adrenal cortex of 5HTT+/- and 5HTT-/- mice (Jiang et al., 2009).

Oxytocin (OXT) is a mammalian neuromodulator that acts through binding to the OXT receptor (OXTR), an important component of the brain's social network located in 'social regions' like the amygdala, prefrontal cortex, insula, and serotonergic raphe nuclei. Besides playing an important role in the neurobiology of intimacy and maternal functions, OXT influences more complex behaviours. OXT is known to enhance (pro)social behaviour and empathy, and reduce anxiety (Bartz, Zaki, Bolger, & Ochsner, 2011; Meyer-Lindenberg, Domes, Kirsch, & Heinrichs, 2011; Striepens, Kendrick, Maier, & Hurlemann, 2011; Van IJzendoorn & Bakermans-Kranenburg, 2012). The anxiolytic potential of OXT has indeed been reflected by decreased cortisol and subjective behavioural responses to social stress (Heinrichs, Baumgartner, Kirschbaum, & Ehlert, 2003). Moreover, rodent studies have reported that OXT receptor-deficient (OXTR-/-) mice show robust social deficits (Takayanagi et al., 2005). Accumulating evidence suggests that oxytocin influences serotonin neurotransmission and reversely (Dölen, Darvishzadeh, Huang, & Malenka, 2013; Emiliano, Cruz, Pannoni, & Fudge, 2007; Jørgensen, Riis, Knigge, Kjaer, & Warberg, 2003; Larsen, Hay-Schmidt, Vrang, & Mikkelsen, 1996; Lee, Garcia, Van de Kar, Hauger, & Coccaro, 2003; Marazziti et al., 2012; Montag, Fiebach, Kirsch, & Reuter, 2011). Interestingly, OXT infusion not only facilitated serotonin release in the median raphe nucleus but also reduced anxiety-related behaviour (Yoshida et al., 2009). Infusion of a 5-HT2A/2C receptor antagonist blocked the anxiolytic effect of OXT, indicating that OXTR activation in serotonergic neurons mediates the anxiolytic effects of OXT (Yoshida et al., 2009). In addition, pharmacological studies suggested that key mediator of the antidepressant effect of selective 5-HT reuptake inhibitors (SSRIs), such as citalopram and fluvoxamine, is OXT release (De Jong, Veening, Olivier, & Waldinger, 2007; Swaab, Fliers, Hoogendijk, Veltman, & Zhou, 2000; Uvnäs-Moberg, Björkstrand, Hillegaart, & Ahlenius, 1999).

1.4 Hypothesis

In our study we assessed the rGE effect as the influence of the individual's genotype (5-HTT-/- or +/+) on the behaviour of the partner, comparing the occurrence (frequency/time) of behaviours expressed by the interaction partners. In addition, the assessment of changes in social behaviour between the 5-HTT-/-homogenotype and heterogenotype interaction provided a measure of the GxE effect. Taking into account previous studies showing that 5-HTT-/- animals negatively changed their social behaviour when switched from homo-genotype to heterogenotype interactions (GxE effect), one of our aims was to test the hypothesis that homogenotype interactions lead to bonds in social interactions, whereas heterogenotype social interactions lead to conflicts (or absence of bonding) in 5-HTT-/- rats. We investigated this hypothesis in different ages, peri-adolescence, adolescence and adulthood, since the nature of social interactions is not static but changes during development (Kiser, Steemers, Branchi, & Homberg, 2012).

Regarding OXT and CRF neurotransmission, we hypothesized that the formation of social bonds (reflected as increased social play and prosocial behaviour) requires higher OXT levels while the lack of social bonding (reflected as reduction or absence of play and prosocial behaviour) is characterized by increased levels of self-grooming, a behaviour considered to reflect anxiety, increased CRF expression and reduced OXT levels.

Taken together, our expectations were that 1) for 5-HTT-/- rats the formation or lack of social bonds is influenced by the type of social interaction (homogenotype vs. heterogenotype) reflecting a GxE effect and this effect is altered across ages, 2) the 5-HTT-/- genotype influences the social environment, which in our case is considered to be the occurrence of social behaviours expressed by the interaction partner, 3) OXT and CRF neurotransmission is affected by the genotype and type of interaction and 4) the formation of social bonds is characterised by increased OXT neurotransmission and lower CRF levels.

2. Materials and methods

2.1 Animals

All experiments were approved by the Committee for Animal Experiments of the Radboud University Nijmegen Medical Centre, Nijmegen, The Netherlands and all efforts were made to minimize animal discomfort and reduce the number of subjects used. In our study we implemented the well characterized 5-HTT-/- rat model. We tested 3 differently aged groups, the juvenile, the adolescent and the adult group, containing 40 male Wistar rats per group. Half of the animals in each age group were 5-HTT-/- and the other half were 5-HTT+/+. 50% of the animals of each genotype (5-HTT+/+ or 5-HTT-/-) interacted with a peer with the same genotype, while the rest interacted with mates of different genotype. Table 1 illustrates the number of animals used in each genotype*type of interaction condition.

The rats were housed in cages of 2 or rarely 3 (macrolon cage type III) under a normal day/light cycle (lights off at 7 P.M.) at controlled room temperature (21±2°C). Food and water were available ad libitum.

2.2 Social Interaction paradigm

We assessed the rat social behaviour by means of 15-minute sessions of interaction between two animals in a phenotyper cage (46x46x56 cm). The interaction had to be unfamiliar, thus the interacting animals were not cage-mates. On day 1, the experimental procedure started with habituating the animals to the testing environment. For that purpose each animal was allowed to explore, alone, the phenotyper cage for 10 minutes. On day 2, a second 10-minute habituation took place. On day 3, the animals had to be isolated for 3.5 hours prior to the social task so that social play and social behaviour were induced (Niesink & Van Ree, 1989). For that purpose, we separated the cage-mates, leaving one animal per cage. The cages that contained the isolated rats were not placed close to each other in order to avoid the possible effect of isolation-induced calls, considering that separation from conspecifics elicits 50-kHz ultra-sonic vocalizations in rats (Schwarting, Jegan, & Wohr, 2007; Wöhr & Schwarting, 2013). After the 3.5 hours of isolation, the unfamiliar, weight-matched homogenotype or heterogenotype couples were allowed to interact for 15 minutes in the phenotyper cage while being recorded.

At the end of the experiment the rats were transferred back to their original cages with the same cage-mate.

On day 4, the isolation and interaction protocol of day 3 was repeated, with the same couples interacting for second time. We introduced this second session of interaction to investigate the effect of the introduced familiarity on social behaviour.

2.3 Immuno-fluorescence and DAB staining

On day 4, directly after the social interaction session, the animals were anesthetized by intraperitoneal (IP) injection of sodium pentobarbital

Table 1.

Age group	Genotype	Type of interaction	Number of animals
Juvenile (PND 28-35)	5-HTT+/+	Homogenotype	10
		Heterogenotype	10
	5-HTT-/-	Homogenotype	10
		Heterogenotype	10
Adolescent (PND 36-42)	5-HTT+/+	Homogenotype	10
		Heterogenotype	10
	5-HTT-/-	Homogenotype	10
		Heterogenotype	10
Adult (PND 70-78)	5-HTT+/+	Homogenotype	10
		Heterogenotype	10
	5-HTT-/-	Homogenotype	10
		Heterogenotype	10

(90 mg/kg) and perfused transcardially with 0.1 M PBS, (pH 7.3), followed by 4% paraformaldehyde (PFA) dissolved in 0.1 M PBS, pH 7.2. The brains were collected and fixated in PFA at 4°C for 24 hours and then stored in 0.1 M PBS at 4°C. Then, the brains were sectioned in slices of 30 μm, using frozen sliding microtome (MICROM HM 440E) and sections were stored in 0.1 M PBS + 0.05% natrium azide (NaN3).

We assessed OXT neurotransmission using double immunofluorescence (IF) staining of oxytocin receptor (OXTR) and c-Fos (marker of recent neuronal activation) by quantifying co-localization of OXTR and c-Fos positive cells. We specifically investigated brain regions that are known to contain high levels of OXTR and are a part of the "social brain network", such as the central amygdala (ceAm), the hypothalamic paraventricular (PVN) and supraoptic (SON) nuclei (bregma -1.88 mm, anatomically defined according to the atlas of Paxinos (Paxinos, 1994) and the bed nucleus of the stria terminalis (BNST, bregma -0.26mm). After washing with phosphate-buffered saline (PBS; 3x10 min), our slices were pre-incubated with PBS-BT (1% BSA, 0.3% Triton X-100) and incubated overnight with primary goat anti-oxytocin receptor antibody (sc-8103, Santa Cruz Biotechnology) and primary rabbit anti c-Fos antibody, diluted in incubation buffer. We tested different concentrations of the anti-OXTR antibody (1:100, 1:200 and 1:400) while the concentration of anti-c-Fos antibody was consistently 1:1000. The following day, after 3x10 min. PBS washing steps, we incubated the slices with the second incubation buffer, containing the secondary antibodies. We used donkey anti-goat Alexa Fluor 488-conjugated antibody for the OXTR and donkey anti-rabbit Cy3-conjugated antibody for the c-Fos, both with 1:100 dilution in PBS-BT. 3 Hours later, after 3x10 min PBS washings the slices were mounted on gelatin coated slides, dried and stored overnight. After being embedded in Fluorsave, they were ready for observation with a confocal microscope. Unfortunately, we did not observe any OXTR-positive cells thus we decided to make modifications in the staining procedure.

Therefore, we followed the same procedure but tested different dilutions of the primary anti-OXTR antibody (1:50, 1:100, 1:200) and replaced the original secondary antibody for OXTR with donkey anti-goat Cy3- conjugated antibody and the original secondary antibody for c-Fos with donkey anti-rabbit Alexa 488-conjugated antibody. Again, the staining was not successful and we decided to proceed to the CRF staining.

We measured CRF levels using DAB staining. After 3 x 10 min washings with PBS we incubated the sections containing PVN (bregma -0.82 mm), amygdala (bregma -2.80 mm) and BNST (bregma -0.26 mm) in endogenous peroxidase blocking solution with 30% of 0.3% H2O2 in 0.1 M PBS. Following 3x10 min washing in PBS, we put the slices in pre-incubation buffer with PBS-BT (1% BSA, 0.3% Triton X-100) for 30 min and then incubated the slices with primary anti-CRF antibody in 1:2000 dilution, overnight. 3x10 min washings in PBS followed and then we incubated the sections with secondary anti-rabbit IgG biotinylated antibody in dilution of 1:1500 at room temperature for 90 min. After 3x10 min PBS washings the sections were incubated with Avidin/Biotyn Complex (ABC) in dilution of 1:800 for 90 min. 3x10 min PBS washings were followed by pre-incubation with 30% DAB dilution without H2O2 for 8 min and incubation with 30% DAB+H2O2 dilution for 8 min until the reaction was stopped by another series of 3x10min PBS washings. The sections were mounted on gelatin-coated slides and dried overnight in the stove at 37°C. The next day the sections were dehydrated in series of alcohol dilutions from 70% to 100% and (2 min in each concentration), then immersed in xylene for 3x2 min and finally covered with coverslips, using Entellan, in order to be ready for observation. The pictures of the areas of interest are obtained with Olympus slide scanner microscope, using 20x magnification.

2.4 Analysis of social behaviour

Social interaction was recorded for 15 minutes and analyzed with the program Keys Version 1.01. Different components of social behaviour were measured for the different age groups, as displayed in the Table 2. The program computed the values of total duration and frequency of each behavioural component. Pouncing describes soliciting another animal by attempting to nose or rub the nape of its neck. Pinning refers to the situation in which the rat that is solicited upon rotates to its dorsal surface on the floor with the other animal standing over it. Active prosocial behaviour describes the behaviour of scratching, grooming or licking the fur of the partner as well as the behaviour of moving over or beneath the body of the partner (crawling over/ under). Extended time of keeping this "over or under" position, as well as the behaviour of adjacent lying are considered as passive prosocial behaviour.

Table 2.

	Social play behavior	**Aggression-like behavior**	**Social non -play (social exploration/ prosocial behavior)**	**Offensive aggression**	**Anxiety-related behavior**
	Pinning, pouncing	Boxing/wrestling Following/ chasing	Naso-nasal/ anogenital/ body sniff (exploration) Allo-grooming (active prosocial) Adjacent lying (passive prosocial)	Attack/biting/ escalated fight	Self-grooming
Age goups					
Juveniles	+	+	+		+
Adolescents		+	+	+	+
Adults			+	+	+

2.5 Statistical analysis

All statistical analyses were performed using SPSS Version 22.0. GxE interactions were analyzed using the univariate general linear model (GLM). The effects introduced by familiarity, which is induced by the second social interaction session, were measured by repeated-measures GLM, using session (first or second) as within-subjects variables. Furthermore, in order to assess the rGE effect, we used the measure of occurrence, defined as the frequency of the behaviour divided by the time (f/t). The measure of occurrence is more informative about the initiation of the behaviour of interest. Then, we performed a bivariate correlation analysis for each component of social behaviour. The result of this procedure was a series of correlation coefficients indicating whether the expression of a specific behaviour by the animals is correlated with the expression of the same behaviour by their partners.

3. Results

3.1 Genotype x Environment interaction (GxE)

The across-age effects of the factor "genotype" (5-HTT+/+ vs 5-HTT-/-) and the factor "type of interaction" (homogenotype vs. heterogenotype) on different behavioural components as well as the interaction between the two factors (genotype x type of interaction) are described on Table 3. A general pattern in our results is that the type of interaction influences social behaviour in a genotype-independent way for adults and in a 5-HTT-/- specific way for adolescents.

With reference to juvenile social play, we found significant genotype x type of interaction effect for pouncing ($F(2, 15) = 4.742$, $p < .05$) and avoidance ($g(2, 37) = 7.736$, $p < .01$). Pouncing and avoidance time is higher in the heterogenotype vs homogenotype condition and the effect is 5-HTT-/- specific (Fig. 1).

The total time that the animals spent not being in any type of contact (non-contact condition) was significantly higher for the 5-HTT-/- heterogenotype adolescent group (genotype x type of interaction: $F(2, 40) = 6.288$, $p < .05$). Regarding adult non-contact, we found significant type of interaction effect ($F(2, 40) = 6.768$, $p < .05$) showing that the heterogenotype interaction is characterized by increased non-contact time compared to homogenotype and a significant genotype effect ($F(2, 40) = 9.240$, $p < .005$) with 5-HTT+/+ rats displaying a higher level of non-contact time than 5-HTT-/- (Fig. 2).

For all the age groups, the passive prosocial behaviour is increased in the homogenotype condition, compared to the heterogenotype and this

Table 3.

	Adults	Adolescents	Juveniles
		Passive prosocial	
Significance	Interaction type sig .001	Genotype*interaction type sig .017	Genotype*interaction type sig .000
Comparison	Homo > hetero	5-HTT-/-homo > 5-HTT-/-hetero	Genotype sig .003 5-HTT+/+ < 5-HTT-/- 5-HTT-/-homo > 5-HTT-/-hetero
		Non contact	
Significance	Genotype sig .003 Interaction type .011	Genotype*interaction type sig .014	Genotype*type of interaction sig. .005 Genotype sig. .001
Comparison	5-HTT+/+ > 5-HTT-/- Homo < hetero	5-HTT-/-homo < 5-HTT-/-hetero	5-HTT+/+ > 5-HTT-/- 5-HTT+/+homo > 5-HTT+/+hetero
		Active Prosocial	
Significance	Interaction type sig .001	Genotype*interaction type sig .033	Interaction type sig .000
Comparison	Homo > hetero	5-HTT-/-homo > 5-HTT-/-hetero	Homo < hetero
		Pouncing	
Significance	Not expressed	Not expressed	Day1: genotype*interaction type sig .037
Comparison			5-HTT-/-homo < 5-HTT-/-hetero
		Avoiding	
Significance	Not expressed	Not measured	Genotype*interaction type sig .007
Comparison			5-HTT+/+homo > 5-HTT+/+hetero 5-HTT-/-homo < 5-HTT-/-hetero t-test .011
		Self-grooming	
Significance	Genotype sig .005	Genotype*interaction type sig .033	Interaction type sig .000
Comparison	5-HTT+/+ < 5-HTT-/-	5-HTT-/-homo < 5-HTT-/- hetero	Homo > hetero

effect is 5-HTT-/- specific in adolescents (genotype x type of interaction: $F(2, 40) = 5.988$, $p < .05$) and juveniles (genotype x type of interaction: $F(2, 39) = 14.717$, $p < .001$) but genotype-independent in adults (type of interaction: $F(2, 40) = 11.993$, $p < .001$). Interestingly, in the juvenile group, the 5-HTT-/- animals generally show increased passive prosocial behaviour compared to the 5-HTT+/+ (genotype

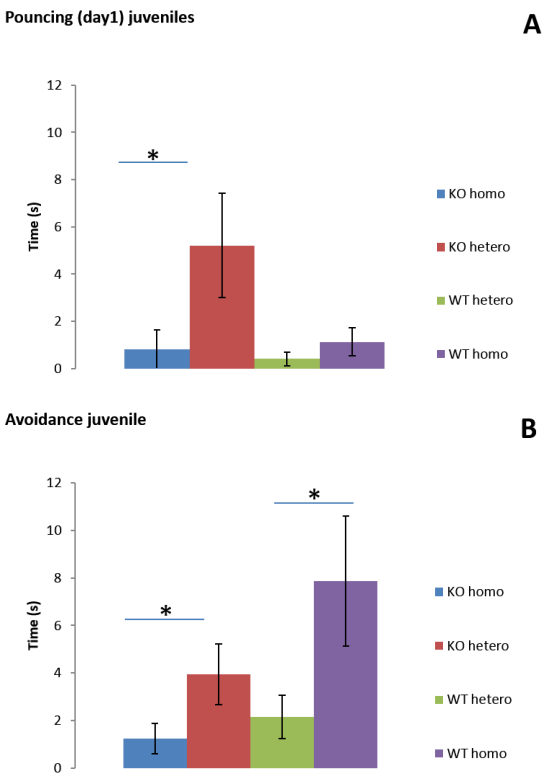

Fig. 1 Social play behavior juveniles, mean ± SEM, * p <.05. **A.** 5-HTT-/- (KO) animals increase their pouncing when interacting with 5-HTT+/+ (WT) peers, p = .037. Similarly, avoiding **B.** is higher for the 5-HTT-/- heterogenotype vs. 5-HTT-/- homogenotype group and higher for the 5-HTT+/+ homogenotype vs. 5-HTT+/+ heterogenotype group, p = .007.

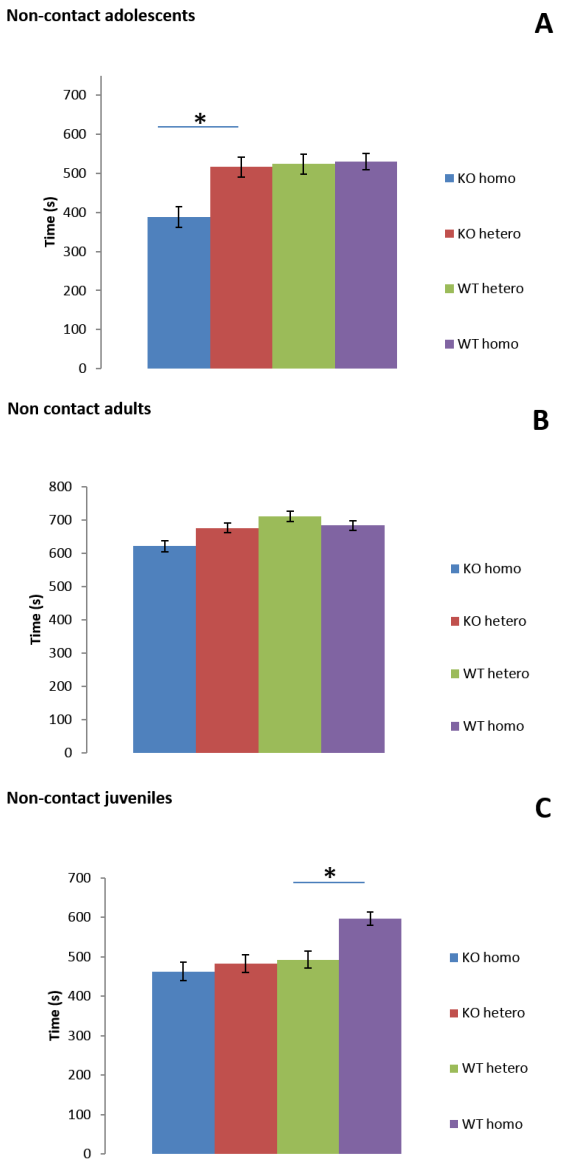

Fig. 2 Non-contact across ages, mean ± SEM, * p < .05. **A.** The time of no contact increased in the adolescent KO group during heterogenotype vs. homogenotype interaction, p = .014. **B.** In adults we saw a significant effect of genotype (less non-contact in 5-HTT-/-), p = .003 and an effect of type of interaction (higher non-contact in hetero- vs. homogenotype), p = .011. **C.** In peri-adolescence, WT rats show higher non contact than KO, p = .001. 5-HTT+/+ homogenotype interaction increased non-contact compared to 5-HTT+/+ heterogenotype interaction, p = .005.

effect: $F(2, 39) = 7.303$, $p < .01$) (Fig. 3).

In juvenile age the self-grooming behaviour is reduced during heterogenotype interaction compared to homogenotype interaction (type of interaction effect: $F(2, 35) = 23.498$, $p < .001$). The opposite effect is observed during adolescence, where the heterogenotype interaction induces higher levels of self-grooming compared to the homogenotype condition, specifically in the 5-HTT-/- group (genotype x type of interaction effect: $F(2, 36) = 7.413$, $p < .05$). During adulthood we see an effect of genotype only ($F(2, 38) = 8.169$, $g < .01$) on grooming behaviour, with 5-HTT-/- animals displaying higher self-grooming levels than the 5-HTT+/+ (Fig. 4).

Regarding active prosocial behaviour, the heterogenotype condition induced higher levels of this behaviour in the juvenile group (type of interaction effect: $F(2, 38)=7.273$, $p < .001$) while in the adult group we observed an opposite pattern, with increased active prosocial behaviour during homogenotype compared to heterogenotype interaction. Similarly, the adolescent group expressed lower levels of active prosocial behaviour in 5-HTT-/- heterogenotype vs. homogenotype condition (genotype x type of interaction effect: $F(2, 37) = 4.697$, $p < .05$) (Fig. 5).

We didn't find any significant differences between groups in social exploration, chasing, pinning and boxing/wrestling levels. Regarding offensive aggression, we didn't observe any attacks or escalated fights during the adolescent and adult social interactions. The repeated-measures general linear model (GLM) did not reveal any significant

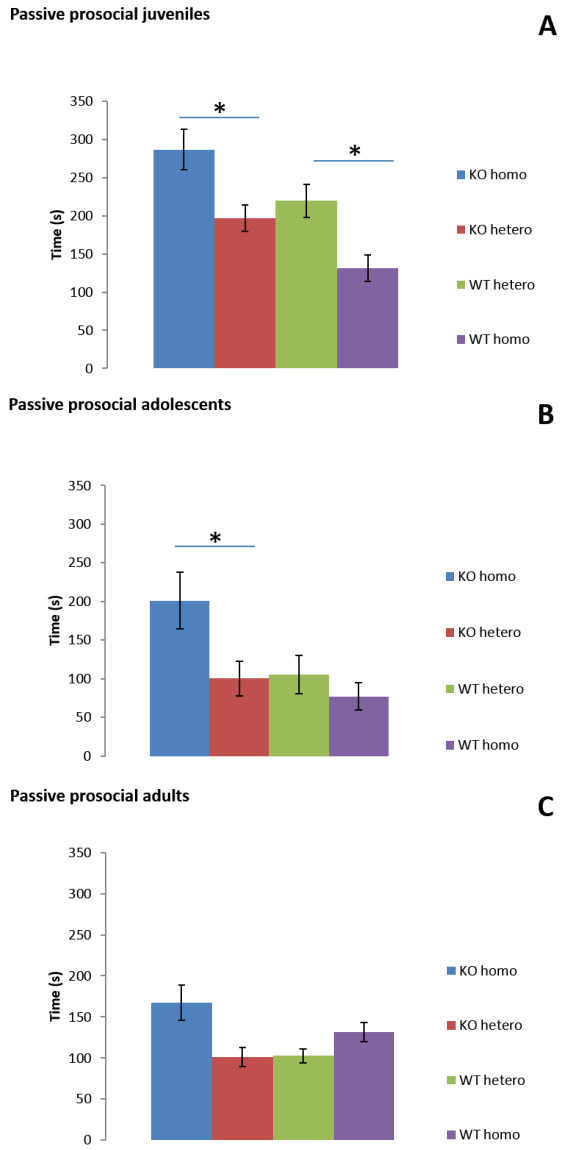

Fig. 3 Passive prosocial behavior across ages, mean ± SEM, * p < .05. **A.** Passive prosocial behavior is higher for the homogenotype condition in KO juveniles (p = .000), **B.** KO adolescents (p = .017) and **C.** adults (p = .001). In the juvenile group (C.) the WT rats increase their passive prosocial behavior during interactions with KO peers (p = .003).

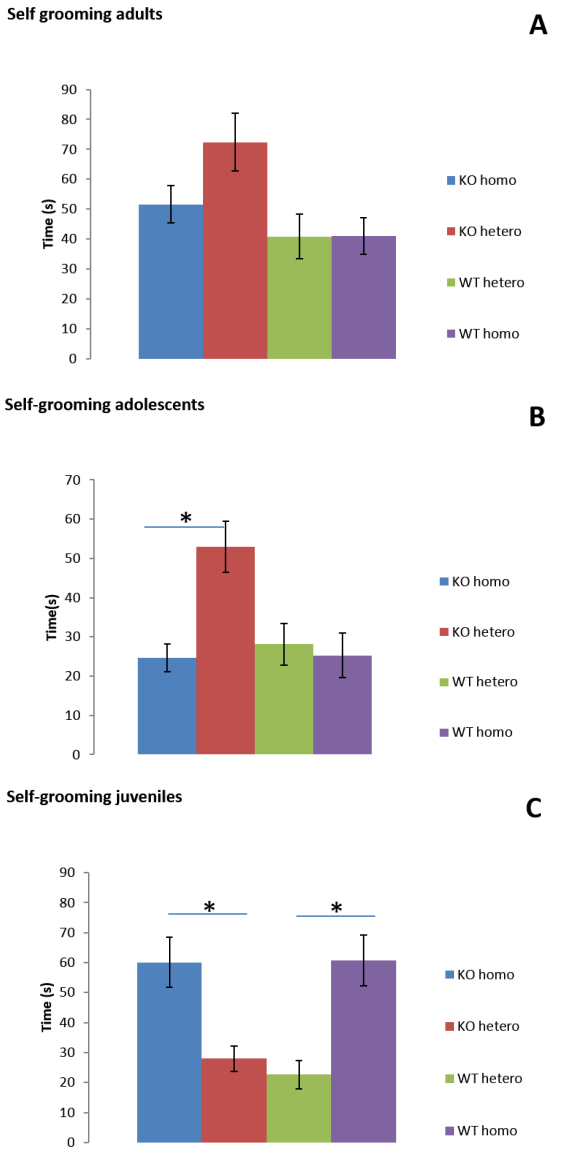

Fig. 4 Self-grooming across ages, mean ± SEM, * p < .05. **A.** During adulthood the 5-HTT-/- (KO) animals display increased self-grooming compared to the 5-HTT+/+ (WT) mice, regardless of the interaction type, p = .005. **B.** During adolescence, heterogenotype interaction induces higher levels of self-grooming, specifically for 5-HTT-/- animals, p = .033. **C.** For juveniles the effect is reversed, with self-grooming being lower in the heterogenotype condition for both genotypes, p = .000.

interaction between the session and the genotype x type of interaction effect.

3.2 Gene-Environment correlation (rGE)

Correlation analysis of occurrence (frequency/time) of social behaviours between the interacting animals revealed significant correlation of active prosocial behaviour between 5-HTT-/- rats and their homogenotype partners during peri-adolescence ($r(10) = -.514$, $p < .05$, 1-tailed) and adolescence ($r(10) = -.636$, $p < .05$, 2-tailed). Correlations between 5-HTT+/+ animals and their homogenotype partners or between 5-HTT+/+ and 5-HTT-/- partners were not significant for occurrences of active prosocial behaviour in any of the age groups. Regarding the occurrence of passive prosocial behaviour, we found significant correlation between all the interacting groups during adolescence and adulthood but not peri-adolescence (Table 4: passive prosocial).

Moreover, during adolescence we didn't find any significant correlation of self-grooming occurrence

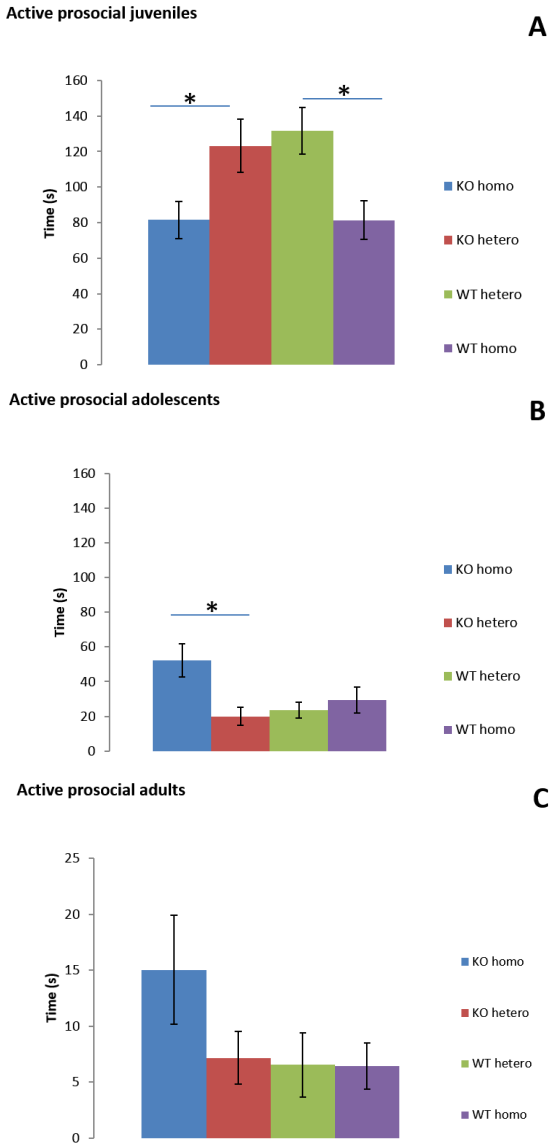

Fig. 5 Active prosocial behavior across ages, mean ± SEM, * p < .05. **A.** In juveniles, heterogenotype interaction increases active prosocial behavior for both genotypes (p = .000) while during adolescence **B.** The effect is reversed, (heterogenotype interaction reduces active prosocial behavior specifically for the 5-HTT-/- group , p = .033). **C.** In adults the heterogenotype condition decreases active prosocial behavior in a genotype-independent way, p = .001.

between the interacting groups, for any of the types of interaction. However, in juveniles the self-grooming occurrence expressed by the 5-HTT-/- rats is significantly correlated with the self-grooming occurrence of their heterogenotype (5-HTT+/+) partners (r(20) = .606, p< .01, 1-tailed). During adulthood, correlation of self-grooming occurrence is significant between 5-HTT+/+ homogenotype partners (r(10) = -.684, p< .05, 2-tailed) as well as between the 5-HTT+/+ animals and their heterogenotype (5-HTT-/-) interaction partners (r(17) = .449, p< .05, 1-tailed). The results of the

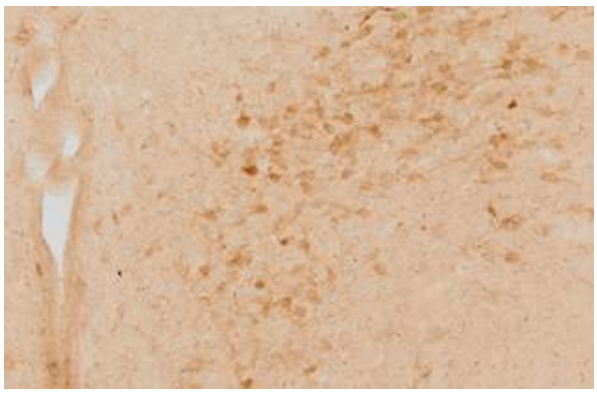

Fig. 6 Paraventricular nucleus (PVN) DAB staining for CRF. Microscope: Olympus slide scanner, 10x magnification.

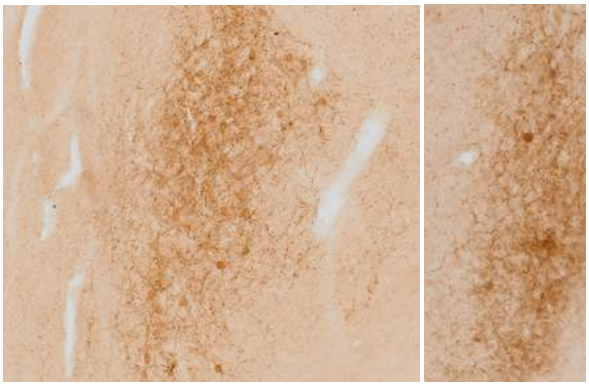

Fig. 7 (left) Central amygdala (ceAm) DAB staining for CRF. Microscope: Olympus slide scanner, 10x magnification.

Fig. 8 (right) Bed nucleus of stria terminalis (BNST). DAB staining for CRF. Microscope: Olympus slide scanner, 10x magnification.

correlation analysis for all the interacting groups, with reference to the occurrence of several behaviours of interest, are illustrated in Table 4. Behaviours that didn't reflect any significant correlations between the interacting groups across ages (like pouncing, chasing, avoidance, non-contact) are not included in the table.

3.3 OXTR-cFos and CRF staining

With reference to the OXT neurotransmission, we did not go further than the test staining because no OXTR-positive cells were identified with neither of the two different secondary antibodies we used. Regarding the CRF staining, it successfully revealed CRF-positive cells in the PVN (Fig. 6), ceAm (Fig. 7) and BNST (Fig. 8). However no significant correlation between CRF and the duration of any of the behavioural components were observed.

Table 4

Behavioral Occurrence correlation between interacting groups	Juveniles	Adolescents	Adults
active prosocial			
5-HTT-/- homogen- 5-HTT-/- homogen partners	*r* = -.514* N=10	*r* = -.636* N=10	*r* = .386 N=9
	sig .044	sig (2-tailed) =.048	sig (2-tailed) .305
5-HTT+/+ homogen- 5-HTT+/+homogen partners	*r* = .479 N=10	*r* = .436 N=9	*r* = -.124 N=10
	sig (2-tailed) .161	sig (2-tailed) .241	sig .366
5-HTT-/- heterogen-5-HTT+/+ heterogen	*r* = .190 N=20	*r* = -.067 N=20	*r* = .214 N=19
	sig (2-tailed) .421	sig (2-tailed) .779	sig (2-tailed) .380
passive prosocial			
5-HTT-/- homogen-5-HTT-/- homogen partners	*r* = .494 N=10	*r*= .601* N=10	*r* = .597* N=10
	sig .073	sig .033	sig .034
5-HTT+/+ homogen-5-HTT+/+ homogen Partners	*r* = .441 N=10	*r* = .583* N=9	*r* = .908** N=10
	sig .075	sig .05	sig (2-tailed) ,000
5-HTT-/- heterogen-5-HTT+/+ heterogen	*r* = .276 N=17	r = .572* N=19	r = .770** N=20
	sig .141	sig (2-tailed) .011	sig (2-tailed) ,000
self grooming			
5-HTT-/- homogen-5-HTT-/- homogen partners	*r* = -.086 N=10	*r* = .008 N=8	*r* = -.192 N=10
	sig (2-tailed) .790	sig .493	sig (2-tailed) .595
5-HTT+/+homogen-5-HTT+/+ homogen partners	*r* = .154 N=10	*r* = -.470 N=10	*r* = -.684* N=10
	sig (2-tailed) .633	sig .085	sig (2-tailed) .042
5-HTT-/- heterogen-5-HTT+/+ heterogen	*r* = .606** N=20	*r* = -.366 N=19	*r* = .449* N=17
	sig .002	sig .062	sig .042

4. Discussion

Our results are in many aspects in line with findings of previous studies. Significant GxE effect was observed for many components of social behaviour. We consider as GxE effect the behavioural differences between 5-HTT-/- rats that interact with heterogenotype peers and 5-HTT-/- rats interacting with homogenotype peers. We found a reduction in duration of active and passive prosocial behaviour in 5-HTT-/- heterogenotype compared to homogenotype group during adolescence (indicating a GxE effect on prosocial behaviour for adolescents). Our study is the first to examine changes in prosocial and social play behaviours between 5-HTT-/- homogenotype and heterogenotype interaction. However, our results are in line with the study of Lewejohann et al. (2010), which compared the ability of 5-HTT-/- mice to establish a dominant position during homogenotype vs. heterogenotype interaction. In that study, the 5-HTT-/- animals expressed a social impairment (were always in subordinate position) when switched from homogenotype to heterogenotype interaction, suggesting that the heterogenotype condition introduces conflict or lack of social bonding for the 5-HTT-/- animals. Similarly, our results indicate lack of bonding, assessed as reduction in

prosocial behaviour, for adolescent 5-HTT-/- rats during heterogenotype interaction. Moreover, the lack of bonding is also expressed as higher amount of time spent in non-contact instead of other social contact during 5-HTT-/- hetero- vs. homogenotype interaction, in adolescence (GxE effect on non-contact). Interestingly, the duration of self grooming, an anxiety-associated behaviour, is significantly increased in 5-HTT-/- heterogenotype vs. homogenotype interaction during adolescence. This effect suggests that the social inhibition that the adolescents 5-HTT-/- displayed when interacting with 5-HTT+/+ rats (as described above) might be caused by increased stress levels. It is possible that interaction with heterogenotype peer is a stressful situation for 5-HTT-/- animals that inhibits their social behaviour and induces more time being spent on self-grooming rather than investing on social bonding. Moreover, during adulthood we saw higher self-grooming expressed by 5-HTT-/- rats compared to 5-HTT+/+, as was also observed by Nonkes and Homberg (unpublished data), and in line with other anxiety-measure studies (Lira et al., 2003; Wellman et al., 2007) However, as previously stated, social behaviour is not static but varies across ages. Indeed, we observed an opposite effect of type of interaction on self-grooming in the juvenile group. Juveniles, independently of their genotype, significantly decreased their self-grooming during heterogenotype compared to homogenotype interactions (E effect), reaching a self-grooming level similar to the 5-HTT+/+ heterogenotype group. This might suggest that having a heterogenotype peer could have an anxiolytic-like effect in juvenile animals. In addition, this observation suggests that peri-adolescence is a very crucial period, during which the social behaviour is still plastic and might determine the social behaviour of the later life.

In accordance with this finding, we also observed a significant GxE effect on the duration of pouncing behaviour, a basic component of social play behaviour in juvenile animals. A previous study of social play (Homberg et al., 2007c) reported that rat play behaviour, like pouncing, was reduced in juvenile 5-HTT-/- rats compared to 5-HTT+/+ rats. Importantly, in our study, interacting with a 5-HTT+/+ partner made 5-HTT-/- juvenile rats more playful (in terms of pouncing), which may support the concept that interaction training of autistic children with typically developing peers improve their diminished social skills (Harper, Symon, & Frea, 2008). The other concept of our interest, the rGE effect, has not yet been extensively investigated by other studies. Therefore, we decided to use a type of statistical analysis that has not been used and verified before but appears to be plausible. In order to investigate the possible rGE effect, or in other words whether the 5-HTT-/- genotype influences the partner's behaviour which constitutes the social environment of the individuals, we used the occurrence (frequency/time) measurement. The correlation analysis between the individuals and their partners yielded different results for occurrences of different components of social behaviour. The occurrence of passive prosocial behaviour expressed by the adolescent and adult animals was significantly correlated with the occurrence of the same behaviour by their partners. This result indicates that the individual's passive prosocial attempts influence the passive prosocial behaviour expressed by the partner, in a genotype- and type of interaction-independent fashion. Thus we can't infer that there is an rGE effect on passive prosocial behaviour occurrence since the individual's influence on the partner's behaviour (and thus the social environment) is not genotype-dependent.

On the other hand, with reference to the active prosocial behaviour occurrence, we found significant correlation between the adolescent as well as the juvenile 5-HTT-/- rats and their 5-HTT-/- partners (5-HTT-/- homogenotype interaction) but not between the 5-HTT+/+ rats and their 5-HTT+/+ partners (5-HTT+/+ homogenotype interaction). This might suggest an rGE effect for active prosocial behaviour: the occurrence of active prosocial behaviour during adolescence & peri-adolescence influences the occurrence of active prosocial behaviour expressed by the partner and this effect is 5-HTT-/- specific. To put differently, the 5-HTT-/- genotype effect on a subject's active prosocial behaviour (G) influences the active prosocial behaviour of the partner which is a component of the social environment of the subject (E).

Regarding the occurrence of self-grooming, we found a significant correlation between the 5-HTT+/+ animals and their homogenotype (5-HTT+/+) partners and between the 5-HTT+/+ heterogenotype group and their 5-HTT-/- partners. This might indicate that, with reference to the rats expressing the 5-HTT gene (5-HTT+/+), their self-grooming habits may influence the self-grooming displayed by their partners, independently of the partner's genotype. Considering the results of our double-staining (c-Fos/OXTR) experiment, although the c-Fos staining in ceAm, PVN, SON and BNST was successful, we didn't see any OXTR-positive staining with neither of the secondary

antibodies that we used. Thus, we concluded that the primary anti-OXTR antibody did not successfully bind to its target. Moreover, the results of the CRF-DAB staining are being analyzed. Taking into account the findings of abnormal CRF levels during stressful conditions in 5-HTT-/- mice (Jiang et al., 2009), our expectation is to reveal a positive correlation between the CRF levels and the high levels of self-grooming found in 5-HTT-/- adult animals. We also expect to see higher CRF levels in the adolescent 5-HTT-/- heterogenotype group which displayed reduced prosocial behaviour, in order to verify that anxiety plays an inhibitory role on the expression of prosocial behaviour in 5-HTT-/- animals.

5. Conclusions and future directions

From the results of our social interaction experiments we could infer that there is a GxE effect on passive and active prosocial behaviour and self-grooming during adolescence, with prosocial behaviours being reduced and self-grooming being increased during 5-HTT-/- heterogenotype interaction, reflecting a lack of social bonding between the interacting peers. On the other hand, there is also a GxE effect on pouncing in the juvenile group, making the 5-HTT-/- rats more playful during the heterogenotype interaction (reflecting social bonding), that might support the use of training with typically developing peers for improving the social deficits of autistic children. Active prosocial behaviour is also higher during heterogenotype interaction for juveniles, which suggests that behaviour during this period is plastic. With reference to the rGE effect, we have an indication that there might be an rGE effect of the 5-HTT-/- genotype on the active prosocial behaviour of the partner and of the 5-HTT+/+ genotype on the self-grooming of the partner.

As a continuation of this study we would like to repeat the CRF-positive cell quantifications and analyze again how the CRF levels correlate with self-grooming and lack of bonding. We are also planning to use a different primary anti-OXTR antibody to repeat the immunofluorescence staining. Our expectation is that higher OXTR-c-Fos co-localization in the social regions of the brain is positively correlated with social bonding (increased prosocial behaviour and social play) and decreased anxiety (lower levels of self-grooming). Furthermore, it would be interesting to make correlations between the occurrence of different behaviours, expressed by the interacting partners, for example, see whether the higher self grooming displayed by the adolescent 5-HTT-/- heterogenotype group is negatively correlated with the occurrence of prosocial behaviour expressed by their 5-HTT+/+ partners. It would be also very informative and important to repeat the social interaction paradigm, with the addition of OXT administration and test how this pharmacological manipulation would affect social bonding in 5-HTT-/- and 5-HTT+/+ rats.

6. References

Altamura, C., Dell'Acqua, M. L., Moessner, R., Murphy, D. L., Lesch, K. P., & Persico, A. M. (2007). Altered neocortical cell density and layer thickness in serotonin transporter knockout mice: a quantitation study. *Cerebral Cortex*, *17(6)*, 1394-1401.

Baenninger, L. P. (1967). Comparison of behavioural development in socially isolated and grouped rats. *Animal Behaviour*, *15(2)*, 312-323.

Bartz, J. A., Zaki, J., Bolger, N., & Ochsner, K. N. (2011). Social effects of oxytocin in humans: context and person matter. *Trends in Cognitive Sciences*, *15(7)*, 301-309.

Belsky, J., Jonassaint, C., Pluess, M., Stanton, M., Brummett, B., & Williams, R. (2009). Vulnerability genes or plasticity genes&quest. *Molecular Psychiatry*, *14(8)*, 746-754.

Bengel, D., Murphy, D. L., Andrews, A. M., Wichems, C. H., Feltner, D., Heils, A., ... & Lesch, K. P. (1998). Altered brain serotonin homeostasis and locomotor insensitivity to 3, 4-methylenedioxymethamphetamine ("Ecstasy") in serotonin transporter-deficient mice. *Molecular Pharmacology*, *53(4)*, 649-655.

Canli, T., Omura, K., Haas, B. W., Fallgatter, A., Constable, R. T., & Lesch, K. P. (2005). Beyond affect: a role for genetic variation of the serotonin transporter in neural activation during a cognitive attention task. *Proceedings of the National Academy of Sciences of the United States of America*, *102(34)*, 12224-12229.

Caspi, A., Sugden, K., Moffitt, T. E., Taylor, A., Craig, I. W., Harrington, H., ... & Poulton, R. (2003). Influence of life stress on depression: moderation by a polymorphism in the 5-HTT gene. *Science*, *301(5631)*, 386-389.

De Jong, T. R., Veening, J. G., Olivier, B., & Waldinger, M. D. (2007). Oxytocin Involvement in SSRI-Induced Delayed Ejaculation: A Review of Animal Studies. *The Journal of Sexual Medicine*, *4(1)*, 14-28.

Dölen, G., Darvishzadeh, A., Huang, K. W., & Malenka, R. C. (2013). Social reward requires coordinated activity of nucleus accumbens oxytocin and serotonin. *Nature*, *501(7466)*, 179-184.

Emiliano, A. B., Cruz, T., Pannoni, V., & Fudge, J. L. (2006). The interface of oxytocin-labeled cells and serotonin transporter-containing fibers in the primate

hypothalamus: a substrate for SSRIs therapeutic effects? *Neuropsychopharmacology*, *32(5)*, 977-988.

Hariri, A. R., & Holmes, A. (2006). Genetics of emotional regulation: the role of the serotonin transporter in neural function. *Trends in Cognitive Sciences*, *10(4)*, 182-191.

Hariri, A. R., Mattay, V. S., Tessitore, A., Kolachana, B., Fera, F., Goldman, D., ... & Weinberger, D. R. (2002). Serotonin transporter genetic variation and the response of the human amygdala. *Science*, *297(5580)*, 400-403.

Harper, C. B., Symon, J. B., & Frea, W. D. (2008). Recess is time-in: Using peers to improve social skills of children with autism. *Journal of Autism and Developmental Disorders*, *38(5)*, 815-826.

Heils, A., Teufel, A., Petri, S., Stöber, G., Riederer, P., Bengel, D., & Lesch, K. P. (1996). Allelic variation of human serotonin transporter gene expression. *Journal of Neurochemistry*, *66(6)*, 2621-2624.

Heinrichs, M., Baumgartner, T., Kirschbaum, C., & Ehlert, U. (2003). Social support and oxytocin interact to suppress cortisol and subjective responses to psychosocial stress. *Biological Psychiatry*, *54(12)*, 1389-1398.

Heinz, A., Braus, D. F., Smolka, M. N., Wrase, J., Puls, I., Hermann, D., ... & Büchel, C. (2004). Amygdala-prefrontal coupling depends on a genetic variation of the serotonin transporter. *Nature Neuroscience*, *8(1)*, 20-21.

Holmes, A., & Hariri, A. R. (2003). The serotonin transporter gene-linked polymorphism and negative emotionality: placing single gene effects in the context of genetic background and environment. *Genes, Brain and Behavior*, *2(6)*, 332-335.

Holmes, A., Murphy, D. L., & Crawley, J. N. (2002). Reduced aggression in mice lacking the serotonin transporter. *Psychopharmacology*, *161(2)*, 160-167.

Homberg, J. R., Olivier, J. D. A., Smits, B. M. G., Mul, J. D., Mudde, J., Verheul, M., ... & Cuppen, E. (2007a). Characterization of the serotonin transporter knockout rat: a selective change in the functioning of the serotonergic system. *Neuroscience*, *146(4)*, 1662-1676.

Homberg, J. R., Pattij, T., Janssen, M. C., Ronken, E., De Boer, S. F., Schoffelmeer, A. N., & Cuppen, E. (2007b). Serotonin transporter deficiency in rats improves inhibitory control but not behavioural flexibility. *European Journal of Neuroscience*, *26(7)*, 2066-2073.

Homberg, J. R., Schiepers, O. J., Schoffelmeer, A. N., Cuppen, E., & Vanderschuren, L. J. (2007c). Acute and constitutive increases in central serotonin levels reduce social play behaviour in peri-adolescent rats. *Psychopharmacology*, *195(2)*, 175-182.

Homberg, J. R., van den Bos, R., den Heijer, E., Suer, R., & Cuppen, E. (2008). Serotonin transporter dosage modulates long-term decision-making in rat and human. *Neuropharmacology*, *55(1)*, 80-84.

Jaffee, S. R., & Price, T. S. (2007). Gene–environment correlations: a review of the evidence and implications for prevention of mental illness. *Molecular Psychiatry*, *12(5)*, 432-442.

Jiang, X., Wang, J., Luo, T., & Li, Q. (2009). Impaired hypothalamic-pituitary-adrenal axis and its feedback regulation in serotonin transporter knockout mice. *Psychoneuroendocrinology*, *34(3)*, 317-331.

Jørgensen, H., Riis, M., Knigge, U., Kjaer, A., & Warberg, J. (2003). Serotonin receptors involved in vasopressin and oxytocin secretion. *Journal of Neuroendocrinology*, *15(3)*, 242-249.

Kalueff, A. V., Aldridge, J. W., LaPorte, J. L., Murphy, D. L., & Tuohimaa, P. (2007a). Analyzing grooming microstructure in neurobehavioral experiments. *Nature Protocols*, *2(10)*, 2538-2544.

Kalueff, A. V., Fox, M. A., Gallagher, P. S., & Murphy, D. L. (2007b). Hypolocomotion, anxiety and serotonin syndrome-like behavior contribute to the complex phenotype of serotonin transporter knockout mice. *Genes, Brain and Behavior*, *6(4)*, 389-400.

Kalueff, A. V., Ren-Patterson, R. F., & Murphy, D. L. (2007c). The developing use of heterozygous mutant mouse models in brain monoamine transporter research. *Trends in Pharmacological Sciences*, *28(3)*, 122-127.

Kalueff, A. V., & Tuohimaa, P. (2004). Grooming analysis algorithm for neurobehavioural stress research. *Brain Research Protocols*, *13(3)*, 151-158.

Kiser, D., Steemers, B., Branchi, I., & Homberg, J. R. (2012). The reciprocal interaction between serotonin and social behaviour. *Neuroscience & Biobehavioral Reviews*, *36(2)*, 786-798.

Larsen, P. J., Hay-Schmidt, A., Vrang, N., & Mikkelsen, J. D. (1996). Origin of projections from the midbrain raphe nuclei to the hypothalamic paraventricular nucleus in the rat: a combined retrograde and anterograde tracing study. *Neuroscience*, *70(4)*, 963-988.

Lee, R., Garcia, F., van de Kar, L. D., Hauger, R. D., & Coccaro, E. F. (2003). Plasma oxytocin in response to pharmaco-challenge to D-fenfluramine and placebo in healthy men. *Psychiatry Research*, *118(2)*, 129-136.

Lewejohann, L., Kloke, V., Heiming, R. S., Jansen, F., Kaiser, S., Schmitt, A., ... & Sachser, N. (2010). Social status and day-to-day behaviour of male serotonin transporter knockout mice. *Behavioural Brain Research*, *211(2)*, 220-228.

Li, Q., Holmes, A., Ma, L., Van de Kar, L. D., Garcia, F., & Murphy, D. L. (2004). Medial hypothalamic 5-hydroxytryptamine (5-HT) 1A receptors regulate neuroendocrine responses to stress and exploratory locomotor activity: application of recombinant adenovirus containing 5-HT1A sequences. *The Journal of Neuroscience*, *24(48)*, 10868-10877.

Li, Q., Wichems, C., Heils, A., Van de Kar, L. D., Lesch, K. P., & Murphy, D. L. (1999). Reduction of 5-hydroxytryptamine (5-HT) 1A-mediated temperature and neuroendocrine responses and 5-HT1A binding sites in 5-HT transporter knockout mice. *Journal of Pharmacology and Experimental Therapeutics*, *291(3)*, 999-1007.

Lira, A., Zhou, M., Castanon, N., Ansorge, M. S., Gordon, J. A., Francis, J. H., ... & Gingrich, J. A. (2003). Altered depression-related behaviors and functional changes in the dorsal raphe nucleus of serotonin transporter-deficient mice. *Biological Psychiatry*, *g*, 960-971.

Marazziti, D., Baroni, S., Giannaccini, G., Betti, L., Massimetti, G., Carmassi, C., & Catena-Dell'Osso, M. (2012). A link between oxytocin and serotonin in humans: Supporting evidence from peripheral markers. *European Neuropsychopharmacology*, *22(8)*, 578-583.

Meyer-Lindenberg, A., Domes, G., Kirsch, P., & Heinrichs, M. (2011). Oxytocin and vasopressin in the human brain: social neuropeptides for translational medicine. *Nature Reviews Neuroscience*, *12(9)*, 524-538.

Montag, C., Fiebach, C. J., Kirsch, P., & Reuter, M. (2011). Interaction of 5-HTTLPR and a variation on the oxytocin receptor gene influences negative emotionality. *Biological Psychiatry*, *69(6)*, 601-603.

Murakami, F., Shimomura, T., Kotani, K., Ikawa, S., Nanba, E., & Adachi, K. (1999). Anxiety traits associated with a polymorphism in the serotonin transporter gene regulatory region in the Japanese. *Journal of Human Genetics*, *44(1)*, 15-17.

Murphy, D. L., Li, Q., Engel, S., Wichems, C., Andrews, A., Lesch, K. P., & Uhl, G. (2001). Genetic perspectives on the serotonin transporter. *Brain Research Bulletin*, *56(5)*, 487-494.

Niesink, R. J., & Van Ree, J. M. (1989). Involvement of opioid and dopaminergic systems in isolation-induced pinning and social grooming of young rats. *Neuropharmacology*, *28(4)*, 411-418.

Olivier, J. D. A., Van Der Hart, M. G. C., Van Swelm, R. P. L., Dederen, P. J., Homberg, J. R., Cremers, T., ... & Ellenbroek, B. A. (2008). A study in male and female 5-HT transporter knockout rats: an animal model for anxiety and depression disorders. *Neuroscience*, *152(3)*, 573-584.

Paxinos, G. (1994). The rat nervous system. Fourth edition, 32 Jamestown Road, London NW1 7BY, UK.

Pezawas, L., Meyer-Lindenberg, A., Drabant, E. M., Verchinski, B. A., Munoz, K. E., Kolachana, B. S., ... & Weinberger, D. R. (2005). 5-HTTLPR polymorphism impacts human cingulate-amygdala interactions: a genetic susceptibility mechanism for depression. *Nature Neuroscience*, *8(6)*, 828-834.

Poole, T. B., & Morgan, H. D. R. (1975). Aggressive behaviour of male mice (Mus musculus) towards familiar and unfamiliar opponents. *Animal Behaviour*, *23*, 470-479.

Reneman, L., Schilt, T., de Win, M. M., Booij, J., Schmand, B., van Den Brink, W., & Bakker, O. (2006). Memory function and serotonin transporter promoter gene polymorphism in ecstasy (MDMA) users. *Journal of Psychopharmacology*, *20(3)*, 389-399.

Schwarting, R.K., Jegan, N., & Wohr, M. (2007). Situational factors, conditions and individual variables which can determine ultrasonic vocalizations in male adult Wistar rats, *Behavioural Brain Research*, *182(2)*, 208-222.

Serretti, A., Mandelli, L., Lorenzi, C., Landoni, S., Calati, R., Insacco, C., & Cloninger, C. R. (2006). Temperament and character in mood disorders: influence of DRD4, SERTPR, TPH and MAO-A polymorphisms. *Neuropsychobiology*, *53(1)*, 9-16.

Smits, B. M., Mudde, J. B., van de Belt, J., Verheul, M., Olivier, J., Homberg, J., ... & Cuppen, E. (2006). Generation of gene knockouts and mutant models in the laboratory rat by ENU-driven target-selected mutagenesis. *Pharmacogenetics and Genomics*, *16(3)*, 159-169.

Striepens, N., Kendrick, K. M., Maier, W., & Hurlemann, R. (2011). Prosocial effects of oxytocin and clinical evidence for its therapeutic potential. *Frontiers in Neuroendocrinology*, *32(4)*, 426-450.

Swaab, D. F., Fliers, E., Hoogendijk, W. J. G., Veltman, D. J., & Zhou, J. N. (2000). Interaction of prefrontal cortical and hypothalamic systems in the pathogenesis of depression. *Progress in Brain Research*, *126*, 369-396.

Takayanagi, Y., Yoshida, M., Bielsky, I. F., Ross, H. E., Kawamata, M., Onaka, T., ... & Nishimori, K. (2005). Pervasive social deficits, but normal parturition, in oxytocin receptor-deficient mice. *Proceedings of the National Academy of Sciences of the United States of America*, *102(44)*, 16096-16101.

Uvnäs-Moberg, K., Björkstrand, E., Hillegaart, V., & Ahlenius, S. (1999). Oxytocin as a possible mediator of SSRI-induced antidepressant effects. *Psychopharmacology*, *142(1)*, 95-101.

Van IJzendoorn, M. H., & Bakermans-Kranenburg, M. J. (2012). A sniff of trust: meta-analysis of the effects of intranasal oxytocin administration on face recognition, trust to in-group, and trust to out-group. *Psychoneuroendocrinology*, *37(3)*, 438-443.

Wellman, C. L., Izquierdo, A., Garrett, J. E., Martin, K. P., Carroll, J., Millstein, R., ... & Holmes, A. (2007). Impaired stress-coping and fear extinction and abnormal corticolimbic morphology in serotonin transporter knock-out mice. *The Journal of Neuroscience*, *27(3)*, 684-691.

Wendland, J. R., Martin, B. J., Kruse, M. R., Lesch, K. P., & Murphy, D. L. (2006). Simultaneous genotyping of four functional loci of human SLC6A4, with a reappraisal of 5-HTTLPR and rs25531. *Molecular Psychiatry, 11(3)*, 224-226.

Wöhr, M., & Schwarting, R. K. (2013). Affective communication in rodents: ultrasonic vocalizations as a tool for research on emotion and motivation. *Cell and Tissue Research*, *354(1)*, 81-97.

Yoshida, M., Takayanagi, Y., Inoue, K., Kimura, T., Young, L. J., Onaka, T., & Nishimori, K. (2009). Evidence that oxytocin exerts anxiolytic effects via oxytocin receptor expressed in serotonergic neurons in mice. *The Journal of Neuroscience*, *29(7)*, 2259-2271.

Functional Parcellation of the Human Entorhinal Cortex

Nestor Israel Zaragoza Jimenez[1]
Supervisor: Tobias Navarro Schroeder[1], Erik van Oort[1], Prof. Dr. Christian Beckmann[1], Dr. Christian Doeller[1]

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

The entorhinal cortex (EC) is the interface between the hippocampus and the neocortex. Since the discovery of grid cells in the rodent medial EC, it has been recognized as a crucial hub for spatial coding and navigation in rodents. Traditionally, the EC has been subdivided into a medial (MEC) and a lateral (LEC) portion. Postmortem comparative anatomical evidence, however, suggests broader subdivisions among rodents, non-human primates and humans. Cytoarchitectonically, up to nine subdivisions of the human EC have been distinguished. Several cytoarchitectonic features also suggest a modular arrangement, such as periodic bundling of pyramidal cell dendrites and axons. Furthermore, tract tracing studies in rodents and non-human primates suggest a wide range of projections along the EC, with medial portions of the EC projecting to the temporal portion of the hippocampus, while lateral and caudal portions connected to the septal hippocampus. Differential connectivity pattern of EC subregions could thus lead to distinct intrinsic signal fluctuations, visible to functional magnetic resonance imaging (fMRI) at high spatial resolution. However, knowledge about the fine-grained, in-vivo fMRI-based parcellation of the EC has been scarce. Here, we aimed to parcellate the human EC based on its intrinsic signal fluctuations during a virtual navigation task combined with high-resolution, submillimetre fMRI at 7 tesla (T). We generated a functional parcellation of the EC by leveraging a novel instantaneous correlation analysis approach, termed instantaneous connectivity parcellation (ICP). The resulting parcellations showed strong correspondence to topographies from postmortem cytoarchitectonic studies. These findings allow us to test connectivity fingerprints of human EC subregions in-vivo and could open up the possibility to investigate neural mechanisms of spatial cognition at an unprecedented level of detail.

Keywords: human entorhinal cortex, parcellation, functional connectivity, fMRI

Corresponding author: N. I. Zaragoza Jimenez; **Email:** nestorizj@gmail.com

1. Introduction

The entorhinal cortex (EC) is commonly perceived as a major input and output structure of the hippocampal formation, entertaining the role of the nodal point of cortico-hippocampal circuits (Canto, Wouterlood, & Witter, 2008). Postmortem comparative anatomical evidence suggests a broad number of different subdivisions of the EC in rodents, non-human primates and humans. Cytoarchitectonically, up to nine subdivisions of the human EC have been distinguished (Krimer, Hyde, Herman, & Saunders, 1997). Brodmann divided this brain region in area 28 and 34. Braak (1972) described only three, with additional regions being defined as transitional zones. More recently, Insausti et al. (1995) distinguished eight entorhinal subareas and Krimer et al (1997) identified even 9 subregions. Nonetheless, animal studies performed mostly in rodents have subdivided the EC into two initial segments, a lateral segment, or lateral EC (LEC) and a medial segment or medial EC (MEC). Both subdivisions have extensive connectivity profiles and it has been shown (Kerr et al., 2007; O'Reilly et al., 2013; Morrissey et al., 2012) that both receive different inputs; while the LEC might be processing information from the interaction with objects, the MEC might be processing information relating to the environment (Van Cauter et al., 2013) Additionally, studies on rodents and nonhuman primates suggest a wide range of projections along the EC, with medial EC portions projecting to the temporal portion of the hippocampus (HPC), thalamus (dorsolateral and dorsocaudal thalamic nuclei), amygdala, olfactory structures, visual, parietal, cingulate and retrosplenial cortices; while lateral EC portions are connected to the septal HPC, basal ganglia, claustrum, amygdala, thalamus, (midline dorsal thalamus) as well as frontal, insular, and cingulate cortices (Kerr et al., 2007; Insausti & Amaral, 2004; Canto et al., 2008; O'Reilly et al., 2013). On top of that, projections to other cortical regions including orbitofrontal regions and the medial temporal lobe (MTL) have been described (Munoz & Insausti, 2005; Ding, 2013).

Recent observations have implicated the medial part of the EC as possible hub in the neural circuitry for spatial representation (Fyn et al., 2004; Witter & Moser, 2006). The finding of grid cells in superficial MEC layers (Boccara et al., 2010) and reports on interactions between deep and superficial layers indicates that its intrinsic organization may show different functional characteristics. It has been suggested (McNaughton et al., 2006; Burgalossi & Brecht , 2014; Witter & Moser, 2006) that such interaction may subserve a neural spatial map like representation of the subject within an environment. Subsequently, if we add the complicated extrinsic input-output organization on top of the functional intrinsic organization we will expect a detailed and widespread functional network of cortical projections. Thus, the commonly used description of EC as merely a layered input-output structure can be considered an oversimplification, given the evidence that it is a crucial region for the global functions of the cortico-hipocampal network.

Although single-cell recording studies of the EC in a variety of species have confirmed its importance for different cognitive processes, (Killian et al. 2012, Van Couter 2013) few studies have tried to understand in-vivo the relationship between its topographical arrangement and its functional characteristics (Killian et al. 2012, Libby et al. 2012). So far, few studies have tried to localize, characterize its borders and functional regions employing probability maps based on cytoarchitectural traits (Kedo et al., 2005; Schultz et al., 2012; Libby et al., 2012). However, a fine-grained characterization is still missing. Surprisingly, up to now, no one has created a parcellation of the human EC that could express the dynamics of the area. Therefore, in our study we aimed to parcellate the human EC based on its intrinsic signal fluctuations during a virtual navigation task by employing high-resolution, submillimetre functional magnetic resonance imaging (fMRI) at 7 T. We have generated a functional parcellation of the EC by leveraging a novel instantaneous correlation analysis approach, termed instantaneous connectivity parcellation (ICP). ICP provides information on how the functional connectivity within a region evolves over time and allows the creation of three-dimensional parcellations of brain regions based on their intrinsic functional properties alone. Here, we applied this novel method to the human EC. The resulting parcellations showed strong correspondence to topographies from postmortem cytoarchitectonic studies (e.g. Krimer et al, 1997). Subsequently to the parcellation, a functional connectivity analysis for every segment was performed. The emerged functional connectivity maps of the corresponding parcels permitted us to assess the functional organization of the region. The consistency and stability of our results opens new avenues for testing the topographical organization in-vivo of the human EC. Additionally, our parcellation may serve as starting point to characterize the intrinsic and functional organization in-vivo of the human EC.

2. Methodology

2.1 Participants

Twenty-six participants (15 males, 11 females) took part in this experiment (aged 19 - 36, mean age 23.3 years). Participants gave written consent and were paid for participating, as approved by the local Research Ethics Committee. All were right-handed with normal or corrected-to-normal vision and reported to be in good health. 4 participants were excluded from all further analyses because the number of frame wise displacements (Power et al., 2012) above the threshold of 0.5 mm exceeded the group average by more than 1 standard deviation.

2.2 Virtual reality environment and task

The data analyzed in our study was previously obtained by Tobias Navarro Schroeder (Abstract: SfN 2014 Grid cell representations in humans align to a common reference frame) from participants performing a virtual navigation object-location memory task. Participants navigated a virtual grassy plane surrounded by a circular wall with a black background and 12 coloured triangles serving as landmarks. During task participants collected and replaced everyday life objects within the environment.

2.3 Image acquisition

Blood-oxygenation-level-dependent (BOLD) T2*-weighted functional images were acquired on a 7T Siemens MAGNETOM scanner using a three-dimensional echo-planar imaging (EPI) pulse sequence (Poser et al., 2010): repetition time (TR) = 2756 ms, echo time (TE) = 20 ms, flip angle = 14°, slice thickness = 0.92 mm, slice oversampling = 8.3%, in-plane resolution = 0.92 mm, field of view (FoV) = 210 mm in each direction, 96 slices per slab, Phase Encoding acceleration factor = 4, 3D acceleration factor = 2.

Five dummy volumes immediately before the main scan were discarded to allow for T1 equilibration. Each scanning session was subdivided into EPI acquisition blocks of 210 volumes. The vast majority of participants performed 5 blocks (n = 20). Due to technical problems or by request of the participant, 4 participants underwent 4 blocks and 2 participants 6 blocks. In addition, a field map using a gradient echo sequence was recorded for distortion correction of the acquired EPI images. We used an MP2RAGE sequence (Marques et al., 2010) for the structural scan (resolution = 0.633 mm).

2.4 Image preprocessing

Image preprocessing was implemented using FEAT (FMRI Expert Analysis Tool, FMRI's Software Library), which is part of FSL analysis package. Functional data for every session from each participant were motion-corrected using MCFLIRT (FLS linear regression tool), correcting for apparent head motion artifacts by aligning each single image into a common reference space. MCFLIRT was also used to plot the six vectors representing the translation and rotation on different dimensions (X, Y and Z). The functional data were moderately smoothed with a full width at half maximum (FWHM) kernel of 2.5 mm to increase the robustness of the data denoising with FSL FIX (see below). This smoothing kernel maintains anatomical precision to investigate subregions of the 25 - 30 mm2 wide and up to 4.7 mm thick human entorhinal cortex (Krimer et al., 1997). A high-pass temporal filtering was applied to remove low frequency artifacts.

Additionally, the images were normalized to a standard space using FNIRT (FMRI non-linear registration tool) which in the present study was the Montreal Neurological Institute (MNI) 1 mm standard template. Finally, an independent component analysis (ICA) decomposition was created with MELODIC (Multivariate Exploratory Linear Optimized into Independent Components) to separate artifactual spatial components from the rest (Beckmann & Smith, 2005).

In order to increase the signal to noise ratio we performed a data denoising through FIX (FMRIB's ICA-based Xnoiseifier- part of the FSL software package). FIX tends to auto classify components into "good" and "bad" components, so the bad components can be removed from the data. For this purpose, noise components were identified manually from the data of 10 subjects following the procedure suggested by Kelly and colleagues (Kelly et al., 2010). This served as training data for a classifier algorithm used by FIX. Once the classifier was created, it was tested within the same ten subjects to verify its efficiency. FIX was then applied to the rest of the preprocessed data to remove artifactual components automatically. The denoised data were used in all subsequent analyses.

2.5 Instantaneous correlations parcellation

The instantaneous correlations parcellation (ICP) is a recently developed method (Van Oort et al., 2014, in preparation) that provides information on how the connectivity within a region evolves over time. For its implementation, firstly a region of interest (ROI) has to be provided, and then the average time course of the ROI is subtracted by combining the original preprocessed data in a spatial regression.

Normally, something like a seed based correlation analysis gives a single (time averaged) value to the connectivity between two voxels/regions. This average hides dynamic information; information that can be used to create a correlation map which may result in a really sensitive and dynamic parcellation.

The ICP method takes the average time course of a region selected for parcellation as the reference time course. A correlation map of this time course will closely resemble the selected region. This means that the dynamic information represented by the instantaneous correlations between this reference and the voxels of the selected region only contains that aspect of the data which contributes to the delineation of the region. This emphasises subtle differences within the region, and de-emphasises outside influences.

For a better understanding, we may start with the definition of the Pearson correlation between two time courses

$$\rho_{x,y} = \frac{E[X(n-\mu_x)Y(n-\mu_y)]}{\sigma_x \sigma_y}$$

with the mean and the standard deviation of each time course. In this example x can be the time course of a seed voxel, and y the time course of some target voxel. If time course x and y are normalised, this changes to

$$\rho_{x,y} = \frac{1}{T}\sum_{n=1}^{T} x(n)y(n)$$

where T is equal to the number of time points in the measurement. This means that the Pearson correlation is equal to the temporal mean of the element-wise product between two normalised time courses. One way to include the dynamic aspect of this is to describe this Pearson correlation as a cumulative process, for instance the vector defined by

$$\rho_{x,y}(1) = x(1)y(1)$$
$$\rho_{x,y}(2) = \frac{1}{2}x(1)y(1) + x(2)y(2)$$
$$\rho_{x,y}(3) = \frac{1}{3}x(1)y(1) + x(2)y(2) + x(3)y(3)$$
$$\vdots$$
$$\rho_{x,y}(T) = \frac{1}{T}x(1)y(1) + x(2)y(2) + x(3)y(3)... + x(T)y(T)$$

shows the evolution of the Pearson correlation across time by including more time points. The correlation at time point T is formed by the data between time points 1 and T. This is of limited use, as the effects at later time points only have a minimal effect on the outcome because of the averaging. We can also express this as

$$\begin{pmatrix} \rho_{x,y}(1) = x(1)y(1) \\ \rho_{x,y}(2) = x(2)y(2) \\ \rho_{x,y}(3) = x(3)y(3) \\ \vdots \\ \rho_{x,y}(T) = x(T)y(T) \end{pmatrix}$$

which contains an instantaneous correlation at each time point between the time course x and the target time course y. This vector contains all the dynamic information which cumulatively results in the static correlation at the end of the scan.

Consider the following situation: we have an area of interest with four different substructures, each of which with different time course coming from a voxel within these four substructures. These four time courses are nearly identical, however, every single one presents transient events at different time points; assuming that each of the substructures showed spatially structured temporal events, which may be represented by these transient events, a simple correlation can not distinguish between time courses, as they are almost identical. A Pearson correlation analysis will not be able to distinguish the differences between time courses because they would probably have the same correlation coefficient with each other and with the mean time course. Now, if we take the same four time courses and instead of applying a Pearson correlation analysis we move towards ICP analysis, the effect of the transient events would be potentiated, allowing us to distinguish and better define information coming from the different time courses.

Thus, this procedure highlights subtle differences among time courses within the ROI. If such differences behave in a structured manner the parcellation is created by implementation of a group independent component analysis (gICA). The number of components set by the gICA is

equivalent to the number of segments created within the parcellation. To identify which of these segmentations is the most reliable, a split-half reproducibility test is performed to diagnose which of the created parcellations is the most stable given the data. Here we refer to the number of segments as the scale of the parcellation.

In such a way, the parcellation created is based on the internal network connectivity between substructures within the ROI, meaning that every voxel contains dynamic information, which contributes to the region's delimitation.

For our study, a binary mask of the human EC was created in the space of the group-specific EPI template based on anatomical landmarks according to Insausti et al. (1998) and Franko et al. (2012).

The resulting chosen parcellation for a seed based correlation analysis was scale 9. This decision was based upon topographical correspondence with previous cytological studies, the results of the split-half reproducibility analysis, and literature reports (Fig 4). The analysis pipeline of the study is presented in Figure 1.

2.6 Seed based correlation analysis

Based on the time series of a seed voxel or ROI, connectivity is calculated as the correlation of time series for all other voxels in the brain in a dual-regression analysis. The result of a seed based correlation analysis (SBCA) is a connectivity map showing z-scores for each voxel indicating how well its time series correlates with the time series of the seed. With this approach statistical maps are generated by calculating the correlation of the time courses of all grey matter voxels with the time courses of all nine segments of the entorhinal cortex parcellation created by the ICP analysis in order to find a temporally consistent active network. We employed SBCA Tool (part of the FSL software) to execute the analysis. Individual SBCA was performed first on every session of every participant for each parcel separately and the resulting maps of correlation coefficients were Fisher-Z transformed for variance stabilisation. Secondly, the obtained correlation maps were fed into the RANDOMISE tool (part of FSL software for MRI data analysis) to test for common functional connectivity maps.

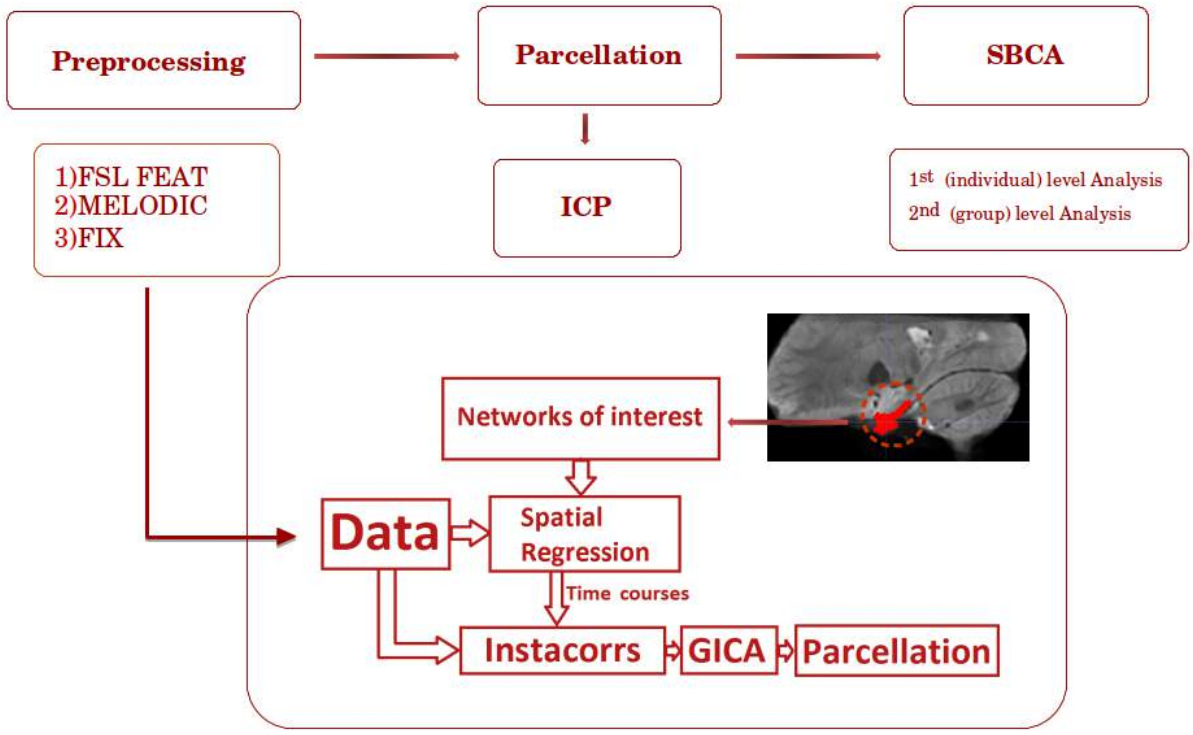

Fig. 1 Instantaneous correlations analysis pipeline: a mask of the region/network of interest along with the data set are fed into the analysis, a spatial regression is performed to obtain the mean time course from the region/networks of interest and the data sets. A group ICA is then used to create the parcellation based on the intrinsic dynamic explicitly expressed by the instantaneous correlation analysis.

RANDOMISE is a permutation program enabling modelling and inference using standard GLM design setup. RANDOMISE produces a test statistic image and sets of p-value images. The p-value images obtained were corrected for multiple comparison employing familywise error (FWE) correction at $\alpha = 0.05$.

3. Results

3.1 Human entorhinal cortex parcellation

In 1909, Brodmann labeled the EC as areas 28 and 34. Braak (1972) describe only three, with additional regions being defined as transitional zones. More recently Insausti et al. (1995) distinguished eight entorhinal subareas and Krimer et al. (1997) identified even 9 subregions.

The results of ICP are shown in Figure 2 and 4 and supplementary material*, split-half results are also shown in Figure 3. The highest peaks are observed at scales 2, 8, and 15. Nonetheless, in order to compare the most recent histological (Krimer et al., 1997) parcellation to our own, we selected scale number 9 of our series of parcellations. This scale widely resembles such topographical subdivision. Images of the parcellation at scale nine can be seen in Figure 2.

For a detailed description of our parcellation we have taken as a comparative point the nomenclature and descriptions made by Krimer et al. (1997), based on cyto-myeloarchitecture. These authors identified nine segments: a rostral area, or prorhinal area (Pr), followed by a rostral-lateral area (L), an intermediate area (I) that is subdivided in three more subareas-a superior area (Is), a rostral area (Ir) and a caudal area (Ic), and finally the more caudal, medial area (M) that is also subdivided into two other subregions, a rostral area (Mr) and a caudal area (Mc). In this division, all the areas bordered laterally with the sulcal subdivision (S). For comparison, we continue to employ this nomenclature to demonstrate the similarity in topographical arrangement found in our parcellation and the parcellation proposed by Krimer et al. (1997). As can be noticed from inspection of Figure 2 and 4, the global structure of the parcellation is highly symmetrical between hemispheres, with exception of an extra segment on the left hemisphere. Nonetheless, the right hemisphere presents a ventrolateral elongated segment that cross from the caudal to the rostral section. In the left hemisphere, this homologous segment runs only until the middle part of the

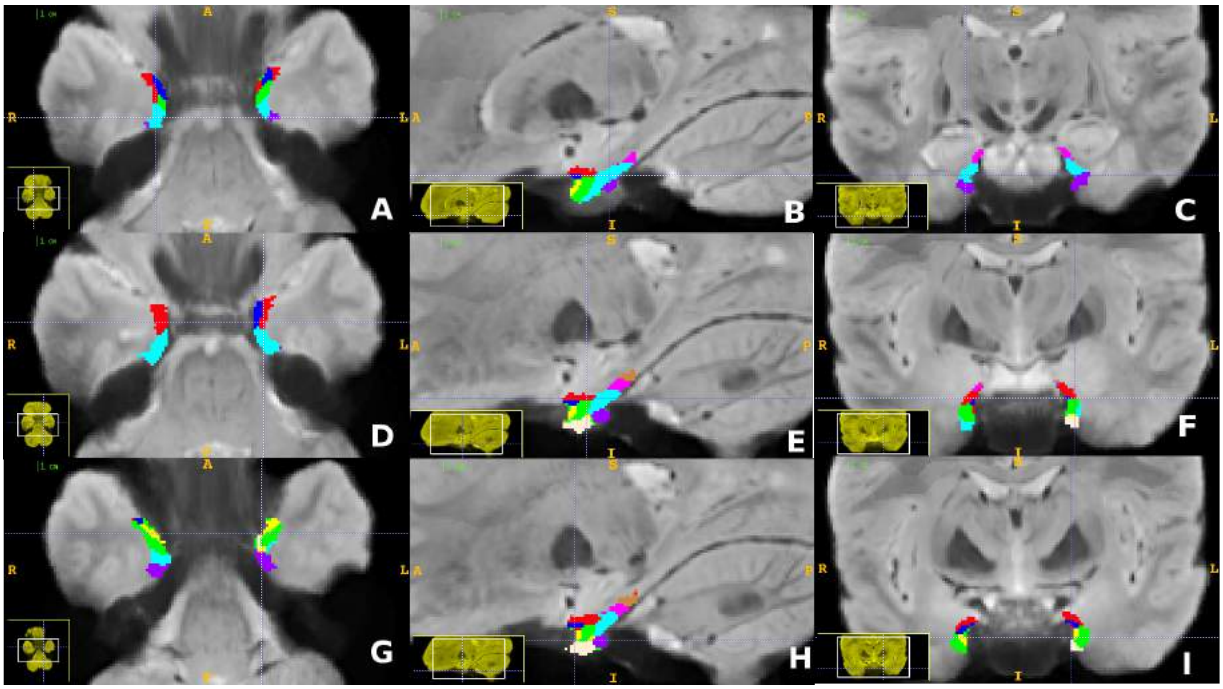

Fig. 2 Parcellation of the Human Entorhinal cortex: Images were taken from scale number 9 of the parcellation at different spatial points. Images **A**, **D**, and **G** are three different axial views. Similarly, images **B**, **E**, and **H** display three different sagittal views; whereas images **C**, **F**, and **I** are three different coronal views. Note the symmetry between hemispheres; this symmetry is conserved at every spatial point of the parcellation. The different segments are colour coded for its better appreciation. The degree of symmetry suggests the internal signal fluctuation in both hemispheres is homologous. This could also suggest that the type of internal information processed by the EC in both hemispheres is organized in a similar way.

* Supplementary material: http://www.ru.nl/master/cns/journal/archive/volume-10-issue-1/print-edition/

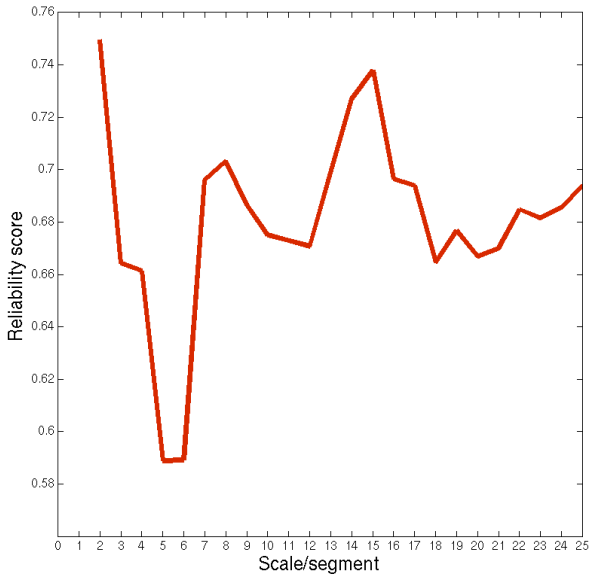

Fig. 3 Split-half reproducibility - the scale corresponds to the number of subdivisions of the ROI. Peaks can be seen at scale numbers 2, 8 and 15.

parcellation; topographically this segment may be the homologue of the lateral sulcal subdivision (light blue segment) of Krimer's parcellation. Furthermore, four other symmetrical segments can be seen dorsal to the ventral elongated segment. These four segments are located in the medial portion of the parcellation, topographically these segments appear to be candidates as homologues of areas L (green segment), Pr (yellow segment) and Ic, Ir (both in blue) and Is (red segment) respectively. Similarly on the most caudal part two other segments can be observed. These two segments appear to be corresponding to segments Mr and Mc. Finally, on the most ventral section of the left hemisphere's parcellation, two more segments can be seen, on the right hemisphere only one can be seen. These segments are likely to have captured signals from the perirhinal cortex, because the precise, histologically defined border between entorhinal and perirhinal cortex varies across subjects by several millimetres on the medio-lateral axis (Amunts et al., 2005). Therefore, we consider these two segments as non entorhinal but rather segments that belongs to perirhinal areas. Figure 4 shows the topographical similarities between the scale 9 ICP and the histological topography described by Krimer et al. (1997). A considerable correspondence between the in-vivo, fMRI-based ICP parcellation and the post mortem, histological staining based segmentation is immediately apparent.

3.2 Human entorhinal cortex functional connectivity

The T-value maps obtained by the RANDOMISE tool revealed different functional connectivity patterns among the different parcellations. Overall, every segment displayed connectivity with a variety of brain regions. Consistency to areas described from the literature (Muñoz, M., & Insausti, R., 2005; Kerr et al., 2007; Kahn et al., 2008) can be appreciated. Clear projection to the hippocampus, thalamus, and orbitofrontal cortex can be appreciated on Figure 5. Remarkably, every connectivity profile demonstrated that each area has extensive networks across the whole neocortex and other subcortical areas.

4. Discussion and conclusions

In the current study, we applied a recently developed data-driven parcellation strategy that exploits the fluctuations on the MRI signal to create an anatomically reliable segmentation of the human EC. To our knowledge, the resulting parcellation is the first fine-grained parcellation of the human entorhinal cortex created employing of 7T fMRI

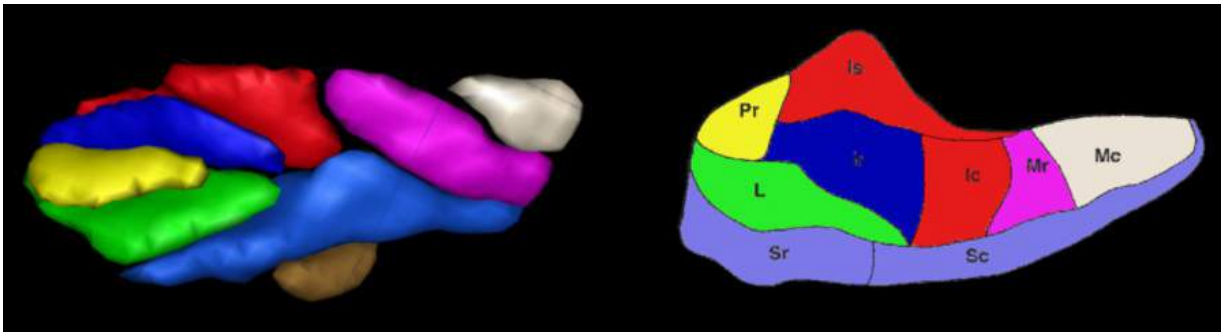

Fig. 4 3D rendering of the parcellation of the human entorhinal cortex. On the left the EC parcellation; on the right a 2D map of cytoarchitectonic subregions of the human EC identified by Krimer et al. (1997). The subregions have been colour coded to highlight the topographical arrangement and the potentially matching parcels of the fMRI-based analysis. Note the correspondence in orientation and shape of most subdivisions.

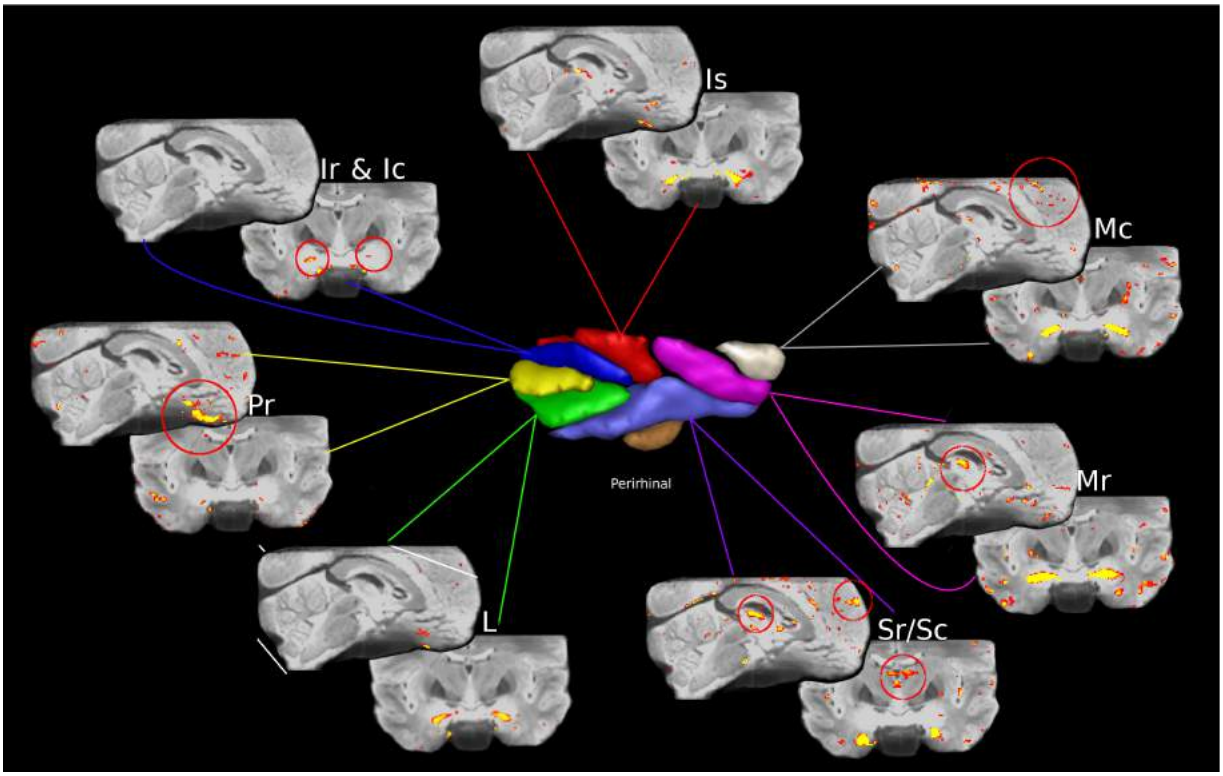

Fig. 5 Examples of the different connectivity profiles emerged from each segment of the parcellation. At the centre of the image, a 3D rendering of the right EC hemisphere parcellation is shown. Sagittal and coronal views of the connectivity profiles are connected to each segment by colour coded lines exemplifying thus the difference observed in functional connectivity profiles across segments of our parcellation. Labels Is, Ir & Ic, Pr, Sr/SC, Mr, Mc exemplify the potential topographical correspondence with Krimer's (1997) parcellation. Consistency of connections to both hippocampus can be seen from segment Ir and Ic, to thalamus from segment Sr/Sc and orbitofrontal cortex from Pr among others. FWE was used for multiple comparison correction. Image threshold $p<.05$.

data. The literature (Krimer et al., 1997; Kedo et al., 2005; Witter & Moser, 2006; Kerr et al., 2007; Cato et al., 2008; Burgalossi & Brecht, 2014) describes the EC as multi segmented unit. The parcellation observed in Figure 2 and 4 present a high topographic overlap from what has been reported by Krimer et al. (1997) employing cytological methods in postmortem human brain samples. Moreover, the organization of the parcellation appears highly structured. For example, over both hemispheres it is largely symmetrical. The conformity of shape, organization and orientation of the segments between the two hemispheres, suggest that they represent homologous brain regions. Additionally, the present subdivision suggests the existence of a broadly defined functional regions in the EC which themselves contain particular functional connectivity profiles as shown in Figure 5. This suggests a high organizational degree of the structure. The organization found in our parcellation and its resemblance to the topographical arrangement described will allows us to explore new approaches to uncover how the EC is interacting with different regions of the brain. The study of our parcellation could open new avenues to comprehend the internal organisation of the EC in-vivo during different cognitive processes.

The functional connectivity results are broadly consistent with postmortem tracing studies in human and non-human primate brains (Insausti et al., 2008; O'Reilly et al., 2013); especially those segments projecting to different sections of the frontal cortex, hippocampus, presubiculum and perirhinal cortex (Kerr et al. 2007). Just as the major airport of a great city welcomes and distributes passengers to and from different destinations, the entorhinal cortex receives and projects information across multiple areas in the human brain. Confirming thus its pivotal role as processor of a variety of information. The functional connectivity patterns encountered may be related to specific functional roles. Normally, studies on the cytoarchitectonic organization of the EC do not investigate functional correlates of the activity in the respective subregions. Hitherto, knowledge about the functional specificity of EC subregions is therefore scarce. Although interest in the EC has increased since the discovery of the grid cells in the MEC of the rat, a comprehensive

understanding of its functional subdivisions is still incomplete. An interesting avenue of future study would be an independent study of every segment of the parcellation. This could provide new insight in how information is processed in the EC and related to other cortices. Additionally, it could also help to enhance our understanding of the neural circuits for spatial representation employing the functional characteristics of every segment and the flux of visual information via the EC towards hippocampal subfields. Moreover, the distinct functional connectivity profiles can be useful to expose the areas that are homologues to the rodent MEC and LEC in the human brain. It has often been assumed that the human brain possess functional homologue areas to MEC and LEC of the rodent brain, interestingly, so far the few studies (Schultz et al., 2012; Libby et al., 2012, Ranganath et al., 2012) have addressed such disambiguation. In these uncharted terrains the parcellation can serve as a starting model for such connotation. With no one-to-one topographical correspondence but with comparable functional overlap, the corresponding areas could be described based on their functional properties. A characterization of the homologue MEC and LEC areas in the human brain is pivotal to increase our understanding of the networks involved in memory processes.

Additionally, given the amount of neurodegenerative diseases in which the EC has been involved (e.g. Alzheimer's disease, Schizophrenia); the wide variety of functional connectivity profiles could also be employed to monitor the possible changes of EC dynamics. Comparative studies between, for example, resting-state-functional connectivity and task fMRI functional connectivity are needed to establish a association between connectivity profiles of the segments in our parcellation and the progress of a particular disease. For example, Li and colleagues (2002) demonstrated that resting-state-functional connectivity between right and left HPC and the rest of the cortex fluctuates in patients with Alzheimer's disease (AD) in comparison with healthy controls. A similar approach could be employed using our parcellation to analyse changes on connectivity patterns within segments of the EC in patients with AD, schizophrenia or another neurodegenerative disease.

To avoid misinterpretation of the results, there are potential limitations that should be considered. Firstly, the precise relationship between the segments of the parcellation and the functional connectivity in task fMRI activity remains unclear. Although the topographical organisation obtained in this study resembles the cytoarchitectonic parcellation made by Krimer et al. (1997), this does not mean that the modules encountered have a linear correspondence with those from postmortem studies. Comparative interspecies studies using the same methodology are needed to characterise the similarities in topography across species. The methods of data preprocessing and analysis described in this study could be adapted for the use in non-human animals. Secondly, we found that a considerable amount of noise can compromise the stability of the parcellation creating thus possible misinterpretations. Whilst split-half analysis is run to determine which scale is the most reliable, we observed considerable switches in shape and distribution of the internal segments when we included participants who presented a high degree of movement during the fMRI data acquisition. Additionally, individual differences in shape, position and boundary of the EC may compromise the optimal performance of the method. For example, the exact position of the lateral border of the histologically defined entorhinal cortex may vary with respect to the fundus of this sulcus by several millimetres (Amunts et al., 2005). This border constitutes the transition to the perirhinal cortex anteriorly. The rather liberal ROI used in the present study is therefore very likely to capture signal from a part of the perirhinal cortex, as appreciated in Figure 4. The extra segment underneath Sr/Sc segment it is probably not part of EC but likely part of perirhinal cortex. Thirdly, with regard to the functional connectivity, despite the fact that we found different connectivity profiles between segments, the lack of the parietal lobe from the scan volume represents a disadvantage. In order to indeed corroborate our findings and also to draw wider conclusions about the functional connectivity, whole brain sampling is required. This would provide a complete overview of the connectivity profiles displayed by the segments of our parcellation. Additionally, there are chances that remained noise has affected our connectivity results. In a detailed examination of the connectivity profiles we noticed activity of voxels on different biological edges, for example between cerebellum and occipital lobe. We attribute this effect to possible physiological noise as e.g. effect arising from arteries and other blood vessel gather in the area, as well as remained artifact noise e.g. movement.

Nonetheless, the fact that we found consistent activity and stability at high thresholds (FWE-corrected $p < 0.005$) across areas (HPC, thalamus, and orbitofrontal cortex) indicates the functional connectivity profiles are highly reliable.

Thus, our parcellation provides information about

the likely underlying topographical organisation of the human entorhinal cortex in vivo. Additionally, the distinct connectivity patterns encountered per segment are suggestive of particular functional roles. The particular study of the interaction among segments would expand our knowledge about the functional dynamics of the human EC.

In conclusion, this parcellation has much strength including the consistency with the topographical arrangement suggested by cytoarchitecture studies; therefore it could have the potential to be a valuable tool for localising effects of fMRI studies in the future.

5. Acknowledgements

I would like to thank all the people involved in this project. Tobias Navarro Schroeder and Erik van Oort for their time, patience and of course for all the knowledge shared during this whole year. Christian Doeller and Christian Beckmann for their supervision and time. Thank you.

The abstract of this work has been accepted for a poster presentation at the annual meeting of the Society of Neuroscience Washington D.C. 2014.

The mathematical principles underlying the ICP method have been reproduced with permission from Van Oort et al., 2014: Instantaneous correlation parcellation; in preparation.

6. References

Amunts, K., Kedo, O., Kindler, M., Pieperhoff, P., Mohlberg, H., Shah, N. J., … Zilles, K. (2005). Cytoarchitectonic mapping of the human amygdala, hippocampal region and entorhinal cortex: intersubject variability and probability maps. *Anatomy and Embryology*, *210*(5-6), 343-52.

Beckmann, C. F., DeLuca, M., Devlin, J. T., & Smith, S. M. (2005). Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *360*(1457), 1001-13.

Boccara, C. N., Sargolini, F., Thoresen, V. H., Solstad, T., Witter, M. P., Moser, E. I., & Moser, M.B. (2010). Grid cells in pre- and parasubiculum. *Nature Neuroscience*, *13*(8), 987-94.

Burgalossi, A., & Brecht, M. (2014). Cellular, columnar and modular organization of spatial representations in medial entorhinal cortex. *Current Opinion in Neurobiology*, *24C*, 47-54.

Canto, C. B., Wouterlood, F. G., & Witter, M. P. (2008). What does the anatomical organization of the entorhinal cortex tell us? *Neural Plasticity*, 2008, 381243.

Ding, S. L. (2013). Comparative anatomy of the prosubiculum, subiculum, presubiculum, postsubiculum, and parasubiculum in human, monkey, and rodent. *Journal of Comparative Neurology*, *521*(18), 4145-4162.

Insausti, R., Tunon, T., Sobreviela, T., Insausti, A. M., & Gonzalo, L. M. (1995). The human entorhinal cortex: a cytoarchitectonic analysis. *Journal of Comparative Neurology*, *355*(2), 171-198.

Insausti, R., Juottonen, K., Soininen, H., Insausti, A. M., Partanen, K., Vainio, P., ... & Pitkänen, A. (1998). MR volumetric analysis of the human entorhinal, perirhinal, and temporopolar cortices. *American Journal of Neuroradiology*, *19*(4), 659-671.

Insausti R, Amaral DG. (2004). Hippocampal formation. In: Paxinos G, Mai JK (eds) T*he human nervous system*, 2nd edn. Elsevier, Amsterdam, pp 871-914.

Insausti, R., & Amaral, D. G. (2008). Entorhinal cortex of the monkey: IV. Topographical and laminar organization of cortical afferents. *Journal of Comparative Neurology*, *509*(6), 608-41.

Kahn, I., Andrews-Hanna, J. R., Vincent, J. L., Snyder, A. Z., & Buckner, R. L. (2008). Distinct cortical anatomy linked to subregions of the medial temporal lobe revealed by intrinsic functional connectivity. *Journal of Neurophysiology*, *100*(1), 129-39.

Kelly, R. E., Alexopoulos, G. S., Wang, Z., Gunning, F. M., Murphy, C. F., Morimoto, S. S., Hoptman, M. J. (2010). Visual inspection of independent components: defining a procedure for artifact removal from fMRI data. *Journal of Neuroscience Methods*, *189*(2), 233-45.

Kerr, K. M., Agster, K. L., Furtak, S. C., & Burwell, R. D. (2007). Functional neuroanatomy of the parahippocampal region: the lateral and medial entorhinal areas. *Hippocampus*, *17*(9), 697-708.

Krimer, L. S., Hyde, T. M., Herman, M. M., & Saunders, R. C. (1997). The entorhinal cortex: an examination of cyto- and myeloarchitectonic organization in humans. *Cerebral Cortex*, *7*(8), 722-31.

Li, S. J., Li, Z., Wu, G., Zhang, M. J., Franczak, M., & Antuono, P. G. (2002). Alzheimer Disease: Evaluation of a Functional MR Imaging Index as a Marker 1. *Radiology*, *225*(1), 253-259.

Libby, L. A., Ekstrom, A. D., Ragland, J. D., & Ranganath, C. (2012). Differential Connectivity of Perirhinal and Parahippocampal Cortices within Human Hippocampal Subregions Revealed by High-Resolution Functional Imaging. *Journal of Neuroscience*, *32*(19), 6550-6560.

Marques, J. P., Kober, T., Krueger, G., van der Zwaag, W., van de Moortele, P.F., & Gruetter, R. (2010). MP2RAGE, a self bias-field corrected sequence for improved segmentation and T1-mapping at high field. *NeuroImage*, *49*(2), 1271-81.

McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M.B. (2006). Path integration and the neural basis of the "cognitive map". *Nature Reviews Neuroscience*, *7*(8), 663-78.

Morrissey, M. D., Maal-Bared, G., Brady, S., & Takehara-Nishiuchi, K. (2012). Functional dissociation within the entorhinal cortex for memory retrieval of an association between temporally discontiguous stimuli. *The Journal of Neuroscience*: The Official Journal of the Society for Neuroscience, *32*(16), 5356-61.

Muñoz, M., & Insausti, R. (2005). Cortical efferents of the entorhinal cortex and the adjacent parahippocampal region in the monkey (Macaca fascicularis). *The European Journal of Neuroscience*, *22*(6), 1368-88.

Navarro Schroeder, T. (2014). Grid cell representations in humans align to a common reference frame. Poster presentation, Society for Neuroscience Meeting, Washington D.C., 2014.

O'Reilly, K. C., Dahl, A. G., Kruge, I. U., & Witter, M. P. (2013). Subicular-parahippocampal projections revisited. The development of a complex topography in the rat. *Journal of Comparative Neurology*, *521*(18), 4284-4299.

Power, J. D., Barnes, K. A., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *NeuroImage*, *59*(3), 2142-54.

Ranganath, C., & Ritchey, M. (2012). Two cortical systems for memory-guided behaviour. *Nature Reviews Neuroscience*, *13*(10), 713-726.

Schultz, H., Sommer, T., & Peters, J. (2012). Direct Evidence for Domain-Sensitive Functional Subregions in Human Entorhinal Cortex. *Journal of Neuroscience*, *32*(14), 4716-4723.

Van Oort et al., (2014). Instantaneous correlation parcellation (in preparation).

Van Cauter, T., Camon, J., Alvernhe, A., Elduayen, C., Sargolini, F., & Save, E. (2013). Distinct roles of medial and lateral entorhinal cortex in spatial cognition. *Cerebral Cortex*, *23*(2), 451-9.

Witter, M. P., & Moser, E. I. (2006). Spatial representation and the architecture of the entorhinal cortex. *Trends in Neurosciences*, *29*(12), 671-8.

Dynamic Auditory Localisation: Head Tracking of Virtual Moving Sounds

Nynke Niehof[1]
Supervisors: Marc M. van Wanrooij[1], A. John van Opstal[1]

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

Human head tracking of moving sound sources was investigated with the use of dynamic sound stimuli. Sound motion was simulated by adjusting the relative intensities of two speakers, thus creating a dynamic head shadow effect. We observed that sound localisation performance is more variable with these virtual sources than with real sources, but people were indeed able to follow both highly predictable and completely unpredictable periodic sound movements. Analysis reveals that the responses to unpredictable periodic stimuli can be described by a linear transformation. Sounds with a sudden jump in position, followed by a constant displacement in the opposite direction (step-ramp stimuli) suggest that head tracking is largely based on sound position, but that sound velocity may also play a role. Head tracking of sounds can be seen as a series of saccades, interspersed with low-velocity, smooth movements.

Keywords: sound localisation, dynamic sound perception, head movement

Corresponding author: Nynke Niehof; **E-mail:** n.niehof@donders.ru.nl

1. Introduction

Humans can localise sounds from a static source accurately and consistently (Blauert, 1997; Makous & Middlebrooks, 1990; Middlebrooks & Green, 1991), and can make fast head movements towards a sound source (Goossens & Van Opstal, 1997). Real-world sound sources, such as people, animals and vehicles, are often in motion. Can people accurately localise and track moving sounds?

Accurate sound localisation is possible, despite the fact that there is no direct representation of space on the basilar membrane. Instead, spatial information is derived from spectral, temporal and intensity characteristics of the incoming sound signal. A sound, when presented to the side of the listener, travels a slightly longer path to one ear than to the other. This creates an interaural time difference (ITD) in the arrival of the sound. Furthermore, the head casts an acoustic shadow, so that the sound is less intense at the far ear compared to the near ear: the interaural level difference (ILD). Sound localisation in the azimuth (horizontal) plane is largely based on ITDs (for frequencies below about 1.5 kHz) and ILDs (for higher frequencies, mostly above 3 kHz; Blauert, 1997). Sound position in elevation (vertical) is determined from spectral changes caused by reflection and diffraction in the pinnae, and on the body and head. Every sound elevation angle causes a sound wave to reflect in the shape of the pinna in a unique way. A phase delay between a direct sound and a reflected sound can cause interference of the superimposed waves (Moore, 2012). Some frequencies are attenuated while others are magnified, and this pattern depends systematically on the incidence angle of the sound in the pinna. These spectral modifications form a unique profile, from which the elevation angle of the sound can be derived. ITDs, ILDs and spectral modifications give cues about static sound position. Spatial information from sound is important, for example, in traffic. In order to react adequately to alarm signals, such as a car horn or an ambulance siren, one needs to be able to quickly know which direction the sound came from. Dynamic sound localisation can be useful to keep track of moving vehicles, for instance, and to be able to direct the eyes and head towards them.

Visual targets that move through the environment, can be tracked with smooth pursuit eye movements (Rashbass, 1961; Robinson, 1965). This tracking keeps the target's retinal image near the fovea, the retina's area of highest visual acuity. Smooth eye movements to auditory targets seem to be practically non-existent, as people make a series of saccades instead (Berryhill, Chiu, & Hughes, 2006). In contrast, people orient their ears (and thus, their heads) towards sounds, as straight-ahead is the area of highest auditory spatial acuity (Makous & Middlebrooks, 1990). Perhaps people can make smooth pursuit movements with the head to auditory targets. Studies with moving sounds have so far mostly focused on only the psychophysics of dynamic sound perception as an isolated process (Getzmann & Lewald, 2007; Grantham, 1997). In this study, we considered the audiomotor system as a whole by looking at how moving sounds can guide head tracking movements. The goal of this study is to determine how well people can follow moving sounds, and to find a way to model and predict head tracking behaviour.

In order to create the percept of a moving sound, we used simulated ILDs. The ILDs can be manipulated in isolation to evaluate azimuth sound localisation without altering the elevation component (Oldfield & Parker, 1984). We created virtual sound movement by playing the same high-pass noise over two speakers, separated in the horizontal plane, and then systematically varying the relative sound intensities (intensity panning). This procedure can simulate the acoustic head shadow by creating artificial ILDs; leading to the percept of a sound that moves between the two speaker locations. We did not simulate ITDs, because ITDs become ambiguous for higher frequencies (> 1.5 kHz), and we did not want to have extra cues for only the lower frequencies. Intensity panning is used because it can provide a flexible way to create moving sounds without having to physically move the speakers in space (Chowning, 1977). Using only two static speaker locations is less costly and does not create any unwanted noise from the motor, compared to moving the speakers around in the experimental room.

Visual smooth pursuit appears to be based on velocity signals, as opposed to position signals (Rashbass, 1961; Robinson, Gordon, & Gordon, 1986). It is currently unknown whether there exists a similar velocity signal in the auditory system. To characterise head tracking of sounds, and to determine whether this behaviour is likely to be based on a velocity signal, or whether it can just as well be explained based on position, or a combination thereof, we used different types of stimuli. The

sound wave itself was the same for all conditions, namely high-pass noise, but the sound envelope was varied to simulate sound source motion. First of all, sinusoid envelopes of different frequencies were used as a sound movement profile. These stimuli were used to see whether people could perceive and track sound movement from simulated sources at all. From the sinusoids, we also determined the transfer characteristic, namely the gain and phase shift of the response with respect to the stimulus. This was done in a linear time-invariant systems (LTI) framework, which assumes that the frequencies in the response are the same frequencies as in the stimulus, with only changes in amplitude and phase. An LTI system can be entirely characterised by its linear stimulus-response relationship, which stays the same over time. However, the head tracking responses to a sinusoid sound movement may not be linear. Sinusoid movement is highly predictable, so that people may start anticipating it after a few cycles, which violates the time-invariance criterion. This can be checked by seeing whether the response delay for predictable stimuli disappears over time (Michael & Jones, 1966). As a second type of stimulus, we used a sum of sines as a more complex sound envelope. This complex stimulus consisted of a linear superposition of five sine components with randomly selected phase shifts to make the simulated sound movement unpredictable to the listener. The sine components were of the same frequencies as the individual sound envelopes in the single sine stimuli, so that the response power of the components in the head movement can be compared. Sum of sines response spectra were compared to the stimulus spectra, to investigate whether dynamic sound localisation can be described as a linear system. The audiomotor tracking system is only truly linear when the response contains no frequencies that were not present in the stimulus. Finally, a step-ramp stimulus (Rashbass, 1961) was used to determine whether the movement tracking is based on the position or the velocity of the simulated sound source. This stimulus jumps to a new position (step), then moves with a constant velocity in the opposite direction (ramp). We predicted the amplitude of the initial saccade (after the step) from the sound position and sound velocity parameters. Our data show that the head response to a complex, unpredictable stimulus (sum of sines) can be reasonably well described as a linear system. Moreover, the head movements to the step-ramp stimuli were best predicted by a combination of position and velocity information.

2. Methods

2.1 Participants

The experiments were conducted with four participants (one female, three male; ages 23-35). All participants reported normal hearing, normal or corrected to normal vision, and no motor problems. Two of them had participated in sound localisation experiments before, and one was aware of the purpose of the experiments at the time of testing. Since the tasks were performed without any kind of feedback, influence of this prior knowledge on the task performance is not to be expected. Visual localisation behaviour in a task where participants made head movements to static sound sources was normal: stimulus-response relationships, defined as coefficients of determination $R^2 > .9$.

2.2 Experimental setup

The experiments were performed in a completely dark and sound-attenuated room, with the participant sitting on a chair in the middle. The walls, floor and ceiling were covered with black sound-attenuating foam (AX2250, Uxem B.V., Lelystad, The Netherlands), effectively removing echoes for frequencies exceeding 500 Hz. The background noise level in the room was 17 dB SPL (sound pressure level), measured with a Brüel & Kjaer type 2240 sound level meter.

On the wall in front of the participant were 15 green light-emitting diodes (LEDs) positioned on a horizontal line, with radial directions at R = [0, 5, 9, 14, 20, 27, 35, 43] degrees relative to the participant's midline, both to the left and to the right. The LEDs were approximately at eye level for sitting participants. Auditory stimuli were presented on two loudspeakers (SC5.9, Visaton GmbH, Haan, Germany), mounted on the wall at the same height as the LEDs. The speakers were positioned just outside the range of the furthest horizontal LEDs, at ±45 degrees from the participant's midline. Measurements were done with a small coil, placed on the participant's head on a lightweight spectacle frame. The experimental room was equipped with a 3-dimensional oscillating magnetic field system (Remmel Labs, Katy, TX, USA) at different frequencies of [80, 60, 48] kHz, for the horizontal, vertical and front-back axis, respectively. The participants were unrestrained, and head movements towards sounds were registered by

measuring voltage in the head coil as it moved through the magnetic field. In this way, head position as a function of time could be recorded. The head coil voltage signal was amplified (EM7, Remmel Labs, Katy, TX, USA) and low-pass filtered with a custom-built, fourth order Butterworth filter at 150 Hz to increase the signal-to-noise ratio before digitising. The analogue signal was sampled at 1017.25 Hz and digitised (TDT3 RA16PA and RA16, Tucker-Davis Technologies, Alachua, FL, USA), then saved to the hard disk of a PC. On top of the head coil spectacle frame was a red laser pointer. Extending from the frame was a 30 cm long rod with a small plate at the end. The head-mounted laser reflects on this platelet, producing a visible red dot as a reference point for participants, but without illuminating any other objects in the room. The laser dot was used as a head-fixed fixation point, which keeps the eye-in-head position constant, so that participants would be forced to make tracking movements with their heads, instead of just with their eyes.

2.3 Procedure

At the beginning of each trial, an LED lit up on the wall, on which the participant had to fixate. The participant then pressed a button to indicate that the trial could begin, after which the LED was extinguished and a sound was presented. The task for the participants was to keep looking in the direction where they heard the sound, and follow the sound direction throughout the trial. They followed the sound movement by pointing towards it with the head-mounted laser. No performance feedback was provided.

2.4 Stimuli

Auditory stimuli were digitally synthesised Gaussian white noise (Matlab, Mathworks, Natick, MA, USA). The sound levels for the left and right speakers were manipulated to induce a percept of sound source movement, by employing a simulated head shadow effect. First, a sound position vector was defined to specify the azimuth angle of the sound, α (in degrees) over time t.

$$\alpha(t) \leq \alpha_{max} \quad (1)$$

Where α max is the maximum deviation angle from the midline, which was kept at 45° in all dynamic sound experiments. The relative sound intensity I_{rel}, needed to create the percept of a sound at position α, was calculated as follows:

$$I_{rel}(t) = \frac{\alpha(t)}{\alpha_{max}} \quad (2)$$

which varies between [-1, +1]. Sound intensities for the left (I_L) and right (I_R) speakers were then calculated as follows:

$$I_L = 0.5 \cdot (1 - I_{rel}(t)) \quad (3)$$

$$I_R = 0.5 \cdot (1 + I_{rel}(t)) \quad (4)$$

The original white noise sound was high-pass filtered at 3 kHz, then multiplied with I_L and I_R, respectively, to determine the signals for the left and right sound channels. The sounds were ramped on and off with 5 ms squared sine and cosine ramps to avoid clicks at the start and end of playback. Acoustic stimuli of 60 dBA were generated from the saved sound files with a real-time processor (RP2.1 System 3, Tucker-Davis Technologies, Alachua, FL, USA), at 48,828.125 Hz sampling rate.

2.5 Experimental paradigms

Single-speaker sound localisation. All participants completed a standard sound localisation task with brief noise bursts (150 ms; 3-20 kHz), presented on single speakers. The goal of this control experiment was to assess whether the participants were able to localise sounds, and to determine baseline reaction times, correlations (Pearson's r), and localisation gains and offsets as measures of accuracy for each individual. The setup consisted of a vertical hoop (1.2 m radius) that can rotate around the participant, with 29 speakers (SC5.9; Visaton GmbH, Haan, Germany) attached at 5° intervals. Each trial (n = 50) started with a fixation light located straight ahead, and a button press, after which a sound was presented. Participants were instructed to move their heads towards the perceived direction of the sound as quickly and as accurately as possible. Participants P1 and P3 were given stimuli in the azimuth range of ± 60° and elevation ± 20°, at an intensity of 25 dBA. Participant P4 was given stimuli in the range of ± 60° azimuth and ± 60° elevation, at an intensity of 55 dBA. Participant 2 took part in both versions of the experiment.

Two-speaker sound localisation. We used a two-speaker sound localisation paradigm to compare accuracy of the localisation performance to the intensity-panned sounds to the results of the free field single-speaker control experiment. Moreover, we verified whether localisation performance would

be affected by the participant's head position at the time of presentation. After all, the stimuli were intensity-panned to simulate a head shadow effect, but the position of the head in space (during the tracking movement) created an additional head shadow on top of the simulated one.

It is not a priori clear whether these head positions may or may not affect the perceived sound locations. We assessed the effect of head position by using the wall-fixed LEDs as 15 different fixation points. Each fixation LED was combined with 20 different relative intensities for the panned sound stimuli (150 ms; 3-20 kHz; 65 dBA), amounting to a total of 300 trials. All simulated sound positions α were confined to the range of ± 45° azimuth. Task and instructions were the same as for the single-speaker control experiment.

Calibration. For each experimental session, participants completed a calibration experiment. The calibration consisted of 55 LEDs that lit up one by one in the frontal hemifield. Participants were instructed to fixate each LED with the head-mounted laser, and then press the button to record their response. After each button press, the next LED lit up. Just like in the other experimental sessions, the voltage in the coil on the participant's head was measured. This recorded coupling of head coil voltage to fixed locations in space was then used to translate voltage measurements from other experiments into spatial coordinates. All head coordinates were translated to a double polar coordinate system, to make it possible to study azimuth and elevation components in head position independently (Knudsen, 1982; Van Wanrooij, n.d.).

Sines. Sinusoidal sound movements were generated in order to see whether people were perceptually and motorically able to track a predictable, simulated sound source in the horizontal plane, and to measure their adaptation to the predictable movement pattern. We expected that participants would have a response delay at the beginning of a sine movement trial, but that this delay would disappear after a few periods of the 20 s long sounds. The instantaneous sound positions $\alpha_{sine}(t)$ were created as:

$$\alpha_{sine}(t) = A \cdot sin(2\pi \cdot t \cdot f) \tag{5}$$

Amplitude A ([10, 20] degrees) is the maximum azimuth deviation of the sound position from the participant's midline, and f gives the frequency ([0.1, 0.15, 0.35, 0.65, 1.05] Hz) at which the sound moves between those endpoints. The peak velocities for amplitude 10° are [2, 3, 7, 13, 21]π (deg/s), and the peak velocities for amplitude 20° are twice that. See Figure 1A for an example stimulus sound envelope.

Sum of sines. Sum of sines component frequencies F_k were the harmonics [2, 3, 7, 13, 21] of the fundamental frequency of 0.05 Hz. Note that the contributing frequencies are the same as those used for the individual sines, but here all frequencies are summed, then divided by the number of frequency components N_{freqs} to limit the amplitude:

$$\alpha_{SoS}(t) = \frac{\Sigma_k \; A \cdot sin(2\pi \cdot t \cdot F_k + \phi_k)}{N_{freqs}} \tag{6}$$

Amplitude A is 10°, and ϕ_k is a random phase shift between [0, 2π] added for each sinusoid component k, to make the sound movement unpredictable. An example sum of sines stimulus is shown in Figure 1B.

Step-ramp and step. In the step-ramp stimuli, the simulated azimuth sound position $\alpha_{step\text{-}ramp}$ first stays at the midline for a period between 1-2 seconds, by playing noise at equal intensities for both speakers. At the time of the step, the sound $\alpha_{step\text{-}ramp}$ suddenly jumped to another azimuth position. This was immediately followed by the ramp phase, in which $\alpha_{step\text{-}ramp}$ moved with a constant velocity in the azimuth direction opposite to the step direction. The simulated sound position was thus defined as follows:

$$\alpha_{step-ramp}(t) = \begin{cases} 0 & t < \tau \\ A & t = \tau \\ A - V(t-\tau) & t > \tau \end{cases} \tag{7}$$

Amplitude A gives the size of the step ([10, 15, 20, 30] degrees), and the constant velocity ramp depends on the velocity V: [10, 15, 20, 30] degrees per second. Step onset time τ was jittered between one and two seconds after sound onset, to avoid stimulus predictability. The sounds ended with another static phase where $\alpha_{step\text{-}ramp}$ stays at the end position of the ramp. Figure 1C and D show two examples of step-ramp sound envelopes. There were 24 step-ramp trials and 8 catch trials that contained only a step, but no ramp. Both types of stimuli were 5.5 seconds in duration. Step stimuli, illustrated in Figure 1E, were created as follows:

$$\alpha_{step}(t) = \begin{cases} 0 & t < \tau \\ A & t \geq \tau \end{cases} \tag{8}$$

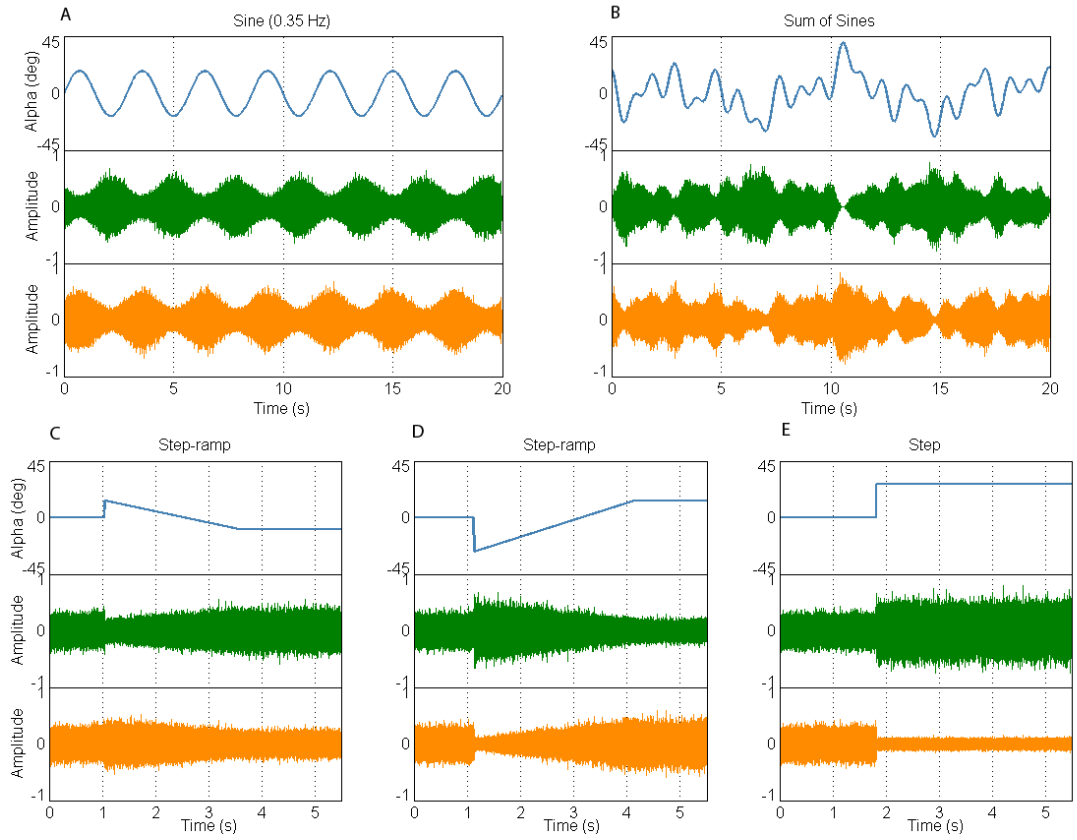

Fig. 1 Examples of sound stimuli used in the dynamic localisation experiments. Top rows: simulated azimuth stimulus position in degrees from the sagittal midline. Middle rows: intensity-panned sound envelopes for the left channel (green). Bottom rows: intensity-panned sound envelopes for the right channel (orange). **A** Sine sound movement, frequency: 0.35 Hz; amplitude: 20 deg. **B** Sum of sines sound movement. **C** Step-ramp sound movement, step size: 15 deg to the left; ramp velocity: 10 deg/s to the right. **D** Step-ramp sound movement, step size: 30 deg to the right; ramp velocity: 15 deg/s to the left. **E** Step sound movement, step size: 30 deg.

3. Results

3.1 Static sound localisation

Individual localisation responses. Sound localisation behaviour was quantified by determining the stimulus-response relationship for each participant. We performed a linear regression analysis on the (simulated) sound location α and the head movement endpoints, as visualised in Figure 2. The regression equation of each participant's response behaviour is listed above the raw data subfigures, giving the response offset and gain. All participants could localise sounds in azimuth well, with stimulus-response coefficients of determination $R^2 > .9$. Participants P1, P2 (in the session with an intensity of 25 dBA) and P3 performed markedly worse on localisation in elevation, with R^2 ranging from .35-.66. However, since auditory localisation in azimuth is independent of elevation localisation (Oldfield & Parker, 1984), and the experiments described here focused on azimuth sound localisation, these participants were included. The bad performance on elevation localisation was possibly because of the low sound intensity, making it hard to pick up spectral cues from the sound. Thus, in subsequent sessions for P2 and P4, the sounds were presented at an intensity of 55 dBA. Participants P1 and P3 were no longer available for the experiment after their first participation, and therefore did not complete a second session with the sounds at 55 dBA. Participant P2 (in the session with 55 dBA intensity) had a cluster of six outliers, all of which were sounds in the bottom left quadrant. The responses of these outliers were flipped to the top left quadrant, possibly because of the sound reflecting on the participant's leg. The second subfigure in the top row of Figure 2 shows the stimulus-response relationship with these outliers removed. Participant P2 showed good elevation sound localisation in the higher intensity session, just as participant P4, with $R^2 > .9$.

The individual localisation responses are more variable in the two-speaker paradigm, when compared with the single-speaker control experiment (Fig. 2, middle row). Upon visual inspection, the stimulus-

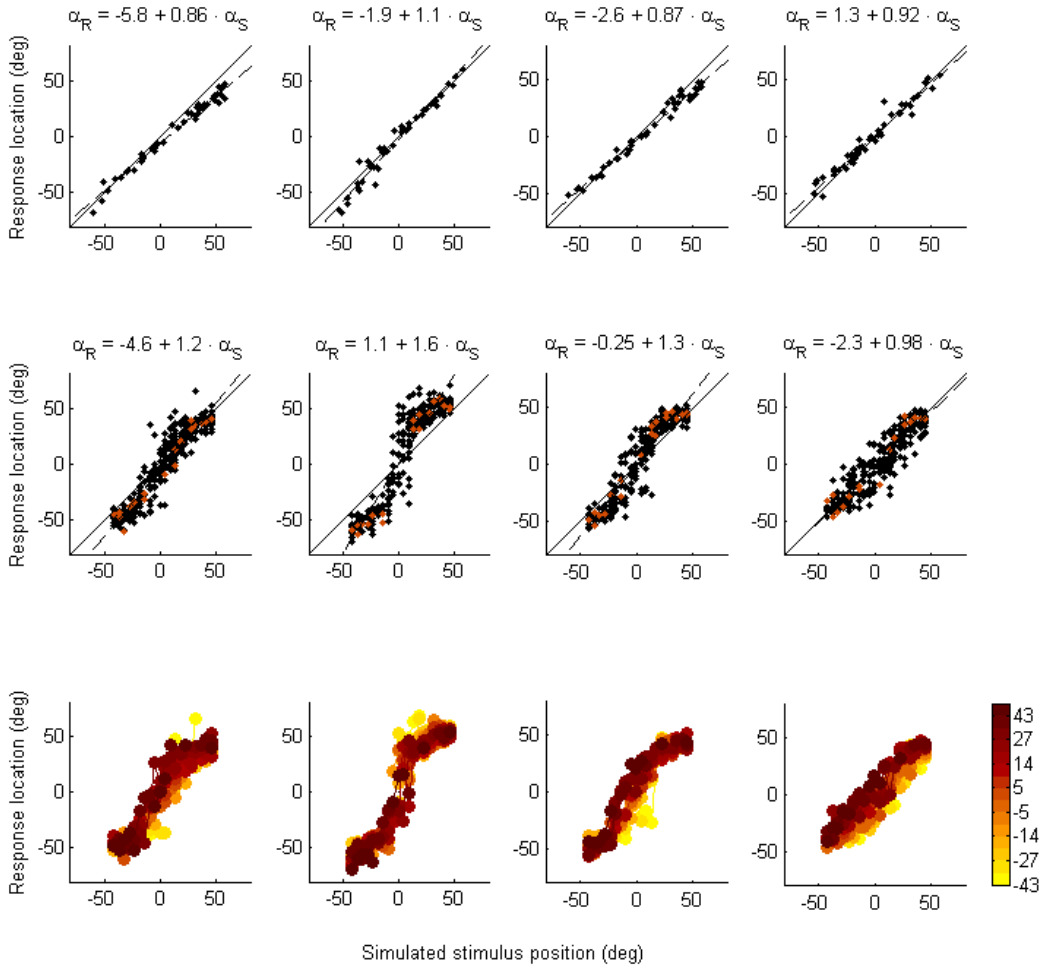

Fig. 2 Auditory localisation responses to static virtual sound positions. Azimuth response locations are plotted against the (simulated) stimulus position in degrees azimuth. Each column shows stimulus-response plots of one participant. Dashed black lines show the regression lines, defined by the equations above the figures, which give the individual response gain and offset. Top row: localisation responses in a single speaker localisation experiment, where stimulus positions on the x-axis are real positions in space, and not simulated. Middle row: azimuth responses with static two-speaker localisation. Orange dots depict measurements with participant head position at the midline (0°). Bottom row: mean localisation responses, with a different colour for each initial head position in a range of ± 43° (head angle-colour mapping in colour bar).

response relations seem more similar to a third order polynomial than to a straight regression line, especially for participants P1, P2 and P3. This seems unlikely to be an effect of head position alone, as the responses with head positions at midline (shown as orange dots) show this same trend. However, model comparison with linearly fitted polynomials (2^{nd} to 5^{th} order) yields negligible improvements in fit, as compared to a first order model. Individual adjusted R^2 values for a first order model were in the range of .84-.88, which increased with at most .05 for third order polynomial fits. Thus, a first order regression line is enough to predict the largest part of the variance in these localisation responses.

Group localisation responses. The pooled data (Fig. 3) show an increased spread on the group level, with higher stimulus-response correlation for single speaker localisation ($r(188) = .98$; $p < .001$) than two-speaker localisation ($r(188) = .92$; $p < .001$). Means and standard deviations are not useful for a comparison between these paradigms, as the (simulated) stimulus position had a different range in these experiments. The larger spread in two-speaker localisation might mean that virtual sound sources created through intensity panning are perceptually less clear-cut than real sound sources. One possible source of variance is small eccentricities of the head, which could distort the intended simulated sound location. We panned the sound intensities to be centred at the midsagittal plane, and an off-center head position during stimulus presentation can distort the simulated head shadow. However, averaged responses per head position (Fig. 3D) showed surprisingly little effect of head eccentricity.

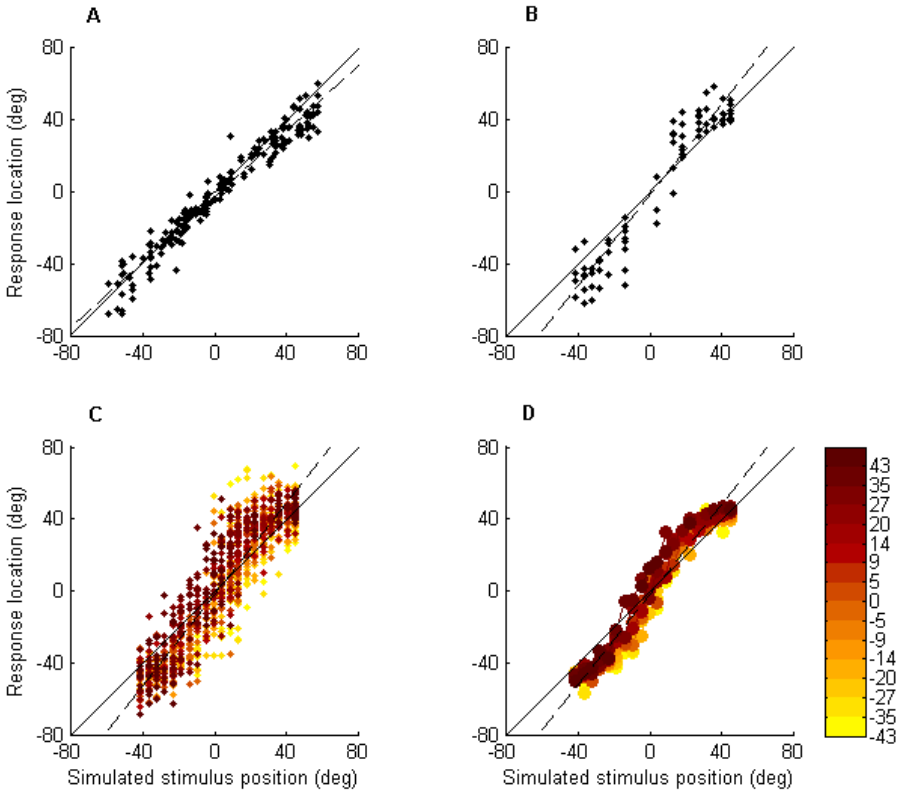

Fig. 3 Stimulus-response relations for static virtual sound positions, pooled over all participants. **A** Single speaker localisation. **B**-**D** Two-speaker localisation. **B** Measurements with initial head position at midline (fixation LED at 0° azimuth). **C** All measurements, with a different colour for each initial head position. **D** Mean response locations for each initial head position. Colour bar: azimuth head fixation positions in degrees. Dashed black lines: regression lines.

Other than a slight bias in the direction of the head angle, the overall distribution of responses is similarly shaped across head positions. Another explanation for the increased variance could lie in expectations of the participants. They were able to see the two speakers on the wall when the lights were on in between experiment blocks. Considering the bias of three participants towards locations around ± 40-50° azimuth (P1, P2, and P3, see Fig. 2, middle and bottom rows), this is a possible explanation for the larger response spread with two speakers. However, linear polynomial fitting shows that on the group level, the responses are also mostly predicted by a first order fit (adjusted R^2 = .84) with an increase in R^2 of only .02 for a third order polynomial fit. Even with a larger spread for two-speaker localisation, the stimulus-response correlations are still high, and show that participants are able to localise the intensity-panned sound sources fairly accurately and consistently.

Head position effect. Fitting a polynomial surface through the responses, with both the relative intensities and the head positions as predictors, showed negligible overall effects of head position

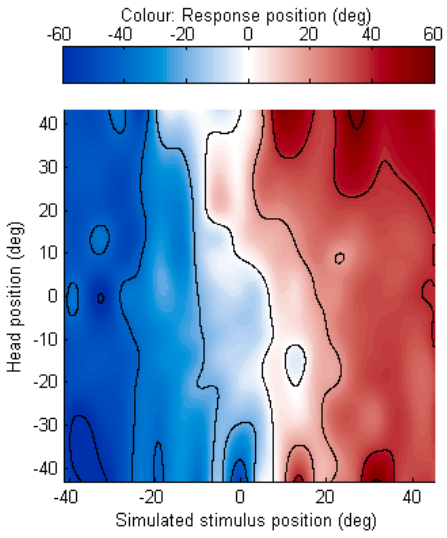

Fig. 4 Example interpolated response locations for each simulated stimulus position, dependent on head position (shown for participant P1). The colours indicate the predicted azimuth response location, with red for positions left of the sagittal midline, and blue for response positions right of the midline. The cusp-like contour line deviations at the top and bottom are interpolation artefacts, and not part of the participant's response pattern.

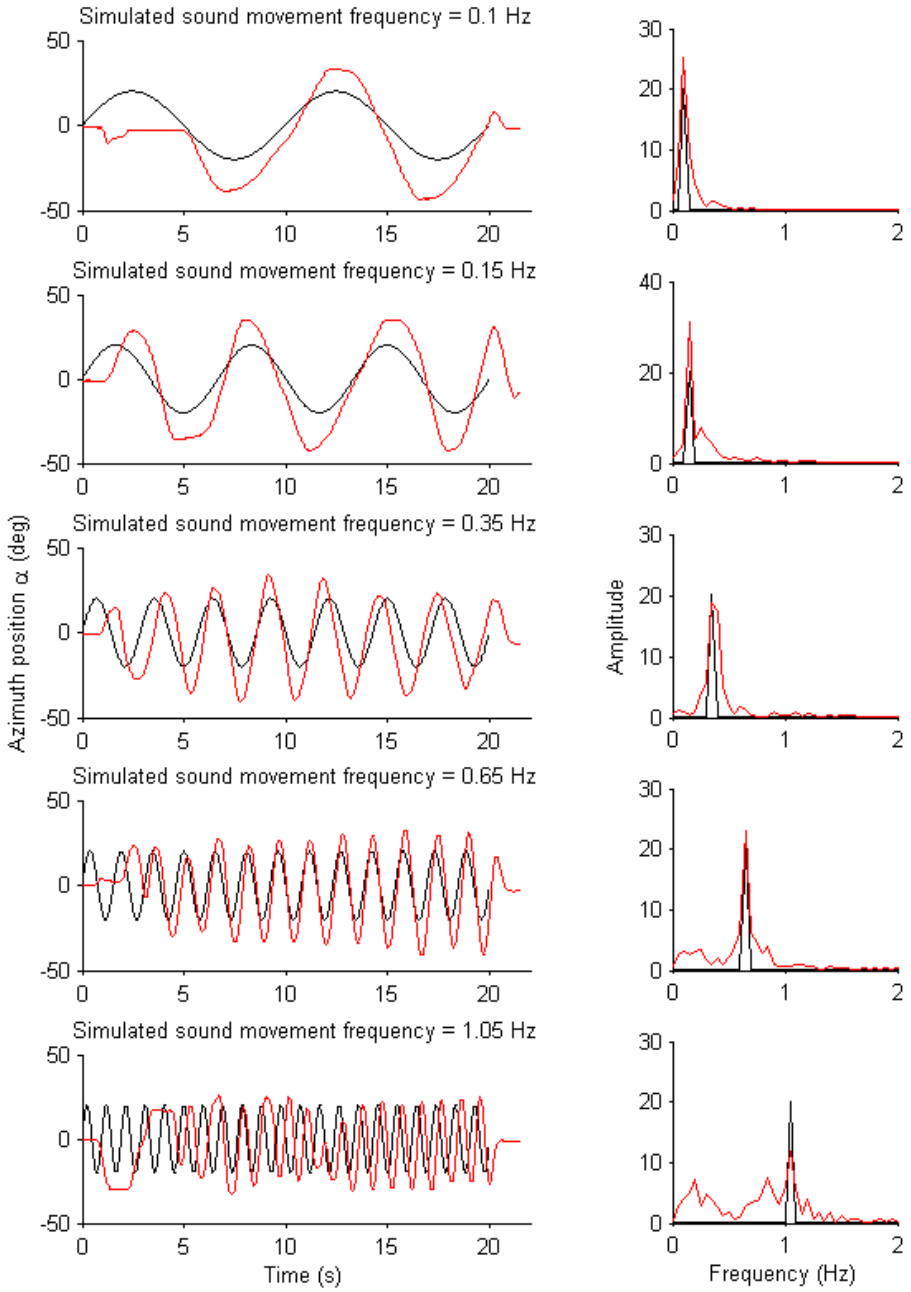

Fig. 5 Responses of participant Pl to sinusoid movements with amplitude 20°. Left column: raw head movement responses in azimuth (red lines) to the simulated sound movement shown in black. Right column: the corresponding frequency spectra of the stimuli (black) and responses (red).

on sound localisation. Changing the head position from a first order to a third order predictor changed the adjusted R^2 with less than .01. We did expect there to be some head position effect, considering that we used a simulated head shadow effect to make moving sound stimuli. To explore this, we determined the effect of head position for each participant separately, from the two-speaker static localisation data. We interpolated a 1000x1000 simulated stimulus position x head position matrix (cubic spline interpolation), with values in the matrix indicating the expected azimuth response position in degrees. An example of this matrix for one participant (P1) is illustrated as a contour line surface in Figure 4. The figure shows only very minor variations in sound localisation over head positions. It should be noted that the cusp-like contour variations at the top and bottom are artefacts from the interpolation, and do not represent actual deviations in localisation behaviour. Outside of the area affected by artefacts, the contour lines are almost vertical, which means no effect of head position. Note that negative values

on the axes are angles to the right, while positive values are angles to the left, so that the *x*-axis is left-right flipped. There does not seem to be a strong, structural localisation bias from head position, judging from the two-speaker response data.

3.2 Dynamic sound localisation

Sines. The frequency spectra of individual head tracking responses to sine movements were taken as a measure of tracking accuracy. The participant was judged as able to track a movement when the frequency spectrum showed a distinctive peak at the frequency of the sinusoid movement. All participants were able to track simulated sound movements at amplitude 20° up to 0.35 Hz with accuracy. One participant was accurate up to 0.65 Hz, and one up to 1.05 Hz. Inaccurate frequency responses still had a clear peak, but it was shifted to a lower frequency than that of the simulated sound movement. Thus, above an idiosyncratic cutoff frequency, participant's responses became nonlinear. For simulated sinusoid movements at amplitude 10°, only one participant was able to accurately follow the sound up to 0.35 Hz. Data of P1 were lost for this condition, due to technical difficulties. Figure 5 shows the time and frequency domain responses of one participant (P1). From the raw traces, it appears that people only lag behind the stimulus movement for the first few periods, then catch up, and even start moving ahead of the stimulus. This is only possible with a predictable stimulus. The predictive responses made the stimulus-response relationship change over time, thus the relationship does not meet the criteria of a time-invariant linear system. Due to buffer limitations of the sound processor, the sound stimuli could not be made longer than 20 seconds. Because of this signal length limitation, it was not possible to accurately determine the phase shift over time for such low frequency movements.

Figure 6 shows the transfer curves (plotted transfer characteristics) of each participant, and the mean transfer curves. Participants were not very well able to follow the simulated movements of amplitude 10° with their heads, which is apparent from the steep drop-off in gain (bottom row). The delay for the amplitude 20° data is remarkably long, but this is mostly due to one participant who fell several cycles behind on the 1.05 Hz stimulus, causing a large phase difference.

Sum of sines. We did a Fourier transform on the sum of sines stimuli and responses, to find the response gain and phase shift for each stimulus component.

Figure 7 shows an example of one participant's head response on a single trial, with the amplitude and phase spectrum, as well as the gain and phase shift. From this, we determined the transfer characteristic for dynamic sound localisation for each participant. The averaged response power

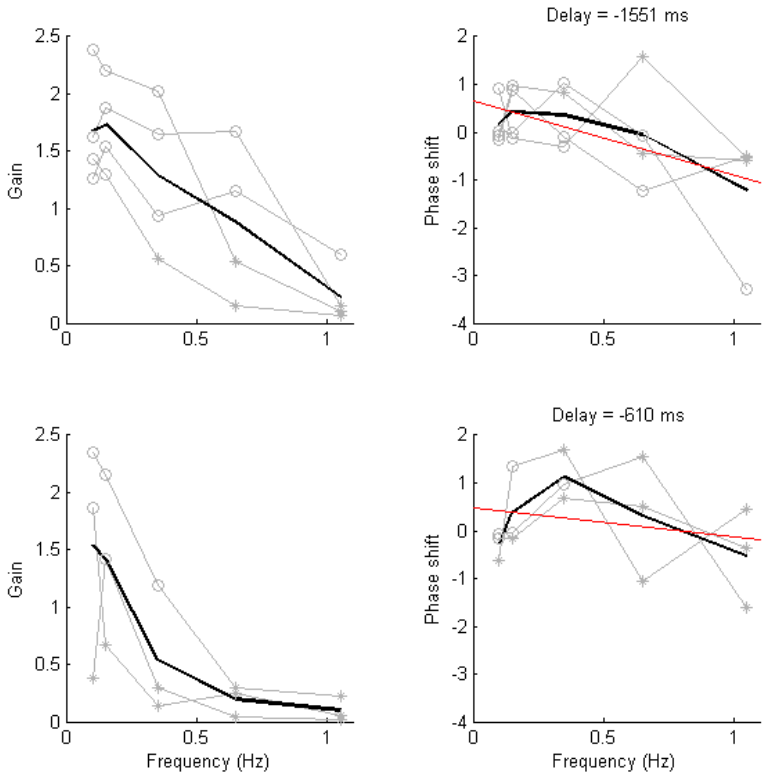

Fig. 6 Transfer curves for sinusoid movements with amplitude 20° (top) and amplitude 10° (bottom). Left column: stimulus-response gain for each participant (grey) and averaged (black). Asterisks denote data points where the peak of the response power spectrum was not at the same frequency as the stimulus power. Right column: phase shift in radians. Red line: regression line over the phase shift, the slope of which gives the delay.

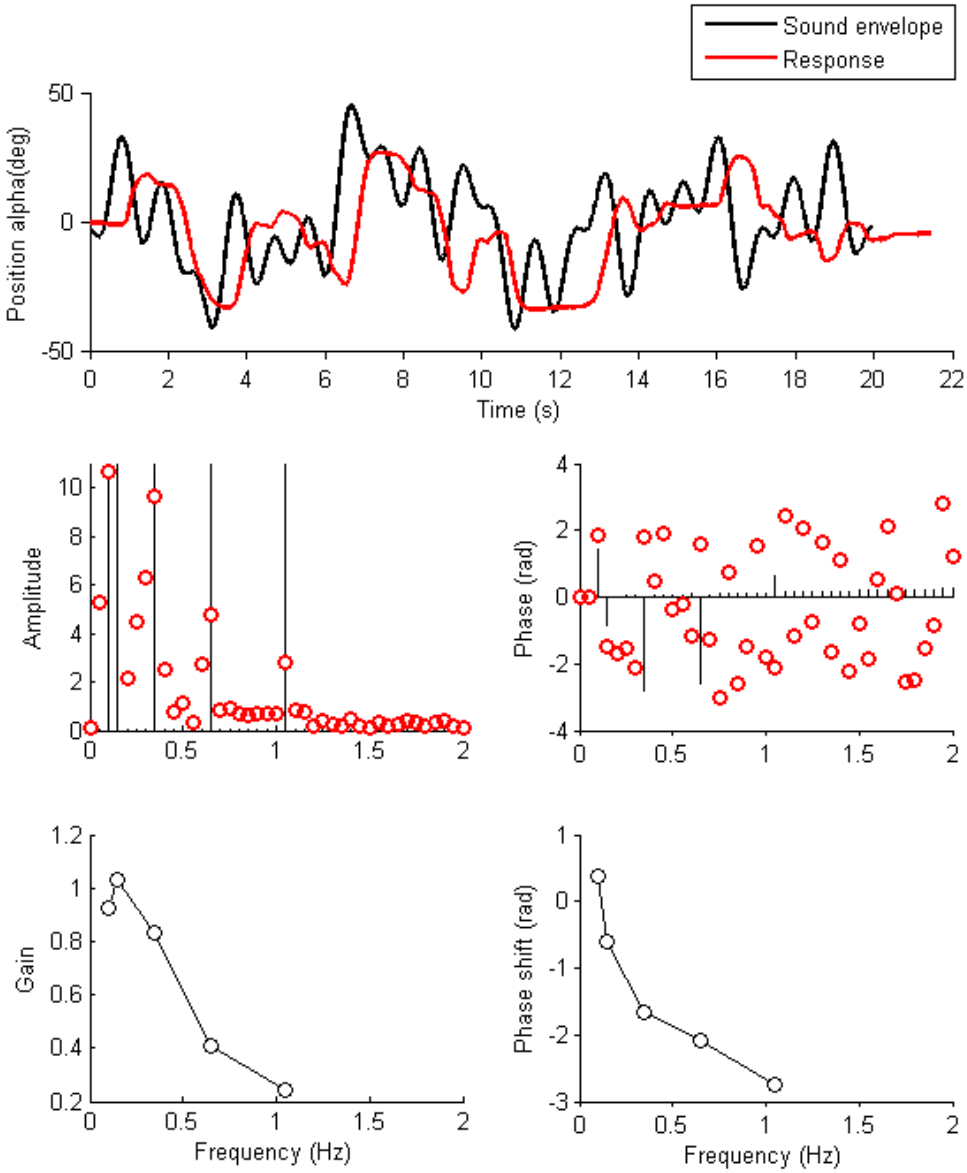

Fig. 7 Example of one sum of sines stimulus (black) and head response (red) from one participant, on the top row. Middle row: power and phase of the stimulus and response. Notably, the largest amplitudes in the response are at the frequency components of the stimulus, indicated by black vertical lines. Other frequencies are present to a lesser extent, indicating nonlinearities in the dynamic sound localisation system. Bottom row: transfer characteristic, given by the gain and phase shift of the response with respect to this particular stimulus.

and transfer characteristic over all participants is shown in Figure 8. The power spectrum shows clear peaks at the five sinusoid frequency components that were incorporated in the sound movement. In a perfectly linear system, the response would only show power at these five frequencies. Part of the power in other frequencies may be explained by the frequency resolution of 0.05. This makes the peaks wider, and in case of the frequencies 0.1 and 0.15, lumps the peaks together into one. Apart from the peak width artefact, there is still some power visible in other frequencies, indicating either nonlinearities in the frequency responses, or noise. The transfer characteristic (Fig. 8, bottom) provides a way to predict the linear components of head responses to any type of simulated sound movement stimulus. The gain and phase shift had the same overall shape for all participants, but the magnitudes varied substantially. The phase difference between stimulus and response becomes larger with stimulus frequency. The phase difference also becomes more variable with frequency, and the gain gets smaller. Accurate head tracking of the sound drops off dramatically above 0.35 Hz, and continuing this line, sound tracking is probably nonexistent above 2 Hz. Only linear components of the frequency responses can be predicted from this transfer characteristic, and nonlinear components would have to be approximated in some other way.

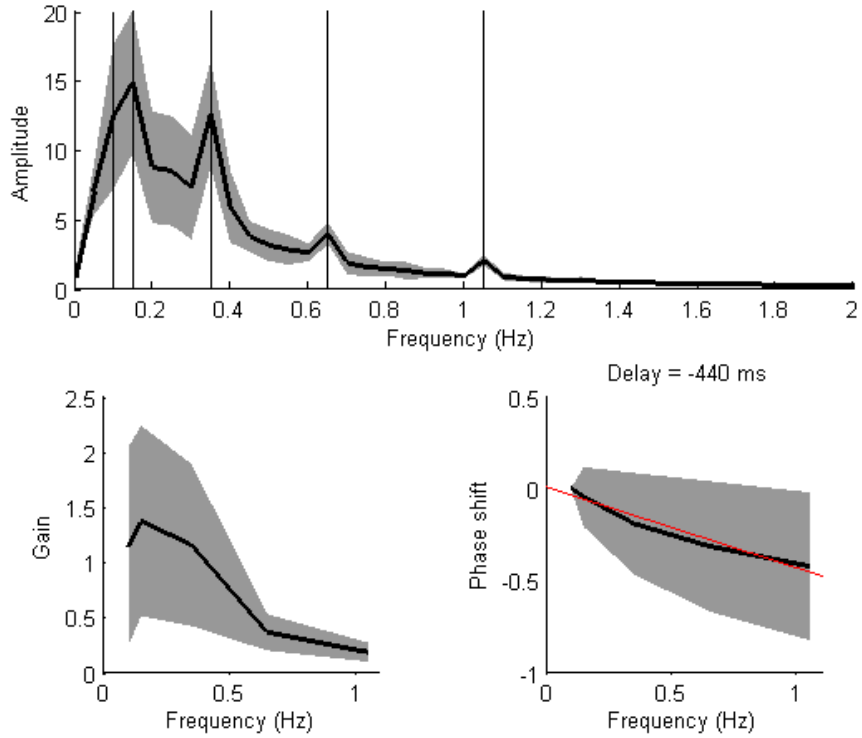

Fig. 8 Mean response characteristics for sum of sines stimuli. Error patches show the standard deviation. Top: response power. Vertical black lines indicate the frequency components of the simulated sound movement. Bottom: transfer curves, with stimulus-response gain on the left, and phase shift between stimulus and response on the right. The slope of the red regression line over phase shifts gives the overall delay.

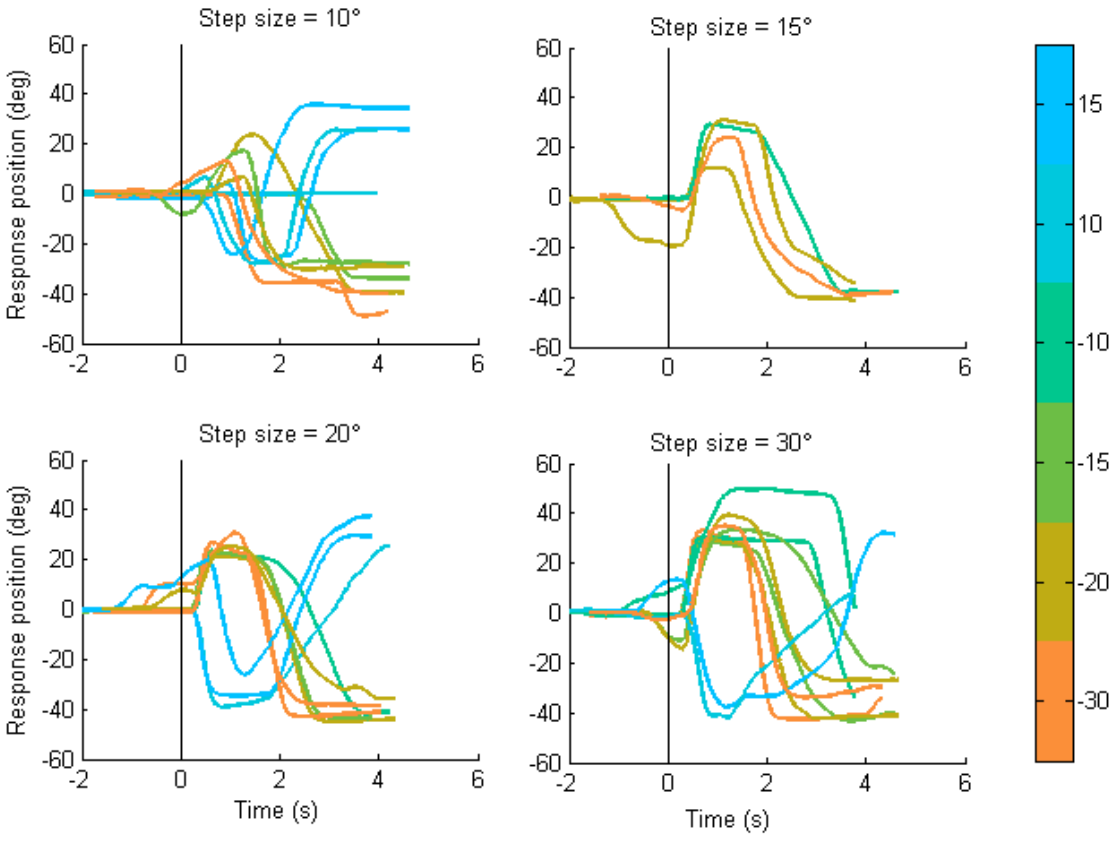

Fig. 9 Step-ramp responses of participants P1 and P2. Each subplot shows azimuth head responses for one step size, listed at the top, and colours signify different ramp velocities (see colour bar).

Step-ramp and step. Figure 9 shows raw responses to different step size x ramp velocity combinations. Data from two participants were excluded for this session, as they were unable to perform the task with these stimuli. They made head movements before the onset of the step in simulated stimulus position, and in many trials they were in the middle of a head movement at the time of the step. For this reason, it was not possible to detect the first saccadic head movement after the step for these participants. This leaves only two participants in this experiment, so results must be interpreted with caution. The initial head saccade to the step in the stimulus was further analysed. We modelled the amplitude of this saccade as a function of the stimulus parameters and the reaction time as follows:

$$A_{saccade} = R_{stepsize} - V_{ramp} \cdot RT \quad (9)$$

Where $R_{stepsize}$ is the simulated angular sound displacement in azimuth, V_{ramp} is the ramp velocity, and reaction time RT is the time between the stimulus step onset and the start of the initial saccade.

This model is graphically depicted in Figure 10, where each line depicts as combination of a step size (intercept) and a ramp velocity (slope).

Five measurements had very long reaction times for the first saccade (RT 600-2400 ms), and were thus unlikely to be a first, fast reaction to the step. For this reason, these responses were excluded. The regression, depicted as the blue line in Figure 11, showed a very good fit (adjusted R^2 = .83) for predicting the saccade amplitude from the reaction time and stimulus parameters.

Figure 12 shows an alternative model, where the initial saccade amplitude is predicted from merely the step size. The step size is also a good model (adjusted R^2 = .80) but this position signal is not as good a predictor as the step (position) and ramp (velocity) together.

4. Discussion

4.1 Summary

We used simulated sound movement through intensity panning to study head tracking of moving sounds. First, we found that people can localise virtual sound sources, both static and moving, but that this static localisation is more variable than for real sound sources. Second, the head tracking behaviour to sounds can be reasonably well modelled by a linear system, but the sound movement should be unpredictable, and there are either some nonlinearities not accounted for, or noise components. Third, sound tracking with the head appears to be based on sound source position signals and perhaps also in part on velocity signals.

4.2 Response variability to panned sounds

All participants displayed a markedly larger spread in responses in the two-speaker panned

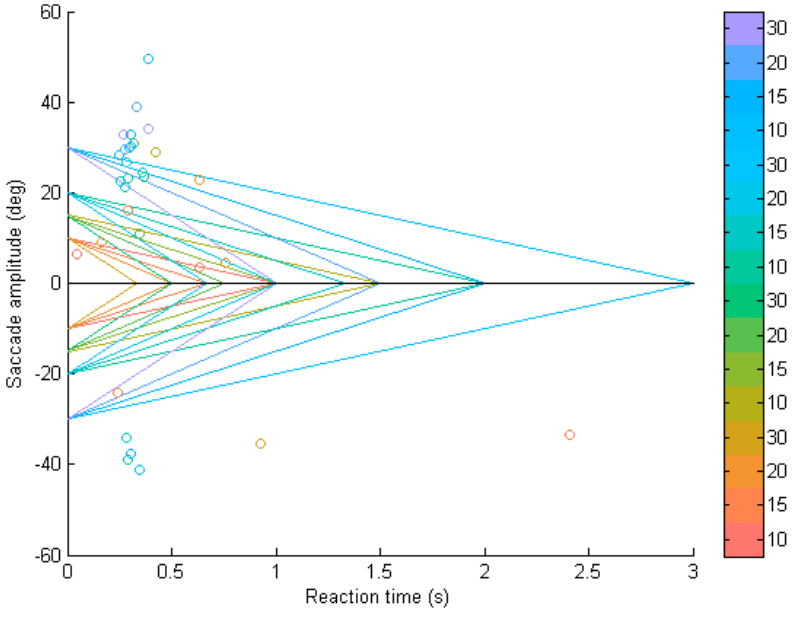

Fig. 10 Modelled initial saccade amplitudes (lines), based on the step size (intercept on y-axis), and ramp velocity (slope) as a function of reaction time. Circles are measured saccade amplitudes of each first saccade after the presentation of the step. Data points have the same colour coding as the model lines. Colour bar shows the ramp velocity in degrees azimuth; colours depict different step size x ramp velocity combinations.

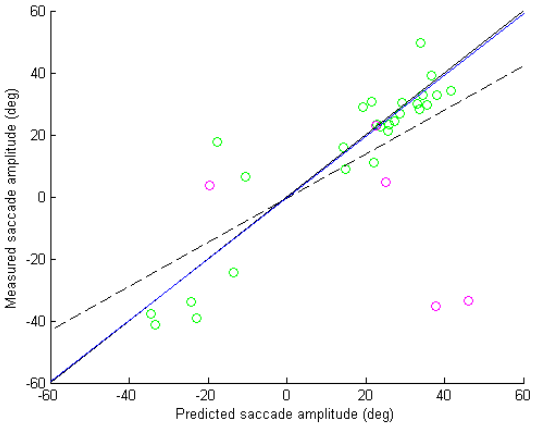

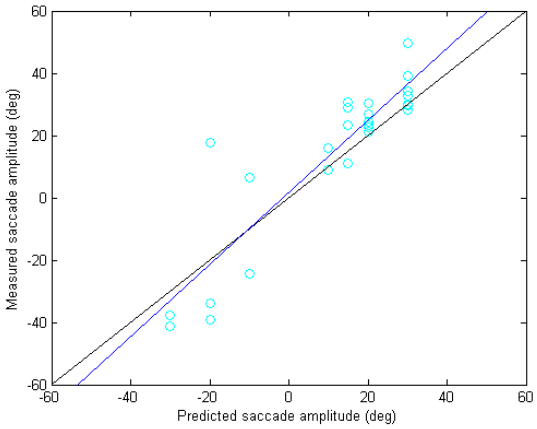

Fig. 11 (left) Model fit for the predicted amplitudes of step-ramp responses in Figure 10 on the measured saccade amplitudes. Circles show the amplitude of the initial saccade as a function of the predicted amplitude. Magenta circles are outliers in reaction time (> 600 ms). All green data points had corresponding reaction times under 500 ms. The dotted line shows the regression with the outliers included, the blue line is the regression without the outliers. R^2 = .83.

Fig. 12 (right) Model fit for the predicted amplitudes of step-ramp responses from the step sizes on the measured saccade amplitudes. Circles show the amplitude of the initial saccade as a function of the predicted amplitude. Only the regression line without the outliers (in blue) is shown. R^2 = .80.

experiment, as compared to the single-speaker control experiment. Even though the correlation is very high for both experiments, intensity panned stimuli clearly yielded more variable responses. The data showed hardly any effect of head position, thus this is probably not the reason for the found spread. Bremen, Van Wanrooij, and Van Opstal (2010) found a bimodal localisation effect in elevation when two speakers were used. When the presented sounds were too far apart, in their case around 45°, response averaging no longer happened, and listeners responded to one of the two sounds (the louder one) instead. Even though sound localisation in elevation is quite different from localisation in azimuth, this seems like the most plausible explanation. Bimodal perception would explain why so many localisation responses are clustered around the real positions of the speakers in space. The participants may not have heard the presented stimulus as a single sound, but as two, and followed the loudest one to its source: the speaker location.

4.3 Transfer characteristic of audiomotor following

People were able to track simulated sound movement reasonably well in the sines and sum of sines experiments. The stimulus movement frequencies were clearly distinguishable in the power spectra of the responses. The responses were generally in the right direction (left or right), albeit with a delay. The delay was fairly stable within participants. People lag behind the sound movement by about 440 ms. Surprisingly, the response amplitudes were often much larger than the stimulus amplitudes, especially for the frequencies under 0.65 Hz. This is apparent from the gains larger than one in Figure 8. However, the variance in response amplitude is also large, showing overshoots as well as undershoots. People were less able to follow sounds with their heads as the movement frequency increased. It is not very likely that this is merely a motor limitation. The mass of the head and the stiffness of the neck muscles impose a limit on the maximum acceleration and jerk of the head. However, people are able to voluntarily shake their heads faster than the maximum frequency of 1.05 Hz used in these experiments. Another possibility is that the limitation lies in internal errors of estimating head position or velocity. We found only very small effects of head position on sound localisation, but we studied this with static head eccentricities. It is possible that head velocity plays a role in sound localisation. Information about sound direction enters the brain in head-centred coordinates; if the internal representation of sound position is kept in head-centred coordinates, then this relative position must be updated across head movements (Goossens & Van Opstal, 1999). On the other hand, a study by Vliegen, Van Grootel, and Van Opstal (2004) showed that eye and head movements are fully accounted for in sound localisation. Even with very fast intervening head movements, localisation of visual and auditory stimuli is still accurate. The

problem is probably not in the estimation of head position, but in the stimuli. Perhaps the percept of a simulated sound source is more easily perturbed by head movements because it is less stable, and can more easily split into a bimodal percept, especially at larger angles between speakers. Perhaps having a larger array of closely spaced speakers would help stabilise the percept of a moving sound.

4.4 Linearity-nonlinearity

From the sum of sines data, it seems that the auditory tracking system can be largely modelled as a linear system. The most prominent peaks in the power spectrum are at the frequencies of the sound envelope components. These peaks overlie a lopsided power distribution with a heavy right tail. The amount of power in neighbouring frequencies because of spectral peak width should be estimated. One way would be to use a setup that can take longer measurements to increase the frequency resolution and get narrower peaks in the power spectrum. The responses to sinusoidal movement stimuli violated linearity in several ways. The relative phase changed during the course of the stimulus presentation: the predictability makes it so that the responses are not time-invariant. Moreover, many of the responses to the sinusoid movements showed frequency components that were not present in the stimulus. Thus, audiomotor following can only be approximated by a linear system when there are no other processes that change the stimulus-response relationship in a nonlinear way.

4.5 Step-ramp responses

The responses in the step-ramp and sum of sines experiments often seemed to jump from point to point, with low-velocity movements in between. The response patterns resembled a sequence of steps, but the slow intervals between saccades were not fixations. The head moved to a point, then showed a period of slow drift, then quickly moved to another point. Interestingly, a model that incorporated both position and velocity information about the sound, fit best to the data. These findings seem to speak against the snapshot theory of sound localisation (Grantham, 1997), which states that sound displacement is derived from a sequence of position signals, rather than from a velocity signal. The alternative theory is that the brain has auditory velocity detectors, specifically for movement (see Middlebrooks & Green, 1991, for a review). However, since these are data from a small data set, some caution is in order. These theories about auditory motion could be further tested with an auditory apparent motion paradigm as in Ingham, Hart, and McAlpine (2001), where the sound jumps from speaker to speaker. If people indeed use a sequence of position signals to estimate sound movement, then the responses should be the same, even though the velocity information is left out. If velocity signals are indeed important for sound tracking, then the responses should look clearly different. The jumps in sound position could be varied in size between conditions, to get an estimate of how frequently people sample the sound position signal.

In conclusion, people are able to follow moving sound sources with their heads, and this is partially a linear transformation. The tracking movement is largely based on sound position, but velocity signals may also play a part. Simulated sound movements through intensity panning are fairly suitable to study this, because the effect is relatively robust across different head positions of the listener. People follow sounds that move through the environment, and do so with a combination of fast saccades and intervals of smooth movement or slow drift.

5. References

Berryhill, M. E., Chiu, T., & Hughes, H. C. (2006). Smooth pursuit of nonvisual motion. *Journal of Neurophysiology, 96*(1), 461–465.

Blauert, J. (1997). Spatial hearing: the psychophysics of human sound localization. *MIT press*.

Bremen, P., Van Wanrooij, M. M., & Van Opstal, A. J. (2010). Pinna cues determine orienting response modes to synchronous sounds in elevation. *Journal of Neuroscience*, *30*(1), 194–204.

Chowning, J. M. (1977). The simulation of moving sound sources. *Computer Music Journal*, *1*(3), 48–52.

Getzmann, S., & Lewald, J. (2007). Localization of moving sound. *Perception & Psychophysics*, *69*(6), 1022-1034.

Goossens, H. H. L. M., & Van Opstal, A. J. (1997). Human eye-head coordination in two dimensions under different sensorimotor conditions. *Experimental Brain Research*, *114*(3), 542–560.

Goossens, H. H. L. M., & Van Opstal, A. J. (1999). Influence of head position on the spatial representation of acoustic targets. *Journal of Neurophysiology*, *81*(6), 2720–2736.

Grantham, D. W. (1997). Auditory motion perception: snapshots revisited. In: R. H. Gilkey & T. R. Anderson (Eds.), Binaural and spatial hearing in real and virtual environments (pp. 295–313). Hillsdale, NJ: Lawrence Erlbaum Assoc.

Ingham, N. J., Hart, H. C., & McAlpine, D. (2001). Spatial receptive fields of inferior colliculus neurons to auditory apparent motion in free field. *Journal of Neurophysiology*, *85*(1), 23–33.

Knudsen, E. I. (1982). Auditory and visual maps of space in the optic tectum of the owl. *Journal of Neuroscience*, *2*(9), 1177–1194.

Makous, J. C., & Middlebrooks, J. C. (1990). Two-dimensional sound localization by human listeners. *Journal of the Acoustical Society of America*, *87*(5), 2188–2200.

Michael, J. A., & Jones, G. M. (1966). Dependence of visual tracking capability upon stimulus predictability. *Vision Research*, *6*(11–12), 707-716.

Middlebrooks, J. C., & Green, D. M. (1991). Sound localization by human listeners. *Annual Review of Psychology*, *42*(1), 135–159.

Moore, B. C. J. (2012). An introduction to the psychology of hearing. 6th ed. Leiden, Netherlands: Brill Academic Publishers.

Oldfield, S. R., & Parker, S. P. (1984). Acuity of sound localisation: a topography of auditory space. II. Pinna cues absent. *Perception*, *13*(5), 601–617.

Rashbass, C. (1961). The relationship between saccadic and smooth tracking eye movements. *Journal of Physiology*, *159*(2), 326.

Robinson, D. (1965). The mechanics of human smooth pursuit eye movement. *Journal of Physiology*, *180*(3), 569.

Robinson, D., Gordon, J., & Gordon, S. (1986). A model of the smooth pursuit eye movement system. *Biological Cybernetics*, *55*(1), 43–57.

Van Wanrooij, M. M. (n.d.). Double polar coordinates. Retrieved 2014-08-15, from http://www.neural-code.com/index.php/tutorials/stimulus/sound/69-double-polar-coordinates

Vliegen, J., Van Grootel, T. J., & Van Opstal, A. J. (2004). Dynamic sound localization during rapid eye-head gaze shifts. *Journal of Neuroscience*, *24*(42), 9291–9302.

Striatal GABAergic Control of Human Reward Anticipation

Annelies J. M. van Nuland[1]
Supervisors: David G. Norris[1], Roshan Cools[1], Hanneke E. M. den Ouden[1]

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

The dopaminergic midbrain (substantia nigra and ventral tegmental area (SN/VTA)) is associated with computation of reward prediction error signals. This reward prediction error is essential for learning. Gamma-aminobutyric acid (GABA) also plays a crucial role in this process, GABAergic neurons in the VTA appear to encode the reward prediction independent of outcome. It is however still unclear what drives this GABA response, i.e. where information encoding this reward prediction originates. Potential candidates are the orbital frontal cortex and striatum. We test for a potential role for the striatum as input to form this reward prediction. We hypothesized that the influence of the striatum occurs indirectly through activation of the GABAergic interneurons in the SN/VTA. We measured GABA concentration bilaterally in the striatum and occipital cortex using magnetic resonance spectroscopy (MRS). Reward anticipation was induced using a monetary incentive delay task. We found significant differential activity in the VTA upon presentation of (potentially) rewarding and neutral cues, which thus reflected a reward anticipation response. We report a trend level correlation between VTA reward response and striatal GABA levels, but nothing from the occipital cortex. This correlation was specific for blood oxygenation level dependent (BOLD) activity during cue outcome presentation. We suggest that this timeframe corresponds with VTA GABAergic signalling of reward anticipation.

Keywords: ventral tegmental area, substantia nigra, striatum, reward, GABA, dopamine, MR-Spectroscopy, reward prediction error, MID

Corresponding author: Annelies van Nuland; **E-mail:** a.vannuland@fcdonders.ru.nl

1. Introduction

To maximize chances of survival, an organism needs to be able to select behavior that leads to the best possible outcome. Prediction of reward along with the ability to flexibly update this prediction is essential for optimizing behavior. Dopamine (DA) is a key neuromodulator involved in the process of reward prediction, notably the computation of reward prediction errors. According to the reward-prediction error hypothesis, phasic activity of dopaminergic neurons in the midbrain signal a discrepancy between the predicted and currently experienced reward of a particular event (Colombo, 2014). The ventral tegmental area (VTA) is one of the main dopaminergic nuclei. Cells in the VTA were found to fire in response to rewarding sensory input and expectation of reward (Cohen, Haesler, Vong, Lowell, & Uchida, 2012; Schultz, 1998). The origin of the reward prediction itself is still unknown, but is thought to be modulated by inputs from either striatum or orbital frontal cortex (Russo & Nestler, 2013; Watabe-Uchida, Zhu, Ogawa, Vamanrao, & Uchida, 2012). In this study we use magnetic resonance (MR) gamma-aminobutyric acid (GABA) spectroscopy in combination with functional magnetic resonance imaging (fMRI) in humans to further investigate the influence of striatal GABA on VTA mediated reward anticipation response.

Based on firing rates from DA neurons in the macaque VTA, Schultz (1998) observed that midbrain dopaminergic firing rates matched exactly what a reward prediction error signal would look like as predicted by computational models of learning (Schultz, 1998). These neurons increase firing upon delivery of an unexpected reward, and decrease their firing when an expected reward is omitted. When an expected reward is presented, their firing rate does not change. In other words, they appear to encode the discrepancy between anticipated and actual reward.

Prediction errors are essential to induce conditioned learning, as shown by the phenomenon of blocking (Holland, 1984; Steinberg et al., 2013). Blocking describes the phenomenon in which the association between a cue and a reward is prevented or blocked, when another cue in the same environment at the same time already fully predicts reward delivery. The new cue does not provide any new information; hence no prediction error is generated. Dopaminergic prediction errors are not only necessary but also sufficient for learning to occur; Eshel, Tian, and Uchida (2013) stimulated DA neurons in vivo using optogenetics in a blocking paradigm. Where formation of a cue reward association was previously blocked, stimulation of DA neurons induced learning of the new cue-reward association, mimicking a prediction error (Eshel et al., 2013). This proved a causal role between dopamine neuron signalling in cue-reward learning, presumably through the signalling of prediction errors.

The same lab showed that there is a significant role for GABAergic neurons in the VTA in calculation of a reward prediction (Cohen et al., 2012). Genetically tagged optical stimulation along with electrophysiological recording allowed specific differentiation of GABAergic and dopaminergic cells. Where dopaminergic neurons fired upon presentation of reward-predicting cues and unexpected rewards, GABAergic neurons fired during the interval from cue onset to reward presentation. The firing rate would increase up to the point of expected reward delivery while the firing rate was proportional to the size of the expected reward. This suggests that while DA is involved in firing in response to reward prediction errors, GABAergic neurons in the same region encode the reward prediction per se. GABAergic firing could serve as an inhibitory drive that counteracts the excitatory drive generated by the primary reward when this reward is presented. The net difference between excitatory drive (reward) and inhibitory drive (expectation) results in a reward prediction error (Cohen et al., 2012). It is however unlikely that these local interneurons are fully responsible for calculating the reward prediction, and input from higher brain regions is suggested to drive these predictions. Relevant input might be provided by the striatum (Stocco, Lebiere, & Anderson, 2010) and the prefrontal cortex or orbital frontal cortex (Takahashi et al., 2011), since all of these regions have been shown to have projections to the VTA and are important in value expectation processing (Russo & Nestler, 2013; Watabe-Uchida et al., 2012).

The ventral striatum is closely associated with reward (Lenz & Lobo, 2013); it receives input from the orbitofrontal cortex, insula, cingulate cortex and amygdala (Haber, 2011). This makes it an ideal site for integrating cognitive and affective information from different areas (Redgrave, Vautrelle, & Reynolds, 2011; Stocco et al., 2010). In a recent study in mice a significant part of striatal GABAergic medial spinal neurons were found to project to the substantia nigra (SN); they were shown to specifically connect to SN GABAergic neurons, whereas they showed no connections to SN DA neurons. This suggests that the influence of the striatum on reward prediction is mediated through modulation of GABAergic inter-neurons in this region (Chuhma, Tanaka, Hen, &

Rayport, 2011). The same population of GABAergic medial spiny neurons (MSNs) that was shown to project to the SN is also innervated by the VTA and SN (Houk, Adams, & Barto, 1995). The GABAergic MSNs account for most of the population of neurons in the striatum (Russo & Nestler, 2013). Any influence of the striatum to the SN and VTA is therefore mediated by activity of GABAergic neurons. We therefore suggest that the influence of the striatum on the SN/VTA is specific for GABAergic activity and only indirectly affected by other cell types such as the cholinergic interneurons.

Evidence of similar activity related to a dopaminergic midbrain and striatum mediated reward prediction is widely reported in humans (Ballard et al., 2011; D'Ardenne, McClure, Nystrom, & Cohen, 2008; Düzel et al., 2009; Iglesias et al., 2013; Klein-Flügge, Hunt, Bach, Dolan, & Behrens, 2011; Niv, Edlund, Dayan, & O'Doherty, 2012; Sescousse, Caldú, Segura, & Dreher, 2013). Blood Oxygenation Level Dependent signal (BOLD) response in the dopaminergic midbrain was both related to the size of the reward (Klein-Flügge et al., 2011) as to whether a reward was predicted or unpredicted (D'Ardenne et al., 2008). Similarly, BOLD responses in the striatum closely corresponded to a reward prediction error (Niv et al., 2012) and were suggested to encode timing information to the dopaminergic midbrain (Klein-Flügge et al., 2011). One of the most reliable tasks to induce a reward anticipation/prediction response in human dopaminergic midbrain has been the monetary incentive delay task (Adcock, Thangavel, Whitfield-Gabrieli, Knutson, & Gabrieli, 2006; Knutson, Westdorp, Kaiser, & Hommer, 2000). In this task participants are presented with a potentially rewarding or neutral cue, and have to respond to a target quickly to be able to win. By using a task with a strong reward anticipation/prediction response we hope to increase sensitivity when measuring in a small region such as the dopaminergic midbrain. Most human reward studies were limited to the VTA and did not include the SN, however we suggest including the SN as region of interest in human studies. In rodents we see a clear distinction in SN and VTA function; VTA activation is strongly associated with reward and the SN with gating of motor activity. However, in humans and primates we find that the SN is anatomically and functionally more continuous with the VTA and there is only evidence for gradual differences between these regions of the dopaminergic midbrain (Düzel et al., 2009).

Both rodent and human studies suggest an influence of striatal activity on VTA reward anticipation computation. Striatal GABAergic cells project on the SN-VTA GABAergic interneurons, not on the DA neurons itself (Chuhma et al., 2011). Since GABA-neurons do not directly drive blood oxygenation level dependent (BOLD) response (Northoff et al., 2007; Sotero & Trujillo-Barreto, 2007), we are most likely to measure an indirect effect on striatal BOLD response. Thus we aimed to test the hypothesis that individual difference in striatal GABA predicts the size of the reward expectation signal in the VTA. To this end we measured basal concentration of GABA in the striatum, and found a hint of evidence that reward prediction related BOLD in the dopaminergic midbrain relates to the levels of GABA in the striatum.

2. Materials and Methods

2.1 Participants

Forty-eight young adults were scanned, gave informed consent prior to participation and were paid €26,- to participate. Consent and procedures were approved by the local institutional review board (CMO2001/095, version 6.2) in accordance with the declaration of Helsinki. Exclusion criteria were left-handedness, psychiatric disorder, pregnancy, history of closed-head injury and presence of metal objects in the body. Two participants were excluded due to technical problems, and in two only partial data could be acquired, which resulted in 44 full datasets. In this article based on limited time-resources we report a subset using 22 participants (11 females), and a subset of the test battery (details below).

2.2 Experimental procedure

Before entering the scanner, participants were fitted with a breathing belt, and a heart rate monitoring clip was adjusted on their index finger or ring finger. They further received a button box, panic button and mirror to be able to see the screen. The scanning procedure started with an anatomical T1 scan followed by 45 minutes of GABA magnetic resonance spectroscopy (MRS). After a short break participants would proceed with the instructions for the first task, while a susceptibility weighted scan was acquired. This was followed by 26 minutes of a probabilistic reward learning task (adapted from (Klein-Flügge et al., 2011; Niv et al., 2012), data will be reported in a different paper). Next, a FieldMap was acquired to map inhomogeneities of the magnetic field, followed by 12-13 minutes

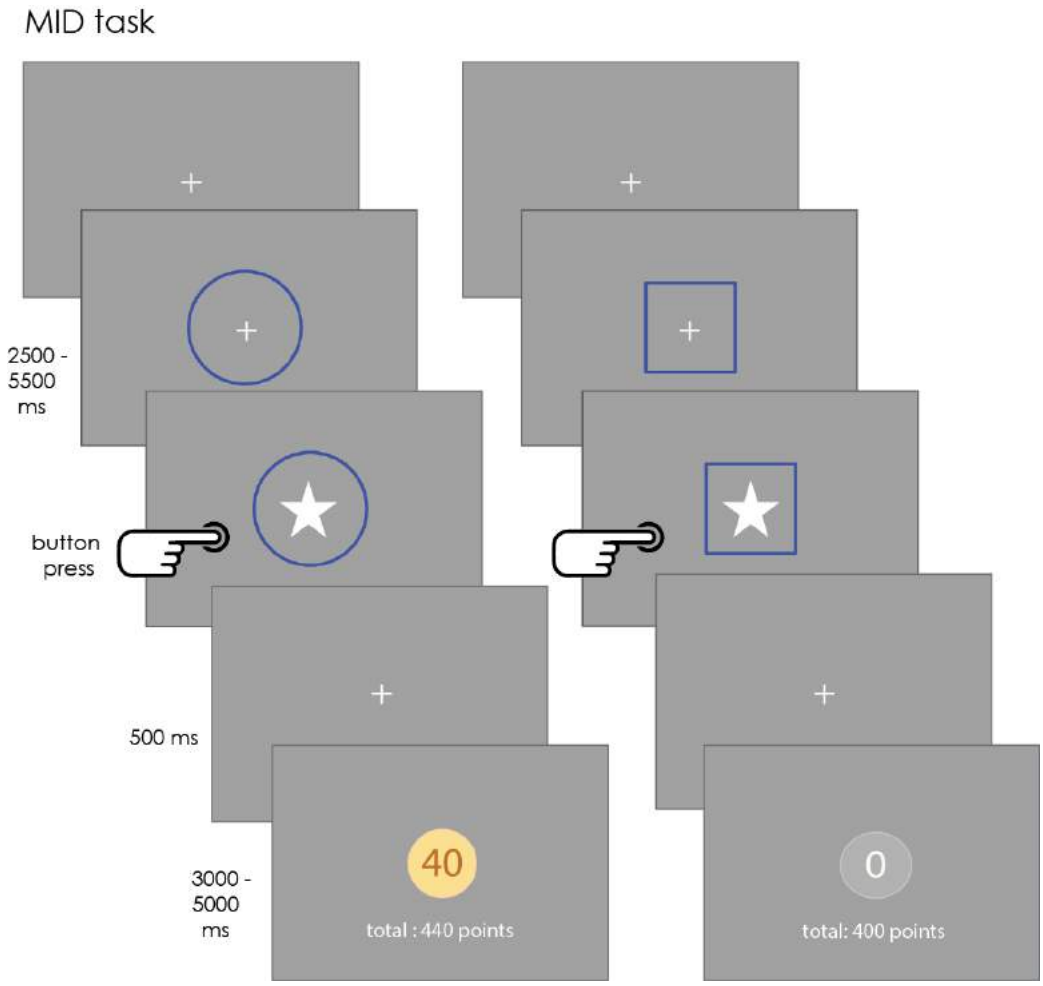

Fig. 1 MID (monetary incentive delay) task; in each trial, participants viewed a cue that indicated whether they had the opportunity to earn a reward or not. After a variable interval, a target appeared to which subjects had to respond before it disappeared. Target presentation time was adjusted so that participants were on time in 66% of the trials. 0.5 seconds after target offset feedback was presented for 1 second, consisting of the current and total number of points earned. After a variable inter-trial interval (3-5 s), the next cue would appear. Participants were instructed that they could NOT press before the target appeared.

of a monetary incentive delay task (depending on reaction time) (Knutson et al., 2000), the results of which are reported in this article. After scanning, participants completed two behavioral tasks; a go/nogo task and a probabilistic reversal learning task (data will be reported in a different paper).

2.2.1 Experimental paradigm

To obtain a neural measure of reward anticipation we used a modified version of the monetary incentive delay task (Knutson et al., 2000), see Fig.1. Every trial began with cue presentation, in which the cue signalled the potential outcomes. A blue circle signalled reward was available, and a blue square signalled that no reward was available. After a variable delay (3-5.8 s, Poisson distributed, average 4 s) a target in the shape of a white star appeared in the middle of the cue. Participants were required to respond to the target by pressing a button with the index finger of the right hand, before the target disappeared. On reward trials, a sufficiently fast reaction would result in 40 points, on neutral trials this would result in 0 points. A miss would always result in 0 points. Feedback was displayed after a 500 ms delay and included both the number of points earned during that round and the total number of points earned so far. The presentation time of the target was determined by an adaptive algorithm that assured that the subject would be successful on approximately 66% of trials. Inter-trial time ranged between 3 and 5.8 s, was Poisson distributed, average 4 s. Trial time varied as a function of reaction time but was on average 9.81 s (standard error of the mean (s.e.m.) = 0.81 s). Each condition (reward and neutral) was presented 40 times. Before starting the task, participants completed a short practice run inside the scanner, in which a separate cue was used.

2.2.2 MR data acquisition

MR data were acquired at 3T on a Siemens Trio using a 32-channel head coil. First a whole-brain high resolution T1-weighted anatomical scan was acquired using an MPRAGE sequence with voxels of 1.0 mm isotropic (TE = 3.03 ms, TR = 2300 ms, 192 slices). This was followed by approximately 35 minutes of GABA MRS (details below). For SN-VTA parcellation, a susceptibility weighted imaging (SWI) scan was acquired with separate images for phase and magnitude information (1 mm isotropic, multi-echo (10.04 ms, 20.27 ms, 30.5 ms, 40.73 ms), TR = 48 ms, flip angle = 15 deg., GRAPPA acceleration factor = 3). To allow for high resolution scanning within reasonable timeframe, a partial brain volume was acquired. The slab positioning was centered on the midbrain, while avoiding the eye region to prevent eye specific artifacts. FOV = 256x192x72 mm.

During the two tasks, 3D echo planar imaging (EPI) scans were acquired (voxel size 1.5 mm isotropic, TR = 57 ms, TE = 26 ms, effective TR = 1.710, FoV = 210 x 210 x 84 mm). The slab positioning and rotation optimally covered both prefrontal, deep brain regions and striatum, while avoiding the eye region to prevent artifacts (average angle of 18 degrees to AC axis). Between the two tasks a standard SIEMENS field map (C Hutton, Deichmann, Turner, & Andersson, 2004; Chloe Hutton et al., 2002) was acquired to allow for off-line correction of EPI data for field inhomogeneities (TR = 1020 ms, TE1 = 10.00 ms, TE2 = 12.46 ms). Breathing and pulse data were recorded on a BrainAmp (ExG amplifier, Brain products GmbH). Total scan duration was about 1.5 hrs.

2.2.3 GABA MRS

Voxel placement was based on the T1 anatomical scan. An in-house automatic aligning procedure provided a rough placement of the voxel, which was then adjusted for each participant individually following the anatomical guidelines described below (see also Fig. 2). We acquired a left and right striatal voxel and a single right lateralized occipital cortex control voxel. Voxel dimensions (1.8x2.5x1.8 cm) were optimized for the slightly elongated shape of the striatum. This included sections of the nucleus accumbens, globus pallidum and caudate nucleus. A small section of the thalamus is also included.

MRS provides an in vivo non-invasive tool to study biochemical properties of brain tissue. MRS detects radiofrequency signals from hydrogen spins from different tissue metabolites, which, caused by their chemical environment, have a specific frequency or chemical shift. By measuring the relative amplitude at this specific frequency one can determine the relative concentration of target brain metabolites. For a detailed explanation see (Puts & Edden, 2012). MRS data were acquired using a modified version of MEGA-PRESS with a TR of 1600 ms; MEGA-PRESS was first developed by Mescher (Mescher, Merkle, Kirsch, Garwood, & Gruetter, 1998; Mescher, Tannus, Johnson, & Garwood, 1996) and has been extensively reviewed by Mullins et al. (2014). It uses known coupling of hydrogen molecules within the GABA molecule to specifically eliminate overlapping signals, which forms one of the main problems in measuring GABA. For the striatal voxels, 128 averages were acquired. Due to the relatively higher signal-to-noise ratio, only 80 averages were acquired for the occipital voxel. Each voxel acquisition was preceded by a Siemens 'Fastestmap' shimming procedure (Xiaodong Zhong, 2013). Shimming reduces the magnetic variability within a single voxel by adjusting the gradients until an optimal configuration is reached. Variability in the magnetic field causes a shift in overall resonance frequency which is strongly detrimental to the final signal strength when averaging over spectra, and results in a broader line-width within the spectra. Since shimming procedures can be sensitive to local optima, shimming was performed over a number of iterations (at least 3-5 is advised, depending on the region) until a spectral line width of maximum 14 Hz was reached. If this line width could not be obtained within 6 iterations, the voxel was moved a few millimeters further from nearby cerebrospinal fluid (CSF). CSF is avoided since inclusion of CSF strongly increases the danger of magnetic inhomogeneity, therefore having a negative influence on line width. Furthermore, it includes a scale of molecules not naturally present in striatal tissue.

The following anatomical guidelines were used for voxel placement: Striatum - adjunct to the rostral/anterior edge of the striatum, medially bordering the lateral ventricle (to avoid inclusion of CSF). Occipital cortex - along the medial edge of the right hemisphere, with the inferior/anterior edge bordering the cerebellum, and the posterior/anterior border just within the cortical tissue. All voxels were aligned with the magnetic field in x and y position to effect optimal shimming, and were transversally orientated (longest along anterior-posterior line).

2.3 Data analysis

The main aim of this study was to test whether striatal GABA levels affect reward anticipation related BOLD responses in the dopaminergic midbrain (SN-VTA). Specifically, we hypothesized that individual differences in striatal GABA levels will predict the cue-induced reward expectation signal in the VTA. In the section below, we describe how GABA MRS data and fMRI data were separately analyzed and then combined to test this hypothesis.

2.3.1 Behavioral data analysis

We analyzed both reaction times and hit rates to test for performance differences in reaction to reward/neutral cues potentially reflecting a difference in motivation. To determine if reaction times between conditions were significantly different we performed a paired t-test. For the hit rates we performed another paired t-test.

2.3.2 Solving potential difficulties with imaging the dopaminergic midbrain

The VTA is a relatively small region consisting of only a few voxels when using regular 3 mm^3 sequencing (D'Ardenne et al., 2008). This prompted us to use a sequencing design with a higher 1.5 mm^3 resolution. Moreover, normalization could misalign the region within participants; it is not visible on normal anatomical scans and its small size makes misalignment more likely. Therefore we used a subject specific design in which we anatomically determined the location of the VTA and SN using susceptibility weighted imaging. Although to our knowledge not used previously for localizing the SN, we show that SWI can function as a reliable localizer of the SN along with the VTA, see below. An anatomically specified region of interest (ROI) analysis for each individual is in this case preferred since it allows for sensitivity for activity that would normally not show up at cluster level because of the small size of this region. Along with its small size, the VTA is positioned in an area strongly influenced by physiological noise from nearby arteries. To overcome this problem we extracted heartbeat and breathing related parameters using physiological measurements of heartbeat and breathing cycle.

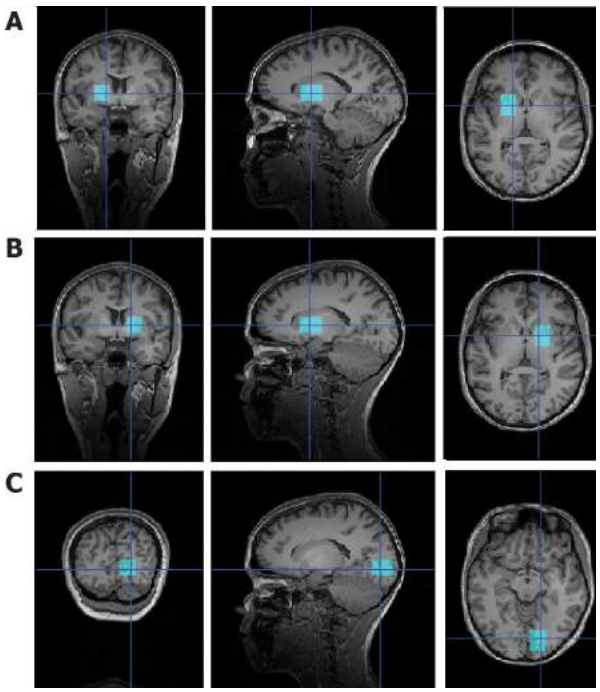

Fig. 2 Example of positioning of MRS voxels for one subject. **A.** Left striatum. **B.** Right striatum. **C.** Right occipital cortex.

2.3.3 ROI definition of the dopaminergic midbrain (SN/VTA)

The dopaminergic midbrain was defined using an optimized SWI scan for each individual using the following procedure. First, the high-resolution T1 scan was co-registered with the SWI to determine the approximate location of the SN based on anatomical landmarks (upper pons). After this rough anatomical localization, SWI images were used for more precise anatomical localization. The SN are visible on the magnitude images of the SWI as an elongated dark structure, positioned ventral to the red nucleus, which is visible as a dark round structure (Fig. 3). The echo time (TE) and the number of echoes were chosen using piloting studies to result in an optimal contrast between SN and non-SN tissue. The VTA was identified as the region between the SN and the red nucleus on the medial side of the pons (Damier, Hirsch, Agid, & Graybiel, 1999). The SN/VTA was then hand drawn for each individual using MRIcron and a BAMBOO graphics tablet. The segmentation of the SN was performed using the fourth echo magnitude image (TE = 22 ms) for each participant individually, based on pilots showing the largest contrast difference at the location of the SN. These individual SN-VTA ROIs were then used to analyze in the functional data.

2.3.4 fMRI data processing and analysis

fMRI data was analyzed using SPM 8 (www.fil.ion.ucl.ac.uk/spm; Wellcome Department of Imaging Neuroscience, London, UK). The images

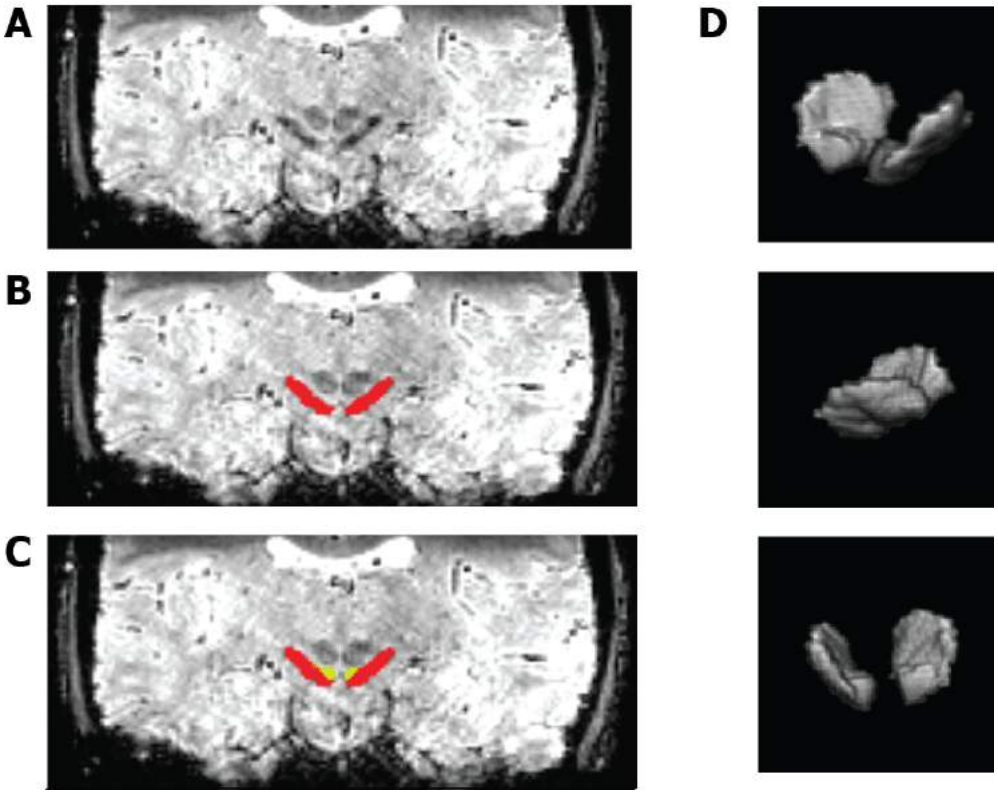

Fig. 3 Parcelation of SN and VTA using SWI. **A.** Original SWI image, red nucleus and SN are visible as dark spots in the midbrain. **B, C.** Examples of parcelation steps; **B.** The Substantia Nigra, which is most apparently visible as the dark area beneath the dot-shaped red nucleus is marked in red. **C.** The VTA is determined as the most medial region between the red nucleus, the SN and 4th ventricle and is added to the ROI area in yellow. **D.** 3D visualization of combined VTA/SN ROI using microgl from different angles.

from each subject were realigned to correct for head movement, after which the magnetic FieldMap was used to create a mean undistorted image; images from each subject were corrected for geometric distortions caused by susceptibility-induced field inhomogeneities. A combined approach was used that corrects for both static distortions and changes in these distortions attributable to head motion (Andersson, Hutton, Ashburner, Turner, & Friston, 2001; Hutton et al., 2002). The static distortions were calculated for each subject by acquiring a B0 field map and processing it using the FieldMap toolbox implemented in SPM8 (Hutton et al., 2004). We analyzed the data twice, once in subject-space, to extract data from the individual, anatomically defined SN-VTA ROIs, and once for a whole-brain analysis in MNI space. For the SN-VTA ROI analysis, scans were co-registered and re-sliced to the SWI image after realignment, with no further preprocessing steps. For the whole-brain analysis, scans were co-registered to the structural image. The structural image was processed using a unified segmentation procedure combining segmentation, bias correction, and spatial normalization (Ashburner & Friston, 2005). The same normalization parameters were then used to normalize the echo planar images. Finally, the functional images were spatially smoothed with a 3-dimensional Gaussian kernel of 5 mm full width half maximum.

Event-related statistical analysis was performed according to the general linear model as implemented in SPM8. The same design matrix was used for the ROI and the whole-brain analyses, and comprised the following regressors: Cue presentation of the reward or no-reward cue (2 regressors; duration = duration of the cue presentation), target presentation (1 regressor, duration = 0 s), feedback (4 regressors: reward/hit, neutral/hit, reward/miss and neutral/miss; duration = 0 s). We also included 3 sets of nuisance regressors: i) 24 movement regressors, including the 6 basic movement parameters (translation and rotation), their first derivative, and the spin history effects and their first derivative, see also (Friston, Williams, Howard, Frackowiak, & Turner, 1996), ii) the average signal of out of brain and CSF voxels, based on the segmented anatomical scan (2 regressors), iii) cardiac and respiratory based regressors using a RETROICOR approach (Glover, Li, & Ress, 2000). RETROICOR uses 10 cardiac phase regressors (5th order Fourier set),

10 respiratory phase regressors (5th order Fourier set), and 6 other nuisance regressors (heart rate frequency, heart rate variability, raw respiration data, respiratory amplitude, respiratory frequency, and frequency times amplitude of respiration).

2.3.5 ROI analysis

The SN/VTA ROI was selected manually for each individual based on SWI image. ROI analysis of the striatum was based on anatomical mask using SPM8 WFU-PickAtlas tool. This mask included the nucleus accumbens, pallidum, caudate nucleus and putamen and was chosen to roughly coincide with the area of our MRS voxel. We extracted contrast images of each ROI based on a first level analysis using SPM8.

Beta extraction of SN/VTA ROI was conducted in native space with the extension of SPM MarsBaR. A paired t-test was used to test whether there was a significant difference between beta estimates for the reward trials versus neutral trials. Beta extraction of striatal voxels was done using second-level analyses of reward-neutral contrast in SPM8.

In addition to a beta analysis we extracted BOLD time series from both ROI for each subject. Time series extraction was performed as an extension to beta correlation analysis to account for time specific effects of the striatal influence on reward anticipation. This GABA-encoded reward anticipation was found to peak during expected outcome delivery. Effects of GABA on dopaminergic midbrain BOLD signal might therefore be similarly specific in time. The nuisance-corrected (mean, retroicor, CSF, OOB and motor) time courses were upsampled to 100 ms resolution using spline interpolation and averaged for each subject within each condition.

2.3.6 Whole-brain analysis

A second level group analysis with the normalized functional images was used to assess effect in the other brain regions. We looked at reward-neutral cue presentation. We report significant activity at voxel-level ($p = .05$ family wise error rate corrected (FWE)) and cluster level (0.05 FWE) threshold.

2.3.7 GABA MRS analysis

MRS analysis was performed using the jMRUI software package version 5 (http://www.mrui.uab.es/mrui). Preprocessing included a zero filling procedure which added 512 data points to the signal. The spectra were smoothed using a 5 Hz Lorentzian curve. The reference frequency of the N-acetyl aspartate (NAA) peak was set to 2.03 ppm, and any residual water signal at 4.7 ppm was filtered out using Hankel Lanczos Singular Value Decomposition (HLSVD) (Pijnappel, Van den Boogaart, De Beer, & Van Ormondt, 1992). Quantification of the spectra was done using Advanced Method for Accurate, Robust, and Efficient Spectral fitting (AMARES) (Vanhamme, van den Boogaart, & Van Huffel, 1997). The peaks were modelled using several Lorentzian curves which were locked to specific ppm values, using prior knowledge to optimally model the underlying spectra. Specific prior knowledge settings were locally developed and encompass expected resonance frequency of GABA, NAA peaks, Glx, Glutamate and Macromolecues. GABA was modelled using two Lorentzian curves of which the amplitudes were constraint to have the same value, and the line width was constraint to that of the NAA resonance peak.

Our assumption was that the effect would be striatum specific, therefore not correlated with GABA in the occipital cortex control voxel. Earlier research shows that, although there is some measure of bilateral asymmetry within GABA levels found in some brain regions (one of which is the nucleus accumbens), a much stronger difference is seen between different brain regions (Starr & Kilpatrick, 1981). A stronger correlation between left and right striatum was therefore expected. A strong correlation between the occipital cortex and the basal ganglia might point to a whole-brain effect, and would suggest that differences in striatal GABA between subjects was driven by an overall GABA increase and effects would therefore not be specific to the striatum.

Each value determined by AMARES was corrected for the level of grey matter, white matter and CSF in this voxel, using a detailed segmentation of the T1 image. GABA signal is assumed to be limited to grey matter only, since conversion of GABA out of its precursor occurs locally in the pre-synapse (Kanaani, Kolibachuk, Martinez, & Baekkeskov, 2010; Rowley, Madsen, Schousboe, & White, 2012; Stagg, Bachtiar, & Johansen-Berg, 2011). White matter corrected data is calculated as the portion of grey matter times the signal intensity.

2.3.8 Correlation analysis of GABA with fMRI BOLD

Basic correlation between GABA and fMRI BOLD signals was done based on beta extraction of ROIs (see above). For SN/VTA the separate reward and neutral contrasts were subtracted to produce

a reward-neutral contrast. Striatal pre-whitened betas were extracted based on second level reward-neutral contrast. GABA concentration measures were averaged over both voxels and resulted in a mean striatal GABA concentration value. A measure of correlation between striatal BOLD and GABA concentrations is added as a control, since we expect correlation of GABA to be with the target region of the striatum, not with the striatum itself. Correlation analysis between striatal GABA and ROI specific beta values was done using MATLAB 2013a robustfit. For illustration purposes we split Midbrain BOLD response based on whether the participant showed high (top half results) or low (lower half results) GABA concentrations.

We next wanted to assess whether any effects of GABA were specific during a particular epoch of the cue presentation. We expected there to be a peak at time of outcome presentation. For this within-trial temporal analysis, BOLD response was averaged over all reward and neutral trials. We used a time window of 2 seconds after cue presentation up till 15 seconds after cue presentation, 15 seconds coincides with the average start of a new cue presentation and end of the previous trial plus 5 seconds to account for slowness of the BOLD response. We chose this time window as it encompasses the relevant task related BOLD response. We correlated mean striatal GABA with ROI BOLD response for each time bin (100 ms resolution) using a Pearson correlation analysis in Matlab. We corrected for multiple comparisons using a permutation based Monte Carlo cluster approach (Maris & Oostenveld, 2007), implemented using the Fieldtrip toolbox (Oostenveld, Fries, Maris, & Schoffelen, 2010). This approach was originally developed for electroencephalography (EEG) & magnetoencephalography (MEG) data, however the nature of the extracted time series makes this analysis comparable to a single source analysis EEG with subject specific set variable (here represented by the GABA measures). This means that GABA measures were randomized during each permutation and correlated over the time series. Permutation distribution was based on the added Monte Carlo based cluster score for each permutation, significance value is now determined by comparing the rareness of our original Monte Carlo based cluster score with this distribution.

3. Results

3.1 Behavioral measures

Participants were faster in responding to the target when they could win points compared to when they could not ($p < .001$, Reward RT = 0.29 s, Neutral RT = 0.32 s standard error of deviation (s.e.d.): 0.02), and were therefore also more likely to respond on time ($p < .001$; Hit rate: Rew = 0.74 Neu = 0.57, s.e.d. : 0.11).

3.2 Validation of striatal GABA measurements

We measured GABA in three different brain regions; left and right striatum and a single right lateralized voxel of the occipital cortex. There

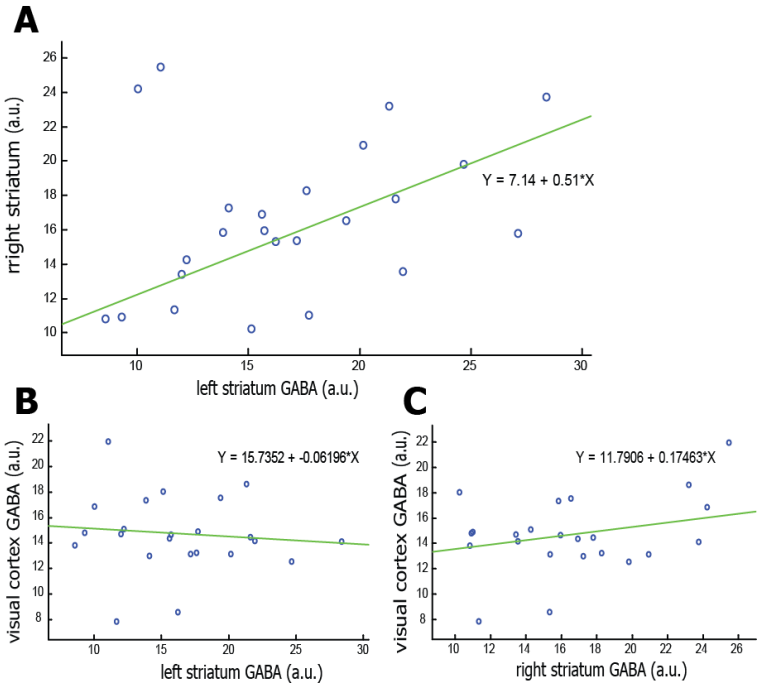

Fig. 4 Correlation distribution between GABA concentration between 3 regions. **A.** A significant correlation was found between the left and right striatum ($r = 0.51$, $p = .0008$). **B.** No significant correlation found between the left striatum and visual cortex ($r = -0.06$, $p = .51$). **C.** No significant correlation between the right striatum and the occipital/visual cortex ($r = 0.17$, $p = .21$).

was a significant correlation between the left and right striatal voxels (r = .51, p < .001, Fig. 4A) but not between the striatum and occipital cortex (left striatum: r = -.06, p = .51, right striatum: r = .17, p = .21, Fig. 4B-C). The presence of the striatal correlation confirmed that the striatal GABA measures were reliable within subjects, as anticipated based on rodent work (Starr & Kilpatrick, 1981). In addition, the absence of a relation between the striatum and occipital cortex confirms that the GABA measure gives a consistent isolated effect and is not driven by a whole-brain / global signal.

3.3 MID fMRI results

3.3.1 ROI analysis

The ROI analysis of the SN-VTA unveiled a significantly larger response to reward compared to neutral trials (p = .004) (see Fig. 5A). No significant difference was found between reward and neutral condition in our striatal region (Fig. 5C). The shape and location of the SN/VTA ROI is depicted in Figure 5E.

In addition to standard BOLD comparisons in both conditions, BOLD signal of ROI analysis were plotted over time and compared during cue presentation of both the reward and neutral condition (Fig. 5B, Fig. 5D). We found increased activity for reward cues compared with neutral cues both in the SN-VTA and in the striatum. Both cue and outcome onsets are marked with a small black line.

3.3.2 Whole brain

We found significantly increased activity in the right dorsal striatum or thalamus and significantly increased activity at the anterior cingulated cortex. There was a wide spread significant response in the striatal system to reward over neutral cues, in a large cluster encompassing bilateral caudate, putamen, ventral striatum and the main striatal output structure, the thalamus. All cluster level significant foci of activity are reported in Table 1 (see also Fig. 5F).

3.4 Relationship of striatal GABA to reward-anticipation induced BOLD

A simple correlation between reward-neutral BOLD in the SN-VTA and striatal GABA resulted in a negative correlation r = -.38 at trend level (p = .08) (see Fig. 6A). There was no significant correlation between striatal GABA concentration and striatal BOLD reward anticipation response (r = -.18, p = .41), (see Fig. 6B). Since this concerns the difference between reward and neutral cues this means that in the dopaminergic midbrain (SN/VTA) there is an increased difference between reward and neutral BOLD response in participants with low striatal GABA (Fig. 6E) when compared with participants with high striatal GABA (Fig. 6G).

For the SN/VTA only one positive cluster survives the initial uncorrected threshold (p > .05), and is shown as a grey bar in Figure 6C. Correction for multiple comparisons gives a corrected p-value of p = .08, and is therefore at trend level. Onset of this cluster is around 12 seconds. When correlating the occipital GABA concentrations; no clusters survive the initial threshold (p > .05). Similarly, neither striatal nor occipital GABA concentration explained any differential response in the striatal differential reward-neutral cue response (p > .05) (Fig. 6D).

Table 1. MNI coordinates and z-values for regions of significant activation.

Foci of activation	MNI coordinates			Z value	# voxels
	x	**y**	**z**		
Reward – neutral cue					
Anterior cingulate cortex	**4**	**5**	**50**	**5.97**	**3481**
Thalamus	**12**	**-18**	**6**	**5.79**	**1495**
Precentral gyrus	**48**	**-6**	**50**	**4.76**	**547**
Temporal lobe	**58**	**10**	**2**	**4.72**	**1548**
Insula	**-34**	**12**	**13**	**4.64**	**1226**
Precentral gyrus	**-40**	**-8**	**55**	**4.47**	**751**
Caudate nucleus	**8**	**6**	**10**	**4.27**	**723**
Globus pallidum	**-15**	**2**	**0**	**4.19**	**522**

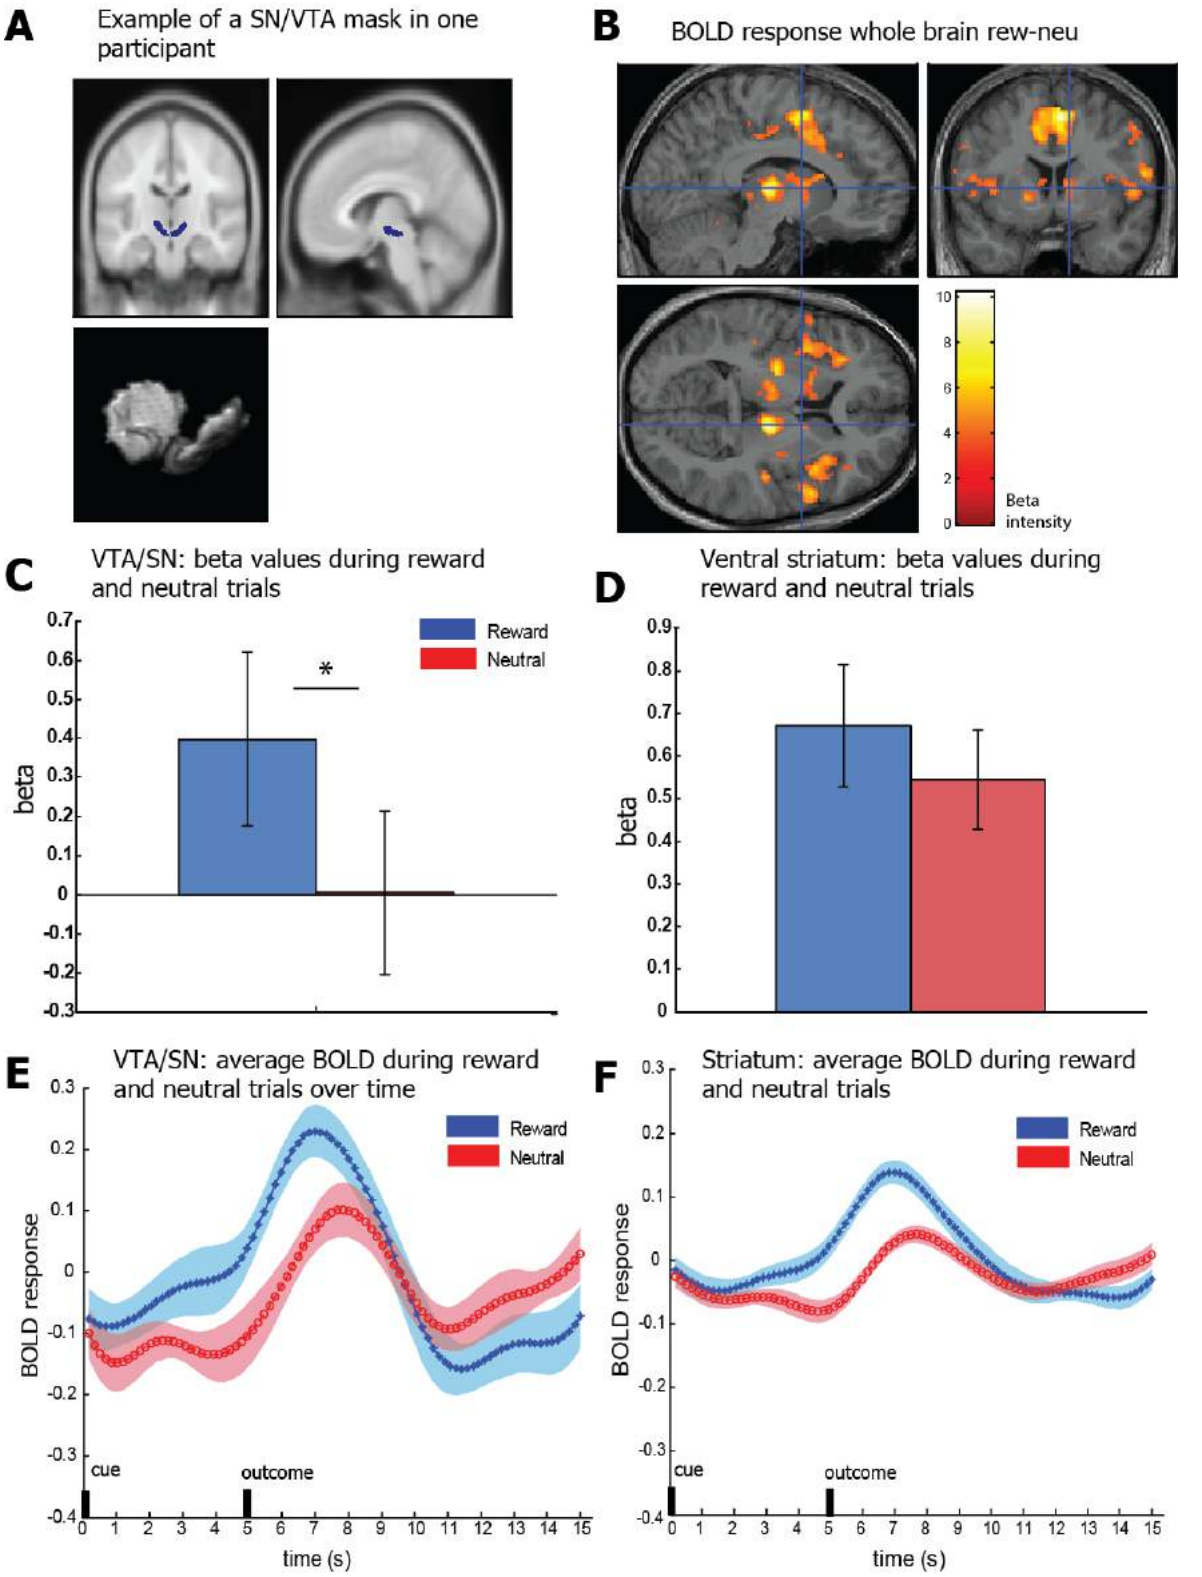

Fig. 5 BOLD responses of the ROI analysis in the SN/VTA and striatum using the prewhitened BOLD response over time for the specified ROI. A, C and E concern the SN/VTA, while figure B, D and F concern the dorsal striatum. **A.** Location and size of SN/VTA cluster in the brain. **B.** BOLD response in reward-neutral contrast at p = .001 uncorrected on a whole brain level analysis. **C, D.** Beta values, a significant difference was found between reward and neutral trials ($p < .05$) for the SN/VTA ROI but not for the striatal ROI. Error bars represent the standard error **E, F.** BOLD timeseries extraction for Reward and Neutral cues for **E.** VTA/SN and **F.** Striatum. Cue and outcome presentation onset are marked by black lines. BOLD response is expected to peak around 7 seconds for cue and around 12 seconds for outcome presentation. Bounded lines represent the standard error. Cue presentation accounts for a much stronger reaction in BOLD than outcome presentation, although this may be partly caused by the jitter between cue and outcome presentation that was present in our MID design.

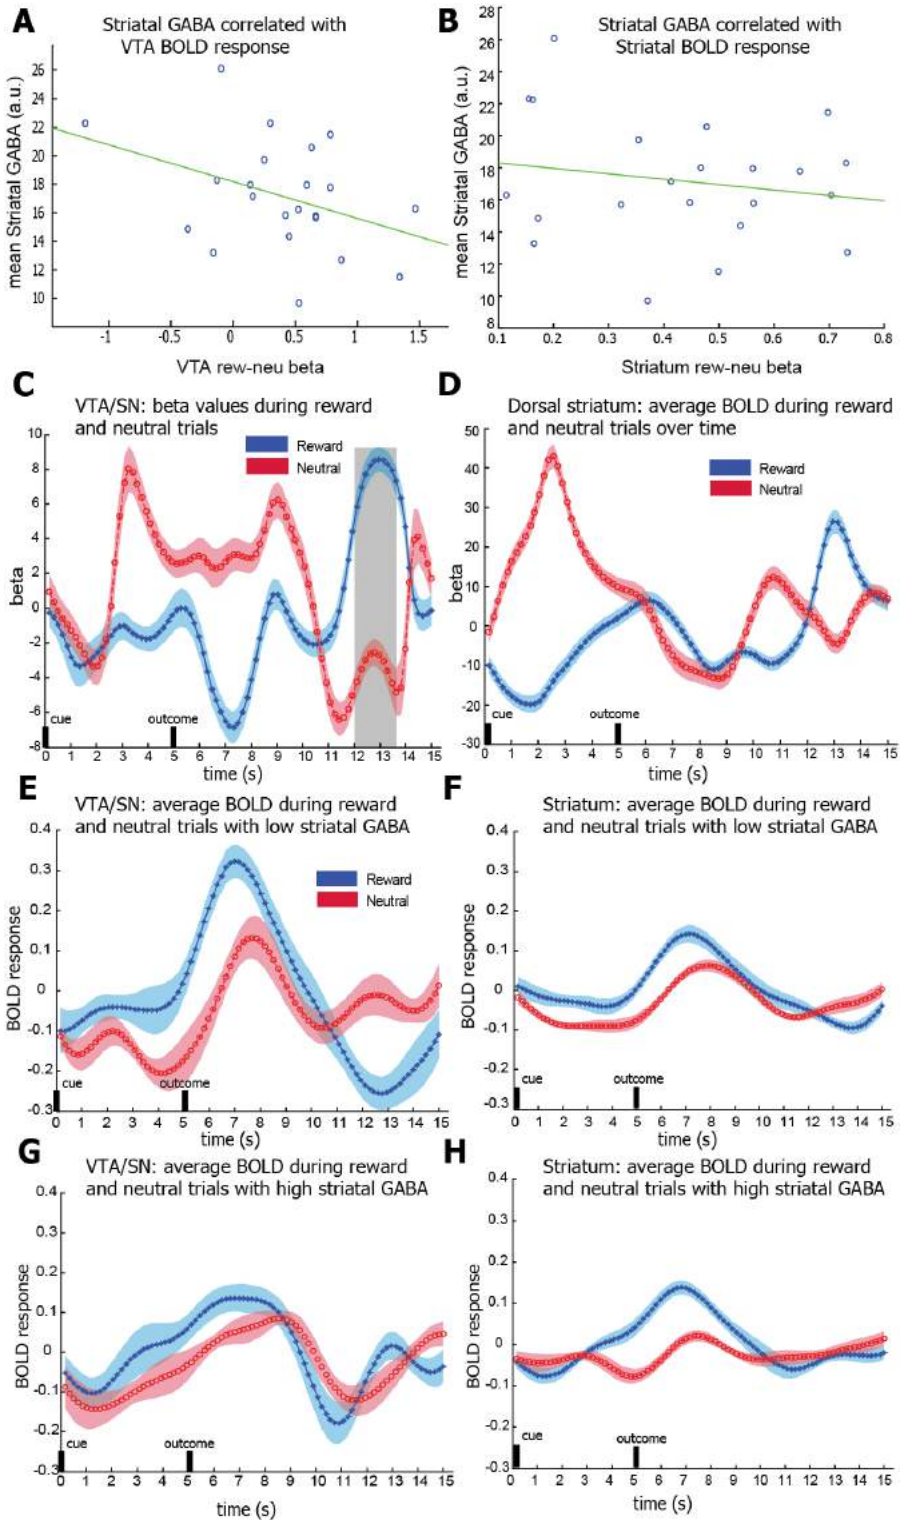

Fig. 6 Correlation between mean striatal GABA and fMRI extracted betas of BOLD response. **A.** A trend level correlation between mean striatal GABA and VTA BOLD ($p = .08$). **B.** No correlation between mean striatal GABA and dorsal striatal BOLD ($p = .41$). **C, D.** Correlation of averaged BOLD signal (within the specified ROI) with GABA over time after cue presentation, using a multiple correction approach. Error bars represent standard error. **C.** VTA/SN: we found a cluster with trend level p-value ($p = .08$) at 12 second after cue presentation. At further inspection this effect was mostly driven by the reward cue for which a positive cluster was found at $p = .11$, negative cue presentation did produce a cluster able to pass initial threshold. **E, F, G, H.** Averaged (over subjects and cue presentation) BOLD responses for reward and neutral cue split up into low GABA (E, F) and high striatal GABA groups (G, H) in the VTA/SN (E, G) and striatum (F, H). In the VTA/SN we see an increased difference in BOLD response between reward and neutral cues for low GABA versus high GABA, this is the same effect as reported in figure 6A.

4. Discussion

In this combined fMRI and MRS study we looked into the relation of reward prediction errors in the dopaminergic midbrain and striatal GABA levels. Dopaminergic midbrain structures (substantia nigra and ventral tegmental area (SN/VTA)) are associated with computation of a reward prediction. This reward prediction error is essential for learning to occur and this dopaminergic system can even be sufficient to induce associative cue reward learning, as is shown by the works of (Cohen et al., 2012). We believe that GABAergic MSNs from the striatum connect to midbrain GABAergic interneurons and provide this region with task-relevant prediction information, hence affecting the anticipation response. We hypothesized that striatal GABA may affect the reward anticipation response in the dopaminergic midbrain (SN-VTA). We used human high resolution fMRI of the dopaminergic midbrain in combination with MRS to measure GABA concentrations in the striatum, and hypothesized that individual differences in striatal GABA levels would predict differences in a cue-induced reward anticipation response in the SN-VTA. In this study we found a hint of evidence that reward prediction related BOLD in the dopaminergic midbrain correlates to the levels of GABA in the striatum. While an answer to this question would benefit from further testing in a larger sample, we can show that it is possible to reliably test GABA-concentrations in the striatum using MRS.

4.1 Correlation of brain activity with GABA concentration

The SN and VTA are responsible for calculating a reward prediction. According to the latest theories this is done by including measures of received reward through dopaminergic response and reward anticipation, through local GABAergic activity. We hypothesized that the GABAergic projections from the striatum into the VTA provide measures of reward expectation into the VTA/SN region. To be able to find a connection between striatal GABA and brain activity we measured whole-brain and local ROI activity during a reward-sensitive monetary-incentive delay task. The results seem promising towards using GABA in combination with reward-related activity, specifically in the VTA and SN.

4.1.1 ROI beta activity

Averaged beta activity in the VTA was found to indeed reflect a prediction-error-like response, in which there was no increase in activity upon presentation of a neutral cue but an increase in response to a reward cue, this effect was found to be significant. Although reward activity seems to be higher than neutral activity, no such significant difference is found in the striatum. This is a surprising, since this is regularly found in previous studies (Knutson & Greer, 2008), a reason for this might be the relatively extensive ROI; this ROI included not only nucleus accumbens, in which according to the meta-analysis by (Knutson & Greer, 2008) most of the reward related activity is found, but also pallidum, caudate nucleus and putamen. A further analysis using only the nucleus accumbens would therefore be interesting in future study. Clusters in these regions were found to be significant at whole-brain analysis. Furthermore, there is a strong significant increase in activity in reaction to reward cues visible in the thalamus which shows up at whole-brain level analysis in a reward-neutral contrast, which is also described below.

4.1.2 Effects of GABA on BOLD activity in the dopaminergic midbrain

Using averaged beta activity extracted from the SN-VTA and ventral striatum we found no significant correlation between GABA and BOLD activity. We did find a trend level correlation between GABA and VTA/BOLD activity which was in line with our expectations. This correlation shows a decreased difference between reward and neutral cue BOLD response when GABA levels in the striatum were increased. The direction in which this effect is found (negative correlation) cannot be explained using a direct connection of striatal GABAergic neurons on GABAergic interneurons in the VTA. Since this system involves increased inhibition on an inhibitory region, the net effect should be positive. A potential explanation for such a negative correlation would be a triple inhibition through an indirect connection of the striatum through the pallidum on the GABAergic interneurons in the VTA. One of the difficulties with reliably predicting the direction of this striatal to VTA connection, are the many potential pathways the signal could follow. The striatum and basal ganglia are known for their complex inhibitory influence (e.g. indirect pathway

and direct pathways) (Haber, 2011), and the exact pathway from striatum to VTA is not precisely known. Our predictions are based on results from tracer studies in rodents (Chuhma et al., 2011; Düzel et al., 2009; Houk et al., 1995). But even within these studies no complete consistent answer is given. The exact nature of the striatal (especially nucleus accumbens) to VTA connection would therefore be an interesting subject of further study.

Based on the work of Cohen et al. (2012) we expect that the influence of striatal GABA on reward prediction formation might differ over time and be more pronounced during a cue-induced anticipation response. This idea is supported by preliminary results from a study by Eshel et al. (2013) in which GABAergic input from the striatum to the VTA was severed. As a result the cue-induced anticipation response disappeared, while the reward prediction error seemed to be intact (Eshel, personal communication). According to Eshel, dissection of GABAergic input from the striatum was responsible for removing this anticipation response, while leaving outcome response intact. Moreover, mice research shows that this effect must occur through GABAergic interneurons, since GABAergic MSNs from the striatum were only found to synapse onto GABAergic and not DA cells (Chuhma et al., 2011). According to Cohen et al. (2012) GABAergic interneurons were found to reflect a reward prediction dependent on the cue, that increases its activity with a maximum around outcome anticipation (Cohen et al., 2012). Therefore highest anticipation should be around the time that this outcome is expected to arrive. We did find trend level correlations specifically around 12 seconds after cue presentation, which correlates with this maximal outcome anticipation response. This would mean that participants showed a larger anticipation response with lower striatal GABA concentrations. We must, however, consider that we did not in fact find a significant correlation; when corrected for multiple comparisons only a trend level effect was found.

This does however give an interesting prospect to future works looking into the correlation between striatal inputs into VTA reward processing. Since the effect seemed to concentrate around the time of outcome presentation, a paradigm that is more sensitive to outcome presentation would be beneficial in finding an association between VTA outcome activity and GABAergic input. This sensitivity can be increased by removing jitter between cue and outcome which would allow for better averaging. If this effect is indeed present, a slight increase in participants and a better equipped task (e.g. no time jitter in between cue and outcome) should be able to elucidate it.

4.1.3 Whole-brain activity

At a whole-brain level we found significant activity in the dorsal striatum or thalamus in reward versus neutral condition. Ventral striatal activity, notably the nucleus accumbens is usually seen in response to rewarding situations. The thalamus contains several output nuclei (Haber, 2011) and has a function in motor and action monitoring (Helmich, 2011; Lenz & Lobo, 2013). Significant activity might reflect the behavioral requirements of the task (participants had to push a button to gain a reward, and could 'lose' the chance when their reaction was too slow, while in the neutral condition reaction speed had no real consequences). Therefore the shift to dorsal activation might indicate an increase in striatal action monitoring of response to the potential reward situation. The second cluster of significant activity was found in the border of the anterior cingulate cortex (ACC). These regions are commonly associated with cognitive control, and conflict monitoring (Botvinick, Braver, Barch, Carter, & Cohen, 2001). This suggests that participants show increased monitoring during rewarding conditions in which their actions have stronger consequences. At lower threshold we did still find cluster level significant activity for the dorsal striatum (nucleus accumbens, caudate and pallidum) showing that these regions are indeed active during reward.

4.2 Striatal GABA can be reliably measured using 3T MRI

We were able to measure GABA in the striatum while maintaining a decent level of reliability. We are to our knowledge the first to successfully perform GABAergic MRS in the striatum. Previous studies using GABAergic MRS have predominantly focused on cortical regions such as the occipital cortex and motor cortex (Jocham, Hunt, Near, & Behrens, 2012; Muthukumaraswamy, Edden, Jones, Swettenham, & Singh, 2009; Northoff et al., 2007). Since these regions have a smaller distance to the receiver coils, cortical regions generally provide a better signal-to-noise ratio and signal and are therefore easier to measure. However by increasing the number of averages for subcortical voxels and by using new Fastestmap shimming sequences (Xiaodong Zhong, 2013) we were able to obtain a sufficient line width and signal-to-noise ratio to get reproducible results

in the striatum. Within-subject reliability was tested by bilaterally measuring GABA in the striatum. Previous mice studies have shown that GABA levels between bilateral regions do not significantly differ in most brain regions (Starr & Kilpatrick, 1981) and for all regions variation within bilateral regions differed substantially less than variation between brain regions. Bilateral GABA measures should therefore be significantly correlated, and can be used to give a measure of within-subject reliability. Such a significant correlation was indeed found. Moreover, our results show that there is a significant correlation between both striatal voxels but not between both striatal voxels and the occipital cortex. This shows that the effect of GABA is not based on a brain wide increase or decrease of GABA but is specific to the striatum itself. We hereby show that, although previously avoided due to increased difficulty, GABA concentration measurements using an MRS in an 8ml voxel are indeed possible and can result in reliable results.

We showed that it is possible to localize the VTA and SN using SWI imaging and use this to make subject specific masks. We suggest a method of using EEG based Monte Carlo cluster analysis using a permutation approach to allow for correction of multiple comparisons when analyzing BOLD activity differences, or correlations over different time points.

5. Conclusions

Our results provide a tantalizing hint of a relation between GABA and VTA BOLD activity but require further testing in a larger sample. We provide a combination of new methods that allow for better specialized research to occur in the future, in this research field as well as in other fields. We are to our knowledge the first to use MRS to measure GABA in the striatum, and show that one can indeed generate reliable spectra using this method.

We found significant activity in the thalamus and ACC during reward versus neutral cue presentation. Although dorsal striatal activity (nucleus accumbens, caudate and pallidum) is usually more prominent, we suggest that this shift in activity is associated with the strong action anticipation difference between the two conditions that causes differences in monitoring activity of the thalamus and ACC. At ROI level we found a significant difference between reward and neutral cues in the SN-VTA. Beta activity during the neutral cue was negligible while this was not the case in the striatum. This promotes the idea of a VTA specific reward prediction signal. Further research is needed to find if striatal GABA is indeed responsible for providing a measure of reward prediction.

6. Acknowledgements

We would like to thank Paul Gaalman, Silvy Collin, Markus Barth, Erno Hermans and Julia Uddén for their technical support.

7. References

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. E. (2006). Reward-motivated learning: mesolimbic activation precedes memory formation. *Neuron*, *50*(3), 507-517.

Andersson, J. L. R., Hutton, C., Ashburner, J., Turner, R., & Friston, K.. (2001). Modeling geometric deformations in EPI time series. *Neuroimage*, *13*(5), 903-919.

Ashburner, J., & Friston, K. J. (2005). Unified segmentation. *Neuroimage*, *26*(3), 839-851.

Ballard, I.C., Murty, V. P., Carter, R. M., MacInnes, J. J., Huettel, S. A., & Adcock, R. A. (2011). Dorsolateral prefrontal cortex drives mesolimbic dopaminergic regions to initiate motivated behavior. *Journal of Neuroscience*, *31*(28), 10340-10346.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*(3), 624.

Chuhma, N., Tanaka, K. F., Hen, R., & Rayport, S. (2011). Functional connectome of the striatal medium spiny neuron. *Journal of Neuroscience, 31*(4), 1183-1192.

Cohen, J. Y., Haesler, S., Vong, L., Lowell, B. B., & Uchida, N.. (2012). Neuron-type-specific signals for reward and punishment in the ventral tegmental area. *Nature*, *482*(7383), 85-88.

Colombo, M.. (2014). Deep and beautiful. The reward prediction error hypothesis of dopamine. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, *45*, 57-67.

D'Ardenne, K., McClure, S. M., Nystrom, L. E., & Cohen, J. D. (2008). BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science*, *319*(5867), 1264-1267.

Damier, P., Hirsch, E. C., Agid, Y., & Graybiel, A. M. (1999). The substantia nigra of the human brain I. Nigrosomes and the nigral matrix, a compartmental organization based on calbindin D28K immunohistochemistry. *Brain*, *122*(8), 1421-1436.

Düzel, E., Bunzeck, N., Guitart-Masip, M., Wittmann, B., Schott, B. H., & Tobler, P. N. (2009). Functional imaging of the human dopaminergic midbrain. *Trends in Neurosciences*, *32*(6), 321-328.

Eshel, N., Tian, J., & Uchida, N.. (2013). Opening the black box: dopamine, predictions, and learning. *Trends in Cognitive Sciences*, *17*(9), 430-431.

Friston, K. J., Williams, S., Howard, R., Frackowiak, R. S. J., & Turner, R.. (1996). Movement-related effects in fMRI time-series. *Magnetic Resonance in Medicine*, *35*(3), 346-355.

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.

Haber, S. N. (2011). Neuroanatomy of reward: a view from the ventral striatum. In: Gottfried JA, editor. Neurobiology of sensation and reward. Boca Raton, FL: CRC.

Helmich, Rick. (2011). Cerebral Reorganization in Parkinson's Disease: Universiteitsbibliotheek Nijmegen [host].

Holland, P. C. (1984). Unblocking in Pavlovian appetitive conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *10*(4), 476.

Houk, J. C., Adams, J. L., & Barto, A. G. (1995). A model of how the basal ganglia generate and use neural signals that predict reinforcement. *Models of information processing in the basal ganglia*, 249-270.

Hutton, C., Deichmann, R., Turner, R., & Andersson, J. (2004). Combined correction for geometric distortion and its interaction with head movement in fMRI. Paper presented at the Proc. of the 12th Scientific Meeting and Exhibition of the ISMRM.

Hutton, C., Bork, A., Josephs, O., Deichmann, R., Ashburner, J., & Turner, R.. (2002). Image distortion correction in fMRI: a quantitative evaluation. *Neuroimage*, *16*(1), 217-240.

Iglesias, S., Mathys, C., Brodersen, K. H., Kasper, L., Piccirelli, M., den Ouden, H. E. M., & Stephan, K. E. (2013). Hierarchical prediction errors in midbrain and basal forebrain during sensory learning. *Neuron*, *80*(2), 519-530.

Jocham, G., Hunt, L. T., Near, J., & Behrens, T. E. J. (2012). A mechanism for value-guided choice based on the excitation-inhibition balance in prefrontal cortex. *Nature Neuroscience*, *15*(7), 960-961.

Kanaani, J., Kolibachuk, J., Martinez, H., & Baekkeskov, S.. (2010). Two distinct mechanisms target GAD67 to vesicular pathways and presynaptic clusters. *Journal of Cell Biology*, *190*(5), 911-925.

Klein-Flügge, M. C., Hunt, L. T., Bach, D. R., Dolan, R. J., & Behrens, T. E. J. (2011). Dissociable reward and timing signals in human midbrain and ventral striatum. *Neuron*, *72*(4), 654-664.

Knutson, B., & Greer, S. M. (2008). Anticipatory affect: neural correlates and consequences for choice. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *363*(1511), 3771-3786.

Knutson, B., Westdorp, A., Kaiser, E., & Hommer, D.. (2000). FMRI visualization of brain activity during a monetary incentive delay task. *Neuroimage*, *12*(1), 20-27.

Lenz, J. D., & Lobo, M. K. (2013). Optogenetic insights into striatal function and behavior. *Behavioural Brain Research*, *255*, 44-54.

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

Mescher, M., Merkle, H., Kirsch, J., Garwood, M., & Gruetter, R. (1998). Simultaneous in vivo spectral editing and water suppression. *NMR in biomedicine*, *11*(EPFL-ARTICLE-177509), 266-272.

Mescher, M., Tannus, A., Johnson, M., & Garwood, M. (1996). Solvent suppression using selective echo dephasing. *Journal of Magnetic Resonance*, Series A, *123*(2), 226-229.

Mullins, P. G., McGonigle, D. J., O'Gorman, R.L., Puts, N. A. J., Vidyasagar, R., Evans, C. J., & Edden, R. A. E. (2014). Current practice in the use of MEGA-PRESS spectroscopy for the detection of GABA. *Neuroimage*, *86*, 43-52.

Muthukumaraswamy, S. D., Edden, R. A. E., Jones, D. K., Swettenham, J. B., & Singh, K. D. (2009). Resting GABA concentration predicts peak gamma frequency and fMRI amplitude in response to visual stimulation in humans. *Proceedings of the National Academy of Sciences*, *106*(20), 8356-8361.

Niv, Y., Edlund, J. A., Dayan, P., & O'Doherty, J. P. (2012). Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *Journal of Neuroscience*, *32*(2), 551-562.

Northoff, G., Walter, M., Schulte, R. F., Beck, J., Dydak, U., Henning, A., . . . Boesiger, P.. (2007). GABA concentrations in the human anterior cingulate cortex predict negative BOLD responses in fMRI. *Nature Neuroscience*, *10*(12), 1515-1517.

Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J. M.. (2010). FieldTrip: open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational intelligence and neuroscience*, 2011.

Pijnappel, W. W. F., Van den Boogaart, A., De Beer, R., & Van Ormondt, D. (1992). SVD-based quantification of magnetic resonance signals. *Journal of Magnetic Resonance* (1969), *97*(1), 122-134.

Puts, N. A.J., & Edden, R. A. E. (2012). In vivo magnetic resonance spectroscopy of GABA: A methodological review. *Progress in nuclear magnetic resonance spectroscopy*, *60*, 29-41.

Redgrave, P., Vautrelle, N., & Reynolds, J. N. J. (2011). Functional properties of the basal ganglia's re-entrant loop architecture: selection and reinforcement. *Neuroscience*, *198*, 138-151.

Rowley, N. M., Madsen, K. K., Schousboe, A., & Steve White, H. (2012). Glutamate and GABA synthesis, release, transport and metabolism as targets for seizure control. *Neurochemistry international*, *61*(4), 546-558.

Russo, S. J., & Nestler, E. J. (2013). The brain reward circuitry in mood disorders. *Nature Reviews Neuroscience*, *14*(9), 609-625.

Schultz, W.. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, *80*(1), 1-27.

Sescousse, G., Caldú, X., Segura, B., & Dreher, J. C.. (2013). Processing of primary and secondary rewards:

a quantitative meta-analysis and review of human functional neuroimaging studies. *Neuroscience & Biobehavioral Reviews*, *37*(4), 681-696.

Sotero, R. C., & Trujillo-Barreto, N. J. (2007). Modelling the role of excitatory and inhibitory neuronal activity in the generation of the BOLD signal. *Neuroimage*, *35*(1), 149-165.

Stagg, C. J., Bachtiar, V., & Johansen-Berg, H.. (2011). What are we measuring with GABA magnetic resonance spectroscopy. *Communicative & Integrative Biology*, *4*(5), 573-575.

Starr, M. S., & Kilpatrick, I. C. (1981). Bilateral asymmetry in brain GABA function? *Neuroscience Letters*, *25*(2), 167-172.

Steinberg, E. E., Keiflin, R., Boivin, J. R., Witten, I. B., Deisseroth, K., & Janak, P. H. (2013). A causal link between prediction errors, dopamine neurons and learning. *Nature Neuroscience*, *16*(7), 966-973.

Stocco, A., Lebiere, C., & Anderson, J. R. (2010). Conditional routing of information to the cortex: A model of the basal ganglia's role in cognitive coordination. *Psychological Review*, *117*(2), 541.

Takahashi, Y. K., Roesch, M. R., Wilson, R. C., Toreson, K., O'Donnell, P., Niv, Y., & Schoenbaum, G.. (2011). Expectancy-related changes in firing of dopamine neurons depend on orbitofrontal cortex. *Nature Neuroscience*, *14*(12), 1590-1597.

Vanhamme, L., van den Boogaart, A., & Van Huffel, S.. (1997). Improved method for accurate and efficient quantification of MRS data with use of prior knowledge. *Journal of Magnetic Resonance*, *129*(1), 35-43.

Watabe-Uchida, M., Zhu, L., Ogawa, S. K., Vamanrao, A., & Uchida, N.. (2012). Whole-brain mapping of direct inputs to midbrain dopamine neurons. *Neuron*, *74*(5), 858-873.

Zhong, X., Lyubich, Y. M., DeVito, T., Shah, S., Knight-Scott, J. (2013). Quantitative Comparison of Shim Algorithms for In Vivo 1H-MRS. *Proceedings of the International Society for Magnetic Resonance in Medicine*, 21.

Predicting Speech: How Semantic Context and Visual Cues Modulate Audiovisual Speech Processing

Heidi Solberg Økland[1]
Supervisors: Ana Todorovic[1], Claudia S. Lüttke[1], James M. McQueen[1,2,3], Floris P. de Lange[1]

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

Spoken language communication usually happens face-to-face. Both the content of what a speaker has already said and her visible mouth movements (visemes) can help us predict which word we will hear next, because both of these cues precede the acoustic onset of the upcoming word. However, it is not clear whether and how these two types of predictions interact when we perceive speech audiovisually. We orthogonally manipulated contextual constraint and viseme saliency to investigate whether a previously found auditory facilitation effect caused by strong visemes would be modulated by semantic context. Our results support a mechanism whereby a strong semantic prediction reduces the predictive benefit of a strong viseme, indicating that predictions based on sentence context and visemes are made on different levels of the processing hierarchy. We conclude that strong visemes facilitate early auditory speech processing only when there is a high amount of uncertainty about the upcoming word.

Keywords: audiovisual speech, magnetoencephalography, hierarchical predictive coding, viseme, N100m, N400

Corresponding author: Heidi Solberg Økland; **E-mail:** heidi.s.okland@outlook.com

1. Introduction

Ever since we started learning our mother tongue as babies, we have heard and said countless sentences while talking to other people. Since the majority of these conversations happen while seeing the person we talk to, we are used to perceiving speech with both auditory and visual components. A growing body of literature suggests that we use our extensive experience with language to anticipate upcoming words (see e.g. Altmann & Mirkovic, 2009; Federmeier, 2007; Pickering & Garrod, 2013). One recent study showed that a strongly constraining sentence context enables us to pre-activate the very form of a word before we read it (DeLong, Urbach, & Kutas, 2005). Participants read sentences which enabled them to predict a certain noun (e.g. "The day was breezy so the boy went outside to fly...,"), and the researchers measured electrical brain activity time-locked to the onset of the article preceding the noun ("a" or "an"). When participants were expecting a consonant-initial noun like "kite", reading "an" (incompatible with "kite") compared to "a" caused enhanced electrical brain activity (the "N400-effect", Kutas & Hillyard, 1980, for recent reviews see Kutas & Federmeier, 2011; Lau, Phillips, & Poeppel, 2008). This result strongly suggests that participants had a specific word ("kite") in mind before they read it (for similar evidence see Altmann & Kamide, 1999). Thus, we seem to make predictions about both the meaning and the pronunciation of upcoming words based on semantic context.

Analogously, the visual articulation of a speech sound (the viseme) can be predictive of the acoustic properties of that speech sound. If you for example see someone closing their lips, you would expect to hear a /b/, /p/ or /m/ sound next. Conversely, if someone makes a /g/ sound, it is less clearly visible from the lip movements what sound it is, since /g/ is articulated in the back of the vocal tract. In sum, visemes involving visible articulations (strong visemes) are more salient than ones involving less visible articulations (weak visemes). Visemes typically precede the onset of vocal cord vibration at word onset (Chandrasekaran, Trubanova, Stillittano, & Ghanzanfar, 2009; Schwartz & Savariaux, 2014), and communicate information about the place of articulation of a speech sound earlier and more efficiently than the acoustic signal alone (Jesse & Massaro, 2010). Therefore, a viseme has the potential to constrain the possibilities for what we expect to hear just like a sentence context can.

Two recent studies on audiovisual syllable recognition show that predictive visemes facilitate auditory speech processing (Arnal, Morillon, Kell, & Giraud, 2009; van Wassenhove, Grant, & Poeppel, 2005). In both studies, participants watched videos of spoken syllables (e.g. "ba" and "ga") while their brain activity was recorded using electroencephalography (EEG) or magnetoencephalography (MEG). The researchers investigated the effect of weak versus strong visemes on auditory processing by means of the auditory N100(m). The auditory N100(m) is an event-related potential/field (ERP/ERF) peaking around 100 ms following a change in the acoustic environment (Hari et al., 1987). Although traditionally seen as a stimulus-driven response to an auditory event, the auditory N100(m) has also been found to be sensitive to higher-level cognitive states like attention and expectation (Hari et al., 1987; Näätänen & Picton, 1987; Todorovic, van Ede, Maris, & de Lange, 2011). It was shown that the auditory N100(m) peaked earlier when visemes were strong (Arnal et al., 2009; van Wassenhove et al., 2005, see also Brunellière, Sánchez-García, Ikumi, & Soto-Faraco, 2013). The shortening of the auditory N100(m) latency by strong visemes, henceforth the 'latency shortening effect', has been interpreted as a sensory anticipation effect, such that seeing someone closing their lips to say "ba" enables the auditory cortex to form a more accurate prediction about the upcoming sound compared to seeing "ga". Thus, due to the auditory prediction the N100(m) peaks earlier for speech sounds with strong visemes. This suggests that the influence of visual speech on our auditory speech perception depends on how salient the visual information is.

So far, the latency shortening effect has only been shown in experiments where participants had no prior expectations about upcoming speech. Since constraining sentence contexts also causes expectations about upcoming words, we wondered whether these two types of predictive processes might interact. Specifically, we asked whether a strong high-level expectation about an upcoming word, like that from a constraining sentence, would influence the low-level audiovisual processing benefit of strong visemes indexed by the shortening of the auditory N100m latency. We hypothesized that a strong prediction about an upcoming word derived from the sentence context would make the auditory prediction caused by the viseme redundant. This hypothesis would be in line with hierarchical predictive coding (Friston, 2005; Friston & Kiebel, 2009; Wacongne et al., 2011), and will be referred to as the 'hierarchical hypothesis' in the following. An alternative hypothesis would be that semantic context

does not influence low-level audiovisual sensory predictions, in which case we would not expect to see any modulation of the latency shortening effect by sentence context. There is ample evidence from behavioural experiments to suggest that contextual information can change or bias our perception of auditory words (see e.g. review by Davis & Johnsrude, 2007). Recent work has also shown that a biasing sentence context influences the perception of audiovisually incongruent words (Windmann, 2004). However, it is not known whether sentence context can influence the very early stage of sensory processing indexed by the auditory N100m, which is what we set out to test.

In order to find out whether semantic prediction would interact with the latency shortening effect, we designed an experiment where sentence-level and viseme-level predictions were manipulated orthogonally. Specifically, we constructed sentences that induced either strong or weak predictions about the final word, and let the sentence-final words start with speech sounds that had either strong or weak visemes. If the N100m measured relative to the auditory onset of these words only show a latency difference when the words are weakly predicted by the sentence context, but not when the words are strongly predicted, this would be evidence in support of the hierarchical hypothesis. Alternatively, the N100m latency effect could be completely insensitive to the sentence context manipulation, in which case we would not observe any differences depending on the strength of the sentence-level prediction. This latter result would indicate that hierarchical predictive coding is not a plausible mechanism in speech perception. In sum, the hierarchical hypothesis predicts an N100m latency difference only in the weakly constraining sentences; whereas the alternative hypothesis predicts that the effect should be present regardless of whether the sentence context is weakly or strongly constraining. In addition to these two hypotheses, we expected strongly constraining sentence contexts to produce an attenuated N400 component regardless of the viseme manipulation, replicating the well-known N400 effect of semantic context (for recent reviews see Kutas & Federmeier, 2011; Lau et al., 2008).

To investigate the question of whether and how sentence predictions and viseme predictions interact, we measured the auditory N100m latency using MEG. Healthy participants watched videos of a person uttering words which started with either a visually salient (strong) or non-salient (weak) viseme in strongly and weakly constraining sentence contexts. We included a condition without any sentence preceding the target words as used in previous studies to be able to compare the results of the current study with previous findings. Here, we expected to replicate the results found in the studies using syllables, namely a latency shortening effect for strong visemes. Our results are consistent with the hierarchical hypothesis, suggesting that the brain relies on the information provided by visual speech only when there is a high amount of uncertainty about the upcoming word.

2. Materials and methods

2.1 Participants

Twenty-four Dutch native speakers (6 male) aged 19-28 (M = 22.25, SD = 2.17) participated in the experiment after signing an informed consent form in accordance with the Declaration of Helsinki. None of the participants reported having any neuropsychiatric disorders. All were right-handed, had normal hearing and normal or corrected-to-normal vision. One additional participant was excluded due to excessive artifacts caused by makeup residue that had been magnetized during magnetic resonance imaging (MRI) earlier that day. The study was approved by the local ethics committee (CMO region Arnhem/Nijmegen), and participants were paid for their participation.

2.2 Experimental materials, design and procedure

Participants watched videos of the target words without any context as well as in both strongly and weakly constraining context, while we recorded brain activity using MEG. To study the hypothesized interaction between sentence context and viseme, we orthogonally manipulated these two parameters to be either strongly constraining (henceforth "strong") or weakly constraining (henceforth "weak"). Our measure of interest was brain activity related to the sentence-final audiovisual words (Table 1, see also timeline in Fig. 1). The four experimental conditions (examples are given in Table 2) were weak viseme/weak context, strong viseme/weak context, weak viseme/strong context and strong viseme/strong context.

2.2.1 Target words

For the target words belonging to the strong and weak viseme conditions, we kept the acoustic

features both across as well as within the conditions constant, while manipulating the saliency of the viseme. We selected two Dutch speech sounds that are phonetically similar but differ in visual saliency: the unvoiced labio-dental fricative /f/ (the same sound as an English f) and the unvoiced velar fricative /x/ (similar to the final sound in 'Bach', the famous German composer). Fricatives are consonant sounds we produce by forcing air through a narrow space. In the case of an /f/, the space is between the upper teeth and lower lip, making it visually salient. For /x/, on the other hand, the narrow space is created by positioning the back of the tongue close to the soft palate, resulting in a less salient viseme. A Dutch auditory gating study has shown that people can correctly distinguish /f/ and /x/ after having heard approximately equal portions of the full sound (Smits, Warner, McQueen, & Cutler, 2003). Furthermore, it was also shown that these two sounds are often confused with each other early in the auditory recognition process. Because /f/ and /x/ are acoustically and perceptually similar, both have the potential to benefit from additional visual information.

We chose four words starting with each of these speech sounds followed by the vowel /I/ to further ensure acoustic similarity, which resulted in eight target words (Table 1): four words starting with /f/ (strong viseme), and four words starting with /x/ (weak viseme). The word frequencies (quantified as occurrence per million words) were extracted from the SUBTLEX-NL database (Keuleers, Brysbaert, & New, 2010). The mean frequencies of the words in the two conditions were not significantly different ($t(6) = .145$, $p = .89$). We determined acoustic duration by visual inspection of the sound clips in Praat (Boersma, 2001) and visual duration (i.e. onset of lip-movements to offset of lip-movements) by visual inspection of the video clips in Adobe Premiere Pro 6 (see Table 1). To minimize the effect of word repetition we also selected 16 filler words that were semantically related but acoustically distinct from the target words (e.g. "verband" (bandage) and "klei" (clay) for the target word "gips" (plaster)). None of the filler words started with /f/ or /x/ in order to increase the number of possible onset sounds and thus reduce the amount of strategic processing.

2.2.2 Sentences

We constructed 24 sentences to go with each of the 8 target words (12 strongly constraining, 12 weakly constraining), leading to a total of 192 initial target sentences. After pilot testing on a group of 40 Dutch students (native speakers with no dyslexia) with a pen-and-paper sentence completion test, we selected the 10 most constraining and 10 least constraining sentences for each target word, which gave us a final set of 160 target sentences (see Table 2 for examples). A given sentence was considered sufficiently constraining towards the target word when the completions were largely in agreement and the target word was the most frequent one, and unconstraining if completions were inconsistent and did not include the target word. We also constructed 10 sentences for each filler word (5 strongly constraining, 5 weakly constraining), so that the proportion of target and filler trials was 50/50.

Table 1. The eight target words.

Target words	English translation	IPA	Frequency	Acoustic word duration	Visual word duration	Total clip duration
fik	fire	[fɪk]	6.49	293 ms	440 ms	1000 ms
film	film	[fɪlm]	174.28	372 ms	720 ms	1160 ms
filter	filter	[ˈfɪltər]	2.04	541 ms	640 ms	1040 ms
fit	in good shape	[fɪt]	5.37	373 ms	600 ms	1040 ms
gif	poison	[xɪf]	13.56	400 ms	960 ms	1320 ms
gil	scream/ shriek	[xɪl]	9.99	330 ms	640 ms	1040 ms
gisteren	yesterday	[ˈxɪstərə(n)]	131.79	513 ms	640 ms	1000 ms
gips	plaster/gyp-sum	[xɪps]	2.36	492 ms	680 ms	1040 ms

IPA = International Phonetic Alphabet. Frequency did not differ significantly between f- and g-words (see 2.2.1).

Table 2. The four experimental conditions with example sentences. The English translations (in italics) are semi-literal.

	Strong viseme	Weak viseme
Strong context	Het brandhout vloog meteen in de FIK. *The firewood suddenly caught FIRE.*	De tanden van een cobra bevatten dodelijk GIF. *The teeth of a cobra contain deadly POISON.*
Weak context	Toen Roel thuiskwam stonden zijn spullen in de FIK. *When Roel came home his stuff was on FIRE.*	De biologiestudenten lezen een artikel over GIF. *The biology students read an article about POISON.*

To validate our semantic context manipulation we checked that weakly constraining sentences elicited the N400 effect (see 3.2 below).

2.2.3 Video recording and editing

A male native Dutch speaker with a neutral dialect was chosen as speaker. He was seated in a chair with the camera directly in front of him, with a soft box lighting device behind the camera to ensure maximal visibility of the facial movements. The speaker was instructed to speak clearly and at a natural speaking rate, but to include a pause before the sentence-final word to mirror the experimental stimulus presentation (see 2.2.4). He said each sentence at least five times. The recording was done in blocks such that all sentences ending with the same target word were recorded "in one go", with short breaks in between. This was done to enable the embedding of one single target word video in all 20 sentences belonging to that word. The recording was done over two consecutive days in a soundproof room by means of a digital HD video camera (JVC HD GY-HM100E) at 1920x1080 resolution and 25 progressive frames per second. Sound was recorded with the camera microphone at a sampling rate of 48 kHz.

First, video clips (i.e. one of each sentence and one of each word) were selected and cut out from the raw recordings in Adobe Premiere Pro 6 such that every sentence and word video started and ended with a neutral face and a closed mouth. The video clips of the lead-in sentences had six frames (240 ms) of neutral face at the beginning and three frames (120 ms) at the end, whereas the word video clips started with three frames of neutral face and ended with six frames of neutral face. This enabled us to control the visual gap between lead-in sentence and target word. The sound was noise-reduced in Audacity and amplitude was root mean square equalized in Praat. We put the edited sound clips back into the video clips in Adobe Premiere Pro 6 without any audiovisual asynchrony or realignment, after which we exported the edited video clips in uncompressed AVI format with a 720x480 frame size. Finally, the video clips were compressed and exported as AVI files in the Indeo 5.10 codec using VirtualDub.

2.2.4 Stimulus presentation

The target words were presented both in a sentence context and without a sentence context. During sentence blocks, a target word was never repeated on two consecutive trials and the proportion of fillers and targets was balanced across blocks. Each trial started with a fixation cross placed approximately between the eyes of the speaker in the video, after which the video of the sentence started playing. When the sentence video ended, there was a gap of 320 ms where a black screen was presented (see Fig. 1 for an illustration), before the target word video was presented. The rationale behind having this gap was that we wanted to use the same word video clip in all 20 sentences belonging to the same word, and to keep the interval between (visual) sentence offset and word onset constant (i.e. at 560 ms, see 2.2.3). Word-only trials were identical to sentence trials, except for the sentence video and the black screen following it. The order of presentation was pseudorandom. Importantly, the auditory onset in the word videos naturally lagged the viseme onset by 50 ms, and was constant within and across conditions.

2.2.5 Procedure

Participants watched videos of a person saying sentences and words while their brain activity was recorded with MEG. The video was projected on a screen at a distance of 70 cm from the participant and subtended a horizontal visual angle of 25.8°

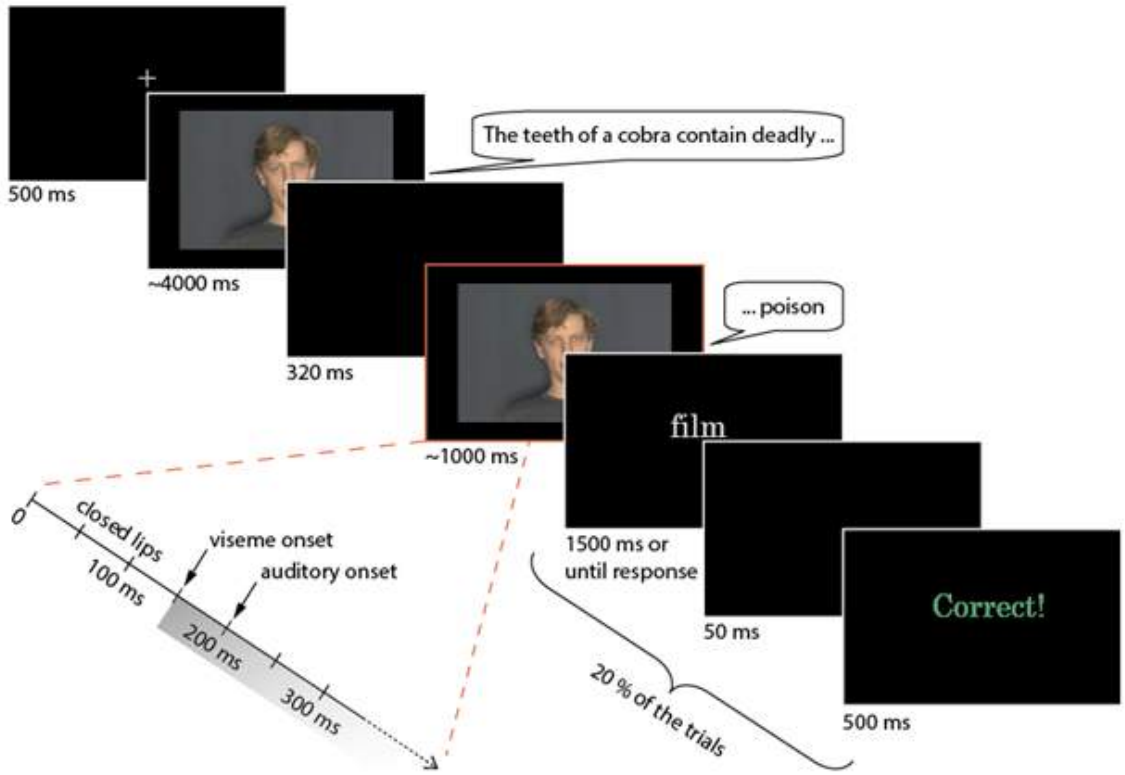

Fig. 1 Example of a sentence trial with task. Participants watched audiovisually presented sentences that were either strongly or weakly constraining with regards to the target words. These words began with speech sounds that had either a visually salient articulation (strong viseme) or a visually non-salient articulation (weak viseme). Target words were presented sentence-finally after a brief gap. The target word video started with the speaker having closed lips, before the lips started moving at 150 ms. The acoustic word onset was 50 ms after viseme onset in both viseme conditions. In some trials a written word appeared immediately after the sentence-final word video, and participants indicated with a button press whether this word matched the one they just heard.

and vertical visual angle of 17.9°. The sound was delivered binaurally through MEG-compatible air tubes at a comfortable sound pressure level. Stimulus presentation was controlled by a PC running Presentation software (Neurobehavioral Systems). There were 8 sentence blocks of 40 trials, of which half were fillers, yielding 160 filler trials and 160 target trials. After the sentence blocks there was one word-only block with 20 trials per target word (i.e. the same number of repetitions as in the sentence part), yielding a total of 160 trials. The inter-trial interval varied randomly between 1 and 2 seconds throughout the experiment. All participants saw the same set of stimuli, but the order of presentation differed per participant.

Participants were told that their primary focus should be on attending to what was being said in the videos, but that they would also have to perform a task on some trials. This task appeared in 20% of the sentence trials and 35% of the word trials. In a task trial, a written word appeared on the screen immediately after the word video, and the participants had to indicate with a left/right button press whether this word was different from the word they just heard. We included this task to ensure that participants attended to the stimuli. The task was evenly distributed both within and across conditions, but as a) it was predetermined which trial would be a task trial, and b) the list of trials was randomized, task trials were not evenly distributed across time. Since we did not intend to analyze reaction times, the participants were not specifically instructed to respond as fast as possible, but rather to respond before the word disappeared from the screen (i.e. within 1500 ms).

Before the start of the experiment, participants were instructed to look at the face in the video as they would when having a face-to-face conversation, but to move their eyes as little as possible. To minimize artifacts in the MEG recordings, the experimenter also asked them to try sit as still as possible and to blink between trials. Before the start of the experiment they were shown two consecutive trials with a task so that they knew what to expect and how to perform the task when the experiment started. After this the experimenter asked them whether

everything was clear, and whether they were ready to start. Participants then started the experiment with a button press. There was a break of 30 seconds after every block of 40 sentence trials and a longer break of 60 seconds before the word-only block. After the built-in breaks participants could press a button to proceed with the experiment whenever they wanted. The experimenter advised participants before the start of the experiment to take longer breaks if they felt tired, and talked to them over the intercom during the breaks to check how they were doing and to motivate them.

2.2.6 MEG acquisition

We recorded magnetic brain activity using whole-head MEG (275 axial gradiometers, VSM/CTF Systems) at a sampling rate of 1200 Hz in a magnetically shielded room. Participants' head position was monitored during the experiment using coils placed at the nasion (i.e. between the eyes) and in both ear canals, and was corrected during breaks if needed. Horizontal and vertical electro-oculogram (EOG) was recorded using 10-mm-diameter Ag–AgCl surface electrodes. We used the vertical EOG to aid offline rejection of blink artifacts (see 2.3.1). We monitored the horizontal EOG during the experiment to make sure that participants were not moving their eyes, but we did not analyze it offline.

2.3 MEG data analysis

The MEG data was preprocessed and analysed using the FieldTrip toolbox developed at the Donders Institute for Brain, Cognition, and Behaviour (Oostenveld, Fries, Maris, & Schoffelen, 2011) in Matlab 7, version 2011b. We calculated event-related fields (ERFs) time-locked to the onset of the target word videos, and used these to analyse both auditory N100m latencies as well as early and late amplitude effects.

2.3.1 Calculation of event-related fields

First, we extracted epochs of 1300 ms (henceforth 'trials') from the data starting 300 ms before the onset of the target word videos. Trials containing jumps in the MEG signal caused by the superconducting quantum interference device (SQUID) electronics were rejected based on visual inspection. Trials containing excessive muscle artifacts were then rejected after using a 110-140 Hz band pass filter and subsequent visual inspection of the amount of variance in an epoch (i.e. outlier trials with a relatively large amount of variance were discarded). Next, we used independent component analysis (Bell & Sejnowski, 1995) to compute and remove variance in the signal pertaining to eye blinks (Jung et al., 2000). Finally, we discarded any remaining trials where the ERF amplitude was more than 4 standard deviations above the mean. This resulted in an average of 12.75 (SD = 4.80) rejected trials per participant.

Before calculating ERFs for each condition of interest, we low-pass filtered the data at 40 Hz and baseline corrected relative to a 100 ms time window before the onset of the word video (i.e. while the screen was blank in the sentence blocks and while the fixation cross was on the screen in the word-only block). Finally, we calculated planar gradient transformed ERFs (Bastiaansen & Knösche, 2000). This simplifies the interpretation of the sensor-level data because it places the maximal signal above its source (Hämäläinen, Hari, Ilmoniemi, Knuutila, & Lounasmaa, 1993).

2.3.2 Left temporal sensors of interest

We constrained our analyses of auditory latency and amplitude effects to temporal sensors. Since the latency difference for words in isolation has already been found by others in the left hemisphere (Arnal et al., 2009; van Wassenhove et al., 2005), we decided to focus our analysis accordingly. To select the sensors of interest, we used the grand averaged ERF data from words presented in isolation to find the 10 most active temporal sensors in the left hemisphere 60-100 ms after auditory onset, i.e. when amplitude was maximal during the time window of the auditory N100m. We restricted the sensor choice to 10 because this would give us an acceptable signal-to-noise ratio (which was especially important for the latency analysis) at the same time as allowing us to get a reasonable coverage of auditory areas. We excluded sensors bordering occipital sensors to avoid contamination by visual activity. The selected sensors are shown in Figure 2.

2.3.3 Analysis of auditory N100M peak latency

In order to calculate the auditory N100M peak latencies for the target words we used the jackknife approach on ERFs averaged over the sensors of interest. Jackknifing is a method that allows robust estimation of latency differences (Miller, Patterson,

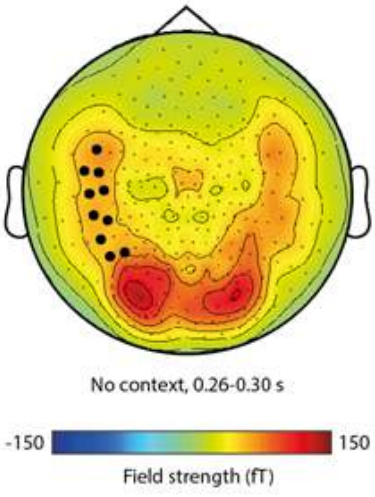

Fig. 2 Sensors of interest selected from the most active left temporal sensors when audiovisual words were presented without a sentence context and the auditory N100m reached its maximum (see also Fig. 4 and 5). The topography shows the magnetic brain activity averaged over the same time window, viewed from the top with the front of the head upwards.

& Ulrich, 1998). Instead of computing one average value per participant, one computes as many averages as there are participants while leaving one participant out each time. The rationale for this is that if the latency is consistent, then the average value will not change substantially depending on which participant is left out (Miller et al., 1998: 101). Thus, we computed 24 jackknife estimates of the auditory N100M latency for each condition, where the first estimate was an average N100M peak latency value in the time window of 60-100 ms post auditory onset of all participants except participant 1, the second was an average of all participants except participant 2, and so forth. To test whether latency differences were significantly different from zero, we first computed the (estimated) standard error of the difference (SED) following Miller et al. (1998). Next, we did paired samples t-tests on peak latency differences (i.e. the estimated average latency difference divided by the SED) for weak/strong visemes in isolation, as well as for weak/strong visemes in weak/strong contexts and the viseme/context interaction.

2.3.4 Analysis of early event-related activity

In addition to early latency effects, we also tested for amplitude differences in early auditory event-related activity caused by the manipulation of context and/or viseme. To this end, we tested ERF data averaged over the left temporal sensors of interest with nonparametric cluster-based permutation t-tests for paired samples (Maris & Oostenveld, 2007). Cluster-based permutation tests (henceforth 'cluster tests') are well suited to test for differences between conditions in multi-dimensional data like MEG measurements, where effects are typically clustered across time, space and/or frequency. Instead of testing each time point/sensor/frequency bin separately, such a cluster test works by finding clusters of significant differences between conditions and thus also controls the false alarm rate. In other words, it does not just detect significant differences between e.g. two amplitude values at the same time point, but also whether several subsequent different amplitude values are next to each other in the data matrix. After searching for clusters in the observed data, the test calculates a pre-defined number of random permutations of the observed data. The p-value for any cluster found in the observed data represents the probability under the null hypothesis of observing a cluster-level statistic (the sum of t-values for a cluster) that is more extreme, i.e. smaller or larger, in the randomly permuted data than in the observed data. In our analysis of early amplitude differences, we tested for clusters of amplitude differences across samples of a 300 ms long time-window starting from 100 ms before auditory onset (i.e. 0.1-0.4 s post target word video onset) with 5000 permutations per test. We first compared weak/strong visemes in the word-only condition, and then made the same comparison for words in sentence context. We also contrasted weak/strong contexts and tested the viseme/context interaction.

2.3.5 Analysis of late event-related activity

To test whether our sentence context manipulation produced the semantic N400-effect (Kutas & Hillyard, 1980), we compared ERFs to the target words presented in weak and strong context by means of a cluster test (see 2.3.4 above for details). Since the N400 typically starts at around 200-300 ms post stimulus onset (and peaks at around 400 ms with a negative polarity in EEG measurements - hence the name), we ran a cluster test with 5000 permutations to search for clusters of significant differences between conditions across all sensors and time points from 200 ms post auditory onset onwards (i.e. 0.4-1 s after target word video onset). We also did cluster tests with the same settings comparing strong/weak visemes and on the viseme/context interaction.

3. Results

3.1 Task accuracy

The behavioral results showed that participants paid attention to the videos on the screen. When asked to indicate whether the written word on the screen after the target word video matched or mismatched the word they just heard, participants performed very well. All 24 participants had an accuracy of between 72% and 100% (M = 95%) on trials where they responded within the time-limit. In sum, all participants performed the task well above chance level, suggesting that participants were paying attention to the stimuli during the experiment and that they were able to accurately recognize the target words.

3.2 Weak sentence context increases the N400m amplitude

As a basis for interpreting possible N100a latency effects we tested whether words embedded in strongly constraining sentences produced an attenuated N400m component compared to words embedded in weakly constraining sentences. Thus we checked for clusters of amplitude differences between strong and weak context from 400 ms (i.e. 200 ms post auditory onset) onwards. We found a reliable amplitude difference ($p < .001$) between weakly and strongly constraining sentences, such that words in weakly constraining sentence contexts elicited a larger neuromagnetic response than words in strongly constraining sentences contexts. The effect, which is illustrated as the shaded area between the two ERFs in Figure 3A, emerged at around 0.4 s and persisted until around 0.8 s. The difference topography in right in Figure 3B suggests that the effect of contextual constraint was present in both hemispheres. Although we did not test lateralization or spatial location statistically, the topography indicates that the effect was strongly left-lateralized and that it originated from enhanced activity in a fronto-temporal network. The time course and spatial location of this amplitude effect of sentence context are thus typical for the N400m (Maess, Herrmann, Hahne, Nakamura, & Friederici, 2006). We found no evidence that the N400 effect depended on viseme condition ($p > .2$). In sum, we found strong evidence for the classical N400 effect in our data, which shows that our sentence context manipulation was successful.

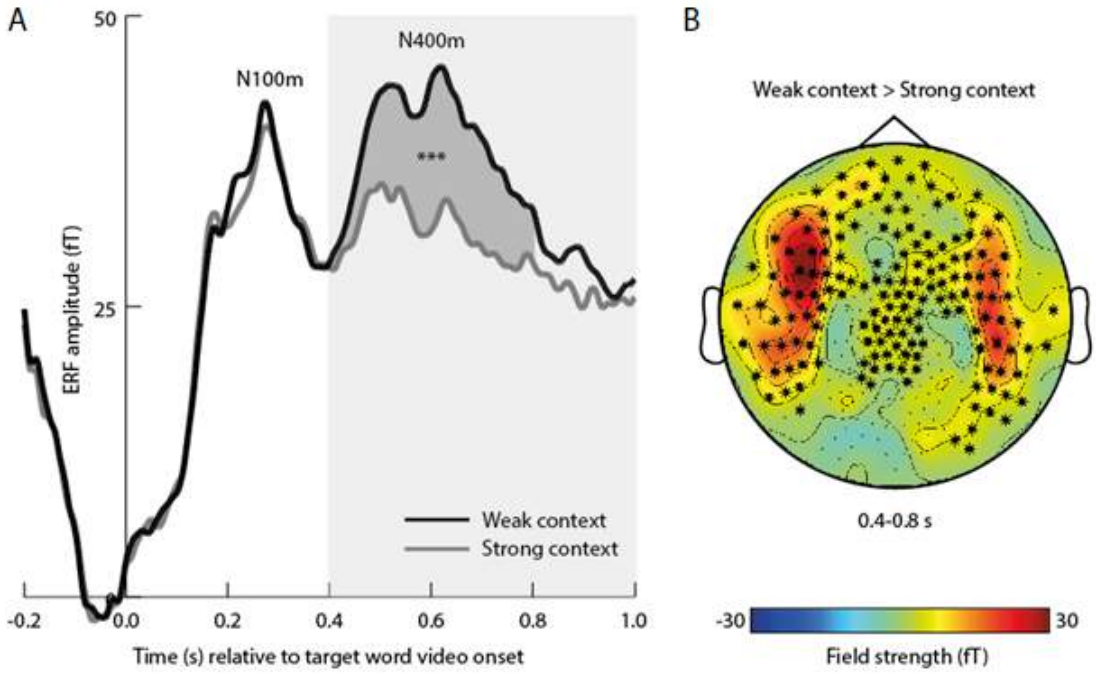

Fig. 3 N400-effect of sentence context. **A.** Event-related fields (ERFs) averaged over significant sensors found in the cluster test for the target words presented in weak and strong context respectively. (Note that the two lines represent brain activity elicited by identical audiovisual word tokens.) The light grey area shows the time window used in the statistical analysis, and the dark grey area between the lines shows the time window of the significant cluster. **B.** Topography showing the difference in activity for weak > strong context in the time window of the significant cluster, with the significant sensors highlighted. (*** = $p < .001$)

3.3 Strong context removes the auditory N100m latency shortening

We first assessed whether the peak latency of the auditory N100m differed depending on whether the word began with a strong or weak viseme. To this end, we compared auditory N100m peak latency differences for strong versus weak viseme in two situations: without context (when words were presented in isolation) and with context (when words ended a sentence). As can be seen in Figure 4A, a strong viseme led to a substantially earlier N100m peak than the weak viseme in the word-only condition (mean difference = 18 ms, $t(23) = 2.25$, $p < .05$, SED = 8 ms). Furthermore, Figure 4B shows that sentence context influenced the effect of viseme on auditory processing. In a weak context, the strong viseme also led to a substantially shorter N100m latency than the weak viseme (mean difference = 21 ms, $t(23) = 2.15$, $p = .02$, SED = 10 ms). However, we did not find such a latency difference for words presented in a strong sentence context (mean difference = 6 ms, $t(23) = .82$, $p = .36$, SED = 7 ms). This led to an interaction between context and viseme that trended towards significance ($t(23) = 1.50$, $p = .07$, SED = 10 ms).

3.4 Weak visemes lead to more auditory activity in a strong context

To test whether there were any amplitude differences from 100 to 400 ms after the onset of the target word video, we started by comparing the amplitude of the activity evoked by words with weak and strong visemes presented without a sentence context. This comparison did not result in any evidence to support amplitude differences for strong or weak visemes ($p > .2$). When we did the same comparison (weak vs. strong viseme) for words presented in a sentence context, however, we detected a reliable amplitude increase caused by weak viseme (278-326 ms, $p < .01$). The time window of this effect corresponded well with that of the auditory N100m. We further compared words presented in weak and strong sentence context, but did not find any clusters of amplitude differences ($p = 1$). However, when we tested for an interaction between context and viseme we found a trending cluster of amplitude differences (331-360 ms, $p < .06$). As can be observed in Figure 5A, post hoc comparisons showed that the interaction was caused by an amplitude increase for weak viseme words compared to strong viseme words in sentences where context was strong (271-352 ms, $p < .004$).

4. Discussion

In this study, we investigated how audiovisual words are processed depending on how well they can be predicted based on sentence context and viseme. Previous work has shown that strong visemes facilitate auditory speech processing. Our aim was to test whether having a strong prediction about a word based on sentence context would modulate this effect. We found a robust N400-effect of sentence context, showing that we successfully induced predictions about upcoming words in our participants when the sentence was constraining. Furthermore, we replicated the finding that strong visemes shorten the auditory N100m peak latency, and extended it by showing that this effect was only present either in the absence of semantic context or when context was weakly constraining. Finally, we found that weak visemes caused increased auditory event-related activity compared to strong visemes, and that this effect was enhanced when sentence context was strongly constraining.

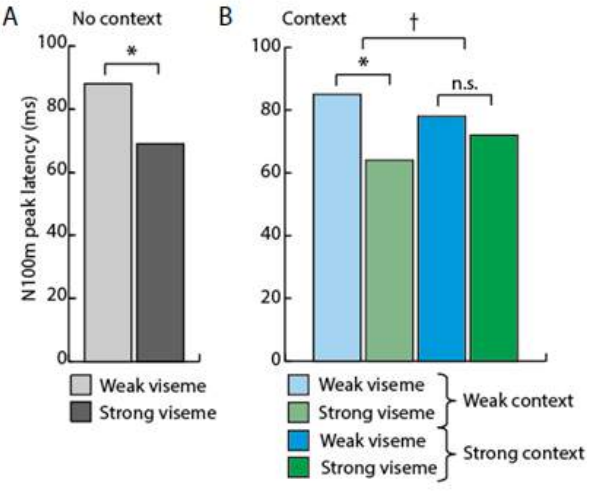

Fig. 4 Auditory N100m peak latency effects. **A.** Auditory N100m peak latencies relative to auditory onset for target words when there was no sentence context. The latency was reliably longer for words with a weak viseme than words with a strong viseme. **B.** Auditory N100m peak latencies relative to auditory onset for target words in weak and strong sentence context. The latency was reliably longer for words with a weak viseme than words with a strong viseme only when the sentence was not predictive of the target word. (* = $p < .05$, † = $p < .08$, n.s. = not significant/ $p > .1$)

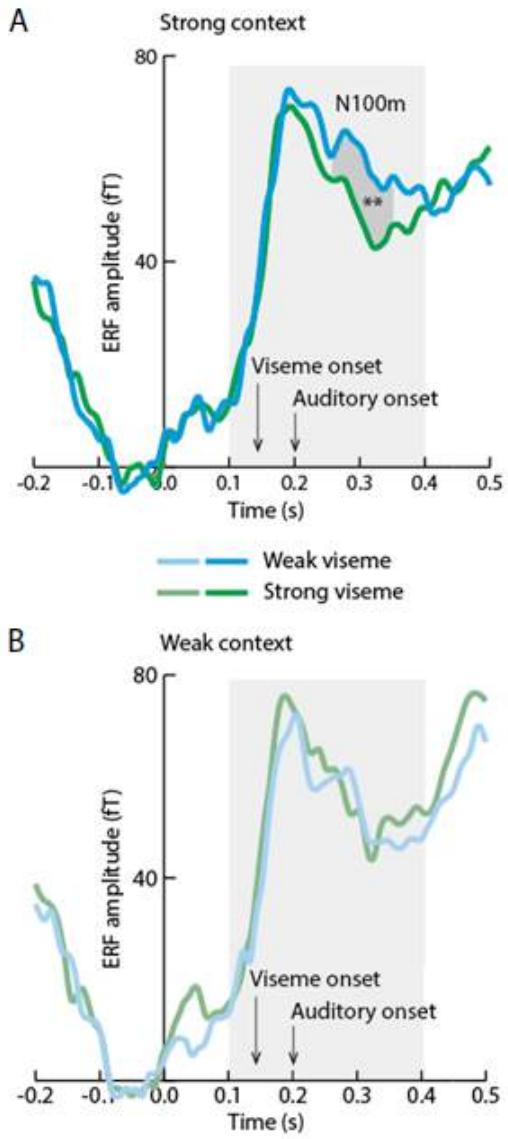

Fig. 5 Early auditory amplitude effects. Both panels show event-related fields (ERFs) averaged over the left temporal sensors of interest for strong (**A**) and weak (**B**) sentence context separately. The light grey area shows the time window used in the statistical analysis, and the dark grey area between the lines in the upper panel shows the time window of the significant cluster of amplitude differences found in the post hoc test. (** = $p < .01$)

4.1 Strong semantic context reduces the early processing benefit of a strong viseme

When the sentence context was predictive of the target word, we found an attenuated N400m (Fig. 5A) (Kutas & Hillyard, 1980). Based on this result, as well as results from previous studies using similar sentence context manipulations (e.g. DeLong et al., 2005), we conclude that our participants were anticipating word forms in the strong context condition. Furthermore, we replicated earlier findings of an auditory N100m latency shortening caused by strong visemes in isolated words (Fig. 4A) (Arnal et al., 2009; Brunellière et al., 2013; van Wassenhove et al., 2005). Interestingly, sentence context modulated the auditory N100m latency shortening effect caused by a strong viseme, such that this effect persisted in a weak context but was absent in a strong context (Fig. 4B). The trending interaction between sentence and viseme constraint suggests that the auditory N100m latency difference depended on the degree of contextual constraint. The current study is, to the best of our knowledge, the first to show such a result.

Firstly, our results indicate that the previously found auditory N100m latency shortening caused by strong visemes generalizes to more ecologically valid sentence comprehension, such that strong visemes speed up auditory processing when the sentence context is weak. Secondly, given that we do not find a reliable latency shortening in the strong context condition, we argue that our results lend support to the hierarchical hypothesis, where higher-level predictions derived from sentence context can influence lower-level predictions derived from e.g. visemes. This indicates that having a weak prediction about the sentence-final word causes us to rely more on low-level cues such as visemes, whereas having a strong prediction about the sentence-final removes the need to rely on such low-level cues as a large part of the upcoming sensory input has already been predicted.

4.2 Expectation suppression and enhanced auditory attention might explain early amplitude effects

In our analysis of early auditory amplitude effects, we found an increased auditory response to weak visemes, and that a strong context enhanced this effect (Fig. 5A). These amplitude differences were only present for words presented in a sentence context. To the best of our knowledge, only one other study has found an increased N100m to highly predicted words compared to less predicted words in a sentence context (Sivonen, Maess, Lattner, & Friederici, 2006). In this study, the authors measured ERPs to word-initial sounds that had be replaced by coughs, a manipulation which often results in an intact percept of the word ("phoneme restoration") when the word can be predicted by the context. However, as in the current study, the enhanced N100m to predicted words was an unexpected result, and was not replicated by a study with similar experimental manipulations (Groppe et al., 2010). However, this study used tones rather than coughs

to replace the word-initial speech sounds, and the difference between strongly and weakly predictive contexts was smaller than in the study by Sivonen et al. Thus, it remains difficult to explain why the effect emerges, although one likely possibility is that it has to do with the strength of the word prediction. We speculate that the enhanced N100m to predicted strong viseme words can be explained by two underlying mechanisms: firstly, we suggest that strong visemes cause an amplitude decrease due to a stronger auditory prediction, and secondly, we suggest that a strong sentence context causes an amplitude increase due to enhanced auditory attention.

Regarding the first mechanism, expecting an upcoming auditory stimulus attenuates the auditory N100m (Todorovic & de Lange, 2012; Todorovic et al., 2011), a phenomenon generally known as 'expectation suppression'. Auditory predictions caused by strong visemes also seem to cause this effect (Arnal et al., 2009). The researchers investigated N100m amplitude reduction effects with MEG in a larger variety of visemes than in our study, and did not find a systematic correlation between visual-only recognition rates of the visemes and N100m amplitude. Nevertheless, their results (Fig. 3C) show an amplitude reduction for /p/ (strong viseme) compared to /g/ (weak viseme). Since these visemes are similar to the ones we used in the present study, we propose that the amplitude decrease found for strong visemes by Arnal and colleagues might correspond to our results. However, it remains difficult to explain why we did not find such a viseme expectation suppression effect for words presented in isolation.

Regarding the second mechanism, enhanced attention is known to increase the amplitude of early auditory ERP/Fs like the auditory N100m (Picton & Hillyard, 1974). We speculate that our participants were more efficient in attending to the task-relevant final word in the sentence when the context was strongly constraining. Indeed, some participants reported that they were sometimes consciously expecting certain words. Results from an EEG study where people read paragraphs that were either embedded in a strong or weak context showed that ERPs to words in strong compared to weak context paragraphs were characterized by both an attenuated N400 and, interestingly, an early amplitude increase 100 ms after word onset (St George, Mannes, & Hoffinan, 1994). This early amplitude increase was interpreted to reflect enhanced visual attention. It is thus possible that the early amplitude increase we find with a strong context reflects increased auditory attention. All in all, we propose that a combination of expectation suppression and increased auditory attention might explain the pattern of our early auditory amplitude effects, although this remains a highly tentative explanation.

4.3 Visemes and sentences: different levels in the prediction hierarchy

Our results are in line with the hierarchical predictive coding framework (Friston, 2005; Friston & Kiebel, 2009; Wacongne et al., 2011), where predictions made on higher levels influence predictions made on a lower level. This theory of how the brain functions has also been proposed to govern language comprehension and production (Pickering & Garrod, 2013; see also Altmann & Mirkovic, 2009). We found that the high-level sentence predictions removed the effect of low-level viseme predictions on the auditory N100m latency, which suggests that a) sentence context is higher up than visemes in the prediction hierarchy and b) that such higher-level predictions interact with lower-level predictions.

Our study has somewhat limited ecological validity due to repetition of words and semantic relatedness between sentences. Furthermore, we directed participants' attention towards sentence-final words by inserting a pause before them and having a task on these words. Nevertheless, given that our data provide evidence for early effects of context, these experimental manipulations may have enhanced our sensitivity to effects of predictive speech processing earlier than the N400 (for discussion, see Penolazzi, Hauk, & Pulvermüller, 2007; Pulvermüller, Shtyrov, & Hauk, 2009). A recent EEG study also investigated the interaction between context and viseme in audiovisual speech processing (Brunellière et al., 2013). However, their results differ from ours in that they did not find any latency effects when sentence context was manipulated. Although there can be many reasons for why our results are different, one reason might be that they had more variation in the stimulus materials.

Both the early auditory latency and amplitude effects we found seem to reflect predictive processing mechanisms. However, it seems plausible that they originate from different neural processes. The latency shortening effect of viseme has been argued to reflect faster auditory processing (van Wassenhove et al., 2005), and is insensitive to whether or not the speech sound matches the viseme (Arnal et al., 2009). Phase-resetting of activity in auditory cortex due to input from visual motion areas has been

suggested as the underlying mechanism of both visual facilitation of auditory speech processing in general and of the latency shortening effect (Arnal et al., 2009; Schroeder, Lakatos, Kajikawa, Partan, & Puce, 2008). Furthermore, it has recently been shown that the auditory cortices track the speech rhythm irrespective of comprehension, but that comprehensible speech enhances this tracking in the left auditory cortex (Peelle, Gross, & Davis, 2013). Given that a helpful context improves comprehension (e.g. St George et al., 1994), we suggest that contextual influence on the latency effect might be due to an improved left-lateralized phase locking of cortical oscillations to the speech rhythm. As we inserted short gaps before the critical words, our study was not designed to look at phase locking of cortical oscillations. Therefore, future studies using connected audiovisual speech might shed more light on the relationship between the latency shortening effect and oscillatory activity.

In contrast to the latency effect, the early amplitude effect is sensitive to mismatch between visual and auditory speech (Arnal et al., 2009). Findings from both MEG and functional magnetic resonance imaging (fMRI) by Arnal and colleagues indicate that this effect reflects a feedback signal from the superior temporal sulcus that encodes the degree of overlap between the prediction based on the viseme and the subsequent acoustic signal. In other words, when visemes are less predictive towards specific speech sounds this causes an enhanced neural response. This explanation suggests that the amplitude of the auditory N100m does not reflect the process of pre-activating certain words or sounds in itself, and fits well with the notion of expectation suppression. Rather, it may reflect the first step in fine-tuning audiovisual processing to improve future predictions. Additionally, it seems plausible that an attended sentence should result in a higher degree of predictive fine-tuning. Thus, the amplitude effect could be linked to the latency effect in that the enhanced N100m response would cause a slightly better predictive tracking of the speech signal during a future face-to-face conversation.

It has recently been argued that visual speech only serves as input to perceptual processing, and not to semantic/higher-level processing (Samuel & Lieblich, 2014). On the one hand, we show that the brain does not use visemes to make predictions when the semantic context is strong, which points to a non-semantic role of visual speech in these situations. On the other hand, the brain does seem to use visemes to make predictions when context is weak or non-existent. Thus, on the basis of our results it seems that we cannot completely dismiss a role for visual speech in speech processing at the level of semantic processing. Since we found an interaction between our semantic and visual manipulations, it seems that the roles of visual speech and semantic context might not be separate, but rather that the role of visual speech changes depending on the strength of the semantic (higher-level) prediction.

Since our goal when talking to someone is ultimately to understand what they are saying, as opposed to determine which speech sounds they are uttering, it seems like an efficient strategy not to let redundant sensory input such as visual speech "interfere" with the process unless it really does improve the brains' hypothesis on what is coming up next. Thus, it might be that when visual speech is not strictly necessary for successful comprehension - such as in our experiment - the brain simply treats the visual speech as an additional perceptual source of information that it can draw on if needed. We suggest that visual speech as has a more direct influence on comprehension in situations where there is a relative high degree of uncertainty about upcoming words, or when the acoustic signal is noisy, unreliable or absent (e.g. MacLeod & Summerfield, 1987). Such situations include talking to someone about an unfamiliar topic or having a conversation in a noisy bar. Future work might further investigate the functional role of visual speech and how it depends on how easily we are able to comprehend what is being said – does visual speech only play a significant role in comprehension of clear speech when higher-level cues are poor or absent? Finally, future studies should also investigate at whether hierarchical predictive coding is a valid framework for the interplay between different types of predictive cues in language processing.

5. Conclusion

The current study sheds light on the relative influence of visual speech cues on our perception of spoken language when additional contextual information is available. Overall, our results show that high-level semantic cues and low-level visual speech cues influence our speech comprehension in a hierarchical manner. A strong high-level prediction about an upcoming word renders low-level predictions concerning the perceptual features of that word largely redundant, which is why we do not see any effect of viseme in this situation. When the high-level prediction is weak and unspecific, on the other hand, the visual cue receives more weight

because it can improve prediction of an upcoming word. Furthermore, we show that low-level predictions like those from visual cues and high-level predictions from semantic context interact at an early stage in audiovisual speech perception. We conclude that predictions derived from semantic context and visual speech are hierarchically organized, which indicates that hierarchical predictive coding provides a useful framework for understanding predictive processes in speech perception.

6. References

Altmann, G. T. M., & Kamide, Y. (1999). Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition*, *73*, 247–264.

Altmann, G. T. M., & Mirkovic, J. (2009). Incrementality and prediction in human sentence processing. *Cognitive Science*, *33*(4), 583-609.

Arnal, L. H., Morillon, B., Kell, C. A., & Giraud, A. L. (2009). Dual neural routing of visual facilitation in speech processing. *The Journal of Neuroscience*, *29*(43), 13445-13453.

Bastiaansen, M., & Knösche, T. R. (2000). Tangential derivative mapping of axial MEG applied to event-related desynchronization research. *Clinical Neurophysiology*, *111*(7), 1300-1305.

Bell, A. J., & Sejnowski, T. J. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, *7*(6), 1129-1159.

Boersma, Paul (2001). Praat, a system for doing phonetics by computer. *Glot International*, *5*(9/10), 341-345.

Brunellière, A., Sánchez-García, C., Ikumi, N., & Soto-Faraco, S. (2013). Visual information constrains early and late stages of spoken-word recognition in sentence context. *International Journal of Psychophysiology*, *89*(1), 136-147.

Chandrasekaran, C., Trubanova, A., Stillittano, S., Caplier, A., & Ghanzanfar, A. A. (2009). The Natural Statistics of Audiovisual Speech. *PLoS Computational Biology*, *5*(7): e1000436.

Davis, M. H., & Johnsrude, I. S. (2007). Hearing speech sounds: top-down influences on the interface between audition and speech perception. *Hearing Research*, *229*(1), 132-147.

DeLong, K. A., Urbach, T. P., & Kutas, M. (2005). Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature Neuroscience*, *8*(8), 1117-1121.

Federmeier, K. D. (2007). Thinking ahead: The role and roots of prediction in language comprehension. *Psychophysiology*, *44*(4), 491-505.

Friston, K. (2005). A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological Sciences*, *360*(1456), 815-836.

Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1521), 1211-1221.

Groppe, D. M., Choi, M., Huang, T., Schilz, J., Topkins, B., Urbach, T. P., & Kutas, M. (2010). The phonemic restoration effect reveals pre-N400 effect of supportive sentence context in speech perception. *Brain Research*, *1361*, 54-66.

Hämäläinen, M., Hari, R., Ilmoniemi, R. J., Knuutila, J., & Lounasmaa, O. V. (1993). Magnetoencephalography-theory, instrumentation, and applications to noninvasive studies of the working human brain. *Reviews of Modern Physics*, *65*(2), 413-497.

Hari, R., Pelizzone, M., Mäkelä, J. P., Hällström, J., Leinonen, L., & Lounasmaa, O. V. (1987). Neuromagnetic responses of the human auditory cortex to on-and offsets of noise bursts. International *Journal of Audiology*, *26*(1), 31-43.

Jesse, A., & Massaro, D. W. (2010). The temporal distribution of information in audiovisual spoken-word identification. *Attention, Perception, & Psychophysics*, *72*(1), 209-225.

Jung, T. P., Makeig, S., Westerfield, M., Townsend, J., Courchesne, E., & Sejnowski, T. J. (2000). Removal of eye activity artifacts from visual event-related potentials in normal and clinical subjects. *Clinical Neurophysiology*, *111*(10), 1745-1758.

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

Kutas, M., & Federmeier, K. D. (2011). Thirty years and counting: finding meaning in the N400 component of the event-related brain potential (ERP). *Annual Review of Psychology*, *62*, 621-647.

Kutas, M., & Hillyard, S. A. (1980). Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science*, *207*(4427), 203-205.

Lau, E. F., Phillips, C., & Poeppel, D. (2008). A cortical network for semantics:(de) constructing the N400. *Nature Reviews Neuroscience*, *9*(12), 920-933.

MacLeod, A., & Summerfield, Q. (1987). Quantifying the contribution of vision to speech perception in noise. *British Journal of Audiology*, *21*(2), 131-141.

Maess, B., Herrmann, C. S., Hahne, A., Nakamura, A., & Friederici, A. D. (2006). Localizing the distributed language network responsible for the N400 measured by MEG during auditory sentence processing. *Brain Research*, *1096*(1), 163-172.

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

Miller, J., Patterson, T., & Ulrich, R. (1998). Jackknife-based method for measuring LRP onset latency differences. *Psychophysiology*, *35*(1), 99-115.

Näätänen, R., & Picton, T. (1987). The N1 wave of the human electric and magnetic response to sound: a review and an analysis of the component structure. *Psychophysiology*, *24*(4), 375-425.

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.

Peelle, J. E., Gross, J., & Davis, M. H. (2013). Phase-locked responses to speech in human auditory cortex are enhanced during comprehension. *Cerebral Cortex*, *23*(6), 1378-1387.

Penolazzi, B., Hauk, O., & Pulvermüller, F. (2007). Early semantic context integration and lexical access as revealed by event-related brain potentials. *Biological Psychology*, *74*(3), 374-388.

Pickering, M. J., & Garrod, S. (2013). An integrated theory of language production and comprehension. *Behavioral and Brain Sciences*, *36*(04), 329-347.

Picton, T. W., & Hillyard, S. A. (1974). Human auditory evoked potentials. II: Effects of attention. *Electroencephalography and Clinical Neurophysiology*, *36*, 191-200.

Pulvermüller, F., Shtyrov, Y., & Hauk, O. (2009). Understanding in an instant: neurophysiological evidence for mechanistic language circuits in the brain. *Brain and Language*, *110*(2), 81-94.

Samuel, A. G., & Lieblich, J. (2014). Visual speech acts differently than lexical context in supporting speech perception. *Journal of Experimental Psychology: Human Perception and Performance*, *40*(4), 1479-1490.

Schroeder, C. E., Lakatos, P., Kajikawa, Y., Partan, S., & Puce, A. (2008). Neuronal oscillations and visual amplification of speech. *Trends in Cognitive Sciences*, *12*(3), 106-113.

Schwartz, J., & Savariaux, C. (2014) No, there is no 150 ms lead of visual speech on auditory speech, but a range of audiovisual asynchronies varying from small audio lead to large audio lag. *PLoS Computational Biology*, *10*(7), e1003743.

Sivonen, P., Maess, B., Lattner, S., & Friederici, A. D. (2006). Phonemic restoration in a sentence context: evidence from early and late ERP effects. *Brain Research*, *1121*(1), 177-189.

Smits, R., Warner, N., McQueen, J. M., & Cutler, A. (2003). Unfolding of phonetic information over time: A database of Dutch diphone perception. *The Journal of the Acoustical Society of America*, *113*(1), 563-574.

St George, M., Mannes, S., & Hoffinan, J. E. (1994). Global semantic expectancy and language comprehension. *Journal of Cognitive Neuroscience*, *6*(1), 70-83.

Todorovic, A., van Ede, F., Maris, E., & de Lange, F. P. (2011). Prior expectation mediates neural adaptation to repeated sounds in the auditory cortex: an MEG study. *The Journal of Neuroscience*, *31*(25), 9118-9123.

Todorovic, A., & de Lange, F. P. (2012). Repetition suppression and expectation suppression are dissociable in time in early auditory evoked fields. *The Journal of Neuroscience*, *32*(39), 13389-13395.

van Wassenhove, V., Grant, K. W., & Poeppel, D. (2005). Visual speech speeds up the neural processing of auditory speech. *Proceedings of the National Academy of Sciences of the United States of America*, *102*, 1181–1186.

Wacongne, C., Labyt, E., van Wassenhove, V., Bekinschtein, T., Naccache, L., & Dehaene, S. (2011). Evidence for a hierarchy of predictions and prediction errors in human cortex. *Proceedings of the National Academy of Sciences*, *108*(51), 20754-20759.

Windmann, S. (2004). Effects of sentence context and expectation on the McGurk illusion. *Journal of Memory and Language*, *50*(2), 212-230.

Identifying Traces of Consciousness in the Process of Intending to Act

Ceci Verbaarschot[1]
Supervisors: Pim Haselager[1], Jason Farquhar[1]

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

In 2008, Matsuhashi and Hallett created a novel version of the experimental design of Libet, Gleason, Wright and Pearl (1983) which made it possible to get real-time estimations of a participant's awareness of consciously willing to act. Using this method, the onset of the intention to act was measured up to 1.42 s prior to movement onset, whereas Libet et al. (1983) found it only about 0.2 s prior to movement. If one takes intentions to be discrete mental states these results seem to be at odds with each other, yet they fit very well within a framework in which intentions are regarded as processes developing over time. While the later stages in this process of intending are available for self-initiated report (similar to the reported intention timings by Libet et al., 1983), early stages appear to be reachable and reportable through external probing only (as used by Matsuhashi & Hallett, 2008). However, to the best of our knowledge, no one has conducted a within-subject comparison between the Libet- and Matsuhashi-task in order to investigate whether the measured onsets of intending indeed differ significantly between the two. In this thesis, we will propose a novel conceptual framework describing intending as a process consisting of multiple stages developing over time. With this framework in mind, we conducted a new experiment, incorporating adapted versions of both the Matsuhashi and Hallett experiment as well as that of Libet and colleagues into a within-subject design. The results indicate that the onsets of intending as measured using external probing indeed occur significantly earlier in time (and quite close to the onset of the neural preparation for action) compared to the onsets of intending as measured using self-initiated reports, thus providing credence to the interpretation of intending to act as a process developing over time.

Keywords: action, consciousness, EEG, intention, LRP, RP

Corresponding author: Ceci Verbaarschot; **E-mail:** c.s.verbaarschot@student.ru.nl

1. Introduction

Through the repetitive experience of willing an act and performing the intended act shortly thereafter, we intuitively assume that it is our conscious intention to act that causes the act to happen. However, we don't actually 'see' our intention causing the act, but can only infer this from the constant experience of the intention being predictive of and prior to the subsequent act (Wegner, 2002). Libet, Gleason, Wright and Pearl (1983) were among the first to use empirical methods to investigate the nature and temporal order of the processes of mind that produce the experience of will and those that produce the action itself. They found that the neural preparatory processes for movement (i.e., Readiness Potential (RP)) preceded the onset of the reported intention to act by about 350 ms and the act itself (a self-paced flexion of the right hand) by 550 ms on average. According to Libet et al. (1983), these results suggest that the brain already starts preparing the upcoming act before it is intended consciously by a person. If this is true, it means our intuition on the causation of our actions is false, for this would mean that our acts are initiated unconsciously rather than consciously.

The results found by Libet et al. (1983) have been replicated several times (Haggard & Eimer, 1999; Keller & Heckhausen, 1990; Trevena & Miller, 2002; Verbaarschot, Farquhar, & Haselager, 2014) and still play an important role in the debate on consciousness and free will (Dennett, 1991, 2003; Haggard, 2008; Hallett, 2007; Mele, 2010; Pockett, 2004; Wegner, 2002). In these experiments, the onset of an intention to act is typically measured through a self-initiated report by the participant. Often, the participant is instructed to watch some variant of a clock and report the spatial configuration of that clock at the time they first feel their intention to act arise. From this it seems that many researchers, including Libet et al. (1983), regard intentions as being discrete mental states. In their view, an intention to act consists of a simple mental state (having the thought of 'wanting to act... now!') that is encoded in brain activity at a single point in time. However, following Dennett (1991) and Uithol (2012), regarding intentions as processes seems much more plausible.

Similar to the neural processes responsible for initiating and performing an act, intending an act can also be viewed as a process developing over time. This process of intending might very well have its onset already at the start of a Libet-type experiment, as a participant agrees to obey the provided experimental instructions. Some phases in this process will leave traces in consciousness and are available for subsequent self-initiated report, while others remain silent (Dennett, 1991). Relying only on self-initiated reports on the onset of a concrete mental state (i.e., the intention to act) might leave various phases of the process of intending unexplored.

In 2008, Matsuhashi and Hallett created a novel experimental design that made it possible to get real-time estimations of a participant's awareness of consciously willing to act. Instead of relying on self-initiated post-action reports, they used auditory probes to initiate an intention report prior to action performance. Using this method, the onset of the intention to act was measured up to 1.42 s prior to movement onset, whereas Libet et al. (1983) found it only about 0.2 s prior to movement. These results suggest that before a participant is able to provide a self-initiated report of intending their act, some form of action-related awareness is already present and can be revealed by using external probes. However, to the best of our knowledge, no one has conducted a within-subject comparison between the Libet- and Matsuhashi-task in order to investigate whether the measured onsets of intending indeed differ significantly between the two. If this is found to be the case, it strongly suggests that intending to act can be better understood when viewed as a process, instead of a mental state. While the later stages in this process of intending are available for self-initiated report, early stages appear to be reachable and reportable through external probes.

Departing from the assumption that intending to act is a process rather than a discrete mental state, we have created a novel conceptual framework describing this process of intending over time and linking it to the different neural preparatory processes identified in voluntary action performance in previous research. With this framework in mind, we developed a new experiment that incorporates an adapted version of both the Matsuhashi and Hallett experiment as well as that of Libet et al. into a within-subject design. The results indicate that the measured onsets of intending to act indeed differ significantly between the Libet- and Matsuhashi-task, thus providing a first exploration of the 'silent' phases in the process of intending close to action performance. Similar to Matsuhashi and Hallett (2008), we found that the probed onset of intending lies much closer to the onset of the neural preparation for action (as measured by the RP) than reported in previous research (Haggard & Eimer, 1999; Keller & Heckhausen, 1990; Trevena & Miller,

2002; Verbaarschot et al., 2014). These results suggest that the process of intending develops alongside the process of acting, leaving traces in consciousness at certain points along the road and ultimately reaching awareness in the form of the conscious thought of 'wanting to act... now!'.

The next section of this paper will provide an overview of relevant previous research on the topic of voluntary action performance. In Section 3, we will sketch the proposed novel conceptual framework describing the different phases in the process of intending and relating it to the neuro-scientific results described in Section 2. The outline of this framework is followed by a detailed description of our experimental design in Section 4 and results in Section 5. In Section 6, we will discuss a number of complexities that we encountered during our experiment. Finally, we will give our conclusions in Section 7 and provide suggestions for future research.

1.1 Background

Although it intuitively feels as though our conscious intentions exert a causal force which leads to the subsequent performance of an intended act, we cannot be certain that this

is indeed the case. Libet et al. (1983) were among the first to use an empirical method to measure the onset of an intention to act and investigate the temporal order between the neural preparation for action, intending to act, and performing the intended act. They encouraged their participants to "let the urge to act appear on its own at any time without any preplanning or concentration on when to act" (Libet et al., 1983, p. 625) while they made spontaneous flexions with the fingers and/or wrist of their right hand. After each movement, the participants were instructed to recall the spatial configuration of a clock at the time they first felt their intention to act arise. After reporting this configuration, participants were asked to confirm whether their movement was made spontaneously or whether it was planned.

Libet et al. used electro-encephalography (EEG) to record their participant's brain activity during voluntary action performance and used the RP as an indicator of the neural preparation for movement. The RP can be divided into an early and a late phase (Kornhueber & Deecke, 1965; Shibasaki & Hallett, 2006). The early phase starts around 1.5-2 s before action and consists of a slowly rising negative potential that can be measured maximally over the central-medial scalp. The late phase starts around 400 ms before action and consists of a more steeply going negative potential that is lateralized over the primary motor area.

Interestingly, Libet et al. (1983) found that the RP preceded the reported intention to act by 350 ms, and the actual movement by 500 ms on average. This implied, according to Libet et al. (1983), that the neural processes required to initiate a movement preceded the onset of the intention to move. This strongly contradicts our intuition on the cause of our actions since it suggests that our intentions lag behind events and therefore cannot play a causal role in action initiation.

In 2008, Matsuhashi and Hallett created a novel experimental design that made it possible to get real-time estimations of a participant's awareness of consciously willing to act. Their participants made self-paced brisk right index finger extensions at intervals of 5-10 s while their brain activity was measured using EEG. At random intervals of 3-20 s, a tone was presented and participants were instructed to:

- ignore the tone if they were not thinking about the next movement and continue their self-paced acts, or
- veto their movement if they had started thinking about their next act and wait for another 5-10s interval before performing their next movement.

This design enabled Matsuhashi and Hallett (2008) to create a model of the RP in relation to the onset of the intention to act and being able to veto an act. The estimated onset of the intention to act was found between 1.42 and 0.69 s prior to movement, the early phase of the RP was found between 2.17 and 0.69 s prior to movement and the late phase was found between 0.57 and 0.20 s prior to movement. Furthermore, the point of no return (i.e., where a participant was no longer able to veto their act since the presented tone occurred very close to movement onset) was estimated between 0.13 and 0.07 s prior to movement onset. These results indicate a period without intention awareness until about 1.42 s prior to movement, after which there is an awareness of intending to act after being probed. A self-initiated report of consciously willing to act 'now' (being meta-aware) arises only up to 0.2 s prior to movement as reported by Libet et al. (1983).

The results of Matsuhashi and Hallett (2008) are not consistent with the traditional conceptual framework in which intentions are viewed as discrete mental states occurring at a single point in time. On the contrary, these results seem to suggest that intending to act is a process consisting of several phases and developing over time. While

participants are able to produce self-initiated reports of intending their act only quite close to action performance, external probes can be used to reveal the action-related awareness that is present earlier in time and closer to the start of the neural preparation for action (as identified by the RP).

In the next section, we will define what we understand 'the process of intending' to be and propose a new conceptual framework describing the process of intending as it develops over time alongside the neural preparatory processes for action.

1.2 The Process of intending

The majority of experimental researchers seem to use the folk psychological notion of 'an intention to act.' They take the intention to act to be reflecting a simple mental state whose onset can be measured empirically (by means of reporting the spatial configuration of a clock at the time the intention to act arises, for instance). This theoretical entity (i.e., the intention to act) represents the feeling of wanting to perform an action and can be captured at a single point in time (i.e., the reported intention onset). However, we do not think that this interpretation maps correctly onto the results found in previous research (Bai et al., 2011; Blankertz et al., 2006; Libet et al., 1983; Pfurtscheller & Aranibar, 1979; Schurger, Sitta, & Deheane, 2012; Soon, Brass, Heinze, & Haynes, 2008; Trevena & Miller, 2002; Trevena & Miller, 2010) since it cannot account for the complexity found in the processes of both acting and intending. It seems much more plausible that, similar to the neural preparatory processes for action which can be measured over a period of time prior to an act, intending to act is also a process developing over time.

Brass and Haggard (2008) have proposed a three-component model of the process of intending in voluntary action performance. The first component relates to deciding what action to perform, the second to deciding when to perform it, and the third to deciding whether or not to execute the act. The existence of these three components (the what, when and whether components) has received support from neuro-scientific research (Brass & Haggard, 2008). For instance, Hoffstaedter et al. (2012) used functional magnetic resonance imaging (fMRI) to associate two different neural networks with the what- and when-components of intending. The what-component employed the preSMA and dorsal premotor cortex bilaterally, whereas the when-component was employed by the superior SMA together with the insula and Area 44 bilaterally as well as bilateral anterior putamen, globus pallidus and left cerebellum subcortically. Moreover, a study by Brass and Haggard (2007) identified a specific region in the fronto-median cortex that was more strongly activated prior to intentionally vetoed acts (i.e., acts that were intended at first, but canceled at a later point in time) than intentionally executed acts. Together, these results support the distinction between the what-, when- and whether components involved in the process of intending since they suggest that these components are realized by distinct neural networks.

We will incorporate the what, when and whether model of Brass and Haggard into a new conceptual framework for the process intending in a Libet-type experiment and relate it to the neuro-scientific results on voluntary action preparation described in previous research (Bai et al., 2011; Blankertz et al., 2006; Libet et al., 1983; Pfurtscheller & Aranibar, 1979; Schurger et al., 2012; Soon et al., 2008; Trevena & Miller, 2002; Trevena & Miller, 2010). Note that this means we will only consider the processes of acting and intending preceding simple and spontaneous left- and right-hand actions. Note furthermore, that the global whether-component will not play a big role on a single trial level in this kind of experiments since a participant agrees to perform an act on each trial by obeying the provided experimental instructions at the start of the experiment. Making this agreement leads to implementing the whether-component of intending by forming a global intention to act and having an overall readiness for action. However, which action to perform and when to perform it has to be decided spontaneously on each single trial.

In between the start of a trial and the moment of deciding what action to perform and when to perform it, there is a period of waiting. At a certain moment in time, the when- and what-decision will be made and can leave a reportable trace in consciousness. At this point, the decision to act enters the participant's awareness which enables the subsequent report of consciously willing to act.

This sequence in the process of intending can be linked to multiple neuro-scientific results on action preparation found in previous research. Between the whether- and what-decision, urges to act might come and go. These urges might be represented by the random fluctuations in neural activity as found by Schurger et al. (2012). These fluctuations might come and go, until one of them crosses a certain threshold and leads to a decision on the what- and/or when-component in the process of intending. Furthermore, the early frontal lobe activity found by

Soon et al. (2008) and Bode et al. (2011) is suggested to be predictive of the subsequent act and might arise during the waiting phase, leading to the what-decision in the process of intending. The activity in the preSMA and SMA, as found by Libet et al. (1983) and Soon et al. (2008), might be more tightly linked to the initiation of the subsequent act, representing the when-decision in the process of intending.

Depicted in Figure 1 are the four proposed stages in the process of intending in a Libet-type experiment:

1. A very global phase that has its onset as soon as a participant decides to participate in the experiment.
2. A phase of doing nothing/waiting during which the urge to act comes and goes.
3. A decision phase in which the what- and when-components of the action are being decided when an urge to act crosses a certain threshold.
4. A phase in which traces of the process of intending might become available to subsequent self-initiated report.

The conceptual framework for the process of intending, as proposed in this section, seems to accommodate the various results from previous research concerning intentions and their relation to voluntary action performance. To underline the applicability of this framework, we have conducted an experiment that incorporates the Libet- and Matsuhashi-task in a within-subject design.

After describing the experimental design and results during the next two sections, we will explain why the results fit well within the new framework on intending and suggest how this framework can be incorporated in future research during the remaining sections.

2. Materials and methods

Whether the measured onsets of intending to act indeed differ significantly between the Libet- and Matsuhashi-task has not yet been investigated. In order to answer this question, we have incorporated the conceptual framework on the process of intending (as described in Section 3) into a novel experiment combining adapted versions of both the Matsuhashi- and Libet-task into a within-subject design. This experiment is a first step towards investigating the different phases in the process of intending in voluntary action performance and will focus on those phases of acting and intending that are close to action onset. Additional neuroimaging techniques are required to investigate the earlier phases of intending as the neural preparatory signals in the frontal cortex associated with the whether- and what-components of intending require a good spatial resolution.

We have made some significant changes to the original experimental designs of both the Libet- and Matsuhashi-task in order to fit them within the

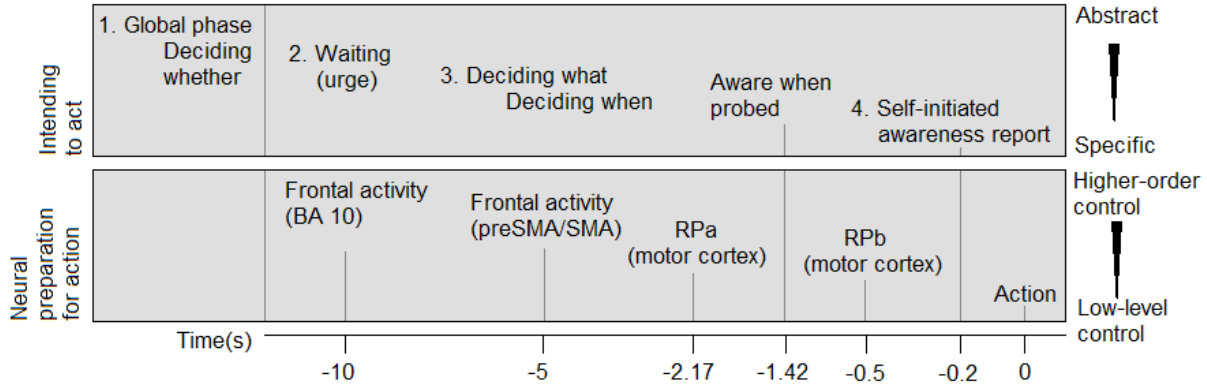

Fig. 1 The process of intending consists of four phases: (1) a very global phase that has its onset as soon as a participant decides to participate in the experiment (implementing the whether-component of intending), (2) a phase of doing nothing/waiting, during which the urge to act comes and goes, (3) a decision phase in which the what- and when-components of the action are being decided and (4) a phase in which traces of the process of intending might become available for subsequent self-initiated report. Phase (1) and (2) including the decision of what action to perform in phase (3) can be represented by frontal activity in BA 10 whereas the when-decision of phase (3) can be represented by activity in the preSMA/SMA and the RP in the motor cortex. The phases in the process of intending run from highly abstract (i.e., global agreement to participate (phase 1)) to quite specific (i.e., knowing which action to perform and when to perform it (phase 3)). Similarly, the neural preparatory processes for action run from activity in higher cognitive areas (BA 10) to lower cognitive areas (motor cortex). RPa corresponds to the onset of the RP as reported by Matsuhashi and Hallett (2008) and RPb corresponds to the onset of the RP as reported by Libet et al. (1983).

framework of intending as described in Section 3. Inspired by the experimental design presented by Grey Walter to the Ostler Society at the Oxford University in 1963 (Dennett, 1991, p. 167), we decided to instruct participants to watch slides on a computer screen. This everyday task creates a more realistic experimental setting than the one that was used in the original Matsuhashi-task, where participants were instructed to make spontaneous movements with their right hand every 5 to 10 s. In other words, participants were performing an action without any reason to act and without observing any effect of their act. By allowing participants to watch slides and press a button to go to a next slide whenever they want to, their act can be made for a reason (where the simplest one is being bored of the current slide) and has an apparent effect (presenting a new slide on the computer screen). We believe that intended actions involve some form of reasoning about the current situation and relevant background information (otherwise we might as well replace the participant by a random number generator).

Furthermore, by allowing participants to press a button with either their left or right thumb in the adapted versions of both the Matsuhashi- and Libet-task, we incorporated the what-component of intending along with the whether- and when-components that were already present in the original designs. Although this experiment will only investigate the subjective experience during the when-component of intending, the design allows for the investigation of the other two components in future research.

In contrast to the original Matsuhashi-task, participants were not instructed to perform their actions with a certain frequency but performed one action on each trial (where a trial is equal to the presentation length of one slide) in order to enhance the level of spontaneity of the act.

The details of the experimental design including the exact experimental procedures of both Libet- and Matsuhashi-task will be discussed in the following subsections.

2.1 Participants

A total of 12 healthy volunteers were tested during the experiment (7 females). The experiments were conducted with the written consent of each subject. The participants were between the age of 18 and 30 (with an average age of 23.25 years old). Ten of the participants were right-handed and two of them were left-handed. All participants had normal or corrected-to-normal vision.

2.2 Apparatus

A standard colour monitor was used to display the instructions and stimuli to the participant. In-ear headphones were used to play the auditory stimuli. Two small button boxes were connected to the computer running the experiment and were placed to the left and right of the participant.

2.3 Recording

The experiments were run in BrainStream (Severens, 2009). Brain activity during the experiment was recorded with 64 Ag/AgCl active electrodes. The electrodes were placed on the scalp according to the International 10/20 System (Klem, Luders, Jasper, & Elger, 1999). The offsets of the electrodes were kept under 25 μV and a Biosemi active 2 amplifier was used to record and amplify the brain signal. The brain signals were sampled at 2048 Hz. Eye blinks and movements were measured using 4 electrooculogram (EOG) electrodes. Bipolar electrodes were attached just above and below the left eye to record blinks and vertical eye movements. Two other electrodes were attached to the outer sides of both the left and right eye to record horizontal eye movements. The recorded EOG activity was used to filter the eye blinks and movements from the recorded brain activity during off-line analysis.

The muscle activity of both left and right arm were measured using electromyogram (EMG). Two bipolar electrodes were placed on the flexor pollicis longus muscle around the center of both right and left forearm to record the activity of the arm muscles moving both right and left thumbs during a button press. Two other electrodes were placed on the wrist bone of both the right and left wrist. The EMG was recorded for each arm as the difference between the electrode on the wrist bone and the forearm. The recorded EMG activity was used to check the accuracy of the timing of the recorded button presses during off-line analysis.

2.4 Procedure

During the experiment, the participant was placed inside an electricity-shielded experiment booth in front of a computer screen on which the instructions and stimuli were presented. The experimenter sat outside this booth in the experiment room and was able to observe the participant by video camera and microphone. The experiment consisted of a Libet-task and a Matsuhashi-task. Both tasks will be described in detail below.

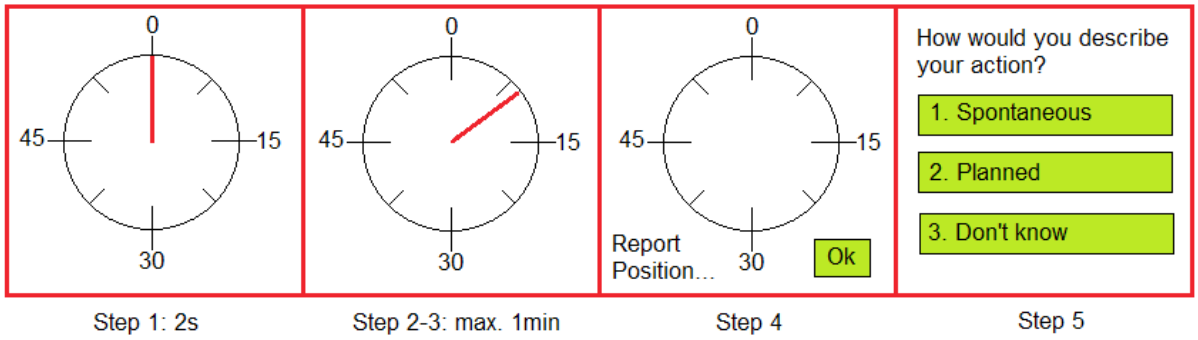

Fig. 2 Stimuli provided to a participant during an intention trial. Step 1. A clock appears on the screen. Step 2-3. The clock starts running and the participant can press either left or right button whenever their intention to do so arises. Step 4. The participant reports the remembered clock position at the time of their intention to act. Step 5. The participant reports whether the act was made spontaneously, planned or don't know. Adapted from Verbaarschot, Farquhar and Haselager (2014).

2.4.1 Libet-task

This task consists of two types of trials: intention trials and sound trials.

Intention trials. During an intention trial, the participant is presented with the following sequence of stimuli (see Fig. 2):

1. A stationary clock is presented for a period of 2 seconds.
2. The clock starts running. The participant is instructed to keep his/her focus on the centre of the clock. After the clock has completed one revolution, the participant is free to decide which button to press and when to press that button. The participant should wait for the feeling of wanting to press either button to arise and remember the position of the clock at the time he/she felt that internal intention to act.
3. When a participant has pressed a button, the clock will stop running after a random interval between 0.5 and 1 s.
4. A clock without a hand along with the words "Report position..." is presented on the screen. The participant should indicate the hand position at the time he/she felt the intention to act by clicking with the computer mouse on that remembered position inside the clock. After the participant has clicked inside the clock, the hand will appear in that position. The participant is free to adapt the hand position by dragging the hand around or clicking on another position. When the participant is satisfied with the indicated clock position, he/she clicks on the "OK" button.
5. The question "How would you describe your action?" is displayed on the screen. The participant can decide by clicking on one of three options: (1) Spontaneous, (2) Planned, (3) Don't know.

Sound trials. Sound trials are used to train a participant in remembering and reporting clock positions and to have a global accuracy measure of the reported positions. The sound trials are very similar to the intention trials, only here the participant should remember the position of the clock and press a button at the time he/she hears an auditory stimulus:

1. A stationary clock is presented on the screen for a period of 2 s (similar to step 1 of an intention trial).
2. The clock starts running. The participant should wait for an auditory stimulus to be presented. This auditory stimulus consists of a simple 'beep' sound. The sound will be played on a random time point within 3 s after the clock started running. The participant should remember the position of the clock at the time they heard the sound and press a button with their right thumb as fast as possible.
3. When a participant has pressed a button, the clock will stop running after a random interval between 0.5 and 1 s (similar to step 3 of an intention trial).
4. A clock without a hand along with the words "Report position…" is presented on the screen. This step is similar to step 4 of an intention trial only now the participant should report the remembered hand position at the time the auditory stimulus was presented.

2.4.2 Matsuhashi-task

This task consists of three types of trials: reaction time (RT) trials, free trials, and no action trials.

Reaction time trials. Similar to the original Matsuhashi-task, here a simple reaction time task was used to assess the elapsed time between the

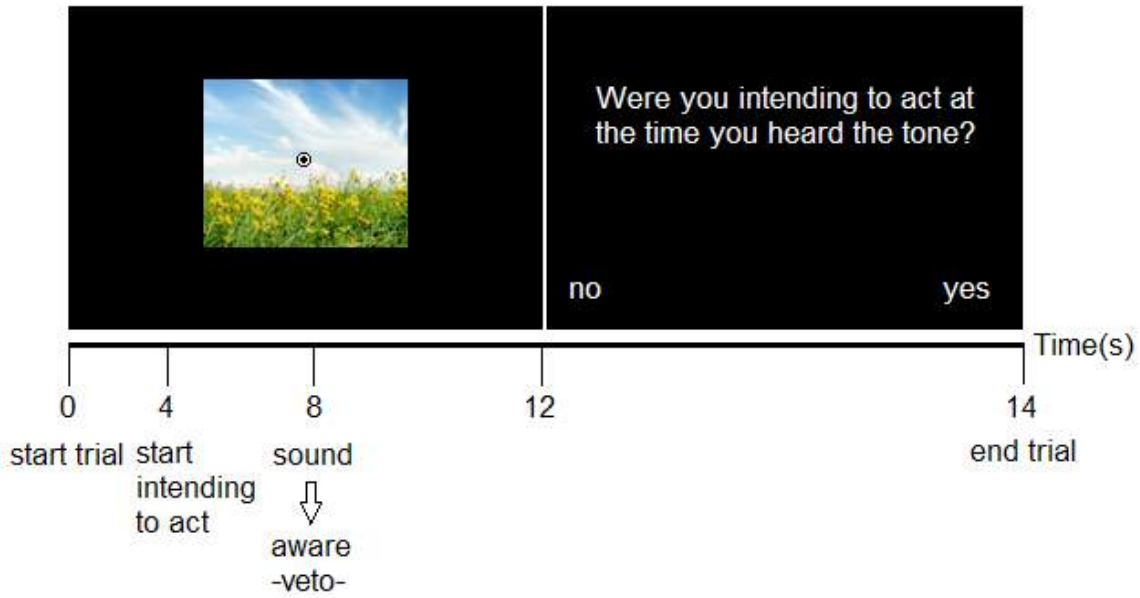

Fig. 3 Sequence of events in a free trial of the Matsuhashi-task in which a participant is aware of intending their act at the moment the auditory probe is presented at 8s and waits for the trial to end at 12s. When the trial ends the participant is asked whether or not they vetoed their act upon hearing the auditory probe.

presentation of an auditory probe and a button press. At the start of each trial, a fixation cross is presented for 1 s followed by a black screen. At a random point in time within an interval of 3 s, an auditory probe consisting of a simple 'beep' sound is presented. Upon hearing the sound, the participant is instructed to press a button with either their left or right thumb as soon as possible. The trials are divided into two blocks, one for left hand responses and one for right hand responses.

Free trials. In these trials, participants watch slides of natural landscape images on a computer screen. Each image is selected manually and consists of a roughly symmetrical image without any animate objects. At the centre of each image, a fixation point is displayed. The participants are instructed to focus on this point and make themselves ready for the upcoming act. Each time a new image is presented on the screen (i.e., the start of a new trial), participants are instructed to empty their thoughts, breathe deeply in and out, and relax for a moment before intending any act. This moment of relaxation is introduced to record a proper baseline period of EEG activity at the start of each trial. While watching slides, participants can press a button with either their left or right thumb to go to the next slide (both actions will trigger a new slide).

Participants are instructed not to plan their actions, but to make them as spontaneously as possible. When participants consistently performed their actions within the first 4 s of each trial, they are asked to act a bit slower in order to maintain a proper baseline period. If participants mostly pressed only one of two buttons, they are instructed to use both hands and try to choose their actions randomly. If participants showed a certain pattern in their acts (alternate between pressing left and right), they are asked to perform their actions a bit more randomly.

Since we introduced separate trials into the adapted task design of Matsuhashi and Hallett (2008), participants no longer acted with a certain frequency which made it difficult to ensure that the auditory probes were presented at an appropriate time in order to probe action-specific intentions. In order to optimize the probe onsets, an individualized probe distribution was calculated for each participant. This distribution is calculated during the training phase of the experiment and updated every 50 trials. It defines a window of time based on the average reaction time and variance of an individual subject within which a probe can be presented. More specifically, this window runs from the size of variance before the average movement onset until the size of variance after movement onset and is shifted by 500 ms back in time to increase the chance that a probe is presented prior to action onset.

If the variance of a participant's reaction time distribution is less than 1.5 s, it is automatically set to 1.5 s in order to keep a minimum window of 3 s within which probes can be presented. Within this window, 70% of the probes occur within the size of the variance before the average movement onset and 30% within the size of the variance after the average movement onset. If the calculated time window falls within the baseline period or outside the maximum

trial length, it is automatically set to a minimum of 3.5 s and a maximum of 10.5 s, respectively. Due to the trial segmentation of the adapted Matsuhashi-task, a maximum trial length was defined in order to continue to the next trial in case of a veto. This trial length was defined semi-randomly such that 80% of the trials lasted 11 s, 10% lasted 12 s and 10% lasted 13 s in order to enhance the feeling of being able to freely decide when to act. Participants were not informed of this maximum trial length, but were told that it could sometimes happen that a next picture appears while they have not yet pressed a button. In that case, the participant is instructed not to worry and continue their normal routine at the start of the new trial.

When a participant presses a button, the current image slides off the screen in the direction of the button that was pressed (i.e., when the left-hand button is pressed, the image slides off to the left of the screen and when the right-hand button is pressed, the image slides off to the right of the screen). As participants are asked to restrain from blinking or moving while the fixation point is present, they can relax for around 2 s when the fixation point disappears and the image slides off the screen. Whenever a trial ends and no button press has been made (caused either by a conscious veto or the absence of an intention to act), a question appears on the screen asking the participant whether or not they had vetoed the act upon hearing the auditory probe (see Fig. 3). This question is asked to check whether the participant has performed a conscious veto or not.

No action trials. These trials are very similar to free trials. In contrast to the free trials, the images slide off automatically to either left or right side of the screen after a random period of time between 5 and 5.5 s. The purpose of these trials is to give the participants a moment of relaxation and bias them towards pressing a button roughly each 5 s without explicitly instructing them to do so.

Questionnaire. At the end of the experiment, the participants filled out a questionnaire on their subjective experience during the free trials of the Matsuhashi-task. The purpose of this questionnaire was to find out whether the acts were made spontaneously, whether participants used a certain strategy to perform an act, and whether the auditory probes had an effect on decision making (e.g., did you want to perform an act prior to the onset of the auditory probe?).

2.4.3 Experimental timeline

In order to optimize the estimates of both RP onset and the onset of intending, we split the experiment in two sessions: a behavioural session and an EEG session. The free trials of the Matsuhashi-task used in the behavioural session contained a lot of auditory probes in order to optimize the intention estimate, whereas they contained very little probes in the EEG session in order to optimize the RP estimate. The participants were not informed of the difference between the free trials of these two sessions and received exactly the same instructions in both sessions.

The behavioural session consisted of three parts: a training, test, and second training session. The first training session consists of 5 RT trials, 10 free trials, and another 20 free trials (which were used to calculate the initial probe distribution for the test session). The test session consists of alternate blocks of 2x20 RT trials (using left or right hand actions only) and 3x50 free trials. The last training session consists of 20 sound trials and 10 intention trials and is used to train the participants on both tasks of the Libet-task. The behavioural session lasted about 60 min (15 min for the first training task and on the intention trials of the Libet-task. session, 30 min for the test session and another 15 min for the second training session).

In the EEG session, EEG data is collected on the free and no action trials of the Matsuhashi-task and on the intention trials of the Libet-task. A total of 6 blocks of 50 free trials is tested. Out of all these blocks, only 2 obey the normal probe distribution at all times. The other 4 blocks obey the probe distribution only on 2 out of 10 trials. The probe onsets of the remaining 8 trials are set to 10 s, thereby highly increasing the chance that a participant will act before the probe is presented (and thus, no auditory probe will be presented on that trial). These free trial blocks are alternated with 4 blocks of 25 intention trials and 2 blocks of 25 action trials. The participant is free to have a break between each block and continues the experiment on a self-paced basis. The EEG session lasted about 75 min. Including 15 min to fill in the questionnaire at the end of the experiment, the total duration of the experiment is about 3 hr.

3. Results

3.1 Action onset

The threshold for muscle activity was set to 20 μV. When the measured EMG exceeded this threshold at a certain point in time, the muscle was assumed to be active, indicating that a button was being pressed. The average difference over all participants between the recorded time of a button press and the onset of EMG activity was around -1.64 ms for both the Libet- and Matsuhashi-task. Since this difference was so small, the timing of a recorded button press, instead of the EMG activity, was used as the onset of an act in both tasks.

3.2 Reaction times (RT)

3.2.1 Matsuhashi-task

The average reaction time over all participants between the moment of probe presentation and the corresponding button press was 295 ms (with a variance of 8 ms) for RT trials (see Table 1).

3.2.2 Libet-task

The average reaction time over all participants between the moment of hearing a simple 'beep' sound and pressing a button was 365 ms (with a variance of 33 ms) for sound trials (see Table 1). The difference between this average reaction time and that of the RT trials of the Matsuhashi-task can be explained by the fact that during sound trials, participants are engaged in a second task: remembering the correct spatial configuration of the clock at the time they hear the beep. The average difference over all participants between the actual and reported sound onsets was 88 ms (with a variance of 1 ms).

3.3 Intention onset

3.3.1 Libet-task

The onset of intending to act (i) was calculated by the following formula:

$$i = \begin{cases} T_1 + \dfrac{a * 2.56}{360} & \text{if } a > b + 40 \\ T_2 + \dfrac{a * 2.56}{360} & \text{otherwise} \end{cases}$$

Where: i = onset of an intention to act, a = angle of the clock hand at the reported intention time, b = angle of the clock hand at the time of the button press, T_1 = time at which the clock hand passed the twelve o'clock position for the next to last time before the button press and T_2 = time at which the clock hand passed the twelve o'clock position for the last time before the button press.

The intention to act is assumed to arrive within

Table 1. Reaction Time (RT) results of the Libet- and Matsuhashi-task. "Recall - sound" indicates the average difference between the recalled and actual sound onset during the sound trials of the Libet-task.

Subject	RT Libet (sec.)	Recall - sound (sec.)	RT Matsuhashi (sec.)
1	0.238	0.116	0.270
2	0.267	0.097	0.344
3	0.239	0.185	0.398
4	0.423	0.072	0.447
5	0.331	0.102	0.223
6	0.292	0.058	0.419
7	0.380	0.071	0.284
8	0.834	0.098	0.278
9	0.260	0.066	0.258
10	0.599	0.055	0.205
11	0.309	0.072	0.229
12	0.206	0.059	0.183
Mean	0.365	0.088	0.295
Var	0.033	0.001	0.008

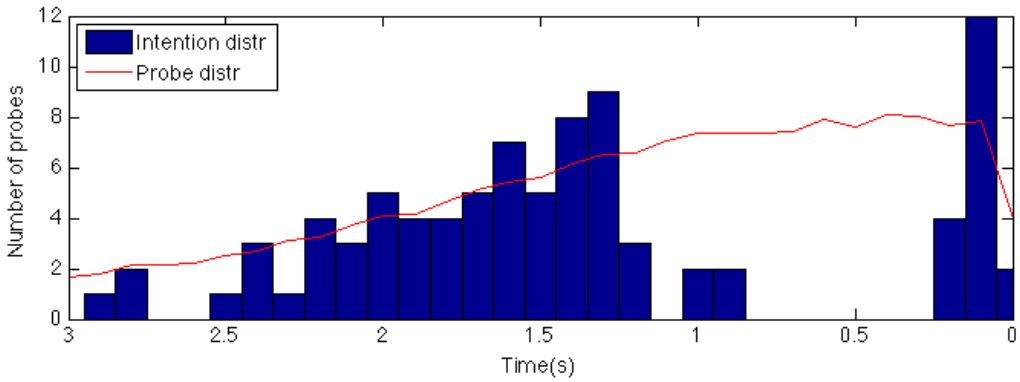

Fig. 4 Intention (blue) versus probe (red) distribution of participant 2. Each bar in the intention distribution represents the average number of times that a probe was presented and ignored at a certain point in time prior to an act. For example, an intention value of 8 at 1.4s prior to action, means that on average, 8 probes were presented, ignored and followed by an action about 1.4s later. In contrast, the red line represents the average theoretical number of probes that have been presented at a certain time prior to action. For instance, around 1s prior to action only 2 probes were presented and ignored (on average), whereas a total of 7 probes has been presented at that point in time (on average). This means that around 1s prior to action around 5 probes have been presented that were not followed by an act since the participant already started intending their act and performed a veto.

2.56 seconds before a button press. However, because on a lot of trials the angle of an intention was reported just a little bigger than the angle of the button press, a small error interval of 40 degrees was introduced. With this error interval, the intention to act can occur one revolution earlier than that of the button press only when the reported angle is more than 40 degrees bigger than that of the button press. It makes more sense that the conscious intention to act is wrongly reported as occurring just a little after the button press than occurring almost 2.5 s prior to the button press, when making spontaneous actions.

3.3.2 Matsuhashi-task

In order to calculate the onset of intending to act, the intention and probe distributions of the free trials of the Matsuhashi-task were calculated for each participant using the following steps:

1. Determine all ignored probe onsets (i.e., a participant heard a beep, ignored it since he/she was not intending any act and pressed a button within this trial) and calculate their onset relative to the subsequent act. This results in the relative intention distribution.
2. Sort the relative intention distribution into bins of 100 ms, running from the earliest ignored probe onset until zero (action onset). This results in the intention distribution as displayed in blue in Figure 4. Each bar in the intention distribution represents the average number of times that a probe was presented and ignored at a certain point in time prior to an action. This means that each of these probes was followed by an act at a later point in time, since the participant was not yet intending their act. As such, the lower the bars in the intention distribution, the higher the chance that a participant was already intending

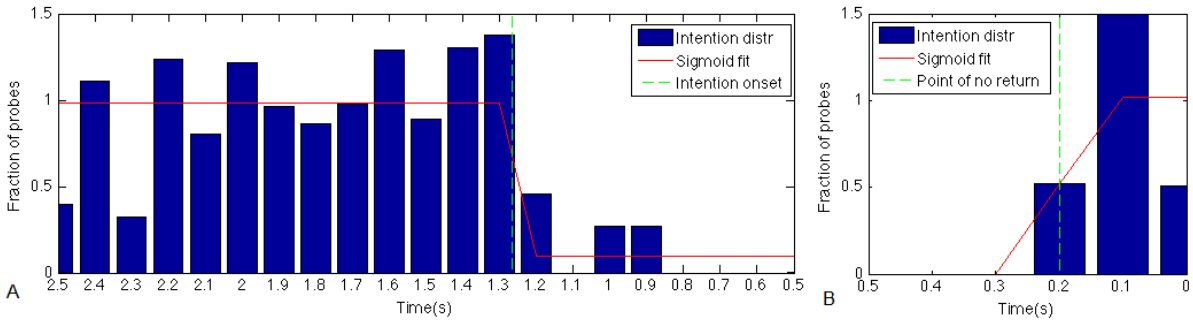

Fig. 5 A. Estimated intention onset (green) of participant 2 using a sigmoid fit (red). The blue bars display the normalized intention distribution, excluding the bins close to action onset. B. Estimated point of no return (green) of participant 2 using a sigmoid fit (red). The blue bars display the normalized intention distribution close to action onset.

their act at that point in time since most probes were followed by a veto and as such did not appear in the intention distribution.

3. Calculate the probe distribution: first of all, delete all probe onsets from the probe distribution that were followed by a non veto act. As the theoretical probe onset on each trial (as described in Section 4) as well as the trials on which a participant performed a veto or a non veto (as reported at the end of each trial without an action) are known, the probe onsets involved in a non veto trial can easily be deleted from the distribution.

4. Calculate the probe distribution relative to the intention distribution: for each act (including acts on trials where no probe was presented), find all probe onsets from the probe distribution which occurred prior to this act (including the probe onsets that were not presented since the participant made an act prior to its pre-determined onset). This results in the relative probe distribution.

5. Bin the relative probe distribution into the same bins that were used for the intention distribution. Normalize the binned probe distribution by dividing it by the total number of observed acts. This results in the probe distribution as displayed in red in Figure 4.

As can be seen in Figure 4, the intention distribution starts to differentiate from the probe distribution close to action onset. All probes that have been followed by a veto will not appear in the intention distribution, thus causing a gap: a period of time close to action onset in which the presented probes were followed by a veto. During this gap, lots of probes were presented, but no longer ignored since the participant was already intending their act. Very close to action onset, the intention distribution becomes similar to the probe distribution again. At this point, although a probe was presented to a participant and they were intending to act, they could no longer veto their act since the probe occurred too close to action onset.

The onset of intending is estimated by fitting a sigmoid to the intention distribution. This sigmoid is of the form:

$$\text{sigmoid} = \frac{a(1) + a(2)}{1 + e^{-\frac{x - a(3)}{a(4)}}}$$

Where a(1): the earliest time point on the horizontal asymptote, a(2): the latest time point on the horizontal asymptote, a(3): the point of inflection and a(4): the width of the inflection, are estimated using the Matlab nlinfit function. In order to fit this sigmoid correctly, the intention distribution is cut off close to action onset. The onset of intending is determined as the turn-over point of the sigmoid (see the dotted green line in Fig. 5A). The point of no return was calculated using a similar sigmoid fit as used for the calculation of the onset of intending.

Table 2. Overview of the intention results, including the point of no return, of both Libet- and Matsuhashi-task.

Subject	Intention Libet (sec.)	Intention Matsuhashi (sec.)	Point of no return Matsuhashi (sec.)
1	-0.325	-1.837	-0.250
2	-0.048	-1.200	-0.200
3	-0.059	-2.478	-0.200
4	-0.241	-3.249	-0.267
5	-0.005	-0.758	-0.080
6	-0.063	-2.672	-0.132
7	-0.135	-2.000	-0.204
8	-0.051	-2.594	-0.199
9	-0.074	-2.946	-0.670
10	-0.514	-1.662	-0.378
11	-0.087	-1.450	-0.199
12	-0.132	-3.196	-0.249
Mean	-0.145	-2.170	-0.252
Var	0.022	0.645	0.023

The difference is that in this case, only the intention distribution close to action onset is used (see Fig. 5B).

3.3.3 Comparison Libet versus Matsuhashi

An overview of the results on the onset of intending in both the Libet and Matsuhashi-task, including the onsets of the point of no return, is provided in Table 2. A one-sided paired-sample t-test confirmed that the intention onsets of the task were significantly earlier than those of the Libet task ($t(11) = 8.55$, $p < .001$, eta2 = .87). Furthermore, a second one-sided paired-sample t-test confirmed that the point of no return had its onset significantly later than the average onset of intending during the Matsuhashi-task ($t(11) = 8.70$, $p < .001$, eta2= .87). However, a third one-sided paired-sample t-test confirmed that the onset of the point of no return did occur significantly earlier than the average onset of intending during the Libet-task ($t(11) = -2.10$, $p = .030$, eta2 = .29).

3.3.4 Questionnaire

Ninety-two percent of the participants reported feeling relaxed or at ease during the free trials of the Matsuhashi-task. None of the participants used a strategy to perform their acts and 92% confirmed that they had made their actions spontaneously or intentionally on each trial. All subjects confirmed that they felt that they were free to choose what action they would like to perform and when they would like to perform it. Fifty percent found it hard to judge whether or not they were already intending their act at the moment that they were presented with a probe. Moreover, multiple participants reported that sometimes, probe presentation induced a feeling of intending to act.

3.4 EEG analysis

In order to analyse the individual trials in the recorded EEG of both free and intention trials of the Matsuhashi- and Libet-task respectively, the raw data was sliced in trials of 8 s surrounding each button press (from seconds before a button press until 2 s after). Each trial was labeled according to the type of button press: a left or a right hand press. Furthermore, the raw EEG data was down sampled from 2048 Hz to 128 Hz. Only the data recorded during the free and intention trials from the EEG session were analysed. For the Libet-task, only intention trials in which the act was reported to be 'spontaneous' (82% of the intention trials over all participants) were used in the analysis. The trials in which the action was reported as 'don't know' (10% of the intention trials over all participants) or 'planned' (8% of the intention trials over all participants) were excluded from the analysis. Furthermore, intention trials in which the button press occurred within the first 4 s from the start of the trial were excluded as well, since in these trials the baseline might include reactions to the stimuli changes at the start of a trial.

For the Matsuhashi-task, only free trials containing a button press were used for EEG analysis. Trials in which the button press occurred within less than 4.5 s from the start of the trial were excluded from the analysis, since in these trials the baseline might include reactions to the stimuli changes at the start of a trial. This period is 0.5 s longer than the 4 s period used in the intention trials of the Libet-task, since the visual stimuli changes at the start of these trials (i.e., the previous image slides off the screen and a new one appears) are slightly larger than those of the intention trials (in which participants are first presented with a stationary clock for 2 s, after which the clock starts running and the trial begins). Furthermore, trials in which a probe was presented within 4.5 s prior to the act were also excluded, since probes cause an additional event related potential in the EEG signal which may interfere with the baseline.

To remove noise from the recorded raw EEG and get a clearer view on the desired RP signal, several pre-processing steps were taken. A linear detrend was used to remove slow drifts. The signal was re-referenced to a common average reference and bad trials were rejected when the data deviated more than 3.5 SDs in variance from the median. Furthermore, eye blinks and movements were removal based on the measured EOG data.

The first 2.56 s of each intention trial of the Libet-task (i.e., the first complete revolution of the clock) were assumed to represent normal brain activity. Since the acts are ensured to occur at minimum 4 s after the clock started running (as previously explained), the EEG signal from 3.5 s until 2.5 s before each button press was used as a baseline in the EEG analysis. For the free trials of the Matsuhashi-task, a similar baseline period was used from 3.5 s until 2.5 s before each act. During this period participants are instructed to relax and breathe deeply in and out, thus ensuring a representation of normal brain activity in the baseline period.

The average signal measured during the baseline

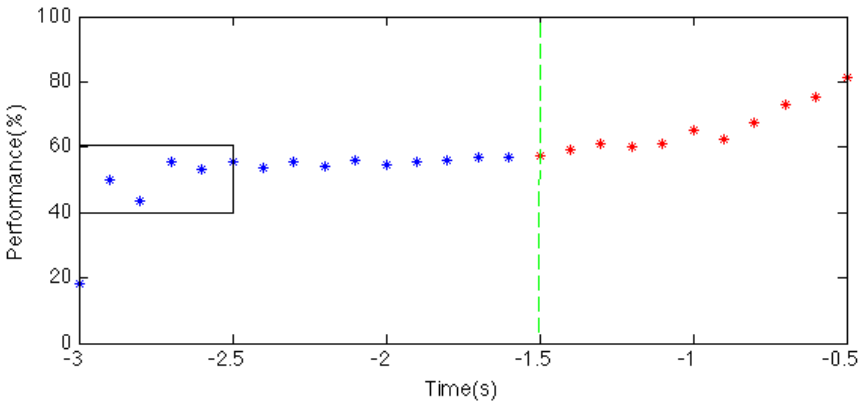

Fig. 6 Example calculation by classifier of the early RP onset of participant 12. The blue and red stars indicate the non significant and significant classification performances over time, respectively. The dotted green line indicates the calculated RP onset. The black rectangle indicates the baseline period, running from -3.5 to -2.5s prior to action.

period was subtracted from the measured signal at each electrode and on each trial. Subtracting the baseline from the EEG data causes the average signal in the baseline period to be zero, leaving only event-related brain activity in the following period. After subtracting the baseline from the EEG data, a band-pass filter with a pass band between 0 and 15 Hz was used to filter out the remaining noise. All frequencies above 30 Hz were completely removed from the data. To finalize these pre-processing steps, the data was further down-sampled to 96 Hz.

3.4.1 RP analysis

To determine the RP onset, the time point at which the RP starts its slow negative wave was calculated using two different measures:

1. RP onset by eye: here the onset of the RP is determined visually by looking backwards from the time point of the button press to the time point where five data points in a row are above the baseline. This method is similar to the method used by Libet et al. (1983) and Keller and Heckhausen (1990). For this measure, only the Cz (vertex) electrode is used since the early RP is most clearly expressed in this electrode. As there are no signs of lateralization within the RP as measured at the Cz electrode, both left and right hand actions were grouped together causing the calculation of the early RP to be based on the average EEG activation over all actions.
2. RP onset by classifier: a linear classifier was trained and tested on baseline versus action data of the Libet- and Matsuhashi-task separately using 10-fold cross-validation. Baseline samples of 500 ms were taken from the baseline period (between -3,000 and -2,500 s prior to action) of each trial. Action samples consisted of 500 ms of data that were collected every 100 ms between -2,500 and 0 s prior to action. Since the classifier is a linear one, it constructs a straight line through the time versus EEG activity space separating action samples from non-action samples. Because of this characteristic, using a linear classifier is a good way to find out when the recorded EEG activity starts to deviate from the baseline consistently over all trials. As such, the early negative shift of the RP can be detected since the correlated action data samples should all be consistently below the baseline. Classification performances above the binomial confidence interval were considered to be significant. The RP onset was calculated as the earliest point in time at which 3 data samples in a row could be classified with a significant performance (see Fig. 6). For this analysis, the activity at the central electrodes (including Cz, C1, C2, C3, C4, FCz, FC1, FC2, FC3, FC4, CPz, CP1, CP2, CP3, CP4) was used.

The RP onset was calculated for each participant using these two measures. An overview of these results can be found in Table 3. Unfortunately, the RPs of the Matsuhashi-task of participant 2 and the Libet-task of participant 3 and 5 were obscured by noise, which made the calculation of the RP by eye quite impossible. However, the onset of the RP could be calculated by classifier for each participant since this measure looks at the consistent activity pattern of single trials. In addition to the RP onsets that were calculated for each participant, the RP onset of the grand average over all participants was also calculated using the RP eye method. The RP onset of the grand average was found to lie at -1,938 s for the Matsuhashi-task and at -2,914 s for the Libet-task. The grand average over Cz, C3 and C4 for the Matsuhashi-task can be viewed in Figure 7.

3.4.2 Lateralized Readiness Potential (LRP) analysis

The LRP was calculated for each participant using the following formula (as described by Trevena and Miller (2002)):

$$LRP = [(C3l - C4l) + (C4r - C3r)]/2$$

Where C3 and C4 are the EEG recordings over the motor cortex of the left and right hemisphere respectively, and l and r indicate the average EEG activity of left or right hand actions respectively. The LRP onset was determined visually by looking backwards from the time point of the button press to the time point where five data points in a row are below the baseline (similar to the estimated RP onsets by eye). An overview of these results can be found in Table 3. Unfortunately, the LRP could not be determined for participant 7 for the Libet-task, since only a single left-hand action was performed. Furthermore, the LRP of the Matsuhashi-task of participant 6 and of the Libet-task of participant 3, 4, and 5 were not clearly visible in the average EEG and could not be calculated by eye.

In addition to the LRP onsets that were calculated for each participant, the grand average of the LRP onset over all participants was also calculated using similar methods. The lateralized nature of the RP as captured by the LRP is clearly expressed in the grand average (see Fig. 7 for the grand average of the free trials of the Matsuhashi-task). The grand average of the LRP onset was found to lie at -0.570 ms for the Matsuhashi-task (see Fig. 8) and at -0.320 s for the Libet-task.

3.4.3 Statistics

Multiple paired-sample and one-sided paired-sample t-tests have been performed in order to establish the temporal order of events prior to action performance. First of all, a paired-sample t-test confirmed that the RP onsets measured by eye and using a classifier did not differ significantly for

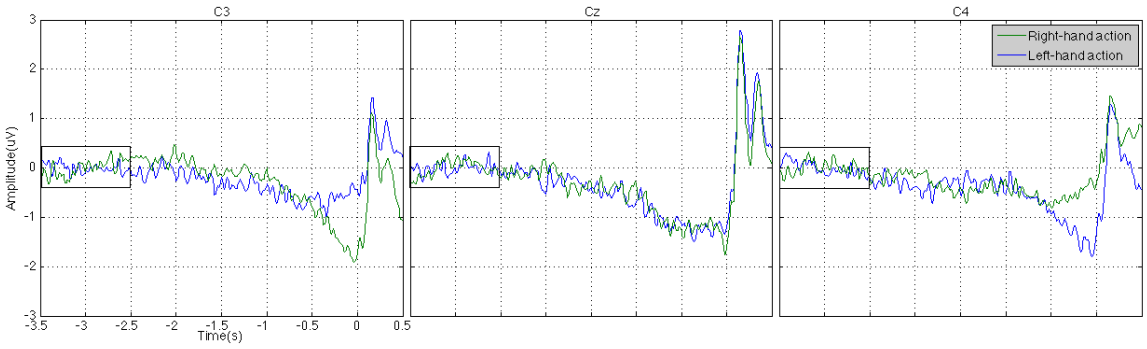

Fig. 7 Grand average of the free trials of the Matsuhashi-task at Cz (vertex), C3 and C4 (left and right motor cortices, respectively). As can be seen in C3 and C4, the lateralized nature of the RP is clearly expressed close to action onset. The black rectangle indicates the baseline period, running from -3.5 to -2.5s prior to action onset.

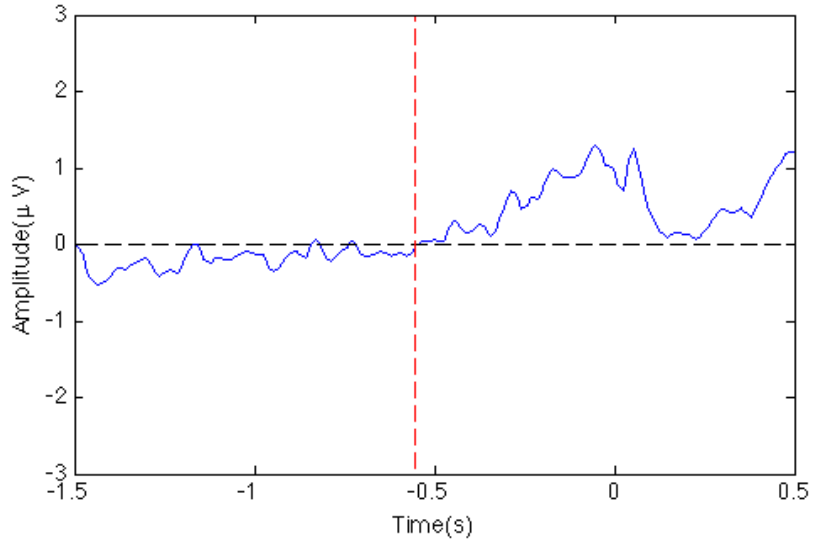

Fig. 8 Estimated onset by eye of the LRP of the grand average over the free trials of the Matsuhashi-task.

Table 3. Overview of the RP and LRP onsets of both the Libet- and Matsuhashi-task. "#Left" and "#Right" indicates the number of left and right hand actions that were used in the EEG analysis. "Min perf." and "Max perf." indicate the minimal and maximal significant performance accuracies of the classifier that was used to calculate the RP onset.

Participant	Task	#Left	#Right	RP eye (sec.)	RP classifier (sec.)	LRP (sec.)	Min perf. (%)	Max perf. (%)
1	Libet	21	75	-2.602	-1.200	-0.273	60.42	68.23
	Matsuhashi	69	78	-2.711	-2.100	-0.563	57.48	72.11
2	Libet	40	40	-2.531	-1.800	-0.320	58.75	75.63
	Matsuhashi	69	50	-	-0.700	-1.367	56.30	62.61
3	Libet	32	43	-	-0.700	-	61.33	66.00
	Matsuhashi	78	82	-2.461	-1.500	-0.398	56.56	69.06
4	Libet	73	12	-2.664	-0.900	-	60.59	66.47
	Matsuhashi	34	51	-1.906	-1.500	-1.352	57.06	70.59
5	Libet	38	57	-	-1.900	-	58.42	68.95
	Matsuhashi	78	54	-0.805	-0.700	-0.375	56.06	57.58
6	Libet	45	39	-2.930	-1.900	-2.305	58.33	66.67
	Matsuhashi	100	110	-1.156	-0.900	-	57.38	62.14
7	Libet	1	76	-2.867	-2.500	-	57.79	70.13
	Matsuhashi	28	138	-0.930	-2.900	-2.813	60.24	65.66
8	Libet	47	29	-3.055	-0.500	-0.313	51.97	67.76
	Matsuhashi	92	76	-1.000	-2.600	-0.414	55.66	68.75
9	Libet	36	39	-2.664	-1.300	-0.352	59.33	76.00
	Matsuhashi	97	108	-2.539	-2.500	-0.914	55.61	67.32
10	Libet	62	23	-2.945	-1.500	-0.820	57.06	76.47
	Matsuhashi	70	91	-1.813	-2.300	-0.477	56.83	79.81
11	Libet	33	28	-2.227	-1.600	-0.383	58.20	77.05
	Matsuhashi	93	99	-2.719	-1.200	-0.531	55.47	71.09
12	Libet	54	39	-0.031	-2.900	-0.539	57.53	81.72
	Matsuhashi	61	65	-3.125	-1.600	-0.227	56.35	72,22
Mean	Libet	40	42	-2.452	-1.558	-0.663	58.31	71.76
	Matsuhashi	72	84	-1.924	-1.708	-0.857	56.75	68.25
Grand average	Libet	482	500	-2.914	-	-0.320	-	-
	Matsuhashi	869	1002	-1.938	-	-0.570	-	-

both the Matsuhashi-task (t(10) = -0.37, p = .722, eta2 = .01) and the Libet-task (t(9) = -1.84, p = .099, eta2 = .27).

Since the RP onsets did not differ significantly between the two analysis methods, the remaining t-tests will be performed with the RP onsets as calculated by classifier only because this method seems to provide the most conservative and reliable results.

A paired-sample t-test confirmed that the average RP onsets did not differ significantly between the Matsuhashi- and Libet-task (RP by eye: t(8) = 0.84, p = .425, eta2 = .08, RP by classifier: t(11) = -0.47, p = .651, eta2 = .02). Another paired-sample t-test confirmed that average LRP on sets did not differ significantly between the Matsuhashi- and Libet-task either (t(6) = 1.16, p = .0291, eta2 = .18). A one-sided paired-sample t-test showed that the average LRP onsets of both the Libet- and Matsuhashi-task occurred significantly later than the average RP onsets of both tasks (LRP Matsuhashi vs. RP Matsuhashi: t(10) = 3.49, p = .003, eta2 = .55, LRP Matsuhasi vs. RP Libet: t(10) = 2:50, p = .016, eta2 = .39, LRP Libet vs. RP Libet: t(7) = 3.15, p = .008,

eta2 = .59, LRP Libet vs. RP Matsuhashi: t(6) = 2.54, p = .019, eta2 = .48).

The temporal order of events between the onsets of intending of both the Libet- and Matsuhashi-task were also compared to the average LRP and RP onsets of both tasks. First of all, a one-sided paired-sample t-test confirmed that the average onset of intending of the Libet-task was found to occur significantly later than the average RP onsets of both the Libet- and Matsuhashi-task (onset intending Libet vs. RP Libet: t(11) = 6.70, p < .0010, eta2 = .80, onset intending Libet vs RP Matsuhashi: t(11) = 7.54, p < .001, eta2=.84). A paired-sample t-test showed that the average onset of intending of the Matsuhashi-task did not differ significantly from the average RP onset of the Matsuhashi-task (t(11) = -1.77, p = .105, eta2 = .22), although a one-sided paired-sample t-test did show that it seemed to occur somewhat earlier than the average RP onset of the Libet-task (t(11) = -1.85, p = .045, eta2 = .24).

Another one-sided paired-sample t-test showed that the average onset of intending during the Libet-task did occur significantly later than the average LRP onsets of both the Libet- and Matsuhashi-task (onset intending Libet vs. LRP Libet: t(7) = 1.98, p = .044, eta2 = .36, onset intending Libet vs. LRP Matsuhashi: t(10) =3.01, p = .007, eta2 = .48). In contrast, a one-sided paired-sample t-test showed that the average intention onset during the Matsuhashi-task did occur earlier than the average LRP onset during both the Libet- and Matsuhashi-task (onset intending Matsuhashi vs. LRP Libet: t(7) = -4.91, p < .001, eta2 = .77, onset intending Matsuhashi vs. LRP Matsuhashi: t(10) = -3.73, p < .001, eta2 = 58).

Finally, the average onset of the point of no return was compared to both the average LRP and RP onsets of the Libet- and Matsuhashi-task. A one-sided paired-sample t-test confirmed that the average point of no return occurred later than the average RP onsets of both Libet- and Matsuhashi-task (point of no return vs. RP Libet: t(11) = 6.08, p < .001, eta2 = .77, point of no return vs. RP Matsuhashi: t(11) = 7.25, p < .001, eta2 = .83). Another one-sided paired-sample t-test showed that the average point of no return occurred later than the average LRP onset of the Matsuhashi-task (t(10) = 2.56, p = .014, eta2 = .40). However, a paired-sample t-test showed that the average onsets of the point of no return did not differ

significantly from the average LRP onsets of the Libet-task (t(7) = 1.14, p = .201, eta2 = .22).

To summarize, the RP onset and the onset of intending as measured during the Matsuhashi-task seem to arise earliest in time and around the same point in time (about -1.911 s prior to action onset). These events are followed by the LRP onset around -0.760 s prior to action onset. The LRP is followed by the point of no return around 0.252 s prior to action. Last but not least, the onset of intending as measured using the Libet-task arises around 0.145 s prior to action. An overview of these results is provided by Figure 9.

4. Discussion

Before we present our conclusions on the results of our experiment in relation to the framework of intending as a process, we will discuss some complexities we encountered during our experiment.

4.1 The onset of intending

4.1.1 Libet-task

Intuitively speaking, most reported intention onsets should occur prior to action onset, since the thought of wanting to act is assumed to come to mind prior to action performance (as can be seen in Fig. 10A, for example). Surprisingly, around 28% of the onsets of intending over all participants were reported to occur after action onset. In contrast to Figure 10A, the reported intentions in these cases seem to surround the button press rather than precede it (see Fig. 10B, for example).

What participants report as the onset of intending seems to be more or less equivalent to action onset, with a certain precision error. Moreover, since the point of no return was found to precede the onset of intending significantly, these intentions seem to be related more strongly to action performance than action initiation. However, the coincidence of the reported onsets of intending and the measured onsets of acting seems quite intuitive since a participant is instructed to perform spontaneous actions without any preplanning or focus on when to act. Thus, although the reported intention onsets still occur prior to action onset on average, it is important to keep in mind that these onsets do not seem to differ much on a single trial level.

4.1.2 Matsuhashi-task

During the free trials of the Matsuhashi-task, multiple participants reported that sometimes, they had the feeling that the probes induced their onset of intending to act. The earlier onsets of intending

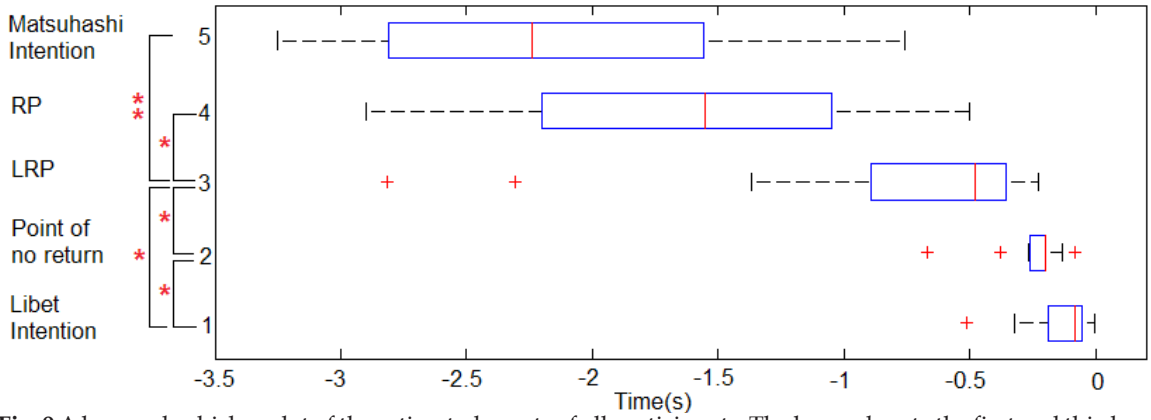

Fig. 9 A box-and-whisker plot of the estimated onsets of all participants. The boxes denote the first and third quartiles of the data and the whiskers extend to the most extreme data points (outliers are denoted by a little red cross). The median is denoted by a vertical line inside a box. The RP box incorporates the estimated RP onsets by classifier for both the Libet- and Matsuhashi-task (since no significant difference was found between the two). Similarly, the LRP box incorporates the estimated LRP onsets by eye for both the Libet- and Matsuhashi-task (since no significant difference was found between the two). Two red stars indicate that the estimated onsets differed significantly between the indicated two groups, with $p < .001$. A single red star indicates a similar significant difference, only with $p < .05$.

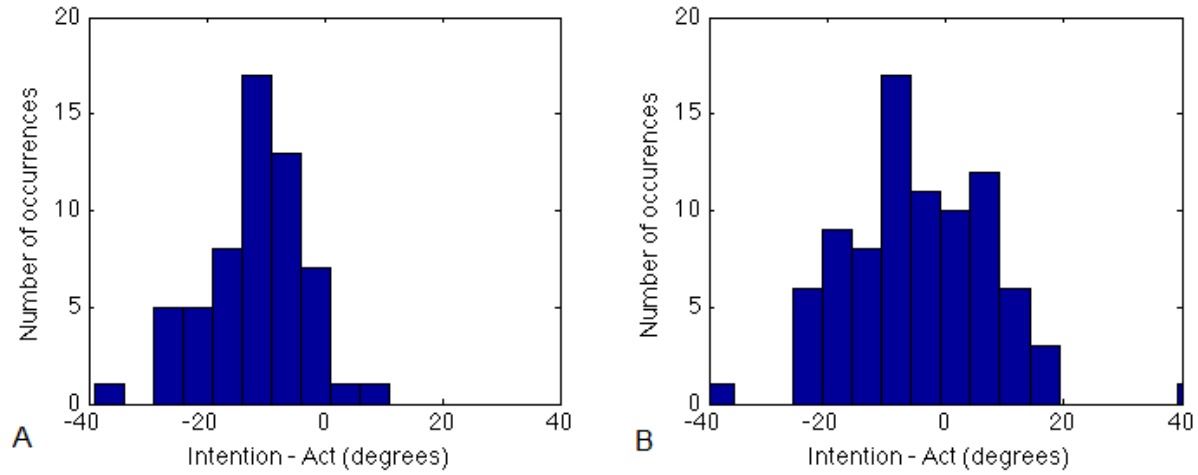

Fig. 10 Action onset versus reported onset of intending during the intention trials of the Libet-task of participant 11 (A) and 6 (B). The x-axis denotes the difference in degrees between the onset of an act and the reported intention onset, where a negative difference indicates that the intention was reported prior to the action.

of the Matsuhashi-task compared to those of the Libet-task could be explained by the learned response mapping between the probe and an act of the reaction time trials of the Matsuhashi-task, causing a vague intention to act after each probe presentation. However, since the reported onsets of intending of the Matsuhashi-task occur consistently and significantly earlier in time compared to those of the Libet-task, it seems unlikely that this difference can be explained as a side-effect of the probes alone. Actually, these results fit rather well in the proposed framework of intending as a process, as the use of external probes seems to enable participants to report their onset of intending before they are be able to do so on their own. In other words, although the presented probes make a participant aware of their intention to act, they do not seem to cause their intention to act.

5. Conclusion

The results of the within-subject comparison between the Libet- and Matsuhashi-task indicate that the measured onsets of intending indeed differ significantly between the two. More specifically, the onset of intending of the Matsuhashi-task preceded that of the Libet task by 2 s on average. It appears that intending has different onsets when it is measured in different ways. Although these results present a problem for those who assume that intentions are discrete mental states occurring

at a single point in time, they support the proposed conceptual framework describing intending to act as a multistage process developing over time. Whether one is engaged in the process of intending seems to become available for self-initiated report during the later stages of this process only, at a time that nearly coincides with action onset. However, using external probes, participants can be made aware of their intention to act during earlier stages of this process.

The average onset of intending measured using the Matsuhashi-task did not differ significantly from the RP onset, suggesting that these two processes have their onset around the same point in time. However, as other neural signatures in the frontal cortex have been found to correlate with action initiation much earlier in time (Bode et al., 2011; Soon et al., 2008), the RP seems to play a part in only the final stages of action preparation. The current experiment, along with previous research (Fried, Mukamel, & Kreiman, 2011; Haggard & Eimer, 1999; Libet et al., 1983; Matsuhashi & Hallett, 2008; Schurger et al., 2012; Trevena & Miller, 2002; Trevena & Miller, 2010), has focused only on those stages of intending that are close to action onset. However, the processes of acting and intending start much earlier in time and seem to be intertwined, starting in the frontal cortex and moving up to the primary motor cortex via the supplementary motor cortex (Bode et al., 2011; Soon et al., 2008). More research is needed to differentiate the neural processes relating to acting from those relating to intending and/or becoming conscious of intending. Our hope is that the new conceptual framework (as proposed in this thesis) can contribute to that by enabling a more robust interpretation of research results and by setting the stage for new experiments.

6. References

Bai, O., Rathi, V., Lin, P., Huang, D., Battapady, H., Fei, D.-Y., & Hallett, M. (2011). Prediction of human voluntary movement before it occurs. *Clinical Neurophysiology*, *122*(2), 364–372.

Blankertz, B., Dornhege, G., Lemm, S., Krauledat, M., Curio, G., & Müller, K. R. (2006). The Berlin Brain-Computer Interface: Machine learning based detection of user specific brain states. *Journal of Universal Computer Science*, *12*(6), 581-607.

Bode, S., He, A., Soon, C., Trampel, R., Turner, R., & Haynes, J. (2011). Tracking the unconscious generation of free decisions using ultra-high field fMRI. *PLoS One*, *6*(6), 1–13.

Brass, M., & Haggard, P. (2007). To do or not to do: The neural signature of self-control. *Journal of Neuroscience*, *27*(34), 9141–9145.

Brass, M., & Haggard, P. (2008). The what, when, whether model of intentional action. *The Neuroscientist*, *14*(4), 319–325.

Dennett, D. C. (1991). Consciousness explained. London, England: *The Penguin Press*.

Dennett, D. C. (2003). Freedom evolves. London, England: *Penguin Books*.

Fried, I., Mukamel, R., & Kreiman, G. (2011). Internally generated preactivation of single neurons in human medial frontal cortex predicts volition. *Neuron*, *69*(3), 548–562.

Haggard, P. (2008). Human volition: towards a neuroscience of will. *Nature Reviews Neuroscience*, *9*(12), 934–946.

Haggard, P., & Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research*, *126*(1), 128–133.

Hallett, M. (2007). Volitional control of movement: The physiology of free will. *Clinical Neurophysiology*, *118*(6), 1179–1192.

Hoffstaedter, F., Grefkes, C., Zilles, K., & Eickhoff, S. B. (2013). The "what" and "when" of self-initiated movements. *Cerebral Cortex*, *23*(3), 520-530.

Keller, L., & Heckhausen, H. (1990). Readiness potentials preceding spontaneous motor acts: voluntary vs. involuntary control. *Electroencephalography and Clinical Neurophysiology*, *76*(4), 351–361.

Klem, G. H., Luders, H. O., Jasper, H., & Elger, C. (1999). The ten-twenty electrode system of the International Federation. *Electroencephalography and Clinical Neurophysiology Supplement*, *52*, 3–6.

Kornhueber, H. H., & Deecke, L. (1965). Hirnpotentialanderungen bei willkurbewegungen und passiven bewegungen des menschen: Bereitschaftspotential und reafferente potentiale. *Plugers Archiv*, *284*(1), 1–17.

Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential): The unconscious initiation of a freely voluntary act. *Brain*, *106*(3), 623–642.

Matsuhashi, M., & Hallett, M. (2008). The timing of the conscious intention to move. *European Journal of Neuroscience*, *28*(11), 2344–2351.

Mele, A. R. (2010). Testing free will. *Neuroethics*, *3*(2), 161–172.

Pfurtscheller, G., & Aranibar, A. (1979). Evaluation of event-related desynchronization (ERD) preceding and following voluntary self-paced movement. *Electroencephalography and Clinical Neurophysiology*, *46*(2), 138–146.

Pockett, S. (2004). Does consciousness cause behaviour? *Journal of Consciousness Studies*, *11*(2), 23–40.

Schurger, A., Sitta, J. D., & Dehaene, S. (2012). An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proceedings of the National Academy of Sciences*, *109*(42), 1–10.

Severens, M. (2009). Brainstream public site. Internet:

http://www.brainstream.nu/.

Shibasaki, H., & Hallett, M. (2006). What is the bereitschaftspotential? *Clinical Neurophysiology*, *117*(11), 2341–2356.

Soon, C. S., Brass, M., Heinze, H.-J., & Haynes, J.-D. (2008). Unconscious determinants of free decisions in the human brain. *Nature Neuroscience*, *11* (5), 543–545.

Trevena, J. A., & Miller, J. (2002). Cortical movement preparation before and after a conscious decision to move. *Consciousness and Cognition*, *11*(2), 162–190.

Trevena, J. A., & Miller, J. (2010). Brain preparation before a voluntary action: Evidence against unconscious movement initiation. *Consciousness and Cognition*, *19*(1), 447–456.

Uithol, S. (2012). Representing action and intention (Unpublished doctoral dissertation). Donders Institute for Brain, *Cognition and Behaviour*, Nijmegen.

Verbaarschot, C., Farquhar, J., & Haselager, P. (2014). Lost in time: The search for conscious intentions and readiness potentials.(Manuscript in preparation).

Wegner, D. M. (2002). The illusion of conscious will. Cambridge, MA: *The MIT Press*.

Abstracts

Proceedings of the Master's Programme Cognitive Neuroscience is a platform for CNS students to publish their Master thesis. 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.

Perceptual Learning Increases Sampling Efficiency

Denise Moerel, Sam Ling, Janneke Jehee

Although visual orientation discrimination has been shown to improve with training, the mechanism(s) underlying this task improvement remain poorly understood. One possibility, that is investigated in this paper, is that an increase in sub-sampling efficiency with practice causes a reduction in orientation discrimination thresholds. That is, a more reliable estimate of the stimulus orientation could be obtained by sampling a larger part of this stimulus. We used the classification image method (Ahumada, 1996), which relates the response of the observer to the orientation variability at different locations in the stimulus, to test whether a change in sub-sampling occurs with perceptual learning on an orientation discrimination task. The results showed that a decrease in orientation thresholds was accompanied by an increase in sub-sample size. An effect for eccentricity was observed; while the decision bias, the part of the stimulus that drives the decision, was restricted to the inner part of the stimulus before training, an outward spread of sampling with training was observed. Furthermore, a radial bias was observed after training, with the largest spread in sub-sampling observed for the top quadrant aligned with the stimulus orientation. These results showed an increase in sub-sampling efficiency with training, making it a likely mechanism to underlie perceptual learning.

On the Specificity of Neuronal Oscillations in Human Visual Cortex

Floris de Vries, Mathilde Bonnefond, Ole Jensen, Floris de Lange, René Scheeringa, Peter Kok

Oscillatory neural activity plays an important role in cognition. In this combined electroencephalography (EEG)/ functional magnetic resonance imaging (fMRI) study we explore the spatial and representational specificity of oscillatory activity in alpha (8-12 Hz) and gamma (30-90 Hz) frequency bands in human visual cortex by means of a localiser-based analysis of fMRI data. These frequency bands are thought to represent the functional inhibition of brain areas and the neural representation of a stimulus, respectively. In line with the functional specialisation of oscillatory frequency bands, it may be expected that alpha activity is more global than gamma activity, which in turn is expected to be representationally specific. In this thesis, an experimental and analysis pipeline is described which allows for testing this hypothesis. Pilot results are shown which demonstrate that stationary gratings can be used for eliciting gamma components in fMRI/ EEG analysis and demonstrate that the parameters determined by means of two EEG pilot experiments can be validly transferred to an EEG/fMRI context.

Expectation Suppression Depends on Working Memory but not Perceptual Load

Christian Utzerath, Elexa St. John-Saaltink, Floris de Lange

To cope with the amount of information in the sensory world, predictive coding suggests that efficient observers suppress irrelevant but predictable information. However, such expectation suppression (ES) occurs inconsistently across studies. Previous research raises the possibility that this might be explained due to the tasks towards which attention is diverted. To test this hypothesis, we manipulated subjects' attention and perceptual expectation during functional magnetic resonance imaging (fMRI). We exposed subjects to predictive and non-predictive audiovisual stimulus pairings while varying their available processing resources. Participants would either attend the (un-)predicted stimuli or their attention would be diverted to a task that increases perceptual or working memory load. We further aimed to explore the whole-brain correlates of perceptual expectations. Several brain regions, among which hippocampus, superior temporal gyrus, and cerebellum were more responsive to predictive than non-predictive pairings. Converging evidence suggests that superior temporal gyrus (STG) in particular responds to the informational value contained in predictive cues. Importantly, we found ES occurs for unattended stimuli, yet only under perceptual but not working memory load. This may explain why previous research failed to replicate ES. It furthermore dissociates expectation and repetition suppression, suggesting that the two are implemented through different physiological mechanisms. Predictive sequences furthermore induced increases in functional connectivity between STG, visual cortex and a thalamic region projecting to prefrontal cortex. We tentatively suggest that top-down modulations of visual perception and working memory share a common neural substrate, possibly in prefrontal cortex.

Investigating the Lexical Representation of Highly Reduced Word Forms With the Mismatch Negativity

Martijn Bentum, Mirjam Ernestus, Kimberley Mulder

Spoken language is challenging to comprehend because the acoustic signal of speech can be highly variable. Research on speech perception has shown that the representation of pronunciation variants is of theoretical importance for models of speech perception. The mismatch negativity (MMN) has been used to investigate the representation of words in long-term memory. The current study aims to extend this research by investigating the representation of highly reduced pronunciation variants in the mental lexicon. Two experiments were conducted and no reliable MMN enhancement effect was found and therefore, the results do not provide evidence concerning the representation of reduced pronunciation variants. Unexpectedly, subsequent exploratory data analysis showed that a P3a response was elicited for contrasts of pseudowords with reduced forms and real words, while no such response was elicited with pseudoword-pseudoword contrasts. This finding is in line with the idea that lexically represented word forms are processed differently than forms that do not have a long-term memory representation. Further research is necessary to investigate whether the P3a is sensitive to the lexical status of spoken word forms.

Development of Impulsivity in Autism Spectrum Disorders: an Event-Related Potential Study

Renske van der Cruijsen, Edita Poljac, Sabine Hunnius

Impulsivity is commonly seen in individuals with autism spectrum disorder (ASD), but its neurocognitive mechanisms have not been investigated so far. In this electroencephalographic (EEG) study, a voluntary task switching (VTS) paradigm was used to investigate impulsivity and its underlying neuronal markers in ASD. We focused on two EEG markers, one of global task preparation – a slowly developing frontal negativity: contingent negative variation (CNV) – and one of stimulus processing – a positive wave that develops around 300ms after stimulus presentation (P3). Participants with ASD and controls were matched on age and IQ. All participants were asked to choose between two simple cognitive tasks, press a space bar to indicate that they had made their choice and respond according to the chosen task. Impulsivity was measured by choice time (CT): the time participants took to press the space bar. Participants with short CTs were assumed to be more impulsive than participants with long CTs. Participants completed a questionnaire as an additional measure of impulsivity. Results indicated that individuals with ASD indeed scored higher on impulsivity as measured by the questionnaire than controls and that impulsivity generally decreases with age. Participants with shorter CTs showed a weaker CNV, so the participants who were assumed to be more impulsive showed less task preparation than participants who were assumed to be less impulsive (with long CTs). The CNV became stronger with age in fast choosers, while in deliberate choosers, the CNV became weaker with age. P3 amplitude was similar for fast and slow choosers and it became weaker with age. The EEG markers did not differ between individuals with ASD and controls. Altogether, this study provided promising first results on impulsivity, its development and its neuronal markers in ASD, but more research is required to better explain and understand these neurocognitive mechanisms.

Alpha and Gamma Band Oscillations Index Differential Processing of Acoustically Reduced and Unreduced Words

Linda Drijvers, Kimberley Mulder, Mirjam Ernestus

Reduced words like yeshay for yesterday often occur in natural conversational speech. Previous behavioural research reported a processing advantage for unreduced words over reduced words. The present study investigated whether this processing advantage is reflected in a modulation of alpha (8-12 Hz) and gamma (30+ Hz) band activity during unreduced and reduced word processing. In three electrophysiological experiments (using electroencephalography; EEG), participants listened to unreduced and reduced infinitives in isolation (1), in sentence-final position (2), or in mid-sentence position (3). Alpha power was larger in response to reduced words than in response to unreduced words in only the first two experiments, whereas gamma power increases were only observed in response to unreduced words, in all experiments. Our findings suggest that alpha power increases reflect a higher auditory cognitive load during lexical processing, while the observed gamma increases in response to unreduced words reflect the activation of semantic representations. This study lends further plausibility to differential processing mechanisms for unreduced and reduced words and gives insight into how neurophysiological resources are allocated to optimize reduced word processing.

Functional Specialization for Case Processing and Phonotactics in LIFG: Improving Tests for Functional Specialization?

Friederike Seyfried, Julia Uddén

The left inferior frontal gyrus (LIFG) is a node in the language network associated with different linguistic processing tasks. A meta-analysis by Bookheimer (2002) suggests that there is functional specialization for different linguistic processing tasks, such as semantics, syntax, and phonology, in the LIFG. This has not been experimentally tested with a within-subject design. To experimentally test whether there is functional specialization in the LIFG for different domains of linguistic processing, we presented German sentences containing either a case violation (syntactic manipulation) or a phonotactically illegal German sound cluster (phonotactic manipulation) to 26 native German speakers in an event-related fMRI experiment. Looking at three regions of interest (ROI) in the LIFG: the pars opercularis, pars triangularis, and pars orbitalis, we found that both manipulations activated the pars opercularis and pars triangularis. We further analyzed our data to test for functional specialization for phonotactics and syntax by comparing the coordinates of cluster peaks activated by the two violations at the single-subject level. We did not find significant differences between the coordinate values for syntactic and phonotactic manipulations. Our data does not support functional specialization in the LIFG for phonotactics and syntax.

Behavioural Rigidity in Adults With ASD: Investigation of the Underlying Neurocognitive Mechanisms

Vincent Hoofs, Edita Poljac

This study investigated the underlying mechanism of behavioural rigidity in Autism Spectrum Disorders (ASD) by means of a Voluntary Task Switching (VTS) paradigm while recording behavioural and electroencephalography (EEG) data. The VTS paradigm requires the participants to perform two simple cognitive tasks, and to choose themselves which task to perform in each trial. At the moment, there are two different views regarding the nature of the underlying mechanism of behavioural rigidity in ASD. The initial view states that persons with ASD exhibit impaired task switching abilities, whereas the modern view states that persons with ASD experience impaired intentional control as the underlying mechanism for behavioural rigidity. The current study aimed to contribute to this debate by comparing the quality of different processes which are required to perform this paradigm correctly between an ASD and a control group. The behavioural data revealed a stronger tendency in the ASD group to repeat tasks (i.e. enlarged repetition bias) as indicated by increased run lengths. Moreover, the task execution data indicated comparable switch costs (reaction times (RT) and errors) between the groups, suggesting intact task switching abilities in the ASD group. Analysis of the task preparation data and of the stimulus repetition data could not provide evidence in support of the view of impaired intentional control in ASD. The EEG data revealed differences between the ASD and the control group in the P3 and Contingent Negative Variation (CNV), two Event Related Potential (ERP) components on which was focused in this study. The P3, previously shown to be modulated by the quality of the task switch preparation, was decreased and delayed in the ASD group on switch trials. The CNV, which was suggested to track the active intentional control when switching from the shape to the location task, was also attenuated in the ASD group. Together, these findings provide evidence supporting the view that behavioural rigidity in ASD is likely to be caused by deviations in intentional control processes.

Subadditive Auditory Cortex Activation to Audiovisual Speech Measured with Functional near-Infrared Spectroscopy

Guus van Bentum, Marc van Wanrooij, John van Opstal

To understand audiovisual speech, multisensory information from different parts of the brain needs to be integrated. Several non-invasive techniques can be used to study the neural processes underlying such integration, and are useful to study speech processing in healthy adults. Most of these techniques, however, are limited when it comes to children or subjects with electro-magnetic cochlear implants (CI). Here, we assess multisensory integration of audiovisual speech in auditory cortex with functional near-infrared spectroscopy (FNIRS). We use unisensory (audio/video only) or multisensory (audiovisual) speech stimuli in a passive experiment. We find weak (0.2 μM concentration change) evoked responses of auditory cortex to both unisensory and multisensory stimuli, in both hemispheres across participants. Modeling the evoked response was done using a canonical hemodynamic response function and a Bayesian approach to calculation regression coefficients. We demonstrate some sub-additive multisensory enhancement for both oxyhemoglobin and deoxyhemoglobin measurements. Our experiments show that FNIRS is able to determine multisensory integration within the brain on a group level for a selection of participants. The possibility that the sub-additive enhancement we measured was actually some form of saturation in the blood oxygenation level dependent (BOLD) response was not accounted for. Problems using FNIRS in the current set up were high levels of measurement noise, high response variance between sessions and unknown physiological artefacts. This led us to conclude that FNIRS in its current state is not a feasible tool for measuring CI-patients and infants.

Neural Time Course of Language-Attention Interactions in Spoken Word Processing

Jana Krutwig, Yury Shtyrov

To investigate how the complex neural interplay between language and attention unfolds in time, we systematically contrasted electrophysiological responses (event-related potentials, ERPs) to spoken words and meaningless pseudowords in a classical oddball paradigm while participants either actively attended to the linguistic input or were distracted from it. The mismatch negativity (MMN) responses to these stimuli differed early on between the stimulus types and conditions. Already at 130-140 ms after word-recognition point, we found a significant interaction between attention, lexicality and topography. Specifically, at left-frontal recording sites, MMN responses were more negative-going for all attended than unattended stimuli. In right-frontal electrodes, however, we found an interaction between attention and lexicality that was driven by the responses to words being relatively immune to attentional modulation, while the pseudoword responses significantly increased with attention. Furthermore, a contrast between the positive-going response to attended words and that to attended pseudowords brought out a typical centro-parietal N400 effect, i.e. more negative-going pseudoword activity. In this N400 interval, attention clearly enhanced responses to both words and pseudowords equally. In the "ignore" condition, however, no significant N400 was observed. These results suggest that while word responses initially appeared more robust, pseudoword ERPs were subject to stronger attention-modulation. Later on, attention effects are more pervasive. Unlike previous lexical MMN studies, we also demonstrate that an N400-like effect (pseudoword>word) can be successfully elicited in an oddball design, suggesting that it is to a degree resilient to multiple stimulus repetition. It is, however, limited to "attend" conditions only and therefore likely reflects later top-down controlled processes in word processing, whilst the first response at ~130 ms more likely reflects initial automatic stages of lexical access.

The Developmental Trajectory of Intentional Control in Individuals With Autism Spectrum Disorder

Rianne Haartsen, Edita Poljac, Sarah Gerson

Individuals with Autism Spectrum Disorder (ASD) show rigid behaviours and prefer sameness. It has recently been suggested these rigid behaviours might be accounted for by less intentional control during the formation of intentions. In addition, ASD is a developmental disorder. The developmental trajectory of intentional control in individuals with ASD might therefore also deviate from the trajectory in their typically developing peers. The developmental trajectory of intentional control and behavioural rigidity in individuals with ASD, however, is currently unknown. This study was designed to investigate whether the developmental trajectory of intentional control in individuals with ASD deviates from the trajectory in typically developing (TD) individuals. We invited participants with ASD and typically developing participants between 12 and 30 years old that were matched on age, IQ, and gender to our study. Participants were asked to voluntarily choose between matching a stimulus on location or on shape in the voluntary task switching paradigm with double registration. The double registration procedure consists of a space bar press that participants perform when they have made their choice and marks the time of the formation of intentions. Reaction times, error rates, task choices, and electroencephalography (EEG) during task preparation were simultaneously recorded. The behavioural results showed no differences in developmental trajectory between individuals with ASD and TD individuals in the formation of task intentions or task execution. EEG markers (contingent negative variation (CNV), preparatory alpha power, and post-stimulus P3) seemed to show a less pronounced or even no development in individuals with ASD compared to their control peers. We suggest that individuals with ASD might experience less top down control later in life and discuss how less top down control might affect different preparatory processes in task switching.

Retaking Control: Third Language Learning Modulates Inhibitory Control in Language Production

Henriette Raudszus, Ton Dijkstra, Ardi Roelofs, Kristin Lemhöfer

Inhibition has been proposed as a mechanism allowing speakers to produce words in one language while preventing or reducing competition from words in the other languages they know. However, little is known about the development of this mechanism in the course of foreign language learning. In a longitudinal production study with two sessions, we examined the development of inhibition ability in German-English bilingual speakers during an intensive Dutch (third language; L3) course. Speakers named pictures in blocked and switched conditions, once in the beginning of the language course when L3 proficiency was very low, and once after 3-4 weeks of instruction when L3 had reached intermediate proficiency. Naming response time (RT) and event-related potentials (ERPs) were measured. We found a switch cost asymmetry with larger switch costs for the second language (L2) than for the first language (L1) in the first session. In the second session, however, switching costs were symmetrical, accompanied by lower L2 switch RT than in the first session. The decrease in switch cost correlated with inhibitory abilities as measured by the Simon task. Furthermore, L1 mixing costs decreased from the first to the second session, reflecting improved global control. In the ERPs, N2 amplitude was larger in the second than in the first session. Taken together, these findings support the notion that inhibitory control changes during L3 learning. We also found evidence that L2 needs to be suppressed more in an L3 than in an L1 context. This can be attributed to effects of foreign language status or to a difference in relative proficiency.

Neurocognitive Mechanisms of Cognitive Flexibility in Dyslexia

Vera van 't Hoff, Edita Poljac

Although it is well established that individuals with dyslexia show several deficits in cognitive control functions in addition to their primary reading disorder, studies investigating cognitive flexibility abilities in this population yielded mixed results. The current study therefore further examined cognitive flexibility capacities in dyslexia as well as the underlying neurocognitive mechanisms. To this aim, we assessed task-switching abilities, as measure of cognitive flexibility, in students with dyslexia and in their control peers. Participants were asked to voluntarily respond either to the location (task A) or shape (task B) of presented stimuli. Importantly, the two tasks differed in difficulty with respect to stimulus-response (S-R) mappings: while the location task had direct S-R mappings, in the shape task these mappings were arbitrarily chosen. Furthermore, to manipulate general task demands, each stimulus was preceded by either a short or long preparatory interval. In addition to participants' response times and accuracy, we recorded their brain activity with electroencephalography (EEG) during task execution. Whereas individuals with dyslexia made more errors immediately after a task-switch compared to the control group, no significant group differences were found in response times. However, EEG results indicated that individuals with dyslexia needed a stronger preparation in shape-specific switch trials, as reflected by an increased contingent negative variation (CNV), to reach the same performance level in behaviour as controls. Altogether, the findings demonstrate that dyslexia is accompanied by cognitive flexibility difficulties.

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
www.ru.nl/neuroimaging/

Donders Institute for Brain, Cognition and Behaviour:
Centre for Neuroscience
Geert Grooteplein Noord 21, hp 126
6525 EZ Nijmegen

P.O. Box 9101
6500 HE Nijmegen
www.ru.nl/neuroscience

Donders Institute for Brain, Cognition and Behaviour:
Centre for Cognition
Montessorilaan 3
6525 HR Nijmegen

P.O. Box 9104
6500 HB Nijmegen
www.ru.nl/cognition/

MAX-PLANCK-GESELLSCHAFT

Max Planck Institute for Psycholinguistics
Wundtlaan 1
6525 XD Nijmegen

P.O. Box 310
6500 AH Nijmegen
http://www.mpi.nl

IN·DEI·NOMINE·FELICITER

Radboud Institute for Molecular Life Sciences
Geert Grooteplein 28
6525 GA Nijmegen

P.O. Box 9101
6500 HB Nijmegen
http://www.ncmls.nl

Universitair Medisch Centrum St Radboud
Geert Grooteplein-Zuid 10
6525 GA Nijmegen

P.O. Box 9101
6500 HB Nijmegen
http://www.umcn.nl/

Baby Research Center
Montessorilaan 3
6525 HR Nijmegen

P.O. Box 9101
6500 HB Nijmegen
http://babyresearchcenter.nl